ANALYSIS METHODOLOGY OF SYNCHRONY IN SIMULTANEOUSLY

RECORDED SINGLE UNIT ACTIVITY IN THE VISUAL CORTEX

By

Melanie Rebecca Bernard

Thesis

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

MASTER OF SCIENCE

in

Biomedical Engineering

May, 2006

Nashville, Tennessee

Approved:

Professor A. B. Bonds, III

Professor Richard Shiavi

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

LIST OF FIGURES

If the brain were so simple we could understand it,
we would be so simple we couldn't.

--Lyall Watson (naturalist and author)

CHAPTER I

REVIEW OF NEURAL CODING WITH POPULATIONS

Introduction

Since the evolution of man, we have been guided by our senses to perceive and

respond to the myriad of sensations that describe our environment. The senses of hearing

and vision allow us to communicate with others and navigate our surroundings. The

senses of taste and smell help us locate food and alert others of sexual preparedness. The

sense of touch allows us to protect our bodies and stay clear of danger. Each one of our

senses is reliable, yet adaptable, juggling such roles for the sake of dependability as well

as allowing for processes such as learning. Each sense may work alone or in concert to

give rise to the perception of our environment. But how is sensory information

transformed into perception?

The brain, as the control center of our bodies, interprets sensory information and

motivates a perception and/or behavioral response appropriate to the environment. The

brain is able to integrate the senses, separate figure from ground, perform invariant

recognition, complete partially occluded objects, and recognize shape from coherent

motion. Modern neuroscience seeks to understand the human brain and determine how

electrochemical interactions among neurons can generate such perceptions and behaviors.

On a molecular level, sensory stimulation induces neurons to relay signals to one another

in a manner that reflects a given stimulus. Neural coding describes how these salient

stimulus features are represented in neuronal responses. A great deal has been learned

about *what* happens in the brain, yet *how* the brain encodes sensory information with

neuronal representations that govern emergent properties like perception or action remains unknown.

Synchronization of neuronal responses is an attractive candidate to play a role in generating perception. Operationally defined, synchrony is the simultaneous response of two or more neurons within some margin of error known as the integration time period. Synchrony exists between visual cortical neurons, but its functional significance is largely unknown. As a hypothetical neural substrate for encoding salient stimulus properties, synchrony enhances the probability of eliciting postsynaptic action potentials, thus ensuring propagation of this information to subsequent levels of the cortical hierarchy. The physical mechanism underlying synchrony would allow for neurons to be "effectively" connected and form dynamic regional circuits to reliably and efficiently transmit information throughout the cortex while minimizing metabolic demands.

Although synchrony could constitute a suitable encoding scheme, proving its importance in signaling is difficult since adequate methods to measure synchrony have not been developed. Current approaches quantify synchrony as a relationship between two neurons. However, synchrony allows for the formation of transient functional groups which could include tens, hundreds, thousands, or even larger numbers of neurons. Pair-wise distance calculations increase exponentially as group size increases and can be computationally exhaustive for large assemblies. Furthermore, very few studies have investigated how to measure the quality of synchrony in an assembly as well as develop some means to display these quantities.

The work presented here derives a measure for synchrony within assemblies of arbitrary size and evaluates synchrony's role as a possible neural substrate for contour integration by investigating dynamic grouping and characteristics of group membership.

The layout of this paper is as follows: in the first chapter, we will follow the development of the theory of synchrony as a possible neural code by discussing the history of single-unit coding schemes and their associated problems. Next, we will cover the subsequent evolution of multi-unit coding schemes including synchronization and oscillation. Previous research involving roles for synchronization will be summarized. Lastly, earlier approaches for quantifying population activity will be discussed. In Chapter 2, we will introduce a novel method to detect and quantify synchrony among large neural assemblies. The method has been applied to multi-electrode array recordings in the visual cortex of paralyzed and anesthetized cats and results concerning dynamic grouping, comparison of this method to the well-known Joint Peri-Stimulus Time Histogram method (Aertsen et al. 1989), and characteristics for group membership are discussed. Finally, the third chapter explores future applications involving natural images and implicating synchrony as a sparse code for natural vision.

Neural Coding

Encoding Sensory Information with Single Units

The function of the brain has been the focus of scientific investigation for thousands of years. In his book, Bear (2001) describes the history of brain study and how its function was eventually realized through the collective works of many researchers. Early thinkers postulated that the heart, not the brain, was the center of man's life force and the seat of mental ability. In ancient Egypt, brains were removed from the deceased while the heart and other organs were preserved through the process of mummification. Ancient Greeks like Aristotle also held the belief that the heart was the seat of

intelligence, and the brain functioned merely as a radiator to cool the heart's blood. It

wasn't until 300 B.C. that Alexandrian biologists Herophilus and Erasistratus, after

dissecting a human body, concluded that the brain was the seat of thought and ability.

The Roman anatomist Galen theorized that movement of fluid through the ventricles

governed brain function. In the Age of Reason, Descartes believed fluid mechanics could

explain brain function in animals, but the higher mental abilities of humans were fulfilled

by the soul. The fluid mechanic approach to brain function was dispelled in favor of an

electrical theory developed by Galvani in the 1700s. Eventually, a modern framework

for the brain was laid out in the 1800s when the collective works of Broca, Wernicke, and

Brodmann showed that different areas of the brain controlled specific functions (Kandel,

2001).

At the turn of the 20^{th} century, the study of individual neurons in the brain became

a reality through the anatomical works of Ramon y Cajal and Golgi. The field of

neurophysiology was pioneered with a novel experimental protocol and recording

technique introduced by Adrian and Zotterman (1926). Using a capillary electrometer

and three-stage amplifier, Adrian and Zotterman recorded the impulses produced in the

plantar digital nerves of the cat when stimulated by contact or pressure. They noted that

the frequency of the impulses (action currents – related to action potentials through

Ohm's Law) varied with the intensity of the stimulus, but the magnitude of individual

impulses did not. These results supported an all-or-none relationship between stimulus

and nerve impulses.

A binary, all-or-none response to a stimulus could signal the presence or absence

of the stimulus, but what property of a cellular response represents stimulus intensity?

Since action potentials are indistinguishable events, the strength of a stimulus cannot be

reflected in the shape or size of an impulse and therefore must be reflected in the temporal characteristics of the neuronal spike train, such as average firing rate, spike count, spike patterns, interspike intervals, or precise spike arrival times. While the relationship between stimulus intensity and firing frequency described by Adrian and Zotterman remains central to our understanding of information transmission by neurons, a great number of more subtle properties have been proposed as graded signalling mechanisms. For instance, Strehler and Lestienne (1986) found distinct patterns in individual spike trains that occurred more often than chance with submillisecond precision. Also, Victor (2000) showed that different stimulus features (contrast, orientation, spatial frequency) could be represented at different temporal resolutions of interspike interval histograms. In this manner, stimulus features can be multiplexed in a spike train.

While temporal properties like spike patterns or interspike intervals have gained some support in the field as a means to encode stimulus information, traditional studies have concentrated on average firing rate. To address the relationship between the firing of single neurons and perceptual experience, Barlow (1972) proposed a single-neuron doctrine emphasizing the role of average firing rate as a stimulus-encoding mechanism. His classic view of the cardinal cell holds that individual neurons, each responsive to a particular set of local features, modulate their firing rates to reflect salient information in the visual field. The most complex features are detected via convergence up the cortical hierarchy.

While feature detection via correlation filtering by discrete receptive field organizations has inspired much work on the neural representation of visual structure, this concept ultimately fails as a foundation for neural coding. One reason is that

response gradations are limited by an action potential's refractory period. A cell therefore has a maximum firing rate, which imposes a limit on the dynamic range for encoding stimulus intensity. Furthermore, average firing rate is highly variable across stimulus repetitions, so reliable responses must be obtained by averaging across repeated stimulus presentations (Gershon et al. 1998). From a biological perspective, this cannot happen as the brain usually recognizes objects with a single viewing. Finally, the most important challenge to a rate code based on cardinal cells is that every feature requires a unique unit, and the dimensionality of the visual world overwhelms even the imposing numbers of neurons in the visual cortex.

Encoding Sensory Information With Multiple Units

Alternative theories for the neural representation of structure involve dynamic assemblies of neurons, which contribute to a population code. Due to the dynamic nature of grouping, the combinatorial possibilities of such a scheme offer a vastly increased dimensional magnitude for encoding visual information and also have implications for learning and plasticity. In 1941, Sherrington proposed that groups of neurons may cooperate synergistically such that the whole is more than the sum of its parts. Hebb (1949) and Hayek (1952) expanded on this theory and suggested that groups of cells could form dynamic regional circuits or spatiotemporal assemblies to represent structures in a visual scene. Indeed, they hypothesized that connections and interactions between neurons, defining a neural network architecture, were more functionally significant than the individual properties of the neurons themselves. In this manner, visual information would not only be inherent in the activity of individual neurons, but could be extracted from the collective activity of the group.

6

Originally, Hebb (1949) and Hayak (1952) proposed that cooperative relationships in cell assemblies were formed based on anatomical connections, perhaps defined by plasticity during brain development. However, noting the adaptive nature of the brain, Hayak (1952) also suggested that the formation of cell assemblies could result from short-term enhancement of synaptic effectiveness initiated by changes in the temporal structure of spike trains (becoming effectively connected) instead of requiring actual anatomical changes in synaptic connections. In this manner, groups of cells could assemble and disassemble during certain tasks and individual cells could belong to more than one functional group. Furthermore, uncorrelated groups could coexist without interference. The temporal binding theory (Milner 1974; von der Malsburg 1981) postulates that dynamic assembly formation is the physical basis for certain perceptual phenomena such as shape perception, figure-ground separation, long and short-term plasticity, and memory. According to this theory, perceptually-related features are linked through correlated firing among subpopulations of cells. Grouping into subpopulations is defined by perceptually-based relationships, e.g. feature proximity, similarity, or motion coherency. In this scheme, only simple feature detectors are required and complex features are extracted through the activities of multiple assemblies. Experimental support for the temporal binding theory was provided independently by Eckhorn et al. (1988) and Gray et al. (1989). In both studies, synchrony between cell pairs was found to depend on the orientation and coherence of the visual stimulus.

In the synfire chain model of the cortex (Abeles et al. 1991), precise spike arrival times from multiple neurons (synchrony) govern the processing and transmission of neural information in a network of converging and diverging connections when postsynaptic neurons act as coincidence detectors (Abeles 1982). Synchronous inputs

lower spike threshold, increasing the sensitivity to these inputs (Azouz and Gray 2000).

In theoretical simulations of the synfire chain model, Diesmann (1999) showed that

assemblies on the order of one hundred neurons could achieve submillisecond precision.

Arguments against the synfire chain model cite the unreliability of synaptic transmission

and propose instead that neurons act as integrators (Shadlen and Newsome 1994).

Confirmation of population theories presents the significant challenge of both

acquiring responses from multiple neurons and understanding the significance of

information derived from joint activity. As early as 1981, a population of 19 neurons was

recorded simultaneously in monkey visual cortex using a 30-electrode array (Kruger and

Bach 1981). In the last two decades, there have been advances in the areas of

microelectrode and tetrode arrays, and the improvements and availability of this

technology have stimulated many recent reviews (Pouget et al. 2000; Milton and Mackey

2000; Buzsaki 2004; Brown et al. 2004). Simultaneous population recordings (> pairs)

have been made in the retina (Warland et al. 1997), LGN (Mehta et al. 2000), visual

cortex (Gray et al. 1995; Nordhausen et al. 1996; Reich et al. 2001), and frontal cortex

(Abeles et al. 1993; Vaadia et al. 1995; Prut et al. 1998) and in general have shown

distributed (multi-task) levels of activity and overlapping and changing levels of

interaction that are highly dependent on stimulus parameters and behavioral states.

Synchrony as a Population Code

Synchronization/Oscillation as a Basis for "Feature Binding"?

A population code (i.e. any coding strategy involving multiple cells – not just

synchronization), requires a specific association among cells. One possibility is that cells

have correlated relative firing rates (oscillation). However, the integration of firing rate must take place over some finite time and given the required broad temporal resolution distinct populations can become confounded. On the other hand, synchronization is believed to be more effective than elevation of firing rates (Abeles 1991; Azouz and Gray 2003) in forwarding information because the transmission efficiency in the cortex is generally low (Nicoll and Blakemore 1993; Thomson and West 1993). The precise timing of synchronization allows expression of unambiguous relationships between cells in an assembly. Simulations suggest that the effective summation interval is less than 10 ms (Softky and Koch 1993), permitting rapid and dynamic assembly and disassembly as well as the coexistence of numbers of independent groups.

Exploring temporal coding began with the identification of local stimulus-specific firing correlations associated with oscillations of firing probability and field potentials (Gray and Singer 1989). Synchronous discharges were strongest when adjacent neurons with similar tuning preferences were optimally stimulated. The oscillations could be found over a large cortical distance and were enhanced by coherent stimuli and destroyed by incoherent stimuli (Gray et al. 1989). Oscillations could even retain synchrony across cortical areas (Engel et al. 1991a). Cells with different orientation preferences could synchronize to single light bars, but when presented with bars of differing orientations, cells segregated into groups according to their orientation preferences, demonstrating dynamic reorganization of assemblies as hypothesized (Engel et al. 1991b). Qualitatively similar results and conclusions, focusing mainly on oscillations, have also been reported (Eckhorn et al. 1988; Bauer et al. 1989; Eckhorn and Obermueller 1993; Eckhorn 1994).

This work has however been met with skepticism, due in part to inconsistencies. Singer's work is based mainly on field potentials or unresolved population recordings and

9

does not necessarily provide an accurate picture of single unit behavior. Much of the emphasis has been on oscillations in the gamma frequency range, but these patterns are by no means ubiquitous and are often hard to detect (Bair et al. 1994; Samonds and Bonds 2005) in recordings from single cells (aside from intrinsically oscillatory cells in layer 5; Gray and McCormick 1996). We find that oscillation is not needed to generate synchrony, although it can help to sustain it (Samonds and Bonds 2005). More recently, several laboratories have been unable to find unambiguous relationships between synchrony and specific visual tasks designed around segmentation or figure/ground discrimination (Lamme and Spekreijse 1998; Thiele and Stoner 2003; Roelfsema et al. 2004; Palanca and DeAngelis 2005), indicating that synchrony is not directly supportive of feature binding, at least in the classic sense. There is even considerable controversy over whether correlated firing (or the measures commonly used for its representation) actually adds information (e.g. Shadlen and Newsome 1998; Nirenberg et al. 2001; Petersen et al. 2001; Averbeck and Lee 2004).

These reports do not however critically examine the dependencies between temporal structure/synchrony and differences in stimulus information. We suggest that cooperation exists when the response of multiple neurons contains *emergent* information, in the form of constructive correlation that is not already represented in the individual responses of the neurons. One example of this is when the synchrony modulates to changes in stimulus features in a way that is independent of firing rate (e.g. Engel et al. 1991a,b; Kreiter and Singer 1996 (but see Palanca and deAngelis 2005); Castelo-Branco et al. 2000; Frien et al. 2000). We have found synchrony/cooperation to be especially important and useful in the cases of fine angular discriminations where firing rate is essentially constant (Samonds et al 2003, 2004).

10

Synchrony as a Substrate for Higher-Order Correlations

The successive hierarchies of the brain (as nonlinear networks) could progressively extract higher-order spatial and temporal relationships within the visual stimulus. At the level of the cortex and beyond, these features become progressively harder to detect as the higher-order correlations become more abstract and global (e.g. Purpura et al. 1994). This might explain why it is typically more difficult to find synchronization and the appropriate stimulation to identify functional properties of single cells as we move up the visual system hierarchy (Ursey and Reid 1999). Therefore, we prefer to think of synchrony as a reliable signal transmission mechanism that extracts higher-order visual relationships as a Gestalt rather than as an active structural binding mechanism that represents a secondary code within a system of simple feature extraction. The importance in the distinction between an active intracortical binding mechanism versus higher-order filtering is that the latter process need not pass every possible test of feature binding (e.g. Shadlen and Movshon 1999).

One conceptual framework that we believe to be consistent with the synchronous activity in V1 is the association field (Field et al. 1993; Hess et al. 2003), which is based on studies of the perception of contours and continuity. Association field theory predicts linking between orientation-tuned cells that is dependent on their joint relative orientation and spatial position. In natural images, contours are predominantly linear with a decreasing probability of greater curvature (Geisler et al. 2001; Sigman et al. 2001; Elder and Goldberg 2002). The probability for linking is strongest between elements with shallower relative angles and closer separations, but greater distances can support greater angular differences. What is important is the relative variance, i.e. the extent to which receptive fields are aligned along notional contours. Synchrony can overcome the

11

ambiguity of firing rate to identify salient contours, which then have the potential to be integrated at subsequent locations in the extrastriate hierarchy.

We have found experimental evidence that supports synchrony's role as a substrate for contour detection (Samonds et al. 2005). By using an experimental protocol consisting of drifting concentric rings, we found that neuron pairs with proper receptive field alignment tangent to the rings would synchronize their responses, despite wholly different orientation preferences. Synchronous responses were more reliable than changes in average firing rate in discriminating between concentric ring and grating stimuli. And as theorized above, group membership was found to be dynamic in that individual cells could belong to more than one functional group, which assembled based on the spatiotemporal properties of the stimulus.

Significance

The extraordinary progress of the last four decades in understanding the process of vision has relied extensively on our knowledge of single neurons in the visual pathways. However, properties of single cells cannot account for perceptual phenomena such as contour detection or shape recognition. The discovery of the strategy by which information is assembled by groups of cells is critical to understanding the overall function of the brain. Synchrony is an attractive candidate for participation in this process. It occurs in the olfactory, auditory, and somatosensory systems and is found in the hippocampus, frontal cortex, and motor system (for review see Engel et al. 1999). Across a population, rate coding can define but a single neural assembly, whereas given timing constraints, synchrony can support numbers of assemblies simultaneously (e.g. Singer and Gray 1995). Synchrony is a code that can be transmitted coherently through

multiple levels (Abeles 1991; Reinagel and Reid 2000) and is readily discriminated by cells acting as coincidence detectors (Azouz and Gray 2003). Encoding of information through populations of neurons defined by synchrony may be key in linking neurophysiology and perception as well as providing a crucial bridge between single unit and global population (EEG, fMRI) studies. Finally, knowing how information is encoded by groups of cells in the visual cortex has numerous applications such as vision prosthetics or understanding and treating certain visual pathologies.

Methods for Population Analysis

Introduction

By the mid-1960's, automatic data processing for neurophysiological experiments was widely available and enabled researchers to record the precise timing of spike events over long periods. In 1967, Perkel et al. introduced a variety of statistical techniques for the analysis of neuronal spike trains. Mathematical descriptions of output behavior could in turn be used to make inferences about input or cell processing behavior. Perkel and his colleagues operated under the assumption that details about the nervous system are inherent in the structure of a neuronal spike train. Since action potentials are essentially identical, all-or-none events, only the timing of such events could relay information about the processes which lead to their generation.

The principle working assumptions (Moore et al. 1966) are: 1) the precise timing of action potentials carries enormous amounts of information about the structure and function of the nervous system and it is in this timing that information is processed; 2) analysis of spike trains can provide insight into the cellular mechanisms responsible for

13

spike production, input to the cell, and the cell's transfer function; and 3) interconnections and functional interactions can be derived from the analysis of simultaneous spike trains. They also acknowledge that all neuronal processes involve a probabilistic element, which must be accounted for in quantitative descriptions of neuron behavior or models of neuron function.

PST and Shuffle Predictor

By treating spike trains as stochastic point processes, where each spike event is instantaneous and indistinguishable, a framework for quantifying spike train properties and the information inherent in the spike trains became available. In a peri-stimulus time (PST) histogram (Gerstein and Kiang 1960), effects of repeated stimulation on one spike train can be quantified. A PST histogram displays the probability of one neuron firing as a function of time around the onset of the stimulus (see Figure 1.1). Since all stimulus presentations are identical and only the timing of such presentations is of significance, a stimulus "train" can also be treated as a point process. A PST histogram is thus effectively a "cross-correlation" of the spike train and a "train" of stimulus presentations. If the stimulus had no effect on the neuron, then the probability of firing will be constant and the PST histogram will be flat. If, on the other hand the stimulus produces a time-locked evoked response, then fluctuations in the PST histogram will be seen. For instance, peaks in the histogram reflect a higher probability of firing at a certain time after the stimulus presentation and usually represent excitatory processes. Conversely, depressions in the histogram reflect a lower probability of firing and usually correspond to inhibitory or refractory processes.

To provide a control with which to statistically compare the PST histogram, the interspike intervals (elapsed time between two spikes) in a spike train can be randomly shuffled to destroy serial independence, but preserve order-independent statistics (Cox and Lewis 1966). The recomputed PST histogram, also known as the shuffle predictor, then represents chance firing events (see Figure 1.1).
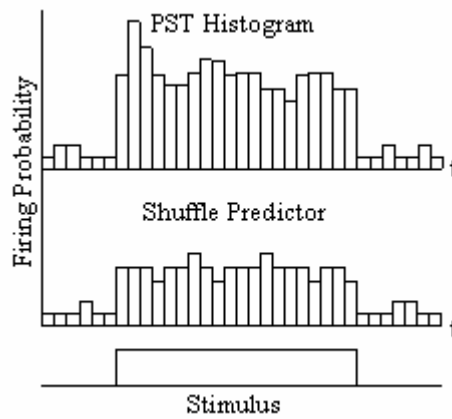
Figure 1.1: *Top*: A peri-stimulus time (PST) histogram displays the effects of repeated stimulation on one spike train. In this diagram, the firing probability of the neuron increased on average after the onset of the stimulus. This reflects an excitatory process. *Bottom*: The shuffle predictor is used to estimate the amount of chance firing events by recomputing the PST histogram after randomly shuffling the interspike intervals to destroy serial independence.

CCH and Shift Predictor

If all data are assumed stationary at all time scales in the absence of stimulation, statistical comparisons between two simultaneously recorded spike trains can reveal information about possible anatomical connections, common sources of activation, responses to stimuli, or synaptic input-output relations. However, moderate non-stationarities do not affect the detection of interactions between two spike trains (Perkel et al. 1967). When comparing spike trains from two neurons, A and B, the cross-

correlation represents the probability of encountering any spike in train B as a function of time before or after a spike in train A. A cross-correlation histogram (CCH, or cross-correlogram) can be constructed of the intervals between all train A events and all train B events, forward and backward, out to some designated time (see Figure 1.3). An observed dependence in the CCH can reflect functional interaction or common input (see Figure 1.2). Functional interaction is when the firing of one neuron affects the probability of firing of the other neuron and these effects can be synaptic (direct or indirect), ephaptic, or arise from field effects. Common input describes any mechanism that influences the firing properties of both neurons at the same time, such as input from a common presynaptic neuron or some other outside source.

Figure 1.2: Functional relationships between stimulus and two neurons A and B. *a.* Functional interaction where the stimulus affects either neuron A or neuron B, who in turn affects the firing probability of the other. *b.* Common input where the firing probabilities of both neurons are affected by the same source (the stimulus in this example). Note that these schematics depict direct relationships, but could also be indirectly activated through intermediate networks.

In the presence of stimulation, the shape of the correlogram may be influenced by the firing rate change of one or both neurons, the direct or indirect input to both cells from a common source that responds to the stimulus, the effect of the stimulus on interaction mechanisms between the cells, or any combination of these effects. It is

assumed that effects produced by the stimulus are independent and additive. However, several arrangements of functional interaction may lead to the same CCH. It is difficult to develop a statistical test of independence for the CCH because successive bins are not independent due to a decreased probability of firing during a neuron's refractory period. The CCH can however be predicted by offsetting stimulus trials by the length of one stimulus presentation, which destroys any temporal relationship, and recomputing the CCH (shift predictor – see Figure 1.3) (Perkel et al. 1967; Gerstein and Perkel 1969; 1972). This shift predictor is based on the concept that most neuronal interactions occur on a shorter time scale than the time between two successive stimulus presentations. By shifting, the spike trains still contain all direct stimulus effects, but the interaction effects are destroyed because the time shift is so large. The shift predictor reflects the null hypothesis that spikes from the two neurons are statistically independent and only the firing probabilities are related to the stimulus.
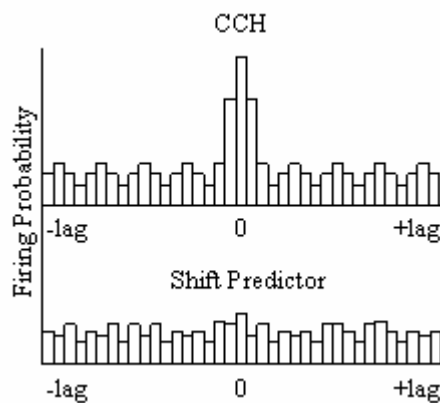
Figure 1.3: *Top*: A cross-correlation histogram (CCH) is constructed of the intervals between all train A events and all train B events, forward and backward, out to some designated time lag. An observed dependence in the CCH can reflect functional interaction or common input. *Bottom*: The shift predictor is used to estimate the amount of chance firing events by recomputing the CCH after shifting one spike train by the length of one stimulus period to destroy serial independence.

While analysis with the CCH has provided insight into the "effective" underlying neuronal architecture and interactions among neurons, there are some disadvantagaes to using this method. Due to high variability in neuronal responses, the CCH must represent an average over many stimulus repetitions. Temporal dynamics of the cellular interactions cannot be resolved. Furthermore, there are several pathways of functional interaction that may lead to the same CCH. Determinations of the underyling network configuration are estimates at best because of the ambiguity inherent in the CCH. Finally, a cross-correlation can only be applied to pairs of neurons and pair-wise calculations increase in a power relation for larger neural assemblies.

JPST Scatter Diagram

To examine the temporal dynamics of neuronal firing correlation, the idea of a joint peri-stimulus time (JPST) scatter diagram was introduced (Gerstein and Perkel 1969; 1972). Each spike train from a target pair form the axes of a scatter plot in which a dot represents the (delayed) coincidence of the spike trains relative to the stimulus onset (see Figure 1.4a). The dot density is built up by carrying out many stimulus repetitions. The temporal dynamics of cooperation can be visualized in the patterns of the dots. For instance, regions of high or low density in rows or columns represent stimulus-locked activation or suppression of either neuron. The variation in density reflects the stimulus-induced modulation of firing rate in the corresponding neuron. Regions of high or low density parallel to the line $y = x$ represent delayed coincidence firing, whereas the line $y = x$ represents synchrony. Note that the JPST scatter diagram does not handle superposition of dots from different stimulus trials.

18

Similar to the CCH, the JPST scatter diagram reflects interactions among only pairs of neurons. In an attempt to resolve this, the JPST scatter diagram has been extended to compare three neurons (Gerstein and Perkel 1972; Perkel et al. 1975), where the resulting Joint Impulse Configuration Scatter Diagram uses a non-traditional triangular coordinate system (see Figure 1.4b). The axes intersect at 120° and represent pair-wise time differences among the three neurons. Synchronous events between pairs of neurons are represented by high (or low) dot densities midway between coordinate axes and synchronous events among all three neurons are concentrated at the origin. Direct stimulus effects are reflected in dot densities along lines parallel to the coordinate axes. Generalizing this method to four spike trains requires a tetrahedrally arranged coordinate axes and so on. Besides poor generalization of this method to incorporate large numbers of neurons, these results only provide a qualitative description of joint interactions.

Figure 1.4: Joint Peri-Stimulus Time Scatter Diagrams. *a*. The (delayed) coincidences of spike events in two spike trains can be visualized with a scatter diagram. *b*. For three neurons, the plot is called the Joint Impulse Configuration Scatter Diagram and uses triangular coordinate axes. Synchronous events from all three neurons are clustered near the origin.

JPSTH

Aertsen et al. (1989) developed methods to quantify and normalize the JPST scatter diagram. These new procedures allowed the separation of raw correlation into correlation caused by direct stimulus effects and correlation caused by interaction between the two neurons. Instead of using a scatter diagram where dot superposition does not occur, the joint peri-stimulus time histogram (JPSTH) uses the PST histograms from each neuron along the coordinate axes to produce a two-dimensional scatter histogram (see Figure 1.5). Joint spike events along the main diagonal comprise the raw PST coincidence histogram, which represents the modulation of raw synchrony over the duration of the stimulus. A CCH can be produced by averaging across bins parallel to the main diagonal. In order to remove direct stimulus modulations of single-neuron firing rates, a predictor computed from the cross-product of the two single-neuron PST histograms can be subtracted from the raw JPSTH. The predictor represents the null hypothesis of independent firing and is mathematically equivalent to an average over the set of all possible shift predictors (Palm et al. 1988). The corrected JPSTH is subsequently scaled by the cross-product of the standard deviations of the single-neuron PST histograms to produce the normalized JPSTH.

The information-theoretical concept of surprise (Legendy 1975; Palm 1981) is used for significance testing. The surprise of a value is the negative natural logarithm of the probability of finding that value or one more deviant and is computed for every bin in the normalized JPSTH. Regions of the normalized JPSTH are determined to be significant if there is an extreme value in most of the bins in that region. This requirement for spatial coherence ensures that isolated bins with extreme values are

considered noise. However, the concept of spatial coherence is fuzzy and somewhat

subjective and requires the additional process of smoothing.
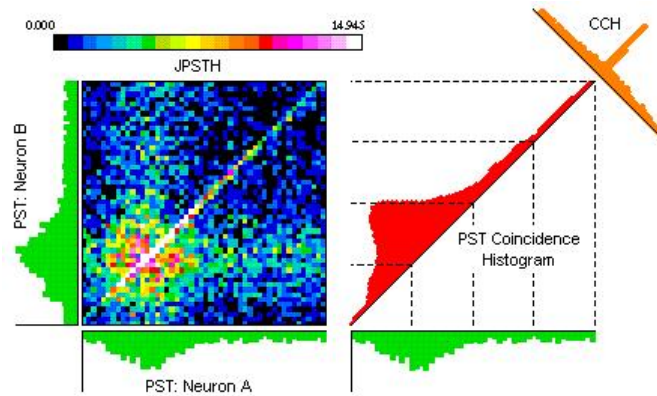
Figure 1.5: Joint Peri-Stimulus Time Histogram (JPSTH). This figure shows the
components of the raw JPSTH. The JPSTH matrix displays the (delayed) coincidences of
spike events in each neuron's PST histogram where color indicates number of spike
events. The PST coincidence histogram shows those events along the main diagonal of
the matrix and represents synchronous events over the duration of the stimulus. A CCH
is computed by averaging across bins parallel to the main diagonal of the matrix. Image
obtained from http://mulab.physiol.upenn.edu/jpst.html.

        The JPSTH method provides a measure of correlation between two neurons

without making any assumptions about the structure of the underlying neuronal

architecture. However, analysis of the normalized JPSTH can recover the "effective

connectivity" of the underlying network, or the simplest neural network that would

replicate the observed patterns in the JPSTH. As in the discussion of the CCH, there are

numerous neuronal configurations that could produce the same JPSTH and any analysis

based on the effective connectivity derived from this method should be done with

caution. Although this method provides a way to quantify neural relations compared to

the JPST scatter diagram, significance testing with the surprise measure is subjective and

makes it difficult to produce definitive conclusions. Numerous spike events are needed

21

and non-stationarities in the data could produce incorrect results. Finally, this procedure

is also designed to measure interactions between pairs and cannot be applied practically

to larger assemblies.

Gravitational Clustering

In order to investigate the cooperative activity of assemblies with more than two

neurons, other approaches were developed. The gravitational clustering algorithm was

introduced to characterize the time-varying organization and extent of neural assemblies

(Gerstein et al. 1985; Gerstein and Aertsen 1985; Aertsen et al. 1986). Gerstein et al.

(1985) acknowledged that characterizing neural relationships meant defining what is

meant by cooperation and developing quantitative criteria to recognize and describe this

cooperation. The gravitational clustering algorithm relates the activity of neurons to the

motion of particles in a multidimensional Euclidean space. A spike generated by a

neuron results in an increment of charge, which can add to or decay before the next spike

depending on how close in time the second spike is fired relative to the first. The force

exerted between particles is proportional to the product of their charges and inversely

related to their separation distance. The force on a particle affects its velocity and results

in clustering of particles that tend to fire together. Separate aggregates of particles

represent independent neural assemblies.

While this method identifies groups of cooperative neurons, we still cannot escape

combinatorial arithmetic since the force on each particle is a vector sum of all pair-wise

relationships. And although each relationship is based on a quantitative measure of

attraction, there is no quantitative description of the cooperation in an entire assembly

besides, perhaps, the aggregation time of these fictitious particles in a hypothetical

multidimensional space. Thus, statistical comparisons are difficult. Furthermore, this

algorithm does not readily identify cells that belong to more than one functional group or

subsets of cells within larger functional groups. Finally, this method cannot be

practically displayed without using projections through hyperspace (Gerstein and Aertsen

1985) or pair-wise distance plots (see Figure 1.6). Although there have been

improvements to the sensitivity of this technique (Baker and Gerstein 2000), the current

problems make it unsuitable for extensive population analysis.
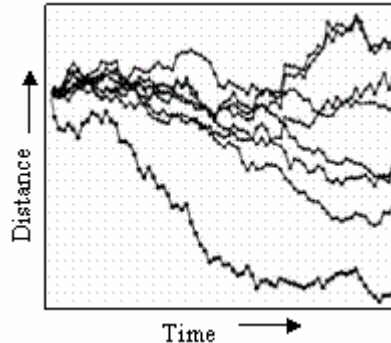
Figure 1.6: Graphical illustration of the gravitational clustering technique. The activity
of cooperative neurons is mapped onto the motion of particles in a multidimensional
space. Particles gravitate towards each other following rules of attraction and repulsion
governed by spike discharges. This plot displays the distance between pairs of neurons
over time. As the distance between particles becomes very small, they are clustered and
considered a neural assembly. Adapted from Gerstein and Aertsen (1985).

Information Theory

In an effort to measure the amount of cooperative activity in an assembly of

arbitrary size, some methods were developed using information theory as a mathematical

foundation. For instance, type analysis (Johnson et al. 2001) considers a point process

spike train as a binary sequence (using an appropriate bin width) and uses the firing

patterns of all the cells in the target population to create a unique response sequence.

After repeated presentations of the stimulus, a probability distribution for each bin can be estimated. Probability distributions from two different stimulus conditions can be compared using a distance measure such as the modified Kullback-Leibler or Resistor Average distance (Johnson et al. 2001). Advantages of this method include that discharge history can be included in the calculations and there are no assumptions about the nature of the neural code. However, drawbacks include the requirement for large amounts of data that increases in a power fashion with the number of neurons in the population and the need for additional bootstrap procedures to estimate the bias inherent in the distance measure (Efron and Tibshirani 1993; Johnson et al. 2001).

Metric space analysis (Victor and Purpura 1996) uses cost-based metrics to calculate a distance between spike trains. The metrics are based either on the spike arrival times or spike intervals and the cost of translating one spike train into another spike train is determined by qt, where t is the time required to move a spike or change the length of an interspike interval and q is the cost scaling parameter. The method has been recently developed for population analysis (Victor 2000; Aronov 2003; Aronov et al. 2003) by including an additional cost scaling parameter k, where k scales the cost of changing the cell-label of a spike (i.e. which cell fired the spike). By clustering metric-based distances over a set of stimuli, the mutual information is calculated to assess the role of population coding. Although quantitative, methods like this are abstract and increasingly more difficult to apply to larger neural assemblies.

Other Methods for Quantitative Comparison of Cooperative Activity

Assemblies could be represented by more complex interactions that require more subtle discrimination. One method of multivariate signal processing that has been

developed for neural ensemble analysis is principle component analysis (PCA) (Chapin and Nicolelis 1999). PCA transforms data into linear combinations by rotating the data into a new coordinate system defined by the direction of greatest variance. PCA does not detect nonlinear interactions, so independent component analysis (ICA) has been used to detect these higher-order interactions (Laubach et al. 1999). The problems with ICA include the requirement of low firing probabilities and relatively high correlations within assemblies. Furthermore, these methods are designed for pair-wise interactions.

Another approach to characterizing population dynamics has been to search for statistically significant patterns of spikes. The approach was initially developed by Dayhoff and Gerstein (1983a,b) and has been refined to account for various levels of jitter and models of significance (Abeles and Gerstein 1988; Tetko and Villa 2001a,b). Spike pattern recognition has also been developed to search for higher-order interactions among assemblies (Martignon et al. 2000).

Grun et al. (2002a), Gal et al. (2003) and Czanner et al. (2005) have introduced several other novel methods for analysis of multispike information but they are largely theoretical and have not been applied to analysis of data from the visual cortex. Despite the creativity of these approaches and the identification of some forms of information in spatially and temporally extended spike trains, there remains no clear identification of the visual features this information might represent or, more critically, how the brain might make use of it.

Other Methods for Qualitative Comparison of Cooperative Activity

Although there are display techniques for population activity (e.g. raster plots), there have been very few methods that specifically display the organization of

cooperative activity within a population. One such method is the synchrony colormap described by Samonds and Bonds (2005). Rooted in the sorting procedures used by Johnson et al. (2001), this technique creates an N-bit firing pattern for every instance in time where N is the number of cells in a neural assembly. By converting these patterns to their decimal equivalents and arranging each number as a square on a checkerboard, a continuum of activation can be created where the upper left portion of the board represents no synchrony and the lower right portion represents all N cells synchronizing. This map can be used for a real-time or post-experiment visualization of synchrony. However, drawbacks of this method include the requirement to identify and monitor a specific assembly and that the identities of cells in cooperative subsets are not readily identified. Furthermore, this display is qualitative and claims of weak, moderate, or strong synchrony may be misleading in that strength is determined by number of cooperative cells and not by interactions among cells.

Conclusion

The previous pages detail the formidable task of describing synchrony among large neural assemblies. Multielectrode array technology allows for the simultaneous single-unit recording of dozens of cells, but current analysis methods are unsuitable to accurately and completely describe cooperative population activity. Therefore, development of a new method is imperative if we are to understand how synchrony and assembly formation contribute to the perception of our environment. The next chapter describes a method that quantifies cooperative activity within a neural assembly of arbitrary size and introduces a new parameter that measures the quality of synchronous relationships. Furthermore, this new procedure can identify dynamic grouping and track

temporal dynamics of cooperative relationships. No binning or smoothing techniques are required and all results are subject to simple significance testing. Computation time increases linearly with the number of cells involved so that exploration of group characteristics and size trends is no longer an unattainable goal.

CHAPTER II

QUANTIFICATION OF MAGNITUDE AND QUALITY OF SYNCHRONOUS
ACTIVITY WITHIN NEURAL ASSEMBLIES OF ARBITRARY SIZE

Introduction

Our brains process and interpret sensory information in order to generate

perceptions of the environment or motivate behavior. However, the underlying

mechanisms by which salient stimulus qualities are represented by neuronal response

patterns remain a mystery. Precise coordination of spike events, or synchrony, is an

attractive candidate to play a role in (visual) coding since it exists among (visual) cortical

neurons, but its functional significance is largely unknown. Recent experimental

evidence demonstrates that synchronous activity in cat visual cortex can discriminate

between co-linear and co-circular contours (Samonds et al. 2005), suggesting that

synchrony may be involved in the representation of contours or shape. As a hypothetical

neural substrate for encoding salient stimulus properties, synchrony enhances the

probability of eliciting postsynaptic action potentials when neurons behave as coincident

detectors (Azouz and Gray 2003), thus ensuring propagation of this information to

subsequent levels of the cortical hierarchy. Synchrony allows for neurons to be

effectively connected and form dynamic regional assemblies to reliably and efficiently

transmit information throughout the cortex while minimizing metabolic demands.

To measure cooperative interactions within neural assemblies, researchers have

traditionally used single electrodes, tetrodes, or small arrays to record simultaneously

from small numbers of neurons. Perkel et al. (1967) introduced the cross-correlogram

method to quantify cooperative relationships between pairs of neurons to determine whether observed dependences were consistent with functional interaction or common input. Current multielectrode array technology allows recording of dozens of neurons simultaneously, but cross-correlograms and other techniques such as the Joint Peri-Stimulus Time Histogram (JPSTH; Aertsen et al. 1989) cannot characterize the synchrony between more than two cells. The cross-correlogram has been extended to include cooperation among three cells (Gerstein and Perkel 1972; Perkel et al. 1975; Abeles and Goldstein 1977), but the resulting display is limited to triangular coordinates and cannot be applied practically to larger assemblies.

Other techniques such as gravitational clustering (Gerstein and Aertsen 1985; Gerstein et al. 1985) identify cells that fire together, but the results are qualitative and still based on pair-wise distance calculations. Correlated group activity quantified through information theory (Johnson et al. 2000; 2001; Pola et al. 2003) or metric-space analysis (Victor and Purpura 1996; 1997) describes the amount but not the nature of cooperative information. Methods based on principle component analysis and independent components analysis (Laubach et al. 1999), those based on separating synchrony and oscillatory firing patterns (Konig 1994), and many others (Lee 2002; Panzeri et al. 1999) aimed at quantifying synchrony in an assembly all apply to pairs and cannot be extended to larger groups. A few methods have been developed to display population activity (Ortega et al. 2004; Samonds and Bonds 2004), but these displays are qualitative and do not quantify the amount of synchrony or quality of membership in a neural assembly.

Although synchrony could constitute a suitable encoding scheme, implicating synchrony in neural coding is challenging because adequate methods to measure

synchrony have not been developed. Most of the methods mentioned above quantify synchrony as some distance between two neurons. However, synchrony allows for the formation of transient functional groups which could include tens, hundreds, thousands, or even larger numbers of neurons. Pair-wise distance calculations increase exponentially as group size increases and can be computationally exhaustive for large groups. Furthermore, very few studies have investigated how to measure the quality of synchrony as well as develop some means to display these quantities. A formal and informative means of analyzing the correlated behavior across larger groups of cells is needed if we are to understand the significance of this information to the visual process.

This chapter describes a method to quantify the magnitude and quality of synchrony within assemblies of arbitrary size ($N \geq 2$), where computation time increases linearly with the number of neurons involved. The results of this method can be used to measure average synchrony or track the temporal dynamics of synchrony within the assembly. This method was applied to multielectrode array recordings in the visual cortex of paralyzed and anesthetized cats to characterize properties of assembly membership.

Method

As synchrony may serve to detect relevant information in the visual field by increasing the probability of eliciting postsynaptic action potentials, our basic algorithm is designed to reflect the relevance of group synchrony to postsynaptic neurons by modeling the temporal and spatial summation of postsynaptic potentials (PSPs). We will first describe the steps in quantifying the magnitude and quality of synchrony and then show how the results can be integrated into a visual display.

30

Quantifying the Magnitude of Synchrony in a Neural Assembly

The purpose of this method is to measure the timing similarity between neuronal spike trains in an assembly of arbitrary size. Neurons with similarly timed events are considered synchronous and the magnitude of synchrony depends on the degree of similarity.

Step 1: Convert Spikes to Trains of Point Events

If we assume that action potentials are instantaneous and indistinguishable, then only the timing of such events is sufficient to describe visual information. The activities recorded simultaneously from all neurons in a target assembly (N neurons with N_{TS} total spikes and N_{CS} coincident spikes, where coincident spikes are defined as events that occur within some integration time period defined below) are preprocessed to retain only spike initiation times, creating point-event spike trains ($S_i(t)$ is the spike train from neuron i). Spike trains for a given evaluation must have the same duration, L_S. Note that this does not imply that all spike trains have the same number of spikes.

Step 2: Generate PSP Trains

The comparison of simultaneous spike trains requires deriving a similarity measure that is conscious of time. This can be accomplished by convolving a point-event spike train, $S_i(t)$, with a PSP waveform, $W(t)$, that has area A and duration L_W.

$$P_i(t) = S_i(t) * W(t) = \int_t S_i(\tau)W(t-\tau)d\tau \qquad (2.1)$$

This yields a PSP train, $P_i(t)$ with length L. A PSP waveform is often approximated using an alpha function, but any waveform can essentially be chosen based on its desired Weight Function (see Figure 2.1). For instance, an alpha waveform nonlinearly weights events so that spikes occurring closely in time are weighted more than those occurring

towards the end of the integration time period. Alternatively, waveforms can be chosen

to reflect linear or constant weighting schemes. For example, the JPSTH method

(Aertsen et al. 1989) effectively uses a constant weighting scheme by binning. For all

subsequent illustrations and analysis, an alpha function will be used to model PSP

waveforms.
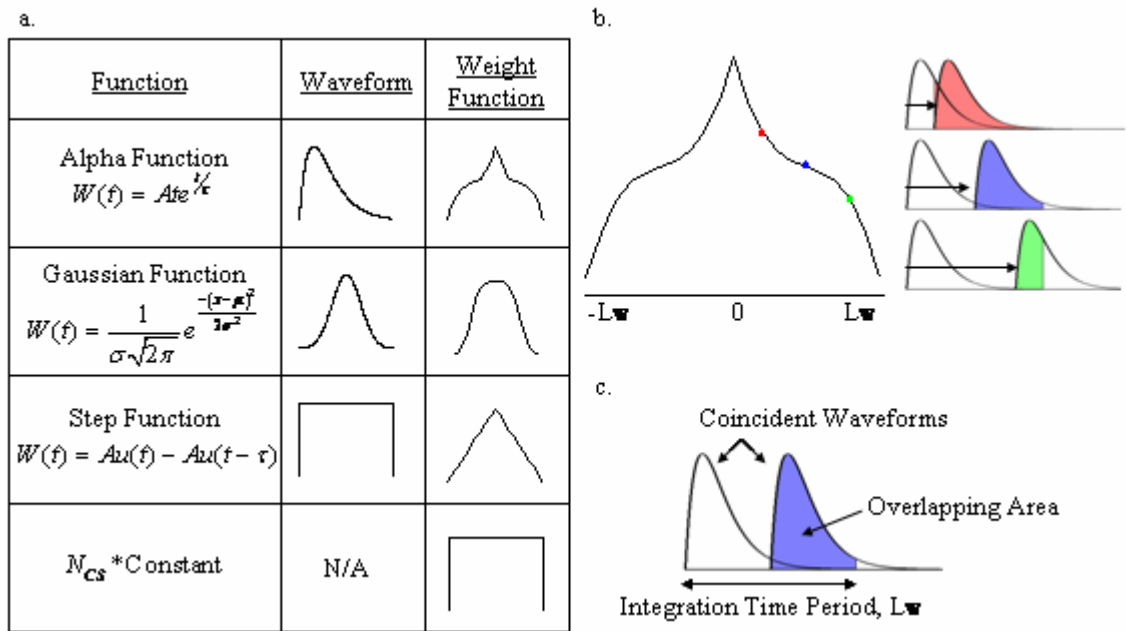
Figure 2.1: *a*. This table shows different functions used to approximate a PSP
waveform. The two-dimensional weight function ($N = 2$) for each waveform is shown in
the last column. In order to produce a constant weight function, no matter how two
waveforms from different neurons are arranged within the integration time period, their
overlap is given the same weight. This can be accomplished without using a PSP train by
simply multiplying the number of coincident waveforms in an assembly by a constant
weight. *b*. A weight function is created by integrating the overlapping area under two
coincident waveforms as the waveforms are shifted by lengths up to the integration time
period, L_W (see *c*.). In the case of an alpha function (*b*), two waveforms that occur at the
exact same time (lag = 0) have the most overlapping area and therefore have the most
weight. Note that the weight function is a conceptual tool and will not be implemented or
discussed further.

Step 3: Filter PSP Trains

The next step is to create a filtered PSP train, $\tilde{P}_i(t)$, by removing parts of the PSP trains that do not overlap with all other PSP trains in the selected assembly. Note that coincident waveforms from different neurons occur when the respective spike events are within L_W of each other. Therefore, L_W is the integration time period (Figure 1c). Here, we will define coincident waveforms as waveforms in which all or part of the waveforms are overlapped in time. The actual overlapping portion will be referred to as the overlapped area of coincident waveforms. A filtered PSP train is created by multiplying each PSP train by a filter, *F(t)*, which has a value of 1 at times when the waveforms are overlapping (synchronous) and 0 elsewhere (see Appendix for how to create *F(t)*).

$$\tilde{P}_i(t) = P_i(t) \cdot F(t) \qquad (2.2)$$

Step 4: Calculate Raw Score

The magnitude of synchrony is computed as the ratio of the area under the overlapped portion of coincident waveforms to the total area under all waveforms in the assembly. To determine the Raw Score, the filtered PSP trains are integrated and the area is divided by the area computed from integrating the original PSP trains. Alternatively, the area under the filtered PSP trains can be normalized by the area under one waveform, *A*, and the total number of spikes in the assembly, N_{TS}. The Raw Score is a number between 0 and 1 and represents the percentage of total waveforms that are overlapped in an assembly. An assembly with a large Raw Score is comprised of neurons whose responses occur at similar times and are thus very synchronous.

$$Score_{raw} = \frac{\int_L \int_N \tilde{P}_i(t)\,di\,dt}{\int_L \int_N P_i(t)\,di\,dt} = \frac{\int_L \int_N \tilde{P}_i(t)\,di\,dt}{AN_{TS}} \qquad (2.3)$$

Step 5: Calculate the Chance Score and Determine Significance

The Chance Score is the percentage of waveforms expected to overlap by chance under the null hypothesis that all neurons in the assembly are firing independently. This value can be predicted by calculating the shift predictor (Perkel et al. 1967). Successive trials of stimulus presentations are shifted in time, which destroys any temporal relationships and preserves order-independent statistics. The Chance Score is computed by completing steps 1-4 with spike trains that are shifted in time by the length of at least one stimulus trial compared to all other trains. This was found to be not significantly different (t-test, $p > 0.05$) than averaging the scores from all possible shift combinations in assemblies where $N = 2$. However, it is beneficial to perform this step a number of times with different shift combinations to build a distribution with which to statistically compare the Raw Score.

Step 6: Compute Normalized Score

A Normalized Score can be computed by subtracting the Chance Score from the Raw Score and renormalizing the resulting value to be between -1 and 1, which represent Scores that are far less than or greater than Chance, respectively, and Chance is assigned a value of 0. A Normalized Score is independent of firing rate and Scores ≥ 0 represent the percentage of waveforms that are synchronous, but not due to chance from firing rate-induced modulation of synchrony. For instance, a Normalized Score of 0.5 means 50% of waveforms are synchronous beyond chance. The absolute values of Scores ≤ 0 represent the percentage of waveforms less than Chance that are synchronous. For instance, a Normalized Score of -0.1 means that the coordinated activity in an assembly is 10% less than activity expected by Chance. Renormalizing the score ensures that assemblies with different spike counts and Chance Scores can be compared and

assemblies with identical spike trains have a Normalized Score of 1. Similarly, spike trains with no coincident spikes have a Normalized Score of -1.

$$Normalized\ Score = \begin{cases} \dfrac{Raw\ Score - Chance\ Score}{1 - Chance\ Score} & for\ Raw\ Score\ \geq\ Chance\ Score \\ \dfrac{Raw\ Score - Chance\ Score}{Chance\ Score} & for\ Raw\ Score\ <\ Chance\ Score \end{cases} \quad (2.4)$$

Figure 2.2 depicts a graphical representation of the steps in quantifying the magnitude of synchrony. Although this method was developed to measure the temporal similarity of spike trains within an assembly and not to determine underlying functional anatomy (e.g. like the JPSTH, Aertsen et al. 1989), this method can reveal some aspects of effective connectivity. Individual Scores can be calculated by integrating the area under each filtered PSP train and dividing by the Score (this is usually done with the Raw Score). The resulting values represent individual contributions to the collective group synchrony. When using an asymmetric waveform like an alpha function in the PSP trains, this information can reveal whether some neurons tend to fire before or after others. Neurons with similar contributions reflect a shared input while uneven contributions reflect direct interactions. Although there are no assumptions about the underlying configuration of the network, these results can be compared to those from a known network configuration to determine whether they are consistent with a certain neuronal architecture.

$$Score_{individual} = \frac{\int_L \tilde{P}_i(t)dt}{Score_{raw}} \quad (2.5)$$
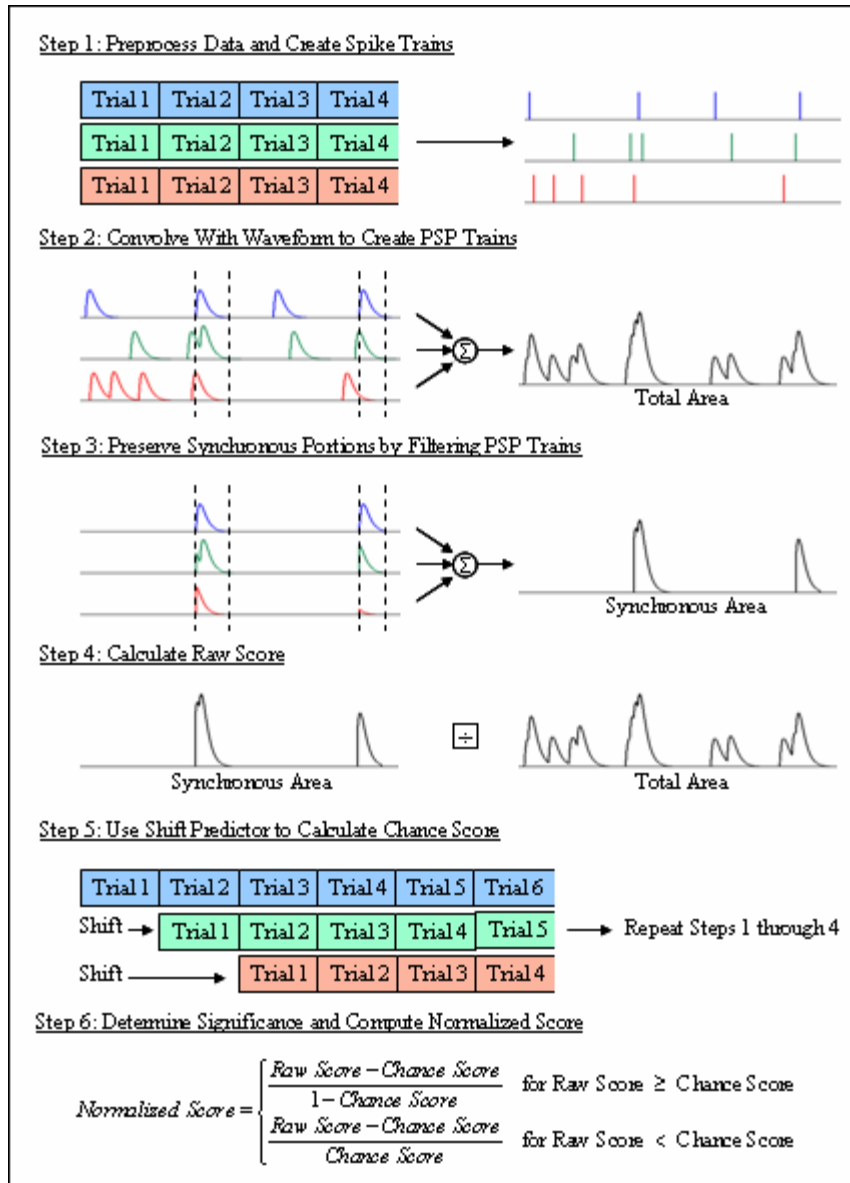
Figure 2.2: Steps involved in calculating the magnitude of synchrony within an assembly of arbitrary size (demonstrated for 3 neurons).

Quantifying the Quality of Synchrony in a Neural Assembly

Consider two neural assemblies: One assembly synchronizes often, but the amount of overlap is small each time. The other assembly synchronizes only a few times, but the amount of overlap is large. Based on the total amount of overlapping area, these two assemblies may have similar magnitudes of synchrony, but the nature of synchrony is

different (see Figure 2.3). To quantify the nature of synchrony, we have developed two measures of quality – one based on the number of coincidences (indirectly measured by duration of synchrony) and the other based on average degree of overlap during synchronous events.

Step 1: Calculate Duration of Synchrony

Determining the number of instances when an assembly is cooperative is not a trivial task. However, we can measure the total amount of time the assembly cooperates. The percentage of time a neural assembly synchronizes is computed as the ratio of duration of synchronous activity to the duration of all activity. This can be calculated by simply integrating the Filter function (value is 1 when waveforms are overlapping and 0 elsewhere) and dividing by the total time covered by all waveforms, L_{TS}, which may not be equal to $N_{TS}L_W$ if spike events from the same neuron occur within the integration time period and those PSP waveforms are superimposed.

$$Q_{time} = \frac{\int_L F(t)dt}{L_{TS}}$$

(2.6)

Step 2: Calculate Average Overlapped Area of Coincident Waveforms

The average percentage of overlapped area per cooperative instant is computed by dividing the area under overlapped portions by the total area under coincident waveforms. Pseudo code for computing the number of coincident spikes, N_{CS}, is provided in the Appendix.

$$Q_{overlap} = \frac{\int_L \int_N \tilde{P}_i(t)didt}{AN_{CS}}$$

(2.7)

37
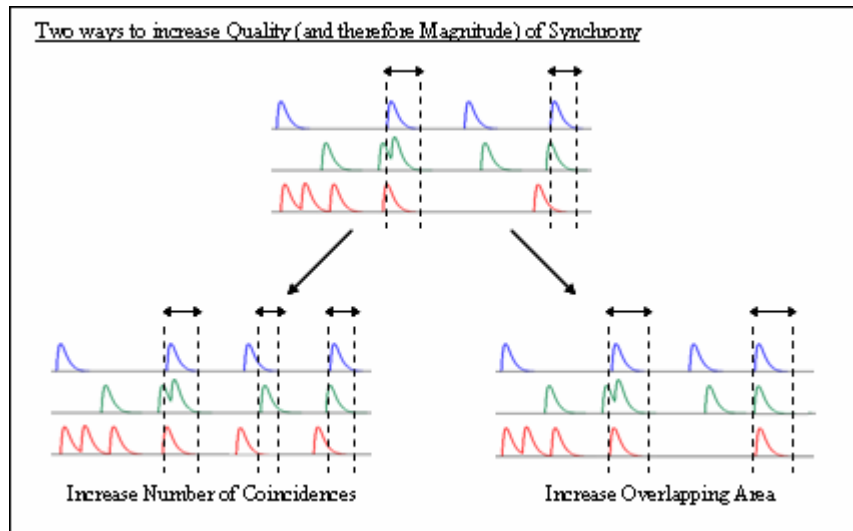
Figure 2.3: Quality of synchrony is determined by the number of coincidences and the average amount of overlapping area per coincidence. *Top*: The responses of three neurons in an assembly would be more synchronized if (*Bottom-Left*) there were more instances of coincident waveforms and/or (*Bottom-Right*) the time lag between synchronous spikes was shorter, increasing the average overlapping area of coincident waveforms. Counting the number of coincidences is not trivial, so we use duration of synchronous activity instead. Note that increasing the overlapping area increases the duration of synchrony, but it is possible to increase the duration of synchrony (by increasing the number of coincidences) without increasing the average overlapping area of coincident waveforms.

The magnitude and quality of synchronous activity are both normalized so that a neural assembly with identical spike trains will yield values of unity. Therefore, each quantity is a percentage of its maximum synchrony potential. Note that maximum potential is a conceptual, mathematical quantity and may not have a real-life correlate. The maximum synchrony potential occurs when all spikes within an assembly are completely synchronized. However, this may not be realistic if for instance, one non-bursting cell (A) has 25 spikes and another non-bursting cell (B) has 75 spikes (N = 2). This pair cannot achieve their maximum potential to obtain a score of 1 (100 spikes synchronized). Therefore, their highest score (with 50 synchronized spikes – 25 from each cell) is a fraction of their maximum potential (50/100 or 0.5). (Note that bursting

could increase the score since two partially overlapping spikes from cell B could coincide with one spike from cell A.)

Visual Display of Synchronous Activity

Understanding the rules (and roles) of group synchrony requires comparison of the spatial properties of the stimulus with the receptive fields of participant cells. To accomplish this, we have developed an informative display of a population's spatial and temporal properties (Figure 2.4). It contains a dynamic image of the stimulus overlaid on receptive fields from all of the cells in the measured population. Non-synchronous receptive fields are transparent. With significant cooperation (e.g. $p < 0.05$) active receptive fields will assume a color that indicates their individual contributions to synchrony (see Figure 2.5). The central bar in a receptive field outline delineates the cell's preferred stimulus orientation.

Temporal information in the display is shown as the image changes in time and also with a plot of the summed PSP trains for the currently active neural assembly. Even if there is no constructive cooperation between the cells in the assembly, occasionally synchronous events can occur simply due to chance collisions resulting from the firing rates of the cells. The plot has a line indicating the chance level of synchronous events as well as a second line indicating when the strength of synchronous events exceeds significance (t-test, $p > 0.05$). Subthreshold activity can be observed below these levels. The solid red bar indicates the current time. The plot of synchronous waveforms can be used to reveal information about the stimulus and the assembly. For instance, the example in Figure 2.4 shows an active assembly ($N = 3$) during concentric ring stimulation. The plot shows synchronous activity during one cycle (black/white) of the

stimulus. Gray dotted lines in the graph denote where synchrony peaks, then decays in accordance with the change in contrast. Peaks in the graph with similar magnitudes (arrows) represent a common relationship among spikes in the assembly (i.e. they coincide in fixed proportion to produce the same amount of overlapping area).

The display also contains an image of the multi-electrode array to show the spatial relationship between the locations of each cell belonging to the assembly on the surface of the cortex. By identifying the stimulus features that are within the receptive fields of a group of cells when synchrony occurs, we can determine the relationships between receptive field properties and stimulus patterns that are conducive to cooperative activity. Finally, the magnitude and quality of synchrony (see Method) are displayed numerically and graphically with colored slide bars.
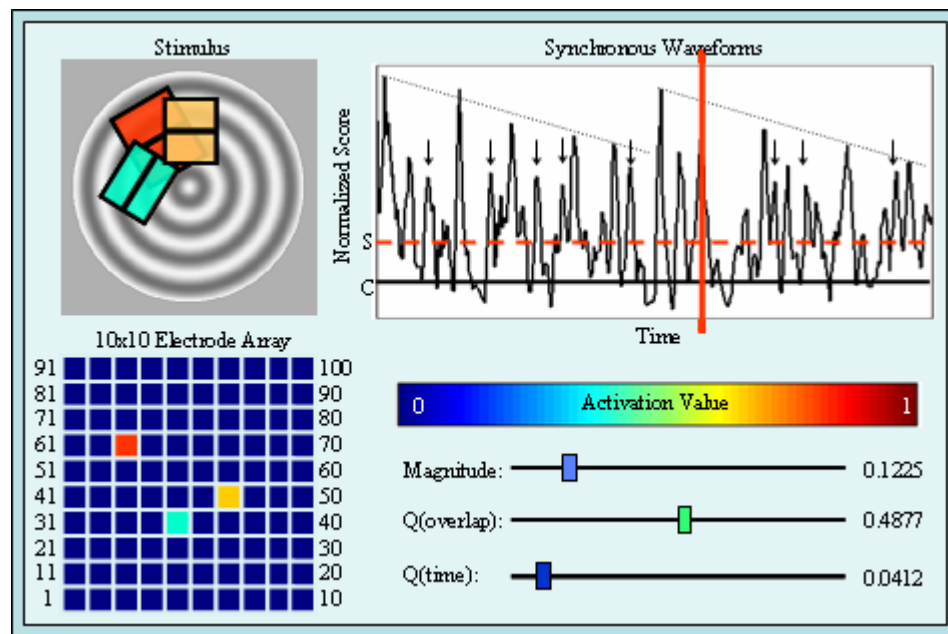

Figure 2.4: Dynamic display of assembly activity. Features include an image of the stimulus with superimposed receptive fields from the cells in the currently active assembly, slide bars to display the magnitude and quality of synchrony in a group, an image of the multi-electrode array, and a plot of the summed PSP trains for the currently active group.
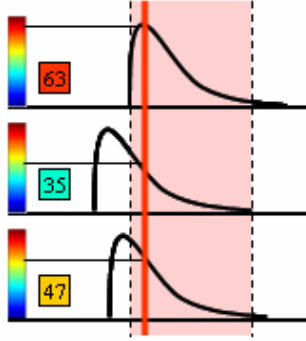
<u>Figure 2.5</u>: Individual contributions to synchrony. Receptive fields and array channels assume colors indicated by the relative area overlapping with other waveforms. Similar colors reflect similar contributions and the actual color indicates the amount contributed.

Applying PSP Method to Experimental Recordings

We applied the PSP method to assemblies in the visual cortex of paralyzed and anesthetized cats when viewing drifting grating and concentric ring stimuli. We wanted to confirm previous results of stimulus-dependence and dynamic grouping (Samonds et al. 2005) and extend these findings to larger neural assemblies. We compared results to those from the JPSTH (Aertsen et al. 1989) to gauge the effectiveness of this method in identifying synchronous pairs and describing the amount of activity within those pairs. We show how magnitude and quality depend on assembly size and demonstrate how the variance of Individual Scores can be used to measure quality. Finally, we show how results of the PSP method can be used to increase experimental efficiency.

Materials and Methods

Preparation

During recording, cats were paralyzed with pancuronium bromide (Pavulon; 0.3 mg/kg hr) and artificially ventilated with a mixture of N_2O, O_2, and CO_2 (75:23.5:1.5) to

41

hold expired pCO_2 at 3.9%. An infusate also contained Propofol (0.3 mg/kg hr) to maintain effective anesthesia. Rectal temperature was maintained at 37.5° C with a servo controlled heat pad. Eyelids were retracted and the natural pupils dilated by instillation of phenylephrine HCl 10% and atropine sulfate 1% in the conjunctival sacs. Contact lenses with 4 mm pupils were placed on the corneas and auxiliary spectacle lenses were added as dictated by direct ophthalmoscopy to render the retinas conjugate with the stimulus plane 57 cm away. At this distance, a visual angle of 1° is equivalent to 1 cm on the screen.

Data Acquisition

Simultaneous single-unit recordings were made via the "Utah Intracortical Electrode Array" (UIEA; Jones et al. 1992) from complex cells in the visual cortex of two cats. This is a square 10 x 10 (100 total) silicon array on 400 micron centers (4 x 4 mm footprint). The electrodes had a length of 1.0 mm inserted to a depth of 0.6 mm with a pneumatic implantation tool (Rousche and Normann 1992) that minimizes tissue damage (Schmidt et al. 1993; Rousche and Normann 1998). The insertion depth concentrated the electrodes in layers II/III, minimized tissue displacement (estimated at 4%) due to the volume of the electrodes, and avoided impact to the cortical surface by the electrode base. Amplification and signal processing were provided for each individual electrode (Guillory and Normann 1999), recordings were displayed in real time, and the waveform of each neural event was stored for later analysis by a comprehensive software system (Bionics, Salt Lake City, UT).

For stimulation, we used both drifting sinusoidal grating and concentric ring stimuli (see Figure 2.6). Gratings were oriented in ten degree increments from 10° to 360°. Both stimuli had the following properties: 0.5 cycles/degree spatial frequency, 2 Hz temporal frequency, and contrast set to 0.5. The center location of drifting concentric rings varied throughout the experiments and depended on manually mapped locations of receptive fields. All stimuli were displayed on a computer monitor with a 120 Hz refresh rate through a circular aperture against a mean luminance background (73 cd/m^2). Each stimulus was presented for 2 seconds followed by a 1 second mean luminance interval and repeated 100 or 200 times for reliability. Further details on the preparation, data acquisition, and stimuli are discussed elsewhere (Samonds et al. 2005).

Figure 2.6: Stimuli. *Left*: Drifting sinusoidal grating. *Right*: Drifting sinusoidal concentric ring. Each stimulus covered the classical receptive fields of the population and was displayed against a mean luminance background.

Results

Although the PSP method can quantitatively describe synchrony within large assemblies, the membership of meaningful assemblies (i.e. all cells that might fire synchronously) must be defined. Because assembly membership is to some extent stimulus-dependent (dynamic grouping; see below), it may be impossible to identify

unambiguously members of a single assembly based solely on spike train properties. Each member may belong to numerous groups and the specific spike train properties relevant to a group of interest are unknown. Therefore, due to the combinatorial arrangement of cells from the population, there are a large number of possible inputs to the PSP algorithm. If we assume that only a small fraction of subsets in the population will synchronize to any given stimulus, then the probability of randomly testing combinations of cells and finding a significant assembly is very low. Fortunately, we have found that members of larger assemblies can, in general, be predicted based on Normalized Scores from smaller assemblies. Therefore, we can compute a relatively low number of small assembly scores and progressively cluster cells that synchronize well to create an assembly of arbitrary size that has a high probability of synchronizing. (Note that cells are "clustered with replacement" so they can belong to more than one assembly.) These large assemblies can then be analyzed using the PSP algorithm to determine if, in fact, their grouping is significant.

We simultaneously recorded single-unit activity from 28 and 23 complex cells in the visual cortex of two paralyzed and anesthetized cats. Using an alpha function (non-linear weighting scheme), where $\tau = 1$ ms and $L_W = 10$ ms (Softky and Koch 1993), we applied the PSP method to compute Normalized Scores for all pair-wise combinations of cells (631 pairs). The scores obtained from these groups were used to estimate larger assemblies (N = 3, 4, 5, 6) for each cell. For instance, with 28 cells, there were 28 assemblies for each size group (N = 3, 4, 5, 6) where the first assembly in each size group was comprised of cell 1 and the cells with which it synchronized the best. Then the method was applied to these larger assemblies to measure their collective synchrony. This procedure is helpful to pinpoint larger assemblies to investigate although not

necessary if promising groups can be identified through other means, e.g. receptive field analysis. As a reminder, pair-wise calculations are not needed to evaluate the score for a larger assembly.

Significance testing was accomplished by comparing the distribution of Raw Scores to the distribution of Chance Scores (1 Raw Score and Chance Score obtained per trial) in a Paired Student's t-test under the condition that the difference of scores was normally distributed. Most score distributions were normal as determined by the Lilliefors test for goodness of fit to a normal distribution. Although the t-test is fairly robust to deviations from normalcy, nonparametric tests like the Wilcoxon signed rank test or the sign test can be used in situations where the distributions are not normal. Alternatively, the data can be transformed into a normal distribution before testing with a Paired Student's t-test (e.g. the natural logarithm transform can correct for skewness to the right).

Stimulus-Dependence

For drifting grating and concentric ring stimulation, we found that 100% of Normalized Scores from significant 2-cell assemblies ($p < 0.01$) exhibited stimulus-dependence. In general, synchrony was highest for stimuli that were optimal for the group. Each neuron synchronizes with others best during its preferred stimulus (determined by maximum average firing rate). In a pair of cells with different preferred stimuli, this tug-of-war results in highest synchrony for the average stimulus between the two. For gratings, this explains why cells with similar orientation preferences tend to synchronize best. This was even true for cells with large differences in average firing rate (> 20 sps). Normalization corrects for synchrony's correlation with firing rate, but

45

these results suggest that significant synchrony is not directly coupled to or caused by firing rate. Otherwise, a highly discharging cell would contribute more to synchrony and the optimal stimulus would be closer to the preferred stimulus for that cell.

These results extended to larger neural assemblies, involving up to five cells for gratings and six cells for rings. These assembly sizes are limited by the nature of the recording sample and its relationship to the form of the stimuli. In samples of 28 and 23 cells, we generally find only up to five cells with the same orientation preference (for gratings) and six cells that are reasonably configured for a concentric ring stimulus. Whether larger assemblies exist is currently unknown; rather larger recorded samples will be required to find this out. Figure 2.7a shows the orientation tuning curves for three cells: X, Y, and Z (preferred orientations at 30°, 40°, and 70°, respectively). These cells form a significant functional assembly (Figure 2.7b) when presented with gratings that are optimal for the group (20° to 70°) and exhibit the maximum synchrony response for a grating near the mean orientation preference of the group (40°). In Figure 2.7c, a tightly tuned assembly (maximum orientation difference of 10°) synchronizes well for a grating at the mean orientation of the group, but synchronizes poorly for an appropriately placed concentric ring. On the other hand, Figure 2.7d shows a loosely tuned group (maximum orientation difference of 60°) that synchronizes well to a ring stimulus when all members are aligned properly, but does not synchronize as well for a grating at the mean orientation for the group.
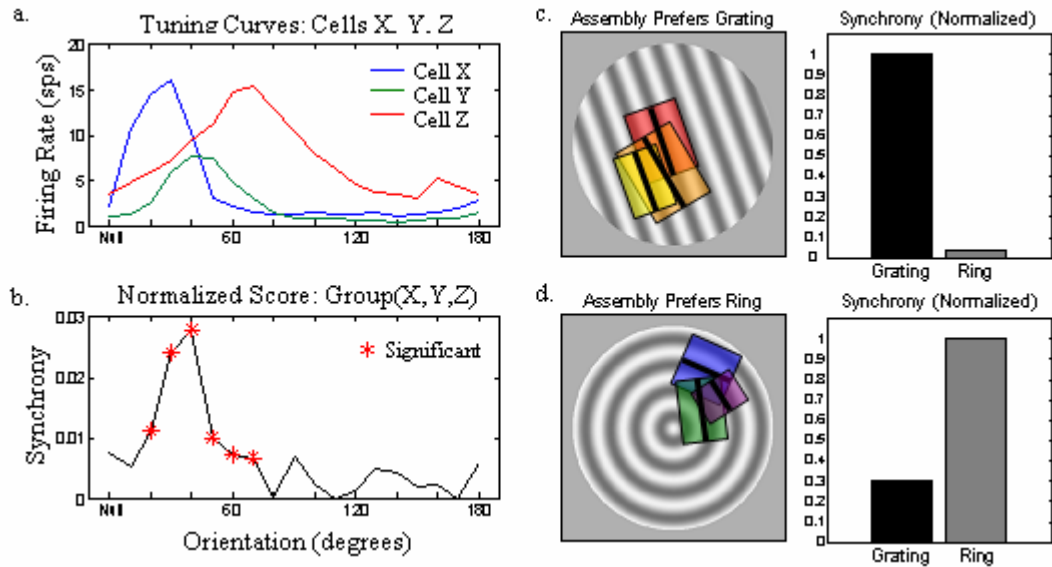
Figure 2.7:   Normalized Synchrony is stimulus-dependent. *a*. Orientation tuning curves for three cells (X – 30°, Y – 40°, Z – 70°) in an assembly. *b*. Normalized Scores for the assembly(X,Y,Z) for each grating stimulus (10° to 360° in 10° increments). The cells form a significant functional group during stimuli that are optimal for the group (20° – 70°) with the most amount of synchrony during a stimulus close to the mean orientation preference for the group (40°). *c*. A tightly tuned assembly (N = 3, largest difference is 10°) has high synchrony for an optimal grating stimulus, but low synchrony for a ring stimulus. *d*. A loosely tuned assembly (N = 3, largest difference is 60°) has high synchrony when configured to the appropriate ring stimulus, but low synchrony for a grating at the mean orientation for the group.

Dynamic Grouping

In general, we found that grouping between pairs of cells was dynamic, in that neurons could belong to more than one assembly and that grouping was stimulus-dependent as mentioned above. Dynamic grouping can occur across different stimuli (see Figure 2.8a), but also during the same stimulus (see Figure 2.8b). For instance, a cell that prefers 40° could form a significant assembly with a cell that prefers 20° when presented with a 30° grating stimulus or the cell that prefers 40° could synchronize with a cell that prefers 60° when presented with a 50° grating stimulus. Alternatively, some of the discharges from one spike train could synchronize with one cell while others

47

synchronize with a different cell during presentation of the same stimulus. Q_{time} values can be very low for pairs ($< 5\%$), which means that cells are synchronizing for less than 5% of the duration of their total activity. This information, together with the finding that approximately 80% of spikes from one spike train are synchronized with at least one other spike from another cell in the population (28 cells), suggests that dynamic grouping during the same stimulus is a common occurrence. Since cells can be involved in more than one task at a time (multitask or multiplex information in their spike trains), this increases the amount of information represented by a subpopulation of cells. This is efficient and beneficial, especially during complex stimuli where a large amount of information can be represented by a relatively few number of assemblies without having to recruit metabolically demanding numbers of cells.

The grouping within larger assemblies was also dynamic as the stimulus configuration changed or stayed the same. For instance, Figure 2.8a (left) shows synchronous contributions across a group of cells ($N = 4$) for a 240° drifting grating. (Here colors are arbitrary and used to distinguish cells.) The red cell only contributes 13% of the overall synchrony in the group. In Figure 2.8a (right), the red cell joins with a different group for a 290° grating and contributes 26%. Synchrony in the second group is more significant ($p < 0.0004$) and more evenly shared than in the first group ($p < 0.05$). The coupling of more even synchrony with stronger synchrony was a noticeable trend. Figure 2.8b graphically depicts how cells can belong to more than one functional assembly during the same stimulus. Cell B can form a significant group with A or C, but a successful group of 3 (ABC) does not exist. Similarly, Cell C forms groups with B and D, but group (BCD) does not exist.
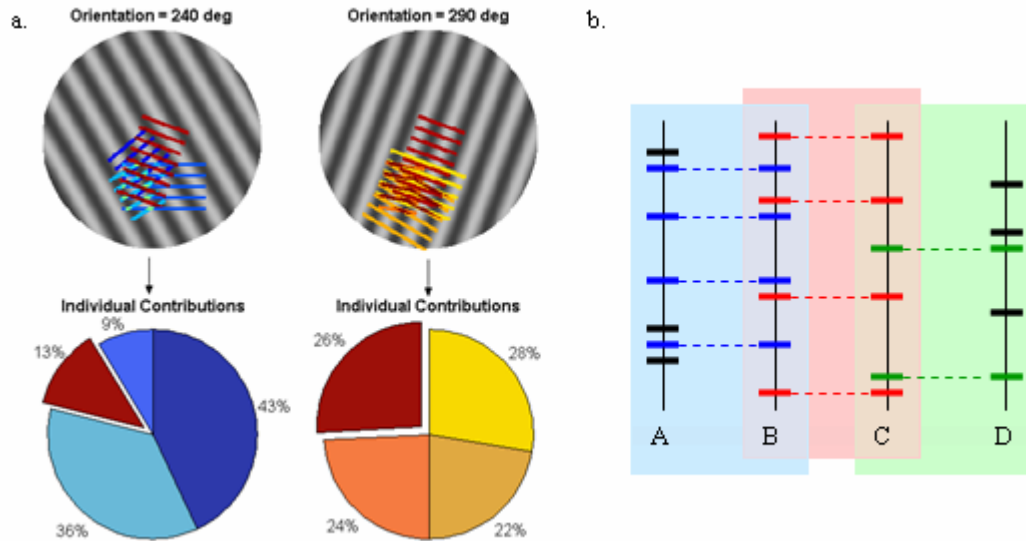
Figure. 2.8:   Dynamic grouping. *a*. (*Top-Left*) Receptive fields (preferred orientations: 230°, 240°, 270°, 290°) superimposed on a 240° drifting grating. Lines are drawn in the direction of preferred motion, which is orthogonal to the preferred orientation. (*Bottom-Left*) Individual contributions to synchrony are resolved in a pie chart. (*Top-Right*) A different group with receptive fields (preferred orientations: 280°, 290°, 290°, 300°) superimposed on a 290° drifting grating. (*Bottom-Right*) Individual contributions to synchrony. Note that the cell in red (290°) contributes more activity when it views a more optimal stimulus and forms a group with individuals that have tuning properties similar to itself. *b*. Dynamic grouping during the same stimulus occurs by sharing spikes with multiple partners. Cell B forms groups with Cells A and C, but group (ABC) does not exist since there are no spikes all three cells have in common.

Compare to JPSTH

To establish our algorithm as a viable method for quantifying synchrony, we compared it to the Joint Peri-Stimulus Time Histogram method (JPSTH, Aertsen et al. 1989). The JPSTH can only quantify synchrony among pairs of neurons, so we focused on assemblies where N = 2. We found that the PSP method identified the same significant neuron pairs as the JPSTH, but relative magnitudes of synchrony between assemblies differed. This was attributed to the non-linear weighting scheme employed by the PSP method compared to the constant (via binning) weighting scheme employed by

the JPSTH. However, selectively emphasizing spike activity that occurred closely in time allowed for the further identification of other significant assemblies that were not recognized by the JPSTH method.

To visualize whether descriptions of synchrony among neurons within the same assembly were consistent with the JPSTH method, we constructed "cross-correlograms" by shifting one spike train relative to the other and calculating the magnitude of synchrony at each lag from -100 ms to 100 ms (see Figure 2.9). This procedure is slightly different than the one used to calculate the cross-correlogram in the JPSTH method, where coincident spikes are simply summed at each time lag and the resulting histogram is smoothed and normalized. The top graph in Figure 2.9 shows the cross-correlogram for a significant neuron pair during grating stimulation constructed via the PSP method and the bottom graph shows the cross-correlogram from the JPSTH method, after further smoothing with a 5-point moving average filter. Peaks in both graphs were normalized to unity. As seen in Figure 2.9, both cross-correlograms are qualitatively similar, with large central peaks indicating significant synchrony, and periodic side lobes resembling gamma oscillation. The smooth appearance of the top graph is most likely due to the graded values available in the PSP method as opposed to discrete spike counts in the JPSTH method, which makes the resulting cross-correlogram noisy and require smoothing techniques.
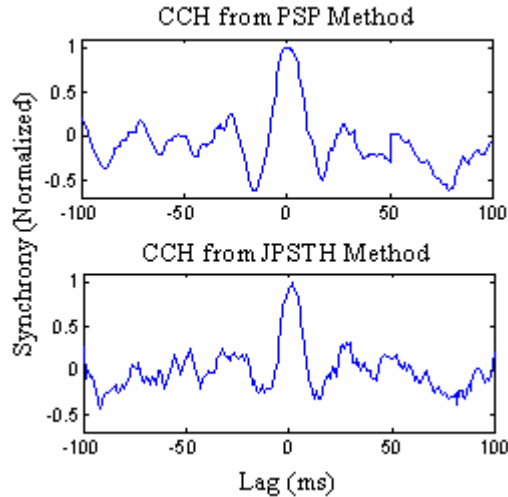
Figure 2.9: Cross-correlograms constructed via the (*top*) PSP method (detrended) and (*bottom*) JPSTH method (after further smoothing with a 5-point moving average filter). The peaks in each graph were normalized to unity for assessment of qualitative features. Each graph contains a large peak at 0 lag indicating significant synchrony. Although the main lobe in the top graph seems wider, the widths of both are comparable at 0 synchrony (normalized) and are approximately 20 ms, reflecting a 10 ms integration time period (forwards and backwards). Gamma oscillation (~ 40 Hz) is seen as side lobes in each graph. The cross-correlogram derived from the JPSTH method is noisier due to the discrete spike counts in each histogram bin as opposed to graded synchrony values at each lag in the PSP method.

Magnitude, Quality & Individual Contributions

As more and more cells are added to an assembly, the probability of having all N cells synchronize within the integration time period (e.g. 10 ms) diminishes rapidly. For example, in an assembly of 6 cells firing independently at 10 sps, the probability of encountering an instance of synchronization is approximately one in a million. However, we observed numerous assemblies (N = 6) during ring stimulation with an average Normalized Score of 0.001. When Chance encounters are negligible, the Normalized Score is approximately equal to the Raw Score and can be interpreted as the percentage of synchronous activity in an assembly. In this case, six-member assemblies are coordinating their responses 1000 times more often than that expected by chance under

51

the null hypothesis that all cells are firing independently. In a few instances, Normalized

Scores were higher than 0.05 for six-member assemblies, demonstrating significant

activity that was more than 50,000 times the expected amount.

In given assemblies, the magnitude of synchrony tends to decrease with group

size. However, the *capacity* of synchrony or the magnitude of *collective* synchrony

within an assembly including synchronous activity generated in all subassemblies

increases in a power relation to group size (see Figure 2.10a). Therefore, if synchrony

increases in proportion to stimulus salience, the amount of information represented by an

assembly including all subsets of assemblies increases in a power fashion with the

number of members. Similar to dynamic grouping, this helps to reduce the number of

cells needed to describe visual scenes, which reduces metabolic demands and increases

efficiency.

The increase in information capacity of a large assembly is intuitive as it reflects

the combinatorial arrangements of a large number of cells. Likewise, intuition tells us

that as synchronous activity increases, the quality of synchrony should increase (per the

definition of quality). As synchrony increased, we found a proportional increase in the

average amount of waveform overlap and not the number of coincidences measured via

duration of synchrony. (Note: as both quality measures involve an increase in the

duration of synchrony, we found an excess in time spent during synchronous activity that

was accounted for by the increase in overlapping area of waveforms.) Recall that

individual contributions to synchrony within an assembly can determine whether cells, on

average, fire together or tend to fire with a time delay. Relatively asynchronous

discharges (but still within the integration time period) are reflected in a large variance of

individual contributions, which indicates a low $Q_{overlap}$ score. Figure 2.10b shows that as

assemblies become more significant (have more synchrony compared to chance), the

variance of individual contributions decreases linearly, which indicates a linear increase

in quality. Therefore, for highly synchronous assemblies (see Figure 2.8a), each member

contributes nearly equal amounts to maximize the amount of overlapping waveforms to

increase synchrony without increasing the number of coincidences and taking spikes that

could be used in other assemblies (see Figure 2.8b). For 1 out of 2 concentric ring

experiments, quality increased with assembly size. As can be seen in Figure 2.10b, the

largest ring assemblies (N = 6) were already producing near-maximal overlap.
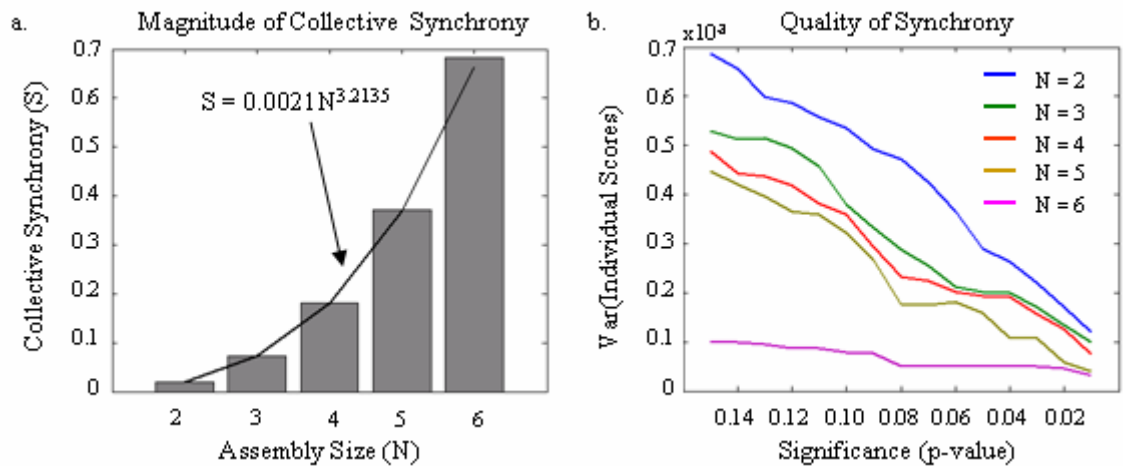
Figure 2.10: *a*. The magnitude of collective synchrony (measured from an assembly and all subsets within the assembly) increases in a power relation to assembly size. *b*. The quality of synchrony (measured inversely as the variance of individual contributions to synchrony) increases as assemblies become more significant and, in 1 out of 2 concentric ring experiments, increases with assembly size.

Synchrony and Reliability

While the PSP method can be used for data analysis, results from the method can

also be used to streamline experimental protocols and improve efficiency. For instance,

we calculated the number of trials needed to reliably produce synchrony. We found that

total spike number in an assembly is a good indicator of reliability. In 3 different

experiments, we measured the significant synchrony among 295 pairs (n = 27 cells) and

found that 90% achieved 90% reliability with less than 1527 spikes (top, Figure 2.11). If

the average neuron fires 15.15 spikes/sec (bottom graph), a pair of average neurons

would fire 30.30 spikes/sec and they would need to generate spikes for 50.4 seconds to

yield a synchrony score that is 90% reliable. Thus for the average neuron, 26 repetitions

of a stimulus duration of 2 seconds would be required to yield this general level of

reliability. We have previously used 100 or 200 repetitions, but future experiments can

test 4 – 8x more stimuli in the same amount of time and yield the same quality of data.
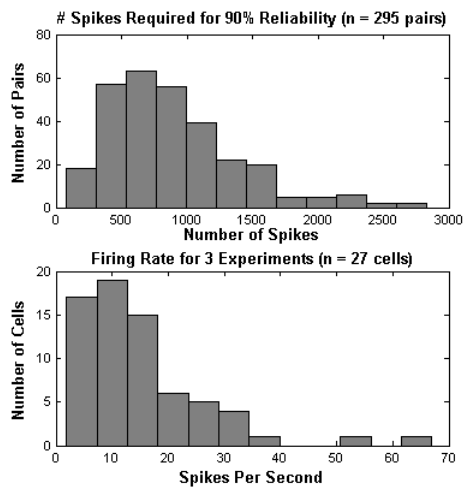

Figure 2.11:   *Top*: Histogram of total spike count required to measure significant
synchrony 90% of the time (n = 295 pairs). *Bottom*: Firing rate in spikes per second
during 3 experiments (n = 27 cells).

Discussion

We have derived a method based on the biological process of postsynaptic

potential integration to quantify the magnitude and quality of synchronous activity within

neural assemblies of arbitrary size. As a substrate for encoding stimulus properties like contours (Samonds et al. 2005), synchrony enhances the probability of eliciting postsynaptic action potentials when neurons behave as coincidence detectors (Azouz and Gray 2003). This ensures the propagation of information to subsequent levels of the cortical hierarchy. Synchrony allows for neurons to be effectively connected and form dynamic circuits to reliably and efficiently transmit information throughout the cortex.

By introducing a method whose output is relevant to the postsynaptic neuron, we have also derived a measure of efficiency. Higher scores reflect more efficient transmission in that threshold can be reached more quickly than with assemblies with lower scores. Synchrony, itself, represents an efficient coding strategy as described in the sections on dynamic grouping and magnitude of collective synchrony. Dynamic grouping allows for the formation of transient functional groups during different stimuli or separate functional groups during the same stimulus. Victor (2000) suggested that stimulus information is multiplexed at different temporal resolutions of the interspike interval histogram, but we suggest that cells can multiplex information in their spike trains by forming separate assemblies during the same stimulus. These multiple assemblies may form a larger assembly whose collective synchrony increases with a power relation to group size. Therefore, synchrony allows for multitasking so that visual information can be processed with a minimum number of cells. Conversely, a larger number of cells can process more complicated stimuli.

This method is advantageous because: (1) It can be applied to an arbitrary number of cells. (2) Any waveform can be chosen and parameters like PSP amplitude, integration time and threshold are adjustable. (3) The spike trains from excitatory and inhibitory cells can both be represented (an IPSP train results from convolution with a

negative alpha function). (4) Simple significance testing can be used. (5) We can normalize synchrony scores by subtracting the shift predictor. (6) Members of large groups can be predicted from the scores of smaller groups, thus drastically reducing the number of group permutations that need to be computed. (7) No binning or smoothing techniques are required. (8) This method is based on biological principles and the results are physically meaningful (i.e. a score of 0.15 means that 15% of all waveforms in a group were coincident). The flexibility of this approach allows us to investigate the effects of increased/decreased integration times, examine group characteristics, and decipher trends correlated with group size. In order to understand group dynamics and document rules governing group membership, we can use this method to measure collective responses and explore temporal characteristics of neural assemblies.

Conclusion

Focus has centered on recording techniques and the invention of novel ways to measure population activity. The assumption is that the more neurons we can record from simultaneously, the more we will understand about population effects in the brain. However, even if we could measure from every neuron in the brain, we still would not know how cells interact on a large scale because we do not have a way to analyze large populations. The PSP method was introduced to remedy this situation. As the technology grows, so with it must grow our analytical resources. By studying pairs of neurons and their interactions, we have learned a great deal about what the brain is doing beyond that known for single cells. Expansion of this study to relationships between tens, hundreds, or thousands of cells or more offers an entirely new frontier in neuroscience.

CHAPTER III

FUTURE EXPLORATIONS

Introduction

For decades, study of cooperative relationships was mostly confined to

interactions among cell pairs because simultaneous recording of the behavior of large

populations was impractical. This in turn discouraged the development of analysis

techniques for comparing more than two simultaneous neuronal responses. In the

previous chapter, we introduced a novel method to describe cooperative activity among

any neural assembly of arbitrary size. By modeling the biological process of postsynaptic

potential integration, both the magnitude and quality of synchrony in an assembly can be

quantified. Introduction of this method is timely and addresses the need to explore the

characteristics of larger neural assemblies. Neither the properties of single neurons nor

pairs of neurons can explain the behavioral and perceptual repertoire of the brain.

Understanding the functional interactions and emergent properties of larger neural

networks is crucial to linking neurophysiology and perception. The basis for our current

models for functioning of the visual system has been developed from the analysis of

individual or pairs of neurons, but we cannot grasp the foundations of perception without

exploring how cells work together in local networks to describe the visual environment.

By stepping outside the abstract frameworks for identifying synchrony, we move

closer to exploring synchrony as a physical neural mechanism for transmission of salient

information. However, as our measures become more realistic, so must our experimental

paradigm. For future research, we suggest an in-depth investigation into the role of

synchrony and assembly formation in the context of natural vision. Furthermore, synchrony should be investigated as a viable sparse code employed by the visual cortex to represent high-order stimulus features in natural scenes. The following pages discuss natural stimulation and natural scene statistics, sparse coding as an information processing strategy, synchrony as a viable sparse code, and some possible experiments to investigate which natural stimulus features induce synchrony. Finally, we will end with some concluding remarks.

Using Natural Stimulation

The historic use of spatially or spectrally pure stimuli (lines, gratings) allowed controlled and systematic manipulation of the visual environment. This was founded on the presumptions (from linear systems theory) that responses to simple stimuli could be used to describe the processing of more complex stimuli and has contributed immensely to our understanding of neuronal response properties. However, more recent descriptions of previously unsuspected interactions such as receptive field reorganization due to stimulation of the "non-classical" periphery (Pettet and Gilbert 1992) suggest that we may not yet have a realistic model for cell behavior in the context of complex, natural scenes. For instance, David et al (2004) found that natural tuning properties are not well-predicted from responses to grating sequences. This could be due in part to the reduced effectiveness of simple stimuli as opposed to natural scenes to drive cortical cells (Rieke et al. 1995). Since responses of V1 neurons to natural stimuli are both qualitatively and quantitatively different from those to simple stimuli (Baddeley et al. 1997; Kayser et al. 2003), it follows that an investigation of natural stimulation is necessary for understanding cell behavior in the context of natural vision.

Unlike simple stimuli, natural scenes have complex high-order spatial correlations (Field et al. 1993; Schwartz and Simoncelli 2001). In an image, neighboring locations have similar color and intensity values and line segments tend to be arranged in a co-linear or co-circular fashion. For example, Figure 3.1 shows three pictures of a lamb with corresponding pixels being averaged from left to right. Averaging neighboring pixel values mostly removes redundancy in the image and results in minimal loss of salient information. Despite the loss in resolution, the lamb can still be identified in all three images. This situation is different from white noise, in which there is no identifiable structure because neighboring pixel intensities are random and uncorrelated. First-order image statistics describe individual pixel intensities, whereas second-order statistics describe correlations between pairs of pixel values. Higher-order statistics are needed to describe local features such as contours, surfaces, and textures. In order to represent the salient stimulus features in a natural scene, neuronal responses must be able to represent high-order image statistics.

Figure 3.1: There are high-order spatial correlations in natural images. This is an image of a lamb rendered at three different resolutions. Left: 64 x 64 pixels, Center: 32 x 32 pixels, Right: 16 x 16 pixels. The lamb can still be identified because neighboring pixel values are correlated and the process of averaging removes this redundancy with only minor information loss.

Sparse Coding as an Information Processing Strategy

But how can neuronal responses represent contours, surfaces, and textures in natural scenes? This is dependent on the more general question of how information about the natural environment is encoded by the early visual system. While the traditional model involving cells acting as independent filters doubtless contributes to this process, it fails to account for any interactions between the filters making up the population and, as discussed above, the previously-unsuspected dynamic nature of the filters. Although current technology allows us to monitor the interactions of hundreds of neurons simultaneously, we cannot understand their contribution to visual perception without a conceptual framework that describes their information processing strategy. Discussions from Chapter 1 present the advantages of employing a population code in which subsets of the neural population form assemblies to accomplish tasks. Two competing theories within this general approach include *compact coding* and *sparse coding*, both of which take advantage of redundancy in natural scenes to produce more efficient representations of the environment.

In a compact coding strategy (see Figure 3.2), the goal of visual coding is to reduce the redundancy in natural scenes by creating a representation with the minimum number of cells (Field 1994). This idea is closely related to Barlow's theories of redundancy reduction (Barlow 1961) as well as dimensionality reduction using Principal Components Analysis. Different objects are represented by the different firing patterns of the same subset of cells. Although a number of studies have suggested that spatial coding by the visual system is consistent with a compact code (e.g. Atick and Redlich 1990, 1992; Barlow and Foldiak 1989; Daugman 1988, 1991; Linsker 1988; Sanger 1989), there are several inconsistencies in accounting for receptive field properties of

cells in the retina and primary visual cortex. By not including phase information due to the stationarity of natural scenes (i.e. features do not prefer any specific spatial location), compact coding cannot account for the localized nature of receptive fields (Field 1994).

On the other hand, the visual system may be optimized for processing the statistics of natural scenes (Barlow 1961; Kersten et al. 1987; Simoncelli 2003) by employing a sparse coding strategy where natural scene information is represented with the minimum number of *active* cells. In this strategy (see Figure 3.2), the dimensionality is not reduced, but the redundancy in the input is transformed into the redundancy in the firing pattern of the cells. All cells have an equal response probability across the class of input images, but have a low response probability for any single image. In this manner, information about the environment is distributed across all cells and objects are represented by which cells are active and not by the relative activity or overall amount of activity of a specific subset. This approach has been found to be consistent with the representations of natural scenes (Field 1987, 1989, 1993, 1994; Zetsche 1990) as well as other sensory information (Barlow 1972, 1985; Palm 1980; Baum et al. 1988).

One of the main advantages of sparse coding is that it can assist in the process of recognition and generalization (Field 1994). Since the response of any one cell is relatively rare, tasks that require matching or detecting corresponding features are more successful. As a code becomes more sparse (i.e. lowered probability of one cell responding), the probability of detecting a correct correspondence increases. Furthermore, if the probability of any cell response is low, the probability of two cells responding is even lower, assuming response independence (this is still true even though cortical responses to natural stimulation are not completely independent – Field 1994). Higher-order relations requiring large neural assemblies are increasingly rare and thus

more informative when present. High-order features in nature can then be represented by

a unique subset of cells.

Figure 3.2: Two competing theories of information processing (adapted from Field 1994). *Left*: In a compact coding strategy, the redundancy in the input is reduced by representing data with the minimum number of units. This reduces the dimensionality and objects in the visual scene are represented by the relative activity across all cells in the subset. *Right*: In a sparse coding strategy, the redundancy in the input is transformed into the redundancy in the firing rate of each unit and only a few units are active for any given input. Different objects are represented by which cells within the broader population are active.

In fact, Hoyer and Hyvarinen (2002) have shown that a multi-layer sparse coding

network is capable of learning contour coding from natural images in an unsupervised

fashion. By modeling complex cell responses in which firing rate response distributions

were both sparse and non-negative (Hyvarinen and Hoyer 2000), contour coding and end-

stopped receptive fields emerged. Neural network models have also shown that when

inputs to the network are sparse, they can store more information and provide more

effective retrieval with partial information (Palm 1980; Baum et al. 1988). Storage and

retrieval with associative memories are biologically plausible and so efficient that

preprocessing in many networks artificially sparsifies data so that learning or classifying

fewer higher-order relations requires less computation.

Although the work mentioned above is theoretical and based on neural network implementations, there is also experimental evidence that supports a biological sparse coding scheme for natural vision. For instance, neurons produce sparse, distributed responses to natural stimulation (Vinje and Gallant 2000; Weliky et al. 2003) and average firing rates are low compared to that from optimal grating stimulation, which reduces the metabolic demands of visual processing (Baddeley et al. 1997; Guo et al. 2005). Furthermore, natural scenes are large and typically cover cells' classical and non-classical receptive fields. Large stimuli produce high sparseness values (Vinje and Gallant 2000), and Guo et al (2005) found that the majority of center-surround interactions in V1 cells are sensitive to high-order structures in natural images. Center-surround interactions enable neurons to integrate information outside their classical receptive fields, which is important for contour integration (Gilbert 1998; Gilbert and Wiesel 1990; Fitzpatrick 2000).

Synchrony as a Viable Sparse Code

So far, we have speculated that the brain may employ a population code and that the population code may be organized as a sparse code. Recall, however, that a population code requires specific rules for association among cells. While sparse coding requires a subset of cells to be active for any given input, what properties of the input define the subsets formed? Is formation guided by intrinsic properties such as anatomical connections among cells with similar tuning preferences, extrinsic properties such as features within the stimulus, or both? Furthermore, how is this property represented by the subset? Since a given subset is particular to certain inputs, logic tells us that this property may be stimulus-dependent and relatively rare across the population of inputs.

As mentioned above, higher-order features (contours, surfaces, textures) may be represented in a sparse code. We propose that this representation is reflected in the cooperative activity of the subset. In other words, synchrony is a way to implement a sparse coding strategy where higher-order stimulus information is represented in the precise temporal pattern of a neural assembly. Precise temporal coordination among neurons preserves the requirements for sparse coding in that information is transmitted efficiently (through cooperation) and by a few neurons (assembly).

We have shown that responses from synchronized assemblies in cat visual cortex can represent curvilinear contours (see Figure 3.3; Samonds et al. 2005). When presented with a concentric ring stimulus, two cells of differing orientation preferences that lie tangent to the same ring can show synchronized responses as measured by the normalized JPSTH (Aertsen et al. 1989). Responses from neuron pairs were compared when stimulated with drifting sinusoidal grating and concentric ring stimuli and the results suggested that synchronous activity was more selective and reliable than changes in firing rate in discriminating between the two stimuli. Extending assembly formation to cells with spatial relationships that are not collinear is vital in establishing whether synchrony might serve as a general mechanism for contour integration and shape detection.


Figure 3.3: Synchrony can represent curvilinear contours. *Left*: Cell pairs with differing orientation preferences are stimulated with a drifting concentric ring stimulus. *Right*: Normalized cross-correlogram (JPSTH, Aertsen et al. 1989) with a large central peak indicating significant synchrony.

Possible Experiments

Given that artificial sparse coding networks can learn contour coding (Hoyer and

Hyvarinen 2002) and our *in vivo* results suggesting that synchrony may encode contours

(Samonds et al. 2005), a logical inference may be that synchrony forms an

implementation of a sparse code in the (visual) cortex. However, more experiments

showing that synchrony is consistent with a sparse coding strategy must be performed

before any definitive conclusions can be drawn. For instance, our results are based on

pair-wise responses to artificial concentric ring stimulation and the neural network results

were only found for collinear contours. Whether our findings extrapolate to larger neural

assemblies under natural stimulation and the neural network results extrapolate to

curvilinear contours are unknown. Furthermore, we do not yet know what properties of

the stimulus affect synchrony and whether these same properties influence sparse codes.

Field (1994) proposed that the property of an image that allows for a sparse code is the

phase spectrum and that sparse codes should have response distributions with high

kurtosis ("peakedness"). Three possible experiments that investigate whether synchrony

fits this criterion are discussed.

Experiment #1: Does synchrony signal contours in natural images?

We can explore which stimuli are most effective for forming neural assemblies by

presenting many natural images with a variety of features. In a method adapted from

Smyth et al (2003), natural image sequences can be presented and images that elicit a

synchronized group response can be identified. Those images can be summed together

and weighted by the magnitude of the response. The resulting composite will synthesize

a stimulus attribute seen in each picture or combine different structural features to create a new attribute that is ideal for the assembly.

We expect that this approach will yield contours that are preserved in the response set or piece segments from different images to produce a contour not present in the stimulus sequence. As a preliminary finding, we have applied this method to two groups of three cells each from V1 of an anesthetized cat. We presented each group with a set of drifting gratings, oriented 10° to180° in 10° increments. The gratings eliciting significant cooperative responses (t-test, $p<0.01$) were weighted and summed. The results are shown in Figure 5. For the tightly tuned group (preferred orientations: 60°, 60°, and 70°), the optimal stimulus appears to be a grating at the mean orientation for the group (60°). For the loosely tuned group (preferred orientations: 90°, 100°, and 110°), however, the optimal stimulus appears to be a ring that is tangent to the orientation and location of each cell in the group. In fact, when presented with a ring stimulus, the synchrony in the loosely tuned group was higher than that measured for a grating in the mean direction for the group (100°). These results indicate that it is possible to find an optimal group stimulus or create one by combining salient features. We note that a limited stimulus set was used in this case and that a truly optimal stimulus may involve higher-order features that cannot be defined by simple combination of gratings. In response to this, many natural images with a variety of structural features and stimulus properties can be presented.

Figure 3.4: Composite stimuli with superimposed receptive fields of the cells in an assembly. *Left*: The composite from a tightly tuned group resembles a grating at the mean orientation for the group. *Right*: The composite stimulus for a loosely tuned group resembles a ring.

Experiment #2: Does synchrony depend on the phase spectrum of a natural image?

Field (1994) suggested that the phase spectrum of a natural image describes the redundancy necessary for sparse coding (when the input has stationary statistics). We can determine whether synchrony depends on the phase spectrum of an image by altering it. For instance, a stimulus (Figure 3.5 left) can be separated into its power and phase components in the frequency domain. By swapping components with, say, a random white noise stimulus (Figure 3.5 left-middle), images can be produced that have a natural phase spectrum and random power spectrum (Figure 3.5 right-middle) or a random phase spectrum and natural power spectrum (Figure 3.5 right)(Yang Dan, pers. comm.).

Figure 3.5: Swapping image components in the frequency domain to create hybrid stimuli. *Left*: Natural image. *Left-middle*: Random white noise stimulus. *Right-middle*: Image created with the phase spectrum from the natural image and the power spectrum from the random stimulus. *Right*: Image created with the phase spectrum from the random stimulus and the power spectrum from the natural image.

If Field's theory on sparse coding is correct and synchrony reflects the implementation of a sparse code in the cortex, we would expect to find that synchronized activity is stronger when viewing hybrid images with natural phase spectra than when stimulated with images that have random phase spectra.

Experiment #3: What other properties of natural images affect synchrony?

One challenge in analyzing neuronal responses to natural stimuli lies in pairing the responses to a specific stimulus attribute. Spatial, temporal, and luminance properties may combine in any number of ways to elicit a given response. To investigate any other specific properties of natural stimuli that encourage assembly formation, differential measurements can be utilized. A general strategy is to measure a response to an original, control image and then systematically vary the properties of that image. Cooperative responses can be compared to those elicited from the original image and the difference represents the consequences of the imposed modification. Such modifications may include: polarity-reversal, noise-addition, edge-enhancement, low-pass filtering, or high-pass filtering.


Figure 3.6: Image modifications. *Top-Left*: Original image serves as a control. *Top-Right*: Polarity-reversed image. *Bottom-Left*: Image with noise added. *Bottom-Right*: Edge-enhanced image.

The use of a difference approach allows us to examine whether the modifications mentioned above affect neural assembly activity and formation without having to know exactly which features are preferred by a group. For instance, if the receptive fields in a neural assembly cover the eye of the toucan (Figure 3.6), we can assess sensitivity to noise-addition without having to analyze which specific feature (curvature, pixel intensity, shape of surroundings) induces synchronization. We expect to find that modifications that increase the ability to detect contours will increase synchrony, while modifications that destroy structure will decrease synchrony. This result will support our hypothesis that synchrony signals high-order stimulus features such as contours. Collectively, these results address whether synchrony depends on specific local structure or more general qualities such as shape or contour.

Conclusion

Encoding of salient stimulus information by the brain has been a topic of discussion for over a century. Similarly, the benefits of sparse coding have been known to the computational world for decades. Speculations on a biological sparse coding strategy have surfaced, but the technology was never available to confirm or dispute such theories. The introduction of a method as described in Chapter 2 allows us to move beyond the analysis of a couple neurons to investigate such theories and determine how populations of neurons may be working together, perhaps with a synchronous sparse code, to generate emergent behavior such as contour integration and perception of our visual environment.

APPENDIX

Determining the filter, F(t).

The filter, F(t), has a value of 1 at times when the waveforms are synchronous and
0 elsewhere. Although there are many ways to create this filter, a simple method is
described here. Each PSP train in an assembly is multiplied together, ensuring that
positive values in the resulting train are synchronous. The resulting train is then
normalized by the largest value to create values between 0 and 1. Finally, the ceiling
function is used to make all positive values equal to 1.

$$F(t) = ceil\left(\frac{\prod_{i=1}^{N} P_i(t)}{max\left(\prod_{i=1}^{N} P_i(t)\right)}\right)$$

Calculating the Number of Coincident Spikes, N_{CS}.

The number of coincident spikes within an assembly can be calculated by
comparing spike initiation times from spike trains with the Filter train. At a certain spike
initiation time, check the next L_W elements of F(t). If any element is a 1, then count that
spike as a coincident spike. This algorithm is demonstrated below in MATLAB.

```
Ncs = 0;
for t = SpikeTimes
    if sum(F(t:t+Lw-1)) > 0
        Ncs = Ncs + 1;
    end
end
```

where *SpikeTimes* is a vector containing spike initiation times for all cells in the
assembly.

REFERENCES

Abeles M., Bergman H., Margalit E., and Vaadia E. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol.*, 70:1629-1638, 1993.

Abeles M. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, New York, 1991.

Abeles M. and Gerstein G.L. Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.*, 60:909-924, 1988.

Abeles M. Role of cortical neurons: integrator or coincidence detector? *Isr. J. Med. Sci.*, 18:83-92, 1982.

Abeles M. and Goldstein M. Jr. Multiple spike train analysis. *Proc. IEEE*, 65:762 -773, 1977.

Adrian E.D. and Zotterman Y. The impulses produced by sensory nerve endings, Pt. 3: Impulses set up by touch and pressure. *J. Physiol. (Lond)*, 61:465-493, 1926.

Aertsen A.M.H.J., Gerstein G.L., Habib M.K., and Palm G. Dynamics of neuronal firing correlation: Modulation of "effective connectivity". *J. Neurophysiol.*, 61:900-917, 1989.

Aronov D., Reich D.S., Mechler F., and Victor J.D. Neural coding of spatial phase in V1 of the macaque monkey. *J. Neurophysiol.*, 89:3304-3327, 2003.

Aronov D. Fast Algorithm for the metric-space analysis of simultaneous responses of multiple single neurons. *J. Neurosci. Meth.,* 124:175-179, 2003.

Atick J.J. and Redlich A.N. What does the retina know about natural scenes? *Neural Comp*. 4:449-572, 1992.

Atick J.J. and Redlich A.N. Towards a theory of early visual processing. *Neural Comp*. 4: 196-210, 1990.

Averbeck B.B. and Lee D. Coding and transmission of information by neural ensembles. *Trends Neurosci.* 27:225-230, 2004.

Azouz R. and Gray C.M. Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. *Neuron,* 37:513-523, 2003.

Azouz R. and Gray C.M. Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *PNAS*. 97(14):8110-8115, 2000.

Baddeley R., Abbott L.F., Booth M.C., Sengpiel F., Freeman T., Wakeman E.A., and Rolls E.T. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc. R. Soc. London Ser. B* 264:1775–1783, 1997.

Bair W., Koch C., Newsome W., and Britten K. Power spectrum analysis of bursting cells in area MT in the behaving monkey. *J. Neurosci.*, 14: 2870-2892, 1994.

Baker S.N. and Gerstein G.L. Improvements to the sensitivity of gravitational clustering for multiple neuron recordings. *Neural Comp.*, 12:2597-2620, 2000.

Barlow H.B. The Twelfth Bartlett Memorial Lecture: The role of single neurons in the psychology of perception. *Q. J. Exp. Psychol.* 37A:121-145, 1985.

Barlow H.B. Single units and sensation: a neuron doctrine for perceptual psychology? *Perception*, 1:371-394, 1972.

Barlow H.B. Possible principles underlying the transformation of sensory messages. In *Sensory Communication*, ed. WA Rosenblith, Cambridge, MA: MIT Press, 217–234, 1961.

Barlow H.B. and Foldiak P. Adaptation and decorrelation in the cortex. In *The Computing Neuron*, R. Durbin, C. Miall, and G. Mitchison, edt., 54-72. Addison-Wesley, Reading, MA. 1989.

Bauer R., Eckhorn R. and Jordan W. Iso- and cross-oriented columns in cat striate cortex: a study with simultaneous single- and multi-unit recordings. *Neuroscience* 30:733-740, 1989.

Baum E.B., Moody J., and Wilczek F. Internal representations for associative memory. *Biol. Cybern.* 59:217-228, 1988.

Bear, M.F., B.W. Conners, and M.A. Paradisa. *Neuroscience: Exploring the Brain*, Baltimore: Lippincott. 2001.

Brown E.N., Kass R.E. and Mitra P.P. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat. Neurosci.* 7:456-461, 2004.

Buzsaki G. Large-scale recording of neuronal ensembles. *Nat. Neurosci.* 7:446-451, 2004.

Castelo-Branco M., Goebel R., Neuenschwander S., and Singer W. Neural synchrony correlates with surface segregation rules. *Nature*, 405:685-689, 2000.

Cox D.R. and Lewis P.A.W. The statistical analysis of series of events. John Wiley and Sons, Inc., New York, 1966.

Chapin J.K. and Nicolelis M.A.L. Principal component analysis of neuronal ensemble activity reveals mulidimensional somatosensory representations. *J. Neurosci. Methods*, 94:121-140, 1999.

Czanner G., Grun S., and Iyengar S. Theory of the snowflake plot and its relations to higher-order analysis methods. *Neural Comput.* 17:1456-1479, 2005.

Daugman J.G. Self-similar oriented wavelet pyramids: Conjectures about neural non-orthogonality. In A. Gorea, ed., *Representations of Vision*. Cambridge University Press, Cambridge, 1991.

Daugman J.G. Complete discrete 2-D Gabor transforms by neural networks for image analysis and compression. *IEEE Transact. Acoustics, Speech, Signal Process.* 36(7):1169-1179, 1988.

David S.V., Vinje W.E., and Gallant J.L. Natural stimulus statistics alter the receptive field structure of v1 neurons. *J. Neurosci.* 24:6991-7006, 2004.

Dayhoff J.E. and Gerstein G.L. Favored patterns in spike trains. I. Detection. *J. Neurophysiol.*, 49:1334-1348, 1983a.

Dayhoff J.E. and Gerstein G.L. Favored patterns in spike trains. II. Application. *J. Neurophysiol.*, 49:1349-1363, 1983b.

Deismann M., Gewaltig M., and Aertsen A. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402:529-531, 1999.

Eckhorn R. Oscillatory and non-oscillatory synchronizations in the visual cortex and their possible roles in associations of visual features. *Prog. Brain Res.* 102:405-426, 1994.

Eckhorn R., Bauer R., Jordan W., Brosch M., Kruse W., Munk M., and Reitboeck H.J. Coherent Oscillations: A mechanism of Feature Linking in the Visual Cortex? *Biol. Cybern.*, 60:121-130, 1988.

Eckhorn R. and Obermueller A. Single neurons are differently involved in stimulus-specific oscillations in cat visual cortex. *Exp. Brain Res.* 95:177-82, 1993.

Efron B. and Tibshirani R. J. *An introduction to the bootstrap*. Chapman and Hall, New York, 1993.

Elder J.H. and Goldberg R.M.. Ecological statistics of Gestalt laws for the perceptual organization of contours. *J. Vis.* 2:324-353, 2002.

Engel A.K., König P., Kreiter A.K., and Singer W. Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science*, 252:1177-1179, 1991a.

Engel A.K., Kreiter A.K., König P., and Singer W. Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proc. Nat. Acad. Sci. USA,* 88:6048-6052, 1991b.

Engel A.K., Fries P., Konig P., Brecht M., and Singer W. Temporal binding, binocular rivalry, and consciousness. *Conscious Cogn.* 8:128-151, 1999.

Field D.J. What is the goal of sensory coding? *Neural Comp.* 6:559-601, 1994.

Field D.J., Hayes A., and Hess R.F. Contour integration by the human visual system: evidence for a local "association field". *Vision Res.* 33:173-193, 1993.

Field D.J. What the statistics of natural images tell us about visual coding. *Proc. SPIE.* 1077:269-276, 1989.

Field D.J. Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Amer.* 4:2379-2394, 1987.

Fitzpatrick D. Seeing beyond the receptive field in primary visual cortex. *Curr. Opin. Neurobiol.*10:438-443, 2000.

Frien A., Eckhorn R., Bauer R., Woelbern T., and Gabriel A. Fast Oscillations display sharper orientation tuning than slower components of the same recordings in striate cortex of the awake monkey. *Eur. J. Neurosci.,* 12:1453-1465, 2000.

Gal V., Grun S., and Tetzlaff R. Analysis of multidimensional neural activity via CNN-UM. *Int. J. Neural. Syst.* 13:479-487, 2003.

Geisler W.S., Perry J.S., Super B.J., and Gallogly D.P. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res.* 41:711-724, 2001.

Gershon E.D., Weiner M.C., Lathan P.E., and Richmond B.J. Coding strategies in monkey V1 and inferior temporal cortices. *J. Neurophysiol.* 79:1135-1144, 1998.

Gerstein G.L. and Aertsen A.M.H.J. Representation of cooperative firing activity among simultaneously recorded neurons. *J. Neurophysiol.* 54:1513-1528, 1985.

Gerstein G.L. and Perkel D.H. Mutual temporal relationships among neuronal spike trains. Statistical techniques for display and analysis. *Biophys. J.* 12:453-473, 1972.

Gerstein G.L., Perkel D.H., and Dayhoff J.E. Cooperative firing activity in simultaneously recorded populations of neurons: Detection and measurement. *J. Neurosci.,* 5:881-889, 1985.

Gerstein G.L. and Perkel D.H. Simultaneously recorded trains of action potentials: analysis and functional interpretation. *Science.* 164(881):828-830, 1969.
Gerstein G.L. and Kiang N.Y.S. *Biophys. J.* 6:15, 1960.

Gilbert C.D. Adult cortical dynamics. *Physiol. Rev.* 78:467-485, 1998.

Gilbert C.D. and Wiesel T.N. The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Res.* 30:1689-1701, 1990.

Gray C.M. and Singer W. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. Nat. Acad. Sci.* USA, 86:1698-1702, 1989.

Gray C.M., Konig P., Engel A.K., and Singer W. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334-337, 1989.

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

Grun S., Diesmann M., and Aertsen A. Unitary Events in Multiple Single-Neuron Spiking Activity: I. Detection and Significance. *Neural Comp.,* 14:43-80, 2002a.

Guillory, K.S. and Normann, R.A. A 100-channel system for real time detection and storage of extracellular spike waveforms. *J. Neurosci. Methods* 91:21-29, 1999.

Guo K., Robertson R.G., Mohamoodi S., and Young M.P. Centre-surround interactions in response to natural scene stimulation in the primary visual cortex. *Eur. J. Neurosci.* 21:536-548, 2005.

Hayek F.A. *The Sensory Order*. Chicago: University of Chicago Press, 1952.

Hebb D.O. *The Organization of Behavior: a Neuropsychological Theory.* New York: Wiley, 1949.

Hess R.F., Hayes A., and Field D.J. Contour integration and cortical processing. *J. Physiol. Paris.* 97:105-119, 2003.

Hoyer P.O. and Hyvarinen A. A multi-layer sparse coding network learns contour coding from natural images. *Vision Res.* 42(12):1593-1605, 2002.

Hoyer P.O. and Hyvarinen A. Independent component analysis applied to feature extraction from colour and stereo images. *Network: Computation in Neural Systems.* 11(3):191-210, 2000.

Johnson D.H., Gruner C.M., and Glantz R.M. Quantifying information transfer in spike generation. *Neurocomp.* 32-33:1047-1054, 2000.

Johnson D.H., Gruner C.M., Baggerly K., and Seshagiri C. Information-theoretic analysis of neural coding. *J. Comp. Neurosci.*, 10:47-69, 2001.

Kandel E., Schwartz J., and Jessel T.. *Principles of Neural Science, Fourth Edition*, McGraw-Hill, New York. 2000.

Kayser C., Salazar R.F., and Konig P. Responses to natural scenes in cat V1. *J. Neurophysiol*. 90:1910-1920, 2003.

Kersten D. Predictability and redundancy of natural images. *J. Opt. Soc. Am. A* 4:2395 – 2400, 1987.

Konig P. A method for the quantification of synchrony and oscillatory properties of neuronal activity. *J. Neurosci. Methods*. 53:31-37, 1994.

Kreiter A.K. and Singer W. Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J. Neurosci.* 16:2381-2396, 1996.

Kruger J. and Bach M. Simultaneous recording with 30 microelectrodes in monkey visual cortex. *Exp. Brain Res.*, 41:191-194, 1981.

Lamme V.A.F. and Spekreijse H. Neural synchrony does not represent texture segregation. *Nature*, 396:362-366, 1998.

Laubach M., Shuler M., and Nicolelis M.A.L. Independent component analysis for quantifying neuronal ensemble interactions. *J. Neurosci. Methods*, 94:141-154, 1999.

Lee D. Analysis of phase-locked oscillations in multi-channel single-unit spike activity with wavelet cross-spectrum. *J. Neurosci. Meth*. 115:67-75, 2002.

Legendy C.R. Three principles of brain function and structure. *Int. J. Neurosci*. 6(5):237-54, 1975.

Linsker R. Self-organization in a perceptual network. *Computer*. 21:105-117, 1988.

Martignon L., Deco G., Laskey K., Diamond M., Freiwald W., and Vaadia E. Neural coding: higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comp.* 12:2621-2653, 2000.

Mehta A.D., Ulbert I., and Schroeder C.E. Intermodal selective attention in monkeys. I: distribution and timing of effects across visual areas. *Cereb. Cortex*, 10:343-358, 2000.

Milner P. A model for visual shape recognition. *Psychol. Rev.,* 81: 521-535, 1974.

Milton J.G. and Mackey M.C. Neural ensemble coding and statistical periodicity: Speculations on the operation of the mind's eye. *J. Physiol.*, 94:489-503, 2000.

Moore G.P., Perkel D.H., and Segundo J.P. *Ann. Rev. Physiol.* 28:493, 1966.

Nicoll A. and Blakemore C. Patterns of local connectivity in the neocortex. *Neural Comp.* 5:665-680, 1993.

Nirenberg S., Carcieri S.M., Jacobs A.L., and Latham P.E. Retinal ganglion cells act largely as independent encoders. *Nature*, 411:698-701, 2001.

Nordhausen C.T., Maynard E.M., and Normann R.A. Single unit recording capabilities of a 100 micorelectrode array. *Brain Res.*, 726:129-140, 1996.

Ortega G. J., Bongard M., Louis E., and Fernandez E. Conditioned spikes: a simple and fast method to represent rates and temporal patterns in multielectrode recordings. *J. Neurosci. Methods.* 133:135-141, 2004.

Palanca B.J. and DeAngelis G.C. Does neuronal synchrony underlie visual feature grouping? *Neuron.* 46:333-346, 2005.

Palm G., Aertsen A.M.H.J., and Gerstein G.L. On the significance of correlations among neuronal spike trains. *Biol. Cybern.* 59:1-11, 1988.

Palm G. Evidence, information, and surprise. *Biol. Cybern.,* 42(1):57-68, 1981.

Palm G. On associative memory. *Biol. Cybern.* 36:19-31, 1980.

Panzeri S., Schultz S.R., Treves A., and Rolls E.T. Correlations and the encoding of information in the nervous system. *Proc. R. Soc. Lond.* 266:1001-1012, 1999.

Perkel, D.H., Gerstein, G.L., and Moore, G.P. Neuronal spike trains and stochastic point processes, II. Simultaneous spike trains. *Biophys. J.* 7:419-440, 1967.

Perkel D.H., Gerstein G.L., Smith M.S., and Tatton W.G. Nerve-impulse patterns: a quantitative display technique for three neurons. *Brain Res*. 100:271-296, 1975.

Petersen R.S., Panzeri S., and Diamond M.E. Population coding of stimulus location in rat somatosensory cortex. *Neuron.* 32:503-514, 2001.

Pettet M.W. and Gilbert C.D. Dynamic changes in receptive field size in primary visual cortex. *Proc. Nat. Acad. Sci.* USA, 89:8366-8370, 1992.

Pola G., Thiele A., Hoffman K.P., and Panzeri S. An exact method to quantify the information transmitted by different mechanisms of correlational coding. *Network: Comput. Neural Syst*. 14:35-60, 2003.

Pouget A., Dayan P., and Zemel R. Information processing with population codes. *Nature Rev. Neurosci.* 1:125-132, 2000.

Prut Y., Vaadia E., Bergman H., Haalman I., Slovin H., and Abeles M. Spatiotemporal structure of cortical activity: properties and behavioral relevance. *J. Neurophysiol.*, 79:2857-2874, 1998.

Purpura K.P., Victor J.D., and Katz E. Striate cortex extracts higher-order spatial correlations from visual textures. *Proc. Nat. Acad. Sci. USA,* 91:8482-8486, 1994.

Reich D.S., Mechler F., and Victor J.D. Temporal coding of contrast in primary visual cortex: When, what, and why. *J. Neurophysiol.*, 85:1039-1050, 2001.

Reinagel P. and Reid R.C. Temporal coding of visual information in the thalamus. *J. Neurosci.,* 20:5392-5400, 2000.

Rieke F., Bodnar D.A., and Bialek W. Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory afferents. *Proc. R. Soc. Lond. B. Biol. Sci.* 262:259-265, 1995.

Roelfsema P.R., Lamme V.A., and Spekreijse H. Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nat. Neurosci.* 7:982-991, 2004.

Rousche P.J. and Normann R.A. A method for pneumatically inserting an array of penetrating electrodes into cortical tissue. *Annals of Biomed. Eng.*, 20:413-422, 1992.

Rousche P.J. and Normann R.A. Chronic recording capability of the Utah Intracortical Electrode Array in cat sensory cortex. *J. Neurosci. Methods* 82:1-15, 1998.

Samonds J.M., Allison J.D., Brown H.A., and Bonds A.B. Cooperation between Area 17 neuron pairs enhances fine discrimination of orientation. *J. Neurosci.* 23:2416-2425, 2003.

Samonds J.M., Allison J.D., Brown H.A., and Bonds A.B. Cooperative synchronized assemblies enhance orientation discrimination. *Proc. Natl. Acad. Sci. U S A.* 101:6722-6727, 2004.

Samonds J.M. and Bonds A.B. Real-time visualization of neural synchrony for identifying coordinated cell assemblies. *J. Neurosci. Methods* 139:51-60, 2004b.

Samonds J.M. and Bonds A.B. Gamma oscillation maintains stimulus structure-dependent synchronization in cat visual cortex. *J. Neurophysiol.* 93:223-236, 2005.

Samonds J.M., Zhou Z., Bernard M.R., and Bonds A.B. Synchronous activity in cat visual cortex encodes collinear and cocircular contours. Submitted for publication, *J. Neurophysiol.,* 2005.

Sanger T.D. Optimal unsupervised learning in a single layer network. *Neural Networks.* 2:459-473, 1989.

Schmidt S., Horch K. and Normann R.A., Biocompatibility of silicon-based electrode arrays implanted in feline cortical tissue. *J. Biomed. Mater.* Res. 27:1393-1399, 1993.

Schwartz O. and Simoncelli E.P. Natural signal statistics and sensory gain control. *Nature Neuroscience*. 4:819-825, 2001.

Shadlen M.N. and Movshon, J.A. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron*, 24:67-77, 1999.

Shadlen M.N. and Newsome W.T. The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neurosci.*, 18:3870-3896, 1998.

Shadlen M.N. and Newsome W.T. Noise, neural codes and cortical organization. *Curr. Op. Neurobiol.*, 4:569-579, 1994.

Sherrington C.S. *Man on His Nature*. London: Cambridge University Press, 1941.

Sigman M., Cecchi G.A., Gilbert C.D., and Magnasco M.O. On a common circle: natural scenes and Gestalt rules. *Proc. Natl. Acad. Sci. USA*. 98:1935-1940, 2001.

Simoncelli E.P. Vision and the statistics of the visual environment. *Curr. Opin. Neurobiol.* 13:144-149, 2003.

Singer W. and Gray C.M. Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.*, 18:555-586, 1995.

Smyth D., Willmore B., Baker G.E., Thompson I.D., and Tolhurst D.J. The receptive-field organization of simple cells in primary visual cortex of ferrets under natural scene stimulation. *J. Neurosci.* 23:4746-4759, 2003

Softky W.R. and Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.*, 13: 334-350, 1993.

Strehler B.L. and Lestienne R. Evidence of precise time-coded symbols and memory of patterns in monkey cortical neuronal spike trains. *Proc. Nat. Acad. Sci*. USA, 83:9812-9816, 1986.

Tetko I.V. and Villa A.E.P. A pattern grouping algorithm for analysis of spatiotemporal patterns in neuronal spike trains. 1. Detection of repeated patterns. *J. Neurosci. Methods*, 105:1-14, 2001a.

Tetko I.V. and Villa A.E.P.A pattern grouping algorithm for analysis of spatiotemporal patterns in neuronal spike trains. 2. Application to simultaneous single unit recordings. *J. Neurosci. Methods*, 105:15-24, 2001b.

Thiele A. and Stoner G. Neural synchrony does not correlate with motion coherence in cortical area MT. *Nature*, 421:366-370, 2003.

Thomson A.M. and West D.C. Fluctuations in pyramid-pyramid excitatory postsynaptic potentials modified by presynaptic firing pattern and postsynaptic membrane potential using paired intracellular recordings in rat neocortex. *Neuroscience,* 54:329-46, 1993.

Usrey W.M. and Reid R.C. Synchronous activity in the visual system. *Ann. Rev. Physiol.*, 61:435-56, 1999.

Vaadia E., Haalman I., Abeles M., Bergman H., Prut Y., Slovin H., and Aertsen A. Dynamics of neuronal interactions in monkey cortex in relation to behavioral events. *Nature*, 373:515-518, 1995.

Victor J.D. How the brain uses time to represent and process visual information. *Brain Res.*, 886:33-46, 2000.

Victor J.D. and Purpura K.P. Nature and precision of temporal coding in visual cortex: A metric-space analysis. *J. Neurophysiol.*, 76:1310-1326, 1996.

Victor J.D. and Purpura K.P. Metric-space analysis of spike trains: theory, algorithms and application. *Network: Comput. Neural Syst.* 8:127-164, 1997.

Vinje W.E. and Gallant J.L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287:1273 –1276, 2000.

von der Malsburg C. The correlation theory of brain function. *Internal Report, Max-Planck Institute for Biophysical Chemistry.* Goettingen, Germany, 1981.

Warland D.K., Reinagel P., and Meister M. Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.*, 78:2336-2350, 1997.

Weliky M., Fiser J., Hunt R.H., and Wagner D.N. Coding of natural scenes in primary visual cortex. *Neuron.* 37:703-718, 2003.

Zetsche C. Sparse coding: The link between low level vision and associative memory. In *Parallel Processing in Neural Systems and Computers*, R. Eckmiller, G. Hartmann, and G. Hauske, eds. North-Holland, Amsterdam, 1990.

mulab.physiol.upenn.edu/jpst.html