

Fractal Features of Dark, Maintained, and Driven Neural Discharges in the Cat Visual System

Steven B. Lowen,^{*,1} Tsuyoshi Ozaki,[†] Ehud Kaplan,[‡]
Bahaa E. A. Saleh,^{*} and Malvin C. Teich^{*,§,2}

**Department of Electrical & Computer Engineering and §Department of Biomedical Engineering, Boston University, 8 Saint Mary's St., Boston, Massachusetts 02215; †The Rockefeller University, 1230 York Avenue, New York, New York 10021; and ‡Department of Ophthalmology, Mt. Sinai School of Medicine, One Gustave Levy Place, New York, New York 10029*

We employ a number of statistical measures to characterize neural discharge activity in cat retinal ganglion cells (RGCs) and in their target lateral geniculate nucleus (LGN) neurons under various stimulus conditions, and we develop a new measure to examine correlations in fractal activity between spike-train pairs. In the absence of stimulation (i.e., in the dark), RGC and LGN discharges exhibit similar properties. The presentation of a constant, uniform luminance to the eye reduces the fractal fluctuations in the RGC maintained discharge but enhances them in the target LGN discharge, so that neural activities in the pair cease to be mirror images of each other. A drifting-grating stimulus yields RGC and LGN driven spike trains similar in character to those observed in the maintained discharge, with two notable distinctions: action potentials are reorganized along the time axis so that they occur only during certain phases of the stimulus waveform, and fractal activity is suppressed. Under both uniform-luminance and drifting-grating stimulus conditions (but not in the dark), the discharges of pairs of LGN cells are highly correlated over long time scales; in contrast discharges of RGCs are nearly uncorrelated with each other. This indicates that action-potential activity at the LGN is subject to a common fractal modulation to which the RGCs are not subjected. © 2001

Academic Press

The sequence of action potentials recorded from cat retinal ganglion cells (RGCs) and lateral geniculate nucleus (LGN) cells is always irregular. This is true whether the retina is in the dark (1, 2), or whether it is adapted to a stimulus of fixed luminance (3–6). It is also true for time-varying visual stimuli such as drifting gratings. With few exceptions, the statistical properties of these spike trains have been investigated from the point of view of the interevent-interval histogram (3), which provides a measure of the relative frequency of intervals of different durations. The mathematical model most widely used to describe the interevent-interval histogram under all of these stimulus conditions derives from the gamma renewal process (7), though point processes incorporating refractoriness have also been considered (3, 8, 9).

However, there are properties of a sequence of action potentials, such as long-duration correlation or memory, that cannot generally be inferred from measures that reset at short times such as the interevent-interval histogram (6, 10). The ability to uncover features such as these demands the use of measures such as the count-based normalized wavelet variance, the periodogram, or rescaled range analysis (R/S), which can extend over time (or frequency) scales that span many events. RGC and LGN spike trains exhibit variability and correlation properties over a broad range of time scales, and the analysis of these discharges reveals that the spike rates exhibit fractal properties.

Received April 14, 1998. Accepted for publication July 14, 1998. Posted on the Los Alamos National Laboratory server October 18, 1999.

¹ Present address: McLean Hospital and Harvard University, 115 Mill Street, Belmont, MA 02478.

² To whom correspondence should be addressed. Fax: (617) 353-6440. E-mail: teich@bu.edu.

Fractals are objects that possess a form of self-similarity: parts of the whole can be made to fit to the whole by shifting and stretching. The hallmark of fractal behavior is power-law dependence in one or more statistical measures, over a substantial range of the time or frequency scale at which the measurement is conducted (11). Fractal behavior represents a form of memory because the occurrence of an event at a particular time increases the likelihood of another event occurring some time later, with this likelihood decaying in power-law fashion. Fractal signals are also said to be self-similar or self-affine.

This fractal behavior is most readily illustrated by plotting the estimated firing rate of a sequence of action potentials for a range of averaging times. This is illustrated in Fig. 1A for the maintained discharge of a cat RGC. The rate estimates are formed by dividing the number of spikes in successive counting windows of duration T by the counting time T . The rate estimates of the shuffled (randomly reordered) version of the data are presented in Fig. 1B. This surrogate data set maintains the same relative frequency of interevent-interval durations as the original data, but destroys any long-term correlations (and therefore fractal behavior) arising from other sources, such as the relative ordering of the intervals.

Comparing Figs. 1A and 1B, it is apparent that the magnitude of the rate fluctuations decreases more slowly with increasing counting time for the original data than for the shuffled version. Fractal processes exhibit slow power-law convergence: the standard deviation of the rate decreases more slowly than $1/T^{1/2}$ as the averaging time increases. Nonfractal signals, such as the shuffled RGC spike train, on the other hand, exhibit fluctuations that decrease precisely as $1/T^{1/2}$. The data presented in Fig. 1 are typical of all RGC and LGN spike trains.

1. ANALYSIS TECHNIQUES

1.1. Point Processes

The statistical behavior of a neural spike train can be studied by replacing the complex waveforms of each individual electrically recorded action potential (Fig. 2, top) by a single point event corresponding to the time of the peak (or other designator) of the action potential (Fig. 2, middle). In mathematical terms, the neural spike train is then viewed as an unmarked

point process. This simplification greatly reduces the computational complexity of the problem and permits use of the substantial methodology previously developed for stochastic point processes (6, 10, 11).

The occurrence of a neural spike at time t_n is therefore simply represented by an impulse $\delta(t - t_n)$ at that time, so that the sequence of action potentials is represented by

$$s(t) = \sum_n \delta(t - t_n).$$

A realization of a point process is specified by the set of occurrence times of the events or, equivalently, of the times $\{\tau_n\}$ between adjacent events, where $\tau_n = t_{n+1} - t_n$. A single realization of the data is generally all that is available to the observer, so that the identification of the point process, and elucidation of the mechanisms that underlie it, must be gleaned from this one realization.

One way in which the information in an experimental sequence of events can be made more digestible is to reduce the data into a statistic that emphasizes a particular aspect of the data, at the expense of other features. These statistics fall into two broad classes which have their origins, respectively, in the sequence of interevent intervals $\{\tau_n\}$ illustrated at the lower left of Fig. 2 or in the sequence of counts $\{Z_n\}$ shown at the lower right of Fig. 2.

1.1.1. Examples of Point Processes

The homogeneous Poisson point process, which is the simplest of all stochastic point processes, is described by a single parameter, the rate λ . This point process is memoryless: the occurrence of an event at any time t_0 is independent of the presence (or absence) of events at other times $t \neq t_0$. Because of this property, both the intervals $\{\tau_n\}$ and counts $\{Z_n\}$ form sequences of independent, identically distributed (iid) random variables. The homogeneous Poisson point process is therefore completely characterized by the interevent-interval distribution (which is exponential) or the event-number distribution (which is Poisson) together with the iid property. This process serves as a benchmark against which other point processes are measured; it therefore plays the role that the white Gaussian process enjoys in the realm of continuous-time stochastic processes.

A related point process is the nonparalyzable fixed-dead-time-modified Poisson point process, a close cousin of the homogeneous Poisson point process that

differs only by the imposition of a dead-time (refractory) interval after the occurrence of each event, during which other events are prohibited from occurring (9). Another cousin is the gamma- r renewal process which, for integer r , is generated from a homogeneous Poisson point process by permitting every r th event to survive while deleting all intermediate events (6). Both the dead-time-modified Poisson point process and the gamma renewal process require two parameters for their description. All the examples of point process presented above belong to the class of renewal point processes, which are defined in Section 1.2.1.

However, spike trains in the visual system cannot be adequately described by renewal point processes; rather, nonrenewal processes are required (6). Of particular interest are fractal-rate stochastic point processes, in which one or more statistics exhibit power-law behavior in time or frequency (11). One feature of such processes is the relatively slow power-law convergence of the rate standard deviation, as illustrated in Fig. 1A. We have previously shown that a fractal, doubly stochastic point process that imparts multiscale fluctuations to the gamma- r renewal process provides a reasonable description of the RGC and LGN maintained discharges (6).

RATE FUNCTIONS λ_n

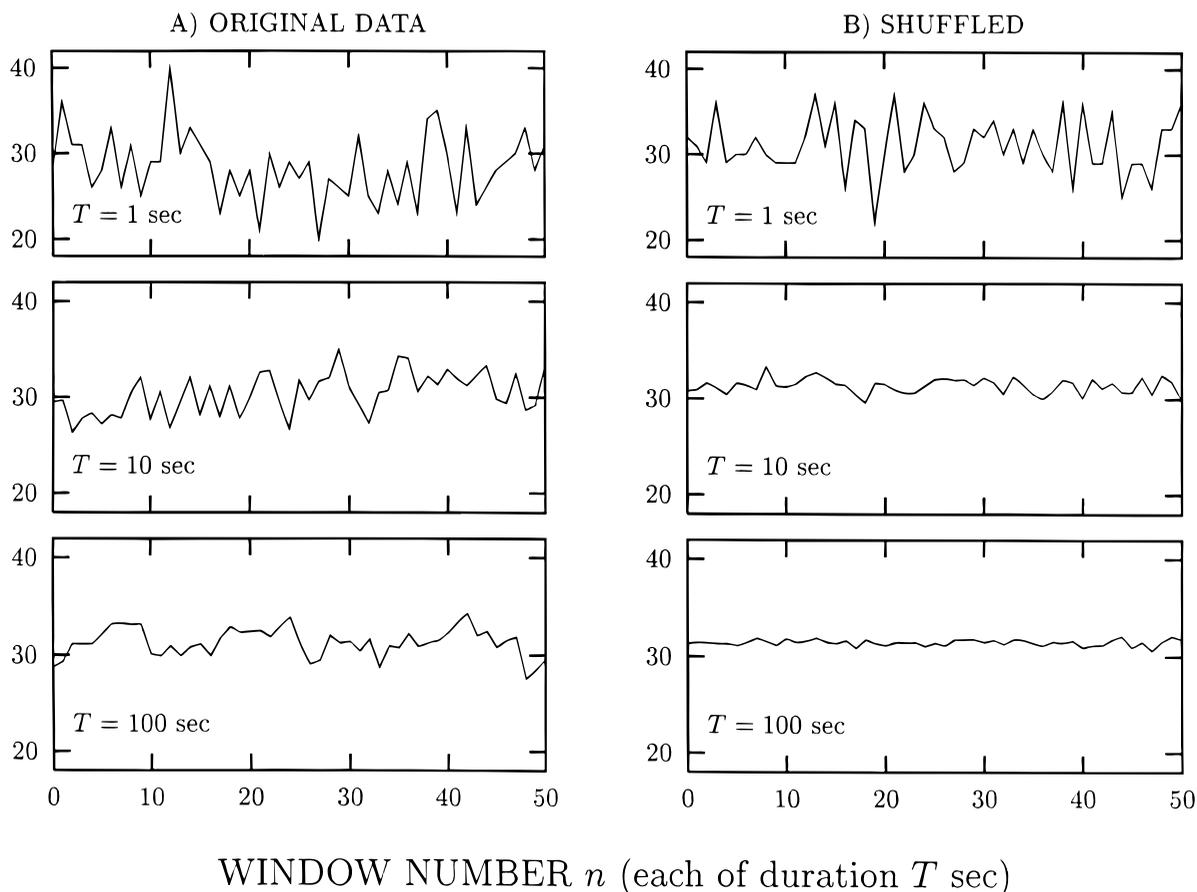


FIG. 1. Rate estimates formed by dividing the number of events in successive counting windows by the counting time T . The stimulus was a uniformly illuminated screen (with no temporal or spatial modulation) of luminance 50 cd/m^2 . (A) Rate estimate for a cat RGC generated using three different counting times ($T = 1, 10,$ and 100 s). The fluctuations in the rate estimates converge relatively slowly as the counting time is increased. This is characteristic of fractal-rate processes. The convergence properties are quantified by measures such as the normalized wavelet variance and periodogram. (B) Rate estimates from the same recording after the intervals are randomly reordered (shuffled). This maintains the same relative frequency of interval sizes but destroys the original relative ordering of the intervals and, therefore, any correlations or dependencies among them. For such nonfractal signals, the rate estimate converges more quickly as the counting time T is increased. The data presented here are typical of the 50 data sets examined.

1.2. Interevent-Interval Measures of a Point Process

Two statistical measures are often used to characterize the discrete-time stochastic process $\{\tau_n\}$ illustrated in the lower left corner of Fig. 2. These are the interevent-interval histogram (IIH) and rescaled range analysis (*R/S*).

1.2.1. Interevent-Interval Histogram

The interevent-interval histogram (often referred to as the interspike-interval histogram or ISIH in the physiology literature) displays the relative frequency of occurrence $p_r(\tau)$ of an interval of size τ ; it is an estimate of the probability density function of interevent-interval magnitude (see Fig. 2, lower left). It is, perhaps, the most commonly used of all statistical measures of point processes in the life sciences.

The interevent-interval histogram provides information about the underlying process over time scales that are of the order of the interevent intervals. Its construction involves the loss of interval ordering and, therefore, dependencies among intervals; a reordering of the sequence does not alter the interevent-interval histogram since the order plays no role in the relative frequency of occurrence.

Some point processes exhibit no dependencies among their interevent intervals at the outset, in which case the sequence of interevent intervals forms a sequence of iid random variables and the point process is completely specified by its interevent-interval histogram. Such a process is called a renewal process, a definition motivated by the replacement of failed parts (such as light bulbs), each replacement

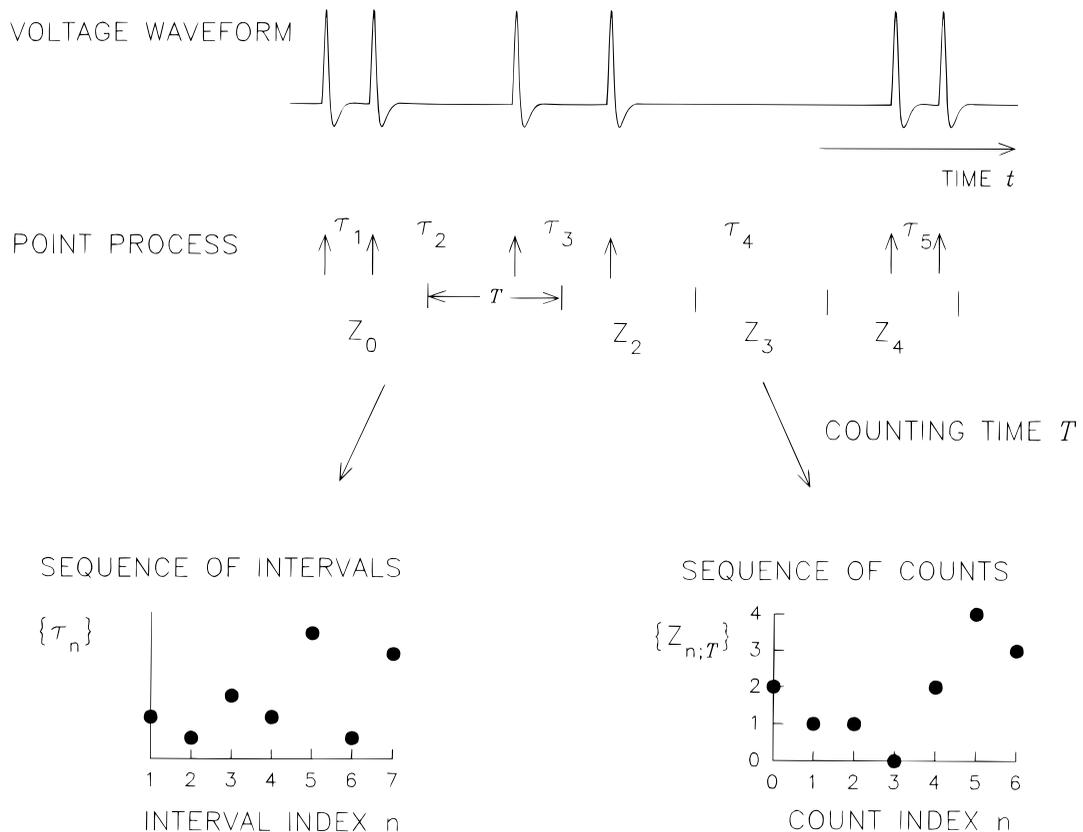


FIG. 2. A sequence of action potentials (top) is reduced to a set of events (represented by arrows, middle) that form a point process. A sequence of interevent intervals $\{\tau_n\}$ is formed from the times between successive events, resulting in a discrete-time, positive, real-valued stochastic process (lower left). All information contained in the original point process remains in this representation, but the discrete-time axis of the sequence of interevent intervals is distorted relative to the real-time axis of the point process. The sequence of counts $\{Z_n\}$, a discrete-time, nonnegative, integer-valued stochastic process, is formed from the point process by recording the numbers of events in successive counting windows of duration T (lower right). This process of mapping the point process to the sequence $\{Z_n\}$ results in a loss of information, but the amount lost can be made arbitrarily small by reducing T . An advantage of this representation is that no distortion of the time axis occurs.

of which forms a renewal of the point process. The homogeneous Poisson point process, dead-time-modified Poisson point process, and gamma renewal process are all renewal processes, but experimental RGC and LGN spike trains are not.

1.2.2. Rescaled Range (R/S) Analysis

Rescaled range (R/S) analysis provides information about correlations among blocks of interevent intervals. For a block of k interevent intervals, the difference between each interval and the mean interevent interval 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 the degree of correlation in the process is obtained by fitting $R(k)$ to the function k^H , where H is the so-called Hurst exponent (12). For $H > 0.5$ positive correlation exists among the intervals, whereas $H < 0.5$ indicates the presence of negative correlation; $H = 0.5$ obtains for intervals with no correlation. Renewal processes yield $H = 0.5$. For negatively correlated intervals, an interval that is larger than the mean tends, on average, to be preceded or followed by one smaller than the mean.

This widely used measure is generally assumed to be well suited to processes that exhibit long-term correlation or have a large variance (12–15), but it is robust since it exhibits large systematic errors and highly variable estimates of the Hurst coefficient for some fractal sequences (16, 17). Nevertheless, it provides a useful indication of correlation in a point process arising from the ordering of the interevent intervals alone.

1.3. Event-Number Measures of a Point Process

It is advantageous to study some characteristics of a point process in terms of the sequence of event numbers (counts) $\{Z_n\}$ rather than via the sequence of intervals $\{\tau_n\}$.

Figure 2 illustrates how the sequence is obtained. The time axis is divided into equally spaced, contiguous time windows (center), each of duration T s, and the (integer) number of events in the n th window is counted and denoted Z_n . This sequence $\{Z_n\}$ forms a random counting process of nonnegative integers (lower right). Closely related to the sequence of counts is the sequence of rates (events/second), λ_n ,

which is obtained by dividing each count Z_n by the counting time T . This is the measure used in Fig. 1.

We describe several statistical measures useful for characterizing the counting process $\{Z_n\}$: the normalized variance (Fano factor), the normalized wavelet variance (Allan factor), and the event-number-based power spectral density estimate (periodogram).

1.3.1. Normalized Variance

The normalized variance, or (Fano factor) is defined as the event-number variance divided by the event-number mean, which is a function of the counting time T :

$$F(T) \equiv \frac{\text{Var} [Z_n(T)]}{\text{E} [Z_n(T)]}.$$

This quantity provides an abbreviated way of describing correlation in a sequence of events. It indicates the degree of event clustering or anticlustering in a point process relative to the benchmark homogeneous Poisson point process, for which $F(T) = 1$ for all T .

The Fano factor must approach unity at sufficiently small values of the counting time T for any regular point process (6, 11). In general, a Fano factor less than unity indicates that a point process is more orderly than the homogeneous Poisson point process at the particular time scale T , whereas an excess over unity indicates increased clustering at the given time scale. This measure is sometimes called the index of dispersion; it was first used by Fano in 1947 (18) for characterizing the statistical fluctuations of the number of ions generated by individual fast charged particles. For a fractal-rate stochastic point process the Fano factor assumes the power-law form T^{α_F} ($0 < \alpha_F < 1$) for large T . The parameter α_F is defined as an estimate of the fractal exponent (or scaling exponent) α of the point-process rate.

Though the Fano factor can detect the presence of self-similarity even when it cannot be discerned in a visual representation of a sequence of events, mathematical constraints prevent it from increasing with counting time faster than $\sim T^1$ (19). It therefore proves to be unsuitable as a measure for fractal exponents $\alpha > 1$; it also suffers from bias for finite-length data sets (20). For these reasons we employ other count-based measures.

1.3.2. Normalized Wavelet Variance

The reliable estimation of a fractal exponent that may assume a value greater than unity requires the

use of a measure whose increase is not constrained as it is for the Fano factor, and which remains free of bias. In this section we present a measure we first defined in 1996 (19), and called the Allan factor, but a more descriptive name for it is the normalized wavelet variance. The Allan factor is the ratio of the event-number Allan variance to twice the mean:

$$A(T) \equiv \frac{E\{[Z_n(T) - Z_{n+1}(T)]^2\}}{2E[Z_n(T)]}.$$

The Allan variance was first introduced in connection with the stability of atomic-based clocks (21). It is defined in terms of the variability of differences of successive counts; as such it is a measure based on the Haar wavelet. Because this measure functions as a derivative, it has the salutary effect of mitigating against linear nonstationarities. More complex wavelet Allan factors can be constructed to eliminate polynomial trends (22, 23).

Like the Fano factor, the Allan factor is also a useful measure of the degree of event clustering (or anticlustering) in a point process relative to the benchmark homogeneous Poisson point process, for which $A(T) = 1$ for all T . In fact, for any point process, the Allan factor is simply related to the Fano factor by

$$A(T) = 2F(T) - F(2T)$$

so that, in general, both quantities vary with the counting time T . In particular, for a regular point process the Allan factor also approaches unity as T approaches zero. For a fractal-rate stochastic point process and sufficiently large T , the Allan factor exhibits a power law dependence that varies with the counting time T as $A(T) \sim T^{\alpha_A}$ ($0 < \alpha_A < 3$); it can rise as fast as $\sim T^3$ and can therefore be used to estimate fractal exponents over the expanded range $0 < \alpha_A < 3$.

1.3.3. Periodogram

Fourier-transform methods provide another avenue for quantifying correlation in a point process. The periodogram is an estimate of the power spectral density of a point process, revealing how the power is concentrated across frequency. The count-based periodogram is obtained by dividing a data set into contiguous segments of equal length \mathcal{T} . Within each segment, a discrete-index sequence $\{W_m\}$ is formed by further dividing \mathcal{T} into M equal bins, and then

counting the number of events within each bin. A periodogram is then formed for each of the segments according to

$$S_W(f) = \frac{1}{M} |\tilde{W}(f)|^2,$$

where $\tilde{W}(f)$ is the discrete Fourier transform of the sequence $\{W_m\}$ and M is the length of the transform. All of the segment periodograms are averaged together to form the final averaged periodogram $S(f)$, which estimates the power spectral density in the frequency range from $1/\mathcal{T}$ to $M/2\mathcal{T}$ Hz. The periodogram $S(f)$ can also be smoothed by using a suitable windowing function (24).

The count-based periodogram, as opposed to the interval-based periodogram (formed by Fourier transforming the interevent intervals directly), provides direct undistorted information about the time correlation of the underlying point process because the count index increases by unity every \mathcal{T}/M s, in proportion to the real time of the point process. In the special case when the bin width \mathcal{T}/M is short in comparison with most interevent intervals τ , the count-based periodogram essentially reduces to the periodogram of the point process itself, since the bins reproduce the original point process to a good approximation.

For a fractal-rate stochastic point process, the periodogram exhibits a power-law dependence that varies with the frequency f as $S(f) \sim f^{-\alpha_S}$; unlike the Fano and Allan factor exponents, however, α_S can assume any value. Thus in theory the periodogram can be used to estimate any value of fractal exponent, although in practice fractal exponents α rarely exceed a value of 3. Compared with estimates based on the Allan factor, periodogram-based estimates of the fractal exponent α_S suffer from increased bias and variance (11). Other methods also exist for investigating the spectrum of a point process, some of which highlight fluctuations about the mean rate (25).

1.3.4. Relationship among Fractal Exponents

For a fractal-rate stochastic point process with $0 < \alpha < 1$, the theoretical Fano factor, Allan factor, and periodogram curves all follow power-law forms with respect to their arguments, and in fact we obtain $\alpha_F = \alpha_A = \alpha_S = \alpha$. For $1 \leq \alpha < 3$, the theoretical Fano factor curves saturate, but the relation $\alpha_A = \alpha_S = \alpha$ still obtains. The fractal exponent α is ambiguously related to the Hurst exponent H , since some

authors have used the quantity H to index fractal gaussian noise whereas others have used the same value of H to index the integral of fractal gaussian noise (which is fractional brownian motion). The relationship between the quantities is $\alpha = 2H - 1$ for fractal gaussian noise and $\alpha = 2H + 1$ for fractal brownian motion. In the context of this paper, the former relationship holds, and we can define another estimate of the fractal exponent, $\alpha_R = 2H_R - 1$, where H_R is the estimate of the Hurst exponent H obtained from the data at hand. In general, α_R depends on the theoretical value of α , as well as on the probability distribution of the interevent intervals. The distributions of the data analyzed in this paper, however, prove simple enough so that the approximate theoretical relation $\alpha_R = \alpha$ will hold in the case of large amounts of data.

1.4. Correlation Measures for Pairs of Point Processes

Second-order methods prove useful in revealing correlations between sequences of events, which indicate how information is shared between pairs of spike trains. Such methods may not detect subtle forms of interdependence to which information-theoretic approaches are sensitive (26), but the latter methods suffer from limitations due to the finite size of the data sets used. We consider two second-order methods here: the normalized wavelet cross-correlation function (NWCCF) and the cross periodogram.

1.4.1. Normalized Wavelet Cross-Correlation Function

We define the normalized wavelet cross-correlation function $A_2(T)$ as a generalization of the normalized wavelet variance (see Sec. 1.3.2). It is a Haar wavelet-based version of the correlation function and is therefore insensitive to linear trends. It can be readily generalized by using other wavelets and can thereby be rendered insensitive to polynomial trends. To compute the normalized wavelet cross-correlation function at a particular counting time T , the two spike trains first are divided into contiguous counting windows T . The number of spikes $Z_{1,n}$ falling within the n th window is registered for all indices n corresponding to windows lying entirely within the first spike-train data set, much as in the procedure to estimate the Allan factor. This process is repeated for the second spike train, yielding $Z_{2,n}$. The difference between the count numbers in a given window in the first spike train ($Z_{1,n}$) and the one after it

($Z_{1,n+1}$) is then computed for all n , with a similar procedure followed for the second spike train. Paralleling the definition of the normalized wavelet variance, the normalized wavelet cross-correlation function is defined as

$$A_2(T) \equiv \frac{E\{[Z_{1,n}(T) - Z_{1,n+1}(T)] [Z_{2,n}(T) - Z_{2,n+1}(T)]\}}{2\{E[Z_{1,n}(T)] E [Z_{2,n}(T)]\}^{1/2}}.$$

The normalization has two salutary properties: (1) it is symmetric in the two spike trains, and (2) when the same homogeneous Poisson point process is used for both spike trains the normalized wavelet cross-correlation function assumes a value of unity for all counting times T , again in analogy with the Allan factor. To determine the significance of a particular value for the normalized wavelet cross-correlation function, we make use of two surrogate data sets: a shuffled version of the original data sets (same interevent intervals but in a random order), and homogeneous Poisson point processes with the same mean rate. Comparison between the value of the normalized wavelet cross-correlation function obtained from the data at a particular counting time T , on the one hand, and from the surrogates at that time T , on the other hand, indicates the significance of that particular value.

1.4.2. Cross Periodogram

The cross periodogram (27) is a generalization of the periodogram for individual spike trains (see Section 1.3.3), in much the same manner as the normalized wavelet cross-correlation function derives from the normalized wavelet variance. Two data sets are divided into contiguous segments of equal length \mathcal{T} , with discrete-index sequences $\{W_{1,m}\}$ and $\{W_{2,m}\}$ formed by further dividing each segment of both data sets into M equal bins, and then counting the number of events within each bin. With the M -point discrete Fourier transform of the sequence $\{W_{1,m}\}$ denoted by $\tilde{W}_1(f)$ (and similarly for the second sequence), we define the segment cross periodograms as

$$\begin{aligned} S_{2,w}(f) &\equiv \frac{1}{2M} [\tilde{W}_1^*(f) \tilde{W}_2(f) + \tilde{W}_1(f) \tilde{W}_2^*(f)] \\ &= \frac{1}{M} \text{Re} [\tilde{W}_1^*(f) \tilde{W}_2(f)], \end{aligned}$$

where $*$ represents complex conjugation and $\text{Re}(\cdot)$ represents the real part of the argument. As with the ordinary periodogram, all of the segment cross

periodograms are averaged together to form the final averaged cross periodogram, $S_2(f)$, and the result can be smoothed. This form is chosen to be symmetric in the two spike trains and to yield a real (although possibly negative) result. In the case of independent spike trains, the expected value of the cross periodogram is zero. We again employ the same two surrogate data sets (shuffled and Poisson) to provide significance information about cross periodogram values for actual data sets.

The cross periodogram and normalized wavelet cross-correlation function will have different immunity to nonstationarities and will exhibit different bias-variance trade-offs, much as their single-dimensional counterparts do (11).

2. RESULTS FOR RGC AND LGN ACTION-POTENTIAL SEQUENCES

We have carried out a series of experiments to determine the statistical characteristics of the dark, maintained, and driven neural discharge in cat RGCs and LGN cells. Using the analysis techniques presented in Section 1, we compare and contrast the neural activity for these three different stimulus modalities, devoting particular attention to their fractal features. The results we present all derive from on-center X-type cells.

2.1. Experimental Methods

The experimental methods are similar to those used by Kaplan and Shapley (28) and Teich *et al.* (6). Experiments were carried out on adult cats. Anesthesia was induced by intramuscular injection of xylazine (Rompun 2 mg/kg), followed 10 min later by intramuscular injection of ketamine HCl (Ketaset 10 mg/kg). Anesthesia was maintained during surgery with intravenous injections of thiamylal (Surital 2.5%) or thiopental (Pentothal 2.5%). During recording, anesthesia was maintained with Pentothal (2.5%, 2–6 (mg/kg)/h). The local anesthetic Novocain was administered, as required, during the surgical procedures. Penicillin (750,000 units intramuscular) was also administered to prevent infection, as was dexamethasone (Decadron, 6 mg intravenous) to forestall cerebral edema. Muscular paralysis was induced and maintained with gallium triethiodide (Flaxedil, 5–15 (mg/kg)/h) or vecuronium bromide (Norcuron, 0.25

(mg/kg)/h). Infusions of Ringer's saline with 5% dextrose at 3–4 (ml/kg)/h were also administered.

The two femoral veins and a femoral artery were cannulated for intravenous drug infusions. Heart rate and blood pressure, along with expired CO₂, were continuously monitored and maintained in physiological ranges. For male cats, the bladder was also cannulated to monitor fluid outflow. Core body temperature was maintained at 37.5°C throughout the experiment by wrapping the animal's torso in a dc heating pad controlled by feedback from a subscapular temperature probe. The cat's head was fixed in a stereotaxic apparatus. The trachea was cannulated to allow for artificial respiration. To minimize respiratory artifacts, the animal's body was suspended from a vertebral clamp and a pneumothorax was performed when needed.

Eyedrops of 10% phenylephrine hydrochloride (Neo-synephrine) and 1% atropine were applied to dilate the pupils and retract the nictitating membranes. Gas-permeable hard contact lenses protected the corneas from drying. Artificial pupils of 3-mm diameter were placed in front of the contact lenses to maintain fixed retinal illumination. The optical quality of the animal's eyes was regularly examined by ophthalmoscopy. The optic discs were mapped onto a tangent screen, by back-projection, for use as a positional reference. The animal viewed a CRT screen (Tektronix 608, 270 frames/s; or CONRAC, 135 frames/s) that, depending on the stimulus condition, was either dark, uniformly illuminated with a fixed luminance level, or displayed a moving grating.

A craniotomy was performed over the LGN (center located 6.5 mm anterior to the earbars and 9 mm lateral to the midline of the skull), and the dura mater was resected. A tungsten-in-glass microelectrode (5–10 μm in tip length) (29) was lowered until spikes from a single LGN neuron were isolated. The microelectrode simultaneously recorded RGC activity, in the form of S potentials, and LGN spikes, with a timing accuracy of 0.1 ms. The output was amplified and monitored using conventional techniques. A cell was classified as Y-type if it exhibited strong frequency doubling in response to contrast-reversing high-spatial-frequency gratings, and X-type otherwise (30, 31).

The experimental protocol was approved by the Animal Care and Use Committee of Rockefeller University, and was in accord with the National Institutes of Health guidelines for the use of higher mammals in neuroscience experiments.

2.2. RGC and LGN Dark Discharge

Results for simultaneously recorded RGC and target LGN spike trains of 4000-s duration are presented in Fig. 3, when the retina is thoroughly adapted to the dark (this is referred to as the “dark discharge”). The normalized rate functions (A) for both the RGC (solid curve) and LGN (dashed curve) recordings exhibit large fluctuations over the course of the recording; each window corresponds to a counting time of $T = 100$ s. Such large, slow fluctuations often indicate fractal rates (6, 11). The two recordings

bear a substantial resemblance to each other, suggesting that the fractal components of the rate fluctuations either have a common origin or pass from one of the cells to the other.

The normalized interevent-interval histogram (B) of the RGC data follows a straight-line trend on a semilogarithmic plot, indicating that the interevent-interval probability density function is close to an exponential form. The LGN data, however, yield a nonmonotonic (bimodal) interevent-interval histogram. This distribution favors longer and shorter

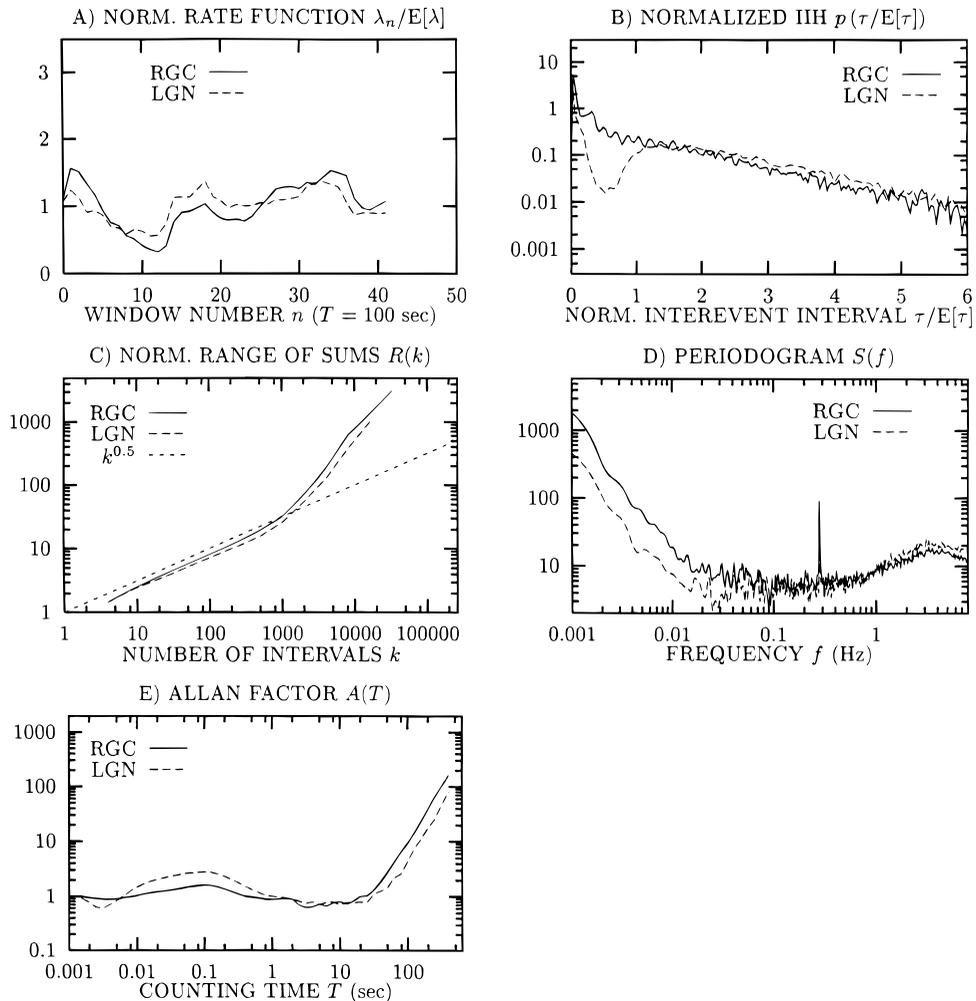


FIG. 3. Statistical measures of the dark discharge from a cat on-center X-type retinal ganglion cell (RGC) and its associated lateral geniculate nucleus (LGN) cell, for data of duration $L = 4000$ s. RGC results appear as solid curves, whereas LGN results are dashed. (A) Normalized rate function constructed by counting the number of neural spikes occurring in adjacent 100-s counting windows, and then dividing by 100 s and by the average rate. (B) Normalized interevent-interval histogram (IIH) versus normalized interevent interval constructed by dividing the interevent intervals for each spike train by the mean, and then obtaining the histogram. (C) Normalized range of sums $R(k)$ versus number of interevent intervals k (see Section 1.2.2). (D) Periodogram $S(f)$ versus frequency f (see Section 1.3.3) (E) Normalized wavelet variance (Allan factor) $A(T)$ versus counting time T (see Section 1.3.2).

intervals at the expense of those near half the mean interval, reflecting clustering in the event occurrences over the short term. Various kinds of unusual clustering behavior have been previously observed in LGN discharges (2, 32).

R/S plots (C) for both the RGC and LGN recordings follow the $k^{0.5}$ line for sums less than 1000 intervals, but rise sharply thereafter in a roughly power law fashion as $k^{H_R} = k^{(\alpha_R+1)/2}$, suggesting that the neural firing pattern exhibits fractal activity for times greater than about 1000 intervals (about 120 s for these two recordings).

Both smoothed periodograms (D) decay with frequency as $f^{-\alpha_S}$ for small frequencies, and the Allan factors (E) increases with times as T^{α_A} for large counting times, confirming the fractal behavior. The 0.3-Hz component evident in the periodograms of both recordings is an artifact of the artificial respiration; it does not affect the fractal analysis. As shown in Table 1, the fractal exponents calculated from the various measures bear rough similarity to each other,

TABLE 1

Neural-Discharge Statistics for Cat Retinal Ganglion Cells (RGCs) and Their Associated Lateral Geniculate Nucleus (LGN) Cells^a

Stimulus	Cell	Moment		Fractal exponent		
		Mean (ms)	CV	α_R	α_S	α_A
Dark	RGC	112	1.54	1.71	1.89	1.96
	LGN	152	1.62	1.66	1.75	1.85
Maintained	RGC	32	0.52	0.53	0.58	0.99
	LGN	284	1.63	0.89	2.01	1.41
Driven	RGC	27	1.21	0.79	0.54	0.74
	LGN	77	1.15	1.35	2.10	1.76

^a Under three stimulus conditions: dark discharge in the absence of stimulation (data duration $L = 4000$ s); maintained discharge in response to a uniform luminance of 50 cd/m² (data duration $L = 7000$ s); and driven discharge in response to a drifting grating (4.2 Hz frequency, 40% contrast, and 50 cd/m² mean luminance; data duration $L = 7000$ s). All cells are on-center X type. The maintained and driven data sets were recorded from the same RGC/LGN cell pair, whereas the dark discharge derived from a different cell pair. Statistics, from left to right, are mean interevent interval, interevent-interval coefficient of variation (CV = SD divided by mean), and fractal exponents estimated by least-squares fits on doubly logarithmic plots of (1) the rescaled range (R/S) statistic for $k > 1000$, which yields an estimate of the Hurst exponent, H_R , and of α_R , in turn, through the relation $\alpha_R = 2H_R - 1$; (2) the count-based periodogram for frequencies between 0.001 and 0.01 Hz, which yields α_S ; and (3) the Allan factor for counting times between $L/100$ and $L/10$, where L is the duration of the recording, which yields α_A .

as expected (11); further, the onset times also agree reasonably well, being in the neighborhood of 100 s. The coherence among these statistics leaves little doubt that these RGC and LGN recordings exhibit fractal features with estimated fractal exponents of 1.9 ± 0.1 and 1.8 ± 0.1 (mean \pm SD of the three estimated exponents), respectively. Moreover, the close numerical agreement of the RGC and LGN estimated fractal exponents suggests a close connection between the fractal activity in the two spike trains under dark conditions (6). Curves such as those presented in Fig. 3 are readily simulated by using a fractal-rate stochastic point process, as described in (6).

With the exception of the interevent-interval distribution, it is apparent from Fig. 3 that the statistical properties of the dark discharges generated by the RGC and its target LGN cell prove to be remarkably similar.

2.3. RGC and LGN Maintained Discharge

Figure 4 presents analogous statistical results for simultaneously recorded maintained-discharge RGC and target-LGN spike trains of 7000-s duration when the stimulus presented by the CRT screen was a 50 cd/m² uniform luminance. The cell pair from which these recordings were obtained is different from the pair whose statistics are shown in Fig. 3. As is evident from Table 1, the imposition of a stimulus increases the RGC firing rate, though not that of the LGN. In contrast to the results for the dark discharge, the RGC and LGN action-potential sequences differ from each other in significant ways under maintained-discharge conditions. We previously investigated some of these statistical measures and their roles in revealing fractal features, for the maintained discharge (6).

The rate fluctuations (A) of the RGC and the LGN no longer resemble each other. At these counting times, the normalized RGC rate fluctuations are suppressed, whereas those of the LGN are enhanced, relative to the dark discharge shown in Fig. 3. Significant long-duration fluctuations are apparently imparted to the RGC S-potential sequence at the LGN, through the process of selective clustered passage (26). Spike clustering is also imparted at the LGN over short time scales; the RGC maintained discharge exhibits a coefficient of variation (CV) much less than unity, whereas that of the LGN significantly exceeds unity (see Table 1).

The normalized interevent-interval histogram (B)

of the RGC data resembles that of a dead-time-modified Poisson point process (fit not shown), consistent with the presence of relative refractoriness which becomes more important at higher rates (9). Dead-time effects in the LGN are secondary to the clustering that it imparts to the RGC S-potentials, in part because of its lower rate.

The R/S (C), periodogram (D), and Allan factor (E) plots yield results that are consistent with, but different from, those revealed by the dark discharge shown in Fig. 3. Although both the RGC and LGN recordings exhibit evidence of fractal behavior, the two spike trains now behave quite differently in the presence of a steady-luminance stimulus. For the RGC recording, all three measures are consistent

with a fractal onset time of about 1 s and a relatively small fractal exponent (0.7 ± 0.3). For the LGN, the fractal behavior again appears in all three statistics, but begins at a larger onset time (roughly 20 s) and exhibits a larger fractal exponent (1.4 ± 0.6). Again, all measures presented in Fig. 4 are well described by a pair of fractal-rate stochastic point processes (6).

2.4. RGC and LGN Driven Discharge

Figure 5 presents these same statistical measures for simultaneously recorded 7000-s-duration RGC and LGN spike trains in response to a sinusoidal stimulus (drifting grating) at 4.2-Hz frequency, 40% contrast, and 50 cd/m^2 mean luminance. The RGC/

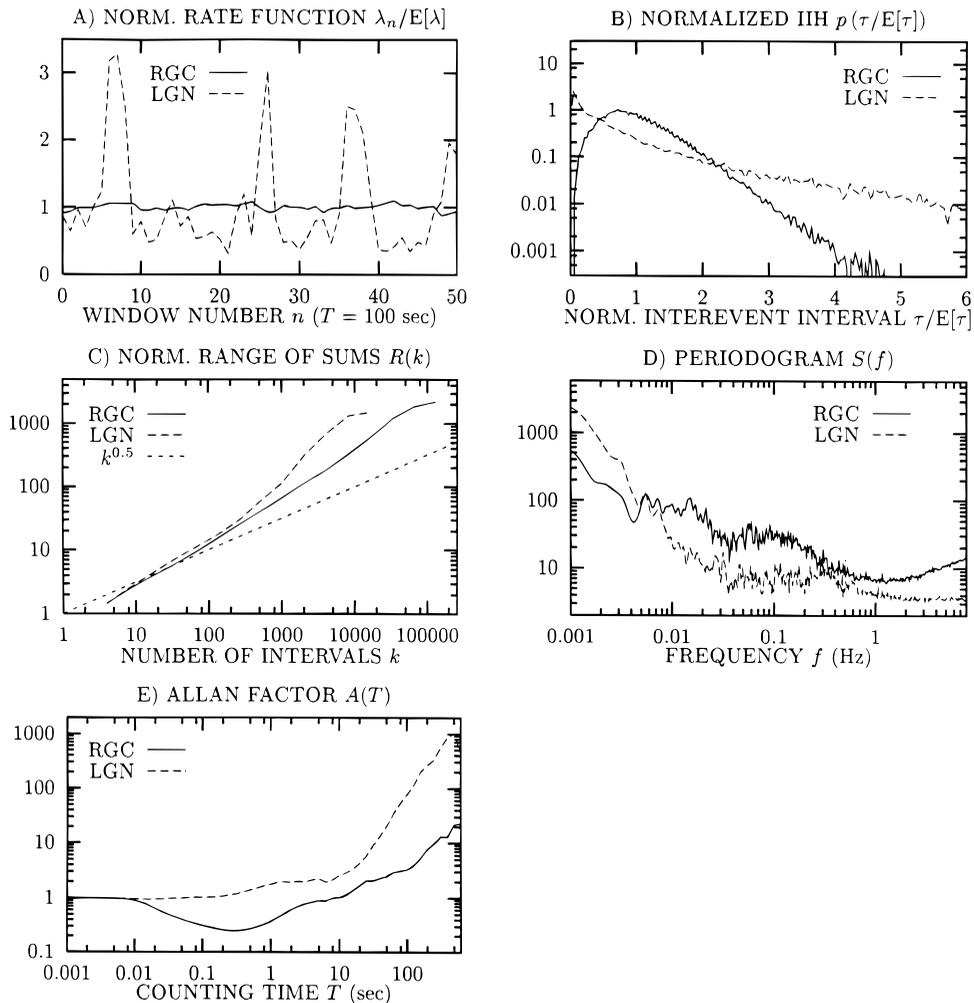


FIG. 4. Statistical measures of the maintained discharge from a cat on-center X-type RGC and its associated LGN cell, at a steady luminance of 50 cd/m^2 , for data of duration $L = 7000$ s. This cell pair is different from that illustrated in Fig. 3. The results for the RGC discharge appear as solid curves, whereas those for the LGN are presented as dashed curves. (A)–(E) as in Fig. 3.

LGN cell pair from which these recording were obtained is the same as the pair illustrated in Fig. 4. The results for this stimulus resemble those for the maintained discharge, but with added sinusoidal components associated with the restricted phases of the stimulus during which action potentials occur. Using terminology from auditory neurophysiology, these spikes are said to be “phase locked” to the periodicity provided by the drifting-grating stimulus. The firing rate is greater than that observed with a steady-luminance stimulus, particularly for the LGN (see Table 1).

Again, the RGC and LGN spike trains exhibit different behavior. The rate fluctuations (A) of the LGN still exceed those of the RGC, but not to as great an extent as in Fig. 4. Both action-potential sequences exhibit normalized interevent-interval histograms (B) with multiple maxima, but the form of the histogram is now dominated by the modulation imposed by the oscillatory stimulus.

Over long times and small frequencies, the R/S (C), periodogram (D), and Allan factor (E) plots again yield results in rough agreement with each other and also with the results presented in Fig. 4. The most

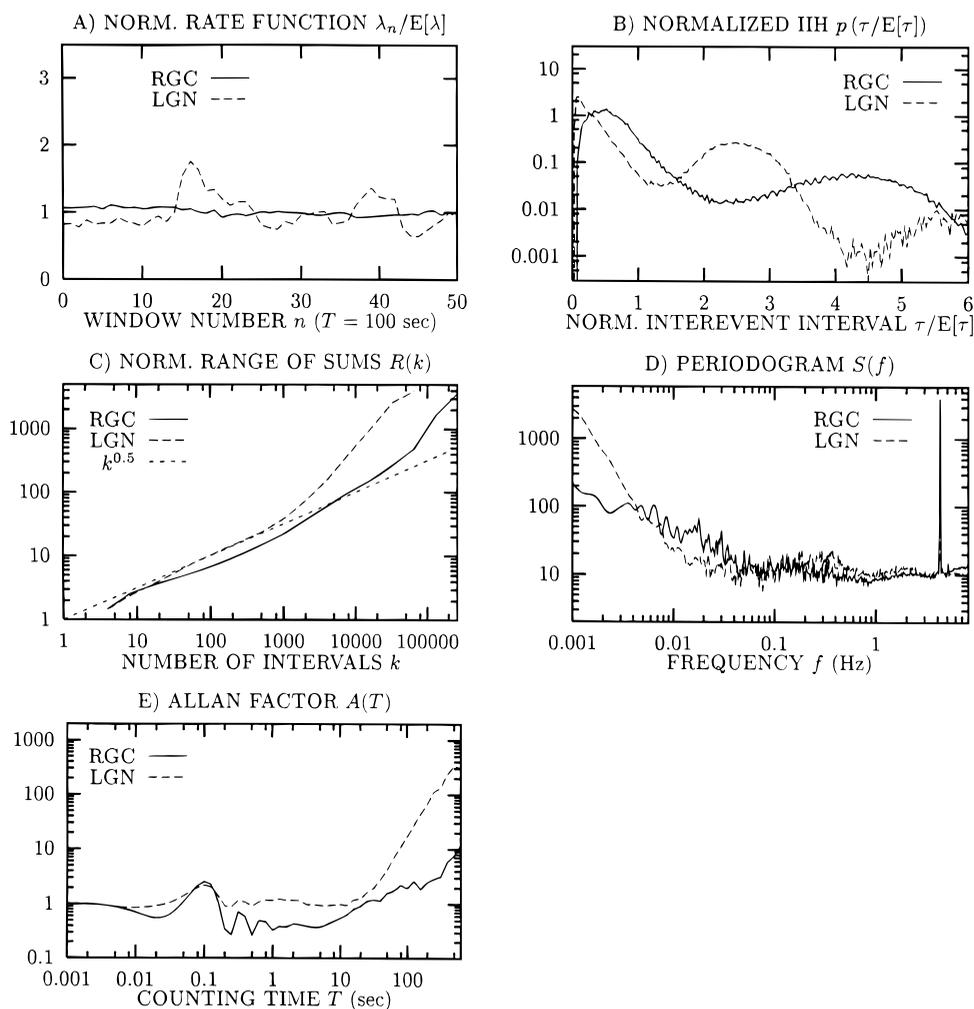


FIG. 5. Statistical measures of the driven discharge from a cat on-center X-type RGC and its associated LGN cell, for a drifting-grating stimulus with mean luminance 50 cd/m^2 , 4.2-Hz frequency, and 40% contrast, for data of duration $L = 7000$ s. This cell pair is the same as that illustrated in Fig. 4. The results for the RGC discharge appear as solid curves, whereas those for the LGN are presented as dashed curves. (A)–(E) as in Figs. 3 and 4.

obvious differences arise from the phase locking induced by the sinusoidal stimulus, which appears directly in the periodogram as a large spike at 4.2 Hz, and in the Allan factor as local minima near multiples of $(4.2 \text{ Hz})^{-1} = 0.24 \text{ s}$.

The RGC results prove consistent with a fractal onset time of about 3 s and a relatively small fractal exponent (0.7 ± 0.1), whereas for the LGN the onset time is about 20 s and the fractal exponent is 1.7 ± 0.4 . For both spike trains fractal behavior persists in the presence of the oscillatory stimulus, though its magnitude is slightly attenuated.

2.5. Correlation in the Discharges of Pairs of RGCs and LGN Cells

We previously examined information exchange among pairs of RGC and LGN spike trains using information-theoretic measures (26). While these approaches are very general, finite data length renders them incapable of revealing relationships between spike trains over time scales longer than about 1 s. We now proceed to investigate various RGC and LGN spike-train pairs in terms of the correlation measures for pairs of point processes developed in Sec. 1.4.

Pairs of RGC discharges are only weakly correlated over long counting times. This is readily illustrated in terms of normalized rate functions such as those presented in Fig. 6A, in which the rate functions of two RGCs are computed over a counting time $T = 100 \text{ s}$. Calculation of the correlation coefficient ($\rho = +0.27$) shows that the fluctuations are only mildly correlated.

Unexpectedly, however, significant correlation turns out to be present in pairs of LGN discharges over long counting times. This is evident in Fig. 6B, where the correlation coefficient $\rho = +0.98$ ($p < 10^{-16}$) for the rates of two LGN discharges computed over the same counting time $T = 100 \text{ s}$.

For shorter counting times, there is little cross-correlation for pairs of RGC or of LGN spike trains (not shown). However, strong correlations are present in the spike rates of an RGC and its target LGN cell as long as the rate is computed over times shorter than 15 s for this particular cell pair.

The cross-correlation can be quantified at all time and frequency scales by the normalized wavelet cross-correlation function (see Section 1.4.1) and the cross periodogram (see Section 1.4.2), respectively.

Figure 6C shows the normalized wavelet cross-correlation function, as a function of the duration of the counting window, between an RGC/LGN spike-train pair recorded under maintained-discharge conditions, as well as for two surrogate data sets (shuffled and Poisson). For this spike-train pair, it is evident that significant correlation exists over time scales less than 15 s. The constant magnitude of the normalized wavelet cross-correlation function for $T < 15 \text{ s}$ is likely associated with the selective transmission properties of the LGN (26). Figure 6D presents the normalized wavelet cross-correlation function for the same RGC/LGN spike-train pair shown in Fig. 6C (solid curve), together with that between two RGC action-potential sequences (long-dashed curve), and between their two associated LGN spike trains (short-dashed curve). Also shown is a dotted line representing the aggregate behavior of the normalized wavelet cross-correlation function absolute magnitude for all surrogate data sets, which resemble each other.

While the two RGC spike trains exhibit a normalized wavelet cross-correlation function value that remains below 7, the two LGN action-potential sequences yield a curve that steadily grows with increasing counting window T , attaining a value in excess of 1000. Indeed, a logarithmic scale was chosen for the ordinate to facilitate the display of this wide range of values. It is of interest to note that the LGN/LGN curve begins its steep ascent just as the RGC/LGN curve abruptly descends. Further, the normalized wavelet cross-correlation function between the two LGN recordings closely follows a power law form, indicating that the two LGN action-potential rates are cofractal. One possible origin of this phenomenon is a fractal form of correlated modulation of the random-transmission processes in the LGN that results in the two LGN spike trains. Some evidence exists that global modulation of the LGN might originate in the parabrachial nucleus of the brain stem; the results presented here are consistent with such a conclusion.

Analogous results for the cross-periodograms, which are shown in Figs. 6E and 6F, provide results that corroborate, but are not as definitive as, those obtained with the normalized wavelet cross-correlation function.

The behavior of the normalized wavelet cross-correlation functions for pairs of driven spike trains, shown in Fig. 7, closely follow those for pairs of maintained discharges, shown in Fig. 6, except for the

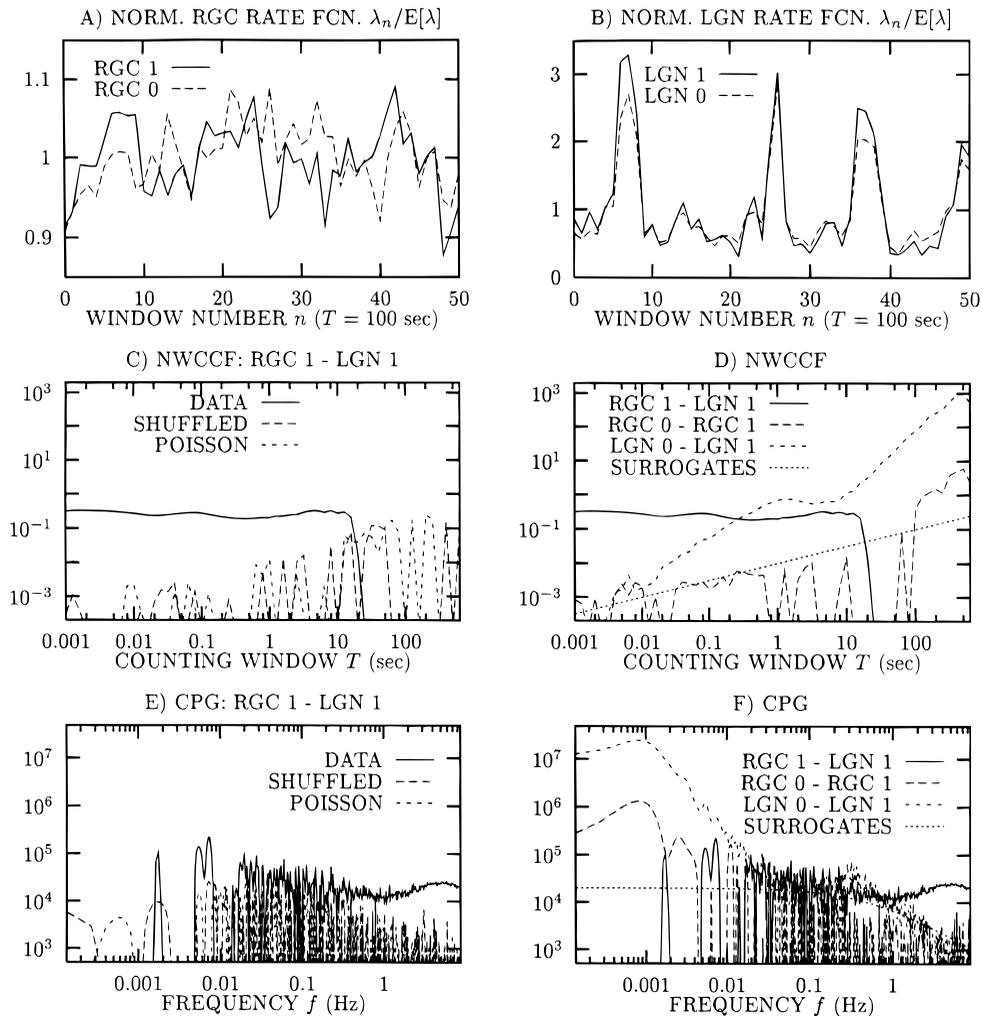


FIG. 6. Statistical measures of the maintained discharge from pairs of cat on-center X-type RGCs and their associated LGN cells, stimulated by a uniform luminance of 50 cd/m^2 , for data of duration $L = 7000 \text{ s}$. RGC and LGN spike trains denoted “1” are those that have been presented in Figs. 4 and 5, while those denoted “0” are another simultaneously recorded pair. (A) Normalized rate functions constructed by counting the number of neural spikes occurring in adjacent 100-s counting windows, and then dividing by 100 s and by the average rate, for RGC 1 and RGC 0. Note that the ordinate scale differs from that in Figs. 4A and 5A. (B) Normalized rate functions for the two corresponding target LGN cells, LGN 1 and LGN 0. (C) Normalized wavelet cross-correlation function (NWCCF) between the RGC 1 and LGN 1 recordings (solid curve), shuffled surrogates of these two data sets (long-dashed curve), and Poisson surrogates (short-dashed curve). Unlike the Allan factor $A(T)$, the normalized wavelet cross-correlation function can assume negative values and need not approach unity in certain limits. Negative normalized wavelet cross-correlation function values for the data or the surrogates are not printed on this doubly logarithmic plot, nor are they printed in (D). Comparison between the value of the normalized wavelet cross-correlation function obtained from the data at a particular counting time T , on the one hand, and from the surrogates at that time T , on the other hand, indicates the significance of that particular value. (D) Normalized wavelet cross-correlation functions between RGC 1 and LGN 1 (solid curve, repeated from (C)), the two RGC spike trains (long-dashed curve), and the two LGN spike trains (short-dashed curve). Also included is the aggregate behavior of both types of surrogates for all three combinations of recordings listed above (dotted line). (E) Cross periodograms of the data sets displayed in (C). (F) Cross periodograms of the data sets displayed in (D).

presence of structure at the stimulus period imposed by the drifting grating.

3. DISCUSSION

The presence of a stimulus alters the manner in which spike trains in the visual system exhibit fractal behavior. In the absence of a stimulus, RGC and LGN dark discharges display similar fractal activity (see Fig. 3). The normalized rate functions of the two recordings, when computed for long counting times,

follow similar paths. The R/S , Allan factor, and periodogram quantify this relationship, and these three measures yield values of the fractal exponents for the two spike trains that correspond reasonably well (see Table 1). The normalized interevent-interval histogram, a measure that operates only over relatively short time scales, shows a significant difference between the RGC and LGN responses. Such short-time behavior, however, does not affect the fractal activity, which manifests itself largely over longer time scales.

The presence of a stimulus, either a constant luminance (Fig. 4) or a drifting grating (Fig. 5), causes

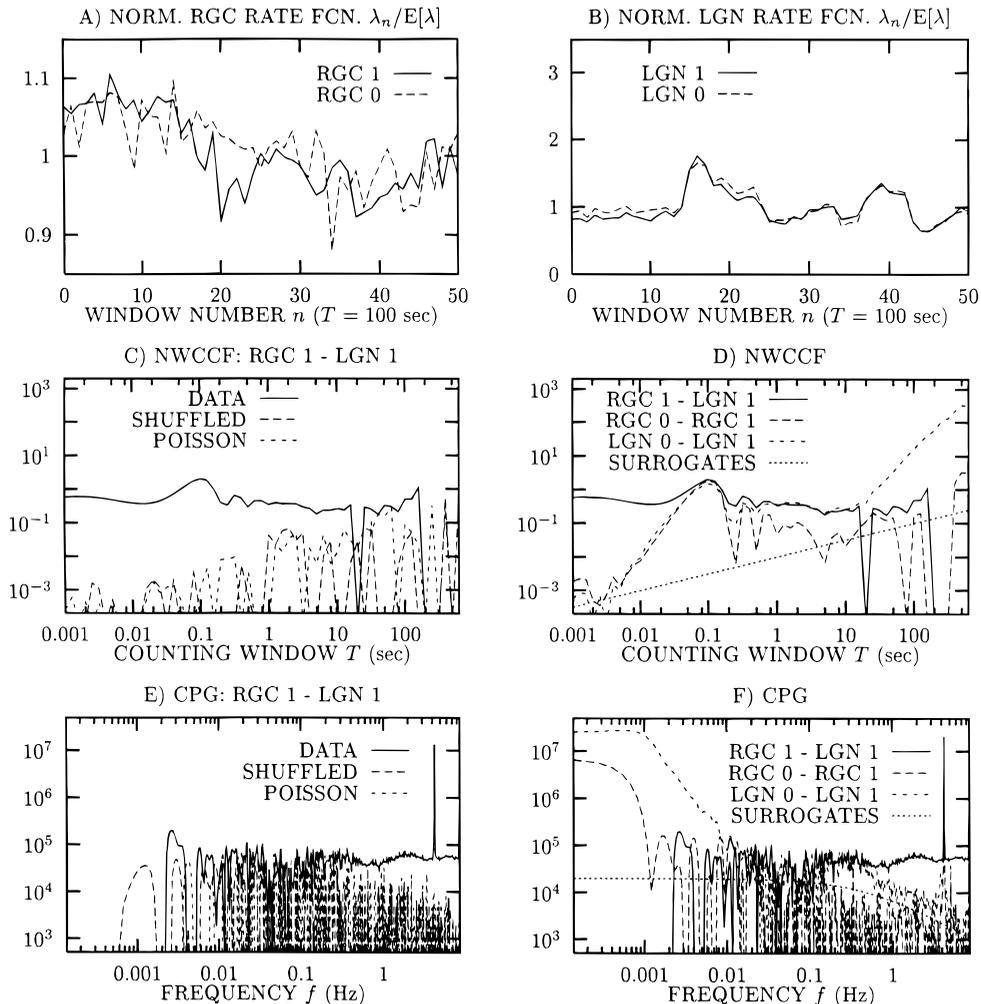


FIG. 7. Statistical measures of the driven discharge from pairs of cat on-center X-type RGCs and their associated LGN cells, stimulated by a drifting grating with a mean luminance of 50 cd/m^2 , 4.2-Hz frequency, and 40% contrast, for data of duration $L = 7000$ s. RGC and LGN spike trains denoted "1" are recorded from the same cell pair that have been presented in Figs. 4–6, while those denoted "0" are recorded simultaneously from the other cell pair, which was presented in Fig. 6 only. (A)–(F) as in Fig. 6.

the close linkage between the statistical character of the RGC and LGN discharges over long times to dissipate. The normalized rate functions of the LGN spike trains display large fluctuations about their mean, especially for the maintained discharge, while the RGC rate functions exhibit much smaller fluctuations that are minimally correlated with those of the LGN. Again, the R/S , Allan factor, and periodogram quantify this difference, indicating that fractal activity in the RGC consistently exhibits a smaller fractal exponent (see also Table 1) and also a smaller fractal onset time (higher onset frequency). Both the R/S and Allan factor measures indicate that the LGN exhibits more fluctuations than the RGC at all scales; the periodogram does less so, apparently because it is the only one of the three constructed without normalization.

In the driven case (Fig. 5), the oscillatory nature of the stimulus phase-locks the RGC and LGN spike trains to each other at shorter time scales. The periodogram displays a peak at 4.2 Hz, and the Allan factor exhibits minima at multiples of $(4.2 \text{ Hz})^{-1} = 0.24 \text{ s}$, for both action-potential sequences. The normalized interevent-interval histogram also suggests a relationship between the two recordings mediated by the time-varying stimulus; both RGC and LGN histograms achieve a number of maxima. Although obscured by the normalization, the peaks do indeed coincide for an unnormalized plot (not shown).

In the presence of a stimulus, RGCs are not correlated with their target LGN cells over the long time scales at which fractal behavior becomes most important, but significant correlation does exist between pairs of LGN spike trains for both the maintained and driven discharges (see Figs. 6 and 7, respectively). These pairs of LGN discharges, exhibiting linked fractal behavior, may be called cofractal. The normalized wavelet cross-correlation function and cross periodogram plots between RGC 1 and LGN 1 remain significantly above the surrogates for small times (Figs. 6C and 6E). The results for the two RGCs suggest some degree of cofractal behavior, but no significant correlation over short time scales for the maintained discharge (Figs. 6D and 6F). Since the two corresponding RGC spike trains do not appear cofractal nearly to the degree shown by the LGN recordings, the cofractal component must be imparted at the LGN itself. This suggests that the LGN discharges may experience a common fractal modulation, perhaps provided from the parabrachial nucleus in the brain stem, which

engenders cofractal behavior in the LGN spike trains. Although similar data for the dark discharge are not available, the tight linkage between RGC and LGN firing patterns in that case (Fig. 3) suggests that overall fractal modulation may not be present in the absence of a stimulus and, therefore, that discharges from nearby LGN cells would in fact not be cofractal; this remains to be experimentally demonstrated. Correlations in the spike trains of relatively distant pairs of cat LGN cells have been previously observed in the short term for drifting-grating stimuli (33); these correlations have been ascribed to low-threshold calcium channels and dual excitatory/inhibitory action in the corticogeniculate pathway (34).

In the context of information transmission, the LGN may modulate the fractal character of the spike trains according to the nature of the stimulus present. Under dark conditions, with no signal to be transmitted, the LGN appears to pass the fractal character of the individual RGCs on to more central stages of visual processing, which could serve to keep them alert and responsive to all possible input time scales. If, as appears to be the case, the responses from different RGCs do not exhibit significant correlation with each other, then the LGN spike trains also will not, and the ensemble average, comprising a collection of LGN spike trains, will display only small fluctuations. In the presence of a constant stimulus, however, the LGN spike trains develop significant degrees of cofractal behavior, so that the ensemble average will exhibit large fluctuations (20). Such correlated fractal behavior might serve to indicate the presence of a stimulus at the visual input, while still maintaining fluctuations over all time scales to ready neurons at later stages of visual processing for any stimulus changes that might arrive. Finally, a similar behavior obtains for a drifting-grating stimulus, but with somewhat reduced fractal fluctuations; perhaps the stimulus itself, though fairly simple, serves to keep more central processing stages alert.

3.1. Prevalence and Significance of Fractal and Cofractal Behavior

Fractal behavior is present in all 50 of the RGC and LGN neural spike-train pairs that we have examined, under dark, maintained-discharge, and drifting-grating stimulus conditions, provided they are of sufficient length to manifest this behavior.

Indeed, fractal behavior is ubiquitous in sensory

systems. Its presence has been observed in cat striate-cortex neural spike trains (35) and in the spike train of a locust visual interneuron, the descending contralateral movement detector (36). It is present in the auditory system (37) of a number of species; primary auditory (VIII nerve) nerve fibers in the cat (19, 38), chinchilla, and chicken (39) all exhibit fractal behavior. It is exhibited at many biological levels, from the microscopic to the macroscopic; examples include ion-channel behavior (40–43), neurotransmitter exocytosis at the synapse (44), and spike trains in rabbit somatosensory cortex neurons (45) and mesencephalic reticular formation neurons (46). In almost all cases, the upper limit of the observed time over which fractal correlations exist is imposed by the duration of the recording itself.

The significance of the fractal behavior is not fully understood. Its presence may serve as a stimulus to keep more central stages of the sensory system alert and responsive to all possible time scales, awaiting the arrival of a time-varying stimulus whose time scale is a priori unknown. It is also possible that fractal activity in spike trains provides an advantage in terms of matching the detection system to the expected signal (37) since natural scenes have fractal spatial and temporal noise (47, 48).

4. CONCLUSION

Using a variety of statistical measures, we have shown that fractal activity in LGN spike trains remains closely correlated with that of their exciting RGC action-potential sequences under dark conditions, but not with stimuli present. The presence of a visual stimulus serves to increase long-duration fluctuations in LGN spike trains in a coordinated fashion, so that pairs of LGN spike trains exhibit cofractal behavior largely uncorrelated with activity in their associated RGCs. Such large correlations are not present in pairs of RGC spike trains. A drifting-grating stimulus yields similar results, but with fractal activity in both recordings somewhat suppressed. Cofractal behavior in LGN discharges under constant-luminance and drifting-grating stimulus conditions suggests that a common fractal modulation may be imparted at the LGN in the presence of a visual stimulus.

ACKNOWLEDGMENTS

This work was supported by the U.S. Office of Naval Research under Grants N00014-92-J-1251 and N0014-93-12079, by the National Institute for Mental Health under Grant MH5066, by the National Eye Institute under Grants EY4888 and EY11276, by the Whitaker Foundation under Grant RG-96-0411, and by the National Science Foundation. E. Kaplan is Jules and Doris Stein Research-to-Prevent-Blindness Professor at Mt. Sinai School of Medicine.

REFERENCES

1. Mastronarde, D. N. (1983) *J. Neurophysiol.* 49, 303–324.
2. Bishop, P. O., Levick, W. R., and Williams, W. O. (1964) *J. Physiol. (London)* 170, 598–612.
3. Kuffler, S. W., FitzHugh, R., and Barlow, H. B. (1957) *J. Gen. Physiol.* 40, 683–702.
4. Levine, M. W., and Troy, J. B. (1986) *J. Physiol. (London)* 375, 339–359.
5. Troy, J. B., and Robson, J. G. (1992) *Vis. Neurosci.* 9, 535–553.
6. Teich, M. C., Heneghan, C., Lowen, S. B., Ozaki, T., and Kaplan, E. (1997) *J. Opt. Soc. Am. A* 14, 529–546.
7. Robson, J. G., and Troy, J. B. (1987) *J. Opt. Soc. Am. A* 4, 2301–2307.
8. Levick, W. R. (1973) in *Handbook of Sensory Physiology. VII/3: Central Processing of Visual Information, Part A* (Jung, R., Ed.), pp. 575–598, Springer-Verlag, New York.
9. Teich, M. C., Matin, L., and Cantor, B. (1978) *J. Opt. Soc. Am.* 68, 386–402.
10. Teich, M. C., and Khanna, S. M. (1985) *J. Acoust. Soc. Am.* 77, 1110–1128.
11. Thurner, S., Lowen, S. B., Heneghan, C., Feurstein, M. C., Feichtinger, H. G., and Teich, M. C. (1997) *Fractals* 5, 565–595.
12. Hurst, H. E. (1951) *Trans. Am. Soc. Civ. Eng.* 116, 770–808.
13. Feller, W. (1951) *Ann. Math. Stat.* 22, 427–432.
14. Mandelbrot, B. B. (1983) *The Fractal Geometry of Nature*, Freeman, New York.
15. Schepers, H. E., van Beek, J. H. G. M., and Bassingthwaighte, J. B. (1992) *IEEE Eng. Med. Biol. Mag.* 11(2), 57–64.
16. Beran, J. (1994) *Statistics for Long-Memory Processes*, Chapman & Hall, New York.
17. Bassingthwaighte, J. B., and Raymond, G. M. (1994) *Ann. Biomed. Eng.* 22, 432–444.
18. Fano, U. (1947) *Phys. Rev.* 72, 26–29.
19. Lowen, S. B., and Teich, M. C. (1996) *J. Acoust. Soc. Am.* 99, 3585–3591.
20. Lowen, S. B., and Teich, M. C. (1995) *Fractals* 3, 183–210.
21. Allan, D. W. (1966) *Proc. IEEE* 54, 221–230.
22. Teich, M. C., Heneghan, C., Lowen, S. B., and Turcott, R. G. (1996) in *Wavelets in Medicine and Biology* (Aldroubi, A., and Unser, M., Eds.), pp. 383–412, CRC Press, Boca Raton, FL.
23. Abry, P., and Flandrin, P. (1996) in *Wavelets in Medicine and Biology* (Aldroubi, A., and Unser, M., Eds.), pp. 413–437, CRC Press, Boca Raton, FL.

24. Oppenheim, A. V., and Schafer, R. W. (1975) *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ.
25. Lange, G. D., and Hartline, P. H. (1979) *Biol. Cybern.* 34, 31–34.
26. Lowen, S. B., Ozaki, T., Kaplan, E., and Teich, M. C. (1998) *in Computational Neuroscience: Trends in Research, 1998* (Bower, J. M., Ed.), pp. 447–452, Plenum, New York.
27. Tuckwell, H. C. (1989) *Stochastic Processes in the Neurosciences*, Soc. for Industrial and Applied Mathematics, Philadelphia.
28. Kaplan, E., and Shapley, R. M. (1982) *J. Physiol. (London)* 330, 125–143.
29. Merrill, E. G., and Ainsworth, A. (1972) *Med. Biol. Eng.* 10, 662–672.
30. Hochstein, S., and Shapley, R. M. (1976) *J. Physiol. (London)* 262, 237–264.
31. Shapley, R. M., and Hochstein, S. (1975) *Nature (London)* 256, 411–413.
32. Funke, K., and Wörgötter, F. (1997) *Prog. Neurobiol.* 53, 67–119.
33. Sillito, A. M., Jones, H. E., Gerstein, G. L., and West, D. C. (1994) *Nature (London)* 369, 479–482.
34. Kirkland, K. L., and Gerstein, G. L. (1998) *Vision Res.* 38, 2007–2022.
35. Teich, M. C., Turcott, R. G., and Siegel, R. M. (1996) *IEEE Eng. Med. Biol. Mag.* 15(5), 79–87.
36. Turcott, R. G., Barker, P. D. R., and Teich, M. C. (1995) *J. Statist. Comput. Simul.* 52, 253–271.
37. Teich, M. C. (1989) *IEEE Trans. Biomed. Eng.* 36, 150–160.
38. Kelly, O. E., Johnson, D. H., Delgutte, B., and Cariani, P. (1996) *J. Acoust. Soc. Am.* 99, 2210–2220.
39. Powers, N. L., Salvi, R. J. (1992) *in Abstracts of the Fifteenth Midwinter Research Meeting of the Association for Research in Otolaryngology* (Lim, D. J., Ed.), Abstract 292, p. 101, Assoc. for Research in Otolaryngology, Des Moines, IA.
40. Läger, P. (1988) *Biophys. J.* 53, 877–884.
41. Millhauser, G. L., Salpeter, E. E., and Oswald, R. E. (1988) *Proc. Natl. Acad. Sci. USA* 85, 1503–1507.
42. Liebovitch, L. S., and Tóth, T. I. (1990) *Ann. Biomed. Eng.* 18, 177–194.
43. Lowen, S. B., and Teich, M. C. (1993) *in Noise in Physical Systems and 1/f Fluctuations*, AIP Conference Proceedings 285 (Handel, P. H., and Chung, A. L., Eds.), pp. 745–748, Am. Inst. of Physics, New York.
44. Lowen, S. B., Cash, S. C., Poo, M.-m., and Teich, M. C. (1997) *J. Neurosci.* 17, 5666–5677.
45. Wise, M. E. (1981) *in Statistical Distributions in Scientific Work 6* (Taillie, C. E. A., Ed.), pp. 211–231, Reidel, Boston.
46. Grüneis, F., Nakao, M., Mizutani, Y., Yamamoto, M., Meesmann, M., and Musha, T. (1993) *Biol. Cybern.* 68, 193–198.
47. Olshausen, B. A., and Field, D. J. (1996) *Network* 7, 333–339.
48. Dan, Y., Atick, J. J., and Reid, R. C. (1996) *J. Neurosci.* 16, 3351–3362.