

Representation and selection of time-varying signals by single cortical neurons

Jonathan Vincent Toups

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Physics and Astronomy.

Chapel Hill
2009

Approved by:

Paul Tiesinga

Rene Lopez

Charles Evans

Reyco Henning

Paul Manis

© 2009
Jonathan Vincent Toups
ALL RIGHTS RESERVED

Abstract

JONATHAN VINCENT TOUPS: Representation and selection of time-varying signals by single cortical neurons.
(Under the direction of Paul Tiesinga.)

I present a theoretical effort to develop tools and statistical analysis of neural responses to repeated presentations of identical time varying stimuli. Such experiments may produce responses characterized by regions of elevated firing separated by longer periods of low firing rate referred to as “events.” Unlike previous methods, which find events based on a firing rate threshold, I present a four parameter, reproducible method which first discovers “spike patterns” (subsets of trials with similar spike timing) using unsupervised clustering and a spike train metric, followed by the use of an interspike interval threshold to detect events. I present results from *in vitro* data showing that the precision of the resulting events is higher than that estimated using a firing rate threshold technique; events within a single spike pattern may be very precise, even if they overlap in time across spike patterns. This analysis provides a model of neural activity which preserves information about spike patterns which can be used to generalize single unit recordings to multi-unit activity. I also present a statistical test to characterize whether events are correlated with one another. A new insight provided is that the choice of time scale for the metric space analysis should maximize the information in the distances between spike trains, regardless of specific timescales in the data. The event finding procedure works well for data sets with more than 20 trials, and with events which are well separated within spike patterns.

I also present a comparison of two methods for selecting one of two stimuli present in the receptive field of a single cortical neuron. In the first, clustering of excitatory

and inhibitory synaptic input in the dendrites of a model layer 2/3 pyramidal cell is demonstrated to be insignificant in the selection of signals by unbalancing inhibition. In the second case, phase locking of excitatory inputs at two different phases of a local γ oscillation is demonstrated to produce statistically significant stimulus selection in output firing rate.

Together these results contribute to the analysis, representation and understanding of neural responses to time-varying signals.

This is dedicated to Jessica Polka. ...

Acknowledgments

I would like to thank Paul Tiesinga, my parents, Danna and Stanley Halpin and Kenneth and Susan Touns; Shelley Rogers, Jessica Polka, and all my friends and colleagues who supported me psychologically and academically while I wrote this thesis.

I would also like to thank Jean-Marc Fellous for performing the *in vitro* experiments which we analyze in Chapters 3 and 4. These experiments were conducted in the Sejnowski Lab at the Salk Institute in La Jolla California, and the Fellous Lab at the University of Arizona in Tucson.

Thanks also to Peter Thomas for assistance with the data analysis.

I would also like to thank the committee members for their time and attention.

This research was supported in part by the National Institutes of Health (R01-MH68481), start-up funds from the University of North Carolina at Chapel Hill and the Human Frontier Science Program.

Table of Contents

Abstract	iii
List of Figures	xi
List of Tables	xiv
1 Introduction	1
2 Neurobiological Background	3
2.1 General Background	3
2.1.1 Neurons and Synapses	3
2.1.2 Gross Anatomy	8
2.1.3 Origins of Feature Selectivity in the Ventral Visual Pathway . .	11
2.1.4 Receptive Field Size and Attention	18
2.1.5 Precision and Reliability of Neural Responses	20
2.1.6 Oscillations	22
2.2 Neural Modeling	22
2.2.1 Integrate and Fire Neurons	22
2.2.2 Hodgkin-Huxley type models	25
2.2.3 Compartmental Modeling	30
2.2.4 Modeling Synapses	35

3	Finding and Validating Spike Patterns	38
3.1	Abstract	38
3.2	Introduction	39
3.3	Experimental procedures	41
3.3.1	General experimental procedures	41
3.3.2	Stimulus generation and experimental design	42
3.3.3	General procedures	43
3.3.4	Calculation of the VP distance	44
3.3.5	Fuzzy clustering algorithm	44
3.3.6	Calculation of entropy and mutual information between classifications	45
3.4	Results: Finding spike patterns and determining the event structure in artificial data	47
3.4.1	Overview and goal of the event finding procedure	47
3.4.2	The interval method for identifying events	50
3.4.3	Selecting the temporal resolution parameter q	52
3.4.4	Selecting the number of patterns/clusters N_c	56
3.4.5	Merging events common across multiple patterns	61
3.4.6	Binary word representation of spike patterns	63
3.5	Finding spike patterns and determining the event structure in experimental data	70
3.6	Discussion	78
4	Spike Patterns Encode Amplitude or Offset Optimally Around Response Bifurcations	81
4.1	Abstract	81
4.2	Introduction	82

4.3	Methods and Experimental procedures	86
4.3.1	General analysis procedures	86
4.3.2	Calculation of the VP distance	88
4.3.3	Overview and goal of the event finding method	88
4.3.4	Calculation of entropy and mutual information between classifications	88
4.3.5	Simulation experiments	88
4.4	Results	90
4.5	Discussion	105
4.5.1	Implications for cortical coding	107
4.5.2	Limitations of the experimental study.	108
4.5.3	Analysis method	110
4.5.4	Future studies	111
5	Mechanisms for Stimulus Selection at the Single Pyramidal Cell	112
5.1	Abstract	112
5.2	Multicompartmental Model	113
5.2.1	Synapses	113
5.2.2	Calibrating the Spontaneous Background Activity	114
5.2.3	Generating Spike Times	115
5.2.4	Generating phase-locked spike times	116
5.2.5	Distribution of synapses in model	116
5.3	Stimulus Selection by localization of inhibitory and excitatory synaptic activity	117
5.4	Selection by inhibitory oscillation and excitatory phase-locking	138
6	Discussion and Conclusion	155

6.1	Overview of Results	155
6.2	Finding Events and Spike Patterns	157
6.2.1	Previous Studies of Events and Spike Patterns	157
6.2.2	Implications of the Current Study for Metric Space Analysis . .	159
6.2.3	Reproducibility, Heuristics and Limitations of Presented Method	160
6.3	Coding Implications of Spike Patterns	163
6.4	Future Spike Patterns Research	165
6.5	Stimulus Selection at the Single Neuron Level	166
6.5.1	Stimulus Selection with Synaptic Clustering	167
6.5.2	Stimulus Selection with Inhibitory Oscillations	170
6.6	Conclusions	173

Bibliography		175
---------------------	--	------------

List of Figures

2.1	A Schematic Neuron	4
2.2	A Schematic Ion Channel.	5
2.3	A Schematic Synapse	6
2.4	Basic Anatomical Terms	9
2.5	The Dorsal and Ventral Visual Pathways.	10
2.6	Gross Anatomy of the Brain.	11
2.7	Retinotopy	13
2.8	Feed-forward Orientation Selectivity.	15
2.9	The Somer’s Model of Orientation Selectivity.	17
2.10	Neurons Can Be Precise.	21
2.11	The Power Spectrum of the Brain.	23
2.12	The Leaky Integrate and Fire Neuron.	24
2.13	The Hodgkin-Huxley Model.	28
2.14	Compartmental Modeling.	31
3.1	Precision and Reliability can be Hidden in Neural Responses.	40
3.2	Comparison of Hand Detected Events and Automatic Detection.	49
3.3	The Interspike Interval Method for Finding Events.	50
3.4	Selecting the VP-metric Parameter q	55
3.5	Evaluating the q Heuristic.	57
3.6	Testing the q and N_c Heuristics On Artificial Data.	60
3.7	Merging Events with an ROC Analysis.	64
3.8	Merging Events with an ROC Analysis (continued)	65
3.9	The Binary Word Representation of Spike Trains.	68

3.10	Evaluating the Significance of Spike Patterns.	69
3.11	In Vivo Responses to Time Varying Signals.	71
3.12	Procedure for Finding Events.	73
3.13	Procedure for Finding Events (continued).	74
3.14	Spike Patterns Persist over Multiple Amplitudes.	77
4.1	Differences between Offset and Amplitude Modulations.	84
4.2	Reliability and Spike Patterns.	94
4.3	Spike Patterns Correspond to Voltage Patterns	96
4.4	Spike Patterns Correspond to Voltage Patterns (continued).	97
4.5	Bifurcations Occur in a Model Neuron.	99
4.6	Bifurcation Points Encode Information Robustly.	102
4.7	Bifurcation Points Encode Information Robustly (continued).	103
5.1	Geometry of the Traub Model.	118
5.2	Typical Distribution Of Synapses.	119
5.3	Three Synaptic Placement Strategies.	121
5.4	Stimulus Selection By Location: Concept.	122
5.5	Inhibition is Not Localizable in the Dendrites.	123
5.6	Contralateral Versus Colateral Inhibition.	126
5.7	CON Inhibition Produces Better Correlations.	129
5.8	CON Inhibition Produces Slightly Better Selection.	130
5.9	Stimulus Selection by Unbalancing Inhibition.	132
5.10	Firing Rate Versus Mean And Standard Deviation.	134
5.11	CON Inhibition Lowers Somatic Standard Deviation.	135
5.12	CON Inhibition Lowers the Firing Rate via the Standard Deviation of the Membrane.	137

5.13 Stimulus Selection by Oscillation and Locking.	139
5.14 Gain versus Incoming Phase.	140
5.15 Phase-Locked Spike Can Still Encode A Signal.	143
5.16 Selection At A 250 ms Time Scale.	146
5.17 Selection At A 100 ms Time Scale.	147
5.18 Selection At A 25 ms Time Scale.	148
5.19 Selection At A 10 ms Time Scale.	149
5.20 Selection At A 10 ms Time Scale, Filtered.	150
5.21 Selection Degrades for Fast Signals, but Persists for Slow Time-Scales. .	153

List of Tables

3.1	Experimental Information:	78
5.1	Synaptic Parameters	113

Chapter 1

Introduction

The brain is constantly presented with a world in motion. Not only must an organism make sense of this vast tumult of information, it must do so dynamically, selecting aspects of the sensory universe to devote resources to depending on its current state and the demands of the world. Neuroscience has made great strides in understanding the basic principles of neural computation which allow organisms to take input, from the retina, for instance, and extract meaningful aspects of visual space from that raw data, such as orientation, shape, and other features. However, most insights into neural computation have been derived from experiments where the stimuli have simple temporal and spatial characteristics. Realistically, brains must be able to process and attend to signals with rich temporal structure, but responses to this kind of stimuli are difficult to study.

This dissertation presents several results structured around the theme of investigating neural responses to time-varying signals. Chapter 2 presents neurobiological background material which provides context for the subsequent investigations. Basic anatomical terms, aspects of neural modeling, and the origins of stimulus selectivity and attention are discussed. Chapter 3 presents a method for investigating the kinds of responses neurons may make to rich, time varying signals. A mostly-unsupervised

method to find so-called “spike patterns,” interesting features of neural responses to repeated presentations of signals with temporal structure, is described and applied to *in vitro* data. Chapter 4 applies these methods to the question of simultaneously encoding information about the amplitude and structure of a time-varying signal. We show that the amplitude of a signal can be encoded in the distribution of spike patterns, and, as a consequence of this fact, regions in the “response space” which have the highest density of possible patterns (termed “bifurcation points”) can encode amplitude with the highest fidelity.

Chapter 5 describes two methods for dynamically selecting one of multiple inputs into a neuron. In the first case, the question of whether the clustering of synapses into specific sub-domains of the dendrites of pyramidal cells might allow for a possible means of stimulus selection is investigated. By adjusting inhibitory and excitatory activity to specific locations in the dendrites, we attempt to dynamically select which excitatory signal is represented in the neurons output. We find that in our model, the location of synapses is not significantly important in the time-varying case, although there are significant differences between different location strategies in the constant firing rate regime which indicate further areas of research. In the second case, the relationship between inhibitory oscillations, which are a ubiquitous feature in the brain, and stimulus selection. Signals are again represented as excitatory synaptic activity in the model, but each signal’s activity is tightly locked to a particular phase with respect to a reference, inhibitory oscillation. We report that by adjusting the phases of incoming signals, the output of the neuron can be strongly modulated to reflect, or neglect, the temporal structure of those signals.

Finally, Chapter 6 discusses the significance of these results in terms of current and future research.

Chapter 2

Neurobiological Background

2.1 General Background

2.1.1 Neurons and Synapses

The human brain is made of around 160 billion cells, divided roughly in half between glial cells and “proper” neurons, with glial cells slightly outnumbering neurons (Johnston and Wu, 1995). Glial cells are usually thought of as taking a “support” role in neural computations (for instance, maintaining the correct biochemistry in the intercellular medium). Since they are not electrically active on short time scales, they are usually neglected when considering problems of neural computation (Fields and Stevens-Graham, 2002). The majority of the rest of the cells in the brain are neurons ($\approx 10^{12}$). Neurons are electrically active cells with a distinctive morphology apparently adapted to the task of processing and transmitting information. Although neurons come in a large variety of shapes, sizes, morphologies and electrophysiologies (Markram et al., 2004; Larkum et al., 1999; Migliore et al., 1995; Golding et al., 1999; Schaefer et al., 2007), they nevertheless share enough in common that a single description can convey shared features. Most, but not all, neurons share the features illustrated in Figure 2.1: a dendritic arbor (usually an extensively branched structure made of cell membrane

attached to the cell body), a soma (the part of the cell most like other cells in the body) and an axon (usually a long, thin extrusion of cell membrane which may branch many times before forming synapses onto other cells). From the simplest perspective, information can be visualized as flowing from the dendrites, towards the soma, and down the axon, as indicated by the arrow in Figure 2.1.




Figure 2.1: A Schematic Neuron. Neurons are electrically active cells that integrate input from other neurons with their dendrites, generate spikes, depending on the input, at the soma, and transmit those spikes to other neurons (or muscle tissue) along their axons. They are the basic computational unit of the brain.

Neurons are electrically active cells. Their membranes contain specialized molecular pores, pumps, and channels (see Figure 2.2) which allow the neuron's transmembrane voltage (also called the membrane potential or voltage) to rapidly change in response to signals from other neurons, and for those rapid changes to propagate from one location in the cell to another. The physics of these transmembrane currents are discussed at greater length in Section 2.2. For the moment, it is sufficient to observe that such transmembrane currents allow the neuron to send signals, called action potentials, from their somas to their axons, where they are transmitted to other neurons.

Neurons connect to one another at electro-chemical junctions called synapses (Johnston and Wu, 1995). Synapses can be either entirely electrical in nature (called gap




Figure 2.2: Ionic channels in the membrane of the neuron selectively pass certain ions in a voltage dependent way. The cell membrane is illustrated with a single ion channel. Conceptually, a charged gating protein acts as a molecular switch, closing or opening the channel as it is pulled by the trans-membrane electric field. Modeling the average behavior of many such channels underlies conductance based modeling.

junctions) or chemical in nature (usually referred to as simply “synapses” unless there is a special need to distinguish them from gap junctions - when there is, the prefix “chemical” is usually added). Gap junctions are places where two cells literally make electrical contact, allowing ions to flow directly between the two cells. At chemical synapses, the “signal” propagating between the two cells is mediated chemically. Although both gap-junctions and chemical synapses are important in the brain, chemical synapses are more relevant to the topics discussed in this dissertation because so-called feed forward and top down inputs are mediated by them. The biophysics of synapses and the simulation thereof are discussed in Subsection 2.2.4.

Like neurons themselves, chemical synapses (hereafter merely synapses) come in a large variety of types, but can be described by a single set of features. Figure 2.3 illustrates a schematic, generic synapse. The presynaptic cell approaches very near the post-synaptic cell, creating a synaptic cleft around 20 nm in width. The presynaptic terminal contains, near the cell wall, vesicles of neurotransmitter. The arrival of an

action potential causes these to merge with the cell membrane and dump their contents into the synaptic cleft. Receptors in the membrane of the post-synaptic cell detect the presence of these neurotransmitters and directly or indirectly cause conductance changes in the post-synaptic neuron, which in turn causes currents to flow, changing the voltage across the membrane of the post-synaptic neuron.




Figure 2.3: Synapses convert the electrical action potential into a chemical signal and back again. The arrival of the action potential at the synaptic junction causes the release of neurotransmitters into the synaptic cleft. These bind to receptors in the postsynaptic neuron's membrane, causing changes in conductance which can depolarize (increase) or hyperpolarize (decrease) the membrane, providing information about the presynaptic neuron to the dendrites of the postsynaptic cell.

Neurons are themselves broadly, if imperfectly, divisible into excitatory neurons and inhibitory types. These names refer to the effect the synapses formed by a presynaptic neuron have on postsynaptic cells when spikes arrive. As it happens, all the synapses a given cell makes onto postsynaptic neurons are either excitatory or inhibitory (this is known as Dale's Principle (Eccles, 1976)). No neurons exist which make synapses of both types on postsynaptic neurons. Excitatory synapses generally cause the membrane potential to increase (depolarize) in the postsynaptic neuron (they are also sometimes

called “asymmetric” because of their morphologies). Inhibitory synapses (sometimes called “symmetric”) generally cause a postsynaptic cell’s membrane potential to decrease (hyperpolarize) (DeFelipe and Farinas, 1992). Within these two broad categories, synapses can fall into a variety of other groups, depending on the neurotransmitter used and the morphology of the junction. In brain areas where high fidelity and accurate timing are needed (such in the auditory system), the postsynaptic neuron’s synaptic bouton is often enclosed almost entirely in an extrusion of the presynaptic cell (Satzler et al., 2002; Joris et al., 2004). In this dissertation, simple model synapses are used and the details of neurotransmitters and synaptic morphology are abstracted.

The dichotomy between excitatory and inhibitory neuron types is strongly associated with another important distinction between neurons: that of interneuron versus non-interneuron cell types (described below). The brain is not one undifferentiated lump of neurons - it is divided into functional sub-units at different size scales. The brain at large is divided between cortical and sub-cortical areas, each of which is divided into areas which perform a common function, and are differentiated by different cell types and distributions of cell types and connections. A repeated organizational principle in complex nervous system is that of layers and mini-columns (Mountcastle, 1997). The mammalian cortex is usually conceptualized as having six layers. The generally accepted view of cortical computation has “columns” of neurons in a given cortical area interacting with both excitatory and inhibitory connections to perform a computation, the output of which is then passed via excitatory connections to another cortical area for further analysis.

Interneurons are neuron types which are limited in their connectivity to a single cortical area. They are usually inhibitory and smooth, although exceptions to both characterizations exist. Connections between cortical areas are generally made by varieties of pyramidal cells, so-called because of the morphology of their cell body. Pyramidal

cells are always excitatory and make both local and non-local connections.

There are estimated to be at around 10,000 “classes” of neurons in vertebrates (Johnston and Wu, 1995), although there is no universally accepted method for categorizing them. Despite this bewildering variability in form and function, some basic principles of neural computation have been described. After describing some basic anatomy of the nervous system, some of these principles will be reviewed (see Subsection 2.1.3).

2.1.2 Gross Anatomy

For readers without a background in biology, it is useful to establish a few common anatomical terms with respect to describing brain anatomy. Figure 2.4 illustrates these basic terms. Figure 2.4A shows the head and brain in profile. Because the head in humans is almost perpendicular to the line of the back (in contrast to quadrupeds), some anatomical terms can be slightly misleading. Dorsal means along the top of the head, while ventral indicates the opposite. Anterior means towards the nose, while posterior indicates towards the back of the head. Occasionally the synonyms rostral and caudal are used, respectively, to describe the same directions. Conveniently, left and right retain their meaning when discussing anatomy, as illustrated in the top view of the brain shown in Figure 2.4B. When referring to neurons, additional terms are useful. Proximal means close to the cell body while distal means the opposite. To describe locations along the axis of the neuron, the word basal indicates the side of the neuron from which the axon emerges, while apical indicates the other direction. Although the neuron illustrated in Figure 2.4C does not have basal dendrites, it is nevertheless a common feature in pyramidal cells, and so the location of the axon is the way to orient oneself to a neuron.

This dissertation is principally concerned neurons in the visual cortex, as attention in

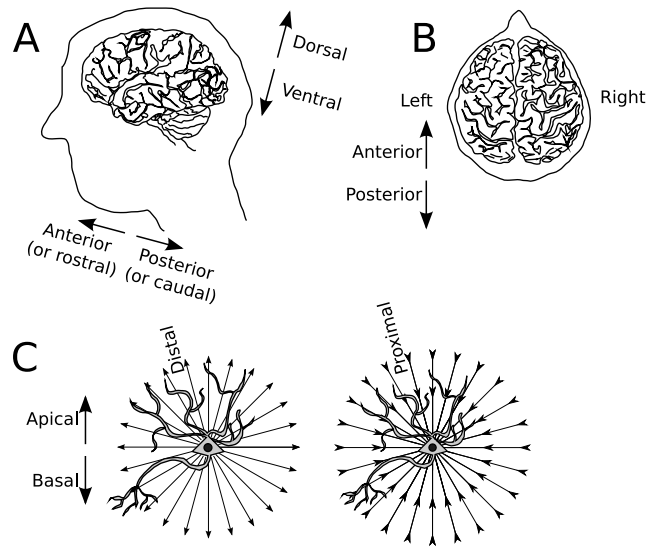


Figure 2.4: Basic anatomical terms for describing locations in the brain. (A) Dorsal refers to the top of the brain, ventral to the bottom. Anterior or rostral indicates the direction towards the nose, while posterior or caudal indicates the back of the brain. (B) The left and right side of the brain are referred to simply as “left” and “right.” (C) When referring to neurons, distal means far from the cell body, proximal means close to it and basal refers to the end of the neuron from which the axon emerges, while apical indicates the opposite end of the neuron.

the visual system is well studied due to the location of the visual cortex and the nature of visual responses there (Koch, 2004b). The occipital lobe (labeled in Figure 2.6) contains the cortical areas responsible for the processing of low-level visual information. Moving towards the front of the brain from the occipital cortex, the visual system splits into ventral and dorsal pathways (Figure 2.5). The ventral pathway is generally associated with object recognition, while the dorsal pathway is associated with motion planning and action (Koch, 2004b).

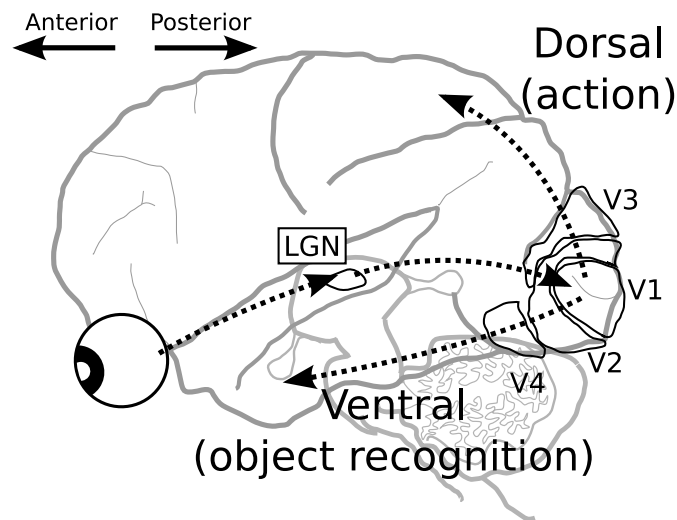


Figure 2.5: Extraction of non-trivial features from visual input begins in the occipital lobe's visual cortex. From there, visual processing splits into two pathways, the ventral, running along the bottom of the brain, and the dorsal, running along the top. Each pathway is hypothesized to be responsible for different aspects of the perception and utilization of visual information (Koch, 2004b). The ventral stream is involved in object recognition - as we move from V1 to V2 to V4, neurons become responsive to more and more complex objects. Further into the ventral pathway, neurons which respond to faces and complex objects are found, for instance, in area IT (Mesulam, 1998). The dorsal pathway is involved in motion planning, and neurons in this pathway respond to location and motion of objects in the visual world. (This brain's folds have been simplified for visual clarity).

The ventral, object recognition pathway is better understood than many brain areas because it is relatively easy to access with electrodes. It has also proved easy to find simple stimuli which selectively activate the neurons there. As a consequence, the visual




Figure 2.6: Gross anatomical divisions of the brain. Although each lobe is too large to associate with a single function, each is traditionally ascribed a general operational area. The occipital lobe is associated with vision, the parietal lobe with some kinds of abstract thought, the frontal lobe with planning, and the temporal lobe with music. The cerebellum is associated with muscle control and balance, and the brainstem connects the brain to the rest of the nervous system and controls some reflexes (reviewed in Diamond et al. (1985), Chapter 1).

system forms the conceptual background to this study.

2.1.3 Origins of Feature Selectivity in the Ventral Visual Pathway

Some neurons can be thought of as having a stimulus preference. For instance, while recording from a neuron in V1, one may find that it responds only when a bar of a particular orientation is placed in a particular location in space, relative to the fixation point of the eyes. This behavior was first reported by Hubel and Wiesel (1962) in the visual cortex of cats. The orientation and spatial location to which the neuron responds is sometimes called its stimulus preference, and we may say that the neuron is “orientation selective”. The spatial component of a visual system neuron’s stimulus preference is often referred to as its receptive field, and its stimulus preference is often described neglecting the spatial location of the receptive field. A neuron might be said, for instance, to have a stimulus preference for bars oriented at 45° without specifying at which position in space the oriented bar needs to be. Some authors include all aspects of the stimulus which have an effect on the response of the neuron in the term

“receptive field,” but for simplicity and clarity, in this dissertation, receptive field will mean the spatial location and extent in which an object must be to produce a response. The stimulus preference will mean the properties that produce a response when present in the receptive field.

With this information, the progress of information through the early visual system and into the ventral visual pathway can be described. A complete, if casual, introduction to the visual system can be found in Koch (2004b) Chapters 3-4. Since the intent here is to produce an understanding of the origins of stimulus selectivity in the visual cortex, a simpler introduction which highlights the most relevant aspects of the visual system will be more appropriate.

Visual information enters the brain at the retina, which is a 2D surface of photoreceptive neurons and support cells located on the back, inside surface of the eyeball. The density of neurons is not constant on the retina (most notably the so-called *fovea* at the center of the retina has a much higher density of photoreceptive cells than other parts of the retina). A crucial feature of the visual system, however, is apparent in the retina: nearby neurons respond to nearby portions of the visual space. This *retinotopic* organization is preserved throughout the visual system, so that two neurons which are near one another physically are likely to have receptive field centers which are also nearby. A gross but functional simplification is that each retinal neuron responds to the presence of either light or darkness in its receptive field and sends this information to the rest of the visual system (Figure 2.7).

The retinal neurons encode information about the visual scene in spike trains (sequences of spikes), which travel down axons to the lateral geniculate nucleus (LGN) (Figure 2.5), located in the thalamus, situated roughly “on top” of the brainstem, beneath the cortex (Figure 2.6). Although in practice the LGN’s role is not completely understood, for our purposes it is sufficient to note that cells projecting from the LGN

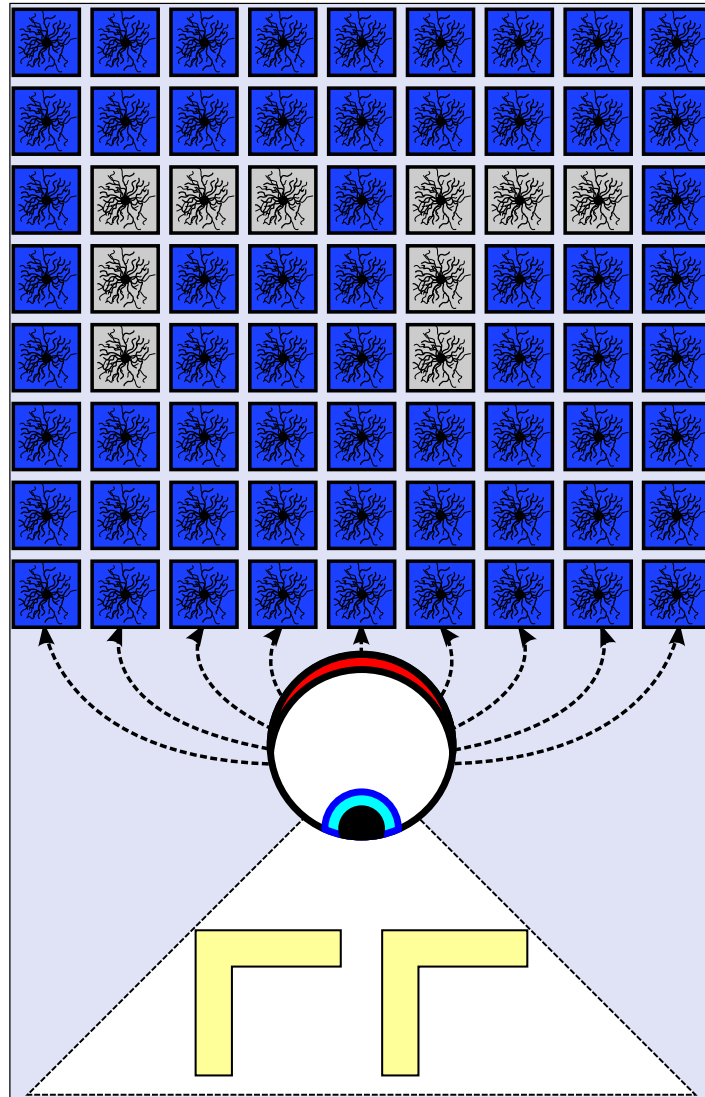


Figure 2.7: The retina reproduces spatial relationships in its pattern of activity. In this schematic representation of a visual scene and the retina, two shapes are reproduced in the response of the retinal cells (indicated by color). Nearby areas of space occupied by the two stimuli are represented by neighboring neurons on the retina. A real retina has large variations in the density of retinal neurons and reflects an inverted representation of space, but the relationship between neighboring points in space and neighboring neurons is retained.

to V1, the next stop in the visual system, have similar receptive fields and stimulus preferences as retinal neurons (Koch, 2004b). That is, LGN cells respond to the presence of light or dark spots in their receptive field (Koch, 2004b, Chapter 3). The retinotopic map is preserved.

Neurons in V1 begin to show the first signs of “interesting” stimulus preferences in that they respond to the presence of specific orientations in their receptive fields. The origins of orientation selectivity are not fully understood in V1, however, several models provide sufficient explanatory power to shed substantial light on the mechanisms which are at work (Teich and Qian, 2006). Models fall on an axis defined by the degree to which feed-forward and recurrent connections are used to create orientation selectivity. Hubel and Wiesel (1962) originally proposed an entirely feed-forward model to explain the origins of orientation selectivity. In such a model, a neuron which has a particular orientation preference receives inputs from a line of neurons along that orientation (Figure 2.8). When neurons have identical receptive field centers but different orientations, there is some overlap in their inputs. For instance, if the blue LGN neurons in Figure 2.8 are activated, then the red V1 neuron receives a small amount of input, despite the fact that it is selective for the opposite orientation. A large threshold for activation ensures that the red neuron remains inactive under such stimulation, and only responds when a large number of its presynaptic neurons are activated, as they would be if a visual stimulus triggered all the red LGN cells. While all models of orientation selectivity presume some degree of selectivity in feed-forward connections (and in fact, there is evidence for this connectivity from anatomical studies (Chapman et al., 1991; Tanaka, 1983; Reid and Alonso, 1995)), pure feed-forward models appear to be untenable due to their inability to reproduce the degree and manner of contrast invariance that experiments have demonstrated (a review of the results in the cat is Ferster and Miller (2000)).

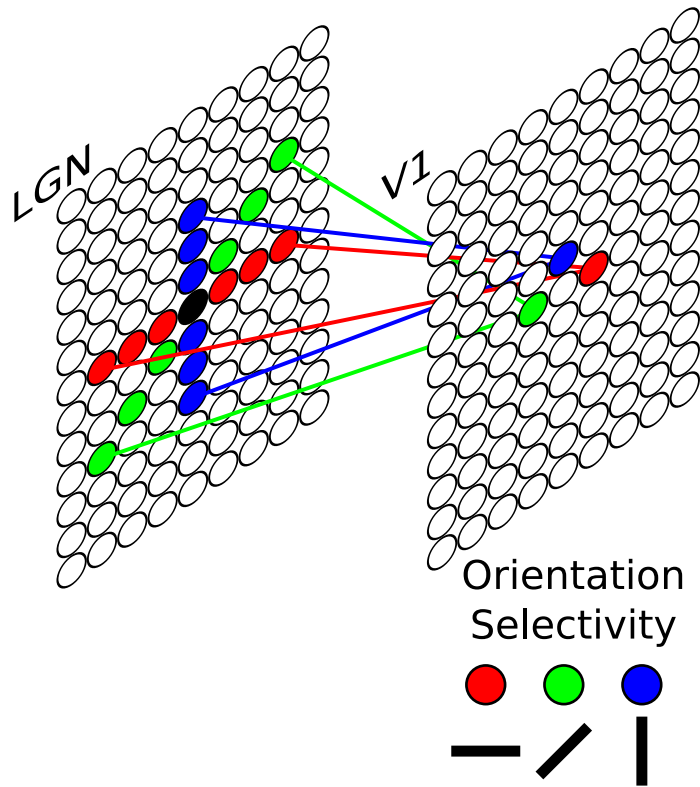


Figure 2.8: Orientation selectivity in V1 can be produced by feed-forward connections from lines of neurons in the LGN. A neuron in V1 which is sensitive to horizontal lines receives input from the LGN neurons which lie on that line, while other orientation selective neurons receive inputs from other lines. In this schematic representation, three neurons receive input from three overlapping sets of LGN neurons. Since the inputs for two distinct orientations with the same receptive field center must overlap, the threshold for firing in a pure feed-forward model must be set so that a neuron only responds when most of its presynaptic cells are activated (Hubel and Wiesel, 1962).

In order to achieve more realistic responses to contrast variation in orientation selectivity, inhibitory connections in V1 have to be added to the model. Exactly how the inhibitory connections are made and what the inputs to the local inhibitory cells look like can vary from model to model (see Teich and Qian (2006) for a review of models of orientation selectivity), but the key feature is that orientation selectivity derives first from feed-forward connections, and then it is refined by competitive interactions mediated by inhibitory interneurons.

The Somers model (illustrated in Figure 2.9) is one such model. Feed-forward connections from the LGN provide the initial selectivity for a neuron in the cortex because they all come from an oriented line. Similar lines in space will also activate this neuron, except that inhibition (which is more broadly tuned) prevents activation unless a stimulus with the right orientation is present.

Stimulus selectivity in deeper areas of the primate ventral visual pathway can be conceptualized by analogy to orientation selectivity. For instance, many V4 neurons respond to specific angles or curvatures in their receptive fields (Pasupathy and Connor, 1999). The inputs to such a neuron might consist of the outputs of orientation selective neurons in V2 whose spatial relationship and orientation selectivity produce the appropriate relative angle or curvature. Local inhibition may then be involved in sharpening or refining the V4 neuron's stimulus preference.

In summary, a neuron's stimulus selectivity is produced by selective feed-forward inputs representing features extracted by the previous cortical area. The selectivity produced by these connections is then refined by local inhibition. In Chapter 5, we examine possible ways that such feed-forward inputs into a neuron might be modulated to bias the response of the cell in favor of one component. Such modulations might be a part of the mechanism by which brains devote attention to specific inputs while ignoring others.




Figure 2.9: Schematic connectivity of the Somers model (adapted from Somers et al. (1995)). In this model, feed-forward excitatory activity from the LGN stimulates both pyramidal cells and inhibitory interneurons in the cortex. Input from a line of LGN cells produces weak orientation selectivity in the cortical pyramidal cells, while inhibition is activated broadly, refining selectivity by inhibiting less strongly activated pyramidal cells with similar, but not identical, orientation selectivity.

2.1.4 Receptive Field Size and Attention

One problem with this view is that receptive field sizes grow quickly as we move deeper into the ventral visual pathway, starting at .3 degrees or less in the LGN and V1 and increasing to 4 degrees by V4. Deep into the ventral pathway, in the inferior temporal cortex (IT), receptive fields can be as large as 25° (Desimone and Duncan, 1995).

Large receptive fields pose a problem, however, in that they may be occupied by multiple stimuli at once. Under many circumstances, an organism is interested in only a small portion of the visual space or objects with specific features, and wants to ignore other inputs. How can a neuron with a large receptive field simultaneously respond to the presence of a stimulus in an attended location, while ignoring another stimulus at an unattended location within its receptive field? This dissertation examines possible mechanisms of spatial attention at the scale of a single neuron in V4.

Although we have discussed neural responses in terms of firing rates up to this point, the brain is replete with oscillatory activity as well (Buzsaki, 2006). Any discussion of attention should address the relationship between neural activity and ongoing or stimulus induced oscillatory activity.

Studies of visual attention have demonstrated the brain's ability to dynamically effect the responses of neurons to the presence of stimuli in their receptive fields (Fries et al., 2001, 2008; Taylor et al., 2005). In Fries et al. (2008), macaque monkeys were trained to respond to a color change in a target stimulus by releasing a lever. A similar distracter stimulus would also undergo a color change at some point during the trial. Significant differences in both the firing rate and γ -band synchrony were observed during these trials depending on whether attention was directed into or out of the recorded neuron's receptive field.

Chapter 5 of this dissertation examines a single pyramidal cell receiving multiple inputs representing different spatial locations. The possible roles of both inhibitory γ

oscillations and location specificity of feed-forward and inhibitory synapses in spatial attention are examined.

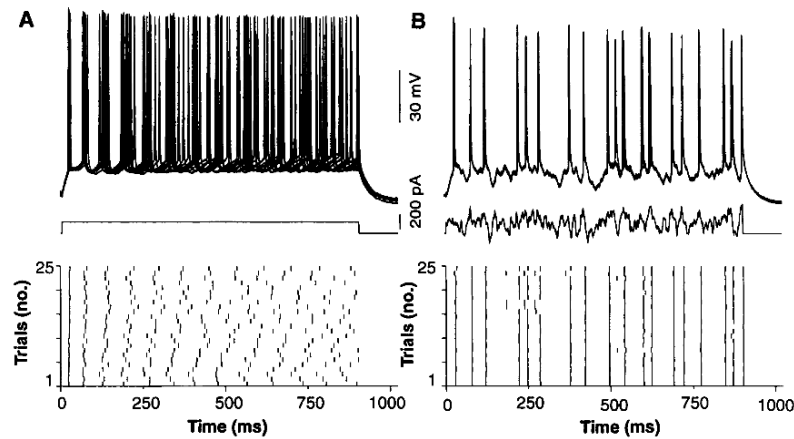
2.1.5 Precision and Reliability of Neural Responses

Up to this point, the activity of the brain has been described in terms of the firing rates of neurons. Stimulus preference, for instance, is measured by showing a neuron a variety of stimuli and measuring its response over an interval of time to each, calculating the total firing rate, and finding the stimulus which produces the most spikes. In nature, it is uncommon for an animal to maintain a visual fixation at some point in space, and then see an isolated white bar which appears in a field of black, maintains its orientation for 500-1000 ms, and vanishes again. During normal behavior humans and other primates almost never fixate for longer than 200 ms - their eyes are instead constantly making rapid movements called saccades, causing the image falling on any particular area of the retina to rapidly vary (Koch, 2004b).

Responses to constant stimulation like those used in many types of experiments often have high variability in both the exact timing of spikes and their total number. This has lead neuroscientists to hypothesize that the cortex is a fundamentally noisy place where neurons are neither precise nor reliable in their spiking and where large numbers of neurons must be averaged over, in some sense, to extract meaningful information about the sensory universe (Shadlen and Newsome, 1998). Until recently, it has been impossible to record from many neurons simultaneously, and so multiple recordings of one neuron responding to the same repeated stimulus presentation were used as a surrogate.

Contrary to this view, some experiments have shown that when stimulated with time varying signals, cortical neurons can respond with high precision and reliability, both *in vivo* and *in vitro* (Mainen and Sejnowski, 1995, 1996; Reinagel and Reid, 2002). In Figure 2.10, the responses of a pyramidal cell *in vitro* are plotted in response to a constant depolarizing current (A) and a “frozen noise” stimulus (a temporally varying

signal with exactly the same structure on each trial) (B). In response to a constant depolarization (A), spikes rapidly become unsynchronized (within about 500 ms), while the responses to the “frozen noise” (B) stay synchronized for the duration of the stimulus.



Mainen and Sejnowski (1995)

Figure 2.10: Neurons respond with high precision to time varying signals (Mainen and Sejnowski, 1995). (A) When stimulated with a constant depolarizing current, this cortical pyramidal cell produces highly variable responses, in terms of spike timing. (B) In contrast, the response of the same neuron to a temporally varying signal is precise and reliable in both spike timing and firing rate.

The high precision of neurons in response to time varying stimuli opens up the possibility that information might be encoded in either the total firing rate of a neuron or its exact spike times.

Neurons firing with high precision can be demonstrated to produce a small number of distinct patterns of activity in response to repeated stimulation by the same signal under certain circumstances (Tiesinga and Toups, 2005). These patterns have been observed in experiment (Fellous et al., 2004), but a full development of mechanisms to detect and characterize them has not been presented. This dissertation examines some ways of accomplishing this task.

Chapters 3 and 4 examine the implications of precise spike timing for coding both the

temporal variations in a stimulus and contrast, and discuss algorithms to automatically extract meaningful patterns from precise data. A statistical measure of the relevance of spike patterns to the neural code is also examined.

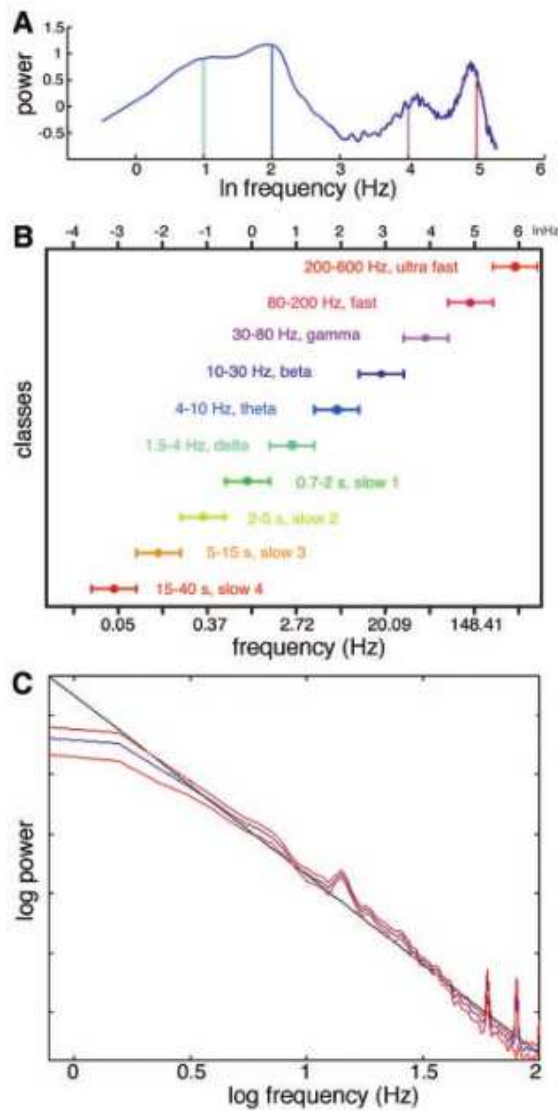
2.1.6 Oscillations

Oscillations are a persistent and ubiquitous aspect of the nervous system (Buzsaki and Draguhn, 2004). Figure 2.11, from Buzsaki and Draguhn (2004) shows the presence of oscillations at various distinct frequencies in mouse (A), rat (B) and human (C) subjects. The hypothetical role of the γ oscillation in particular is addressed in this dissertation (see Section 5.4).

2.2 Neural Modeling

2.2.1 Integrate and Fire Neurons

Among the simplest methods for modeling the behavior of neurons is the “leaky integrate and fire” (LIF) model (Figure 2.12). The model reproduces the intuition that a neuron is accumulating inputs from its dendrites and summing them up until a threshold is reached, upon which it emits a spike, resets, and the process starts anew. The integrative part of the neuron is simulated as a capacitance, which accumulates charge (and therefore the voltage across the capacitance increases). The “leaky” aspect of the model refers to the fact that there is a transmembrane leakage, R_L which slowly drains the charge off the “membrane.” This reflects the observation that old inputs are less important than new ones (or, equivalently, that the neuron tends to return to a “resting” state after a period of time) (Koch, 2004a, Chapter 1), so that the same amount of input over a long time may not produce a spike, even though it can produce one if it is received within a short time period. In a real neuron, an “action potential,” or spike,



Buzsaki and Draguhn (2004)

Figure 2.11: The Power spectrum of the brain (Buzsaki and Draguhn, 2004). (A) The power spectrum from a sleeping mouse shows peaks at various distinct oscillation frequencies. (B) Different oscillation regimes and their names in the rat. (C) The power spectrum for a sleeping human (taken from an electroencephalographical (EEG) recording) shows similar peaks in frequency corresponding to different oscillations. The large difference between A and C is due to a different scale and because power spectrum in A was whitened by having the log-trend subtracted.

is a non-linear process due to the coordinated action of multiple active conductances across the neural membrane. It manifests itself as a rapid rise and fall in the membrane potential which results in the membrane returning to near the resting membrane potential (or below it). In the LIF neuron this is represented by an ad-hoc “spiking mechanism” which emits a spike and resets the membrane potential after it passes a set threshold. The LIF neuron is an RC circuit with time constant $\tau = R_m C_m$ and is often simulated with a voltage range of $[0 \ 1]$, (corresponding to removing the battery in Figure 2.12, making the resting membrane potential zero, and then adjusting $R_m C_m$ to give the desired time constant).

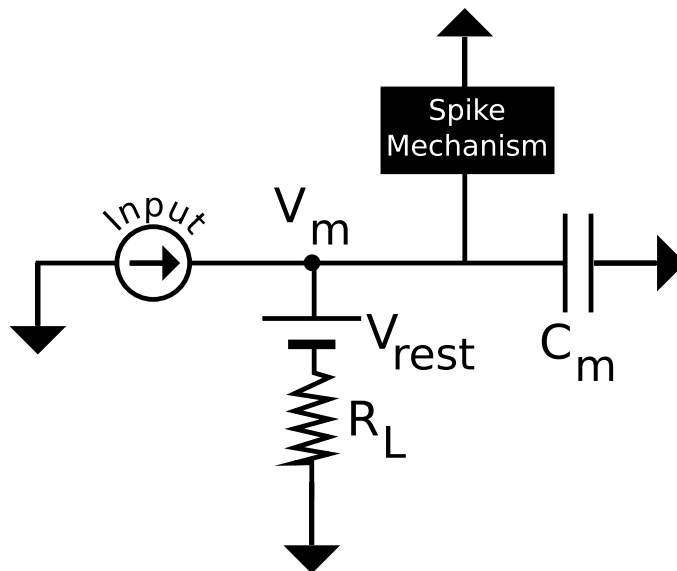


Figure 2.12: The leaky integrate-and-fire (LIF) neuron. This abstract model of a neuron contains only a membrane resistance and capacitance, an abstract representation of input as a current source, and an ad-hoc “spiking mechanism” which issues a “spike” and then resets the membrane potential to zero.

When the input is conceptualized as coming in bursts from multiple sources, in analogy with synaptic input in a real neuron, the LIF neuron can be conceptualized as a coincidence detector if it has a short time constant: many inputs must arrive in a short interval to trigger a spike. LIF neurons with long time constants are more like

accumulators - their responses reflect the presence of a certain level of activity over a long period (see Section 3.4).

In real neurons a spike is often followed by a period during which the neuron is unable to fire again - a refractory period. LIF neurons are sometimes given an artificial refractory period by simply preventing the voltage from changing for a length of time after a spike or by resetting the membrane potential to a value lower than the resting potential, making a spike immediately after firing relatively unlikely.

2.2.2 Hodgkin-Huxley type models

The Hodgkin-Huxley (HH) approach for modeling the behavior of a neuron's membrane was originally applied to the experimentally tractable giant squid axon (Hodgkin et al., 1952). Like the simple LIF model, the HH method begins with the conceptualization of a neuron as a capacitive membrane through which charge flows. Rather than modeling only input and leakage currents, and simulating the spike generating mechanism as an instantaneous reset of the membrane potential, HH, or conductance based models, as they are also called, attempt to characterize the voltage gated ion channels in a cell's membrane and simulate their behavior (averaged over the many channels of a given type in the membrane) as a function of time. No ad-hoc spike generating mechanism is necessary - a conductance based model with the proper conductances will spike by virtue of its voltage dependent conductances.

We first concern ourselves with modeling a single small patch of membrane. A neuron is electronically compact when the membrane potential does not vary much across the surface of the neuron. Under these circumstances, a "single compartment" model (in which only a patch of the membrane is simulated) may be sufficiently accurate. For neurons with a large and extensive dendrites compartmental modeling is needed

(Carnevale and Hines, 2006, Chapters 2, 3). However, since such compartmental modeling begins with the equations for modeling a small patch of membrane, we discuss that process here first. In a subsequent Section, the generalizations necessary to model a spatially extended neuron will be described.

Conservation of charge forms the basis for modeling the membrane potential, since the capacitive part of the conservation law gives an expression for the time derivative of the membrane potential $\frac{dV}{dt}$:

$$\begin{aligned}I_C + I_{other} &= 0, \\C_m \frac{dV}{dt} &= -I_{other}.\end{aligned}$$

We proceed by writing expressions (possibly dependent on voltage) for the currents entering and exiting the membrane. The general form of a membrane conductance is:

$$I = g(V - E),$$

where g is the conductance of the “channel” being modeled and E is the reversal potential of the channel in question. The value of g is, in general, a time varying function of V . A simple case in which it is not is the leak conductance:

$$I_L = g_L(V - E_L),$$

which represents the passive leakage of ions across the membrane. A model neuron with only a leak conductance will settle from any initial condition to the reversal potential of the leak conductance, where $\frac{dV}{dt} = 0$. The dynamics of such a patch of membrane is identical to that of the sub-threshold behavior of the LIF neuron with $g_L = \frac{1}{R_L}$ and the battery determining the neuron's resting potential set to E_L .

The reversal potential of an ionic channel can be calculated by finding the potential for which the ionic flux across the membrane due to the concentration gradient of the ion in question equals that due to the electric field across the membrane. In terms of the concentrations inside and outside the membrane ($[C]_{in}$ and $[C]_{out}$) and the charge of the ion, z , the reversal potential is (Johnston and Wu, 1995):

$$E = \frac{RT}{zF} \ln\left(\frac{[C]_{out}}{[C]_{in}}\right).$$

The HH paper provided a model for two additional channels. These are sufficient to produce an action potential followed by a refractory period. More complex models include more conductances, but the technique is the same.

Figure 2.13 illustrates the basic set up of a conductance based model (in this case, the original HH model is shown). For this model, the equation for the derivative of the membrane potential is:

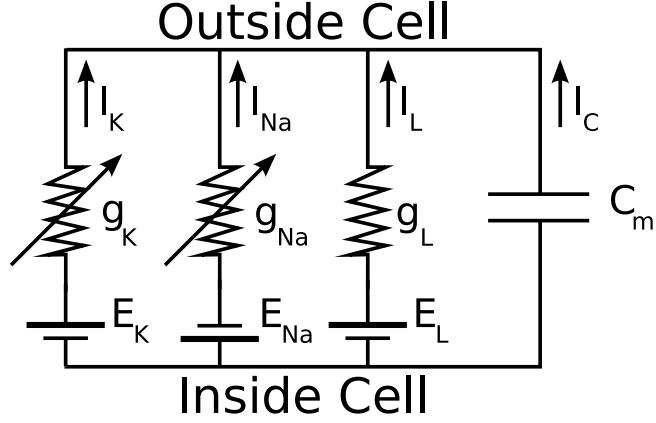


Figure 2.13: The HH model is an electronic circuit with a capacitor and voltage-gated-conductances. (Figure adapted from Johnston and Wu (1995)). The membrane itself is represented by a simple capacitance C_m while conductances g_K and g_{Na} correspond to potassium and sodium channels in the cell membrane, with associated reversal potentials E_k and E_{Na} . The leakage conductance g_L and its associated reversal potential E_L model non-specific, passive conductances across the membrane. Conductance based modeling generally follows this example, but may use many more conductances of various types.

$$\begin{aligned}
 -C_m \frac{dV}{dt} &= I_K + I_{Na} + I_L, \\
 -C_m \frac{dV}{dt} &= g_K(V, t)(V - E_k) + g_{Na}(V, t)(V - E_{Na}) + g_L(V - E_L), \quad (2.1)
 \end{aligned}$$

where g_K and g_{Na} are also dynamical variables which need to be integrated to find the time course of V . In the original HH model,

$$g_K(V, t) = n^4 \bar{g}_K, \quad (2.2)$$

$$g_{Na}(V, t) = m^3 h \bar{g}_{Na}, \quad (2.3)$$

where n , m and h are the *gating variables*, and where \bar{g}_K and \bar{g}_{Na} are the maximum

conductances of each channel.

Conceptually, a gating variable models the behavior of a protein in the cell membrane through which ions of a specific species may or may not pass, depending on its state (Figure 2.2). Gating variables are derived from the assumptions that an individual channel will transition from an open to a closed state (or from a closed to an open state) with a voltage dependent probability $\beta(V)$ ($\alpha(V)$) and that the reaction governing this transition is first order (Johnston and Wu, 1995). Under these assumptions, the rate coefficients $\alpha(V)$ and $\beta(V)$ can be estimated by considering the energy barrier necessarily overcome to transition from a closed to an open state and *vice versa* (Johnston and Wu, 1995).

$$\alpha(V) = \alpha_0 e^{\frac{-\Delta G_0 - \delta z F V}{RT}}, \quad (2.4)$$

$$\beta(V) = \beta_0 e^{\frac{(1-\delta)z F V}{RT}}, \quad (2.5)$$

wherein α_0 and β_0 are constants determined by experiment, $-\Delta G_0$ is the “height” of the activation barrier in zero field (this constant is absorbed into β_0 in Equation 2.5), δ is the scaling factor representing the effect of the transmembrane potential on the activation barrier, R is Boltzmann’s constant, z is the valence of the gated ion, and F is Faraday’s constant.

If we call the open state y , then the time derivative of the probability for this state is given by first order kinetics in terms of α and β :

$$\frac{dy}{dt} = \alpha(V)(1 - y) - \beta(V)y. \quad (2.6)$$

In words, the change in y is the transition probability from the closed state to the

open state times the probability we are in the closed state, minus the probability we are in the open state times the probability of transitioning to the closed state. Each gating variable has α and β functions fit from experimental results. In the case of the h variable in the HH equations, a transposition of δ and $(1 - \delta)$ in the expressions for α and β reflect the fact that the blocking gating variable is inactivated by depolarization.

Conceptually, the HH model imagines that there are four particles governing the transition of the K channel (Equation 2.2) from closed to open, and so the gating variable is raised to the fourth power. The Na channel (Equation 2.3) is “transient”, which means that its state is governed by both a probability of opening and a separate “blocked” probability. In the HH model this is reflected by the use of two gating variables, m and h , where m is the “opening” gating variable, and h is the “blocking” variable. In contrast to the Na, which can become blocked, the K channel is called “persistent”.

Using these formulas, the expression for the time rate of the change of the membrane potential can be written in terms of n , m and h and then integrated by an appropriate numerical integration algorithm. Where single compartment models are used in this dissertation, a fourth order Runge-Kutta algorithm was used to integrate these equations (Press et al., 1992; Gerald and Wheatley, 1999).

2.2.3 Compartmental Modeling

Neurons are spatially extended objects, and for some purposes it is appropriate to simulate not just the activity of one piece of membrane, but a whole spatially extended cell. Doing so efficiently and stably is a complex problem – software packages such as NEURON (Carnevale and Hines, 2006) and GENESIS (Bower and Beeman, 1998) (the former of which is used here for modeling compartmental neurons) exist to abstract away the complexity from the average neuroscientist. However, the basic principles

behind compartmental modeling will be described here for completeness.

Once again, the neuron is modeled as an electronic circuit and conservation of charge and Ohm's law form the basis of model. However, rather than simulating a single patch of membrane, we simulate a series of joined patches (which may branch at the ends). Neglecting branching, the circuit for a compartmental model looks like Figure 2.14.




Figure 2.14: Compartmental modeling extends single compartment modeling by connecting each segment of membrane via an internal resistance (adapted from Johnston and Wu (1995)). Each sub-circuit above represents an electronically discrete patch of membrane (see Figure 2.13), where all the conductances have been visually represented as a single symbol.

We have already written the expression for g_m above. If we discretize the neuron spatially into segments which should be approximately iso-potential, then the expression for the current flowing along the inside of the neuron between section k and j is (Carnevale and Hines, 2006, chap. 3):

$$i_{kj} = \frac{(v_k - v_j)}{r_{jk}},$$

where r_{jk} is the axial resistance between the j th and k th compartment. Here we follow the convention of Carnevale and Hines (2006), and write i_{kj} instead of I_{kj} , the lowercase indicating that i_{kj} is just one of the currents under consideration.

In general, a section of membrane may be connected to many other sections, so the total conservation of charge gives for a compartment j gives

$$(i_m)_j A_j = \sum_k \frac{v_k - v_j}{r_{jk}},$$

where i_{m_j} is the capacitive and ionic current calculated for the single compartment case. Because single compartment modeling is usually done per unit area of membrane (the area can be factored out), the factor of A_j scales the per-unit-area current to the total membrane current. When the situation is that of Figure 2.14, and the section in question has only two neighbors, to the left and right, and writing out the capacitive and ionic components of the membrane out directly, the previous equation becomes:

$$c_j \frac{dv_j}{dt} + i_{ion}(v_j, t) = \frac{v_{j-1} - v_j}{r_{j-1,k}} - \frac{v_{j+1} - v_j}{r_{j+1,k}}.$$

If the neighboring segments have the same surface area and resistance, this is

$$c_j \frac{dv_j}{dt} + i_{ion}(v_j, t) = \frac{v_{j-1} - 2v_j + v_{j+1}}{r \times dx},$$

Where r is the internal resistance per unit length area squared and dx is the spacing between sections. If we let $dx \rightarrow 0$, then we find

$$c_m \frac{\partial v}{\partial t} + i_{ion}(v, t) = \frac{1}{r} \frac{\partial^2 v}{\partial x^2}.$$

If we multiply by the resistance per unit area of the membrane, then we get the following

for a uniform segment of neuron without voltage gated channels:

$$r_m c_m \frac{\partial v}{\partial t} + v_m(t) = r_m \frac{1}{r} \frac{\partial^2 v}{\partial x^2},$$

In order to get a sense for what these values mean, it is useful to write them in terms of specific membrane capacitances and resistances. That is, the capacitance per unit area C_m and the resistance-area (for membrane resistance R_m , which scales with the area) and resistance-length (for the axial resistance R_a , which scales with the diameter of the section in question, d).

$$R_m C_m \frac{\partial v}{\partial t} + v = \frac{d R_m}{4 R_a} \frac{\partial^2 v}{\partial t^2}.$$

This is the so-called “cable equation.” The cable equation is useful for estimating the length and time scales of changes in the membrane potential:

$$\tau_m = R_m C_m, \tag{2.7}$$

$$\lambda = \frac{1}{2} \left(\frac{d R_m}{R_a} \right)^{\frac{1}{2}}, \tag{2.8}$$

where τ_m is, as before, the time constant of the membrane, and λ is its characteristic length scale.

These values are in turn useful for selecting temporal and spatial discretization parameters for numerical modeling. In the presence of channel based conductance models, it is generally not possible to analytically solve the cable equation. Under these circumstances, τ_m provides an upper bound on the simulation time step, and λ

provides one for the spatial discretization Δx . Using typical values of R_m , R_a , and C_m (Johnston and Wu, 1995, Chapter 3):

$$C_m \approx 1 \frac{\mu F}{cm^2} \quad (2.9)$$

$$R_m \approx 10 K\Omega cm^2 \quad (2.10)$$

$$R_a \approx .025 K\Omega cm. \quad (2.11)$$

The membrane time constant can be calculated without reference to the geometric properties of the of the cell, since the units of area cancel out:

$$\tau_m = 10ms. \quad (2.12)$$

We have to provide an estimate for the diameter and area of the membrane in order to estimate λ . A pyramidal cell in the neocortex might have a axonal diameter of $\approx 5 \mu m$ (DeFelipe and Farinas, 1992), so:

$$\lambda = 0.022cm. \quad (2.13)$$

The NEURON environment provides a front end allowing the user to specify directly the biophysics and morphology of an extended neuron, for which it then constructs the partial differential equations in time and space which need to be integrated to model the neuron (Carnevale and Hines, 2006). It also provides tools to statically or adaptively estimate the time step and spatial discretization parameters. The environment automatically calculates the appropriate boundary conditions for terminal sections and

branches and has a specialized numerical integrator for solving the resulting type of differential equations (Carnevale and Hines, 2006).

2.2.4 Modeling Synapses

Although synapses come in two broad categories, excitatory and inhibitory, there are a large variety of synapse types. Electrical synapses, created by gap junctions, are places where two neurons exchange ions directly, via a passage between the two cells (Johnston and Wu, 1995). Chemical synapses, on the other hand, transform the electrical impulse of the action potential in the presynaptic neuron into a neurotransmitter release into a synaptic cleft. The neurotransmitter then binds to receptors in the post-synaptic neuron’s membrane, either directly causing the opening of ion channels or triggering a chemical pathway in the cell which leads to the opening of other channels. Since feed-forward connections are mediated by chemical synapses in the mammalian cortex (Johnston and Wu, 1995, chap. 11) hereafter “synapse” alone will indicate “chemical synapse.”

Among the simplest conductance based models for synaptic activity is the exponential synapse. In this model, the synaptic conductance at a time t is given by the following differential equation:

$$\frac{dg_{syn}(t)}{dt} = -\frac{g_{syn}(t)}{\tau_{syn}},$$

where τ_{syn} is the time constant of the synapse, governing the speed at which the conductance decays. The current provided to the neuron by such a synapse is:

$$i(t) = g_{syn}(t, t_0)(V_{syn} - E_{syn}),$$

where V_{syn} is the voltage at the post-synaptic site and E_{syn} is the synaptic reversal potential. Integration of this synapse numerically proceeds normally except when a presynaptic spike arrives, at which point $g_{syn}(t, t_{0,previous})$ is instantaneously updated to $g_{syn}(t, t_0) + w$, where w can be thought of as the strength of the synapse (Carnevale and Hines, 2006). The time course of the conductance during a period when no spikes are received is simply:

$$g_{syn}(t, t_0) = g_0 e^{(t-t_0)/\tau_{syn}},$$

where g_0 is the conductance at the last received spike time, t_0 . By adjusting the values of w , τ_{syn} and E_{syn} , the synapse can be made to emulate slow, fast, and excitatory or inhibitory synapses.

Real synapses have a variety of other properties not reproduced by this model, even apart from the time course of the synaptic conductance. Neurotransmitters may fail to release in response to an action potential, or release spontaneously (Meir et al., 1999), and the conductance may have medium or long time scale variations, referred to as synaptic depression, facilitation or adaptation (Johnston and Wu, 1995, chap. 11). For the purposes of this dissertation, which examines the response of a neuron over the course of at most 6 seconds, during which the input is fluctuating, the medium and long-time scale variations in synaptic weight are neglected. The issue of synaptic reliability and spontaneous firing is partially dealt with by calibrating the synaptic properties to produce the desired output firing rate and membrane potential standard

deviation, adjusting the spike times which drive each synapses (for instance, lowering the firing rate to reflect synaptic failure).

Chapter 3

Finding and Validating Spike Patterns¹

3.1 Abstract

Neurons in sensory systems must convey information about the temporal structure of stimuli. In vitro, single neurons respond precisely and reliably to the repeated injection of the same fluctuating current, producing regions of elevated firing rate, which are termed events. Further analysis reveals spike patterns, which can be measured as trial-to-trial patterns of correlations between spike times (Fellous et al., 2004). Finding events in data with realistic spiking statistics is challenging because events may overlap. Overlapping events typically belong to different spike patterns. Therefore, we developed a method for finding spiking events that uses information about which pattern a trial belongs to. The procedure can be applied to spike trains of the same neuron across multiple trials to detect and separate responses obtained during different brain states. The procedure can also be applied to spike trains from multiple simultaneously recorded neurons in order to identify volleys of near synchronous activity or to

¹with Jean-Marc Fellous, Peter Thomas, Terrence J. Sejnowski and Paul Tiesinga

distinguish between excitatory and inhibitory neurons. The procedure was tested using artificial data as well as responses recorded *in vitro* in response to fluctuating current waveforms.

3.2 Introduction

The question of how information is encoded in the spiking activity of populations of neurons ranks among the most fundamental in systems neuroscience (Ermentrout et al., 2008; Tiesinga et al., 2008). A commonly used measure for neural activity is the firing rate (Hubel and Wiesel, 1962): the number of spikes per second that the neuron produces, but under some circumstances, cortical neurons produce precise spike times across repeated stimulation by the same stimulus (Bair and Koch, 1996; Buracas et al., 1998; Reinagel et al., 1999; Reinagel and Reid, 2000, 2002). These events are presumably encoding some aspects of the stimulus because they are produced reliably from trial to trial, and their exact timing depends on the specific stimulus. A representative example rastergram and histogram obtained from a model is shown in Figure 3.1A & B. The rastergram is characterized by vertical bands of aligned spike times, which occur almost on each trial at approximately the same time with respect to stimulus onset and which correspond to peaks in the histogram. Each such line/peak is referred to as an event. To better characterize these events, we use measures that quantify the likelihood of occurrence of a spike during an event (reliability) and the precision of this spike when it occurs (precision).

Algorithms have been proposed to determine events (Mainen and Sejnowski, 1995; Tiesinga et al., 2002), but they often fail to separate overlapping events (Figure 3.1B, double-headed arrow). We hypothesize here that this separation can be achieved by taking into account the spiking history, because the refractory period and afterhyperpolarization currents lead to correlations in the form of spike patterns. An example

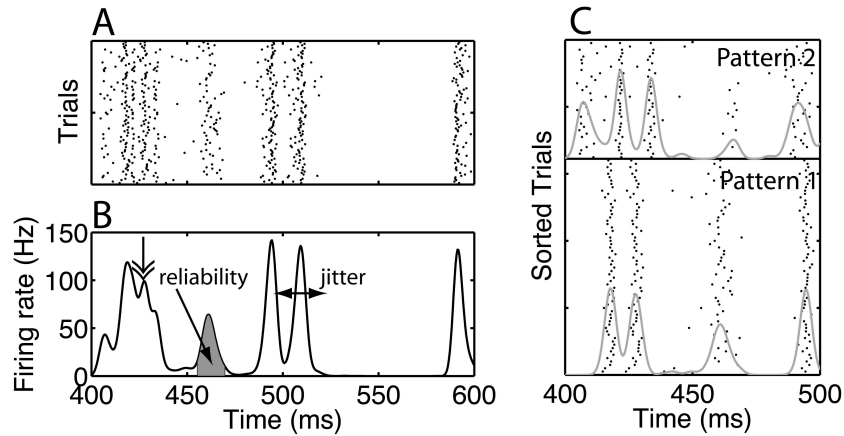


Figure 3.1: Responses can conceal patterns which have high precision and reliability. (A) A rastergram where each tick represents a spike recorded on a given trial, and trials are indexed on the y-axis. (B) “Events” in the rastergram can have precision (inverse of the event standard deviation) and reliability (the probability of a spike occurring during the event, and the area underneath the portion of the firing rate curve corresponding to the event).. (C) In the presence of spike patterns, however, estimating these values can be complicated because events overlap.

with two spike patterns is shown in Figure 3.1C. Although the events were overlapping when the histogram was calculated across all trials (Figure 3.1B, double-headed arrow), they were clearly non-overlapping when trials were re-ordered according to the spike pattern they belonged to (Figure 3.1C). Using spike patterns makes the procedure more powerful, but also more computationally expensive because it involves additional steps: spike patterns have to be detected and events common to multiple patterns have to be merged. We introduce heuristics for the event-based analysis that reduce the computational load and that consequently better allow for an automatic analysis of large volumes of spike train data.

Using this method across trials, events are characterized by their time of occurrence, their precision and their reliability. In the following we will use precision and jitter interchangeably to refer to the temporal resolution of spike times. Numerically, the precision is equal to one over the jitter. An advantage of the event-based analysis is

that it does not rely on fitting a specific parametric model for neural dynamics (Jolivet et al., 2004, 2006; Keat et al., 2001; Pillow et al., 2005), rather it models the data directly.

The method was applied to data we recorded in vitro from layer 5 cortical neurons driven by a fluctuating current waveform with an amplitude that was varied systematically. The neurons locked to the current drive, which means that they produced events at specific times relative to stimulus onset. We find that spike trains change with amplitude such that the information about the stimulus time course is preserved. Hence, the information about the amplitude is implicitly encoded in the trial-to-trial variability. Specifically, we found that the probability of occurrence of spike patterns varied as a function of amplitude.

3.3 Experimental procedures

3.3.1 General experimental procedures

The voltage response of cortical neurons was measured in a rat slice preparation as described previously (Fellous et al., 2001). The Salk Institute Animal Care and Use Committee approved protocols for these experiments; the procedures conform to USDA regulations and NIH guidelines for humane care and use of laboratory animals. Briefly, coronal slices of rat pre-limbic and infra limbic areas of prefrontal cortex were obtained from 2 to 4 weeks old Sprague-Dawley rats. Rats were anesthetized with isoflurane and decapitated. Their brains were removed and cut into 350 μm thick slices on a Vibratome 1000 (EB Sciences, Agawam, Mass.). Slices were then transferred to a submerged chamber containing standard artificial cerebrospinal fluid (ACSF, mM: NaCl, 125; NaH_2CO_3 , 25; D-glucose, 10; KCl, 2.5; CaCl_2 , 2; MgCl_2 , 1.3; NaH_2PO_4 , 1.25) saturated with 95% O_2 / 5% CO_2 , at room temperature. Whole cell patch clamp recordings

were achieved using glass electrodes containing (4-10 M Ω : μ M: KmeSO₄, 140; Hepes, 10; NaCl, 4; EGTA, 0.1; Mg-ATP, 4; Mg-GTP, 0.3; Phosphocreatine, 14). Patch-clamp was performed under visual control at 30-32 °C. In most experiments Lucifer Yellow (RBI, 0.4%) or Biocytin (Sigma, 0.5%) was added to the internal solution for morphological identification. In all experiments, synaptic transmission was blocked by D-2-amino-5-phosphonovaleric acid (D-APV; 50 μ M), 6,7-dinitroquinoxaline-2,3,dione (DNQX;10 μ M), and bicuculine methiodide (Bicc; 20 μ M). All drugs were obtained from RBI or Sigma, freshly prepared in ACSF and bath applied. Data were acquired with Labview 5.0 and a PCI-16-E1 data acquisition board (National Instrument, Austin Tex.) at 10 kHz, and analyzed with MATLAB (The Mathworks).

3.3.2 Stimulus generation and experimental design

To test the event finding method we used data collected to study the effect of varying the amplitude and offset of a repeated frozen noise stimulus. For all experiments the same frozen noise waveform was used. A white noise waveform (sampling rate 10 kHz, with samples uniformly distributed on the unit interval) was generated using the matlab function `rand` with the state of the random number generator set to zero. It was twice filtered using the matlab routine *filter*. First, we applied a low-pass filter with $a = [1 -0.99]$ and $b = 1$. Second, we performed a 50-sample (5 ms) running average ($a = 1$, b has fifty elements equal to $1/50$). The first 500 samples (i.e. 50 ms) were discarded as a transient and the resulting waveform was centered around zero by subtracting the mean and normalized to have unit variance by dividing by the standard deviation, then an offset was added. Depending on the cell and the quality of the seal, the waveform was multiplied by a factor representing the maximum amplitude. Waveforms with fractional amplitudes from zero to one were presented to the cell. The characteristics of the experimental data set is listed in Table 3.1. Trials were separated by at least 15

seconds of zero current, to let the membrane return to its resting state. Throughout the experiment, a few hyperpolarizing pulses were injected to monitor the access resistance of the preparation. These pulses were clearly separated from other stimuli.

3.3.3 General procedures

Spike times were detected from recorded voltage traces as the time the membrane potential crossed 0 mV from below. The firing rate was the number of spikes recorded during a trial, averaged across all similar trials and normalized by the duration of the trial in seconds.

In the rastergram, each row represented a spike train from a different trial. Each spike is represented as a tick or a dot, with the x-ordinate being the spike time and the y-ordinate being the trial number. Often we group trials together based on the stimulus amplitude or re-order trials based on which pattern they belong too. This is indicated in the corresponding figure caption.

The spike time histogram is an estimate for the time-varying firing rate. It was obtained by dividing the time range of a trial into bins (typically 1 or 2 ms wide) and counting the number of spikes that fell in each bin across all trials. The bin count was normalized by the number of trials and by the bin width in seconds. The latter was to ensure that a bin entry had the dimensions of a firing rate, Hz. The histogram was subsequently smoothed by a gaussian filter with a standard deviation equal to 1 bin size.

Events were detected using the procedure detailed in the Results Section. At the end of this procedure, all spikes were either assigned to an event or were classified as noise. The event-reliability is the fraction of trials on which a spike was observed during that event, the event-jitter is the standard deviation of the spike times belonging to the event. The event-precision is the inverse of the event-jitter. For a given condition

(amplitude, offset or initial current step) the reliability, precision and jitter are defined as the event-reliability, event-jitter and event-precision averaged across all events.

In the event finding procedure we use three techniques that are described below: the Victor-Purpura distance; the fuzzy clustering method and classification entropy and mutual information.

3.3.4 Calculation of the VP distance

Briefly, the Victor-Purpura (VP) metric (Victor and Purpura, 1996) calculates the distance between two spike trains A and B by calculating the cost of transforming A into B (or B into A - the measure is symmetric). This distance is obtained as the minimum cost of transformation under the following rules: adding or removing a spike from A costs +1 point, while sliding spikes forward or backward in time by an interval $|dt|$ costs q times $|dt|$. The variable q represents the sensitivity of the metric to the timing of spikes and is expressed in units of 1/ms. For large q values it is frequently cheaper to add and remove spikes than to move them. Hence, for large q , the metric is simply the number of spikes with different times between the two trains. For small q values, spike moving transformations are cheap, leaving the majority of the metric's value to the difference in the number of spikes which must be added or removed to produce train B ; in the limit, the metric becomes the difference in the number of spikes in each spike train.

3.3.5 Fuzzy clustering algorithm

Fuzzy c-means (FCM) was used to cluster trials into groups that had similar spike timings. FCM can be understood by first considering K-means clustering (also, but less commonly, referred to as c-means). In a K-means clustering, a number of clusters is chosen (N_c) and the objects to be clustered are assigned on a random basis to

each of the potential clusters (Duda et al., 2001). Using these assignments, the mean of each cluster is found. Then, using these means, objects are re-assigned to each cluster based on which cluster center to which they are closest. This process repeats until the cluster centers have converged onto stable values or a maximum iteration count is reached. This type of clustering minimizes the sum of the squared distances of the clustered objects from their cluster means. FCM functions in the same way, but rather than belonging to any particular cluster, each object i is assigned a set of normalized probabilities u_{ij} of belonging to cluster j (Bezdek, 1981). This is equivalent to minimizing a non-linear objective function of the distances of the objects from the cluster centers, characterized by the fuzzifier parameter, which is set to two. A more complete description is given in (Fellous et al., 2004). We use FCM on the columns of the pair-wise distance matrix (see Subsection 3.3.4 - **Calculation of the VP distance**) because similar trials will have a similar distance from all other trials and are thus represented by similar columns (Fellous et al., 2004). The computational effort of FCM increases with the number of vectors (N_{trial}) as well as the dimensionality of the vectors (also N_{trial}). Hence, we reduced the dimensionality from N_{trial} to 10 components using principal component analysis (PCA) (Jolliffe, 2002). These components accounted for at least 80% of the variance and resulted in clusterings that were similar to those obtained using all principal components.

3.3.6 Calculation of entropy and mutual information between classifications

The outcome of the FCM procedure is that each spike train is assigned to a cluster. Formally, a set of trials has a classification c_i , where i is the trial index between 1 and N_{trial} ; and the classification c is a number between 1 and the number of clusters N_c . The class distribution p_j is the fraction of trials which were classified as class j ,

$p_j = \frac{1}{N_{trial}} \sum_{i=1}^{N_{trial}} \delta(j - c_i)$, where the δ denotes the Kronecker delta. The diversity of the classification was characterized by the entropy $S = - \sum_{u=j}^{N_c} p_j \log_2(p_j)$, where the sum was over all nonzero p_j values because $0 \log 0$ was defined to be zero (Cover and Thomas, 1991). The entropy S was zero (minimal diversity) when all trials were assigned to the same class and was maximal at $S = \log_2(N_c)$ when all classes had the same probability of occurring (maximal diversity). The entropy was also calculated for continuous values, such as VP distances between pairs of spike trains, in which case the same formulas were applied to a binned probability distribution. The estimation of entropy was biased (Panzeri and Treves, 1995; Strong et al., 1998). However, because these values were used for comparative purposes no bias correction was applied.

The mutual information was used to measure the similarity between two classifications. The joint distribution between two classifications c_i and d_j with N_c and N_d classes, was computed as: $p_{ij} = \frac{1}{N_{trial}} \sum_{k=1}^{N_{trial}} \delta(c_k - i) \delta(d_k - j)$. The mutual information was then expressed as $I = S_c + S_d + \sum_{i=1}^{N_c} \sum_{j=1}^{N_d} p_{ij} \log_2(p_{ij})$, with $S_c = - \sum_{i=1}^{N_c} p_i^c \log_2(p_i^c)$ and $S_d = - \sum_{j=1}^{N_d} p_j^d \log_2(p_j^d)$. In these formulas, the class distributions for c and d have a different subscript in order to distinguish them. Hence, here, i refers to the classification rather than being a trial index. To obtain a measure between 0 and 1 we normalized the mutual information by the maximum entropy, yielding the normalized mutual information $I_N = I / \max(S_c, S_d)$.

3.4 Results: Finding spike patterns and determining the event structure in artificial data

3.4.1 Overview and goal of the event finding procedure

In a recent study we uncovered multiple spike patterns in trials obtained in response to repeated presentation of the same stimulus (Fellous et al., 2004; Reinagel and Reid, 2002). Spike patterns are present when trials, or at least short segments thereof, can be separated in two (or more) distinct groups of spike sequences. As an example consider the case where on some trials the neuron spikes at 10 and 35 ms (relative to stimulus onset), whereas on other trials it spikes at 15 and 30 ms. This group of trials would be considered comprised of two distinct patterns as long as there are no trials with spikes at 15 and 35 ms or 10 and 30 ms. Hence, spike patterns correspond to a within-trial correlation, because, in the example, a spike at 10 ms implies that a spike will be found at 35 ms with a high probability. The problem of finding spike patterns is made harder by the presence of spike-time jitter and trial-to-trial unreliability.

We designed a method to uncover patterns independently of the event structure and then used the patterns to construct the event structure, which can subsequently be used to validate the extracted patterns. The method itself is unsupervised, but four parameters need to be provided. The four parameters are: the threshold for finding events (parameter t_{ISI} , Subsection 3.4.2 - **The interval method for identifying events**); the temporal resolution for which two spike times are considered similar (parameter: q , Subsection 3.4.3 - **Selecting the temporal resolution parameter q**); the number of patterns the clustering algorithm looks for (parameter N_c , which stands for the number of clusters, Subsection 3.4.4 - **Selecting the number of patterns/clusters N_c**); a threshold t_{ROC} which determines which events need to be merged because they are common to multiple spike patterns (Subsection 3.4.5 - **Merging events common**

across multiple patterns). For the parameter settings used here, each cluster corresponds to a spike pattern, hence we will use these designations interchangeably. The procedure has been tested on short temporal segments with on average approximately 2 or 3 spikes per trial. These segments can be found by cutting spike trains at times with a low or zero spike rate in the spike time histogram.

The basic premise is that two trials on which the same pattern was produced (in the following text we restrict the discussion to one segment) are more similar to each other than two trials on which different patterns were produced. We use the Victor-Purpura distance (VP, Victor and Purpura (1996) see Subsection 3.3.4 - **Calculation of the VP distance**) to quantify the similarity. The distance between spike trains i and j is represented as a matrix d_{ij} . The distance matrix generally appeared unstructured when the trials were arranged in the order in which they were recorded (look ahead to Figure 3.12A for an example). The goal is to re-order the matrix such that it becomes block-diagonal (Figure 3.12F). The block diagonals correspond to trials that have a small distance among themselves and are more distant to trials outside the block. That is, blocks on the diagonal correspond to spike patterns. This goal is achieved using the fuzzy c-means (FCM) method (Bezdek, 1981) applied to the columns of the distance matrix (see Subsection 3.3.5 - **Fuzzy clustering algorithm**). FCM finds N_c clusters and assigns to each trial a probability of belonging to a cluster. If the clustering is good each trial belongs to only one cluster with a high probability, if it is bad a trial has similar probabilities of belonging to two or more different clusters. Once patterns have been uncovered a preliminary event structure is determined using on each pattern the method outlined in the Subsection 3.4.2 **The interval method for identifying events** (Tiesinga et al., 2002). The common events that occur in multiple spike patterns were found and merged using an Receiver Operating Characteristic (ROC) analysis (Green and Swets, 1966). After this analysis a cluster-assisted event

structure is available. For the purpose of this paper the spike pattern detection and event assignments were manually checked and corrected using an in-house interactive software tool. A comparison of intuitively correct event structures is shown in Figure 3.2. All 474 segments were run through the event finding procedure with default parameter values, and the results compared to the intuitively correct event clustering as determined by this author. Forty-three percent of the segments were “Acceptable”, “Good” or “Excellent”. Eight-nine percent were equal or better than “Mediocre,” implying a small number of errors. “Poor” data was excluded from analysis as lacking clear event features, and mediocre data was corrected by hand, either by adjusting the parameters of the algorithm for each set, or by manually assigning spikes to events.

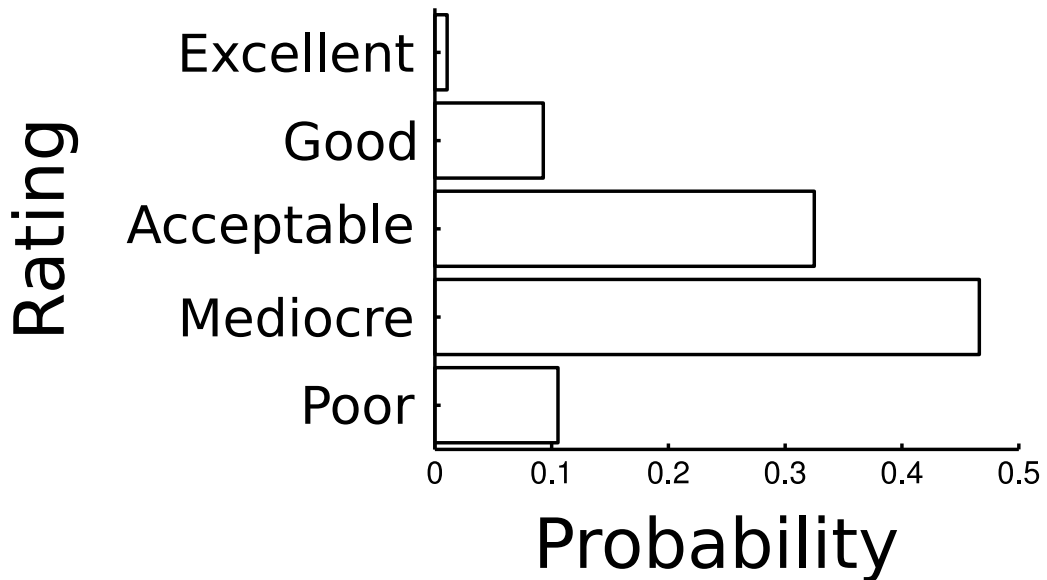


Figure 3.2: Summary of a comparison between automatically detected events and events categorized by eye. Four-hundred and seventy four segments of data were analyzed from the data sets described in this paper, all with one set of parameter values. Each was run through the detection algorithm described in 3.3.1, and the output event finding was displayed. The author then rated each event-clustering based on his intuitive estimate of the correct number of events. “Excellent” indicated an exact agreement, “Good” indicated one or two mergers of nearby events. “Acceptable” indicated mergers between near, but more clearly distinct events, “Mediocre” indicated at least one large merger error or that part of the data set was not properly divisible into events, and “Poor” indicated that the data set consisted of spike times not properly divisible into events.

3.4.2 The interval method for identifying events

Our event-finding procedure is based on the interspike intervals (ISI) of an aggregate spike train, which is obtained by combining spikes across all trials into one set and sorting them, with the earliest spikes first (Tiesinga et al., 2002). Intuitively, consecutive spikes in the same event are close together, and have a small ISI, whereas those in different events are further apart (Figure 3.3A). We can thus separate spikes in different events using a threshold t_{ISI} : only spikes separated by an interval less than t_{ISI} are considered to be in the same event. In addition, we require that each event contains at least m_s spikes. Taken together, this leads to the following algorithm.

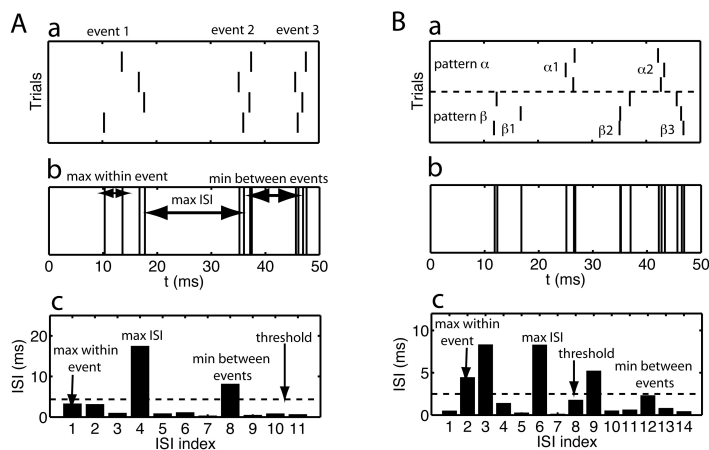


Figure 3.3: It is not always possible to choose an appropriate threshold for the interval method. (A-B) In each panel, we show, from top to bottom, (a) the rastergram, (b) the aggregate spike train and (c) the ISI time series. The relevant time scales are indicated in the graph: the maximum ISI between two spikes in the same event, the minimum and the maximum ISI between two spikes in different events. An estimate for the threshold, 0.1 times the maximum ISI, is shown as a dashed line in panel c. (A) An example with one spike pattern where an appropriate threshold can be chosen that is higher than the maximum ISI within an event but less than the minimum ISI between events. (B) An example with two spike patterns (labeled by α and β) where there is no such threshold that separates within-event intervals from those between events.

1. Take all spikes from all trials and place them in an aggregate, time ordered list;
2. Beginning with the first spike, collect spikes (and remove them from the list) until

the difference between the last collected spike and the next spike is greater than t_{ISI} ;

3. If the collected spikes exceed the minimum number of spikes allowed in an event (m_s) assign them to a new event. Otherwise, categorize the spikes as noise.
4. If any spikes remain in the time ordered list, return to step 2.

We set a minimum spike number for two reasons. First, to avoid calculating the precision based on too few spikes. Second, to avoid interpreting a near coincidence of two or more noise spikes as an event. To see how this could happen, consider the following hypothetical situation where the data set contains a Poisson process with a fixed rate of R_{noise} . The aggregate spike train is then a Poisson process with a rate $R = N_{trial}R_{noise}$. The probability p of obtaining two spikes with a distance of less than t is $p = R \int_0^t e^{-Rs} ds = 1 - e^{-Rt}$ (Rieke et al., 1997). Because R increases with the number of trials, for a large enough number of trials there likely will be near-coincident noise spikes. Specifically, for $R_{noise} = 1Hz$ and $N_{trial} = 100$, the probability of a within 3 ms coincidence is 26%. In the analysis performed for this paper, between 10 and 150 trials were used and the m_s was set to two.

The correct threshold for interspike intervals is determined by three distinct time scales in the data, which are shown in Figure 3.3. An appropriate threshold is greater than all interspike intervals within an event but less than all ISIs between spikes in different events (Figure 3.3A). Because two consecutive spikes in the aggregate spike train come from different trials, the inter-event interval can be smaller than the refractory period, which means that in some cases there is no value for the threshold that satisfies the constraint (Figure 3.3B). This problem is addressed by grouping trials according to the spike patterns they express and analyzing each group separately with the interval method. We found that a value of 10% of the maximum ISI often provided results

similar to events detected by hand. As a rule of thumb the appropriate value for the threshold decreases as the number of available trials increases. For the data sets in this paper, we used a t_{ISI} between 1 and 3 ms.

3.4.3 Selecting the temporal resolution parameter q

The first step in uncovering spike patterns is to select an appropriate value for the temporal-resolution-parameter q . The defining characteristic of spike patterns is a significant difference in event timing. Furthermore, spikes in spike patterns could be unreliable, which means there could be a difference in spike count between trials expressing the same spike pattern. Hence, low q values are not appropriate because the distance is dominated by the difference in spike count, which would split up patterns with unreliable events. For low q values the elements d_{ij} of the distance matrix are approximately integer. High q values are also not appropriate because the distance is dominated by small differences in spike timing, which splits trials expressing the same pattern when events are imprecise. Because the distance is the difference in spike count plus twice the number of spike pairs that are different by more than $\frac{2}{q}$, it again takes values close to integers. (Note that it is not exactly an integer because there still are some spike pairs separated by less than $\frac{2}{q}$, although their number indeed goes to zero in the limit where q goes to infinity - in practice this occurs when the time resolution of spike recording is reached, at which point spikes are either identical or deleted and added, and the metric takes on integer values only). The optimal q value is in between these two extremes. We developed a procedure to find this value using only information on how the distances d_{ij} behave as a function of q .

The procedure is illustrated using artificially patterned spike trains mimicking experimental data (Figure 3.4A). The first step is to combine all the matrix elements (Figure 3.4B) into one set and calculate a histogram (Figure 3.4C). A new matrix is

obtained by stacking the histograms for each q value as columns (Figure 3.4D). The image of this matrix shows a transition from integer-dominated distances for low q values (Figure 3.4D, left ellipse) to integer-dominated distances for high q values (Figure 3.4D, right ellipse) via a broader and a less peaked transition state. We find that the best results are obtained by using $\log(q)$ values uniformly distributed between 10^{-4} and 2 (ms^{-1}). The transition can be visualized by determining the fraction of distances within 0.05 of an integer as a function of q (Figure 3.4E, black line, left scale). As expected a trough emerges at intermediate q values. Because the entropy of a distribution with a sharp peak is lower than that of a broader distribution, the transition is signaled by a peak in the entropy of the distance distribution as a function of q (Figure 3.4E, gray line, right scale). The entropy was calculated using the formula $S_d = \sum h_i \log_2(h_i)|_{h_i \neq 0}$, where subscript d stands for distance and h_i is the fraction of distance values falling in the i th bin of distance histogram (see Subsection 3.3.6 - **Calculation of entropy and mutual information between classifications**). We achieved good results using 200 bins to cover distance values obtained across the entire range of q values studied. A more sensitive measure for detecting changes in the distinguishability of patterns is based on the coefficient of variation of the distances (CV_d), which, unlike the entropy measure, does not require binning of the data. The CV_d generally decreased with q (Figure 3.4F, black line, left scale), but there were modulations in the rate of change which reflected the structure in the data that we sought to uncover. These modulations were clearest in the differenced CV_d (dCV_d , Figure 3.4F, gray line, right scale). This measure corresponded to the derivative of CV_d with respect to $\log_{10}(q)$, because the $\log(q)$ values were uniformly distributed. We found that the best clusterings were obtained at the q value corresponding to the deepest trough after the highest peak. However, this location varied across independent stochastic realizations of spike trains generated using the same spike-pattern parameters (see Subsection 3.4.6 **Binary word**

representation of spike patterns). Hence, to make the procedure more robust we choose as q value the mean of the position of the peak in the entropy S_d (less sensitive but robust) and the position of the deepest trough after the highest peak in the dCV_d (sensitive, but less robust). Note that the mean of the peak and trough is not exactly in the middle when the data are plotted on a log scale for q as in Figure 3.4E and F.

To gain more insight in the sensitivity of the dCV_d modulation to features of the spike patterns, we analyzed the clusterings obtained for different q values. For this analysis, we used the best clustering (optimal N_c value) determined by the differenced gap statistic defined in the following Subsection 3.4.4 **Selecting the number of patterns/clusters** N_c . A clustering assigns to each trial i a class c_i , which is also referred to as a classification. Two clusterings can be compared by computing the normalized mutual information I_N between two classifications (see Subsection 3.3.6 - **Calculation of entropy and mutual information between classifications**). When the two clusterings are identical, $I_N = 1$. This is true even when the class labels are different (for instance when class 1 in classification c corresponds to class 2 in classification d), because the mutual information is insensitive to a permutation of class labels.

We constructed a similarity matrix SM_C using the normalized mutual information I_N between clusterings obtained for different q values. Specifically, the element in the i th row and j th column is the I_N between the clustering for the i th q value and the clustering for the j th q value (both with the correct number of clusters N_c). The resulting matrix (Figure 3.5A) consists of diagonal blocks, indicating that the same clustering was obtained over a range of q values: the larger the block, the more robust the clustering. Because the data was generated from a known parameter set, the correct class assignments were known. Hence, the mutual information between the best clustering for a specific q value and the correct classification was determined. The resulting graph of I_N versus q (Figure 3.5B) consisted of a sequence of steps, each

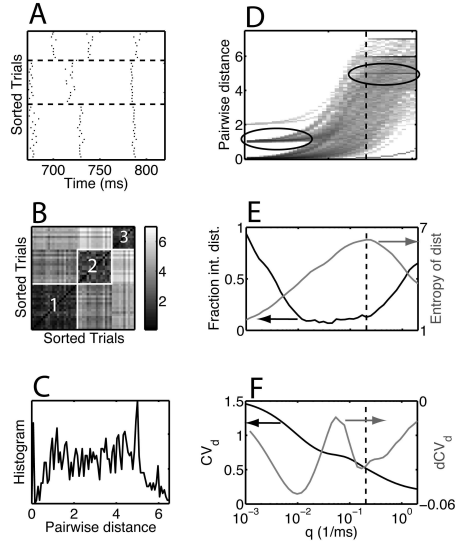


Figure 3.4: Heuristic for selecting the temporal resolution parameter q . (A) Rastergram of 45 trials containing three patterns (separated by horizontal dashed lines). (B) The distance matrix for the below-determined q value. The matrix has a diagonal structure with blocks of small distance values corresponding to the patterns (as labeled by the numbers inside the blocks). (C) The histogram of the matrix elements of the distance matrix shown in panel B. (D) Density plot of the distance distribution as a function of q . Each column represents a histogram just as the one shown in panel C, the location of which is indicated by a vertical dashed line. The ellipses highlight examples of enhanced density for integer values for the distance for small and large values of q . Panels D-F have a common x-axis scale, which is shown in panel F. (E) We show (black curve, left-hand-side scale) the fraction of distances within 0.05 from an integer and (gray curve, right-hand-side scale) the entropy of the distance histogram as a function of q . (F) We show (black curve, left-hand-side scale) CV of the distances and (gray curve, right-hand-side scale) the differenced CV .

corresponding to a distinct clustering that was obtained across a range of q values. Transitions between clusterings are related to modulations in the dCV_d (Figure 3.5C). The correct clustering was obtained just after the highest peak in dCV_d (Figure 3.5C, white arrow in panel A) but only for a small range of q values. The next best clustering was obtained in a broad range of q values around the deepest trough following the highest peak in the dCV_d (Figure 3.5C, asterisk in panel A). This clustering was selected by the heuristic. Although for experimental data the correct clustering is not available, the similarity matrix between clusterings for different q values can still be determined in order to select a small number of distinct and robust clusterings for more detailed evaluation.

3.4.4 Selecting the number of patterns/clusters N_c

Once an appropriate q value is selected, the optimal number of clusters N_c needs to be determined. This is a difficult problem that has attracted significant attention (Bouguessa et al., 2006; Pal and Bezdek, 1995; Rezaee et al., 1998; Zahid et al., 1999). Intuitively, the best clustering minimizes the distance between trials within a cluster, while maximizing the distance between trials in different clusters. The output of the FCM algorithm is the matrix u_{ij} , which is the probability of trial i belonging to cluster j and the cluster centers, which are obtained by averaging across all vectors corresponding to trials in the cluster. A large number of cluster-validation measures have been proposed based on u_{ij} and the cluster centers that quantify these intuitive expectations (Bouguessa et al., 2006). An estimate for the number of clusters is obtained as the number that minimizes (or maximizes, as appropriate) these validation measures. In previous work, these measures were successfully tested on data generated from Gaussians mixture models (Bouguessa et al., 2006). When we applied these measures to our data we did not achieve success because our experimental data did not correspond to a




Figure 3.5: The q value chosen by the heuristic corresponds to a robust close-to-optimal clustering. We analyzed an artificial data set with three patterns and a known classification. The FCM algorithm was applied to distance matrices for different q values, with the number of clusters N_c set to three. (A) The pair-wise similarity between clusterings (classifications) for different q values was quantified using the normalized mutual information (I_N) and is shown as an image. The range of $I_N = 0$ to $I_N = 1$ is represented on a gray scale from white to black. The clustering indicated by the arrow was optimal, the clustering picked by the q -value heuristic is part of the block labeled by an asterisk. The blocks of similar clustering are delimited by dashed boxes, which were obtained by clustering the classification similarity matrix. (B) The similarity between the known classification and those obtained by the FCM algorithm as a function of q . The black arrow indicates the correct clustering. (C) We show (black curve, left-hand-side scale) the dCV_d and (gray curve, right-hand-side scale) the entropy of the distances as a function of q . The dashed vertical lines in (B-C) correspond to the dashed boxes in panel A. The vertical gray line is the q value chosen by the heuristic. The x-axis scale is the same for panels A to C and is shown in panel C.

Gaussian mixture model (see below). Instead, our criterion for picking the number of clusters is based on the gap statistic (Tibshirani et al., 2001; Yan and Ye, 2007), which is calculated as follows. The data is divided into N_c clusters with cluster assignments c_k . The within-cluster distance D_r for the r th cluster is the Euclidean distance $d_{kk'}$ averaged across pairs of trials in the same cluster, $D_r = \sum_{k,k'} \delta(r - c_k)\delta(r - c_{k'})d_{kk'}$. Note that the Euclidean distance is based on the first ten principal component coordinates of columns in the distance matrix (see Subsection 3.3.5 - **Fuzzy clustering algorithm**), because the spike trains themselves are not vectors in an Euclidean space (Victor and Purpura, 1996). The weighted distance for N_c clusters is $W_{N_c} = \sum_{r=1}^{N_c} \frac{D_r}{2n_r}$, here n_r is the number of elements in cluster r , that is, $n_r = \sum_{k=1}^{N_{trial}} \delta(r - c_k)$. The quantity W is compared to values of clusters obtained from clustering surrogate data that is uniformly distributed in the range spanned by the original data. This yields $\tilde{W}_{N_c}^i$ for the i th surrogate data set, the gap statistic then is: $G_{N_c} = \sum_{i=1}^B \log(\tilde{W}_{N_c}^i) - \log(W_{N_c})$, where B is the number of surrogates used. When data with C clusters was generated using a Gaussian mixture model, this measure correctly peaked at $N_c = C$. For the artificial data generated based on our experimental recordings this was generally not the case, because a cluster with only one pattern but with an unreliable event can always be broken up in two clusters, yielding smaller within cluster distances. For instance, consider the case for which event 2 is unreliable. The cluster containing the pattern can be split into a cluster without a spike during event 2 and a second cluster with a spike during event 2. We found that for an appropriate q value the biggest change in the gap statistic occurred when a cluster containing two distinct patterns was split in two (Figure 3.6A). We therefore used the following heuristic: we pick the number of clusters obtained after the largest increase in the gap, that is the one yielding the largest value of the gap difference, $dG_{N_c} = G_{N_c} - G_{N_c-1}$ (asterisk in Figure 3.6A). For our data set we added two more rules to the heuristic. We found them to be appropriate for our data,

but we did not evaluate their performance on a broader set of data. Hence, for other data sets the percentages quoted in the following should be adjusted if necessary. First, when there are two or more nearly equal consecutive increases (here to within 20%), we pick the cluster number obtained after the latest increase, even if it is the smaller increase. Second, if the first dG value was the highest, but there was another peak a certain fraction (here: 80%) of the highest dG value, this peak was chosen instead.

The heuristic for choosing q and N_c was studied using surrogate data sets with varying numbers of spike patterns and trials. For Figure 3.6, we generated ten independent sets of 45 trials containing the same three spike patterns. For each q value the distance matrix was first calculated, then clustered ten times for N_c values between 2 and 10. For each N_c , the clustering with the lowest W_{N_c} value was chosen. Subsequently, the gap-statistic was determined by clustering $B = 10$ surrogate data sets. The number of clusters that maximized dG was picked (for these data there was only one peak). The clusterings so chosen were compared to the correct clustering in terms of the normalized mutual information I_N . In Figure 3.6B, the number of clusters chosen are shown as a function of q . In the same graph the heuristic choice of q for each set of trials is indicated by a tick and the mean q across ten sets is indicated by the dashed vertical line. This mean q value is in the middle of the range of q values for which the dG statistic picks the correct number of three clusters/patterns with 100% accuracy. For other q -values there is a high variability in the number of clusters selected and the mean is different from the correct value. In Figure 3.6C, I_N is plotted as a function of q . It has a peak value near 1, whose location is correctly predicted by the heuristically chosen value of q . To evaluate how this performance generalizes to other patterns, we determined three statistics: (1) How often the correct number of clusters is picked, (2) The mean value of the normalized mutual information I_N at the selected q and N_c value, (3) The relative height of the peak in the difference gap statistic, dG . The relative height was

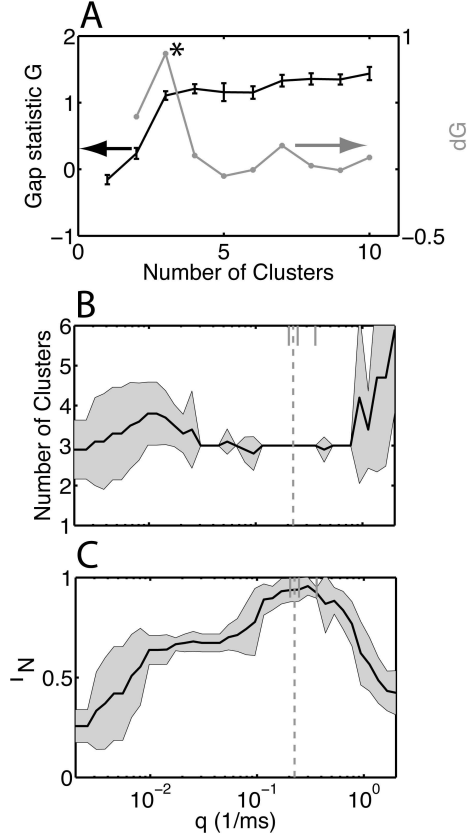


Figure 3.6: The heuristics for finding the q value and the number of clusters perform well for artificial data sets. (A) We show (black curve, left-hand-side scale) the gap-statistic G and (gray curve, right-hand-side scale) its difference dG as a function of the number of clusters. The number of clusters N_c reached after the highest increase in the gap-statistic (at the peak of dG indicated by the asterisk) is selected as a heuristic for the number of clusters in the data. (B) We generated ten independent artificial data sets of 45 trials with three patterns all with the same statistics as those shown in Figure 3.4. The average N_c value picked by the heuristic is plotted as a function of q , the shaded area represents the standard deviation. (C) The similarity I_N between the clustering obtained with the selected N_c value and the known classification as a function of q . The black line is the mean and the shaded area represents the standard deviation. The dashed line in B and C is the mean value of the q value picked by the heuristic, the choice for each of the ten data sets is indicated by the ticks at the top of the graph (these ticks are restricted to be one of the log-uniform distributed q values). The heuristically picked q value leads to an optimal similarity and an N_c that is equal to the actual number of clusters present in the data.

obtained by comparing the peak dG to the mean and standard deviation of the other dG values with the peak value excluded. It was numerically equal to the peak in dG minus the mean, after which the result was divided by the standard deviation. Values exceeding 2 were considered good. We used two data sets that were based on analyzed experimental data (the first one, E1, has been used for the data in Figures 3.4, 3.5 and 3.6) and an artificial set constructed to have four patterns. The procedure performed well on all these sets as measured by the statistics (1) to (3). Specifically, the correct number of clusters was picked between 90 and 100% of the time; I_N exceeded 0.90 and the peak in the gap was higher than 3. We also reduced the precision of set E1 by doubling the jitter and in another simulation we halved the number of trials. We found that the reduction in precision affected the performance of the procedure the most: the correct number of clusters was picked only in 60% of the cases, I_N was reduced to 0.6, but the relative peak height of the difference gap-statistic remained at about 3.

3.4.5 Merging events common across multiple patterns

When for each individual spike pattern the events are detected using the interval method, there are events common to multiple patterns. These common events have to be merged prior to assessing the significance of the patterns. The t-test is a statistical procedure for testing whether the means of two groups of data are equal under the assumption that their standard deviations are the same (Larsen and Marx, 1986; Rohatgi, 2003). The means (events) are considered different when the probability of obtaining a t-value higher than the measured value is less than the significance level. Within the context of the experimental data it is not possible to use this statistic to design an automatic method for merging events. *First*, the t-statistic tests whether the samples are drawn independently from one normal distribution. For the experimental data, events are often split across patterns in a non-symmetric way. For instance, for

spike times in a common event it is often the case that the spikes belonging to one pattern always appear before those of the other pattern. This can be modeled by drawing spikes from a normal distribution, sorting them and assigning the lowest half to one group and the remainder to the second group. Because the spikes in these groups are not independently drawn, the difference in means between the two groups would be highly significant according to the t-statistic even though all spike times come from the same distribution. *Second*, even if there is a well-defined criterion for pair-wise merging of clusters, we found that the following problem arose. Consider three groups of spikes. As groups 1 and 2 are merged, the standard deviation of the aggregate group will likely increase. This increases the likelihood of the combined group merging with the third group, even though by itself group 1 would not merge with group 3. Hence, the result depended on the order in which the pair comparisons were made. Our goal was to design a procedure that was characterized by only a few parameters that would yield reproducible results, while addressing these preceding two issues.

The t-statistic measures the significance of the difference in means over the standard deviation ($\frac{\Delta m}{\sigma}$). We use a ROC analysis for the same purpose (Green and Swets, 1966), because it has the advantage that it does not assume an underlying Gaussian distribution. Assume that there are two sets of spike times described by a probability distribution $p_1(t)$ and $p_2(t)$, respectively, with the mean of the first being below the mean of the second (see Figure 3.8A). A possible decision rule is to assign an observation x to p_1 if $x < \mu$ and to p_2 , otherwise (μ is the threshold parameter). The true positive rate of this rule is $1 - \epsilon_1 = \int_{-\infty}^{\mu} p_1(t) dt$ and the false positive rate is $\epsilon_2 = \int_{-\infty}^{\mu} p_2(t) dt$. The ROC curve is $1 - \epsilon_1$ plotted versus ϵ_2 for all possible choices of the threshold μ . Because the theoretical distribution is not available, we use the empirical distribution with a delta function at the location of each observation: $p_1(t) = \frac{1}{N} \sum_{i=1}^N \delta(t - t_i)$, here N is the number of observations x_i . An analogous definition holds for $p_2(t)$. The area

under the ROC curve ranges from 0.5 for distributions that cannot be distinguished to 1 for distributions that can be perfectly distinguished. We show two distributions in Figure 3.8A, together with 30 samples drawn from them. The resulting ROC curve is in panel B, with the area above the diagonal shown in gray. We use a scaled ROC value defined as $sROC = (\frac{ROC_{area}-0.5}{0.5})^4$ because we found it to have a larger dynamical range (between 0 and 1) and to be more linear around $\frac{1}{2}$. The ROC and sROC measures are compared in Figure 3.8C as a function of $\frac{\Delta m}{\sigma}$. The data set shown in Figure 3.8D contains more than two events. For each pair of events the sROC distance is determined and placed in a distance matrix (Figure 3.8E). We cluster this matrix using a hierarchical clustering method (using the Matlab routines *linkage* followed by *cluster*), to merge all groups of events with a pair-wise sROC distance less than the threshold t_{ROC} . Note that events are not merged pair-wise, all events are merged at once after which their statistics (standard deviation, reliability) are recalculated. We analyzed the performance of this procedure for multiple independent realizations of the data sets. The procedure performed well even for sets with as few as eleven samples per event and it was relatively insensitive to the value of the sROC threshold (Figure 3.8F). Typically, we use sROC thresholds of about $\frac{1}{2}$.

3.4.6 Binary word representation of spike patterns

Once the event structure has been determined, the data can be represented as binary words for the purpose of validating the presence of spike patterns. In this analysis the spike time jitter is ignored. Each trial is a binary word with a length (the number of bits) equal to the number of events that were detected, a bit is one when there was a spike during the event and zero otherwise. The entire experiment is reduced to a matrix b_{ij} , with i being the trial index (from 1 to N_{trial}) and j being the event index (from 1 to N_E). The null hypothesis is that each event is independently occupied

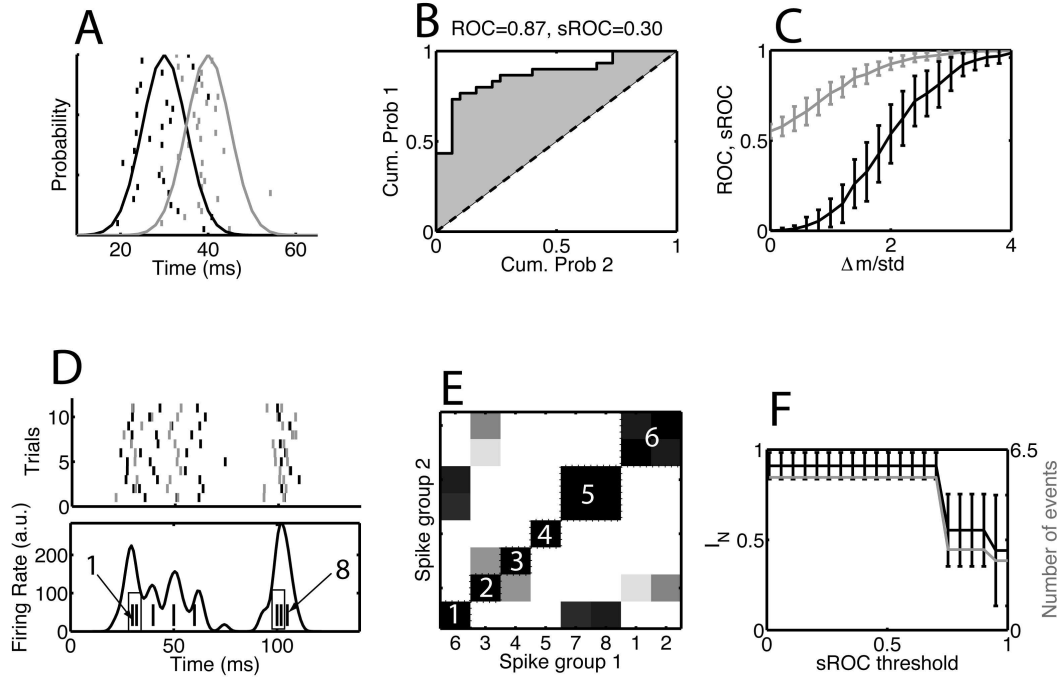


Figure 3.7: The ROC-based criterion for merging events common across multiple patterns. (A) Thirty trials of (black ticks) spikes drawn from a (black curve) Gaussian with a mean of 30 ms and thirty trials of (gray ticks) spikes drawn from a (gray curve) Gaussian with mean 40 ms. The standard deviation of the Gaussian was 5 ms for both cases. (B) The ROC curve was estimated as the cumulative distribution of the first group of spikes plotted versus that of the second group of spikes (see Subsection 3.4.5 - **Merging events common across multiple patterns**). The ROC area was 0.87. The scaled ROC (sROC) was 0.30, which was obtained by normalizing the ROC area above the diagonal by 0.5 and taking the fourth power. (C) The same procedure was repeated for different values of the difference in means between the two distributions. We show the (gray curve) ROC and (black curve) sROC area versus the difference of theoretical means over the standard deviation ($\frac{\Delta m}{\sigma}$). We use the sROC because it is more linear near $\frac{\Delta m}{\sigma} = 2$. (continued) ...

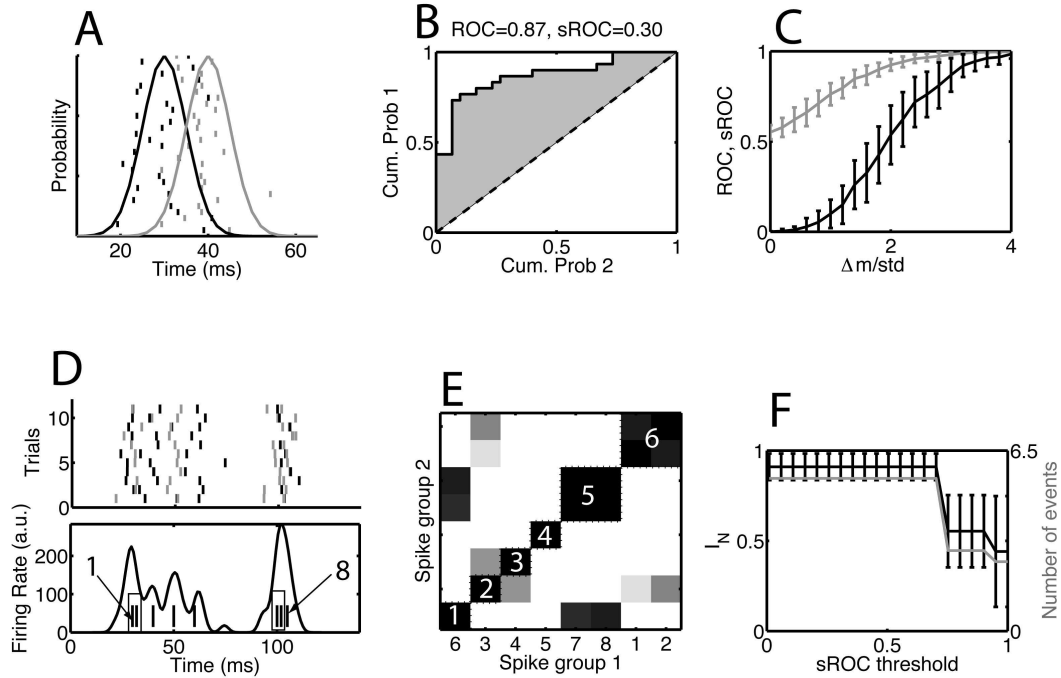


Figure 3.8: (continued) ... (D) Eleven trials of an example data set consisting of 8 events, of which two sets should be merged based on the criterion $\frac{\Delta m}{\sigma} < 1$. In (top) the rastergram the events are sorted by their theoretical means, with the corresponding spikes rendered alternately in gray and black. (Bottom) the histogram obtained from one realization of the data set. The ticks at the bottom represent the theoretical means of the events and are labeled 1 to 8, with the events to be merged shown in a box. (E) Matrix of sROC distances between the events. The events are ordered (the numbers on the x-axis correspond to those in panel D) according to a hierarchical clustering with an sROC threshold of $t_{ROC} = \frac{1}{2}$. (F, black curves) The normalized mutual information between the classification using clustering and the desired cluster and (gray curves) the number of events remaining after merging. The error bars on the mutual information represent the standard deviation across 100 independent realizations of the data set. For the event distribution shown here, the desired merges are not always achieved, because sometimes events 6 and 7 are merged and sometimes events 7 and 8 are merged.

with a probability p_i , which is the reliability of the event. The event-reliability is estimated as the fraction r_i of trials on which there is spike during the event, specifically $r_i = \frac{1}{N_{trial}} \sum_{i=1}^{N_{trial}} b_{ij}$ and $p_j = \langle r_j \rangle$, where the average is taken across experiments with the same number of trials. The probability for obtaining a word $w = [b_1 \dots b_{N_E}] = \sum_{j=1}^{N_E} b_j 2^{j-1}$ under the null hypothesis is a product of Bernoulli probabilities, $P_{null} = \prod_{j=1}^{N_E} p_j^{b_j} (1 - p_j)^{1-b_j}$. When there are two (or more) spike patterns, the experiment is described as a mixture model of two (or more) word distributions. For simplicity we describe the procedure assuming that there are only two patterns: A and B. This yields $P_{patt}(w) = c_A P_A(w) + c_B P_B(w)$ with c_A being the fraction of trials with pattern A, and $c_B = 1 - c_A$, $P_A(w) = \prod_{j=1}^{N_E} p_{A,j}^{b_j} (1 - p_{A,j})^{1-b_j}$ and $P_B(w) = \prod_{j=1}^{N_E} p_{B,j}^{b_j} (1 - p_{B,j})^{1-b_j}$. In this expression $p_{A,j}$ and $p_{B,j}$ are the reliabilities for event j in pattern A and B, respectively.

Certain events in pattern A are never occupied because they belong to pattern B. For a word containing such an event $P_A(w) = 0$. In order to estimate the parameters for this pattern distribution, each trial needs to be assigned to a pattern (for instance using the clustering procedure described in the preceding text). The reliabilities $p_{A,j}$ are then estimated as $R_{A,j} = \frac{1}{N_A} \sum_{i \in A} b_{ij}$, where the average is across all trials i on which pattern A was obtained. The pattern probability \hat{C}_A is estimated as the fraction of trials on which pattern A was obtained. Similar expressions hold for $r_{B,j}$ and \hat{C}_B .

The pattern distribution is not necessary in order to accept or reject the null hypothesis, but we use it to determine how many trials are required to reliably detect a given pattern using estimated reliabilities. Each experiment b_{ij} corresponds to N_{trial} words w_j , which are represented by the word distribution $n_w = \sum \delta(w - w_i)$, where δ , as before, denotes the Kronecker delta. The significance of the pattern is determined in terms of the empirical χ^2 computed using the estimated reliabilities r_i :

$$\chi_{emp,null}^2 = \sum_w \frac{(n_w - N_{trial} P_{null}(w|r_j))^2}{N_{trial} P_{null}(w|r_j)}.$$

When the probability to obtain this χ^2 value is less than the significance level, the null-hypothesis is rejected, which means there is a significant spike pattern.

We examine this using a distribution with parameter values $p_A = [0.4 \ 0 \ 0.8 \ 0]$, $c_A = 0.5$, $p_B = [0 \ 0.6 \ 0.6 \ 0.8]$, and $c_B = 0.5$. The corresponding parameter value for the null-hypothesis distribution is $p = [0.2 \ 0.3 \ 0.7 \ 0.4]$. The rastergram and binary word representation for an example experiment with pattern A and B each expressed on ten trials is shown in Figure 3.9A and B. The corresponding theoretical reliability model is shown in panel C. In panel D the empirical word distribution (black vertical lines) is compared to the theoretical pattern distribution P_{patt} (gray vertical bars), whereas in panel E it is compared to the theoretical null-hypothesis distribution P_{null} (gray vertical bars). The key observation is that the empirical distribution is different from both theoretical distributions. As the number of trials increases the χ^2 difference of the pattern spike trains with the theoretical pattern distribution is expected to decrease, whereas the difference with the null-hypothesis distribution should increase.

To examine this we generated a thousand word distributions n_w from the pattern distribution and the null-hypothesis distribution and calculated the χ^2 between the respective word distributions and the exact null-hypothesis distribution. For 20 trials (Figure 3.10A) the two distributions of χ^2 -values obtained across 1000 realizations of the empirical word distribution were overlapping, whereas for 100 trials (Figure 3.10 B) they were clearly separated. The ability to distinguish these two χ^2 values based on a single experiment was quantified using an ROC analysis based on the difference between the two distributions. The area under the ROC curve increased from 0.5 (not distinguishable) as the number of trials increased. For more than 50 trials the empirical word distribution could be perfectly distinguished from the null-hypothesis distribution (area under ROC curve close to 1), but the distinguishability was better than chance for more than ten trials.




Figure 3.9: Binary word representation for a data set with patterns. (A) An example rastergram with twenty trials and two patterns (pattern 1 for the bottom ten trials and pattern 2 for the top ten trials). (B) In the corresponding binary representation, the presence of a spike during an event is indicated by a one (dark rectangle) and the absence by a zero. (C) The event model from which the data in panel A was obtained. The event occupation is a mixture. The event reliability for pattern 1 is indicated by the dark bars, whereas the event reliability of pattern 2 is indicated by the white-filled bars. The event model corresponding to the null hypothesis is the sum of both bars. Each binary word can be represented as a number, for four events this number is between 1 and 16 (the binary value plus one). Each binary word occurs with a certain probability, n_w , which is different for (D) the pattern model than for (E) the null-hypothesis model. In (D-E) the theoretical probabilities are represented by the gray bars, the empirical distribution corresponding to the data in panel B is indicated by the stem graphs.




Figure 3.10: The difference between the patterned spike trains and the null-hypothesis model can be determined perfectly when more than fifty trials are available. The χ^2 statistic is calculated for the difference between an empirical and theoretical distribution (see panel D and E in Figure 3.9). We determined the χ^2 distribution based on a thousand realizations of the empirical distribution for (A) twenty trials and (B) hundred trials. The data is drawn from (dark bars) the pattern model or (gray bars) the equivalent null-hypothesis model and is compared to the corresponding theoretical null-hypothesis model. In (B) pattern models are perfectly distinguishable from the null-hypothesis model because the distributions do not overlap. (C) The distinguishability is quantified using the area under the ROC curve which varies between 0.5 (not distinguishable) and 1 (perfectly distinguishable).

3.5 Finding spike patterns and determining the event structure in experimental data

We applied our event finding procedure to an experimental data set comprised of 11 experiments on 10 different cells with between 18 and 51 trials per condition (Table 3.1). These data were collected for the purpose of a different analysis, which will be published elsewhere. A current comprised of a constant current step (offset) and a fluctuating drive was injected in the soma of a layer 5 pyramidal cell. The fluctuating drive was exactly the same on each trial, but we used eleven different amplitudes (the factor by which the fluctuating drive was multiplied before it was added on top of the current step), which were expressed as percentages. Because the injected waveform was prepared off-line and stored in a file, at the time of recording we could only adapt the overall gain to the properties of the neuron we recorded from. The overall gain was chosen such that the neuron produced spikes for the lowest value of the amplitude, which was achieved for 9/10 cells.

In Figure 3.11 we show rastergrams for four example cells, with at the bottom of each panel the template of the injected current waveform (without the current step and scaling factor). The rastergram consists of blocks of constant amplitude, with the highest amplitude on top. Within each block the trials are in the order they were collected, with the earliest at the bottom. In the rastergram of Figure 3.11A, events are visible as lines of spike alignments that appear for low to intermediate amplitudes and become sharper when the amplitude is increased further. Overall this graph suggests that the precision and the reliability improve with amplitude (see Section 3.3.1 for information on calculating the reliability).

A second example is shown in Figure 3.11B. The overall behavior is similar, except that the cell shows more adaptation. For zero amplitude, this cell stopped spiking after

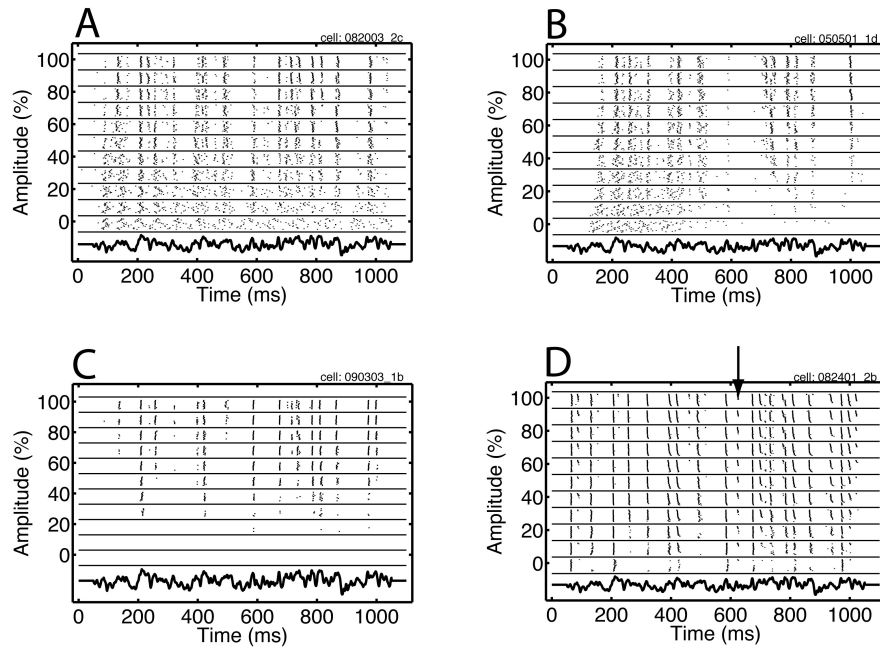


Figure 3.11: Rastergrams obtained in response to injection of the same fluctuating current waveform for four different neurons. Each line in the rastergram represents a spike train obtained on a trial, with the x ordinate of each tick being the spike time. The spike trains are ordered in blocks (delineated by horizontal lines) based on the amplitude of the injected current, expressed as a percentage, with the highest amplitude on top. Within each block the trials are ordered as they were recorded, with the earliest trial at the bottom. The injected current is shown for one amplitude at the bottom. Each panel shows the response of an example cell, the details of which are in the main text.

about 500 ms, whereas the cell in panel A continued spiking to the end of the trial. We characterized this in terms of the ratio of the mean of the first interspike interval (ISI) divided by the mean of the second ISI. This was 0.92 for panel A, and 1.1 for panel B (mean across 5 cells: 1.2). For the example in panel C, new events emerged as the amplitude increased, but the position and precision of the other events appeared not to be affected. The rastergram in panel D shows a response that did not strongly depend on amplitude. However, this example is instructive because it shows the type of nonstationary responses that were encountered in the data. Consider the time interval centered around 650 ms (indicated by the arrow), a spike was present on the later trials, but not during the earlier trials. In general we attempted to record as many trials as possible, which means that we stopped only when the cell could not maintain a stable membrane potential or input resistance between the stimulation periods. Therefore, the last few trials usually were different and were discarded prior to further analysis. The event finding algorithm is robust against the effects of nonstationarity.

We illustrate the procedure to find events using the data shown in Figure 3.11A at an amplitude of 70% and in the time segment starting at 650 ms and ending at 850 ms (Figure 3.12B). Because the events were overlapping they could not be separated using an algorithm based on an ISI threshold. First, the similarity matrix was determined (Figure 3.12B). The pixel in the i th row and j th row represents the Victor-Purpura (VP) distance between trial i and j (Victor and Purpura, 1996). The VP metric requires a temporal resolution parameter q , which represents a trade-off between reliability and precision: a pair of spikes between the two spike trains are considered different when their difference exceeds $\frac{2}{q}$. The q should be chosen such that spike patterns are most distinguishable. Our heuristic is based on how the distribution of distance values (the elements in the distance matrix) varies as a function of q . We utilize the entropy obtained by binning the distribution in 200 bins (Figure 3.12C, gray line)

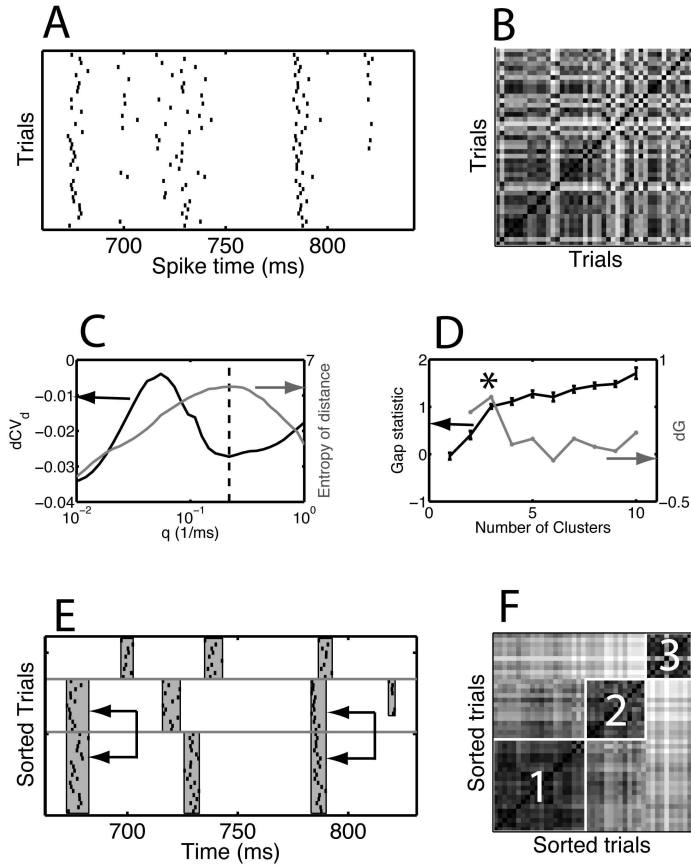


Figure 3.12: The procedure for detecting and characterizing events. (A) Rastergram of the data in Figure 3.11A for an amplitude of 70% and the time segment between 650 ms and 850 ms. (B) Distance matrix for the trials shown in panel A. The pixel on the i th row and j th column represents the Victor-Purpura distance between the spike train on trial i and the one on trial j using the selected value for the temporal resolution parameter q . The range from small to large distances is mapped onto a gray scale going from dark to white. (C) The heuristic for determining q is based on the distribution of the elements of the distance matrix, referred to as the distances. We show (black curve, left-hand-side scale) the dCV_d and (gray curve, right-hand-side scale) the entropy of the distances as a function of q . The coefficient of variation was calculated as the ratio of the standard deviation of the distances over their mean at q -values whose \log_{10} values were uniformly distributed. The dCV_d is the difference between consecutive CV_d values and thus corresponds to a logarithmic derivative. The entropy is obtained by binning the distances and using the standard $p \log(p)$ expression (see Section 3.3.6 - **Calculation of entropy and mutual information between classifications**). The q value chosen by the heuristic is the mean of the q -value at which the entropy has a maximum and the location of the deepest trough in the dCV_d that occurs after the highest peak. (continued) ...

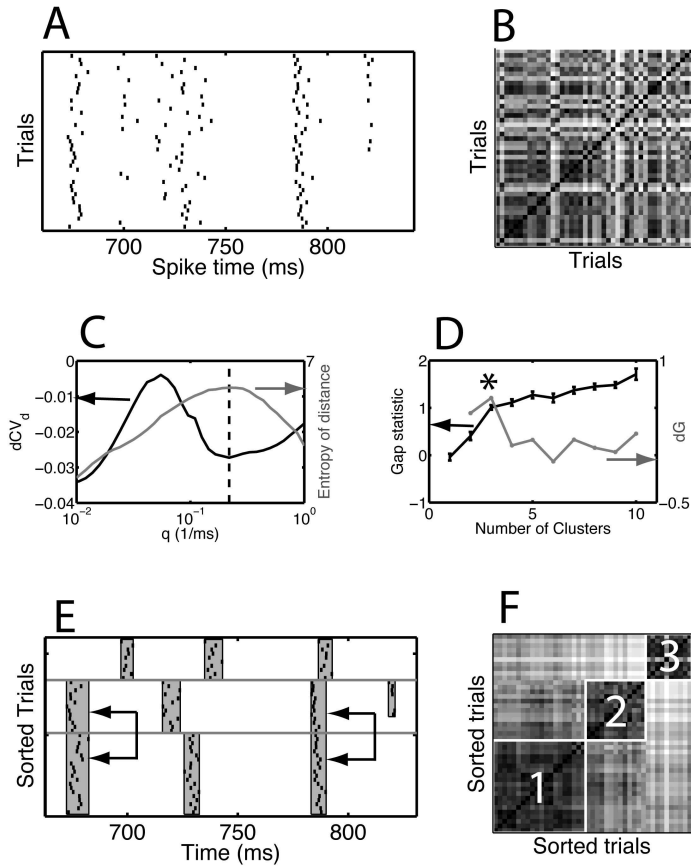


Figure 3.13: (3.12 continued) ... (D) The heuristic for determining the number of spike patterns. The FCM algorithm is used to find N_c clusters in the data, after which the gap-statistic $G(N_c)$ is computed. Each cluster is hypothesized to be a spike pattern. The gap-statistic measures the reduction of within cluster variance achieved by clustering relative to a similarly clustered surrogate data set with points uniformly distributed in the hypercube spanned by the range of the original data. The derivative of the G is $dG(N_c) = G(N_c) - G(N_c - 1)$. We show the (black curve, left-hand-side scale) $G(N_c)$ and (gray curve, right-hand-side scale) $dG(N_c)$. The errorbars are the standard deviation obtained across 50 surrogates. The value of N_c chosen by the heuristic is the one for which the dG is maximal (asterisk). (E) Rastergram and (F) distance matrices with the trials reordered according to their cluster membership. The horizontal lines in (E) separate the clusters and the vertical gray bands are the events. On two instances events common to more than one cluster were merged as indicated by the arrows.

and the logarithmic derivative of the coefficient of variation of the distances (dCV_d). The CV_d is the standard deviation of the distance divided by its mean. The log-derivative is approximated by differencing the CV at logarithmically-spaced q values, this corresponds to taking the derivative with respect to $\log_{10} q$ and is referred to as the dCV_d (Figure 3.12C, black line). For small q , the distance between two spike trains is equal to the difference in spike counts, which means it can only take close-to-integer values yielding a low value for the entropy. For large q , each spike in the two spike trains is considered different. Hence, the distance is again close-to-integer and of low entropy because it is the sum of spike counts. For intermediate q values there is a broad peak, for which spike patterns are distinguishable. The dCV_d is more sensitive to the specific q value as indicated by the peak and troughs, each of which reflects a change in the spike-pattern feature the measure is sensitive to. We found that the most robust estimate was produced by taking the mean of these two values (dashed vertical line in Figure 3.12C).

The second step was to cluster the distance matrix at the chosen q value. Each column was interpreted as a vector in a high-dimensional space. For computational efficiency, the dimensionality was reduced by applying principal component analysis and retaining only the 10 components with the highest variance. (We found 10 components to be adequate for our dataset, but for different data sets a procedure is necessary to pick the number of components based on the desired fraction of the data variance to be retained in the components (Jolliffe, 2002)). These data were then analyzed using the fuzzy c-means (FCM) procedure with the fuzzifier parameter kept at its default value of 2. This procedure requires the number of clusters to be specified, which we determined using the differenced gap-statistic dG (Tibshirani et al., 2001), which is a function of the number of clusters. The gap-statistic measures the improvement of the within cluster variance relative to data, clustered by the same method, in which there are no clusters

because the data points fill the space uniformly within the range of the original data. When used on data for which the gap-statistic was developed, the gap-statistic had a peak at the actual number of clusters. For our data the gap-statistic initially increased quickly with the number of clusters, followed by a smaller rate-of-increase. We found that the number of clusters reached after the largest increase, which is indicated by the peak in the differenced gap-statistic (asterisk in Figure 3.12D), in most cases led to the same number of clusters that would be selected manually. This heuristic indicated that there were 3 clusters in the data (Figure 3.12F). In Figure 3.12E, the clusters (spike patterns) are separated by vertical lines. Within each cluster, the events are well separated so they can easily found by the ISI-based method (Tiesinga et al., 2002).

Across clusters there are events with similar spike times, which need to be merged (example: first spike time in cluster 1 and 2 in Figure 3.12E). An ROC analysis is used to determine how distinguishable the two groups of spikes are. This is a standard procedure to evaluate the results of the two-alternative-force-choice task (Green and Swets, 1966). The ROC was rescaled to yield the sROC (see Subsection 3.4.5 - **Merging events common across multiple patterns**), which takes values between 0 and 1. If the two groups are non-overlapping the sROC is equal to 1. The merging procedure is characterized by a threshold for sROC below which groups of events are considered the same. The merged events in Figure 3.12E, indicated by arrows, were obtained using a threshold of 0.5.

The events visible in the multi-amplitude rastergrams (Figure 3.11) persist across amplitudes. This means that the precision, reliability and mean spike time of events can be compared across amplitudes. One strategy would be to use procedure outlined in Figure 3.12 to find events for each amplitude separately, after which common events across amplitudes are merged. In Figure 3.14 we show the results of an alternative

strategy, where the five highest amplitudes are analyzed at the same time. This approach is better because the clustering procedure works better when there are more trials in each cluster. Our analysis revealed the presence of four clusters (Figure 3.14 C, D), which led to 8 events, some of which were common to multiple patterns.

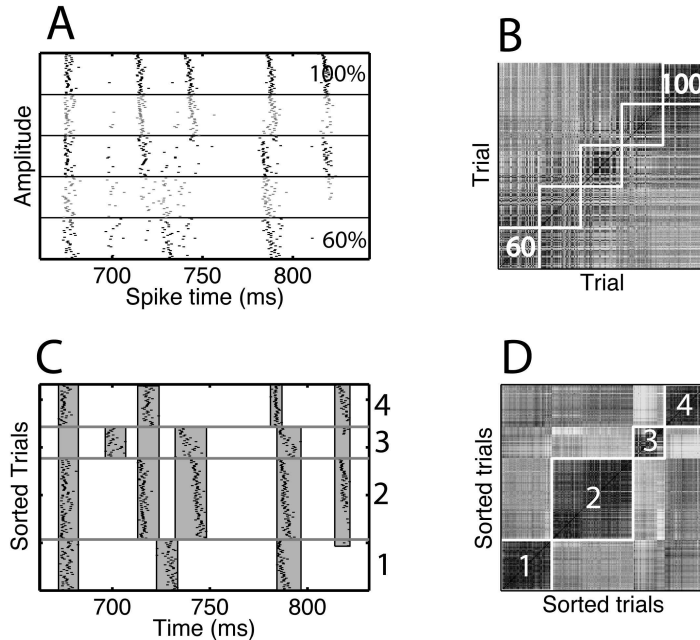


Figure 3.14: Spike patterns persisted across multiple amplitudes. (A) The rastergram for the data shown in Figure 3.11A for amplitudes between 60% and 100% and during the time segment between 650 ms and 850 ms. (B) The corresponding distance matrix obtained for the q values selected by the heuristic. (C, D) The gap-statistic suggested that there were four patterns. We show the (C) rastergram and (D) distance matrix with the trials sorted according to their cluster membership. The numbers are the cluster index. The gray vertical bands show the detected events that remained after applying the merging procedure.

We emphasize that we used these experimental data sets for the sole purpose of illustrating the analysis procedure. Our goal is to use this method in a future study to determine how the precision and reliability, calculated as independent variables, vary with the amplitude of the injected current waveform. Within the context of this goal we do find that the procedure is less efficient for the lower amplitude data sets, where the precision is reduced and the response is more sensitive to drifts, which was predicted

Table 3.1: Experimental Information:

Data set	Number of trials	Normalized amplitudes	Max amplitude	Offset	Online gain (nA)
050501_1c	18	0:10%:100%	0.1	0.03	0.5
050501_1d	51	0:10%:100%	0.1	0.15	0.5
082003_2c	45	0:10%:100%	0.1	0.15	0.8
082503_3b	41	0:10%:100%	0.1	0.15	1.0
090203_1b	21	0:10%:100%	0.1	0.15	0.9
090303_1b	29	0:10%:100%	0.1	0.15	0.9
091803_1c	40	0:10%:100%	0.1	0.15	2.2
082301_1c	22	20%:8%:100%	0.1	0.15	0.4
082401_2b	18	20%:8%:100%	0.1	0.15	1.0
012007_2b	36	0:10%:100%	0.1	0.15	5
012607_1d	40	0:10%:100%	0.1	0.15	4

based on the effect of reduced precision in the simulated data sets.

3.6 Discussion

We developed a four-step procedure with a few parameters to determine the event structure of a set of spike trains. This procedure is broadly applicable. The histogram needs to be peaked, which is indicative of events, and if there are overlapping peaks, they should be due to multiple spike patterns. The procedure was tested using artificial data with embedded spike patterns. It was also applied to spike trains recorded in response to the same fluctuating current injected across multiple trials. The method uncovered evidence for spike patterns in these data, the relevance of which is discussed below. The procedure can also be applied to groups of neurons, for instance a set of inhibitory and excitatory neurons. It can then be used to separate the inhibitory and excitatory responses and determine the precision and relative lag, which is the subject of experimental and theoretical studies (Buia and Tiesinga, 2006; Mishra et al., 2006; Womelsdorf et al., 2007).

Our results provide evidence for within-trial correlations (spike patterns) and thereby confirm previous results obtained in various experimental preparations. For instance, reverse correlation experiments were conducted to determine which stimulus features generate spikes. For experiments where the stimulus was a current injected at the soma they usually consisted of a hyperpolarization followed by a depolarization. The depth and duration of the hyperpolarization depended on when the preceding spike occurred (Jolivet et al., 2004; Powers et al., 2005), which means that the reverse correlation reflects both the spiking history as well as the stimulus features. These effects were effectively captured using leaky integrate-and-fire (Pillow et al., 2005) or spike-response model neurons (Jolivet et al., 2004, 2006). When the spike history and adaptation effects were incorporated explicitly, these models could account for the precision and reliability of in vitro neurons driven by current injection at the soma. In addition, with the proper stimulus filter, these spiking models could also account for the response of retinal ganglion cells driven by visual stimuli (Keat et al., 2001; Pillow et al., 2005). These experiments show evidence for the same type of spike patterns as analyzed here (see Figure 8 in (Pillow et al., 2005)) and in our previous work (Fellous et al., 2004; Tiesinga and Tjauw, 2005).

Although the type of dynamics exposed here can be captured in terms of relatively simple models, the analysis method presented here has a number of advantages, because it does not depend on a fitting procedure, nor does it require knowledge of the explicit stimulus waveform other than its starting time. The analysis method requires less data than would be necessary to fit a model. This is an advantage for in vivo data obtained from awake, behaving subjects, where the cell is only held for a limited amount of time and there is no direct control of the current drive at the spike generator. In vivo data also shows that the same neuron can change its spiking dynamics from regular spiking to intrinsically bursting (Steriade, 2004), which is significant because it had

been thought that these two spiking dynamics represent two different types of neurons (McCormick et al., 1985). A consequence of this result is that the parameters of the fit will change depending on cortical state, which is where the power of the clustering method becomes clear. It will be able to detect the difference in dynamics and place the spike trains in different clusters.

Recent studies as well as this study, show that neurons can become synchronized to common input (de la Rocha et al., 2007; Ermentrout et al., 2008; Markowitz et al., 2008; Tiesinga et al., 2008; Tiesinga and Toups, 2005). These neurons form an ensemble, whose collective postsynaptic impact is enhanced because of synchrony (Salinas and Sejnowski, 2000, 2001). In our data, the most precise and reliable responses are obtained when the neurons produce one pattern, which is at the highest value for the amplitude. Hence, this interpretation suggests that the single pattern state is best from the viewpoint of information transmission. However, what state provides the most information about the temporal fluctuations in the stimulus waveform? Reverse correlation studies show that spikes are most often elicited in response to a transient depolarization (Powers et al., 2005), that is: peaks in the stimulus waveform. When because of the spike rate, which is determined by the afterhyperpolarization currents, the neuron produces only a spike on half of the stimulus peaks, information about the stimulus waveform could be lost. However, when there are two spike patterns, each of which spikes on half of these peaks, all the peaks are represented in the ensemble discharge. In the next Chapter, we show using this analysis that this can happen at intermediate amplitudes, with the precision remaining high. The multiple spike pattern case thus represents an optimal compromise with a relatively high precision for information transmission and a good coverage of the temporal features of the stimulus waveform.

Chapter 4

Spike Patterns Encode Amplitude or Offset Optimally Around Response Bifurcations¹

4.1 Abstract

We characterized how the trial-to-trial variability of spike trains in response to a somatically injected aperiodic current varied with the amplitude and pedestal (offset) of the stimulus waveform. Under these conditions, the neuron produced precise and reliable spike trains. The spike times reflected specific features of the stimulus and they were robust against changes in amplitude and offset, except at points termed bifurcation points (BP), because there large shifts in the spike times were obtained in response to small changes in the stimulus. We applied an unsupervised method to further analyze the response at BPs, which revealed evidence for multiple spike patterns.

When the response obtained across trials is interpreted as the response of an ensemble of similar neurons, precise spike times lead to synchronous volleys that are effective

¹with Jean-Marc Fellous, Peter Thomas, Terrence J. Sejnowski and Paul Tiesinga

in driving postsynaptic neurons. The gain and sensitivity of a neuron is modulated by the state of the network it is embedded in and can be represented as changes in the amplitude and offset of the injected current. The *in vitro* experiment predicts that the information about the stimulus time course that is encoded by spike volleys is preserved across changes in gain and sensitivity, whereas the strength of the volleys is modulated by network state. Although the reliability measured through conventional means is reduced at BPs, the information about the stimulus time course is increased because the volleys in each pattern contain independent information, which is faithfully transmitted to postsynaptic neurons due to dendritic nonlinearities.

Taken together, the results predict a specific signature of multi-neuron discharge *in vivo* due to BPs, which can be uncovered using the analysis methods applied here.

4.2 Introduction

Neural activity recorded *in vivo* is often analyzed in terms of the peristimulus time histogram, which is a measure of the slow (50-500 ms) increases or decreases in firing rate in response to stimulus onset (Richmond et al., 1987). Recordings at the sensory periphery, such as in the retina or the lateral geniculate nucleus (LGN), indicate that the spiking response can be tightly locked to stimulus features with a temporal resolution as high as 1 ms (Butts et al., 2007; Reinagel et al., 1999; Reinagel and Reid, 2000, 2002). Therefore, it is expected that cortical neurons also respond precisely, because they receive inputs from this precise population of neurons (Kara et al., 2002; Kara and Reid, 2003; Kara et al., 2000). Nevertheless, only a few experiments have reported precise stimulus-locked responses (Bair and Koch, 1996; Buracas et al., 1998; Elhilali et al., 2004). A hypothesis is that internally generated temporally coherent synaptic inputs, which are uncorrelated with stimulus onset, reduce the precision of spikes relative to stimulus onset (Tiesinga et al., 2008). According to this hypothesis, cortical

neurons *in vivo* can respond as precisely as measured *in vitro*. There have to be many cortical neurons that receive common inputs, because there are 10-100 times more neurons in the primary visual cortex than there are in the retina or LGN (Churchland and Sejnowski, 1996), and each spiny stellate cell receives inputs from tens of LGN cells (Ahmed et al., 1994) with a similar situation in barrel cortex (Bruno and Sakmann, 2006). This group of neurons with common input is referred to as a neural ensemble, which may overlap with the anatomical concept of a minicolumn/microcolumn comprised of neurons with approximately the same stimulus preferences and response properties (Horton and Adams, 2005; Hubel and Wiesel, 1962; Lubke and Feldmeyer, 2007; Mountcastle, 1997).

We performed *in vitro* experiments to determine how information about the time-course of the stimulus is encoded in the spike times and how it is modulated by the strength of the stimulus. We interpret the results of these experiments in order to shed light on the role of spike timing *in vivo* (de la Rocha et al., 2007; Ermentrout et al., 2008; Markowitz et al., 2008; Reyes, 2003; Tiesinga et al., 2008; Tiesinga and Troup, 2005). The feedforward synaptic inputs a neuron receives *in vivo* are represented *in vitro* by an aperiodic fluctuating current injected at the soma. The *in vivo* neuron also receives recurrent inputs from the cortical network it is embedded in, which are modulated by top-down inputs (Murayama et al., 2009) such as those representing the effect of attention (Desimone and Duncan, 1995; Reynolds and Chelazzi, 2004). We approximately represent these effects *in vitro* as a change in amplitude and offset of the stimulus waveform.

The response obtained across multiple trials for different amplitudes or offset is interpreted as the response of a neural ensemble. According to this interpretation precise spike times (across trials) result in synchronous volleys of the neural ensemble (Ly and Doiron, 2009; Tiesinga and Troup, 2005), which are effective in driving postsynaptic

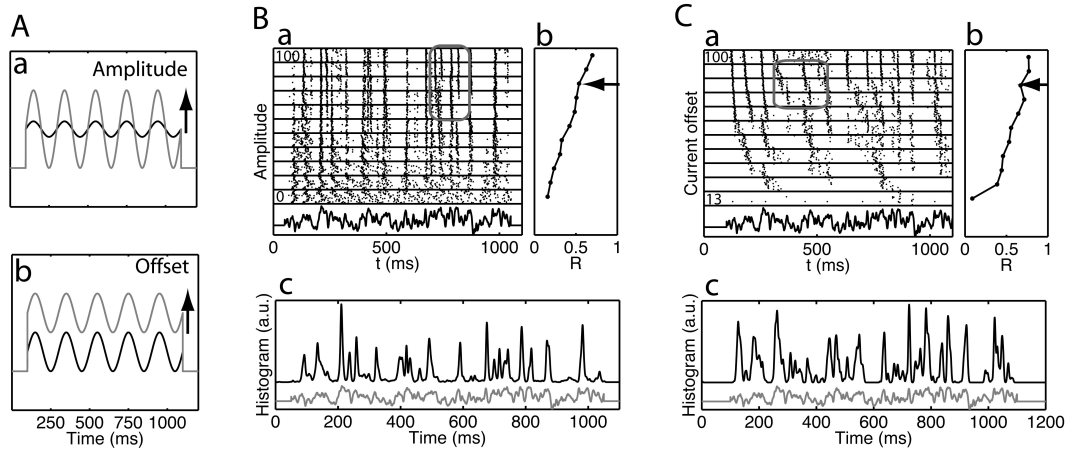


Figure 4.1: Spike timing in response to a fluctuating current is robust against changes in amplitude and offset. We show the response of two Layer 5 pyramidal cells in a slice preparation of rat prefrontal cortex, in (B) the amplitude (Aa) of the fluctuating current was varied, whereas in (C) the current offset was varied (Ab). In each panel, we plot (a) the rastergram, (b) the R-reliability (Schreiber measure with $\sigma=3$ ms, see Methods) versus amplitude or current offset and (c) the spike time histogram across all values of either the amplitude or offset. The stimulus waveform is shown for reference at the bottom of sub-panel (a) and (c). Each line in the rastergram represents a spike train obtained on a trial, with the x ordinate of each tick being the spike time. The spike trains are ordered in blocks (delineated by horizontal lines) based on the amplitude or offset of the injected current, expressed as a percentage, with the highest amplitude or offset on top. Within each block the trials are ordered as they were recorded, with the earliest trial at the bottom. The arrow in sub-panel b indicates the dip in the R-reliability, which is related to the spike train dynamics highlighted by the corresponding gray box in sub-panel a. This behavior is related to the presence of so-called bifurcation points.

cells both *in vivo* (Bruno and Sakmann, 2006) and *in vitro* (Gasparini and Magee, 2006). The strength of a volley depends on the number of cells in the ensemble that produce a spike, which is related to the reliability across trials, and their degree of synchrony, which is related to the precision. The postsynaptic effect of reliability is different than that of precision. However these effects are hard to characterize, because conventional measures of reliability represent a combination of precision and reliability (Mainen and Sejnowski, 1995; Schreiber et al., 2004). To address this issue we developed an event-based analysis, using which the across-trial spike trains are represented as a set of events. The events can then be characterized in terms of their occurrence time, precision and reliability and it can be determined how these quantities varied with amplitude and offset. In the following we will use precision and jitter interchangeably to refer to the temporal resolution of spike times. Numerically, the precision is equal to one over the jitter and it is expressed in 1/ms.

We report on three key results. First, we find that spike trains change with amplitude in such a way that the information about the stimulus time course is preserved, which also means some information about the amplitude is only reflected in the trial-to-trial variability and thus needs to be reconstructed based on multiple trials. Second, the general behavior as a function of drive parameter (amplitude or offset current) can be characterized using the concept of spike patterns and bifurcation points (BPs). Spike patterns are within-trial spike correlations, which are due to afterhyperpolarization currents (Powers et al., 2005) and other slower currents activated by action potentials. At BPs the spike times change rapidly in response to a small change in parameter value. Multiple spike patterns diversity are in some sense optimal for representing information about the stimulus time course. We find that the diversity is highest for the intermediate amplitudes used in our experiments, which is a consequence of the

presence of BPs. Third, we utilize an event-based analysis that makes it easy to automatically analyze spike train data. The main advantage of the event-based analysis is that it does not rely on fitting a specific parametric model for the neural dynamics given the stimulus (Jolivet et al., 2004, 2006; Keat et al., 2001; Pillow et al., 2005), rather it models the data directly.

4.3 Methods and Experimental procedures

General experimental procedures and stimulus generation and experimental design were the same as those described in Sections 3.3.2 and 3.3, respectively.

4.3.1 General analysis procedures

Spike times were detected from recorded voltage traces as the time the membrane potential crossed 0 mV from below. The firing rate was the number of spikes recorded during a trial, averaged across all similar trials and normalized by the duration of the trial in seconds.

In the rastergram, each row represented a spike train from a different trial. Each spike is represented as a tick or a dot, with the x-ordinate being the spike time and the y-ordinate being the trial number. Often we group trials together based on the stimulus amplitude or re-order trials based on which pattern they belong too. This is indicated in the corresponding figure caption.

The spike time histogram is an estimate for the time-varying firing rate. It was obtained by dividing the time range of a trial into bins (typically 1 or 2 ms wide) and counting the number of spikes that fell in each bin across all trials. The bin count was normalized by the number of trials and by the bin width in seconds. The latter was to assure that a bin entry had the dimensions of a firing rate, Hz. The histogram was

subsequently smoothed by a Gaussian filter with a standard deviation equal to 1 bin size. The spike-triggered average (STA) was obtained by for each spike selecting the 25 ms stimulus segment prior to the spike and average across all spikes.

Events were detected using the procedure summarized below. At the end of this procedure, all spikes were either assigned to an event or were classified as noise. The event-reliability is the fraction of trials on which a spike was observed during that event, the event-jitter is the standard deviation of the spike times belonging to the event. The event-precision is the inverse of the event-jitter. For a given condition (amplitude, offset or initial current step) the reliability, precision and jitter are defined as the event-reliability, event-jitter and event-precision averaged across all events.

The R-reliability is calculated based on all spike times without detecting events. The spike trains are first transformed into a continuous waveform, where each spike is convolved with a Gaussian distribution with a standard deviation σ (Schreiber et al., 2003; Wiskott et al., 1997). The cosine of the angle between the two waveforms, when considered as vectors, is computed as the inner product between the waveforms of trial i and j , normalized by the product of the square roots of the inner product of each trial with itself. This quantity is a number between 0 and 1 (the waveforms are positive valued) and is called the similarity S_{ij} . The reliability estimate R is the mean of S_{ij} across all distinct pairs (i, j) . Intuitively, the inner product measures the degree of overlap between spike times: the closer two spike times are, the larger the overlap and thus their contribution to the inner product. The parameter σ sets the time scale of the reliability measure and determines which spike times between the pairs are considered overlapping. For σ approaching zero, all spike times are considered different (except when the spike trains are identical to machine precision), hence $R=0$. For σ much larger than the trial length, all spikes overlap and $R=1$. In the first case, the emphasis is on precise spike times; in the second case, the emphasis is on the global amount of

spikes (spike rate). We used a more efficient method for calculating R by summing, for each pair of trials (i, j) , the following expression across all spike pairs (k, l) that are separated by less than six σ s, $S_{ij} \propto \sum_{k,l} e^{-(t_k^i - t_l^j)^2 / 4\sigma^2}$ (here t_k^i is the k th spike on the i th trial and for simplicity the normalization was omitted, see (Kruskal et al., 2007) for details).

4.3.2 Calculation of the VP distance

Refer to Section 3.3.4.

4.3.3 Overview and goal of the event finding method

Refer to Section 3.4.1

4.3.4 Calculation of entropy and mutual information between classifications

Refer to Section 3.3.6.

4.3.5 Simulation experiments

The neuron was modeled as a single compartment with Hodgkin-Huxley-type voltage-gated sodium and potassium currents and a passive leak current (Tiesinga and Jose, 2000; Wang and Buzsaki, 1996). The equation for the membrane potential of the model neuron is:

$$C_m \frac{dV}{dt} = -I_{Na} - I_k - I_L + I_{inj} + C_m \xi \quad (4.1)$$

where $I_L = g_L(V - E_L)$ is the leak current, $I_{Na} = g_{Na}m_\infty^3h(V - E_{Na})$ is the sodium current, $I_K = g_Kn^4(V - E_K)$ is the potassium current, I_{inj} is the injected fluctuating current, which is the same on each trial and $C_m\xi$ is a noise current that is different on each trial. The values for the maximum conductance and reversal potential are listed in Table 5.1. The gating variables are m , n , and h and they satisfy the equation

$$\frac{dx}{dt} = \zeta(\alpha_x(1 - x)\beta_x x).$$

Here the label x stands for the kinetic variable, and $\zeta = 5$ is a dimensionless time scale that can be used to tune the temperature dependent speed with which the channels open or close. The rate constants are:

$$\begin{aligned} \alpha_m &= \frac{-0.1(V + 35)}{e^{-0.1(V+35)} - 1}, & \beta_m &= 4e^{-(V+60)/18} \\ \alpha_h &= 0.07e^{-(V+58)/20}, & \beta_h &= \frac{1}{e^{-0.1(V+28)} + 1} \\ \alpha_n &= \frac{-0.1(V + 34)}{e^{-0.1(V+34)} - 1}, & \beta_n &= 0.125e^{-(V+44)/80}. \end{aligned}$$

and the asymptotic values of the gating variables are:

$$x_\infty(V) = \frac{\alpha_x}{\alpha_x + \beta_x}$$

where x stands for m , n , or h . We made the approximation that m follows the asymptotic value $m_\infty(V)$ instantaneously (Wang and Buzsaki, 1996). The noise ξ_i in the current of neuron i is chosen such that $\langle \xi_i(t) \rangle = 0$ and $\langle \xi_i(t)\xi_i(t') \rangle = 2\lambda\delta(t - t')\delta_{ij}$.

On each integration time step, the noise was drawn independently from a uniform distribution between $-12\lambda/dt$ and $12\lambda/dt$, where dt was the time step. For Figure 4.5 we used $\lambda = 0.00025 \text{ mV}^2/\text{ms}$, whereas in Figure 4.7 we used $\lambda = 0.0001 \text{ mV}^2/\text{ms}$ (low noise) and $\lambda = 0.025 \text{ mV}^2/\text{ms}$ (medium noise). For I_{inj} we either used the same 1050 ms long fluctuating drive as in the experiment (Figure 4.5, amplitude between 0 and 1, offset $0.2 \mu\text{A}/\text{cm}^2$), but without the constant depolarizing current pulses preceding the fluctuating drive in vitro; or we used a sinusoidal drive with time-varying frequency (illustrated in Figure 4.7).

The initial values of the membrane potential at the beginning of the simulation were set to a fixed value, usually -70 mV. The gating variables were set to their asymptotic stationary values, $x_\infty(V)$, corresponding to the starting value, V , of the membrane potential. The differential equations were integrated using a second-order Runge-Kutta method with a time step of $dt=0.05 \text{ ms}$ (Gerald and Wheatley, 1999; Press et al., 1992).

4.4 Results

Our goal was to determine how the amplitude and current offset of a fluctuating drive was reflected in the structure of neural spike trains. Therefore, a current comprised of a constant current step (offset) and a fluctuating drive was injected in the soma of a layer 5 pyramidal cell in a slice of rat prefrontal cortex. The fluctuating drive was exactly the same on each trial, but for the first experiment we used eleven different amplitudes (the factor by which the fluctuating drive was multiplied before it was added on top of the current step), which were expressed as percentages. We performed experiments on 8 different cells with between 18 and 51 trials per amplitude. For 8/10 cells we used 0% to 100% in steps of 10%, for 2/10 cells we used 20% to 100% in steps of 8%. Because the injected waveform was prepared off-line and stored in a file, at the time of recording we could only adapt the overall gain to the properties of the neuron we recorded from.

The overall gain (between 0.4 and 5) was chosen such that the neuron produced spikes for the lowest value of the amplitude, which was achieved for 8/10 cells.

In Figure 4.1B we show rastergrams for one example cell, with at the bottom of the panel the template of the injected current waveform (without the current step and scaling factor). The rastergram consists of blocks of constant amplitude, with the highest amplitude on top. Within each block the trials are in the order they were collected, with the earliest at the bottom. In the rastergram (Figure 4.1Ba), events are visible as lines of spike alignments that appear for low to intermediate amplitudes and become sharper when the amplitude is increased further. Overall this graph suggests that both the precision and the reliability improve with amplitude. The trial-to-trial variability was characterized using the Schreiber measure (R-reliability, see Section 4.3.1 - **General Analysis Procedures**) with a σ value of 3 ms (Schreiber et al., 2003). Simply put, spikes in two different spike trains with a time difference of less than 3 ms are considered coincident. Note that this measure does not distinguish between changes in reliability and changes in precision, for that an event-based analysis is necessary. Overall, R increased with amplitude (Figure 4.1Bb), but there were dips, one of which is indicated by the arrow. The model analysis performed below indicates that this is due to a so called bifurcation point (BP).

The spike times do not vary much with amplitude. To demonstrate this we determined the spike time histogram across all amplitudes (Figure 4.1Bc), which displays a sequence of peaks. These peaks were broader than those obtained using only the spike trains for the highest amplitude, whereas the responses for the lowest few amplitudes were not strongly stimulus-locked and were imprecise. However, for the latter, because the spike times were spread out, they contributed predominantly to the flat background rather than broadening the peaks. To interpret these results, consider a neural ensemble comprised of different neurons for which the effect of the feedforward

drive varies in effective strength. For instance, because the number of synaptic inputs with these temporal dynamics is different across neurons, or because neurons have different densities of leak and other voltage-gated channels that lead to variability in the input conductance (Koch, 1997). The in vitro experiments show that despite this diversity the ensemble would produce volleys that are effective in driving postsynaptic neurons and transmitting information about the time course of the input.

For the second set of experiments (Figure 4.1C), the current offset (before the overall gain factor) was varied between 0.05 and 0.3 nA in 10 equal steps, while the amplitude of the stimulus waveform was held constant. We presented 7 such stimulus sets to 5 different cells and recorded the responses on 17 to 36 trials, with overall gain factors between 2.7 and 3. The overall behavior was similar: precise spiking was obtained across current offsets (Figure 4.1Ca and c), the reliability measure R increased with amplitude (Figure 4.1Cb) and also displayed a dip (arrow in 4.1Cb). There were also differences. First, the firing rate increased faster than for the case where the amplitude was varied. This is because the overall level of depolarization increased, whereas for increasing amplitude not only the peaks got higher, the troughs also got deeper, which meant that some spikes would appear and other spikes would disappear. Second, for a nonzero firing rate the neurons immediately locked, which was indicated by the higher reliability (we excluded the first current offset, for which the neuron only spiked on a few trials) compared to the amplitude case. In the low-current-offset blocks the first spike times seem to drift with current offset, but they actually shift to earlier events.

The events visible in the multi-amplitude rastergrams (Figure 4.1) persist across amplitudes. This means that the precision, reliability and mean spike time of events can be compared across amplitudes (see Subsection 3.3.3). One strategy would be to find events for each amplitude, after which events common across amplitudes are merged. In Figure 3.14 we show the results of an alternative strategy, where the five

highest amplitudes are analyzed at the same time. This approach is better because the underlying clustering procedure works better when there are more trials in each cluster. Our analysis revealed the presence of four clusters (Figure 3.14C, D), which led to 8 events, some of which were common to multiple patterns.

Because there were more amplitudes (five) in the data set than there were patterns (four), a given pattern had to persist across multiple amplitudes. To further study this we re-plotted the rastergram in blocks of constant amplitude, but within each block the trials were ordered based on the pattern they expressed (Figure 4.2A). The pattern composition which patterns were present on what fraction of the trials varied across amplitude and was quantified in Figure 4.2B. When the presence of the second pattern (Figure 4.2B, line 2) increased, the fraction on trials on which the first pattern was present decreased (Figure 4.2B, line 1). Hence, the reliability of events in the second pattern increased (black ellipses in Figure 4.2A), whereas the reliability of events in the first pattern decreased (gray ellipses in Figure 4.2A). Hence, changes in pattern occupation account for the, at first glance counterintuitive, decrease in reliability with amplitude.

The diversity of patterns present for a given amplitude is quantified using the entropy of the pattern distribution (Figure 4.2C). For this data, the entropy was maximal at a specific amplitude (see Section 3.3.6, arrow in Figure 4.2C). Across data sets and segments a decrease in entropy with increasing amplitude was the more robust observation, because for higher amplitudes only one pattern survives. This behavior is relevant for the amount of information that can be extracted about the temporal dynamics of the injected current and is discussed below.

During extracellular recordings in vivo the neurons membrane potential is not accessible, which is unfortunate because it would provide information about the dynamics of excitatory and inhibitory inputs to the neuron (Monier et al., 2003). An important

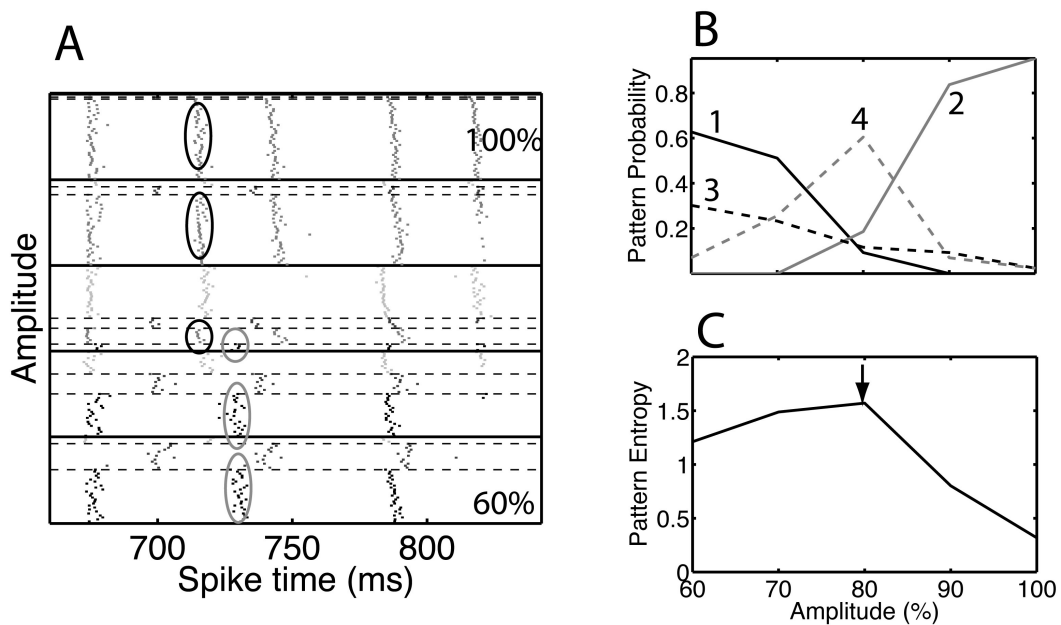


Figure 4.2: Decreasing spike time reliability was associated with changes in the pattern occupation. (A) Rastergram of the clustered data shown in Figure 3.14A. Each block (separated by thick black lines) corresponds to a different amplitude, with the lowest amplitude at the bottom and the highest amplitude at the top. Within each block, the trials are ordered based on their cluster membership. The clusters are separated by thin dashed lines. Two events are highlighted: the ones in the black ellipses, whose reliability increased with amplitude and the ones in the gray ellipses, whose reliability decreased with amplitude. (B) The pattern occupation (or probability) for a given amplitude is the fraction of trials on which that pattern is obtained. We show the pattern occupation as a function of amplitude for the four patterns that were detected, as indicated by the numbers in the graph. (C) The diversity of patterns observed for a given value of the amplitude is quantified as the entropy of the pattern distribution. The entropy as a function of amplitude has a peak at 80% (arrow), indicating that the pattern diversity is largest for that amplitude.

question is what aspects of the membrane potentials are reflected in, or can be reconstructed from, the detected spike times. To address this issue we analyzed recordings where on each trial the same waveform with the same amplitude was injected because for these experiments the most trials were available. For these experiments, the fluctuating waveform was extended to 1700 ms and was preceded by a constant current offset of 200 ms (in addition to 50 ms zero current at the start of all current injections). Because these datasets were taken for a different experiment, the initial current offset took eleven different values, the influence of which is discussed below. We injected this drive on nine experiments using eight cells, with between 10 and 35 trials (which means 110 to 385 trials when ignoring the current offset) and an overall gain with values between 0.9 to 4.7.

We applied the spike pattern analysis to all the segments in a representative data set. In Figure 4.4A, within each segment the trials were ordered according to the pattern they expressed in that segment. The spikes on a row of the rastergram most likely correspond to different trials, because the trial ordering on each segment was determined independently using the clustering algorithm. Patterns correspond to within-trial correlation between spike times. We determine how long this correlation persists by determining how well the pattern expressed on a trial during one time segment predicts which pattern is expressed in a preceding or following segment. In Figure 4.4B, the trials are ordered in each segment based on the pattern they were on in the last (fourth) segment (indicated by the asterisk). This panel shows that even though a group of trials expressed the same pattern during the fourth segment, the same group expressed a mixture of patterns during segment 3. The association between different segments can be expressed as the normalized mutual information between the pattern classification of a trial in two segments (I_N , see Section 3.3.6). Because it is normalized, the maximum value is one, which is obtained if the classifications are identical. The I_N

between the classification in segment 4 and that in segments 3, 2 or 1 was 0.31, 0.006, or 0.039, respectively.

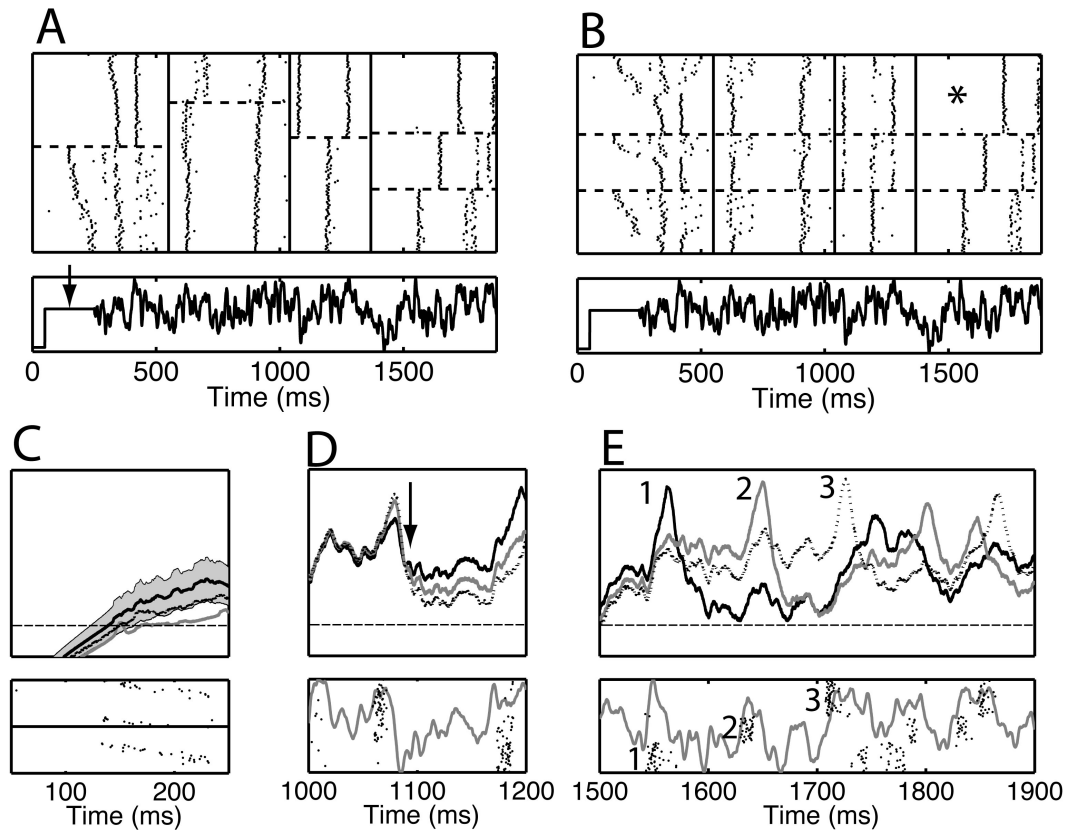


Figure 4.3: Spike patterns corresponded to voltage patterns. The single-amplitude data set was divided into four time segments. (A) Segment-by-segment rastergrams. In each segment trials were ordered according to the cluster membership in that segment. The clusters are separated by horizontal dashed lines, whereas segments are indicated by vertical lines. Because the trial order varies from segment-to-segment, spikes on the same row but in different segments are not necessarily obtained in the same trial. (B) Rastergram with trials in each segment ordered based on their cluster membership on the fourth segment (asterisk). At the bottom of A and B the current waveform is repeated for reference. There was a 200 ms long constant current step (arrow), whose amplitude took eleven different values (only one is shown). ...

We further analyzed the three patterns uncovered in the time interval between 1500 and 1900 ms. For each pattern, the voltage traces were averaged across all corresponding trials and the standard deviation was used as an estimate for the trial-to-trial variability. The mean voltage traces were not only different because the neurons spiked

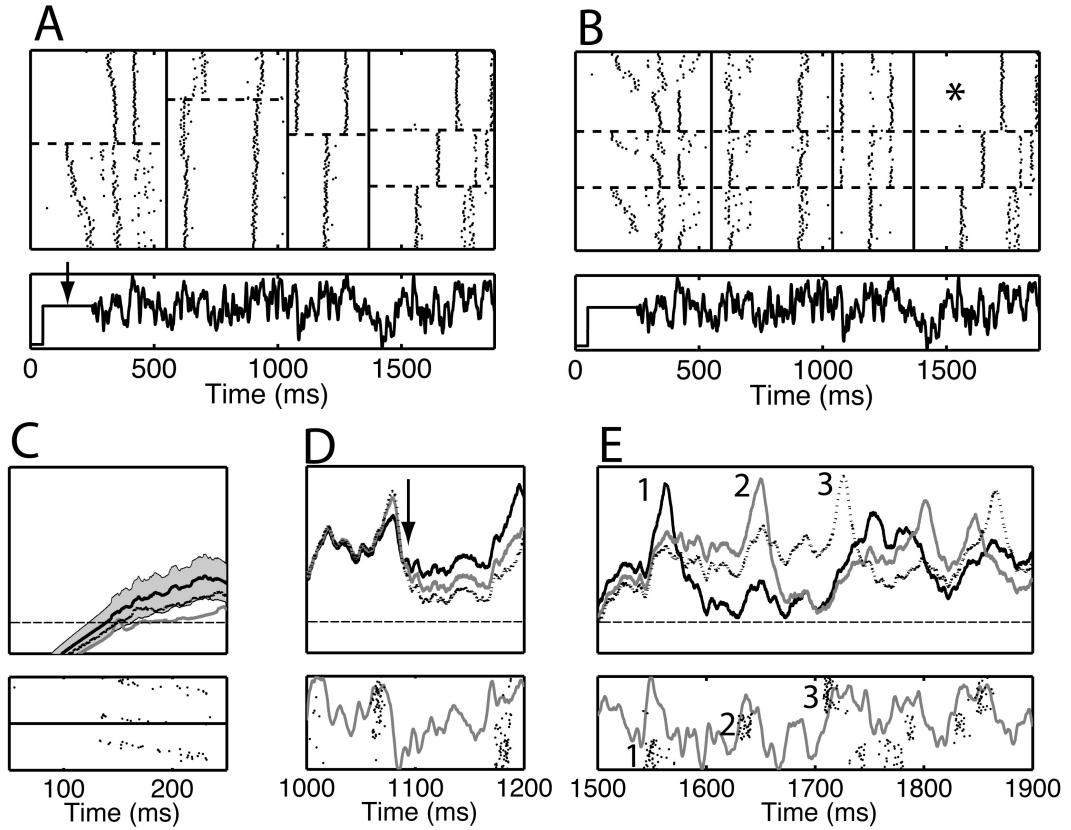


Figure 4.4: ... (continued) (C-E) The analysis procedure found 3 spike patterns in the fourth segment (between 1500 ms and 1900 ms), labeled 1 to 3. In each of the panels C-E, we show (top) the voltage traces averaged across all trials expressing that pattern (the y-axis covers the range from -65 to -35 mV) and (bottom) the current waveform together with a rastergram where the trials were ordered based on the cluster membership in the fourth segment. The spikes were shifted to the left by 12 ms so that they were approximately aligned with an upswing in the injected current. In (C) the gray bands indicate the plus or minus two standard error range for the black curve. The arrow in (D) indicates where hyperpolarization from a spike exposes differences in the membrane potential of the three different patterns.

at different times, but also because the conductances associated with the afterhyperpolarization following the spike altered the response to the current injection (Figure 4.4E, labels 1 to 3). The spikes reflected periods where the injected current had a large positive slope, but each pattern was triggered by a different subset of these upswings (Figure 4.4E, bottom). Once a spike was produced the neuron did not spike during an otherwise viable upswing shortly thereafter, even though it had produced a spike there on other trials during which it expressed an alternative pattern. For instance (Figure 4.4E, bottom), on trials labeled 1, the neuron did not spike in response to the upswing that caused the neuron to spike on trials labeled 2.

This is not the complete story because there can be correlations that persist across an entire trial, which are not visible in the membrane potential unless they are unmasked by a deep hyperpolarization. For each pattern expressed on segment 4, we averaged the corresponding voltage traces across the entire duration of the trial. For clarity, we only show three time intervals (Figure 4.4C to E). In the first interval there is a small difference in the mean membrane potential (Figure 4.4C), which has disappeared at $t=1000$ ms (Figure 4.4D), but reappears after a deep hyperpolarization (Figure 4.4D). This difference then leads to three clearly distinct voltage patterns in the last interval (Figure 4.4E). Often, whether or not a neuron spikes at a given upswing determines which pattern a neuron is on for a hundred milliseconds or more, a time scale comparable to fixation between saccades (McIlwain, 1996). Our hypothesis is that the spiking probability depends on the state of the neuron, which may not necessarily be visible in the membrane potential. The state on a given trial could be correlated with the height of the depolarizing step that preceded the fluctuating current or it could be due to nonstationarity, which means it depends on whether the trial is at the beginning or end of the experiment. The normalized mutual information was calculated between the trial number and the pattern on segment 4 (the degree of nonstationarity)

and between the height of the offset and the pattern, for which we obtained the values $I_N=0.13$ and 0.11 , respectively. This analysis shows that only a small fraction of the variability can be explained by these two factors.

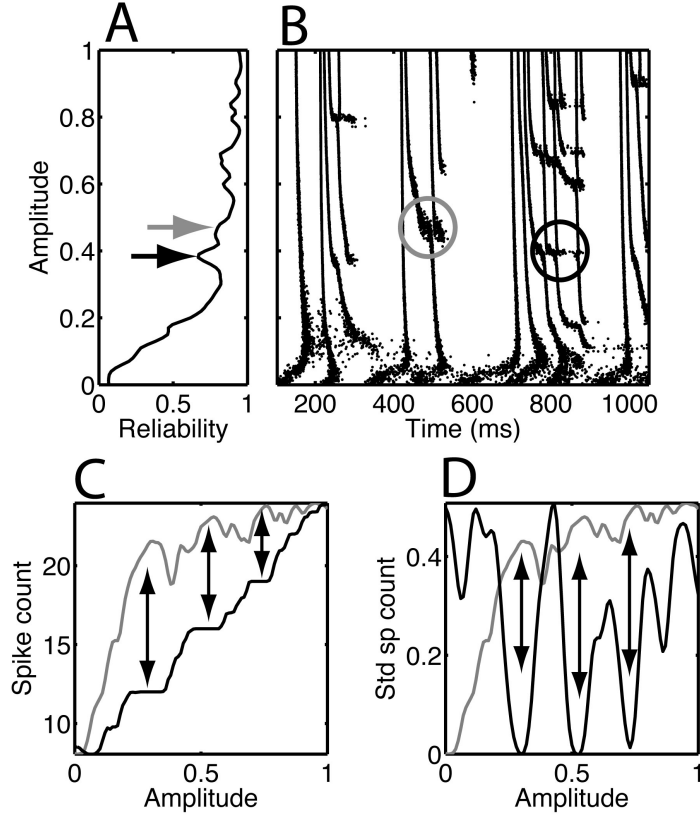


Figure 4.5: Bifurcation points were observed in model simulations for amplitudes at which the spike count changed. We show the (A) R-reliability ($\sigma=1$ ms) and (B) rastergram as a function of amplitude. The dip in reliability indicated by the black and gray arrow in (A) corresponds to the bifurcation in the black and gray circle in (B), respectively. We plot (C) mean spike count and (D) standard deviation of the spike count across trials versus the amplitude. The gray curve is the R-reliability replotted from panel A, the full range for R, 0 to 1, is represented in the graph. Peaks in the reliability, indicated by the double-headed arrows, correspond to (C) plateaus in the spike count, for which (D) the trial-to-trial variability in the spike count is small. The spike trains were obtained from simulations of the WB model neuron (Tiesinga and Toups, 2005; Wang and Buzsaki, 1996), driven by the same drive as used experimentally.

We performed simulations using the Wang-Buzsaki (WB) model neuron (Wang and Buzsaki, 1996) in order to clarify the link between spike patterns, bifurcation points and

the reduced reliability measured in terms of the R-reliability. We simulated 50 trials for each of 101 different amplitudes between 0 and 100% (the stimulus was normalized so that the mean was zero and the standard deviation was one). Panel A of Figure 4.5 shows the reliability curve, which was smoothed by a three-point running average, and panel B shows the corresponding rastergram with matching y-ordinates. For clarity we only used half of the trials and half of the amplitude values in the rastergram. The R-reliability generally increased with amplitude but it had dips. Two of the dips are highlighted by the arrows and correspond to the spike train features inside the circle of matching color in Figure 4.5B. Two spike times enter the gray circle for low amplitudes and three leave at higher amplitudes. This is a bifurcation because the dynamics changed rapidly for small changes in the parameter value. Inside the circle there are fuzzy clouds of spikes that correspond to multiple patterns for a given amplitude. As the amplitude increased, the fraction of patterns with two spikes decreased, whereas those with three spikes increased. A similar transformation takes place within the black circle, where three spikes enter and four spikes leave.

This relationship was further examined by comparing the mean spike count and the trial-to-trial variability of spike count to the modulation of the reliability (Figure 4.5C and D). Local maxima in the reliability corresponded to plateaus in the spike count, for which the variability of the spike count was minimal. A similar behavior was obtained when the current offset was varied, with the same two differences as observed in experiment. First, the increase in firing rate was higher with offset than with amplitude. Second, the reliability started out high even for low firing rates. By contrast, for zero or low amplitudes, there was nothing to lock to, and a low R-reliability similar to that in response to a current step was obtained.

Computational models give insight into what are bifurcation points and how they are generated. Consider, for instance, a leaky integrate-and-fire neuron. A constant,

above threshold current drive causes a spike train with regular interspike intervals, because the membrane potential increases over time, crosses the spike threshold, after which a spike is generated and the membrane potential is reset and the process starts over again. Previous analyses have shown that jitter is proportional to the rate of change of the membrane potential at the threshold (Cecchi et al., 2000), which is in turn proportional to the current offset. This noise stability also predicts the stability of the spike times against parameter changes. For a fluctuating drive, the spike times are generated by upswings in the current drive, yielding a specific pattern of spike times. Usually there are more upswings than there are spikes produced by the neuron. Thus, when one of the upswings is missed, spikes are produced at a new sequence of upswings a new pattern. Nevertheless the spikes in the new pattern still reflect the stimulus.

Consider the case where one upswing brings the neuron close to threshold, but when it does not actually cross it. With a small increase in amplitude or current offset, the neuron will cross threshold, and a big change in the response has occurred: a bifurcation. This sensitivity to parameter changes also implies a sensitivity to noise. On some trials the noise can induce a spike, whereas on others it prevents a spike, which means that at least two patterns are produced. For a leaky integrate-and-fire model neuron a bifurcation is discontinuous because of the discontinuous voltage reset after a spike, but for Hodgkin-Huxley type models, such as the WB model used here, it is continuous.

Near a bifurcation in the presence of noise there is, at the single neuron level, confusion about which amplitude gave rise to the observed spike train. This confusion does not exist across multiple trials or an ensemble: the fraction of trials on which the below-bifurcation spike-pattern is observed will decrease as the amplitude is increased.

Each precise event is due to, and thus represents, an upswing in the stimulus waveform. Away from BPs, there is only one spike pattern and the neuron spikes with a

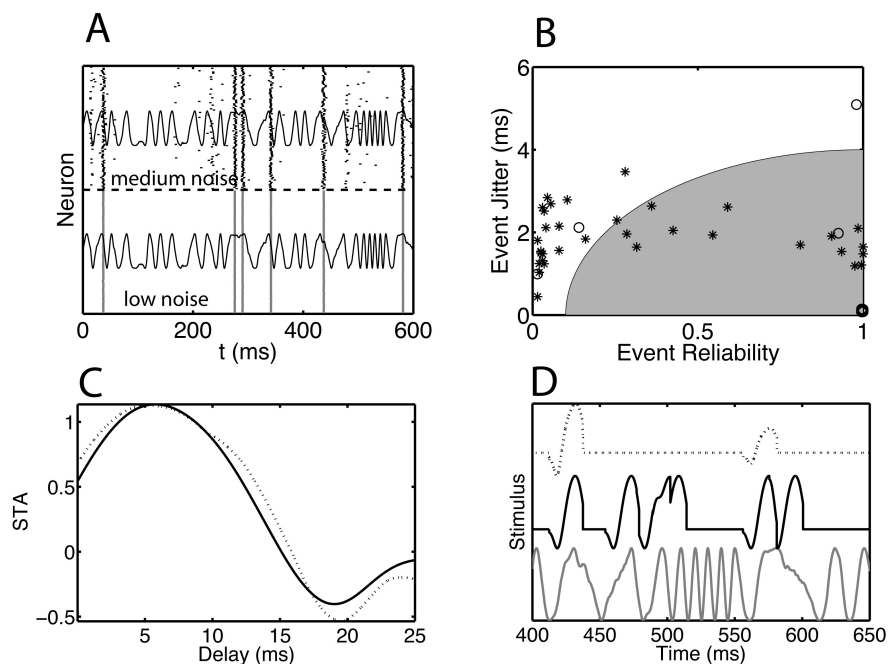


Figure 4.6: Information about the time course of the stimulus waveform is improved at bifurcation points because of the presence of multiple spike patterns. Data is from an example model neuron described in (Tiesinga and Toups, 2005; Wang and Buzsaki, 1996). (A) We show the rastergram for a short time segment across 100 trials for a (bottom) low-noise and (top) medium-noise model neuron. The noise level refers to the magnitude of a white noise current that varies from trial-to-trial relative to the amplitude of the repeated fluctuating current waveform (shown as a thin solid line on top of each rastergram). For low noise, the neuron spiked only at six events, whereas for medium noise there were additional events. (continued) ...

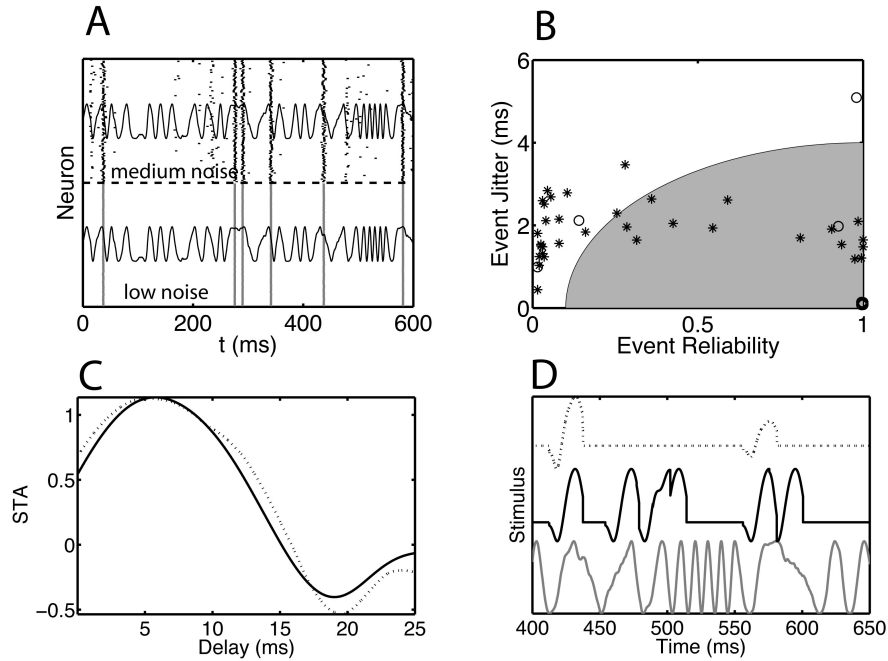


Figure 4.7: (continued) ... (B) We calculated the reliability and jitter for each event for the entire stimulus duration (1100 ms). The open circles represent the low-noise, and the asterisks represent the medium-noise result. The gray-filled region represents the combination of jitter and reliability for which a putative postsynaptic neuron would generate a spike. (C) The spike-triggered average obtained across the entire stimulus period for (solid line) the medium-noise neuron and (dotted line) the low-noise neuron. (D) The stimulus waveform reconstructed using the (dotted line) low-noise and (solid line) medium-noise spike trains was compared to (gray solid line) the actual stimulus waveform. We used an event-based reconstruction, where each extracted event contributed equally to the reconstruction regardless of reliability and jitter, as long as the reliability exceeded 5%. For clarity the three curves are offset from each other.

high precision and reliability. However, when there are more upswings than there are spikes, information about the time course of the stimulus is lost. By contrast, near a BP the dynamics is more noise sensitive and multiple spike patterns are obtained with non-overlapping event times, each of which provides information about a stimulus feature. To illustrate this point, we drove the model neuron with a sinusoidal waveform with a frequency that fluctuated over time. We found a BP and compared the response at this point under two circumstances: low noise and medium noise. In this way we could compare responses at the same amplitude with a similar, but not exactly the same, spike rate. In the low-noise case there were six events during the time interval displayed in the graph (Figure 4.7A, bottom), but for the medium-noise case there were additional events with a reduced reliability and precision (Figure 4.7A, top). We applied the event-based analysis to the entire simulation time interval (1100 ms) and determined for each event the reliability and precision (Figure 4.7B). We also determined the spike-triggered average (STA) for both cases (Figure 4.7C). Using all detected events with a reliability exceeding 5% and the measured STA, we reconstructed the stimulus waveform (Figure 4.7D). The medium-noise spike trains, with their multiple spike patterns, yielded a better reconstruction than the low-noise case with a single spike pattern.

To interpret these results, consider the response of a neuron to a volley with a specific number of spikes (reliability times the size of the ensemble) and jitter. In order for an event to contribute to a reconstruction it should be detectable by a postsynaptic neuron. For a subthreshold background drive, the volley will result in an output spike when the reliability and precision is high enough. An example of the effective range is indicated by the gray shading in Figure 4.7B. Another constraint is that each volley should generate only one output spike. That is, no matter how deep you are in the gray shaded area, only one spike should be generated. Preliminary model studies at

the single compartment level indicate that this is hard to achieve, except when there is a correlated inhibitory input that follows the excitatory volley at an appropriate delay. Recent studies using two-photon uncaging of glutamate show that dendritic action potentials may be an appropriate mechanism (Gasparini and Magee, 2006), which requires studying the impact of synchronous ensembles using multicompartment models.

4.5 Discussion

Previous studies have shown that neurons in the slice preparation produce precise and reliable spike trains in response to fluctuating currents injected at the soma (Bryant and Segundo, 1976; Mainen and Sejnowski, 1995). Furthermore, arguments have been presented that suggest neurons in vivo can respond as precise and reliable as in vitro (Tiesinga et al., 2008). This raises the possibility that information is encoded in precise spike times, that this information can be transmitted from neuron to neuron and that it is utilized by the nervous system. We addressed the following two issues by conducting in vitro experiments and simulating computational models. First, to what extent is information about the time-course of the stimulus conserved under common modes of response modulation. Second, is there an optimal operating regime for representing this information in ensembles of uncoupled, similarly tuned neurons.

We find that precise spike times – events across trials or synchronous volleys in a neural ensemble – are generated by upswings in the fluctuating drive. Because upswings remain upswings when the amplitude or the offset is changed, precise spikes will still represent information about those upswings. However, how many and which upswings are represented in the spike trains could change. We find that these changes are discrete, which means that event times change little with changing amplitude or offset, until a bifurcation point is encountered at which they suddenly change. Because

of afterhyperpolarization currents and other slower currents, a sudden change at one spike time leads to changes in subsequent spike times. Near BPs spike trains are more noise sensitive and the R-reliability is reduced. We mean by noise the component of the neurons input that varies from trial-to-trial. In the context of models this corresponds to an additional pseudorandom current, with a different seed on each trial (Tiesinga and Toups, 2005). Noise sensitivity is often considered detrimental to coding. However, here it is beneficial at the ensemble level, because near BPs multiple precise patterns are generated, each of which provides information about a different set of upswings. Furthermore, because across a range of amplitudes/offsets a range of spike patterns is obtained, each representing upswings in the stimulus, the overall spiking activity is still precise and informative.

The bifurcation structure is determined by stimulus characteristics. This is clearest when comparing the structure obtained in response to a periodic drive to that for an aperiodic drive. For the former, the density of bifurcations is highly non-uniform: there are large ranges without any bifurcations separated by a range full of bifurcations (see Figure 3A in Tiesinga (2004) for an example). The bifurcation cause changes in the spike trains for the entire time range. In contrast, for an aperiodic drive, for a given parameter, there is almost surely a bifurcation, because a random drive will come close to threshold at some time during the infinite stimulus presentation (see Figure 4.5B). However, a large hyperpolarization or depolarization can often reset the neuron. Because such events are expected quite regularly during *in vivo* dynamics (Haider and McCormick, 2009), the bifurcation structure should only be considered for finite intervals, just as studied here experimentally. Taken together, this implies that dynamics at BPs are likely to be relevant *in vivo*.

4.5.1 Implications for cortical coding

In experiment the response of a single neuron, uncoupled from other neurons, was measured across multiple trials. However, the overall goal was to shed light on the dynamics of neural ensembles in vivo. This raises two issues discussed below: what is the nature of the neural ensemble and how is the information represented in neural ensembles?

There are two alternative hypotheses about how information is represented in an ensemble. The first hypothesis holds that slow modulations (50-500 ms) in the firing rate code for stimulus properties. This firing rate can be estimated by averaging across all neurons in the ensemble, from which the time course of the stimulus can be reconstructed. The second hypothesis is that ensembles of neurons produce precisely timed spikes, which lead to synchronous volleys that are effective in driving postsynaptic neurons. Experimental results, reviewed in (Tiesinga et al., 2008), suggest that a combination of these two strategies is at work at the level of the cortex. First, volleys represent the fast fluctuations in the inputs (either sensory related or internally generated). Second, slow modulations change the number of volleys and the number of spikes per volley. This does not change the information about the temporal fluctuations, but alters the strength with which they are expressed.

Physiological studies provide support for this idea, because they show that spatially clustered, temporally precise synaptic inputs are very effective at eliciting reliable and precise spikes via dendritic action potentials (Gasparini and Magee, 2006), in a way which is decoupled from the spikes generated by slow input modulations. Furthermore, this response has an all-or-none character. If, for instance, 50 spikes in a 2 ms long interval are enough to elicit an output spike, then increasing the number of spikes or their precision will generally not increase the number of spikes produced in response to this volley. Any volley meeting the minimum requirement will elicit a spike and

will be able to transmit information about the upswing that generated it. Hence, when there are multiple precise spike patterns, the ensemble provides more information than a single pattern, despite the reduced reliability at the single neuron level. This would suggest that it is beneficial for the nervous system to keep ensembles close to bifurcations so that they are most informative.

Ensembles of similar neurons are a theoretical construct. However, there is experimental evidence that supports the functional existence of ensembles. First, in a large number of animals there are cortical micro/minicolumns (Horton and Adams, 2005; Hubel and Wiesel, 1962; Lubke and Feldmeyer, 2007; Mountcastle, 1997) composed of neurons with similar stimulus preferences. Second, recent *in vivo* recordings in barrel cortex (Poulet and Petersen, 2008) show that under certain circumstances nearby neurons have correlated membrane potentials, indicating that they receive similar inputs. Furthermore, their spikes are preceded by large, sharp deflections of the membrane potential, indicating the presence of synchronized volleys. Third, because there are less thalamic neurons providing input to the primary sensory cortex than there are cortical neurons, convergence occurs (Kara and Reid, 2003), which means that cortical neurons must receive common inputs. Nevertheless, even neurons in the same cortical column are diverse in terms of their morphology, input conductance and so on. Our analysis shows that the temporal information is robust across a range of parameter values and can be combined into precise volleys even across a moderately heterogeneous ensemble.

4.5.2 Limitations of the experimental study.

The experimental results presented here apply to uncoupled neural ensembles receiving feedforward inputs. This clearly is an approximation to *in vivo* dynamics because: (1) There are recurrent synaptic connections that generate coherent activity in various frequency ranges and across different spatial scales (Buzsaki, 2006); (2) There

is feedforward inhibition that follows feedforward excitatory volleys at a small delay (Cruikshank et al., 2007) and (3) There are recurrent loops between the different cortical layers (Callaway, 1998; Douglas and Martin, 2004; Foss and Milton, 2000). Nevertheless, at the soma/spike generating zone the sum of these inputs leads to a fluctuating drive. This fluctuating drive consists of fast fluctuations riding on top of slower modulations. We show here that spike timing is generated by these fast fluctuations and modulated by slow fluctuations. Experiments show that the relative phase between gamma oscillations in two brain areas can modulate the effectiveness of communication (Womelsdorf et al., 2007). Hence, an important issue for further study is how fast feedforward fluctuations interact with internally generated oscillations in the gamma frequency range.

In the in vitro experiments the drive is injected as a current at the soma. According to our interpretation this represents the excitatory and inhibitory synaptic inputs that are spatially distributed across the soma and dendritic tree. Nevertheless, somatic current injection can not account for the nonlinear integration that takes place in the dendrites (Poirazi et al., 2003; Polsky et al., 2004). Furthermore, since the conductance effect due to opening of synaptic channels is not included in the current drive, its effect has to be included separately. We have done this by varying the amplitude and offset. This is motivated by model analyses of the effects on the firing rate versus current (f-I) characteristic. Adding a constant conductance shifts the entire f-I curve of a leaky integrate-and-fire neuron to the right, thus acting as a hyperpolarizing current (Holt and Koch, 1997), which can be accounted for by a current offset. Synaptic inputs add a fluctuating conductance, which in models can change the gain of the f-I multiplicatively and increase the impact of other fluctuations. In vitro experiments show random background activity produced by a network can cause changes in the gain or sensitivity to other inputs (Chance et al., 2002; Fellous et al., 2003). This

not exactly the same as changing the amplitude, but shares some commonalities in its effect.

The offset and amplitude approximate a number of effects. First, because neurons in the ensemble are not identical, there is diversity of offset and amplitude values across the ensemble. Our results show that despite this diversity, the ensemble can produce synchronous volleys. Second, when the network state changes, as it might do in response to top-down modulation, the overall offset and amplitude changes. In vivo experiments have documented corresponding changes in gain and sensitivity in response to top-down activation of cortical networks (Reynolds and Chelazzi, 2004; Reynolds and Heeger, 2009).

Interestingly, the in vitro experiment across trials maybe a better approximation of a neural ensemble on one trial than the response of one in vivo neuron across multiple trials because the drive component that varies from trial-to-trial has a large shared component among different neurons in the ensemble (Deweese and Zador, 2004).

4.5.3 Analysis method

Precision and reliability are independent properties of neural dynamics, although they may be modulated in a correlated fashion. There is both a need for an easy way to characterize the overall variability as well as for parsing out the reliability and precision separately. A reliability measure such as the Schreiber measure (Schreiber et al., 2003) is appropriate for the former, because it tells you that there is less overlap between pairs of spike trains, but it does not directly tell you whether this is due to a reduced reliability or precision or whether there are multiple patterns. Because reliability and precision are event properties, the event structure needs to be extracted for the latter. The problem is that distinct events may be overlapping, so they cannot be easily separated. To see this more clearly, consider sampling from two Gaussian densities

with a mean of -0.5 and 0.5, respectively, and both with a standard deviation of 1. It is not possible to say to which of the two distributions a sample point at 0 belongs to. However, by using the spiking history the overlapping events can be separated, which means using the fact that there are spike patterns. Hence, we utilized an analysis that used the knowledge of spike patterns. Because this analysis procedure only depends on a few well-defined parameters, for which heuristics are available, it is reproducible from lab to lab. It is described in Chapter 3. Briefly, it determines the temporal resolution at which the spike patterns can be optimally distinguished; it then finds the number of patterns; subsequently it determines the events for each pattern and it finally merges events common to multiple patterns.

4.5.4 Future studies

New technologies utilizing light-activated channels and pumps together with the read-out of neural activity via two-photon microscopy, offers the opportunity to alter and follow neural ensembles in vivo (Cardin et al., 2009; Han et al., 2009; Sohal et al., 2009; Zhang et al., 2007a,b). This makes it possible to move neural ensembles away or towards BPs. Furthermore, in vitro, at the single neuron level, rapid spatially-distributed glutamate uncaging (Gasparini and Magee, 2006; Iyer et al., 2006) can be used to determine how neurons respond to the feedforward inputs generated by a neural ensemble and how these inputs interact with pharmacologically generated fast oscillations. Taken together, these technologies offer the opportunity to test in vivo predictions about the functional role of neural ensembles positioned at BPs.

Chapter 5

Mechanisms for Stimulus Selection at the Single Pyramidal Cell

5.1 Abstract

Humans are able to selectively process different aspects of the sensory universe around them. This ability, called attention, requires that the brain have a mechanism to rapidly modulate the neural circuits which process sensory information depending on the current goals of the organism. Although some mechanisms of neural computation are well understood, the exact method by which the brain accomplishes attention modulation is not well understood. Here we explore two methods which operate at the level of a single pyramidal cell which allow for the selection of one of two signals, both of which stimulate the neuron simultaneously. In the first, synapses are clustered into sub-domains of the dendrites, and inhibition to sub-domains is adjusted to facilitate stimulus selection. We find that selection is difficult to accomplish in our model by manipulating inhibitory firing in sub-domains of the dendrites.

The second mechanism uses the γ oscillations, which are known to be experimentally associated with attentional modulation. In this method, the phase spike times

Table 5.1: Synaptic Parameters

Type/Param.	τ_{syn} (ms)	E (mV)	g_{syn} (mS)
Excitatory	3	0	0.00034
Inhibitory	11	-85	0.00044486

in incoming feed-forward excitation representing sensory signals is adjusted relative to the phase of an ongoing inhibitory γ oscillation local to the modulated neuron. We find that by adjusting the phase of an input, its contribution to the output firing rate can be strongly modulated. Hence, only phase adjustments are required to selectively enhance or inhibit the representation of a signal in the output of the neuron.

5.2 Multicompartmental Model

A geometrically simplified, biophysically inspired multicompartmental model of a layer 2/3 pyramidal cell originally described by Traub et al. (2003) was used for all multicompartmental simulations in this study. Modeling was performed in the NEURON environment Carnevale and Hines (2006) (see Chapter 2). The Traub model was initialized with default settings as per its ModelDB entry. All simulations were performed with a timestep of .025 ms.

5.2.1 Synapses

Synaptic inputs were provided using simple exponential synapses (*ExpSyn* in NEURON). Spike times were generated using the *Matlab* software environment and loaded into a custom neuron MOD which triggered *ExpSyn* activity. Synapses were either excitatory or inhibitory, with the parameters for each enumerated in Table 5.1. Synaptic parameters were adjusted in order to produce the expected relationship between input and output firing rates, but the final values were within the range of experimental results (Pare et al., 1998).

In large simulations synapses become the dominant computational bottleneck - the number of dynamical variables associated with 10,000 synapses (even very simple ones) is much larger than the approximately 60 compartment model itself. Anatomical estimates indicate that a single pyramidal cell in the cortex can receive on the order of 10,000 synapses (the model in this dissertation has 18,000) (DeFelipe and Farinas, 1992). It is both computationally inefficient and redundant to simulate each synapse individually. The characteristic length of membrane potential dynamics in the Traub model is around 25 - 50 μm (see Subsection 2.2.3). One branch of the basal dendrites contains around 650 excitatory synapses (DeFelipe and Farinas, 1992) in about 200 μm of dendrite, giving a synaptic density of 81 synapses every 25 μm . Accordingly, the number of *actual* excitatory synapses simulated by the model was reduced by a factor of 50. Each synapse was activated with spike times from 50 “virtual” presynaptic neurons. Inhibitory synapses make up around 15-20% of the synapses on the cell – there are substantially fewer inhibitory synapses. In order to maintain a uniform distribution of synapses in the dendrites, the number of actual inhibitory synapses was reduced only by a factor of 10. The calibration of the model was repeated for different scaling factors of inhibition and excitation. Changes in the neuron’s response for different scaling factors were small, and could be adjusted for by altering the maximum conductances of both excitatory and inhibitory synapses. All results come from modeling with the number of excitatory synapses scaled back by a factor of 50 and inhibitory synapses scaled by a factor of 10.

5.2.2 Calibrating the Spontaneous Background Activity

One goal of this study was to examine selection in the presence of reasonable cortical “background” activity. What constitutes cortical background activity, however, is controversial, if it is even well defined at all. Experiments indicate that excitatory neurons

in the cortex may rest at around 1 Hz and inhibitory neurons at around 5 Hz, so these values were used for the purposes of calibration (Fujisawa et al., 2008). They define both the input and output rate of the neuron, since if the cortex is resting, we expect our neuron’s firing rate to equal the background excitatory firing rate of 1 Hz. This was accomplished by adjusting the conductances of the excitatory and inhibitory synapses until the output firing rate was 1 Hz. There are many pairs of values g_e and g_i for which the output firing rate is 1 Hz, (these tend to lie on a line in g_e versus g_i space such that the ratio between g_e and g_i is about $g_e/g_i \approx 0.76$). One point on this line was selected (see Table 5.1) by finding the point where the membrane potential fluctuations were between 5 and 10 mV (excluding spike times). These parameters were used for all results reported here.

5.2.3 Generating Spike Times

Spike times driving the synapses were generated in Matlab using correlated Poisson processes with a given rate. Oscillations were then added to the spike trains (when appropriate) as described after this Section. Spike trains are not independent between all synapses on a section of a neuron. In order to simulate correlations between spike trains arriving at synapses, we used an algorithm which, given a desired firing rate and parameter dictating correlation, produces a given number of spike trains.

Begin with the number of synapses required n_{syn} , the desired firing rate f , the correlation factor $0 < c_f \leq 1$, the reliability r and jitter σ of events. The algorithm begins by producing a number of seed trials equal to the ceiling of $n_{seed} = n_{syn}c_f$. Then, n_{seed} spike trains are generated from a Poisson Process with rate $\frac{f}{r}$. Each process serves as the seed for the ceiling of $\frac{n_{seed}}{c_f}$ spike trains. For each such train the seed train is visited spike by spike. Whether the seed spike appears in the generated train depends on the reliability, and the output spike is shifted by a value drawn from a Gaussian

with the given jitter. All generated spike trains are gathered together, shuffled, and then n_{syn} are taken as the generated spike trains. For $c_f = 1$, the output spike trains are all correlated (drawn from a single “type” of spike train, with given precisions and reliabilities for each event). When $c_f < \frac{1}{n_{syn}}$, then each spike train is independently generated. All simulations in this study use a $c_f = 4/10$. This value was selected in order to increase the firing rate while keeping synaptic conductances within reasonable ranges, since correlated spikes are more likely to produce output spikes.

In cases where there were time-varying firing rates, the above algorithm was used, but the stimulus was generated in blocks of different firing rates and then the output spikes were concatenated.

5.2.4 Generating phase-locked spike times

Phase locking of spikes was achieved by first generating either constant or rate-variable spikes via the above method and then locking those spikes to specific phases with a given jitter. Each spike in a data set was visited in turn, moved to the nearest phase (specified with a frequency and phase pair) and jittered by an amount drawn from a Gaussian distribution with a specified phase locking value. Since spikes are rarely moved further than one period \pm the phase locking jitter, for small locking jitters, the resulting spike train retains the firing rate variation of the input spike trains at frequencies below than the oscillation frequency.

5.2.5 Distribution of synapses in model

For the purposes of distributing synapses in the model, the neuron was divided into 17 logical sections. The Traub Model geometry is illustrated in Figure 5.1. Each logical section has a group of inhibitory and excitatory synapses distributed uniformly throughout the section using custom written Matlab code. Logical sections may be

made up of several NEURON sections, in which case the uniform distribution is normalized by the area of each section, so that the number of synapses per unit area is constant.

Eighteen-thousand (18,000) synapses were modeled in this study. Eight-six (86) percent were excitatory and the remainder (14 percent), were inhibitory. The exact number and distribution of synapses onto a layer 2/3 pyramidal cell in the visual cortex of any particular animal is unknown. Therefore, general observations about pyramidal cells were used to distribute the synapses (DeFelipe and Farinas, 1992). Thirty-three percent of the synapses were distributed onto the basal dendrites (both inhibitory and excitatory). The remaining 66% of the synapses were placed onto the apical trunk and apical tufts. This distribution is consistent with the observation that the largest density of synapses are located on the apical trunk (DeFelipe and Farinas, 1992). Excitatory synapses were excluded from the somatic section, while inhibitory synapses were not (DeFelipe and Farinas, 1992).

5.3 Stimulus Selection by localization of inhibitory and excitatory synaptic activity

The output firing rate of a neuron increases when its excitatory input increases, and decreases when its inhibitory input increases. In the simplest case, the firing rate of the neuron can be approximated as a weighted, linear sum of the excitatory and inhibitory firing rates activating its synapses (with the weight on the inhibitory rate necessarily negative). In real neurons, however, synapses are located throughout the dendrites, and specific locations in the dendrites may be enervated by synapses which follow a common signal (Petreanu et al., 2009). What effect does this spatial clustering have on the relationship between the level of inhibitory and excitatory input and the output

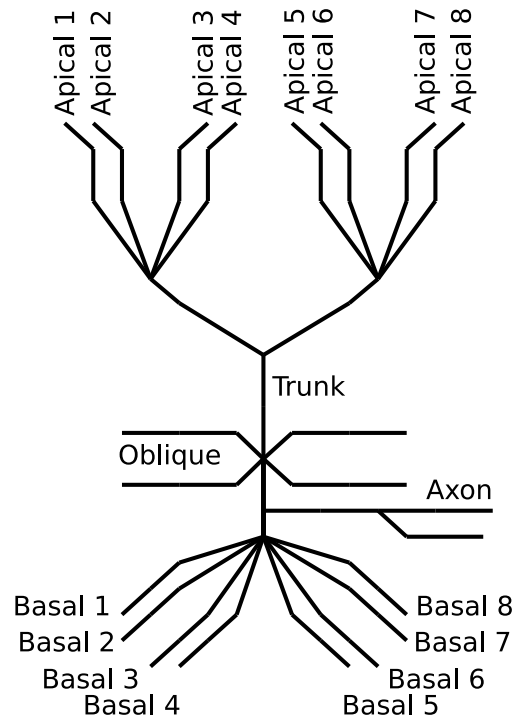


Figure 5.1: Geometry of the Traub Model (Traub et al., 2003). Eighteen-thousand synapses were distributed throughout the dendrites of the Traub model: one third in the basal dendrites, another in the apical tufts, and a third on the trunk. The ratio of number of inhibitory to excitatory synapses was 0.16. Firing rates to each basal branch were varied for this study. Each such branch has 650 excitatory and 105 inhibitory synaptic inputs.

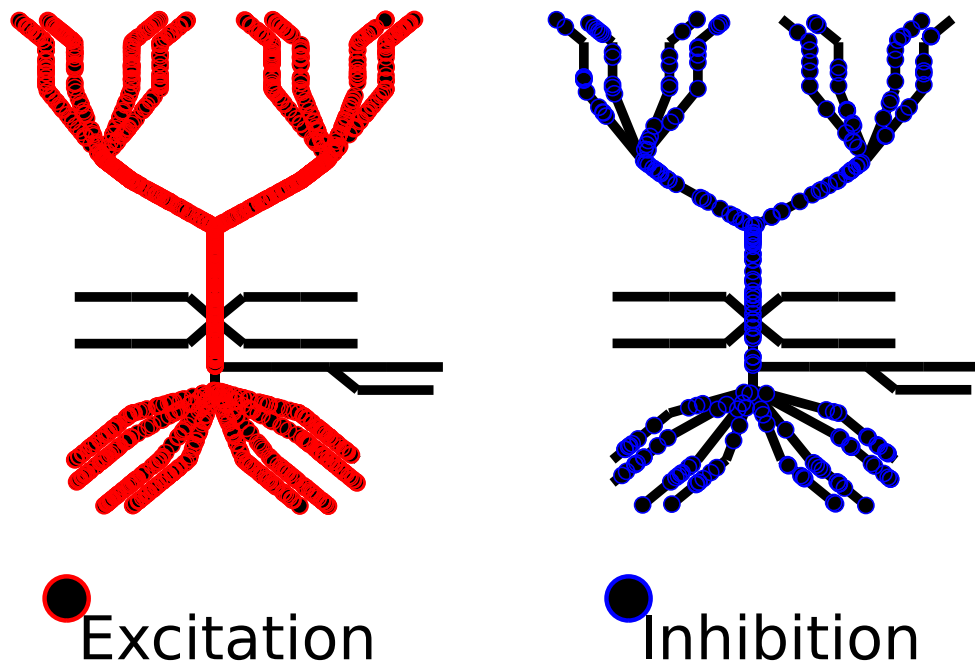


Figure 5.2: Typical distribution of synapses (1/10 of the total synapses are shown). Excitatory (right) and inhibitory (left) synapses are plotted at their locations on the Traub model. There are 18,000 synapses in total, so only a tenth are plotted here for visual clarity. During actual modeling, only 1/50th of the excitatory and 1/10th of the inhibitory synapses were placed on the neuron. Each was driven at 50 and 10 times the desired firing rate.

firing rate? In particular, we aim in this Section to discover if any such effect can be used to enhance the ability of a neuron to select one from many incoming signals.

We examined the feasibility of stimulus selection at the single pyramidal cell level by simulating the effect of location specific inhibitory and excitatory activity. In this scenario, feed-forward, excitatory (FFE) inputs representing a given signal are clustered in sub-domains of the basal dendrites Figure 5.1 and 5.3. Inhibition is also assumed to be distributed in a stimulus-specific manner, although the exact relationship between the level of activation of the inhibitory and excitatory synapses at a given location is different depending on the specific scenario being investigated (these are described below and in Figure 5.3).

We begin by examining the effect of moving the location of excitatory and inhibitory constant firing rate “signals” relative to one another in the basal dendrites and examining the effect on the output firing rate.

The simplest scenario we can examine for modulating stimulus selection is illustrated in Figure 5.4. This mechanism applies inhibition to the location at which the signal we wish to block arrives (Figure 5.4A). Ideally, this local inhibition blocks only the excitation it is specific to (Figure 5.4B), allowing other excitatory drives to “pass through” the neuron unaffected (Figure 5.4C). To shed light on this process, we compare the effect of inhibition on a single dendrite on the firing rate of the neuron when excitation is *colocated* (COL) on that dendritic branch versus when it is located on another dendritic branch (that is, *contralocated*, CON). In particular, if excitatory firing above the background rate is added to a single dendritic branch, how much colocated inhibition is necessary to return the neuron to the baseline firing rate? How does this amount of inhibition effect the neuron when the excess excitation is absent? How does it effect the neuron when the excess excitation is located on another basal branch? The answers to these questions will inform the models of selective mechanisms examined

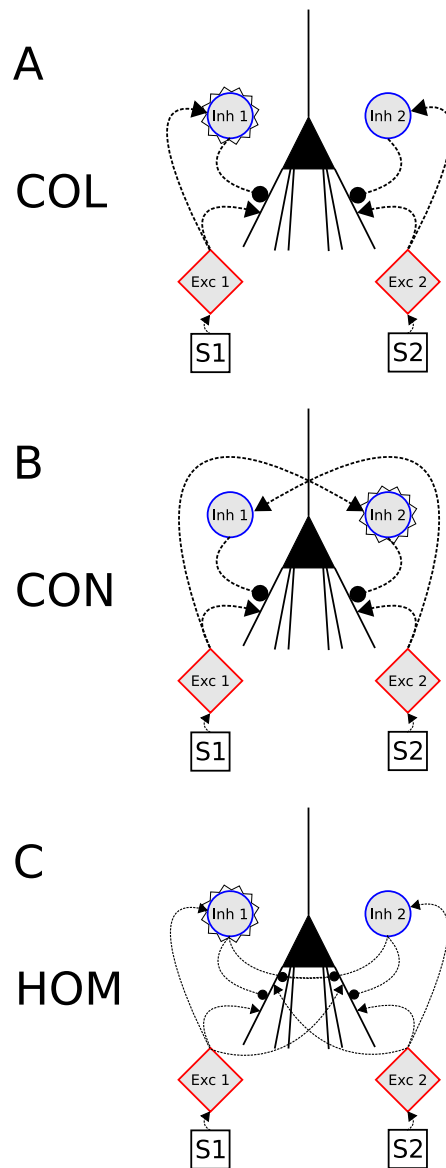


Figure 5.3: Three synaptic placements strategies for stimulus selection. In each case, the inhibitory and excitatory populations of presynaptic neurons follow either $S1$ or $S2$. In the COL (A) case, excitation and inhibition following a given signal are colocated. In the CON (B) case, they are located on opposite dendritic branches. In the HOM (C) case, synapses are distributed throughout the basal dendrites. In each case an inhibitory population's firing rate is reduced as a way of selecting either $S1$ or $S2$. This population is indicated by a starburst pattern (see text).

below.

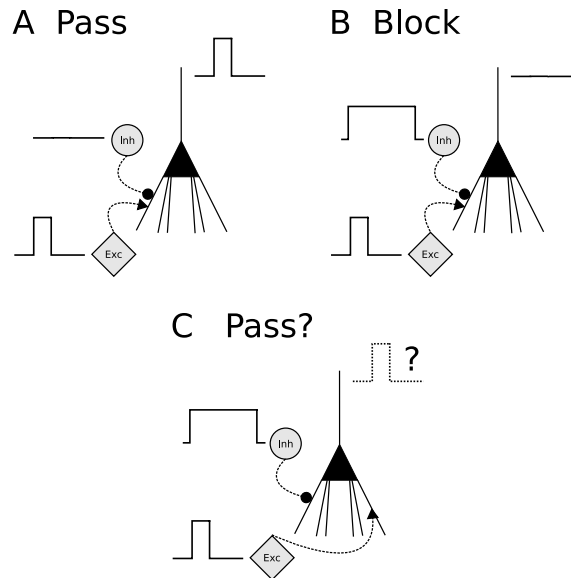


Figure 5.4: A naive approach to stimulus selection: inhibit the location where the excitatory stimulus arrives. Three conditions are illustrated (A-C). In the top two conditions, excitation arrives on the left most dendrite. In the “Pass” condition (A), there is no inhibition (above background) and the excitatory input produces a response. In the “Block” condition (B), inhibition is applied to the location of the stimulus we wish to block, with the intention of preventing the excitation from reaching the soma and producing spikes. The hope is that excitation applied to the right-most dendrite would be unaffected (C).

In Figure 5.5 the output firing rates for the model neuron are shown in 8 conditions, described below. The left-most wide, light bar indicates the background firing rate condition (BG). All other conditions are additions to this basic simulation. In the BG case, each of the eight basal dendritic sub-branches receive 2 Hz of excitatory activity in their 650 synapses (see Section 5.2.5 for the synaptic distribution scheme). The remainder of the excitatory synapses (10,400) are attached to the neuron’s apical dendrites and trunk and were activated homogeneously at 1 Hz. The inhibitory synapses, of which there are 2520 distributed onto the basal and apical dendritic arbors, are activated at 5 Hz. In order to minimize the effect of the specific location of excitatory and inhibitory synapses in the basal dendritic branches, which were all similar in length, the same seed

was used to distribute synapses in all eight basal dendritic branches. Spike times for both excitatory and inhibitory populations were generated in groups for each dendritic branch. We tuned the excitatory inputs so that an output of 20 Hz in the BG case was attained in order to allow for stronger modulation of the output firing rate.




Figure 5.5: Inhibition can “cancel” the effects of excitation, but its effect is not local to the dendritic branch to which it is applied. Eight experimental conditions are shown in the Figure, where each rectangle’s height represents an output firing rate. Beginning with the white rectangles from left-most to right-most these conditions are: background (BG) where each of the 8 dendritic branches receives 5 Hz inhibition and 2 Hz excitation; left-only (L) where the left-most dendritic branch receives an additional 6 Hz of excitatory activity; the right-only case (R), where the right-most dendritic branch receives this additional 6 Hz; and in the right and left case (R+L), where this excess excitation is delivered to both dendrites. The gray bars indicate the same conditions with an additional 35 Hz inhibitory signal delivered to the right-most dendrite. The effect of inhibition is to lower the output firing rate to the BG case in the R case, and to lower rate to the L case in the R+L case, but in the other cases the effect of inhibition is still apparent. If the effect of inhibition was isolated to the branch on which its synapses were located, the BG and R cases would be unaffected by inhibition on L, but this is not the case.

The bar graph (white bars) in Figure 5.5 shows responses to the background case described above and to three other conditions. In the Left (L) condition, an additional $F_{e,L} = 6$ Hz of excitatory activity is delivered to the synapses on the left-most branch

of the basal dendrites, raising the output firing rate to 40 Hz. In the Right (R) case, the same increase is applied to the right-most branch, also leading to an increase in output firing rate to 40 Hz. In the Left plus Right case (L+R) both inputs are added at the same time. In this case the output firing rate was about 54 Hz, less than twice the R or L case alone (≈ 94 Hz). These results show that the neuron responds equally to excitation of the same magnitude on either dendrite, but responds sub-linearly to the presence of two inputs at once.

How much inhibition is necessary to return the neuron’s output firing rate to the *BG* rate? Does it matter whether this inhibition is colocated or contralocated, and if so, how and to what degree? We addressed this question by first using simple bisection (Arfken, 1985) to obtain the amount of colocated inhibition needed to lower the output firing rate to the *BG* case. This value turns out to be a constant multiple of the excitatory firing rate (for rates between 5 and 40 Hz), so that $F_{i,R}^{F_{out} \rightarrow BG} \approx 5F_{e,R}$. The result of applying this amount of inhibition (approximately 30 Hz) to the *R* dendritic sub-branch results in the firing rate being decreased to the *BG* rate as indicated by the bar graph (filled bars) in Figure 5.5.

What happens in the other cases? In particular, when the “counterbalancing” inhibition is located on *R* but the extra excitation is located on *L* (that is, the contralocated case; CON), what is the relative effect on the output rate compared to the colocated (COL) case? The filled bars in Figure 5.5 show the output firing rate in the four conditions, but with the “counterbalancing” inhibition located on *R*. These results show a substantial reduction in the output firing rate for all conditions, indicating that while an appropriate amount of inhibition can cancel a given amount of excitation in the firing rate regime. Nevertheless, the firing rate is primarily determined by the total balance of inhibition and excitation because the location of the two components of the input at these firing rates is less important.

In our model of stimulus selection (to be described below), the total level and inhibition and excitation may be “unbalanced” as the selective mechanism increases or decreases the amount of inhibition associated with a signal. Under these circumstances, does it matter where the “unbalanced” inhibition is in relation to the various excitatory inputs? In order to more closely examine the effect of location specific inhibition on the firing activity of the neuron when inhibition and excitation are *not balanced* we examined the F_{out} (output firing rate) versus $F_{e,R}$ (input excitatory firing rate) curve under different inhibitory firing rates in the COL and CON conditions (Figure 5.6).

We found that for some combinations of firing rates contralocated inhibition was more effective than colocated inhibition at reducing the output firing rate. Two different views of the effect of COL versus CON inhibition are shown in Figure 5.6. The top-most panels (A and B) show the output firing rate of the neuron against the input rate of a single dendritic group located on the left-most dendritic branch of the neuron. Each line, starting from the highest (and darkest) and going to the lowest (and lightest), shows the F_{out} versus F_e curve for increasing inhibitory firing rates. The location of inhibition relative to excitation (COL, Figure 5.6A versus CON, Figure 5.6B) effects the shape and modulation of the F_{out} versus F_e curve. The modulation of the firing rate is more pronounced for the CON case, and saturates for large inhibitory firing rates. COL inhibition shifts the curve down rather than changing its shape, and does not appear to saturate at these firing rates. Viewing the data as a function of the level of inhibition, rather than excitation (Figure 5.6C, D, and E) illustrates that the difference between COL and CON firing rates is largest when the excitatory input is large (40-50 Hz), but eventually vanishes for large enough inhibition. These differences can be used to enhance the effectiveness of stimulus selection by balancing excitation and inhibition, as described below.

Given that for most firing rates, the output rate of the neuron depends only on

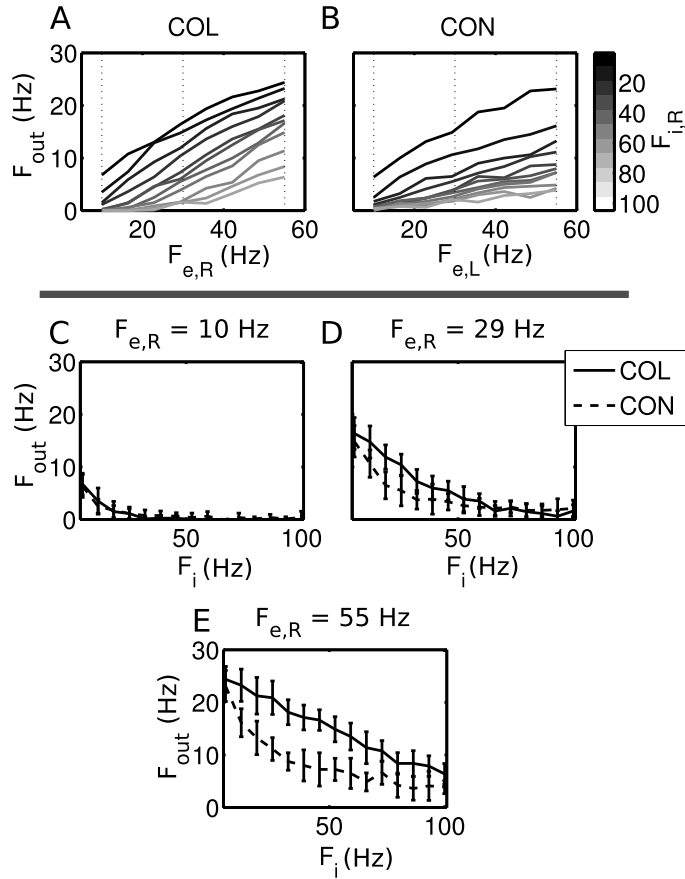


Figure 5.6: Contralocated inhibition decreases the output firing rate F_{out} more effectively than collocated inhibition for some firing rates. (A) and (B) show the output firing rate versus input firing rate for a variety of inhibitory firing rates applied to either the left-most dendritic sub-branch (where the excitation is) in (A, COL) or to the right-most dendritic sub-branch (opposite to the excitation) in (B, CON). The inhibitory firing rate is encoded by the color of each line, with lighter colors indicating higher firing rates. In (C)-(E) three cross-sections of this data (at $F_{e,R} = 10, 30$ and 55 Hz, indicated by dashed vertical lines in (A) and (B)) are shown (error bars indicate standard deviation of the firing rate across 30 trials), with inhibitory rate on the x-axis and firing rate on the y-axis. At low excitatory firing rates (C) the difference between CON and COL inhibition is not significant, but at higher rates the difference increases - (D) and (E).

the total level of excitation and inhibition and not on the location of the synapses of either, the expectation is that stimulus selection should be attainable by selectively unbalancing inhibition and excitation without utilizing any effect of location.

Consider a neuron on which excitatory and inhibitory inputs are dynamically balanced, such that each signal arriving on a dendrite is canceled by a similarly varying inhibitory signal (weighted by c_1 or c_2):

$$\begin{aligned} R_{out}(t) &= R_{e,L}(t) - R_{i,L}(t) + R_{e,R}(t) - R_{i,R}(t) + R_{rest} \\ &= R_{e,L}(t) - c_1 R_{e,L}(t) + R_{e,R}(t) - c_2 R_{e,R}(t) + R_{rest}, \end{aligned}$$

where we make the simplifying assumption that inhibition and excitation have the same effect on the output firing rate per unit input firing rate. When c_1 and c_2 are equal to one, the output firing rate is just R_{rest} ,

$$\begin{aligned} R_{out}(t) &= R_{e,L}(t) - R_{e,L}(t) + R_{e,R}(t) - R_{e,R}(t) + R_{rest} \\ &= R_{rest}, \end{aligned}$$

If, however, one of these constants is lower than one, the balance of excitation and inhibition is upset, and the associated excitatory drive appears in the output (with $c_1 = .25$ and $c_2 = 1$):

$$\begin{aligned} R_{out}(t) &= R_{e,L}(t) - .25R_{e,L}(t) + R_{e,R}(t) - 1R_{e,R}(t) + R_{rest} \\ &= .75R_{e,L}(t) + R_{rest}. \end{aligned}$$

In other words, selection is accomplished by decreasing the coefficient relating the excitatory and inhibitory firing rates which represent or follow the incoming signals, unbalancing inhibition and excitation. In terms of the connectivity of the cortical circuit, feed-forward connections arriving on the neuron from the previous cortical area or layer carry the signal directly and correspond to $R_{e,L}$ or $R_{e,R}$ above, while the inhibitory neurons which receive similar feed-forward excitation, and immediately project inhibition onto the local pyramidal cells, are represented by $R_{i,L}$ or $R_{i,R}$. The fact that the excitatory and inhibitory firing rates are related by a constant represents the fact that the inhibitory population are being driven by the same signal that drives the pyramidal cells (Vogels and Abbott, 2009).

The selective ability of three different strategies (illustrated in Figure 5.3) are characterized in Figures 5.8 and 5.7. In the first strategy (Figure 5.3A), inhibition is colocated, which we will continue to call the COL case. In the second (Figure 5.3B), inhibition is contralocated (CON) and in the final case (Figure 5.3C), both excitation and inhibition are spread throughout the basal dendrite of the neuron (HOM case). Regardless of which strategy is examined, selection is always achieved by lowering the appropriate inhibitory firing rate.

A one second long sample of the input firing rates for the three cases is shown in Figure 5.9. Each panel is divided horizontally between selecting $S1$, whose excitation arrives on the left side, and selecting $S2$, whose excitation arrives on right side. In the COL and CON cases, Figure 5.9A & B, excitation and inhibition are divided between the Left dendritic branch and the Right dendritic branch (left and right columns, respectively). In the COL case, the *shape* of excitation and inhibition is the same on a given dendrite (Figure 5.9A). When selecting the input arriving on a dendrite, the firing rate of the inhibition there is decreased, unbalancing the inhibition and allowing the excitatory signal to produce changes in the neuron's firing rate. In the CON case,

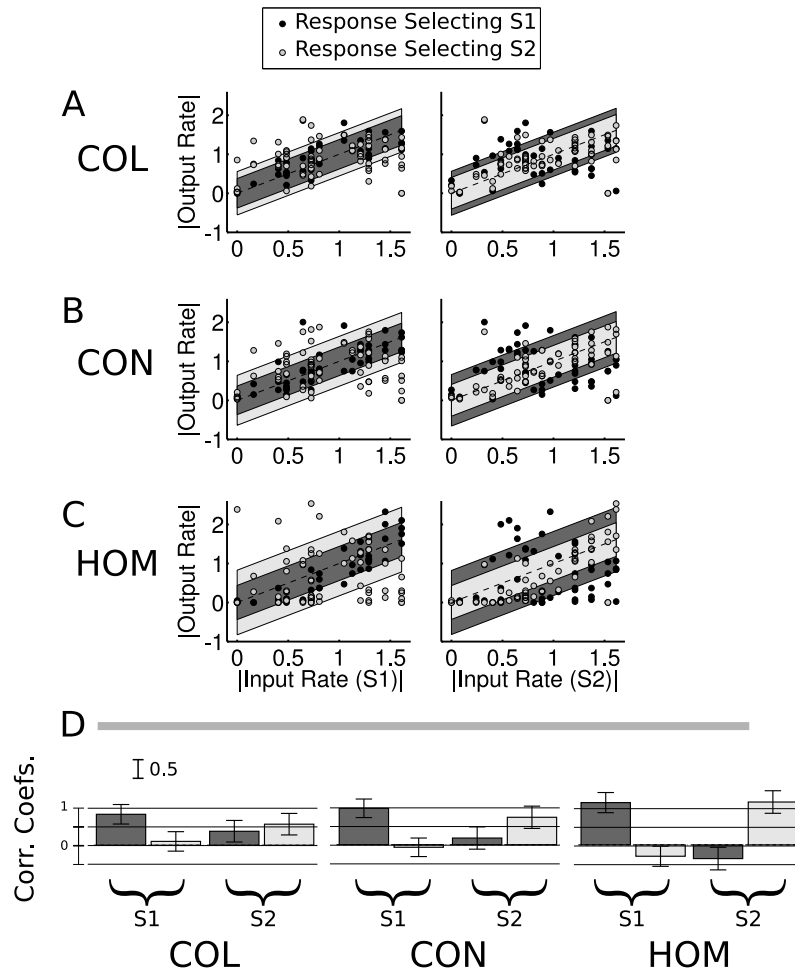


Figure 5.7: CON inhibition produces slightly better correlations between selected input and output, while minimizing correlations between the output and unselected input. In each case (A, COL; B, CON; C, HOM) the output rate is plotted against the input rate $S1$ or $S2$ (left, right sub-panels, respectively) when selecting either $S1$ or $S2$ (light circles, dark circles, respectively). Dark or light shading indicate the remaining standard deviation (selecting $S2$ or $S1$, respectively) when the linear trend relating input and output is subtracted. In Panel D, bar graphs show the regression coefficients between either input and the output (same color codes as above), relative to either $S1$ or $S2$ (as labeled). For these plots firing rates were divided by a constant $c = std(r \cup -r)$ (see text) where r is the set of firing rates. Error bars represent the confidence intervals on the linear regressions (see text).

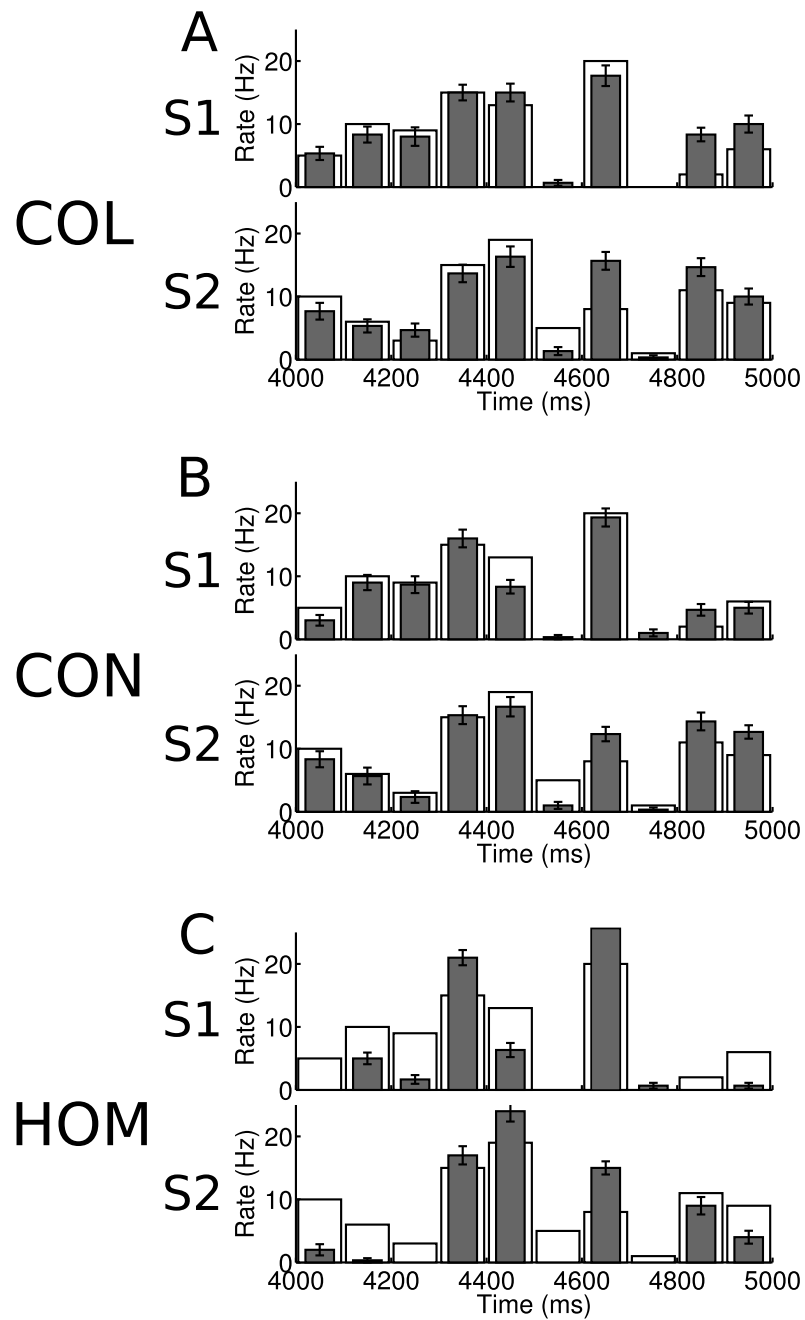


Figure 5.8: CON inhibition produces slightly better selection. Output firing rates for 100 ms bins of time are shown for the COL (A), CON (B) and HOM (C) cases. Outputs and inputs are shown for the cases selecting $S1$ or $S2$ (top & bottom graphs in each panel respectively). Input rates are shown as empty bars, output rates are plotted as filled bars.

the inhibitory firing rates have their locations reversed: the right-branch inhibition has the same shape as the left-branch excitation and vice-versa (Figure 5.9B). In this case, to select the signal represented by the excitation on the left-most branch, the inhibition on the right-most branch is scaled back. *In both cases, the inhibition which follows the selected signal is reduced.* In the HOM case, synapses throughout the basal dendrite reflect only the total amount of inhibition and excitation. Selection is affected in the same way, however, as in the COL and CON cases: the inhibition which cancels the selected signal is reduced.

The CON inhibition strategy produces the best relationship between input and output firing rates, shown in the results plotted in Figure 5.8 and 5.7. By comparing the output firing rate to the input signals, $S1$ and $S2$, this difference can be visualized as a correlation between the input and the selected and unselected output (Figure 5.7). In the scatter plots, points lying on the line $y = x$ indicate a strong correlation between input and output, while points distributed uniformly in the plot indicate the absence of correlation. The linear regression coefficients (calculated using Matlab's *regress* function on $S1$, $S2$ and the output firing rate, adjusted so that distribution of the firing rate and its negative values has a standard deviation of one), are shown in the lower portion of the figure.

Although the HOM selection case shows a better correlation between selected input and output, negative correlations coefficients for unselected inputs indicate that the unselected signal is “leaking” into the output – when the unselected signal is high, the output tends to be low and vice-versa. In an ideal situation, the unselected signal would not affect the output firing rate. The CON case produces good correlation between selected input and output, and the unselected input's coefficient is near zero, indicating no inverse correlation in the output signal. Taken together, these results indicate that using a CON strategy for stimulus selection leads to a slightly cleaner

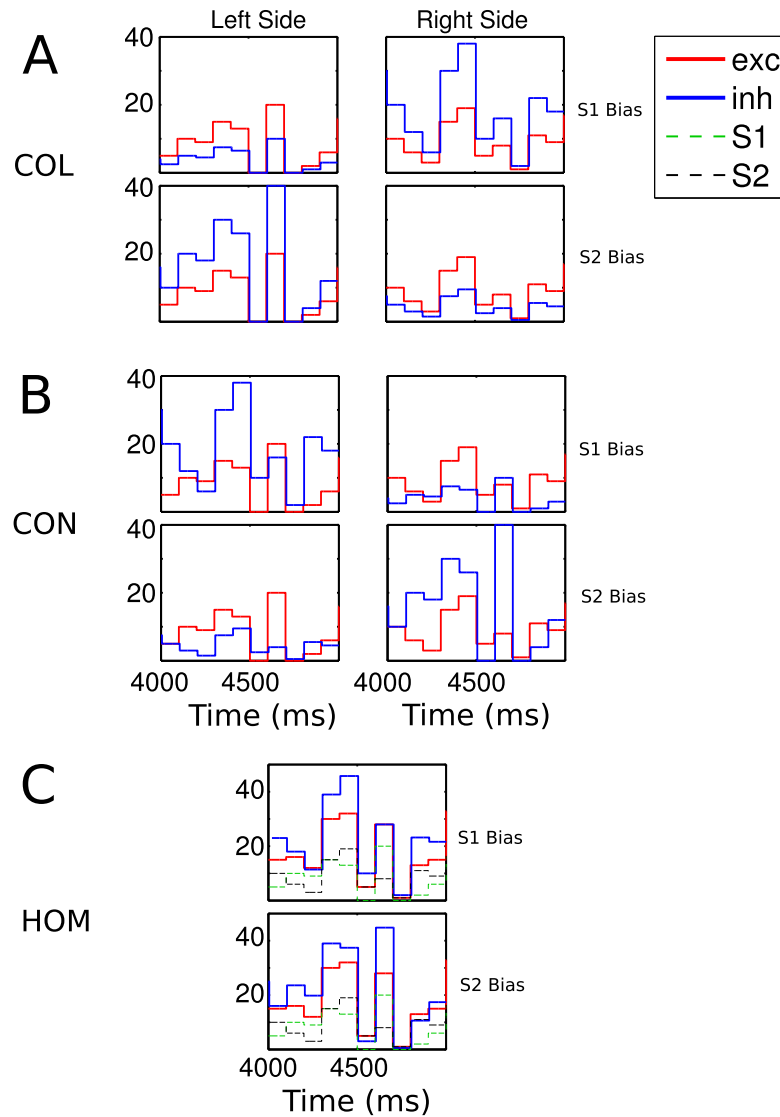


Figure 5.9: Reducing the firing rate of inhibition responding to a particular signal “selects” that stimulus. The inhibitory and excitatory input firing rates are plotted in blue and red, respectively for the COL (A), CON (B), and HOM (C) cases. For the COL and CON case, firing rates are location dependent, but all inputs are spread out evenly in the HOM case, therefore only one plot is shown. In the HOM case $S1$ and $S2$ are plotted as dashed lines.

representation of the desired output from the neuron, although the difference between COL and CON strategies is not statistically significant (p -value = .23, null hypothesis that selection is equal for both COL and CON, p -values calculated by using a mixing term with linear regression with Matlab's *regstats* function (Matlab, 2009)).

CON selection is most effective at lowering the firing rate because it reduces the standard deviation of the membrane potential at the soma. A neuron's firing rate can depend on both the mean membrane potential and its standard deviation (Tiesinga and Jose, 2000). Figure 5.11 shows views of the somatic membrane potential and standard deviation thereof along with the output firing rate of the neuron. A simple conceptual model of a neuron is useful in understanding these results. Such a model is illustrated in Figure 5.10. In this figure, a "membrane potential" (drawn from a Gaussian distribution) is plotted as a function of time. Each time it crosses the spiking threshold (at 1 unit), a spike is recorded (for simplicity the refractory period is ignored here). To increase the firing rate of the neuron we can either increase the mean of the membrane potential or increase the standard deviation - either approach produces more threshold crossings, and therefore more spikes. Decreasing one parameter can, for some values of the other, "cancel out" the effect of changing the second value. Going from panel C to B in Figure 5.10, the membrane potential has been increased, but the standard deviation has decreased, resulting in a net decrease in firing rate. From B to D the standard deviation, and therefore firing rate, increases. In summary, a neuron needs to have the right combination of mean membrane potential and standard deviation to fire at a high rate.

When firing rates are large, locating inhibition on an unexcited dendrite (CON) produces a lower standard deviation of the somatic membrane potential than inhibition located on an excited dendrite (COL) (Figure 5.11B, E). This is accompanied by a lower firing rate in the CON case, despite the fact that the mean of the membrane potential is

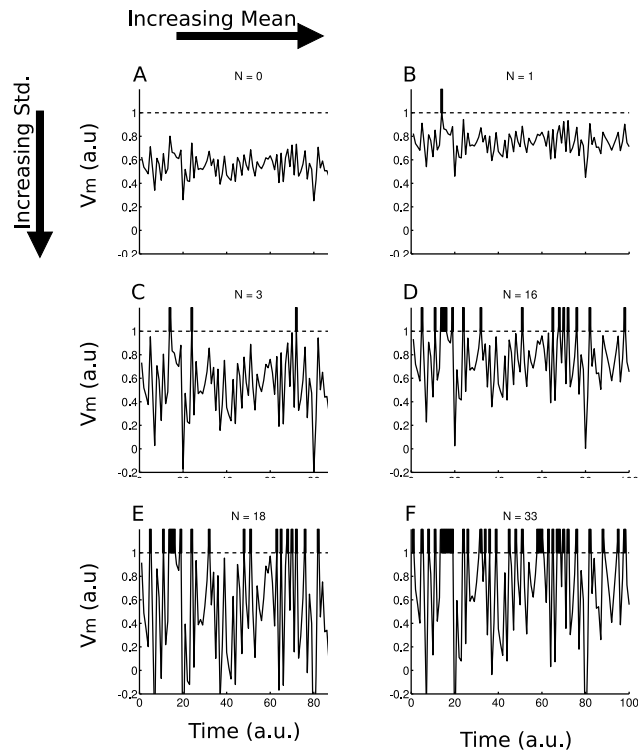


Figure 5.10: Firing rate depends both on the mean and standard deviation of a neuron’s membrane potential. A schematic “voltage trace” is plotted as a function of time, with moments for which it crosses the spiking threshold marked by black bars at the top of each panel. Moving down the figure, the standard deviation of the membrane potential is increased, while moving from right to left shows an increase in the mean of the membrane potential. Both trends lead to an increased firing rate as more points cross the spiking threshold. Note that with appropriate combinations of mean and standard deviation, a neuron can be very depolarized (high membrane potential) but have a very low firing rate because the standard deviation is so small that there are few threshold crossings (B). Only when both the membrane potential and the standard deviation thereof are large, as in D or F, are high firing rates found.

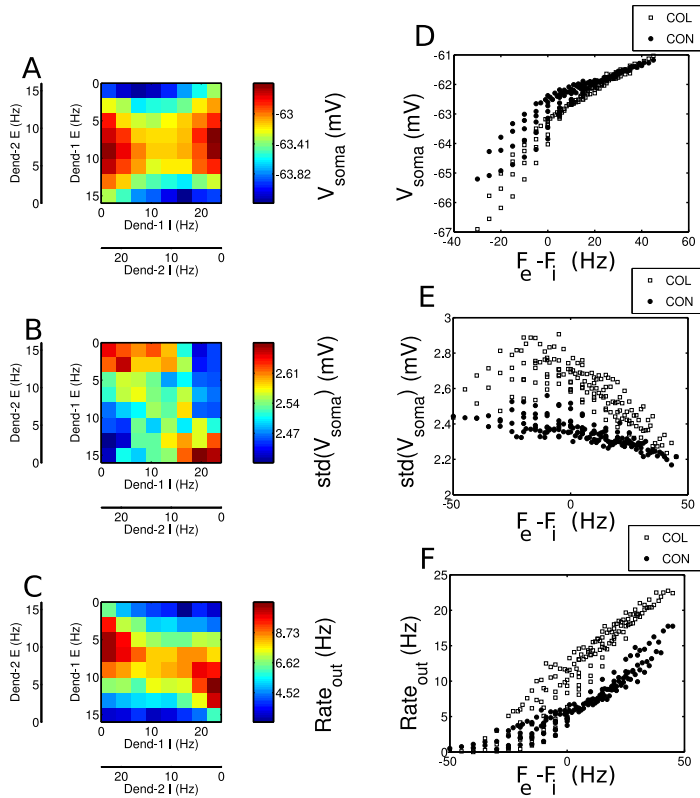


Figure 5.11: CON inhibition lowers somatic standard deviation, and therefore firing rate, better than COL. In the left column A, B and C, the membrane potential, its standard deviation, and the firing rate (respectively) are plotted as a function of location asymmetry in both excitation and inhibition. Along each y-axis the excitatory firing rate on the left-most dendrite decreases while the firing rate on the right-most dendrite increases. On each x-axis, the inhibitory firing rates similarly vary. The total excitatory and inhibitory firing rates are constant for all points on the grid. In D, E and F, the same values are examined for non-constant sums of excitation and inhibition. On the x-axis the excitatory firing rate minus the inhibitory firing rate is plotted, while the y-axes show the same values as in A, B and C respectively. Dark points show the CON case, light points the COL case, as indicated in the legend.

similar for both cases (Figure 5.11A). Figure 5.11A shows the mean membrane potential as a function of the “effective input rate”, that is $F_{e,R} - F_{i,R}$ in the COL case or $F_{e,R} - F_{i,L}$ in the CON case. When the excitation is smaller than the inhibition, the COL case tends to have the lower average membrane potential. This does not translate into a lower firing rate (Figure 5.11F), however, because at these input rates the standard deviation of the somatic membrane potential is much larger than that in the CON case (Figure 5.11E). CON inhibition lowers firing rate more effectively predominantly due to the effect it has on the membrane potential standard deviation, rather than on the mean.

The key to understanding the increased effectiveness of CON inhibition over COL or HOM is to recognize that the modulation strategy we are employing depends on canceling inputs to the neuron until one of them is “selected.” In the hypothetical null state, all the incoming signals are canceled by similarly varying inhibition, and so the output firing rate of the neuron does not vary from the background. When a signal is selected, the inhibition associated with it is decreased. Because inhibition cancels a signal more effectively when it is located on *another* dendritic sub-branch, it is advantageous to locate the canceling inhibition for a signal on other branches of the neuron than the one it arrives on, particularly since that inhibition, by virtue of the fact the COL inhibition is less effective at filtering out a feed-forward signal, will less effectively modify whatever feed-forward signals are arriving at its location. Changes to the level of inhibition in the CON strategy have a larger effect on the cancellation of the signals they follow than they do on the signals they are colocated with, which is what makes it a good strategy. In an optimal world, changes to the inhibitory firing rate blocking *S1*, for instance, would not change the degree to which *S2* is passed or blocked. Although that ideal is not attainable in the electrically compact neurons we examined here, contralocating inhibition is slightly better in this respect than colocated

inhibition, producing the improved correlations and rejections we see in the comparisons of input and output firing rates (Figures 5.8, 5.7).

In conclusion, we find that although a neuron's total firing rate is largely determined by the total amount of inhibition and excitation present in the cell, regardless of location, that a CON inhibitory/excitatory layout can improve selection for some firing rates over either a HOM or COL strategy.




Figure 5.12: Firing rate depends on both standard deviation and mean of the membrane potential – the CON strategy leads to lower standard deviations. Standard deviation appears on the Y axis, mean on the X axis. Size and lightness represent firing rate (see side panel). Note that for the same mean membrane potential, the CON case has consistently lower firing rates and standard deviations.

5.4 Selection by inhibitory oscillation and excitatory phase-locking

The attractiveness of modulating ongoing oscillations as a means of stimulus selections follows from several observations. The first is that signals from the outside world are not

predictable in time – the pitcher in a baseball game does not cue the batter before he throws the ball. Ongoing oscillations provide a brain or area-wide clock to synchronize and select stimuli. Additionally, neural oscillations are a persistent element of neural systems, and have been experimentally observed to correlate with attention and arousal (see Chapter 2).

The method of utilizing oscillations for selection examined in this Chapter is illustrated in Figure 5.13. Rather than using specific locations to encode whether a stimulus ought to be passed or not passed, this strategy uses specific *times* (relative to an ongoing γ oscillation) to encode the same information. In the location specific strategy, inhibition and excitation were located at the same position in the dendrites, and to inhibit a specific signal, inhibition located with that signal was altered. In the oscillatory case, excitation is instead located at specific *times*, and to select a signal, either the phase of the excitatory signal or the phase of the inhibitory signal is adjusted to cause signals to arrive at either the high or low point of inhibition in the γ cycle (at the level of the single neuron, shifting the excitatory phase is indistinguishable from shifting the inhibitory phase by the same amount, in the opposite direction).

In order to characterize the effectiveness of this method of stimulus selection, simulations were carried out to examine the effect of arrival phase and input firing rate on the output rate and jitter of the model neuron. These results are illustrated in Figure 5.14.

Inhibitory inputs (firing at 10 Hz) to the basal dendrites were entrained to a 40 Hz γ -style oscillation as described in Section 5.2.3. All excitatory synapses fired spikes drawn from a Poisson process with a rate of 3 Hz, except for the 650 excitatory synapses on one dendritic sub-branch of the basal dendrites, which were also entrained to a 40 Hz oscillation. The phase of the inhibitory oscillation was varied in 10 steps between 0 and 2π . Both inhibitory and excitatory spike times were aligned to their particular




Figure 5.13: Stimulus selection by oscillation and phase locking. In this diagrammatic representation of stimulus selection by oscillation, each presynaptic source of excitatory input (illustrated in “V2” on the lower part of the figure), carries information about a signal. Spike times coming from each neuron are synchronized at a particular phase relative to a “reference” γ oscillation. The ongoing gamma oscillation at the post-synaptic cell prevents the neuron from responding to an input unless the spikes arrive at the correct phase. Selection is accomplished by aligning the phases of the pre- and post-synaptic neurons so that excitation is able to produce a response. “Unselected” or inhibited signals arrive at unfavorable phases with respect to the post-synaptic γ .

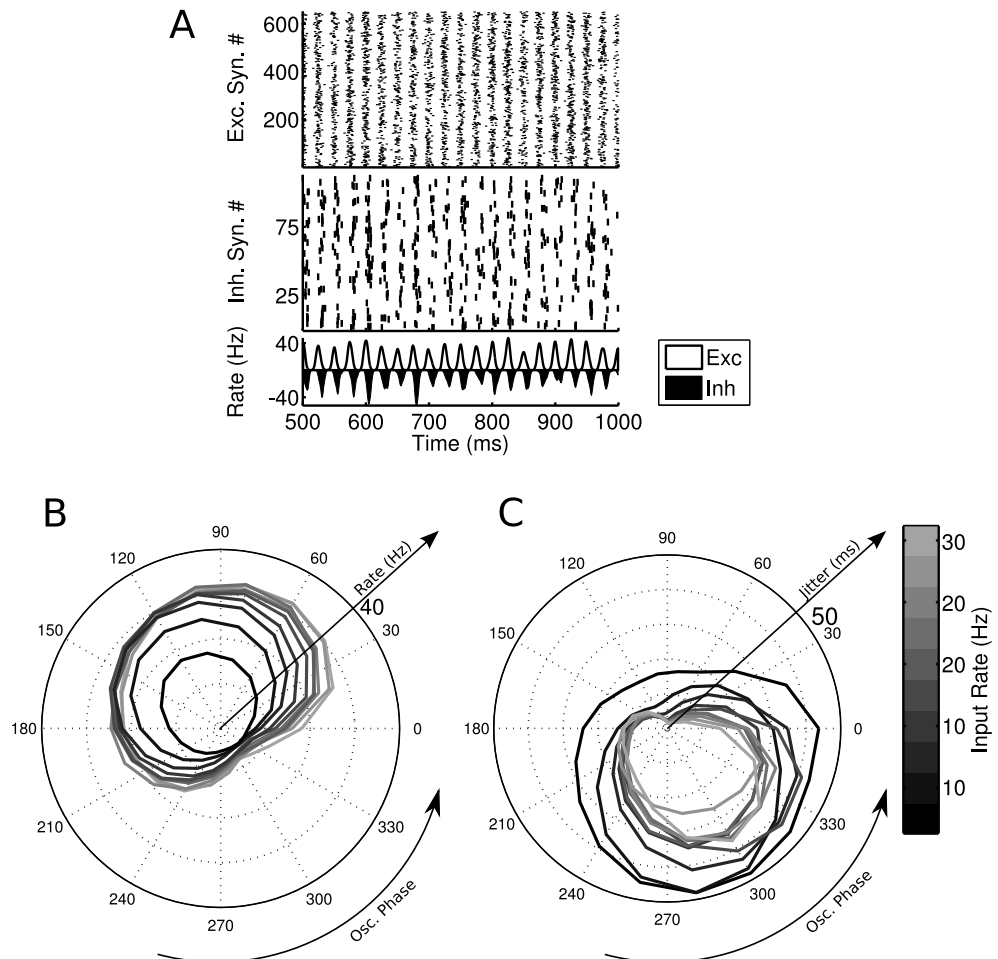


Figure 5.14: Relationship between incoming rate and phase and output firing rate and jitter. (A) shows typical input spike trains for the experimental case, along with the firing rate in the lowest panel. (B) shows the output firing rate as a function of both input rate (color) and phase. (C) shows the output firing jitter (calculated using the jitter measure described in Subsection 4.5.3). When excitation has no phase locking, the firing rate is radially symmetric (not shown).

phases with a 3 ms jitter. Output firing rates were measured in a sliding window 800 ms wide, aligned with the phase of the inhibitory oscillation. Jitters were calculated as described in Section 4.5.3, on the spikes contained in the sliding window.

FFE phase locking can substantially modulate the output firing rate, the gain of the firing rate, and the jitter of the output spikes (Figure 5.14). Just below a phase of 2π , the output firing rate is almost unaffected by changes in the input firing rate. At this phase, spikes are arriving at the peak of the inhibitory volleys which make up the inhibitory *gamma* oscillation. Jitter shows a similar effect - when the excitation arrives at a non-preferred phase, the spikes that are produced have high jitter. At the point 180° away from the lowest firing rate (approximately $\frac{6}{5}\pi$), the Jitter is ~ 2 ms. Taken together, this implies that locking the incoming inputs to a particular phase can produce effective modulations of the output firing rate and jitter, while preserving information about the incoming rate on at least very long time scales.

Real stimuli have temporal structure, and locking feed-forward spikes to particular phases interferes, at least in principle, with the temporal information encoded in the variations of the feed-forward firing rate. The typical timescales of visual stimuli, however, are below the γ range, and the integration time of retinal neurons is also slower - suggesting that it is possible to produce stimuli which preserve temporal variation at frequencies below gamma while still spiking at times restricted to certain γ phases. In this scheme, the temporal variations of the firing rate encodes multiple pieces of information - low frequencies encode the (slowly varying) signal, while high frequencies encode attentional or selective state.

This mechanism requires that oscillations occur in the inhibitory neurons (both at the selected cell, and in the previous cortical area), and that excitatory feed-forward spikes are entrained to those oscillations, so that they spike at particular phases. A complete examination of synchronization of different networks of neurons is outside of

the scope of this investigation, but it is worth examining whether entrained spikes can be produced in a reasonable model neuron. To that end, we performed simulations with a single compartment pyramidal cell to demonstrate the feasibility of generating phase-locked spike times. (N.B. - the output spike trains of the multi-compartmental model in this Chapter are also strongly phase locked, providing additional assurances that phase locking in the presence of inhibitory oscillation is reasonable).

In a single compartment neuron, adding an oscillation to a time-varying input produces phase-locked output spikes (Figure 5.15). A single-compartment model neuron after Golomb and Amitai (1997) was injected with a signal created by superposing a time varying signal (scaled, filtered Gaussian “frozen” noise) and a 40 Hz sine-wave oscillation (the oscillation is substantially less effective at driving spikes in the Golomb model, and so has a standard deviation of roughly 7 times that of the frozen noise in order to produce a similar number of spikes) (Figure 5.15A). The superposition was adjusted so that the output firing rate was approximately constant for all drives. This resulted in a larger standard deviation as the oscillation became the dominant component of the input, since the frozen noise was more effective at eliciting spikes from the model. The current was injected along with noise sampled from a Gaussian distribution with standard deviation $10 \mu A$ sampled at 40 kHz, that is, at each step of the simulation.

The degree of phase locking can be adjusted by changing the relative weight of the oscillation and the signal in the inputs. The results are displayed in two different ways – the first is a standard rastergram (Figure 5.15B), where each group of horizontal line-delimited trials is aligned with the current that produced it in Figure 5.15A. Since it is difficult to see the degree to which spikes are phase-locked in Figure 5.15B, a second view of the same data is provided in Figure 5.15C. These “phasograms” show all the spikes in response to a given input (from all trials). Distance of a spike from the origin

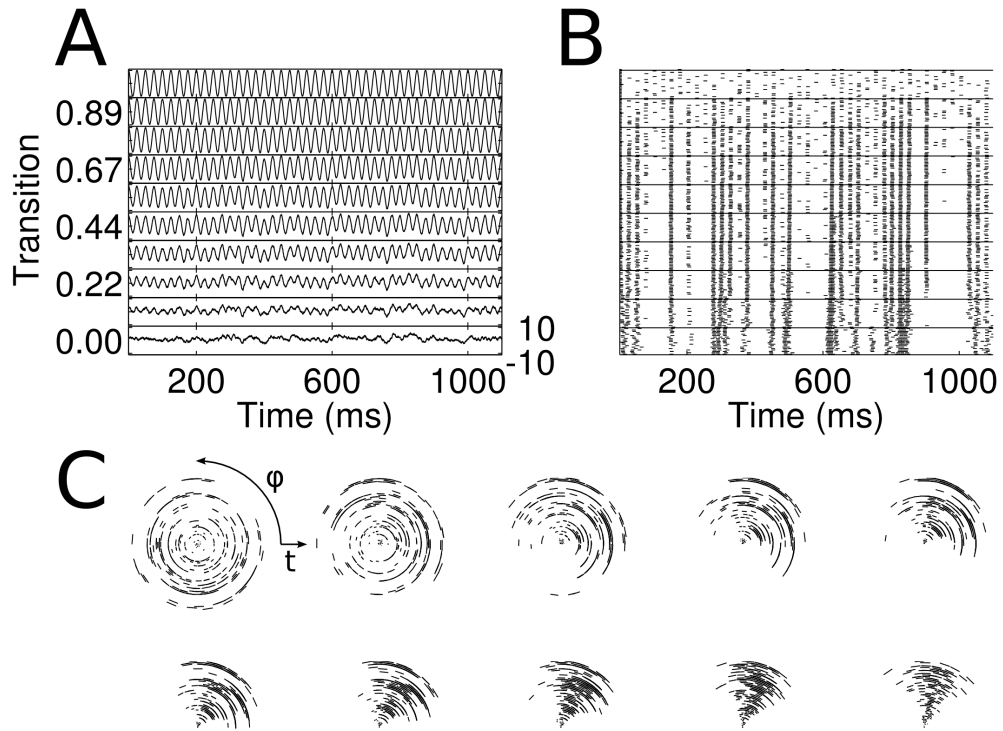


Figure 5.15: Phase-locked spike trains can retain temporal information lower than at γ frequency. A Golomb model neuron (Golomb and Amitai, 1997) was driven with a series of current injections made up of a super-position of a frozen-noise stimulus and an oscillation (A). The superposition was gradually adjusted between pure oscillation and pure signal, and the output spike times were recorded (B). (C) Phasograms, in which the angle around the origin represents the γ phase of a spike and distance from the origin indicates spike time relative to the start of each trial. The output spike trains are shown, proceeding left-to-right and top-to-bottom. Intermediate responses show strong phase locking (jitter $\sim 3ms$ around a fixed phase, and slower time-scale firing structure reflective of the frozen noise input).

indicates its absolute time (since the beginning of the trial), while its phase relative to the gamma oscillation is plotted as the angle around the origin. The responses proceed from left-to-right and from top-to-bottom, and from having an output rate reflecting the signal's variations, to being essentially completely oscillation locked. Around the 5th or 6th current injection, the neuron's output spikes are strongly locked to the oscillation (standard deviation around the peak phase of about 3 ms) but spiking activity also reflects the temporal structure of the input at timescales larger than the γ -range. This indicates that it is feasible to drive a neuron in such a way as to both lock the spikes to an ongoing oscillation *while preserving* temporal modulations in the input.

Although this establishes the feasibility of producing phase-locked inputs with “real” neurons, it is convenient to be able to “dial up” phase-locking of any particular strength in a predictable and simple way. In this Chapter, this is accomplished by taking a spike train and visiting each spike, moving it to the nearest phase to which spikes are to be locked, and then adding a random Gaussian jitter. The standard deviation of the Gaussian determines the strength of the phase-locking (see the detailed explanation given in Section 5.2.3).

How well can this mechanism select between two different time-varying signals arriving into the neuron at one time? Simulations were performed in which two time-varying signals were transformed into excitatory spike times and used to trigger synapses in the model. Each signal was locked either to the phase most-likely to produce output firing (selected) or the phase least-likely to produce output firing (inhibited). These phases were chosen by examining the firing rates in the constant firing rate, single signal case. The correct phases cannot be chosen only by knowing the shape of the inhibitory oscillation, presumably because the dynamics of the synapses and membrane make predicting the phase which minimized the effect of inhibition difficult. The output firing was calculated and then compared to the input signals to quantify the degree of selection.

Because the performance of this mechanisms is expected to depend on the temporal character of the input signals, simulations were run on time varying signals with distinct timescales, from 250 ms all the way to 10 ms (where the method is expected to fail, since this is below the γ -period of 25 ms). The results are shown in Figures 5.16 - 5.20.

In each figure except for Figure 5.20, similar simulations were performed and analyzed. A pair of stimuli ($S1$ and $S2$) were generated with selected characteristic timescales. A timescale for variation was chosen between 250 and 10 ms, and then time-varying firing rates were generated by sampling uniformly from integers between 1 and 25 Hz. For instance, for a 25 ms time scale, a stimulus might consist of (in part), four 25 ms blocks with firing rates 1, 23, 13 and 25 Hz. With a large block-width, the time scale of the signal is slow, and with a small block-width, the signal varies very quickly. For any given model run, both $S1$ and $S2$ had the same block-width to facilitate comparison.

Effectiveness of selection was measured by comparing the output firing rate with the input firing rates for both $S1$ and $S2$ while selecting either $S1$ or $S2$. Correlation between $S1$ or $S2$ and the output was calculated using a multiple linear regression using least squares (Matlab function *regress* (Matlab, 2009)). Such an analysis takes N components and one signal which is hypothesized to be a sum of the components and estimates the coefficients needed on each component to sum to the output. In this case, $S1$ and $S2$ are our components hypothesized to contribute to the output firing rate in a linear way. If selection were perfect, and we were selecting $S1$, for instance, then the coefficient between $S1$ and the output firing rate would be one (neglecting differences in total firing rate and offset) and the coefficient for $S2$ would be zero. The hypothesis underlying these simulations is that selection between $S1$ and $S2$ will be better the higher above the γ period the time scale of the stimulus is, and then selection

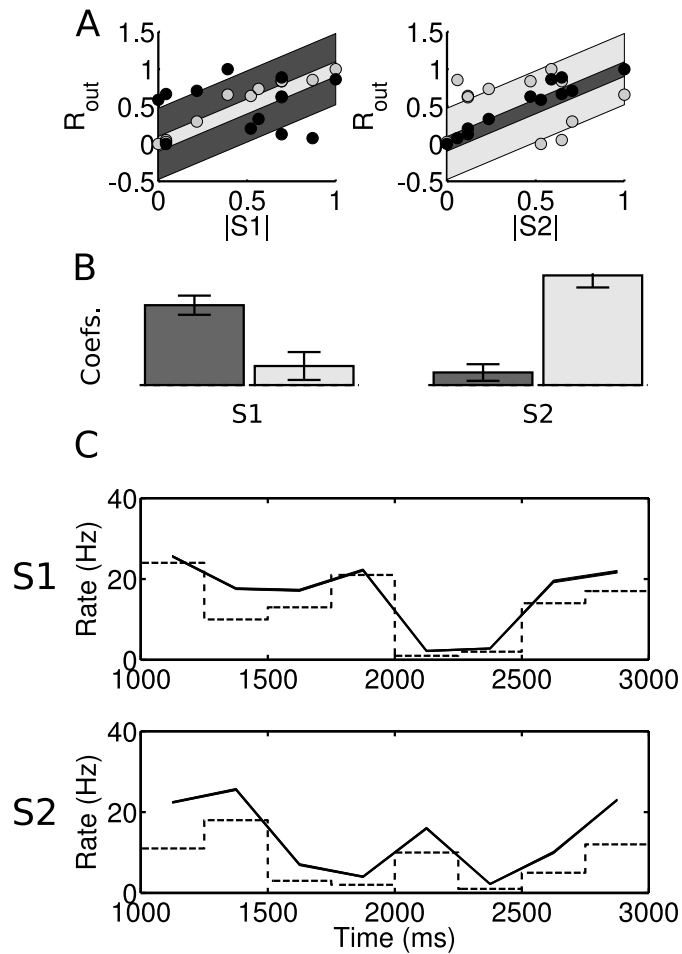


Figure 5.16: Summary of selection by oscillation with the signal varying on a 250 ms timescale. From top to bottom: (A), scatter plot of incoming firing rates ($S1$ and $S2$) versus outgoing firing rate selecting either $S1$ (left, light points) or $S2$ (right, dark points). Unexplained standard deviation is indicated by colored bars around each set of points; (B), linear regression coefficients and confidence intervals relating $S1$ and $S2$ to the output firing rate while selecting $S1$ (left pair) and selecting $S2$ (right pair); (C), input (dotted lines) and output (solid lines) firing rates plotted together for comparison. The top panel shows the output selecting $S1$ with $S1$, and the bottom panel shows the same for $S2$.

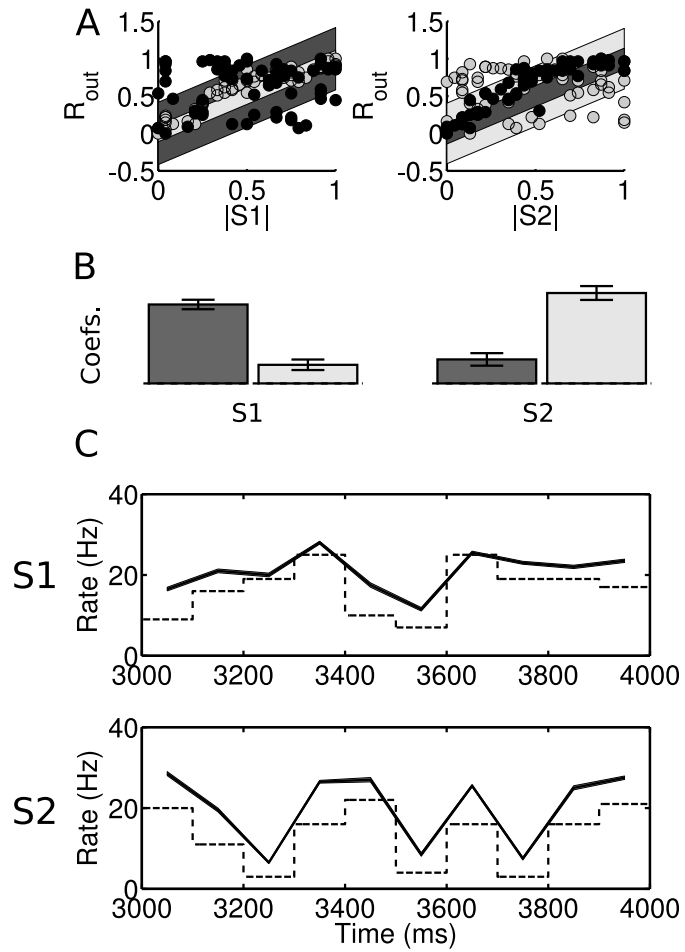


Figure 5.17: Same as Figure 5.16 except for a signal with a timescale of 100 ms.

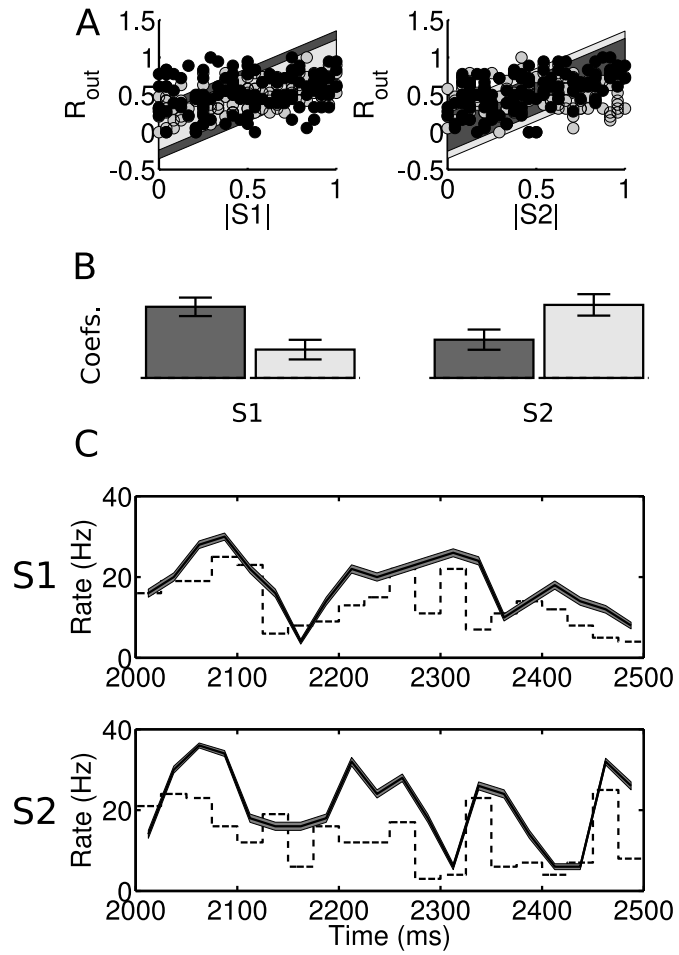


Figure 5.18: Same as Figure 5.16 except for a signal with a timescale of 25 ms.

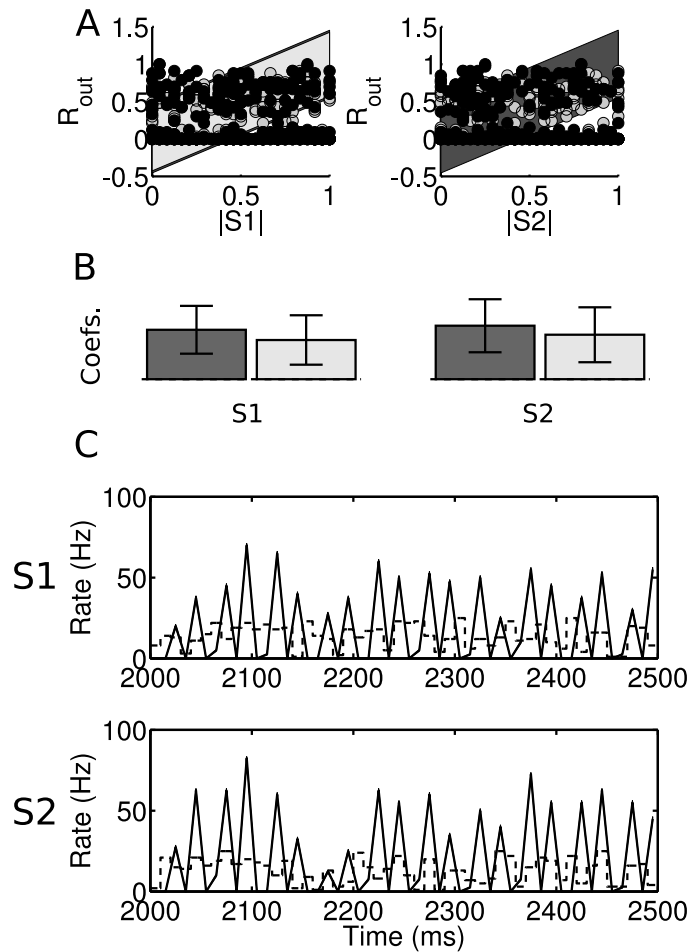


Figure 5.19: Same as Figure 5.16 except for a signal with a timescale of 10 ms. When the stimulus varies at a time-scale faster than a γ period, firing rate changes associated with the oscillation obscure the selection.

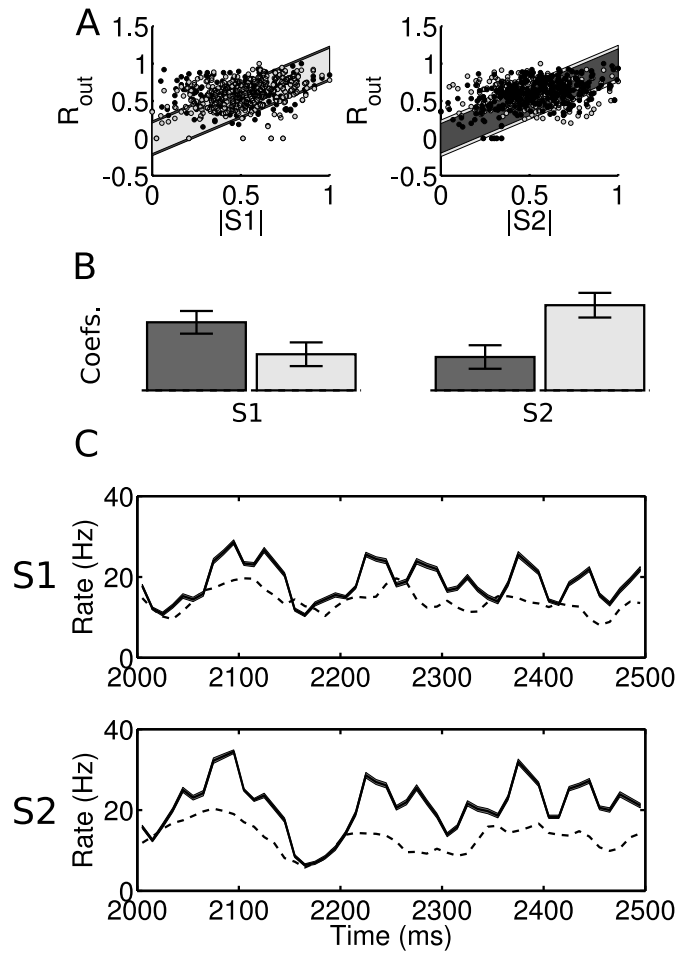


Figure 5.20: Same as Figure 5.16 except with a timescale of 10 ms. In order to determine whether variations in the input slower than the γ range were still selected in the presence of faster changes, the stimulus and the output were filtered with a window of approximately 30 ms (see text). Features with longer timescales in the input appear in the output.

will rapidly deteriorate for variations on faster timescales. This prediction is largely born out in Figures 5.16 - 5.20.

In Figure 5.16A scatter plots are shown for two cases. To the left the firing rates during selection of $S1$ are plotted in gray dots as a function of the $S1$ value which produced them. Both the input and output firing rates have been normalized so that their values are between zero and one. The gray points lie near the $y = x$ line, indicating a good correlation between $S1$ and the output while selecting $S1$. Points during the selection of $S2$ are plotted in the same panel against $S1$. These points evidence less correlation between input and $S1$, showing that when $S2$ is selected, the output is not strongly locked to $S1$. In the panel in Figure 5.16A on the right, the same data are plotted against $S2$ instead of $S1$. The same color coding is in effect, showing that when $S2$ is selected, the output is correlated with $S2$ and when $S1$ is not selected, the output is not correlated with $S2$. Figure 5.16B summarizes the data in 5.16A by showing the regression coefficients between $S1$ (left panel) and the output selecting $S1$ (right bar) and the output selecting $S2$ (left bar). The right panel shows the regression coefficients between $S2$ and the output selecting $S1$ (right bar) and $S2$ (left bar). For illustrative purposes the firing rates and input are shown (unnormalized) in Figure 5.16C. The case selecting $S1$ is shown in the top panel, with the dotted line representing the incoming $S1$ firing rate, and the solid line representing the output rate. In the bottom panel, $S2$ and the output are shown. For all data shown in this figure, both $S1$ and $S2$ are being received by the neuron. Only the arrival time phase of each changes during selection.

Despite the fact that changes on the scale of 10 ms are not selectable by this technique, slower changes “on top” of a 10 ms signal are still selected by this method, as evidenced in Figure 5.21. In this figure, the input and output of the model were compared after being filtered to remove components at frequencies in the γ range and higher (Matlab *filter* function, with $B = [1 \ 1 \ 1 \ 1]$ and $A = 4$, applied twice, once

forward, and once backward. Since the data has a time resolution of ten ms, this amounts to a filter width of about 30 ms). After filtering, both the shape of the input and output curves and the regression coefficients show selection comparable to slower varying signals. In other words, although selection fails for signals which vary quickly, longer time scale features of the signal are selected and appear in the output. This constitutes a “graceful failure” of the method in the presence of quickly changing firing rates.

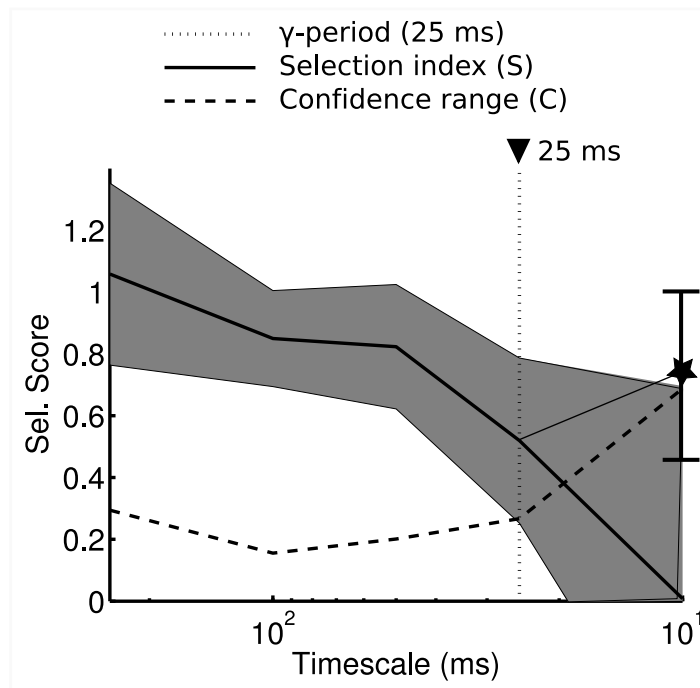


Figure 5.21: Selection degrades for time scales above that of the γ oscillation. On the x-axis is the characteristic time scale of a time varying stimulus. On the y-axis, the “selection index” (see Text), reflecting the degree to which temporal selection effectively filters incoming stimuli, is plotted, along with the “confidence range” (larger corresponds to less confidence) of this measure, both as a gray area around the score and as its own line (dotted) for comparison. A vertical dotted line indicates the γ period. Stimulus selection fails when the time scale of the input becomes faster than γ (indicated where the solid and dotted lines cross, around 20 ms) The point marked with a star shows the selection index for a 10 ms signal when the comparison between input and output is made after filtering (filter width \approx 30 ms), showing that correlations on time scales above that of the γ oscillation are preserved.

The results are summarized in Figure 5.21. A “selection index” was devised to measure the quality of selection for both $S1$ and $S2$ at each of the five time scales examined above. Each linear regression produces a coefficient for $S1$ and $S2$ as well as confidence intervals around such coefficients. For time scales for which selection works effectively, we expect several things to be true about these coefficients. When selecting $S1$ ($S2$), the difference between the $S1$ ($S2$) coefficient and the $S2$ ($S1$) coefficient should be large and positive and the confidence interval around each coefficient should be small. Therefore a selection index can be defined as:

$$S = \frac{(C_{S1}^{S1} - C_{S2}^{S1}) + (C_{S2}^{S2} - C_{S1}^{S2})}{2},$$

and a “general range” as:

$$C = \frac{CI_{S1}^{S1} + CI_{S1}^{S2} + CI_{S2}^{S1} + CI_{S2}^{S2}}{4},$$

where C_x^y indicates the coefficient weight for stimulus x while selecting y , and CI_x^y is the magnitude of the confidence interval around that coefficient. When the “selection index” becomes approximately equal to the “confidence range”, then it is difficult to determine which, if either, of the signals, is dominating the output of the model. In other words, the model is no longer strongly selective.

In Figure 5.21, the “selection index” is plotted along with the “confidence range” for each of the time scales displayed in this Section. Selection is good for values larger than the γ period, but falls off precipitously for smaller time scales, with the intersection between the selection and the confidence measure happening at around 20 ms.

While selection fails at time scales of 10 ms, slower variations are observable in

Figure 5.20 and the resulting selection index is plotted as the single star at 10 ms in Figure 5.21.

Taken together, these results show that selection by phase-locking to an inhibitory oscillation can feasibly select stimuli with bio-physically plausible time scales of variation.

Chapter 6

Discussion and Conclusion

6.1 Overview of Results

I present results in two related areas. In the first, the analysis of neural responses to repeated, time varying signals is examined. In Chapter 3 I present a novel combination of algorithms and heuristics which reproducibly find regions of elevated firing separated by regions of relative quiescence called events, in this kind of spike train data. The presence of events in spike trains was originally described by Mainen and Sejnowski (1995), wherein the trial-averaged firing rate and a threshold was used to determine where to demark events in time.

The results presented here improve upon the results reported in Mainen and Sejnowski (1995) in that they can discover events which overlap in time, but exist in separate spike patterns. Information about correlations between events is also preserved by these methods. An important new result presented here is that the events produced in response to a time varying stimulus can be correlated with one another in such a way that, on a given trial, only a particular subset of events will occur. This result was predicted by (Tiesinga and Toups, 2005) based on the approximately linear nature of the subthreshold membrane potential dynamics together with the nonlinearity

of the spike generating mechanism. These methods not only discover the boundaries of events, but also the spike patterns which produce correlations in events. The presence of events impacts the estimation of the precision and reliability of neural responses, as well as the ability of a downstream neuron to “decode” the structure and amplitude of an incoming signal, as described in detail below. The methods of finding events and spike patterns also provide new insight into the relationship between spike train clustering and the time scale used for spike train metric space calculations. These results have implications for neural coding and represent a set of useful tools and heuristics which will facilitate future research by allowing more accurate statistical models of single neuron and ensemble activity.

I also present an analysis of *in vitro* spike trains consistent with the concept of sparse representation of time varying signals, such as those reported in Jadhav et al. (2009).

In the second area, I examine two methods for producing stimulus selection at the level of a single cortical pyramidal neuron when each stimulus varies in time.

This investigation is motivated by the fact that attention is necessary at the level of a single V4 neuron, and it must operate on visual signals which are rapidly varying in time. Because cortical neurons in V4 and higher visual cortical areas can have receptive fields large enough to include multiple stimuli, the brain may require a mechanism to enhance the response of the neuron to inputs representing a subset of the receptive field in order to implement spatial attention.

I present an analysis of two possible mechanisms of stimulus selection. In the first, the existence of clustering of synapses into sub-domains (motivated by, for instance, Petreanu et al. (2009)) of the dendrites is examined, and I find that, although there are differences in the ability of inhibition to modulate firing rate depending on whether it is located in the same dendritic sub-region or another one, relative to excitation, that

in the presence of multiple time varying signals, there is no significant difference in the selection based on synaptic location. I also present simulations and an analysis of using phase locking of feed-forward, excitatory “signal” spikes to specific phases of an ongoing, inhibitory γ oscillation. In this case, statistically significant stimulus selection was achieved for stimuli which varied more slowly than the γ period (25 ms) and when both stimuli and the inhibition had strong phase locking (jitter less than 3 ms). These results make specific predictions about the phase locking and phase relationships of inhibition and feed-forward excitation. This method of stimulus selection predicts that the inputs to a neuron comprised of selected and unselected stimuli will differ only in their phase relative to the local γ - they both require strong phase locking in order to accomplish selection. Because the phase range which results in significant attenuation of the feed-forward signal is relatively small, amounting to around 30° (see Figure 5.14), these simulations predict that unselected stimulus components may be even more strongly tuned to γ than selected components, which can be located at a larger range of phases. Both of these predictions should be re-evaluated in future network simulations.

6.2 Finding Events and Spike Patterns

6.2.1 Previous Studies of Events and Spike Patterns

Mainen and Sejnowski (1995) re-introduced the issue of spike time precision in neuroscience. In their study of pyramidal cells in vitro, a single, time varying “frozen noise” type current was repeatedly injected into a neuron and the resulting spikes were recorded. The resulting spike trains had repeatable spike times, with very low jitter, in contrast to the spikes produced to a constant current injection. Because the spike times occurred at roughly the same time on each trial, they could be correlated with aspects

of the stimulus. Additionally, spikes could be conceptualized as belonging to “events,” roughly corresponding to all spikes from all trials which are produced in response to the same component of the stimulus, and manifesting themselves as regions of elevated firing surrounded by regions of lower firing rate. In the original paper, events were detected and characterized using a firing rate threshold. Hence, continuous regions of time during which the firing rate exceeded a set value were counted as a single event.

The study of spike patterns in this kind of data was initiated by (Fellous et al., 2004), which combined the use of a spike train metric presented by Schreiber et al. (2003) with a fuzzy-c-means clustering to reveal that neural responses to repeated current injection could be characterized in terms of spike patterns. Other studies document a variety of alternative ways to cluster or sort spike trains by temporal structure (Schreiber et al., 2004; Victor and Purpura, 1996; Lindsey and Gerstein, 2006; Van Rossum, 2001). The approach described here is an extension of the preceding results, but is novel in several ways. First, the results described here go beyond merely finding spike patterns: spike patterns are discovered as one part of a multistep system meant to more accurately characterize neural responses. Unlike the Fellous et al. (2004) results, we use the Victor-Purpura spike train metric rather than the metric described in Schreiber et al. (2003) because of the ease with which the algorithm’s q parameter is interpreted, the computational efficiency of the measure, which does not quantize time as part of the calculation, and because the measure has a simple extension described by Aronov (2003), allowing for a straightforward generalization to multineuron spike patterns. The finding of events within individual spike patterns, and the statistical procedure for merging events across patterns, constitute a new contribution to the study of spike patterns in neural data.

The entire method constitutes a way of discovering a statistical model of the response which includes trial to trial correlations. This information is lost when averaging

over all trials.

An important new result when compared to the original Mainen and Sejnowski method is that events from different spike patterns may overlap in the spike time histogram, and will therefore appear in a pattern-agnostic analysis of neuronal data as a single, low precision event (see Figure 3.1 for an example). The implications of this result are discussed in Subsection 6.3.

6.2.2 Implications of the Current Study for Metric Space Analysis

An important result presented here is in the area of finding the most meaningful metric space parameter q for the Victor-Purpura metric. Previous studies of spike train clustering and metric space analysis have suggested that the parameter q should be chosen to reflect time scales in spike trains of interest (Fellous et al., 2004; Victor and Purpura, 1996). Choosing this timescale is difficult in our case because data sets can have a variety of inter-event intervals and jitters (the two time scale categories of interest). Additionally, we report that events from different spike patterns can overlap in the full data set, implying that there is an overlap between the inter-event and event jitter timescales. As a consequence, I pursued a more heuristic approach for selecting the spike train metric parameter based on maximizing the correctness of clustering of known data sets. I present results suggesting that the optimal timescale parameter is only indirectly related to the jitter and event time scales, and that the parameter value which maximizes the entropy of the resulting spike train distances is a more principled choice of q , at least when the final goal is finding spike patterns.

Although we work with the Victor-Purpura spike train metric in this dissertation, this insight should generalize to any measure which has a tunable sensitivity to exact spike timing, as is the case in the “Schrieber” and “Van Rossum” spike train measures

(Schreiber et al., 2004; Van Rossum, 2001), and may be indirectly applicable to more complex methods like that of Lindsey and Gerstein (2006), which use spike correlated changes in the effective “mass” of a spike train and a dynamical simulation to calculate spike train similarity.

6.2.3 Reproducibility, Heuristics and Limitations of Presented Method

I present a novel method for finding events and spike trains in neural data. Although this problem has been examined before, I present a complete analysis method with a statistical examination of parameter choices for each step. Using these algorithms, spike train clustering can be standardized and repeatable across labs, improving collaboration and verification of results. In addition to specifying a procedure for characterizing the spike patterns and events in spike train data, I provide heuristics for each step of the analysis which provide good results for data with more than 20 trials and where events are separated by larger gaps than their jitters. The only step where the algorithms are not simply deterministic is the fuzzy-c-means, where cluster membership is initialized with random values. If these random initial values are provided, the analysis presented here is completely reproducible.

Finding q , the metric timing parameter

The selection of the Victor-Purpura metric space timing parameter is discussed above in Subsection 6.2.2, but, in summary, I show that the choice of the parameter q is most effective at finding the correct known clustering when the entropy in the distances between spike trains is maximal. We also present a more precise but less reliable heuristic for q , which examines peaks in the change in the coefficient of variation of the spike train distances, with the benefit that it can be calculated without binning the

spike train distances.

Choosing the number of Clusters, N_c

It is difficult to choose the correct number of clusters in any clustering problem where it is not known in advance. The “Gap Statistic,” which compares clusterings of real data to clusterings of surrogate data created in such a way as to be cluster free, is used in this dissertation (Tibshirani et al., 2001). Whether this measure is appropriate can be addressed in several ways. Heuristically, we address this question by comparing calculated clusterings (and numbers of clusters) to many ensemble data sets which have a known cluster structure (see Figure 3.6). Although this demonstrates that the Gap Statistic produces good estimates of the actual number of clusters for the types of data we examine here, it does not establish the significance of spike train clusters in real data. We address this weakness by creating a precise definition of a spike pattern: a subset of trials from a data set for which events are uncorrelated. This implies a statistical test on events which can verify the presence of spike patterns (given the events in a set) and also one for testing whether a particular trial cluster corresponds to a real spike pattern or not. See Figure 3.9.

In a broader sense the question of what any spike pattern means depends on the nature of the neural system in which it is observed. From this perspective, spike train clustering, as presented here, is useful for retaining information in statistical models of neural activity which may be correlated with some experimentally observable quantity. These results can be thought of as a means of retaining more information about spike train responses - the question of whether that information is meaningful in terms of a particular neural system is left to be determined by future experiments and analysis.

Finding and Merging Events

We find that using an ISI threshold equal to a large fraction of the largest interspike interval in a spike train cluster produces a good estimate of the events in that spike train cluster. This simple choice of ISI threshold is only possible because overlapping events have been separated by the clustering step (see Figure 3.1). When events overlap in time between different spike patterns, there is no single interspike interval threshold which is shorter than the interevent-distances but longer than all interspike distances. When spike patterns are isolated, the refractory period of the neuron ensures that within a single pattern, such a value can be found.

However, because a single event may occur in multiple spike patterns, I present an ROC analysis and heuristic for merging events across spike patterns. I present that using an sROC value of .5 as a threshold for merging nearby events produces correct event assignments for artificial data sets with specified events.

The underlying reason for these observations is that, within a given spike pattern, the refractory period of the neuron assures that events will not be arbitrarily close in time. The insight provided from this study is that the minimum interspike interval expected from this observation does not generalize to the entire data set - unless care is taken to examine the trials closely, events will appear to be arbitrarily close in time, making naive characterization of events difficult. For this kind of data, a full characterization of event precision requires finding and isolating the spike patterns.

As in the case of the number of meaningful clusters in the data set, the question of whether particular events are really distinct or identical is impossible to objectively say in the abstract. These methods represent a way of capturing neural variability at a finer grain than the spike time histograms. Only experimental work and insight into specific neural systems can answer the question of whether a particular conceptualization of spike train data is meaningful.

6.3 Coding Implications of Spike Patterns

This dissertation examines one impact of spike patterns and event reliability on the question of neural coding in Chapter 4. The approach presented sidesteps the question of whether information is represented in the interspike intervals of the spike trains in favor of examining the impact of spike patterns and the associated event unreliability on the ability of a downstream neuron to decode features of the stimulus based on the activity of an ensemble of neurons which may or may not fire in different spike patterns. Figure 4.7 shows the result of driving a highly idealized neuron with two different groups of spikes, one representing a very synchronous population of neurons, the other representing one which is less synchronous, and sampling from a broader set of spike patterns. Using a spike triggered average on the latter data produces a better estimate of the input. This is a kind of stochastic resonance (see Rizzo (1997) for a short review of the concept). The primary novelty in the result presented here is that noise induces different kinds of changes in the variability of the output spike times. Our analysis of data from pyramidal cells in slices of rat prefrontal cortex shows that responses can have variability in spike pattern and concomitant high unreliability while, within one spike pattern, spikes continue to be precise to within a few milliseconds. One way of viewing these results is that, by quantifying the different kinds of variabilities in neural data separately (jitter, reliability and spike patterns), we are able to show that a neuron can be unreliable without becoming imprecise. Methods which average over trials to estimate variability mix up these two kinds of variability, and therefore obscure the potential of encoding time varying signals with precise, but unreliable, spike trains.

Spike patterns can be thought of as sequences of spikes with specific interspike intervals, so they can at least theoretically be thought of as temporally distinct symbols. The view I present in this dissertation is that spike patterns are merely different samples representing pieces of the same underlying temporal structure, so the implications for

temporal coding, in which distinct patterns would encode different states, are not clear. Recent results presented by Riehle et al. (1997) show some evidence for correlation based coding, where the transient synchrony between neurons in a population reflects something about the inner state of the brain. Although it is difficult to compare my results to theirs, there is some tension between the mechanism I propose - that unreliability lowers synchronous firing across a population to enhance coding - and theirs - where increased synchrony is actually used to represent the internal state of some aspect of the system. Because I hypothesize that unreliability is useful, but binding by synchrony would increase reliability, there is a trade off between synchrony of a neural population and the ability to encode information *in* synchrony (see Tiesinga and Jose (2000)). Without further modeling and experimentation on spike patterns, however, it is difficult to predict if they make a difference in neural systems in practice.

In Chapter 4, I report that *in vitro* neurons driven with a time varying current injection at a variety of amplitudes encode the structure of that injection most robustly when medium amplitudes are considered, and the reliability of individual neurons is small or moderate. This coincides with their being a large number of spike patterns. Finding the spike patterns in such a response shows that the neuron responds in many unreliable but temporally precise events, suggesting that if a population of neurons is to encode the structure of a time varying stimulus, then the reliability of each neuron in the ensemble should be low. When the reliability of all the neurons is high, then all the neurons respond to the same stimulus features, and, due to the presence of a refractory period, *miss* similar features. When reliabilities are low, different sets of neurons respond to different features of the stimulus, allowing the ensemble to reflect more features. These results are consistent with a recent experimental study conducted by Jadhav et al. (2009), which reported that neurons sensitive to the rapid accelerations produced by whisker slips used a low reliability, high precision encoding scheme

where individual neurons had high precision responses on only a few trials. This study hypothesizes that if such whisker cells were driven with repeated stimuli rather than random slips, results similar to the ones presented here would be found.

6.4 Future Spike Patterns Research

Future research in this area could progress in two directions, both of which focus on the challenges of analyzing and understanding ensembles of neurons.

Although we chose the Victor-Purpura metric for the simplicity of choosing the parameter q , an additional benefit of the measure is a simple generalization of the method which allows for the calculation of distances between ensembles of neurons with tunable importance given to the specific identity of each cell. Multi-electrode recordings and spike sorting, pioneered by Georgopoulos et al. (1988), and two photon microscopy (Denk et al., 1990) are now allowing simultaneous recording from a large number of neurons. Future research should examine whether patterns of activity analogous to those described here occur in networks of neurons, and whether the identity of neurons in the network is important. Because the only difference between single unit and multiunit data in terms of the analysis method presented here is the metric space step, it should be easy to adapt the clustering steps to multiunit data. The question of how an event should be defined across an ensemble of neurons and over repeated trials is more complex, especially if there are predictable but non trivial changes in which neurons participate, on a given trial, in the representation of a single event. For instance, if the selection of neurons from the ensemble which spike in a given event changes as a function of time due to some ongoing process, like synaptic adaptation, so that spikes belonging to a particular event appear at the same time and with the same reliability on all trials, but in different subsets of the population.

A more simple direction for future research lies in the association between single

unit recordings over many trials and the behavior of ensembles. I present a single, highly idealized downstream neuron responding to ensembles of neurons with different reliabilities and show that the reliability of individual events, and therefore spike patterns, has an effect on the ability of the neuron to encode the shape of a time varying signal. This question should be examined in more realistic systems. A simple experiment I hope to carry out in the future will be the construction of models of ensemble activity based on single neuron recordings. Once the models of ensemble activity are built, they can be used to generate conductance waveforms which can then be injected into another neuron *in vitro*, varying the degree of “pattern synchrony” and analyzing the output spike trains both in terms of spike patterns and in terms of reconstruction of the original, time varying input signal. This experiment would reveal important information about the capabilities of spike patterns to enhance representation of stimuli in multi-layered networks.

6.5 Stimulus Selection at the Single Neuron Level

The ability to attend - that is, to direct cognitive resources to processing only parts of the sensory input - is crucial to behaving complex organisms, which receive more information from their sensory systems than could possibly be consciously processed (Desimone and Duncan, 1995). Additionally, deficits in attention constitute or accompany a large number of mental illnesses from attention deficit disorder to schizophrenia (Siever and Davis, 2004; Vale, 2008; Hutchinson and Mathias, 2007; Austin et al., 2001). Consequently, understanding the mechanisms of attention is an important goal of neuroscience with many implications for understanding and treating mental illnesses.

Current conceptual models of attention are based on the idea of biasing competition between cortical columns (Reynolds and Desimone, 2003; Ogawa and Komatsu, 2004). I present two methods for modulating the activity of a single cortical neuron which may

be relevant to biased competition models, but also function in the absence of explicit competition.

I find that in a model of a layer 2/3 pyramidal cell with active dendrites originally described in Traub et al. (2003), the exact relationship between the locations of inhibitory and excitatory synapses does not significantly impact the ability of a neuron to select between two stimuli by unbalancing the ratio of inhibition and excitation associated with one signal. In this case, stimulus selection is possible, but does not depend on how synapses are distributed in the dendrites.

I also examine the ability of a neuron to select between two signals in the presence of an ongoing γ (40 Hz) oscillation when each signal is represented by excitatory activity with spike times locked to specific phases. In this case, I show there is significant stimulus selection when there is strong phase locking (≈ 3 ms) and regular oscillatory activity at 40 Hz.

6.5.1 Stimulus Selection with Synaptic Clustering

It is becoming technically feasible to map the detailed pattern of synaptic enervation in the dendrites of pyramidal cells (Petreanu et al., 2009), and this has revealed that there is clustering of synapses from presynaptic neurons onto sub-domains of the dendrites. The question examined in this thesis is whether the clustering of synapses can be utilized to produce or enhance stimulus selection for rapidly varying signals in a model of a layer 2/3 pyramidal neuron (Traub et al., 2003).

This dissertation examines three variations on the theme of stimulus selection exploiting the location of synapses. The simplest is the collocation case, where inhibitory and excitatory neurons driven by the same signal are located in the same dendritic sub-domain. In the baseline case, inhibition and excitation are balanced so that the neuron does not produce any spikes, and to “select” a stimulus, the inhibitory firing

rate is reduced, allowing the neuron to fire. In the second case, inhibition and excitation following a given signal are contralocated - that is, excitation is located on one dendritic sub-branch, while inhibition is located on another. In the presence of two signals, inhibition following signal 2, for instance, will be located with the excitatory activity representing signal 1 and vice versa. I finally examined the case where there is no specificity at all in either inhibitory or excitatory synaptic activity. Although I do present a significant difference in the ability of inhibition to modulate the output firing rate of a neuron depending on whether it is colocated or contralocated with excitation (see Figure 5.6), these results depend on the inhibited branch being unexcited, and only occur with numerical significance when averaged over long time intervals. When driven simultaneously by two signals, there is not a statistically significant difference between the selection in the colocated or contralocated cases. These results indicate that while receiving time varying signals and background firing, the Traub model neuron is too electronically compact for synaptic location to be a significant factor in selecting signals varying on a timescale of 250 ms or below.

Comparison to Previous Simulation Results and Limitations of This Study

The implications of synaptic clustering in dendritic sub-domains of the dendrites are not well characterized. Archie and Mel (2000) examined the question in a similar model neuron for constant firing rate stimuli and found that using a contralaterally arranged set of synaptic inputs, modulations were observed that were consistent with experimental studies of firing rate modulations and attention. The results of this present study are consistent with those reported in Archie and Mel (2000), with the novel addition that stimulus selection is accomplished for time varying signals. Alternative synaptic layout strategies were not examined in the Archie and Mel (2000) study. This dissertation also contributes the novel result that, at least for time varying signals, the

layout of synapses in dendritic subdomains is unimportant for selection, which depends only on maintaining or selectively perturbing the balance of inhibition and excitation.

The small effect of the location of synapses on the response properties of the neuron during stimulation by multiple sources of input raises the question of why clustering of synapses in the dendrites is observed at all (Petreanu et al., 2009). One answer may be that the model neuron used in this dissertation, originally described by Traub et al. (2003), is not appropriate for the examination of this question. The Traub model was constructed to reproduce fast and slow dendritic Ca^{2+} spikes, with the goal of studying experimentally observed bursting behavior. Neither behavior is examined explicitly in this dissertation. The choice of the Traub model was motivated by a desire to use a biophysically motivated Layer 2/3 pyramidal cell model. Nevertheless, there are two limitations to this model. First, the nature and distribution of ion channels in the dendrites of Layer 2/3 pyramidal neurons is not well understood. In pyramidal cells in L5 or CA1 where the distribution of channels is better studied, the character of the dendrites varies between cell types (Gasparini and Magee, 2006; Schaefer et al., 2007; Migliore et al., 1995; Prescott and De Koninck, 2003)), suggesting that the results from other cell types cannot be easily generalized to Layer 2/3 cells. Secondly, the morphology of the basal dendrites is highly idealized in the Traub model. Either of these limitations may restrict the usefulness of the results reported here. Future research on the subject of selection by modulation of neural populations with selective synaptic locations should examine the question of dendritic electrophysiology and morphology more carefully.

Future Work

Future work in this area would be advanced by examining the model of the Layer 2/3 pyramidal cell and by examining the same question in the context of other pyramidal

cell varieties. Although there are other mechanisms which explain the discrepancy between the results reported for time varying signals, and the ability of contralocated inhibition to lower the firing rate of a neuron more effectively than colocated inhibition, one factor may be the relatively high variability of the output spike trains in these results. Specifically, when the input frequency varies rapidly in time, the firing rate variability in a short time interval may be large enough to obscure differences in rate which are obvious at longer time scales. As a consequence, one possible mechanism which may illuminate differences between synaptic layout strategies is the addition of an oscillatory component to some or all of the inhibitory inputs arriving at the neuron. Although this would introduce a time scale limit on the possible frequency content of the selected stimuli similar to the one reported for oscillatory selection, it would also reduce spike count variability by synchronizing spikes to the oscillation, which may reveal significant differences in selection strategies. On a related note, a combination of this method with oscillatory selection may reveal that the two combined produce better selection than either mechanism in isolation.

6.5.2 Stimulus Selection with Inhibitory Oscillations

Much experimental evidence exists for a link between neural oscillations and attention or stimulus selection (see Chapter 2) (Buzsaki and Draguhn, 2004). This study pursues the idea that the brain dynamically adjusts the effective connection strengths between groups of neurons by changing the relative phases of different groups of neurons (this theory was proposed by Fries (2005) and tested in Womelsdorf et al. (2007), by correlating the local field potential with multi-unit activity). This theory hypothesizes that inhibitory interneurons prevent pyramidal cells from firing during certain phases of their oscillation (in particular, those immediately after an inhibitory volley) while allowing firing for other phases. To prevent two cortical areas from communicating

without changing the local firing rates, the phases of the two areas are adjusted such that spikes from the first area arrive in the second just as the peak of the inhibitory oscillation has arrived, so that the second area is unable to respond. Positive selection is also accomplished in a similar way, except the arrival phase is calibrated so that excitation arrives at the low (rather than high) point of the inhibitory oscillation.

This dissertation examines selection by oscillatory firing in the context of a single neuron receiving time varying signals. I report that stimulus selection can be accomplished for reasonable background firing rates when there is strong phase locking (jitter on the order of ≈ 3 ms).

Previous Results, Limitations of this Study and Future Work

Selection by oscillation at the single neuron level has been examined before in the constant firing rate case. In (Mishra et al., 2006) the relative phase between inhibitory oscillations and excitatory inputs was varied to simulate the effects of attention for constant firing rates, and the results were consistent with experimental examinations of attention like those done in Reynolds et al. (1999). My study differs from previous work primarily in the use of time varying, rather than constant firing rate, stimuli. The results here indicate that attentional modulation by phase locking can be accomplished even in the context of multiple, time varying signals with a large range of firing rates (between 5 and 25 Hz). Since neural systems experience temporally varied stimuli during natural behavior, these results represent an important validation of the theory of attentional selection using inhibitory oscillations.

The selection demonstrated here is consistent with the case of attention being directed into the receptive field of the neuron (see Figure 5.14) and Reynolds et al. (1999) in that attention results in both higher firing rates and lower spike time jitter. When no signals are selected in this model, all inputs are located at the phase least likely to

produce spikes (around 300°). It is at this phase that the firing rate is lowest and the jitter is highest. However, the precision of the inputs is still very high - all inputs in our models were phase locked with a jitter of only 3 ms, and for jitters larger than 4 ms, selection was reduced or completely eliminated. Therefore, although selection works in our model, when a neuron is not receiving any attended inputs, its output is very imprecise (see Figure 5.14C). In other words, in the model presented here, attention can lead to a substantial increase in precision. The effect of such imprecise spike times on attentional selection in the *next* cortical area was not examined here, but should be an aspect of future work.

An important aspect of attention is its effect on the firing rate of neuron which is responding to both a preferred and non-preferred stimulus at the same time when spatial or feature-based attention is directed to one of the attended stimuli. The results presented here correspond to the case wherein two identical, preferred stimuli appear at two different locations in the receptive field of the neuron, and where the phase of the incoming spike is a means of representing spatial, rather than feature based attention. An important result not reproduced here but reported in Reynolds and Chelazzi (2004) is that feature based attention to a non-preferred stimulus when both a preferred and non-preferred stimulus are present in a receptive field *reduces* the firing rate of the neuron. Whether feature based and spatial attention are both subserved by phase locking to the γ oscillation, and whether such a scheme can reproduce the experimentally observed modulations in firing rate are important questions which need to be addressed by future research.

Our modeling was done with very strong phase locking (3 ms) for excitation and inhibition. In addition, phases did not drift as a function of time, and there were no so-called “noise” spikes not involved in the oscillation. While a 3 ms phase locking strength is not far out of the range of observed experimental phase locking to γ (Fries

et al., 2008), the rapid decrease in selectivity for phase lockings larger than 3 ms implies that the mechanism explored here is insufficiently robust to function in more realistic or noisy settings. One possibility is that despite selection decreasing rapidly in the single neuron case, ensembles of coupled neurons with similar preferences might have better selective profiles because of recurrent connectivity, or that selection might still be detectable in the activity of the ensemble by averaging. Generalizing these results to networks of neurons will help answer these questions. Since these results do not depend critically on the use of a compartmental model, studying them at the network level with single compartment neurons may be sufficient.

6.6 Conclusions

We present several results based on the analysis of time varying data. First, methods for finding and quantifying spike patterns are developed. These methods are then applied to the problem of encoding the amplitude or offset of time varying signals, where we find that regions in the response plane where multiple response-types are possible encode the amplitude or offset of a time varying signal with the greatest fidelity.

I also analyze simulations of two different strategies for selecting between time-varying stimuli simultaneously driving a single neuron. Hypothesizing that clustering of synaptic inputs in sub-domains of the dendritic arbor of pyramidal cells might be evidence of selection by selective application of inhibition, we found that the input-output relation for a pyramidal neuron is mostly determined by the total balance of inhibition and excitation, and that the location specificity of inputs is generally a small effect which does not significantly affect the selection of time-varying signals. Despite this result, we report that locating inhibition on dendritic sub-branches which are not being stimulated by excitation lowers the output firing more effectively over long time periods than locating it on the same dendrite as excitation.

Selection by phase-locking of excitatory input to an ongoing local inhibitory oscillation was also examined. In contrast to location specific inhibition, phase locking of excitatory inputs produced strong modulations in the effect of a time varying excitatory drive to control the output firing rate of a neuron. However, precise phase locking was necessary to achieve good modulation.

Taken together, these results and techniques represent a significant contribution to the analysis of variability in neural systems and to the understanding of the roles of spike timing and time varying signals in the neural code.

Bibliography

- Ahmed, B., Anderson, J. C., Douglas, R. J., Martin, K. A., and Nelson, J. C. (1994). Polyneuronal innervation of spiny stellate neurons in cat visual cortex. *J Comp Neurol*, 341(1):39–49.
- Archie, K. and Mel, B. (2000). Dendritic compartmentalization could underlie competition and attentional biasing of simultaneous visual stimuli. *NIPS*, pages 82–88.
- Arfken, G. (1985). *Mathematical Methods for Physicists*. Academic Press, third edition.
- Aronov, D. (2003). Fast algorithm for the metric-space analysis of simultaneous responses of multiple single neurons. *Journal of neuroscience methods*, 124(2):175–180.
- Austin, M., Mitchell, P., and Goodwin, G. (2001). Cognitive deficits in depression: possible implications for functional neuropathology. *The British Journal of Psychiatry*, 178(3):200.
- Bair, W. and Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput*, 8(6):1185–202.
- Bezdek, J. (1981). *Pattern recognition with fuzzy objective function algorithms*. Plenum, New York.
- Bouguessa, M., Wang, S. R., and Sun, H. J. (2006). An objective approach to cluster validation. *Pattern Recognition Letters*, 27(13):1419–1430.
- Bower, J. M. and Beeman, D. (1998). *The book of GENESIS (2nd ed.): exploring realistic neural models with the GEneral NEural SIMulation System*. Springer-Verlag New York, Inc., New York, NY, USA.
- Bruno, R. M. and Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science*, 312(5780):1622–7.
- Bryant, H. L. and Segundo, J. P. (1976). Spike initiation by transmembrane current: a white-noise analysis. *J Physiol*, 260(2):279–314.
- Buia, C. and Tiesinga, P. (2006). Attentional modulation of firing rate and synchrony in a model cortical network. *Journal of Computational Neuroscience*, 20(3):247–264.
- Buracas, G. T., Zador, A. M., DeWeese, M. R., and Albright, T. D. (1998). Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron*, 20(5):959–69.

- Butts, D. A., Weng, C., Jin, J., Yeh, C. I., Lesica, N. A., Alonso, J. M., and Stanley, G. B. (2007). Temporal precision in the neural code and the timescales of natural vision. *Nature*, 449(7158):92–5.
- Buzsaki, G. (2006). *Rhythms of the brain*. Oxford University Press, Oxford.
- Buzsaki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304(5679):1926–1929.
- Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annu Rev Neurosci*, 21:47–74.
- Cardin, J. A., Carlen, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L. H., and Moore, C. I. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, 459(7247):663–7.
- Carnevale, N. and Hines, M. (2006). *The NEURON Book*. Cambridge University Press, Cambridge, UK.
- Chapman, B., Zahs, K., and Stryker, M. (1991). Relation of cortical cell orientation selectivity to alignment of receptive fields of the geniculocortical afferents that arborize within a single orientation column in ferret visual cortex. *J. Neurosci.*, 11(5):1347–1359.
- Churchland, P. and Sejnowski, T. (1996). *The computational brain*. MIT press.
- Cover, T. and Thomas, J. (1991). *Elements of Information Theory*. Wiley, New-York.
- Cruikshank, S. J., Lewis, T. J., and Connors, B. W. (2007). Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex. *Nat Neurosci*, 10(4):462–8.
- de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K., and Reyes, A. (2007). Correlation between neural spike trains increases with firing rate. *Nature*, 448(7155):802–6.
- DeFelipe, J. and Farinas, I. (1992). The pyramidal neuron of the cerebral cortex: Morphological and chemical characteristics of the synaptic inputs. *Progress in Neurobiology*, 39:563–607.
- Denk, W., Strickler, J., and Webb, W. (1990). Two-photon laser scanning microscopy. *Science*, 248(1):73–76.
- Desimone, R. and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu Rev Neurosci*, 18:193–222.
- Deweese, M. R. and Zador, A. M. (2004). Shared and private variability in the auditory cortex. *J Neurophysiol*, 92(3):1840–55.

- Diamond, M., Scheibel, A., and Elson, L. (1985). *The Human Brain Coloring Book*. Harper Collins, New York, New York.
- Douglas, R. J. and Martin, K. A. (2004). Neuronal circuits of the neocortex. *Annu Rev Neurosci*, 27:419–51.
- Duda, R., Hart, P., and Stork, D. (2001). *Pattern classification*. Wiley, 2nd edition.
- Eccles, J. (1976). From electrical to chemical transmission in the central nervous system. *Notes and Records of the Royal Society of London*, pages 219–230.
- Elhilali, M., Fritz, J. B., Klein, D. J., Simon, J. Z., and Shamma, S. A. (2004). Dynamics of precise spike timing in primary auditory cortex. *J Neurosci*, 24(5):1159–72.
- Ermentrout, G., Galán, R., and Urban, N. (2008). Reliability, synchrony and noise. *Trends in Neurosciences*, 31(8):428–434.
- Fellous, J. M., Houweling, A. R., Modi, R. H., Rao, R. P., Tiesinga, P. H., and Sejnowski, T. J. (2001). Frequency dependence of spike timing reliability in cortical pyramidal cells and interneurons. *J Neurophysiol*, 85(4):1782–7.
- Fellous, J. M., Tiesinga, P. H., Thomas, P. J., and Sejnowski, T. J. (2004). Discovering spike patterns in neuronal responses. *J Neurosci*, 24(12):2989–3001.
- Ferster, D. and Miller, K. D. (2000). Neural mechanisms of orientation selectivity in the visual cortex. *Annu Rev Neurosci*, 23:441–71.
- Fields, R. and Stevens-Graham, B. (2002). Neuroscience: new insights into neuron-glia communication. *Science*, 298(5593):556.
- Foss, J. and Milton, J. (2000). Multistability in recurrent neural loops arising from delay. *J Neurophysiol*, 84(2):975–85.
- Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10):474–480.
- Fries, P., Reynolds, J., Rorie, A., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291(5508):1560–1563.
- Fries, P., Womelsdorf, T., Oostenveld, R., and Desimone, R. (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *Journal of Neuroscience*, 28(18):4823.
- Fujisawa, S., Amarasingham, A., Harrison, M., and Buzsaki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience*.
- Gasparini, S. and Magee, J. C. (2006). State-dependent dendritic computation in hippocampal ca1 pyramidal neurons. *J Neurosci*, 26(7):2088–100.

- Georgopoulos, A., Kettner, R., and Schwartz, A. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience*, 8(8):2928.
- Gerald, C. and Wheatley, P. (1999). *Applied Numerical Analysis*. Addison-Wesley, Reading, California, 6th edition.
- Golding, N., Jung, H., Mickus, T., and Spruston, N. (1999). Dendritic calcium spike initiation and repolarization are controlled by distinct potassium channel subtypes in CA1 pyramidal neurons. *Journal of Neuroscience*, 19(20):8789.
- Golomb, D. and Amitai, Y. (1997). Propagating neuronal discharges in neocortical slices: computational and experimental study. *Journal of Neurophysiology*, 78(3):1199–1211.
- Green, D. and Swets, J. (1966). *Signal Detection theory and psychophysics*. John Wiley and sons, New York.
- Haider, B. and McCormick, D. (2009). Rapid Neocortical Dynamics: Cellular and Network Mechanisms. *Neuron*, 62(2):171–189.
- Han, X., Qian, X., Bernstein, J., Zhou, H.-h., Franzesi, G., Stern, P., Bronson, R., Graybiel, A., Desimone, R., and Boyden, E. (2009). Millisecond-timescale optical control of neural dynamics in the nonhuman primate brain. *Neuron*, 62:191198.
- Hodgkin, A., Huxley, A., and Katz, B. (1952). Measurement of current-voltage relations in the membrane of the giant squid axon of loligo. *Journal of Physiology*, 116:424–448.
- Horton, J. C. and Adams, D. L. (2005). The cortical column: a structure without a function. *Philos Trans R Soc Lond B Biol Sci*, 360(1456):837–62.
- Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J Physiol*, 160:106–54.
- Hutchinson, A. and Mathias, J. (2007). Neuropsychological deficits in frontotemporal dementia and Alzheimer’s disease: a meta-analytic review. *British Medical Journal*.
- Iyer, V., Hoogland, T. M., and Saggau, P. (2006). Fast functional imaging of single neurons using random-access multiphoton (ramp) microscopy. *J Neurophysiol*, 95(1):535–45.
- Jadhav, S., Wolfe, J., and Feldman, D. (2009). Sparse temporal coding of elementary tactile features during active whisker sensation. *Nature Neuroscience*, 12:792–800.
- Johnston, D. and Wu, S. (1995). *Foundations of Cellular Neurophysiology*. MIT Press, Cambridge, MA, 2nd edition.

- Jolivet, R., Lewis, T. J., and Gerstner, W. (2004). Generalized integrate-and-fire models of neuronal activity approximate spike trains of a detailed model to a high degree of accuracy. *J Neurophysiol*, 92(2):959–76.
- Jolivet, R., Rauch, A., Luscher, H. R., and Gerstner, W. (2006). Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *J Comput Neurosci*, 21(1):35–49.
- Jolliffe, I. (2002). *Principal component analysis*. Springer, New York, 2nd edition.
- Joris, P., Schreiner, C., and Rees, A. (2004). Neural processing of amplitude-modulated sounds. *Physiological reviews*, 84(2):541–577.
- Kara, P., Pezaris, J. S., Yurgenson, S., and Reid, R. C. (2002). The spatial receptive field of thalamic inputs to single cortical simple cells revealed by the interaction of visual and electrical stimulation. *Proc Natl Acad Sci U S A*, 99(25):16261–6.
- Kara, P. and Reid, R. C. (2003). Efficacy of retinal spikes in driving cortical responses. *J Neurosci*, 23(24):8547–57.
- Kara, P., Reinagel, P., and Reid, R. C. (2000). Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, 27(3):635–46.
- Keat, J., Reinagel, P., Reid, R. C., and Meister, M. (2001). Predicting every spike: a model for the responses of visual neurons. *Neuron*, 30(3):803–17.
- Koch, C. (1997). Computation and the single neuron. *Nature*, 385(6613):207–10.
- Koch, C. (2004a). *Biophysics of Computation: Information Processing in Single Neurons (Computational Neuroscience Series)*. Oxford University Press, Inc. New York, NY, USA.
- Koch, C. (2004b). *The Quest for Consciousness*. Roberts and Company, Englewood, Colorado, first edition.
- Kruskal, P. B., Stanis, J. J., McNaughton, B. L., and Thomas, P. J. (2007). A binless correlation measure reduces the variability of memory reactivation estimates. *Stat Med*, 26(21):3997–4008.
- Larkum, M., Zhu, J., and Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature*, 398(6725):338–341.
- Larsen, R. and Marx, M. (1986). *An introduction to mathematical statistics and its applications*. Prentice Hall, Englewood Cliffs, New Jersey, 2nd edition.
- Lindsey, B. and Gerstein, G. (2006). Two enhancements of the gravity algorithm for multiple spike train analysis. *Journal of neuroscience methods*, 150(1):116–127.

- Lubke, J. and Feldmeyer, D. (2007). Excitatory signal flow and connectivity in a cortical column: focus on barrel cortex. *Brain Struct Funct*, 212(1):3–17.
- Ly, C. and Doiron, B. (2009). Divisive gain modulation with dynamic stimuli in integrate-and-fire neurons. *PLoS Comput Biol*, 5(4):e1000365.
- Mainen, Z. F. and Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–6.
- Mainen, Z. F. and Sejnowski, T. J. (1996). Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature*, 382(6589):363–6.
- Markowitz, D. A., Collman, F., Brody, C. D., Hopfield, J. J., and Tank, D. W. (2008). Rate-specific synchrony: using noisy oscillations to detect equally active neurons. *Proc Natl Acad Sci U S A*, 105(24):8422–7.
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*, 5(10):793–807.
- Matlab (2009). The MathWorks Inc. *Natick, MA*.
- McCormick, D. A., Connors, B. W., Lighthall, J. W., and Prince, D. A. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol*, 54(4):782–806.
- Mellwain, J. (1996). *An introduction to the biology of vision*. Cambridge Univ Pr.
- Meir, A., Ginsburg, S., Butkevich, A., Kachalsky, S., Kaiserman, I., Ahdut, R., Demirgoren, S., and Rahamimoff, R. (1999). Ion channels in presynaptic nerve terminals and control of transmitter release. *Physiological reviews*, 79(3):1019–1088.
- Mesulam, M. (1998). From sensation to cognition. *Brain*, 121(6):1013.
- Migliore, M., Cook, E., Jaffe, D., Turner, D., and Johnston, D. (1995). Computer simulations of morphologically reconstructed CA3 hippocampal neurons. *Journal of neurophysiology*, 73(3):1157–1168.
- Mishra, J., Fellous, J. M., and Sejnowski, T. J. (2006). Selective attention through phase relationship of excitatory and inhibitory input synchrony in a model cortical neuron. *Neural Netw*, 19(9):1329–46.
- Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain*, 120 (Pt 4):701–22.
- Murayama, M., Perez-Garci, E., Nevian, T., Bock, T., Senn, W., and Larkum, M. E. (2009). Dendritic encoding of sensory stimuli controlled by deep cortical interneurons. *Nature*, 457(7233):1137–41.

- Ogawa, T. and Komatsu, H. (2004). Target selection in area V4 during a multidimensional visual search task. *Journal of Neuroscience*, 24(28):6371.
- Pal, N. R. and Bezdek, J. C. (1995). On cluster validity for the fuzzy c-means model. *Ieee Transactions on Fuzzy Systems*, 3(3):370–379.
- Panzeri, S. and Treves, A. (1995). Analytical estimates of limited sampling biases in different information measures. *Network: Comput. Neural Syst*, 7:87–107.
- Pare, D., Shink, E., Gaudreay, H., Destexhe, A., and Lang, E. (1998). Impact of spontaneous activity on the resting properties of cat neocortical pyramidal neurons in vivo. *Journal of Neurophysiology*, 79:1450–1460.
- Pasupathy, A. and Connor, C. (1999). Responses to contour features in macaque area V4. *Journal of Neurophysiology*, 82(5):2490–2502.
- Petreaunu, L., Mao, T., S.M., S., and Svoboda, K. (2009). The subcellular organization of neocortical excitatory connections. *Nature*, 457:1142–1146.
- Pillow, J. W., Paninski, L., Uzzell, V. J., Simoncelli, E. P., and Chichilnisky, E. J. (2005). Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *J Neurosci*, 25(47):11003–13.
- Poirazi, P., Brannon, T., and Mel, B. W. (2003). Arithmetic of subthreshold synaptic summation in a model cal pyramidal cell. *Neuron*, 37(6):977–87.
- Polsky, A., Mel, B. W., and Schiller, J. (2004). Computational subunits in thin dendrites of pyramidal cells. *Nat Neurosci*, 7(6):621–7.
- Poulet, J. F. and Petersen, C. C. (2008). Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature*, 454(7206):881–5.
- Powers, R. K., Dai, Y., Bell, B. M., Percival, D. B., and Binder, M. D. (2005). Contributions of the input signal and prior activation history to the discharge behaviour of rat motoneurons. *J Physiol*, 562(Pt 3):707–24.
- Prescott, S. and De Koninck, Y. (2003). Gain control of firing rate by shunting inhibition: roles of synaptic noise and dendritic saturation. *Proceedings of the National Academy of Sciences*, 100(4):2076–2081.
- Press, W., Teukolsky, S., Vetterling, W., and Flannery, B. (1992). *Numerical Recipes*. Cambridge University Press, Cambridge.
- Reid, R. C. and Alonso, J. M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378(6554):281–4.
- Reinagel, P., Godwin, D., Sherman, S. M., and Koch, C. (1999). Encoding of visual information by lgn bursts. *J Neurophysiol*, 81(5):2558–69.

- Reinagel, P. and Reid, R. C. (2000). Temporal coding of visual information in the thalamus. *J Neurosci*, 20(14):5392–400.
- Reinagel, P. and Reid, R. C. (2002). Precise firing events are conserved across neurons. *J Neurosci*, 22(16):6837–41.
- Reyes, A. D. (2003). Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nat Neurosci*, 6(6):593–9.
- Reynolds, J., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, 19(5):1736–1753.
- Reynolds, J. and Desimone, R. (2003). Interacting roles of attention and visual salience in V4. *Neuron*, 37(5):853–863.
- Reynolds, J. H. and Chelazzi, L. (2004). Attentional modulation of visual processing. *Annu Rev Neurosci*, 27:611–47.
- Reynolds, J. H. and Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61(2):168–85.
- Rezaee, M. R., Lelieveldt, B. P. F., and Reiber, J. H. C. (1998). A new cluster validity index for the fuzzy c-mean. *Pattern Recognition Letters*, 19(3-4):237–246.
- Richmond, B. J., Optican, L. M., Podell, M., and Spitzer, H. (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. i. response characteristics. *J Neurophysiol*, 57(1):132–46.
- Riehle, A., Grun, S., Diesmann, M., and Aertsen, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278(5345):1950.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., and Bialek, W. (1997). *Spikes: exploring the neural code*. MIT press, Cambridge.
- Rizzo, M. (1997). stochastic Resonance: A New Concept of Neural Coding. *The Neuroscientist*, 3(4):211.
- Rohatgi, V. (2003). *Statistical Inference*. Dover Publications, Mineola, New York.
- Salinas, E. and Sejnowski, T. (2000). Impact of correlated synaptic input on output variability in simple neuronal models. *J Neurosci*, 20:6193–6209.
- Salinas, E. and Sejnowski, T. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2(8):539–550.

- Satzler, K., Sohl, L., Bollmann, J., Borst, J., Frotscher, M., Sakmann, B., and Lubke, J. (2002). Three-dimensional reconstruction of a calyx of Held and its postsynaptic principal neuron in the medial nucleus of the trapezoid body. *Journal of Neuroscience*, 22(24):10567.
- Schaefer, A., Helmstaedter, M., Schmitt, A., Bar-Yehuda, D., Almog, M., Ben-Porat, H., Sakmann, B., and Korngreen, A. (2007). Dendritic voltage-gated K⁺ conductance gradient in pyramidal neurones of neocortical layer 5B from rats. *The Journal of Physiology*, 579(3):737.
- Schreiber, S., Fellous, J. M., Tiesinga, P., and Sejnowski, T. J. (2004). Influence of ionic conductances on spike timing reliability of cortical neurons for suprathreshold rhythmic inputs. *J Neurophysiol*, 91(1):194–205.
- Schreiber, S., Whitmer, D., Fellous, J. M., Tiesinga, P., and Sejnowski, T. J. (2003). A new correlation-based measure of spike timing reliability. *Neurocomputing*, 52-54:925–931.
- Shadlen, M. and Newsome, W. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.*, 18:3870–3896.
- Siever, L. and Davis, K. (2004). The pathophysiology of schizophrenia disorders: perspectives from the spectrum. *American Journal of Psychiatry*, 161(3):398.
- Sohal, V. S., Zhang, F., Yizhar, O., and Deisseroth, K. (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, 459(7247):698–702.
- Somers, D., Nelson, S., and Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, 15(8):5448–5465.
- Steriade, M. (2004). Neocortical cell classes are flexible entities. *Nat Rev Neurosci*, 5(2):121–34.
- Strong, S., Koberle, R., de Ruyter van Stevenick, R., and Bialek, W. (1998). Entropy and information in neural spike trains. *Phys Rev Letts*, 80:197–2000.
- Tanaka, K. (1983). Cross-correlation analysis of geniculostriate neuronal relationships in cats. *J Neurophysiol*, 49(6):1303–18.
- Taylor, K., Mandon, S., Freiwald, W., and Kreiter, A. (2005). Coherent oscillatory activity in monkey area v4 predicts successful allocation of attention. *Cerebral Cortex*, 15(9):1424–1437.
- Teich, A. and Qian, N. (2006). Comparison among some models of orientation selectivity. *J. Neurophys.*, 96:404–419.

- Tibshirani, R., Walther, G., and Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society Series B-Statistical Methodology*, 63:411–423.
- Tiesinga, P. (2004). Chaos-induced modulation of reliability boosts output firing rate in downstream cortical areas. *Physical Review E*, 69(3):31912.
- Tiesinga, P., Fellous, J., and Sejnowski, T. (2008). Regulation of spike timing in visual cortical circuits. *Nature Reviews Neuroscience*, 9(2):97–107.
- Tiesinga, P. H., Fellous, J. M., Jose, J. V., and Sejnowski, T. J. (2002). Information transfer in entrained cortical neurons. *Network*, 13(1):41–66.
- Tiesinga, P. H. E. and Jose, J. V. (2000). Robust gamma oscillations in networks of inhibitory hippocampal interneurons. *Network-Computation in Neural Systems*, 11(1):1–23.
- Tiesinga, P. H. E. and Toups, J. V. (2005). The possible role of spike patterns in cortical information processing. *Journal of Computational Neuroscience*, 18(3):275–286.
- Traub, R., Buhl, E., Gloveli, T., and Whittington, M. (2003). Fast rhythmic bursting can be induced in layer 2/3 cortical pyramidal neurons by enhancing persistent Na^+ conductance or by blocking K^+ channels. *J. Neurophysiol.*, 89:909–921.
- Vale, S. (2008). Current management of the Cognitive Dysfunction in Parkinson’s disease: how far have we come? *Experimental Biology and Medicine*, 233(8):941.
- Van Rossum, M. (2001). A novel spike distance. *Neural Computation*, 13(4):751–764.
- Victor, J. D. and Purpura, K. P. (1996). Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J Neurophysiol*, 76(2):1310–26.
- Vogels, T. and Abbott, L. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nature Neuroscience*, 12(4):483–491.
- Wang, X. J. and Buzsaki, G. (1996). Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J Neurosci*, 16(20):6402–13.
- Wiskott, L., Fellous, J. M., Kruger, N., and von der Malsburg, C. (1997). Face recognition by elastic bunch graph matching. *Ieee Transactions on Pattern Analysis and Machine Intelligence*, 19(7):775–779.
- Womelsdorf, T., Schoffelen, J. M., Oostenveld, R., Singer, W., Desimone, R., Engel, A. K., and Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. *Science*, 316(5831):1609–12.
- Yan, M. and Ye, K. (2007). Determining the number of clusters using the weighted gap statistic. *Biometrics*, 63(4):1031–1037.

- Zahid, N., Limouri, N., and Essaid, A. (1999). A new cluster-validity for fuzzy clustering. *Pattern Recognition*, 32(7):1089–1097.
- Zhang, F., Aravanis, A. M., Adamantidis, A., de Lecea, L., and Deisseroth, K. (2007a). Circuit-breakers: optical technologies for probing neural signals and systems. *Nat Rev Neurosci*, 8(8):577–81.
- Zhang, F., Wang, L. P., Brauner, M., Liewald, J. F., Kay, K., Watzke, N., Wood, P. G., Bamberg, E., Nagel, G., Gottschalk, A., and Deisseroth, K. (2007b). Multimodal fast optical interrogation of neural circuitry. *Nature*, 446(7136):633–9.