

**IDENTIFICATION AND CONTROL OF NEURAL
CIRCUIT DYNAMICS FOR NATURAL AND
SURROGATE INPUTS IN-VIVO**

A Thesis
Presented to
The Academic Faculty

by

Daniel C Millard

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Department of Biomedical Engineering

Georgia Institute of Technology
May 2014

Copyright © 2014 by Daniel C Millard

IDENTIFICATION AND CONTROL OF NEURAL CIRCUIT DYNAMICS FOR NATURAL AND SURROGATE INPUTS IN-VIVO

Approved by:

Professor Garrett Stanley, Advisor
Department of Biomedical Engineering
Georgia Institute of Technology

Professor Chris Rozell
Department of Electrical and
Computer Engineering
Georgia Institute of Technology

Professor Steve Potter
Department of Biomedical Engineering
Georgia Institute of Technology

Professor Robert Gross
Departments of Neurosurgery and
Neurology
Emory University

Professor Warren Grill
Department of Biomedical Engineering
Duke University

Date Approved: December 18, 2013

To Darcy and Theo.

ACKNOWLEDGEMENTS

Any significant body of work must depend on a host of individuals, and this thesis was no exception. Without the knowledge and support of the wonderful people listed below, this journey would have been nigh impossible.

First, I would like to thank my family. My parents provided guidance when prompted, but otherwise let me find my own path, while always providing support to push me along. My brother and sister have always kept me grounded, and made life fun along the way.

I would also like to thank some important people at UNC-CH that greatly influenced my decision to pursue graduate research. Sean Washburn provided the initial spark as the professor of my first engineering course and was a superb mentor throughout the rest of my undergraduate tenure. Hugon Karwowski displayed a passion and enthusiasm for teaching and research that was truly inspirational. Rich Goldberg provided many wonderful discussions and taught me the meaning of dedication and hard work. Finally, my undergraduate research advisor, Michael Falvo, played a crucial role in my development as a scientist, and more importantly as a person.

While at Georgia Tech, I was fortunate to have the opportunity to work and learn from an equally influential group of people. Even though my thesis committee was only formally in place for a year, each of the members provided a lasting impact. I would like to thank Steve Potter for his collaborative spirit and technical expertise, Chris Rozell for keeping me honest with the technical portions of this work and emphasizing the importance of professional skill development, Robert Gross for keeping me on my toes and forcing me to evaluate my work from a patient's perspective, and Warren Grill for always putting in the effort from a distance and providing an

exquisite example of how science and industry can work together to improve health.

I think the best part of graduate school was that I learned as much, if not more, from spending time with my fellow graduate students in the trenches as I did from the esteemed faculty. I truly appreciate every conversation, insight, experiment, and escape that I shared with everyone at Tech. To name a few: Jon, Jeff, Akhil, Yogi, Neal, Alex, J.P., Chris, Brock, Melissa, Becca, Abbie, Ming-fai, Mengchen, and John.

The same can be said for the members of the Stanley lab: Qi, Gaelle, Doug, Sean, Clare, Lauren, He, Peter, Clarissa, Bilal, and Will. There wasn't a boring day in the lab with these guys around, and much of the work in this thesis was entirely dependent on their contributions. I truly can't wait to see what the future holds for this talented group of individuals. I would like to specifically thank Qi Wang. In my early years at Georgia Tech he was like a second advisor to me, and for that I am truly grateful. He is a fantastic engineer and I wish him the best as a new professor at Columbia University.

And I must thank my advisor, Garrett Stanley. There are no words to fully describe the influence Garrett has had on me as a scientist, engineer, future-parent, and all around human being. I truly can't thank him enough. Throughout my time in his lab, he would always remind us that after discovering the structure of DNA Francis Crick decided to work on a truly difficult problem - studying the brain. I'd like to think that Garrett and I made a small dent in understanding the brain with the work presented here. And if the world had a few more people like Garrett, I think we might actually have a shot at figuring out how the brain works.

Finally, I need to thank my wife, Darcy Rice Millard. She has been an absolute rock throughout my time in graduate school, providing me something to grab onto when life got rough. She put up with many late night experiments, long conference trips away from home, my thoughts constantly drifting away to my research, and endless weekend data crunching. She is an amazing woman and I know that together

we can overcome all of life's challenges. I look forward to the next stage of this journey called life with her by my side.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
LIST OF TABLES	xii
LIST OF FIGURES	xiii
SUMMARY	xv
I INTRODUCTION	1
1.1 Writing the neural code	1
1.2 Cochlear implant as a model for sensory prostheses	2
1.3 Electrical and optogenetic stimuli as surrogate sensory stimuli	4
1.4 Natural and artificial stimuli activate neural circuits in distinct ways	6
1.5 Direct neural readout of nonlinear propagation in neural circuits	8
1.6 Experimental framework for the design of functionally relevant surrogate sensory inputs	10
1.7 Organization of the thesis	15
II CHARACTERIZATION OF THE NONLINEAR SENSITIVITY OF THE CORTICAL RESPONSE TO SURROGATE THALAMIC DRIVE	17
2.1 Introduction	17
2.2 Methods	19
2.2.1 Experimental Preparation	19
2.2.2 Voltage Sensitive Dye Imaging	19
2.2.3 Whisker Stimulation	20
2.2.4 Electrical Stimulation	20
2.2.5 Optical Stimulation	21
2.2.6 Computational Network Model	22
2.3 Results	26

2.3.1	The cortical response is highly nonlinear for whisker, electrical, and optical stimuli	26
2.3.2	Whisker stimuli are more sensitive to trial to trial variability	27
2.3.3	The cortical response to microstimulation exhibits the greatest spatial spread	29
2.3.4	Computational model of thalamocortical circuit allows direct control of thalamic input	32
2.3.5	Magnitude and synchrony of population activity control nonlinear response in network model	35
2.4	Discussion	40
III	CHARACTERIZATION OF THE NONLINEAR DYNAMICS OF CORTICAL ACTIVATION IN RESPONSE TO SURROGATE THALAMIC DRIVE	45
3.1	Introduction	45
3.2	Methods	47
3.2.1	Surgery and preparation	47
3.2.2	Thalamic microstimulation	47
3.2.3	Whisker stimulation	47
3.3	Results	48
3.3.1	Nonlinear paired pulse suppression of the cortical response to whisker stimuli	48
3.3.2	Stimulus-dependent bimodal nonlinear dynamics of the cortical response to thalamic microstimulation	51
3.3.3	Facilitation from sub-threshold thalamic microstimulation is a general phenomenon	54
3.3.4	Nonlinear facilitation sharpens the spatial cortical response .	56
3.3.5	Optogenetic stimulation of the thalamus also exhibits bimodal nonlinear dynamics	60
3.4	Discussion	63
IV	SYSTEM IDENTIFICATION OF THE NONLINEAR DYNAMICS IN THE CORTICAL RESPONSE TO THALAMIC MICROSTIMULATION	68
4.1	Introduction	68

4.2	Methods	70
4.2.1	Experimental Preparation	70
4.2.2	Electrical Stimulation	71
4.2.3	System Identification	72
4.2.4	Extension of the phenomenological model to include space	80
4.2.5	Model Simulations	81
4.3	Results	82
4.3.1	The cortical response to thalamic microstimulation is highly nonlinear	82
4.3.2	Development of a phenomenological model based on experimental observations	84
4.3.3	Error residuals illustrate the improved performance of the phenomenological model	85
4.3.4	Trial to trial variability and feedback	87
4.3.5	Activity propagation through electrical microstimulation	93
4.3.6	Linear point spread function captures spatial spread in cortex	95
4.4	Discussion	96
V	DESIGN OF SURROGATE INPUTS TO IMPROVE DISCRIMINABILITY OF THE DOWNSTREAM CORTICAL RESPONSE	104
5.1	Introduction	104
5.2	Methods	107
5.2.1	Surgery and preparation	107
5.2.2	Thalamic microstimulation	108
5.2.3	Whisker stimulation	108
5.2.4	VSD response analysis	109
5.2.5	Functional Coding Analysis	110
5.3	Stimulus waveform design augments spatial response properties	114
5.3.1	Functional validation of thalamic electrode placement.	114
5.3.2	Thalamic microstimulation results in a systematic topographical bias in cortical activation.	115

5.3.3	Cathode-leading asymmetric microstimulation significantly improved the topographic activation of cortex.	116
5.3.4	The topographic alignment of the cortical activation was enhanced with increasing asymmetry of the microstimulation waveform.	119
5.3.5	Cathode-leading asymmetric waveform also improved the specificity of thalamic microstimulation.	119
5.4	Optimal signal set design maximized response discriminability . . .	123
5.4.1	Signal set design accounts for distinct static nonlinear properties	124
5.4.2	Functional consequences of nonlinear properties of activity propagation	127
5.5	Discussion	129

VI CONTROL OF THE NONLINEAR DYNAMICS OF CORTICAL ACTIVATION THROUGH OPTOGENETIC MANIPULATION OF THALAMIC STATE 136

6.1	Introduction	136
6.2	Methods	138
6.2.1	Surgery and preparation	138
6.2.2	Optogenetic expression	138
6.2.3	Optical stimulation	138
6.2.4	Thalamic microstimulation	139
6.2.5	Electrophysiology analysis	139
6.2.6	VSD analysis	140
6.3	Results	140
6.3.1	Thalamic depolarization leads to desynchronized state in-vivo	141
6.3.2	Depolarization reduces spontaneous and stimulus-driven bursting	141
6.3.3	State control of the cortical response to thalamic microstimulation	147
6.3.4	Depolarization of the thalamus reduces trial-to-trial variability in the cortical response	147

6.3.5	Depolarization of the thalamus eliminated nonlinear facilitation for thalamic micrstimulation	149
6.4	Discussion	150
VII CONCLUSION AND FUTURE WORK		156
7.1	Application to sensory prostheses	163
7.2	Applications towards microstimulation for other disorders	166
7.3	Interfacing with neural circuits: a framework for control	167
APPENDIX A — GENERAL VSD METHODS		170
APPENDIX B — ANATOMICALLY BASED BAYESIAN DECODING OF THE CORTICAL RESPONSE TO INTRACORTICAL MICROSTIMULATION		177
REFERENCES		187
VITA		206

LIST OF TABLES

2.1 Parameters for Izhikevich cortical neurons	23
--	----

LIST OF FIGURES

1.1	Block diagram of a central nervous system sensory prosthesis	4
1.2	Architecture of the rodent vibrissa somatosensory system	11
1.3	Functional validation of thalamic electrode placement	13
1.4	Voltage sensitive dye imaging of cortical network activity	14
2.1	Whisker stimuli are more sensitive to trial-to-trial variability	28
2.2	Electrical stimulation activates a significantly larger region of cortex .	31
2.3	Computational network model of activity propagation in the thalamo-cortical circuit	34
2.4	Extreme synchrony ensures robust response in the face of trial-to-trial variability	36
2.5	Axonal stimulation increases spatial spread by overcoming feedforward inhibition	38
3.1	Paired pulse suppression of the cortical response to whisker stimuli .	50
3.2	Bimodal nonlinear dynamics of the cortical response to thalamic microstimulation	52
3.3	Thalamic microstimulation facilitates whisker and electrical inputs . .	55
3.4	Nonlinear facilitation sharpens spatial cortical response	57
3.5	Facilitated cortical response to microstimulation activates same area as whisker stimulation	59
3.6	Optogenetic stimulation of the thalamus exhibits bimodal nonlinear dynamics	61
3.7	Facilitation occurs at the level of VPM for optogenetic stimuli	62
4.1	Microstimulation of the thalamus produces a nonlinear cortical response	73
4.2	Nonlinear modeling architecture	77
4.3	Phenomenological model based on experimental observations accurately predicted the cortical response to patterned microstimulation	86
4.4	Phenomenological model improved error residuals compared to the Volterra model	88
4.5	Phenomenological model with feedback reproduced single trial variability in facilitation and suppression	91

4.6	Model parameters predict linear local response properties, but nonlinear propagation of activity	94
4.7	Extension of phenomenological model captures spatial spread	97
5.1	Thalamic microstimulation with symmetric waveforms results in a systematic topographical bias in cortical activation	117
5.2	Cathode-leading asymmetric microstimulation significantly improved the topographic activation of cortex	120
5.3	The topographic alignment of the cortical activation was enhanced with increasing asymmetry of the microstimulation waveform	121
5.4	Cathode-leading asymmetric waveform also improved the specificity of thalamic microstimulation	122
5.5	Optimal signal set design can account for distinct static nonlinear properties	126
5.6	Functional consequences of trial-to-trial variability on response discriminability	128
6.1	Optogenetic control of thalamic state	142
6.2	Depolarization of the thalamus eliminates bursting	144
6.3	Continuous transition from burst to tonic firing modes	146
6.4	Thalamic depolarization reduces cortical trial to trial variability in response to thalamic microstimulation	148
6.5	Thalamic depolarization eliminates paired-pulse facilitation	151
7.1	Nonlinear propagation of activity across neural circuits	157
A.1	Voltage Sensitive dye imaging of cortical activation in response to thalamic activation and whisker deflection	172
A.2	Functional validation of thalamic electrode placement	173
B.1	Voltage sensitive dye imaging of intracortical microstimulation	179
B.2	Decoding of stimulus location	183
B.3	ICMS produces distinct activation compared to whisker inputs	184

SUMMARY

A principal goal of neural engineering is to control the activation of neural circuits across space and time. The ability to control neural circuits with surrogate inputs is needed for the development of clinical neural prostheses and the experimental interrogation of connectivity between brain regions. Electrical stimulation provides a clinically viable method for activating neural tissue and the emergence of optogenetic stimulation has redefined the limitations on stimulating neural tissue experimentally. However, it remains poorly understood how these tools activate complex neural circuits.

The goal of this proposed project was to gain a greater understanding of how to control the activity of neural circuits in-vivo using a combination of experimental and computational approaches. Voltage sensitive dye imaging was used to observe the spatiotemporal activity within the rodent somatosensory cortex in response to systematically varied patterns of sensory, electrical, and optogenetic stimulation. First, the cortical response to simple patterns of sensory and artificial stimuli was characterized and modeled, revealing distinct neural response properties due to the differing synchrony with which the neural circuit was engaged. Then, we specifically designed artificial stimuli to improve the functional relevance of the resulting downstream neural responses. Finally, through direct optogenetic modulation of thalamic state, we demonstrate control of the nonlinear propagation of neural activity within the thalamocortical circuit.

The combined experimental and computational approach described in this thesis provides a comprehensive description of the nonlinear dynamics of the thalamocortical

circuit to surrogate stimuli. Together, the characterization, modeling, and overall control of downstream neural activity stands to inform the development of central nervous system sensory prostheses, and more generally provides the initial tools and framework for the control of neural activity in-vivo.

CHAPTER I

INTRODUCTION

1.1 Writing the neural code

Neurons communicate with each other through an electrical language of sorts, often referred to as the neural code, comprised of patterns of action potentials across diverse populations of cells. Through the neural code, populations of neurons process and transmit information to downstream structures, ultimately resulting in perception [86], movement [63], memory [135], or decision making [158]. Even the activity of a single cell can contain significant information about the surrounding environment [86] or an ensuing movement [63], but more likely it is the activity across entire neural populations that encodes information reliably [24, 177, 62, 5]. Our ability to extract such information from the activity of neural populations is often described as “reading” the neural code [177, 12].

If patterns of neural activity underlie function and behavior, it should not be surprising that the majority of disorders of the central nervous system are the product of aberrant patterns of activity within and across cells, resulting in undesirable or deleterious effects on neural function [80, 35]. However, if we, through our understanding of the neural code, knew the difference between physiological and pathological neural activity, we might be able to design therapies to “write in” the normal pattern of activity. To “write” the neural code, we must already know how to “read” the neural code to a degree, and, in addition, we need tools to augment ongoing neural activity in the brain, such as drugs, electrical stimulation, or more recently optogenetic stimulation. The overarching goal of this thesis was to develop methods to “write” the neural code in the central nervous system, or more specifically to control patterns of

activity in the brain.

Currently, the treatment of neurological disorders primarily consists of the use of pharmaceutical agents to modulate neural activity. While the region of pathological activity may be small, pharmaceuticals non-specifically bathe the brain, producing diffuse modulation of neural activity in pathological and normal functioning circuits, which can lead to significant side effects. Other non-invasive neuromodulatory techniques, such as transcranial magnetic stimulation or even physical rehabilitation, operate on similarly slow timescales with low spatial specificity.

In stark contrast is the use of invasive macroelectrodes inserted into the brain to deliver electrical stimuli, providing the ability to modulate local neural activity on the timescale of milliseconds. Electrical neuromodulation has principally been used in the deep structures of the brain to treat the symptoms of Parkinson's disease, depression, and tremor, leading to the term "Deep Brain Stimulation". However, the most successful neurological intervention using electrical stimulation to date is the cochlear implant.

1.2 Cochlear implant as a model for sensory prostheses

The cochlear implant is a sensory prosthesis that uses patterns of electrical impulses to restore functional audition to the deaf [202]. The device truly "writes" the neural code as sound, picked up by a microphone worn by the individual, is mapped into a pattern of electrical stimuli across a number of electrodes interfacing with the auditory nerve in the cochlea. The cochlear implant is specifically designed to compensate for loss of function in the auditory hair cells within the cochlea, whereas more generally the goal of a sensory prosthesis is to restore sensory function that has been lost due to trauma or disease.

A block diagram of a generic sensory prosthesis is shown in figure 1.1. Under

normal sensory function, receptors in the peripheral nervous system transduce environmental cues into neuronal signals, which then propagate throughout the sensory pathway and eventually result in a percept. Damage at any point along the sensory pathway obstructs the normal propagation of neural activity, requiring the introduction of surrogate sensory information through artificial stimulation downstream of the lost function. Surrogate sensory information is produced by sampling the environment with sensors, extracting salient features through online processing, and encoding the sensory information into spatiotemporal patterns of artificial stimulation in the peripheral or central nervous system.

The success of the cochlear implant has been crucially dependent on the advancement in the “encoding” stage of the prosthesis, where the sounds are mapped into patterned stimulation across the electrode array within the cochlea. Early iterations provided proof of concept in human patients using only a single channel and very simple encoding schemes. However, after decades of animal [125, 126] and human experiments [13] assessing the mapping of auditory nerve stimulation to the downstream neural response and resulting percept, more efficient and effective encoding schemes have been developed [201].

More recently, a retinal prosthesis is under development to restore vision to those suffering from blindness caused by retinitis pigmentosa [87]. But, implementation of a sensory prosthesis using electrodes implanted in the central nervous system has not yet been achieved, likely due to the increased difficulty of interfacing with neural circuits in the brain, but also because of the heightened, and largely uncharacterized, complexity of the mapping from artificial stimuli to neural response. Whether the aim is to reproduce natural neural activity or merely to deliver discriminable inputs to the brain, the advancement of sensory prostheses requires a greater understanding of the mapping from electrical stimuli to neural response within complex circuits and the resulting propagation along neural pathways.

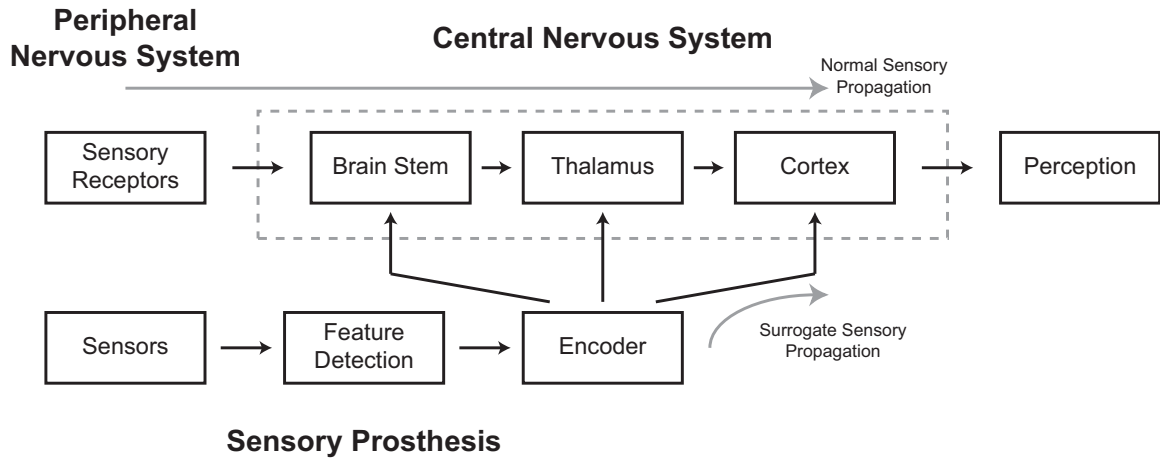


Figure 1.1: Block diagram of a central nervous system sensory prosthesis. Under normal conditions, receptors transduce sensory signals into neural activity that propagates along the sensory pathway. A sensory prosthesis supplies surrogate sensory information by encoding salient sensory features into patterned stimulation of the peripheral or central nervous system.

Towards this end, this thesis thoroughly characterizes the propagation of neural activity generated by artificial stimulation in the central nervous system. Subsequent stimulus design is then used to augment the propagation of neural activity. Finally, control of the propagation of neural activity is demonstrated through modulation of the underlying brain state.

1.3 Electrical and optogenetic stimuli as surrogate sensory stimuli

Artificially activating neural cells has a long history, pre-dating even the recording of electrical activity from neurons. As early as the late 1800s, electrical stimulation was used to activate neurons in the central nervous system [58, 167]. In fact, the methods used today are largely unchanged since stimulation through a microelectrode, termed microstimulation, became common in the 1960s. Early studies with microstimulation began to elucidate the minimal currents necessary to stimulate single axons or cell bodies [179, 74], quantified through chronaxie and rheobase measurements, and how

this relationship changed with distance from the electrode [162]. This research continues today, seeking to answer the questions of which neural elements are stimulated in the micro-environment around the electrode and how specific can this stimulation be [140, 141, 123, 83].

Concurrently, microstimulation has been used as an input to behavioral studies at the macroscopic scale, with the initial goal to investigate the arrangement of motor and sensory maps on the cortical surface [150], but more recently to augment or drive perception and motor output in behavioral tasks [166, 165, 155, 143]. Most recently, simple patterns of microstimulation were used to provide tactile feedback for a monkey performing an active exploration motor task through a brain machine interface device [142]. Impressive behavioral tasks such as this demonstrate the potential for electrical microstimulation to deliver surrogate sensory signals for sensory prosthesis applications [70, 64]. However, it remains to be seen how the stimulation paradigms using only a few simple patterns can generalize to produce natural [82], or at least functional sensations, especially considering that behavioral tasks using animals cannot readily provide the requisite level of detail in the perceptual report and studies within humans are relatively rare [71, 82]. For this reason, animal studies must focus on directly measuring the downstream neural activation from electrical microstimulation to quantify and enhance the perceptual degrees of freedom of patterned stimulation [130].

Exactly how electrical stimulation activates a population of neurons within a complex neural circuit, which ultimately gives rise to behavioral percepts, is far less well understood than for single cell or behavioral readouts. Recent work has pushed towards recording population responses downstream of the delivery of patterned microstimulation in vivo [33, 29, 38, 83, 112, 22, 200]. Through this approach, it will ultimately be possible to generate a model of the downstream neural response to electrical stimulation, allowing the design of stimuli to produce optimal neural activations

that mimic the naturalistic neural responses or maximize the degrees of freedom of a prosthesis [110, 44].

Much like electrical microstimulation, in the short existence of the optogenetics technique, the majority of the focus has been on tool development, using intracellular recordings as the readout [16, 213, 120], or impressive behavioral experiments designed to attack long standing hypotheses within the field of neuroscience [99, 212, 105]. Meanwhile, the network level response properties under optogenetic stimulation are not known. Taken together, electrical and optical stimulation comprise a class of artificial stimulation techniques that can activate the brain with much higher spatial and temporal resolution than pharmacological agents or other non-invasive techniques like transcranial magnetic stimulation, and ultimately be used to investigate information processing across populations of neurons within large-scale neural circuits and deliver patterns of activation for sensory prosthesis applications.

1.4 Natural and artificial stimuli activate neural circuits in distinct ways

Much of the research on electrical and optical stimuli has been based in the early sensory pathways. Across all major sensory systems, the anatomy of the early stages in the pathway is well known and adheres to a strict topological map [204, 51, 86], such that propagation can be studied across multiple stages [92, 196]. Further, the response properties of individual cells to various types of sensory stimuli have been rigorously mapped, and sensory driven behavioral experiments have become increasingly common. Altogether, sensory systems provide an ideal testbed for artificial stimuli by allowing the comparison of the ways in which sensory and artificial stimuli activate neural circuits, and how that activation propagates to downstream structures.

From the wealth of behavioral experiments utilizing interleaved sensory and artificial stimuli, and the success of sensory prostheses like the cochlear implant, sensory and artificial stimuli would seem to activate neural circuits in a similar manner. Indeed,

behavioral work in primates has shown that artificial stimuli can bias the perception of sensory stimuli [166], or, as Romo et al. showed, functionally replace sensory stimuli in certain tasks [165]. However, neither of these studies, or the myriad others, can comment on how the artificial stimuli truly feel or how they activate the neural circuit as compared to sensory stimuli.

Human studies with artificial stimuli, however, do allow for the patient to report the nature and quality of the percept. With the exception of the cochlear implant, the majority of human studies have indicated a significant difference between the percept generated by artificial stimulation and a normal sensory stimulus. The term paresthesia is often used to describe the sensation resulting from an artificial stimulus in the somatosensory pathway. Paresthesias have been reported as being discordant (i.e. not topologically matched to the receptive field of the stimulated neurons)[71, 46] or unnatural altogether [145, 148, 82, 81]. This suggests that artificial stimuli either: 1) activate distinct neural elements from sensory stimuli, or 2) propagate to downstream structures differently than sensory stimuli. Electrical stimuli are known to activate axons preferentially over cell bodies [83], which may contribute to the discordant and unnatural sensations, but the differences, if any, in the propagation of neural activity generated by sensory and artificial stimuli remain largely unknown.

Two recent studies highlight the fact that neural signals generated by sensory and artificial stimuli propagate differently through neural circuits. In an anesthetized primate preparation, Logothetis et al [112] used single unit recordings in visual cortex, and fMRI throughout the visual pathway, to measure the propagation of neural activity generated by visual stimuli delivered to the eye and electrical stimuli delivered to the lateral geniculate nucleus of the thalamus. Using fMRI, visual stimuli were found to propagate through all stages in the visual pathway, from the thalamus to higher association areas in cortex, whereas electrical stimuli ($200\mu\text{A}$, 200Hz) induced positive levels of activation in primary visual cortex, but not in areas further downstream.

While both visual and electrical stimuli activated the thalamus, the propagation of activity downstream was profoundly different. One explanation provided for these results was the differential recruitment of inhibitory circuitry within the network, which may have prevented propagation for the electrical stimuli.

Another study, by Masse and Cook [119], tested the behavioral consequences of neural circuit dynamics elicited by subthreshold visual and electrical stimuli (delivered to the cortex in this case). While each was tuned to a subthreshold amplitude, the relative timing of the stimuli to a suprathreshold test stimulus augmented the probability that the animal would detect the stimulus, with shorter timescales producing a greater decrease in detectability. The performance of the animal recovered over a shorter time period for the subthreshold visual stimuli than the subthreshold electrical stimuli. In this case the behavioral performance of the animal was the readout, suggesting that there is a functional/behavioral significance to the distinct properties of neural propagation for sensory and artificial stimuli.

1.5 Direct neural readout of nonlinear propagation in neural circuits

To better understand how artificial stimulation activates the brain, we must directly record from multiple nodes within neural circuits simultaneously. Intra-operative stimulation of humans provides a qualitative report of the nature of the artificially generated sensation [82, 46, 145, 71]; however, the opportunity for these experiments is scarce. Animal experiments, while affording higher throughput, do not provide a description of the nature of the sensations [130]. Thus, the direct recording of the downstream neural response is needed to inform the development of sensory prostheses. The goal of this thesis was to use a direct readout of neural activity to assess the functional differences in signal propagation for sensory and artificial stimuli, and use this information to develop methods to control spatiotemporal patterns of neural activity in the brain.

A number of recent studies have focused on directly measuring the downstream neural response to electrical stimulation [33, 112, 93, 29, 83, 38, 22, 200]. The measurement location has principally been the primary sensory cortex, regardless of whether the stimulation was in the periphery [115], spinal cord [200], thalamus [33, 93], or cortex [29, 27, 83, 22]. The decision to record in cortex is two-fold in that 1) it's location on the surface of the brain allows for easier neural interfacing, and 2) the cortical response is believed to be the first stage of emergent sensory perception [133]. Typically the experiments involved simple patterns of stimuli, such as in Butovas et al., where a multi-electrode array was used to record the cortical spiking response to pairs of intracortical microstimulation [29]. But, the direct recording of spiking activity in cortex can be plagued by the artifact generated by electrical stimulation, and the spatial resolution of multi-electrode arrays is typically sparse.

The incorporation of various imaging techniques has expanded the spatial resolution of neural circuit recordings, allowing many nodes to be captured at once, while also reducing the affect of the stimulation artifact. For instance, optical imaging methods, such as voltage sensitive dye imaging [38], calcium dye imaging [11], and two photon techniques [83], are invariant to the electrical artifacts generated by microstimulation and provide high spatial resolution over large regions of the cortical surface. The gains in spatial extent of the neural recording are at the expense of the temporal resolution of the signals and the indirect relationship between the recorded signals and spiking activity [72, 104]. Functional magnetic resonance imaging (fMRI) extends this by sacrificing temporal resolution to gain recordings across the entire brain, allowing analysis across multiple stages of a neural pathway [112]. However, the fMRI BOLD signal is even more indirectly related to neural activity, by virtue of the metabolic demands of spiking [113].

These tools are critical to the understanding of surrogate sensory inputs. By directly observing the neural response to patterns of sensory and artificial stimuli, we

can begin to understand and model the mapping from artificial stimulation to neural activation to inform the development of sensory prostheses.

1.6 Experimental framework for the design of functionally relevant surrogate sensory inputs

The mapping from artificial stimulation to neural activation will likely depend on the exact neural circuit under investigation. However, many neural circuits within the brain share a common canonical structure, such that there may be emergent properties that can be generalized across brain structures. The thalamocortical circuit is one such canonical circuit [171]. The architecture of the thalamocortical circuit is largely conserved across sensory modalities and mammalian species, and contains both recurrent excitation and feedforward inhibition, as depicted in figure 1.2(b). Recurrent inhibition occurs at the level of the thalamus, with the neighboring reticular thalamus providing direct inhibitory feedback. The interaction between the thalamus and reticular thalamus is strongly modulated by state, as the two structures dynamically shift between tonic and burst firing modes [170]. The pyramidal neurons of the thalamus then project to layer 4 of cortex, synapsing on both excitatory and inhibitory neurons within the population. The feedforward drive on the inhibitory population quickly shunts excitability in cortex, leaving only a short window of opportunity for thalamic drive to activate the excitatory population [59]. The resulting dynamics of the thalamocortical circuit are well known in response to sensory stimuli, both under anesthesia [173, 174, 185, 23, 176, 107, 31, 106, 57] and in the awake animal [161, 55, 122].

In this thesis, we delivered artificial stimuli to the thalamus, allowing a direct neural readout of the downstream cortical response. And, while there is some evidence that the thalamus may be a legitimate target for sensory prosthesis implementation [155, 154, 82, 81], the only assertion here is that it is an ideal canonical circuit through which we can study the relative behaviors of natural and artificial stimulation.

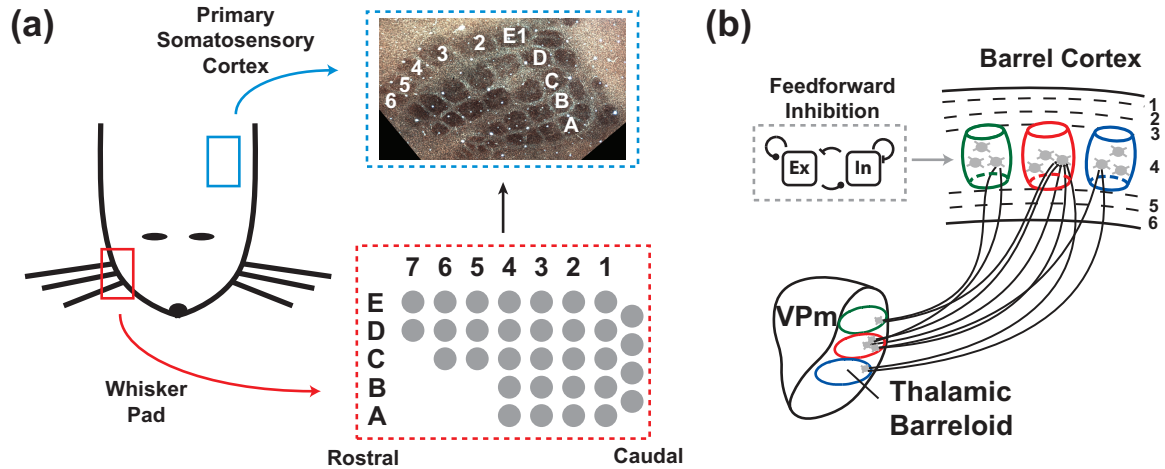


Figure 1.2: Architecture of the rodent vibrissa somatosensory system. (a) There is a one-to-one topographic mapping from the whiskers on the face and the cortical columns in cortex, as shown in the tangential section stained with cytochrome oxidase. (b) The one-to-one mapping occurs in the thalamus as well and maintains the topography in the projection to layer 4 of primary somatosensory cortex. Each cortical column contains a mixture of excitatory and inhibitory neurons, producing canonical circuit phenomena like recurrent excitation and feedforward inhibition.

Specifically, the lemniscal pathway within the rodent vibrissa somatosensory system was utilized for all experiments in this body of work [51]. The detailed anatomical literature on the rodent vibrissa system makes it an ideal system for studying the propagation of neural activity. The rodent vibrissa system is characterized by its discrete nature, such that the mechanical deformation of a given whisker on the face activates, with high specificity, a somatotopically defined collection of cells [51, 206, 188]. The discretization is illustrated by the cartoon in figure 1.2(a), where the topographic mapping of the whiskers on the face exactly matches the arrangement of the cortical columns in cortex, as revealed by cytochrome oxidase staining.

In this way, the location of the artificial stimulation can be determined and matched to the deflection of the corresponding whisker on the face, as in figure 1.3. The electrode location was determined electrophysiologically by measuring the spiking activity of a single unit (example in figure 1.3(a)) in response to the deflection of neighboring whiskers on the face. An example in figure 1.3(b) presents the spiking

response to the deflection of seven neighboring whiskers. The cell responded with the greatest number of spikes (summarized in figure 1.3(c)), and shortest latency, to the deflection of whisker D2. For this example and all others throughout the thesis, the location of the electrode was determined in this manner. In addition to determining the primary whisker, the electrophysiological data was used to verify the location of the electrode in the ventro-posterior medial (VPM) nucleus of the thalamus, comprising the lemniscal portion of the whisker somatosensory thalamus, as opposed to the posterior medial (POM) nucleus that forms the paralemniscal portion of the pathway [51]. VPM neurons have been shown previously to respond with much shorter latency to whisker deflection (3-5ms) than POM neurons, enabling electrophysiological classification [50, 4].

While recent advances in electrode technology, based on silicon probes or tetrode arrays, have enabled high density electrophysiological recordings of neural activation [32], these techniques are not ideally suited for studying the response to electrical microstimulation due to the large stimulus artifact that occurs upon stimulation and sparse spatial sampling. Methods exist for the minimization of information loss during the stimulus artifact [79, 192, 137], but we instead used voltage sensitive dye imaging (VSDI) to record the cortical response to patterns of natural and artificial stimulation of the thalamus, which, as an optical signal, does not suffer from electrical microstimulation artifacts. VSDI fluorescence provides a spatiotemporal signal that is linearly proportional to the membrane potential of neurons [72], and has been used extensively to measure the propagation of neural activity in functional circuits [152, 55]. Due to the light scattering of the tissue, voltage sensitive dye imaging principally measures the activity in the superficial layers of cortex [73]. Further, the dynamics and imaging of the signal are too slow to resolve individual action potentials, restricting the interpretation of the signal to sub-threshold activation [72]. However, the onset of sub-threshold activation in layer 2/3 is likely highly correlated with the

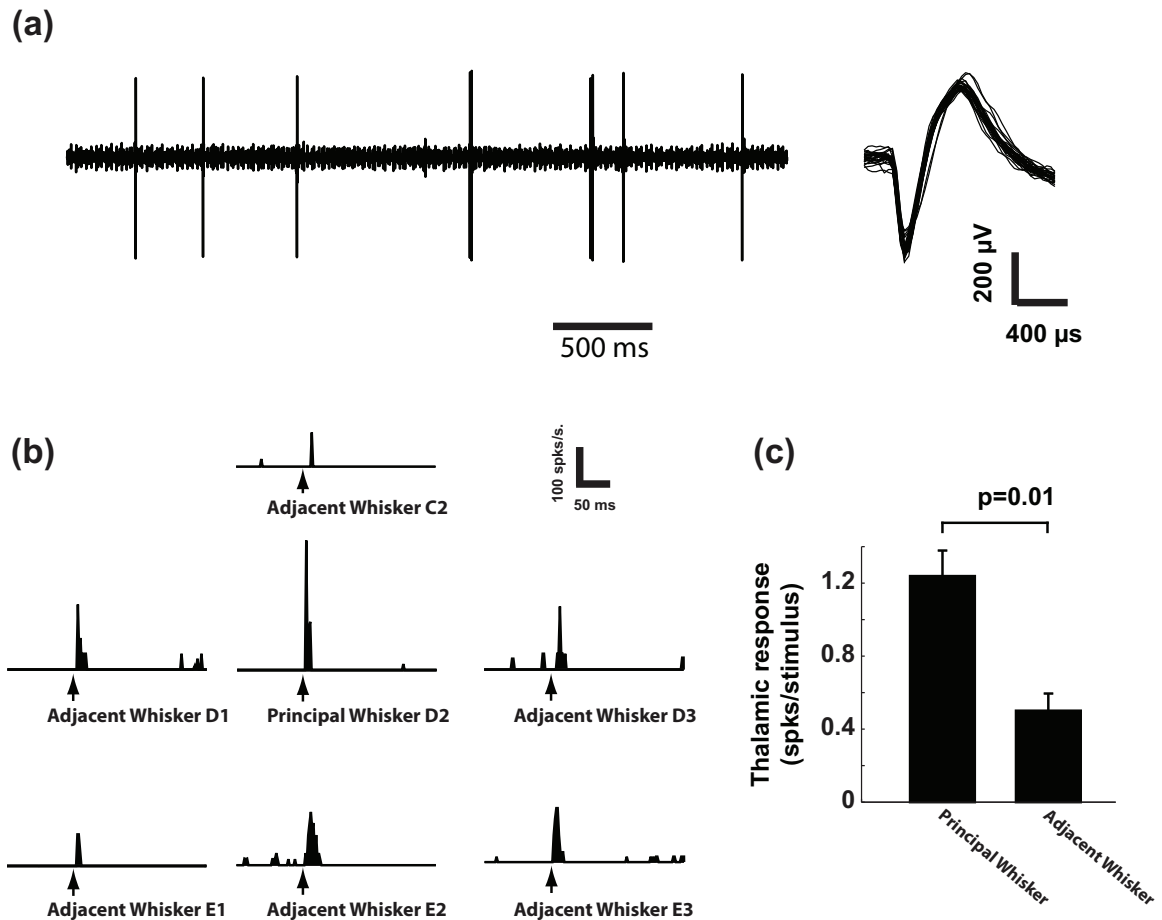


Figure 1.3: Functional validation of thalamic electrode placement. A) A typical example of single-unit activity recorded in VPM (left) and spike waveforms (right). B) The peri-stimulus time histograms (PSTHs) of a VPM cell in response to punctate deflection of whisker C2, D1, D2, D3, E1, E2, and E3 (binsize = 2ms). C) Selection of principal whisker (PW) was confirmed by the fact that PW evoked stronger spiking activity than adjacent whiskers (AW) ($p=0.01$, Mann-Whitney U-test).

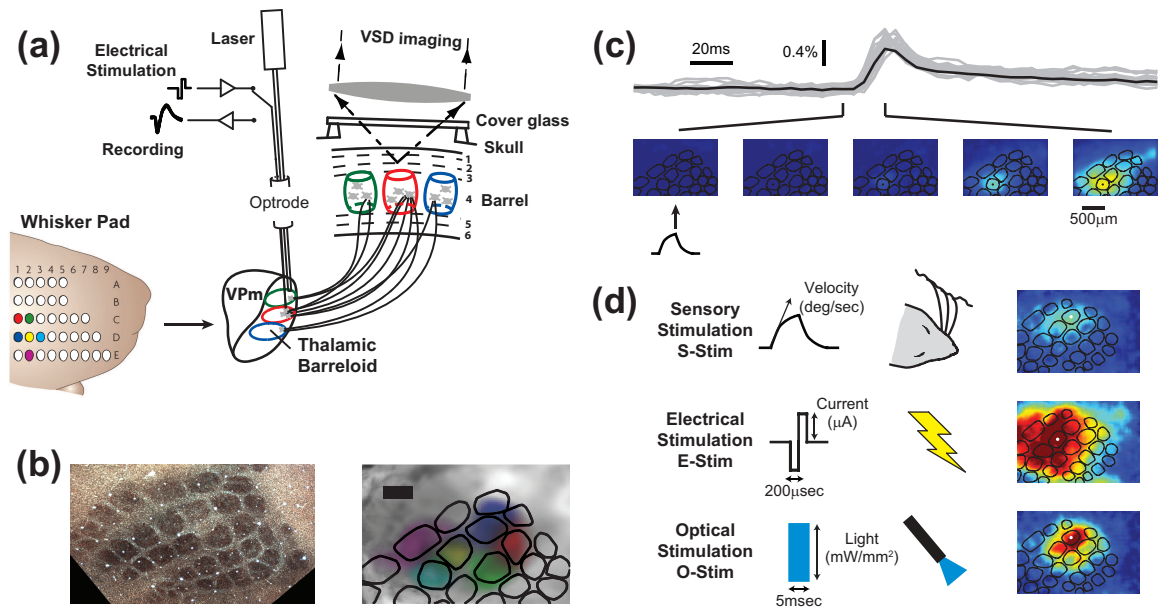


Figure 1.4: Voltage sensitive dye imaging of cortical network activity. (a) Diagram of the experimental setup. The electrode and/or optical fiber was positioned in a barreloid in the thalamus, a collection of cells that respond most vigorously to a common whisker. Imaging in cortex captured the response of each cortical column. (b) Histological analysis provides an anatomical map of the cortical column structure (left), which was then fit to the functional imaging data (right). (c) The cortical response to whisker stimulation begins 10-15 milliseconds after stimulation, in a focal location before growing in amplitude and spreading spatially. (d) Examples of the cortical response to whisker, electrical, and optical stimuli delivered within a single experiment. Scale bars in (b) and (c) are 500 micrometers.

supra-threshold activity in layer 4 of cortex [152]. Given these advantages and limitations, we have developed VSDI techniques in our laboratory to measure circuit dynamics across functional cortical columns with fast temporal resolution [195].

An example of the VSDI setup is shown in figure 1.4(a). The VSDI data is presented and analyzed in two different ways throughout the project, both illustrated in figure 1.4(c). Individual frames provide information about the spatial spread of the cortical response with sub-columnar resolution (bottom of figure 1.4(c)). Alternately, to analyze the temporal evolution of the cortical response on single trials (top of figure 1.4(c)), we average spatially within the topographically matched cortical column to

the position of the electrode and/or optical fiber in the thalamus or the whisker on the face. Histological analysis was performed through cytochrome oxidase staining to identify the cortical columnar map, which was then registered functionally to the VSDI frames, as in figure 1.4(b). Examples of the cortical response to whisker, electrical, and optogenetic inputs from a single experiment are shown in figure 1.4(d) for comparison. Taken together with the goals of this project, VSDI provides a test-bed that allows us to identify and model the nonlinear dynamics of the spatiotemporal response of a neural circuit in vivo and informs the design of patterned stimuli for the control of neural circuit activation.

1.7 Organization of the thesis

In the following chapters, this experimental framework is used to characterize the cortical response to sensory, electrical, and optogenetic inputs to the thalamocortical circuit. Chapter 2 explores the nonlinear response properties for single stimuli with varying intensity and concludes with computational model to support the development of a hypothesis about how sensory and artificial stimuli differentially engage neural circuits. Chapter 3 extends the analysis to patterns of stimulation in time, and in Chapter 4 a model is developed to describe the nonlinear dynamics of the cortical response to patterns of electrical stimulation, which further supports the hypothesis outlined in Chapter 2. In Chapter 5, stimulus design is explored to overcome the different neural responses to natural and artificial stimuli detailed in Chapter 2. Specifically, the electrical stimulation waveform is varied in an effort to improve spatial specificity in the cortical response, and optimal signal set design is utilized to maximize discriminability amongst cortical responses. Chapter 6 generalizes the previous results by quantifying the effect of the state of the thalamus on the nonlinear properties of the cortical response to thalamic microstimulation by using optogenetics to control the depolarization of the thalamic population. Finally, Chapter 7 details

the combined impact of these experimental and computational results on the use of artificial stimuli as surrogate sensory signals, and more generally on the field of neural engineering as a whole, before concluding with a generalized framework for the control of neural activity.

CHAPTER II

CHARACTERIZATION OF THE NONLINEAR SENSITIVITY OF THE CORTICAL RESPONSE TO SURROGATE THALAMIC DRIVE

Portions of this work were presented in poster form at the following conference:

Millard, DC, Gollnick, CA, Hendry, WJ, Rozell, CJ, & Stanley, GB. The role of magnitude and synchrony of population activity in nonlinear circuit processing in the thalamocortical circuit of the rodent vibrissa system. Society for Neuroscience Annual Meeting, San Diego, CA, November 2013.

2.1 Introduction

Artificial stimulation techniques afford the ability to replace lost function or interrogate neural circuit operation by acting as surrogate inputs to the brain. Electrical stimulation remains the only clinically viable means of artificial stimulation on fast timescales, while optogenetic stimulation has revolutionized the ability to dissect neural circuits experimentally through genetic cell type specificity. However, behavioral [119] and electrophysiological [112] evidence suggests that neural activity generated by natural and artificial stimuli differentially propagate to downstream structures. Yet, the extent to which natural and artificial stimuli differentially activate neural circuits, the mechanism, and how this affects the functional relevance of the downstream signals remains unknown.

The thalamocortical circuit has proven an ideal model system for studying the propagation of neural activity in sensory pathways experimentally. The anatomy of the circuit is well-known across all sensory modalities [171] and is characterized by

a set of highly convergent/divergent connections from the thalamus to cortex. The thalamocortical synapses are known to be weak [26] and noisy, yet robustly transmit information to the downstream cortical neurons, indicating a role for synchrony in activity propagation [2, 24, 177]. Indeed, across sensory modalities, the synchrony of early stage sensory neurons is modulated by the properties of the stimulus [156, 185, 196]. In contrast, artificial stimulation of neuronal tissue through electrical or optogenetic means induces an extreme degree of timing precision and synchronization across neuronal populations [191, 16]. Importantly, synchrony has been shown in computational models to play a key role in the propagation and transformation of activity across brain structures [100]. Thus, the degree of synchronization, and the properties of the ensuing downstream neural response, for sensory and artificial stimuli will have profound coding consequences on the design of surrogate signals.

In this study, we directly measured the propagation of neural activity generated by sensory and artificial stimuli, built a computational network model to investigate the mechanistic differences in circuit engagement, and established a theoretical framework for the optimal design of artificial stimuli. Specifically, we quantified the amplitude, variability, and spatial spread of the cortical response to whisker, electrical, and optogenetic stimuli using voltage sensitive dye imaging in the vibrissa region of the rodent primary somatosensory cortex. Then, using a model of the thalamocortical circuit, we explained the distinct cortical response properties for whisker, electrical, and optogenetic inputs through systematic variations of the magnitude and synchrony of the input population activity. The model attributed the highly nonlinear properties of electrical and optogenetic stimulation to the high degree of synchronization, while the fixed, focal spatial activation of whisker and optical stimuli was due to the activation of cell bodies and not axons. Finally, we explored the functional ramifications of the distinct neural response statistics associated with natural and artificial stimulation in the context of generating maximally discriminable neural responses in downstream

structures.

2.2 Methods

2.2.1 Experimental Preparation

All procedures were approved by the Georgia Institute of Technology Institutional Animal Care and Use Committee and followed guidelines established by the National Institutes of Health. The same experimental preparation described in Chapter 1 was used here, with the detailed methods described in Appendix A. Briefly, female sprague-dawley rats (250-300g) were initially anesthetized with 4% isoflurane before intraperitoneal injection of Nembutal (50mg/kg weight) for long term anesthesia. Subsequent doses of Nembutal were used to maintain a surgical level of anesthesia.

Animals were mounted in a stereotactic device and a craniotomy was performed over the left parietal cortex (coordinates: 1-4mm posterior to bregma, 4-7mm lateral to midline) to expose the barrel representation of the primary somatosensory cortex [149]. Another craniotomy was performed to allow access to the ventral postero-medial (VPM) region of the thalamus (coordinates: 2-4mm posterior to bregma, 1.5-2.5mm lateral to midline, 4.5-5.5mm depth at a 12 degree angle to brain surface).

A subset of the animals underwent an initial survival surgery, during which the viral vector (AAV2-CaMKIIa-hChR2(H134R)-mCherry, UNC Viral Vector Core, Chapel Hill, NC) was delivered to the thalamus using stereotactic coordinates. The injection was delivered at 0.2 μ L/min for 5min for a total of 1 μ L. The animals were allowed to recover for 3-4 weeks, providing time for the ChR2 expression to reach functional levels.

2.2.2 Voltage Sensitive Dye Imaging

Voltage sensitive dye imaging (VSDI) was used to monitor cortical activation in response to thalamic microstimulation. The VSDI data were acquired at five millisecond interframe intervals beginning 200 milliseconds preceding stimulus presentation.

Multiple trials of VSDI data were collected for each stimulus. For each trial, the 40 frames (200ms) collected before the presentation of the stimulus were averaged to calculate the background fluorescence, against which the activation was measured. For each frame, the background fluorescence was subtracted to produce a differential signal ΔF . Additionally, each frame was divided by the background image to normalize for uneven illumination and staining to produce the signal $\Delta F/F_0$. For presentation purposes only, the individual trials were averaged together and then filtered with a 9x9 pixel spatial averaging filter.

The anatomical mapping, acquired through cytochrome oxidase histology, was registered with the functional cortical column mapping from VSDI by solving a linear inverse problem, the details of which have been described previously [195]. Following the functional image registration, the cortical response was discretized, where each signal corresponds to a single functional cortical column. In so doing, the VSDI signal was averaged spatially within the contour of the cortical column.

2.2.3 Whisker Stimulation

Whiskers were trimmed at approximately 12mm from the face, and were inserted into a glass pipette fixed to the end of a calibrated multi-layered piezoelectric bimorph bending actuator (range of motion, 1 mm; bandwidth, 200 Hz; Physik Instrumente (PI), Auburn, MA) positioned 10 mm from the vibrissa pad. Vibrissae were always deflected in the rostral-caudal plane. Punctate deflections consisted of exponential rising and falling phases (99% rise time, 5 ms; 99% fall time, 5 ms), with varying angular deflection velocity (75-1200 deg/s).

2.2.4 Electrical Stimulation

A glass coated tungsten microelectrode (impedance = 1-2 M Ω at 1kHz) was advanced to the ventral postero-medial (VPM) region of the thalamus using a precision microdrive (Knopf Instruments, Tujunga, CA). The principal vibrissa was determined

by manually deflecting individual whiskers and confirmed using the latency and spike count of single unit recordings in response to controlled whisker deflection using a piezo-electric actuator. In the event that single unit recordings could not be achieved, multi-unit activity was used.

Following electrophysiological determination of the electrode position and its associated principal vibrissa, the electrode was used to deliver microstimulation to the surrounding tissue. The stimulus waveforms were designed using a digital stimulus generator (WPI Inc, Sarasota, Florida) and delivered using a current controlled, optically isolated stimulator (WPI Inc, Sarasota, Florida). Individual electrical stimuli were charge-balanced, cathodal-first, biphasic waveforms of 200 microsecond duration per phase. A series of single electrical stimulation pulses with varying amplitude between 10 and 100 microamperes was used to test the static nonlinearity of the neural circuit. The current range was chosen to elicit the full range of sub-threshold to maximal cortical responses.

2.2.5 Optical Stimulation

For optical stimulation, an optrode was used. The optrode consisted of a multimode optical fiber (105 μ m core diameter, 125 μ m coating diameter, 0.22NA, Thorlabs, Newton, NJ) and one or two quartz coated platinum-tungsten microelectrodes (80 μ m diameter, Thomas Recording, Giessen, Germany). The microelectrodes were pulled and ground to an impedance of 1-2 M Ω at 1kHz. The optical fiber was also ground to a fine point, producing a spherical, rather than conical, pattern of light delivery.

A DPSS laser (Laserglow Technologies, Toronto, Canada) was used to deliver blue (473nm) light to the VPM thalamus and stimulate the ChR2 expressing cells. The stimuli were square pulses of 5ms duration and varying light intensity. The maximum light delivered during the experiments was approximately 150 mW/mm², but typically only 50 mW/mm² was needed to drive neural activity in ChR2 expressing neurons.

According to Aravanis et al, 2007, the light intensity at 200 micrometers from the optical fiber tip was $1/10^{th}$ of the calibrated light at the tip, such that the majority of neurons were activated at $5\text{mW}/\text{mm}^2$ in this case [3].

2.2.6 Computational Network Model

A network model was used to investigate the hypothesis that sensory and artificial stimuli differentially activate thalamic neurons, leading to distinct response properties in the downstream cortex. Specifically, the development of a model allowed direct, and independent, manipulation of the magnitude and synchrony of input population activity, whereas we do not have control of these variables experimentally. Towards this goal, a simple cortical network model was developed and the thalamic input activity was systematically controlled.

The network was modeled after the thalamocortical circuit of the rodent vibrissa pathway and the extensive anatomical and computational literature for this model system [101, 51]. The neural circuit consisted of 400 thalamic neurons that project to a downstream cortical population of 800 excitatory neurons and 200 inhibitory neurons. The relative size of the thalamic and cortical populations and the balance of excitation and inhibition in cortex were based off of previous computational studies [101], which were scaled down from the true anatomy. The thalamic neurons made direct synapses on both the excitatory and inhibitory neurons in cortex, allowing for feedforward inhibition. Additionally, the connectivity to the inhibitory neurons from the thalamus, and from the inhibitory neurons to the rest of cortex, were more spatially diffuse to be consistent with anatomical studies [95].

The cortical population was modeled as a collection of quadratic integrate and fire neurons [89] according to the following set of equations:

$$v = 0.04v^2 + 5v + 140 - u + I \quad (1)$$

Table 2.1: Parameters for Izhikevich cortical neurons. The excitatory and inhibitory cortical neurons in the model have different parameters. Parameters are randomized across neurons through r_i , which is a random variable uniformly distributed on the interval $[0, 1]$ with the subscript i indicating the cell index.

Parameter	Excitatory	Inhibitory
a	0.02	$0.02+0.08r_i$
b	0.2	$0.25-0.05r_i$
c	$-65+15r_i^2$	-65
d	$8-6r_i^2$	2

$$u = a(bv - u) \quad (2)$$

where v is a 1000×1 vector of the membrane potential of the cortical neurons and u is a 1000×1 vector describing the recovery variable instituting the nonlinear dynamics of the cortical neurons. Spiking was determined by the membrane potential passing a fixed threshold of 30mV, at which point $v \rightarrow c$ and $u \rightarrow u + d$. The parameters were set according to previous work [89] and are summarized in table 2.1. While u and v describe the membrane dynamics of the individual neurons, the input, I , delivered to the cells is responsible for generating the network dynamics. The input delivered to each cell was composed of a membrane noise term, I_m , and summed post-synaptic potentials from the thalamocortical and intracortical activity, I_S , according to the following set of equations:

$$I = I_m + I_S \quad (3)$$

$$I_S = W^{TC} \cdot S^{TC} + W^{CC} \cdot S^{CC} \quad (4)$$

where W^{TC} is a 1000×400 matrix describing the thalamic synaptic weights, such that W_{ij}^{TC} is the strength of the synapse from the j^{th} thalamic neuron on the i^{th} cortical neuron, and S^{TC} (a 400×1 vector) is the thalamic activity from the preceding time step, where S_j^{TC} is one if the j^{th} thalamic cell spiked in the preceding time sample and

zero otherwise. W^{CC} and S^{CC} describe the intracortical synaptic activity in a similar manner. I_m is drawn from a Normal distribution with mean of zero and covariance K , producing spatially correlated membrane potential fluctuations in the cortical population [102]. The membrane potential fluctuations were also filtered temporally over a 10ms window. The connectivity and synaptic weights were extrapolated from Kyriazi, et al, 1993 and updated according to more recent in-vivo work [101, 26] and are also summarized in table one. This type of network has previously been shown to exhibit many of the common response features and rhythms of in-vivo cortical circuits [89].

The input spikes were drawn from an inhomogeneous Poisson process with rate:

$$\lambda(t) = \lambda_{stim}(t) + \lambda_{spont} \quad (5)$$

where

$$\lambda_{stim}(t) = t \cdot \exp((-t)/\tau_{sync}) \quad \text{such that} \quad \int_T \lambda_{stim}(t)dt = M \quad (6)$$

such that M gave the average number of spikes per neuron per trial in the stimulus window, T , and τ_{sync} determined the synchrony of those spikes. Both τ_{sync} and M were functions of neuron location, x , such that neurons near the center of the stimulus responded with the greatest synchrony and magnitude, and neurons far from the center of the stimulus responded with low synchrony and magnitude. The fall off of these variables with space was dictated by the following Gaussian shaped functions:

$$M(x) = M_{peak} \cdot \exp(-(x - x_0)^2/\sigma^2) \quad (7)$$

$$\tau_{sync}(x) = S_{peak} + 10 * [1 - \exp(-(x - x_0)^2/\sigma^2)] \quad (8)$$

where x gives the location along the linear array of thalamic cells, with x_0 at the

center of the stimulus. M_{peak} and S_{peak} determine the magnitude and synchrony, respectively, at the center of the stimulus where it is maximal, and σ determines how quickly the magnitude and synchrony fall off with distance from the center of the stimulus. S_{peak} was systematically varied, while M_{peak} was held constant, in the top portion of figure 2.4, and vice versa for the bottom portion of figure 2.4. Meanwhile, the σ parameter is controlled in figure 2.5.

While τ_{sync} modulated the synchrony of the input spiking activity, this variable was not explicitly related to commonly used measures of synchrony. Therefore, in all analyses, the synchrony was calculated according to methods used previously in the literature [185, 196]. Briefly, the cross-correlogram was computed for pairs of neurons using the input spikes contributing to the stimulus. The area under the cross-correlogram within a synchrony window (± 5 milliseconds), normalized for the number of spikes used to create the cross-correlogram, was calculated for each pairwise combination of the central twenty thalamic neurons and averaged across the permutations to give the synchrony measure. Under this calculation, a synchrony value of one means that all stimulus-related spikes occur within the synchrony window.

The amplitude, variability, and spatial spread metrics were measured from the population of excitatory cortical neurons within the simulation. The amplitude and spatial spread were calculated by fitting a Gaussian shape function to the spatial distribution of spiking activity in the twenty milliseconds following the stimulus, using the following equation:

$$M(x) = A \cdot \exp(-(x - x_0)^2 / \sigma_{spread}^2) + B \quad (9)$$

where again x gives the location along the linear array of cortical excitatory cells, with x_0 at the center of the response. A and σ_{spread} give the magnitude (spikes/stimulus) and spread (number of cells) of the simulated cortical response, where B accounts for the baseline firing rate of the neurons.

2.3 Results

Natural sensory stimuli and artificial stimuli both drive neural activation, allowing the dissection of neural circuits and complex psychophysical tasks aimed at understanding how the brain works at a systems level. However, previous literature suggests that sensory and artificial stimuli activate neural circuits in distinctly different ways [112, 119]. Here, we specifically compared the propagation of neural activity generated by whisker deflection sensory inputs and electrical and optogenetic stimuli delivered to the thalamus. Optogenetic stimuli were found to exhibit a mixture of the properties associated with whisker and electrical inputs, and ultimately the experimental results were well-predicted by systematic variations of the magnitude and synchrony of population neural activity in a computational model of the thalamocortical circuit. Finally, we explored the functional consequences of the distinct activation patterns between sensory and artificial stimuli using a theoretical model to estimate the number of discriminable cortical responses that can be generated by each stimulus type. Taken together, these results provide insight into the neural activation patterns and information content that can be generated by natural and artificial stimuli.

2.3.1 The cortical response is highly nonlinear for whisker, electrical, and optical stimuli

All experiments utilized in-vivo voltage sensitive dye imaging (VSDI) of layer 2/3 in the whisker representation of the primary somatosensory cortex with whisker stimuli delivered to individual whiskers on the face, and electrical and optical stimuli delivered to the topologically matched ventral postero-medial (VPM) portion of the thalamus in the anesthetized rodent, according to the methods mentioned in Chapter 1 and described in detail in Appendix A.

First, the amplitude of the cortical response was quantified for whisker, electrical, and optical stimuli of increasing intensity. To do so, the cortical response was averaged across trials and then averaged spatially within the principal cortical column. The

mean and standard deviation of the maximum response amplitude was then calculated and plotted as a function of whisker deflection velocity, electrical stimulation current, and optical stimulation light intensity.

A representative example of the cortical response to each stimulus type is shown in figure 2.1(a). For all stimulus types, the response amplitude increased monotonically as a function of stimulus intensity, and was well approximated by a sigmoidal function. From the sigmoidal fit, the average threshold whisker deflection velocity was determined to be $229 \pm 139 \text{ deg sec}^{-1}$ (N=12), consistent with previous literature [180, 146]. Similarly, the average threshold current was $54 \pm 20 \mu\text{A}$ (N=18) and the threshold light intensity was $87 \pm 31 \text{ mW mm}^{-2}$ (N=6). Figure 2.1(b) presents a compilation of the mean response amplitude across all experiments and all stimulus types. For each experiment, the stimulus intensities are plotted relative to the threshold stimulus and the response amplitudes are normalized with respect to the amplitude of the sigmoidal fit. The compilation of all the data illustrates the consistency and similarity of the nonlinear relationship between cortical response amplitude and the strength of sensory and artificial stimuli.

2.3.2 Whisker stimuli are more sensitive to trial to trial variability

While the mean cortical response showed similar trends for sensory and artificial stimuli, the trial to trial variability did not. The trial-by-trial variability was quantified as the standard deviation of the single trial maximum response amplitudes for each stimulus type. Figure 2.1(c) presents the summary data for all experiments and all stimulus types. Because there is no clear relationship between deflection velocity, microstimulation current, and light intensity, the variability was analyzed as a function of the normalized mean response amplitude. For whisker stimuli, the trial to trial variability increased linearly as a function of the response amplitude. At low response amplitudes, the trial by trial variability matched the variability in noise

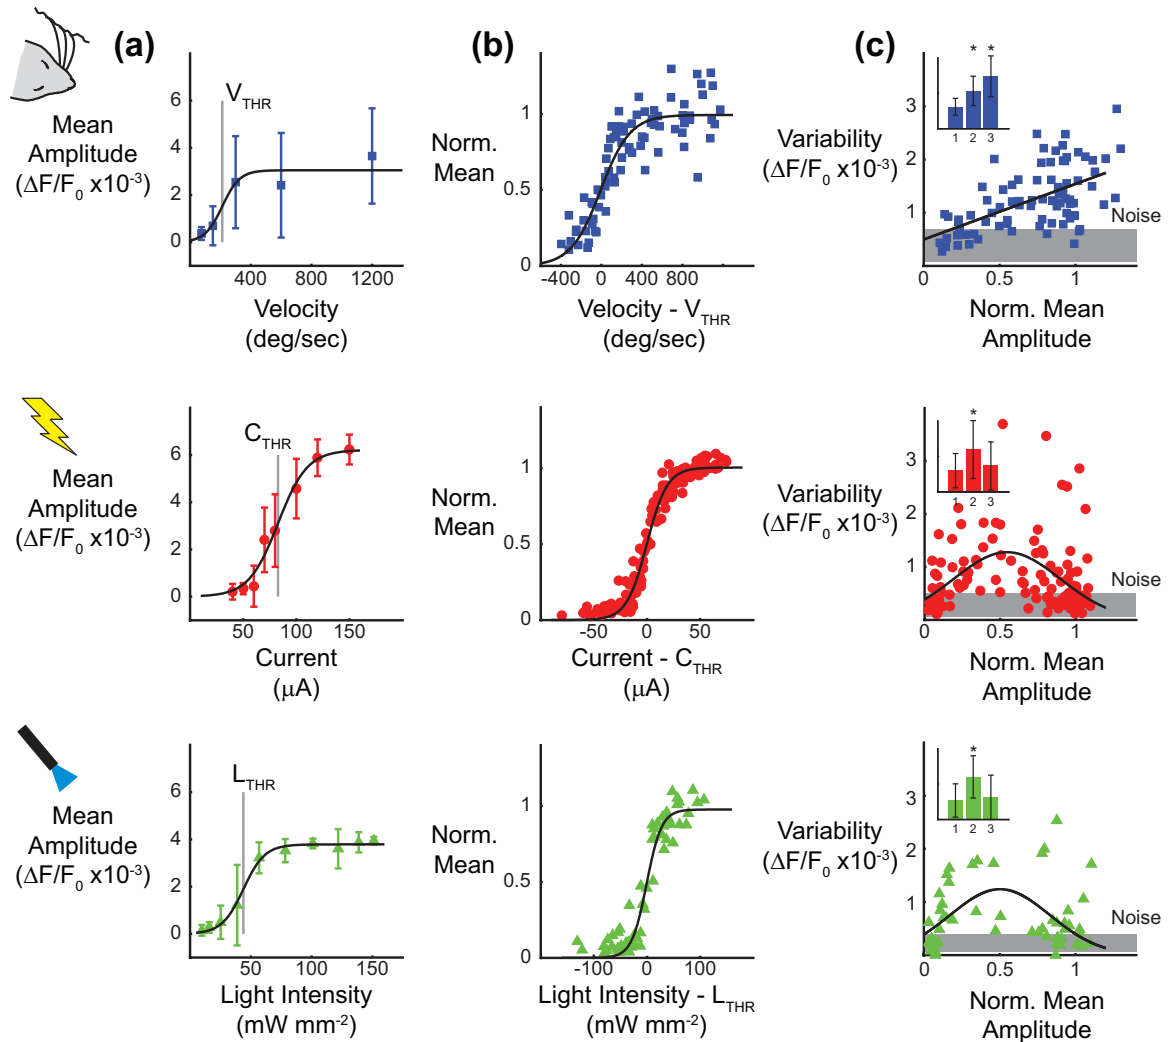


Figure 2.1: Whisker stimuli are more sensitive to trial-to-trial variability. (a) The response amplitude was quantified as the maximum VSDI signal, after averaging spatially within a single cortical column and across trials, for whisker, electrical, and optical stimuli. (b) Tuning curves for velocity (top, $N=12$), current (middle, $N=18$), and light intensity (bottom, $N=6$) across all experiments. Each experiment was normalized relative to the threshold velocity. (c) The variability in the response amplitude was quantified as the standard deviation of the maximum amplitude across trials, averaged spatially within a single cortical column. The response amplitude variability for whisker stimuli (top) increases as a function of mean response amplitude ($N=12$). The response amplitude variability for electrical (middle) and optical (bottom) stimuli peak at the threshold response amplitude ($N=18$ for electrical, $N=6$ for optical).

trials (gray band, mean of variability in the noise \pm one standard deviation across N=12 experiments), whereas high stimulus intensities exhibited a marked increase in variability. The variability was binned into the lower, middle, and upper response amplitudes, corresponding to sub-threshold, peri-threshold, and supra-threshold, respectively. The variability in the peri- and supra-threshold bins were significantly different from the sub-threshold response amplitudes.

The variability in the cortical response to thalamic microstimulation and optogenetic stimulation followed a different trend. At sub-threshold and supra-threshold response amplitudes, the trial by trial variability was comparable to that of noise trials (gray band, mean of variability in the noise \pm one standard deviation across N=18 experiments). However, near the threshold response amplitude, the trial by trial variability was significantly higher. When binned in a similar manner to the whisker stimuli, only the peri-threshold bin was significantly different from the sub-threshold variability, indicating that the variability returned to baseline at strong response amplitudes for electrical and optical stimuli. Importantly, the distinct trends of trial to trial variability for sensory and artificial stimuli may have implications for the capacity to generate discriminable cortical responses (see Chapter 5) and influence behavior (see discussion).

2.3.3 The cortical response to microstimulation exhibits the greatest spatial spread

In addition, the voltage sensitive dye imaging modality enabled the quantification of the spatial spread of the cortical response to whisker, electrical, and optical stimuli. Examples of the spatial cortical response to each stimulus type are presented in figure 2.2(a) for varying stimulus intensities. The initial frame of activation was used for all stimuli (see methods). The spatial spread was quantified as the length of the major axis, minor axis, and equivalent radius of the contour at 70% of the maximum amplitude. In this way, the spread measurement was normalized for differences in

the maximum amplitude across variations in stimulus intensity and stimulus type. Also, the contour was taken from the sum of Gaussians fit to the raw data in order to provide a smooth estimate to the contour. Qualitatively, for increasing velocity of whisker stimuli (figure 2.2(a), top) and light intensity of optical stimuli (bottom), the area of activation remains restricted to a small region of cortex. For increasing current amplitude of the electrical stimuli (middle), the area of activation is initially constrained, but spreads farther across cortex for higher current intensities.

The individual contours for all stimulus intensities are overlaid in figure 2.2(b) for each stimulus type (light gray is the lowest stimulus intensity, black is the strongest). Again, it is clear that the spatial spread grew slightly across varying whisker deflection velocities and optical light intensities, while the spatial spread grows more steeply for increasing current of electrical stimulation. The average radius of the contours is plotted as a function of the maximum response amplitude for all experiments in figures 2.2(c)-(e), for whisker, electrical, and optical stimuli, respectively, along with the linear fit of the relationship (black line). The linear fits to the trends for the major and minor axis radii are plotted on the same graph (dashed lines). As in the single experiment example, the average radius did not increase as a function of the response amplitude for whisker (slope: $15 \mu\text{m} / \text{unit of activity}$; 95% confidence interval: [-18, 48]) and optical (slope: $38 \mu\text{m} / \text{unit of activity}$; 95% confidence interval: [-11, 87]) stimuli. Electrical microstimulation, however, exhibited an increasingly large spatial spread as the response amplitude increased (slope: $108 \mu\text{m} / \text{unit of activity}$; 95% confidence interval: [68, 149]). For comparison, the radius of a cortical column in the rodent primary somatosensory cortex is approximately $400 \mu\text{m}$. For a strong response ($\Delta F/F_0$ of 0.4%), the average radius of the 70% contour for whisker ($474 \mu\text{m}$) and optical ($558 \mu\text{m}$) stimuli only extends in nearest neighbor cortical columns, whereas the 70% contour for electrical stimuli ($770 \mu\text{m}$) reaches two cortical columns outside of the principal column, encompassing approximately 2.6 times as much cortical surface

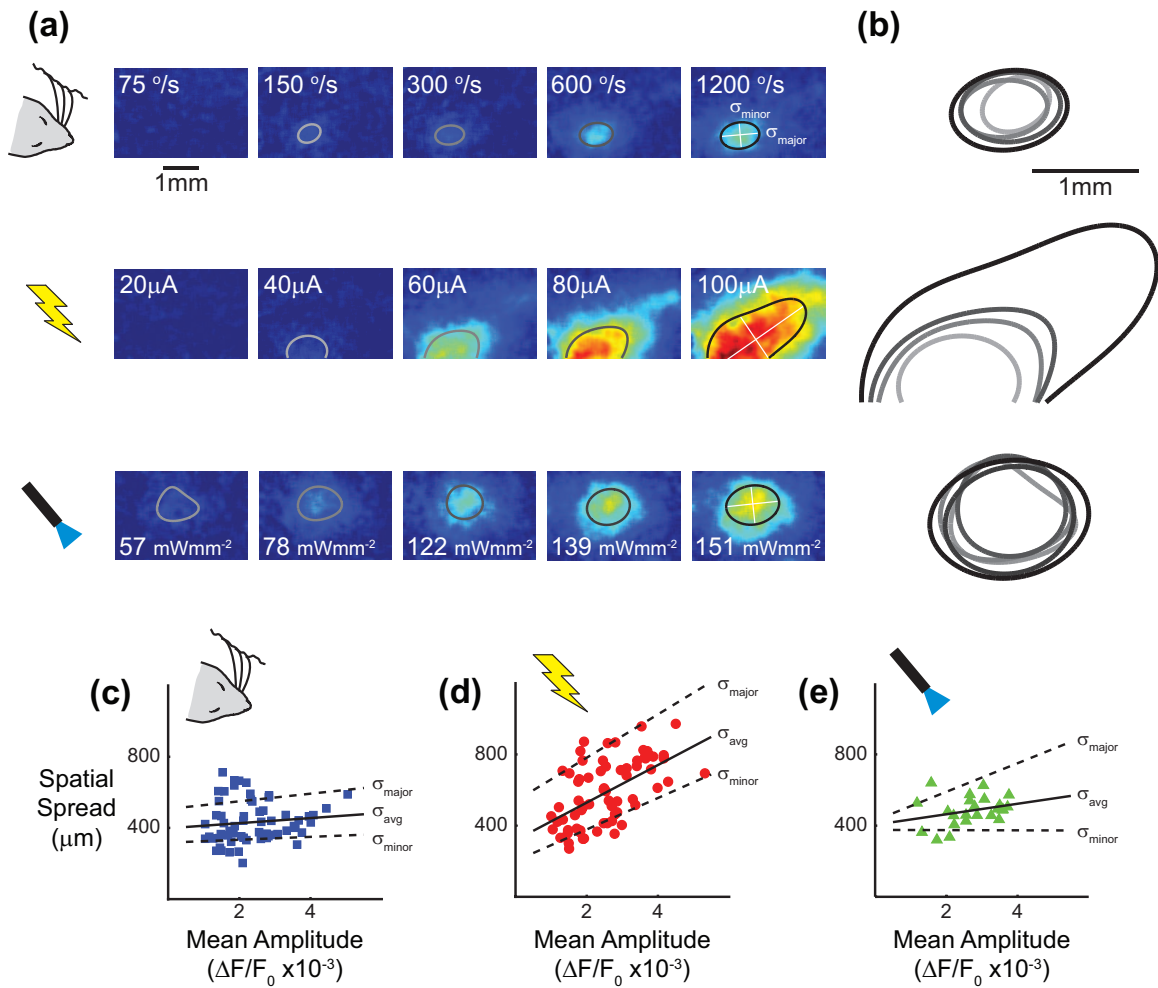


Figure 2.2: Electrical stimulation activates a significantly larger region of cortex. (a) The spatial spread was quantified as the average radius of the 70% amplitude contour. The initial frame of activation was used and the contour was taken from a parametrized sum of Gaussians fit to the VSDI data. (b) The size of the contours is similar across whisker deflection velocities (top) and optical light intensity (bottom), but increases dramatically for increasing electrical stimulation current (middle). (c) The spatial spread increases slightly as a function of response amplitude for whisker stimuli (N=12). (d) The spatial spread increases quickly as a function of response amplitude for electrical stimuli (N=17). (e) The spatial spread increases weakly as a function of response amplitude for optical stimuli (N=4).

as the response to whisker stimuli. The trends for the major and minor axis suggest a continuum between whisker, optical, and electrical stimuli, as the optical stimuli exhibited a significant increase in the major axis length, but an insensitivity of the minor axis as the response amplitude increased. These observations are consistent with the hypothesis that activation of axons locally in the thalamus leads to increased spatial spread in cortex given that electrical stimuli (Histed et al., 2009), and optical stimuli to a lesser extent (Cruikshank et al., 2010), are known to activate axons near the site of stimulation.

2.3.4 Computational model of thalamocortical circuit allows direct control of thalamic input

Striking differences were observed in the properties of the cortical response to whisker, electrical and optical stimuli, likely the result of the stimuli activating the thalamus in fundamentally different ways. Previous studies have shown that the timing precision and synchrony of population activity in the thalamus are modulated by variations in the strength of a sensory stimulus. Conversely, literature suggests that electrical and optical stimuli may synchronize the neurons they activate [191, 16], with higher current and light intensity causing the recruitment of more neurons. However, as the current and light spread through the tissue, they recruit different neural elements. Namely, electrical stimuli have been shown to preferentially activate axons [83], while optical stimuli are believed to stimulate cell bodies at a lower threshold [211]. While the literature suggests that whisker, electrical, and optical stimuli may produce different thalamic population activity, the relative contributions of the magnitude, identity, and synchrony to the static nonlinearity of the cortical response, and how this compares to the experimental results presented above, remains unknown.

In order to explicitly control which thalamic neurons were activated, and the degree of synchronization, we built a computational model of the thalamocortical projection. The network connectivity is illustrated in figure 2.3(a) and was based

off of previous anatomical and computational studies of the thalamocortical circuit [101, 26, 51]. The network was a canonical feedforward model, similar to what has been used extensively to explore neural activity propagation [100]. A population of thalamic neurons formed the input to the network and the cortical population of interconnected excitatory and inhibitory neurons was the output layer. The synaptic weights and relative connectivity were drawn from previous experimental and computational work in the rodent vibrissa system (see methods). However, to retain generality, the neurons in the model were arranged in a linear array, as opposed to the discrete nature of the cortical columns in the rodent barrel cortex. In this way, neighboring thalamic neurons synapsed on overlapping cortical populations. As the distance between thalamic neurons increased, the probability of shared connections to a single cortical neuron also decreased. The individual synaptic weights were also drawn according to a probability distribution, the details of which are found in the methods section.

Each individual cortical neuron was modeled as a nonlinear dynamical system [89]. Specifically, each neuron was described by two state variables, one being the membrane potential of the neuron, and the other being a recovery variable. The membrane potential evolved according to a quadratic integrate and fire model, while the recovery variable enforced the relative refractory period of the neuron and overall excitability. Example traces of the membrane potential of a single neuron in response to a simulated stimulus are shown in top portion of figure 2.3(a) for multiple trials. Noise was added to the membrane potential of each neuron in order to mimic the in vivo high conductance state (Destexhe et al., 2003; Ratt et al., 2013). The noise was correlated across the cortical neurons, such that neighboring neurons experienced similar fluctuations. The thalamic population activity acted as the input to the cortex and was generated by controlling the overall spike probability and the degree of synchrony across the thalamic neurons as illustrated by the inhomogeneous Poisson rate




Figure 2.3: Computational network model of activity propagation in the thalamo-cortical circuit. (a) The population activity of the VPM thalamus is controlled as the input to the cortical population. The neurons are arranged spatially in a linear array, such that neurons are more likely to be connected to their neighbors. The network promotes feedforward inhibition in that thalamic VPM neurons synapse on both excitatory and inhibitory neurons within the cortical population. Example membrane potential traces for a single neuron across multiple trials are shown in the top portion of the panel. (b) Cartoon illustrating the inhomogeneous rate function used to generate the thalamic spikes. The magnitude and synchrony of the thalamic spiking activity was systematically modulated as the input to the model. (c) Spatial PSTH and averaged PSTH for the excitatory (top left) and inhibitory (top right) cortical population in response to the thalamic input (bottom).

functions in 2.3(b). An example of the input thalamic population activity (bottom) and output cortical population activity (excitatory middle, inhibitory top) is shown in figure 2.3(c). For each cell type, the image shows the average activity across the entire population. The peristimulus time histogram (PSTH) for the central twenty neurons is shown above and the spatial cross-section, averaged across the first twenty milliseconds following the stimulus, is on the right. Using this modeling architecture, we evaluated the same response metrics as the experimental VSDI data, namely the response amplitude, variability, and spatial spread.

2.3.5 Magnitude and synchrony of population activity control nonlinear response in network model

The identity, magnitude, and synchrony of the input thalamic population activity were systematically varied to investigate the relative contributions of these fundamental variables on the response amplitude (figure 2.4(c)), variability (figure 2.4(d)), and spatial spread (figure 2.5(c)) within the model. For the amplitude and variability, two different inputs were used. In the top portion of figures 2.4(c) and 2.4(d), the average number of spikes per trial per neuron was fixed, but the synchrony of those spikes varied to modulate the intensity of the stimulus, as in the cartoon in figure 2.4(a). In the bottom portion of figures 2.4(c) and 2.4(d), the synchrony of the population spikes was held fixed at a high level, and the average number of spikes per trial per neuron was varied to modulate the intensity of the stimulus, as in the cartoon in figure 2.4(b). The response amplitude was measured as the average number of spikes in the twenty milliseconds following the stimulus for the central twenty excitatory cortical neurons, and grew monotonically as a function of the input synchrony (top) and number of spikes (bottom). Both were well approximated by a sigmoidal function, similar to the experimental results of figure 2.1(a).

The trial to trial variability was measured as the standard deviation across trials

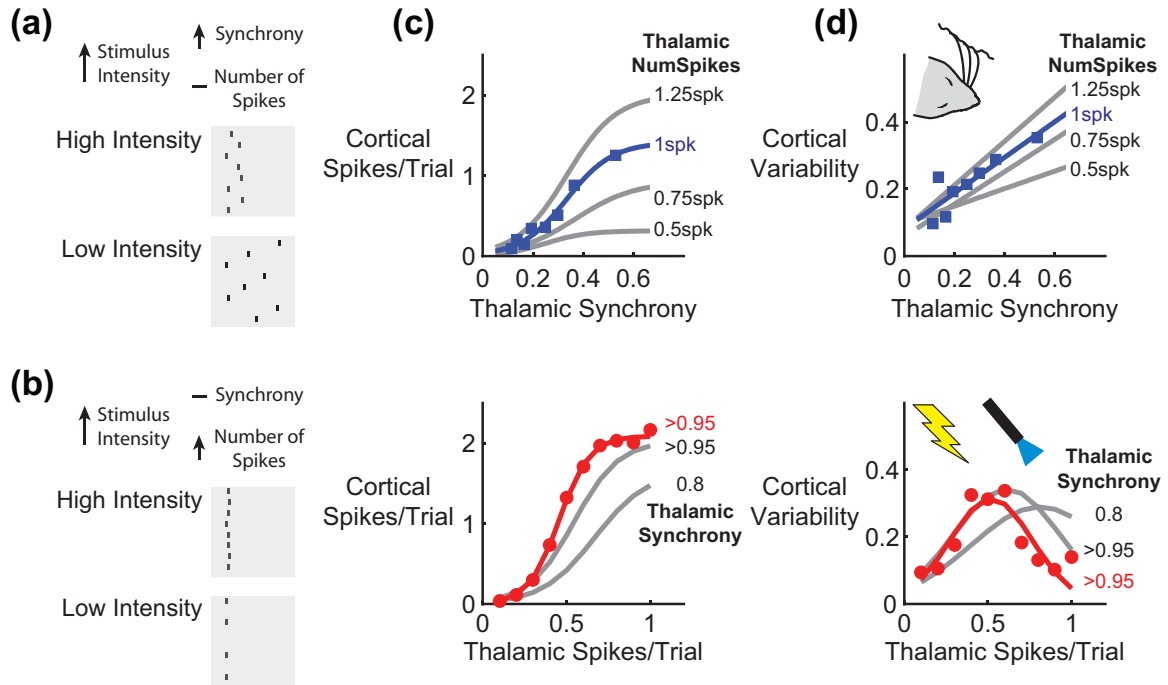


Figure 2.4: Extreme synchrony ensures robust response in the face of trial-to-trial variability. (a) Cartoon of the input population activity with a fixed number of spikes per trial, but at varying degrees of synchrony. This input stimulus was used for the top portion of (c) and (d). (b) Cartoon of the input population activity with a fixed, high synchrony, but an increasing number of spikes per trial per neuron. This input stimulus was used for the bottom portion of (c) and (d). (c) The response amplitude of the central excitatory units in response to systematic variation in the number and timing of the input spikes. The gray lines indicate alternate values of precision and spike magnitude. The cortical response is nonlinear with respect to both the number and timing of the spikes. (d) The variability in the response amplitude of the central excitatory units in response to systematic variation in the number and timing of the input spikes. Again, the gray lines indicate alternate values of precision and spike magnitude. The variability increased with synchrony at a fixed number of spikes (top), whereas the variability peaked and returned to baseline for a fixed, high synchrony (bottom).

of the mean number of spikes following the stimulus across the central twenty excitatory cortical neurons. As a function of the input synchrony (figure 2.4(d), top), the variability increased linearly, with a greater slope for larger input population activity magnitude. This trend in variability matches the experimental observations for whisker stimuli in figure 2.1(c), providing support for the hypothesis that synchrony is an important variable for coding sensory stimuli. Meanwhile, the variability at a fixed, high synchrony, but increasing number of input spikes (figure 2.4(d), bottom) peaked at the threshold input intensity. For lower levels of fixed synchrony (gray lines), the peak in the variability shifted to the right and approached the linear trend observed for input synchrony (figure 2.4(d), top). The peak variability at threshold matched the experimental observations for electrical and optical stimuli in figure 2.1(c), suggesting that the extreme synchronization generated by electrical and optical stimuli was responsible for the quenched variability at supra-threshold stimulus intensities.

In exploring the spatial spread of the cortical response in the model, the identity of the thalamic population proved to be the most important variable. Again, two different stimulus profiles were used. To model the spread associated with natural recruitment of sensory neurons, the number and synchrony of the input spikes were maximal at the center of the stimulus and fell off for distant neurons, as in figure 2.5(a). The relative contribution of the stimulus surround was modulated by changing the rate at which the number and synchrony of the input spikes fell off across space. Alternately, the recruitment of axons by electrical stimulation was modeled by uniformly recruiting neurons in the surround of the stimulus, regardless of the distance from the stimulus center, as in figure 2.5(b). The proportion of axons was varied in order to modulate the surround in this case. Importantly, the timing of stimulation of axons and cell bodies in figure 2.5(b) was fixed at a high synchrony.

In figure 2.5(c), the spread measurement was calculated from a Gaussian fit to the




Figure 2.5: Axonal stimulation increases spatial spread by overcoming feedforward inhibition. (a), (b) The spatial spread of the input was modulated in two different ways. (a) The number of spikes and synchrony were a function of the spatial location of the cell. The spatial distribution of the number of spikes and synchrony varied simultaneously for the top portion of (c). (b) The number of spikes was a function of the spatial location of the cell, but the synchrony was fixed at a high level. To simulate the activation of axons, the surrounding neurons were recruited uniformly, as opposed to (a) where the nearby neurons were activated first as the spread increased. (c) For a given spread of the number of spikes and synchrony of the input, the spatial spread of the simulated excitatory cortical neurons remained constant, or slightly decreased, as a function of the response amplitude (top). The gray lines indicate alternate values of input spread. When axons were stimulated, according to the description in (b), the spread in the output instead increased as a function of the response amplitude (bottom). The gray lines indicate alternate values for the proportion of axons stimulated.

spatial cortical response, similar to the analysis in figure 2.3, and was independent of the response amplitude. For a fixed input spread (top), the output spread was constant, or decreased, as a function of the response magnitude when the number and synchrony of spikes decreased with distance from the center of the stimulus, as in figure 2.5(a). This trend was consistent across a wide range of the input spread and was similar to the experimental results for whisker and optical stimuli in figures 2.2(c) and 2.2(e), respectively. Mechanistically, the focal spatial activation in the model was dependent on spatially diffuse feedforward inhibition. The feedforward inhibition recruited by the synchronous and low latency response in the center of the stimulus prevented excitation in the surround by the less synchronous and longer latency peripheral input.

Axonal stimulation (figure 2.5(c), bottom), however, led to an increase in the spatial spread as a function of the response amplitude, similar to the experimental results observed for electrical stimuli in figure 2.2(d). The gray lines display the trend for varying ratios of axon to cell body activation. Increasing the proportion of axons stimulated accentuated this effect, while stimulating only the cell bodies eliminated it. The increasing spread in cortex was mediated by the spatially diffuse, but temporally precise, activation of axons within the input stimulus surround, which, due to the high synchrony, activated the cortical surround before feedforward inhibition could begin to restrict the spatial spread.

The computational results support the hypothesis that the static nonlinear properties, including the amplitude, variability, and spatial spread, of the cortical response are determined by the identity, magnitude, and synchrony of the input population. Specifically, the model predicts that the extreme synchronization believed to occur with electrical and optical stimuli caused the highly nonlinear transformation in cortex, whereas the stimulation of cell bodies, and not axons, is responsible for the spatially constrained cortical response to whisker and optical stimuli.

2.4 Discussion

A combined experimental and computational approach was employed to assess the nonlinear propagation of sensory- and artificially-evoked neural activity to downstream structures. We specifically studied this in the context of neural propagation in the thalamocortical circuit, with voltage sensitive dye imaging used to measure the cortical response to whisker, electrical, and optogenetic inputs within the rodent vibrissa somatosensory system. Using this technique we found that whisker stimuli of increasing velocity produce a weakly nonlinear increase in the mean response amplitude, with high trial-to-trial variability, in a fixed, focal spatial location. Electrical stimuli, on the other hand, exhibited a highly nonlinear modulation of the response amplitude, with little trial-to-trial variability except at the threshold current, over an increasingly large region of the cortical space. Optical stimuli displayed a mixture of the properties observed with whisker and electrical stimuli, where the nonlinearity of the amplitude and variability closely matched that of electrical stimulation, and the spatial response properties were more akin to that of whisker stimuli.

A computational model of the thalamocortical circuit was then used to test the hypothesis that the differential response properties in cortex to whisker, electrical, and optical inputs were due to the distinct ways in which the stimuli modulated the identity, magnitude, and synchrony of the neurons activated in the thalamus. The model attributed the quenched variability at supra-threshold intensities of electrical and optical stimuli to the extreme synchronization of the thalamic population, with increasing stimulus intensity recruiting a growing number of neurons. Meanwhile, the model indicated that the spatially focused cortical activation caused by whisker and optical stimuli was due to the preferential stimulation of cell bodies as opposed to fibers of passage. Taken together, the experimental and computational manipulations of the identity, magnitude, and synchrony of thalamic population activity, combined with measurement of the resulting downstream cortical activity, reveal the relative

contributions of these fundamental variables of population activity to the nonlinear processing of the thalamocortical circuit and the functional relevance for delivering information to the brain.

The thalamocortical circuit was used for this study due to the extensive literature detailing the anatomy of the network [51]. Thalamic neurons are known to make weak, noisy synapses onto downstream cortical neurons, and these connections comprise a relatively small percentage of the incoming synapses compared to intracortical connections [26]. This anatomical wiring alone has led many to suggest that synchronization of the thalamic neurons be necessary for the reliable activation of downstream cortical neurons [24, 177]. The experimental and computational results of this study further support the importance of thalamic synchronization in driving cortical activity. Within the model, high synchrony was found to reduce trial-by-trial variability as the number of synchronously active neurons increased above threshold. Experimentally, the electrical and optical stimuli, which are believed to synchronize neurons [191, 16], produced a similar reduction in the trial-by-trial variability in the cortical response as measured with voltage sensitive dye imaging. However, outside of the context of the anatomical wiring in the thalamocortical projection, the relative importance of synchronization on neural propagation may be substantially different.

Indeed, an entire computational field has been built on exploring the consequences of neural circuit anatomy on propagation within networks [100]. The work has primarily been based on the use of the “feedforward network” architecture, where the rate of convergence and divergence of synaptic connections from one stage to the next, and the strength of these connections, are the key parameters defining the anatomy of a neural circuit [100]. For a network with sparse connectivity and high synaptic weights, synchrony does not play a significant role in activity propagation, and can even reduce the information content of the firing rate code [121]. However, for dense connectivity and weak synaptic strengths, synchrony emerges and ensures

the reliable propagation of activity to downstream structures [52]. To study the role of precise timing in propagation within simulated neural networks, the concept of a “pulse packet” was introduced [1]. A pulse packet describes the population activity input that initiates the propagation downstream, where the number of spikes and their relative timing are the only two parameters. Similar to the results in this study, previous simulation work has systematically varied the pulse packet parameters and quantified the transformation across a single stage of the feedforward network, finding that the timing precision determines the sensitivity of the nonlinearity in spiking of the downstream neural population [1].

While not a generalized feedforward network, the model used in this study contains dense, weak synaptic connections and a phenomenologically accurate neuron model. Previous work has utilized the integrate and fire neuron, whereas this study used a quadratic integrate and fire neuron with a recovery variable, which has been shown to accurately describe spike initiation dynamics for a broad class of neurons [89]. Additionally, the properties of the input were varied systematically, similar to the pulse packet parameters described above. The model in this study expanded the concept of the pulse packet spatially, such that the location/identity of the neurons activated was controlled in concert with the number and precision of spikes, allowing direct comparison to the spatial cortical measurements afforded by the VSDI technique. The model only contained a single projection, from the thalamus to layer 4 in cortex, whereas VSDI primarily represents subthreshold layer 2/3 activity [72]. However, the stimulus-driven onset of cortical activity, as measured through VSDI, reflects the supra-threshold spiking activity projecting from layer 4 [152, 11].

The agreement in the computational and experimental results supports a conceptual model for the distinct ways in which natural and artificial inputs activate neural circuits. Sensory stimuli modulate thalamic synchrony to drive downstream cortical activity, whereas electrical and optical stimuli activate a population of neurons at a

fixed, high synchrony. This dichotomy has direct relevance to the use of electrical and optical stimuli as surrogates for sensory stimuli in driving neural circuits. An obvious example is the development of sensory prosthesis applications where the goal is to produce functionally relevant patterns of activation in sensory neural circuits to replace missing afferent input. The cochlear implant [202], and more recently the retinal prosthesis [87], has been effective at restoring functional levels of sensation through electrical stimulation of the peripheral nervous system. However, delivery of artificial inputs to the central nervous system has proven more difficult [19, 17, 82, 81, 71], potentially due to the extreme synchronization and mixed activation of axons and cell bodies caused by electrical stimulation. Behavioral experiments in non-human primates have indicated that sensory and electrical stimuli produce functionally different neuronal dynamics [119], while anesthetized work has revealed complex signal propagation generated by electrical stimuli in the primate visual system [112]. These differences may arise due to the extreme synchronization of electrical stimuli, and specifically through the differential recruitment of inhibitory elements in the circuit [29, 27].

More generally, the distinction between sensory and artificial stimuli applies to the use of surrogate stimuli in dissecting sensory neural circuits. Recently, optogenetic stimuli have become commonplace in exploring neural circuit function across a variety of brain regions [99, 211, 212, 105]. Pulsed excitation via Channelrhodopsin (ChR2), or other excitatory opsins, may lead to unnaturally high levels of synchrony and engage the circuit in a pathological, or at least non-physiological, manner. Principled design of the stimulation profiles and parameters for electrical and optical stimuli may be able to overcome the limitations described here, generating neural activity patterns more akin to those generated by sensory stimuli. For instance, ramp stimuli have been used for optogenetics on long timescales in order to avoid synchronization at the onset of the stimulus [159]. Additionally, it remains unknown how the underlying

brain state may affect the results described here. Under anesthesia, the thalamic state in these experiments can be characterized by low spontaneous firing rates and a high degree of bursting. In this way, the neurons in the thalamus may have been overly susceptible to the synchronization caused by electrical and optical stimuli. However, in a desynchronized thalamic state, characterized by high spontaneous firing rates and few bursts, intermittent and temporally dispersed refractory periods across the thalamic population may limit the overall degree of stimulus driven synchronization, leading the neural response properties induced by artificial stimuli to approach those of natural sensory stimulation. Ultimately, knowledge of the mechanisms by which sensory and artificial stimuli activate neural circuits informs the design of stimulus patterns or signal sets that optimize the information capacity of artificial stimuli and utility in conveying surrogate sensory signals.

CHAPTER III

CHARACTERIZATION OF THE NONLINEAR DYNAMICS OF CORTICAL ACTIVATION IN RESPONSE TO SURROGATE THALAMIC DRIVE

Portions of this work were presented in poster form at the following conference:

Millard, DC, Wang, Q, Gollnick, CA, & Stanley, GB. Characterization of the dynamically varying spatiotemporal cortical response to patterns of sensory and electrical stimulation in the thalamocortical network. Society for Neuroscience Annual Meeting, New Orleans, LA, October 2012.

3.1 Introduction

In the previous chapter, sensory and artificial stimuli were found to activate the thalamocortical circuit with distinct neural response properties, even for single stimuli. However, the delivery of spatiotemporal patterns of activity will be required to faithfully represent the dynamically changing sensory experience in a prosthetics applications. Further, the use of patterned stimuli is an important tool for revealing neural circuit dynamics in an experimental setting. Here, in this chapter, we investigated the nonlinear dynamics of the thalamocortical circuit in response to simple patterns of sensory and artificial stimuli.

Sensory stimuli are known to be coded by dynamic patterns of action potentials across neural circuits [86, 51, 204]. Furthermore, the majority of experimental evidence suggests a nonlinear dynamical relationship. The simplest example is the use of the “paired pulse” stimulus, in which two stimuli are delivered in quick succession, separated in time by a varying inter-stimulus interval. By comparing the response of

the second stimulus to the first, the nonlinear dynamics of the circuit can be characterized as suppressive (i.e. the second response is smaller than the first) or facilitative (i.e. the second response is stronger than the first). This paradigm has been particularly well-studied in the rodent vibrissa system [173, 33, 15, 14, 197, 198]. However, because the majority of previous studies have utilized single electrode recordings, the spatial properties associated with the nonlinear dynamics remains unknown.

Additionally, the comparative nonlinear dynamics for sensory and artificial stimuli have not been studied in detail. Electrical stimulation has been used extensively in vitro to study synaptic and neural circuit dynamics when sensory or other external stimulation was not possible [190, 191]. Similarly, electrical stimulation has been used in vivo to probe the nonlinear dynamics of hippocampal circuits [175, 10], among others, for which natural stimuli are not clearly defined or easily presentable within an experimental setting. Yet, recent work has provided behavioral [119] and electrophysiological [112] evidence that the nonlinear dynamics, within the same neural circuit, are different for sensory and electrical stimuli. The extent to which this is generally true across neural circuits, and the mechanism by which this occurs, is not well understood and has implications for the exploration of neural circuit dynamics and the delivery of time varying surrogate sensory inputs to the brain.

In this chapter, we experimentally investigate the nonlinear dynamics of the thalamocortical circuit of the rodent vibrissa somatosensory system in response to whisker, electrical, and optogenetic inputs using voltage sensitive dye imaging to record the spatiotemporal stimulus-evoked cortical activity. We find that while both whisker deflections and thalamic microstimulation produced paired pulse suppression for supra-threshold stimuli, thalamic microstimulation activated an additional dynamical mode characterized by paired pulse facilitation of sub-threshold inputs. Additionally, the facilitation dynamics caused by thalamic microstimulation induced a significant spatial

sharpening of the cortical response. Through a series of further experiments and computational results, we suggest that the additional dynamical mode of facilitation was caused by the activation of thalamic bursting through the extreme synchronization of the reticular thalamus by electrical and optogenetic stimulation of the thalamus.

3.2 Methods

3.2.1 Surgery and preparation

The same methods described previously (see Appendix A for more detail) were used for the experiments in this chapter. Briefly, voltage sensitive dye imaging was used to record the response of primary somatosensory cortex in the anesthetized rodent to whisker deflections on the face and thalamic microstimulation.

3.2.2 Thalamic microstimulation

An electrode positioned in the thalamus was used to deliver single electrical current pulses to evoke cortical responses in the somatosensory pathway. The electrical stimuli were created using a digital stimulus generator (Model: DS8000, WPI Inc., Sarasota, Florida) and delivered using a digital linear stimulus isolator (Model: DLS 100, WPI Inc., Sarasota, Florida) acting in current source mode. All individual electrical stimuli were charge balanced, cathode-leading, symmetric biphasic waveforms of 200 microseconds duration per phase. A paired pulse stimulus paradigm was used, where two stimuli with the same current amplitude were delivered separated by an inter-stimulus interval. Current amplitudes (30-150 microamps) and inter-stimulus intervals were varied (50, 100, 150, 200, 250, 500 milliseconds) were systematically varied.

3.2.3 Whisker stimulation

Sensory stimulation (S-Stim) was applied through computer controlled whisker deflections. Whiskers were trimmed at approximately 12mm from the face, and were

inserted into a glass pipette fixed to the end of a calibrated multi-layered piezoelectric bimorph bending actuator (range of motion, 1 mm; bandwidth, 200 Hz; Physik Instrumente (PI), Auburn, MA) positioned 10 mm from the vibrissa pad. Vibrissae were always deflected in the rostral-caudal plane. Punctate deflections consisted of exponential rising and falling phases (99% rise time, 5 ms; 99% fall time, 5 ms) and angular deflection velocities of 75, 150, 225, 300, 450, 600, 900, and 1200 deg/s were used as mechanical probe stimuli (S-Stim). The same inter-stimulus intervals listed above were used for whisker stimuli in the paired pulse paradigm.

3.3 Results

Neural activity varies on fast timescales within the early stages of sensory pathways. In this way, surrogate sensory inputs must contain temporal patterns of stimulation to faithfully represent functionally relevant signals in the brain. Here, we experimentally measure the nonlinear dynamics of the thalamocortical circuit to simple patterns of whisker and thalamic microstimulation inputs using voltage sensitive dye imaging (VSDI). As with the static nonlinearity in Chapter 2, we find that the dynamic nonlinearity exhibits distinct properties for whisker and thalamic microstimulation inputs, but that the additional mode observed for thalamic microstimulation may provide functional benefits by spatially sharpening the response in cortex.

3.3.1 Nonlinear paired pulse suppression of the cortical response to whisker stimuli

Temporal patterns of sensory input are known to produce nonlinear responses even in the early sensory pathways. A significant portion of this literature comes from experiments in the thalamocortical circuit of the rodent whisker system, using pairs or triplets of individual whisker deflections to map out the nonlinear dynamics of neurons in the primary somatosensory cortex [173, 33, 15, 14, 197, 198]. These studies have principally detailed “paired-pulse suppression”, where the response to the second

stimulus in a pair is suppressed due to the presence of a previous response occurring immediately prior. Using the voltage sensitive dye imaging setup described earlier (and in Appendix A), we investigated the spatiotemporal nonlinear properties of the cortical response to pairs of whisker stimuli.

We used pairs of whisker stimuli at varying inter-stimulus intervals and varying velocity to directly probe the second order nonlinear dynamics of the system. Panels (a) and (b) of figure 3.1 show an example of the temporal response in cortex, averaged across trials, to pairs of whisker deflections with low and high velocity, respectively. For the low velocity stimuli in figure 3.1(a), the response to the second stimulus was equal in strength to the response to the first stimulus, except at the 50ms inter-stimulus interval, indicating a very short duration for the nonlinear dynamics. On the other hand, the high velocity stimuli in figure 3.1(b) exhibited profound paired pulse suppression. The response to the second stimulus was strongly suppressed relative to the first. The suppression relaxed for longer inter-stimulus intervals, but at a much slower time constant than for the low velocity whisker stimuli in figure 3.1(a). This suggests an activity dependent modulation of the nonlinear dynamics, consistent with previous reports in the barrel cortex [14]. Whereas Boloori et al. found the cortical response suppression lasted for ~ 100 ms, the cortical response did not fully recover until >500 ms after the first stimulus in our experiments. The increased duration of suppression may be related to inhibitory mechanisms specific to layer 2/3, from which the majority of the VSDI signal originates. But in general, across both studies, the suppression was more profound for stronger sensory inputs.

To determine whether the nonlinear suppression was consistent on a trial-to-trial basis, we analyzed single trials of the VSDI signal in figure 3.1(c) for the 150ms inter-stimulus interval. On a single trial basis, the peak response to the second stimulus was plotted relative to the peak response to the first stimulus. Each data point is in response to a single presentation of the paired pulse stimulus, with the color

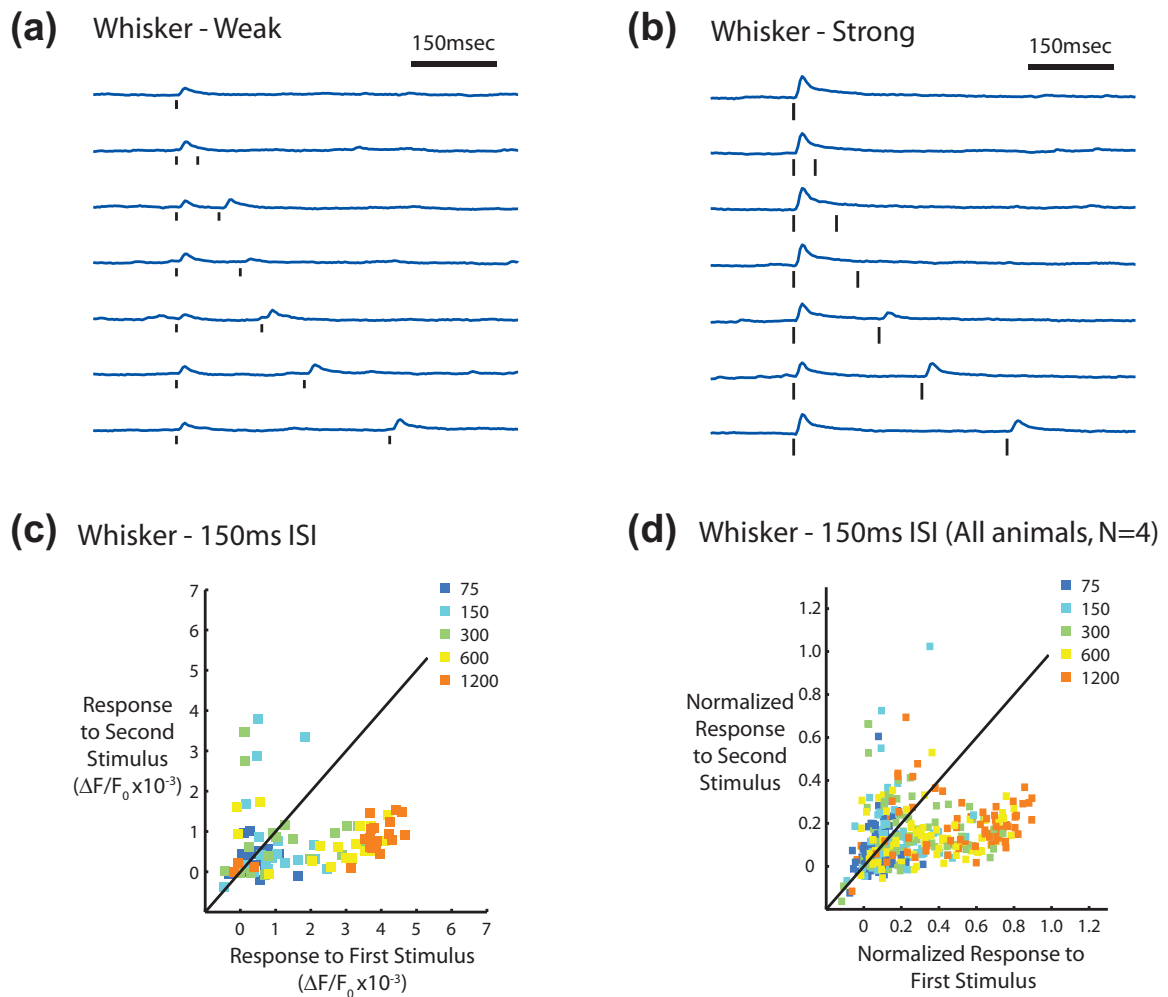


Figure 3.1: Paired pulse suppression of the cortical response to whisker stimuli. (a) Cortical response to pairs of weak whisker stimuli. The response to the second stimulus was suppressed only for the 50ms inter-stimulus interval. (b) Cortical response to pairs of high velocity (1200 deg/sec) whisker deflections at varying inter-stimulus interval. The response to the second stimulus was suppressed relative to the first, with the suppression relaxing for longer inter-stimulus intervals. The response to the first and second stimulus are plotted for all single trials across all whisker deflection velocities at the 150ms inter-stimulus interval for a single experiment (c) and all experiments (d). When plotted across all experiments, the response amplitudes were normalized to the mean response amplitude to the strongest stimulus.

indicating the velocity of the whisker deflections. As the velocity of the whisker deflections increased, the response to the first stimulus also increased, consistent with the results in Chapter 2. However, for all velocities, the response to the second stimulus was suppressed relative to the first for the majority of trials, as most of the data points fell below the unity line (black). A few trials at low velocity showed strong responses to the second stimulus, with no response to the first, but this occurred at low probability. This is consistent with the response to the second stimulus being independent of the first, and responding probabilistically in the same manner as the first response. Figure 3.1(d) summarizes the data across four experiments, where the data from each experiment has been normalized. Again, the majority of trials lie below the unity line, indicating a prevalence for paired pulse suppression amongst the cortical population in response to pairs of whisker deflections.

3.3.2 Stimulus-dependent bimodal nonlinear dynamics of the cortical response to thalamic microstimulation

Within the same experiments, we also recorded the cortical response to pairs of thalamic microstimulation pulses, in order to compare the nonlinear dynamics between sensory and artificial stimuli. Figure 3.2 presents the results from these experiments. Panels (a) and (b) of figure 3.2 show an example of the temporal response in cortex, averaged across trials, to pairs of thalamic microstimuli with low and high current, respectively. While the high current stimuli in figure 3.2(b) exhibited paired pulse suppression similar to high velocity whisker deflections in figure 3.1(b), the low current stimuli in figure 3.2(a) produced a dramatically different nonlinear response. For inter-stimulus intervals of 100-200ms, the response to the second stimulus was strongly facilitated even though the first stimulus was sub-threshold and produced no cortical response.

The facilitation and suppression were consistent across trials for sub- and supra-threshold stimuli, respectively. The response to the second stimulus is plotted relative

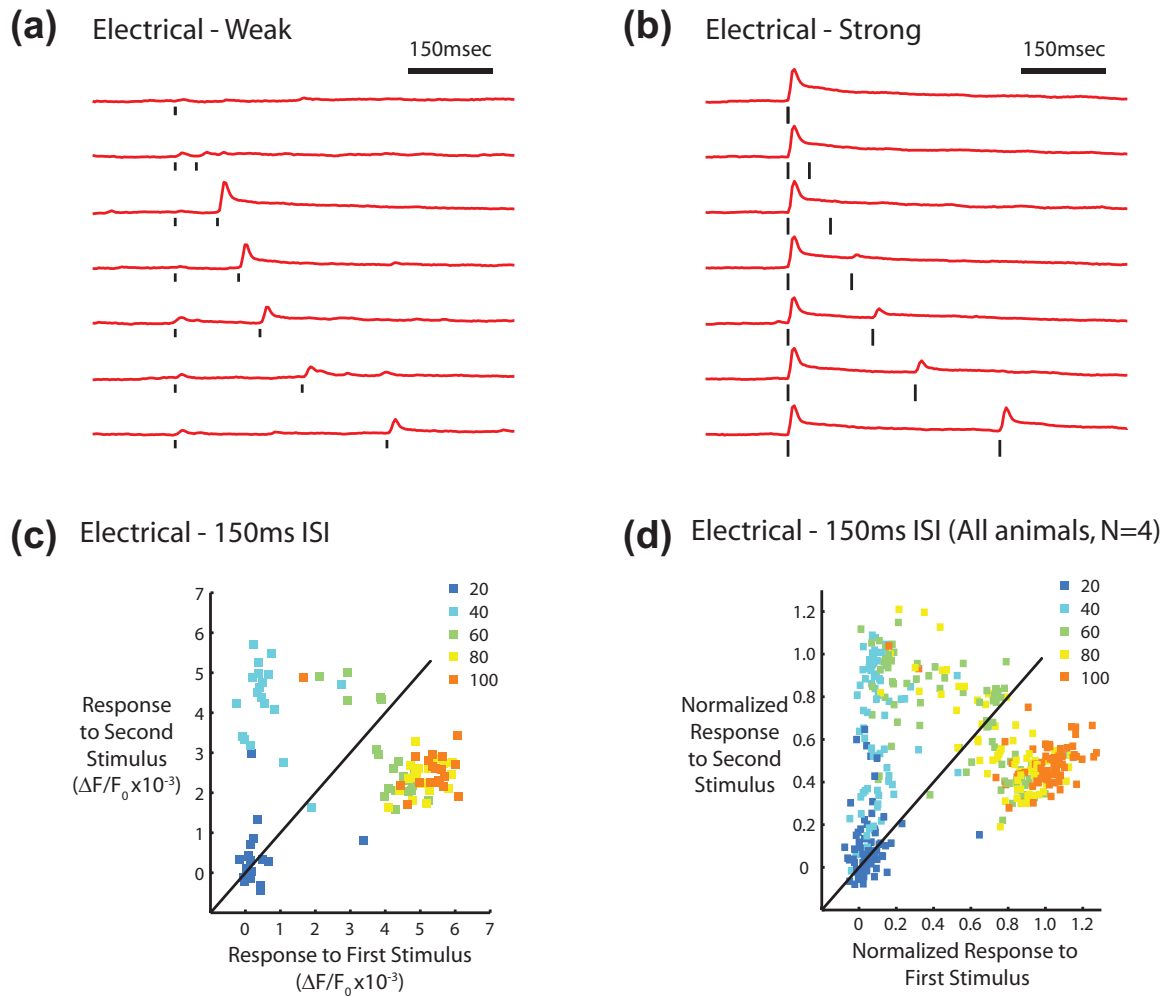


Figure 3.2: Bimodal nonlinear dynamics of the cortical response to thalamic microstimulation. (a) Cortical response to pairs of weak electrical stimuli. The response to the second stimulus was facilitated relative to the first, but only for the 100-200ms inter-stimulus intervals. (b) Cortical response to pairs of strong ($100\mu\text{A}$) electrical stimuli at varying inter-stimulus interval. The response to the second stimulus was suppressed relative to the first, with the suppression relaxing for longer inter-stimulus intervals. The response to the first and second stimulus are plotted for all single trials across all electrical stimulus currents at the 150ms inter-stimulus interval for a single experiment (c) and all experiments (d). When plotted across all experiments, the response amplitudes were normalized to the mean response amplitude to the strongest stimulus.

to the response to the first for single trials and varying current amplitude in figure 3.2(c) for the 150ms inter-stimulus interval. For the lowest current intensity (dark blue), there was no response to either stimulus presentation. At a slightly higher current intensity (light blue), the response to the second stimulus was strongly facilitated for each trial, leading to a cluster of points in the upper left portion of figure 3.2(c). Importantly, the clustering in the upper left portion of the axes distinguished the reliable facilitation of electrical stimulation from the lack of nonlinear dynamics for sub-threshold whisker inputs in figure 3.1(c),(d). At the highest current intensities (yellow and orange), the response to the second stimulus was consistently suppressed relative to the first, leading to a cluster in the lower right portion of the axes. At a threshold stimulus intensity (green), the response to the first stimulus varied significantly on a trial-to-trial basis, as seen in Chapter 2, while the response to the second stimulus was strongly negatively correlated with the response to the first stimulus. In this way, a weak response to the first stimulus was associated with facilitation of the second response, while a strong response to the first stimulus elicited suppression of the second response, all for the same stimulus intensity.

The data is summarized in figure 3.2(d) across animals, where the responses were normalized within experiments. Again, the cluster in the top left indicates facilitation, whereas the cluster in the bottom right represents suppression. Threshold stimuli occupy the space between the two clusters, where the response to the second stimulus was negatively correlated with the response to the first stimulus on a trial-to-trial basis. The consistency of the pattern across animals (N=4) established the distinct bimodality of the cortical nonlinear dynamics in response to thalamic microstimulation as compared to sensory stimuli.

3.3.3 Facilitation from sub-threshold thalamic microstimulation is a general phenomenon

Thalamic microstimulation produced activity-dependent bimodal nonlinear dynamics in the cortical response, while whisker stimuli only produced suppression. Here we directly tested whether the facilitation dynamics were exclusive to pairs of electrical stimuli, or if combined patterns of whisker and electrical inputs could cross-facilitate. Specifically, we explored the ability of whisker stimuli to facilitate thalamic microstimulation, and vice versa.

First, we identified a sub-threshold stimulus intensity of thalamic microstimulation that was known to cause the nonlinear facilitation described in the previous section at a 150ms inter-stimulus interval. We refer to this stimulus as A_F . Figure 3.3(a) shows the temporal response (top) to thalamic microstimuli of increasing current intensity when preceded by A_F . The maximum amplitude of the cortical response to the test stimuli is plotted as a function of the stimulus current (bottom), with (gray) and without (black) the conditioning stimulus, A_F . The conditioning stimulus dramatically facilitated the cortical response to the test electrical stimulation, shifting the stimulus-response curve to the left by $50\mu A$.

Now, in figure 3.3(b), we preceded whisker deflections of varying velocity by the same A_F conditioning stimulus. The temporal response, averaged across trials, is shown in the top portion of the figure, while the maximum response as a function of the whisker deflection velocity is shown in the bottom portion of the figure. When the whisker stimuli were preceded by the conditioning electrical stimulus, A_F , the stimulus-response curve was again shifted to the left, indicating that electrical stimuli were capable of facilitating the response to whisker stimuli.

By comparison, figure 3.3(c) presents the cortical response to A_F when preceded by whisker deflections of varying velocity. From the temporal response, averaged across trials, in the top portion of figure 3.3(c) it is clear that whisker stimuli, of

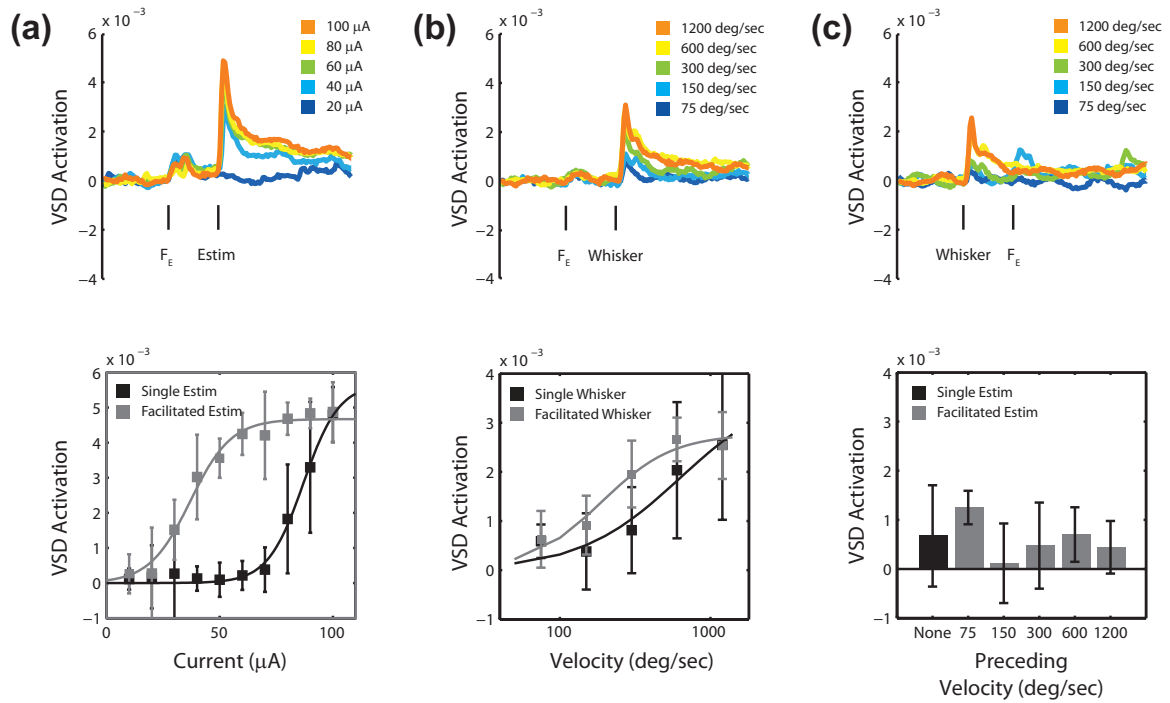


Figure 3.3: Thalamic microstimulation facilitates whisker and electrical inputs. (a) Top: A sub-threshold electrical stimulation current was delivered 150ms before a second stimulus of varying stimulus intensity. Bottom: The sub-threshold electrical stimulus shifted the current-response relationship to the left compared to when only a single stimulus was delivered. (b) Top: A sub-threshold electrical stimulation current was delivered 150ms before a whisker deflection of varying velocity. Bottom: The sub-threshold electrical stimulus also shifted the velocity-response relationship to the left, as compared to a single whisker deflection. (c) Top: A whisker deflection of varying velocity was delivered 150ms before an electrical stimulus with the same current that produced facilitation in (a) and (b). Bottom: While the current caused facilitation in (a) and (b), no whisker velocity was able to cause facilitation of the electrical stimulus.

any velocity, were not capable of facilitating an electrical stimulus. The data are summarized in the bottom portion of figure 3.3(c), where the black bar represents the mean response and variability to the sub-threshold electrical stimulus without any preceding whisker stimulus. The gray bars indicate the mean response and variability to the sub-threshold electrical stimulus with preceding whisker stimuli of increasing deflection velocity.

These results suggest that the conditioning electrical stimulus activates a state trajectory in the thalamocortical circuit capable of facilitating any subsequent input into the circuit. Whisker stimuli, on the other hand, were not capable of activating this state trajectory in the thalamocortical circuit. In this way, the facilitation was a general phenomenon of the thalamocortical circuit, but could only be engaged by microstimulation.

3.3.4 Nonlinear facilitation sharpens the spatial cortical response

Using the spatiotemporal VSDI signal, we analyzed the spatial properties of the nonlinear dynamics detailed above. Figure 3.4 presents an example of the spatiotemporal evolution of the cortical response to supra-threshold whisker and electrical inputs, in panels (a) and (b), respectively, along with the cortical response to a facilitated electrical stimulus in panel (c). Consistent with the results in Chapter 2, the cortical response to a whisker deflection in figure 3.4(a) was constrained to a focal location in cortex, while the response to thalamic microstimulation in figure 3.4(b) spread across a large region of cortex. The cortical response to the facilitated electrical stimulus, however, was more similar to the response to a whisker deflection than to an electrical stimulus. The activation was spatially localized to a small region of cortex topographically aligned to the position of the electrode in the thalamus.

The initial frames of activation in response to the three conditions are directly compared in figure 3.4(d), where the cortical columns within the row containing the

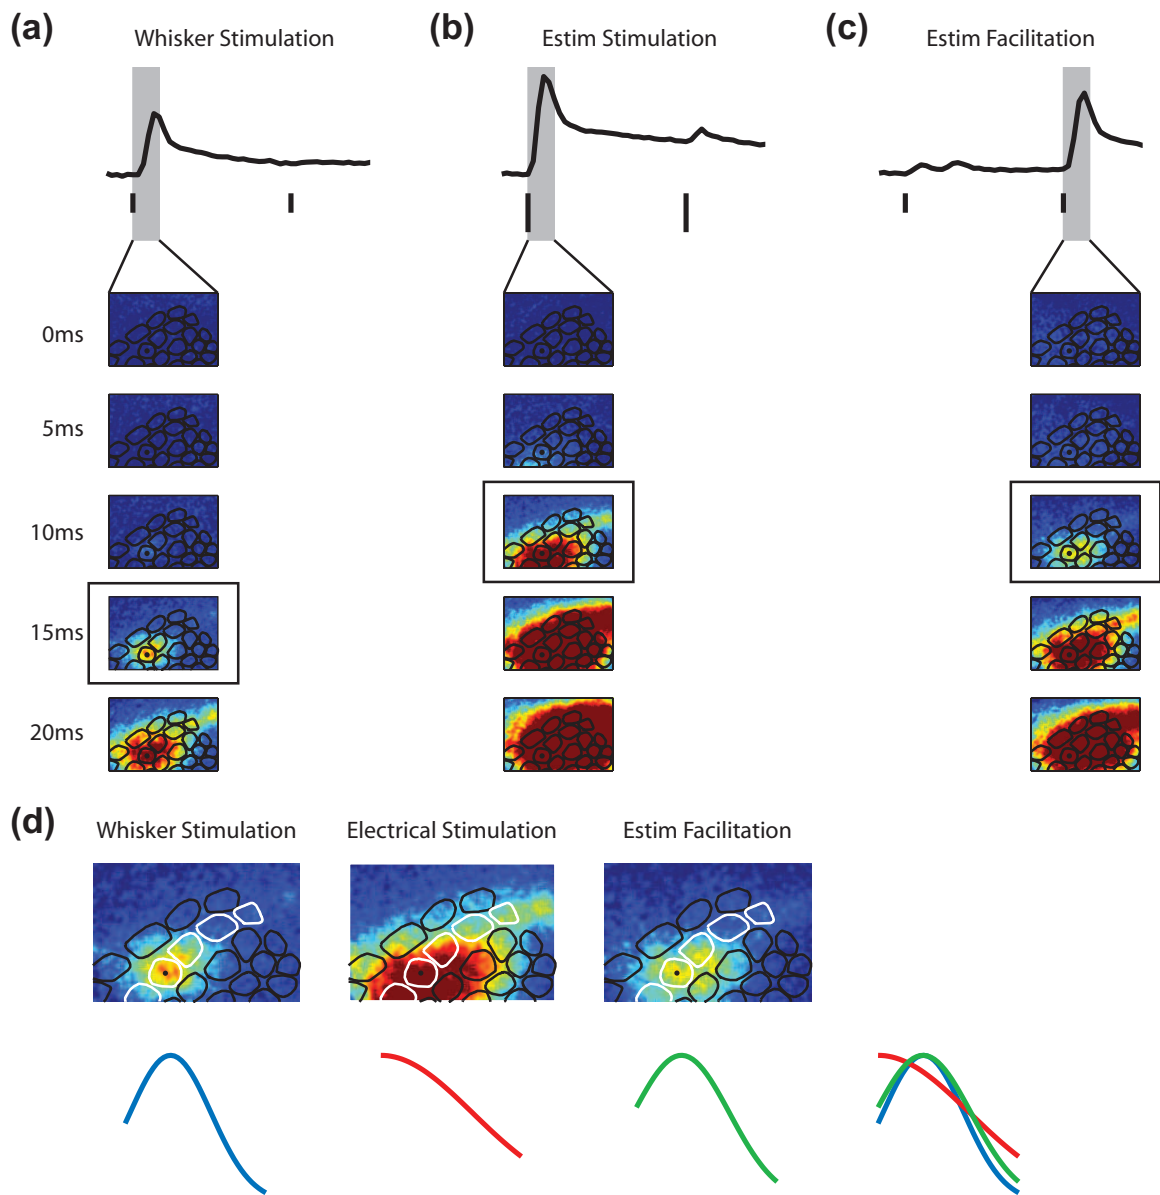


Figure 3.4: Nonlinear facilitation sharpens spatial cortical response. (a) The cortical response to a supra-threshold whisker stimulus was localized in a focal region of cortex. (b) The cortical response to a supra-threshold electrical stimulus spread throughout a large region of cortex, even in the initial frame. (c) The facilitated response to a second, normally sub-threshold, microstimulation pulse was localized in a focal region of cortex. (d) The initial frame of activation illustrated the differences across the stimuli. A Gaussian fit to the activity with the row containing the principal whisker (white cortical columns) is below each image. When overlaid (right), the response to the facilitated electrical stimulation and the whisker stimulation activated a similar region of cortex.

principal cortical column are outlined in white. A Gaussian fit to the average activation within the cortical row provided a one-dimensional representation of the spread across cortex. The Gaussian fits to the supra-threshold whisker (blue) and electrical inputs (red), along with the facilitated electrical stimulus (green) are presented below the respective initial frames of activation, and are overlapped to the right portion of the panel. When overlapped, the whisker and facilitated electrical stimulus are qualitatively much more spatially constrained than the supra-threshold electrical stimulus.

A quantitative analysis of the spatial spread for the three stimulus types is shown in figure 3.5. The example initial frames for the whisker, electrical, and facilitated electrical stimuli are repeated for clarity in panels (a), (c), and (e) of figure 3.5, respectively. In figure 3.5(b), the one-dimensional profile of the VSDI signal is shown for the row containing the principal cortical column, as in figure 3.4(d). The mean and standard deviation of the response within each cortical column, across trials, are given by the data point and error bars, and the gray shadow indicates the principal cortical column matched to the deflected whisker. With increasing velocity, the maximum amplitude of the response increased, but the activation remained focused in the principal cortical column and the immediate neighboring columns. Panels (d) and (f) of figure 3.5 present the data in a similar format for supra-threshold and facilitated electrical stimuli, respectively. In the example in figure 3.5(d), the response amplitude and spread increased as the current amplitude increased. For the highest current, the activation in cortex spread across the entire row of cortical columns, and the peak was located outside the principal cortical column. For the facilitated electrical stimulation in figure 3.5(f), the spread was comparatively small and the peak occurred in the principal cortical column.

Figure 3.5(g) plots the width of the Gaussian fit against the amplitude of the response for each stimulus type. For whisker stimuli and the facilitated electrical

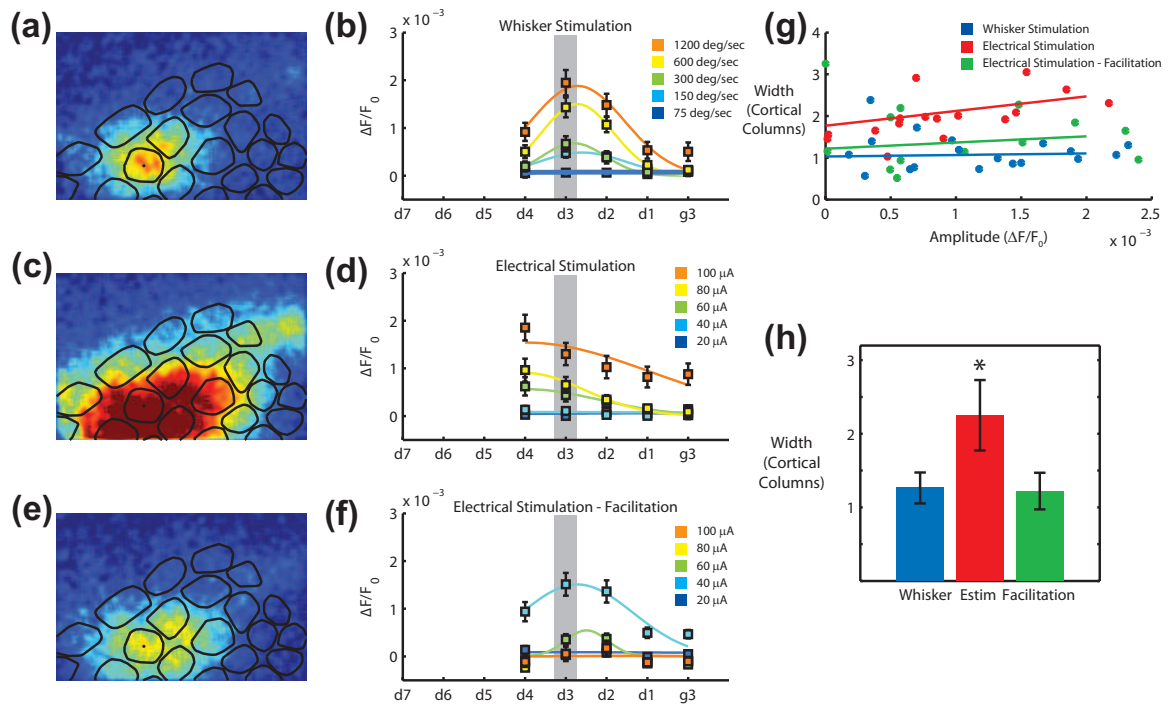


Figure 3.5: Facilitated cortical response to microstimulation activates same area as whisker stimulation. (a) Example of initial frame in response to a whisker stimulus. (b) Activation across the row containing the principal cortical (indicated by gray bar) for whisker deflections of varying velocity. (c) Example of initial frame in response to an electrical stimulus. (d) Activation across the row containing the principal cortical (indicated by gray bar) for electrical stimulation of varying current. (e) Example of initial frame in response to a facilitated electrical stimulus occurring 150ms after an initial sub-threshold electrical stimulus. (b) Activation across the row containing the principal cortical (indicated by gray bar) for an electrical stimulus of varying current delivered 150ms after a previous electrical stimulus of the same current. (g) For increasing response amplitude, the spread of a supra-threshold electrical stimulus increased, while the response to whisker inputs and facilitated electrical stimuli had a constant spread. (h) Electrical stimulation from a supra-threshold current activates a significantly larger region of cortex than whisker inputs or facilitated electrical stimuli.

stimuli, the slope of the best fit line was not significantly different from zero, similar to the spatial results in Chapter 2, whereas the width of the response to electrical stimulation increased as the amplitude of the response increased, again similar to the results of Chapter 2. We averaged the width measurements across all supra-threshold stimulus intensities ($> 0.1\%$ VSDI amplitude) for each stimulus type and found that the width of the response to electrical stimulation was significantly larger than for whisker and facilitated electrical stimuli, which were not significantly different from each other (figure 3.5(h)).

3.3.5 Optogenetic stimulation of the thalamus also exhibits bimodal nonlinear dynamics

In a separate set of experiments ($N=6$), we performed an identical analysis for optogenetic stimulation of the thalamus and found a similar bimodality in the nonlinear dynamics as with thalamic microstimulation. Figure 3.6(a) presents an example of the cortical response to a sub-threshold light intensity (top) and a supra-threshold light intensity (bottom), each at a 150ms inter-stimulus interval. For the sub-threshold light intensity, the majority of trials exhibited strong facilitation, whereas suppression was observed on all trials at the supra-threshold light intensity. In figure 3.6(b), the response to the second stimulus is plotted relative to the response to the first for each trial and each stimulus intensity. Similar to the thalamic microstimulation results described above, three clusters are apparent in this representation of the data. At very low stimulus intensities, there was no response to either of the two stimuli. At intermediate stimulus intensities, the majority of trials exhibited facilitation, forming a cluster in the upper left portion of the axes. At the highest stimulus intensities, the response to the second stimulus was suppressed relative to the response to the first, forming a cluster in the lower right portion of the axes. The same three clusters are present in the compilation of data across animals, in figure 3.6(c).

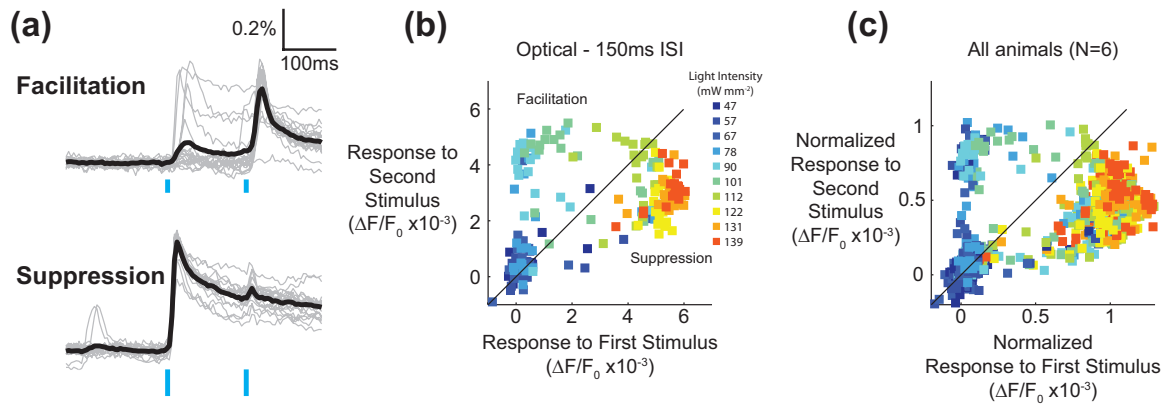


Figure 3.6: Optogenetic stimulation of the thalamus also exhibits bimodal nonlinear dynamics. (a) Example of the cortical response, averaged within a single cortical column, to a sub-threshold (top) and supra-threshold (bottom) optogenetic stimulus. Both single trials (gray) and the mean (black) are presented. (b) The response of the second stimulus is plotted against the response of the first stimulus for single trials of varying light intensity. Points above the unity line (black) were facilitated, while points below the unity line were suppressed. (c) Compilation of data in (b) across all animals (N=6). The responses were normalized, for each animal, relative to the mean response for the strongest light intensity during that experiment. Optogenetic stimulation exhibited a mixture of facilitation and suppression, with the facilitation less reliable on a trial-to-trial basis.

However, as compared to the thalamic microstimulation results in figure 3.2, optogenetic stimuli were less reliable in producing facilitation. This can be seen qualitatively by observing the data points spanning between the lower left cluster at the origin and the suppression cluster. On these trials, the second response was suppressed relative to the first, even though the response to the first stimulus was relatively small. In figure 3.2(d), no data points were observed between the lower left cluster and the suppression cluster across four animals. This slight distinction between the thalamic microstimulation and optogenetic stimulation may provide evidence towards a mechanism for the facilitation (see discussion).

Additionally, the optogenetic stimulation produced a substantially smaller stimulation artifact than thalamic microstimulation, allowing the direct recording of VPM units. In all cells recorded (N=4), facilitation existed at the level of the thalamus, while suppression occurred in only one of the four cells. An example cell is shown

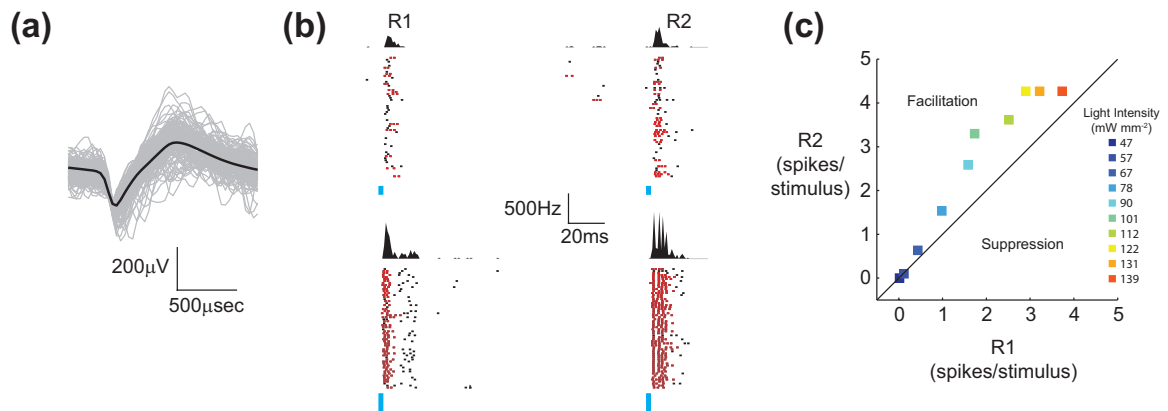


Figure 3.7: Facilitation occurs at the level of VPM for optogenetic stimuli. (a) Action potential waveform for an example VPM cell. Single action potentials (gray) and the mean across all action potentials (black) are shown. (b) Rastergram and peri-stimulus time histogram for the single-unit VPM cell in (a) in response to paired optogenetic stimuli of intermediate (top) and strong (bottom) intensity. Within the rastergram, red spikes belong to an identified burst. (c) The response of the second stimulus is plotted against the response of the first stimulus for various light intensities. The response was measured as the average number of spikes elicited within 20ms of the stimulus onset.

in figure 3.7, with the single-unit waveform in panel (a). The single cell was driven by the same paired optogenetic stimulus used in figure 3.6, consisting of two five millisecond pulses of light at varying light intensity, separated by a 150 millisecond inter-stimulus interval. The rastergram and peri-stimulus time histogram in the top and bottom portions of figure 3.7(b) show the response of the cell to an intermediate (top) and strong (bottom) light intensity. In response to the intermediate light intensity (top), the cell responded reliably across trials, but only with a single spikes or an occasional burst (red spikes in the rastergram belong to a burst). The response to the second stimulus was facilitated, responding more often with two spikes per stimulus, and often in the form of a burst. Similarly, the response was facilitated for the strong light intensity (bottom) as well. The cell responded reliably with a burst for the first stimulus, however, the bursts contained more spikes on average for the second stimulus.

Similar to the VSD analysis above, the response to the second stimulus was plotted

against the response to the first stimulus for varying light intensity. Contrary to the VSD results, no suppression occurred for this example cell at any light intensity, as all data points were above the unity line (black). Only one of the four cells exhibited a transition from facilitation to suppression as the light intensity increased. Again, this provides evidence to the mechanism of the facilitation, suggesting that it occurs at the level of the thalamus, whereas the suppression is likely a cortical phenomenon (see discussion).

3.4 Discussion

We used voltage sensitive dye imaging to study the nonlinear dynamics of the cortical response to sensory and artificial stimuli. Specifically, pairs of stimuli, separated by inter-stimulus intervals of various duration, were delivered either to the whiskers on the face, or to the thalamic via electrical or optogenetic inputs. The cortical response to pairs of whisker stimuli exhibited classical suppressive dynamics, consistent with previous reports [173, 33, 15, 14, 197, 198]. Thalamic microstimulation and optogenetic stimulation, however, produced an additional nonlinear dynamical mode. For strong inputs, artificial stimuli generated suppressive responses, whereas for sub-threshold inputs, the cortical response demonstrated profound facilitation. The facilitation dynamics also contained a spatial component, resulting in a spatial sharpening of the cortical response. Finally, the facilitation was observed for single unit activity within VPM during optogenetic stimulation, suggesting that the mechanism was thalamic in origin.

Based on previous literature, there are three primary candidate mechanisms through which the artificial stimuli likely recruited the additional nonlinear dynamics: 1) simultaneous activation of axons from the postero-medial (POM) nucleus in the thalamus, which has been associated with the thalamocortical augmenting response [33], 2) enhanced activation of the reticular thalamic (RT) nucleus through the extreme

synchrony induced by artificial stimuli, resulting in thalamic bursts [76], or 3) preferential activation of class II, facilitating synapses that extend directly from VPM to layer 2/3 cells [190].

The thalamocortical augmenting response was originally described in 1943 [48], but has more recently been studied in the rodent whisker system by Castro-Alamancos and Barry Connors [33]. The augmenting response is characterized by progressive facilitation of the cortical response to thalamic microstimulation for interstimulus intervals between 50 and 200 milliseconds, and has been shown to occur in the awake animal [33]. The exact mechanism of the augmenting response is disputed, however, with some pointing to cortical mechanism, while others propose the dynamics are thalamic in origin. Principally, though, the augmenting response is believed to be a product of bursting thalamocorecipient cells within layer 5 of cortex. In the rodent vibrissa system, this has primarily been considered through stimulation of the POM nucleus, which projects to layer 5 of cortex as a part of the paralemniscal system, whereas VPM primarily projects to layer 4 (although recent evidence has demonstrated direct projections from VPM to layer 5 [40]). Due to the close proximity of POM and VPM, it is possible that thalamic microstimulation recruited the augmenting response by activating POM axons passing through/near VPM. However, given the observance of facilitation in the VPM units and the facilitation caused by optogenetic stimuli, which are believed to stimulate axons to a lesser extent, the classical augmenting response likely was not the mechanism of the facilitation.

More likely, the facilitation was due to the engagement of bursting mechanisms within VPM. Bursting in the thalamus is caused by the recruitment of T-type calcium channels in the thalamus [170]. When thalamic cells are hyperpolarized for a long period of time ($> 100\text{ms}$), the calcium channels de-inactivate, allowing a large influx of calcium during any subsequent depolarization. Calcium bursts are characterized by multiple spikes occurring in a single cell with inter-spike intervals less than four

milliseconds, immediately following an inter-spike interval greater than 100ms [164]. The time period of de-inactivation of calcium channels is consistent with the inter-stimulus intervals (100-200ms) for which we observed facilitation. Also, we directly observed an increase in bursting for VPM cells in the response to the second stimulus of a facilitating pair. Finally, the spatial sharpening of the cortical response also suggests a bursting mechanism in the thalamus. The large spatial spread in cortex in response to a single supra-threshold electrical stimulus was likely due to the activation of axons, as opposed to cell bodies, near the electrode tip in the VPM (as discussed in Chapter 2) [83]. However, the facilitated response was spatially constrained and more closely matched to the topographic electrode position in VPM. Activation of calcium channels would reduce the threshold to stimulation of cell bodies, but would not affect the threshold of axonal stimulation due to the prevalence of calcium channels in the cell soma and proximal dendrites [214]. In this way, the second stimulus would activate a greater proportion of cell bodies near the electrode tip, as opposed to passing axons, and produce a cortical response with spatial spread more similar to that of whisker stimuli.

Finally, direct synaptic connections from VPM to layer 2/3 have recently been reported, with facilitation properties on a consistent timescale to the experimental results we have observed here. Namely, for 10Hz stimuli delivered to VPM in vitro, the post synaptic potential in the layer 2/3 neurons was facilitated [190]. This was in contrast to the synaptic response in layer 4 neurons, which demonstrated suppression. While this may have played a role in the facilitation, the observation of facilitation in the VPM units again suggests a bursting mechanism.

Regardless of the exact mechanism, the facilitation dynamics were only recruited by artificial stimuli. Any of the above mechanistic explanations derive from neural circuit properties that could have been similarly engaged by whisker stimuli, but were not. In this way, the extra dynamical mode highlights the distinct engagement of

neural circuits by artificial stimuli, much in the way of the different variability trends in Chapter 2. More subtle differences in nonlinear dynamics have been observed in the early sensory pathways, both experimentally [112] and behaviorally [119]. Logothetis et al found that the timescale of suppression was significantly longer for thalamic microstimulation than visual stimulation, a finding born out in behavior as Masse and Cook found that the timescale of suppression was longer for intracortical microstimulation than visual stimuli [112, 119]. The recruitment of suppression was significantly different for downstream neural structures as well, as Logothetis and colleagues used fMRI to track the nonlinear dynamics in response to thalamic microstimulation across multiple stages of the visual pathway [112]. Drawing on the modeling results in Chapter 2, it may be the case that the extreme synchrony induced by artificial stimuli enhances the strength of feedforward inhibition within recurrent neural circuits.

Whether due to the activation of axons or the extreme synchronization, the distinct dynamics for sensory and artificial stimuli reported here have implications for the study of non-sensory neural circuits. In these circuits, “natural stimuli” are less clearly defined (i.e. in the case of memory), thus artificial stimulation techniques are often used to map the dynamics of circuits like the hippocampus [175] or pre-frontal cortex [77]. Using artificial stimulation in this context could lead to misrepresentation of the neural dynamics that are important to normal function or characterize pathological states.

Perhaps more directly, the nonlinear dynamics in response to artificial inputs must be known in order to design encoding schemes to map constantly changing sensory inputs into patterns of stimulation for sensory prostheses. The most successful implementation of a sensory prosthesis, the cochlear implant, does not include any nonlinear dynamics into the stimulation encoding scheme [201], and yet has

demonstrated remarkable success [202]. However, the neural circuits of the peripheral nervous system are decidedly less nonlinear than those in the central nervous system, likely one of the reasons why central nervous system sensory prostheses have failed to date [19, 17]. Therefore, the nonlinear dynamics must be known in order to design stimuli to generate functionally relevant, time-varying neural responses. In this work, we characterized and quantified the spatiotemporal nonlinear dynamics of a canonical neural circuit to sensory and artificial stimuli, but future work must develop quantitative and predictive dynamical models to advance the field of sensory prostheses.

CHAPTER IV

SYSTEM IDENTIFICATION OF THE NONLINEAR DYNAMICS IN THE CORTICAL RESPONSE TO THALAMIC MICROSTIMULATION

This chapter was originally published as an article in the Journal of Neural Engineering and is presented with only minor stylistic changes:

Millard, DC, Wang, Q, Gollnick, CA, & Stanley, GB. System identification of the nonlinear dynamics in the thalamocortical circuit in response to patterned thalamic microstimulation *in vivo*. Journal of Neural Engineering, 10(6), 2013. [128]

Portions of this work were presented in poster form at the following conferences:

Millard, DC, Wang, Q, Gollnick, CA, & Stanley, GB. Optimization of patterned thalamic microstimulation for control of cortical activation. Society for Neuroscience Annual Meeting, Washington, DC, November 2011.

Millard, DC, Wang, & Stanley, GB. Nonlinear dynamical modeling of the cortical response to thalamic microstimulation. Biomedical Engineering Society Annual Meeting, Atlanta, GA, October 2012.

4.1 Introduction

Sensory prostheses seek to use electrical stimulation to deliver information to the brain about the sensory environment when the native neural pathways have been damaged due to trauma or disease. While peripheral sensory prostheses, like the cochlear or retinal implants, have been successful [87, 202], attempts at delivering

information directly to the central nervous system have proven difficult. Whether the aim is to reproduce natural neural activity or merely to deliver discriminable inputs to the brain, the advancement of sensory prostheses requires a greater understanding of the mapping from electrical stimuli to neural response within complex circuits and the resulting propagation along neural pathways.

Recent work has pushed towards recording population responses downstream of the delivery of patterned microstimulation in-vivo [33, 93, 29, 38, 83, 112, 22, 200]. In all but the simplest scenarios, the neural response to electrical stimulation is highly nonlinear, ranging from paired stimulus facilitation in the thalamocortical augmenting response [48, 33] to paired stimulus suppression at the level of the cortex [93, 29]. Furthermore, the nonlinear effects of natural sensory stimuli and electrical stimuli are behaviorally and electrophysiologically different, indicating that electrical stimuli activate neural circuits in a manner distinct from the natural physiological recruitment [112, 119]. In order to design patterns of stimulation to faithfully represent ongoing changes in the sensory environment for prosthesis applications, particularly in the central nervous system, we must develop predictive models of these dynamical nonlinear mappings in-vivo. Here we perform nonlinear system identification within the central nervous system, specifically using the thalamocortical circuit, to model the system dynamics in response to patterns of microstimulation.

System identification has a long history of application in sensory neuroscience for creating nonlinear dynamical models [98, 116, 88, 207]. Typically, system identification takes advantage of nonparametric model structures capable of estimating complicated system dynamics with few assumptions; however, these types of models often require a large amount of data to fit [118]. In-vivo experimental models are typically limited in the amount of data that can be realistically collected during an experimental session, often precluding the high-order modeling that is necessary to adequately capture the complexity of the system. However, through a combination

of non-parametric modeling and empirical observations of the system, it is possible to constrain the model subspace, greatly reducing the number of parameters needed, while only minimally restricting the generality of the model [132, 178].

Through both nonparametric and parametric system identification techniques, we develop a model of the nonlinear system dynamics in the thalamocortical circuit *in vivo*. Specifically, using voltage sensitive dye imaging (VSDI) techniques, we recorded the cortical layer 2/3 response to upstream microstimulation in the ventral posteromedial (VPM) region of the thalamus in the anesthetized rat. The canonical network architecture of the thalamocortical circuit [171], along with the extensive literature on the anatomy of the rodent vibrissa system [206, 51], makes this an ideal model system for studying the network level neural response to electrical stimuli. Systematic probing of the input-output relationship enabled the development of a nonlinear phenomenological model based upon experimental observations that is highly predictive of the cortical response to patterns of thalamic microstimulation. The ultimate structure of the model revealed complex interactions within a multi-stage architecture composed of canonical facilitative and suppressive dynamics, and a sensitivity to noise that mediates trial-by-trial bimodality in the facilitative/suppressive dynamics. Finally, from simulations with the model, we suggest that electrical stimulation activates local circuitry through linear recruitment, but that this activity propagates in a highly nonlinear fashion to downstream targets. More generally, the nonlinear dynamical model developed in this study informs future encoding schemes that map sensory signals to patterns of microstimulation for sensory prosthesis implementation.

4.2 *Methods*

4.2.1 Experimental Preparation

The animal preparation was identical to that described in section 1.4, with VSDI used to monitor the cortical response to sensory, electrical, and optical stimulation

in vivo. The temporal evolution of the VSDI signal, as measured through spatially averaging within the principal cortical column, was used throughout this chapter.

4.2.2 Electrical Stimulation

A glass coated tungsten microelectrode (impedance = 1-2 M at 1kHz) was advanced to the ventral postero-medial (VPm) region of the thalamus using a precision microdrive (Knopf Instruments, Tujunga, CA). The principal vibrissa was determined by manually deflecting individual whiskers and confirmed using the latency and spike count of single unit recordings in response to controlled whisker deflection using a piezo-electric actuator [196]. In the event that single unit recordings could not be achieved, multi-unit activity was used.

Following electrophysiological determination of the electrode position and its associated principal vibrissa, the electrode was used to deliver microstimulation to the surrounding tissue. The stimulus waveforms were designed using a digital stimulus generator (WPI Inc, Sarasota, Florida) and delivered using a current controlled, optically isolated stimulator (WPI Inc, Sarasota, Florida). Individual electrical stimuli were charge-balanced, cathodal-first, biphasic waveforms of 200 microsecond duration per phase. A diagram of the stimulus waveform is displayed in figure 4.1(a). More complicated stimuli were generated through temporal patterns of this base stimulus unit with varying amplitudes. Although we have recently shown a topographic displacement of the spatial cortical response to symmetric and asymmetric waveforms of thalamic microstimulation [195], we observed no difference in the dynamics presented in this study between the two stimulus waveforms (data not shown).

Three different stimulation protocols were used in this study. A series of single electrical stimulation pulses with varying amplitude between 10 and 100 microamperes was used to test the static nonlinearity of the system (shown in Figure 4.1(b), discussed in the results section). The current range was chosen to elicit the full range

of sub-threshold to maximal cortical responses. To sample the nonlinear dynamics of the system, pairs of electrical stimulation pulses were delivered with varying inter-stimulus intervals between 50 and 500 milliseconds. For the system identification procedure, random impulse trains were generated, where the event times were determined by a Poisson distribution and the intensities of the events were drawn randomly from a set of current amplitudes (30, 40, 50, 60, 80, 100 microamperes), all with equal probability. Throughout the remainder of the document, these stimuli will be referred to as random amplitude Poisson (RAP) impulse trains, consistent with previous literature [208]. In four experiments, twenty different random instances of RAP impulse trains were used with a Poisson homogeneous rate of 10Hz, while the remaining three experiments used thirty different random instances with a Poisson homogeneous rate of 3Hz. The results were not qualitatively different using the two different stimulus rates. While previous studies have used high frequency pulse trains as the fundamental unit for stimulating the brain [165, 155, 112, 142], we treated each single stimulus pulse as its own event. Given the inter-event intervals were generated from a Poisson process, the RAP stimulus presented a wide bandwidth (up to 500Hz) to the experimental system.

4.2.3 System Identification

Two different model structures were used in this study: a second order Volterra series and a custom phenomenological model. Each model was fit using the cortical response, $r[n]$, averaged within the principal cortical column, to the RAP impulse trains, $s[n]$, where n are the discrete timepoints sampled by the VSD imaging for a given trial of length T . A block diagram of the experimental system is shown in figure 4.2(a). Each electrical stimulus was considered a discrete Dirac delta function with amplitude in units of microamperes. The second order Volterra series was fit according to the cross-correlation based methods of Wu and Scialabassi [208]. The

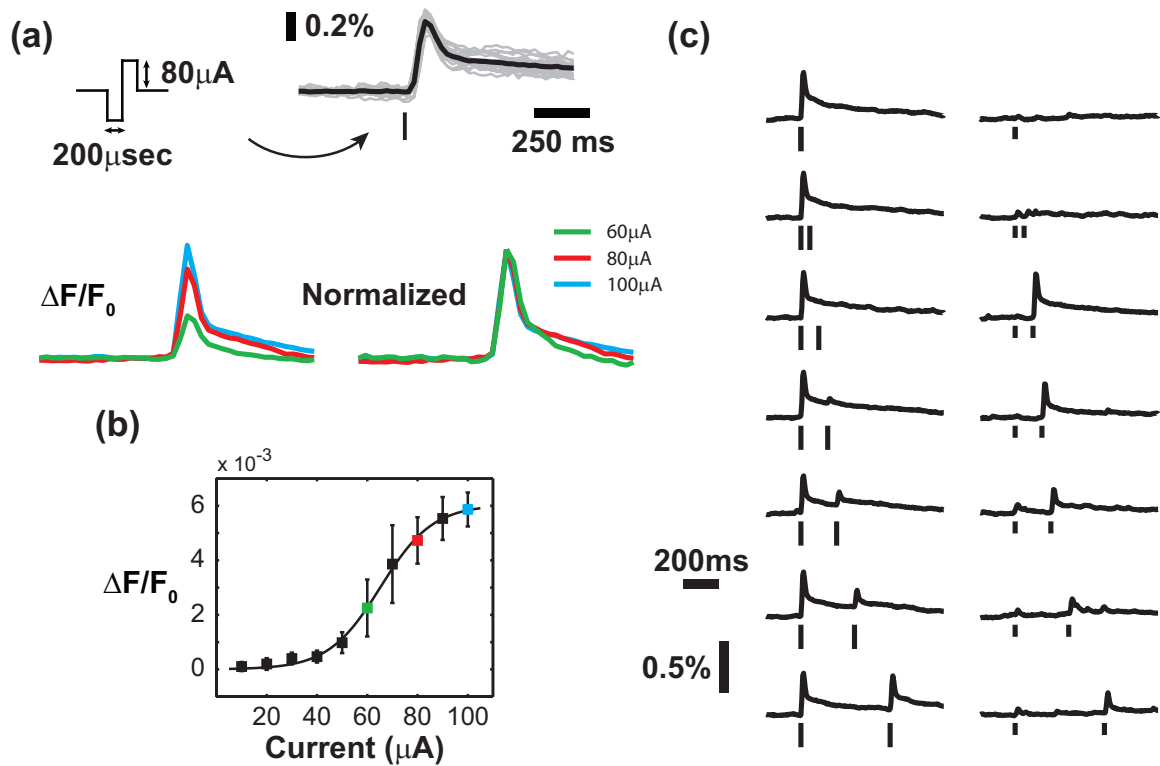


Figure 4.1: Microstimulation of the thalamus produces a nonlinear cortical response. (a) Top: By averaging within a single cortical column, we obtain a timecourse of cortical activation with high signal to noise ratio. The tick represents the presentation time of the stimulus, which is a symmetric biphasic current waveform. The height of the ticks throughout the paper indicates the current amplitude of the stimulus. Microstimulation elicited a characteristic timecourse in the cortical response (single trials in gray, trial-averaged response in black). Bottom: When normalized by the amplitude of the response, the timecourse was consistent across a wide range of stimulus intensities and response amplitudes. (b) The amplitude of the cortical response displayed a nonlinear relationship with the current intensity of the single electrical stimulus. The currents used in (a) are color-coded in (b) for reference. (c) Microstimulation of the thalamus engaged two sets of nonlinear dynamics. A strong electrical stimulus suppressed the response to a second electrical stimulus, with the suppression decreasing for long inter-stimulus intervals. A weak electrical stimulus, however, caused profound facilitation of the response to the second stimulus. This facilitation principally occurs for inter-stimulus intervals of 100-200 milliseconds.

parameters of the phenomenological model were fit simultaneously through a least squares regression algorithm in MATLAB. Confidence intervals on the parameters were estimated by fitting the models with shuffled versions of the stimuli, where each instance of the RAP impulse train was randomly reassigned to a cortical response generated by a different instance of the stimulus.

The models were cross-validated using the leave one out method [94]. The models were fit using all but one instance of the RAP impulse train and then tested on the remaining instance of the RAP impulse train stimulus. This procedure was then repeated for all instances of the stimulus. Performance was measured as the percent of variance accounted for (VAF) by the model:

$$VAF = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2} \quad (10)$$

where y_i is the experimental response and \hat{y}_i is the predicted response to the i^{th} stimulus of the RAP impulse train and \bar{y} is the mean experimental response across all stimuli. The VAF in a given experiment was calculated as the median of the VAF across the set of test stimuli. To compare the Volterra and phenomenological models, the VAF was then averaged across experiments (N=7). The model parameters displayed throughout the paper are the average across all test stimuli for a single experiment and then across all animals (N=7). A VAF of 100% would indicate that the model prediction exactly matches the timeseries of the cortical response. However, due to noise in the biological signal, the variance that the models could realistically be expected to explain was bounded below 100%. For instance, a single trial of the cortical response predicts 85% of the variance in the mean cortical response, indicating a degree of stochasticity that a deterministic model could not be expected to account for. However, each of the models is affected identically, such that the variance accounted for measure can be used to compare the two models.

4.2.3.1 Volterra Model

The goal of this study was to identify the nonlinear system dynamics that govern the cortical response to patterns of electrical stimulation delivered to the thalamus. Traditionally, system identification has been performed using nonparametric black box methods, such as Volterra kernel estimation, which have a long history in neuroscience applications (for review see [118, 207]). A Volterra series model describes the n -th order nonlinear dynamics of a system with input, $s[n]$, and output, $r[n]$, through a series of kernel functions, k_0, k_1, \dots, k_n , according to the following equation:

$$\begin{aligned} r[n] = & k_0 + \sum_{m=0}^{M-1} k_1[m] \cdot s[n - m] \\ & + \sum_{m_1=0}^{M-1} \sum_{m_2=0}^{M-1} k_2[m_1, m_2] \cdot s[n - m_1] \cdot s[n - m_2] + \dots \end{aligned} \quad (11)$$

where M is the memory of system, indicating how far in the past that previous inputs will still affect the output, and m_i represent each time index of the VSDI signal (5ms), such that a memory of 100 timepoints corresponds to 500ms. A diagram of the Volterra series model is shown in figure 4.2(b). Here we used a second order Volterra series that, with a memory of 100 timepoints, contained 5150 parameters. Many of these parameters were not important descriptors of the system dynamics, but this information was not available a priori. In this study the second order Volterra system was estimated according to the methods of Wu and Sciabassi [208]. Briefly, a set of RAP-kernels were orthogonalized with respect to the RAP stimulus, such that they could be identified through cross-correlation techniques. These RAP kernels were then mapped into the true Volterra kernels of the system. All kernels presented here are the true Volterra kernels.

Extension to a third order Volterra system would require a prohibitively long data record due to the extremely large number of additional parameters it would require. There exist methods for parameterizing the Volterra kernels for more computationally

and experimentally efficient estimation of higher order kernels. The most common approach in neuroscience applications uses Laguerre polynomials as basis functions for the kernels, providing support for high frequency content at short time lags and exponentially vanishing support at long time lags [117, 175]. However, these approximations are only useful if the chosen basis set can adequately represent the structure of the true Volterra kernels, which likely is not known a priori. These methods did not perform better than the cross-correlation based approach for our data set (data not shown), leading to the development of the phenomenological model.

4.2.3.2 Phenomenological Model

The phenomenological model used many fewer parameters by including specific elements into the model that were derived from experimental observations. First, the temporal structure of the VSDI response to a single electrical stimulus was not affected by the strength of the stimulus or the amplitude of the response, as shown in figure 4.1(a). In this way, the experimental system in figure 4.2(a), with input $s[n]$ and output $r[n]$, can be modeled as a nonlinear mapping of the discrete input followed by a linear filter to convert the discrete inputs to the continuously varying VSDI output. This is depicted in the top portion of figure 4.2(c), where the input, $s[n]$, undergoes a nonlinear mapping to the intermediate variable, $s'[n]$, before passing through the linear filter to produce the output, $r[n]$. This simplification is common for systems classified as same response shape (SRS) systems, where only the amplitude of the response is affected by the discrete inputs and system dynamics [98, 169, 178].

The model for the discrete nonlinear mapping was also derived from experimental observations. The system exhibited two different sets of nonlinear dynamics, as illustrated by figure 4.1(c). For the high current intensity in the left portion of figure 4.1(c), the response to the second stimulus was suppressed relative to the response to the first, and this suppression relaxed for long inter-stimulus intervals. For the low

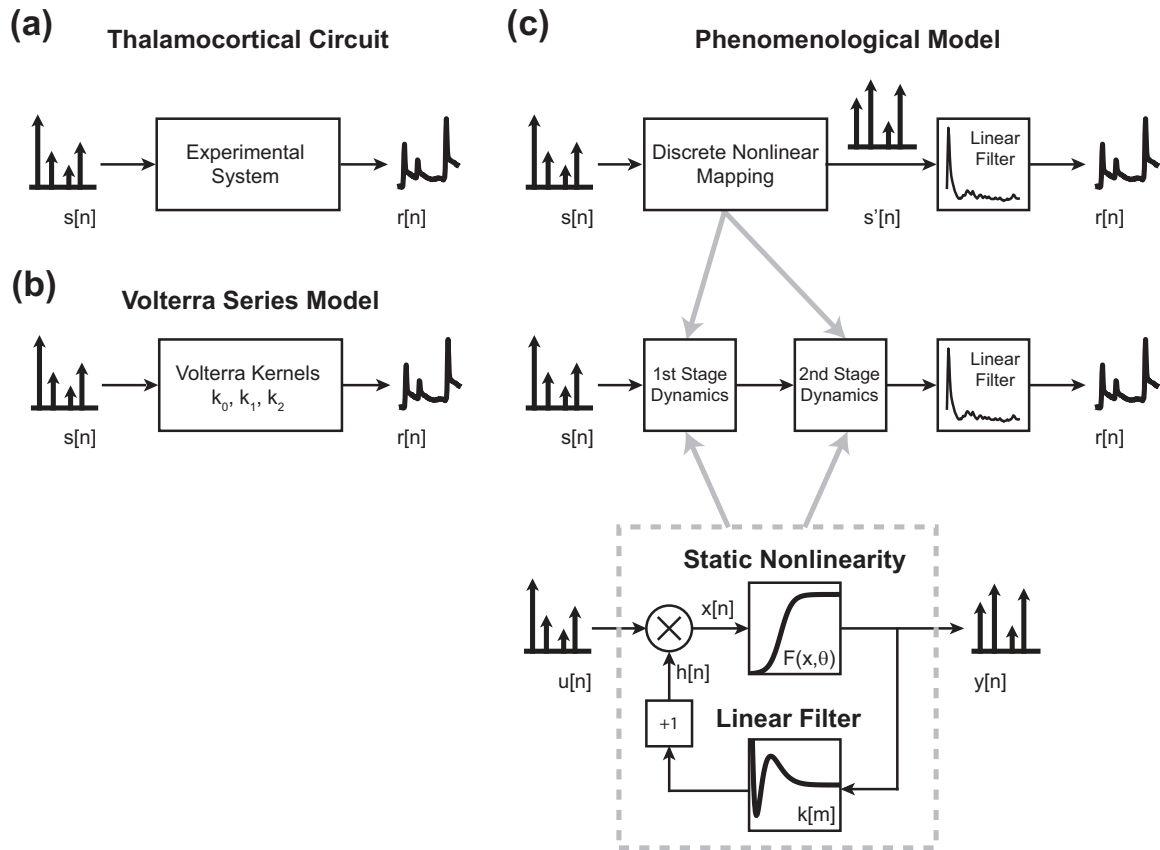


Figure 4.2: Nonlinear modeling architecture. (a) The goal of the study was to create a nonlinear dynamical model of the response of the thalamocortical circuit to patterns of thalamic microstimulation. A train of microstimulation pulses was delivered as the input, while a continuously varying signal, generated by averaging within a single cortical column, was used as the output. (b) A second order Volterra series model was developed in an attempt to capture the dynamics of the system. The kernels of the model mapped trains of discrete inputs to continuously varying signals. (c) Top: The phenomenological model was developed according to experimental observations described in Figure 4.1. The response of the system was similar in shape, regardless of stimulus or response amplitude, allowing separation of the model into a nonlinear mapping of discrete impulses followed by a linear filter. Middle: Further, two distinct sets of dynamics were observed and directly incorporated into the model architecture. Bottom: Each of the two stages within the model was comprised of a canonical unit, in which the static and dynamic nonlinearity were independently modeled and parameterized.

current intensity in right portion of figure 4.1(c), the response to the second stimulus was strongly facilitated relative to the response to the first, but only for interstimulus intervals of 100-200 milliseconds.

These two sets of nonlinear dynamics were explicitly included into the structure of the phenomenological model as a cascade, shown in the middle portion of figure 4.2(c). This was done for two reasons. First, the separate dynamics outlined in figure 4.1(c) were each second order, indicating that the total order of nonlinearity in the system was at least third order. However, a serial cascade simplified the high order nonlinear system into two second order systems, each with a simpler description of the dynamics. Second, the cascade mirrors the multi-stage biology of the neural circuit between the thalamus and layer 2/3 of cortex.

The architecture of the nonlinear mapping within each cascade was identical and is shown diagrammatically in the bottom portion of figure 4.2(c) as a canonical unit. The canonical model separated the static and dynamic nonlinearity within each stage. The static nonlinearity was modeled as a sigmoidal function, F , with input argument, x , and parameters, θ , according to (12) below and in line with the experimental observations in figure 4.1(b):

$$F(x; \theta) = \frac{\theta_1}{1 + \exp \frac{-(x-\theta_2)}{\theta_3}} \quad (12)$$

where θ_1 , θ_2 , and θ_3 are the amplitude, threshold, and sensitivity of the sigmoid, respectively. The output of the canonical unit, $y[n]$, was equal to the output of the sigmoidal function, $F(x; \theta)$. The dynamic nonlinearity was modeled through a history term, $h[n]$, that scaled the input of the canonical unit, $u[n]$, before the static nonlinearity, according to (13) below:

$$y[n] = F(u[n] \cdot h[n]; \theta) \quad (13)$$

In this way, the history term modified how the static nonlinearity acted on the

input, such that when the history term was greater than one, the output was facilitated, and when the history term was less than one, the output was suppressed. The history term in the standard phenomenological model was generated using feedback according to (14):

$$h[n + 1] = 1 + \sum_{m=0}^{M-1} k[m] \cdot y[n - m] \quad (14)$$

where $y[n - m]$ gives the previous outputs within the memory, M , of the canonical unit, and $k[m]$ is a linear filter. The convolution between the previous output and the linear filter was zero when there were no previous outputs within the memory of the system. For this reason, a value of one was added to the convolution such that the history term was equal to one when there had been no previous inputs or outputs for the system and thus had no scaling effect on future inputs. In this way, a positive result from the convolution creates a history term greater than one, leading to facilitation, and a negative result from the convolution produces a history term less than one, leading to suppression.

The static and dynamic nonlinearity were separated to allow for a high order description of the static nonlinearity, while restricting the dynamic nonlinearity to second order. The output of the first canonical unit was the input to the second canonical unit. Finally, the output of the second canonical unit was passed through a linear filter to convert the delta functions into a continuously varying signal. The sigmoidal static nonlinearity functions had 3 parameters each, while the dynamical filters were parameterized through a basis set composed of the first five Laguerre polynomials [117]. The final linear filter that produced the characteristic shape of a VSDI signal had a basis set composed of the first eight Laguerre polynomials. Thus, the phenomenological model contained 24 parameters in total.

The standard phenomenological model described above was compared with a feed-forward implementation of the model. Whereas the phenomenological model with

feedback used the previous responses to implement the dynamic nonlinearity, the feedforward model used the previous stimuli according to (15) below:

$$h[n + 1] = 1 + \sum_{m=0}^{M-1} k[m] \cdot u[n - m] \quad (15)$$

Otherwise, the two implementations of the phenomenological model were identical. The parameters were fit specifically for the feedforward model. A diagram of the canonical unit for the feedforward architecture is shown in figure 4.5(d).

4.2.4 Extension of the phenomenological model to include space

The phenomenological model was extended spatially in an effort to model the entire spatiotemporal cortical response measured using VSDI. In this way, the cortical response was described by the set of $r^{(i)}[n]$, where i indicates the i^{th} cortical column. The objective was to minimize the mean squared error between the neural response, $r^{(i)}[n]$, and the predicted neural response, $\hat{r}_{\theta}^{(i)}[n]$, with model parameters θ and across all i :

$$\arg \min_{\theta} \sum_{i=1}^{32} \sum_{n=1}^N (r^{(i)}[n] - \hat{r}_{\theta}^{(i)}[n])^2 \quad (16)$$

With 24 parameters required to fit a single instance of the phenomenological model, a total of 768 parameters would need to be estimated to fit the phenomenological model for each of the individual cortical column outputs. Instead, we employed a point spread function, $A(x, y; \phi)$, modeled as a two dimensional Gaussian function that mapped the output of the phenomenological model into an image. The Gaussian function was defined by the parameters ϕ , where ϕ_1 and ϕ_2 determine the center of mass of the Gaussian in the two dimensional image, ϕ_3 and ϕ_4 give the width of the Gaussian along the major and minor axis, respectively, and ϕ_5 represents the angular orientation of the major and minor axes with respect to the x and y axes of the image, and given by the following equation:

$$A(x, y; \boldsymbol{\phi}) = \exp(-(a \cdot (x - \phi_1)^2 + b \cdot (x - \phi_1) \cdot (y - \phi_2) + c \cdot (y - \phi_2)^2)) \quad (17)$$

where

$$a = \frac{\cos \phi_5^2}{2 \cdot \phi_3^2} + \frac{\sin \phi_5^2}{2 \cdot \phi_4^2} \quad (18)$$

$$b = \frac{-\sin 2 \cdot \phi_5^2}{4 \cdot \phi_3^2} + \frac{\sin 2 \cdot \phi_5^2}{4 \cdot \phi_4^2} \quad (19)$$

$$c = \frac{\sin \phi_5^2}{2 \cdot \phi_3^2} + \frac{\cos \phi_5^2}{2 \cdot \phi_4^2} \quad (20)$$

The point spread function simplified the model in that the phenomenological model only needed to be fit for a single cortical column and then the point spread function determined the relative activation levels for all cortical columns. This reduced the number of required parameters to 29, with 24 corresponding to the phenomenological model of a single cortical column and 5 parameters for the point spread function. Ultimately, the parameters for the model were determined using the following augmentation of the objective function in (21):

$$\arg \min_{\boldsymbol{\theta}} \sum_{i=1}^{32} \sum_{n=1}^N (r^{(i)}[n] - \hat{r}_{\boldsymbol{\theta}}[n])^2 \quad (21)$$

where x_i and y_i are the coordinates for the center of the i^{th} cortical column and $\hat{r}_{\boldsymbol{\theta}}[n]$ is the predicted response of the principal cortical column. The above optimization was performed for parameters $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ both serially and simultaneously with identical results.

4.2.5 Model Simulations

The models were fit on input-output response data that have been averaged across trials and thus would only be expected to predict trial-averaged responses. However,

to simulate single trial responses from the deterministic phenomenological model, Gaussian white noise was injected at the output of the first and second stages. The variance of the injected noise for the second stage was tuned to reproduce the variance of the background noise in the VSDI signal. The variance of the injected noise in the first stage was scaled according to the relative amplitudes of the static nonlinearity in the first and second stage, such that the noise inputs in the two stages were equally weighted. The simulations were performed with the standard feedback model presented in figures 4.2(c) and 4.3, and also the feedforward implementation of the model.

4.3 Results

In this study, we created a phenomenological nonlinear dynamical model of the cortical response to patterns of microstimulation in the thalamus. We compared this model to a non-parametric model fit with traditional system identification techniques and derived predictions about the circuit level activation caused by electrical microstimulation. Finally, we extended the phenomenological model to capture the spatial properties of the cortical response.

4.3.1 The cortical response to thalamic microstimulation is highly nonlinear

All experiments utilized in-vivo voltage sensitive dye imaging (VSDI) of layer 2/3 in the whisker representation of the primary somatosensory cortex with electrical microstimulation delivered to the topologically matched ventral postero-medial (VPm) portion of the thalamus in the anesthetized rodent, as described earlier. For the purposes of this study, we averaged the VSDI response spatially within the area outlined by the cortical column topologically matched to the position of the electrode in thalamus, as verified by electrophysiological recordings. By averaging within the cortical column, we obtained a high signal to noise ratio for single trials, as shown in figure

4.1(a) where the gray traces are the single trials and the mean across trials is given by the black trace. The response in cortex over time exhibited a consistent trajectory for a wide range of stimulus intensities and response amplitudes. This is illustrated by the example in figure 4.1(a) by normalizing the response to three different supra-threshold current intensities, such that the temporal profile of activation was qualitatively similar across the stimulus range. While the temporal response shape was consistent, suggesting a simple linear filter as a model, the response amplitude demonstrated a nonlinear relationship with stimulus intensity. Figure 4.1(b) shows an example of the nonlinear relationship, which was well approximated by a sigmoidal function. This observation, along with previous literature, suggests a nonlinear projection from the thalamus to cortex for thalamic microstimulation.

Pairs of electrical stimuli were then delivered to the thalamus with varying current intensity and varying inter-stimulus interval to sample the dynamics of the system. Figure 4.1(c) shows the response to pairs of stimuli with high (left) and low (right) current amplitudes. For the high current intensity in the left portion of figure 4.1(c), the response to the second stimulus was suppressed relative to the response to the first, and this suppression relaxed for long inter-stimulus intervals. For the low current intensity in right portion figure 4.1(c), the response to the second stimulus was strongly facilitated relative to the response to the first, but only for interstimulus intervals of 100-200 milliseconds. This observation alone points to dynamics higher than second order. For a second order system, the nonlinear contribution between two pairs of stimuli will have the same shape, with only the amplitude scaled by the stimulus intensity. Here two different sets of second order dynamics were observed: suppression for high current intensities and facilitation for low current intensities with an inter-stimulus interval between 100 and 200 milliseconds.

4.3.2 Development of a phenomenological model based on experimental observations

We devised a phenomenological model using the experimental observations described above. Two stages with identical architecture were used in series, with each stage being a modified version of models used previously to study synaptic physiology [169, 178]. The model is shown diagrammatically in figure 4.2(c) and the details are described in the methods section. Briefly, the model implemented a cascade approach with multiple elements in series. The initial portion of the model was a nonlinear mapping of the discrete stimulus, and the final portion of the model transforms the impulses from the stimulus into a continuous VSD signal. This separation was based on the observation that the VSDI temporal response exhibited a consistent shape for a wide range of current intensities and response amplitudes, as in figure 4.1(a). The nonlinear dynamical mapping contained two stages because of the two distinct sets of dynamics observed in the figure 4.1(c). Each stage consisted of a static nonlinearity, modeled as a sigmoid function as in figure 4.1(b), and a history term that scaled the input to the static nonlinearity (see Methods). The history term was created by feeding back the output from the static nonlinearity through a linear filter and adding one. In this way, the history term only scaled the input when previous stimuli had occurred within the memory of the system. The architecture was the same for the second stage of the model, with the output of the first stage being the input into the second.

Each stage was initialized with the same linear feedback filter and all of the parameters were fit simultaneously. In every case, although not constrained to do so, the first stage filter captured the time course of the facilitative dynamics and the second stage captured the time course of the suppressive dynamics. The average linear filter, first stage dynamical filter, and second stage dynamical filter are shown in figures 4.3(a), 4.3(b), and 4.3(c), respectively. The actual and predicted responses for an

example train of stimuli are shown in figure 4.3(d), with the predicted response of the Volterra model presented in figure 4.3(e) for comparison purposes. The responses labeled with an asterisk illustrate typical examples of the primary improvement of the phenomenological model over the Volterra model. The Volterra model was not able to capture the facilitative dynamics causing significant under-predictions of the response, whereas these dynamics were explicitly built into the phenomenological model and resulted in fewer under-prediction errors. The variance accounted for by the phenomenological model was $58 \pm 12\%$ across animals, whereas the second order Volterra model accounted for $28 \pm 18\%$ of the variance in the cortical response. The improvement of the phenomenological model over the Volterra model was statistically significant ($p = 0.002$, $N=7$, two-sided paired Students t-test).

4.3.3 Error residuals illustrate the improved performance of the phenomenological model

The mechanism for the increased performance of the phenomenological model is illustrated by examining the residuals. In the top portion of figure 4.4, the error residuals from the Volterra model are presented. In figure 4.4(a), the actual response was plotted against the predicted response for each stimulus within the Poisson train. At low values of the actual response, the Volterra model consistently over-predicted the response, whereas at high values of actual response, the model under-predicted the response. The probability distribution of the errors was calculated by subtracting the actual from the predicted, as in the figure 4.4(a) inset and figure 4.4(b). The distribution has a non-zero median and heavy tails, especially towards under-prediction. By examining the heavy tails we can determine if there was a bias towards the model incorrectly predicting certain features of the stimulus. In comparing the prevalence of over-prediction vs. under-prediction as a function of the interstimulus intervals in the stimulus, we found that the large errors made by the Volterra model were predominantly under-predictions for stimuli that occur within 50-200 milliseconds of a

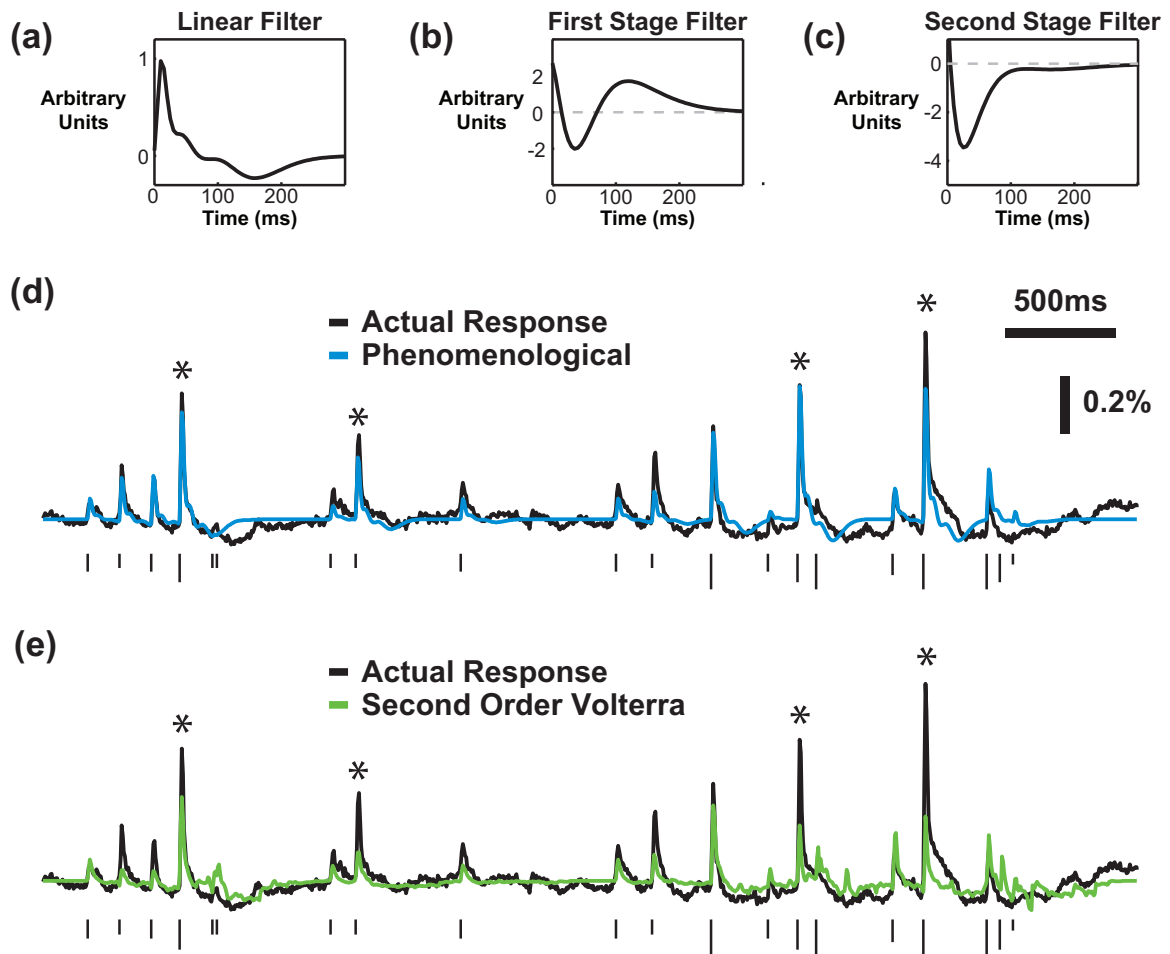


Figure 4.3: Phenomenological model based on experimental observations accurately predicted the cortical response to patterned microstimulation. The average (a) linear kernel, (b) first stage feedback kernel, and (c) second stage feedback kernel of the phenomenological model fit on Poisson train stimuli and VSDI response data ($N=7$). In each case, the first stage kernel implemented the facilitation dynamics and the second stage implemented the suppression dynamics. (d) An example of the performance of the phenomenological model (blue), with the actual response shown in black. (e) An example of the second order Volterra model performance (green), with the actual response shown in black. In (d) and (e), the height of the ticks indicates the current intensity and the asterisks mark stimuli for which the Volterra model severely under-predicted the response. The phenomenological model performed better for these responses.

previous stimulus. These particular errors indicate a failure of the Volterra model to account for the facilitation dynamics presented in the right portion of figure 4.1(c). The skewness of the error distribution was also used to quantify the bias of under-prediction vs. over-prediction. The skewness of the error distribution was -0.77 ± 0.44 , indicating a significantly heavy left tail for under-prediction ($p = 0.003$, $N=7$, two-sided Students t-test).

The error residuals for the phenomenological model are presented in the lower panels of figure 5. In figure 4.4(d), the residuals were closer to the unity line, but still exhibited trends of over-prediction for weak responses and under-prediction for strong responses. However, the collapsed error distribution showed no systematic bias in the residuals towards over-prediction or under-prediction (figure 4.4(e)) and fewer large errors overall due to the increased variance accounted for by the model. Furthermore, there was no systematic bias in the residuals as a function of interstimulus interval. The skewness of the error distribution was -0.02 ± 0.79 , indicating a symmetric distribution, and was not statistically significant from a lack of skewness ($p = 0.94$, $N=7$, two-sided Students t-test). Furthermore, the lack of skewness was statistically different from the skewness for the Volterra model ($p = 0.05$, $N=7$, two-sided paired Students t-test). By including experimental observations explicitly in the phenomenological model, the predictive capabilities were significantly increased while using many fewer parameters. For the remainder of the study, only the phenomenological model was considered.

4.3.4 Trial to trial variability and feedback

While the model described above was highly predictive, it is also purely deterministic, such that it will always predict the same response for the same stimulus. In this way, it cannot strictly reproduce the trial to trial variability in the cortical response. Here we explore the ability of the model to account for single trial response properties




Figure 4.4: Phenomenological model improved error residuals compared to the Volterra model. (a) Error residuals for the second order Volterra model fit via cross-correlation. Each data point is the response to a single stimulus in the RAP stimulus train. (b) The residual distribution had a heavy tail towards under-prediction. (c) The responses were systematically under-predicted for stimuli separated in time by 50-250 milliseconds. (d) Error residuals for the phenomenological model. (e) The residual distribution shows few large errors (indicated by the gray regions) and no bias towards under-prediction or over-prediction. (f) The responses were not systematically under-predicted or over-predicted for any interstimulus intervals.

when noise is added to the system through simulation. While the model was only designed to accurately predict the mean cortical response to a pattern of electrical stimuli, we will show that small changes in the input can lead to large changes in the output experimentally and that the architecture of the model is ideally suited to capture these dynamics on a trial by trial basis.

Two sets of nonlinear dynamics were observed in the trial averaged responses in figure 4.1(c), leading to the explicit inclusion of two stages in the model. Although the dynamics tended to be either suppressive or facilitative, there existed a range of currents for which the occurrence of facilitation or suppression varied on a trial by trial basis. Figure 4.5(a) plots the single trial responses to an initial stimulus against the responses to a second stimulus delivered 150 milliseconds later, with the different colors indicating varying current intensity and the unity line in gray. There are three clear clusters within this plot. At very low current intensities, no response to either stimulus was observed. For intermediate currents, the response to the second stimulus was facilitated relative to the first, creating a cluster in the upper left portion of the axes. At very high current intensities, the response to the second stimulus was suppressed relative to the first, creating a cluster in the lower right portion of the axes. Interestingly, for a select range of current intensities, the occurrence of facilitation and suppression varied on a trial by trial basis. The 60 microampere current intensity in this example spanned the regime connecting the facilitation cluster (upper left) to the suppression cluster (lower right). Plotting the single trial data in this way creates a characteristic pattern, extending vertically from the origin to the facilitation cluster and then traversing across the unity line to the suppression cluster. This characteristic pattern was consistent across animals. Of the ten paired pulse experiments, one data set showed weak facilitation and one showed no facilitation at all, while the remaining eight showed consistent and robust facilitation. For two of the ten paired stimulus experiments there was a systematic

shift from suppressive responses to facilitative responses for the threshold current as the experiment went on; however, this trend was not observed in the other data sets. It should be noted that in these examples, sub-threshold and supra-threshold currents consistently produced facilitative and suppressive responses, respectively, throughout the duration of the experiment. This suggests that the threshold current amplitude is likely a function of the underlying brain state, which was not measured systematically in this study (see discussion).

Examining the trial to trial variability of facilitation and suppression, it is clear that the covariance of the response to the first stimulus and the second stimulus changes dramatically depending on the location within the axes. For very small currents, the individual trials cluster around the origin and have little covariance, creating a circular cloud of points. This is also the case for the facilitative cluster in the top left (30 microampere) and the suppressive cluster in the bottom right (80 microampere). However, for the 40 microampere stimulus, exhibiting roughly half facilitation and half suppression, there is a strong negative correlation between the response to the first stimulus and the response to the second stimulus. Figure 4.5(b) plots the covariance, $\sigma_{1,2}$, between the first and second response as a function of current. For sub-threshold and supra-threshold stimulus intensities, $\sigma_{1,2}$ was small. For the threshold current, the responses to the first and second stimulus strongly co-varied, such that knowledge of the response to the first stimulus was a strong predictor of the magnitude of the response to the second stimulus. This phenomenon was consistent across animals, with the average data presented in figure 4.5(c).

The high variability, σ_1^2 , of the response to a threshold stimulus and the strong covariance, $\sigma_{1,2}$, of the response to a subsequent stimulus indicate the presence of a feedback element within the neural circuit. The phenomenological model presented in this study contains feedback elements to implement the dynamic nonlinearity of the response. Given this, we sought to determine if the phenomenological model,

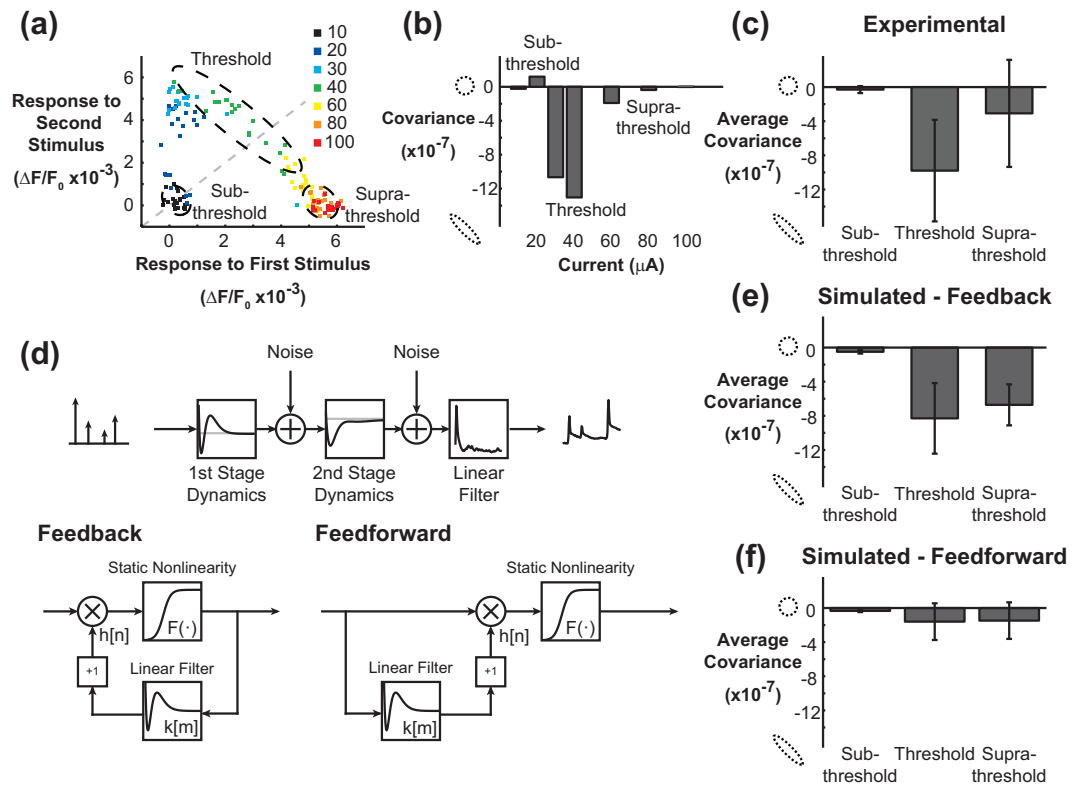


Figure 4.5: Phenomenological model with feedback reproduced single trial variability in facilitation and suppression. (a) The response to the second stimulus (ISI = 150ms) is plotted against the response to the first stimulus for experimental data. The color of each point indicates the current intensity in units of microamperes. The unity line is shown in gray. (b) The trial by trial covariance in the response to the two stimuli was calculated for each stimulus intensity. At sub-threshold and supra-threshold currents, the covariance between the first and second response was low. At a threshold current, there was a strong negative covariance between the first and second response. (c) The negative covariance at the threshold current was consistent across animals. (d) The phenomenological model was used to simulate an identical experiment. The simulated data was created by injecting noise into the model at the output of the first and second stages during presentation of stimuli. A feedback model (same as in Figure 5, displayed in bottom-left of the panel) and a feedforward model (fit specifically for this analysis, displayed in the bottom-right of the panel) were used. Both models utilized a two-stage model architecture, where each stage consisted of a canonical unit. The canonical unit for the feedback (left) and feedforward (right) model are shown in this panel. (e) The feedback model reproduced the strong negative covariance at the threshold current and recovers for supra-threshold currents, and was not significantly different from the experimental data in (c) ($p = 0.89$, $N=7$, two-sided paired Students t-test). (f) The feedforward model did not reproduce the negative covariance for threshold currents, and the difference from the experimental data was statistically significant ($p = 0.003$, $N=7$, two-sided paired Students t-test).

which was built entirely on averaged data, could reproduce single trial response properties. By injecting noise at the output of each stage in the model, the relationship in Figure 4.5(a) was simulated and the covariance as a function of current was extracted. The average sub-threshold, threshold, and supra-threshold $\sigma_{1,2}$ are presented in figure 4.5(e), demonstrating that the phenomenological model can reproduce the strong negative covariance observed experimentally for the responses to two threshold stimuli. However, when the phenomenological model was re-fit using a feedforward architecture (illustrated in the bottom-right portion of figure 4.5(d) and described in the methods), the negative covariance at the threshold current intensity was not observed (figure 4.5(f)). The feedforward model was implemented in a two-stage architecture in the same way as the feedback model described in the previous sections. Only the structure of the canonical units was different, with the feedforward model using the previous inputs to the system to model the dynamics, as opposed to the feedback model that uses previous outputs to model the dynamics.

These simulations indicate that the feedback architecture within the model is important for producing the experimentally observed trial-by-trial variability. However, the result could also be due to the specific parameters fit in the feedback and feedforward cases. As a control, the dynamical filters and static nonlinearities from the feedforward model were implemented in the feedback architecture and the strong covariance, $\sigma_{1,2}$, was still not reproduced (data not shown). Vice versa, when the parameters from the feedback model were implemented in the feedforward architecture, half of the covariance was recovered as compared to the feedback model. This indicates that feedback is necessary, but not sufficient, for creating the covariance observed experimentally, pointing towards a role for either the static nonlinearities or the dynamical filters.

4.3.5 Activity propagation through electrical microstimulation

The dynamical filters did not vary significantly across the feedforward and feedback models in the first or second stage, as shown in figure 4.6(a) and 4.6(b), respectively, nor did the overall static nonlinearity differ in figure 4.6(c). But, while the overall static nonlinearities were similar, the individual static nonlinearities within each of the two stages were dramatically different. For the first stage, shown in figure 4.6(d), the feedback model was very insensitive to the current of the input, modeled as a sigmoid function elongated along the horizontal, whereas the feedforward model was moderately sensitive. In the second stage, depicted in figure 4.6(e), the feedback model was highly sensitive to current intensity, modeled as a sigmoid function compressed along the horizontal, whereas the feedforward model was again moderately sensitive. These model parameters were consistent across animals, as shown in figure 4.6(f).

Due to the high sensitivity of the second stage in the feedback model, small perturbations from the noise caused large changes in the output of the model near the threshold current leading to a high response variability, σ_1^2 . Meanwhile, the highly variable response was fed back into the model and augmented the response to subsequent stimuli resulting in a high $\sigma_{1,2}$ value. The moderately sensitive stages of the feedforward model were more robust to noise and thus did not exhibit high σ_1^2 at the threshold current, and without feedback the model could not account for $\sigma_{1,2}$. In summary, the relative sensitivity in the static nonlinearity of each stage dramatically alters the propagation of activation through a cascade system with feedback. Meanwhile, the similarity of the phenomenological model architecture and the anatomy and physiology of the thalamocortical circuit maps the results of this section to general predictions about information propagation in neural circuits (see discussion).

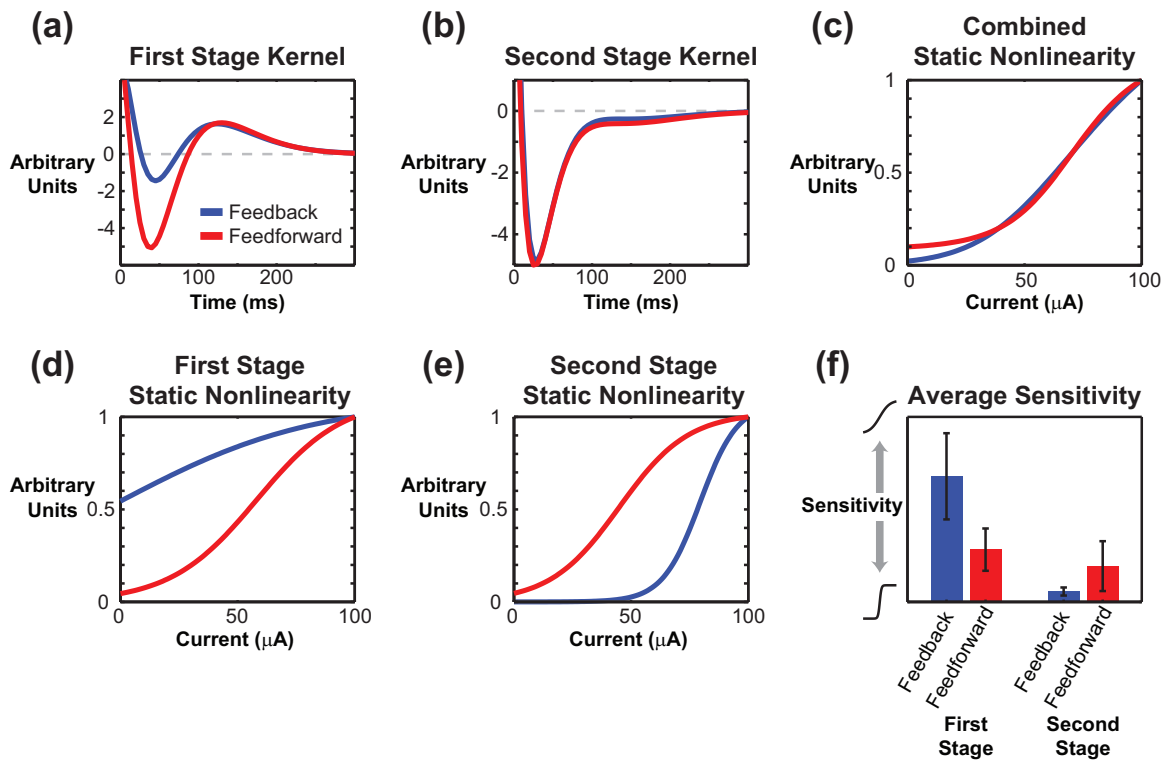


Figure 4.6: Model parameters predict linear local response properties, but nonlinear propagation of activity. (a) The average first stage kernel from the phenomenological model fit with feedback (red) and feedforward (blue) dynamics. (b) The average second stage kernel. (c) An example of the full static nonlinearity created by combining both stages. (d) and (e) The static nonlinearity at the first and second stages, respectively, for the feedback and feedforward models. (f) The average sensitivity across animals ($N=7$) is distinctly different for the feedback and feedforward models. The feedback model was weakly sensitive in the first stage and highly sensitive in the second stage. The feedforward model was moderately sensitive in both stages.

4.3.6 Linear point spread function captures spatial spread in cortex

The phenomenological model was extended to capture the spatial properties of the cortical response to patterns of thalamic microstimulation. In the previous sections, the VSDI signal was discretized according to the anatomical map of the cortical columns, and only the signal from the principal cortical column was used to fit the phenomenological model. The phenomenological model performed equally well for the principal cortical column and adjacent cortical columns. However, fitting an instance of the model for each cortical column would require a large number of parameters, while ignoring the correlation structure of the spatial cortical response. As a first order approximation, we appended a point spread function, $A(x, y; \phi)$, to the output of the phenomenological model, as in figure 4.7(a), mapping a scaled version of the timeseries response from the principal cortical column to the entire set of cortical columns within the barrel cortex. The details of the point spread function (PSF) are described in the methods. Briefly, the PSF was modeled as a two dimensional Gaussian function, with parameters ϕ describing the center of mass, spread along the major and minor axis, and orientation with respect to the coordinate axis of the VSDI images. The resulting PSF was identical whether fit simultaneously with, or immediately following, the identification of the phenomenological model.

Figure 4.7(b) presents an example of the actual and predicted cortical response to a train of electrical stimuli for the principal cortical column (top, blue) and a distant cortical column (bottom, red). The PSF function accurately maps the output of the phenomenological model to each of the cortical columns. The total variance in the response accounted for by the spatial model, across all cortical columns, was 45+/-8% (N=7). The PSF linearly maps the output of the phenomenological model to the various cortical columns. To assess the validity of this assumption, we analyzed the relative response of the primary and adjacent cortical column for neighboring cortical columns (AW, black data points) and a distant cortical column (DW, gray data

points) in figure 4.7(c). The relative locations of the cortical columns are shown in the left portion of figure 4.7(c), with the actual cortical response at the top and the PSF at the bottom. The white contour indicates the primary cortical column, the black contours demarcate the eight nearest neighbor cortical columns, and the gray contour represents the distant cortical column. Each point on the scatterplot in the right portion of figure 4.7(c) came from the peak cortical response to each stimulus in the RAP impulse train used to fit the model, while the red lines represent the relative response amplitudes defined by the PSF. For the entire range of cortical responses, and for adjacent cortical columns near and far, the linear approximation of the relative response magnitudes in the primary and adjacent cortical columns effectively captured the experimentally observed relationship for trial averaged responses.

4.4 Discussion

Here we have demonstrated a phenomenological model capable of accurately predicting the response of the thalamocortical circuit to temporal patterns of thalamic stimulation in-vivo. By explicitly including nonlinear elements modeled after experimental observations, such as the combination of facilitative and suppressive dynamics, into the model architecture, we significantly increased performance with respect to the Volterra series model architecture while using many fewer parameters. Additionally, although the model was fit using trial-averaged data, it was able to reproduce single trial response properties observed experimentally, lending credence to the physiological significance of the model architecture. Finally, from these single trial simulations, we predict that electrical microstimulation activates neurons in its local environment through linear recruitment, but that this activity propagates to downstream structures in a highly nonlinear manner.

The models in this study were fit using input-output data from the thalamocortical

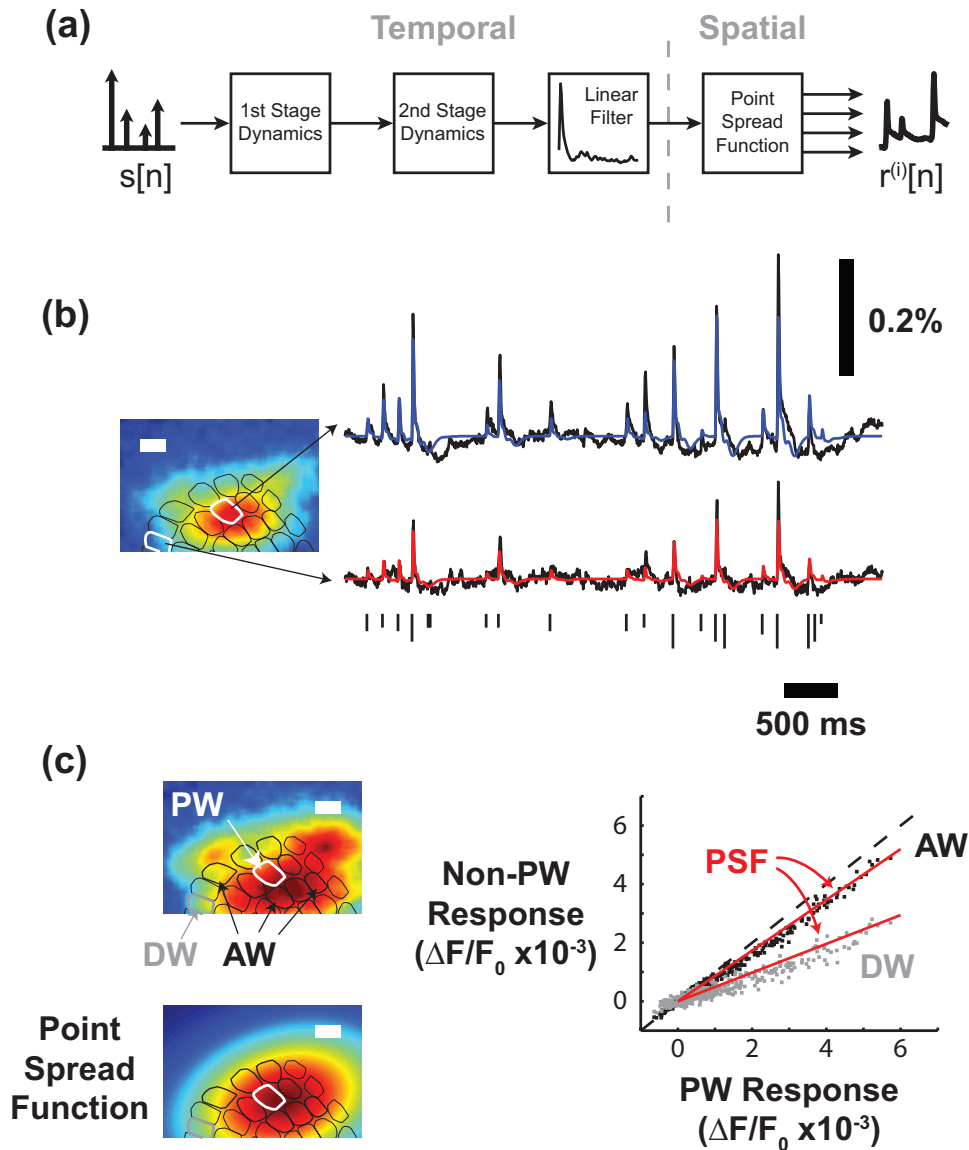


Figure 4.7: Extension of phenomenological model captures spatial spread. (a) A linear point spread function, modeled as a two dimensional Gaussian, was used to extend the model spatially across all cortical columns. (b) The spatial extension to the model captured the dynamics for the principal cortical column as well as distant cortical columns. (c) A representative cortical response (top left) and the point spread function (bottom left) were similar. In the scatter plot, the point spread function (red) captures the relative response properties of the principal cortical column and adjacent cortical column for neighboring (black) and distant (gray) cortical columns. Each data point indicates the response to a single impulse within the RAP impulse train. Scale bars in (b) and (c) are 500 micrometers.

circuit. The input was a train of symmetric biphasic electrical stimuli, Poisson-distributed in time and uniformly varying in amplitude, while the output was the spatially averaged cortical response as measured by voltage sensitive dye imaging. Due to the light scattering of the tissue, voltage sensitive dye imaging principally measures the activity in the superficial layers of cortex [73]. The change in fluorescence of the voltage sensitive dye increases linearly as a function of membrane potential, but the imaging of the signal is too slow to resolve individual action potentials, restricting the interpretation of the signal to sub-threshold activation [72]. Even though the absolute amplitude of the VSDI response in layer 2/3, and the related probability of action potential generation, is known to be strongly modulated by brain state [151], the spatial distribution of sub-threshold activation in layer 2/3 during the onset of activation is likely highly correlated with the supra-threshold activity in layer 4 of cortex [152]. All details considered, VSDI provides a high spatial and temporal measure of the cortical response.

Both models were fit using the same input-output data, but the philosophy and architecture of each was very different. The Volterra model is an example of a black box model, which requires little previous information about the system due to its flexibility, making it an ideal starting point for system modeling. The incredible flexibility of a black box model, however, requires a large number of parameters, and it is difficult to determine which will be important a priori. Also, in order to fit a large numbers of parameters, a large amount of data is needed. For these reasons, we explicitly included certain elements into the phenomenological model, increasing the order of nonlinearity in the model to a degree greater than could be estimated with a Volterra series given the data and time limitations. By effectively fitting the model with small amounts of data, we open the possibility of pseudo real time model construction and experimentation and potentially avoid the slow timescale non-stationarities that may exist in an in-vivo anesthetized biological preparation.

The structure of the phenomenological model invites physiological interpretations due to its unique structure. The observation of two distinct second order dynamics motivated the two-stage architecture. Two sets of second order dynamics in series would imply that a fourth order Volterra series may have captured the dynamics of the system. However, the intentional separation of the dynamics into two stages allowed the dynamics to be described in a simpler form. From this, we would hypothesize that the distinct dynamics are carried out by different elements of the neural circuit, and further that these elements may be acting in series. The model parameters consistently show the facilitation dynamics occurring in the first stage, while the second stage employs the suppression dynamics. As the two-stage model mirrors the disynaptic pathway from the thalamus to layer 2/3 of cortex, the model would predict that the facilitation dynamics occur upstream of the suppression dynamics within the neural circuitry. The suppression dynamics can likely be accounted for by recurrent inhibition within the cortex [93], and the long timescale of inhibition implicates the involvement of GABA-B receptors [27]. With facilitation occurring before suppression in the model, we speculate that the facilitation is sub-cortical in origin. This is in agreement with literature describing the facilitation of successive stimuli termed the thalamocortical augmenting response, which acts through thalamic and cortical structures [48, 33, 8]. The timescale of facilitation is also consistent with the timescale of calcium T-channel mediated bursts in the thalamus [114].

A separate static nonlinearity was included within each stage of the model as well and might be interpreted as the transformation across synapses within the thalamo-cortical circuit. In this way, the static nonlinearity of the first stage would represent the transformation of electrical current into the number of activated thalamic cells. Similarly, the static nonlinearity of the second stage would represent the transformation of the number of activated thalamic cells to the number of activated cortical

cells. While feedback within the model was found to contribute to the single trial response properties, it was the relative sensitivities of the first and second stage static nonlinearities that were most important. Specifically, the first stage was found to be linear with respect to current intensity and exhibit a shallow slope, such that small amounts of noise produced small fluctuations in the output of the first stage. This leads to the prediction that the number of neurons activated in the immediate environment around the electrode are recruited in a nearly linear relationship with the stimulus intensity, and that this recruitment is robust to random fluctuations in membrane potential of the neurons. This interpretation is supported biophysically in that the radius of activation from the electrode tip increases as the square root of the current intensity, while the number of neurons within that sphere increases with the cube of the radius, leading to a weakly nonlinear relationship between the number of neurons activated and the current intensity [183]. Further, when stimulating a fiber bundle, the relationship between current intensity and the number of axons activated becomes linear [210]. Recent work is also in agreement, as Histed et al, report a linear increase in the number of cells activated by increasing microstimulation intensity as measured by calcium imaging of the cortical population [83].

Meanwhile, the output of the second stage, or number of cortical cells activated, was found to be quite sensitive to the output of the first stage, which is interpreted as the number of activated thalamic cells. We hypothesize that the cortical activation is highly sensitive to the number of activated thalamic cells due to the high convergence and divergence of the thalamocortical circuit [171] and the extreme synchrony with which electrical stimulation recruits the thalamic neurons [191, 168]. This is supported by previous work demonstrating that synchronous activation of thalamic neurons drives downstream neural activation in a highly nonlinear manner [2]. The nonlinearity discussed above refers to the static nonlinear relationship between the

neural response and the stimulus intensity; however, we believe this phenomenon extends to the dynamic nonlinearity. The extreme synchrony of the activation caused by electrical stimulation may explain how electrical stimulation and natural sensory stimulation could recruit distinct nonlinear responses [119].

This hypothesis applies directly to the implementation of sensory prostheses, the goal of which is to transduce signals from the sensory environment into patterns of stimulation that create surrogate sensory signals in the peripheral or central nervous system. Ideally a sensory neuroprosthesis would generate neural responses similar to those created by natural stimuli; however, this may prove problematic if electrical microstimulation engages circuits in a fundamentally different manner as compared to natural stimuli. Given this, some degree of plasticity, ranging from interpreting unnatural patterns of electrical stimulation [56] to grasping new coordinate transformations with cross-modal sensory substitution [6, 7], will be required for the successful implementation of a sensory prosthesis. At a minimum, a sensory prosthesis must be capable of producing perceptually distinct neural activations that the patient could learn to interpret functionally. This motivates the mapping of electrical stimuli to downstream neural responses, such that patterns of stimuli can be designed, in real time, with the high spatial resolution and fast timescales that will be needed in the limit of faithfully representing the sensory experience [64]. While implementation of this system identification approach may be possible in a human patient, the immediate impact lies in quantitatively describing the neural response generated by patterns of electrical stimulation within sensory pathways of the central nervous system. The development of the cochlear implant followed a similar trajectory, with many early animal studies characterizing the encoding of electrical stimuli into neural activity [125, 78] to aid in the development of encoding algorithms for early cochlear implants. This work forms an initial step in replicating the cochlear implant development trajectory in the central nervous system.

In addition to the system identification approaches described here, future inclusion of both recording and stimulating electrodes in prosthetic applications will enable control-theoretic approaches to optimize and control the neural activation, and resulting percept, induced by surrogate sensory signals [110, 44]. In order to operate in real-time, the model must contain a description of the stochasticity of the system. While the model presented in this study reproduces single trial response properties, it lacks the ability to predict single trial responses. A feedback signal, such as the electroencephalogram or local field potential, would be needed to determine the instantaneous state of the circuit for improving single-trial predictions [22], in so far as the stochasticity in the single trial responses derive from an underlying state variable [152]. Additionally, the underlying uncertainty could be built directly into the model parameters in the context of a robust control framework. Importantly, the pathological circuit may function differently from normal [46], further emphasizing the importance of closed loop stimulation and recording paradigms.

Precise delivery of information to the brain requires control of the spatial and temporal properties of neural activation. In this study, we modeled the spatial cortical response through a linear point spread function. This simple, linear spatial model performed nearly as well for the entire barrel cortex as the non-spatial phenomenological model did for a single cortical column, indicating the point spread function is a good approximation of the spatial cortical response. This conclusion is consistent with previous studies that coined the term cortical point spread function, referring to the region of cortical space activated by a point source stimulus [73, 45]). The point spread function was effective and efficient in our modeling study, yet a significant portion of the variance in the spatial signal remained unexplained. Future work is needed to determine if the spatial and temporal response properties in cortex are separable or co-dependent in order to control the cortical response with high spatial

and temporal resolution. There exists some evidence from the literature the patterned stimuli can dynamically shape the spatial properties of the cortical response. Brumberg et al demonstrate a spatial sharpening of the cortical response in the barrel cortex to a single whisker deflection when a noise stimulus is applied to the adjacent whiskers [23]. More generally, adapting stimuli are thought to dynamically shape the cortical response, and have been shown to increase spatial acuity in a two-point discrimination task in humans [182]. While space-time separability allows for simpler modeling of the cortical response, co-dependence, if fully understood, could enable complex shaping of the spatiotemporal cortical response and consequent perception.

In this study, through the use of system identification techniques, we have developed a highly predictive phenomenological model of the cortical response to patterns of thalamic microstimulation. Simulations suggest that electrical stimulation may recruit neighboring neurons in a linear manner, but that the resulting activity projects to downstream structures through a highly nonlinear relationship. Future work will extend to a spatiotemporal model of the cortical response and the application of the model as a stimulus design tool for controlling the cortical response. More generally, this framework describes the nonlinear mapping from electrical stimulation to neural response in order to transform environmental cues into surrogate sensory signals at high spatial resolution and fast timescales for the advancement of central nervous system sensory prostheses.

CHAPTER V

DESIGN OF SURROGATE INPUTS TO IMPROVE DISCRIMINABILITY OF THE DOWNSTREAM CORTICAL RESPONSE

Portions of this chapter were originally published as an article in the Journal of Neural Engineering:

Wang, Q, Millard, DC, Zheng, HJV, & Stanley, GB. Voltage-sensitive dye imaging reveals improved topographic activation of cortex in response to manipulation of microstimulation parameters. Journal of Neural Engineering, 9(5), 2012. [195]

Portions of this chapter were presented in poster form at the following conferences:

Wang, Q, Millard, DC, Zheng, HJV, & Stanley, GB. Enhanced cortical specificity of thalamic microstimulation through manipulation of charge delivery. Society for Neuroscience Annual Meeting, Washington, DC, November 2011.

Millard, DC, Gollnick, CA, Hendry, WJ, Rozell, CJ, & Stanley, GB. The role of magnitude and synchrony of population activity in nonlinear circuit processing in the thalamocortical circuit of the rodent vibrissa system. Society for Neuroscience Annual Meeting, San Diego, CA, November 2013.

5.1 Introduction

Electrical stimulation of neural tissue has been utilized for more than a century to either probe neural circuitry [58, 166, 181, 68, 131, 186, 27, 209, 60] or create desired percepts [165, 87, 155, 28, 85, 147, 189] (for a review, see [183, 39]). Following the

success of cochlear implant technology, numerous efforts have been carried out to electrically stimulate various sensory cortical regions to provide surrogate sensory signals for applications in visual and somatosensory prostheses [139, 18, 184, 154], to restore sensory function and aid in sensori-motor learning for closed loop brain computer interfaces and motor prosthetics [143, 90]. Although recent advancement in technology has enabled cell-type selective activation of neurons through optical stimulation [16], the barriers to clinical implementation for these techniques make electrical microstimulation the most likely route in the near term for clinical means of controlling neural circuitry on relatively fast time scales for sensory prosthetics in humans.

Despite intense investigation of the anatomy and physiology of sensory pathways in the nervous system, beyond the sensory periphery little is known about how to generally and systematically control activity in these pathways. Although recent work has successfully demonstrated that it is possible for electrical microstimulation to produce simple tactile and visual percepts that contain the same information as those generated by natural sensory stimuli [155, 165], it is currently unclear how to control neural circuitry on a temporal and spatial scale that relates to more complex, naturalistic sensory inputs. In the context of sensory prostheses, given the prevalence and preservation of well-organized topographic maps in all sensory modalities, it is likely that the similarity between neuronal activation patterns evoked by electrical microstimulation and those evoked by natural sensory stimuli plays a key role in minimizing training time and mental load in disabled individuals who use the sensory prostheses.

The previous chapters have characterized the distinct neural response properties of the thalamocortical circuit to sensory and artificial stimuli, setting the stage for the design of artificial stimuli to overcome these differences. Specifically, chapter 2 demonstrated that thalamic microstimulation activated a significantly larger region

of cortex than topographically matched sensory stimuli. These results are hypothesized to derive from the activation of axons residing close to the electrode tip, but originating from cell bodies far from the electrode, instead of direct stimulation of cell bodies in the immediate vicinity of the electrode. Single cell electrical stimulation studies lend support to this hypothesis [140, 141].

Definitive support has been supplied by a recent study in which the composite neural activity of an in-vivo network was recorded in response to direct electrical stimulation of the circuit, indicating that axonal elements near the electrode are activated to produce a sparse and distributed activity within the network [83]. In this study and others, neuronal excitability as a function of electrical microstimulation parameters such as pulse duration, current amplitude, pulse frequency, and polarity of the leading phase has been experimentally investigated [183]. However, recent modeling work has suggested that microstimulation waveform shape may play a significant role in the preferential activation of cell bodies versus axons [123], and this hypothesis has not been experimentally validated [70].

Additionally, the results of chapters 2 and 4 highlighted the distinct trends for trial-to-trial variability in the amplitude of the cortical response to varying intensity of sensory and artificial stimuli. For sensory stimuli, the variability increased linearly with the response amplitude, while for thalamic microstimulation and optogenetic stimulation, the variability peaked at the threshold response amplitude. This has profound implications for the ability to use surrogate sensory inputs to deliver information to downstream neural structures. This is illustrated by the simple example of encoding information in the strength of the cortical response to a single stimulus, for which the mean amplitude and trial-to-trial variability determine the discriminability between any two response distributions. Typically, the performance is highest at the steepest slope along the stimulus-response curve, which is at the threshold for a sigmoidal relationship. However, the threshold is the exact location of the highest

variability for thalamic microstimulation, limiting overall performance.

Here in this chapter, we use stimulus design techniques to overcome the differences in spatial activation and trial-to-trial variability in thalamic microstimulation, as compared to natural sensory inputs. Specifically, we used voltage sensitive dye (VSD) imaging to quantify in-vivo, network level spatiotemporal cortical activation in response to electrical microstimulation of the thalamic region that provides direct input to cortex in the rat vibrissa pathway. First, we evaluated the location and spread of the cortical response to thalamic microstimulation of varying waveform shapes, which were predicted in the literature to preferentially stimulate cell bodies instead of axons. We find that utilizing electrical microstimulation with cathode-leading asymmetric waveforms significantly reduced the bias of cortical activation and also increased the spatial specificity of cortical activation. And second, we developed a theoretical framework for the optimal design of artificial stimulation signal sets to maximize the discriminability of downstream cortical responses based upon the experimental results of chapter 2. The findings here provide some initial steps toward the design and optimization of microstimulation of neural circuitry, and open the door to more sophisticated engineering tools to develop technologies for more effective control of activity in the nervous system.

5.2 *Methods*

5.2.1 Surgery and preparation

The same methods described previously (see Appendix A for more detail) were used for the experiments in this chapter. Briefly, voltage sensitive dye imaging was used to record the response of primary somatosensory cortex in the anesthetized rodent to whisker deflections on the face and thalamic microstimulation.

5.2.2 Thalamic microstimulation

An electrode positioned in the thalamus was used to deliver single electrical current pulses to evoke cortical responses in the somatosensory pathway. The electrical stimuli were created using a digital stimulus generator (Model: DS8000, WPI Inc., Sarasota, Florida) and delivered using a digital linear stimulus isolator (Model: DLS 100, WPI Inc., Sarasota, Florida) acting in current source mode. Additionally, a fast switching relay was used to prevent charge accumulation on the electrode tip. All individual electrical stimuli were charge balanced. Three type of stimulus waveforms were used in this study: 1) a cathode-leading, symmetric biphasic waveform of 200 microseconds duration per phase (Symm), 2) a cathode-leading, asymmetric biphasic waveform with 1 millisecond of cathodal duration and 200 microseconds of anodal duration (ASymC), and 3) an anode-leading, asymmetric biphasic waveform with 1 millisecond of anodal duration and 200 microseconds of cathodal duration (ASymA). In the last part of this study, the asymmetry of the cathode-leading asymmetric, i.e. ASymC, waveform was systematically changed to more precisely determine the effect of the asymmetry on the specificity of electrical stimulation. Asymmetry was defined as the ratio of the duration of the first phase to the duration of the second phase minus one, such that the Symm waveform had an asymmetry of zero and the ASymC waveform had an asymmetry of four. Waveforms were delivered over a range of current amplitudes (30, 35, 40, 45, 50, 60, 70, 80, 90, 100, 120, 150 microamps). In each experiment, eight amplitudes were randomly chosen.

5.2.3 Whisker stimulation

Sensory stimulation (S-Stim) was applied through computer controlled whisker deflections. Whiskers were trimmed at approximately 12mm from the face, and were inserted into a glass pipette fixed to the end of a calibrated multi-layered piezoelectric bimorph bending actuator (range of motion, 1 mm; bandwidth, 200 Hz; Physik

Instrumente (PI), Auburn, MA) positioned 10 mm from the vibrissa pad. Vibrissae were always deflected in the rostral-caudal plane. Punctate deflections consisted of exponential rising and falling phases (99% rise time, 5 ms; 99% fall time, 5 ms) and angular deflection velocities of 75, 150, 225, 300, 450, 600, 900, and 1200 deg/s were used as mechanical probe stimuli (S-Stim).

5.2.4 VSD response analysis

All analyses were conducted using Matlab (Mathworks, Natick, MA). Voltage sensitive dye responses to mechanical deflection of whiskers (S-Stim) and to thalamic electrical stimulation, i.e. Symm, ASymC and ASymA, were averaged over 20 trials. The averaged responses at each frame were then fitted with 2D Gaussians [37]. The center of the 2D Gaussian was considered as the center of cortical activation.

A vectorized method was used to capture both the magnitude and direction of the offset between the VSD center of mass and the center of the barrel electrophysiologically matched to the electrode location (principal barrel). The center of the principal barrel was defined as the origin of the coordinate system. Because the barrel map orientation within the VSD imaging window was not strictly the same for each experimental session, another reference location was needed to ensure consistency of the coordinate system from day to day. The vector, \vec{r}_0 (black vectors in Figure 5.1(c)), given by the whisker immediately caudal to the principal whisker is used to define the zero degree direction. Within this coordinate system the magnitude and direction of the VSD response relative to the principal barrel are given by the vector, \vec{r}_{VSD} (blue and red vectors in Figure 5.1(c)), connecting the origin to the center of mass of the VSD response. The distance metric used in Figures 5.1 and 5.2 is given by the average of the magnitude of \vec{r}_{VSD} across experimental sessions. The topographic bias measurement used in Figure 5.1 is given by the angle between \vec{r}_{VSD} and \vec{r}_0 . The coherence in this bias across experiments is calculated using the vector strength, V ,

given by the equation below, where θ is the angle from i th experimental session [66]:

$$V = \frac{1}{n} \sqrt{\left(\sum_{i=1}^n \cos \theta_i\right)^2 + \left(\sum_{i=1}^n \sin \theta_i\right)^2}, \quad (22)$$

5.2.5 Functional Coding Analysis

Two theoretical models of the nonlinear input-output relationship were generated based upon the experimentally observed data to investigate the functional consequences of the disparate patterns of trial-to-trial variability for sensory and artificial stimuli. The ability of an ideal observer of cortical activity to classify response patterns is dependent on the statistical properties of the cortical response, such that the different patterns of trial-to-trial variability influence the amount of information that can be transmitted. Our goal was to determine the optimal set of cortical response amplitudes (referred to as the signal set, ϕ that would maximize the classification of an ideal observer, along with the stimuli that would generate optimal response amplitudes.

A standard observer model was used to set up the classification paradigm, where noisy observations of the cortical response, x_i , were classified as one of the signal classes, $\hat{\phi}_k$, given the true signal class, ϕ_k , that the cortical response belonged to in the signal set ϕ :

$$x_i = \phi_k + n_i \quad (23)$$

$$n \sim \mathcal{N}(0, \sigma(\phi_k)) \quad (24)$$

where the stimulus-dependent noise is Gaussian with zero mean and standard deviation, $\sigma_{PVM}(\phi_k)$, that is a function of the signal, ϕ_k . The signal set, ϕ , of size K contained K different scalar signals, $\{\phi_1 \phi_2 \dots \phi_K\}$, representing cortical response

amplitudes. The overlap across the elements of the signal set, given by the relationship between the signals and the variability, $\sigma(\phi_k)$, determined the overall performance of the classifier. As a simple example, consider a set of five different cortical response amplitudes, $\{\phi_1 \dots \phi_5\}$. Upon observing the noisy observation of the cortical response, x_i , the ideal observer must infer which response actually occurred, ϕ_i . For a small amount of noise, the observer will have little uncertainty about which ϕ_i produced x_i . However, for larger variability, the ideal observer may incorrectly classify x_i as ϕ_j , when it was really ϕ_i . Based on the relationship between the response amplitudes, ϕ , and the variability, $\sigma(\phi)$, our goal was to choose the specific ϕ_k that comprise the signal set, ϕ , such that the classification performance of the ideal observer was maximal.

Because the cortical response amplitude was related to the input through a monotonic, invertible sigmoidal function, as in 5.5(a), the notation was simplified by optimizing the set of cortical response amplitudes, rather than the input stimulus itself. However, it is important to note that the optimal input stimulus set, \mathbf{S} , can be computed directly from the optimal signal set, ϕ , according to:

$$\mathbf{S} = f^{-1}(\phi), \quad (25)$$

where f describes the invertible, nonlinear relationship between the input stimulus and the cortical response amplitudes.

We based the two theoretical models on the distinct trends in trial-to-trial variability for sensory and artificial stimuli detailed in Chapter 2. For sensory stimuli, the variability increased linearly as a function of the mean response amplitude (increasing variance model, IVM). We modeled this as a linear function, $\sigma_{IVM}(\phi_k)$ in figure 5.5. For artificial stimuli, the variability peaked at the threshold response amplitude. We modeled this as a piece-wise linear function of the mean response amplitude, with a peak in the variability at the threshold response amplitude (peak variance model,

PVM), given by $\sigma_{PVM}(\phi_k)$ in figure 5.5. The variability increases linearly with the mean response amplitude in the sub-threshold regime and decreases linearly with the mean response amplitude in the supra-threshold regime.

To maximize the performance of the classifier, we optimized the elements of the signal set by maximizing the distance between neighboring signals within the set. A symmetric version of the Kullbeck-Liebler (KL) divergence, referred to as the resistor-average distance (RAD), was used as the information theoretic distance between pairs of Gaussian probability distributions given by the models described above. The KL divergence, $D_{1,2}$, between the probability distributions for two neighboring signals within the signal set was calculated analytically according to the following equation:

$$D_{1,2} = \frac{((\phi_1 - \phi_2)^2 + \sigma^2(\phi_1) - \sigma^2(\phi_2))}{2\sigma^2(\phi_2)} + \log \frac{\sigma(\phi_2)}{\sigma(\phi_1)}, \quad (26)$$

where ϕ_i and $\sigma(\phi_i)$ give the mean and standard deviation of the Gaussian probability distribution i^{th} class in the signal set. The value $D_{1,2}$ is not symmetric. However, the RAD, which is a computationally tractable approximation to the Chernoff distance, is made symmetric by taking the harmonic mean of $D_{1,2}$ and $D_{2,1}$ [91], according to the following equation:

$$\frac{1}{R_{1,2}} = \frac{1}{D_{1,2}} + \frac{1}{D_{2,1}}. \quad (27)$$

Through its close approximation to the Chernoff distance, the RAD is related to the exponential decay rate of the total probability of error for an optimal classifier [41].

A custom optimization routine was used to maximize the pair-wise RAD between nearest neighbor probability distributions along the static nonlinearity. When the size, K , of the signal set, ϕ , was two, the maximization was trivial:

$$\arg \max_{\phi} R_{1,2} \quad (28)$$

In other words, the two elements, ϕ_k , that comprise ϕ were chosen to maximize the RAD distance between their underlying probability distributions. Intuitively, this was achieved by maximizing the difference between the response amplitudes and minimizing the variability. For $K > 2$, the pair-wise RAD of the nearest neighbor distributions, $R_{k,k+1}$ and $R_{k-1,k}$, were simultaneously maximized by iteratively adjusting each ϕ_k .

We evaluate the classification performance following the pair-wise RAD optimization scheme by computing the average probability of correct classification, P_C , by an ideal observer of the output performing maximum likelihood classification. Explicitly, this was the probability that the classified signal, $\hat{\phi}_i$, was in fact the signal delivered, ϕ_i . This value can be computed analytically from the Gaussian probability density function according to the following equation:

$$P_C = P(\hat{\phi} = \phi) = \int_{x_1}^{x_2} \mathcal{N}(\phi_k, \sigma(\phi_k)) dx \quad \text{for the interval } [x_1 \ x_2]$$

$$\text{such that } \mathcal{L}(x|\phi_i) > \mathcal{L}(x|\phi_j) \quad (29)$$

$$\text{for all } i \neq j \text{ and } x \in [x_1 \ x_2]$$

where $N(\phi_k, \sigma(\phi_k))$ is the Gaussian probability density function and $\mathcal{L}(x|\phi_k)$ is the likelihood of having observed x given signal ϕ_k . In this way, performance of the ideal observer was computed as the average P_C across all signals, ϕ_k , within the signal set, ϕ . Alternately, we also utilized a global optimization scheme, as opposed to the pair-wise RAD scheme above, based upon maximizing the minimum P_C across the signal set. Such an optimization scheme is typical in the communications literature as it emphasizes utilization of the entire signal set [103].

5.3 Stimulus waveform design augments spatial response properties

All measurements were conducted in the thalamocortical circuit of the rodent vibrissa/whisker pathway in-vivo. The basic setup for whisker stimulation, thalamic recording and microstimulation, and imaging of cortical activation was identical to that described previously (see Appendix A for more detail). The spatiotemporal dynamics of whisker-driven cortical responses were measured using voltage-sensitive dye (VSD) imaging in-vivo. Deflection of a single mystacial whisker attached to a computer controlled multilayer piezoelectric actuator (see methods) evoked a robust cortical response 15-25ms following the facial whisker deflection, consistent with previous studies [152]. Although the cortical activation was initially topographically restricted in and around the column/barrel associated with the PW, the cortical response quickly spread across a large number of adjacent columns following a single whisker deflection.

5.3.1 Functional validation of thalamic electrode placement.

A critical analysis here was that of the comparison of the cortical response to stimulation of a single whisker (Sensory Stimulation, or S-Stim) with that induced by thalamic microstimulation (Electrical Stimulation, or E-Stim). It was thus necessary to precisely map the position of the stimulating electrode within the thalamic VPM relative to the whisker array on the face. This was achieved by recording whisker-driven spiking activity in response to a punctate deflection applied to individual whiskers in the array. The principal whisker was defined as the whisker which evoked the strongest spiking activity in terms of the mean spike count within a 30ms window following the whisker deflection [196]. On average, the PW elicited significantly more spikes per deflection as compared to the adjacent whiskers (AWs) ($p=0.01$, Mann-Whitney U-test, $n=7$), consistent with previous studies [25].

5.3.2 Thalamic microstimulation results in a systematic topographical bias in cortical activation.

The locus of activation of the primary somatosensory cortex carries important information about the somatotopy of the body surface. Activation of cortex through thalamic microstimulation revealed a systematic bias in the locus of activation that was not seen for the sensory stimulus. Figure 5.1(a) shows a typical VSD response, where the top row is the sequence of VSD images taken following a punctate deflection of a single whisker (S-Stim), and the bottom row is the sequence of VSD images taken after a single microstimulation pulse was delivered to the corresponding topographically aligned region in the VPM thalamus (E-Stim). The E-Stim was a conventional symmetric biphasic current pulse waveform ($60 \mu\text{A}$ in amplitude, $200 \mu\text{s}$ per phase). It is qualitatively apparent from this example that the locus of cortical activation was different between the S-Stim and E-Stim cases, and that the initial onset of activation from E-Stim was not topographically aligned with the cortical column/barrel to which this region of VPM thalamus projects (denoted with black dot in center of overlaid columnar map) as it is for the S-Stim response (top row).

The centroid of the Gaussian fit of the VSD signal was then used to quantify this effect, specifically through the distance between the centroid measured from the VSD signal (at onset frame) and the actual center of the column. On average, the centroid of cortical activation in response to the sensory stimulus was within the cortical column/barrel (distance less than $250 \mu\text{m}$, which is approximately the radius of a cortical column/barrel, denoted by dashed line in Figure 5.1(b)). In contrast, the centroid of the Gaussian fit of the VSD signal in response to thalamic microstimulation diverged significantly from the center of the actual cortical column (distance of centroid to actual center $>850 \mu\text{m}$, Figure 5.1(b)). The mis-alignment of the cortical activation with the actual barrel/column corresponding to the whisker/electrode location was

significantly greater for the thalamic microstimulation than for the whisker stimulation ($p < 0.001$, Wilcoxon signed-rank test, Figure 5.1(b), $n=9$). Furthermore, we have found that there was no correlation between topographic bias and stimulus intensity ($p = 0.539$, Mann-Whitney U test).

The discrepancy between the centroid of cortical activation in response to thalamic microstimulation and the centroid of the principal barrel column might be due to increased variability in the measurement of the centroid of the VSD in the case of the thalamic microstimulation. In this scenario, we would expect the centroid of the cortical activation to be randomly located relative to the actual center of the barrel/column. Figure 5.1(c) shows the distribution of vectors computed as the difference between the centroid of cortical activation and the actual center of the cortical column/barrel, capturing both magnitude and direction. While the vectors measuring discrepancy for the S-Stim case exhibited a more dispersed distribution and an overall difference well within the actual column (left plot, light blue line is average, relative to column denoted by dashed circle), the vectors associated with the thalamic microstimulation exhibited a strong bias along the cortical columns associated with a row of whiskers (right plot). To quantify this notion, we used vector strength to measure how well the vectors were aligned along a common direction, as opposed to randomly dispersed [66] (see Methods). Indeed, the vectors in the thalamic microstimulation case were more locked to the direction along the barrel row (vector strength: 0.91, Figure 5.1(c): right) than in the whisker deflection case (vector strength: 0.63, Figure 5.1(c): left), pointing to a systematic bias in the cortical activation in response to thalamic microstimulation.

5.3.3 Cathode-leading asymmetric microstimulation significantly improved the topographic activation of cortex.

Modeling studies have suggested that the conventional symmetric, charge-balanced current waveform pulse (utilized in Figure 5.1) may preferentially activate axons as

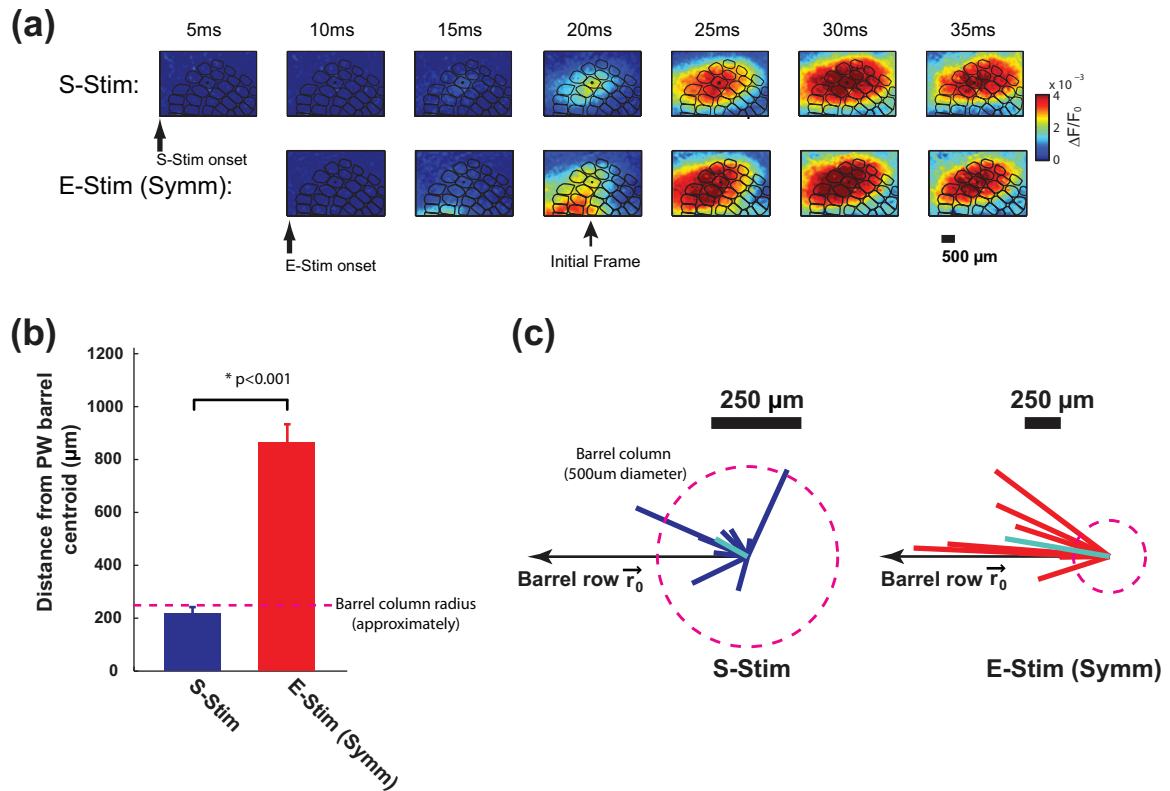


Figure 5.1: Thalamic microstimulation with symmetric waveforms results in a systematic topographical bias in cortical activation. (a) Top: typical VSD response to whisker stimulation. Bottom: typical VSD response to thalamic microstimulation with a symmetric waveform. (b) The centroid of the VSD signal in response to thalamic microstimulation (Symm) diverged significantly more from the center of the actual cortical column than the centroid of the VSD signal in response to sensory stimulation (S-Stim). (c) The distribution of vectors computed as the difference between the centroid of cortical activation and the actual center of the cortical column/barrel for both S-Stim and E-Stim (Symm) ($p < 0.001$, Wilcoxon signed-rank test, $n = 9$). The vectors from the E-stim (Symm) response show a consistent topographical bias.

opposed to cell bodies [123]. In contrast, it has been asserted that using biphasic current pulses with asymmetric waveforms can preferentially activate cell bodies. If the bias in cortical activation exhibited in Figure 5.1 was due to the activation of fibers of passage rather than cell bodies local to the electrode tip, then an asymmetric pulse design may alleviate some of this problem. Using a cathode-leading asymmetric waveform for the thalamic microstimulation significantly decreased the discrepancy between the centroids of cortical activation and the actual center of the barrel/column. This is first shown qualitatively in Figure 5.2(a), where the top row is again the sequence of VSD images following whisker stimulation, the middle row following thalamic microstimulation with a symmetric current injection waveform, and the third row following thalamic microstimulation with an asymmetric, cathode-leading current injection waveform. Qualitatively, we see that the asymmetric microstimulation results in an initial onset of cortical activation that is more topographically aligned with the response to the whisker stimulation, as compared to the symmetric microstimulation. The quantification in Figure 5.1(b) was repeated for the asymmetric, cathode-leading case, and resulted in a significant decrease in the topographic discrepancy (from $>850\mu\text{m}$ to $617\mu\text{m}$, $p=0.003$, Wilcoxon signed-rank test, Figure 5.2(b), $n=9$). Therefore, cathode-leading asymmetric microstimulation significantly diminished the discrepancy between the centroids of cortical activation and the centroid of the principal barrel, improving the topographic activation of cortex.

To rule out the possibility that this improvement was due to the difference of the derivative of charge delivered (the total charge per phase was the same for both AsymC and Symm, and the amplitudes of the anodal phase was the same), we also tested the anode-leading asymmetric waveform, which is the inverse of the cathode-leading asymmetric waveform. In contrast to the cathode-leading asymmetric waveform, the cortical activation induced by the anode-leading asymmetric waveform diverged even further from the correct topographic location than that for the symmetric

microstimulation waveform ($p < 0.001$, Wilcoxon signed-rank test, Figure 5.2(b), $n=9$). Interestingly, the difference in the topographic localization of the cortical responses induced by the different microstimulation waveforms was most prominent at the onset of cortical activation (initial frame), and was reduced as the VSD signal propagated beyond the initial frame (Figure 5.2(c)).

5.3.4 The topographic alignment of the cortical activation was enhanced with increasing asymmetry of the microstimulation waveform.

To measure the effects of temporal asymmetry in the charge delivery of the microstimulation on the localization of the cortical response, the degree of asymmetry in the microstimulation pulse was systematically varied. Specifically, a measure of asymmetry was defined as the duration of the cathodal phase over the duration of the anodal phase minus one, where the perfectly symmetric waveform has an asymmetry measure of zero, and increasing the duration of the initial cathodal component of the waveform increases the measure of asymmetry, as shown in Figure 5.3(a). Note that the waveform in each case was charge balanced, yielding a net charge delivery of zero over the duration of the waveform. The discrepancy between the centroid of cortical activation and the center of the actual column/barrel exponentially decayed with the asymmetry of the waveform, and saturated at approximately $650\mu\text{m}$ (Figure 5.3(b)).

5.3.5 Cathode-leading asymmetric waveform also improved the specificity of thalamic microstimulation.

In addition to affecting the locus of activation following thalamic microstimulation, the waveform properties also had a significant effect on the magnitude and area of cortical activation following microstimulation. The magnitude of the cortical response was defined as the magnitude of the two-dimensional Gaussian fitted to the cortical VSD signal, and the area of cortical activation was defined as the size of the cortical region spanned by the two-dimensional Gaussian at a magnitude of three times the standard deviation of the background noise. This is in contrast to the spatial spread

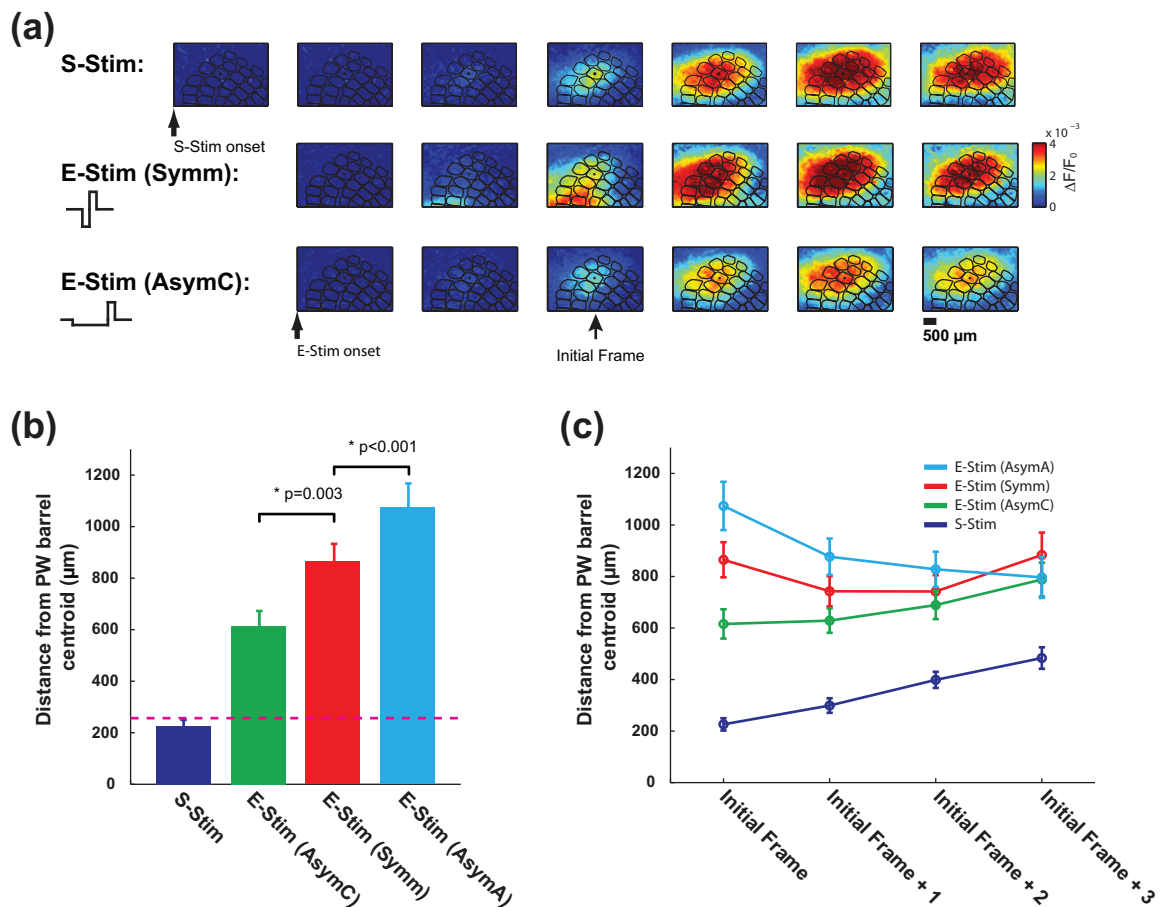


Figure 5.2: Cathode-leading asymmetric microstimulation significantly improved the topographic activation of cortex. (a) Top: typical VSD response to whisker stimulation. Middle: typical VSD response to thalamic microstimulation with a symmetric waveform (Symm). Bottom: typical VSD response to thalamic microstimulation with a cathode-leading asymmetric waveform (AsymC). The examples shown here are from same experimental session. (b) Topographic discrepancy of VSD signals in response to S-Stim, E-Stim (AsymC), E-Stim (Symm), and E-Stim (AsymA). Cathode-leading asymmetric microstimulation significantly decreased the topographic discrepancy ($p=0.003$, Wilcoxon signed-rank test, $n=9$). In contrast, anode-leading asymmetric microstimulation significantly enlarged the topographic discrepancy ($p<0.001$, Wilcoxon signed-rank test, $n=9$). (c) The difference in the topographic localization of the cortical responses induced by the different microstimulation waveforms was most prominent at the onset of cortical activation (initial frame), and was reduced as the VSD signal propagated beyond the initial frame.

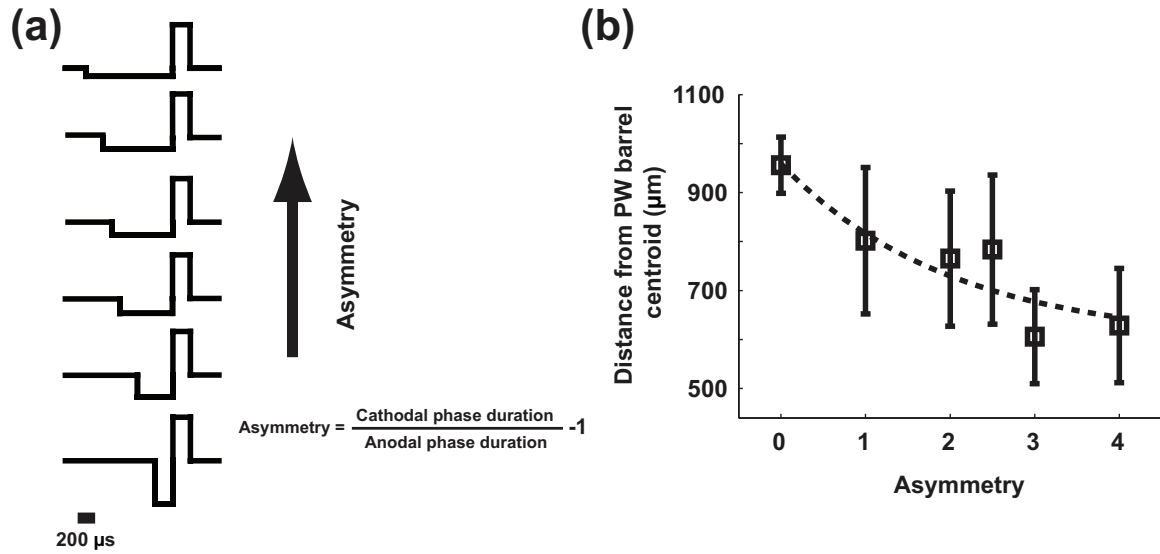


Figure 5.3: The topographic alignment of the cortical activation was enhanced with increasing asymmetry of the microstimulation waveform. (a) Asymmetric waveforms with different asymmetry. Asymmetry was defined as the duration of the cathodal phase over the duration of the anodal phase minus one, where the perfectly symmetric waveform has an asymmetry measure of zero. (b) The topographic discrepancy exponentially decayed with the asymmetry of the waveform, and saturated at approximately $650\mu\text{m}$.

measurement used in chapter 2, which was computed as the average radius of the 70% contour, thus normalizing for the amplitude of the response.

In general, the area and magnitude of the cortical activation increased with increasing stimulus strength. Figure 5.4(a) shows the relationship between increasing angular velocity of the punctate whisker deflection and the magnitude (left) and area (middle) of cortical activation, with a monotonic increase in each. The area of activation was coupled to the magnitude, with a fairly linear relationship (Figure 5.4(a), right). Similarly, with the thalamic microstimulation, there was an increase in the magnitude and area of cortical activation with increasing amplitude of the current pulse, shown for the symmetric and asymmetric cases in the left and middle panels of Figures 5.4(b) and 5.4(c). Again, in both of these cases, the magnitude and area of cortical activation were coupled, exhibiting a fairly linear relationship for each (Figures 5.4(b) and 5.4(c), right panels).

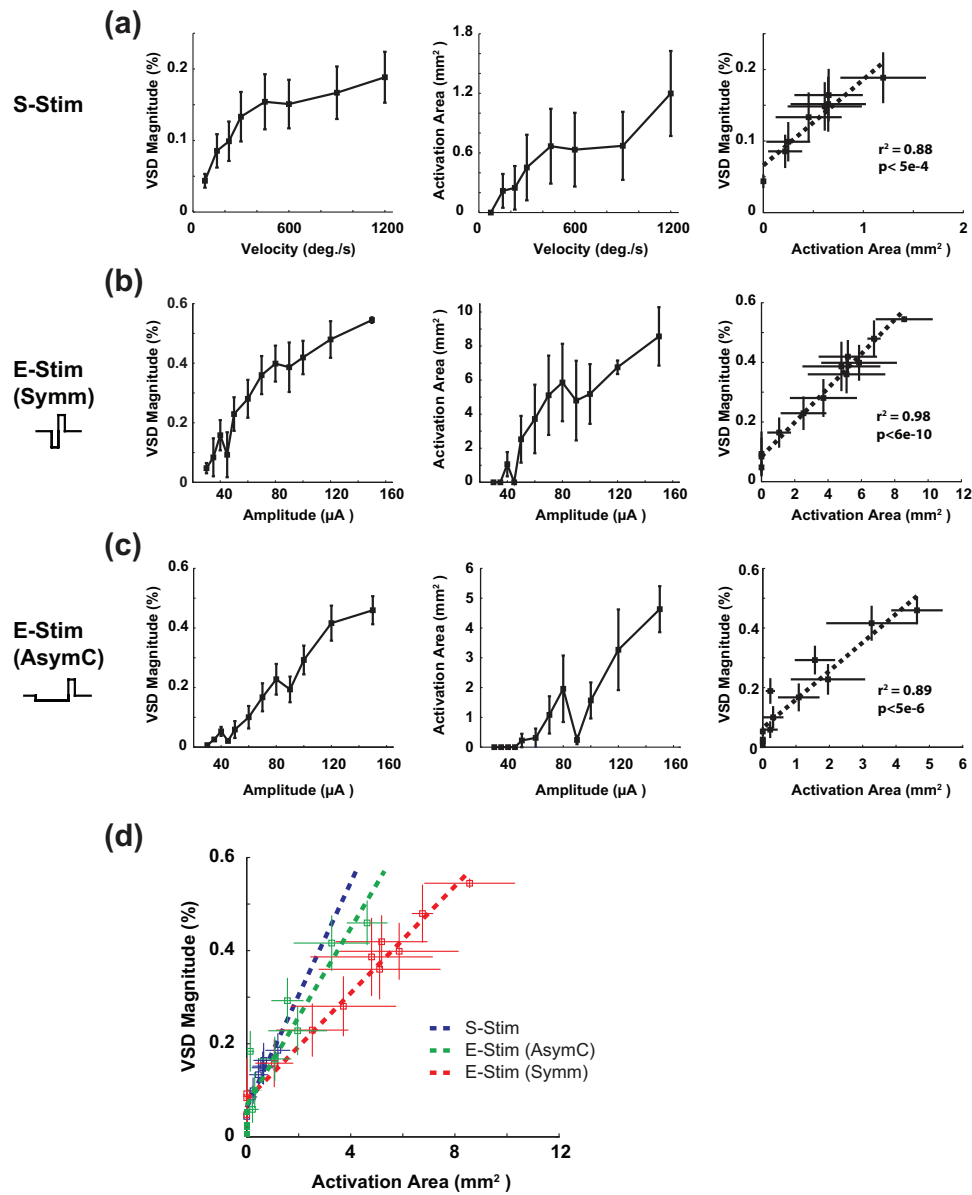


Figure 5.4: Cathode-leading asymmetric waveform also improved the specificity of thalamic microstimulation. (a) The relationship between increasing angular velocity of the punctate whisker deflection and the magnitude (left) and area (middle) of cortical activation. Right: the approximately linear relationship between the area of activation and the magnitude. (b) The relationship between increasing amplitude of the current pulse with symmetric waveform and the magnitude (left) and area (middle) of cortical activation. Right: the approximately linear relationship between the area of activation and the magnitude. (c) Same as in (b) for E-Stim (AsymC). (d) The slopes of the relationships between the area and magnitude of activation for S-Stim, E-Stim (AsymC) and E-Stim (Symm). Using the asymmetric waveform resulted in a slope much closer to that of the sensory stimulation (S-Stim).

In utilizing the thalamic microstimulation as a surrogate input to the pathway, the nominal goal is to activate cortex in a manner that is consistent with that of sensory stimulation (S-Stim, Figure 5.4(a)). To activate the appropriate area of cortex associated with a particular angular velocity of the tactile input, for example, the current amplitude of the microstimulation pulse can be adjusted. However, for the symmetric waveform, the cortical response was quite widespread, resulting in an area of activation that would be substantially greater than that corresponding to the sensory stimulation, even for relatively low values of current near threshold. Furthermore, the coupling of the area and magnitude of activation was such that if the magnitude of activation were matched to that of the sensory stimulation, the area of activation would be too great, and if the area of activation were matched to that of the sensory stimulation, the magnitude would be too low. This was not the case for the asymmetric microstimulation waveform. This is well captured by comparing the slopes of the relationships between the area and magnitude of activation for each case, as shown in Figure 5.4(d), where the slope for the symmetric case was significantly lower than that for sensory stimulation, but using an asymmetric waveform resulted in a slope much closer to that of the sensory stimulation.

5.4 Optimal signal set design maximized response discriminability

The distinct ways in which sensory and artificial stimulus-evoked activity propagate through neural circuits has profound implications for the ability to use artificial stimuli as surrogates for sensory stimulation. Specifically, the ability to discriminate between a set of different input strengths depends upon the mean amplitude and trial-to-trial variability of the cortical response. Similarly, the spatial spread would determine the ability to resolve various sensory inputs or stimulation on adjacent electrodes. We explored the simplest example of this using a theoretical model mimicking the nonlinear input-output relationship observed for sensory and artificial

stimuli described in Chapter 2. Using an observer model (see methods), we assessed the performance of a maximum likelihood classifier on a signal set, ϕ , with K elements, such that $\phi = \{\phi_1 \ \phi_2 \ \dots \ \phi_K\}$. Each signal, ϕ_k , was described by a Gaussian probability density function with coupled mean and standard deviation given by the nonlinear input-output relationship. Importantly, the goal was to determine the optimally discriminable cortical response amplitudes, given by the signal set ϕ , but the optimal stimulus set, \mathbf{S} , was easily extracted through the invertible nonlinear relationship between the input and output shown in the left portion of figure 5.5(a).

5.4.1 Signal set design accounts for distinct static nonlinear properties

In figure 5.5(a), we describe two models of trial to trial variability: 1) the variability increased linearly as a function of the mean response amplitude, $\sigma_{IVM}(\phi_k)$, similar to whisker stimuli, and 2) the variability peaks at the threshold response amplitude, $\sigma_{PVM}(\phi_k)$, as for electrical and optical stimuli. Both the peak variability and baseline variability were equal across the two models, as was the case experimentally. From a traditional population coding perspective, and under constant variability, the large slope in the input-output relationship at the threshold would provide the best discriminability between the response amplitudes. However, this was exactly where the stimulus dependent variability, $\sigma_{PVM}(\phi_k)$, in the model of electrical and optical stimuli was highest, which would significantly impair classification performance. Thus, through the following signal set design analysis, we directly investigated the functional coding consequences of the distinct trial-to-trial variability.

The relationship between the mean and standard deviation establishes the probability of discriminating between the various signals. We used the resistor-average distance (RAD), a symmetric version of the Kullbeck-Liebler (KL) divergence, to quantify the distance between neighboring signals [91]. The RAD is a computationally tractable approximation to the Chernoff distance, which is related to the

exponential decay rate of the total probability of error for an optimal classifier as more data accumulates [41]. In this way, a large RAD between two distributions indicates better classification performance. Thus, we can maximize the classification performance between two stimulus strengths by maximizing the RAD between the corresponding distributions of the cortical response. For both σ_{IVM} and σ_{PVM} in figure 5.5(b), the best classification performance, for $K = 2$, occurred between the extrema of the saturating nonlinearity as it maximized the RAD. For greater than two inputs, however, the optimal choice of the signal set, ϕ , was non-trivial. To find the optimal signal sets, we maximized the pair-wise RAD between the nearest neighbor signal distributions. The optimal signal sets for $K = \{2, 3, 6\}$ signals are shown for σ_{IVM} and σ_{PVM} in the right portion of figure 5.5(b).

After performing the optimization, we calculated the average probability of correct classification by an ideal observer as a function of the size of the signal set, as shown in figure 5.5(c). The average performance for σ_{IVM} and σ_{PVM} was found to be nearly identical. As the size of the signal set increased, the average performance decreased due to the increasing overlap between the nearest neighbor distributions, eventually approaching chance performance (dashed line). The results indicate that although activity propagates differently in the two models, knowledge of the relationship allows for the identification of optimal signal sets that maximize the performance of an ideal observer of the output.

However, evaluating the performance of each class within the signal set illuminates some key differences between the two models. Figure 5.5(d) presents the classification matrix for σ_{IVM} and σ_{PVM} . Each location in the matrix indicates the probability of an ideal observer classifying a trial as any of the signals, ϕ_i , given the true signal, ϕ_j . For both σ_{IVM} and σ_{PVM} , the minimum, ϕ_1 , and maximum, ϕ_K , signal had the highest performance because they only had one nearest neighbor to compete with. But, for the intermediate classes the trends were quite different. For σ_{IVM} ,

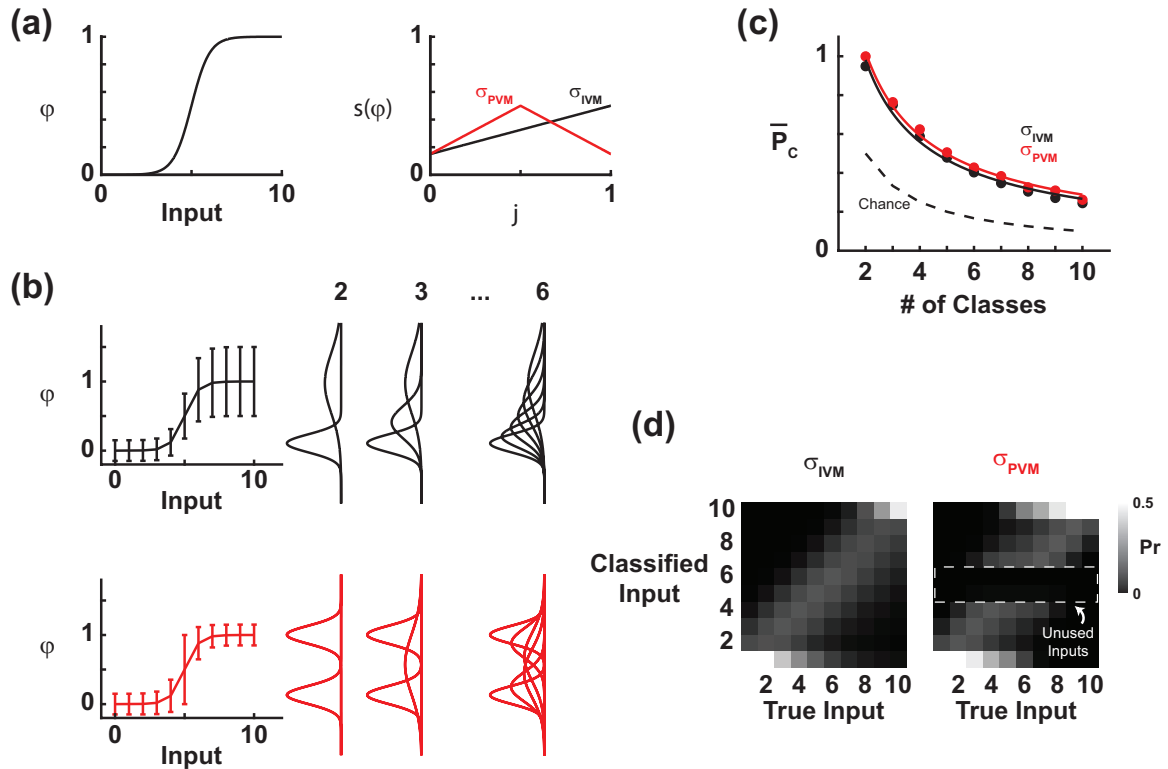


Figure 5.5: Optimal signal set design can account for distinct static nonlinear properties. (a) Two models of trial-to-trial variability. Model one (black) has increasing variability, σ_{IVM} , as function of the mean output. Model two (red) has a peak in variability, σ_{PVM} , at the threshold output. The mean output as a function of the input (arbitrary units) was the same for both models. (b) The variability made it non-trivial to estimate the output response on an individual trial, and this became more difficult as more response levels are to be encoded (right: examples of 2, 3, and 6 signals for σ_{IVM} and σ_{PVM}). (c) Average probability of correct classification by an ideal observer when the pairwise RAD was optimized for σ_{IVM} and σ_{PVM} . The classification performance was identical for the two models and decreases for larger signal sets. The dashed line indicates chance performance. (d) Discrimination matrices for classification performance of σ_{IVM} (left) and σ_{IVM} (right). Each index in the matrix gives the probability of classifying an input as x_i given that it was actually x_j , for all i and j .

the probability of correctly identifying the signal (diagonal values in the classification matrix) was equal for all intermediate signals, and errors were made to the nearest neighbor distributions. For σ_{PVM} , the same trend held, with the exception of the two middle signals, for which the probability of correctly identifying the signal was nearly zero. Thus, although the two models had identical average performance, σ_{PVM} conveyed fewer input signals accurately under the maximization of the pair-wise RAD.

5.4.2 Functional consequences of nonlinear properties of activity propagation

To penalize for ignoring portions of the signal set, we determined the optimal sets through a different optimization scheme. Rather than maximizing the pair-wise distances, we instead maximized the minimum probability of correct classification by an ideal observer across all signals within the signal set, as illustrated in figure 5.6(a). In this way, we forced the optimal signal sets to use all of the input classes, as is evident from the classification matrices in figure 5.6(b). Once again, the distributions with the means at the extrema, ϕ_1 and ϕ_K , had the highest performance in each of the models. However, under this optimization all signals were used for both of the models. The minimum probability of correct across all classes is shown as a function of the signal set size in figure 5.6(c). We now observe different trends across the two models. σ_{PVM} performed better for a signal set with only two classes because of the low variability at the minimum and maximum response amplitudes, whereas σ_{IVM} performed better for signal sets with more than two classes. These results suggest that a static nonlinearity similar to σ_{IVM} , with monotonically increasing variability, should provide better discrimination performance across multiple input classes. Conversely, a static nonlinearity similar to σ_{PVM} , with peak variability at threshold, provides better detection information.

More importantly, these results highlight that comparable levels of performance

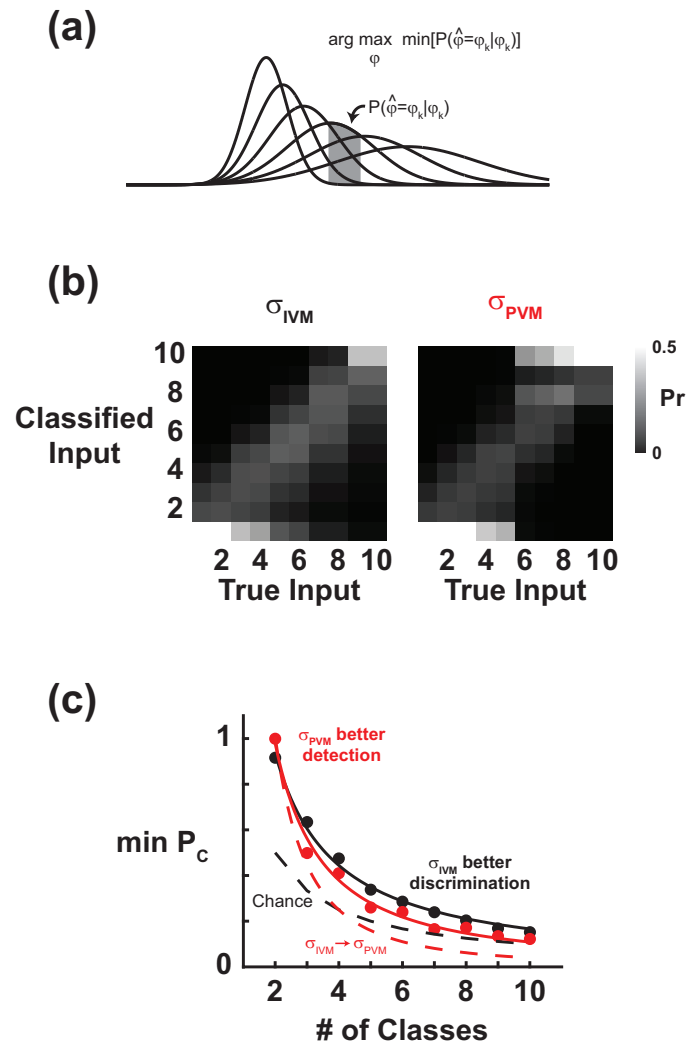


Figure 5.6: Functional consequences of trial-to-trial variability on response discriminability. (a) An alternate optimization scheme was to maximize the minimum probability of the ideal observer correctly classifying an input. The highlighted region indicates when that input would have the highest likelihood and thus be correctly classified. The area under the curve in the shaded region gives the probability of the ideal observer correctly classifying the input. (b) Discrimination matrices for classification of σ_{IVM} (left) and σ_{PVM} (right) under the max-min optimization scheme. Both of the models used all input classes. (c) Minimum probability of the ideal observer correctly classifying the input under the max-min optimization scheme. σ_{PVM} provided better performance for two inputs, while σ_{IVM} provided better performance for larger signal sets. The black dashed line indicates chance, while the red dashed line is the performance when using the optimal signal set from σ_{IVM} as the input to σ_{PVM} .

can be achieved across distinctly different response properties as long as the appropriate signal set is used. To illustrate this, we used the optimal signal set from σ_{IVM} as the input to σ_{PVM} and evaluated the performance (dashed red line in figure 5.6(c)). Using the mismatched signal set, the performance was considerably worse than the optimal signal set. As the size of the signal set increased, the discrepancy became greater and performance fell below chance. Ultimately, our theoretical results suggest that the optimal signal set will not be the same for sensory and artificial stimuli, but by using knowledge of the input-output response transformations, we may be able to design signal sets to maximize performance for artificial stimulation of neural circuits.

5.5 Discussion

Here, we characterized the response of primary somatosensory cortex to both peripheral tactile stimulation and electrical stimulation of the thalamic region providing direct input to barrel cortex. Conventional symmetric biphasic electrical stimulation activated a broad region of the cortex centered at a region that was multiple cortical columns away from the cortical column which was topographically aligned with the electrode location. Although not directly assessed here, the consequences of this topographic mismatch would likely produce undesired percepts because previous work with brain imaging has suggested that there is a strong correlation between cortical activation and perception [36, 9]. Previous studies using thalamic stimulation in humans have reported that percepts induced by electrical stimulation are often unnatural or discordant with relation to the anatomical electrode position [145, 148, 71], potentially linked to the phenomena we describe here.

Although it has been hypothesized that the aberrant sensations resulting from microstimulation are due to activation of axons passing by the electrode tip in addition to the nearby cell bodies, the downstream neural response causing these sensations

has not been studied in detail for functional validation. Beyond the sensory periphery, the anatomy of the circuitry places neuronal cell bodies and axonal fiber tracts within close proximity of one another. As a result, microstimulation can activate neurons with cell bodies that are quite far from the site of stimulation [83, 29], beyond that predicted by conventional understanding of how electric fields grow with the magnitude of the current injection. What has been observed is that neurons with cell bodies very distal to the stimulation site can still be activated for very small currents, attributed to the crossing of fiber tracts through the volume of tissue affected by the electric field, resulting in what Butovas and Schwarz refer to as a spatiotemporal “blurring” of the activation. In a more recent study, two-photon microscopy was utilized to observe the neural responses to nearby electrical stimulation in-vivo. Through this technique, the authors provide strong evidence that electrical stimulation principally activates axons within a very short radial distance from the electrode [83].

Previous modeling work predicts that alterations in the stimulation waveform change the relative stimulation thresholds for cell bodies near the electrode and passing axons [123]. This is hypothesized to occur through differential inactivation of sodium channels at the cell body and nearest node of Ranvier caused by the long subthreshold cathodal phase of the asymmetric stimulus waveform [69, 123]. The sodium channels at the node of Ranvier are inactivated by the depolarizing cathodal pulse and then the subsequent short duration, high current anodal phase preferentially stimulates the nearby cell bodies. Although this study does not provide observation at this level of detail, functionally we observed a distinct difference in cortical response when using symmetric vs asymmetric stimulation that is consistent with this notion. Whereas the response to symmetric stimulation was centered multiple cortical columns away from the cortical column topographically matched to the electrode position, the response to asymmetric stimulation was shifted more closely

to the anatomically expected cortical column. Additionally, recent work in awake behaving rats has shown that asymmetric waveforms had higher detection thresholds than symmetric waveforms for intracortical microstimulation [96]. This finding is consistent with our results and together they indicate that asymmetric and symmetric stimulation likely activate different neural elements. Further investigation is necessary to more directly determine whether asymmetric pulses provide increased selectivity of cell bodies over axons.

Although the exact anatomy of the projecting fibers from the VPM thalamus to layer 4 of the primary somatosensory cortex has not been fully understood, histological evidence suggests that fibers from VPM regions (barreloids) associated with more rostral tactile input pass by the VPM barreloids associated with more caudal input on the way to primary somatosensory cortex [75, 203]. Given this anatomical arrangement, we would expect that if symmetric stimulation preferentially activates axons as opposed to cell bodies, the cortical activation would be centered in the cortical columns corresponding to whiskers more rostral than the whisker corresponding to the electrode location. Indeed, we observed a systematic topographic bias in the center of mass of the cortical response towards cortical columns representing rostral inputs when using symmetric stimulation. The systematic topographic bias, in concert with previous histological evidence, strengthens the support for the hypothesis that asymmetric stimulation increases the selectivity stimulation of cell bodies vs. axons.

While the results presented here are specific to stimulation of the VPM region of the somatosensory thalamus in the rodent, we expect the phenomenon to be consistent in other regions of the brain, and thus the findings here are general. We use thalamic stimulation due to the extensive literature detailing the anatomical and functional characterization of the thalamocortical circuit in the rodent vibrissa system [174, 171, 129], and the fact that it provides direct, mono-synaptic input to

cortical layer 4 in the vibrissa region. Furthermore, stimulation in the thalamus paired with recording in cortex allows us to observe the direct neural response to electrical stimulation on a network level by recording only two synapses downstream of the stimulation, with voltage sensitive dye imaging of layer 2/3 cortical activity. Importantly, the site of the stimulating electrode in the VPM thalamus is approximately four millimeters from the site of the VSD imaging in cortex, implying that the observed cortical activation is due to engagement of the thalamocortical circuit, rather than passive conduction through the tissue. The voltage sensitive dye imaging proves ideal for this study due to its high spatial and temporal resolution, allowing a full characterization of the spatiotemporal cortical response to a point source of current in the thalamus [72, 109]. While the voltage sensitive dye signal is limited to layer 2/3 activity due to light scattering, it has been shown that early periods of activation in layer 2/3 reflect the activation patterns within layer 4, suggesting that the VSD imaging at the activation onset is likely very similar to the activity of the input layer 4 [152]. It should be noted that the topographic misalignment of the cortical activity in response to thalamic microstimulation is most pronounced at the onset of cortical activation. However, the later stage of activation spreads across a large region of cortex and even so, some degree of disrupted somatotopy does persist following the initial activation. The coding of any sensory stimuli more complex than the unitary impulses we utilize here could thus not rely on later stages of cortical activation to alleviate the problem. Initial activation of areas not topographically aligned with the desired percept would preclude subsequent or concurrent activation of that region in the context of more rich, spatiotemporal sensory stimuli. Finally, although the VSD activity has been linked to sub-threshold membrane potentials, and thus the measured signals do not directly reflect action potentials, action potentials do obviously accompany strong depolarizations and the hard-thresholding we apply to the VSD imaging likely captures the distinction.

Our findings have strong implications for delivering information to the brain for sensory prosthetics applications. Microstimulation disrupts the somatotopy that the brain strives to maintain throughout development, possibly through the activation of passing axon fiber tracts, which would be a problem in almost all brain regions including the most simply arranged anatomy of the periphery. This could lead to nearby electrodes producing similar or overlapping perceptions if the same axons pass by each electrode. For instance, Grill, et al, reported that humans often experience paresthesias in response to thalamic microstimulation that are localized to a different somatotopic region than the mapped sensory receptive field of the stimulating electrode, termed a discordant paresthesia [71]. While it is clear that the capabilities for plasticity and learning in the brain improve the likelihood of success in interpreting this surrogate signaling to the brain, there is a fundamental limit in that the sensations must be discriminable to be functionally useful [193]. This problem could be alleviated by ensuring that a one-to-one topographic map is preserved, even if different from the one of the normal physiology. By more closely adhering to the normal somatotopy, asymmetric stimulation pulses ensure that sensations produced by nearby electrodes are less affected by the underlying axonal anatomy and more likely to be distinct and natural sensations. Additionally, asymmetric stimulation pulses result in cortical responses that have a higher spatial specificity, and much closer to normal sensory stimulation, than symmetric biphasic pulses. Given that the discriminability of information delivered to the brain is likely affected by the overlap of the cortical responses to neighboring inputs, asymmetric cathodal stimulation would thus likely increase the functional resolution of a sensory prosthetic.

In addition to the spatial discriminability, the distinct nonlinear response properties for sensory and artificial stimuli pose a problem for encoding information in the strength of the neural response. A functional treatment of the distinct patterns of trial-to-trial variability suggests that sensory and artificial stimuli may provide

different amounts of information to the neural circuit. Through optimal signal set design, knowledge of the trial-to-trial variability can be used to maximize information throughput, but ultimately the pattern of increasing variability observed with whisker stimuli provides greater discrimination information to downstream structures.

Previous psychophysical experiments involving the detection or discrimination of sensory or artificial inputs provides insight into the functional consequences of the theoretical results described above. Two common task designs for exploring the sensation caused by sensory and electrical stimuli are the “go-no go” paradigm and the two alternative forced choice task. The “go-no go” paradigm requires the animal to respond to a particular stimulus and not respond otherwise, and has been used extensively to study the detection of sensory stimuli and artificial inputs delivered at various stages of the sensory pathways. The psychometric curve describing the probability of the animals response to varying strength of the stimulus is typically sigmoidal [180, 146], much like the amplitude-response curves in figure 5.5(a) of this study. The high variability at threshold described for electrical and optical stimuli in chapter 2 and similar to σ_{PVM} suggests that the animal would experience all-or-none sensations preventing the discrimination of stimulus strengths. The “go-no go” paradigm does not produce the requisite information to determine the sensitivity of the psychometric curve because the response is all-or-none itself. To test the predictions from this study, a two alternative forced choice task design (2AFC) would be required, where the animal can respond in two different ways, one response for one stimulus and another for a different stimulus. In this way, the animal could be presented with two stimuli simultaneously, or in succession, and respond in one of two ways to indicate which stimulus was stronger. However, such a task would require simultaneous stimulation on two electrode sites, which likely would not have the same static nonlinearity, or successive stimulation on a single electrode site, requiring the animal to remember the strength of the first.

Taken together, the results here provide some initial directions in the design of charge delivery for producing functionally relevant activation in downstream circuits through electrical microstimulation. Future experimental and modeling studies may further optimize the stimulus design and expand the problem to more complex patterns of microstimulation across electrode arrays, which are requisite to deliver the high amount of information embedded in the sensory environment. For example, the optimal signal set design could be expanded to simultaneously optimize the performance of an ideal observer in classifying the spatial location of the response in addition to the amplitude. Just as this study sought to characterize and control the nonlinearities involved in the amplitude and spatial activation profile of a single stimulus pulse, future work must also address the dynamic nonlinearity, described in chapters 3 and 4, that governs the neural response to temporal sequences of stimulation in order to gain precise control of complex neural circuits, which is essential for high fidelity sensory prosthetic technology.

CHAPTER VI

CONTROL OF THE NONLINEAR DYNAMICS OF CORTICAL ACTIVATION THROUGH OPTOGENETIC MANIPULATION OF THALAMIC STATE

Portions of this chapter were presented in poster form at the following conference:

Millard, DC, Gollnick, CA, Hendry, WJ, Rozell, CJ, & Stanley, GB. The role of magnitude and synchrony of population activity in nonlinear circuit processing in the thalamocortical circuit of the rodent vibrissa system. Society for Neuroscience Annual Meeting, San Diego, CA, November 2013.

6.1 Introduction

The thalamus has traditionally been considered a “relay station”, merely passing along sensory information from the periphery to the cortex. However, due to the unique circuit properties of the thalamus, others have proposed that the thalamus is ideally positioned to gate and/or modulate information transmission to cortex [42]. Specifically, thalamic neurons are known to operate in two distinct firing modes: tonic firing and burst firing [170]. Sherman proposed that, by dynamically switching between these two states, the thalamus may modulate not only “how much”, but also “what type” of information is transmitted to cortex [65, 157, 196, 122, 177, 170]. Here, we directly control the transition between the bursting and tonic firing modes in the sensory thalamus and determine the extent to which the thalamic state modulates the propagation of neural activity to cortex.

While the neural circuit in the thalamus is composed of a canonical relay with

feedforward inhibition [171], the individual neurons contain specialized membrane dynamics enabling diverse firing statistics. In the tonic firing mode, thalamic neurons fire action potentials at high rates with Poisson-like statistics. The burst firing mode, on the other hand, is characterized by long periods of quiescence and brief volleys of action potentials within a short time window [164], mediated by T-type calcium channels [170]. Switching between these two state occurs regularly in concert with shifts in arousal [54, 111, 161] or attention [122], and is linked through these neuromodulatory mechanisms to the “desynchronized” thalamocortical state [43, 160]. Through the use of optogenetics, recent work has demonstrated that the control of thalamic polarization is sufficient to drive cortex into the synchronized [76] or desynchronized [159] states. And yet, the effect of thalamic state on the propagation of stimulus-evoked activity to downstream structures remains poorly understood.

Under direct optogenetic control of thalamic depolarization, we quantified the propagation of stimulus-evoked activity from the thalamus to cortex. Voltage sensitive dye imaging was used to measure the spatiotemporal cortical response to thalamic microstimulation, while the thalamic state was controlled through optogenetic means. First, at the level of the thalamus, depolarization eliminated bursting activity. However, through precise and systematic modulation of the thalamic firing rate, we find that tonic and burst firing occupy two ends of a continuum, rather than discrete states. Within this continuum, the propagation of neural activity generated by thalamic microstimulation was significantly altered. The trial-to-trial variability in the cortical response was reduced under depolarization, potentially increasing the discriminability in cortex. Further, the elimination of bursting prevented the paired pulse facilitation elicited by sub-threshold thalamic microstimulation, such that the dynamics were more similar to those of natural sensory stimuli. Ultimately, by directly modulating the thalamic state, we demonstrate control over the nonlinear propagation of activity to downstream structures, with potential application towards the development of

contextual sensory prostheses and more generally for understanding the role of state in the information processing capabilities of neural circuits.

6.2 Methods

6.2.1 Surgery and preparation

The same methods described previously (see Appendix A for more detail) were used for the experiments in this chapter. Briefly, voltage sensitive dye imaging was used to record the response of primary somatosensory cortex in the anesthetized rodent to whisker deflections on the face and thalamic microstimulation.

6.2.2 Optogenetic expression

All animals used for the experiment in this chapter underwent an initial surgery for the injection of a viral vector to induce expression of either channelrhodopsin (ChR2) or the stabilized step function opsin (SSFO) [16, 211, 212]. The injection was made in the ventral postero-medial thalamus according to coordinates described previously (see Appendix A). Viral expression was verified experimentally in all animals through light-responsive electrophysiological recordings, and using histology in select animals.

6.2.3 Optical stimulation

Light emitting diodes (LED) were used to excite the ChR2 and SSFO in vivo. An “optrode” was positioned in the VPM thalamus, with the fiber optic directly attached to the LED to minimize light loss. A 465nm LED was used for animals expressing ChR2. Because the ChR2 channel quickly closed when the light was removed [120], the light was delivered continuously to maintain the long timescale depolarization involved in this study. A closed loop current source was used to drive the LED [138]. For animals expressing the SSFO, a 465nm LED (LED Engin Inc, San Jose, CA) was used in combination with a 590nm LED (LED Engin Inc, San Jose, CA) through a wavelength combiner (Doric Lenses Inc, Quebec, Canada). The SSFO channel has

a long closing time, such that the channel remains open long past the duration of the light stimulus, so brief pulses of light ($\sim 5\text{-}50\text{ms}$) from the 465nm LED were used to open the channel over long timescales. To close the channel, 15 seconds of yellow light from the 590nm LED was delivered through the fiber optic.

6.2.4 Thalamic microstimulation

The same methods described previously and in Appendix A were used in this chapter. Paired stimuli, with an inter-stimulus interval of 150ms, were delivered to the thalamus with stimulus intensities of 5-150 μA . The response to the first stimulus in the pair was used to analyze the trial-to-trial variability, while the relative responses to the two stimuli were used to study the nonlinear dynamics. The 150ms inter-stimulus interval was chosen because it most reliably activated the facilitation described in Chapter 3.

6.2.5 Electrophysiology analysis

Single-unit activity was recorded in the thalamus to quantify the effects of depolarization on the firing statistics of the thalamic neurons. The electrophysiology methods are described in more detail in Appendix A. Bursts were identified as a collection of two or more spikes with inter-spike intervals less than four milliseconds and preceded by a 100 millisecond inter-spike interval [164]. Only well isolated units were used for the bursting analysis to prevent confusion of bursting with coincident multi-unit activity.

To quantify the activity level of the cells, event rate was used in place of firing rate, where a lone spike and a burst were each considered a single event. This was to distinguish between a cell firing five tonic spikes per second and another that produced a single burst of five spikes once per second, both of which would have an average firing rate of five Hz. Using the event rate to quantify activity levels, the cell firing tonic spikes would produce five events per second, while the bursting cell would

produce one event per second.

Finally, the bursting percentage was calculated as the proportion of events that were bursts. A simple upper bound can be placed on the burst percentage as a function of the event rate when modeled as a Poisson process. An event can only be a burst if the inter-event interval is greater than 100 milliseconds. In this way, the maximum possible burst percentage, B , is given by the proportion of inter-stimulus intervals greater than 100 milliseconds, as follows:

$$B = \int_{0.1}^{\infty} \lambda \exp(-\lambda t) dt = \exp(-0.1\lambda) \quad (30)$$

where λ is the event rate of the homogeneous Poisson process in units of Hertz. This upper bound is plotted in figure 6.3(c) for comparison with the experimentally determined relationship between burst percentage and event rate.

6.2.6 VSD analysis

The VSD analysis in this chapter was performed identically to that of previous chapters. Only the temporal VSD signal, averaged spatially within the topographically matched cortical column, was used to characterize cortical response in this work.

6.3 Results

The desynchronized state in cortex has been classically identified from the local field potential (LFP) or electro-encephalogram (EEG), as a transition from high amplitude, low frequency oscillations to a low amplitude, high frequency signal [43, 160]. Previous work in the rodent vibrissa system has shown that thalamic depolarization is sufficient to generate a desynchronized state in cortex [159]. Here, we experimentally measure the changes in thalamic firing statistics in the transition to a desynchronized thalamocortical state, and quantify the effect of state on the nonlinear dynamics of the cortical response to thalamic microstimulation.

6.3.1 Thalamic depolarization leads to desynchronized state in-vivo

In line with previous work in the rodent vibrissa system, we first verified that optogenetic depolarization of the thalamus produced the desynchronized state in cortex. An example is shown in figure 6.1. An “optrode” was positioned in the VPM thalamus and a separate electrode was positioned in layer 4 of the downstream primary somatosensory cortex. Under anesthesia, the spontaneous firing rate of the VPM neurons was low, as in figure 6.1. However, the delivery of blue light, indicated by the ramp in figure 6.1(a), depolarized the thalamic neurons by opening the ChR2 channels, driving an increase in the firing rate. There was a concomitant increase in the firing rate of the downstream multi-unit activity in cortex and decrease in power of the low frequency component of the cortical LFP. This is quantified in 6.1(b) at varying light intensities. As the thalamic firing rate increased, the low frequency power in the cortical LFP decreased (left) while the high frequency power (middle) remained constant. This resulted in an increase in the ratio of high frequency power to lower frequency power in the cortical LFP, which is characteristic of the desynchronized state in cortex and consistent with previous work in this pathway [159, 160, 43] and others [65, 157]. However, the associated firing properties of the thalamus in the desynchronized state have not been explicitly characterized.

6.3.2 Depolarization reduces spontaneous and stimulus-driven bursting

The thalamus is known to operate in two different modes: tonic firing and burst firing. The distinct firing properties in these two states are believed to have significant impact on the functional neural coding of sensory stimuli [177, 111] and relate to cortical state. The tonic firing mode in the thalamus is associated with the desynchronized state in cortex, and characterized by relatively high firing rates ($>10\text{Hz}$) and Poisson timing statistics [54]. In the burst firing mode, however, the neurons are likely to fire bursts of action potentials, defined as a cluster of two or more spikes in a single




Figure 6.1: Optogenetic control of thalamic state. (a) Using a ramp light stimulus (top) in the thalamus, ChR2 neurons are driven at an increased firing rate. This drives an increased firing rate in the simultaneously recorded cortical neuron and a desynchronization in the cortical LFP. (b) The desynchronization is characterized by a drop in the low frequency content in the cortical LFP (left), but no change in the high frequency content (middle). This leads to an increase in the ratio of high frequency power to low frequency power as a function of the thalamic firing rate (right).

neuron having inter-spike intervals less than four milliseconds [164].

Examples of the tonic and burst firing modes are shown in figure 6.2. Under anesthesia, the VPM neurons were principally in the burst firing mode, as in the top portion of figure 6.2(a) where the neuron fires classically defined bursts of spikes. Using the same experimental setup as before, a fiber optic attached to the electrode was used to deliver a constant amount of light to chronically depolarize the thalamic neuron and push it into the tonic firing mode, as in the bottom portion of figure 6.2(a). To further illustrate the difference in firing statistics between these two modes, we analyzed the inter-spike interval distributions. For each spike, the time since the previous spike (previous ISI) and the time until the next spike (next ISI) are plotted against each other in figure 6.2(b). Within these axes, the red boxes indicate a classically defined burst, where the box in the lower right corner signifies the start of a burst with a long ($>100\text{ms}$) previous ISI followed a short ($<4\text{ms}$) ISI, and the box in the lower left corner contains any subsequent spikes in the burst. In the burst firing mode, a large portion of the spikes (52% in this example) fall within the delimited burst regions. The tonic firing mode, however, has a much smaller proportion of spikes in the burst regions, and instead the majority of the spikes lie in a cloud along the diagonal of the axes, which is typical of Poisson firing statistics.

Through precise optogenetic control of thalamic depolarization, we tested whether the tonic and burst firing modes formed two distinct states, or rather if they represented two ends of a continuum. To do so, we used the stabilized step function opsin (SSFO) to precisely and systematically vary the depolarization of the thalamic neurons. The SSFO channel has a long time constant (approximately 30 minutes) to close, making it a nearly perfect integrator of light that directly mapped the light exposure to a depolarizing inward current [212]. In this way, we delivered a train of brief, low intensity light pulses through the optical fiber at 0.2Hz, or one pulse every five seconds. Each pulse of light slightly increased the depolarization of the neuron,

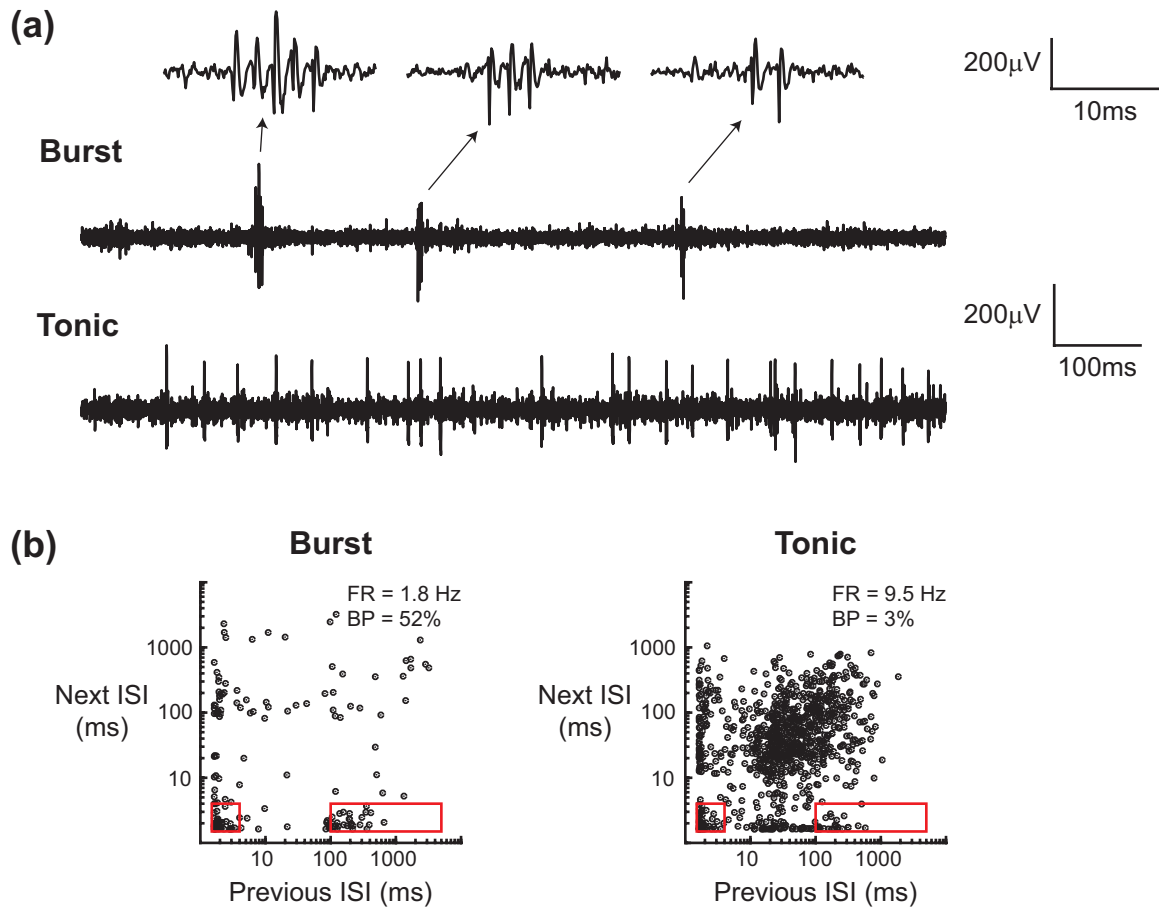


Figure 6.2: Depolarization of the thalamus eliminates bursting. (a) Top: Under anesthesia, thalamic neurons spontaneously burst, firing multiple action potentials in quick succession. Bottom: When depolarized through optogenetic manipulation, the thalamic neurons switch to a tonic firing mode. (b) For each spike, the interspike interval to the previous spike is plotted against the interspike interval to the next spike. The red boxes outline spikes that could make up a burst. In the anesthetized state (control, left), the proportion of spikes that make up bursts is 52%. When depolarized (right), the majority of spikes occur in tonic firing, and very few (3%) occur during bursts.

with the firing rate and burst percentage computed during each five second interval. An example is shown in figure 6.3, with the raw voltage trace on the bottom of panel (a) and a histogram (one second bins) of the spiking activity on the top. As the light pulses continued, and more of the SSFO channels in the cell opened, the firing rate steadily increased.

Within each five second interval, the event rate and burst percentage were calculated. The event rate was used instead of the firing rate, such that a burst of spikes, regardless of the number of spikes in the burst, and a single tonic spikes each counted as a single event. The burst percentage was computed as ratio of the number of burst events to the total number of events. An example of the relationship between burst percentage and event rate is shown in figure 6.3(b). As the cell received more pulses of light and the event rate increased, the burst percentage decreased with an exponential relationship. This trend was consistent across cells for both ChR2 and SSFO stimulation (ChR2: 6 cells, SSFO: 2 cells, N=7 animals total), with the exponential fit to each cell (gray) and the fit to the compiled data across all cells (black) shown in figure 6.3(c). For comparison, the red line indicates the upper bound on burst percentage, as a function of event rate, assuming Poisson event statistics (i.e. every event with inter-event interval $>100\text{ms}$ produced a burst and inter-event intervals were exponentially distributed). The average relationship across cells was far below the upper bound, such that, even for the lowest event rate, the occurrence of a burst was not certain. And further, the continuous relationship between burst percentage and event rate suggests that the tonic and burst firing modes occupy the two extremes of a continuum. Even so, the effect of the tonic and burst firing modes on the propagation of neural activity in the thalamocortical circuit remains unknown.

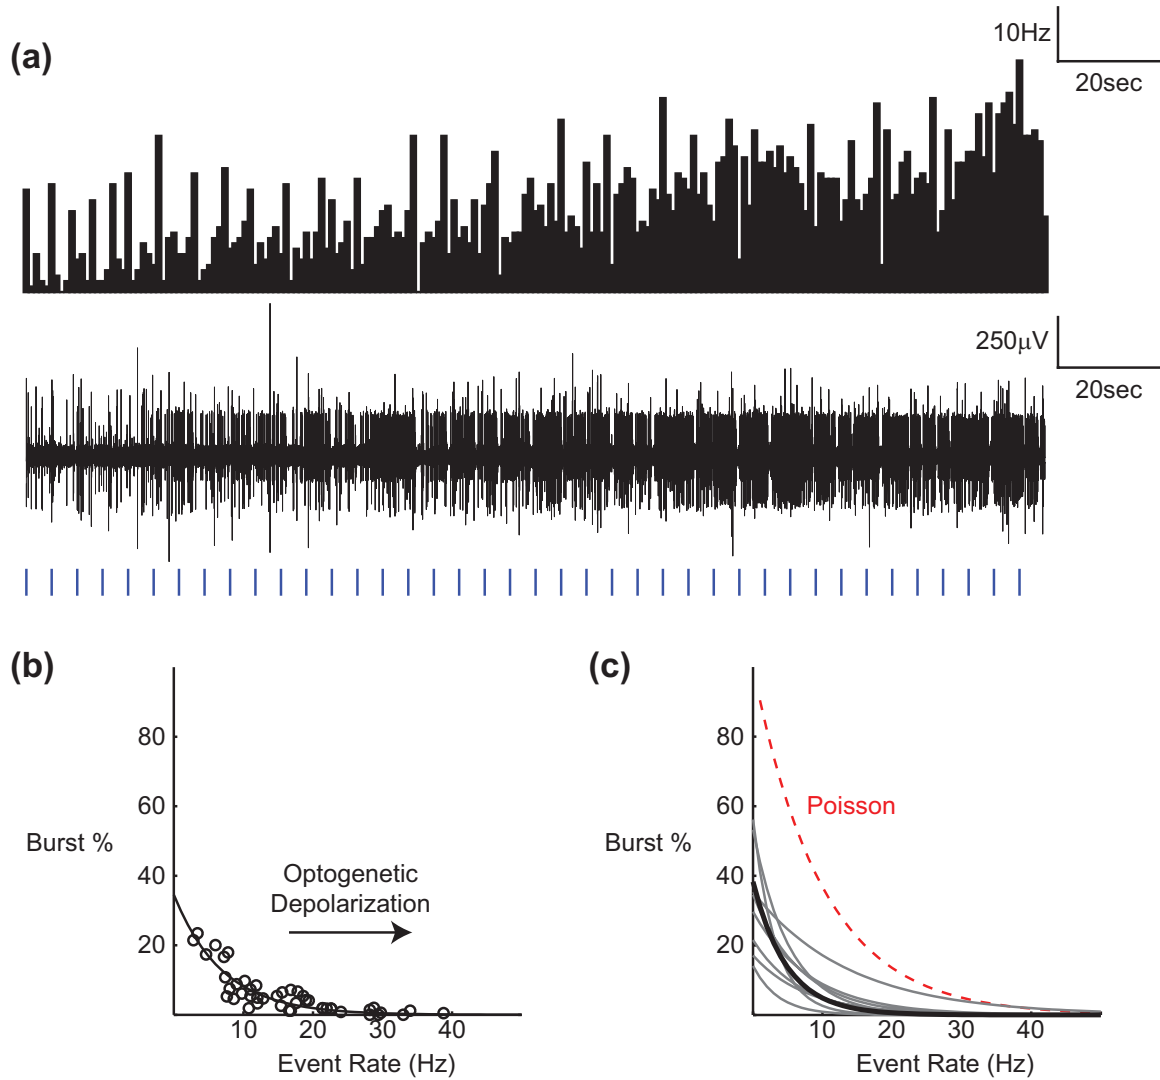


Figure 6.3: Continuous transition from burst to tonic firing modes. (a) An example spike histogram (top) and raw trace (bottom) of a neuron with a SSFO illustrate the transition from burst mode to tonic firing. The long time constant of the SSFO enabled the brief, low intensity light pulses to be perfectly integrated and gradually increase the depolarization of the cell. (b) The burst percentage decreases according to an exponential relationship with the firing rate of the neuron. (c) This trend is consistent across cells.

6.3.3 State control of the cortical response to thalamic microstimulation

Using ChR2 and the SSFO, we controlled the relative depolarization of the thalamus to directly investigate the effect on activity propagation. Voltage sensitive dye imaging (VSDI) was used to record the cortical response to thalamic microstimulation in the baseline anesthetized state (control) and the optogenetically depolarized state. From the work in the previous chapters, it has been shown that thalamic microstimulation induces a strongly nonlinear response in cortex, characterized by high trial-to-trial variability at threshold (Chapter 2) and amplitude-dependent, bimodal nonlinear dynamics for pairs of stimuli (Chapters 3 and 4). The cause of the variability was due to the extreme synchronization of thalamic microstimulation, while evidence suggested that the paired pulse facilitation was caused by thalamic bursting. By shifting the thalamus into a depolarized state, we directly modified the prevalence for bursting and the baseline level of synchronization, and investigated the effect on the propagation of microstimulation-induced neural activity.

6.3.4 Depolarization of the thalamus reduces trial-to-trial variability in the cortical response

First, the input-output relationship was determined by varying the stimulus intensity of a single microstimulation pulse delivered to the thalamus and recording the amplitude of the downstream cortical response on a trial-to-trial basis. By eliminating bursting in the depolarized state, we expected a reduction in the “all-or-none” response in cortex, leading to a lower sensitivity in the nonlinear input-output response curve and lower trial-to-trial variability at threshold. Instead, as shown by the example in figure 6.4(a), the input-output relationship was not changed between the control (black) and depolarized (blue) states. In each case, the cortical response amplitude was a saturating nonlinear function of the input stimulus intensity. There was a slight trend for the depolarized state exhibiting a slightly lower threshold current, but this was not significant across experiments ($N=4$, $p=0.65$), as displayed in figure 6.4(b).

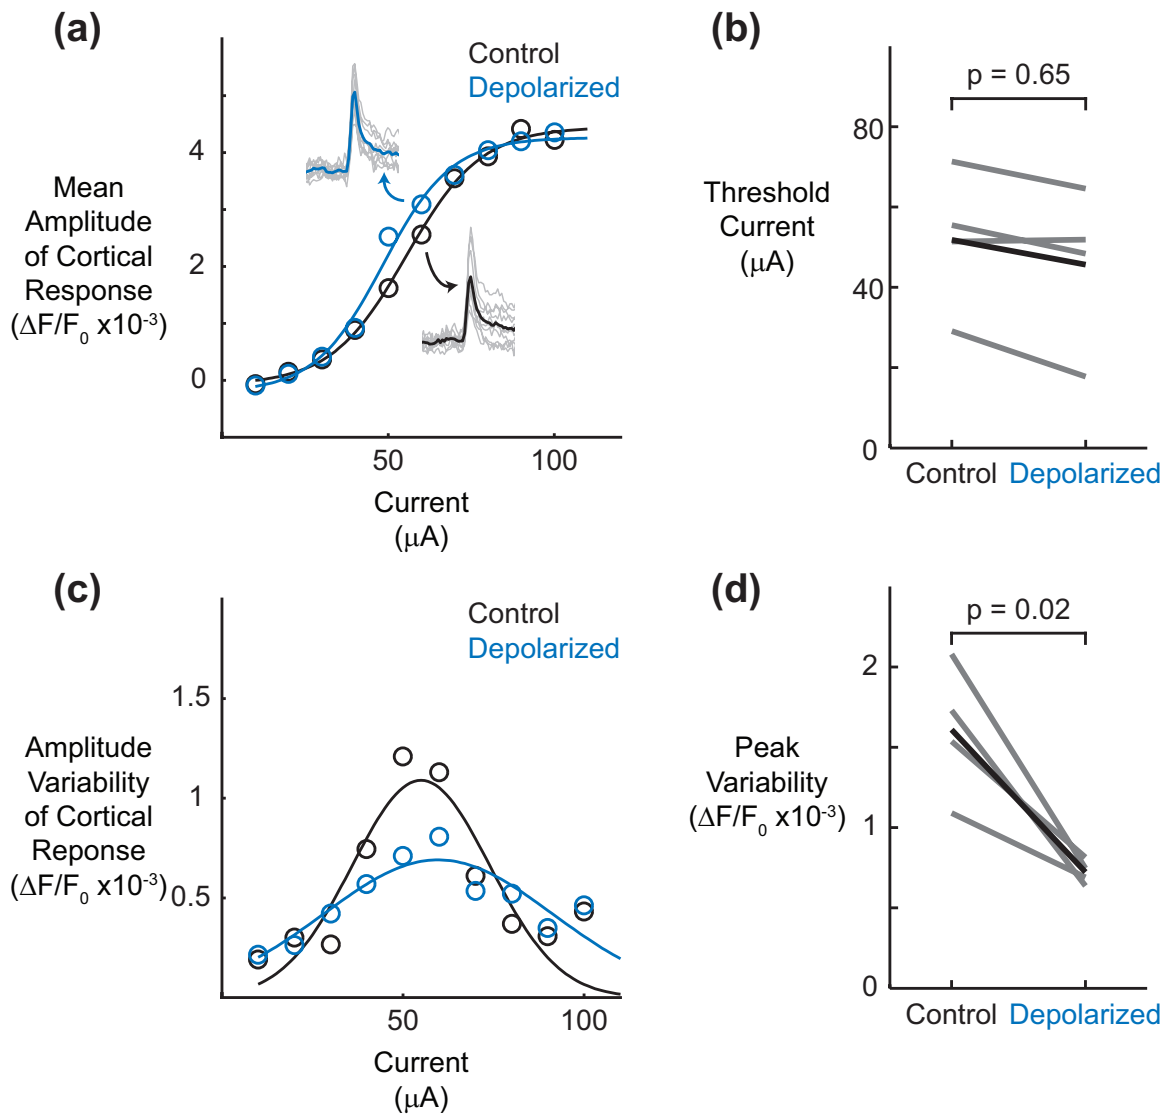


Figure 6.4: Thalamic depolarization reduces cortical trial to trial variability in response to thalamic microstimulation. (a) Example of the mean amplitude of the cortical response, as measured through VSDD, to increasing strength of thalamic microstimulation. The static nonlinearity was unchanged by optogenetically mediated depolarization of the thalamus. (b) The threshold stimulus intensity was not significantly changed between the control and depolarized states. (c) The trial-to-trial variability peaked at the threshold response amplitude, but was decreased in the depolarized condition. (d) The peak variability was significantly reduced in the depolarized condition.

The trial-to-trial variability, however, was changed across the control and depolarized states. In figure 6.4(c), the standard deviation of the response amplitude across trials is plotted as a function of the input stimulus intensity. Consistent with the results of Chapter 2, there was a large increase in variability at the threshold stimulus intensity in the control state (black). In the depolarized state (blue), the variability still peaked at the threshold stimulus intensity, but the overall amount of variability was significantly reduced. This decrease in the peak variability was consistent across animals ($N=4$, $p=0.02$), as shown in figure 6.4(d). Per the results in Chapter 5, the reduction in trial-to-trial variability will impact the design of surrogate sensory inputs, and may improve the overall discriminability of the cortical response amplitudes evoked by thalamic microstimulation.

6.3.5 Depolarization of the thalamus eliminated nonlinear facilitation for thalamic micrstimulation

In addition to modulating the reliability of the cortical response, thalamic state is believed to play a large role in the nonlinear dynamics of the cortical response [170]. Here we explicitly measured the nonlinear cortical response to pairs of thalamic microstimulation in the burst and tonic firing modes. Consistent with previous work (see Chapter 3 and 4), the cortical response to thalamic microstimulation exhibited bimodal nonlinear dynamics (figure 6.5(a)-(c)). In response to simple pairs of stimuli with increasing stimulus intensity, the nonlinear dynamics shifted from facilitation to suppression. For instance, in the example in figure 6.5(a), the cortical response to the second $40\mu A$ stimulus was facilitated relative to the response to the first $40\mu A$ stimulus, whereas the response to the second $100\mu A$ stimulus was strongly suppressed relative to the first. By plotting the response to the second stimulus relative to the response to the first on a single trial basis, as in figure 6.5(b), the evolution of the nonlinear dynamics becomes clear. In the lower left corner of the axes, there were no significant dynamics as the response to each stimulus was negligible. For slightly

higher stimulus intensities, the response to the second stimulus was consistently facilitated across trials, forming a cluster in the upper left portion of the axes. For the highest stimulus intensities, the response to the second stimulus was consistently suppressed across trials, resulting in a cluster in the lower right portion of the axes. These trends were again reliable across experiments. All data across three animals is presented in figure 6.5(c), with the data from each experiment normalized relative to the maximum response amplitude within that experiment.

If the facilitation was caused by thalamic bursting, mediated by T-type calcium channels, then the depolarization associated with shifting the thalamocortical circuit into the desynchronized state should eliminate or reduce the facilitation dynamics. Indeed, this was the case. Figures 6.5(d) and (e) present the matched experiment to panels (a) and (b), but under optogenetic depolarization. In this case, regardless of the stimulus intensity or response amplitude, the response to the second stimulus was suppressed relative to the response to the first. This occurred reliably across trials and experiments, as the overwhelming majority of data points fell below the unity line in figures 6.5(e) and (f). In summary, by manipulating the state of the thalamus, and the associated bursting and synchronization, the nonlinear propagation of neural activity within the thalamocortical circuit was significantly altered, and has implications for the design of surrogate stimuli and the control of neural activity in complex neural circuits.

6.4 Discussion

Using voltage sensitive dye imaging and optogenetic control of the thalamic depolarization, we characterized the propagation of neural activity in response thalamic microstimulation during the tonic and burst firing modes of the thalamus. First, we characterized the firing statistics of the tonic and burst firing modes. Consistent with previous reports, the tonic firing mode consisted of high firing rates and few

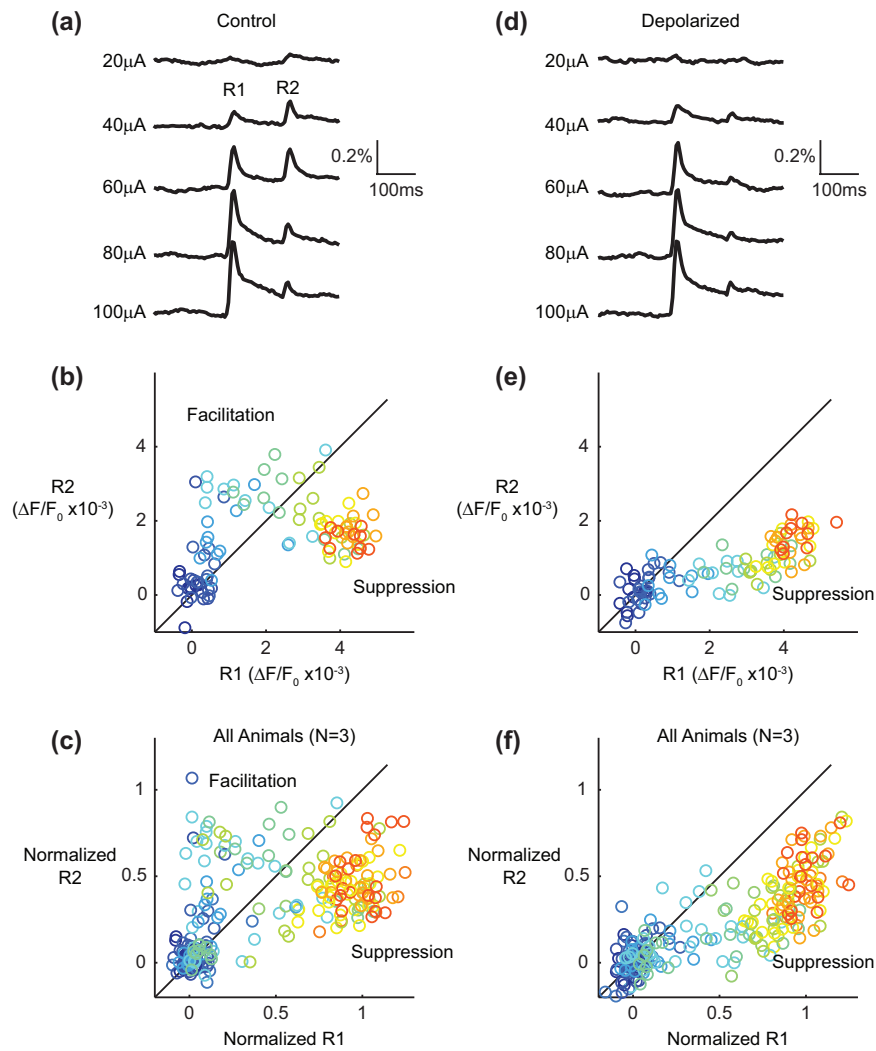


Figure 6.5: Thalamic depolarization eliminates paired-pulse facilitation. (a) Example of the mean cortical response, as measured through VSDI, to pairs of thalamic microstimuli with increasing strength. Facilitation occurs at sub-threshold currents, whereas suppression characterizes supra-threshold currents in the control state. (b) As the current is increased (colors from blue to red), the dynamics transition from facilitation (upper left) to suppression (lower right) for single trial responses. (c) Same as in (b), but with data across three animals. The data are normalized with respect to the maximum response amplitude within that experiment. (d) Example of the mean cortical response, as measured through VSDI, to pairs of thalamic microstimuli with increasing strength. Facilitation occurs at sub-threshold currents, whereas suppression characterizes supra-threshold currents in the control state. (e) As the current is increased (colors from blue to red), the cortical response undergoes paired-pulse suppression for all current intensities. (f) Same as in (e), but with data across three animals. The data are normalized with respect to the maximum response amplitude within that experiment.

bursts, whereas the burst firing mode had a lower overall firing rate, but with a high burst percentage. However, through systematic and precise modulation of thalamic depolarization, we found that the tonic and burst firing modes form two extremes of a continuum, such that the neurons smoothly transitioned from burst firing to tonic firing. Upon demonstrating optogenetic control of thalamocortical state, we characterized the trial-to-trial variability and nonlinear dynamics of the cortical response to thalamic microstimulation in the tonic and burst firing modes. In the burst firing mode (i.e. the control anesthetized state), thalamic microstimulation activated cortex as described in the previous chapters, exhibiting peak trial-to-trial variability at the threshold response amplitude and bimodal nonlinear dynamics with facilitation of sub-threshold inputs and suppression of supra-threshold inputs. In the optogenetically mediated tonic firing mode, however, the peak in trial-to-trial variability was significantly reduced and the facilitation dynamics were eliminated.

Classically, the statistics of neural firing have been extensively explored under naturally occurring burst and tonic states in the anesthetized [164, 108, 107] and awake [54, 161, 43, 160, 159] animal. In these study, dynamic transitions from the burst to tonic firing modes occurred naturally or at the onset of a salient stimulus. This dynamic transition has lead many to believe that the thalamus controls not only “how much”, but also “what type” of information is transmitted to cortex [177, 170]. Specifically, the burst firing mode has been shown to be strongly detectable and highly reliable, whereas the tonic firing mode may be responsible for encoding the more fine details of the stimulus [108, 107].

Prior to optogenetics, the only way to control thalamocortical state was through the activation of the natural neuromodulatory arousal mechanisms in the brain [34, 65, 111]. In this way, the transition from burst to tonic firing could be made quickly and transiently, allowing the study of stimulus-evoked cortical activity on short timescales. However, with the advent of optogenetics, it has become possible to quickly and easily

shift the thalamocortical state bidirectionally. By activating the reticular thalamus, or directly hyperpolarizing the thalamus, burst percentage was increased in the thalamus and spindles were generated in the local field potential in cortex [76]. On the other hand, direct depolarization of the thalamus facilitated the transition from burst to tonic firing rate, producing the classically described “desynchronized” state [159]. Yet, in neither case was the effect on the propagation of stimulus-evoked activity quantified.

In the previous work investigating stimulus-evoked activity under control of state, the cortical response was significantly reduced in amplitude [153], but more reliable across trials [65]. From this, we hypothesized that the tonic firing mode would significantly shift the neurometric tuning curve to the right, while extending the dynamic range and increasing reliability across trials. From the perspective of an ideal observer, these results would be consistent with an improved discriminability of the cortical response. However, we did not observe an appreciable change in the neurometric tuning curve of cortical response amplitude as a function of the thalamic microstimulation stimulus intensity. This may have been due to the extreme synchrony with which thalamic microstimulation is believed to activate neurons locally around the electrode tip (recall Chapter 2). But we did, in fact, find a significant reduction in the peak trial-to-trial variability, consistent with our prior hypothesis. The peak in trial-to-trial variability for the cortical response to thalamic microstimulation was likely due to spontaneous transitions between the up and down states known to occur in the anesthetized [57] or quiescent [153, 61] states. By depolarizing the thalamus, the time spent in the down state was likely significantly reduced, leading to an overall increase in reliability of the cortical response.

Further, we hypothesized that the depolarizing the thalamus and driving the tonic firing mode would eliminate the nonlinear paired pulse facilitation observed previously (recall Chapter 3) for thalamic microstimulation, as evidence suggested that it was

mediated by thalamic bursts. Indeed, we found that thalamic depolarization eliminated the facilitation dynamics for sub-threshold thalamic microstimulation. This is consistent with previous work, as Castro-alamancos showed a reduction in the paired pulse facilitation of the cortical response to thalamic microstimulation during natural transitions of an awake animal between quiescent and active states [33].

Throughout the literature, the burst and tonic firing states are referred to as separate and discrete states (although Mukherjee and Kaplan report a range of burstiness [134]). However, our electrophysiological results in the thalamus suggest that classically described burst and tonic firing states form the extremes of a continuum under the direct control of thalamic depolarization. As the thalamus was gradually depolarized, we observed a smooth, and exponentially decreasing, relationship between burst percentage and event rate. From the perspective of the thalamic neuron, this may not be surprising as the burst percentage should be a function of the relative amount of time spent in the hypopolarized state required to de-inactivate the T-type calcium channels. And yet, the previous literature seems to indicate two predominant states. If tonic and burst firing modes truly do represent two distinct states, the transition between states must be controlled outside of the thalamus, such as in the neuromodulatory centers of the brain stem [34] or forebrain [65], and likely involves the modulation of multiple brain regions simultaneously.

The overall shift in trial-to-trial variability and nonlinear dynamics of the cortical response to thalamic microstimulation has profound impact on the development of sensory prostheses. Peripheral implementations of sensory prostheses, such as the cochlear implant and retinal prosthesis, can use a single encoding rule for mapping sensory stimuli to patterns of artificial stimulation as fluctuations of state do not affect these peripheral structures strongly. A central nervous system prosthesis, however, would require a dynamic modulation of the encoding rule, as the state of the neural

circuit would continuously change based on the arousal of the individual and significantly alter the propagation of information to downstream structures. This suggests the use of online monitoring of brain state [22] in order to switch between encoding rules. While certainly more complicated, dynamically switching the encoding rule based upon the underlying brain state may allow the delivery of contextual surrogate inputs, such as specifically designed “detection” or “discrimination” cues. More generally, the results described here are the most complete description of stimulus-evoked activity under the control of thalamocortical state and inform the understanding of neural circuit information processing, with future work addressing the functional relevance of state transitions for sensory processing of naturalistic sensory stimuli in the awake animal.

CHAPTER VII

CONCLUSION AND FUTURE WORK

The goal of this project was to characterize, model, and control the neural circuit response to artificial stimuli and inform the design of surrogate inputs for sensory prosthesis applications. Our approach is summarized in figure 7.1. For patterns of sensory stimulation, the dynamics of the neural circuit, $H(x)$, map these inputs to spatiotemporal patterns of activity in cortex. However, given that previous work has suggested that sensory and artificial inputs drive neural circuits with distinct nonlinear properties [112, 119], we first characterized the difference between the neuronal dynamics in response to sensory, $H(x)$, and artificial stimuli, $H^*(x)$.

In Chapter 2, the trial-to-trial variability and spatial spread of the cortical response to whisker, electrical, and optogenetic inputs were quantified using voltage sensitive dye imaging (VSDI). Importantly, both the trial-to-trial variability and spatial spread are crucial variables for quantifying the discriminability of information encoded by the downstream neural response. All stimuli produced monotonically increasing response amplitudes as the stimulus intensity increased. However, the response properties diverged from here. We found that the trial-to-trial variability of the cortical response to whisker stimuli increased linearly with the stimulus intensity and response amplitude, such that strong responses were highly variable across trials. The response to strong electrical and optogenetic inputs, on the other hand, was highly reliable across trials. Instead, the peak in the trial-to-trial variability occurred at the threshold stimulus intensity and response amplitude. Whereas the trend in variability was similar for the two types of artificial stimuli, the spatial spread was more similar between whisker and optogenetic inputs, which each produced focal

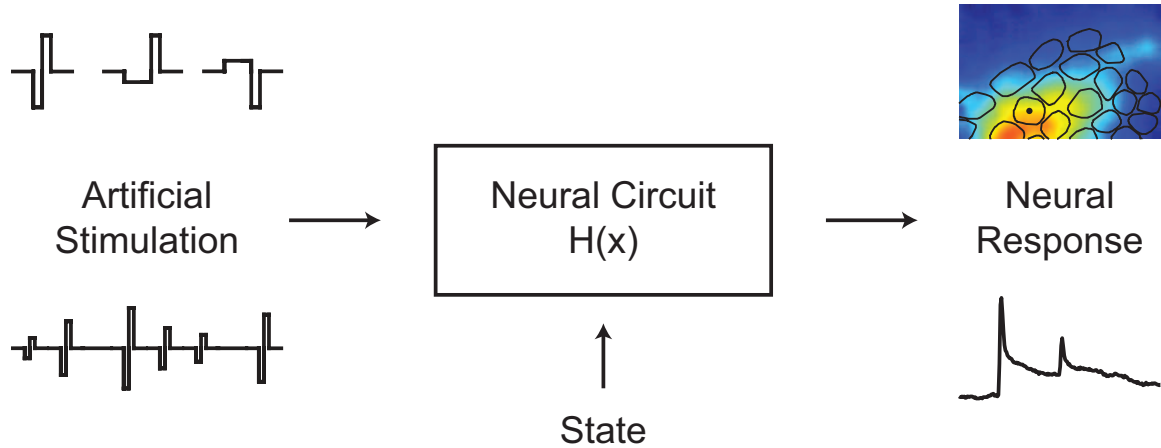


Figure 7.1: Nonlinear propagation of activity across neural circuits. Diagram illustrating the goals and achievements of the thesis. First, the neural circuit response to microstimulation was characterized and modeled. Then, stimulus design was employed to overcome the distinct neural response properties to thalamic microstimulation, as compared to natural sensory stimuli. Finally, the role of state in governing the nonlinear dynamics was explored through optogenetic control.

spatial activation in cortex. Electrical stimulation produced a comparatively large activation in cortex, with the spread increasing for larger response amplitudes.

In this way, the three stimulus types each differentially activated the neural circuit. Informed by previous literature, we sought to determine the mechanism by which the different stimulus types engaged the neural circuit. Specifically, a large body of literature suggests that the synchrony of population neural activity is dynamically modulated during the processing and encoding of sensory stimuli [24, 177], whereas electrical and optogenetic stimuli are believed to strongly synchronize the neurons that they activate and merely recruit more neurons as the electric field or light spreads through the tissue at higher stimulus intensities. Thus, we used a computational network model of the thalamocortical circuit to allow direct manipulation of the magnitude and synchrony of population neural activity, two important variables that were not independently accessible experimentally. Through model simulations, we found that a fixed, high synchrony in the population activity reproduced the peak trial-to-trial variability observed with electrical and optogenetic inputs, but variations

in synchrony for a fixed magnitude of population activity mimicked the trends observed with whisker stimuli. Spatially, the extreme synchrony and mixed stimulation of axons and cell bodies produced increasing spread in the cortical response, similar to the observations for electrical stimulation of the thalamus. Through the combined experimental and computational approach, we demonstrated, using two simple case studies, that the neural response to artificial inputs is distinct from the neural response to natural sensory inputs, and further hypothesize that the mechanism is the extremely high, and potentially pathological, synchronization of population activity induced by electrical and optogenetic stimuli.

The results in Chapter 2 described the nonlinear response to single stimuli of increasing intensity, whereas Chapters 3 and 4 extend the characterization of the neural response to temporal patterns of sensory and artificial stimuli. First, in Chapter 3, a simple “paired pulse” paradigm was used to explore the nonlinear dynamics of the cortical response, and later in Chapter 4 more complicated patterns were used. In the simple paired pulse paradigm, the response to a second stimulus is compared with the response to the first. Under conditions of linearity, the response to the second stimulus could be predicted through the superposition of the response to the first stimulus and a temporally shifted copy. With nonlinear dynamics, however, the response to the second stimulus would be facilitated or suppressed relative to the linear expectation from the response to the first stimulus. Classically, the cortical response to sensory stimuli has been characterized as nonlinear suppression in the paired pulse paradigm. Using VSDI to record the spatiotemporal cortical response, we found that whisker stimuli elicited paired pulse suppression in cortex, with the duration of the suppression increasing for higher velocity whisker deflections. For electrical stimulation, similar paired pulse suppression was observed for high current intensities. But at sub-threshold stimulus intensities, a profound facilitation of the response to the second stimulus occurred. Additionally, the timescale of the facilitation was distinct

from the suppression observed with whisker stimuli and supra-threshold electrical stimuli, as facilitation only occurred for inter-stimulus intervals between 100 and 200 milliseconds. Finally, the spatial spread associated with the facilitation dynamics was dramatically sharpened relative to the response to a single supra-threshold electrical stimulus. This evidence suggests that electrical stimulation of the thalamus recruits a dynamical mode that is distinct from sensory stimuli, and results in amplitude-dependent, bimodal nonlinear dynamics in the cortical response.

To investigate the mechanism of the facilitation, we used optogenetic stimuli to allow simultaneous recording of single unit activity in the thalamus and the spatiotemporal cortical response. Optogenetic stimuli delivered to the thalamus produced a similar bimodal nonlinear response in cortex to electrical stimuli, with suppression occurring for supra-threshold inputs and facilitation for sub-threshold inputs. In response to the same pairs of optogenetic stimuli, the response of single units in the thalamus exhibited paired pulse facilitation at all stimulus intensities. From this, we hypothesize that the facilitation was a thalamic phenomenon, mediated by post-inhibitory calcium bursts, and that the suppression was a cortical phenomenon caused by the engagement of feedforward inhibition.

In Chapter 4, we performed nonlinear system identification on the cortical response to temporal patterns of stimulation to further model the bimodal nonlinear dynamics, $H^*(x)$. The simple paired stimulus paradigm used in Chapter 3 was only able to probe second order nonlinear dynamics. In order to probe higher order nonlinearities, and more efficiently identify the nonlinear dynamics, we used a random amplitude Poisson (RAP) stimulus train [208], where the stimulus times were drawn from a homogeneous Poisson process, and the stimulus intensities varied uniformly. After training on stimulus-response data, a traditional Volterra series black box model performed poorly in predicting the neural response to an out-of-sample stimulus-response

pair, indicating a failure to accurately capture the system dynamics. We then developed a phenomenological model in order to specifically include certain nonlinear features of the cortical response that we had observed experimentally in Chapters 2 and 3. By specifically including two dynamical stages in series, the phenomenological model had significantly improved predictive capability as compared to the Volterra model. Simulations with the phenomenological model further supported the previous conclusions in Chapters 2 and 3, suggesting that thalamic microstimulation of increasing intensity recruits a linearly increasing number of thalamic neurons, but that these neurons project in a highly nonlinear manner to downstream cortex through the extreme synchrony of the population activity.

Ultimately, through the combined experimental and modeling approaches, Chapters 2, 3, and 4 provide an extensive characterization of the nonlinear neural circuit dynamics, $H^*(x)$, in response to thalamic microstimulation and how this deviates from the neural circuit response to natural sensory stimuli. Two crucial differences in the neural response were the trial-to-trial variability and the spatial spread of the cortical response. In Chapter 5, we specifically designed thalamic microstimuli to overcome these differences in circuit engagement. The previous modeling work in Chapter 2 suggested that the increased spatial spread observed in cortex was due to thalamic microstimulation activating axons in addition to cell bodies, consistent with previous evidence in the literature [83]. Using a biophysical model, McIntyre and Grill hypothesized that variations in waveform shape could shift the relative recruitment of axons vs. cell bodies, leading electrical stimuli to activate neural circuits more specifically [123]. We directly tested this prediction by delivering symmetric and asymmetric, cathode-leading biphasic electrical stimuli to the thalamus and recording the spatial response in downstream cortex. We found that symmetric stimuli produced a topographically misaligned cortical response, consistent with the preferential stimulation

of axons in the thalamus. Using asymmetric stimuli, the topographic alignment between the cortical response and the electrophysiologically identified location of the thalamic electrode was significantly improved, but at the expense of requiring greater current intensity to achieve the same neural response amplitude.

By specifically activating cell bodies, as opposed to axons, electrical stimulation surrogate inputs more faithfully adhered to the underlying topography of the early sensory pathways, thus improving the spatial discriminability of an ideal observer of the cortical response. Information could also be encoded in the amplitude of the cortical response, where the trial-to-trial variability, as a function of the response amplitude, would determine the performance of an ideal observer. The distinct trends in trial-to-trial variability for sensory and artificial stimuli made the design of surrogate sensory inputs non-trivial. In the latter portion of Chapter 5, we performed optimal signal set design to maximize the discriminability of the cortical response amplitude as a function of the stimulus intensity. Optimal signal set design is used in the communications literature to maximize information throughput of a noisy channel by choosing the signal set that best separates the output distributions of the channel [103]. Using a theoretical model of the cortical response to sensory and artificial stimuli based on the experimental data in Chapter 2, we optimized the signal set (i.e. response amplitudes) by maximizing an information theoretic distance between nearest neighbor pairs of response distributions. The optimal signal sets for sensory stimuli and artificial stimuli had equivalent average performance, suggesting that knowledge of the nonlinear mappings can be used to overcome the distinct trends in the cortical response. However, while the average performance was equivalent, portions of the signal set were unused in the theoretical model of artificial stimuli, such that less overall information was delivered to downstream cortex. Under a different optimization scheme, where we maximized the minimum performance across all members of the signal set [103], all signals within the set were used across each model. In this

case, the performance for the model of artificial stimuli had higher performance for a small signal set, but larger signal sets were decoded with greater performance under the model of sensory stimuli. Under the assumptions of the theoretical framework, this suggested that artificial stimuli provide highly detectable signals, but an ideal observer of the downstream cortical activity would not be able to discriminate amongst varying stimulus intensities. More generally, the results motivate the extension of optimal signal set design to spatiotemporal patterns of surrogate inputs.

Through the design of the input stimuli, we exploited the properties of the neural circuit to improve the discriminability of an ideal observer of the downstream neural response. However, recent work [65], and longstanding hypotheses in the field [42], suggest that neural circuits in the brain already dynamically switch states to optimize not only “how much”, but also “what type” of information is transmitted to downstream structures [177, 170], allowing the transmission of both “detection” and “discrimination” signals. In Chapter 7, we systematically and precisely controlled the state of the thalamocortical circuit and characterized the propagation of artificially-evoked neural activity. Optogenetic depolarization of the thalamus pushed the thalamus from a burst firing mode to a tonic firing mode, which is classically associated with the “desynchronized” cortical state. Moreover, through precise modulation of the thalamic depolarization, we showed a smooth transition from burst to tonic firing, suggesting a continuum rather than discrete states. Shifting to the depolarized state significantly modified the nonlinear propagation of activity induced by thalamic microstimulation. In the depolarized state, the peak in trial-to-trial variability described in Chapter 2 was significantly decreased and the facilitation portion of the bimodal nonlinear dynamics was suppressed. In this way, the depolarized state made the thalamocortical dynamics elicited by thalamic microstimulation, $H^*(x)$, more similar to those produced by natural sensory stimuli, $H(x)$.

Together, the experimental and computational approach described above provides

a complete description of the nonlinear dynamics of the thalamocortical circuit to artificial stimuli. More specifically, the results highlight the distinct neural response properties for sensory and artificial stimuli and identify the extreme synchrony of artificially generated neural responses as the primary mechanism of the differences. Stimulus design was used to overcome the specific differences in the evoked response to sensory and artificial stimuli and improve the functional relevance of the downstream surrogate sensory signals. Ultimately, the characterization, modeling, stimulus design, and overall control of downstream neural activity documented in this thesis closely mimics the process undertaken in the early, and continuing, development of the cochlear implant [202], and stands to inform the development of sensory prostheses with intervention sites in the central nervous system.

7.1 Application to sensory prostheses

While the cochlear implant has enjoyed great success, other sensory prostheses are significantly less well developed. The retinal prosthesis is in clinical trials and has restored limited vision to those that are otherwise blind [87], albeit not at the same level of performance as the cochlear implant. Beyond the retinal prosthesis, few other sensory prosthesis applications are beyond the concept stage.

A major reason for the relative success of the cochlear implant is the anatomy of the neural interface. Within the cochlea, the auditory nerve interfaces with the hair cells along a linear track, with high frequencies encoded near base and low frequencies mapped to the apex. The one-dimensional topography along the coil of the cochlea enables a relatively simple one-dimensional electrode interface and bio-mimetic encoding model, where sounds are separated into discrete frequency bands and encoded by individual electrodes along the length of the array [201]. By comparison, the topographic alignment of the retina is two-dimensional. In this way, an extra dimension

is needed for the electrode array, but it can still be designed in a planar array. Interfacing directly with the somatosensory periphery, however, is substantially more difficult. Distal from the spinal cord, information from the sensory afferents travels along a complicated network of nerves, each containing multi-modal information (touch, pressure, temperature, pain, motor output, etc.). Thus, intervention in the somatosensory periphery would require a complicated three-dimensional neural interface, spanning a large portion of the human body. Further, the peripheral interface must be designed to be robust to movement of the limbs/body that have retained, or regained, function and movement.

For this reason and others, a significant portion of sensorimotor prosthesis development has been focused in the central nervous system. In this way, the sensory feedback neural interface can be co-located with the motor decoding neural interface, creating a combined sensorimotor prosthesis [199]. However, the neural interface would still need to be at least two-dimensional for a layered structure such as the cortex, and three-dimensional in a globular region such as the somatosensory thalamus. In addition to the increased dimensionality of the interface, the work in this thesis, along with previous work in the field, emphasizes the nonlinearities associated with interfacing in the central nervous system. All encoding models currently in use for the cochlear implant utilize linear dynamics, where the power in a particular frequency band is used as the envelope to modulate the stimulus intensity on a given electrode. From our results, however, the response of neural circuits in the central nervous system would be highly nonlinear for temporal patterns of activity.

Our results specifically point toward the extreme synchrony generated by artificial stimulation as accentuating the nonlinear properties of central nervous system neural circuits. Whether the goal of the sensory prosthesis is to mimic natural sensory-evoked neural responses, or merely to produce maximally discriminable neural responses, the distinct ways in which artificial stimuli activate neural circuits must be included

within the encoding model.

In general, two opportunities exist for the optimal design of the encoding model: 1) the base unit of stimulation and 2) the encoding rule mapping sensory signals to patterns of the base stimulation unit. In this work, we performed very simple design of the base unit of stimulation to overcome the limitations of the cortical response to thalamic microstimulation; however, a much larger stimulus design space remains as an open field of study. In particular, the design of a base unit of stimulation aimed at reducing the stimulus-induced synchronization may simplify the encoding rule by better mimicking natural sensory inputs. High frequency, multi-electrode stimulation has been shown to desynchronize a neural population *in vitro* [191]. Using a high frequency pulse train of increasing stimulus intensity may provide sufficiently strong neural activation, but with reduced synchronization. Further, the rate of increase in the stimulus intensity could be used to titrate the stimulus-induced synchronization. Also, if optogenetics becomes a clinically viable solution, the stimulus design space for optogenetic stimuli is far less restricted, since there is no analogous limitation to charge balancing electrical stimuli. The optogenetic stimulus design problem would additionally extend to the design of the kinetics of the light-sensitive channel.

In addition to the design of the base unit of stimulation, the encoding rule is responsible for mapping the sensory information into a pattern of stimulation to produce functionally relevant surrogate information in downstream neural structures. In the central nervous system, the encoding rule will likely need to include knowledge of the nonlinear response of the target neural circuit and the constantly changing state. However, additional opportunities may exist to exploit beneficial properties of the nonlinear dynamics or even directly control the state. For instance, the spatial sharpening of the facilitation dynamics described in Chapter 3 illustrates a simple example of how the dynamics may be exploited to improve information transmission. If the facilitation dynamics could be engaged by a subthreshold “noise” stimulus

provided through the electrode array, then the propagation dynamics could be shifted dynamically without triggering a confounding sensation. Similar use of “noise” stimuli have been used in sensory neurophysiology to achieve modulation of the downstream cortical response [23]. Finally, through modulation/control of the underlying state, it may be possible to multiplex, or at least dynamically switch between, “detection” and “discrimination” information by modulating neural circuit properties.

7.2 Applications towards microstimulation for other disorders

While the cochlear implant is the only currently successful sensory prosthesis, electrical stimulation is used clinically for a variety of other neuromodulation therapies. The goal of any neuromodulatory therapy is to replace an aberrant pattern of activity with a physiological pattern, or at very least a synthetic pattern that eliminates the pathological symptoms. The most well-known example is the use of deep brain stimulation (DBS) to treat motor [49] and psychological [84] disorders.

Deep brain stimulators consist of a large macroelectrode inserted deep into the brain and a stimulator that delivers high frequency ($\sim 130\text{Hz}$) stimulation across the electrode. Recent work supports the hypothesis that DBS entrains the neural population to the stimulation rate, with each stimulus pulse synchronizing the population of neurons [124, 67]. However, through instilling this synthetic pattern, the stimulation “regularizes” local neural activity, allowing normal motor control information to propagate through the pathway [53].

The design of stimulation patterns for DBS is an active area of research, with irregular patterns of stimulation potentially providing added benefits of greater symptomatic relief or reduced power consumption [20]. Further improvement may be obtained through global state modification of the pathological circuit. Physiological patterns of neural activity in the basal ganglia are characterized by desynchronized,

tonic firing [80], similar to the tonic firing mode of the somatosensory thalamus described in Chapter 6. Extrapolating from the results in this thesis, chronic depolarization of the basal ganglia may achieve the desynchronization necessary to restore normal motor function. Thus, multichannel high frequency microstimulation (within existing power constraints), or chronic depolarization through optogenetic control [212], may enhance the therapeutic effect.

The nonlinear engagement of central nervous system circuits described in this thesis may also have applications for neuromodulatory therapies. For instance, electrical stimulation of the thalamus produced profound suppression over long timescales. Other studies, in concert with the work in this thesis, suggest that microstimulation strongly engages feedforward inhibitory circuitry, which is a common motif across neural circuits. A strategy for treating disorders of hyperactivity, such as epilepsy, could be to engage the pathological circuit upstream of the aberrant activity and utilize the anatomy to recruit long timescale inhibition through electrical stimulation [97, 163, 187].

7.3 Interfacing with neural circuits: a framework for control

More generally, applications of artificial stimulation for sensory prostheses and neuromodulatory therapies rely on the ability to precisely control neural activity. This thesis has performed the initial requisite characterization of the neural circuit dynamics to facilitate the generation of arbitrary patterns of neural activity through artificial stimulation. Further, optogenetic techniques were used to control the underlying state of the neural circuit under investigation. However, future directions of this work should rely on the robust and quantitative methods of modern control theory for controlling neural activity within the brain.

The field of neural control has been intensely investigated since the 1980s, where

experimentalists and theorists alike have used a control theory framework for exploring the control of the human body by the brain. A prime example is the study of motor control [21], or how the brain activates distributed muscle groups in a coordinated fashion to achieve complex movements of the arms, legs, or eyes. The physical nature of the state variables, such as the position of the eyes or the velocity of the arm, have aided the adoption of classical and modern control theory to study the neural control of movement.

On the contrary, the principled control of neural activity on fast timescales is far less well developed. With the advent of optogenetics, it is now possible to bidirectionally control neural activity on fast timescales within the brain. Similarly, the maturity of electrical stimulation techniques has ushered in a new era of clinical neuromodulation. Given these actuators, how can we design a framework for the control of neural activity in the brain? What are the state variables we have access to? Here, we propose three levels of abstraction for controlling neural activity in the brain: 1) control of neuronal variables, 2) control of information content, and 3) control of operating modes.

The most obvious and fundamental variable of population neural activity is the mean firing rate. At various stages within the brain, neuronal firing rates correlate with motor activity [63], sensory information [86], and decision making [158]. Closed loop control of population firing rates have previously been demonstrated in vitro using multichannel microstimulation [191], and more recently through the use of optogenetics in vitro and in vivo [138]. By incorporating mathematical descriptions of neuronal dynamics, such as those identified in this thesis, into a closed loop controller, it may be possible to precisely control dynamic trajectories of population firing rate with defined constraints on the response properties such as overshoot, rise time, etc. Outside of population firing rate, higher order statistical properties could be controlled within a closed loop paradigm. For instance, population synchrony would be

a natural variable of population activity to control for experimental or therapeutic applications.

More abstractly, closed loop control of neural activity could be also used to maximize information content in population activity. For instance, in Chapter 5, the trial-to-trial variability largely determined the discriminability between cortical response amplitudes within the optimal signal set design. Theoretically, closed loop proportional and integral control could be used to significantly reduce steady state error. Reducing the trial-to-trial variability through closed loop control of neural response trajectories would further facilitate optimal signal set design, increasing the amount of information that could be delivered to downstream neural circuits.

Finally, the underlying state of a neural circuit offers another variable for control. In this work, we demonstrated simple open loop control of state by optogenetically depolarizing the thalamus. However, global control of state would require a more direct measurement of state to place under closed loop control. The goal of most neuromodulatory therapies is to control the state of a neural circuit, and thus a principled framework for closed loop control would have significant clinical impact. And experimentally, the closed loop control of neural circuits would allow the systematic interrogation of the various information processing states a neural circuit can assume.

In conclusion, the characterization, modeling, and design of artificial stimulation techniques presented in this thesis directly informs the development of central nervous system sensory prostheses. More generally, the results highlight the distinct ways in which sensory and artificial stimuli activate neural circuits, mediated by the differential synchronization of neural populations. Future directions of this work will incorporate knowledge of the nonlinear dynamics of neural circuits for the closed loop control of functionally relevant activity within the brain.

APPENDIX A

GENERAL VSD METHODS

A.1 Surgery and preparation

Fourteen female adult albino rats (220-350g; Sprague Dawley, Charles River Laboratories, Wilmington, MA) were used in the study. All procedures were approved by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology, and were in agreement with guidelines established by the National Institutes of Health. Briefly, female albino rats were sedated with 2% vaporized isoflurane and anesthetized with sodium pentobarbital (50 mg/kg, i.p., initial dose); supplemental doses were given as needed to maintain a surgical level of anesthesia, confirmed by measurements of heart rate, respiration and eyelid/pedal reflexes to aversive stimuli (toe or tail pinch). In all experiments, body temperature was maintained at 37 °C by a servo-controlled heating blanket (FHC, Bowdoinham, ME). After initial anesthesia, the animal was mounted on a stereotactic device (Kopf Instruments, Tujunga, CA) in preparation for the surgery and subsequent recordings. Atropine (0.5 mg/kg, s.c.) was injected, and Lidocaine was applied to the tissue on top of the head. After the initial midline incision on the head, tissue and skin were resected, and connective tissue was carefully removed. A craniotomy ($\sim 3 \times 4\text{mm}$) was made on the left hemisphere over the barrel cortex (stereotactic coordinates: 1.0-4.0 mm caudal to the bregma, and 3.5-7.0 mm lateral to the midline) and over the ventroposterior medial nucleus (VPm) of the thalamus (2.0-4.0 mm caudal, 2.0-3.5mm lateral to the midline [149]). A dam was constructed with dental acrylic around the craniotomy over the barrel cortex to contain the voltage sensitive dye solution (RH1691, 1.5mg/ml, Optical Imaging Ltd, Rehovot, Israel) for staining. Mineral oil was periodically applied to the cortical

surface over VPm to keep the brain moist. After the recording session, the animal was sacrificed with an overdose of sodium pentobarbital. A subset of the animals were transcardially perfused with 4% paraformaldehyde, and their brains removed for histologically producing the anatomical barrel cortex map.

A.2 VSD imaging

Voltage sensitive dye imaging was achieved by using a high speed, low noise camera coupled with a tandem lens (MiCAM 2, SciMedia, Tokyo, Japan). After the craniotomy over the barrel cortex, the dura mater was allowed to dry for 15 minutes [109]. The cortex was stained with dye RH1691 solution (1.5mg/mL; Optical Imaging) for two hours, during which the dye solution was circulated every 5 minutes to prevent the cerebral spinal fluid from impeding the staining. After staining, saline was applied generously to wash off the dye residue. Then the dam was filled with saline and a glass cover slide was placed on top of the dam to prevent the saline from vaporizing. The dye was excited by a 150W halogen lamp filtered to pass wavelengths only in the 615-645nm band. In all experiments, a 1.0x magnification lens was used as the objective lens in conjunction with a 0.63x condenser lens to provide 1.6x magnification, forming a tandem lens as shown in Figure A.1(a). Twenty trials of VSD data were collected for each stimulus and they were averaged offline for the data analysis (see response analysis section below).

A.3 Electrophysiological recordings

Extracellular recordings in the VPm were obtained by using single tungsten microelectrodes ($\sim 1\text{M}\Omega$, 75 μm in diameter, FHC, Bowdoinham, ME). The detailed procedure was described previously [196]. Briefly, after the craniotomy, a tungsten microelectrode was slowly advanced into VPm using a hydraulic micropositioner (Kopf Instruments, Tujunga, CA). During electrode advancement through VPm, individual whiskers were stimulated manually to identify the principal whisker (PW), i.e. the

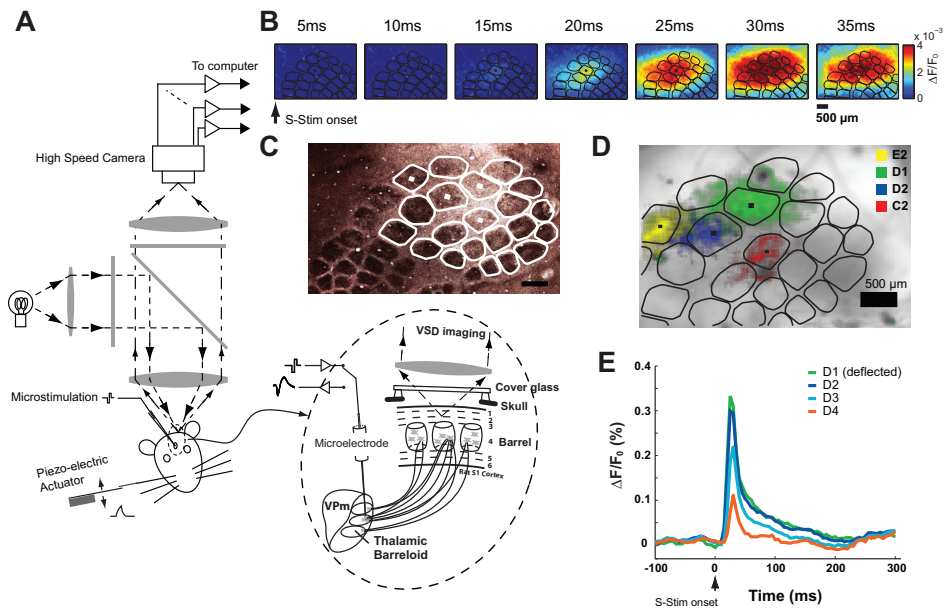


Figure A.1: Voltage Sensitive dye imaging of cortical activation in response to thalamic activation and whisker deflection. (a) Diagram of optical imaging setup. (b) Example of cortical response to a punctate deflection of whisker D2 (sensory stimulation, S-Stim). (c) Example of barrel columns shown in a CO stained brain slice. White traces are contours of the barrel columns shown in (d). (d) Barrel mapping by registering the histologically identified barrel map with functional columns. (e) Mean activation within nearby cortical columns D1, D2, D3, and D4 following the punctate deflection of the D1 whisker. Scale bars in (b), (c), and (d) are $500\mu\text{m}$.

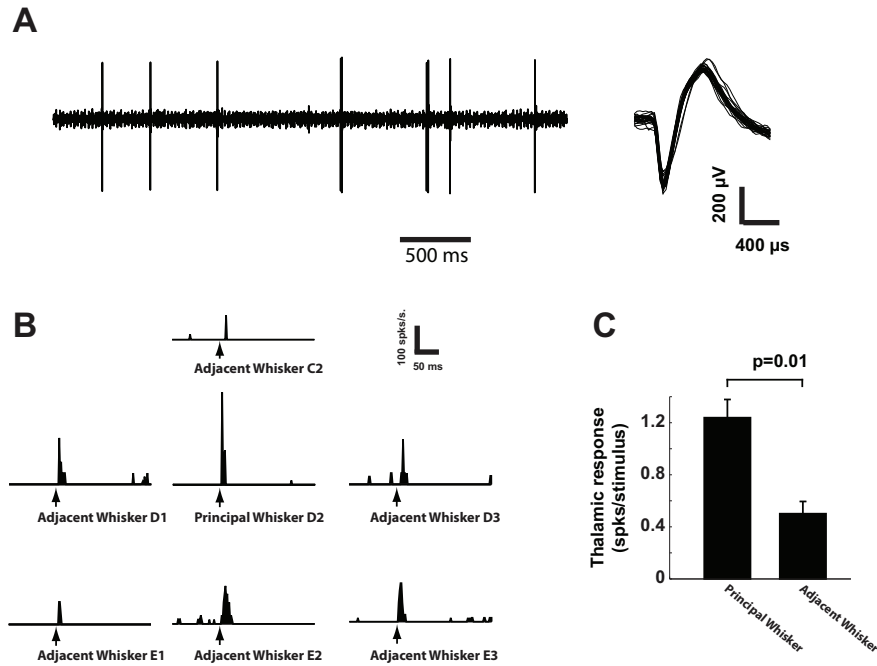


Figure A.2: Functional validation of thalamic electrode placement. (a) A typical example of single-unit activity recorded in VPM (left) and spike waveforms (right). (b) The peri-stimulus time histograms (PSTHs) of a VPM cell in response to punctate deflection of whisker C2, D1, D2, D3, E1, E2, and E3 (binsize = 2ms). (c) Selection of principal whisker (PW) was confirmed by the fact that PW evoked stronger spiking activity than adjacent whiskers (AW) ($p=0.01$, Mann-Whitney U-test).

whisker that evokes the strongest response. In a subset of experiments, the single-unit responses to the principal whisker and the adjacent whiskers were recorded, a typical example of which is shown in Figure A.2. We aimed to recruit barrel fields in the E and D rows, which are located close to the center of imaging field in our preparation. In 14 experiments, the D1 barreloid was accessed 6 times, the D2 barreloid was accessed once, the D3 barreloid was accessed 5 times, the E1 barreloid was accessed once, and the E3 barreloid was accessed once. Neuronal signals were amplified, band-pass filtered (500-5kHz), digitized at 30 kHz/channel, and collected using a 96-channel data-acquisition system (Blackrock Microsystems, Salt Lake city, UT, USA).

A.4 Thalamic microstimulation

After the principal whisker was identified, the thalamic electrode was used to deliver single electrical current pulses to evoke cortical responses in the somatosensory pathway. The electrical stimuli were created using a digital stimulus generator (Model: DS8000, WPI Inc., Sarasota, Florida) and delivered using a digital linear stimulus isolator (Model: DLS 100, WPI Inc., Sarasota, Florida) acting in current source mode. Additionally, a fast switching relay was used to prevent charge accumulation on the electrode tip. All individual electrical stimuli were charge balanced. Three type of stimulus waveforms were used in this study: 1) a cathode-leading, symmetric biphasic waveform of 200 microseconds duration per phase (Symm), 2) a cathode-leading, asymmetric biphasic waveform with 1 millisecond of cathodal duration and 200 microseconds of anodal duration (ASymC), and 3) an anode-leading, asymmetric biphasic waveform with 1 millisecond of anodal duration and 200 microseconds of cathodal duration (ASymA). In Chapter 5, the asymmetry of the cathode-leading asymmetric, i.e. ASymC, waveform was systematically changed to more precisely determine the effect of the asymmetry on the specificity of electrical stimulation. Asymmetry was defined as the ratio of the duration of the first phase to the duration of the second phase minus one, such that the Symm waveform had an asymmetry of zero and the ASymC waveform had an asymmetry of four. Waveforms were delivered over a range of current amplitudes (30, 35, 40, 45, 50, 60, 70, 80, 90, 100, 120, 150 microamps). In each experiment, eight amplitudes were randomly chosen.

A.5 Whisker stimulation

Sensory stimulation (S-Stim) was applied through computer controlled whisker deflections. Whiskers were trimmed at approximately 12mm from the face, and were inserted into a glass pipette fixed to the end of a calibrated multi-layered piezoelectric bimorph bending actuator (range of motion, 1 mm; bandwidth, 200 Hz; Physik

Instrumente (PI), Auburn, MA) positioned 10 mm from the vibrissa pad. Vibrissae were always deflected in the rostral-caudal plane. Punctate deflections consisted of exponential rising and falling phases (99% rise time, 5 ms; 99% fall time, 5 ms) and angular deflection velocities of 75, 150, 225, 300, 450, 600, 900, and 1200 deg/s were used as mechanical probe stimuli (S-Stim).

A.6 Histology

Methods were adapted from Wong-Riley (1979)[205], with modifications adapted from Silverman and Tootell (1987) [172]. Briefly, the cortex was flattened to 1mm and fixed in 4% paraformaldehyde for 12 hours. Tangential sections of 70um were washed for 10 minutes in 10% sucrose in 0.1M phosphate buffer, 10 minutes in 10% sucrose in 0.05M Tris with 275mg/L cobalt chloride, then 3 times in PBS before being incubated at 37 for 5 hours in the staining solution (0.5g/L DAB, 50g/L sucrose and 75mg/L cytochrome c in 0.1M PB).

A.7 Barrel mapping

The VSD data collected in response to whisker stimulation was functionally registered to the histological map of the barrel cortex [194]. The contours of the barrel cortex columns, shown in Figure A.1(c), were outlined using the NeuroLucida software (MBF Bioscience, Williston, VT) and imported into Matlab. The cortical columns were determined in the VSD data by deflecting a single whisker using a piezoelectric actuator and recording the cortical response. An example of the spatiotemporal VSD response to a single whisker deflection is shown in Figure A.1(b). The initial frame of cortical activation, which has previously been shown to be restricted to a single cortical column [152], was captured for individual deflection of 4-6 different whiskers during each experiment. An example of the VSD response to the deflection of four different whiskers is overlaid in Figure A.1(d). A least-squares algorithm was used to map the histologically identified anatomical barrel cortex map to the functional

barrel cortex map measured through the VSD imaging. The histologically identified barrel map was registered with the functional column mapping from VSD by solving the following linear inverse problem [127]:

$$\begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_n \\ v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix} = \begin{bmatrix} x_1 & y_1 & 1 & 0 \\ x_2 & y_2 & 1 & 0 \\ \vdots & \vdots & 1 & 0 \\ x_n & y_n & 1 & 0 \\ y_1 & -x_1 & 0 & 1 \\ y_2 & -x_2 & 0 & 1 \\ \vdots & \vdots & 0 & 1 \\ y_n & -x_n & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} A \cos \theta \\ A \sin \theta \\ t_x \\ t_y \end{bmatrix} \quad (31)$$

where the set of u_i and v_i are the coordinates of the center of mass of VSD activation following each whisker deflection, x_i and y_i are the coordinates of the centroid of each barrel column in the histology frame of reference, and A , θ , t_x , and t_y represent the scale, rotation, and translation, respectively, to register the two coordinate axes.

If histology was not available for a given experimental day, the histological map from other experimental sessions that produced the best fit was used. We have found that histologically identified anatomical barrel cortex maps were consistent across animals in terms of structure, and there was no qualitative difference in the results when histological maps were interchanged across experiments.

APPENDIX B

ANATOMICALLY BASED BAYESIAN DECODING OF THE CORTICAL RESPONSE TO INTRACORTICAL MICROSTIMULATION

B.1 Introduction

Sensory prostheses aim to restore function when it has been damaged due to trauma or disease by delivering surrogate sensory signals through patterns of electrical microstimulation. Implementations of sensory prostheses in the peripheral nervous system, such as the cochlear implant and retinal prosthesis, have demonstrated remarkable success [87, 202]. Meanwhile, attempts to deliver sensory information directly to the brain have proved more challenging. In order to leverage the success of peripheral sensory prostheses, a unified framework for evaluating stimulus encoding models and the resulting percepts delivered to the brain is needed.

Microstimulation has long been known to modulate behavior [166], and recent work has even demonstrated the feasibility of closed loop sensorimotor prostheses in the awake animal [144]. Yet, current behavioral paradigms are limited to a few patterns of stimuli, prompting questions on the generalizability of the patterned stimulation and ultimately leading to inefficient evaluation of encoding model performance. Further, it remains unclear whether the surrogate sensory signals need to exactly mimic naturally occurring sensory evoked activity, or if merely presenting discriminable stimulus patterns is sufficient. In either case, the ability to decode the surrogate sensory signals generated by patterned microstimulation is needed in order to establish performance metrics for the design of stimulation encoding models.

Here we use voltage sensitive dye imaging (VSDI) to record the spatiotemporal

cortical activity in response to intracortical microstimulation (ICMS) as an experimental paradigm upon which an anatomically based Bayesian decoder is constructed. The decoder performance is evaluated as a function of stimulus intensity for stimulation across multiple electrodes in cortex. The results presented here form the beginnings of an anatomically based Bayesian decoding framework for evaluating the performance of stimulation encoding models designed to deliver surrogate sensory signals through spatiotemporal patterns of microstimulation.

B.2 Methods

B.2.1 Experimental Paradigm

All procedures were approved by the Georgia Institute of Technology Institutional Animal Care and Use Committee in accordance with NIH guidelines. Female sprague-dawley rats (250-300g) were initially anesthetized with 4% isoflurane before intraperitoneal injection of sodium pentobarbital (50mg/kg weight) for long term anesthesia.

Animals were mounted in a stereotactic device and a craniotomy was performed over the left parietal cortex (coordinates: 1-4mm posterior to bregma, 4-7mm lateral to midline) to expose the barrel representation of the primary somatosensory cortex. A diagram of the in vivo experimental preparation is shown in Fig. B.1A.

B.2.2 Voltage Sensitive Dye Imaging

Voltage sensitive dye imaging was used to monitor cortical activation. The detailed methods are described elsewhere [195]. Briefly, the cortex was stained with dye RH1691 (1mg/mL; Optical Imaging). A 1.0x magnification lens was used in conjunction with a 0.63x condenser lens to provide 1.6x magnification (48 pixels/mm) to image the cortical surface. The VSD data were acquired at five millisecond interframe intervals beginning 200 milliseconds preceding stimulus presentation and processed offline in Matlab (Mathworks Inc, Natick, MA).

The VSD data was functionally registered to the histological map of the barrel

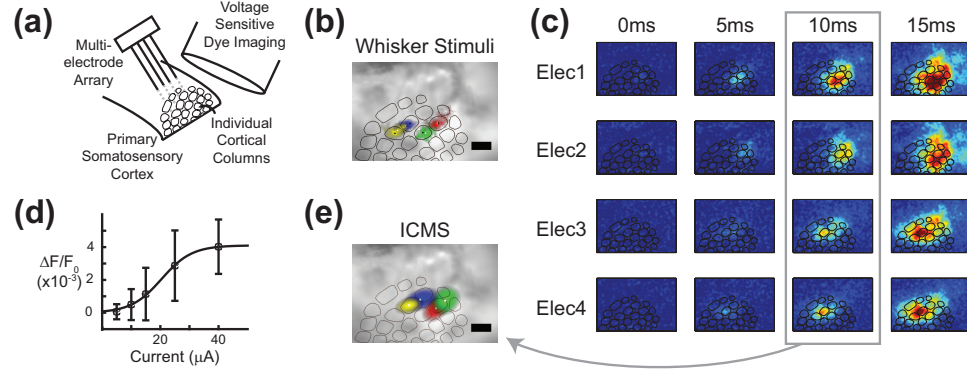


Figure B.1: Voltage sensitive dye imaging of intracortical microstimulation. (a) Experimental setup. (b) Registration of anatomical map with cortical response evoked by deflection of individual whiskers. (c) Spatiotemporal cortical response to intracortical microstimulation across four electrodes. (d) The response amplitude in cortex increases monotonically as a function of the current intensity of the stimulus. (e) Overlay of the onset frame of activation (10ms) for the four electrodes. Error bars in (b) and (e) are $500\mu\text{m}$.

cortex in order to discretize the spatiotemporal cortical signal with regard to well defined cortical columns according to methods described previously [195]. An example of the registration of the anatomical map is shown in Fig. B.1(b). Following the functional image registration, the cortical response is discretized, y_i , where each signal corresponds to the i^{th} functional cortical column. In so doing, the VSD signal is averaged spatially within the contour of the cortical column.

B.2.3 Intracortical Microstimulation

A custom microelectrode array, comprised of four glass coated tungsten microelectrodes (impedance = $0.5\text{-}1\text{ M}\Omega$), was advanced to a depth of $750\mu\text{m}$ in vibrissa region of the primary somatosensory cortex at a 45 degree angle to the cortical surface using a precision microdrive (Knopf Instruments, Tujunga, CA). The principal vibrissa for each electrode was determined by manually deflecting individual whiskers and observing the resulting multi-unit activity.

The electrode was used to deliver microstimulation to the surrounding tissue. Individual electrical stimuli were charge-balanced, cathodal-first, biphasic waveforms

of 200 microsecond duration per phase. The amplitude of the electrical stimuli were varied systematically between $5-80\mu\text{A}$ for each electrode and all stimuli were interleaved.

B.2.4 Decoding

We perform Bayesian decoding on the cortical response to ICMS. The functional architecture of the cortex is used as the basis for decoding according to the following equation:

$$y = \mathbf{A}x + n \quad \text{for } n \sim \mathcal{N}(0, \Sigma) \quad (32)$$

where x , y , and n are $N \times 1$ vectors, and \mathbf{A} is an $N \times N$ matrix, with N equal to the number of cortical columns within the anatomical registration to the VSDI data. This formulation removed the aspect of time to only deal with single frames of VSDI data, and reduced the two-dimensional anatomy to a vector of cortical columns. This decoding framework modeled the stimulus, x , as a point source such that $x_i = 1$ and $x_j = 0$ for all $i \neq j$ and i indicating the cortical column that was stimulated. The matrix \mathbf{A} defines the mapping of the discrete stimulus, x , to the spatial VSDI data, y , where each row of \mathbf{A} is a translated version of a point spread function. The point spread function is modeled as a two dimensional Gaussian function fit to whisker evoked VSDI activity. The noise, n , is spatially correlated and modeled from the covariance across the cortical columns in the pre-stimulus VSDI activity. The Bayesian decoding is performed with a flat prior, such that the maximum likelihood estimate (MLE) and the maximum a posteriori (MAP) are identical. For each single trial of VSDI data, the MLE, \hat{x}_{MLE} , was calculated by maximizing the log-likelihood function:

$$\hat{x}_{MLE} = \arg \max_x \ln \mathcal{L}(y|x) \quad (33)$$

where

$$\ln \mathcal{L}(y|x) = (y - \mathbf{A}x)^T \boldsymbol{\Sigma}^{-1} (y - \mathbf{A}x) \quad (34)$$

The maximum likelihood estimate for each trial, \hat{x}_{MLE} , is determined to be correct if it matches the electrophysiologically identified cortical column location of the electrode stimulated on that trial. Because the MLE was chosen from a discrete set, the decoding framework specifically performed a classification problem.

B.3 Results

B.3.1 ICMS activates localized regions in cortex

A custom 4-electrode array was inserted into the vibrissa portion of the primary somatosensory cortex of the rodent with VSDI used to measure the cortical activity, as illustrated by the diagram in Fig. B.1(a). First, individual whiskers were deflected, and the resulting cortical responses were used to functionally register an anatomical map of the cortical columns to the VSDI data. An overlay of the response to the deflections of four different whiskers, along with the registered anatomical map, is shown in Fig. B.1(b).

The registration was then used to interpret the cortical response to ICMS. Single biphasic stimuli of varying current intensity were delivered on individual electrodes inserted into putative layer 4 of cortex. Figure B.1(c) shows the spatiotemporal evolution of the cortical response to stimulation on each of the four electrodes in the array. The responses were initially localized to a single cortical column, but quickly grew in amplitude and spread spatially. The response amplitude in cortex increased monotonically as a function of the current intensity of the stimulus, as illustrated by Fig. B.1(d). Gaussian fits to the initial frame of activation (10ms) for stimulation of each electrode are overlaid in Fig. B.1(e).

B.3.2 Anatomically based Bayesian decoder accurately estimates stimulus location

A Bayesian decoding framework was used to estimate the location of the stimulus on a single trial basis. The decoding framework is described in detail in the methods. Briefly, a point spread function was fit from whisker evoked activity and used to formulate \mathbf{A} , which describes the blurred VSDI activation, y , given the activation of a single cortical column in x . Within this anatomically based framework, the goal of the decoder was to infer the stimulus location given the blurred, noisy VSDI data. The set of possible stimulus locations available to the decoder corresponded to the number of cortical columns within the VSDI data, and an additional possibility of no stimulus. For each single trial of VSDI, the MLE was evaluated and the decoded stimulus location was determined to be correct if it matched the electrophysiologically identified location of the electrode.

The performance is summarized in Fig. B.2. As the stimulus intensity was increased, the performance of the decoder improved in a systematic way, as in Fig. B.2(a). For a weak stimulus ($5\mu\text{A}$), the response in cortex was non-existent, leading the decoder to estimate that no stimulus was delivered (not shown). For a slightly stronger stimulus ($15\mu\text{A}$), the majority of the trials were still classified as “no stimulus”, but a few trials were identified as the correct cortical column (D2, shown in red). This is shown for a single example in the left portion of Fig. B.2(c), where the majority of the trials were classified as “no stimulus” given by the extra square outside of the anatomical map. For a strong stimulus ($40\mu\text{A}$), the response in cortex was robust, leading to each trial being detected. The majority of the trials were classified correctly, while the mistakes were attributed to nearby cortical columns, as in the right portion of Fig. B.2(c).

The decoding performance for each individual electrode ($N=8$) across all animals is presented in Fig. B.2(b) as the thin gray lines. The median across all electrodes is

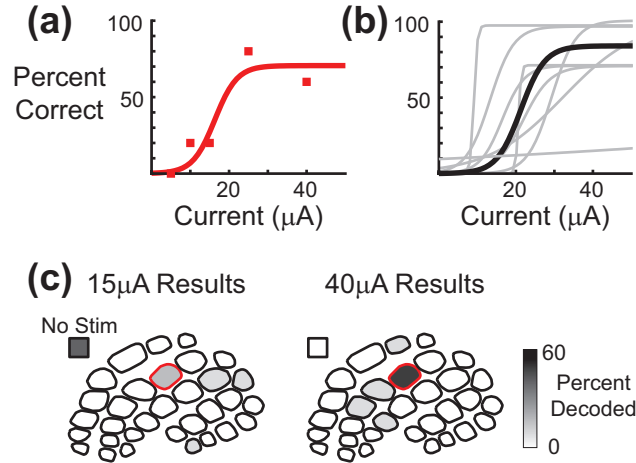


Figure B.2: Decoding of stimulus location. (a) Typical example of decoding performance as a function of stimulus intensity. (b) Summary of decoding performance. Each individual electrode across all animals is shown in gray, with the median performance in black. (c) Pictorial histogram of single trial estimates for the example in panel A. The shading of the cortical columns indicates the number of trials for which it was selected as the MLE, where the cortical column in red was the stimulated column. Weak currents are un-detected the majority of trials (left), while strong currents are correctly estimated with few errors in neighboring cortical columns (right).

given by the thick black line. On average, the 50% decoding performance occurred at $21\mu\text{A}$, while the maximum performance was 84%.

B.3.3 ICMS cortical activation distinct from whisker evoked activity

The decoder was used in the previous section to decode the stimulus location within the context of the anatomical map of the cortical columns. This analysis demonstrated that relevant sensory information was delivered to the cortex through the multi-electrode ICMS. Here we used the same decoding framework to quantify the relative similarity of the neural response to whisker stimuli and ICMS. In this case, the goal of the decoder was to classify the VSDI activation as having been caused by a whisker stimulus, ICMS, or no stimulus. In this case, \mathbf{A} from (32) was a $N \times 2$ matrix, where the two rows corresponded to the decoding filter for a whisker stimulus and ICMS. The whisker stimulus decoding filter was the same as that used in the previous section for the cortical column where the electrode was positioned (left portion

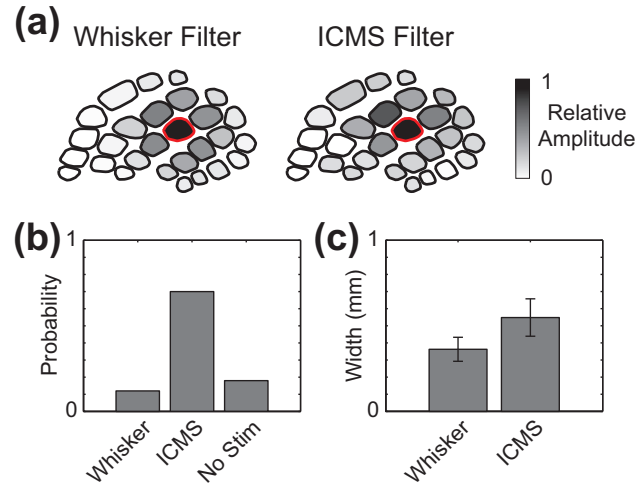


Figure B.3: ICMS produces distinct activation compared to whisker inputs. (a) Decoding filters calculated for whisker stimuli (left) and ICMS (right). (b) Performance of the decoder in determining if the neural response came from a whisker stimulus or ICMS. (c) ICMS generates more spatial spread in the cortical response than whisker stimuli, underlying the high performance of the decoder in separating whisker stimuli and ICMS.

of Fig. B.3(a)). To generate the ICMS decoding filter, 200 trials of VSDI data were collected at a single current intensity on a single electrode. Half of the electrical stimulation data was used to fit a decoding filter (right portion of Fig. B.3(a)), while the other half was used to test the decoder. The whisker and ICMS decoding filters were very similar, suggesting that whisker stimuli and ICMS might be indistinguishable.

Maximum likelihood estimation was performed on the remaining 100 trials of VSDI data collected from ICMS. The average results across the two animals are shown in Fig. B.3(b). The decoder was able to correctly identify the ICMS on 70% of the trials. Importantly, of the two types of errors (i.e. classifying as whisker stimulus or as no stimulus), the “no stimulus” condition occurred most often. This indicates that while the anatomically based Bayesian decoder was effective at extracting behaviorally relevant sensory signals from the ICMS, the neural activation caused by ICMS is distinctly different from whisker evoked activity. This can be seen in the slight differences between the decoding filters for the whisker stimulus and ICMS in

Fig. B.3(a). However, it is more apparent when returning to the raw data and calculating the width of the Gaussian fit to the spatial spread of activation. Figure B.3(c) shows that the width of the spatial spread of the ICMS response in cortex is significantly greater than that of whisker evoked activity ($p=0.002$, Student's T-test).

B.4 Conclusion

Voltage sensitive dye imaging was used as an experimental paradigm for analyzing the cortical response to ICMS and developing a decoding framework for quantifying the information delivered through artificial stimuli. The cortical response to ICMS was initially constrained to the size of a single cortical column. However, quickly the response spread spatially across cortex. This is consistent with previous ICMS studies that report a large spread of activation, both in extracellular recordings [29] and calcium imaging [83].

A Bayesian decoder was built from the underlying anatomical arrangement of the functional cortical columns in the rodent barrel cortex in order to classify the stimulus location generated by ICMS. The decoder accurately estimated the stimulus location of the ICMS. Importantly, one of the available classes was the absence of a stimulus. Therefore, the decoder was required to both detect a stimulus and discriminate between competing stimuli. In our results, the threshold for the decoder detecting the stimulus was consistent with previous behavioral work in rodents and primates [136, 30]. Also, when the stimulus was detected (i.e. not estimated as “no stimulus”), the correct stimulus location was classified with few errors, which is again similar to previous work [47].

The performance of the anatomically based decoder indicates that behaviorally relevant surrogate sensory signals were delivered by the ICMS to the cortex. However, when the decoder was modified in order to classify the neural response as having come from a whisker stimulus or ICMS, it correctly distinguished between the stimuli with

high probability. This suggests that the ICMS does not produce the same sensory perception as that of whisker stimuli. This was further evident from the increased spatial spread of activation in cortex due to ICMS as compared to whisker stimuli. This result is consistent with previous studies that have reported significant spatial spread caused by ICMS [29], and may be due to the stimulation of fibers of passage [83, 195].

The decoder in this study only operated on single frames, restricting the decoding problem to location identification. An immediate extension of this work is to include time into the decoding problem, allowing for the decoding of neural activity evoked by spatiotemporal patterns of stimulation. Whereas a flat prior was used in this study, the extension to decoding patterns of activity would utilize the Bayesian framework to define the ethological relevance of activity patterns in the brain. The framework provided here, paired with the experimental approach, serves as an initial step in the development of an anatomically based spatiotemporal decoder of neural activity for the design and evaluation of stimulation encoding models for sensorimotor prostheses.

REFERENCES

- [1] AERTSEN, A., DIEMANN, M., and GEWALTIG, M. O., "Propagation of synchronous spiking activity in feedforward neural networks.," *Journal of physiology, Paris*, vol. 90, pp. 243–7, Jan. 1996.
- [2] ALONSO, J. M., USREY, W. M., and REID, R. C., "Precisely correlated firing in cells of the lateral geniculate nucleus," *Nature*, vol. 383, no. 6603, pp. 815–819, 1996.
- [3] ARAVANIS, A. M., WANG, L.-P., ZHANG, F., MELTZER, L. A., MOGRI, M. Z., SCHNEIDER, M. B., and DEISSEROTH, K., "An optical neural interface: in vivo control of rodent motor cortex with integrated fiberoptic and optogenetic technology.," *Journal of neural engineering*, vol. 4, pp. S143–56, Sept. 2007.
- [4] ARMSTRONG-JAMES, M. and CALLAHAN, C. A., "Thalamo-cortical processing of vibrissal information in the rat. II. spatiotemporal convergence in the thalamic ventroposterior medial nucleus (VPM) and its relevance to generation of receptive fields of S1 cortical "barrel" neurones.," *The Journal of comparative neurology*, vol. 303, pp. 211–24, Jan. 1991.
- [5] AVERBECK, B. B., LATHAM, P. E., and POUGET, A., "Neural correlations, population coding and computation.," *Nature reviews. Neuroscience*, vol. 7, pp. 358–66, May 2006.
- [6] BACH-Y RITA, P., COLLINS, C. C., SAUNDERS, F. A., WHITE, B., and SCADDEN, L., "Vision substitution by tactile image projection," *Nature*, vol. 221, no. 5184, pp. 963–964, 1969.
- [7] BARROS, C. G., BITTAR, R. S., and DANILOV, Y., "Effects of electrotactile vestibular substitution on rehabilitation of patients with bilateral vestibular loss," *Neurosci Lett*, vol. 476, no. 3, pp. 123–126, 2010.
- [8] BAZHENOV, M., TIMOFEEV, I., STERIADE, M., and SEJNOWSKI, T. J., "Computational Models of Thalamocortical Augmenting Responses," *The Journal of Neuroscience*, vol. 18, no. 16, pp. 6444–6465, 1998.
- [9] BENOIT, M. M., RAIJ, T., LIN, F.-H., JÄÄSKELÄINEN, I. P., and STUFFLEBEAM, S., "Primary and multisensory cortical activity is correlated with audiovisual percepts," *Human Brain Mapping*, vol. 31, no. 4, pp. 526–538, 2010.
- [10] BERGER, T. W., HAMPSON, R. E., SONG, D., GOONAWARDENA, A., MARMARELIS, V. Z., and DEADWYLER, S. A., "A cortical neural prosthesis for restoring and enhancing memory.," *Journal of neural engineering*, vol. 8, p. 046017, Aug. 2011.

- [11] BERGER, T., BORGDORFF, A., CROCHET, S., NEUBAUER, F. B., LEFORT, S., FAUVET, B., FERZOU, I., CARLETON, A., LUSCHER, H.-R., and PETERSEN, C. C. H., “Combined Voltage and Calcium Epifluorescence Imaging In Vitro and In Vivo Reveals Subthreshold and Suprathreshold Dynamics of Mouse Barrel Cortex,” *J Neurophysiol*, vol. 97, no. 5, pp. 3751–3762, 2007.
- [12] BIALEK, W., RIEKE, F., DE RUYTER VAN STEVENINCK, R. R., and WARLAND, D., “Reading a neural code,” *Science (New York, N.Y.)*, vol. 252, pp. 1854–7, June 1991.
- [13] BILGER, R. C., BLACK, F. O., HOPKINSON, N. T., MYERS, E. N., PAYNE, J. L., STENSON, N. R., VEGA, A., and WOLF, R. V., “Evaluation of subjects presently fitted with implanted auditory prostheses,” *Annals of Otolaryngology, Rhinology, and Laryngology*, no. 86, pp. 1–176, 1977.
- [14] BOLOORI, A. R., JENKS, R. A., DESBORDES, G., and STANLEY, G. B., “Encoding and decoding cortical representations of tactile features in the vibrissa system,” *Journal of Neuroscience*, vol. 30, no. 30, p. 9990, 2010.
- [15] BOLOORI, A. R. and STANLEY, G. B., “The dynamics of spatiotemporal response integration in the somatosensory cortex of the vibrissa system,” *Journal of Neuroscience*, vol. 26, no. 14, p. 3767, 2006.
- [16] BOYDEN, E. S., ZHANG, F., BAMBERG, E., NAGEL, G., and DEISSEROTH, K., “Millisecond-timescale, genetically targeted optical control of neural activity,” *Nat Neurosci*, vol. 8, no. 9, pp. 1263–1268, 2005.
- [17] BRADLEY, D. C., TROYK, P. R., BERG, J. A., BAK, M., COGAN, S., ERICKSON, R., KUFTA, C., MASCARO, M., MCCREERY, D., SCHMIDT, E. M., TOWLE, V. L., and XU, H., “Visuotopic Mapping Through a Multichannel Stimulating Implant in Primate V1,” *Journal of Neurophysiology*, vol. 93, no. 3, pp. 1659–1670, 2005.
- [18] BRADLEY, D. C., TROYK, P. R., BERG, J. A., BAK, M., COGAN, S., ERICKSON, R., KUFTA, C., MASCARO, M., MCCREERY, D., SCHMIDT, E. M., TOWLE, V. L., and XU, H., “Visuotopic Mapping Through a Multichannel Stimulating Implant in Primate V1,” *Journal of Neurophysiology*, vol. 93, no. 3, pp. 1659–1670, 2005.
- [19] BRINDLEY, G. S. and LEWIN, W. S., “The sensations produced by electrical stimulation of the visual cortex,” *The Journal of physiology*, vol. 196, pp. 479–93, May 1968.
- [20] BROCKER, D. T., SWAN, B. D., TURNER, D. A., GROSS, R. E., TATTER, S. B., KOOP, M. M., BRONTE-STEWART, H., and GRILL, W. M., “Improved efficacy of temporally non-regular deep brain stimulation in Parkinson’s disease,” *Experimental neurology*, vol. 239, pp. 60–7, Jan. 2013.

- [21] BROOKS, V. B., *The neural basis of motor control*. Oxford University Press, 1986.
- [22] BRUGGER, D., BUTOVAS, S., BOGDAN, M., and SCHWARZ, C., “Real-time adaptive microstimulation increases reliability of electrically evoked cortical potentials.,” *IEEE transactions on bio-medical engineering*, vol. 58, pp. 1483–91, May 2011.
- [23] BRUMBERG, J. C., PINTO, D. J., and SIMONS, D. J., “Spatial gradients and inhibitory summation in the rat whisker barrel system,” *Journal of Neurophysiology*, vol. 76, no. 1, pp. 130–140, 1996.
- [24] BRUNO, R. M., “Synchrony in sensation.,” *Current opinion in neurobiology*, vol. 21, pp. 701–8, Oct. 2011.
- [25] BRUNO, R. M., KHATRI, V., LAND, P. W., and SIMONS, D. J., “Thalamocortical Angular Tuning Domains within Individual Barrels of Rat Somatosensory Cortex,” *J. Neurosci.*, vol. 23, no. 29, pp. 9565–9574, 2003.
- [26] BRUNO, R. M. and SAKMANN, B., “Cortex is driven by weak but synchronously active thalamocortical synapses.,” *Science (New York, N.Y.)*, vol. 312, pp. 1622–7, June 2006.
- [27] BUTOVAS, S., HORMUZDI, S. G., MONYER, H., and SCHWARZ, C., “Effects of electrically coupled inhibitory networks on local neuronal responses to intracortical microstimulation,” *Journal of Neurophysiology*, vol. 96, no. 3, pp. 1227–1236, 2006.
- [28] BUTOVAS, S. and SCHWARZ, C., “Detection psychophysics of intracortical microstimulation in rat primary somatosensory cortex,” *European Journal of Neuroscience*, vol. 25, no. 7, pp. 2161–2169, 2007.
- [29] BUTOVAS, S. and SCHWARZ, C., “Spatiotemporal effects of microstimulation in rat neocortex: a parametric study using multielectrode recordings.,” *Journal of neurophysiology*, vol. 90, pp. 3024–39, Nov. 2003.
- [30] BUTOVAS, S. and SCHWARZ, C., “Detection psychophysics of intracortical microstimulation in rat primary somatosensory cortex.,” *The European journal of neuroscience*, vol. 25, pp. 2161–9, Apr. 2007.
- [31] BUTTS, D. A., WENG, C., JIN, J., YEH, C. I., LESICA, N. A., ALONSO, J. M., and STANLEY, G. B., “Temporal precision in the neural code and the timescales of natural vision,” *Nature*, vol. 449, no. 7158, pp. 92–95, 2007.
- [32] BUZSÁKI, G., “Large-scale recording of neuronal ensembles,” *Nature Neuroscience*, vol. 7, no. 5, pp. 446–451, 2004.

- [33] CASTRO-ALAMANCOS, M. A. and CONNORS, B. W., “Short-term plasticity of a thalamocortical pathway dynamically modulated by behavioral state,” *Science (New York, N.Y.)*, vol. 272, pp. 274–7, Apr. 1996.
- [34] CASTRO-ALAMANCOS, M. A. and OLDFORD, E., “Cortical sensory suppression during arousal is due to the activity-dependent depression of thalamocortical synapses,” *Journal of Physiology-London*, vol. 541, no. 1, pp. 319–331, 2002.
- [35] CHANG, B. S. and LOWENSTEIN, D. H., “Epilepsy,” *N Engl J Med*, vol. 349, no. 13, pp. 1257–1266, 2003.
- [36] CHEN, L. M., FRIEDMAN, R. M., and ROE, A. W., “Optical Imaging of a Tactile Illusion in Area 3b of the Primary Somatosensory Cortex,” *Science*, vol. 302, no. 5646, pp. 881–885, 2003.
- [37] CHEN, Y., GEISLER, W. S., and SEIDEMANN, E., “Optimal decoding of correlated neural population responses in the primate visual cortex,” *Nat Neurosci*, vol. 9, no. 11, pp. 1412–1420, 2006.
- [38] CIVILLICO, E. F. and CONTRERAS, D., “Comparison of responses to electrical stimulation and whisker deflection using two different voltage-sensitive dyes in mouse barrel cortex in vivo,” *The Journal of membrane biology*, vol. 208, pp. 171–82, Nov. 2005.
- [39] CLARK, K. L., ARMSTRONG, K. M., and MOORE, T., “Probing neural circuitry and function with electrical microstimulation,” *Proceedings of the Royal Society B: Biological Sciences*, vol. 278, no. 1709, pp. 1121–1130, 2011.
- [40] CONSTANTINOPLE, C. M. and BRUNO, R. M., “Deep cortical layers are activated directly by thalamus,” *Science (New York, N.Y.)*, vol. 340, pp. 1591–4, June 2013.
- [41] COVER, T. M. and THOMAS, J. A., *Elements of information theory*. John Wiley & Sons, 2012.
- [42] CRICK, F., “Function of the thalamic reticular complex: the searchlight hypothesis,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 81, no. 14, p. 4586, 1984.
- [43] CROCHET, S. and PETERSEN, C. C. H., “Correlating whisker behavior with membrane potential in barrel cortex of awake mice,” *Nature neuroscience*, vol. 9, pp. 608–10, May 2006.
- [44] DALY, J., LIU, J., AGHAGOLZADEH, M., and OWEISS, K., “Optimal space-time precoding of artificial sensory feedback through multichannel microstimulation in bi-directional brain-machine interfaces,” *Journal of Neural Engineering*, vol. 9, p. 065004, Dec. 2012.

- [45] DAS, A. and GILBERT, C. D., “Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex,” *Nature*, vol. 375, no. 6534, pp. 780–784, 1995.
- [46] DAVIS, K. D., KISS, Z. H., LUO, L., TASKER, R. R., LOZANO, A. M., and DOSTROVSKY, J. O., “Phantom sensations generated by thalamic microstimulation,” *Nature*, vol. 391, pp. 385–7, Jan. 1998.
- [47] DELIANO, M., SCHEICH, H., and OHL, F. W., “Auditory Cortical Activity after Intracortical Microstimulation and Its Role for Sensory Processing and Learning,” *Journal of Neuroscience*, vol. 29, pp. 15898–15909, Dec. 2009.
- [48] DEMPSEY, E. W. and MORISON, R. S., “The Electrical Activity of a Thalamocortical Relay System,” *American Journal of Physiology – Legacy Content*, vol. 138, no. 2, pp. 283–296, 1943.
- [49] DEUSCHL, G., SCHADE-BRITTINGER, C., KRACK, P., VOLKMANN, J., SCHÄFER, H., BÖTZEL, K., DANIELS, C., DEUTSCHLÄNDER, A., DILLMANN, U., EISNER, W., GRUBER, D., HAMEL, W., HERZOG, J., HILKER, R., KLEBE, S., KLOSS, M., KOY, J., KRAUSE, M., KUPSCH, A., LORENZ, D., LORENZL, S., MEHDORN, H. M., MORINGLANE, J. R., OERTEL, W., PINSKER, M. O., REICHMANN, H., REUSS, A., SCHNEIDER, G.-H., SCHNITZLER, A., STEUDE, U., STURM, V., TIMMERMANN, L., TRONNIER, V., TROTTEBERG, T., WOJTECKI, L., WOLF, E., POEWE, W., and VOGES, J., “A randomized trial of deep-brain stimulation for Parkinson’s disease,” *The New England journal of medicine*, vol. 355, pp. 896–908, Aug. 2006.
- [50] DIAMOND, M. E., ARMSTRONG-JAMES, M., and EBNER, F. F., “Somatic sensory responses in the rostral sector of the posterior group (POm) and in the ventral posterior medial nucleus (VPM) of the rat thalamus,” *The Journal of comparative neurology*, vol. 318, pp. 462–76, Apr. 1992.
- [51] DIAMOND, M. E., VON HEIMENDAHL, M., KNUTSEN, P. M., KLEINFELD, D., and AHISSAR, E., “Where’ and ‘what’ in the whisker sensorimotor system,” *Nature reviews. Neuroscience*, vol. 9, pp. 601–12, Aug. 2008.
- [52] DIESMANN, M., GEWALTIG, M. O., and AERTSEN, A., “Stable propagation of synchronous spiking in cortical neural networks,” *Nature*, vol. 402, pp. 529–33, Dec. 1999.
- [53] DORVAL, A. D., KUNCEL, A. M., BIRDNO, M. J., TURNER, D. A., and GRILL, W. M., “Deep brain stimulation alleviates parkinsonian bradykinesia by regularizing pallidal activity,” *Journal of neurophysiology*, vol. 104, pp. 911–21, Aug. 2010.
- [54] FANSELOW, E. E., SAMESHIMA, K., BACCALA, L. A., and NICOLELIS, M. A., “Thalamic bursting in rats during different awake behavioral states,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 98, pp. 15330–5, Dec. 2001.

- [55] FERZOU, I., HAIS, F., GENTET, L. J., ARONOFF, R., WEBER, B., and PETERSEN, C. C. H., "Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice," *Neuron*, vol. 56, no. 5, pp. 907–923, 2007.
- [56] FITZSIMMONS, N. A., DRAKE, W., HANSON, T. L., LEBEDEV, M. A., and NICOLELIS, M. A. L., "Primate Reaching Cued by Multichannel Spatiotemporal Cortical Microstimulation," *Journal of Neuroscience*, vol. 27, no. 21, pp. 5593–5602, 2007.
- [57] FRIEDBERG, M. H., LEE, S. M., and EBNER, F. F., "Modulation of receptive field properties of thalamic somatosensory neurons by the depth of anesthesia," *Journal of neurophysiology*, vol. 81, pp. 2243–52, May 1999.
- [58] FRITSCH GUSTAV HITZIG, E., "Uber die elektrische Erregbarkeit des Grosshims," *Arch. Anat. Physiol. wiss Med*, vol. 37, pp. 300–332, 1870.
- [59] GABERNET, L., JADHAV, S. P., FELDMAN, D. E., CARANDINI, M., and SCANZIANI, M., "Somatosensory integration controlled by dynamic thalamocortical feed-forward inhibition," *Neuron*, vol. 48, pp. 315–27, Oct. 2005.
- [60] GAUNT, R. A., HOKANSON, J. A., and WEBER, D. J., "Microstimulation of primary afferent neurons in the L7 dorsal root ganglia using multielectrode arrays in anesthetized cats: thresholds and recruitment properties," *Journal of Neural Engineering*, vol. 6, no. 5, 2009.
- [61] GENTET, L. J., AVERMANN, M., MATYAS, F., STAIGER, J. F., and PETERSEN, C. C. H., "Membrane potential dynamics of GABAergic neurons in the barrel cortex of behaving mice," *Neuron*, vol. 65, pp. 422–35, Feb. 2010.
- [62] GEORGOPOULOS, A., SCHWARTZ, A., and KETTNER, R., "Neuronal population coding of movement direction," *Science*, vol. 233, pp. 1416–1419, Sept. 1986.
- [63] GEORGOPOULOS, A. P., KALASKA, J. F., CAMINITI, R., and MASSEY, J. T., "On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex," *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 2, pp. 1527–37, Nov. 1982.
- [64] GILJA, V., CHESTEK, C. A., DIESTER, I., HENDERSON, J. M., DEISSEROTH, K., and SHENOY, K. V., "Challenges and opportunities for next-generation intracortically based neural prostheses," *Biomedical Engineering, IEEE Transactions on*, vol. 58, no. 7, pp. 1891–1899, 2011.
- [65] GOARD, M. and DAN, Y., "Basal forebrain activation enhances cortical coding of natural scenes," *Nature neuroscience*, vol. 12, pp. 1444–9, Nov. 2009.

- [66] GOLDBERG, J. M. and BROWN, P. B., "Response of binaural neurons of dog superior olivary complex to dichotic tonal stimuli: some physiological mechanisms of sound localization," *Journal of Neurophysiology*, vol. 32, no. 4, pp. 613–636, 1969.
- [67] GRADINARU, V., MOGRI, M., THOMPSON, K. R., HENDERSON, J. M., and DEISSEROTH, K., "Optical deconstruction of parkinsonian neural circuitry.," *Science (New York, N.Y.)*, vol. 324, pp. 354–9, Apr. 2009.
- [68] GRAZIANO, M. S. A., TAYLOR, C. S. R., and MOORE, T., "Complex Movements Evoked by Microstimulation of Precentral Cortex," *Neuron*, vol. 34, no. 5, pp. 841–851, 2002.
- [69] GRILL, W. M. and MORTIMER, J. T., "Inversion of the current-distance relationship by transient depolarization," *IEEE Transactions on Biomedical Engineering*, vol. 44, no. 1, pp. 1–9, 1997.
- [70] GRILL, W. M., NORMAN, S. E., and BELLAMKONDA, R. V., "Implanted neural interfaces: biochallenges and engineered solutions," *Annual Review of Biomedical Engineering*, vol. 11, pp. 1–24, 2009.
- [71] GRILL, W. M., SIMMONS, A. M., COOPER, S. E., MIOCINOVIC, S., MONTGOMERY, E. B., BAKER, K. B., and REZAI, A. R., "Temporal excitation properties of paresthesias evoked by thalamic microstimulation," *Clinical Neurophysiology*, vol. 116, no. 5, pp. 1227–1234, 2005.
- [72] GRINVALD, A. and HILDESHEIM, R., "VSDI: a new era in functional imaging of cortical dynamics," *Nat Rev Neurosci*, vol. 5, no. 11, pp. 874–885, 2004.
- [73] GRINVALD, A., LIEKE, E. E., FROSTIG, R. D., and HILDESHEIM, R., "Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex," *The Journal of Neuroscience*, vol. 14, no. 5, pp. 2545–2568, 1994.
- [74] GUSTAFSSON, B. and JANKOWSKA, E., "Direct and indirect activation of nerve cells by electrical pulses applied extracellularly," *J Physiol*, vol. 258, no. 1, pp. 33–61, 1976.
- [75] HAIDARLIU, S. and AHISSAR, E., "Size gradients of barreloids in the rat thalamus," *The Journal of Comparative Neurology*, vol. 429, no. 3, pp. 372–387, 2001.
- [76] HALASSA, M. M., SIEGLE, J. H., RITT, J. T., TING, J. T., FENG, G., and MOORE, C. I., "Selective optical drive of thalamic reticular nucleus generates thalamic bursts and cortical spindles," *Nat Neurosci*, vol. 14, no. 9, pp. 1118–1120, 2011.

- [77] HAMPSON, R. E., GERHARDT, G. A., MARMARELIS, V., SONG, D., OPRIS, I., SANTOS, L., BERGER, T. W., and DEADWYLER, S. A., "Facilitation and restoration of cognitive function in primate prefrontal cortex by a neuroprosthesis that utilizes minicolumn-specific neural firing," *Journal of neural engineering*, vol. 9, p. 056012, Oct. 2012.
- [78] HARTMANN, R., TOPP, G., and KLINKE, R., "Discharge patterns of cat primary auditory fibers with electrical stimulation of the cochlea," *Hear Res*, vol. 13, no. 1, pp. 47–62, 1984.
- [79] HASHIMOTO, T., ELDER, C. M., and VITEK, J. L., "A template subtraction method for stimulus artifact removal in high-frequency deep brain stimulation," *J Neurosci Methods*, vol. 113, no. 2, pp. 181–186, 2002.
- [80] HEIMER, G., BAR-GAD, I., GOLDBERG, J. A., and BERGMAN, H., "Dopamine replacement therapy reverses abnormal synchronization of pallidal neurons in the 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine primate model of parkinsonism," *J Neurosci*, vol. 22, no. 18, pp. 7850–7855, 2002.
- [81] HEMING, E. A., CHOO, R., DAVIES, J. N., and KISS, Z. H., "Designing a thalamic somatosensory neural prosthesis: consistency and persistence of percepts evoked by electrical stimulation," *IEEE Trans Neural Syst Rehabil Eng*, vol. 19, no. 5, pp. 477–482, 2011.
- [82] HEMING, E., SANDEN, A., and KISS, Z. H. T., "Designing a somatosensory neural prosthesis: percepts evoked by different patterns of thalamic stimulation," *Journal of neural engineering*, vol. 7, p. 064001, Dec. 2010.
- [83] HISTED, M. H., BONIN, V., and REID, R. C., "Direct Activation of Sparse, Distributed Populations of Cortical Neurons by Electrical Microstimulation," *Neuron*, vol. 63, pp. 508–522, Aug. 2009.
- [84] HOLTZHEIMER, P. E. and MAYBERG, H. S., "Deep brain stimulation for psychiatric disorders," *Annual review of neuroscience*, vol. 34, pp. 289–307, Jan. 2011.
- [85] HOUWELING, A. R. and BRECHT, M., "Behavioural report of single neuron stimulation in somatosensory cortex," *Nature*, vol. 451, no. 7174, pp. 65–68, 2008.
- [86] HUBEL, D. H., *Eye, brain, and vision*. Scientific American Library/Scientific American Books, 1995.
- [87] HUMAYUN, M. S., WEILAND, J. D., FUJII, G. Y., GREENBERG, R., WILLIAMSON, R., LITTLE, J., MECH, B., CIMMARUSTI, V., VAN BOEMEL, G., DAGNELIE, G., and DE JUAN, E., "Visual perception in a blind subject with a chronic microelectronic retinal prosthesis," *Vision Research*, vol. 43, pp. 2573–2581, Nov. 2003.

- [88] HUNTER, I. W. and KORENBERG, M. J., “The identification of nonlinear biological systems: Wiener and Hammerstein cascade models,” *Biological Cybernetics*, vol. 55, no. 2, pp. 135–144, 1986.
- [89] IZHIKEVICH, E. M., “Simple model of spiking neurons,” *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council*, vol. 14, pp. 1569–72, Jan. 2003.
- [90] JAMES, M. R. and LEE, E. M., “Enhanced detection threshold for in vivo cortical stimulation produced by Hebbian conditioning,” *Journal of Neural Engineering*, vol. 8, no. 1, p. 16011, 2011.
- [91] JOHNSON, D. H. and SINANOVIC, S., “Symmetrizing the kullback-leibler distance,” *IEEE Transactions on Information Theory*, vol. 1, no. 1, pp. 1–10, 2001.
- [92] KARA, P., REINAGEL, P., and REID, R. C., “Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons,” *Neuron*, vol. 27, pp. 635–46, Sept. 2000.
- [93] KARA, P., PEZARIS, J. S., YURGENSON, S., and REID, R. C., “The spatial receptive field of thalamic inputs to single cortical simple cells revealed by the interaction of visual and electrical stimulation,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, pp. 16261–6, Dec. 2002.
- [94] KEARNS, M. and RON, D., “Algorithmic stability and sanity-check bounds for leave-one-out cross-validation,” *Neural Computation*, vol. 11, no. 6, pp. 1427–1453, 1999.
- [95] KISVÁRDAY, Z. F., TÓTH, E., RAUSCH, M., and EYSEL, U. T., “Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat,” *Cerebral cortex (New York, N.Y. : 1991)*, vol. 7, no. 7, pp. 605–18, 1997.
- [96] KOIVUNIEMI, A. S. and OTTO, K. J., “Asymmetric Versus Symmetric Pulses for Cortical Microstimulation,” *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, vol. 19, no. 5, pp. 468–476, 2011.
- [97] KOUBEISSI, M. Z., KAHRIMAN, E., SYED, T. U., MILLER, J., and DURAND, D. M., “Low-frequency electrical stimulation of a fiber tract in temporal lobe epilepsy,” *Annals of neurology*, Apr. 2013.
- [98] KRAUSZ, H. I., “Identification of nonlinear systems using random impulse train inputs,” *Biological Cybernetics*, vol. 19, pp. 217–230, Sept. 1975.
- [99] KRAVITZ, A. V., FREEZE, B. S., PARKER, P. R. L., KAY, K., THWIN, M. T., DEISSEROTH, K., and KREITZER, A. C., “Regulation of parkinsonian

motor behaviours by optogenetic control of basal ganglia circuitry,” *Nature*, vol. 466, no. 7306, pp. 622–626, 2010.

- [100] KUMAR, A., ROTTER, S., and AERTSEN, A., “Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding,” *Nature Reviews Neuroscience*, vol. 11, no. 9, pp. 615–627, 2010.
- [101] KYRIAZI, H. T. and SIMONS, D. J., “Thalamocortical response transformations in simulated whisker barrels,” *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 13, pp. 1601–15, Apr. 1993.
- [102] LAMPL, I., REICHOVA, I., and FERSTER, D., “Synchronous membrane potential fluctuations in neurons of the cat visual cortex,” *Neuron*, vol. 22, pp. 361–74, Feb. 1999.
- [103] LEANG, C. and JOHNSON, D., “On the asymptotics of M-hypothesis Bayesian detection,” *IEEE Transactions on Information Theory*, vol. 43, no. 1, pp. 280–282, 1997.
- [104] LEE, J. H., DURAND, R., GRADINARU, V., ZHANG, F., GOSHEN, I., KIM, D.-S., FENNO, L. E., RAMAKRISHNAN, C., and DEISSEROTH, K., “Global and local fMRI signals driven by neurons defined optogenetically by type and wiring,” *Nature*, vol. 465, pp. 788–92, June 2010.
- [105] LEE, S.-H., KWAN, A. C., ZHANG, S., PHOUMTHIPPHAVONG, V., FLANER, J. G., MASMANIDIS, S. C., TANIGUCHI, H., HUANG, Z. J., ZHANG, F., BOYDEN, E. S., DEISSEROTH, K., and DAN, Y., “Activation of specific interneurons improves V1 feature selectivity and visual perception,” *Nature*, Aug. 2012.
- [106] LESICA, N. A., JIN, J., WENG, C., YEH, C. I., BUTTS, D. A., STANLEY, G. B., and ALONSO, J. M., “Adaptation to stimulus contrast and correlations during natural visual stimulation,” *Neuron*, vol. 55, no. 3, pp. 479–491, 2007.
- [107] LESICA, N. A., WENG, C., JIN, J., YEH, C. I., ALONSO, J. M., and STANLEY, G. B., “Dynamic encoding of natural luminance sequences by LGN bursts,” *PLoS Biol*, vol. 4, no. 7, p. e209, 2006.
- [108] LESICA, N. A. and STANLEY, G. B., “Encoding of Natural Scene Movies by Tonic and Burst Spikes in the Lateral Geniculate Nucleus,” *The Journal of Neuroscience*, vol. 24, no. 47, pp. 10731–10740, 2004.
- [109] LIPPERT, M. T., TAKAGAKI, K., XU, W., HUANG, X., and WU, J.-Y., “Methods for Voltage-Sensitive Dye Imaging of Rat Cortical Activity With High Signal-to-Noise Ratio,” *Journal of Neurophysiology*, vol. 98, no. 1, pp. 502–512, 2007.

- [110] LIU, J., KHALIL, H. K., and OWEISS, K. G., “Neural feedback for instantaneous spatiotemporal modulation of afferent pathways in bi-directional brain-machine interfaces.” *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 19, pp. 521–33, Oct. 2011.
- [111] LLINÁS, R. R. and STERIADE, M., “Bursting of thalamic neurons and states of vigilance.” *Journal of neurophysiology*, vol. 95, pp. 3297–308, June 2006.
- [112] LOGOTHETIS, N. K., AUGATH, M., MURAYAMA, Y., RAUCH, A., SULTAN, F., GOENSE, J., OELTERMANN, A., and MERKLE, H., “The effects of electrical microstimulation on cortical signal propagation,” *Nat Neurosci*, vol. 13, no. 10, pp. 1283–1291, 2010.
- [113] LOGOTHETIS, N. K. and PFEUFFER, J., “On the nature of the BOLD fMRI contrast mechanism.” *Magnetic resonance imaging*, vol. 22, pp. 1517–31, Dec. 2004.
- [114] LU, S. M., GUIDO, W., and SHERMAN, S. M., “Effects of membrane voltage on receptive field properties of lateral geniculate neurons in the cat: contributions of the low-threshold Ca²⁺ conductance,” *Journal of Neurophysiology*, vol. 68, no. 6, pp. 2185–2198, 1992.
- [115] LU, Y., CAO, P., SUN, J., WANG, J., LI, L., REN, Q., CHEN, Y., and CHAI, X., “Using independent component analysis to remove artifacts in visual cortex responses elicited by electrical stimulation of the optic nerve.” *Journal of neural engineering*, vol. 9, p. 026002, Apr. 2012.
- [116] MARMARELIS, P. Z. and MARMARELIS, V. Z., *Analysis of physiological systems : the white-noise approach*. New York: Plenum Press, 1978.
- [117] MARMARELIS, V. Z., “Identification of nonlinear biological systems using Laguerre expansions of kernels,” *Ann Biomed Eng*, vol. 21, no. 6, pp. 573–589, 1993.
- [118] MARMARELIS, V. Z., *Nonlinear dynamic modeling of physiological systems*. Wiley-IEEE Press, 2004.
- [119] MASSE, N. Y. and COOK, E. P., “Behavioral time course of microstimulation in cortical area MT.” *Journal of neurophysiology*, vol. 103, pp. 334–45, Jan. 2010.
- [120] MATTIS, J., TYE, K. M., FERENCZI, E. A., RAMAKRISHNAN, C., O’SHEA, D. J., PRAKASH, R., GUNAYDIN, L. A., HYUN, M., FENNO, L. E., GRADINARU, V., YIZHAR, O., and DEISSEROTH, K., “Principles for applying optogenetic tools derived from direct comparative analysis of microbial opsins.” *Nature methods*, vol. 9, pp. 159–72, Feb. 2012.

- [121] MAZUREK, M. E. and SHADLEN, M. N., “Limits to the temporal fidelity of cortical spike rate signals,” *Nature neuroscience*, vol. 5, pp. 463–71, May 2002.
- [122] MCALONAN, K., CAVANAUGH, J., and WURTZ, R. H., “Guarding the gateway to cortex with attention in visual thalamus,” *Nature*, vol. 456, no. 7220, pp. 391–394, 2008.
- [123] MCINTYRE, C. C. and GRILL, W. M., “Selective microstimulation of central nervous system neurons,” *Annals of Biomedical Engineering*, vol. 28, no. 3, pp. 219–233, 2000.
- [124] MCINTYRE, C. C., GRILL, W. M., SHERMAN, D. L., and THAKOR, N. V., “Cellular effects of deep brain stimulation: model-based analysis of activation and inhibition,” *Journal of neurophysiology*, vol. 91, pp. 1457–69, Apr. 2004.
- [125] MERZENICH, M. M., MICHELSON, R. P., PETTIT, C. R., SCHINDLER, R. A., and REID, M., “Neural encoding of sound sensation evoked by electrical stimulation of the acoustic nerve,” *Ann Otol Rhinol Laryngol*, vol. 82, no. 4, pp. 486–503, 1973.
- [126] MERZENICH, M. M. and REID, M. D., “Representation of the cochlea within the inferior colliculus of the cat,” *Brain Res*, vol. 77, no. 3, pp. 397–415, 1974.
- [127] MILLARD, D. C., WANG, Q., and STANLEY, G. B., “Nonlinear system identification of the thalamocortical circuit in response to thalamic microstimulation,” in *Neural Engineering (NER), 2011 5th International IEEE/EMBS Conference on*, pp. 1–4, 2011.
- [128] MILLARD, D. C., WANG, Q., GOLLNICK, C. A., and STANLEY, G. B., “System identification of the nonlinear dynamics in the thalamocortical circuit in response to patterned thalamic microstimulation invivo,” *Journal of neural engineering*, vol. 10, p. 066011, Dec. 2013.
- [129] MILLER, K. D., PINTO, D. J., and SIMONS, D. J., “Processing in layer 4 of the neocortical circuit: new insights from visual and somatosensory cortex,” *Current Opinion in Neurobiology*, vol. 11, no. 4, pp. 488–497, 2001.
- [130] MILLER, L. E. and WEBER, D. J., “Brain training: cortical plasticity and afferent feedback in brain-machine interface systems,” *IEEE Trans Neural Syst Rehabil Eng*, vol. 19, no. 5, pp. 465–467, 2011.
- [131] MOORE, T. and ARMSTRONG, K. M., “Selective gating of visual signals by microstimulation of frontal cortex,” *Nature*, vol. 421, no. 6921, pp. 370–373, 2003.
- [132] MORRISON, A., DIESMANN, M., and GERSTNER, W., “Phenomenological models of synaptic plasticity based on spike timing,” *Biological cybernetics*, vol. 98, pp. 459–78, June 2008.

- [133] MOUNTCASTLE, V. B., *Perceptual Neuroscience: The Cerebral Cortex*. Harvard University Press, 1998.
- [134] MUKHERJEE, P. and KAPLAN, E., “Dynamics of neurons in the cat lateral geniculate nucleus: in vivo electrophysiology and computational modeling.,” *Journal of neurophysiology*, vol. 74, pp. 1222–43, Sept. 1995.
- [135] MULLER, R. U., KUBIE, J. L., and RANCK, J. B., “Spatial firing patterns of hippocampal complex-spike cells in a fixed environment.,” *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 7, pp. 1935–50, July 1987.
- [136] MURPHEY, D. K. and MAUNSELL, J. H. R., “Behavioral detection of electrical microstimulation in different cortical visual areas.,” *Current biology : CB*, vol. 17, pp. 862–7, May 2007.
- [137] NAM, Y., BROWN, E. A., ROSS, J. D., BLUM, R. A., WHEELER, B. C., and DEWEERTH, S. P., “A retrofitted neural recording system with a novel stimulation IC to monitor early neural responses from a stimulating electrode.,” *J Neurosci Methods*, vol. 178, no. 1, pp. 99–102, 2009.
- [138] NEWMAN, J. P., *Optogenetic feedback control of neural activity*. PhD thesis, Georgia Institute of Technology, 2013.
- [139] NORMANN, R. A., GREGER, B. A., HOUSE, P., ROMERO, S. F., PELAYO, F., and FERNANDEZ, E., “Toward the development of a cortically based visual neuroprosthesis,” *Journal of Neural Engineering*, vol. 6, no. 3, p. 35001, 2009.
- [140] NOWAK, L. G. and BULLIER, J., “Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter I. Evidence from chronaxie measurements,” *Experimental Brain Research*, vol. 118, no. 4, pp. 477–488, 1998.
- [141] NOWAK, L. G. and BULLIER, J., “Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter II. Evidence from selective inactivation of cell bodies and axon initial segments,” *Experimental Brain Research*, vol. 118, no. 4, pp. 489–500, 1998.
- [142] O’DOHERTY, J. E., LEBEDEV, M. A., LI, Z., and NICOLELIS, M. A., “Virtual active touch using randomly patterned intracortical microstimulation,” *IEEE Trans Neural Syst Rehabil Eng*, vol. 20, no. 1, pp. 85–93, 2012.
- [143] O’DOHERTY, J. E., LEBEDEV, M. A., HANSON, T. L., FITZSIMMONS, N. A., and NICOLELIS, M. A. L., “A brain-machine interface instructed by direct intracortical microstimulation,” *Front Integr Neurosci*, vol. 3, p. 20, 2009.
- [144] O’DOHERTY, J. E., LEBEDEV, M. A., IFFT, P. J., ZHUANG, K. Z., SHOKUR, S., BLEULER, H., and NICOLELIS, M. A. L., “Active tactile exploration using a brain-machine-brain interface.,” *Nature*, vol. 479, pp. 228–31, Nov. 2011.

- [145] OHARA, S., WEISS, N., and LENZ, F. A., “Microstimulation in the region of the human thalamic principal somatic sensory nucleus evokes sensations like those of mechanical stimulation and movement.,” *Journal of neurophysiology*, vol. 91, pp. 736–45, Mar. 2004.
- [146] OLLERENSHAW, D. R., BARI, B. A., MILLARD, D. C., ORR, L. E., WANG, Q., and STANLEY, G. B., “Detection of tactile inputs in the rat vibrissa pathway.,” *Journal of neurophysiology*, vol. 108, pp. 479–90, July 2012.
- [147] PARKER, R. A., DAVIS, T. S., HOUSE, P. A., NORMANN, R. A., and GREGER, B., “Chapter 11 - The functional consequences of chronic, physiologically effective intracortical microstimulation,” in *Brain Machine Interfaces: Implications for Science, Clinical Practice and Society* (JENS SCHOENBORG, M. G. and DANIELSEN, N., eds.), vol. 194 of *Progress in Brain Research*, pp. 145–165, Elsevier, 2011.
- [148] PATEL, S., OHARA, S., DOUGHERTY, P. M., GRACELY, R. H., and LENZ, F. A., “Psychophysical Elements of Place and Modality Specificity in the Thalamic Somatic Sensory Nucleus (Ventral Caudal, Vc) of Awake Humans,” *Journal of Neurophysiology*, vol. 95, no. 2, pp. 646–659, 2006.
- [149] PAXINOS, G. and WATSON, C., *The rat brain in stereotaxic coordinates*. Academic Pr, 2007.
- [150] PENFIELD, W. and JASPER, H. H., *Epilepsy and the functional anatomy of the human brain*. Boston,: Little, [1st ed., 1954.
- [151] PETERSEN, C. C. and CROCHET, S., “Synaptic computation and sensory processing in neocortical layer 2/3,” *Neuron*, vol. 78, no. 1, pp. 28–48, 2013.
- [152] PETERSEN, C. C. H., GRINVALD, A., and SAKMANN, B., “Spatiotemporal dynamics of sensory responses in layer 2/3 of rat barrel cortex measured in vivo by voltage-sensitive dye imaging combined with whole-cell voltage recordings and neuron reconstructions.,” *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 23, pp. 1298–309, Feb. 2003.
- [153] PETERSEN, C. C. H., HAHN, T. T. G., MEHTA, M., GRINVALD, A., and SAKMANN, B., “Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex.,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 100, pp. 13638–43, Nov. 2003.
- [154] PEZARIS, J. S. and ESKANDAR, E. N., “Getting signals into the brain: visual prosthetics through thalamic microstimulation,” *Neurosurgical focus*, vol. 27, no. 1, p. E6, 2009.
- [155] PEZARIS, J. S. and REID, R. C., “Demonstration of artificial visual percepts generated through thalamic microstimulation,” *Proceedings of the National Academy of Sciences*, vol. 104, no. 18, pp. 7670–7675, 2007.

- [156] PINTO, D. J., BRUMBERG, J. C., and SIMONS, D. J., "Circuit dynamics and coding strategies in rodent somatosensory cortex," *J Neurophysiol*, vol. 83, no. 3, pp. 1158–1166, 2000.
- [157] PINTO, L., GOARD, M. J., ESTANDIAN, D., XU, M., KWAN, A. C., LEE, S.-H., HARRISON, T. C., FENG, G., and DAN, Y., "Fast modulation of visual perception by basal forebrain cholinergic neurons.," *Nature neuroscience*, vol. 16, pp. 1857–1863, Oct. 2013.
- [158] PLATT, M. L. and GLIMCHER, P. W., "Neural correlates of decision variables in parietal cortex.," *Nature*, vol. 400, pp. 233–8, July 1999.
- [159] POULET, J. F. A., FERNANDEZ, L. M. J., CROCHET, S., and PETERSEN, C. C. H., "Thalamic control of cortical states.," *Nature neuroscience*, vol. 15, pp. 370–2, Mar. 2012.
- [160] POULET, J. F. A. and PETERSEN, C. C. H., "Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice.," *Nature*, vol. 454, pp. 881–5, Aug. 2008.
- [161] RAMCHARAN, E. J., GNADT, J. W., and SHERMAN, S. M., "Burst and tonic firing in thalamic cells of unanesthetized, behaving monkeys," *Vis Neurosci*, vol. 17, no. 1, pp. 55–62, 2000.
- [162] RANCK, J. B., "Which elements are excited in electrical stimulation of mammalian central nervous system: a review," *Brain Research*, vol. 98, no. 3, pp. 417–440, 1975.
- [163] RASHID, S., PHO, G., CZIGLER, M., WERZ, M. A., and DURAND, D. M., "Low frequency stimulation of ventral hippocampal commissures reduces seizures in a rat model of chronic temporal lobe epilepsy.," *Epilepsia*, vol. 53, pp. 147–56, Jan. 2012.
- [164] REINAGEL, P. and REID, R. C., "Temporal coding of visual information in the thalamus.," *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 20, pp. 5392–400, July 2000.
- [165] ROMO, R., HERNÁNDEZ, A., ZAINOS, A., and SALINAS, E., "Somatosensory discrimination based on cortical microstimulation.," *Nature*, vol. 392, pp. 387–90, Mar. 1998.
- [166] SALZMAN, C. D., BRITTEN, K. H., and NEWSOME, W. T., "Cortical microstimulation influences perceptual judgements of motion direction.," *Nature*, vol. 346, pp. 174–7, July 1990.
- [167] SCHAFER, E. A., "Experiments on the electrical excitation of the visual area of the cerebral cortex in the monkey," *Brain*, vol. 11, no. 1, pp. 1–6, 1888.

- [168] SEKIRNJAK, C., HOTTOWY, P., SHER, A., DABROWSKI, W., LITKE, A. M., and CHICHILNISKY, E. J., “High-resolution electrical stimulation of primate retina for epiretinal implant design,” *The Journal of Neuroscience*, vol. 28, no. 17, p. 4446, 2008.
- [169] SEN, K., JORGE-RIVERA, J. C., MARDER, E., and ABBOTT, L. F., “Decoding Synapses,” *The Journal of Neuroscience*, vol. 16, no. 19, pp. 6307–6318, 1996.
- [170] SHERMAN, S. M., “Tonic and burst firing: dual modes of thalamocortical relay,” *Trends in neurosciences*, vol. 24, pp. 122–6, Feb. 2001.
- [171] SHERMAN, S. M. and GUILLERY, R. W., “Functional organization of thalamocortical relays,” *Journal of Neurophysiology*, vol. 76, no. 3, pp. 1367–1395, 1996.
- [172] SILVERMAN, M. S. and TOOTELL, R. B. H., “Modified technique for cytochrome oxidase histochemistry: increased staining intensity and compatibility with 2-deoxyglucose autoradiography,” *Journal of Neuroscience Methods*, vol. 19, no. 1, pp. 1–10, 1987.
- [173] SIMONS, D. J., “Temporal and spatial integration in the rat SI vibrissa cortex,” *Journal of neurophysiology*, vol. 54, pp. 615–35, Sept. 1985.
- [174] SIMONS, D. J. and CARVELL, G. E., “Thalamocortical response transformation in the rat vibrissa/barrel system,” *Journal of Neurophysiology*, vol. 61, no. 2, pp. 311–330, 1989.
- [175] SONG, D., CHAN, R. H. M., MARMARELIS, V. Z., HAMPSON, R. E., DEADWYLER, S. A., and BERGER, T. W., “Nonlinear modeling of neural population dynamics for hippocampal prostheses,” *Neural networks : the official journal of the International Neural Network Society*, vol. 22, pp. 1340–51, Nov. 2009.
- [176] STANLEY, G. B., LI, F. F., and DAN, Y., “Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus,” *J Neurosci*, vol. 19, no. 18, pp. 8036–8042, 1999.
- [177] STANLEY, G. B., “Reading and writing the neural code,” *Nature neuroscience*, vol. 16, pp. 259–63, Mar. 2013.
- [178] STERN, E., GARCÍA-CRESCIONI, K., MILLER, M. W., PESKIN, C. S., and BREZINA, V., “A method for decoding the neurophysiological spike-response transform,” *Journal of neuroscience methods*, vol. 184, no. 2, pp. 337–356, 2009.
- [179] STONEY JR., S. D., THOMPSON, W. D., and ASANUMA, H., “Excitation of pyramidal tract cells by intracortical microstimulation: effective extent of stimulating current,” *J Neurophysiol*, vol. 31, no. 5, pp. 659–669, 1968.

- [180] STÜTTGEN, M. C. and SCHWARZ, C., “Psychophysical and neurometric detection performance under stimulus uncertainty,” *Nature neuroscience*, vol. 11, pp. 1091–9, Sept. 2008.
- [181] SWADLOW, H. A. and LUKATELA, K., “Cross-correlation and microstimulation: complementary tools in the extracellular analysis of synaptic interactions,” *Journal of Neuroscience Methods*, vol. 64, no. 2, pp. 219–225, 1996.
- [182] TANNAN, V., WHITSEL, B. L., and TOMMERDAHL, M. A., “Vibrotactile adaptation enhances spatial localization,” *Brain research*, vol. 1102, pp. 109–16, Aug. 2006.
- [183] TEHOVNIK, E. J., TOLIAS, A. S., SULTAN, F., SLOCUM, W. M., and LOGOTHETIS, N. K., “Direct and indirect activation of cortical neurons by electrical microstimulation,” *Journal of Neurophysiology*, vol. 96, no. 2, pp. 512–521, 2006.
- [184] TEHOVNIK, E. J., SLOCUM, W. M., SMIRNAKIS, S. M., and TOLIAS, A. S., “Microstimulation of visual cortex to restore vision,” in *Progress in Brain Research* (JOOST VERHAAGEN, E. M. H. I. H. J. W. A. B. G. J. B. and DICK, F. S., eds.), vol. Volume 175, pp. 347–375, Elsevier, 2009.
- [185] TEMEREANCA, S., BROWN, E. N., and SIMONS, D. J., “Rapid changes in thalamic firing synchrony during repetitive whisker stimulation,” *J Neurosci*, vol. 28, no. 44, pp. 11153–11164, 2008.
- [186] TOLIAS, A. S., SULTAN, F., AUGATH, M., OELTERMANN, A., TEHOVNIK, E. J., SCHILLER, P. H., and LOGOTHETIS, N. K., “Mapping Cortical Activity Elicited with Electrical Microstimulation Using fMRI in the Macaque,” *Neuron*, vol. 48, no. 6, pp. 901–911, 2005.
- [187] TOPRANI, S. and DURAND, D. M., “Long-lasting hyperpolarization underlies seizure reduction by low frequency deep brain electrical stimulation,” *The Journal of physiology*, vol. 591, pp. 5765–90, Nov. 2013.
- [188] VAN DER LOOS, H., “Barreloids in mouse somatosensory thalamus,” *Neuroscience letters*, vol. 2, pp. 1–6, Mar. 1976.
- [189] VENKATRAMAN, S. and CARMENA, J. M., “Active Sensing of Target Location Encoded by Cortical Microstimulation,” *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, vol. 19, no. 3, pp. 317–324, 2011.
- [190] VIAENE, A. N., PETROF, I., and SHERMAN, S. M., “Synaptic properties of thalamic input to layers 2/3 and 4 of primary somatosensory and auditory cortices,” *Journal of neurophysiology*, vol. 105, pp. 279–92, Jan. 2011.
- [191] WAGENAAR, D. A., MADHAVAN, R., PINE, J., and POTTER, S. M., “Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation,” *The Journal of Neuroscience*, vol. 25, no. 3, pp. 680–688, 2005.

- [192] WAGENAAR, D. A. and POTTER, S. M., “Real-time multi-channel stimulus artifact suppression by local curve fitting,” *Journal of neuroscience methods*, vol. 120, no. 2, pp. 113–120, 2002.
- [193] WALL, J. T. and KAAS, J. H., “Long-term cortical consequences of reinnervation errors after nerve regeneration in monkeys,” *Brain Research*, vol. 372, no. 2, pp. 400–404, 1986.
- [194] WALLACE, D. J. and SAKMANN, B., “Plasticity of representational maps in somatosensory cortex observed by in vivo voltage-sensitive dye imaging,” *Cerebral Cortex*, vol. 18, no. 6, pp. 1361–1373, 2008.
- [195] WANG, Q., MILLARD, D. C., ZHENG, H. J. V., and STANLEY, G. B., “Voltage-sensitive dye imaging reveals improved topographic activation of cortex in response to manipulation of thalamic microstimulation parameters,” *Journal of Neural Engineering*, vol. 9, p. 059601, Oct. 2012.
- [196] WANG, Q., WEBBER, R. M., and STANLEY, G. B., “Thalamic synchrony and the adaptive gating of information flow to cortex,” *Nature neuroscience*, vol. 13, pp. 1534–41, Dec. 2010.
- [197] WEBBER, R. M. and STANLEY, G. B., “Nonlinear encoding of tactile patterns in the barrel cortex,” *Journal of Neurophysiology*, vol. 91, no. 5, p. 2010, 2004.
- [198] WEBBER, R. M. and STANLEY, G. B., “Transient and steady-state dynamics of cortical adaptation,” *Journal of Neurophysiology*, vol. 95, no. 5, p. 2923, 2006.
- [199] WEBER, D. J., FRIESEN, R., and MILLER, L. E., “Interfacing the Somatosensory System to Restore Touch and Proprioception: Essential Considerations,” *Journal of Motor Behavior*, vol. 44, no. 6, pp. 403–418, 2012.
- [200] WEBER, D. J., LONDON, B. M., HOKANSON, J. A., AYERS, C. A., GAUNT, R. A., TORRES, R. R., ZAAIMI, B., and MILLER, L. E., “Limb-state information encoded by peripheral and central somatosensory neurons: implications for an afferent interface,” *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, vol. 19, no. 5, pp. 501–513, 2011.
- [201] WILSON, B. S., FINLEY, C. C., LAWSON, D. T., WOLFORD, R. D., EDINGTON, D. K., and RABINOWITZ, W. M., “Better speech recognition with cochlear implants,” *Nature*, vol. 352, pp. 236–8, July 1991.
- [202] WILSON, B. S. and DORMAN, M. F., “Cochlear implants: a remarkable past and a brilliant future,” *Hearing research*, vol. 242, pp. 3–21, Aug. 2008.
- [203] WIMMER, V. C., BRUNO, R. M., DE KOCK, C. P. J., KUNER, T., and SAKMANN, B., “Dimensions of a Projection Column and Architecture of VPM and POm Axons in Rat Vibrissal Cortex,” *Cerebral Cortex*, vol. 20, no. 10, pp. 2265–2276, 2010.

- [204] WINER, J. A., MILLER, L. M., LEE, C. C., and SCHREINER, C. E., “Auditory thalamocortical transformation: structure and function,” *Trends in neurosciences*, vol. 28, pp. 255–63, May 2005.
- [205] WONG-RILEY, M., “Changes in the visual system of monocularly sutured or enucleated cats demonstrable with cytochrome oxidase histochemistry,” *Brain Research*, vol. 171, no. 1, pp. 11–28, 1979.
- [206] WOOLSEY, T. A. and VAN DER LOOS, H., “The structural organization of layer IV in the somatosensory region (SI) of mouse cerebral cortex. The description of a cortical field composed of discrete cytoarchitectonic units,” *Brain Res*, vol. 17, no. 2, pp. 205–242, 1970.
- [207] WU, M. C., DAVID, S. V., and GALLANT, J. L., “Complete functional characterization of sensory neurons by system identification,” *Annual Review of Neuroscience*, vol. 29, pp. 477–505, 2006.
- [208] WU, Y. T. and SCLABASSI, R. J., “Identification of nonlinear systems using random amplitude Poisson distributed input functions,” *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 27, no. 2, pp. 222–234, 1997.
- [209] YANG, Y., DEWEESE, M. R., OTAZU, G. H., and ZADOR, A. M., “Millisecond-scale differences in neural activity in auditory cortex can drive decisions,” *Nat Neurosci*, vol. 11, no. 11, pp. 1262–1263, 2008.
- [210] YEOMANS, J. S., *Principles of brain stimulation*. New York: Oxford University Press, 1990.
- [211] YIZHAR, O., FENNO, L. E., DAVIDSON, T. J., MOGRI, M., and DEISSEROTH, K., “Optogenetics in neural systems,” *Neuron*, vol. 71, pp. 9–34, July 2011.
- [212] YIZHAR, O., FENNO, L. E., PRIGGE, M., SCHNEIDER, F., DAVIDSON, T. J., O’/SHEA, D. J., SOHAL, V. S., GOSHEN, I., FINKELSTEIN, J., PAZ, J. T., STEHFEST, K., FUDIM, R., RAMAKRISHNAN, C., HUGUENARD, J. R., HEGEMANN, P., and DEISSEROTH, K., “Neocortical excitation/inhibition balance in information processing and social dysfunction,” *Nature*, vol. 477, no. 7363, pp. 171–178, 2011.
- [213] 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., “Multimodal fast optical interrogation of neural circuitry,” *Nature*, vol. 446, no. 7136, pp. 633–639, 2007.
- [214] ZHOU, Q., GODWIN, D. W., O’MALLEY, D. M., and ADAMS, P. R., “Visualization of calcium influx through channels that shape the burst and tonic firing modes of thalamic relay cells,” *Journal of neurophysiology*, vol. 77, pp. 2816–25, May 1997.

VITA

Daniel Christopher Millard was born in Raleigh, North Carolina, to his father, Thomas Alan Millard, and mother, Rebecca Ann Deal Millard. He was later joined by his brother, Michael Alan Millard, and sister, Amy Marie Millard. During his youth, he spent the majority of his time on the basketball court. However, realizing that a career in the NBA was not in his cards, he decided to attend the University of North Carolina at Chapel Hill to study Biomedical Engineering. He quickly became engrossed in the academic setting, excelling in the classroom and in the laboratory. He performed his undergraduate research under the tutelage of Michael R Falvo, studying the elastic properties of fibrin fibers, the primary mechanical constituent of blood clots. After graduating from UNC-CH, Daniel moved to Atlanta to pursue a PhD in Biomedical Engineering from the Georgia Institute of Technology, and his wonderful girlfriend Darcy Rene Rice, whom he would later marry, came along for the ride. At Georgia Tech, Daniel performed research in the laboratory of Garrett B Stanley, with the goal of understanding how artificial inputs activate the brain, and application of this knowledge toward the development of sensory prostheses and mapping neural circuits in the brain. At the time of writing, he and his wife are expecting their first child, Theodore Burke Millard, and look forward to a life full of love, happiness, and hard work.