

Temporal Encoding in Nervous Systems: A Rigorous Definition

FRÉDÉRIC THEUNISSEN

Department of Molecular and Cell Biology, 195 LSA, University of California, Berkeley, CA 94720
ft@cicada.berkeley.edu

JOHN P. MILLER

Department of Molecular and Cell Biology, 195 LSA, University of California, Berkeley, CA 94720
jpm@cicada.berkeley.edu

Received November 9, 1994; Revised February 10, 1995; Accepted February 16, 1995

Action Editor: John Rinzel

Abstract. We propose a rigorous definition for the term *temporal encoding* as it is applied to schemes for the representation of information within *patterns* of neuronal action potentials, and distinguish temporal encoding schemes from those based on window-averaged *mean rate encoding*. The definition relies on the identification of an *encoding time window*, defined as the duration of a neuron's spike train assumed to correspond to a single symbol in the neural code. The duration of the encoding time window is dictated by the time scale of the information being encoded. We distinguish between the concepts of the *encoding time window* and the *integration time window*, the latter of which is defined as the duration of a stimulus signal that affects the response of the neuron. We note that the duration of the encoding and integration windows might be significantly different. We also present objective, experimentally assessable criteria for identifying neurons and neuronal ensembles that utilize temporal encoding to any significant extent. The definitions and criteria are made rigorous within the contexts of several commonly used analytical approaches, including the *stimulus reconstruction* analysis technique. Several examples are presented to illustrate the distinctions between and relative capabilities of rate encoding and temporal encoding schemes. We also distinguish our usage of *temporal encoding* from the term *temporal coding*, which is commonly used in reference to the representation of information about the *timing* of events by rate encoding schemes.

Keywords: neural code, rate encoding, temporal encoding, temporal coding

Introduction: Definition of the Encoding Problem

An animal's continually evolving perception of its surrounding environment, its awareness of its own internally regulated homeostatic balance within its environment, and its behavioral responses to dynamic sensory stimuli must ultimately be derived from information contained within relatively brief segments of neuronal spike trains. The computations underlying all aspects of the operation of its nervous system are carried out within the context of the neural code with which the relevant information is represented in those spike trains. A determination of the information coding schemes used within nervous systems is an extremely important goal, due not only to the intrinsic interest in the nature of the neural code itself but due also to the very valuable and important constraints a knowledge of the code can place on the development

of physiological models for the mechanisms underlying neural computation.

Deciphering the neural code at any particular location within a neural system can be reduced to three interconnected tasks, representing a quantitative characterization of the observed stimulus-response characteristics of the neurons under study. The first task is to determine the quantity and qualitative nature (such as the specific parameters characterizing a complex sensory stimulus) of the information encoded in the spike trains of the neuron or neuronal ensemble under study. The second task is to determine the nature of the neural symbols with which that information is encoded—that is, is all of the information encoded in the *mean firing rates* of the cells, or is some significant proportion of the information encoded in more complex statistical features of the spike train *patterns*? The third task is to define relevant, objective measures of significance

with which the information and the associated neural symbols are correlated.

For the purposes of our discussion, we assume that the first task has been accomplished by some means and that we have explicit knowledge of the nature of the significant information to be encoded in the neural activity patterns. Specifically, for the sake of simplicity, we cast our discussions in terms of sensory physiology and consider the encoded information to be the values of one of the parameters describing a sensory stimulus that is known to affect the activity of the neuron or ensemble of neurons under study. We deal exclusively with problems associated with the second task listed above. This second task can be termed the *encoding problem* to distinguish it from the larger coding problem encompassed by all three tasks. We address the encoding problem for the case of spiking neurons, in which the neural symbols can be reduced to binary strings in time, where 0s and 1s can be used to mark the absence or occurrences of spikes.

Although these binary strings would seem to be relatively simple objects on which to base any analysis of neural coding, the problem is actually quite complex due to the possibility that the code might be of high dimensionality. The problem is as follows. Any analysis of neural encoding or decoding schemes forces a determination (or assumption) of the associated *encoding time window*—that is, the duration of the spike train assumed to correspond to a single symbol in the code. The longer the encoding time window, the greater is the number of digits in the binary string corresponding to a single symbol of the code, and the greater is the number of possible symbols in the code (that is, the number of possible patterns of spikes within the time window). Considering the fact that action potentials are approximately 1 millisecond in duration, a rough estimate of the number of bits in a binary string corresponding to a single symbol of neural code would be on the order of the duration of the encoding time window measured in msec. For a 100 msec time window, the potential information imbedded in a binary string is on the order of 100 bits, corresponding to the 2^{100} possible symbols. If the possibility of ensemble encoding of information across multiple-cell ensembles is allowed, then the theoretically limiting number of possible symbols in the code strings would grow to even more astronomical numbers.

This large dimensionality of the encoding problem poses two significant problems. First, from a practical standpoint, the task of experimentally characterizing the statistics of occurrence of each of the

different symbols becomes increasingly difficult with longer encoding time windows. Second, from a conceptual standpoint, determining the significance of the analysis becomes increasingly problematic, since (1) many of the different neural symbols might conceivably correspond to the same input information (that is, some differences in the different spike response patterns might be due to some type of noise intrinsic to the system rather than to an actual variation in the input), (2) some of the different neural symbols might occur so infrequently as to be experimentally uncharacterizable, and (3) the postsynaptic neuronal decoder circuit might not be capable of extracting all of the information that *is* available in the transmitted symbols. In other words, even though a 100 msec encoding window might theoretically allow the use of a binary neural “alphabet” having 2^{100} symbols, the neural decoder (cell or ensemble) might lump together huge subsets of these different possible spike patterns as being indistinguishable, reducing the effective number of symbols used for information representation to a small fraction of the theoretical limiting number.

A determination of the neural code therefore requires the reduction of the binary strings representing spike occurrence times to the set of biologically significant symbols within the particular context being investigated. The *neural code* with which information is represented within the nervous system can be defined objectively: the neural code is the minimum set of symbols capable of representing all of the biologically significant information. Thus, in order to define the neural encoding scheme implemented in any particular situation, a measure must be derived that corresponds to the correlation between an observed set of neural activity patterns and the encoded information (such as those sensory stimulus parameters) represented by those activity patterns. One measure of correlation that has been shown to be of particular utility within the context of the neural encoding problem is the measure of *transinformation* or *mutual information*, as defined by Shannon in his development of information theory (Shannon, 1948; Pierce, 1961). The mutual information is the most precise measure of correlation as it is based on the complete form of the *joint probability* between the symbols representing the biologically significant information and the symbols representing the neural code (Eckhorn and Pöpel, 1974; Theunissen, 1993). Within an experimental electrophysiological context, where a set of stimulus-response measurements can be recorded, the joint probability of occurrence of individual neural responses within the

data set and the set of electrical or sensory stimuli that were used to evoke those responses can be characterized and used to compute the mutual information. The minimum set of neural symbols can be determined, at least in theory, by systematically reducing the complete set of binary strings representing the most complete description of the neural responses to smaller and simpler sets composed of the putative neural symbols, until the measure of mutual information begins to decrease.

Many encoding schemes have been considered by neurobiologists, and all schemes can be classified according to the degree to which the dimensionality of the spike-train response data is reduced for the analysis (for an early comprehensive review, see Perkel and Bullock, 1968). At one extreme, an analysis allowing no reduction at all would necessitate the consideration of the full binary code, where the precise temporal placement of every spike within the encoding time window is considered to be capable of conveying significant information. At the other extreme, a coding scheme is imaginable in which a full reduction of the complexity of the binary string within the encoding time window to a single bit, where the presence of one or more spikes in any temporal pattern whatsoever within the encoding time window would convey the significant information.

In many cases, the reduction from the full binary code to a code of intermediate complexity is accomplished by counting the number of spikes within the assumed encoding time window. Any scheme based on such a reduction is generally referred to as a *rate code*. Beginning with the earliest studies of neural encoding, rate codes have been shown to encode significant amounts of information in many different experimental preparations (see, for example, Adrian, 1928; Werner and Mountcastle, 1963; Perkel and Bullock, 1968) and are generally considered as the best first-order assumption for the neural encoding scheme. However, recent studies have shown that aspects of the fine temporal structure or *patterns* of spike trains within the time scale of the encoding time window can also carry significant information about the stimulus (Gray et al., 1989; Richmond et al., 1987; McClurkin et al., 1991; Engel et al., 1992; Middlebrooks et al., 1994; Kjaer et al., 1994). Such a code is generally referred to as a *temporal code*.

One might think that a temporal encoding scheme could be considered as being equivalent to a special case of a rate encoding scheme, in the limit where the encoding time window is of short enough duration to take into account the fine temporal structure in the observed spike train responses. This would seem

to make the distinction between rate encoding and temporal encoding somewhat arbitrary. We refute that idea and show that the distinctions between rate encoding and temporal encoding can be defined clearly and rigorously. Moreover, the distinctions are of significant biological relevance, since (1) the distinctions emerge from a consideration of the intrinsic time scale of the dynamics of the encoded information, and (2) the extremely different nature of these two encoding schemes would have significant implications for biologically plausible *decoding* mechanisms. We also show that the definition has a strict mathematical correlate when the encoding-decoding task is expressed in terms of stochastic systems analysis and in particular within the context of the *stimulus reconstruction* methodology recently derived by Bialek and his colleagues (Bialek et al., 1991; Bialek and Rieke, 1992).

The Encoding Time Window

Determining the Characteristic Time Scale of a Coding Process

Any definitions of temporal encoding and rate encoding must involve some consideration of the nature and dynamics of the encoded information. Specifically, the *duration of the encoding window* used for any analysis of the neural code is not arbitrary but depends on the dynamical nature of the information being encoded. For stimuli that are approximately stationary (such as those for which the time scale of the variation in a stimulus parameter is much longer than the time scale of the dynamics of the animal's associated computational task or behavioral response), the duration of the encoding window can be arbitrarily long, and may include the entire neuronal response to a prolonged stimulus. Examples of relatively stationary stimulus encoding tasks might include the coding of the shape of stationary objects in the visual field in the primate visual system (Richmond et al., 1987) and the coding of the location of a continuous sound source in space (Middlebrooks et al., 1994; Moiseff and Konishi, 1981). The only constraints influencing the duration of the encoding window arise, presumably, from the evolutionary forces acting in opposition to one another to (1) *increase acuity* (by increasing the encoding time, allowing a longer period over which to "signal average") and (2) *decrease reaction time* (by decreasing the encoding time, allowing a shorter delay between the stimulus onset and the computed response). Indeed, one practical approach

toward estimating the duration of the encoding window for stationary stimuli has been to measure the behavioral reaction time following presentation of test stimuli.

However, for many encoding tasks, the time scale of the variation in a relevant stimulus parameter is on the order of (or shorter than) the time within which decisions or behavioral responses must be made by the animal. Examples from sensory physiology include (1) the tracking of a constantly moving visual stimulus and (2) the analysis and identification of spectrally complex sound sources. For such tasks, the duration of the presumed encoding window cannot be arbitrarily long, but must be limited by the rate at which the relevant stimulus parameter is changing. In other words, dynamic variations in the stimulus signal can never be encoded faster than the rate at which the neural code can be updated to represent those variations. This rate is itself limited, presumably, by the intrinsic time scale for the neural computations underlying the *encoding task*, and can be used as another means by which to define the encoding time window: the encoding window is the limiting intrinsic period of the encoding process, corresponding to the inverse of the limiting frequency with which the neural code is updated to represent dynamic variation in the stimulus signal.

For example, in a hypothetical system in which a saccadic eye movement presents a “new” stimulus to the visual system every 300 msec, the upper limit on the duration of the appropriate time window for the encoding of some aspect of the stimuli is equal to 300 msec. The actual duration of the encoding window might be shorter, if behavioral or psychophysical experiments suggest a reaction time less than 300 msec. That is, the limiting duration of the encoding window is uniquely determined by the dynamical characteristics of the relevant aspect of the stimulus being encoded.

For tasks involving continuously varying stimuli, constraints on the encoding time window can be understood from a consideration of a Fourier decomposition of the stimulus signal. For example, if one of the tasks of a particular auditory neuron were to encode continuously the amplitude and phase of the 100 Hz component of the signal, then the duration of the appropriate encoding window for that task would be 5 msec. This is essentially a restatement of the Nyquist theorem: in order to reconstruct the waveform of the 100 Hz component, a decoder would need a minimum of two sampling points per 10 msec cycle period. A simple form of rate decoding of the phase and amplitude of the 100 Hz stimulus component estimated within a specific time

interval of 10 msec centered at a particular time t could be achieved by the following procedure: (1) count the number of spikes within each of the two sequential 5 ms time windows subdividing the 10 msec time interval of interest, (2) scale the sum of those spike counts with the appropriate (experimentally derived) factor to determine the amplitude of the 100 Hz component of the signal within that particular 10 msec time interval, and (3) ratio the spike counts within each of the two sequential 5 ms encoding windows to determine the phase of the 100 Hz component with respect to time t . It is clear from this example that implementation of an encoding time window having a duration longer than 5 msec would prevent continuous representation of the encoded information, since the duration of the window would prevent the minimum number of sampling points (that is, one per sampling window) needed to decode the spike train in order to reconstruct the amplitude and phase of the 100 Hz component.¹

Encoding Window and Integration Window

It is important to distinguish between the concepts of the *encoding time window* and the *integration time window*, which are often confused and used interchangeably. In particular, the term *integration time window* is often used in place of *encoding time window* in analyses of rate encoding schemes, where the spikes within the encoding window are summed to obtain the estimate of a stimulus parameter value. In general, however, the integration window refers to the net duration of all physiological processes that contribute to the integration and transformation of synaptic inputs (and any intrinsically generated conductances) into the spiking output of a neuron. One practical definition of the *integration window* is the period of time preceding any particular timepoint in a neuron’s (or ensemble’s) response pattern during which a variation in the input could have significantly affected the response at that timepoint. Although the durations of the encoding and integration windows might be *equal* in some cases, it is also possible that the integration time is of longer duration than the encoding time window.

To illustrate, consider an auditory neuron that encodes the waveform of a 100 Hz signal component, whose activity is affected by *multiple cycles* of the 100 Hz component. In such cases, the reverse correlation filter, obtained by determining the average stimulus waveform preceding each elicited spike, would show many cycles of a sinusoidal waveform at the frequency to which the cell is maximally tuned. This is

typically the case for auditory neurons with very sharp frequency tuning: see Eggermont (1993) for examples. The duration of this reverse correlation filter would correspond to the neuron's *integration window*, as defined above. However, to be able to continuously represent the amplitude and phase of that 100 Hz component, a new sample point would have to be obtained from the spike train every 5 msec, which is the duration of the neuron's *encoding window*.

Although the integration time window of a neuron (or ensemble) might be of longer duration than its encoding window, it could *not* be of short duration than the encoding time window. This is because the integration time can, in a sense, be taken as a measure of the limiting electrophysiological reaction time—that is, the upper bound on the time limit over which a change in the stimulus can be seen to have significant correlation with the response. Since the duration of the encoding time window is set by the shorter of the reaction time and the dynamic time course of the stimulus, the encoding time window must be shorter than or equal in duration to the integration time window.

Determining the Characteristic Time Scale of a Signal Using the Stimulus Reconstruction Method

In many encoding situations, the relevant stimuli are not stationary, nor can they be characterized as having only a single frequency component (nor, therefore, a single characteristic time scale). Rather, a stimulus may be composed of many frequency components. Such complex dynamic stimuli are referred to as *signals* in the engineering literature, and a neural system encoding the signals would be referred to be as a *dynamical encoder*. An auditory neuron with a broad frequency tuning curve is an example of a dynamical encoder. Since each of the different frequency components must be encoded simultaneously within such a neuron's spike-train response, the notion of a single encoding time window becomes inadequate. Each frequency component requires a corresponding encoding window having a duration equal to half the period of that frequency component. Each spike in a given spike train will fall within the encoding windows associated with all the different frequency components and could in fact be correlated with multiple frequency components of the signal.

The methods of stochastic systems analysis are well suited for the analysis and quantification of such complex encoding scenarios. The goal of stochastic systems analysis is to characterize the deterministic

relationship between two stochastic signals (in our case, the sensory stimulus signal and the neural response pattern), thereby separating the aspects of the two signal variables that are correlated from those that are random. The deterministic relationship can then be expressed in the frequency domain, allowing the encoding of the different frequency components to be "untangled" from one another. In a stochastic systems analysis, the deterministic relationship between the stimulus signal and the response pattern is expressed as a transformation that operates on one of the random variables to generate a prediction of the other one. Such a transformation can be expressed as a functional expansion, similar to the way a Taylor series expansion (using successively higher-order polynomial terms) can be used to approximate a complex function. In a functional expansion, the terms are called *kernels* or *filters* and correspond to the successively higher-order nonlinear terms needed to achieve an adequate approximation of the transformation. The kernels can be derived analytically from a statistical analysis of the data, by requiring that the difference between the predicted and the actual random variable are minimal with respect to some set of objectively defined criteria. Such analyses are known as *Volterra* or *Wiener analyses* (depending on the precise form of the functional expansion) or *white-noise analyses* (after the most frequently used stimuli) and have been applied extensively in sensory physiology (Marmarelis and Marmarelis, 1978; Korenberg and Hunter, 1990; Eggermont, 1993).

The conventional approach in applying this type of analysis to neurobiological systems has been to derive a transformation that operates on the stimulus signal to obtain a best estimate of the spike-train response. In the frequency domain, such a forward transformation describes how any single frequency component of the spike-train response can be predicted from combinations of multiple-frequency components of the signal. The natural inclination is to interpret such a forward transformation as the operation that transforms the raw stimulus signal into the significant information about the signal. That information would then be represented by simply equating the amplitude of the transformed signal to the probability of obtaining a spike. It is important to note that this interpretation forces the assumption that all significant information is encoded with a rate code.

If the nature of the encoding scheme has not yet been determined and the analysis of the code is an explicit goal of the study, then the reverse approach must be taken. That is, an expansion must be derived that

transforms the observed spike train into an optimal estimate of the stimulus that elicited that spike train. This approach allows an analysis of the correlation of different frequency components of the spike train with individual frequency components of the signal (each of which has its own characteristic time scale). To do so, an expansion is derived in the form of a set of filters that, when convolved with the observed spike train, reconstructs an optimal estimate of the stimulus signal.

This innovative approach, referred to as the *reverse reconstruction* or *stimulus reconstruction method*, was defined and first applied for an analysis of sensory coding by Bialek and his colleagues (Bialek et al., 1991; Bialek and Rieke, 1992).² The stimulus reconstruction method allows us to analyze the encoding of each of the frequency components of the signal independently from all other frequency components and resolves the apparent problems presented by the necessity for considering multiple overlapping encoding time windows. In essence, the reconstruction filter is decomposable into its constituent frequency components (see Appendix). Each frequency component will have an associated encoding time window equal to one-half of the associated cycle period. A study of the reconstruction filters in such an analysis thus provides a wealth of quantitative, model-independent information about the intrinsic time scales of neural computation and encoding.

Rate Encoding and Temporal Encoding

Rate Encoding

With the duration of the encoding window defined in terms of the intrinsic time scale of the biologically relevant encoded information, we can distinguish objectively between rate encoding and temporal encoding schemes. It is generally accepted that a *rate encoding scheme* is one in which the relevant information encoded about the stimulus is correlated only with the number of elicited spikes within the encoding window and is *not* correlated with any aspect of the temporal pattern of the spikes within the encoding window. In our definitions, we use the term *correlation* in the loose sense to signify any probabilistic dependence between the two random variables. To be rigorous, this statistical dependence must be tested with a measure that takes into account all the moments the joint probability distribution between the stimulus and the response (such as the mutual information).

The general definition of rate encoding can be recast more specifically in terms of several different analytical approaches that are commonly used for the analysis of the input-output properties of a nerve cell or ensemble, as follows:

1. For neurons encoding stimuli with a single time scale, the definition is this: a rate encoding scheme is one in which there is significant correlation between the relevant stimulus parameter and the mean number of spikes in the elicited response within the encoding window (or between the stimulus parameter and some weighted average derived, for example, from the value of the first principal component of the pattern of spikes) and no additional correlation between the stimulus parameter and any higher-order moments (or higher-order principal components) of the elicited spike pattern within the encoding window.
2. For tasks in which the stimuli are dynamic in time, this definition can be formulated in the frequency domain: a rate encoding scheme is one in which there is significant correlation between identical frequency components of the stimulus signal and elicited spike pattern, and no additional independent correlation between any frequency component of signal with any other higher-frequency components of the elicited spike pattern. Correlations with lower-frequency components are allowed, since these can be determined by averaging the values from multiple encoding windows corresponding to the frequency being encoded.
3. For the analysis of complex signals within the context of the *stimulus reconstruction technique* definition 2 translates to the following: a rate encoding scheme is one in which the transformation used to obtain the best estimate of a given frequency component of a signal is composed of analytical terms that involve only the same (or lower) frequency components of the spike-train response pattern. In other words, in a rate-encoding scheme, no significant improvement could be achieved in the estimation of the stimulus signal at a given frequency by adding expansion terms involving components of the response pattern at higher frequencies. In particular, if all relevant information about a dynamic stimulus signal can be recovered through the application of a first-order (that is, linear) filter term in the reconstruction operation, then the code can automatically be classified as a rate code. In fact, such a code can be classified even more restrictively as a *linear rate code*.

As shown by Abbott (1994), a linear filter obtained via application of the stimulus reconstruction method can be normalized to obtain an estimate of the mean firing rate. Within the context of our definition, this mean firing rate is obtained by a convolution of the spike train with an analytically derived filter that takes into account the time frames of all the frequency components of the significant information being encoded. If this mean firing rate can be linearly scaled to match the stimulus (as it is done in the stimulus reconstruction methodology), and the resulting approximation is the best that can be achieved from the information imbedded in the spike train, the code can be classified as a linear rate code.

Note that our variously worded definitions of rate encoding do not necessarily require the code to be *linear*. The function characterizing the correlation between the information being encoded and the number of spikes within the window can represent a highly nonlinear transformation of the number of spikes. Within the context of the stimulus reconstruction technique, we note also that encoding schemes could exist in which the best stimulus reconstruction filter would include nonlinear terms in the same frequency. Such a scheme would be a *nonlinear rate encoding scheme*. As is discussed below, the encoding scheme is disqualified as being a rate code *only if* particular frequency components of the stimulus signal are correlated with higher-frequency components of the response pattern *and* that this correlation yields additional information beyond what can be obtained from the correlation with the frequency component of the response pattern at the frequency of the stimulus signal.

Another important point is that in a rate code, the amount of information conveyed by a particular response is not necessarily proportional to the number of spikes elicited within the encoding window. Indeed, responses that consist of decreases in the spike activity below the baseline rate can convey considerable amounts of information. Rather, as shown by our work and by that of others (Theunissen and Miller, 1991; Eckhorn and Pöpel, 1975), the mutual information conveyed by a particular symbol (in this case, the spike count within an encoding window) depends on the extent to which that symbol is distinguishable from all other symbols that might have occurred.

Temporal Encoding

In a *temporal encoding scheme*, the relevant information is correlated with the timing of the spikes within the encoding window, over and above any information that

might be correlated with the number of spikes within the window. Specifically, a temporal encoding scheme is one in which there is significant additional correlation between a stimulus parameter and any aspect of the pattern of elicited spikes on a time scale less than the duration of the encoding window. As was the case for our definition of rate encoding, this definition can be recast more specifically in terms of several analytic approaches that are commonly used for the analysis of the input-output properties of a nerve cell or ensemble, as follows:

1. For neurons encoding stimuli with a single time scale: a temporal encoding scheme is one in which there is significant additional correlation between the relevant stimulus parameter and any moments of the elicited spike pattern having *higher order than the mean* (or having any principal components beyond the first).
2. In the frequency domain: a temporal encoding scheme is one in which there is significant additional correlation between a frequency component of a dynamic stimulus signal and a *higher-frequency component* of the corresponding elicited spike pattern.
3. Within the context of the stimulus reconstruction technique: a temporal encoding scheme is one in which the transformation used to obtain the best estimate of a given frequency component of a signal is composed of analytical terms that involve higher-order nonlinear filters involving *higher-frequency components* of the spike-train response pattern. (See the Appendix for the mathematical formulation of this definition.)

Some Examples and Common Misconceptions

Weighted-Average Rate Codes. In our scheme for the classification of rate encoding versus temporal encoding schemes presented above, we have included a *weighted-average-rate code* as a particular example of a rate-encoding scheme. Although a weighted-average code is ultimately based on a count of the number of spikes occurring within the encoding time window, spikes within different subintervals of the encoding window might be weighted differently in the final estimation of the stimulus parameter value. For example, as indicated in our definitions, if all the information about the stimulus can be extracted by considering only the coefficients of the first component in a principal component decomposition of the response,

the code would be classified as a rate code. That principal component might have a much higher value near the beginning of the encoding window than at the end of the window—for example, indicating that a differential weighting of the early spikes within the encoding window would allow a better estimate of the stimulus signal than would an unweighted or “flat” count of the spikes over the whole encoding window.

This generalization of rate encoding to include weighted rate codes is consistent with the fundamental distinction between temporal and rate codes. For an encoding scheme to qualify as a temporal scheme, at least two different weighted spike-count average values would have to be obtained from statistically independent aspects of the spike-train pattern within a single encoding window in order to extract all of the information imbedded in that spike train about the stimulus. The information from these two different measures of the spike-train pattern could then, presumably, be combined together through some *nonlinear* operation to obtain a single value describing the deterministic relationship between the stimulus and the response that was better than either of the weighted-average values alone. If an encoding scheme is based on a derivation of one encoding value per encoding window through a *linear operation on the spike train*, as in a weighted-average code, then it classifies as a rate code. Researchers testing for temporal encoding should therefore compare the accuracy of their presumed temporal encoding scheme (such as the degree of correlation between their decoded response and the actual stimulus signal) with the accuracy attainable using the optimal weighed-average rate code.

From a practical standpoint, the consideration of weighted rate codes allows for an objective determination of the duration of the encoding time window. In particular, when considering the encoding of stationary or pseudostationary stimuli, the duration of the encoding time window can be considered as being equivalent to the integration time of the neural layer being considered. Note that the use of an *unweighted* (that is, flat) rate code could lead to severe problems if the length of the integration window were not known and if an inappropriately long encoding window were chosen for an analysis of neural encoding. In this case, erroneous conclusions could be drawn about the nature of encoding. For example, errors due to spikes counted in the later part of the chosen window (that is, outside of the real encoding window) could significantly degrade information obtained from the spikes counted in the early part of the chosen window (corresponding to

the real encoding window). In that situation, the encoding scheme could be misclassified as being temporal when, in fact, limiting the spike count to the actual encoding time window would result in the correct classification of the scheme as rate encoding. The derivation of weighted average rate codes prevents this kind of mistaken interpretation, even if the duration of the encoding time window is erroneously chosen to be too long. Therefore, this general definition of rate encoding is of practical as well as of theoretical merit.

Temporal Encoding Versus Rate Encoding?

It is important to note that the classification of an encoding scheme as either a rate or a temporal scheme *cannot* be based solely on an observation of the frequency or timing precision of spikes within the response patterns. The fact that a set of spike patterns elicited by repeated presentations of a complex stimulus might show extremely low intertrial variability (for example, many spike placements might be reproducible from trial to trial with little “jitter” on a time scale longer than 1 msec) might only reflect the encoding of the phase of high-frequency components of a signal. On the other hand, a temporal code could be implemented that required relatively low precision in spike timing, as long as different discriminable patterns *could* be produced within the time scale of the encoding window. Thus, one might expect a greater probability of finding temporal encoding in situations where the duration of the encoding window for the relevant information is long relative to the action potential refractory period (or to the mean interspike interval). In such cases, one might even expect to find temporal encoding employed at the level of single neurons. For neurons having encoding time windows with durations approaching the refractory period, the imposed restriction on the minimal interspike interval would prevent the realization of many different spike patterns within the time scale of the encoding window.

Experimental results suggesting the involvement of temporal encoding in single neurons have, in fact, been reported. Examples in which single neurons temporally encode significant amounts of information about relatively *stationary* stimuli can be found in (1) cutaneous receptors encoding location on the skin of the stimulus (Fuller and Looft, 1984), (2) neurons in primate visual system pathways encoding complex visual patterns (Richmond et al., 1987; Richmond et al., 1990; McClurkin et al., 1991; Kjaer et al., 1994; but see also

Tovee et al., 1993), and (3) interneurons in the auditory cortex of cats encoding sound source localization (Middlebrooks et al., 1994). In each of these cases, the same aspect of the stimulus was encoded by the restricted rate code *and* the temporal code. Thus, consideration of the temporal code in these cases allowed an increase in the coding accuracy.

To our knowledge, however, no definitive demonstrations have yet been presented of single neurons encoding *dynamic* signals using temporal encoding schemes. The stimulus reconstruction method has now been applied to the analysis of encoding in several different sensory systems, but in most cases, the linear filter term recovered all significant mutual information, implying the operation of a rate encoding schemes (Bialek et al., 1991; Bialek and Rieke, 1992; Warland et al., 1991). Only in one case reported to date was significant additional mutual information recovered with a second-order (nonlinear) filter term (Rieke, 1991). However, the amount of additional information recovered with the second-order filter was relatively small, and it was not determined in that study whether or not the second-order term involved a correlation between any signal component with higher-frequency components of the spike-train response pattern.

Ensemble Temporal Encoding

When information is encoded by an ensemble of cells, however, the possibilities for temporal encoding become much greater. This is because the number of different discrete spike patterns expressible within the encoding time window is not restricted to the same extent *across an ensemble* as they are in any individual neuron making up the ensemble. To illustrate, if the encoding window for a particular class of neuron were 2 msec, the duration of an action potential 1 msec, and the refractory period 1 msec, then the number of spikes per encoding window would never exceed one for any individual cell. Thus, the number of patterns expressible by one spike per encoding window per cell would be severely restricted. However, the number of patterns expressible across an ensemble of such cells might be substantially greater, since different relative placements of the cells' spikes *within* their respective encoding windows could theoretically be used much more reliably to encode information about the variation of dynamic signal parameters. It is important to note, however, that the mere existence of ensemble coding does not necessarily imply the operation of a temporal encoding scheme. If the information about amplitude

or phase is correlated only with the total spike count within the net encoding window across the ensemble, then the operation of a rate-encoding scheme is implied.

One interpretation of several recent studies of ensemble encoding in the mammalian visual system is that (1) there is a significant amount of information encoded in temporal spike patterns across a neural ensemble and (2) the temporally encoded information is of a completely different nature than the information that is encoded by the rate of firing of the neural ensemble. One hypothesis is that information about the binding of the different components of a visual object is encoded by aspects of the synchronization of neuronal responses across neurons (Engel et al., 1992). For neurons in the cat visual cortex, the synchronization of spike firing has a time scale on the order of 20 msec, whereas the time scale of the encoding window in those tasks is on the order of 100 msec. In this case, essentially no information about linkage could be recoverable from the ensemble firing rate (averaged over the correct 100 msec window) of the neurons.

Temporal encoding across neuronal ensembles has also been demonstrated by Abeles et al. (1993), who found very precise patterns (with time scales on the order of a few msecs) across small ensembles of neurons in the frontal cortical areas of the primate brain. The patterns were shown to correlate with specific behavioral information-processing situations, even though the exact nature the information carried by the patterns could not be determined. As in the case for the synchronous firing of neurons of the primary visual cortex, the temporal patterns in these neurons might encode the results of some higher-order computation occurring in the brain.

Considering even larger ensembles of neurons, Freeman has shown that the complex dynamical response pattern of the ensemble of neurons in the olfactory bulb of the rabbit (represented by a multi-dimensional EEG recording) were correlated with the perception of different odorants (Freeman, 1991). Similarly, complex dynamical patterns have recently been found by recording simultaneously from many neurons in the olfactory system of locusts (Laurent and Davidowitz, 1994). In both of these cases, it still remains to be determined whether or not any additional information can be obtained from the consideration of the spatiotemporal response patterns, beyond the information that would be recoverable from the mean activity of all the neurons in consideration and, if so, what the nature of that information might be.

Temporal Coding Versus Temporal Encoding

It is important to clarify the difference between our definition of temporal *encoding* and the term *temporal coding*, which is often used in the literature to describe the encoding of information about the temporal aspects of dynamic stimuli (for example, contrast the use of *temporal coding* in Engel et al., 1992, with the use in Gooler and Feng, 1992). The time course of information processing in any sensory part of the nervous system is necessarily coupled to the time course of external physical stimuli. In other words, the times of occurrence of relevant stimulus features are often of considerable behavioral relevance to the animal and are almost always encoded within spike-train responses. This encoding of temporal aspects of a stimulus signal is commonly referred to as *temporal coding* and could theoretically be implemented through either a rate-encoding scheme *or* a temporal-encoding scheme.

Many examples of temporal coding have been presented in the literature, for both static (or slowly changing) stimuli and for rapidly changing dynamic signals. Within the context of auditory physiology, the extent of temporal coding is reflected by the precision with which the neural response is *phase locked* to the time of occurrence of relevant stimulus features. The ability of neurons to phase lock to a stimulus signal is exemplified by high-frequency auditory fibers: single neurons can phase lock to stimulus frequencies of up to 7 kHz (Moiseff and Konishi, 1981) (although no single neuron could fire at every cycle). Another remarkable example of phase-locking neurons is found in bats, where the variance in the latency of spiking neurons to the onset of the stimulus is in the submillisecond range (Covey and Casseday, 1991). In those cases, the timing of action potentials is modulated with extreme precision, on a time scale much smaller than the duration of an action potential.

However, it is important to note that the representation of such temporal information, even where such precise timing of spikes can be observed, does not necessarily require a temporal encoding scheme but could be achieved through rate encoding.³ Conversely, the implementation of a temporal encoding scheme does not necessarily require high temporal resolution in a decoder's sensitivity to the exact timing of spikes within the encoding window. Temporal encoding and temporal coding can both produce distinct and reliable spike patterns in time, but the nature of the patterns as it relates to the encoded information is fundamentally different. It is important to remember, however, that a

scheme could be imagined in which additional information about the phase of the signal were temporally encoded. This would be the case if differences in the precise timing of the spikes within the encoding window were correlated with differences in the phase of the signal. Considering the confusing nature of the terminology, we propose the strict use of *temporal coding* when referring to the encoding of temporal aspects of the significant information, such as the phase of a frequency component of a signal, and the use of *temporal encoding* when referring to schemes for encoding information in which spike patterns on a time scale less than the duration of the encoding window correlate significantly with stimulus patterns.

Conclusions

As we have shown, the distinctions between rate encoding and temporal encoding are far from arbitrary and can be defined rigorously by considering the natural time scale of the information being represented. The distinctions are significant for two important reasons. First, from a practical standpoint, the extent by which spike-train data can be statistically reduced for subsequent analysis may be very different for these two different schemes. Any time-window averaging of spike trains containing temporally encoded information would essentially filter out and eliminate some or all of the temporally encoded information, and analysis of such degraded data would lead to miscalculations of the quantity or quality of information represented in neuronal spike trains.

Second, temporal encoding schemes are significantly different than rate-encoding schemes from a conceptual standpoint. In temporal encoding schemes, the dimension of time is used by a neuron or neuronal ensemble to encode information that might not be temporal in nature (that is, not related to either the *timing* of events or the *temporal dynamics* of continuously varying functions). This transformation of the encoding regime presumably requires innovative architecture and sophisticated biophysical encoding mechanisms. On the other hand, the nervous system is an excellent processing machine for such a transformation of encoding domains, since spatial patterns of activity can easily be transformed into temporal patterns and vice versa. Another fascinating possibility is the potential for information that *is* temporal in nature to be temporally encoded. But this, again, requires a significant transformation of the encoding regime. In this case, the

Table 1.

Time windows		
	Integration window	Encoding window
Definition	The period of time preceding any particular point in a neuron's or ensemble's activity pattern during which a variation in the input could have significantly affected that activity	The duration of the spike train assumed to correspond to a single symbol of the code; also the limiting intrinsic period of the encoding process, corresponding to the inverse of the limiting frequency with which the neural code is updated to represent dynamic variation in the stimulus signal
Depends on	All physiological processes that contribute to the integration and transformation of synaptic inputs (and any intrinsically generated conductances) into the spiking output of a neuron	The time scale of the stimuli being encoded; and also on the reaction time of the system. Must be less than or equal to the duration of the integration window.

Encoding schemes		
	Rate encoding	Temporal encoding
Definition	The encoded information is correlated only with the average number of spikes within the encoding time window, or with any weighted average across that window	Additional information is correlated with some aspect(s) of the temporal pattern of spikes within the encoding window

Temporal coding <i>versus</i> Temporal encoding		
	Temporal coding	Temporal encoding
Definition	An implementation of a rate or temporal encoding scheme to represent a continuously varying sensory signal, characterized by a one-to-one correspondence between the time of occurrence of a sensory event and the time of occurrence of the corresponding neural response	An encoding scheme in which information about static or dynamic signals is encoded in some aspect(s) of the temporal pattern of action potentials within the encoding window

transformation is between two different time scales: a time scale for the significant information being encoded and a *faster* time scale for the encoding symbols. It will be of great interest to identify such cases and to determine the cellular and network mechanisms underlying the implementation of temporal encoding schemes for temporal information.

The stimulus reconstruction technique is of particular utility in analyzing the encoding of dynamic stimuli. This analysis technique allows the investigator to analyze the encoding for each intrinsic time scale of the different frequency components of a dynamic stimulus. It also allows the researcher to imagine

biologically plausible schemes for *decoding* the spike-response patterns within the necessary time scale. Further, it facilitates the practical definition of (and distinction between) rate-encoding and temporal encoding schemes and offers a rigorous way to characterize coding accuracy. To our knowledge, the stimulus reconstruction technique has been applied so far only to single neurons, but the same concepts and mathematical formalism can be applied to neuronal ensembles.

Only a very few instances of temporal encoding by single neurons have been demonstrated conclusively to date. Most of those cases were limited to situations

in which the information being encoded by the neuron concerned slowly varying or static stimulus parameters. The possibility for temporal encoding by ensembles of neurons is, however, much greater. An understanding of temporal encoding may be an important key to understanding many functions of nervous systems (and in particular of higher-level functions), and further examples need to be identified and analyzed (see Table 1).

Appendix

Through the stimulus-reconstruction methodology, a best estimate of the stimulus waveform (and any linear or nonlinear transformation of the stimulus waveform) can be obtained from a transformation of the spike-train response waveform. This method is appropriate for the characterization of the nature of the encoding. In that case, the transformation operating on the spike train can be interpreted as the decoding operation, and the form of the transformation will determine whether or not the stimulus is being encoded with a rate or a temporal scheme.

In the time domain, the transformation can be expressed as a functional expansion in terms of Volterra kernels, such that the best estimate of the stimulus $S_{\text{est}}(t)$ can be obtained by operating on the spike-train response waveform $R(t)$:

$$S_{\text{est}}(t) = h_0 + \int_{-\infty}^{\infty} h_1(\tau)R(t - \tau)d\tau + \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h_2(\tau_1, \tau_2)R(t - \tau_1)R(t - \tau_2)d\tau_1d\tau_2 + \dots \quad (1)$$

In order to determine whether $S(t)$ is being encoded with a rate-encoding scheme or a temporal encoding scheme, the Volterra expansion can be rewritten in the frequency domain. The kernels $h_1(\tau)$, $h_2(\tau_1, \tau_2)$... of the Volterra expansion can be expressed in the frequency domain by taking the Fourier transform of the corresponding order. For example,

$$H_2(w_1, w_2) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h_2(\tau_1, \tau_2)e^{-i(w_1\tau_1 + w_2\tau_2)}d\tau_1d\tau_2 \quad (2)$$

The convolutions in the Volterra expansion can then be rewritten in terms of products in the frequency domain, such that

$$S_{\text{est}}(t) = F^{-1}\{H_0\delta(w) + H_1(w)R(w) + H_2(w_1, w_2)R(w_1)R(w_2) + \dots\} \quad (3)$$

where F^{-1} stands for the inverse Fourier transform. In the above transformation, the inverse Fourier transform of the correct order is first applied to each term, and the multiple time variables that are obtained are then all set equal to t . For example, if

$$I_2(t_1, t_2) = \frac{1}{(2\pi)^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{+i(w_1t_1 + w_2t_2)} H_2(w_1, w_2)R(w_1)R(w_2)dw_1dw_2, \quad (4)$$

then

$$G_2(t) = I_2(t, t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h_2(\tau_1, \tau_2)R(t - \tau_1)R(t - \tau_2)d\tau_1d\tau_2. \quad (5)$$

Equation 3 can be rewritten by grouping all the terms corresponding to each frequency:

$$S_{\text{est}}(w) = H_0\delta(w) + H_1(w)R(w) + \oint_{w=w_1+w_2} H_2(w_1, w_2)R(w_1)R(w_2)dw_1dw_2 + \iint_{w=w_1+w_2+w_3} H_3(w_1, w_2, w_3)R(w_1)R(w_2)R(w_3)dw_1dw_2dw_3 + \dots \quad (6)$$

Or, by substitution,

$$S_{\text{est}}(w) = H_0\delta(w) + H_1(w)R(w) + \int_{-\infty}^{\infty} H_2(w_1, w - w_1)R(w_1)R(w - w_1)dw_1 + \dots \quad (7)$$

In the above equation, frequency components of the spike pattern other than w are combined in the second- and higher-order terms in order to obtain the estimate of the stimulus at frequency w . In a rate encoding scheme, these frequency components can only be less in magnitude than w . We can write (7) to separate the

frequency components below and above w :

$$S_{\text{est}}(w) = \left\{ \begin{array}{l} H_0 \delta(w) + H_1(w)R(w) \\ + 2 \int_0^w H_2(w_1, w - w_1)R(w_1)R(w - w_1)dw_1 \\ + 4 \int_0^w \int_0^{w-w_1} \{H_3(w_1, w_2, w - w_1 - w_2) \\ \times R(w_1)R(w_2)R(w - w_1 - w_2)dw_1dw_2\} \\ + \dots \end{array} \right. + \left\{ \begin{array}{l} 2 \int_w^\infty H_2(w_1, w - w_1)R(w_1)R(w - w_1)dw_1 \\ + 4 \int_w^\infty \int_{w-w_1}^0 \{H_3(w_1, w_2, w - w_1 - w_2) \\ \times R(w_1)R(w_2)R(w - w_1 - w_2)dw_1dw_2\} \\ + \dots \end{array} \right. \quad (8)$$

If the Volterra functional expansion relating the response spike pattern to the stimulus can be written solely in terms expressed in the first group, then the encoding can be classified as a rate code. If terms from the second group are involved, then the encoding can be classified as a temporal code.

Notes

- Note that the task of encoding the amplitude and phase of one frequency component of a signal within a particular, discrete time window (equal in duration on one cycle period) is formally equivalent to the task of continuously encoding the waveform of that frequency component over longer time periods. The actual formulation of the neural encoding problem will depend on the memory of the system (see the definition of *integration time window* below). Also note that a neuron can be responsive to a 100 Hz frequency component without necessarily encoding the phase and amplitude of that component. For example, the response of the neuron might be correlated only with the amplitude of an envelope of the 100 Hz component. In such a case, the appropriate encoding time window would be correlated with the longer time scale characteristic of the dynamics of the envelope.
- Note that the forward and reverse transformations are not necessarily inverses of one another and that the difference between the two approaches is not just a matter of interpretation. In cases requiring the incorporation of nonlinear terms in the functional expansions, the forward and reverse transformations could be significantly different. A functional expansion can only describe deterministic relationships between two random variables, and only a single estimate for the output variable can be obtained from operation on the input variable. Therefore, if the deterministic output function to be represented by the expansion is not single-valued, then the transformation is not an adequate

representation of the relationship between the input and output functions.

- Note that the rate code for such a high-frequency encoder is restricted to a simple binary code, where the only two possible symbols are 0 (no spike) or 1 (spike). It is only when an ensemble of such high-frequency neurons are taken into account that an analog rate code capable of encoding the amplitude of the signal would be obtained.

References

- Abbott LF (1994) Decoding neuronal firing and modeling neural networks. *Quarterly Rev. of Biophysics* 27:291–331.
- Abeles M, Bergman H, Margalit E, and Vaadia E (1993) Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophys.* 70:1629–1638.
- Adrian ED (1928) The basis of sensation. Norton, New York.
- Bialek W, Rieke F, Van Steveninck RR, and Warland D (1991) Reading a neural code. *Science* 252:1854–1857.
- Bialek W and Rieke F (1992) Reliability and information transmission in spiking neurons. *TINS* 15:428–434.
- Covey E and Casseday JH (1991) The monaural nuclei of the lateral lemniscus in an echolocating bat: Parallel Pathways for analyzing temporal features of sound. *J. Neurosci.* 11:3456–3470.
- Eckhorn R and Pöpel B (1974) Rigorous and extended application of information theory to the afferent visual system of the cat. I. Basic concepts. *Biol. Cybernetics* 16:191–200.
- Eckhorn R and Pöpel B (1975) Rigorous and extended application of information theory to the afferent visual system of the cat. II. Experimental results. *Biol. Cybernetics* 17:7–17.
- Eggermont JJ (1993) Wiener and Volterra analyses applied to the auditory system. *Hearing Research* 66:177–201.
- Engel AK, Konig P, Kreiter AK, Schillen TB, and Singer W (1992) Temporal coding in the visual cortex: New vistas on integration in the nervous system. *TINS* 155:218–226.
- Freeman WJ (1991) The physiology of perception. *Scientific American* 264:78–88.
- Fuller MS and Looft FJ (1984) An information theoretic analysis of cutaneous receptor responses. *IEEE Trans. on Biomed. Eng.* 314:377–383.
- Gooler DM and Feng AS (1992) Temporal coding in the frog auditory midbrain: The influence of duration and rise-fall time on processing of complex amplitude-modulated stimuli. *J. Neurophys.* 67:1–2.
- Gray CM, Konig P, Engel AK, and Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338:334–337.
- Kjaer TW, Hertz JA, and Richmond BJ (1994) Decoding cortical neuronal signals: network models, information estimation and spatial tuning. *J. Comput. Neurosci.* 1:109–139.
- Korenberg MJ and Hunter IA (1990) The identification of nonlinear biological systems: Wiener kernel approaches. *ANN. of Biomed. Eng.* 18:629–654.
- Laurent G and Davidowitz W (1994) Encoding of olfactory information with oscillating neural assemblies. *Science* 265:1872–1875.
- Marmarelis P and Marmarelis V (1978) Analysis of Physiological Systems: The White Noise Approach. Plenum Press, New York.

- McClurkin JW, Optican LM, Richmond BJ, and Gawne TJ (1991) Concurrent processing and complexity of temporally encoded neuronal messages in visual perception. *Science* 253:675–677.
- Middlebrooks JC, Clock AE, Xu L, and Green DM (1994) A panoramic code for sound location by cortical neurons. *Science* 264:842–844.
- Miller JP, Jacobs GA, and Theunissen FE (1991) Representation of sensory information in the cricket cercal sensory system. I. Response properties of the primary interneurons. *J. Neurophys.* 66:1680–1689.
- Moiseff A and Konishi M (1981) Neuronal and Behavioral sensitivity to binaural time differences in the owl. *J. Neurosci.* 1:40–48.
- Optican LM and Richmond BJ (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophys.* 57:163–178.
- Perkel DH and Bullock TH (1968) Neural Coding. *Neurosci. Res. Prog. Bull.* 63:221–348.
- Pierce JR (1961) Symbols, Signals, and Noise: The Nature and Process of Communication. Harper, New York.
- Richmond BJ, Optican LM, Podel M, and Spitzer H (1987) Temporal encoding of two-dimensional patterns by single units in the primate inferior temporal cortex. I Response characteristics. *J. Neurophys.* 57:132–146.
- Richmond BJ, Optican LM, and Spitzer H (1990) Temporal encoding of two-dimensional patterns by single units in the primate primary visual cortex. I. Stimulus-response relations. *J. Neurophys.* 64:351–369.
- Rieke FM (1991) Physical Principles Underlying Sensory Processing and Computation. Doctoral thesis, University of California at Berkeley.
- Shannon CE (1948) A mathematical theory of communication. *AT&T Bell L. Tech. J.* 27:379–423.
- Theunissen FE (1993) An Investigation of Sensory Coding Principles Using Advanced Statistical Techniques. Doctoral thesis, University of California at Berkeley.
- Theunissen FE and Miller JP (1991) Representation of sensory information in the cricket cercal sensory system. II Information theoretic calculation of system accuracy and optimal tuning curve width of four primary interneurons. *J. Neurophys.* 66:1690–1703.
- Tovée MJ, Rolls ET, Treves A, and Bellis RP (1993) Information encoding and the responses of single neurons in the primate visual cortex. *J. Neurophys.* 70:640–654.
- Werner G and Mountcastle VB (1963) The variability of central neural activity in a sensory system, and its implications for the central reflection of sensory events. *J. Neurophys.* 26:958–977.
- Warland D, Landolfa MA, Miller JP, and Bialek W (1991) Reading between the spikes in the cercal filiform hair receptors of the cricket. In FH Eckman, ed. *Analysis and Modeling of Neural Systems*. Kluwer, Boston, MA. pp. 327–333.