in: Neural Codes and Distributed Representations
Foundations of Neural Computation
Edited by L. Abbott and T. Sejnowski
MIT Press (1999)

Introduction

Laurence F. Abbott
Brandeis University, Waltham, Massachusetts 02254 U.S.A.

Terrence J. Sejnowski
Howard Hughes Medical Institute and The Salk Institute, University of California, San Diego, La Jolla, CA 92037 U.S.A.

Research into how neurons represent the world is proceeding rapidly. Most of the pioneering studies of sensory and motors systems were based on single unit recordings. Although our current understanding of the response properties of neurons in the brain is largely based on this approach, there is likely to be additional information that can only be accessed by examining the firing patterns of large populations of neurons. New experimental techniques, such as simultaneous recordings of more than 100 neurons (Gray et al., 1995; Meister, 1996), provide a glimpse of large-scale population codes and distributed representations in different parts of the brain. The papers in this volume are a selection of those that have appeared in *Neural Computation* and represent a sample of diverse approaches to the problem of neuronal coding. This introduction gives some background for these papers as well as additional references for further reading.

Neural Coding

The study of neural coding is a central issue in the investigation of nervous system function. Before we can understand how neural circuits process information, we must understand how they represent it. Information is conveyed to and processed within the brain primarily in the form of action potentials. In the retina, for example, all of the information about the spatiotemporal pattern of photoreceptor activity is ultimately encoded in the action potential firing patterns of ganglion cells whose axons project though the optic nerve to the thalamus. How does the sequence of action potentials fired by a neuron, or a group of neurons, represent the information that is encoded and conveyed to other neurons? The papers collected here address this basic question from a variety of directions and with a range of viewpoints.

Action potentials are typically all-or-none stereotyped waveforms, so all the information represented by a sequence of action potentials is encoded in their timing. The neural code is particularly complex because both the properties of a stimulus (light, touch or sound intensity, for example), and the way that the stimulus changes over time are encoded in the temporal pattern of neural spiking. Information about stimulus value and about changes in the stimulus are both represented in the time domain by the neural code. As a result, it is important to distinguish between the structure imposed on the code by the nature of the stimulus, and the structure that arises from the nature of the encoding process itself.

The papers in this volume cover three broad issues in neural coding: neuronal response variability, the nature of the neural code, and the encoding of information by populations of neurons. A closely related issue is neural decoding, the process by which information encoded in a population of neurons is read out and deciphered. A final section on temporal sequences of spikes provides insights into how dynamical interactions in populations of neurons can be modeled.

Neuronal Response Variability

The responses of most neurons, even when the same stimulus is presented repeatedly, show a high degree of variability. Response variability has been characterized in a number of cortical areas (Werner and Mountcastle, 1965; Tomoko and Crapper, 1974; Schiller et al., 1976; Hegglund and Albus, 1978; Dean, 1981; Tolhurst et al., 1981 and 1983; Vogels et al., 1989; Snowden et al., 1992; Britten et al., 1993; O'Keefe et al., 1997, Gur et al., 1997). Several descriptive stochastic models of response variability have been developed (see for example Gerstein and Mandelbrot, 1964; Stein, 1967; Teich, 1989; and the review in Tuckwell, 1988). The simplest and most widely used model describes the generation of neuronal spike trains as a Poisson process. In a Poisson model each action potential is generated independently with a probability per unit time proportional to the firing rate of the neuron. The Poisson description provides a fairly good first approximation of the statistics of many cortical spike trains evoked by unchanging stimuli.

A recent study indicates that spike trains recorded from area MT in awake monkeys are more accurately described by a Poisson model than those from anaesthetized animals, which tend to show higher variability (O'Keefe, Bair, and Movshon, 1997). Two common features that are not described by a Poisson model are bursts of spikes (although the occurrences of the bursts can sometime be described by a Poisson model) and refractory effects following action potential firing (Bair, Koch, Newsome, and Britten, 1994). Refractory effects can be included with a relatively minor extension of the basic model, and, of course, more elaborate stochastic descriptions can be constructed as well (Berry and Meister, 1998). The Poisson firing rate can itself be a stochastic variable and may vary rapidly on a scale comparable to the typical interspike interval (Sejnowski, 1976; Buracas et al., 1998).

While stochastic models may describe neuronal response variability, they do not account for it in any mechanistic way. The basic spike generation process in neurons appears to be precise and reliable (Mainen and Sejnowski, 1995; Holt, Softky, Koch, and Douglas, 1996; Nowak et al., 1997; Tang et al., 1997; de Ruyter van Steveninck et al., 1997), and thus it cannot account for the variability in the spike sequences evoked by repeated presentations of the same stimulus. Response variability must arise instead from fluctuations in the input that drives neuronal firing. Softky and Koch (1992) (see also Softky and Koch, 1994) point out that this view is difficult to reconcile with the fact that neurons integrate large numbers of synaptic inputs (on the order of 10,000 for a cortical pyramidal cell). By the central limit theorem, one would expect that the variability arising from summing such a large number of inputs would be much smaller than the variability for any single afferent. As Softky and Koch note, such a high degree of variability is inconsistent with the idea that neurons integrate large numbers of excitatory synaptic inputs over any reasonable length of time.

Troyer and Miller (1997) and Bugmann et al. (1997) show how a highly variability response can be generated in a neuron model that integrates a large number of inputs. The basic idea is that the firing of the neuron must not be due to the average level of synaptic input, for this quantity obeys a central limit theorem and thus has greatly reduced fluctuations. Rather, we must imagine that the average level of excitation is subthreshold for action potential firing, and that spiking is the result of supratheshold fluctuations in the total synaptic drive. This idea has been implemented by balancing the excitatory and inhibitory inputs to the neuron (Amit and Tsodyks, 1992; Tsodyks and Sejnowski, 1995; Shadlen and Newsome, 1994, 1998; van Vreeswijk and Sompolinsky, 1996). The rational for balancing excitation with inhibition is that a neuron near threshold is extremely sensitive to small fluctuations of the membrane potential, such as those be caused by correlated inputs (Sejnowski, 1976), and it can respond to these extremely rapidly (Tsodyks and Sejnowski, 1995; Troyer and Miller, 1997).

The Nature of the Neural Code

There is an intense debate within the neuroscience community concerning whether neurons employ rate or temporal codes. Adrian (1928) first noted the relationship between neuronal firing rate and stimulus intensity that forms the basis of rate coding. There is no doubt that firing rate is a major correlate of the stimulus attributes encoded by most neurons, and that large amount of stimulus-related information are carried by firing rates (figure 1A). Nevertheless, numerous studies have attempted to go beyond firing rate to probe other types of neural coding.

As discussed in Bair and Koch (1996), the division between rate and temporal coding is based on the precision and reproducibility of spike timing. Generally, a temporal code is one in which the precise timing of spikes carries significant amounts of information about a stimulus. Of course, this requires a definition of the term "precise," and here some complications can

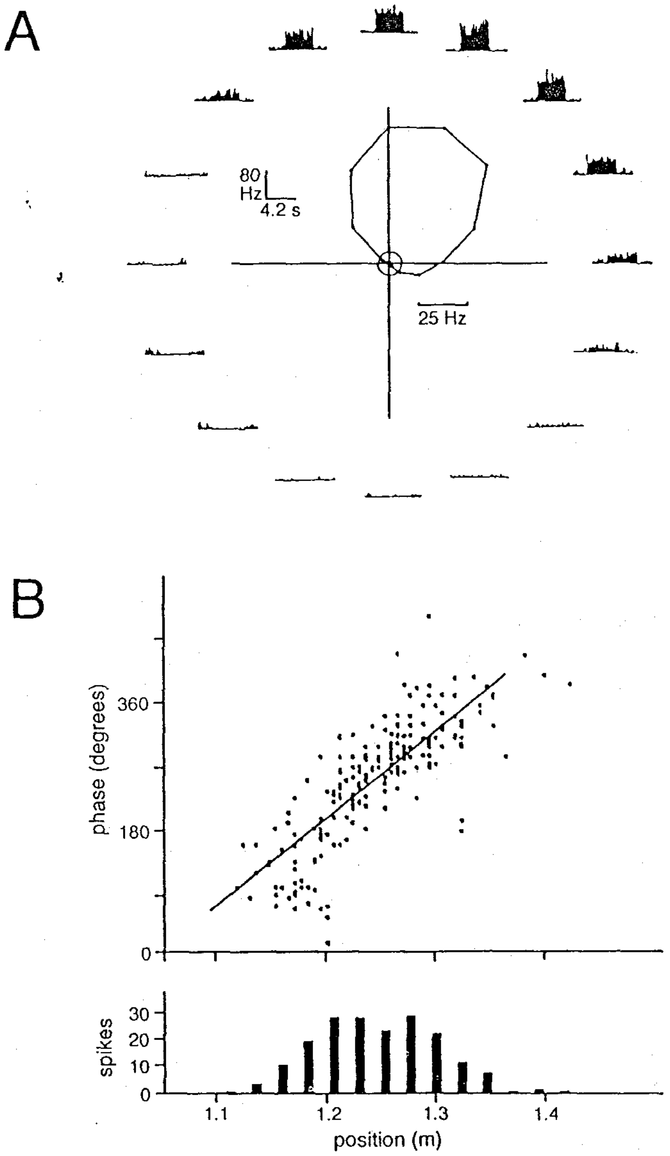

Figure 1: Rate versus correlation coding. (A) Coding of the direction of motion of a visual image by the firing rate of a neuron in area MT (medial temporal) of a monkey. The histograms show the firing rate (measured by counting spikes in discrete time bins) for motion of an image in the directions indicated around the circle. The central plot summarizes the effect of movement direction angle on firing rate (plotted radially). (From Albright 1984). (B) Correlation coding of location by hippocampal place cells. The lower histogram shows that the location of a rat is encoded in the firing rate of this hippocampal place cell. The upper panel shows that more information about location is encoded in the phase of an action potential relative the theta rhythm, a 7–12 Hz oscillation in the overall activity of hippocampal neurons. Each dot in the upper panel shows the phase of the theta rhythm plotted against the position of the animal at the time when a spike was fired. The linear relation shows that the phase of action potential firing encodes information about position. (From O'Keefe and Recce 1993.)

arise due to the dual nature that time plays in the neural code. A number of different issues are tangled up in the debate between rate and temporal codes, and confusion has arisen because the terms may mean different things to different people. The complexity of this debate is well illustrated by the fact that the data discussed in Bair and Koch (1996) within the context of a temporal code have been interpreted as representing rate coding by Shadlen and Newsome (1998).

Spike timing reflects the nature of the neural encoding process, but it is also affected by the temporal properties of the stimulus. It is not clear that a code should be identified as temporal merely because the stimulus contains precisely timed rapid onsets that evoke spikes reliability at specific times. Theunissen and Miller (1995) have proposed that a code should only be called temporal if information is carried by details of spike timing on a scale shorter than the fastest time scale characterizing variations of the stimulus. They thus suggest that the stimulus should be used to set the scale for what is meant be the term "precise timing." However, this definition is far from universally accepted.

Rate codes are as amorphously defined as temporal codes. On any single trial, the firing rate of a neuron can only be measured by examining the firing pattern of the neuron over some finite interval of time. The duration of the spike integration interval used to define the rate plays a significant role in any analysis of rate coding. A reasonable choice for the interval used to define the firing rate is the longest duration over which the stimulus can be approximated as taking a constant value. Thus, the rate of change of the stimulus determines the integration time used to define the firing rate. Some neurons only fire about one spike during such an interval and, as a result, they may appear ill-suited to encode and convey information through their firing rates (Rieke et al., 1997). However, the situation changes if a large number of neurons respond in this manner, because the information can then be encoded and conveyed as a population firing rate. Although individual neurons may fire only about one spike per integration time, the population of neurons can collectively fire many spikes over the same period.

Information theory provides a powerful tool for analyzing the nature and quality of a neural code (Gabbiani and Koch, 1996, and see Rieke et al., 1997 for a general discussion). A number of researchers have stressed the need to consider realistic or "natural" stimuli when considering the character and efficiency of a neural code (Barlow, 1961; Laughlin 1981; Field, 1987 and 1994; Atick, 1992; Rieke et al., 1997). Ruderman and Bialek (1992) show an interesting effect of including knowledge of the statistics of stimuli in the analysis of neural coding accuracy. Information theoretic or other measures of coding efficiency allow us to construct optimal encoding strategies and to test whether neurons use or approximate such strategies. The response characteristics of neurons in a number of different systems appear consistent with optimal coding strategies (Laughlin 1981; Atick, 1992; Rieke et al., 1997) and efficient neural codes are discussed in Levy and Baxter (1996) (see also Barlow et al., 1989; Baddeley et al., 1997).

While much of the debate surrounding the nature of the neural code has focused on rate versus temporal coding, there are other important issues to consider. A major question concerns the information carried by correlations between different action potentials fired by the same neuron, or by different neurons. The simplest assumption is that a negligible amount of information is carried by such correlations, and thus neurons code information into spikes independently of each other. This is called an "independent-spike code," and such coding is assumed in the majority of analyses of neural coding. It is important to realize that assuming independent spikes does not imply that the presence or absence of other spikes is irrelevant to the neural code. The assumption is that the presence of a second spike (either from the same or from a different neuron) does not dramatically alter the significance of the first spike with respect to the stimulus.

A code in which one spike can alter the meaning of another is called a "correlation code" to indicate that correlations between different spikes carry information. These spikes may be fired by one neuron or by different neurons. The difference between independent-spike and correlation codes can be illustrated by considering the representation of numbers by the digits 0 through 9. Our normal notation is a correlation code. In the number 12, the meaning of the digit 1 is completely altered by the presence of the 2; it signifies a 10 not a 1. If this was an "independent digit code," 12 would signify the number three. There are some, but not many, examples of correlation coding in neuroscience. Traub, Whittington, and Jeffreys (1997) discuss mechanisms that would allow neurons to encode information in the phase of their firing relative to an underlying collective rhythm (see also Hopfield, 1996; Sejnowski, 1996). Coding of information about spatial location in the hippocampus of the rat by such a mechanism has been observed by O'Keefe and Recce (1993, see figure 1B). There is also evidence that coincident spikes in retinal ganglion cells (Meister, 1996; Berry, Warland, and Meister, 1997) and the lateral geniculate nucleus (Dan, Atick, and Reid, 1996) can carry information beyond that conveyed independently by each neuron.

A neural coding strategy, especially one that involves precise temporal or coincidence discriminations, is useless unless other neurons can decipher and respond to elements of the code. Thus, the study of the biophysics of neuronal spike integration is an important aspect of research on neural coding, as illustrated in Murthy et al. (1994), Rapp, Yarom, and Segev (1992), and Mel (1992). These authors discuss the role of synchrony, the effect of background activity, and local computations performed in dendritic trees by voltage-dependent membrane conductances.

Cortical neurons must respond in some useful way to action potentials received across thousands of synapses at rates ranging from a zero to several hundred Hz. If these afferent spike sequences are uncorrelated and the neuron simply averages its synaptic inputs, it is difficult to see how a neuron could differentiate one set of afferent spike trains from another. One solution to this dilemma is that the neuron performs local computations involving smaller numbers of synaptic inputs on parts of its dendritic tree (Mel, 1992). Another is that correlations between sets of afferent spike trains play an important role in generating responses. For example, sets of afferents with highly synchronized action potentials might be particularly effective at generating a postsynaptic response (Murthy and Fetz, 1994). Short-term synaptic depression can make a postsynaptic neuron more sensitive to correlated inputs by reducing the impact of the average firing rates (Markram and Tsodyks, 1996; Abbott et al., 1997). It is rather sobering that, at present, we have no clear consensus on what aspects of the collective presynaptic activity actually make a cortical neuron fire.

Population Coding

In most nervous systems, information about a stimulus is encoded in the activity of a large number of responding neurons. Population coding is reviewed in Konishi (1991) (see also Knudsen et al., 1987; Churchland and Sejnowski, 1992; Abbott, 1994). There has been considerable discussion of the theoretical issues involved in interpreting out the information encoded by the activities of a large neuronal population, including Sanger (1994); Snippe (1996); Pouget et al. (1998); Buonomano and Mauk (1994); and others (Baldi and Heiligenberg, 1988; Paradiso, 1988; Altes, 1989; Lehky and Sejnowski, 1990; Vogels, 1990; Zhang and Miller, 1991; Snippe and Koenderink, 1992; Zohary, 1992; Seung and Sompolinsky, 1993; Földiák, 1993; Salinas and Abbott, 1994; Sanger, 1996; Abbott and Dayan, 1998; Zhang et al., 1998; Oram et al., 1998). Much of this work has focused on how the information encoded in the activity of a population of neurons can be readout or decoded. Decoding is an effective way of determining the accuracy, efficiency, and information carrying capacity of a neural code (Rieke et al., 1997).

Two systems, the superior colliculus, involved in the coding of saccadic eye movements (Van Gisbergen et al., 1987; Lee et al., 1988), and neurons in the primary motor cortex of the monkey encoding the direction of reaching movements of the arm (Humphrey et al., 1970; Georgopoulos et al., 1982, 1986, 1988; Kalaska et al., 1989; Caminiti et al., 1991; Scott and Kalaska, 1995; Sanger, 1994; Lin, Si, and Schwartz, 1997; Snippe, 1996) have played an important role in our understanding of population coding. Individual neurons in a population responding to a stimulus or active during a particular motor task are selective to particular attributes of the stimulus or task. For example, as discussed in Sanger (1994), certain neurons in the motor cortex of the monkey fire at rates that are correlated with the direction of the movement of the monkey's arm (Schwartz et al., 1988). The average firing rate of a given neuron is maximal when the arm movement is in a particular direction known as the preferred direction for that neuron. The vector decoding scheme reconstructs the direction of arm movement from a set of neuronal firing rates by summing the preferred direction vectors weighted by the firing rate of each neuron. If enough neurons are included in the sum, this decoding method produces results in fairly good agreement with the actual arm movement directions (Georgopoulos et al., 1982, 1986, and 1988). Vector decoding has been applied successfully to a number of other systems (Gilbert and Wiesel, 1990; Steinmetz et al., 1987; Van Gisbergen et al., 1987; Lee et al., 1988; Young and Yamane, 1992; Salinas and Abbott, 1994; Lewis and Kristan, 1998; Abbott, 1998) and it is used in several papers in this volume that consider temporal sequences and discrimination (Lukashin and Georgopoulos, 1994; Blum and Abbott, 1996; Buonomano and Mauk, 1994). The accuracy and efficiency of vector and related "center of mass" decoding schemes is analyzed in Snippe (1996).

Many researchers have studied the problem of reconstructing a motor response or a stimulus attribute from the activities of a number of neurons using more efficient methods than the simple vector decoding scheme (Snippe, 1996; Paradiso, 1988; Snippe and Koenderink, 1992; Seung and Sompolinsky, 1993; Földiák, 1993; Salinas and Abbott, 1994; Sanger, 1996; Abbott and Dayan, 1998; Zhang et al., 1998; Oram et al., 1998). Lin, Si, and Schwartz (1997) use a self-organizing feature map for this purpose. Other approaches use Bayesian or maximum likelihood arguments to reconstruct the stimulus from the neural response in a way that is optimal according to some predefined criterion. These methods are particularly useful for establishing limits on the accuracy of a neural code. The Fisher information, which is the variance of the maximum likelihood estimate of a stimulus, plays an especially important role because it sets a limit for any unbiased estimator of a stimulus based on the neural responses it evokes (Paradiso, 1988; Seung and Sompolinsky, 1993; Abbott and Dayan, 1998; Zhang et al., 1998).

Establishing a theoretical limit on the accuracy of a neural code is interesting, but it may be irrelevant if there are no biophysically reasonable schemes for implementing an optimal, or near optimal, decoding method using real neural circuitry. Pouget et al. (1998) show how maximum likelihood decoding, which achieves the maximum possible accuracy for any unbiased estimator, can be performed using the highly recurrent architecture of cortical circuits, and thus establishes that the theoretical limit corresponding to the Fisher information is achievable. The posterior probability distribution given by a Bayesian analysis can be used for efficient probabilistic reconstruction. Zhang et al. (1998) show how a feedforward network with one layer of weights could in principle read out a Bayesian code. These results show that optimal decoding is within the capability of the network mechanisms known to exist in the cortex. However, the explicit readout of a population code may not be needed until a motor response is generated, since projections between cortical areas may simply perform transformations between different population codes (Salinas and Abbott, 1995; Snippe, 1996; Lehky and Sejnowski, 1999).

What are the advantages of encoding information across large neuronal populations? The average response tuning curves in such coding arrays can provide a set of basis functions for approximating other functions of stimulus parameters (Girosi and Poggio, 1990; Poggio and Girosi, 1990; Pouget and Sejnowski, 1995). This feature can be used to do some interesting and highly relevant computations; for example, coordinate transformation needed to convert information about an object's location expressed in a retinal coordinate system to body-centered coordinates for visually guided reaching tasks (Zipser and Andersen, 1988; Salinas and Abbott, 1995; Pouget and Sejnowski, 1997). Population codes can simultaneously represent information about several different stimulus attributes, and it may have broader interpretations as well (Hinton, 1992; Anderson, 1994; Anderson and Van Essen, 1994). Anderson (1994) interpreted population codes as providing an estimate of an entire probability distribution, and Zemel, Dayan, and Pouget (1998) explore this idea in an alternative formulation.

Temporal Sequences

Much of the work on neural coding by populations of neurons, both theoretical and experimental, has focused on static behavior. One of the great challenges we face in understanding neural information processing is extending our ability to measure, describe and model network behavior into the dynamic and temporal domains. Several papers in this volume provide examples. Lukashin and Georgopoulos (1994) and Blum and Abbott (1996) discuss how temporal sequences can be learned and recalled by a neural network (see also Minai and Levy, 1993; Montague and Sejnowski, 1994; Abbott and Blum, 1996; Gerstner and Abbott, 1997). Griniasty and Tsodyks and Amit (1993) show how temporal sequences that occur during learning are reflected in correlations among the attractors of a neural network (see also Amit et al., 1994). Buonomano and Mauk (1994) consider the ability of a neural circuit to store and recall information about time intervals between events.

Sources

Bair, W. and Koch, C. (1996) Temporal position of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Computation* 8(6):1185–1202.

Blum, K. I. and Abbott, L. F. (1996) A model of spatial map formation in the hippocampus of the rat. *Neural Computation* 8(1):85–93.

Bugmann, G., Christodoulou, C., and Taylor, J. G. (1997) Role of temporal integration and fluctuation detection in the highly irregular firing of a leaky integrator neuron model with partial reset. *Neural Composition* 9(5):985–1000.

Buonomano, D. V and Mauk, M. D. (1994) Neural network model of the cerebellum: Temporal discrimination and the timing of motor responses. *Neural Computation* 6(1):38–55.

Gabbiani, F and Koch, C. (1996) Coding of time-varying signals in spike trains of integrate-and-fire neurons with random threshold. *Neural Computation* 8(1):44-66.

Griniasty, M., Tsodyks, M. V., and Amit, D. J. (1993) Conversion of temporal correlations between stimuli to spatial correlations between attractors. *Neural Computation* 5(1):1–17.

Konishi, M. (1991) Deciphering the brain's codes. *Neural Computation* 3(1):1–18.

Levy, W. B. and Baxter, R. A. (1996) Energy-efficient neural codes. *Neural Computation* 8(3):531–543.

Lin, S., Si, J., and Schwartz, A. B. (1997) Self-organization of firing activities in monkey's motor cortex: Trajectory computation from spike signals. *Neural Computation* 9(3):607–621.

Lukashin, A. V., Georgopoulos, A. P. (1994) A neural network for coding of trajectories by time series of neuronal population vectors. *Neural Computation* 6(1):19–28.

Mel, B. (1992) NMDA-based pattern discrimination in a modeled cortical neuron. *Neural Computation* 4(4):502–517.

Murthy, V. N. and Fetz, E. E. (1994) Effects of input synchrony on the firing rate of a three-conductance cortical neuron model. *Neural Computation* 6(6):1111–1126.

Pouget, A., Zhang, Z., Deneve, S., and Latham, P. (1998) Statistically efficient estimation using population coding. *Neural Computation* 10(2):373–401.

Rapp, M., Yarom, Y., and Segev, I. (1992) The impact of parallel fiber background activity on the cable properties of cerebellar purkinje cells. *Neural Computation* 4(4):518–533.

Ruderman, D. L. and Bialek, W. (1992) Seeing beyond the Nyquist limit. *Neural Computation* 4(5):682–690.

Sanger, T. D. (1994) Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation* 6(1):29–37.

Snippe, H. P. (1996) Parameter extraction from population codes: A critical assessment. *Neural Computation* 8(3):511–529.

Softy, W. P. and Koch, C. (1992) Cortical cells should fire regularly, but do not. *Neural Computation* 4(5):643–646.

Traub, R. D., Whittington, M. A., and Jefferys, J. G. R. (1997) Gamma oscillation model predicts intensity coding by phase rather than frequency. *Neural Computation* 9(6):1251–1264.

Troyer, Todd W. and Miller, Kenneth D. (1997) Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Computation* 9(5):971–983.

Zemel, R., Dayan, P., and Pouget, A. (1998) Probabilistic interpretation of population codes. *Neural Computation* 10(2):403–430.

References

Abbott, L. F. (1994) Decoding neuronal firing and modeling neural networks. *Quarterly Review of Biophysics* 27:291–331.

Abbott, L. F. (1998) The analytical bend of the leech. *Nature* 391:18–19.

Abbott, L. F. and Blum, K. I. (1996) Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex* 6:406–416.

Abbott, L. F. and Dayan, P. (1998) The Effect of correlated variability on the accuracy of a population code. *Neural Computation*.

Abbott, L. F., Varela, J. A., Sen, K., and Nelson, S. B. (1997) Synaptic depression and cortical gain control. *Science* 275:220–224.

Adrian, E. D. (1928) *The Basis of Sensation: The Action of the Sense Organs*. New York: W. W. Norton.

Albright, T. D. (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology* 52:1106–1129.

Altes, R. A. (1989) Ubiquity of hyperacuity. *Journal of the Acoustical Society of America* 85:943–952.

Amit, D., Brunel, N., and Tsodyks, M. (1994) Correlations of cortical Hebbian reverberations: Experiment versus theory. *Journal of Neuroscience* 14:6435–6445.

Amit, D. and Tsodyks, M. (1992) Effective neurons and attractor neural networks in cortical environments. *Network* 3:121–137.

Anderson, C. H. (1994) Basic elements of biological computational systems. *International Journal of Modern Physics* C 5:135–137.

Anderson, C. H. and Van Essen, D. C. (1994) Neurobiological computational systems. In *Computational Intelligence Imitating Life*. New York: IEEE Press, 213–222.

Atick, J. J. (1992) Could information theory provide an ecological theory of sensory processing? *Network: Computation in Neural Systems* 3:213–251.

Baddeley, R., Abbott, L. F., Booth, M. J. A., Sengpiel, F., Freeman, T., Wakeman, E. A., and Rolls, E. T. (1997) Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society of London* B264:1775–1783.

Bair, W. and Koch, C. (1996) Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Computation* 8(6):1185–1202. Reprinted in this volume.

Bair, W., Koch, C., Newsome, W. T., and Britten, K. H. (1994) Power spectrum analysis of bursting cells in area MT in the behaving monkey. *Journal of Neuroscience* 14:2870–2892.

Baldi, P. and Heiligenberg, W. (1988) How sensory maps could enhance resolution through ordered arrangements of broadly tuned receivers. *Biological Cybernetics* 59:313–318.

Barlow, H. B. (1961) Possible principles underlying the transformation of sensory messages. In W. A. Rosenblith (ed.), *Sensory Communication*. Cambridge, MA: MIT Press.

Barlow, H. B., Kaushal, T. P., and Mitchison, G. J. (1989) Finding minimum entropy codes. *Neural Computation* 1:412–423.

Berry, M. J., Warland, D. K., and Meister, M. (1997) The Structure and precision of retinal spike trains. *Proceedings of the National Academy of Sciences of the United States of America* 94:5411–6.

Berry M. J., Meister, M. (1998) Refractoriness and neural precision. *Journal of*

Neuroscience 18:2200–2211.

Blum, K. I. and Abbott, L. F. (1996) A Model of spatial map formation in the hippocampus of the rat. *Neural Computation* 8(1):85–93. Reprinted in this volume.

Britten, K. H., Shadlen, M., Newsome, W. T., and Movshon, J. A. (1993) Response of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience* 10:1157–1169.

Bugmann, G., Christodoulou, C., and Taylor, J. G. (1997) Role of temporal integration and fluctuation detection in the highly irregular firing of a leaky integrator neuron model with partial reset. *Neural Computation* 9(5):985–1000. Reprinted in this volume.

Buonomano, D. V. and Mauk, M. D. (1994) Neural network model of the cerebellum: Temporal discrimination and the timing of motor responses. *Neural Computation* 6(1):38–55. Reprinted in this volume.

Buracas, G., Zador, A., DeWeese, M., and Albright, T. (1998) Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20:959–969.

Caminiti, R., Johnson, P. B., Galli, C., Ferraina, S., and Burnod, Y. (1991) Making arm movements within different parts of space: The premotor and motor cortical representations of a coordinate system for reaching to visual targets. *Journal of Neuroscience* 11:1182–1197.

Churchland, P. S. and Sejnowski, T. J. (1992) *The Computational Brain*. Cambridge, MA: MIT Press.

Dan, Y., Atick, J. J., and Reid, R. C. (1996) Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *Journal of Neuroscience* 16:3351–3362.

Dean, A. (1981) The Variability of discharge of simple cells in the cat striate cortex. *Experimental Brain Research* 44:437–440.

Field, D. J. (1987) Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America* A4:2379–2394.

Field, D. J. (1994) What is the goal of sensory coding? *Neural Computation* 6:559.

Földiák, P. (1993) The "Ideal homunculus": Statistical inference from neural population responses. In F. H. Eeckman and J. Bower (eds.), *Computation and Neural Systems*. Norwell, MA: Kluwer Academic Publishers, 55–60.

Gabbiani, F. and Koch, C. (1996) Coding of time-varying signals in spike trains of integrate-and-fire neurons with random threshold. *Neural Computation* 8(1):44–66. Reprinted in this volume.

Gerstein, G. L. and Mandelbrot, B. (1964) Random walk models for the spike activity of a single neuron. *Biophysical Journal* 4:41–68.

Gerstner, W. and Abbott, L. F. (1997) Learning navigational maps through potentiation and modulation of hippocampal place cells. *Journal of Computational Neuroscience* 4:79–94.

Georgopoulos, A. P., Kalaska, J. K., Caminiti, R., and Massey, J. T. (1982) On the relations between the directions of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience* 2:1527–1537.

Georgopoulos, A. P., Kettner, R. E., and Schwartz, A. B. (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space, part II: Coding of the direction of movement by a neuronal population. *Neuroscience* 8:2928–2937.

Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986) Neuronal population coding of movement direction. *Science* 243:1416–1419.

Gilbert, C. D. and Wiesel, T. N. (1990) The Influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research* 30:1689–1701.

Girosi, F. and Poggio, T. (1990) Networks and the best approximation property. *Biological Cybernetics* 63:169–176.

Gray, C. M., Maldonado, P. E., Wilson, M., and McNaugton, B. (1995) Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *Journal of Neuroscience Methods* 63:43–54.

Griniasty, M., Tsodyks, M. V., and Amit, D. J. (1993) Conversion of temporal correlations between stimuli to spatial correlations between attractors. *Neural Computation* 5(1):1–17. Reprinted in this volume.

Gur, M., Beylin, A., and Snodderly, D. M. (1997) Response variability of neurons in primary visual cortex (V1) of alert monkeys. *Journal of Neuroscience* 17:2914–2920.

Hegglund, P. and Albus, K. (1978) Response variability and orientation discrimination of single cells in striate cortex of cat. *Experimental Brain Research* 30:275–296.

Hinton, G. E. (1992) How neural networks learn from experience. *Scientific American* 267:145–151.

Holt, G. R., Softky, G. W., Koch, C., and Douglas, R. J. (1996) Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *Journal of Neurophysiology* 75:1806–1814.

Hopfield, J. J. (1996) Transforming neural computations and representing time. *Proceedings of the National Academy of Sciences of the United States of America* 93:15440–15444.

Humphrey, D. R., Schmidt, E. M., and Thompson, W. D. (1970) Predicting measures of motor performance from multiple corticle spike trains. *Science* 170:758–761.

Kalaska, J. F., Cohen, D. A. D., Hyde, M. L., and Prud'Homme, M. A. (1989) A Comparison of movement direction-related versus load direction-related activity of primate motor cortex using a two-dimensional reaching task. *Journal of Neuroscience* 9:2080–2102.

Knudsen, E., DuLac, S., and Esterly, S. D. (1987) Computational maps in the brain. *Annual Review of Neuroscience* 10:41–65.

Konishi, M. (1991) Deciphering the brain's codes. *Neural Computation* 3(1):1–18.

Laughlin, S. B. (1981) A Simple coding procedure enhances a neuron's information capacity. Z. *Naturf.* C. 36:910–912.

Lee, C., Rohrer, W. H., and Sparks, D. L. (1988) Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature* 332:357–360.

Lehky, S. R. and Sejnowski, T. J. (1990) Neural model of stereoacuity and depth interpolation based on a distributed representation of stereodisparity. *Journal* *of Neuroscience* 10:2281–2299.

Lehky, S. R. and Sejnowski, T. J. (1999) Seeing white: Qualia in the context of decoding population codes. *Neural Computation,* 1261–1280.

Levy, W. B. and Baxter, R. A. (1996) Energy efficient neural codes. *Neural Computation* 8(3):531–543. Reprinted in this volume.

Lewis, J. E. and Kristan, W. B. (1998) A Neuronal network for computing population vectors in the leech. *Nature* 391:76–79.

Lin, S., Si, J., and Schwartz, A. B. (1997) Self-organization of firing activities in monkey's motor cortex: Trajectory computation from spike signals. *Neural Computation* 9(3):607–621. Reprinted in this volume.

Lukashin, A. V. and Georgopoulos, A. P. (1994) A Neural network for coding of trajectories by time series of neuronal population vectors. *Neural Computation* 6(1):19–28. Reprinted in this volume.

Mainen, Z. F. and Sejnowski, T. J. (1995) Reliability of spike timing in neocortical neurons. *Science* 268:1503–1506.

Markram, H. and Tsodyks, M. (1996) Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382:807–810.

Meister, M. (1996) Multineuronal codes in retinal signaling. *Proceedings of the National Academy of Sciences of the United States of America* 93:609–614.

Mel, B. (1992) NMDA-based pattern ciscrimination in a modeled cortical neuron. *Neural Computation* 4(4):502–517. Reprinted in this volume.

Minai, A. A. and Levy, W. B. (1993) Sequence learning in a single trial. *INNS World Congress on Neural Networks* II:505–508.

Montague, P. R. and Sejnowski, T. J. (1994) The Predictive brain: Temporal coincidence and temporal order in synaptic learning mechanisms. *Learning and Memory* 1:1–33.

Murthy, V. N. and Fetz, E. E. (1994) Effects of input synchrony on the firing rate of a three-conductance cortical neuron model. *Neural Computation* 6(6):1111–1126. Reprinted in this volume.

Nowak, L. G., Sanchez-Vives, M. V., and McCormick, D. A. (1997) Influence of low and high frequency inputs on spike timing in visual cortical neurons. *Cerebral Cortex* 7:487–501.

O'Keefe, J. and Recce, M. L. (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.

O'Keefe, L. P., Bair, W., and Movshon, J. A. (1997) Response variability of MT neurons in macaque monkey. *Society for Neuroscience Abstracts* 23:1125.

Oram, M. W., Földiák, P., Perrett, D. I., and Sengpiel, F. (1998) The "Ideal homunculus": decoding neural population signals. *Trends in Neuroscience* 21:259–265.

Paradiso, M. A. (1988) A Theory for the use of visual orientation information which exploits the columnar structure of striate cortex. *Biological Cybernetics* 58:35–49.

Poggio, T. and Girosi, F. (1990) Regularization algorithms for learning that are equivalent to multilayer networks. *Science* 247:978–982.

Pouget, A. and Sejnowski, T. J. (1995) Spatial representations in the parietal cortex may use basis functions. In: G. Tesauro, D. S. Touretzky, and T. K. Leen, eds., *Advances in Neural Information Processing Systems* 7:157–164. Reprinted in this volume.

Pouget, A. and Sejnowski, T. J. (1997) Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience* 9.

Pouget, A., Zhang, K., Deneve, S., and Latham, P. (1998) Statistically efficient estimation using population coding. *Neural Computation* 10(2):373–401.

Rapp, M., Yarom, Y., and Segev, I. (1992) The Impact of parallel fiber background activity on the cable properties of cerebellar purkinje cells. *Neural Computation* 4(4):518–533. Reprinted in this volume.

Rieke, F. M., Warland, D., de Ruyter van Steveninck, R., and Bialek, W. (1997) *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press.

Ruderman, D. L. and Bialek, W. (1992) Seeing beyond the Nyquist limit. *Neural Computation* 4(5):682–690. Reprinted in this volume.

Ruyter van Steveninck, R. R. de, Lewen, G. D., Strong, S. P., Koberle, R., and Bialek, W. (1997) Reproducibility and variability in neural spike trains. *Science* 275:1805–1808.

Salinas, E. and Abbott, L. F. (1994) Vector reconstruction from firing rates. *Journal of Computational Neuroscience* 1:89–107.

Salinas, E. and Abbott, L. F. (1995) Transfer of coded information from sensory to motor networks. *Journal of Neuroscience* 15:6461–6474.

Sanger, T. D. (1994) Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation* 6(1):29–37. Reprinted in this volume.

Sanger, T. D. (1996) Probability density estimation for the interpretation of neural population codes. *Journal of Neurophysiology* 76:2790–2793.

Schiller, P. H., Finlay, B. L., and Volman, S. F. (1976) Short-term response variability of monkey striate neurons. *Brain Research* 105:347–349.

Schwartz, A., Kettner, R. E., and Georgopoulos, A. P. (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space, I: Relations between single cell discharge and direction of movement. *Neuroscience* 8:2913–2927.

Scott, S. H. and Kalaska, J. F. (1995) Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *Journal of Neurophysiology* 73:2563–2567.

Sejnowski, T. J. (1976) On the stochastic dynamics of neuronal interaction. *Biological Cybernetics* 22:203–211.

Sejnowski, T. J. (1996) Pattern recognition: Time for a new neural code? *Nature* 376:21–22.

Seung, H. S. and Sompolinsky, H. (1993) Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences* 90:10749–10753.

Shadlen, M. N. and Newsome, W. T. (1994) Noise, neural codes, and cortical organization. *Current Opinion in Neurobiology* 4:569–579.

Shadlen, M. N. and Newsome, W. T. (1998) The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience* 18:3870–3896.

Snippe, H. P. (1996) Parameter extraction from population codes: A critical assessment. *Neural Computation* 8(3):511–529.

Snippe, H. P. and Koenderink, J. J. (1992) Discrimination thresholds for channel-coded systems. *Biological Cybernetics* 66:543–551. Reprinted in this volume.

Snowden, R. J., Treue, S., and Andersen, R. A. (1992) The Response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Experimental Brain Research* 88:389–400.

Softky, W. P. and Koch, C. (1992) Cortical cells should fire regularly, but do not. *Neural Computation* 4(5):643–646. Reprinted in this volume.

Softky, W. R. and Koch, C. (1994) The Highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience* 13:334–350.

Stein, R. B. (1967) Some models of neuronal variability. *Biophysical Journal* 7:37–68.

Steinmetz, M. A., Motter, B. C., Duffy, C. J., and Mountcastle, V. B. (1987) Functional properties of parietal visual neurons: Radial organization of directionalities with the visual field. *Journal of Neuroscience* 7:177–191.

Tang, A. C., Bartels, A. M., and Sejnowski, T. J. (1997) Effects of cholinergic modulation on responses of neocortical neurons to fluctuating input. *Cerebral Cortex* 7:502–509.

Teich, M. C. (1989) Fractal character of the auditory neural spike train. *IEEE Transactions on Biomedical Engineering* 36:150–160.

Theunissen, F. E. and Miller, J. P. (1995) Temporal encoding in nervous systems. A rigorous definition. *Journal of Computational Neuroscience* 2:149–162.

Tomoko, G. and Crapper, D. (1974) Neuronal variability: non-stationary responses to identical visual stimuli. *Brain Res* 79:405–418.

Tolhurst, D. J., Movshon, J. A., and Thompson, I. D. (1981) The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Experimental Brain Research* 41:414–419.

Tolhurst, D. J., Movshon, J. A., and Dean, A. F. (1983) The Statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research* 23:775–785.

Traub, R. D., Whittington, M. A., and Jefferys, J. G. R. (1997) Gamma oscillation model predicts intensity coding by phase rather than frequency. *Neural Computation* 9(6):1251–1264. Reprinted in this volume.

Troyer, T. W. and Miller, K. D. (1997) Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Computation* 9(5):971–983. Reprinted in this volume.

Tsodyks, M. and Sejnowski, T. J. (1995) Rapid switching in balanced cortical network models. *Network* 6:1–14.

Tuckwell, H. C. (1988) *Introduction to Theoretical Neurobiology*. Cambridge: Cambridge University Press.

van Gisbergen, J. A. M., Van Opstal, A. J., and Tax, A. M. M. (1987) Collicular ensemble coding of saccades based on vector summation. *Neuroscience* 21:541–555.

van Vreeswijk, C. and Sompolinsky, H. (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274:1724–1726.

Vogels, R. (1990) Population coding of stimulus orientation by cortical cells. *Journal of Neuroscience* 10:3543–3558.

Vogels, R., Spileers, W., and Orban, G. A. (1989) The Response variability of striate cortical neurons in the behaving monkey. *Experimental Brain Research* 77:432–436.

Werner, G. and Mountcastle, V. B. (1965) The Variability of central neural activity in a sensory system, and its implications for the central reflection of sensory events. *Journal of Neurophysiology* 26:958–977.

Young, M. P. and Yamane, S. (1992) Sparce population coding of faces in the inferotemporal cortex. *Science* 256:1327–1331.

Zemel, R., Dayan, P., and Pouget, A. (1998) Probabilistic interpretation of population codes. *Neural Computation* 10(2):403–430. Reprinted in this volume.

Zhang, J. and Miller, J. P. (1991) A Mathematical model for resolution enhancement in layered sensory systems. *Biological Cybernetics* 64:357–364.

Zipser, D. and Andersen, R. A. (1988) A Back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679–684.

Zhang, K., Ginzburg, I., McNaughton, B. L., and Sejnowski, T. J. (1998) Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology* 79:1017–1044.

Zohary, E. (1992) Population coding of visual stimuli by cortical neurons tuned to more than one dimension. *Biological Cybernetics* 66:265–272.