FIRING RATES OF A RETINAL NEURON ARE NOT PREDICTABLE FROM INTERSPIKE INTERVAL STATISTICS

MICHAEL W. LEVINE, *Department of Psychology, University of Illinois at Chicago Circle, Chicago, Illinois 60680 U.S.A.*

ABSTRACT The intervals between successive action potentials (impulses, or "spikes") produced during the maintained firing of a neuron (ISIs) are often treated as if they were independent of each other; that is, an impulse train is considered as a stationary renewal process. If this is so, the variability of the mean rate of firing impulses in a sequence of temporal windows should be predictable from the distribution of ISIs. This was found not to be the case for the maintained firing of retinal ganglion cells in goldfish. Although some evident nonstationarity sometimes resulted in greater variability of the observed rate distributions than those predicted (for relatively long temporal windows), as a general rule the observed rate distributions were considerably less dispersed than would be predicted by sampling of the ISI distributions. This was taken as evidence of long-term serial dependency between successive ISIs; however, two standard tests for dependency (autocorrelations and serial correlograms failed to reveal structure of sufficiently long duration to account for the effect noted.

INTRODUCTION

Studies of the variability of neural firing may concentrate upon either of two variables: the intervals between successive nerve impulses or "spikes" (interspike intervals: ISIs), or the mean firing rate, defined as the number of impulses that occur in a specified period of time. The simplest representation of the ISI variability is the ISI histogram, which describes the distribution of ISIs of various durations. The equivalent measure of the variability of firing rate is the rate distribution, the distribution of the numbers of impulses occurring in a sequence of temporal windows divided by the total duration of the windows.

It seems reasonable that if there is a single source of variability in the firing of a neuron, the rate distribution should be directly related to the ISI histogram. In particular, if the impulse train is stationary (so that a sample taken late in an experiment is statistically equivalent to one taken early) and generated by a renewal process (so that successive ISIs are independent of each other), one would expect the variability of the rate to be a result of the sampling of finite numbers of ISIs drawn from the ISI distribution. Given these restrictions, a mathematical description of the form of the ISI distribution should be sufficient for deriving the mathematical form of the rate distribution. Barlow and Levick (1969) felt it was a shortcoming of their model for ISI variability that the form of the ISI distribution was not consonant with the form of the rate distribution; the model proposed by Teich et al. (1978) was specifically justified by the fact that rate distributions were derivable from the ISI distributions.

This report demonstrates that in at least one case—the maintained firing of impulses by a

retinal ganglion cell in the goldfish—the rate distribution is not explicable from the statistics of sampling the observed ISI distribution. The observed rate distribution is in fact considerably less variable than would be predicted. This is in the direction opposite to what one would expect if the process were nonstationary, so any nonstationarity must be counteracted by a potent process that tends to stabilize firing rates. Serial correlations or temporal structure (i.e., nonrenewal processes) could have this effect, but long-term effects are not revealed by the usual tests for such higher-order structure.

METHODS

Nerve impulses were recorded extracellularly from ganglion cells in goldfish retinae. Either the retinae were removed under dim red lighting and maintained with a stream of moist, humid oxygen at 25°C (Levine and Shefner, 1977a; Shefner and Levine, 1979), or the impulses were recorded from the optic tracts of whole, self-respiring fish (Shefner and Levine, 1976). There was virtually no ambient or background lighting during an experiment, and cells were not stimulated. Data were collected in periods (gates) of up to 30 s duration, with about the same length of time intervening between successive gates. During each gate, the time of occurrence of each impulse was recorded to the nearest millisecond. An experiment consisted of at least 20 gates.

Histogram

The ISI histograms for all gates in an experiment were computed with a 1-ms binwidth. For clarity, successive bins were combined to present histograms with binwidths of 2–6 ms. However, the 1-ms resolution was retained for all calculations based on the ISI histogram. Histograms were normalized by dividing the numbers of ISIs in each bin by the total number of ISIs in the histogram; the ordinate thus represents an estimate of the probability of obtaining ISIs of each duration.

Rate Distributions

The rate distribution represents the distribution of the total numbers of impulses in each gate divided by the gate duration. The variability of the numbers of impulses/s in each gate depends upon the duration of the gate; to examine this dependency, "gates" of shorter duration than the original gates were generated by discarding the final portion of each gate.

Rate distributions were predicted from ISI distributions according to two methods. The first method was the shuffling of ISIs, in which the order of occurrence of ISIs was randomized. Randomization was based on the output of a random number generator that produced a uniformly distributed sequence on the interval (0, 1) (supplied as part of the Fortran IV software by Digital Equipment Corp., Maynard, Mass.).

The procedure began by multiplying the mean ISI by the first random number in the sequence to compute a "latency" to the first impulse in the first "gate" to be generated from the ISI histogram. The second random number was used to select an interval from the ISI histogram. The interval duration was added to the running time in the "gate," and the count in the selected bin of the ISI histogram decremented by one. The probability of choosing a particular interval was therefore proportional to the number of intervals of that duration, without regard to their actual sequence of occurrence. Decrementing the count deleted the interval so it could not be chosen again (although others of the same duration, if there were any, could be selected). This process was repeated until the running time exceeded the gate length; at this point, the number of impulses added into the gate was counted and divided by gate duration to give a rate, and one was added to the appropriate bin of a generated rate distribution. The excess by which the running time exceeded the gate duration was taken as the initial latency for the next gate: selection of ISIs continued until that gate was filled. This continued until all the intervals were exhausted. The ISI histogram of the "gates" produced was thus identical to the actual ISI histogram except for the 20 or so intervals lost as latencies, and those remaining whose sum was not enough to comprise an additional gate.

The second method of producing a predicted rate distribution was by drawing ISIs at random from a statistically equivalent distribution. The observed ISI histogram was taken as a representation of the underlying probability density function for ISIs. The procedure for generating "gates" by the equivalence method was the same as in the shuffling method, except that intervals were not deleted after being selected. The equivalence method was terminated when the same number of "gates" had been generated as the number of gates in the original data. The intervals were selected with a probability proportional to the number of intervals of each duration actually observed; this probability did not change throughout the procedure. Given a large enough number of intervals picked, the ISI distribution of the generated "gates" would approach the observed ISI distribution, and there would be virtually no difference between this method and the shuffling method. It is worth noting that the seed numbers for the random number generator in the equivalence method were the numbers left by the last operation of the shuffling method, so the same sequence was not taken in both. The sequence also varied from analysis to analysis, as the original seed numbers for the shuffling method included the number of ISIs observed.

Higher Order Statistics

One way to look for temporal patterning in the impulse train is with the autocorrelation (renewal density of Gerstein and Kiang, 1960; Perkel et al. 1967; it is also a normalized version of the autocorrelogram described by Rodieck, 1967; 1973). The autocorrelation is a histogram of the number of impulses occurring within a time window (binwidth) at various delay times after the occurrence of each impulse. (The autocorrelation can also be applied to times preceding an impulse, but this "negative" portion of the function must exactly mirror the positive portion, and so adds no new information.) The autocorrelation represents the rate of firing impulses at t_{i+j}, given the occurrence of an impulse at time t_i.

Autocorrelations were computed with 1-ms resolution for delay times up to 500 ms, and with 5-ms resolutions for delay times from 0.5 to 2.5 s. When the entire function was considered, the portion representing the first 500 ms was collapsed into 5-ms bins. The autocorrelation was taken as the sum of the histograms of impulses after each impulse that occurred from the start to within 2.5 s of the end of each gate; it was normalized by dividing by the number of impulses within those periods and by the binwidth. The units of the autocorrelation were impulses/s, and its expected (asymptotic) value was the mean firing rate (inverse of the mean interval; see Perkel et al., 1967).

The autocorrelation reveals structure in the time domain; specifically, the frequencies represented here range from ~0.4 Hz (because autocorrelations were carried out to $\tau = 2{,}500$ ms) to 500 Hz (limited by the 1-ms time resolution). However, not all temporal structure revealed by the autocorrelation implies a nonrenewal process. If the firing is regular (little variability of the ISIs compared to their mean) a regular structure will result, even though the intervals may be chosen entirely at random from the given ISI distribution (Perkel et al., 1967; Rodieck, 1967; Schellart, 1973). To separate effects imposed by the particular ISI distribution from nonrenewal temporal structure, functions were derived from the ISI histogram that would show the expected autocorrelation on the assumption of a renewal process.

Two distinct methods were used to generate these functions. In the first, 2,560 intervals were drawn at random from the ISI histogram using the same algorithm as the shuffling of intervals described above, and an autocorrelation generated from the shuffled impulse train. This method provides a Monte Carlo approximation to the expected autocorrelation for a renewal process; it is also as uneven as an actual autocorrelation containing the same number of intervals. A less variable approximation was obtained by deriving the autoconvolution of the ISI histogram, where the autoconvolution, $h(\tau)$, is approached by:

$$h_n(\tau) = \sum_{i=1}^{n} k_i(\tau), \tag{1}$$

with $k_i(\tau)$ defined by:

$$k_i(\tau) = k_{i-1}(\tau) * f(t). \tag{2}$$

Here, $f(t)$ is the ISI distribution, n represents the number of terms included, and $*$ denotes convolution. k_0 is defined as a delta function, so k_1 is simply f. Each term may be viewed as the expected distribution of groups of n intervals (the higher-order histograms) on the assumption of a renewal process. (The observed autocorrelation may be computed by summing the actual distribution of various orders, and this was done with no noticeable effect upon the autocorrelations.)

Autoconvolutions were computed with 1-ms resolution for $0 < \tau < 500$ ms; they were then collapsed into 5-ms bins for comparison with the autocorrelations. As a general rule, 16 terms were included in the sum ($n = 16$); in some cases, more terms were required for the function to reach its asymptote at the longest delay times. The autoconvolution also approaches an asymptote which is the mean firing rate (inverse of the mean ISI); it should be noted that this definition of the mean firing rate approached by autocorrelations and autoconvolutions is slightly different from the "rate" used in the rate distributions. There, the rate was defined as the number of impulses divided by total gate time (including the latencies and the dead times between the final impulse and the end of the gate). This creates only a very slight discrepancy, which never lead to a difference in the estimates of mean rate of >0.3%.

Another way of looking for structure within an impulse train is by serial correlation coefficients (Hagiwara, 1949, cited by Moore et al., 1966; Kuffler et al., 1957; Stein, 1965; Rodieck, 1967; 1973; Schellart, 1973). Serial correlation coefficients are a measure of the dependency of the duration of an interval upon the duration of the preceding interval (1st order), or the interval preceding that (2nd order), etc. Serial correlation coefficients do not measure exactly the same thing as the autocorrelation, although they are clearly related measures (see Perkel et al., 1967). The autocorrelation examines dependence upon time of occurrence, and the serial correlation coefficient examines dependence upon order of occurrence.

The serial correlation coefficient is the product-moment correlation between two lists of ISIs: $(t_i - t_{i-1})_{i=2 \text{ to } N-m}$ and $(t_i - t_{i-1})_{i=m+2 \text{ to } N}$. In this representation, N is the number of intervals in a consecutive data collection period, and m is the order of the coefficient. The first 15 orders were computed and presented as a function of order; this function is called a serial correlogram (Perkel et al., 1967).

Since the data were collected in several discrete gates, the correlation coefficient computations cannot extend across the entire experiment. Two alternative methods of computing the serial correlogram were used. In the first, the data were treated as a single series, but any terms that extended across more than one gate were discarded. In the second, the mth coefficient was computed for each gate separately, and the values averaged to give the reported value. (The principal difference between the methods is that in the first, the correlation is based on a global mean and standard deviation of ISIs from all the gates, while in the second a local mean and standard deviation is used within each gate.) Insofar as the impulse train is stationary, the two estimates of the mth coefficient should be the same.

RESULTS

Analyses were completed in 15 experiments on 13 different cells in 12 isolated retinae; an additional five analyses were completed on experiments from five whole fish preparations. Qualitatively similar results were obtained with both preparations. Results from the experiments in which the gate duration was at least 10 s are summarized in Table I.

In general, the observed rate distribution was less variable than the rate distribution produced from the ISI histogram, either by the method of shuffling or of equivalence. Figure 1 shows a typical rate distribution (top) and the corresponding distribution produced from the ISI histogram by the shuffling method (middle). It is evident that the rate distribution produced by shuffling is broader than the actual rate distribution. The two distributions can be characterized by their first two moments, the mean and standard deviation (which is the square root of the second moment, variance). In Fig. 1 the means are virtually identical; in all the analyses reported here, the means never differed by more than a few percent. The

TABLE I
RESULTS OF EXPERIMENTS IN WHICH DURATION WAS ≥ 10 s

	$\sigma_{isi} \cdot \mu_{isi}^{-3/2}$ Predicted $\sigma_{N/T}$ at 1 s	Observed $\sigma_{N/T}$ at 1 s (b)‡	μ_{isi}	σ_{isi}/μ_{isi}	Slope of short duration segment (m)	Break time	Expected $\Delta a(T)$ @ 0.5 s (from Eq. 5)
	(*impulses/s*)	(*impulses/s*)	(*ms*)			(*s*)	(*impulses/s*)
Isolated retina preparations							
	4.15	2.00	44.77	0.88	−0.84	2.0	−0.06
	4.88	3.50	40.54	0.98	−0.75	4.5	−0.18
*	3.42	6.90 (3.60)‡	26.73	0.56	−0.56	0.35	−0.04
*	4.81	5.00 (3.80)‡	46.76	1.04	−0.70	0.70	−0.21
	2.50	1.40	64.69	0.64	−1.02	1.0	+0.01
	4.31	3.70	17.31	0.57	−0.49	1.7	+0.05
§	4.17	2.22	41.05	0.84	−0.69	2.2	−0.06
	5.25	2.00	23.56	0.81	−0.71	1.2	−0.03
	2.72	3.70 (2.55)‡	38.56	0.53	−0.50	0.50	0.00
	1.95	1.80	82.64	0.56	−0.48	1.7	+0.01
‖	6.61	2.50	23.56	1.01	−0.65	10.0	−0.04
‖	6.17	3.70 (2.32)‡	23.41	0.94	−0.90	0.77	−0.04
Whole fish preparations							
	5.13	4.70	34.36	0.95	−0.64	5.0	−0.19
	4.61	3.80	48.89	1.02	−0.62	3.4	−0.15
	6.32	4.30	35.93	1.20	−0.74	1.6	−0.23
	5.33	3.20	37.08	1.03	−0.75	1.9	−0.13
	4.95	4.33	44.12	1.04	−0.65	8.5	−0.21

*These two experiments were performed on different neurons within the same retina.
‡Values given for observed $\sigma_{N/T}$ at 1 s are taken from the line fit to the data, and thus represent b. For those cases in which the breakpoint occurred at a duration <1 s, the observed value at 1 s is given, and the extrapolated value of b given in parentheses.
§This is the experiment that provided the data in Figs. 2–5.
‖These two experiments were performed upon the same neuron.

standard deviations, however, were quite different; the standard deviation of the shuffled rate distribution in the figure is nearly three times as large as the standard deviation of the actual rate distribution. In the analyses reported here, the standard deviation of the shuffled rate distributions could be as much as four times the standard deviation of the corresponding actual distribution.

The standard deviation of a rate distribution (either actual or shuffled) depends upon the gate duration. The longer the gate, the larger the number of impulses in each gate, and hence the better the estimate of mean firing rate. For a renewal process, the variance of the total number of events in a series of sampling periods of long duration approaches

$$\sigma_N^2 = T\sigma_{isi}^2 \mu_{isi}^{-3}, \tag{3}$$

where σ_N^2 is the variance of the numbers of events, T is the duration of each sampling period, σ_{isi}^2 is the variance of the inter-event intervals (ISIs), and μ_{isi} is the mean ISI (Feller, 1966). The standard deviation of the rate is therefore predicted by

$$\sigma_{N/T} = \sigma_{isi} T^{-1/2} \mu_{isi}^{-3/2}. \tag{4}$$

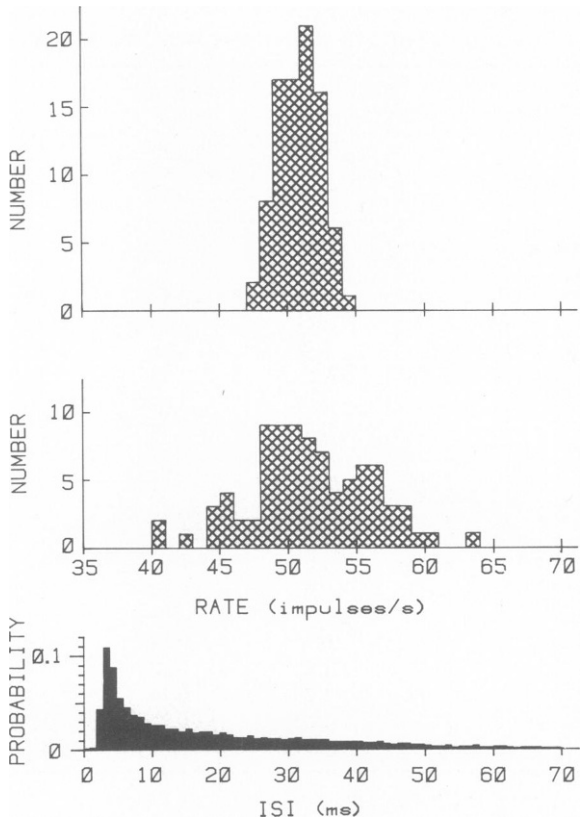

FIGURE 1. Rate distribution of 88 consecutive 2-s gates (top) and 86 "gates" of 2 s duration derived by shuffling the ISIs (middle). 1-impulse/s binwidths. The mean impulse/s of the 88 gates was 51.15, with a standard deviation of 1.53; the mean for the gates derived by shuffling was 51.81, with a standard deviation of 4.44. (Bottom) ISI distribution corresponding to the upper rate distribution. 9,003 ISIs in 1-ms bins. Mean ISI = 19.32 ms; standard deviation = 16.89 ms.

Thus, the standard deviation of rates should decline as the inverse of the square root of the gate duration. This relationship may be examined by plotting log standard deviation vs. log gate duration, as was done for the experiment shown in Fig. 2. The standard deviations of actual rate distributions are shown as solid circles; the standard deviations of rate distributions generated from the ISI histogram are represented by triangles (equivalence method) and open circles (shuffling). The straight line drawn through the triangles and open circles is the prediction made by Eq. 4. It is evident from the figure that the standard deviation of the actual distribution is lower than that of distributions predicted by either Monte Carlo method or by Eq. 4 at each gate duration, as in Fig. 1. There is no noteworthy difference between the standard deviations of the distributions by equivalence and those predicted by shuffling; either prediction is fit satisfactorily by Eq. 4. This indicates validation of the randomization procedures.

The features of Fig. 2 are typical of a similar analysis of the 17 experiments in which the gate duration was 10 s or greater. The standard deviation of the predicted rate distributions (by equivalence and by shuffling) were always well fit by the straight line of Eq. 4. The actual

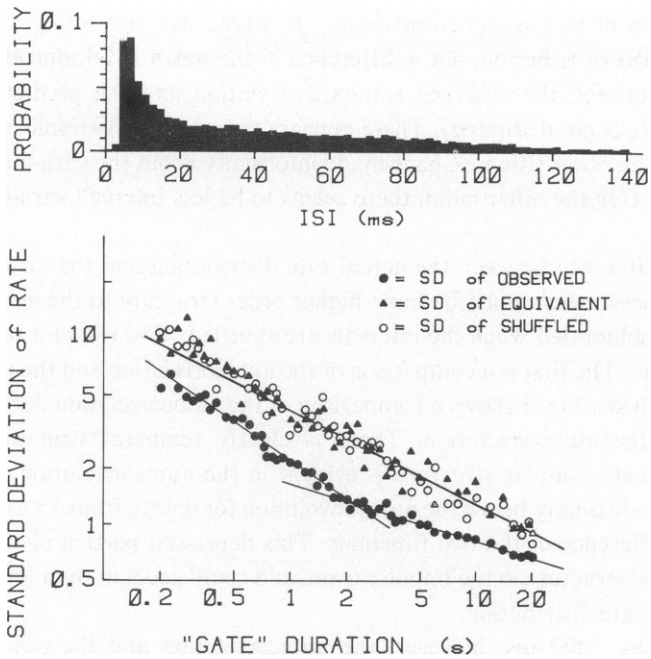

FIGURE 2. Standard deviation of rate distributions as a function of gate duration. Points corresponding to rate distributions generated from the ISI distribution by the method of shuffling (open circles) and by equivalence (triangles) are fit with the straight line predicted by Eq. 4 ($\sigma = 4.17T^{-1/2}$). Actual rate distribution data (closed circles) are fit with two straight line segments, whose slopes are -0.69 and -0.30. 25 gates. (Top) ISI distribution corresponding to 20-s gates. 12,143 ISIs in 2-ms bins. Mean ISI = 41.05 ms; standard deviation = 34.66 ms.

rate distributions, however, cannot be fit by a single line. There were always at least two distinct line segments, generally showing a remarkably sharp breakpoint at their intersection. The segments corresponding to durations longer than the first breakpoint typically had small negative slopes (the observed range was 0 to -0.32), while the segment representing shorter durations typically had a negative slope somewhat steeper than $-\frac{1}{2}$ (the range was -0.48 to -1.02, but most values were close to -0.7). By extrapolation, one would expect each of these segments to intersect the line representing the distributions generated from ISI histograms. In fact, several of the experiments show convergence at short durations, and in nearly half of the experiments there is an intersection at moderate durations so that at the longest durations the actual distributions are more variable than the predicted. It should be noted that the failure of the line segment representing long durations to maintain a slope of $-\frac{1}{2}$ is an indication of nonstationarity in the impulse train (see Discussion).

The standard deviations of the rates at 1-s durations predicted from Eq. 4, the observed standard deviations of the rates at 1 s, the slopes of the line segments representing short durations, and the durations at which the slopes became less steep are listed in Table I. The values in Table I allow a rough comparison of the properties of cells in the isolated retina and those in the whole, self-respiring fish, although the numbers of cells are far too small to allow a quantitative comparison. It appears that there is little qualitative difference between the two preparations, with the possible exception of greater variability in the whole fish. The expected

standard deviation of rate is somewhat higher in whole fish (reflecting a greater standard deviation of the ISI distribution, not a difference in the mean ISIs), and there is apparently less difference between the observed standard deviation and the predicted (although the two-segment curve is not disrupted). There appears to be greater variability in the firing rate in the whole fish, an observation we have made informally about the variability of responses in that preparation. (On the other hand, there seems to be less intercell variability in the whole fish.)

If there is a difference between the actual rate distribution and the distribution predicted for a renewal process, there must be some higher order structure to the actual impulse train; this structure is obliterated when the intervals are shuffled. Two tests for internal order were applied to the data. The first is a comparison of the autocorrelation and the autoconvolution as described in Methods. Fig. 3 shows a comparison of the autocorrelation derived from the data of Fig. 2 with the autoconvolution. There is clearly temporal structure evident in the autocorrelation, but a similar structure is evident in the autoconvolution. Nevertheless, the autocorrelation is distinctly below the autoconvolution for delays from 15 to 150 ms, as can be seen from the difference of the two functions. This depressed portion of the autocorrelation indicates temporal structure in the impulse train, and could account for a lower than predicted variability of the rate distribution.

For delay times >200 ms, however, the autocorrelation and the autoconvolution have clearly approached a level approximating the mean firing rate (indicated on the ordinate and to the right of the graph). To test for the closeness of the approach to mean rate, groups of 15 bins (representing 75 ms) were compared to the mean rate. Comparisons were made starting every 100 ms from 225 ms through 2,425 ms (i.e., 225–300 ms, 325–400 ms, etc.). The mean value of the autocorrelation in each of these periods rarely differed from the mean rate by as

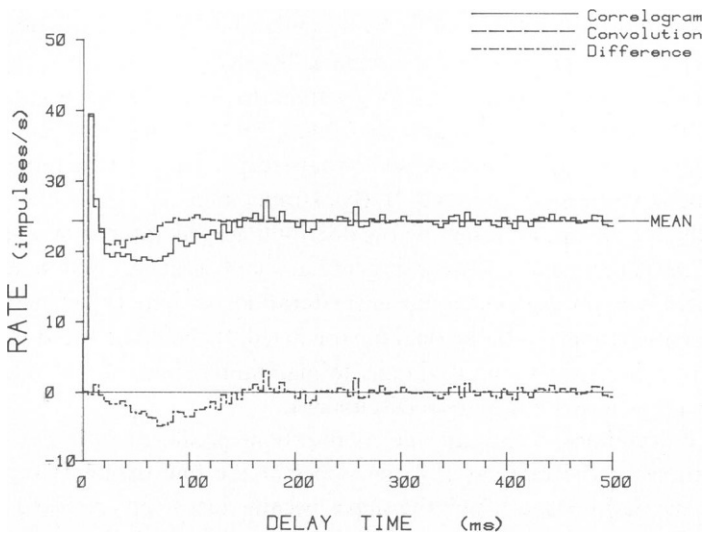

FIGURE 3. Comparison of autocorrelation (solid) derived from 25 gates with the autoconvolution (dashed). The autocorrelation minus the autoconvolution is shown as a dot-dash function. The mean rate (24.36 impulses/s) is indicated on the ordinate. Same data as were used to generate Fig. 2. Binwidth = 5 ms.

much as 0.1 impulse/s, and student t-tests of each comparison almost never reached significance at the 0.2 level. It thus seems safe to conclude that the autocorrelation reveals no temporal structure for times >200 ms; there also was no significant difference between the autocorrelation and the autoconvolution (which approached to within 0.008 impulses/s of the mean rate by 150 ms).

Similar results were obtained from the other experiments; the temporal patterning in the first ~100 ms of the autocorrelation was at least partially present in the autoconvolution although the autocorrelation often was distinctly different from the autoconvolution for ~150 ms. For delays >200 ms, the autocorrelations were not significantly below the mean rate.

The second method of evaluating the fine structure of the impulse train was by looking at serial correlograms. Serial correlograms were computed two ways: an average of the serial coefficients within each gate, and the serial coefficients for all gates combined (see Methods). The first ten terms of the serial correlogram computed from the combined data of Fig. 2 are shown in Fig. 4. The two methods agree very closely, as they did for all but two experiments; in this case the values of the serial correlograms computed each way never differed by >0.002. There is very little serial correlation between intervals; only the first three terms ar noticeably different from 0, and the largest of these is only −0.17. (The 11th–15th terms were all <0.05.) This was the case for all the experiments: only the first few terms of the serial correlogram were significantly different from 0. The mean value of the first term for all 20 experiments was −0.126 with a standard deviation of 0.089. The mean value of the second term was −0.065 with a standard deviation of 0.066.

The negative serial correlation coefficients are presumably another manifestation of the depression of the autocorrelation for delays <150 ms. In fact, they represent about the same delay time; three intervals with a mean ISI of 41.1 ms represents ~120 ms, roughly the delay at which the difference between the autocorrelation and autoconvolution becomes negligible.

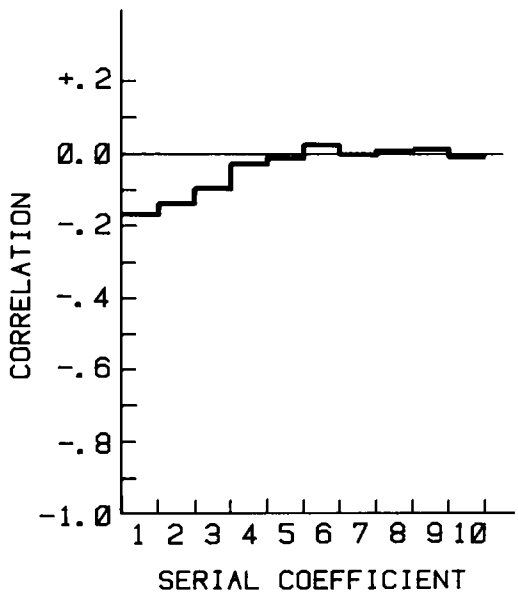

FIGURE 4. First ten terms of the serial correlogram derived from the data used to generate Fig. 2.

DISCUSSION

The rate distributions derived from maintained firing of goldfish ganglion cells do not agree with what would be predicted from the ISI distributions. Even in cases in which the dispersion of the rates is commensurate with that of the ISI distribution, this is a coincidence that depends upon the particular gate duration selected. There were two basic assumptions built into the prediction of rate distributions from ISI histograms: the impulse train represents a process that is stationary and is described as a renewal process. The sections below discuss the implications of a failure of each of these assumptions.

Assumption of Stationarity

The assumption that an impulse train is stationary means that there is no difference in the statistics of data recorded at different times during an experiment. If the data were truly stationary, each gate would be a sample drawn from the identical population; any slow changes that take place with a time course long compared to the gate duration are violations of this assumption. That there is slow variability is well recognized in neurophysiology; in the goldfish retina, slow variability served as a tool for the demonstration of independent on and off response processes (Levine and Shefner, 1975; 1977*a*). Slow variability may appear as an irregular variation in firing from gate to gate (which is virtually indistinguishable from the variability predicted from sampling the ISI distribution), or it may be highly structured (and therefore patent). Structured changes include drifts, in which the firing rate consistently rises or falls throughout the course of an experiment, or it may include cyclic behavior, in which a repeated pattern of increases and decreases in firing may be seen (as described in the cat by Rodieck and Smith, 1967).

In any case, nonstationary behavior should lead to a larger standard deviation of the rate distribution than the ISI histogram would indicate. As an example, consider a consistent drift in firing rate such that the cell fires rapidly in the early part of the experiment, but gradually decreases its mean rate. Within any given gate, there need be little change in rate, as the gate duration is short compared to the duration of the experiment. The ISI distribution contains intervals from all the gates: short intervals from the early gates, long intervals from the last gates. The range of rates can be extreme if the first gate contains only short intervals (and therefore many), and the last gate is made entirely of a few long intervals. But the rate distribution predicted from the ISI histogram does not allow this selectivity; each gate is equally likely to have some long ISIs and some short ISIs, so that either extreme is relatively unlikely.

This effect of nonstationarity explains those cases in which the actual rate distributions were more dispersed than predicted for all gate durations. In all four cases, a plot of the number of impulses in each gate as a function of order of occurrence of the gates showed a clear temporal structure. In two of the cells there was a noticeable drift, one had a strong cyclic behavior, and one showed a cycle superimposed upon a drift. These four experiments also included the only three in which the autocorrelation and autoconvolution were qualitatively different and the only two in which the first term of the serial correlogram (taken as the mean across gates) differed from the first term (all intervals) by >0.1.

Nonstationarity presumably accounts for the actual rate distributions that are more dispersed than predicted at long gate durations. Assuming the longest gate duration to be

short compared to the time course of the nonstationarity, there should be a component of the variability of the rate distribution due to the nonstationarity; this component should be relatively independent of gate duration. Another component of the variation of the rate distribution is that due to variability of the ISIs; this component declines as the square root of gate duration. At some point, the latter component will be negligible compared to the variability due to the nonstationarity, and the standard deviation will be independent of gate duration. The standard deviation of the predicted rate distribution, which cannot take the nonstationarity into account, declines below the level of the actual distribution at the longest gate durations.

Assumption of a Renewal Process

The assumption that the impulse train reflects a renewal process means that the length of any given interval does not depend upon the lengths of the intervals preceding it; each interval is drawn at random from an underlying probability density function. Insofar as this is not true, there is a pattern to the sequence in which intervals are drawn from the probability density function. This pattern might be expected to take one of three related forms: it could be a time-locked sequence due, perhaps, to an oscillatory input to the firing neuron, or it could be interval-locked, such as a threshold for firing that depends upon the length of the interval preceding the last impulse. Alternatively, there could be some more intricate or higher-order relationship involving ensembles of impulses.

Any type of fine structure would have the effect of decreasing the standard deviation of the rate distribution relative to what would be predicted for a renewal process. Whether there is a rapid oscillation in which strings of short intervals follow strings of long intervals or a rule that a long interval should be followed by a short interval, there will be a counterbalancing of long and short intervals within each gate.

A time-locked sequence would be ideally detectable by the autocorrelation, but none of the autocorrelations reported here showed significant evidence of long-term temporal patterns that did not survive the shuffling procedure. The serial correlograms also do show some evidence of coupling between adjacent intervals, but show no evidence of coupling across more than a few intervals. The small negative value for the first coefficients reported here is consistent with (though slightly larger than) values previously reported for the goldfish by Schellart (1973), and for the cat by Rodieck (1967) and Kuffler et al., (1957).

In the remainder of this discussion, I shall refer principally to the autocorrelation as an indicator of long-term variability, rather than to the serial correlogram. I do this for two reasons: first, the autocorrelation is a better indicator of long-term processes for a highly variable signal (Perkel et al., 1967), and second, a process with temporal dependencies extending more than ten times the mean ISI is more likely to be time-dependent than interval-dependent. In any case, the high order terms of the serial correlogram represent correlations across a large number of intervals; given the large variability of the intervals, the serial correlogram and autocorrelation must become equivalent. For delays on the order of ½ s, corresponding in most cases to at least the 10th serial coefficient, the serial correlogram is a scaled version of the autocorrelation.

It is difficult to say whether the small depression of the autocorrelation (or slight negative first few serial coefficients) can account for the lower than predicted variability of the rate

distributions observed. However, the temporal organization of the impulse train revealed by the autocorrelation clearly cannot account for the continuing relative improvement of the actual rate distributions for duration times on the order of seconds. The data representing the standard deviation of observed rate distributions and those representing the theoretical (e.g., Fig. 2) continue to diverge for durations up to ~10 s. Once the gate duration has comfortably exceeded the duration of the temporal structure in the impulse train (~150 ms in Fig. 3) no further relative improvement should be obtained by increasing gate duration, and the lines should be parallel. Once the autocorrelation has reached the mean rate (and the serial correlogram has approached zero), the standard deviation of the rate distribution should decline as the inverse square root of duration. (The exception to the inverse square root rule for long durations is that nonstationarity can make the descent less steep, as mentioned above.) The continuing divergence of standard deviation of the observed rate distributions from the predicted cannot easily be reconciled with autocorrelations that reach asymptote at times shorter than the shortest duration in Fig. 2, a full order of magnitude less than those at which divergence is evident. The serial correlogram also reaches its asymptotic value within a number of terms commensurate with the time at which the autocorrelation reaches its asymptote.

The fact that the autocorrelations do not show a concomitant of the effect noted in plots such as Fig. 2 does not mean that the temporal structure is of some special type to which autocorrelations are immune. There is a mathematical relationship between the autocorrelation and the standard deviation of the rate that is independent of the kind of temporal structure; it is simply that the expected effect upon the autocorrelation is so slight as to be virtually undetectable. If the standard deviation of the rate declines as a power of the duration (that is, there is a straight line on double-logarithmic coordinates with a slope of m and a value b at $T = 1$ s), the expected difference between the autocorrelation and the mean rate is given by

$$\Delta a(T) = \mu_{isi} b^2 (1 + 3m + 2m^2)\, T^{2m}, \tag{5}$$

where μ_{isi} is the mean ISI (this equation is derived in the Appendix). The expected difference was always small; for durations >200 ms, the expected difference was never as great as 1 impulse/s. Such small differences would be lost in the noisiness of the function. (Values of the expected differences at 500 ms are given in the right-most column of Table I; approximating the serial correlogram by the autocovariance scaled by an estimate of the variance of the rate at the mean ISI indicates an expected 10th coefficient that is <0.005.)

The implication is that the impulse train has some temporal structure to which the autocorrelation and serial correlograms are quite insensitive. This structure, which must be effective over relatively long time spans (that is, must include low temporal frequencies), might involve some kind of ordering of impulses into patterned bursts. Nerve impulses do seem to occur in bursts, although it is hard to say to what extent the patterning is actually ordered and to what extent it is fortuitous. To make this point, an actual impulse train (one of the gates in the data of Figs. 2–4) was displayed as if it were a long single oscilloscope sweep. The intervals in this same gate were then shuffled (as described in Methods), and the shuffled impulse train presented in a similar sweep. The two sweeps covered the length of an entire

wall; visitors (not all naive) were unable to point to defining characteristics of the two sweeps, or decide with any confidence which was the (structured) actual record, and which the shuffled impulse train. Segments from each are shown in Fig. 5.

A model for impulse generation in the eccentric cells of the lateral eye of *Limulus* (Shapley, 1971) has the properties required to explain my data. Shapley used a measure proportional to the autocovariance to demonstrate temporal dependencies; he applied Fourier analysis to show that his finding was consistent with filtering of Gaussian noise by a linear filter with a low frequency cutoff, presumably due to self-inhibition. The Gaussian noise in the generator potential was attributed to quantal events in the receptor; it was limited to low and middle frequencies. Shapley showed that the variance spectrum of the impulse rate is well predicted by the product of the variance spectrum of the generator potential times the squared transfer function of the voltage-to-frequency converter; the transfer function was estimated by injecting modulated current into the cell.

This model will also account for the data presented here. A Gaussian input could originate in the receptors or in bipolar cells. A Gaussian input would be necessary to yield a hyperbolic normal distribution of ISIs (Levine and Shefner, 1977*b*). A form of long-lasting self-inhibition would be required to provide attenuation of the low frequencies (and thereby cause the autocorrelation to be below the mean rate for delay times less than the duration of the self-inhibitory process). Self-inhibition could take the form of a cumulative slow after-hyperpolarization of the ganglion cell after each impulse, or it could be due to feedback by the amacrine cells in dyads (Dowling, 1970). The Gaussian input would tend to cause bunching of impulses due to the low frequency components of the input; the self-inhibition would tend to break up the clusters and smooth out the firing rate.

Whatever the exact mechanism by which temporal structure is introduced, the temporal patterning has the interesting and useful property of reducing the variability of the firing rate of the neuron. Despite the considerable variability of the ISIs, the variability of the rate distributions, particularly at moderate durations, is less than would be expected. Variability of the ISIs can be a useful property in that it avoids spontaneous synchronization when a temporally modulated signal is applied (Knight, 1972); on the other hand, if the neurons receiving input from the ganglion cell are monitoring some kind of averaged rate, the signal-to-noise ratio will be improved if the variability of the rate is minimized. The temporal structure of the impulse train mitigates the trade-off between these two aims.

FIGURE 5. (Top) 3½ of an actual impulse train (5-ms resolution). (Bottom) 3½ s of an "impulse train" generated by shuffling the data of which the top record form a portion. Full data set included 481 impulses in 20 s, and was itself a subset of the data used in Fig. 2.

APPENDIX

We wish to state the relationship between the autocorrelation and the standard deviation of the firing rate. Let us begin by following Shapley (1971) in defining the autocovariance, $c(\tau)$:

$$c(\tau) = E[\delta u(t)\, \delta\, u(\tau + t)], \tag{A1}$$

where $E(x)$ denotes the expected value of x (mean); $u(t)$ is the observed rate of firing, defined as $1/(t_{n+1} - t_n)$ in the interval between consecutive impulses at t_n and t_{n+1}; $\delta u(t)$ is the difference between the observed rate at time t and the mean firing rate, $E[u(t)] \equiv \bar{u}$. Thus

$$u(t) = \delta\, u(t) + \bar{u}. \tag{A2}$$

The mean rate is given by the inverse of the mean ISI, μ_{isi}; this is approximately the total number of impulses, N, divided by the total time, T:

$$\bar{u} = 1/\mu_{\text{isi}} \simeq N/T. \tag{A3}$$

Similarly,

$$u(t) \simeq n(t)/\Delta t, \tag{A4}$$

where $n(t)$ is the number of impulses within a bin of width Δt extending from t to $t + \Delta t$.

The autocorrelation reported in this paper is the normalized sum of the numbers of impulses in bins of 1-ms width (the counts in each bin are either 0 or 1):

$$a(t) = \frac{1}{N\Delta t} \sum_{j-1}^{k} n_j \cdot n_{j+t}, \tag{A5}$$

where k is an index on the bin number (time), and $k\,\Delta t = T$. This may be contrasted with the expanded version of Eq. A1, in which the approximation of Eq. A4 is used for $u(t)$:

$$c(t) \simeq \frac{1}{k} \sum_{j-1}^{k} \frac{n_j}{\Delta t} \frac{n_{j+t}}{\Delta t} - \bar{u}^2, \tag{A6}$$

or

$$c(t) \simeq \frac{1}{T\Delta t} \sum_{j-1}^{k} n_j \cdot n_{j+1} - \bar{u}^2. \tag{A7}$$

From Eqs. A3, A5, and A7 we see that

$$a(t) \simeq \frac{1}{\bar{u}}\, c(t) + \bar{u} \simeq \mu_{\text{isi}} c(t) + \bar{u}. \tag{A8}$$

Eq. A8 defines a relationship between the autocovariance and autocorrelation.

We seek a relationship between $c(t)$ and the standard deviation of the firing rate, $\sigma_{N/T}$. The total number of impulses in an interval T is given by the integral of the rate within that interval:

$$N_T = \int_0^T u(t)\mathrm{d}t. \tag{A9}$$

(Note that if $u(t) = \bar{u}$, Eq. A3 follows directly.) The expected value of N_T^2 is then

$$E(N_T^2) = E\left\{\left[\int_0^T u(t)\mathrm{d}t\right]\left[\int_0^T u(t')\mathrm{d}t'\right]\right\}. \tag{A10}$$

Expanding $u(t)$ according to Eq. A2 gives

$$E(N_T^2) = E\left\{\int\int_0^T (\bar{u}^2 + \bar{u}\cdot\delta u(t) + \bar{u}\cdot\delta u(t') + \delta u(t)\cdot\delta u(t'))\mathrm{d}t'\mathrm{d}t\right\}. \tag{A11}$$

The first term on the right hand side is the double integral of the (constant) mean rate squared, which is the square of the total count:

$$\int\int_0^T \bar{u}^2\mathrm{d}t'\mathrm{d}t = \bar{u}^2T^2 = N_T^2. \tag{A12}$$

The second and third terms on the right hand side of Eq. A11 represent integrals of the deviations from the mean, which are 0; the last term will be recognized as the double integral of the autocovariance in Eq. A1, so

$$E(N_T^2) = E(N_T)^2 + \int_0^T\int_0^T c(t'-t)\,\mathrm{d}t'\mathrm{d}t. \tag{A13}$$

Changing the integration variable:

$$E(N_T^2) = E(N_T)^2 + \int_0^T\int_{-t}^{T-t} c(\tau)\,\mathrm{d}\tau\mathrm{d}t. \tag{A14}$$

The inner integral may be taken in two parts:

$$E(N_T^2) = E(N_T)^2 + \int_0^T\left[\int_{-t}^0 c(\tau)\,\mathrm{d}\tau + \int_0^{T-t} c(\tau)\,\mathrm{d}\tau\right]\mathrm{d}t. \tag{A15}$$

The autocovariance is symmetric about the origin (that is, $c[-x] = c[x]$), so the limits of integration may be reversed; this gives

$$E(N_T^2) = E(N_T)^2 + 2\int_0^T\int_0^x c(\tau)\,\mathrm{d}\tau\mathrm{d}x. \tag{A16}$$

The variance of the total counts, σ_N^2 is given by the difference between the expected value of the square and the square of the expected value:

$$\sigma_N^2 = E(N_T^2) - E(N_T)^2. \tag{A17}$$

Eqs. A16 and A17 give the desired result:[1]

$$\sigma_{N/T} = \sqrt{\frac{2}{T^2}\int_0^T\int_0^x c(\tau)\,\mathrm{d}\tau\mathrm{d}x}. \tag{A18}$$

Eq. A18 relates the autocovariance (which is related to autocorrelation by Eq. A8) to the standard deviation of firing rate. We now wish to know how much the autocorrelation must differ from the mean rate to explain the observed relationship between standard deviation of rate and duration, T. In

[1] I am indebted to Professor Bruce Knight of the Rockefeller University for the derivation of this equation

particular, we are concerned with the initial straight segment. A straight segment on double logarithmic coordinates means

$$\log(\sigma_{N/T}) = m \cdot \log(T) + \log(b), \tag{A19}$$

where m is the slope of the line segment, and b is its intercept with T = 1.0 (the value of $\sigma_{N/T}$ at $T = 1$). From Eq. A18

$$\log \sigma_{N/T} = \frac{1}{2} \log \left[\frac{2}{T^2} \int_0^T \int_0^x c(\tau) d\tau \, dx \right]. \tag{A20}$$

From Eq. A19 and A20 we arrive at

$$\frac{b^2}{2} T^{2(1+m)} = \int_0^T \int_0^x c(\tau) d\tau dx. \tag{A21}$$

Differentiating twice leads to an expression for $c(T)$ in terms of the parameters of the observed straight line segment:

$$c(T) = b^2(1 + 3m + 2m^2) T^{2m}. \tag{A22}$$

The autocorrelation is related to the autocovariance by Eq. A8; if we define $\Delta a(\tau)$ as the difference between the autocorrelation and the mean rate:

$$\Delta a(\tau) = a(\tau) - \bar{u} = \mu_{isi} \cdot c(\tau), \tag{A23}$$

we obtain

$$\Delta a(T) = \mu_{isi} b^2(1 + 3m + 2m^2) T^{2m}, \tag{A24}$$

which is Eq. 5 in the text.

I am grateful to my colleagues Dr. Jeremy Shefner and Dr. Leland Wilkinson for helpful advice and discussion. I am especially appreciative to Dr. Shefner, who tolerated spending valuable experimental time to record cells in the absence of stimulation before proceeding with the experiments at hand. This work was supported by grant EYO1951 from the National Eye Institute of the National Institutes of Health.

Received for publication 22 December 1978 and in revised form 9 November 1979.

REFERENCES

BARLOW, H. B., and W. R. LEVICK. 1969. Changes in the maintained discharge with adaptation level in the cat retina. *J. Physiol. (Lond.)*. **202**:699.

DOWLING, J. E. 1970. Organization of vertebrate retinas. *Invest. Ophthal.* **9**:655.

FELLER, W. 1966. An Introduction to Probability Theory and Its Applications. Vol. II. John Wiley & Sons, New York.

GERSTEIN, G. L., and N. Y.-S. KIANG. 1960. An approach to the quantitative analysis of electrophysiological data from single neurons. *Biophys. J.* **1**:15.

KNIGHT, B. W. 1972. Dynamics of encoding in a population of neurons. *J. Gen. Physiol.* **59**:734.

KUFFLER, S. W., R. FITZHUGH, and H. B. BARLOW. 1957. Maintained activity in the cat's retina in light and darkness. *J. Gen. Physiol.* **40**:683.

LEVINE, M. W., and J. M. SHEFNER. 1975. Independence of "on" and "off" responses of retinal ganglion cells. *Science (Wash. D. C.)*. **190**:1215.

LEVINE, M. W., and J. M. SHEFNER. 1977*a*. Variability in ganglion cell firing patterns; implications for separate "on" and "off" processes. *Vision Res.* **17**:765.

LEVINE, M. W., and J. M. SHEFNER. 1977*b*. A model for the variability of interspike intervals during sustained firing of a retinal neuron. *Biophys. J.* **19**:241.

MOORE, G. P., D. H. PERKEL, and J. SEGUNDO. 1966. Statistical analysis and functional interpretation of neuronal spike data. *Ann. Rev. Physiol.* **28**:493.

PERKEL, D. H., G. L. GERSTEIN, and G. P. MOORE. 1967. Neuronal spike trains and stochastic point processes. I. The single spike train. *Biophys. J.* **7**:391.

RODIECK, R. W. 1967. Maintained activity of cat retinal ganglion cells. *J. Neurophysiol.* **30**:1043.

RODIECK, R. W. 1973. The Vertebrate Retina. W. H. Freeman, Co., San Francisco.

RODIECK, R. W., and P. S. SMITH. 1966. Slow dark discharge rhythms of cat retinal ganglion cells. *J. Neurophysiol.* **29**:942.

SCHELLART, N. A. M. 1973. Statistical properties of the ganglion cell discharge; a model for the spike generation. *In* Dynamics and Statistics of Photopic Ganglion Cell Responses in Isolated Goldfish Retina. Mondeel-Offsetdrukkerij, Amsterdam. 57.

SHAPLEY, R. 1971. Fluctuations of the impulse rate in *Limulus* eccentric cells. *J. Gen. Physiol.* **57**:539.

SHEFNER, J. M., and M. W. LEVINE. 1976. A method for obtaining single cell responses in the optic tract of self-respiring fish. *Behav. Res. Methods. Instrum.* **8**:453.

SHEFNER, J. M., and M. W. LEVINE. 1979. A comparison of properties of goldfish retinal ganglion cells as a function of lighting conditions during dissection. *Vision Res.* **19**:83.

STEIN, R. B. 1965. A theoretical analysis of neuronal variability. *Biophys. J.* **5**:173.

TEICH, M. C., L. MATIN, and B. I. CANTOR. 1978. Refractoriness in the maintained discharge of the cat's retinal ganglion cell. *J. Opt. Soc. Am.* **68**:386.