

Discrimination Performance of Single Neurons: Rate and Temporal-Pattern Information

WILSON S. GEISLER, DUANE G. ALBRECHT, RICHARD J. SALVI, AND SAMUEL S. SAUNDERS
*Department of Psychology, University of Texas at Austin, Austin, Texas 78712; and Hearing Research Lab,
State University of New York, Buffalo, New York 14214*

SUMMARY AND CONCLUSIONS

1. A new method of measuring the performance of neurons in sensory discrimination tasks was developed and then applied to single-neuron responses recorded in the auditory nerve of chinchilla and in the striate visual cortex of cat.

2. Most previous methods of measuring discrimination performance have employed decision rules that involve comparing the total counts of action potentials (spikes) produced by two different stimuli. Such measures ignore response pattern and hence may not reflect all the information transmitted by a neuron. The proposed method attempts to measure all (or most) of the transmitted information by constructing descriptive models of the neuron's response to each stimulus in the discrimination experiment; these descriptive models consist of measured probability distributions of the spike counts in small time bins. The measured probability distributions are then used to define an optimal decision rule (an ideal observer) for discriminating the two stimuli. Finally, discrimination performance is measured by applying this decision rule to novel presentations of the same two stimuli.

3. Intensity and temporal-phase discrimination were measured for three neurons in the auditory nerve of chinchilla. The discrimination stimuli were low-frequency pure tones of 70-ms duration. Intensity thresholds were found to be 5–20 dB lower at low intensities using the new *pattern* method compared with the traditional *counting* method. The *pattern* method led to better performance because it utilized both rate and temporal pattern information. Phase discrimination performance using the *counting* method was at chance because the average spike rate did not change with phase. On the other hand, using the *pattern* method, phase discrimination thresholds were found to decrease with intensity, often reaching values equivalent to 30–40 μ s of temporal offset. These thresholds are as good as or better than behavioral thresholds in chinchilla.

4. Contrast and temporal-phase discrimination were measured for three neurons in the striate visual cortex of cat. The discrimination stimuli were drifting sine-wave gratings of 100- to 160-ms duration. Contrast discrimination functions measured by the *pattern* method and the *counting* method were found to be essentially identical. Phase discrimination using the *counting* method was at chance. However, using the *pattern* method, phase thresholds were found to decrease with contrast, reaching values equivalent to 7 ms of temporal offset for the two simple cells.

5. Our results suggest that temporal response pattern carries substantial information for intensity and phase discrimination in the auditory nerve and for phase discrimination in the striate visual cortex. There are likely to be other tasks, such as temporal frequency discrimination and velocity discrimination, for which consideration of only rate information will substantially underestimate discrimination performance.

INTRODUCTION

A long-standing question in neurophysiology concerns the nature of the neural code (Perkel and Bullock 1969); specifically, how is information carried in the responses of individual neurons, and how is this information utilized? Assumptions made regarding the code and how it should be measured can have a significant impact on estimates of the amount of information carried and of the potential level of neural performance (Moore et al. 1966). Consider assessing the discrimination capabilities of a given sensory neuron (e.g., Barlow and Levick 1969a,b; Werner and Mountcastle 1963). To measure the minimum change in the stimulus required to evoke a "reliably different response," the sensory physiologist is required to define (given a certain set of explicit or implicit assumptions) an index of response (e.g., the mean response rate) and an index of reliability (e.g., the trial-to-trial variance associated with the mean rate). If the index of response does not adequately reflect all of the relevant information, and if the index of reliability does not adequately reflect the probabilistic character of the responses, then the estimate of performance will be erroneous.

Most previous studies of discrimination performance in single neurons have not attempted to utilize all of the available information. In the vision literature, studies of detection and discrimination performance have usually considered only the total number of action potentials (spikes) generated in some fixed time period during or after the stimulus presentation (e.g., Barlow and Levick 1969a,b; Barlow et al. 1971; Cohn et al. 1975; Fitzhugh 1958; Shapley and Victor 1986; Tolhurst et al. 1983). A similar approach has been taken by a number of hearing researchers (e.g., Geisler et al. 1985; Relkin and Pelli 1987; Sinex and Havey 1986; Young and Barta 1986). Although some of these studies have demonstrated good discrimination performance, they may have underestimated sensitivity because they did not consider the temporal pattern of the responses. Single neurons may transmit much information in terms of a temporal code; thus it is important to consider the temporal pattern of responses to stimuli when assessing discrimination performance. Indeed, there are some discrimination tasks, such as temporal-phase, temporal-frequency, and velocity discrimination, for which considering only total spikes will greatly underestimate the sensitivity of a neuron.

There have been several attempts to consider the temporal structure of single-neuron responses. Siebert (1970)

developed a model of auditory-nerve responses to pure tones and used this model to derive an appropriate ideal observer for frequency discrimination. The model takes into account both the response rate and temporal pattern. Siebert succeeded in showing that there is a great deal of information for frequency discrimination available in the temporal pattern of auditory-nerve responses. Following Siebert, several recent investigations have also used statistical models of auditory-nerve responses to examine the possible role of temporal pattern information in various auditory discrimination tasks (Goldstein and Srulovicz 1977; Miller et al. 1987). Optican and Richmond (1987) have used Shannon's information theory (Shannon and Weaver 1949) and principle-component analysis to compare the amount of rate and temporal-waveform information carried by single neurons in inferior temporal cortex of monkey.

The present paper describes a new method of measuring the discrimination performance of spike-generating neurons that utilizes both rate and temporal-pattern information and makes minimal assumptions about the neural response properties. The method is based on the concept of the ideal observer from signal detection theory (e.g., Green and Swets 1974). (An ideal observer is a device that performs optimally by using all the available information.) Specifically, we attempt to develop ideal observers that can perform optimal discriminations of single-neuron responses to pairs of stimuli. If this can be achieved, the discrimination performance of the ideal observer provides a single number (proportion correct) that is a precise index of the amount of information available for discrimination in the neuron's response. (See APPENDIX B for more discussion of this measure of information.)

To minimize potentially misleading *a priori* assumptions about the neuron's response properties, the present strategy involves empirically measuring stochastic descriptions of the individual neuron's responses to each stimulus. These descriptive models are then used to construct an ideal observer (optimal decision rule) for each pair of discrimination stimuli. With this approach, it is possible to apply an ideal-observer analysis to arbitrary discrimination stimuli without making overly specific and restrictive assumptions.

There is much to be gained if one can measure all the information transmitted by single neurons in discrimination tasks. To begin with, measurement of all the transmitted information would provide a basis for evaluating the relationship between single-unit activity and sensory performance (e.g., Barlow and Levick 1969a; Tolhurst et al. 1983). Second, such measurements would be of obvious value in testing hypotheses concerning the physiological mechanisms responsible for the neuron's behavior. Third, once the transmitted information has been measured, it can be partitioned to determine how much is carried by changes in rate and how much is carried by changes in temporal pattern (Siebert 1970). Finally, the analysis might be combined usefully with ideal-observer analyses developed for peripheral stages of sensory processing [Barlow 1958; de Vries 1943; Rose 1942, 1948; Tanner and Clark-Jones 1960; see Cohn and Lasley (1986) or Geisler (1989) for a review of the vision literature; Peterson et al. 1954; Van Meter and Middleton 1954; see Green and Swets (1974) for

a review of the early audition literature]. For example, it may be possible to deduce what information is lost between the periphery and the recorded neuron.

METHODS OF ANALYSIS

The present ideal-observer analysis was developed for binary discrimination and detection tasks, in which the stimuli are presented in discrete trials in a single-interval, forced-choice paradigm with equal presentation probabilities and payoffs. Thus the performance of the ideal observer is described completely by its percentage of correct responses, $P(C)$.¹

The transmission of information from one level in a sensory system to the next is often done solely by means of action potentials propagating along the axons of single neurons. Thus electrophysiological recording of the action potentials generated by a single sensory neuron should, in principle, be sufficient for a complete analysis of the information that it transmits. The response on a particular trial is completely represented by the list of the times (relative to stimulus onset) when each spike occurred during a temporal *analysis interval*, or trial, stretching from the stimulus onset until no further information is available from the spike train (or until a decision is forced). With little loss of generality, we assign the spikes into time bins of finite width.² Thus the response of a neuron on a trial is described by a list of integers giving the number of spikes in each bin, $N_1, N_2, \dots, N_{n-1}, N_n$, where N_i is the number of spikes in the i th bin and n is the total number of bins in the analysis interval.

In the discrimination experiment, one of two stimuli (α or β) is presented at random. In the experiments presented and analyzed here, the stimuli were continuous (or near continuous) sine waves (i.e., pure tones or drifting sine wave gratings). These stimuli were chosen to measure discrimination in the absence of onset and offset transients; however, the present methods of analysis apply equally well to transient stimulation.

Figure 1A illustrates hypothetical responses to continuous sinusoidal stimuli at two different intensity levels. Figure 1B shows hypothetical responses to discrete stimuli that are of different intensity (although we will not consider this case further in this paper). Both figures are meant to illustrate the possibility that increasing stimulus level may affect both the rate and the temporal pattern of response.

To apply the discrimination analysis when the stimuli are continuous and periodic, it is necessary to divide the continuous response into discrete trials. To do this, we let the duration of the response (in units of bins) be represented by m (see Fig. 1A). If the nervous system is assumed to have some uncertainty about when the stimulus was presented (which is a reasonable possibility to consider), it becomes necessary to define an analysis interval that is larger than the response duration. Thus let n represent the full analysis interval, which consists of the response duration, m , plus

¹ Note that, once an ideal observer has been developed for this task, it can be generalized to conditions with unequal presentation probabilities and payoffs and to other tasks, such as the 2-interval forced-choice task (e.g., see Green and Swets 1974).

² Spike arrival times are real numbers; however, estimation of the present descriptive models from a finite number of stimulus presentations requires that the arrival times be assigned to time bins of finite width. This imposes little loss of generality because the time bins can be made small enough to ensure that no more than one spike can occur in a bin and that the positions of spikes falling within any given bin are well described by a uniform probability density. Under these circumstances no information would be gained by making the binwidths smaller. However, in the derivations that follow we allow the possibility of large bins, because it is often necessary to use larger bins when there is a shortage of data on which to build the descriptive models.




FIG. 1. Schematic illustration of stimuli and responses used in the ideal-observer analysis of 2-alternative discrimination tasks. The purpose of this figure is to define parameters used in analysis and to illustrate that changing stimulus intensity may affect both response magnitude and response pattern. *A*: continuous periodic stimuli. Stimuli α and β are continuous sine waves that differ in amplitude. For the analysis, the continuous response is divided into discrete responses of length m . When temporal uncertainty is assumed, then the entire response analysis interval (n) is the sum of the response duration, m , and the uncertainty interval, θ . Because responses are periodic, the response period (l) may be less than the response duration (m). *B*: discrete stimuli. Stimuli α and β are energy pulses of fixed duration presented in discrete trials. Because responses are aperiodic, the response period (l) equals the response duration (m). Summary of parameters: n , analysis interval; m , response duration; l , response period; and θ , temporal uncertainty interval. All these parameters are expressed in units of temporal bins.

an uncertainty interval, θ (also in units of bins). When the nervous system is assumed to have no uncertainty, then $n = m$.

When the response to a periodic stimulus is also periodic (as illustrated in Fig. 1*A*), it is possible to make use of the response period (l) to reduce the amount of data that must be collected for an accurate analysis. Specifically, it is necessary to model only one response period (see below).

The task of the ideal observer is to decide whether the response occurring in a trial was produced by stimulus α or by stimulus β . It

is well known that the optimal performance in this situation is obtained by applying the *likelihood decision rule*: the observer computes the probability that the response resulted from stimulus α , computes the probability that the response resulted from stimulus β , and then picks the stimulus with the higher probability. Equivalently, the observer computes the likelihood ratio (L) of these probabilities

$$L = P(N_1, \dots, N_n | \beta) / P(N_1, \dots, N_n | \alpha) \quad (1)$$

and then picks alternative β if the ratio is >1.0 or α if the ratio is <1.0 . The $P(C)$ obtained when this decision rule is used is given by the following formula

$$P(C) = 0.5 + 0.25 \sum |P(N_1, \dots, N_n | \beta) - P(N_1, \dots, N_n | \alpha)| \quad (2)$$

where the summation is over all possible values of N_1, \dots, N_n . This equation is derived in APPENDIX A. As mentioned above, the ideal observer's $P(C)$ is a precise measure of the information available in the neuron's response to perform the discrimination task. (See APPENDIX B for more discussion of this point.)

To understand the basic logic of the present analysis, consider measuring the responses of a neuron to pairs of stimuli in a discrimination experiment. If the neuron's responses were stable enough, and if the neuron could be monitored for a long enough period of time, then each of the two stimuli could be presented many times; from the responses to these repeated presentations, we would learn the probability of obtaining any and all possible responses, no matter how complicated the neuron's behavior. These measured probabilities would constitute a complete descriptive model of the neuron's behavior (with minimal theoretical assumptions). They could then be used to assess optimal discrimination performance by applying Eq. 2 or, equivalently, by applying Eq. 1 (the likelihood decision rule) to novel presentations of the same two stimuli.

However, it is not feasible to measure a complete descriptive model because there are far too many probability values. Thus, for the present empirical approach to work, some *a priori* assumptions must be made. At the same time, it is important to evaluate the degree to which these assumptions allow all the significant information in the neuron's output to be represented. If all the information is not represented, discrimination performance will be underestimated.

The proposed strategy for meeting both demands (limiting probability values while maximizing information represented) is to develop descriptive models based on successively less restrictive *a priori* assumptions. As the assumptions become less restrictive, the amount of information represented increases, and the performance of the ideal observer increases. At some point the gains in performance become minimal. When this point is reached, it is reasonable to conclude that the descriptive model accurately represents most of the discrimination information.

Four levels of descriptive model (with successively less-restrictive assumptions) are described briefly below. This is followed by a discussion of procedures for determining when the models have accurately represented all the information in a neuron's response. It should be kept in mind that, as the assumptions become less restrictive, greater amounts of data are required for accurate measurement of the probabilities in the models. In fact, as we will see later, the fourth and least restrictive model requires more data than is practical to obtain with current electrophysiological procedures.

Counting model

In the first level of descriptive model, we make the following conventional assumption (e.g., Barlow and Levick 1969a): All relevant discrimination information is contained in the sum of the

spikes across all bins. Under this assumption and the assumption of no temporal uncertainty, the likelihood ratio (Eq. 1) becomes

$$L = P(N|\beta)/P(N|\alpha) \quad (3)$$

In this equation, N is the total number of spikes in the response interval on a particular trial (i.e., $N = N_1 + N_2 + \dots + N_n$), and $P(N|\alpha)$ and $P(N|\beta)$ are the probabilities for the total number of spikes in the response interval. (Note that in this case the analysis interval, n , equals the response interval, m .) To apply the maximum-likelihood decision rule, the probability distributions $P(x|\alpha)$ and $P(x|\beta)$, also known as *counting distributions* or *pulse-number distributions* (cf. Teich and Khanna 1985), must be measured. These distributions constitute the descriptive models (one for each stimulus). Figure 2A shows counting distributions measured from a single neuron for two different stimulus contrasts. As can be seen, any spike counts >1 are more likely to have been produced by stimulus β (solid bars). Thus the likelihood decision rule is equivalent to placing a criterion between 1 and 2 spikes per trial.

When uncertainty is introduced, a somewhat more complicated decision rule is required

$$L = \sum_{u=0}^{\theta} P(N(u)|\beta, u) / \sum_{u=0}^{\theta} P(N(u)|\alpha, u) \quad (4)$$

where $N(u)$ is the total number of spikes in the response interval for a starting temporal position of u within the analysis interval. In other words, $N(u) = N_{1+u} + \dots + N_{m+u}$. (Note that in Eq. 4 and elsewhere in this paper, we are assuming that temporal uncertainty has a uniform probability distribution. Other uncertainty functions could easily be substituted, if necessary.)

Pattern model

The major limitation of the *counting* model is that it cannot represent information carried in the temporal response pattern. The following assumption defines a descriptive model that is capable of representing temporal pattern information: All relevant discrimination information is contained in the sum and in the pattern of spikes across bins, and the number of spikes observed in any bin is probabilistically independent of the number observed in all other bins. This class of descriptive model requires measuring counting distributions for each temporal bin.

Consider first the case in which there is no temporal uncertainty. If there is probabilistic independence between bins, then the probability of a particular temporal pattern, that is, a particular list of spike counts (N_1, \dots, N_n), is the product of the probabilities of the count in each bin. In other words

$$P(N_1, \dots, N_n|\beta) = P_1(N_1|\beta)P_2(N_2|\beta) \dots P_n(N_n|\beta)$$

and

$$P(N_1, \dots, N_n|\alpha) = P_1(N_1|\alpha)P_2(N_2|\alpha) \dots P_n(N_n|\alpha)$$

Thus the likelihood ratio (Eq. 1) becomes

$$L = \prod_{i=1}^m P_i(N_i|\beta) / \prod_{i=1}^m P_i(N_i|\alpha) \quad (5)$$

where $P_i(x|\alpha)$ and $P_i(x|\beta)$ are the probability distributions for the i th bin given stimulus α and stimulus β , respectively. Note that these probability distributions are just counting distributions for the individual bins; thus, construction of this descriptive model requires measuring the counting distributions for each bin in the analysis interval. Figure 2B shows the counting distributions for 10 bins, obtained from the responses of a single neuron at two stimulus contrasts.

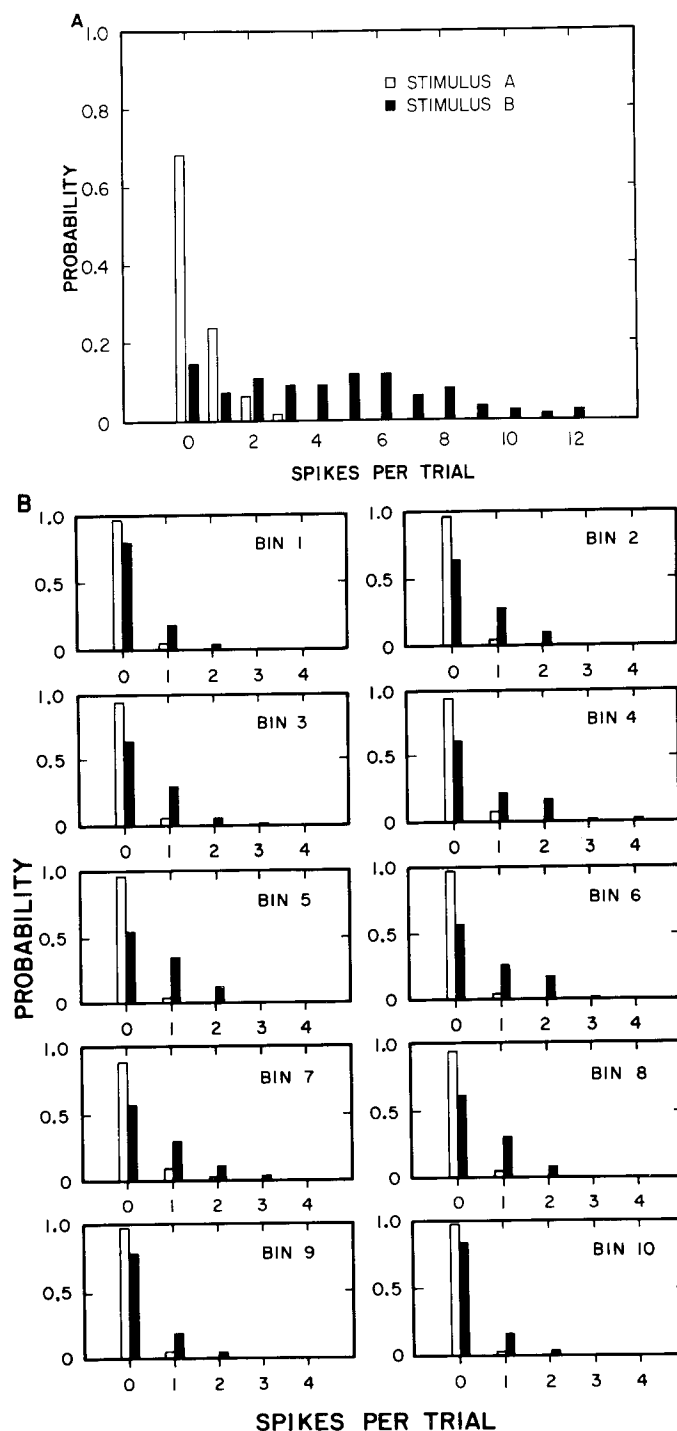


FIG. 2. Examples of measured probability distributions used in the ideal-observer analyses. These particular examples are for a simple cell in cat visual cortex. *A*: *counting* analysis. Probability distributions of the number of spikes generated per trial for 2 contrasts of a drifting sine wave grating. Because spike counts >1 are more likely to have been produced by the higher-contrast stimulus (β), the optimal decision rule is to place a criterion between 1 and 2 spikes/trial. *B*: *pattern* analysis. Probability distributions of the number of spikes generated per trial, in each of 10 temporal bins, for 2 contrasts. The optimal decision rule involves computing the product of the probabilities of the spike counts in each bin using the distributions for the lower-contrast stimulus (α) and then using the distributions for the higher-contrast stimulus (β). Stimulus α is picked if its probability is higher, otherwise stimulus β is picked. Because the *pattern* analysis measures counting distributions for each temporal bin, it measures both rate and temporal pattern information.

The likelihood decision rule would be applied in this case by the use of the 20 counting distributions (10 for each stimulus) to compute the probabilities of the spike counts in the bins. These probabilities would then be substituted into Eq. 5 and the resulting likelihood ratio checked to see if it is greater than or less than 1.0.³ For example, suppose the response on a trial (the list of spike counts in the bins) was 0, 1, 1, 0, 0, 0, 2, 1, 0, 1. The distributions in Fig. 2B show that the probability of obtaining 0 spikes in the first bin is 0.79 if the stimulus were β and is 0.96 if the stimulus were α . Similarly, the probability of obtaining 1 spike in the second bin is 0.27 given stimulus β and 0.018 given stimulus α . Proceeding in this fashion, we find that the numerator in Eq. 5 is $0.79 \times 0.27 \times 0.29 \times 0.61 \times 0.54 \times 0.56 \times 0.118 \times 0.31 \times 0.77 \times 0.154$, and the denominator is $0.96 \times 0.018 \times 0.055 \times 0.94 \times 0.96 \times 0.97 \times 0.0091 \times 0.045 \times 0.96 \times 0.0091$. Dividing these products gives a likelihood ratio of 16,610. Because this ratio is >1.0 , the ideal observer would pick stimulus β on this trial.

When temporal uncertainty is introduced, the likelihood rule becomes

$$L = \sum_{u=0}^{\theta} \prod_{i=1}^n P_i(N_i|\beta, u) / \sum_{u=0}^{\theta} \prod_{i=1}^n P_i(N_i|\alpha, u) \quad (6)$$

Application of this decision rule is very similar to the no-uncertainty rule (Eq. 5); the only difference is that the products of the probabilities must be computed for each possible response position within the uncertainty interval and then summed across all possible response positions. (Again, the temporal-uncertainty distribution is assumed to be uniform.)

Application of the *pattern* model often requires more data than the *counting* model because there are more probability distributions to measure. Thus it is important to make optimal use of the data when applying the *pattern* model. If the responses to the stimuli are periodic, it is possible to reduce the number of probability distributions that need to be measured. Specifically, it is only necessary to measure distributions for the bins within one period of the response, because the distributions within all the other periods are the same. More precisely, if l is the period of the response (see Fig. 1B), then for each bin (i) in the analysis interval

$$P_i(x|\alpha) = P_j(x|\alpha) \quad \text{and} \quad P_i(x|\beta) = P_j(x|\beta) \quad (7)$$

where $j = \text{mod}(i, l)$, if $\text{mod}(i, l) \neq 0$; otherwise, $j = l$. (Note that "mod" represents the modulo function; i.e., j is the remainder of i/l .)

The *counting* model is a special case of the *pattern* model. This is easily seen by noting that the *pattern* model reduces to the *counting* model when the bin size is set to be the entire response interval (i.e., so that there is just 1 bin).

Another important special case of the *pattern* model is obtained by assuming that spike trains are described by a Poisson process (more precisely, an inhomogeneous Poisson process). The Poisson process (which also describes the statistical properties of light) satisfies the independence assumption that defines the *pattern* model. This special case, which we call the *Poisson pattern* model, is developed in APPENDIX A. The *Poisson pattern* model is of interest because it often yields discrimination performances similar to that of the full *pattern* model, but it requires less data to measure

the probability distributions. Siebert's (1970) model and some others in the auditory literature are special cases of the *Poisson pattern* model obtained by assuming a particular intensity function for the Poisson process.

Pattern-1 model

The *pattern* model is based on the assumption that the number of spikes observed in a time bin is statistically independent of the number observed in other bins. This independence assumption must be violated, at least to some extent, by all spike-generating neurons simply because of refractory effects. If refractoriness were the only source of probabilistic dependence, then it would be reasonable to expect that the dependence would only be a function of the elapsed time from the most recent spike (e.g., Teich et al. 1978; Young and Barta 1986). The next level of descriptive model is based on the following assumption: All relevant discrimination information is contained in the sum and in the pattern of spikes across bins; and the number of spikes observed in a bin is, at most, probabilistically dependent on when the previous spike occurred. If the number of spikes observed in a bin is only dependent on when the previous spike occurred, then the probabilities in Eq. 1 are given by the following products

$$P(N_1, \dots, N_n|\beta) = P(T_1|\beta)P_1(N_1|T_1, \beta) \dots P_n(N_n|T_{n-1}, \beta)$$

and

$$P(N_1, \dots, N_n|\alpha) = P(T_1|\alpha)P_1(N_1|T_1, \alpha) \dots P_n(N_n|T_{n-1}, \alpha)$$

where T_i is the waiting time (in bins) from the current bin (the i th bin) to the previous spike (nonempty bin). Because T_1 refers to events before the response, its probability does not depend on which stimulus was presented; therefore $P(T_1|\beta) = P(T_1|\alpha)$. Thus the likelihood ratio Eq. 1 becomes




FIG. 3. Example of Monte Carlo simulations used to determine the amount of data and binwidths required for accurate descriptive models (i.e., probability distributions like those in Fig. 2) of a neuron's response. This example is for a *pattern* model. The procedure begins by selecting parameters of a synthetic spike generator (an inhomogeneous Poisson process) so that it produces spike trains that roughly match those of the neurons to be analyzed. In this example, the spike generator was mimicking an auditory neuron responding to a 100-Hz pure tone at a mean response rate of 150 spikes/s in a temporal phase-discrimination task. A time delay was picked to produce ~77% correct by the mathematically exact ideal observer (solid horizontal line). Curves show the performance of the *pattern* model as a function of the number of trials used to measure the probability distributions for various numbers of bins per response period. (The error bars represent \pm SD.) As can be seen, under these conditions, 100 trials and 10 bins/period are adequate for the *pattern* analysis to approach the exact ideal-observer performance.

³ The probabilities in the bins were obtained by counting the numbers of spikes observed in the bin and dividing by the number of stimulus presentations. If no spikes are observed in a bin it can be problematic (and inaccurate) to assume the probability is literally zero. This is because a spike occurring in that bin during a test trial will completely dominate the decision. Specifically, the spike will force the likelihood ratio to be 0, infinite, or undefined. To avoid this problem, we took the probability for empty bins to be 0.5 divided by the number stimulus presentations (or response periods). The Monte Carlo simulations described later show that this assumption is reasonable.

$$L = \prod_{i=1}^m P_i(N_i|T_i, \beta) / \prod_{i=1}^m P_i(N_i|T_i, \alpha) \quad (8)$$

where $P_i(x|T_i, \alpha)$ and $P_i(x|T_i, \beta)$ are the probability distributions of the spikes counts for the i th bin given a waiting time of T_i , for stimulus α and stimulus β , respectively. Construction of this descriptive model requires measuring the counting distributions for each bin in the analysis interval, for each possible waiting time. To measure all these distributions requires more data than for the *pattern* model.

When there is temporal uncertainty the likelihood ratio is

$$L = \sum_{u=0}^{\theta} \prod_{i=1}^n P_i(N_i|T_i, \beta, u) / \sum_{u=0}^{\theta} \prod_{i=1}^n P_i(N_i|T_i, \alpha, u) \quad (9)$$

Pattern-2 model

The least restrictive assumption that will be considered here is one in which the stochastic dependence is a function of the elapsed times to the two most recent spikes: All relevant discrimination

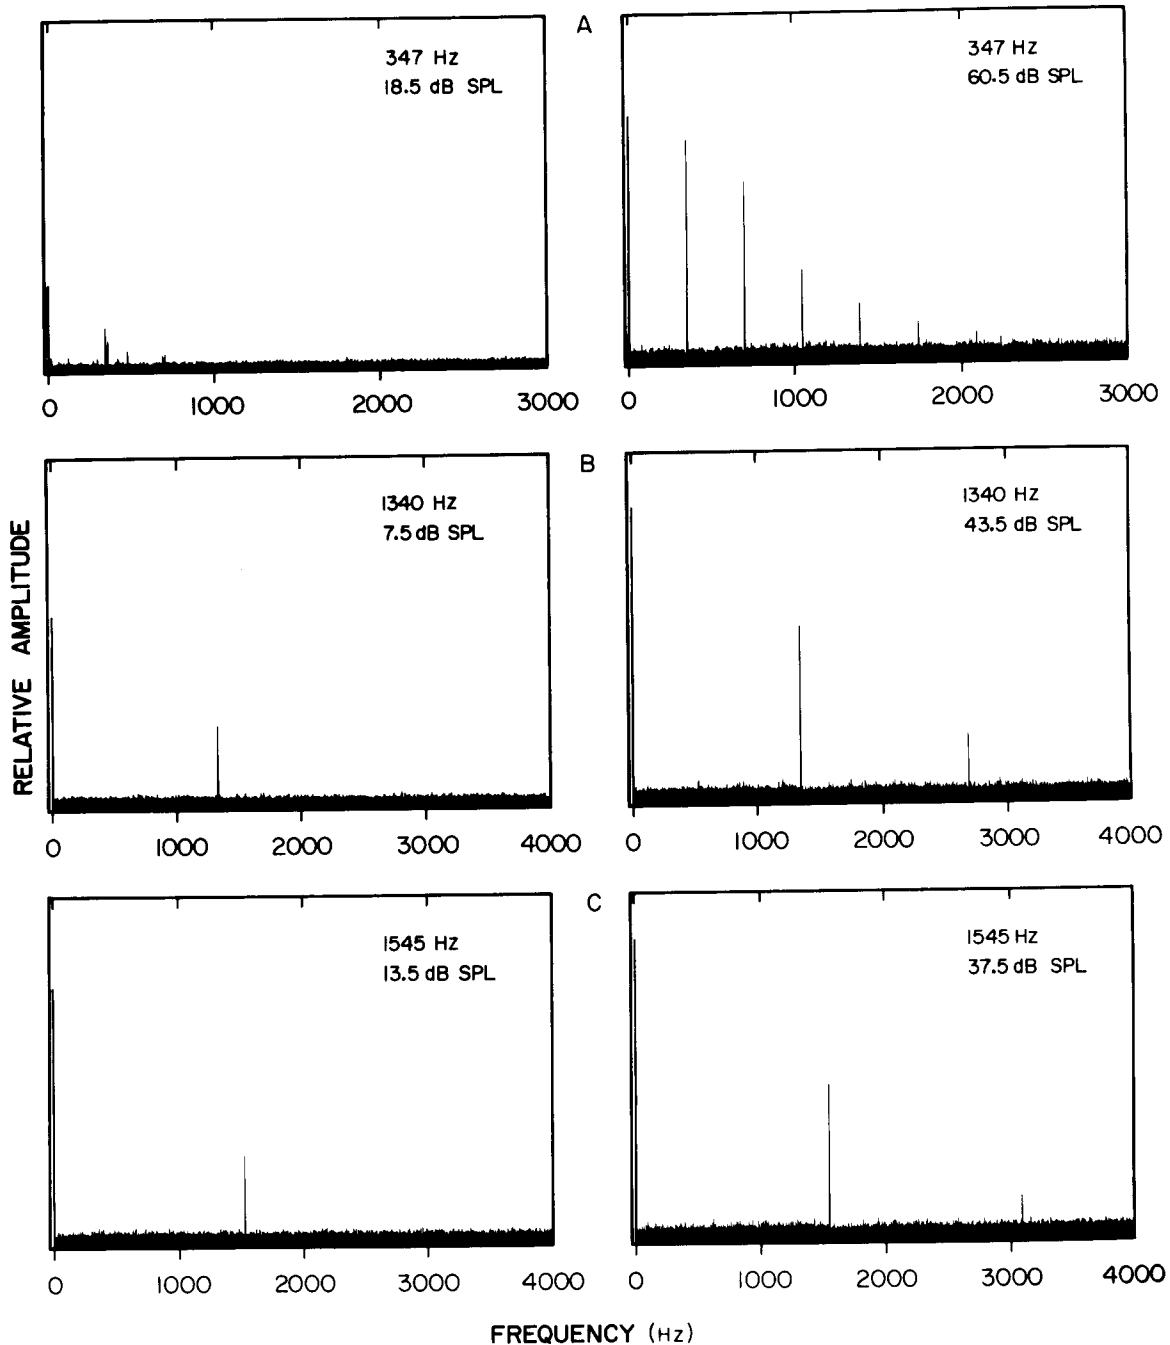


FIG. 4. Amplitude spectra of spike trains recorded from 3 auditory-nerve neurons that were responding to continuous pure tones of near-threshold amplitude (*left*) and near-saturation amplitude (*right*). (Because this is only the right half of the symmetric Fourier transform, the amplitudes of the harmonics should be doubled when comparing them with the dc component.) If the responses were statistically stable and had no other systematic variability, then the only prominent components in the spike trains would be at 0 Hz, at the stimulating frequency, and/or at its higher harmonics. This was true except for the 347-Hz unit at the lowest intensities (*A, left*).

information is contained in the sum and in the pattern of spikes across bins; and the number of spikes observed in a bin is, at most, probabilistically dependent on when the previous two spikes occurred. Under this assumption the probabilities in Eq. 1 are given by the following products

$$P(N_1, \dots, N_n | \beta) = P(T_1, S_1 | \beta) P(N_1 | T_1, S_1, \beta) \dots P(N_n | T_{n-1}, S_{n-1}, \beta)$$

and

$$P(N_1, \dots, N_n | \alpha) = P(T_1, S_1 | \alpha) P(N_1 | T_1, S_1, \alpha) \dots P(N_n | T_{n-1}, S_{n-1}, \alpha)$$

where T_i is the elapsed time (in bins) from the previous spike to the i th bin and S_i is the elapsed time (in bins) from the second most previous spike to the i th bin. Because T_1 and S_1 refer to events before the response, their joint probability does not depend on




FIG. 5. Average response rate of 3 auditory-nerve neurons responding to continuous pure tones of near-threshold amplitude (*left*) and near-saturation amplitude (*right*). Averages were computed with a 0.5-s running integrator. (The dashed line indicates the spontaneous rate.) As can be seen, the responses of the neurons did not drift systematically over the 15-s stimulus duration. Sizes of the fluctuations in average response rate are what would be expected from a stable random process (see text and APPENDIX C).

which stimulus was presented; therefore $P(T_1, S_1|\beta) = P(T_1, S_1|\alpha)$. Thus the likelihood ratio (Eq. 1) becomes

$$L = \prod_{i=1}^m P_i(N_i|T_i, S_i, \beta) / \prod_{i=1}^m P_i(N_i|T_i, S_i, \alpha) \quad (10)$$

where $P_i(x|T_i, S_i, \alpha)$ and $P_i(x|T_i, S_i, \beta)$ are the probability distributions of the spikes counts for the i th bin given waiting times of T_i and S_i . Construction of this descriptive model requires measuring the counting distributions for each bin in the analysis interval for each possible pair of waiting times.

When temporal uncertainty is present, the likelihood ratio becomes

$$L = \sum_{u=0}^{\theta} \prod_{i=1}^n P_i(N_i|T_i, S_i, \beta, u) / \sum_{u=0}^{\theta} \prod_{i=1}^n P_i(N_i|T_i, S_i, \alpha, u) \quad (11)$$

Higher-level models

It is not difficult to extend the descriptive models to higher and less restrictive levels by allowing the possibility that the number of spikes observed in a bin is dependent on even more of the previous spike times. Indeed, if this process were continued, then all the

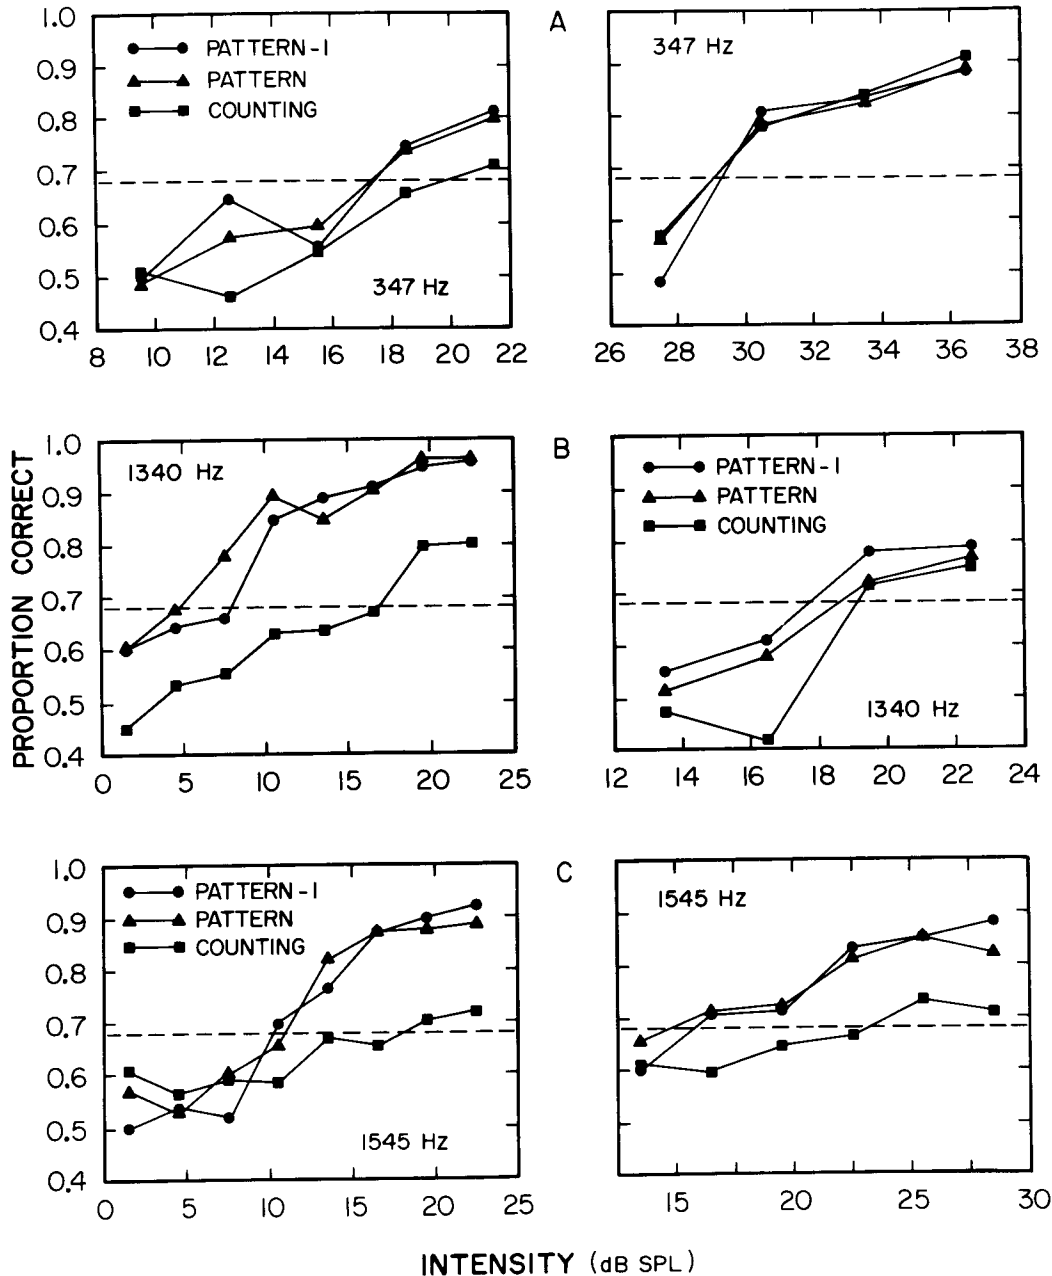


FIG. 6. Neurometric functions (proportion correct as a function of total intensity in dB SPL) of 3 auditory-nerve neurons in an intensity-discrimination task for 3 levels of ideal-observer analysis. *Left*: near-threshold reference intensities. *Right*: near-saturation reference intensities. Trial durations were ~ 70 ms. Frequency indicated in each panel was both the CF of the neuron and the frequency of the stimulus. Results show that temporal response pattern contributes substantial discrimination information, especially at lower intensities (*left*).

discrimination information contained in spike trains could be measured (with no assumptions), no matter how complex the temporal dependencies. This most general case is presented in APPENDIX A. However, as mentioned earlier, measurement of the probability distributions for the *pattern-2* model requires more data than is practical to obtain in a single-unit experiment. Higher-level models would require even more data and hence are not considered here. It is worth noting again that, when the responses are periodic, the number of distributions that must be measured is reduced to some extent.

MEASUREMENT OF IDEAL-OBSERVER PERFORMANCE

Once the probability distributions (which constitute the descriptive models) have been measured, the next step is to evaluate performance [$P(C)$] of the model in the discrimination task. One method would be to compute the theoretical value of $P(C)$ using Eq. 2. Although this approach is elegant, it could be misleading because it does not actually test performance but rather only provides a theoretical estimate of performance. A more empirical

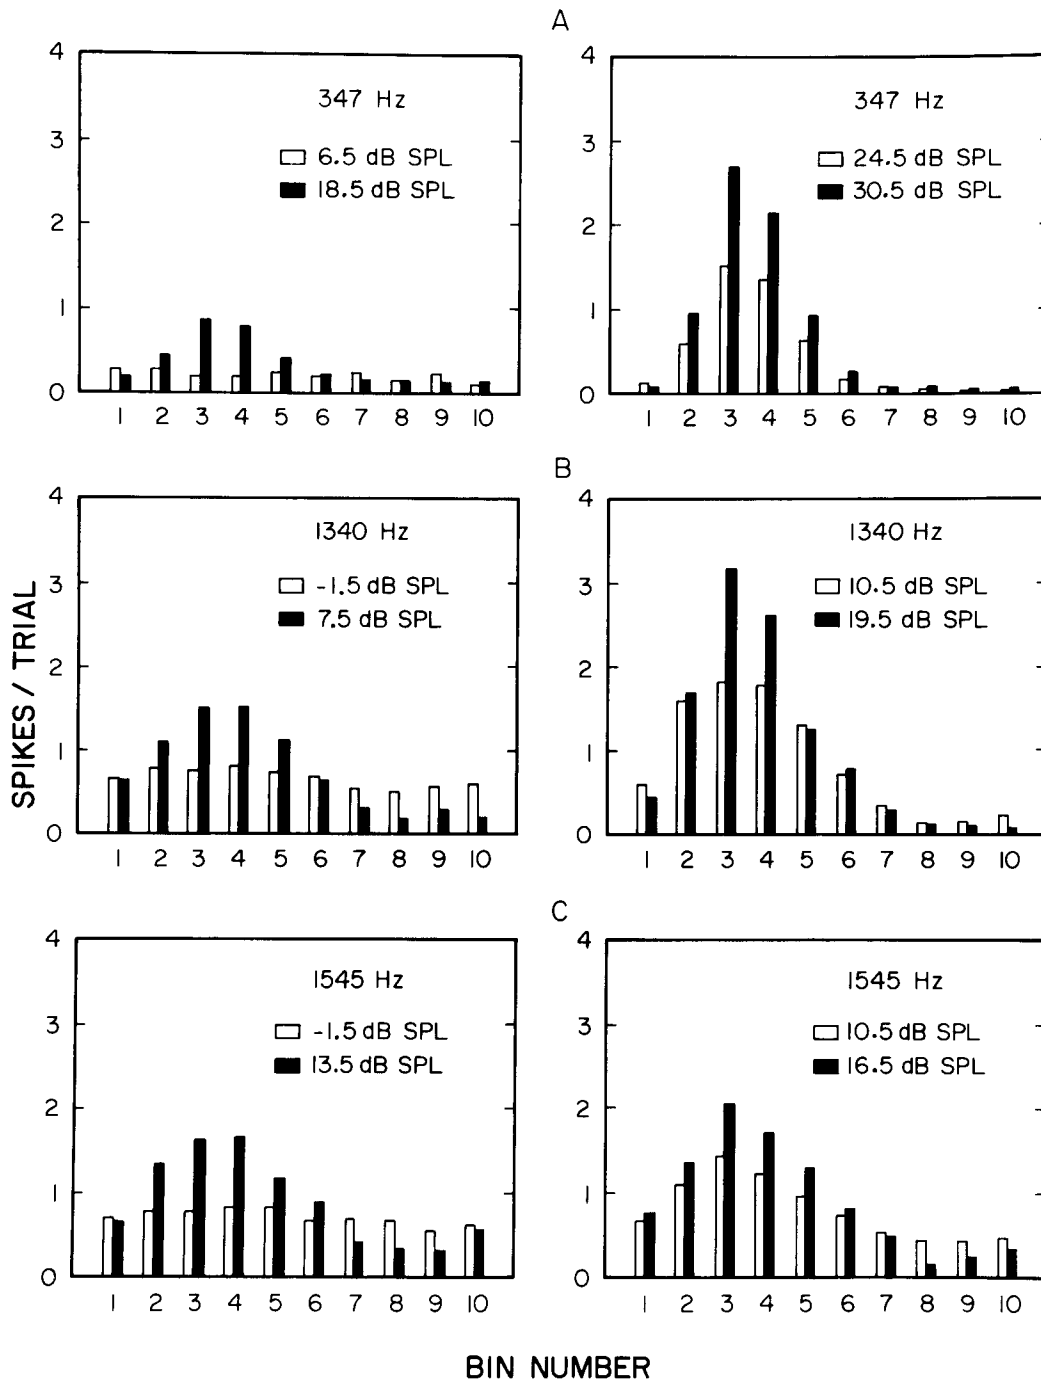


FIG. 6. Neurometric functions (proportion correct as a function of total intensity in dB SPL) of 3 auditory-nerve neurons in an intensity-discrimination task for 3 levels of ideal-observer analysis. *Left*: near-threshold reference intensities. *Right*: near-saturation reference intensities. Trial durations were ~ 70 ms. Frequency indicated in each panel was both the CF of the neuron and the frequency of the stimulus. Results show that temporal response pattern contributes substantial discrimination information, especially at lower intensities (*left*).