Published in final edited form as:
Biol Theory. 2006 ; 1(3): 302–316.

Approaches to Information-Theoretic Analysis of Neural Activity

Jonathan D. Victor

Department of Neurology and Neuroscience Weill Medical College of Cornell University New York, NY, USA

Abstract

Understanding how neurons represent, process, and manipulate information is one of the main goals of neuroscience. These issues are fundamentally abstract, and information theory plays a key role in formalizing and addressing them. However, application of information theory to experimental data is fraught with many challenges. Meeting these challenges has led to a variety of innovative analytical techniques, with complementary domains of applicability, assumptions, and goals.

Keywords

direct method; information theory; metric space method; neural coding; sensory; systems; spike trains

The goal of this review is to identify some of the questions in neuroscience for which information-theoretic techniques provide useful insights and approaches, and to survey the variety of techniques that are applicable to the analysis of neurophysiologic data.

How neurons represent, process, and transmit information is of fundamental interest in neuroscience. The basic biophysics that underlies neuronal action potential generation is well established, as is the biophysics underlying many aspects of synaptic physiology and dendritic information processing. Nevertheless, the features of neuronal activity that convey and manipulate information are not well understood. Among the possibilities are relatively straightforward features, such as the number of spikes fired by a population of neurons (Shadlen and Newsome 1998), but also more subtle ones, such as, their precise times of occurrence (Softky 1994; Théunissen et al. 1996; Berry et al. 1997; Gawne 2000), the pattern of intervals (Sen et al. 1996), the presence or absence of correlations and synchrony (Meister et al. 1995; Dan et al. 1998; Rodriguez et al. 1999; Samonds et al. 2006), oscillations (Gray and Singer 1989), or other patterns of activity (Abeles and Prut 1996).

Questions related to neural coding are intrinsically abstract, since, at a minimum, they seek a description of a mapping from events, percepts, and actions to something very different: patterns of neural activity. Although it may be tempting to assume that a common set of principles governs neural coding, it is more reasonable to anticipate that there is a diversity of biological solutions to the coding problem. That is, we anticipate that neural coding will differ greatly according to the pressures under which a system has evolved. Such “design” criteria likely include minimizing the number of neurons or their connections, minimizing energy utilization, minimizing response latency, maximizing robustness in the face of injury, or maximizing the capacity for learning. We anticipate that coding strategies may differ across brain regions, even within a single “system.” For example, cortical regions early in visual processing (V1, V2) are tightly topographically organized, while visual regions at the “top” of

the inferotemporal stream, which interact extensively with polysensory areas and the hippocampus, have little topographic organization. Even within the early stages of visual processing, there is a qualitative change between coding in V1, and coding in V2—with temporal multiplexing of multiple visual submodalities much more prominent in V2 (Victor and Purpura 1996a). Finally, strategies for representing information, even within a particular cell type, are likely task dependent and subject to top-down influences. For example, attention modulates firing rate (Luck et al. 1997; Reynolds et al. 2000) and synchrony (Roelfsema et al. 2004; van der Togt et al. 2006). However, it is as yet unclear what is the primary neural correlate of attention.

Need for joint experimental and theoretical/computational approach

A purely experimental approach to these questions is not likely to succeed, in that manipulation of one feature of neural activity (e.g., increasing firing rate by electrical stimulation) is certain to change other aspects as well (e.g., interval structure and degree of correlation). Thus, while such experiments (Salzman and Newsome 1994) are critical in demonstrating that a particular brain region is relevant to a particular function, they provide little insight into neural coding.

An appropriate theoretical infrastructure is needed to disentangle these confounds, and also to compare results across a range of modalities, preparations, brain areas, and species. Shannon's groundbreaking work in information and communication theory (Shannon and Weaver 1949) is the natural starting point for this theoretical infrastructure (Rieke et al. 1997). But, while application of Shannon's ideas to man-made communication channels is relatively straightforward, difficulties arise in attempting to apply information measures to biologic systems. Fundamentally, the Shannon theory was designed for characterizing communication systems whose principles were understood, not for the "inverse problem" of determining the principles by which a system works from observations of its behavior.

To make full use of information theory (and to avoid assuming answers to the above questions), one would want to begin with as few assumptions as possible about the nature of the neural code. A minimal assumption is that each possible configuration of neural activity (i.e., each arrangement of spikes across time and a set of neurons) is a candidate for a code word. Ideally, the formalism of information theory would then determine the actual set of words (and hence the structure of the neural code) from this starting point. Unfortunately, this program rapidly runs into practical difficulties. Experimental estimates of information are biased by finiteness of datasets, and the extent of this bias is directly proportional to the size of the a priori set of words (Carlton 1969). Moreover, the Shannon theory does not attempt to describe the relationship between a sensory or motor domain and neural activity (i.e., the nature of the neural representation) but merely provides an index of how faithful this representation is. As we will see below, these considerations motivate a variety of approaches to the analysis of neural coding. These approaches share the goal of quantification of information. However, they differ substantially in the scope of the assumptions concerning neural coding, by the extent to which they yield a description of the representation provided by the code, and the kinds of data to which they may be applied.

Correlation and Causation

Correlation of a behavior or stimulus with a statistical feature of the neural response does not imply that this feature of the neural response is used by the nervous system. Some of the approaches described below, coupled with appropriate experimental design, may be useful in determining causal relationships. For example, a multichannel recording of neural activity (e.g., field potential activity at different locations: Schiff et al. 2000) or multiple neurons within a cluster (Reich et al. 2001b) can be partitioned into two subsets of channels, one considered

as the “input,” and one considered as the “output.” One can then determine whether statistical features in the “input” activity can predict later activity in the “output” channels. A positive answer demonstrates that the statistical features of the input are indeed used at later times in neural processing, thus going a substantial step beyond merely demonstrating the presence of these features.

Alternatively, because information cannot be created *de novo* within the nervous system, it may be possible to rule out a candidate neural code, by showing that it cannot support the sensory performance of the organism. This strategy has demonstrated the importance of spike timing in retinal coding (Nirenberg et al. 2006).

Information-Theoretic Tools Applicable to Neural Data

General Comments: A Wide Variety of Approaches

Many strategies for the application of information-theoretic tools to neural data have been proposed (Table 1). As seen in the table, these strategies have diverse, and to some extent complementary, domains of applicability, limitations, conceptual underpinnings, and questions that can be addressed. We precede our survey by some general comments on these interrelated axes.

Experimental Design—A typical experiment in classical sensory neurophysiology consists of recording neural responses to a large number of presentations of a small set of sensory stimuli. The set of sensory stimuli is generally chosen to be “simple,” with elements that vary along some perceptually salient parameter, or set of parameters. For example, in characterizing neurons in primary visual cortex, a typical stimulus set consists of gratings of varying contrast and/or orientation. Responses to such stimuli can be analyzed (without information-theoretic tools) to provide measures of neural “tuning” to these parameters. The information-theoretic viewpoint considers the neuron to be a communication channel. The “transmitted information” is a natural measure of to what extent an observer of the neural response can reduce uncertainty about which stimulus was presented. There is no pretense that this kind of experiment can fully characterize the response properties of the neuron. Nor can it hope to determine its information-transmitting capacity, since the set of stimuli is intentionally restricted to a tiny subset of all possible stimuli. Rather, the goal of information-theoretic analysis of this kind of experiment is to determine which aspects of the response are responsible for coding a perceptual parameter of interest, and the extent to which this coding is reliable.

An alternative experimental design, especially popular in vision, is based on the rapid presentation of a large set of stimuli (Wu et al. 2006), repeated a small number of times if at all. The stimuli might be chosen for analytic convenience (e.g., white noise, m-sequences), or in the hope that they represent ethologically important stimuli (e.g., real-world movies). The goal of this kind of experiment is to build a model for the functional relationship between a neuron’s input and its output. Such a model can then be tested by its ability to predict responses to other stimuli. Information-theoretic tools can then be applied to determine the information rate for the neuron’s output under the conditions of the particular experiment. Moreover, if a believable model for the neuron’s behavior can be constructed, then, at least in principle, the maximal information-transmitting capacity of the neuron (across all possible stimuli) can be calculated.

One might argue that the distinction between these two kinds of experiments is not very meaningful, since an information-theoretic analysis method that is intended to be applied to one kind of experiment can be forced to apply to the other. However, such application is unlikely to be practical, or to achieve its intended goal, even though there is nothing in the formalism of these approaches that prevents such attempts. The basic issue is that, like any

other application of mathematical concepts to laboratory data (see Slepian [1976] for an elegant discussion), a rigorous implementation of information-theoretic analyses requires evaluation of limits that cannot be achieved in the laboratory. Short of these limits, there is no *guarantee* that values estimated from laboratory data are close to their values at these limits. This difficulty typically persists even if one goes through the efforts of analyzing exactly how rapidly the limits are approached—since this analysis is also only an asymptotic one.

Thus, although the distinctions between the methods we discuss have clear-cut and rigorous theoretical foundations, their practical domains of applicability are distinguished by qualitative terms and fuzzy borders (Table 1). But this should not be taken as an excuse to ignore the philosophical differences between these approaches. At a concrete level, such differences can be recovered by an analysis of how two kinds of procedures differ in simple test cases, whose behavior can be determined analytically. More fundamentally, ignoring the distinctions between these approaches would deny one of the important contributions of the mathematical biologist—namely, creation of formalisms that allow testing, refinement, and extensions of biological intuition.

Response Types—All information-theoretic methods discussed here can be applied to experiments in which the responses are the sequences of stereotyped action potentials (“spike trains”) produced by a single neuron—the substrate for information transmission over large distances. Many of the methods are also applicable to neural signals other than action potentials. For example, subthreshold fluctuations of membrane voltage carry information within neurons. Some small neurons, such as the interneurons of the retina, do not generate action potentials, and use these continuously varying voltage fluctuations for transmission of electronic signals between neurons. Another signal that is appropriate for information-theoretic analysis is the “local field potential,” an extracellularly recorded voltage that represents a combination of synaptic activity, subthreshold fluctuations of membrane voltage, and, to a lesser extent, summed spiking activity, in a neighborhood of approximately 1 mm or less.

A spike train is most naturally represented as a point process, while intracellular and extracellular voltages are most naturally represented as a continuous real-valued function of time. As we will see below, some information-theoretic approaches are directly applicable to the point process itself. Other approaches have functions of time as their primary object of analysis. They can also be applied to spike trains, but only after the latter are converted into functions of time. Methods for making this conversion include convolution with a standard template, such as a Gaussian, or simply considering the spike trains to be a train of delta functions. The latter approach can only be used for methods that do not require that the signals be continuous. Finally, the methods that are the most directly tied to Shannon’s ideas (Shannon and Weaver 1949) have a discrete sequence of symbols drawn from a finite set, typically $\{0,1\}$, as their primary object of analysis. These methods can be applied to spike trains by dividing the data record into narrow time bins, and keeping track of how many spikes occurred in each. They can also be applied to continuous signals, by sampling them in time and discretizing them in amplitude. The utility of these approaches depends critically on how information estimates vary with bin width, which in turn depends on the biological system and the amount of data available.

Understanding neural coding requires not only a characterization of the behavior of individual neurons, but also of their joint activity. Datasets in which many channels of simultaneously recorded neural activity (spikes, continuous signals, and combinations) are increasingly available. All of the methods we will consider have immediate formal extensions from single channels to multiple channels, but these extensions differ widely in practicality. The “multichannel” regime deserves to be broken into two regimes—that of “few” channels and “many” channels. Some methods effectively require estimation of the number of parameters

that grows exponentially with the number of channels; these methods are likely to break down in the “few” channel regime. For other methods, the effective number of parameters to be estimated grows more slowly if at all, but these methods may have computational demands that may limit application when many channels are present.

A Survey of Methods for Information Estimation

The Direct Method—The “direct method” (de Ruyter van Steveninck et al. 1997; Strong et al. 1998) for the estimation of information in spike trains is closest to a literal implementation of Shannon’s ideas, and makes only minimal assumptions about the nature of the code. Thus, it provides a rigorous estimate of information, provided that sufficient data are available.

The primary data consist of records of a single neuron’s response. These records are first partitioned into segments of length L . Each segment is converted into a discrete sequence of symbols (0 or 1) by subdividing it into successive bins of width ΔT , and forming an integer sequence in which each entry indicates the number of spikes within one of these bins. ΔT is typically taken to be sufficiently short so that each bin contains at most one spike. For each integer sequence s , the probability of its occurrence, $p(s)$, is estimated from experimental data. Two entropies are then calculated. The “total entropy,” $H_{total} = -\sum p(s) \log_2 p(s)$, expresses the entropy of the entire repertoire of the observed behavior of the neuron, *for all stimuli*. The noise entropy H_{noise} is a corresponding sum but restricted to responses to a single stimulus. The estimated information is $I = H_{total} - H_{noise}$.

The estimated information I depends on the binning parameters L and ΔT . Strong et al. (1998) provide a procedure for extrapolating to the limits of $\Delta T = 0$ and $L = \infty$, as is required for a rigorous true information estimate.

The direct method has been used at several levels of the visual system, including mammalian retina (Nirenberg et al. 2001), lateral geniculate nucleus (Reinagel and Reid 2000), primary visual cortex (Reich et al. 2001b), and extrastriate visual cortex (Buracas et al. 1998). In each of these settings, the stimulus consisted of a rapidly varying temporal sequence, often constructed from a pseudorandom sequence but occasionally derived from natural images (Nirenberg et al. 2001; London et al. 2002). However, the method can also be applied to data derived from discrete presentation of a small set of stimuli (Reich et al. 2001a).

Limitations—The main limitation of the direct method is that it is simply not possible to make a rigorous extrapolation to the limits of $\Delta T = 0$ and $L = \infty$. These limits of course cannot be attained experimentally, but biologic considerations can provide guidelines for values of ΔT and L beyond which one can assume that an asymptotic regime is reached. Unfortunately, this regime may be inaccessible in practice.

For mammalian cortex, a reasonable choice of ΔT is 1 ms (an upper limit for the intrinsic precision of a neuron), while a reasonable choice for L is 100 ms (a lower limit for the duration of a response). Consequently, the number of possible sequences whose probabilities must be estimated is very large ($2^{L/\Delta T}$), and the probability distribution is necessarily undersampled by laboratory data. In this regime, entropy estimates are unreliable and highly biased—the bias is proportional to the number of probabilities that must be estimated, and inversely proportional to the total number of observations. As described below, debiasing techniques are available, but these procedures are ineffective when most bins are not even sampled at all. Consequently, the direct approach is limited to situations in which responses are highly reproducible, such as insect systems or the retina (so that only a very small number of the possible spike train configurations occur), or, to estimates of instantaneous information rate (artificially limiting L).

The direct method may be extended (Johnson et al. 2001; Nirenberg et al. 2001; Reich et al. 2001b) to simultaneous recordings from multiple neurons. In an M -neuron experiment, the response within each bin of length ΔT is described by an M -tuple of bits, in which each bit represents the firing of one neuron. Otherwise, the estimation of information proceeds exactly as for single-neuron responses. However, the undersampling of the space of all possible sequences is even more severe, since the number of possible sequences is given by $2^{ML/\Delta T}$.

In sum, the philosophy that keeps the direct method closest to Shannon's ideas is also its main limitation. Since minimal assumptions are made about the nature of the code, the probability of each response (as represented by a discrete sequence) is an independent quantity to be estimated from data. That is, the tradeoff for an approach that is free of a priori assumptions is one that, for rigorous implementation, requires an impracticably large amount of data in many circumstances. Moreover, the direct method provides little insight into *how* information is carried—since how information is carried is explicitly a statement about the relationships among the response sequences.

Estimators of Entropy of a Discrete Distribution—A key component of the “direct method,” as well as of many of the approaches described below, is that the entropy of a discrete distribution must be estimated from a finite set of observations. This seemingly simple problem is surprisingly subtle. The entropy of a discrete distribution with J bins and a probability p_j in

each bin is $H = - \sum_{j=1}^J p_j \log_2 p_j$. The naive approach is to estimate this by setting $p_j = n_j/N$, where n_j is the number of times that the j th outcome is observed, and N is the total number of observations. This “plug-in” estimator is well known to be biased—fundamentally, because of the curvature of the log function. A standard fix is to add a bias correction (Miller 1955; Carlton 1969; Treves and Panzeri 1995). This bias correction is asymptotically exact for large N but requires knowledge of the number of kinds of categories (or bins), J , that are occupied with nonzero probability. Moreover, typical datasets are not in the “asymptotic” regime, which requires that even the least likely outcome has been observed several times. An alternative correction is the jackknife (Efron 1982; Efron and Tibshirani 1998), but this has similar asymptotic behavior. More sophisticated estimators have recently been introduced, with clear advantages in regimes relevant to laboratory data. These include Paninski's estimator (Paninski 2003), which is provably the least biased of all polynomial estimators, the “KT” (Krichevsky and Trofimov 1981) and “SG” (Schurmann and Grassberger 1996) estimators, which are based on single Dirichlet priors (Wolpert and Wolf 1995), the “NSB” estimator (Nemenman et al. 2004), which considers a family of Dirichlet priors, and the Chao-Shen estimator, recently introduced in ecology (Chao and Shen 2003). However, none of these estimates succeed in the severely undersampled regime characteristic of cortical datasets.

Metric Space Method

The direct method, though virtually assumption free, can have prohibitive data requirements, and does not attempt to characterize the manner in which information is represented. The metric space method (Victor and Purpura 1997; Victor 2005) represents an alternative viewpoint. By making assumptions as to the nature of a neural code, it can provide useful estimates of information in settings in which the direct method will fail (limited amounts of data, and especially high firing precision but low firing rate).

The metric space method considers several generic families of neural codes, each of which is designed to test a particular hypothesis of *how* information is carried, such as via spike counts, or via the timing of spikes, or via the interval structure of the spike trains. Each of these hypotheses is then formalized in terms of a family of metrics— notions of distance (i.e., dissimilarity) between spike trains. The metrics have a common structure, which allows comparison of the hypotheses on a level playing field. Because the metrics explicitly recognize

that neural responses are point processes and their structure respects the continuity of time, the binning process that limits the use of the direct method is avoided. However, the metric space method typically underestimates the total information that is present, since only a few stereotyped (but interpretable!) hypotheses for neural codes are considered. Also, because of the way that information is calculated, the approach is limited to analysis of episodic responses to a discrete set of stimuli.

Many neurons can be considered to behave like coincidence detectors (Bourne and Nicoll 1993; Mel 1993; Softky and Koch 1993; Cline 1997; Markram et al. 1997; Usrey et al. 1998). This suggests that the meaning of a spike train is determined by the timing of the individual spikes, since it is those timings that determine how the multiple inputs onto a dendritic tree interact to determine a postsynaptic neuron's behavior. To assess the extent to which spike times carry information, the approach uses a family of metrics denoted by $D^{spike}[q]$, parameterized by a quantity q (see below) that describes the role of temporal pattern. According to the metric $D^{spike}[q]$, the distance between two spike trains is the minimum total "cost" to transform one spike train into the other via any sequence of insertions, deletions, and time shifts of spikes. The cost of moving a spike by an amount of time t is set at qt , and the cost of inserting or deleting a spike is set at 1. Thus, in the sense of $D^{spike}[q]$, spike trains are considered similar if they have approximately the same number of spikes, and these spikes occur at approximately the same times, i.e., within $1/q$ or less. A neuron that behaves like a coincidence detector with temporal precision $1/q$ would see incoming spike trains as similar or different, according to the metric $D^{spike}[q]$.

A second family of metrics, denoted by $D^{interval}[q]$, is motivated by the notion that a synaptic response depends on its recent history, and thus, the intervals between successive spikes may also carry information (Bliss and Collingridge 1993; Sen et al. 1996; Abbott et al. 1997; Usrey et al. 1998). In metric $D^{interval}[q]$, the distance between two spike trains is defined as the minimum total cost to transform one spike train into the other via any sequence of insertions of spikes, deletions of spikes, and expansions or contractions of interspike intervals. The parameter q specifies the cost qt of changing an interspike interval by an amount t . In the limit that q approaches 0, both $D^{spike}[q]$ and $D^{interval}[q]$ approach a metric D^{count} , which is sensitive only to the number of spikes, and not to any aspect of their timing.

Each metric is then evaluated by the extent to which it distinguishes the responses to each of the stimuli—namely, the transmitted information between stimulus and response clusters. The dependence of the transmitted information on q for $D^{spike}[q]$ and $D^{interval}[q]$ characterizes the importance of spike timing and interspike intervals, across a range of temporal precisions.

Applications of this approach to neural data, including visual cortex (Victor and Purpura 1996a; Reich et al. 2001c; Samonds and Bonds 2004), chemical senses (Stopfer et al. 1997; Di Lorenzo and Victor 2003), and electric sense (Kreiman et al. 2000) are reviewed in Victor (2005).

The metric space approach is readily extended to the multineuronal context. A multiunit recording is a sequence of labeled events, with the label representing the neuron of origin. To assess the importance of *which* neuron fires each spike, multineuronal metrics add an additional transformation between spike trains: changing the label associated with a neuron. The cost of this transformation is assigned the quantity k . The extreme $k = 0$ corresponds to a code in which the neuron of origin is irrelevant (since it is free to change the label associated with each spike). The other extreme, $k = 2$, corresponds to a labeled-line code (since it costs as much to change the label on a spike as it does to remove it from one neuron, and insert it into another). The above analyses can then be carried out for the two-parameter family $D^{spike}[q, k]$.

By introducing a *single* parameter to explore the continuum between codes in which neuron of origin is irrelevant and labeled-line codes, the explosion of parameters that might otherwise hobble attempts to analyze multineuronal data is circumvented. We have applied this approach to simultaneously recorded neural pairs in V1 (Aronov et al. 2001), and have found that responses are best decoded by keeping track of which neuron fired which spike, but only a modest amount of information is lost by ignoring the neuron of origin. This is in keeping with our analysis of multineuronal recordings in V1 via the direct method (Reich et al. 2001c), but is complementary to it: the direct method can analyze recordings of up to six neurons (the limits of our recording), but only looks at information rates over brief time intervals (e.g., 15 ms). In contrast, the metric space method can examine responses over extended periods.

For multineuronal responses, algorithms for the calculation of distances via straightforward extension of the Sellers algorithm (Sellers 1974) (see below) yield a calculation time proportional to c^{2M} , where M is the number of neurons and c is the typical number of spikes in a spike train. An improved dynamic programming algorithm that drops the exponent from $2M$ to $M + 1$ was recently found (Aronov 2003). This dramatic improvement makes calculations on triplets of neurons practical on a desktop, and enables analysis of four to eight neurons (for firing rates typical of cortical neurons) with a parallel processor array.

Limitations—One important limitation of the metric space approach is that there is no guarantee that the manner of information transmission is similar to either of these caricatures. For example, the informative precision of a spike may be greater during the transient part of a response than during a later period in which firing occurs at a lower rate. In the multineuronal situation, it may be appropriate to distinguish among some neurons within the population and not others, rather than to have a single omnibus cost for changing the label of a neuron. One can augment the metric-space method by including these (and other) variations. Consequently, the maximal value of the transmitted information obtained with any of the candidate metrics is necessarily an underestimate of the total amount of information. Since there are also coding strategies that do not readily fit into the metric structure, it is difficult to place rigorous bounds on the extent of this underestimate.

A second major limitation of the metric-space method is a consequence of the clustering stage, in which distances between responses to the same stimulus and distances between responses to different stimuli are compared. For the clustering stage to be effective, the number of samples collected in response to each stimulus must be somewhat larger than the number of stimuli. This makes it impractical to apply the metric space method to responses elicited to long, rich sequences of continuously presented stimuli.

Relation to Comparison of Genetic Sequences—The above metrics for spike trains have a common structure: distance is defined as the minimum cost of a transformation of one sequence into another, via a sequence of prescribed elementary transformations. This structure is formally identical to that of the distances used to compare genetic sequences (Sellers 1974). For genetic sequences, the elementary transformations include insertion, deletion, and alteration of a discrete element. The spike train metrics operate on point processes in continuous time, while the distances for genetic sequences operate on discrete sequences. Despite this topological difference, the highly efficient dynamic programming algorithms developed by Sellers (1974) for genetic sequences can be adapted to spike train metrics, so that the calculations described above can be carried out efficiently.

Not Just Information—The metric space approach, and others to be described below, goes beyond traditional information-theoretic analysis in an important way. One can determine whether the presumptive code provides for a *representation* of the stimulus domain, and not just for faithful discrimination of distinct stimuli. One way to accomplish this is to use the

pairwise distances as the starting point for multidimensional scaling (Victor and Purpura 1997; Aronov et al. 2001). For example, reanalysis of the auditory data of Middlebrooks et al. (1994) demonstrated that the temporal aspects of the spike trains not only identify the azimuth of origin of a sound, but also that these temporal aspects *represent* the azimuth: they map the responses into a circular locus in an abstract response space (Victor and Purpura 1997). Moreover, the coordinates within the multidimensional scaling space are the temporal features that distinguish and represent the stimuli. Such an analysis of V1 recordings (Aronov et al. 2001) demonstrated a consistent temporal representation of spatial phase across neurons, with one coordinate consisting of the sustained portion of the response, and a second coordinate consisting of a transient component.

Embedding Method

The “embedding method” is an approach that combines many of the advantages of the two approaches discussed above (Victor 2002). Like the metric space method, it exploits the continuity of time and avoids binning. But in contrast to the metric space method, it makes no assumptions concerning the nature of the code, other than that it respects the continuity of time. Consequently, it is provably unbiased (Kozachenko and Leonenko 1987), at least when sufficient data are available. It can be extended to multichannel data, but its behavior is intermediate between that of the metric space method (a single parameter is added) and the direct method (exponential growth in number of parameters to be estimated). While the approach cleanly separates information carried by spike counts from information carried by spike times, it does not provide as detailed a parsing of temporal information as does the metric space method. In contrast to both the metric space method and the direct method, this approach is immediately applicable to continuous responses and spike trains.

The key idea behind this approach is a formalization of a basic attribute that a coding scheme must have in order to be biologically plausible. A sufficiently small change in the time of occurrence of a spike cannot result in a change in the meaning of a spike train, and spike trains that differ by only an infinitesimal change in a spike time must have nearly identical probabilities. Thus, like the metric space method, the continuity of time is used explicitly. But unlike the metric space method, there is no assumption made concerning the relationship of spike trains that differ by small or large displacements of a spike. Also, in contrast to the metric space method, the approach does not assume a relationship between the two spike trains that differ by insertion or deletion of a spike. These ideas are naturally formalized in terms of the topology of spike trains (McFadden 1965). That is, the space of spike trains of finite duration can be considered to consist of a discrete set of strata, one for each number of spikes. Spike trains with n spikes form an n -dimensional manifold (parameterized by the time of each spike). A neuron’s output is described by a probability distribution on this set of strata. Within each stratum, the probability distribution is assumed to vary smoothly, but between strata no assumptions are made.

Thus, to determine the amount of transmitted information in an experimental dataset, spike trains are stratified according to the number of spikes n in the response. This partitioning generates one component of the information, I_{count} , reflecting the extent to which the total number of spikes in the response can distinguish between the stimuli. Since I_{count} is determined from a relatively small number of response categories, a standard discrete calculation may be used, and standard bias corrections are effective. Then, the n th stratum is analyzed to determine a contribution of spike timing $I_{timing}(n)$. The total information is $I_{count} + \sum_n I_{timing}(n)$, where the second term is the total information due to spike timing.

The calculation within the n th stratum crucially exploits the assumption that the probability distribution is a continuous function of the spike times. To determine $I_{timing}(n)$, the spike trains in the n th stratum are embedded into a Euclidean space of dimension $r \leq n$. The coordinates

assigned to a response are determined by inner products with a set of functions $f_1, \dots, f_r$: a spike train x with spikes at times $\tau_1, \tau_2, \dots, \tau_n$ is mapped to coordinates $c_h(x) = \sum_{k=1}^n f_h(\tau_k)$.

For continuous signals, there is no discrete component corresponding to the number of spikes, and all responses are embedded into a space of the same dimension. A reasonable choice for the embedding is the natural extension of the above linear map to continuous signals: a signal $v(t)$ is mapped into the coordinates $c_h(v) = \int f_h(t) v(t) dt$.

As in the direct method, transmitted information is calculated as a difference between a “total entropy” determined from all responses considered together, and a “noise entropy” determined within the responses to each stimulus. However, in contrast to the direct method, these entropies are determined by examining the statistics of the nearest-neighbor distances (Kozachenko and Leonenko 1987). In particular, the contribution of spike timing to the information within the n th stratum is estimated by

$$I_{\text{timing}}(n) \approx \frac{r}{N(n)} \sum_{j=1}^{N(n)} \log_2 \left(\frac{\lambda_j}{\lambda_j^*} \right) - \sum_{k=1}^s \frac{N(n, a_k)}{N(n)} \log_2 \frac{N(n, a_k) - 1}{N(n) - 1},$$

where $N(n)$ is the number of spike trains with n spikes, $N(n, a_k)$ is the number of spike trains with n spikes elicited by the k th stimulus, λ_j is the distance between the j th spike train and its nearest neighbor, and λ_j^* is the distance between the j th spike train and its nearest neighbor elicited by the same stimulus. For quantities of data typically available in an experiment, this nearest-neighbor estimator (of entropy or of information) is substantially more efficient than binned methods. Demonstration that this estimator is unbiased (Kozachenko and Leonenko 1987) relies critically on the assumption of smoothness of the probability distribution.

Limitations—The limitations of the embedding approach relate chiefly to the partitioning of the entropy estimate. When the range of the number of spikes in responses is large, there are many discrete partitions. In this regime, the bias estimates for I_{count} may be ineffective. Moreover, at the tails of the distributions of spike counts, there are only a few responses, so that the estimate of I_{timing} may be ineffective. These difficulties may be mitigated by lumping together partitions with similar numbers of spikes, but this compromises the unbiased nature of the estimator. The practical difficulties of the discrete component are exacerbated when the method is applied to multineuronal data, since a separate partition is required for each combination $(n_1, n_2, \dots, n_M)$ of spike counts on each of the M -neurons. This rate of growth of the number of partitions that must be separately analyzed, though high, is much lower than in the direct method, since it is independent of (rather than exponential in) temporal resolution.

Relation to General Dynamical Systems Approaches—Estimation of entropy from the statistics of nearest neighbors is related to estimation of dimension of a dynamical system's trajectory or attractor set. Grassberger and Procaccia (1983) describe several versions of such procedures, wherein dimension is determined from the relationship between the number of points within a given radius and the radius. When plotted on log-log coordinates, the slope of this relationship is the sought-after dimension. But in the present situation, the slope is known (the dimension of the space in which we have embedded spike trains), and the quantity of interest, the entropy, is essentially the intercept of this line. Grassberger's (1988) finite-sample debiasing procedure applies specifically to the slope (dimension); the Kozachenko and Leonenko (1987) estimator debiases the intercept (entropy).

Grassberger and colleagues (Kraskov et al. 2004) have recently described a related approach to estimating mutual information via a nearest-neighbor approach that avoids explicit estimates of dimension. However, this approach requires that the response variable has a definite dimension. Thus, for application to spike trains, a procedure such as stratification by spike count is required to obtain an unbiased estimator, as in Victor (2002).

Context Tree Method

The context tree method is a promising new approach both for entropy estimation (London et al. 2002; Kennel et al. 2005) and for estimation of mutual information applicable to the “many-presentation” experimental design (Shlens et al. 2006). Like the direct entropy estimator (de Ruyter van Steveninck et al. 1997; Strong et al. 1998), it is based on a discrete representation of spike trains, but, it also makes crucial use of the dynamic nature of spike trains—namely, that a spike train is a temporal sequence in which the recent past influences the probability of spiking. This dynamic process is modeled as a “context tree” (Rissanen 1989), which differs from a Markov process in that the depth of the history dependence can be nonuniform. This model form is intuitively appealing for neural data, and results in a substantial increase in efficiency compared with approaches (see “Compression method”, below) that make use of dynamics, but do not postulate a model form.

In essence, the method has two components: estimation of a context tree model from the spike train data, and then calculation of entropy from the context tree itself (e.g., by a Wolpert-Wolf estimator: Wolpert and Wolf 1995). However, rather than choose a single context tree model (cf. Hirata and Mees 2003), the approach considers many context tree models. Each model’s contribution is discounted (Willems et al. 1995) by a factor that considers both the complexity of the model (its “codelength”: Solomonoff 1964) and the extent to which the model is a poor fit to the data. An advantage of this approach is that confidence limits on the entropy estimates can be determined via a Monte Carlo method that explores the range of estimates that would result from alternative context tree models (Kennel et al. 2005).

Other Methods

Below we describe several other approaches that may be usefully applied to estimation of information in neural data. Our goal is to emphasize the variety of viewpoints that may be taken, rather than to present an exhaustive review.

Principal Components—The procedures used by Richmond and Optican (Optican and Richmond 1987; Richmond and Optican 1987; Chee-Orts and Optican 1993) are based on principal-components analysis of rate functions estimated from single-trial neural response. The hypothesis underlying this approach is that information is coded as a firing rate envelope, and that individual spike trains serve as estimators of this envelope. This approach can also be viewed as a kind of embedding method, in that the rate-coding hypothesis leads to embedding of all responses in a space of the same dimension, regardless of the number of spikes. Within this space, information is estimated by parceling this space into multidimensional bins. A regularization procedure based on an additive noise model and an assumed Gaussian shape of the response cluster were used to improve performance (Chee-Orts and Optican 1993). To the extent that neural codes indeed conform to the rate envelope hypothesis, the principal-components approach will provide a good description of the code, with limited sample sets of the size achievable in typical experiments (Optican and Richmond 1987; Richmond and Optican 1987; McClurkin et al. 1991). However, by design, it will overlook any other forms of coding. Additionally, the Gaussian regularization for estimation of entropy, rather than the nearest-neighbor estimator used in the embedding method, is tantamount to adding an assumption about the manner in which responses vary across trials.

Reconstruction Method—The reconstruction method of Bialek and coworkers (Bialek et al. 1991) was the first information-theoretic approach successfully applied to decoding dynamic neural activity. It provides another way of avoiding the difficulties associated with estimating a large number of probabilities, as is required by the direct method. The basic strategy is to identify a transformation of the observed neural response that best reproduces the known stimulus sequence. The transmitted information in the neural response is then known to be at least as high as the mutual information between the actual stimulus and the stimulus reproduced by this transformation rule. In some settings, a priori calculations allow for an independently calculated upper bound on the amount of information in the neural response, based on the theoretical limits of a sensory system (Bialek et al. 1991). When the upper bound provided by these considerations is close to the lower bound provided by a reconstruction, this approach is particularly powerful and elegant.

To seek a transformation between the neural response and the stimulus, a functional form must be chosen. This functional form is typically linear, though nonlinear extensions via the Volterra formalism (Marmarelis and Marmarelis 1978) can be used. The kernels that describe the transformation can then be interpreted as a recipe for “reading” the neural code (Bialek et al. 1991). The approach is typically applied to the spiking activity of single neurons (Théunissen et al. 1996), but the concept readily extends to multiple channels and/or continuously varying data. One limitation of the approach is that the stimulus must be represented as a time series, rather than as discrete elements of a space. More fundamentally, the approach may be impractical for highly nonlinear transformations, such as are likely to be present within the mammalian central nervous system, since the fitting of second order (or higher) terms in a Volterra series will not be robust.

Power Series Method—Panzeri and Schultz (Panzeri and Schultz 2001; Schultz and Panzeri 2001) introduced another strategy for overcoming many of the shortcomings of the direct method by exploiting the continuity of time. Here, the basic assumption is that information is an analytic function of the length of the analysis interval L . Under this assumption, information can be expanded as a power series in L . Very short intervals are likely to contain at most one spike. The probability that a pair of spikes occurs within the analysis interval increases with the square of the length of the interval. Thus, an advantage of this approach is that the terms of the Taylor series expansion separate the contributions of firing rate, pairwise correlation between spikes, and higher order correlations. This parsing of temporal information, which is explicitly order-by-order, is intrinsically limited to spike trains. However, it is distinct from (and more detailed than) the kind of parsing provided by the metric space method. Additionally, this approach bypasses the construction of a response space, so there is no attempt to determine whether stimuli are “represented” by the temporal patterns of activity.

In contrast to the reconstruction method, it is not assumed that the relationship between a spike train and what it represents has a low-order power series expansion. Rather, a power series is used to represent the information content of a spike train as a function of the duration of the interval (i.e., order-by-order in the number of spikes). Thus, the power series method will have no trouble with highly nonlinear transformations such as thresholds and saturations that might lead to difficulties with the reconstruction method.

The power series approach is readily extended to multiple spike trains, but at any fixed order of approximation, the number of cross-terms grows as a polynomial in the number of neurons. The second-order terms can be further separated into auto- and cross-correlation terms, providing insight into how information is coded across a population of neurons. On the other hand, when the spike trains have structure such as regularity or bursts, there is no guarantee

that the power series converges rapidly, or even at all. This may prevent successful application to such spike trains, or to large analysis intervals.

This approach has been used successfully to study somatosensory encoding in rat barrel cortex. Temporal analysis of single spike trains demonstrated an important role for timing of the first spike (Panzeri et al. 2001), with a smaller role for subsequent multispike patterns. Analysis of multichannel data demonstrated the practicality of the approach for studying coding by correlated activity across neurons, initially with a limited temporal analysis (Panzeri et al. 1999) and later with a full temporal analysis (Petersen et al. 2001).

Compression Method—The entropy of a spike train can be measured by how susceptible it is to lossless data compression, via the Lempel-Ziv algorithm (Wyner and Ziv 1989; Farach et al. 1995; Kontoyiannis et al. 1998; Amigo et al. 2004). As in the direct approach, spike trains are segmented and discretized into a sequence of symbols, and no assumptions are made as to the nature of the code, or the statistical structure of spike trains.

In essence, the Lempel-Ziv algorithm seeks to compress a sequence of symbols by rewriting the sequence in terms of a hierarchy of repeating substrings. The substrings that occur frequently thus provide a characterization of the statistical structure of the neural activity. Additionally, the behavior of the compression algorithm as a function of bin width could be used to characterize the temporal precision of the code. One anticipates that this approach should be highly adept at dealing with high-order statistical patterns of spikes, such as bursts (or even runs of bursts), because the compression algorithm intrinsically seeks recursive layers of structure. Another consequence of the avoidance of an explicit estimate of spike train probabilities is that multineuronal data per se should not be an obstacle.

While in principle this approach is exact, convergence of the entropy estimates is difficult to bound and appears sensitive to the details of the compression algorithm, such as the choice of the initial dictionary of strings. Nevertheless, it can result in efficient, meaningful entropy estimates when applied to neural data (Amigo et al. 2004). Determination of algorithmic complexity (Rapp et al. 1994) is a related approach, as are the context tree methods described above.

Spectrotemporal Methods—Spectrotemporal (or time-frequency) analysis is a general exploratory method that is particularly suitable for neural data, both spiking and continuous (Mittra and Pesaran 1999). It is not typically considered an information-theoretic tool, but we mention it here because it also can be used to identify meaningful statistical structure in spike trains.

Spectrotemporal analysis is a natural extension of spectral analysis. Spectral analysis formally requires that the signals to be analyzed be “stationary” (i.e., have statistical properties that do not change in time). Neural signals, especially those influenced by external stimuli, do not have this property; rather, this evolution in time may be specifically of interest. The straight-forward way to deal with this problem is simply to segment the data into periods that are sufficiently brief so that within each period, the signals can be assumed stationary. Standard spectral analysis applied to each segment can then reveal how the frequency characteristics of a signal evolve over time. As is well known, the length of the analysis segment and the achievable frequency resolution limit are reciprocally related. Sophisticated spectrotemporal techniques based on multitaper estimates (Thomson 1982; Mittra and Pesaran 1999) and wavelets (Schiff et al. 1994; Quiroga et al. 2001), while of course unable to circumvent limits on simultaneous resolution in time and frequency, represent a principled way to approach them.

Spectrotemporal analysis can identify stimulus-dependent changes in neural activity that would escape ordinary averaging techniques, such as event-related synchronization and desynchronization (Pfurtscheller and Andrew 1999). Spectral analysis has a natural extension to the multichannel context: calculation of coherences (or cross-spectra) between channels that characterize their correlations within each frequency band. Spectrotemporal analysis has a directly analogous extension, which provides a description of how the coherence between signals evolves over time. The phase relationships between activities in different channels (e.g., different neurons or field potentials in different brain regions) provide another way to identify the direction of information transfer. The frequency bands at which coherence is present can suggest how information is transferred. For example (Schiff et al. 2000, 2001), coherence between activities in distant cortical areas and between cortex and thalamus is present at particular frequency bands at specific times during a behavioral task, and is correlated with behavioral performance.

Another contact with information-theoretic approaches is that regions of the time-frequency spectrum can be used as classifiers of the neural response (Jarvis and Mitra 2001). Under fairly general assumptions, the logs of the power in nonoverlapping regions of a time-frequency spectrum are approximately independently-distributed Gaussian variables. Thus, reduction of a set of responses into measures of power in multiple time-frequency regions can serve as a first step in calculation of transmitted information. The amount of information, as well as the time-frequency regions that are critical in transmitting it, can thus be readily determined. Note that this approach to estimating information not only exploits the continuity of time, but also the intuition that neural coding is smooth in the frequency domain.

Wavelet methods (Schiff et al. 1994; Tallon et al. 1995; Quiroga et al. 2001) and multitaper methods, in essence, are complementary strategies for parceling the spectrotemporal domain into rectangular tiles. In multitaper methods, the tiles are uniform, and thus optimized for detecting features of a given temporal duration or frequency bandwidth. In contrast, wavelets tile the spectrotemporal domain with regions whose dimensions are reciprocally related, and thus optimized for detecting features whose durations and bandwidths have a given ratio.

Surrogate Datasets—Since many hypotheses concerning neural coding can be phrased in terms of comparisons between the observed data and surrogate datasets, procedures for surrogate data generation are important adjuncts to the procedures described above. The use of surrogate datasets for testing hypotheses concerning the dynamics of continuous neurophysiologic data is widely appreciated (Theiler et al. 1991; Schiff et al. 1996; Theiler and Rapp 1996). The approach is at least as relevant to testing and refining hypotheses concerning information transmission in spike trains.

Shuffling

Perhaps the simplest hypothesis that one might want to test is whether the amount of information in an experimental dataset is nonzero. As mentioned above, analytic estimates of the bias in information estimates are available. However, these estimates may not be applicable for at least two reasons: the asymptotic regime may not be reached because the dataset size is too small, or, the analysis method (e.g., the metric space approach) does not treat each response independently. But even in these circumstances, use of shuffled datasets can determine whether the estimated amount of information, viewed as a nonparametric measure of correlation between input and output, is greater than chance (Victor and Purpura 1996b).

For multichannel datasets, additional simple surrogate datasets are useful. To determine whether correlations between responses can be explained on the basis of common driving by a stimulus, rather than neuronal interconnections, the “shift-predictor,” or more generally the

“shuffle-correction” (Perkel et al. 1967), can be used. Here, the individual channels of the responses to a particular stimulus are regrouped to form surrogate responses to that stimulus.

Maximum-Entropy Methods: Single Neurons

For continuous signals, it is often of interest to determine whether observed dynamical features of a neural signal are fully explained by its second-order correlation properties. If so, then the signals are consistent with a (perhaps multichannel) Gaussian white noise that has been linearly filtered. If not, nonlinear dynamics must be present. This kind of question can be addressed by reanalyzing surrogate data that are constrained to have the same second-order correlation structure as the original data, and have higher order correlations determined by the maximizing the entropy under these constraints. Such surrogate data are conveniently created by randomizing phases but preserving amplitudes (Theiler et al. 1991; Schiff et al. 1996; Theiler and Rapp 1996).

The maximum-entropy idea is readily extended to spike trains, providing natural “coordinates” for response distributions in an elegant formal framework (Amari 2001; Nakahara and Amari 2002).

This approach can be used to formalize questions related to the important notion of “temporal coding” (Théunissen and Miller 1995). Informally, “temporal coding” means that the time course of neural activity, and not just the number of spikes, carries information. Here, the term “time course” includes not only the time-dependent firing rate, but also more subtle features of the firing pattern, such as interval structure or highly reproducible “triplets” of spikes (Lestienne and Tuckwell 1997). These aspects of firing pattern can be distinguished by comparing the information-theoretic analysis of the original data with analysis of surrogate datasets that match the observed responses in terms of the time-dependent firing rate, but are otherwise unconstrained. Such surrogates are inhomogeneous Poisson processes, whose firing rate is determined by the observed poststimulus histogram, and are thus examples of constrained maximum-entropy processes.

Surrogate datasets can be further constrained to match the original data in terms of spike counts on each trial. Such datasets can easily be created by “exchange resampling” (Victor and Purpura 1996b). A further refinement constrains the interspike interval distribution as well (Oram et al. 1999). These strategies have been used to show that precisely timed triplets of spikes do not contribute to information transfer (Oram et al. 1999; Baker and Lemon 2000).

Maximum-Entropy Methods: Multiple Neurons

Application of maximum-entropy principles to analysis of multineuronal activity can lead to substantial insights. It is impossible to determine the stimulus-response distribution empirically for an entire neuronal population, since the dimensionality of this distribution is very large. However, a practical approach is to measure the individual stimulus-conditioned response probabilities of each neuron, and to assume that the full stimulus-conditioned population response distribution is its maximum-entropy extension. This approach is equivalent to approximating the stimulus-conditioned population response distribution as a product of individual stimulus-conditioned response distributions. In the retina—an important model system—the error incurred by this approximation appears to be quite small (Nirenberg et al. 2001; Nirenberg and Latham 2003).

Maximum-entropy methods can provide a compact and comprehensible representation of the correlation structure of the spontaneous activity of neuronal populations. In two recent studies (Schneidman et al. 2006; Shlens et al. 2006), maximum-entropy extension from measured pairwise correlations accounted for the bulk of high-order multineuronal correlations.

Combining these strategies (i.e., fashioning maximum-entropy distributions from a combination of stimulus-conditioned single-neuron distributions and low-order response correlations) may provide a powerful way to analyze and understand population coding.

Summary

Understanding how neurons and neural populations represent information requires a combined experimental and theoretical approach. Shannon's information theory provides the appropriate theoretical framework. In the Shannon approach, no assumptions are made concerning the relationships of the coding elements to each other, or to the objects being represented. This generality is a fundamental aspect of the strength and elegance of the Shannon approach. However, its generality also engenders challenges to its use in experimental neuroscience, for two reasons. First, neural activity is characterized by a wide range of time scales, from the submillisecond range (e.g., the intrinsic precision of spike generation) to times on the order of a second (e.g., inhibitory synaptic potentials). Thus, without the imposition of additional hypotheses as to the nature of the code, the number of codes that need to be explored is far too great for a direct experimental attack. Second, the relationship of the neural activity to the objects being represented *is* of interest. This relationship is important to understand the mechanism of coding, and because neural activity must not only convey information but also manipulate it.

These considerations provide both a (retrospective) rationale for, and a unified view of, many approaches that have recently been advanced for the analysis of neural coding. The approaches described here vary in the assumptions made concerning the neural code, ranging from virtually no assumptions, to merely exploiting the continuity of time, to positing very specific forms for the relationship between coding elements. Making such assumptions allows analysis to be carried out on datasets that are typically available from experiments. By assuming that the codes have structure, these approaches also allow for identification of a systematic relationship between the objects and the code; i.e., a representation. However, imposition of assumptions necessarily increases the risk that the relevant neural codes are simply not being considered. At present, neuroscientists can grapple with this problem by exploring a variety of approaches, each with its own set of assumptions, and hoping that the biological conclusions are relatively independent of the methodology chosen. It remains to be seen whether a more systematic and fundamentally satisfying theoretical approach can be fashioned.

Acknowledgments

The author thanks Daniel Gardner, Simon Schultz, Chip Levy, Alex Casti, and Keith Purpura for helpful comments and suggestions. This work is supported in part by NIH EY9314 to JV and MH68012 to Daniel Gardner.

References

- Abbott LF, Varela JA, Sen K, Nelson SB. Synaptic depression and cortical gain control. *Science* 1997;275:220–224. [PubMed: 8985017]
- Abeles M, Prut Y. Spatio-temporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Physiology Paris* 1996;90:249–250.
- Amari S-I. Information geometry on hierarchy of probability distributions. *IEEE Transactions on Information Theory* 2001;47:1701–1711.
- Amigo JM, Szczepanski J, Wajnryb E, Sanchez-Vives M. Estimating the entropy rate of spike trains via Lempel-Ziv complexity. *Neural Computation* 2004;16:717–736. [PubMed: 15025827]
- Aronov D. Fast algorithm for the metric-space analysis of simultaneous responses of multiple single neurons. *Journal of Neuroscientific Methods* 2003;124:175–179.
- Aronov D, Reich DS, Mechler F, Victor JD. Multidimensional representation of spatial phase in V1. *Investigations in Ophthalmology and Vision Science* 2001;42:405.

- Baker SN, Lemon RN. Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance levels. *Journal of Neurophysiology* 2000;84:1770–1780. [PubMed: 11024069]
- Berry MJ, Warland DK, Meister M. The structure and precision of retinal spike trains. *Proceedings of the National Academy of Sciences USA* 1997;94:5411–5416.
- Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D. Reading a neural code. *Science* 1991;252:1854–1857. [PubMed: 2063199]
- Bliss TV, Collingridge GL. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 1993;361:31–39. [PubMed: 8421494]
- Bourne H, Nicoll R. Molecular machines integrate coincident synaptic signals. *Cell* 72/Neuron 1993;10 (Suppl):65–85.
- Buracas GT, Zador AM, DeWeese MR, Albright TD. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 1998;20:959–969. [PubMed: 9620700]
- Carlton AG. On the bias of information estimates. *Psychological Bulletin* 1969;71:108–109.
- Chao A, Shen TJ. Nonparametric estimate of Shannon's index of diversity when there are unseen species in a sample. *Environmental and Ecological Statistics* 2003;10:429–443.
- Chee-Orts MN, Optican LM. Cluster method for analysis of transmitted information in multivariate neuronal data. *Biological Cybernetics* 1993;69:29–35. [PubMed: 8334188]
- Cline H. Coincidence detection in the nervous system. *Trends in Neurosciences* 1997;19:566–567.
- Dan Y, Alonso JM, Usrey WM, Reid RC. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neuroscience* 1998;1:501–507.
- de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W. Reproducibility and variability in neural spike trains. *Science* 1997;275:1805–1808. [PubMed: 9065407]
- Di Lorenzo PM, Victor JD. Taste response variability and temporal coding in the nucleus of the solitary tract of the rat. *Journal of Neurophysiology* 2003;90:1418–1431. [PubMed: 12966173]
- Efron, B. *The Jackknife, the Bootstrap and Other Resampling Plans*. SIAM; Philadelphia: 1982.
- Efron, B.; Tibshirani, RJ. *Monographs on Statistics and Applied Probability*. Vol. 57. Chapman and Hall/CRC Press; Boca Raton, FL: 1998. *An Introduction to the Bootstrap*; p. 436
- Farach, M.; Noordewier, M.; Savari, S.; Shepp, L.; Wyner, A.; Ziv, J. On the entropy of DNA: Algorithms and measurements based on memory and rapid convergence; *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms*; 1995; p. 48-57.
- Gawne TJ. The simultaneous coding of orientation and contrast in the responses of V1 complex cells. *Experimental Brain Research* 2000;133:293–302.
- Grassberger P. Finite sample corrections to entropy and dimension estimates. *Physics Letters A* 1988;128:369–373.
- Grassberger P, Procaccia I. Measuring the strangeness of strange attractors. *Physica D* 1983;9:189–208.
- Gray CM, Singer W. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences the USA* 1989;86:1698–1702.
- Hirata Y, Mees AI. Estimating topological entropy via a symbolic data compression technique. *Physical Review E* 2003;67(2 Pt 2):026205.
- Jarvis MR, Mitra PP. Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Computation* 2001;13:717–749. [PubMed: 11255566]
- Johnson DH, Gruner CM, Baggerly K, Seshagiri C. Information-theoretic analysis of neural coding. *Journal of Computational Neuroscience* 2001;10:47–69. [PubMed: 11316339]
- Kennel MB, Shlens J, Abarbanel HD, Chichilnisky EJ. Estimating entropy rates with Bayesian confidence intervals. *Neural Computation* 2005;17:1531–1576. [PubMed: 15901407]
- Kontoyiannis I, Algoet PH, Suhov YM, Wyner AJ. Nonparametric entropy estimation for stationary processes and random fields, with applications to English text. *IEEE Transactions on Information Theory* 1998;44:1319–1327.
- Kozachenko LF, Leonenko NN. Sample estimate of the entropy of a random vector. *Problemy Peredachi Informatsii* 1987;23:9–16.
- Kraskov A, Stogbauer H, Grassberger P. Estimating mutual information. *Physical Review E* 2004;69(6 Pt 2):066138.

- Kreiman G, Krahe R, Metzner W, Koch C, Gabbiani F. Robustness and variability of neuronal coding by amplitude-sensitive afferents in the weakly electric fish *eigenmannia*. *Journal of Neurophysiology* 2000;84:189–204. [PubMed: 10899196]
- Krichevsky R, Trofimov V. The performance of universal coding. *IEEE Transactions on Information Theory* 1981;27:199–207.
- Lestienne R, Tuckwell HC. The significance of precisely replicating patterns in mammalian CNS spike trains. *Neuroscience* 1997;82:315–336. [PubMed: 9466445]
- Levy, WB. Experiences, thoughts, and conjectures on implementing a Lempel-Ziv-type algorithm to measure information in a spike train; *Neural Information Processing Systems Workshop on Information and Statistical Structure in Spike Trains*; Breckenridge, CO. December 1-2, 2000; 2000.
- London M, Schreibman A, Hausser M, Larkum ME, Segev I. The information efficacy of a synapse. *Nature Neuroscience* 2002;5:332–340.
- Luck SJ, Chelazzi L, Hillyard SA, Desimone R. Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology* 1997;77:24–42. [PubMed: 9120566]
- Markram H, Lubke J, Frotscher M, Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 1997;275:213–215. [PubMed: 8985014]
- Marmarelis, PZ.; Marmarelis, VZ. *Computers in Biology and Medicine*. Plenum; New York: 1978. Analysis of physiological systems: The white-noise approach; p. 487
- McClurkin JW, Optican LM, Richmond BJ, Gawne TJ. Concurrent processing and complexity of temporally encoded neuronal messages in visual perception. *Science* 1991;253:675–677. [PubMed: 1908118]
- McFadden JA. The entropy of a point process. *Journal of the Society for Industrial and Applied Mathematics* 1965;13:988–994.
- Meister M, Lagnado L, Baylor DA. Concerted signaling by retinal ganglion cells. *Science* 1995;270:1207–1210. [PubMed: 7502047]
- Mel BW. Synaptic integration in an excitable dendritic tree. *Journal of Neurophysiology* 1993;70:1086–1101. [PubMed: 8229160]
- Middlebrooks JC, Clock AE, Xu L, Green DM. A panoramic code for sound location by cortical neurons. *Science* 1994;264:842–844. [PubMed: 8171339]
- Miller, GA. Note on the bias on information estimates. In: Quastler, H., editor. *Information Theory in Psychology: Problems and Methods*. Vol. II-B. Free Press; Glencoe, IL: 1955. p. 95–100.
- Mitra PP, Pesaran B. Analysis of dynamic brain imaging data. *Biophysical Journal* 1999;76:691–708. [PubMed: 9929474]
- Nakahara H, Amari S. Information-geometric measure for neural spikes. *Neural Computation* 2002;14:2269–2316. [PubMed: 12396564]
- Nemenman I, Bialek W, de Ruyter van Steveninck R. Entropy and information in neural spike trains: Progress on the sampling problem. *Physical Review E* 2004;69(5 Pt 2):056111.
- Nirenberg S, Carcieri SM, Jacobs AL, Latham PE. Retinal ganglion cells act largely as independent encoders. *Nature* 2001;411:698–701. [PubMed: 11395773]
- Nirenberg S, Jacobs A, Fridman G, Latham P, Douglas R, Alam N, Prusky G. Ruling out and ruling in neural codes. *Journal of Vision* 2006;6:889a.
- Nirenberg S, Latham PE. Decoding neuronal spike trains: how important are correlations? *Proceedings of the National Academy of Sciences USA* 2003;100:7348–7353.
- Optican LM, Richmond BJ. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: III. Information theoretic analysis. *Journal of Neurophysiology* 1987;57:162–178. [PubMed: 3559670]
- Oram MW, Wiener MC, Lestienne R, Richmond BJ. Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *Journal of Neurophysiology* 1999;81:3021–3033. [PubMed: 10368417]
- Paninski L. Estimation of entropy and mutual information. *Neural Computation* 2003;15:1191.
- Panzeri S, Petersen RS, Schultz SR, Lebedev M, Diamond ME. The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron* 2001;29:769–777. [PubMed: 11301035]

- Panzeri S, Schultz SR. A unified approach to the study of temporal, correlational, and rate coding. *Neural Computation* 2001;13:1311–1349. [PubMed: 11387048]
- Panzeri S, Schultz SR, Treves A, Rolls ET. Correlations and the encoding of information in the nervous system. *Proceedings of the Royal Society of London B* 1999;266:1001–1012.
- Perkel DH, Gerstein GL, Moore GP. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophysical Journal* 1967;7:419–440. [PubMed: 4292792]
- Petersen RS, Panzeri S, Diamond ME. Population coding of stimulus location in rat somatosensory cortex. *Neuron* 2001;32:503–514. [PubMed: 11709160]
- Pfurtscheller G, Andrew C. Event-related changes of band power and coherence: methodology and interpretation. *Journal of Clinical Neurophysiology* 1999;16:512–519. [PubMed: 10600019]
- Quiroga RQ, Rosso OA, Basar E, Schürmann M. Wavelet entropy in event-related potentials: A new method shows ordering of EEG oscillations. *Biological Cybernetics* 2001;84:291–299. [PubMed: 11324340]
- Rapp PE, Zimmerman ID, Vining EP, Cohen N, Albano AM, Jimenez-Montano MA. The algorithmic complexity of neural spike trains increases during focal seizures. *Journal of Neuroscience* 1994;14:4731–4739. [PubMed: 8046447]
- Reich DS, Mechler F, Victor JD. Formal and attribute-specific information in primary visual cortex. *Journal of Neurophysiology* 2001a;85:305–318. [PubMed: 11152730]
- Reich DS, Mechler F, Victor JD. Independent and redundant information in nearby cortical neurons. *Science* 2001b;294:2566–2568. [PubMed: 11752580]
- Reich DS, Mechler F, Victor JD. Temporal coding of contrast in primary visual cortex: When, what, and why. *Journal of Neurophysiology* 2001c;85:1039–1041. [PubMed: 11247974]
- Reinagel P, Reid RC. Temporal coding of visual information in the thalamus. *Journal of Neuroscience* 2000;20:5392–5400. [PubMed: 10884324]
- Reynolds JH, Pasternak T, Desimone R. Attention increases sensitivity of V4 neurons. *Neuron* 2000;26:703–714. [PubMed: 10896165]
- Richmond BJ, Optican LM. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: II. Quantification of response waveform. *Journal of Neurophysiology* 1987;57:147–161. [PubMed: 3559669]
- Rieke, F.; Warland, D.; de Ruyter van Steveninck, R.; Bialek, W. *Spikes: Exploring the Neural Code*. MIT Press; Cambridge, MA: 1997.
- Rissanen, J. *Stochastic Complexity in Statistical Inquiry*. World Scientific; Singapore: 1989.
- Rodriguez E, George N, Lachaux JP, Martinerie J, Renault B, Varela FJ. Perception's shadow: long-distance synchronization of human brain activity. *Nature* 1999;397:430–433. [PubMed: 9989408]
- Roelfsema PR, Lamme VA, Spekreijse H. Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nature Neuroscience* 2004;7:982–991.
- Salzman CD, Newsome WT. Neural mechanisms for forming a perceptual decision. *Science* 1994;264:231–237. [PubMed: 8146653]
- Samonds JM, Bonds AB. From another angle: Differences in cortical coding between fine and coarse discrimination of orientation. *Journal of Neurophysiology* 2004;91:1193–1202. [PubMed: 14614106]
- Samonds JM, Zhou Z, Bernard MR, Bonds AB. Synchronous activity in cat visual cortex encodes collinear and cocircular contours. *Journal of Neurophysiology* 2006;95:2602–2616. [PubMed: 16354730]
- Schiff ND, Kalik SF, Purpura KP. Episodic dynamics of cortical processing in the ventral stream during free-viewing: Analysis of local field potentials in striate/extrastriate and inferotemporal cortices. *Society for Neuroscience Abstracts* 2000;26:1199.
- Schiff ND, Kalik SF, Purpura KP. Sustained activity in the central thalamus and extrastriate areas during attentive visuomotor behavior: Correlation of single unit activity and local field potentials. *Society for Neuroscience Abstracts* 2001;27:1910.
- Schiff SJ, Aldroubi A, Unser M, Sato S. Fast wavelet transformation of EEG. *Electroencephalography and Clinical Neurophysiology* 1994;91:442–455. [PubMed: 7529683]

- Schiff SJ, So P, Chang T, Burke RE, Sauer T. Detecting dynamical interdependence and generalized synchrony through mutual prediction in a neural ensemble. *Physical Review E* 1996;54:6708–6724.
- Schlens J, Kennel M, Abarbanel H, Chichilnisky EJ. Estimating information rates in neural spike trains with confidence intervals. *Neural Computation*. 2006
- Schneidman E, Berry MJ II, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 2006;440:1007–1012. [PubMed: 16625187]
- Schultz SR, Panzeri S. Temporal correlations and neural spike train entropy. *Physical Review Letters* 2001;86:5823–5826. [PubMed: 11415367]
- Schurmann T, Grassberger P. Entropy estimation of symbol sequences. *Chaos* 1996;6:414–427. [PubMed: 12780271]
- Sellers P. On the theory and computation of evolutionary distances. *SIAM Journal of Applied Mathematics* 1974;26:787–793.
- Sen K, Jorge-Rivera JC, Marder E, Abbott LF. Decoding synapses. *Journal of Neuroscience* 1996;16:6307–6318. [PubMed: 8815910]
- Shadlen MN, Newsome WT. The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience* 1998;18:3870–3896. [PubMed: 9570816]
- Shannon, CE.; Weaver, W. *The Mathematical Theory of Communication*. University of Illinois Press; Urbana: 1949.
- Shlens, J.; Field, GD.; Gauthier, JL.; Grivich, MI.; Petrusca, D.; Sher, A.; Litke, AM.; Chichilnisky, EJ. Probing the structure of multi-neuron firing patterns in the primate retina using maximum entropy methods. *CoSyNe*; Salt Lake City, UT: 2006.
- Shlens J, Kennel MB, Abarbanel HD, Chichilnisky EJ. Estimating information rates with confidence intervals in neural spike trains. *Neural Computation*. 2006in press
- Slepian D. On bandwidth. *Proceedings of the IEEE* 1976;64:292–300.
- Softky W. Sub-millisecond coincidence detection in active dendritic trees. *Neuroscience* 1994;58:13–41. [PubMed: 8159289]
- Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience* 1993;13:334–350. [PubMed: 8423479]
- Solomonoff R. A formal theory of inductive inference: Part I. *Information and Control* 1964;7:1–22.
- Stopfer M, Bhagavan S, Smith BH, Laurent G. Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 1997;390:70–74. [PubMed: 9363891]
- Strong SP, Koberle R, Ruyter van Steveninck RR, Bialek W. Entropy and information in neural spike trains. *Physical Review Letters* 1998;80:197–200.
- Tallon C, Bertrand O, Bouchet P, Pernier J. Gamma-range activity evoked by coherent visual stimuli in humans. *European Journal of Neuroscience* 1995;7:1285–1291. [PubMed: 7582101]
- Theiler, J.; Galdrikian, B.; Longtin, A.; Farmer, J. Testing for nonlinearity in time series: The method of surrogate data. *Los Alamos National Laboratory Preprint*; 1991. LA-UR-91-3343
- Theiler J, Rapp PE. Re-examination of the evidence for low-dimensional, nonlinear structure in the human electroencephalogram. *Electroencephalography and Clinical Neurophysiology* 1996;98:213–222. [PubMed: 8631281]
- Théunissen F, Miller JP. Temporal encoding in nervous systems: A rigorous definition. *Journal of Computational Neuroscience* 1995;2:149–162. [PubMed: 8521284]
- Théunissen F, Roddey JC, Stufflebeam S, Clague H, Miller JP. Information theoretic analysis of dynamical encoding by four identified primary sensory interneurons in the cricket cercal system. *Journal of Neurophysiology* 1996;75:1345–1364. [PubMed: 8727382]
- Thomson DJ. Spectrum estimation and harmonic analysis. *Proceedings of the IEEE* 1982;70:1055–1096.
- Treves A, Panzeri S. The upward bias in measures of information derived from limited data samples. *Neural Computation* 1995;7:399–407.
- Usrey WM, Reppas JB, Reid RC. Paired-spike interactions and synaptic efficacy of retinal inputs to the thalamus. *Nature* 1998;395:384–387. [PubMed: 9759728]
- van der Togt C, Kalitzin S, Spekrijse H, Lamme VA, Super H. Synchrony dynamics in monkey V1 predict success in visual detection. *Cerebral Cortex* 2006;16:136–148. [PubMed: 15843628]

- Victor JD. Binless strategies for estimation of information from neural data. *Physical Review E* 2002;66:51903.
- Victor JD. Spike train metrics. *Current Opinion in Neurobiology* 2005;15:585–592. [PubMed: 16140522]
- Victor JD, Purpura KP. Nature and precision of temporal coding in visual cortex: A metric-space analysis. *Journal of Neurophysiology* 1996a;76:1310–1326. [PubMed: 8871238]
- Victor JD, Purpura KP. Nature and precision of temporal coding in visual cortex: A metric-space analysis. *Journal of Neurophysiology* 1996b;76:1310–1326. [PubMed: 8871238]
- Victor JD, Purpura KP. Metric-space analysis of spike trains: theory, algorithms and application. *Network* 1997;8:127–164.
- Willems FMJ, Shtarkov YM, Tjalkens TJ. The context-tree weighting method: basic properties. *IEEE Transactions on Information Theory* 1995;41:653–664.
- Wolpert DH, Wolf DR. Estimating functions of probability distributions from a finite set of samples. *Physical Review E* 1995;52:6841–6854.
- Wu M, David SV, Gallant J. Complete functional characterization of sensory neurons by system identification. *Annual Review of Neuroscience* 2006;29:477–505.
- Wyner AD, Ziv J. Some asymptotic properties of entropy of a stationary ergodic data source with applications to data compression. *IEEE Transactions on Information Theory* 1989;35:1250–1258.

Table 1

Characteristics of several methods of information-theoretic analysis of neural data

Method	Stimulus Types		Response Types			Conceptual Aspects			
	Few, Discrete Presentation	Many, Continuous Presentation	Formal	Spikes	Continuous Signals	Multiple Channels	Assumed Nature of Code	Response Space Construction	Code Description
Direct	Yes	Yes	Discrete	Yes	No	Unfavorable	Virtually none	No	No
Metric Space	Yes	No	Point process	Yes	No	Very favorable	Specific models	Yes	Yes
Embedding	Yes	Maybe	Function of time	Yes	Yes	Favorable	Continuous	Yes	Minimal
Context Tree	No	Yes	Discrete	Yes	No	Favorable	Context tree	No	Minimal
Principal Components	Yes	Maybe	Function of time	Yes	Yes	Favorable	Rate envelope	Yes	Yes
Reconstruction	Maybe	Yes	Function of time	Yes	Yes	Favorable	Volterra	No	Yes
Power Series	Yes	Maybe	Point process	Yes	No	Favorable	Order-by-order	No	Yes
Compression	No	Yes	Discrete	Yes	No	Very favorable	Virtually none	No	Minimal
Spectrotemporal	Yes	Yes	Function of time	Yes	Yes	Favorable	Continuous (time and frequency)	No	Yes

Victor