

Foreword

Neural Pulse Coding

Terrence J. Sejnowski

Neurons use action potentials to signal over long distances, as summarized in Chapter 1 by Gerstner. The all-or-none nature of the action potential means that it codes information by its presence or absence, but not by its size or shape. In this respect, an action potential can be considered a pulse. This is an important fact about how brains are built, but it is equally important as a theoretical challenge to understanding the function of the brain. How do action potentials represent sensory states? How is information contained in the firing patterns of action potentials stored and retrieved? These are old questions that have been the focus of much research, but recent advances in experimental techniques are opening new ways to test theories for how information is encoded and decoded by spiking neurons in neural systems [Rieke et al., 1997]. The papers in the collection provide a window into the current state of theoretical and computational thinking based on spikes.

Spike Timing

The timing of spikes is already well established as a means for coding information in the electrosensory system of electric fish [Heiligenberg, 1991], in the auditory system of echolocating bats [Kuwabara, 1993], and in the visual system of flies [Bialek et al., 1991]. The relative spike timing between auditory inputs to the brain can be used to determine the location of a sound source [Agmon-Snir, Carr and Rinzel, 1998] and Chapter 14 by Gerstner et al. provides a model for how the pulse timing could be learned. There are many possible ways that spike timing could be used in the nervous system, and papers in the collection explore some of the theoretical possibilities. Does the timing of spikes mean anything in cerebral cortex? If so, then there must be a sophisticated system in the cortex to organize the coding and decoding of spike timing.

Although some of the earliest theoretical models of neural networks, such as the McCulloch-Pitts model based on binary units, captured the all-or-none character of signaling by action potentials, many of the recent network models rely on continuous variables such as the average firing rate and input-output functions such as sigmoids. Although the average firing

rate for a single neuron is well defined under stationary conditions and over a long enough time period to achieve statistical significance, none of these conditions obtain in the brain during most behaviors. Visual object recognition takes around 100 ms and the motor system has an even faster time scale. As the time interval for averaging gets shorter, the coding becomes discrete and a more general statistical framework is needed that takes into account joint probability distributions in the population of neurons. One exciting possibility is that the timing of the spikes in a population of neurons represents the probability distributions directly [Anderson, 1994; Zhang et al., 1998; Zhang and Sejnowski, 1999; Hinton, personal communication]. A concrete example will be given below for neurons in the rat hippocampus.

Another issue that is closely tied to spike coding is the issue of spike reliability. Averages of spike trains are presented because of the variability observed in recordings from neurons in the central nervous system. For example, the spike trains of sensory neurons elicited by the same stimulus presented repeatedly display a high degree of variability [Schiller et al., 1976; Britten et al., 1993; O'Keefe et al., 1997; Gur et al., 1997]. For cortical neurons, the intervals between spikes has an approximately Poisson distribution in response to a constant stimulus [O'Keefe, Bair and Movshon, 1997]. The Poisson firing rate can itself be a stochastic variable and may vary rapidly on the scale of the interspike interval [Sejnowski, 1976; Buracas et al., 1998]. This modulated Poisson model is a good compromise between models that are entirely based on average quantities like the firing rate and spike timing variables.

There is growing evidence that the variability observed in cortical spike trains cannot be attributed to unreliability in spike initiation [Mainen and Sejnowski, 1995; Berry et al., 1997; Nowak et al., 1997; Tang, 1997; de Ruyter van Steveninck et al., 1997]. Another possible explanation is that the variability reflects fluctuation in the inputs. This might occur if the dynamic operating point of a neuron is kept near its threshold, which is a region where a neuron is most sensitive to input correlations. This condition can be maintained 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]. When the inputs to a neuron are balanced, its firing rate can be increased by increasing the variance of the fluctuations without increasing the net excitation to the neuron.

In thinking about the distinctions between rate codes and pulse codes, it is important to keep in mind that the time a spike occurs can be used both to represent external time-varying stimuli and internal states. Often, many trials are presented and the poststimulus time (PST) histogram is computed to average out the variability, a form of ensemble averaging. On any given trial, however, the precise timing of a spike is an additional degree of freedom that could be used for many different purposes. For example, spike timing could be used to encode additional information about the sensory stimulus [Dan et al., 1997; Berry et al., 1997; de Ruyter van Steveninck et al., 1997].

The relative timing of spikes in a population of neurons could also encode information in addition to that conveyed by each spike train independently [Ritz and Sejnowski, 1997]. It has been suggested that the synchronous firing of neurons in a population can carry information about the global significance of the stimulus for the animal [Gray et al., 1989] or to organize information together in packets [Jeffreys, Traub and Whittington, 1996]. Learning may also depend on relative spike timing [Stopfer et al., 1997]. Changing the relative timing of presynaptic and postsynaptic spikes in a cortical neuron by as little as 10 msec can determine whether a synapse is potentiated or depressed [Markram et al., 1997]. This suggests that the relative timing of spikes may be highly regulated in order to control the conditions when synaptic plasticity occurs.

A neural coding strategy that depends on precise temporal coincidences requires precise mechanisms for decoding the temporal code. The biophysics of neuronal spike integration, especially in dendrites, is therefore as important as the neural coding [Murthy et al., 1994; Mel, 1992]. Local computation in dendritic trees is accomplished by active membrane conductances. Cortical neurons receive information at thousands of synapses at rates ranging from zero to several hundred hertz (Hz). It has been suggested that a neuron performs computations involving smaller numbers of synaptic inputs on localized regions of its dendritic tree [Mel, 1992]. Correlations between sets of afferent spike trains could also 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 et al., 1994]. Short-term synaptic depression could make a postsynaptic neuron more sensitive to correlated inputs by reducing the impact of the average firing rates of the inputs [Markram and Tsodyks, 1996; Abbott et al., 1997]. Furthermore, single synapses can be quite unreliable, since on average an excitatory synapse in the cortex or hippocampus releases less than a single vesicle in response to a stimulus. Thus, decoding schemes are likely to be probabilistic. The consequences of synaptic unreliability for pulse coded systems is explored by Maass and Zador in Chapter 12.

Population Codes

A sensory stimulus gives rise to action potentials in a large number of cortical neurons, which represent different aspects of the stimulus. The central question how different aspects of the stimulus are represented in the population of responding neurons [Lehky and Sejnowski, 1990; Konishi, 1991; Seung and Sompolinsky, 1993; Foldiak, 1993; Abbott, 1994; Salinas and Abbott, 1994; Sanger, 1996; Abbott and Dayan, 1999; Zhang et al., 1998]. New experimental techniques, such as simultaneous recordings from the spike trains of over 100 neurons [Gray et al., 1995; Meister, 1996], provide a glimpse of the large-scale population codes in different parts of the brain. The goal is to decipher the coding schemes by recording spike trains from a large number of neurons and try to reconstruct from those spike trains physical and psychological properties of the stimulus. In the case of the

motor system, the goal is to predict the movement of the animal from the pattern of spikes recorded from the motor system.

Many neurons in the motor cortex of the monkey fire at rates that are correlated with the direction of arm reaching [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. A popular reconstruction method is called the population vector scheme, which estimates the direction of arm movement by summing the preferred direction vectors weighted by the firing rate of each neuron (Georgopoulos et al., 1988). A similar coding strategy has been found among the interneurons responsible for the bending reflex of the leech (Lewis and Kristan, 1998; Abbott, 1998). A more general approach to reconstruction is to allow the neurons to represent more general basis functions of the physical variables [Girosi and Poggio, 1990; Pouget and Sejnowski, 1997; Zhang et al., 1998]. Each neuron contributes a basis function in this space of variables whenever it fires, and the best estimate of the physical variables can be computed from the sum of these functions weighted by the number of spike occurring in each neuron during a time interval.

An alternative method for decoding a population code is based on Bayesian reconstruction and maximum likelihood estimation. These are probabilistic methods that take into account prior probabilities and attempt to reconstruct the entire probability distribution. Instead of adding together the kernels, as in the basis function method, the probabilistic approach multiplies them, assuming that the spikes in each neuron are independent. This method can be used to establish limits on the accuracy of a neural code. The Fisher information, which is the variance of the maximum likelihood estimate of a stimulus, 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, 1999; Zhang et al., 1998].

Although these reconstruction methods may be useful for telling us what information could be available in a population of neurons, it does not tell us what information is actually used by the brain. In particular, it is not clear whether these reconstruction methods could be implemented with neurons. Pouget [Pouget et al., 1998] show how maximum likelihood decoding 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. Zhang [Zhang et al., 1998] show how a feedforward network with one layer of weights could in principle read out a Bayesian code. Thus, optimal decoding is well within the capability of the network mechanisms known to exist in the cortex. However, an explicit readout of a population code may not be needed until the final common pathway of the motor system since projections between cortical areas may simply perform transformations between different population codes.




Figure 1. True X and Y positions of a rat running on an elevated figure-8 maze as compared with the positions reconstructed by different methods using 25 place cells and a sliding time window of 0.5 sec. The same 60-second segment is shown in all plots. The probabilistic or Bayesian methods were especially accurate, and the erratic jumps in the reconstructed trajectory were reduced by a continuity constraint using information from 2 consecutive time steps. By contrast, the population vector method often yielded implausible positions. From [Zhang et al., 1998].

Hippocampal Place Fields

An example of how the timing of spikes in a population of neurons can be used to reconstruct a physical variable is the reconstruction of the location of a rat in its environment from the place fields of neurons in the hippocampus of the rat. In the experiment reported here, the firing patterns of 25 cells were simultaneously recorded from a freely moving rat [Zhang et al., 1998]. The place cells were silent most of the time, and they fired maximally only when the animal's head was within restricted region in the environment called its place field [Wilson et al., 1993]. The reconstruction problem was to determine the rat's position based on the spike firing times of the place cells.




Figure 2. All reconstruction methods became more accurate when more cells were used, shown here for a rat running on a rectangular maze. Each data point represents the mean error of 40 repetitive trials in which a subset of cells were drawn randomly from the whole sample. The shaded region represents reconstruction errors excluded by the Cramér-Rao bound based on Fisher information. From [Zhang et al., 1998].

Bayesian reconstruction was used to estimate the position of the rat in the figure-8 maze shown in Figure 1 [Zhang et al., 1998]. Assume that a population of N neurons encodes several variables (x_1, x_2, \dots) , which will be written as vector \mathbf{x} . From the number of spikes $\mathbf{n} = (n_1, n_2, \dots, n_N)$ fired by the N neurons within a time interval τ , we want to estimate the value of \mathbf{x} using the Bayes rule for conditional probability:

$$P(\mathbf{x}|\mathbf{n}) = P(\mathbf{n}|\mathbf{x})P(\mathbf{x})/P(\mathbf{n}) \quad (1)$$

assuming independent Poisson spike statistics. The final formula reads

$$P(\mathbf{x}|\mathbf{n}) = kP(\mathbf{x}) \left(\prod_{i=1}^N f_i(\mathbf{x})^{n_i} \right) \exp \left(-\tau \sum_{i=1}^N f_i(\mathbf{x}) \right) \quad (2)$$

where k is a normalization constant, $P(\mathbf{x})$ is the prior probability, and $f_i(\mathbf{x})$ is the measured tuning function, i.e., the average firing rate of neuron i for each variable value \mathbf{x} . The most probable value of \mathbf{x} can thus be obtained by finding the \mathbf{x} that maximizes $P(\mathbf{x}|\mathbf{n})$, namely,

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} P(\mathbf{x}|\mathbf{n}) \quad (3)$$

By sliding the time window forward, the entire time course of \mathbf{x} can be reconstructed from the time varying-activity of the neural population.

A comparison of several different reconstruction methods is shown in Figure 1. The Bayesian reconstruction method was the most accurate. As the

number of neurons included in the reconstruction is increased, the accuracy of all the methods increased, as shown in Figure 2. The best mean error was about 5 cm, in the range of the intrinsic error of the infrared position tracking system. There are thousands of place cells in the hippocampus of a rat that respond in any given environment. However, it is not known how this information is used by the rat in solving navigational problems.

There is evidence for information in the relative timing of neurons in the rat hippocampus. As a rat moves through the preferred place of a hippocampal neuron, the timing of the spikes relative to a background 4-6 Hz theta rhythm changes from phase lag to phase lead [O'Keefe and Recce, 1993]. Thus, the relative timing of spikes in the population of hippocampal place neurons carry information about relative location (Chapter 5 by Recce provides more details about this form of phase coding). Hopfield [Hopfield, 1996] has suggested a general method by which vectors can be encoded by the relative timing of impulses and decoded by neurons using time delays.

Hardware Models

The technology that makes possible digital computers can also be used for implementing large networks of spiking neurons. Several of the chapters in this book illustrate how very large scale integrated (VLSI) technology can be used to emulate the integration in dendrites (Chapter 5 by Northmore and Elias) and recurrent cortical networks (Chapter 6 by Whatley, Deiss and Douglas). Scaling up from a single chip to a system of chips requires a complex communications system that corresponds to the major tracts that connect different brain areas. Most of the brain volume is composed of fiber bundles called white matter, which are composed of axons that carry spikes over long distances. Even within the neuropil of the cortex, a significant fraction of the volume is taken up with axons used for local communication. Thus, it should be expected that any hardware system based on the brain would make a major commitment of resources to communication.

Computing with spikes is not easy. The precision of the hardware required for most spike timing codes is much greater than that required for schemes based on average firing rates. Another important consideration is the amount of power required to generate pulses and charge up long wires. Finally, the high precision must be maintained over varying conditions which, in the case of the cortex, means a range of temperatures and ionic concentrations. Homeostatic mechanisms could be used to stabilize timing circuits, which suggests that calibration is also an important function that needs to be implemented in the cortex.

As more evidence is found for the importance of spike timing in the cortex, the question shifts from whether spike timing carries information to how it is used. This volume provides a rich source of ideas that will serve as the starting point for many research directions.

References

- [Abbott, 1994] Abbott, L. F. (1994). Decoding neuronal firing and modeling neural networks. *Quarterly Review of Biophysics*, 27:291-331.
- [Abbott, 1998] Abbott, L. F. (1998). The analytical bend of the leech. *Nature*, 391: 18-9.
- [Abbott and Dayan, 1999] Abbott, L. F., and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Computation*, 11 (1).
- [Abbott et al., 1997] Abbott, L. F., Varela, J. A., Sen, K., and Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, 275:220-224.
- [AgmonSnir et al., 1998] AgmonSnir, H., Carr, C. E., and Rinzel, J. (1998). The role of dendrites in auditory coincidence detection. *Nature*, 393: 268-272.
- [Amit and Tsodyks, 1992] Amit, D. and Tsodyks, M. (1992). Effective neurons and attractor neural networks in cortical environments. *Network*, 3:121-137.
- [Anderson, 1994] Anderson, C. H. (1994). Basic elements of biological computational systems. *International Journal of Modern Physics C*, 5:135-137.
- [Berry et al., 1997] Berry, M. J., Warland, D. K., 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.
- [Bialek et al., 1991] Bialek, W., Rieke, F., de Ruyter, R., van Steveninck, R.R. and Warland, D. Reading a neural code. *Science* 252, 1854-7 (1991).
- [Britten et al., 1993] 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.
- [Buracas et al., 1998] 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.
- [Dan et al., 1996] Dan, Y., Atick, J. J., Reid, R. C. (1996). Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *Journal of Neuroscience* 16:3351-62.
- [Foldiak, 1993] Foldiak, P. (1993). The 'ideal homunculus': Statistical inference from neural population responses. In *Computation and Neural Systems*, F. H. Eeckman and J. Bower, eds., Norwell, MA: Kluwer Academic Publishers, 55-60.
- [Georgopoulos et al., 1988] 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. II. Coding of the direction of movement by a neuronal population. *Neuroscience*, 8:2928-2937.

- [Girosi and Poggio, 1990] Girosi, F., and Poggio, T. (1990). Networks and the best approximation property. *Biological Cybernetics*, 63:169-176.
- [Gray et al., 1989] Gray, C. M., König, P., Engel, A. K., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334-337.
- [Gray et al., 1995] Gray, C. M., Maldonado, P. E., Wilson, M., McNaughton, 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.
- [Gur et al., 1997] 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.
- [Heiligenberg, 1991] Heiligenberg, W. (1991). *Neural Nets in Electric Fish*, MIT Press, Cambridge, MA.
- [Hopfield, 1996] Hopfield, J. J. (1996). Transforming neural computations and representing time. *Proceedings of the National Academy of Sciences of the United States of America*, 93:15440-4.
- [Jefferys et al., 1996] Jefferys, J. G., Traub, R. D., Whittington, M. A. (1996) Neuronal networks for induced '40 Hz' rhythms. *Trends in Neurosciences*, 19(5):202-8.
- [Konishi, 1991] Konishi, M. (1991). Deciphering the Brain's Codes. *Neural Computation*, 3(1), 1-18
- [Kuwabara and Suga, 1993] Kuwabara, N., and Suga, N. (1993). Delay lines and amplitude selectivity are created in subthalamic auditory nuclei: the brachium of the inferior colliculus of the mustached bat. *Journal of Neurophysiology* 69, 1713-1724.
- [Lehky and Sejnowski, 1990] Lehky, S. R., and Sejnowski, T. J. (1990). Neural model of stereoacuity and depth interpolation based on a distributed representation of stereo disparity. *Journal of Neuroscience*, 10:2281-2299.
- [Lewis and Kristan] Lewis, J. E., and Kristan, W. B. (1998). A neuronal network for computing population vectors in the leech. *Nature*, 391: 76-9.
- [Mainen and Sejnowski, 1995] Mainen, Z. F., and Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268:1503-1506.
- [Markram and Tsodyks, 1996] Markram, H., and Tsodyks, M. (1996). Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382: 807-10.

- [Markram et al., 1997] Markram, H., Lübke, J., Frotscher, M., and Sakman, B. (1997). Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex., *Journal of Physiology*, 500: 409-40.
- [Meister, 1996] Meister, M. (1996). Multineuronal codes in retinal signaling. *Proceedings of the National Academy of Sciences of the United States of America*, 93:609-14.
- [Mel, 1992] Mel, B., (1992). NMDA-Based Pattern Discrimination in a Modeled Cortical Neuron, *Neural Computation*, 4(4), 502-517.
- [Murthy and Fetz, 1994] 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.
- [O'Keefe and Recce, 1993] 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 et al., 1997] 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.
- [Paradiso, 1988] 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.
- [Pouget, 1997] Pouget, A. and Sejnowski, T. J. (1997). Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience* 9(2), 222-237.
- [Pouget et al., 1997] Pouget, A., Zhang, K., Deneve, S., and Latham, P. (1997). Statistically Efficient Estimation Using Population Coding, *Neural Computation*, 10: 373-401, 1998.
- [Rieke et al., 1997] Rieke, F., Warland, D., de Ruyter van Steveninck, R. R., and Bialek, W. (1997). *Spikes – Exploring the Neural Code*. MIT Press, Cambridge, MA.
- [Ritz and Sejnowski, 1997] Ritz, R., and Sejnowski, T. J. (1997). Synchronous oscillatory activity in sensory systems: new vistas on mechanisms. *Current Opinion in Neurobiology*, 7(4):536-46.
- [de Ruyter van Steveninck et al., 1997] de Ruyter van Steveninck, R. R., Lewen, G. D., Strong, S. P., Koberle, R., Bialek, W. (1997). Reproducibility and variability in neural spike trains. *Science*, 275:1805-8.
- [Salinas and Abbott, 1994] Salinas, E., and Abbott, L. F. (1994) Vector reconstruction from firing rates. *Journal of Computational Neuroscience*. 1:89-107.
- [Sanger, 1996] Sanger, T. D. (1996). Probability density estimation for the interpretation of neural population codes. *Journal of Neurophysiology*, 76:2790-2793.

- [Schiller et al., 1976] 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 et al., 1988] 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.
- [Sejnowski, 1976] Sejnowski, T. J. (1976). On the stochastic dynamics of neuronal interaction. *Biological Cybernetics*, 22:203-11.
- [Sejnowski, 1996] Sejnowski, T. J. (1996). Pattern recognition. Time for a new neural code? *Nature*, 376:21-2.
- [Seung and Sompolinsky, 1993] 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 and Newsome, 1994] Shadlen, M. N., and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4:569-579.
- [Shadlen and Newsome, 1998] 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(10):3870-96.
- [Stopfer, 1997] Stopfer, M., Bhagavan, S., Smith, B. H, Laurent, G. (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature*, 390:70-4.
- [Tang et al., 1997] Tang, A. C., Bartels, A. M., Sejnowski, T. J. (1997). Effects of cholinergic modulation on responses of neocortical neurons to fluctuating input. *Cerebral Cortex*, 7:502-9.
- [Tsodyks and Sejnowski, 1995] Tsodyks, M., and Sejnowski, T. J. (1995). Rapid switching in balanced cortical network models. *Network*, 6:1-14.
- [van Vreeswijk and Sompolinsky, 1996] van Vreeswijk, C., and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274:1724-1726.
- [Wilson and McNaughton, 1993] Wilson, M. A., and McNaughton B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055-8.
- [Zhang et al., 1998] 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*, 79:1017-44.
- [Zhang and Sejnowski, 1999] Zhang, K., and Sejnowski, T. J. (1999). Neuronal tuning: to sharpen or broaden? *Neural Computation*, 11 (1).