

CLOSED LOOP OPTOGENETIC CONTROL AND THALAMIC STATE

A Dissertation
Presented to
The Academic Faculty

By

Michael F Bolus

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Wallace H. Coulter Department of Biomedical Engineering

Georgia Institute of Technology
Emory University

August 2020

Copyright © Michael F Bolus 2020

CLOSED LOOP OPTOGENETIC CONTROL AND THALAMIC STATE

Approved by:

Dr. Garrett B. Stanley, Advisor
Wallace H. Coulter Department of
Biomedical Engineering
Georgia Institute of Technology

Dr. Christopher J. Rozell
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Dr. Magnus B. Egerstedt
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Dr. Robert E. Gross
Department of Neurosurgery
Emory University

Dr. Peter A. Wenner
Department of Physiology
Emory University

Date Approved: July 17, 2020

Carissimae familiae sine qua, quamquam beatus, ego nihil sum.

ACKNOWLEDGEMENTS

Throughout my graduate studies I have been buoyed by a large support system of individuals, without whom I would likely not be here today. For the past six years, I have had the fortune to work with Dr. Garrett Stanley, who has not only advised my doctoral project but has also become a great mentor, friend, and an unrelenting source of positivity. I am similarly indebted to Dr. Chris Rozell for his guidance and support over the years. Throughout my time here, both he and Dr. Stanley have always shown great interest in my work, which encouraged me to keep going. Similarly, I must acknowledge the input and support I have received from the rest of my thesis committee. I have known Dr. Robert Gross since I started here, and his interest in my work as well as his clinical perspective has always been a source of great encouragement. Similarly, I first met Dr. Magnus Egerstedt early in my graduate studies, and discussions with him over the years often gave me a fresh perspective on the problems at hand. Last but certainly not least, while I only met him a couple of years ago, I am similarly indebted to Dr. Peter Wenner for his enthusiasm for my work, his expertise, and his advice on prioritizing pushing the techniques to the awake animal. It is certainly not lost on me that in addition to six years of my time and effort, this graduate degree took quite a lot of theirs as well. As busy professors with their own research groups and other responsibilities, I am truly thankful for their support.

Over the years, I have also enjoyed immense support from the members of the Stanley Laboratory, both in research but also in friendship. Namely, I would like to thank Dr. Clarissa Whitmire, Dr. Christian Waiblinger, Dr. Aurélie Pala, Dr. Nathaniel Caleb Wright, Dr. Audrey Sederberg, Dr. Peter Borden, Alex Ortiz, Adam Willats, Yi Juin Liew, and Elaida Dimwamwa. They all have helped me in too many ways to recount, and sharing this time with them has been a blessing. I cannot imagine a better group of people with whom to have spent the past six years.

In addition to the members of the Stanley Lab, my work could not have been carried

out without the help of the Physiological Research Lab and its wonderful staff. While some of the faces have changed over the years, I have known many of them my entire time here. They have not only provided great care for the animals, but they also took great care of me through instruction and friendly conversation. In addition, I must acknowledge the immense amount of help I have received from Kristen Laquidara over the past couple of years. As it came time for me to propose and then later defend my doctoral research, she patiently fielded question after question and has helped guide me to this end.

Finally, I would like to thank my family. My parents, David and Sallie, have been great role models whom I have and will always look up to and can only hope to measure up to. The extent to which they put others first at their own expense always amazes me. Without their love and support over the years, I would certainly not have made it to this point. Similarly, my Sitty has gone above and beyond the call of duty for any grandmother by keeping her doors open to me as well as providing seemingly endless amounts of food. In addition, my brothers have been a constant source of support and encouragement: David, Jr., always provides me inspiration both through his actions and his words; Vincent always finds a way to make me smile; and Nicholas has been my constant life companion, with whom I have had the pleasure to share the same roof for more than 26 years. I am thankful to have spent six more years with my best friend, and I would not be here without him or any one of them.

TABLE OF CONTENTS

Acknowledgments	iv
List of Figures	xi
Chapter 1: Introduction and Background	1
1.1 Thalamocortical sensory pathways	2
1.2 Rodent vibrissa system	5
1.3 Thalamic state	6
1.3.1 Thalamic, cortical, and behavioral correlates of brain state	6
1.3.2 State-dependent sensory processing and factors affecting state	8
1.3.3 Thalamic state in disease: Generalized epilepsy	9
1.4 Closed-loop control of neuronal activity	9
1.4.1 Modeling neuronal systems for control and estimation	11
1.4.2 Optogenetics	15
1.4.3 Optogenetic control	17
1.5 Overview	20
Chapter 2: Model-free feedback optogenetic control for entraining time-varying patterns of neural activity	21
2.1 Introduction	21

2.2	Methods	23
2.2.1	Experimental preparation	24
2.2.2	Reference trajectories	26
2.2.3	Firing rate estimator design	27
2.2.4	Linear-nonlinear Poisson model	29
2.2.5	Controller design	31
2.2.6	Open-loop stimulus design	33
2.2.7	Disturbance	34
2.2.8	Offline firing rate estimation	34
2.2.9	Fano factor	35
2.3	Results	35
2.3.1	Examples of closed- vs. open-loop stimulation	36
2.3.2	Firing rate estimator design	38
2.3.3	Controller design	42
2.3.4	Model accuracy and closed-loop performance	44
2.3.5	Robustness of control to model inaccuracy	46
2.3.6	Control summary: Tracking a sinusoidal trajectory	50
2.3.7	Control summary: Tracking a non-sinusoidal trajectory	52
2.4	Discussion	54
2.5	Conclusions	59

Chapter 3: State-space model-based feedback control of optogenetically driven neural activity 61

3.1	Introduction	61
-----	------------------------	----

3.2	Methods	63
3.2.1	Animal preparation	63
3.2.2	Experimental setup	64
3.2.3	Offline spike sorting	65
3.2.4	Mathematical modeling	66
3.2.5	Estimator	70
3.2.6	Controller	73
3.2.7	Performance measures	77
3.3	Results	78
3.3.1	GLDS captures optical noise-driven responses	82
3.3.2	Parameter-adaptive Kalman filtering provides robust online estimation	85
3.3.3	State-space control performs well in SISO clamping applications	89
3.3.4	Multi-electrode recordings reveal effects of SISO control on simultaneously recorded neurons	92
3.3.5	GLDS models generalize to multi-output datasets	96
3.3.6	Online estimation methods generalize to multi-output applications	99
3.3.7	Simulated SIMO control more robust to population heterogeneity than a SISO “antenna” approach	101
3.4	Discussion	104
3.5	Conclusions	109

Chapter 4: Optogenetic control of thalamic state and its effects on thalamocortical activity 110

4.1	Introduction	110
-----	------------------------	-----

4.2	Methods	112
4.2.1	Animal preparation	112
4.2.2	Experimental setup	114
4.2.3	Quantification of whisker motion	115
4.2.4	Thresholded multi-unit activity	115
4.2.5	Offline spike sorting	116
4.2.6	Cortical local field potential	116
4.2.7	Thalamic state estimation	117
4.2.8	Cortical LFP frequency ratio	119
4.2.9	Closed- and open-loop optogenetic control	119
4.2.10	Dataset inclusion criteria	120
4.3	Results	120
4.3.1	Optogenetic control and multi-unit thalamic activity	121
4.3.2	Thalamic state estimation	124
4.3.3	Effect of optogenetic control on thalamic state	127
4.3.4	Whisking as a natural thalamic state disturbance	129
4.3.5	Closed-loop optogenetic stimulation resists whisking thalamic state disturbance	132
4.3.6	Optogenetic control of thalamus pushes cortex towards activated state	135
4.3.7	Effect of thalamic state elevation on thalamic and cortical sensory responses	138
4.4	Discussion	144
4.5	Conclusions	146

Chapter 5: Conclusions and Future Directions	148
5.1 Robustness of closed-loop optogenetic control	148
5.2 Variability and feedback control of single-neuron vs. population activity	150
5.3 Opportunities for dual-opsin bidirectional control	150
5.4 Choice of mathematical model for optically-driven neural activity	152
5.5 Trajectory tracking vs. constant set point	154
5.6 Dynamical systems approach to brain state definition and estimation	154
5.7 Applicable measurement modalities for CLOC	156
5.8 Thalamic state and burst vs. tonic firing modes	157
5.9 Across-region feedback control	158
5.10 Closing the loop around behavior	158
5.11 Closed-loop optogenetic control for testing role of thalamic state in seizure susceptibility in absence epilepsy	159
Appendix A: Non-sinusoidal reference firing rate: Measured thalamic firing rate modulation	161
Appendix B: Inspection of electrical recordings for optical stimulation artifacts	163
Appendix C: Targeting optogenetics to VPM/VPL: Ai32;NR133	166
References	187

LIST OF FIGURES

1.1	Thalamocortical pathways	4
1.2	Rodent Whisker System	5
1.3	Coordinated thalamic, cortical, behavioral states	7
1.4	Modeling spiking responses to stimulation	14
1.5	Optogenetics	16
1.6	Closed-loop optogenetic control in vitro	18
2.1	Closed loop optogenetic control of firing rate	24
2.2	Closed- vs. open-loop optogenetic control of dynamic firing rate trajectories	39
2.3	Firing rate filter design: Choosing filter bandwidth	40
2.4	Controller Design: Tuning the controller around an LNP model neuron . . .	43
2.5	LNP Model Performance: Open-loop vs. closed-loop	45
2.6	Robustness of Control to Model Inaccuracy	48
2.7	Sinusoidal tracking performance	51
2.8	Non-sinusoidal tracking performance	53
3.1	Closed-loop optogenetic control using state-space linear dynamical systems models	80
3.2	State-space models of SISO optogenetic responses	83

3.3	Kalman filtering for online estimation in SISO applications	88
3.4	Experimental SISO control and estimation	91
3.5	Effects of SISO control on local population	95
3.6	State-space models of SIMO optogenetic responses	98
3.7	Kalman filtering for online estimation in SIMO applications	100
3.8	Simulated SIMO control and estimation	102
4.1	Closed-loop optogenetic control and multi-unit activity	123
4.2	Estimating thalamic state	126
4.3	Effect of optogenetic control on thalamic state	128
4.4	Whisking as a natural thalamic state disturbance	131
4.5	Closed-loop optogenetics resists whisking thalamic state perturbation	133
4.6	Effect of thalamic optogenetic stimulation on cortical activity	137
4.7	Effect of pre-stimulus thalamic activation on thalamic sensory responses . .	141
4.8	Effect of pre-stimulus thalamic activation on cortical sensory responses . .	143
A.1	Firing rate modulation during putative whisking	162
B.1	Inspection of electrical recordings for optical stimulation artifacts	164
B.2	Inspection of spiking activity for optical stimulation artifacts	165
C.1	Coronal section of Ai32;NR133 mouse	167

SUMMARY

Dating as far back as the eighteenth century with Luigi Galvani's seminal studies in bioelectricity ([1]) and beyond, interfacing with the nervous system at fast timescales has proven invaluable for scientific investigation as well as clinical interventions in diseases such as Parkinson's [2]. Until somewhat recently, electrical stimulation has been the most common technique for neuronal control at fast timescales. Over the past fifteen years, the advent of optogenetics, a technique whereby optical excitation or inhibition of neural activity can be targeted genetically, has ushered in a new wave of experimental approaches to dissecting circuit function ([3, 4]). To date, most optogenetic control of neural activity has been limited to open-loop stimulation or event-triggered closed-loop stimulation, in which previously-determined optical inputs are used. However, activity in the brain changes in a state-dependent fashion such that stimuli (*e.g.*, sensory inputs from the periphery) can elicit variable responses and ultimately manifest in variable percepts. Given the ever-changing nature of neuronal activity and the ability to simultaneously record and stimulate targeted cell types with optogenetics, there is growing interest in the intersection of feedback control and optogenetics ([5]). In contrast to open-loop stimulation, closed-loop optogenetic control seeks to achieve target activity by updating optical input as a function of recorded neuronal activity. In this thesis, engineering approaches to feedback control and state estimation are used to tackle the problems of controlling neuronal firing activity *in vivo*, with the goal of developing a set of methods that are general enough that they may be applied to manipulation of other types of neuronal activity or even animal behavior. Specifically, we apply closed-loop optogenetic control (CLOC) to manipulate the thalamus, a deep brain region that serves as a central gateway for conducting sensory information to the cerebral cortex. Given the importance of brain state in health and disease, we end by investigating the effects of optogenetic control on the state of the thalamus and its implications for sensory response properties in the somatosensory thalamocortical pathway. First, Chapter 2

develops a design methodology for using a previously described model-free optogenetic control scheme ([6]) to entrain naturally relevant patterns of rate modulation such as observed in the rodent somatosensory thalamus during active movement of facial whiskers, rather than holding firing rate constant over long timescales as was done previously. In order to ensure the optogenetic control scheme generalizes more gracefully to future multi-input/multi-output control problems, Chapter 3 applies state-space model-based control and estimation to the problem of manipulating thalamic firing rates. Importantly, we investigate the effectiveness of CLOC in the awake animal for the first time, as well as the response of local populations of neurons to optical stimulation rather than recording from single neurons at a time. Finally, in Chapter 4 we investigate the effect of CLOC on thalamic state more generally, analyze the robustness of control to a naturally-occurring disturbance (animal movement), and look at its consequences for downstream cortical activity and sensory response characteristics. As part of this analysis, a broadly-applicable state-space model based notion of thalamic state is put forth, marrying previously distinct neuroscientific and engineering notions of “state”.

CHAPTER 1

INTRODUCTION AND BACKGROUND

Dating as far back as the eighteenth century with Luigi Galvani's seminal studies in bioelectricity ([1]) and beyond, interfacing with the nervous system at fast timescales has proven invaluable for scientific investigation as well as clinical interventions in diseases such as Parkinson's [2]. Until somewhat recently, electrical stimulation has been the most common technique for neuronal control at fast timescales. Over the past fifteen years, the advent of optogenetics, a technique whereby optical excitation or inhibition of neural activity can be targeted genetically, has ushered in a new wave of experimental approaches to dissecting circuit function ([3, 4]). To date, most optogenetic control of neural activity has been limited to open-loop stimulation or event-triggered closed-loop stimulation, in which previously-determined optical inputs are used. However, activity in the brain changes in a state-dependent fashion such that stimuli (*e.g.*, sensory inputs from the periphery) can elicit variable responses and ultimately manifest in variable percepts. Given the ever-changing nature of neuronal activity and the ability to simultaneously record and stimulate targeted cell types with optogenetics, there is growing interest in the intersection of feedback control and optogenetics ([5]). In contrast to open-loop stimulation, closed-loop optogenetic control seeks to achieve target activity by updating optical input as a function of recorded neuronal activity. In this thesis, engineering approaches to feedback control and state estimation are used to tackle the problems of controlling neuronal firing activity *in vivo*, with the goal of developing a set of methods that are general enough that they may be applied to manipulation of other types of neuronal activity or even animal behavior. Specifically, we apply closed-loop optogenetic control to manipulate the thalamus, a deep brain region that serves as a central gateway for conducting sensory information to the cerebral cortex. Given the importance of brain state in health and disease, we end by investigating

the effects of optogenetic control on the state of the thalamus and its implications for sensory response properties in the somatosensory thalamocortical pathway. What follows is the pertinent background necessary for understanding this work, including the anatomical pathway and model system used for these studies, the state-dependent nature of neuronal activity, the types of mathematical models commonly used to represent these systems, and, finally, previous work in the area of optogenetic control.

1.1 Thalamocortical sensory pathways

We rely heavily on our senses for everything ranging from navigation of environments to communication with those around us. Nestled deep in the brain, the dorsal thalamus serves as the information conduit to cerebral cortex in almost all sensory pathways (Figure 1.1(a)). In the context of vision, somatosensation, and audition, sensory signals travel from peripheral sensory organs through the brainstem to the corresponding primary relay nuclei of the dorsal thalamus, the lateral geniculate nucleus (LGN), ventrobasal complex (VB), and medial geniculate nucleus (MGN), respectively, before making their way to primary sensory cortices. The VB of somatosensory thalamus is composed of the ventral posterolateral (VPL) and ventral posteromedial (VPM) nuclei, which receive somatic (*i.e.*, spinothalamic) sensory input and orofacial input conveyed via the trigeminal nerve, respectively ([7]). From these primary thalamic relay nuclei, thalamocortical (TC) neurons project to primary visual, somatosensory, and auditory cortices, chiefly synapsing in the so-called input layer of cortex, layer 4 (LIV, Figure 1.1(b)). Conversely, corticothalamic (CT) cells in the deepest layer of cortex (LVI) project back to these relay nuclei. In addition, TC cells project to the purely inhibitory (*i.e.*, GABA-ergic) reticular nucleus of the thalamus (RT), which provides inhibitory feedback to TC cells as well as disynaptic feedforward inhibition from CT cells ([8]). Together, these constituents make up the thalamocortical loop.

Aside from its position in the pathway as the final way-station before cortex, the thalamus and thalamocortical loop more generally have a number of physiological properties

that make it a gatekeeper for information flow. At one extreme, in the context of slow wave sleep, thalamocortical oscillations serve to block the propagation of sensory inputs onward to cortex ([9, 10]). This is achieved thanks to the thalamo-reticular loop and the expression of low threshold voltage-gated T-type calcium channels, which, as the name suggests, are activated at low membrane potentials, and allow a slow Ca^{++} current, on top of which multiple action potentials may fire in a rapid succession, a phenomenon known as “bursting”. In slow wave sleep, this hyperpolarization followed by bursting happens rhythmically, effectively blocking sensory responses ([9]). While T-type Ca^{++} channel-mediated bursting does occur in wakefulness, the occurrence of burst action potentials is more rare ([10, 11, 12, 13]); however, when they do occur and are driven by sensory inputs, thalamic bursting results in a salient drive to cortex sometimes termed a “wakeup call” ([14, 15]). This large response is thought to be an optimal coding scheme for binary detection of the presence of a simple stimulus ([16]) or features in complex stimuli such as edges in a scene ([17]). Just as bursting of a single neuron provides multiple inputs to cortex in rapid succession, the degree of synchronization of multiple thalamic neurons in their response to stimulation results in more or less potent drive to cortex ([18, 19, 16]). Moreover, [15] found that the “wake up call” behavior observed with bursting also held true for single action potentials that were preceded by periods of thalamic silence. This is because the thalamocortical synapse is susceptible to depression, such that when the thalamus is actively firing preceding a stimulus, the response downstream is less salient ([20]), but less binary and better suited for discrimination of graded features of a stimulus such stimulus strength or position in space ([21, 22]). Clearly, the activity of the thalamus at the time of sensory stimulation can have profound effects on sensory information conveyed downstream and what is ultimately perceived.




Figure 1.1: **Thalamocortical pathways.** (a) Thalamus serves as a central conduit for sensory information processing in almost all sensory pathways. (b) The thalamocortical loop is composed of thalamocortical relay nuclei (*e.g.*, lateral geniculate nucleus, ventrobasal complex, medial geniculate nucleus), reticular thalamic nucleus (RT), and the cerebral cortex. Adapted from [8].

1.2 Rodent vibrissa system

Rodents rely on their large facial whiskers to navigate through and investigate their environments. Along the so-called lemniscal pathway, sensory information is conveyed from the whiskers to the brainstem, VPM thalamus and eventually primary somatosensory cortex (S1) in a discrete, topographically-aligned fashion, such that there are populations of neurons at each level that respond principally to a single whisker (Figure 1.2) ([23]). Dense clusters of LIV cells result in a distinct “barrel” appearance to the cortical column, giving this region of S1 its name “Barrel Cortex” ([24, 25]). Experimentally, this model system provides unparalleled control over sensory inputs and the ability to trace information flow through the primary sensory pathway. In contrast, the medial posterior nucleus (POm), a higher order thalamic nucleus in the somatosensory pathway that is immediately adjacent to the VPM (Figure 1.2(b)), receives more spatially diffuse whisker inputs ([23]) and displays distinct sensory response and movement-related properties ([26, 27]).

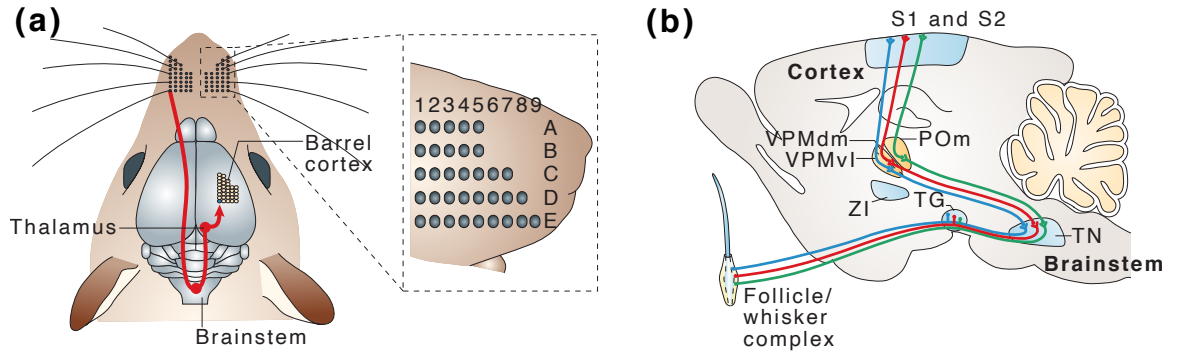


Figure 1.2: Rodent Whisker System. (a) Topographically aligned, discrete populations of cells primarily receptive to single principal whisker from the face, brainstem, thalamus, to cortex where dense clusters of LIV cells form the “Barrel” cortex. Adapted from [23]. (b) Sagittal view of parallel pathways from whisker to cortex. The primary lemniscal (red) pathway to cortex is from the trigeminal nucleus of the brainstem, through the dorsomedial aspect of the VPM (VPMdm), to S1. Adapted from [23].

1.3 Thalamic state

Throughout this work, the term “state” will be used multiply to describe neuronal responses to inputs mathematically (Chapter 3) as well as to describe notions of brain state used in neuroscience (Chapter 4). In engineering, the state of a dynamical system is the variable or collection of variables whose knowledge, along with knowledge of any stimuli/covariates, is sufficient to predict future activity ([28]). While not normally thought of in this way, the neuroscientific attempts at labeling neuronal (*e.g.*, thalamic, cortical) and behavioral activity as readouts of underlying “state” share this same goal: extracting a small number of signals that can, for example, help predict otherwise variable responses to sensory stimuli or that can serve as “biomarkers” for the onset of disease conditions like seizures in epilepsy. Therefore, while the manner in which these notions of state are derived and discussed may vary, they are actually consistent in spirit. To this end, an attempt has been made in Chapter 4 to marry these mathematical and neuroscientific notions of state in the context of changing thalamic activity.

1.3.1 Thalamic, cortical, and behavioral correlates of brain state

Thalamic and cortical activity is state-dependent, ranging from “Up” and “Down” cortical activity and thalamocortical spindle oscillations during slow wave sleep and under certain anesthetics to rapid, small variations in cortical subthreshold membrane potential and elevated thalamic firing rates during periods of active arousal (for reviews, [29, 30, 31, 32]). Previous work defines the state of the brain during wakefulness in discrete extremes such as quiet vs. active/aroused ([33, 34, 35, 27]), inattentive vs. alert [36], or passive vs. engaged in a task ([37]). Despite this discretization, wakefulness is likely a continuum ([29]), albeit with some well-studied extrema. On one end of the spectrum, when animals are quiet/inactive the neuronal activity of cortex as measured by membrane potential or extracellular modalities like local field potential (LFP) or electroencephalog-

raphy (EEG) is characterized by large amplitude, low frequency fluctuations in activity, whereas the thalamic activity is characterized by low firing rates and some degree of burst firing (Figure 1.1(a), “inattentive”; (b) “quiet”). At the other end of this spectrum, in an active state, the cortex is characterized by lower amplitude, high-frequency fluctuations in subthreshold potential, while the thalamus exhibits elevated tonic firing (Figure 1.1(a), “alert”; (b) “active”/whisking). Such state changes can also be behaviorally measured by factors such as pupil diameter ([36, 38, 39, 40]) and movement including locomotion ([41]) and, in the context of the rodent somatosensory system, whisking ([33]).

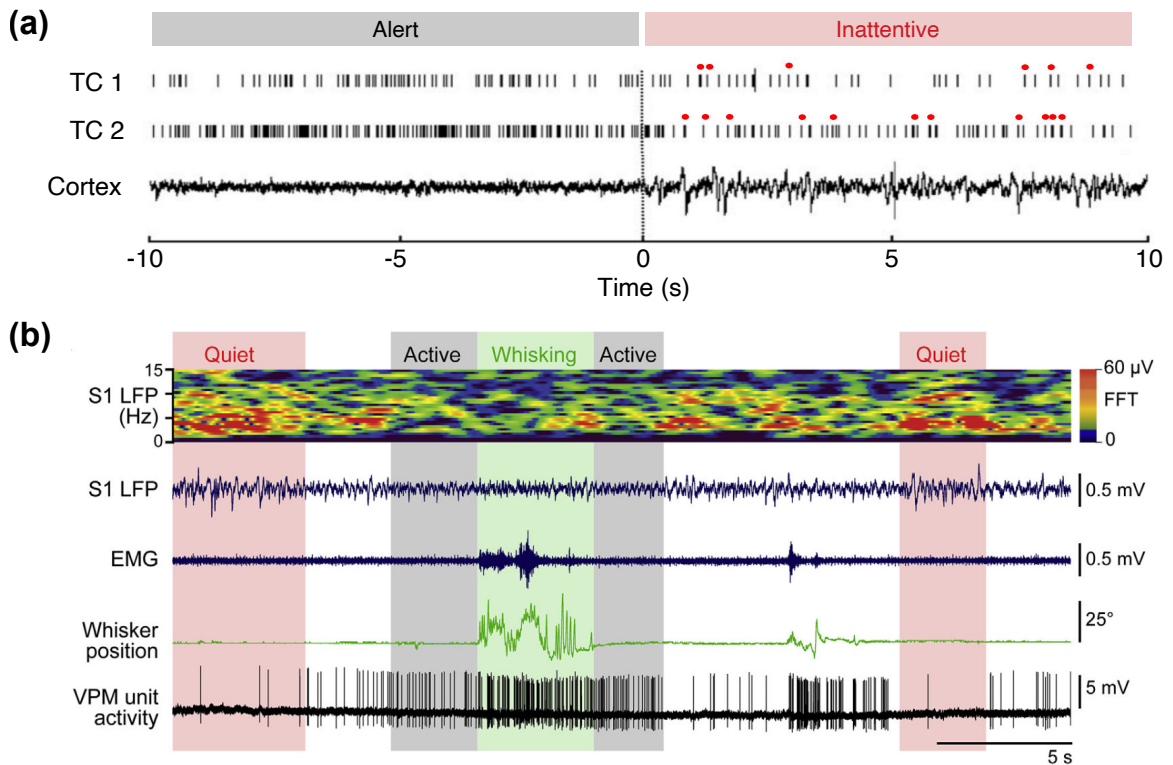


Figure 1.3: Coordinated thalamic, cortical, behavioral states. (a) Transition from “alert” to “inattentive” states in the visual thalamocortical (TC) cells and cortex (LFP) of an awake rabbit. Red dots denote bursts of action potentials. Adapted from [36]. (b) Transition between “quiet” and “active” states across somatosensory thalamus and cortex in an awake, head-fixed mouse. VPM unit activity recorded juxtacellularly; whisker position measured using high-speed videography; electromyography (EMG) recorded from the neck muscles; LFP recorded in S1 and its frequency spectrum analyzed as a function of time. Changes in thalamic firing and cortical LFP spectrum correlate with whisking behavior. Adapted from [27].

1.3.2 State-dependent sensory processing and factors affecting state

While the nature of the effects are apparently different in the visual versus auditory and somatosensory pathways, these state changes in thalamus and cortex have been shown to have consequences for sensation both in terms of neuronal responses as well as in the perception of sensory stimuli across modalities ([42, 41, 43, 44, 37, 38, 45, 46, 47, 48, 49]). In the case of vision, periods of active states of wakefulness coinciding with animal locomotion result in enhanced sensory response gain and signal-to-noise ratios in primary visual cortex and improved stimulus detection performance ([41, 43]). Conversely, in the case of both somatosensation and audition, it appears that a quiet state of wakefulness is optimal for stimulus detection. McGinley *et al* ([38]) report that detection of an auditory tone in noise was best when the stimulus was preceded by quiet, stably hyperpolarized cortical activity and that “hyper-vigilance” that occurred during locomotion was deleterious. Similarly, it has been reported that when animals are engaged in an auditory task, thalamic firing rates are elevated and cortical responses to stimuli are suppressed ([37]). Consistent findings have been seen in the rodent whisker system, in that an active state associated with elevated thalamic firing and movement of whiskers has been shown to reduce cortical sensory responses and reduce detection of a stimulus ([20, 50, 49]).

These states are likely governed by changing tone of neuromodulators such as acetylcholine and norepinephrine in thalamus and cortex ([20, 51, 52, 38, 39]); however, it has been shown that direct manipulations of thalamic firing have been sufficient to control cortical state measures ([51, 35]). Conversely, it is also known that corticothalamic feedback to thalamus is able to modulate thalamic activity [53, 54, 55], such that the cortex could “wake up”, or activate, a quiet thalamus in a similar manner. Whatever the source, changes in thalamic firing rate over longer timescales seem, if not in control, at least predictive of cortical state. Similarly, because of the dependence of TC bursting mechanism on prolonged quiet and the dependence of TC synapse depression on baseline activity, thalamic firing rate is also predictive of changes in sensory response characteristics. Taken

together with the engineering definition of “state” above, these factors suggest that coordinated change in thalamic population firing rate is an ideal candidate for the state of the thalamus (Chapter 4).

1.3.3 Thalamic state in disease: Generalized epilepsy

Beyond its role as a gatekeeper of sensory information, the thalamus may also serve as an effective choke point in generalized epilepsies ([56]). In a photothrombotic ischemic stroke rat model of post-stroke generalized epilepsy, Paz *et al* ([57]) demonstrated that silencing the thalamus at seizure onset effectively aborted it. Similarly, in two genetic rodent models of absence epilepsy (GAERS rats, stargazer mice) it was found that preventing bursts by depolarization of TC neurons halted thalamocortical oscillations characteristic of the absence seizure ([58]). Finally, there is evidence that changing thalamic state may be an effective biomarker for prediction of absence seizures in that there is a dip in thalamic firing rate before seizure onset (GAERS rat, [59]; Scn8a-mutant mice, [60]).

1.4 Closed-loop control of neuronal activity

The ability to control neuronal activity has deep clinical and scientific significance, ranging from treating disorders such as epilepsy to understanding fundamental operating principles of single cells and intricately interconnected networks such as the thalamocortical loop. Indeed, it was the use of a feedback controller—Cole’s voltage clamp—that enabled the pioneering studies of Hodgkin and Huxley and gave rise to our understanding of the ionic currents underlying the action potential [61]. The key to this experimental work was the use of a feedback controller to “clamp” the trans-membrane voltage by injecting current to counter-act naturally occurring changes in ionic currents. This functional decoupling of constituent ionic and capacitive currents led to a quantitative description of the nonlinear dynamics of the action potential.

Today, the neuroscience community is faced more acutely than ever with the task of

dissecting the functions of neurons in interconnected networks and in the context of ever-changing activity during wakefulness. Taking the thalamocortical sensory system as an example, thalamic activity is capable of controlling the ongoing state of cortex as well as its responses to sensory stimuli, and direct neuromodulatory input from the brainstem is able to do the same (Section 1.3.2). Therefore, under normal conditions it is difficult to ascertain the relative roles of thalamic vs. intracortical origins of sensory response variability. A control methodology that could keep the circuit intact while holding components of the circuit invariant to spontaneous changes would be invaluable in this context as well as the ubiquitous problem in neuroscience of decoupling linked variables.

Closed-loop, or feedback, control methods represent an engineering solution to such problems. As opposed to a so-called open-loop control approach where stimulation is designed either through previous experimentation or using mathematical models, closed-loop systems make use of feedback to adjust stimulation in realtime by comparing the value of a measured/estimated signal to a desired setpoint and taking action according to some control logic. General reasons to use closed-loop control may include achieving a target or target trajectory despite model inaccuracy (*e.g.*, Chapters 2-3), maintaining a system at target activity in the face of disturbances (*e.g.*, voltage clamp, Chapter 4), or minimizing energy usage (*e.g.*, closed-loop deep brain stimulation). The control logic may be as simple as toggling control on or off on either side of a setpoint (on-off control), or it may use the error signal to grade stimulation more continuously. Hereon, the terms closed-loop or feedback control will be used to refer specifically to the latter.

In general, such closed-loop systems are composed of multiple elements: sensors and estimators that measure outputs of the system being manipulated and estimate variables important for control; feedback and feedforward controller gains that use estimates of current system state/output and target activity to generate a control signal; and actuators that physically interact with the system to effect control. Sensors for measuring neural activity range from intra- or extracellular electrical recordings by way of single ([15, 42, 19, 27, 62])

or arrayed electrodes ([20, 63, 64]) to widefield or cellular-resolution optical microscopy using (genetically-encoded) fluorescent indicators of voltage and calcium signals ([65, 50, 66, 21, 67, 68, 69, 70]). On the other end of the feedback loop, for control of neural activity at fast timescales, actuation is often achieved by electrical stimulation (*e.g.*, [71, 72, 20, 73]). Over the past 15 years, however, optogenetic approaches to stimulation [74, 75, 76] have seen widespread adoption and use in neuroscience due to the ability to excite or inhibit neural activity in cell-type specific manner (Section 1.4.2).

1.4.1 Modeling neuronal systems for control and estimation

Whether used simply as design tools (Chapter 2) or online during implementation of a control scheme (Chapter 3), mathematical models that describe the relationship between control signal(s) and the controlled variables are invaluable for feedback control. While there are not agreed-upon canonical models in neuroscience such those for describing mechanical systems, there are a set of model types which have seen varying levels of use in describing neuronal responses to stimulation. What follows is a brief overview that emphasizes phenomenological model types that have been specifically used to describe spiking responses, since the control objective in this work will be related to thalamic firing.

Linear receptive field, linear-nonlinear-Poisson and generalized linear models of stimulus-driven neuronal spiking

In sensory systems, feed-forward (*i.e.*, non-dynamical) models have long been used to relate stimuli to neuronal spiking. The simplest form is the linear “receptive field” model ([77, 78, 79, 80]). Often identified using variations on spike-triggered analysis ([81, 82, 83]), they comprise spatiotemporal finite impulse response (FIR) filter(s) that capture features in stimuli to which neurons are “tuned” to respond (“stimulus filter”, Figure 1.4(a)). To respect the statistical nature of spiking measurements, a simple extension to the linear receptive field model is the addition of rectifying nonlinearity and a Poisson assumption

on spike generation, yielding the linear-nonlinear-Poisson (LNP) model (Figure 1.4(a)) ([81, 84, 85, 86, 87, 88, 89, 82]). These models are further extended to include multi-neuron interactions (through “coupling filters”) and a spike history filter to account for the non-Poisson refractory nature of neurons, resulting in the generalized linear model (GLM) (Figure 1.4(b)) ([90, 91]). As the simplest of the spiking models, Chapter 2 will use the LNP to design a feedback control loop for manipulation of firing rate *in vivo*. Contrary to transfer function and state-space models, LNP models have no dynamics and, while they are widely used in neuroscience to describe spiking responses to stimuli, they are not as amenable to common control design and implementation (although, see [92] for example simulated control of GLM).

Dynamical systems state-space models

While the above model types have seen widespread use in describing neuronal responses to stimulation, models with dynamics (*e.g.*, transfer functions or state-space models) are widely used for design and implementation of controllers ([28]). At one extreme, neuroscientists have developed nonlinear dynamical systems models that describe the evolution of membrane potential using differential equations: notably, the seminal work of Hodgkin Huxley ([61]). A simpler, albeit more abstract, model is the leaky-integrate-and-fire (LIF, [94, 93, 95]), in which membrane potential is modeled using a first order differential equation, and action potential generation is a simple spike and post-spike reset following threshold-crossing (Figure 1.4(c)). Notably, LIF models with multiple outputs have independent, rather than shared, dynamics (*e.g.*, [95]), with coupling occurring through current injection as a result of simulated synaptic transmission.

The final form of model considered here is the state-space dynamical system, which is not as commonly used for describing neuronal response to stimulation. That said, variations of this model type are seeing increasing use in neuroscience, albeit mostly applied to motor systems ([96, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106, 64, 107]). In con-

trast to LIF and GLM models, dynamical system model outputs are driven by a common set of underlying latent dynamics, which in the state-space form are captured by a first order differential/difference equation of a multi-dimensional “state” variable, \mathbf{x} . In general, these dynamics could be nonlinear; however, this work (Chapters 3-4) will be limited to linear systems (Figure 1.4(d)). The measured outputs (here, spiking activity) are assumed to be observed either after corruption with additive Gaussian noise (Figure 1.4(d), top) or after element-wise rectification and Poisson spike generation (bottom). These model types are termed Gaussian or Poisson linear dynamical systems ([101, 102]): GLDS and PLDS, respectively. Note that the GLDS and PLDS are dynamical systems equivalents of the linear receptive field and the LNP models above. Given the spiking nature of the measured neuronal activity, PLDS models are more widely used than GLDS ([97, 101, 102]), and estimators equivalent to the Kalman filter for linear systems ([108]) have been developed that allow online estimation of this latent state variable, \mathbf{x} , under the Poisson output assumption ([96, 98, 109, 110]). However, many commonly used linear control and estimation approaches revolve around the use of transfer functions or state-space GLDS models ([28]). Moreover, the state-space model structure in combination with optimal control provide generality to multi-input/multi-output control and estimation problems. Therefore, the use of a GLDS model rather than the more common stimulus-filter-based models above will be evaluated in Chapter 3, as well as the necessity of a spiking nonlinearity (LNP/PLDS).

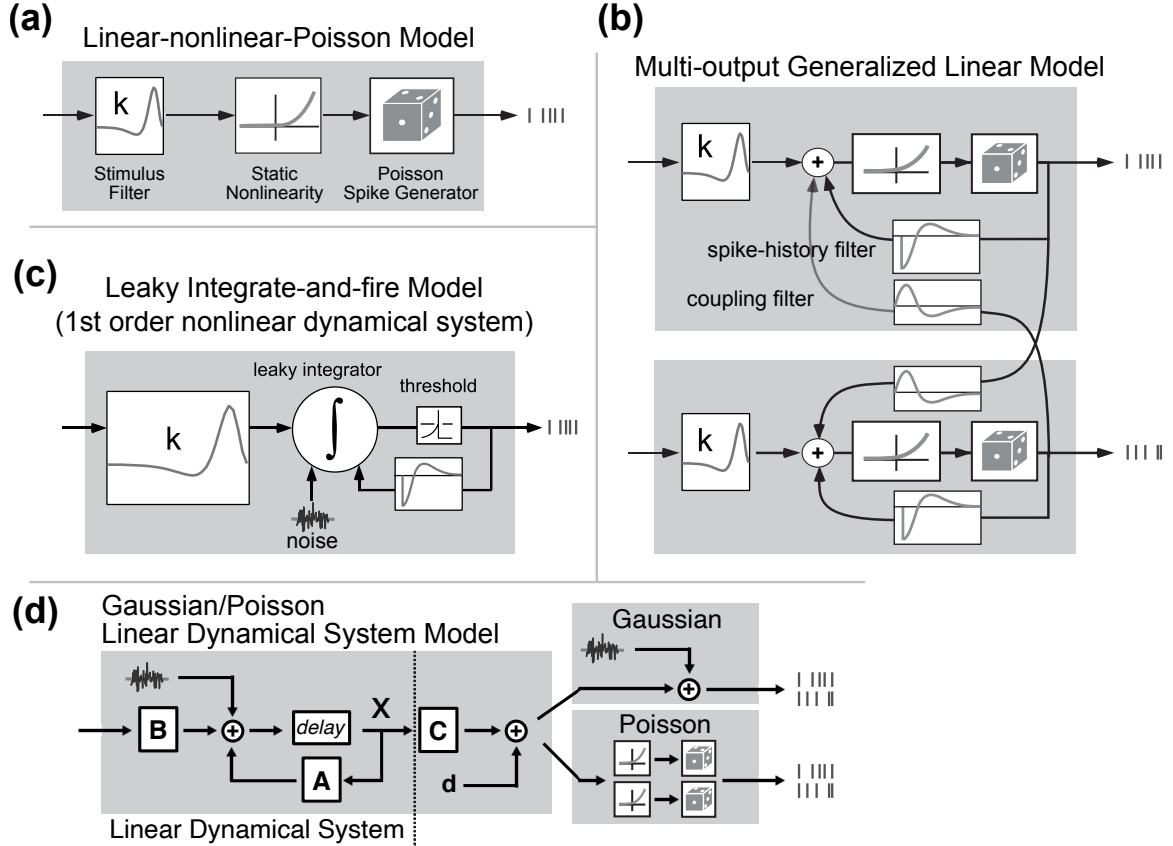


Figure 1.4: **Modeling spiking responses to stimulation.** (a) Linear-nonlinear-Poisson (LNP) cascade model, composed of stimulus filter, a rectifying static nonlinearity, and assumed Poisson spike generator. Adapted from [90]. (b) Multi-output Generalized Linear Model (GLM), which is additionally composed of spike history filter and coupling filters between outputs. Adapted from [90]. (c) Leaky integrate-and-fire (LIF) model. Effectively a 1st-order nonlinear dynamical system, where the model membrane potential is reset to a baseline upon crossing a spike threshold. Adapted from [93]. (d) Gaussian or Poisson Linear Dynamical System (G/PLDS), composed of a linear dynamical system, whose outputs are either assumed to be Gaussian noise-corrupted linear combinations of underlying dynamical states (x), or are assumed to undergo element-wise rectification and Poisson spike generation.

1.4.2 Optogenetics

Until somewhat recently, electrical stimulation has been the chief method for controlling neuronal activity at fast timescales (*e.g.*, [71, 72, 20, 73]), a practice dating as far back as Luigi Galvani’s work in the eighteenth century ([1]). Over the past 15 years, however, optogenetic approaches to stimulation have seen widespread adoption and use in neuroscience due to the ability to excite or inhibit neural activity in a cell-type specific manner (for reviews/protocols, [111, 112, 3, 113, 114, 115, 4]). As the name suggests, optogenetics is a method by which cells can be genetically modified to express “opsins”, membrane bound light-sensitive ion channels or pumps, such that different wavelengths of light can depolarize (excite) or hyperpolarize (inhibit) the cells. The most common opsins are variants of the following: channelrhodopsin (ChR), an excitatory non-specific cation channel (Figure 1.5(a) left, [74]); halorhodopsin, an inhibitory chloride pump (Figure 1.5(a) right, [75]); archaerhodopsin, an inhibitory proton pump ([76]). While there are red-shifted variants of ChR (*e.g.*, [116, 117]), most are blue-light activated (Figure 1.5(b), blue); conversely, halorhodopsin is sensitive to amber light (Figure 1.5(b), yellow). This separation in activation spectra means that at least in theory bidirectional control could be achieved by using different wavelengths to excite or inhibit activity, which has been achieved *in vitro* ([118, 6]). Notably, with the possible exception of [119], this does not appear to have been reported *in vivo*. A survey of the many variants of these opsins is beyond the scope of this work. While the methods developed in Chapters 2-3 apply generally to any such variants, a single ChR has been used throughout: specifically, the H134R mutant of ChR2 ([120, 121]).

Besides the ability to use light to excite or inhibit neural activity at fast timescales, perhaps the biggest reason for the success of optogenetics in neuroscience is the wide array of techniques whereby specific cell types or populations can be targeted for stimulation ([114]). Their expression may be confined to specific cell types under the control of appropriate promoter sequences: *e.g.*, excitatory vs. inhibitory cell types (eg, parvalbumin-

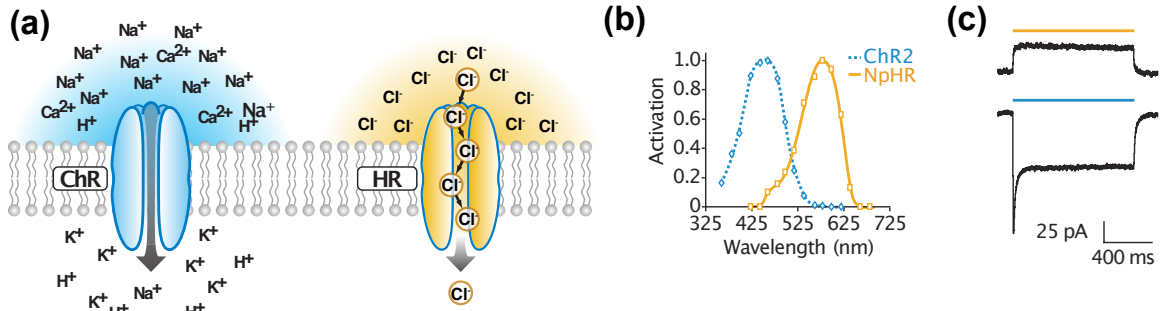


Figure 1.5: **Optogenetics.** (a) Examples of excitatory and inhibitory opsins: channelrhodopsin (ChR, non-specific light-gated cation channel), halorhodopsin (HR, chloride pump). Adapted from [3]. (b) Optical activation spectra for ChR and HR variants. Adapted from [75]. (c) Light-driven hyperpolarizing (*i.e.* outward, top) and depolarizing (inward, bottom) currents for HR and ChR variants recorded intracellularly in cultured hippocampal neurons. Adapted from [75].

expressing interneurons [122, 112]); different lamina of cortex (*e.g.* layer VI CT cells [54]); different subcortical nuclei (*e.g.*, VPM/VPL [123]). Cells can also be targeted as a function of anatomical projections (*e.g.*, retrogradely labeled [124] or synaptic terminal stimulation [125]), and expression can even be confined to a single cell ([62]). Moreover, in the context of rodents and mice especially, there are numerous transgenic lines that allow for brain-wide expression of opsins as a function of cell type ([113, 123]), in addition to the broadly-applicable viral injection approach to transfection.

Finally, with optical stimulation techniques ranging from optical fibers ([72, 6, 16, 126]), scanning lasers ([69, 127]), and micro-LEDs embedded on electrode arrays ([128, 129, 130]), optogenetics is ideally suited for combination with electrophysiology. This is particularly true in the context of feedback control applications because there is less concern for stimulation artifact, allowing continuous stimulation and recording without the sorts of large artifacts seen with electrical stimulation (*e.g.*, [20]).

1.4.3 Optogenetic control

For all the reasons discussed above, there has been recent interest in applying control theory to optogenetic stimulation (for review/perspective, [5]). To date, most examples of closed-loop optogenetics are limited to event-triggered or on-off control approaches, where previously determined optical stimuli are triggered on activity of interest: in the context of two different epilepsy models, Paz *et al* and Krook-Magnuson *et al* triggered optical stimulation on the onset of detected seizures ([57, 131]); O'Connor *et al* achieved illusory touch via direct optical stimulation of barrel cortex when an animal's whisker crossed the path of virtual pole (*i.e.*, interrupted the beam of an infrared laser) ([132]); Latchoumane *et al* used optical stimulation triggered on the “Up” phase of cortical slow wave oscillations to evoke phase-locked thalamic sleep spindles and causally investigate their role in memory consolidation ([133]); and Zhang *et al* used an on-off controller to gate optical pulses to control Ca^{++} activity in individual neurons recorded via two-photon microscopy ([69]). That said, in addition to the feedback optogenetic control that laid the groundwork for this thesis (below), there has been at least one previous report of using continuously-graded closed-loop optical stimulation in which a proportional-integral controller (Chapter 2) was used to change the ankle joint angle in rodents ([134]).

More advanced control-theoretic approaches have been applied to the manipulation of neural spiking activity, but most of these have been limited to simulation studies. For example, Ullah *et al* applied an unscented Kalman filter and feedback control to the problem of manipulating the activity of a single Hodgkin-Huxley model neuron ([135]). Ahmadian *et al* developed an approach for optimal stimulus design for achieving target spike times in a soft-threshold LIF model ([136]); similarly, Nandi *et al* designed optimal stimuli to achieve target spike times in a network of GLM model neurons. Ching and Ritt developed a notion of spike “sequence controllability” in an underactuated network of LIF model neurons with only one light source ([95]). Finally, Iolov *et al* have applied stochastic optimal control to the problem of achieving target spike times in a single LIF model neuron ([137]).

Somewhere between these two extremes of experimentally applied event-triggered optical stimulation versus more advanced approaches confined to simulation lies the *in vitro* work that laid the groundwork for closed-loop optogenetic control (CLOC) as it is described hereon. In an effort to disentangle the relative contributions of firing rate and synaptic transmission on synaptic strength rescaling that occurs during neural network homeostasis, Fong *et al* applied a feedback controller to maintain an average firing rate among neurons cultured in a dish. Activity of these cultured neurons was recorded with an embedded multi-electrode array and they were stimulated optically as shown in Figure 1.6(a). The authors were able to tease apart the contributions of AMPA-ergic synaptic transmission and overall spiking on synaptic rescaling by using CLOC to keep cells at 100% of baseline activity, despite AMPA-receptor blockade (Figure 1.6(b)).

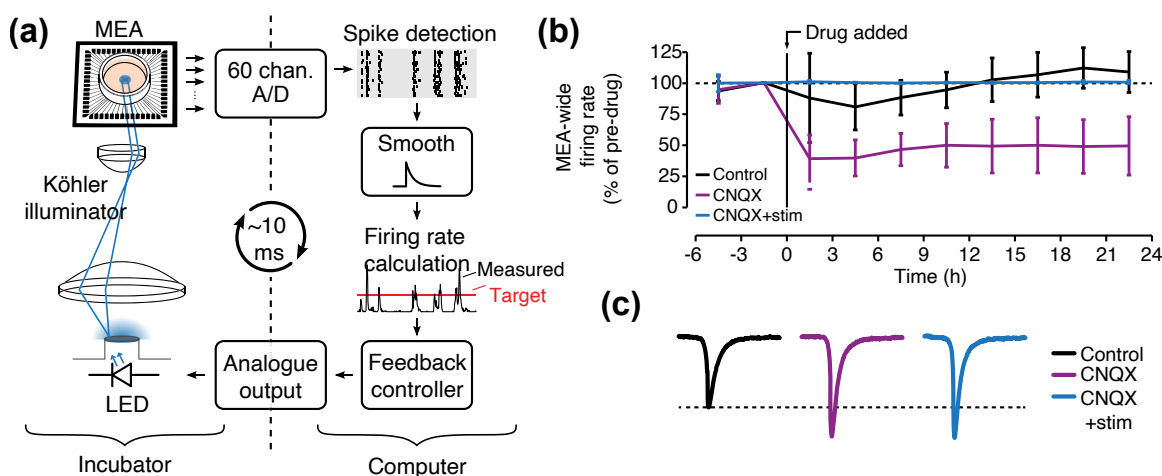


Figure 1.6: Closed-loop optogenetic control in vitro. (a) In vitro closed-loop optogenetic control (CLOC) setup. Adapted from [138]. (b) Multi-hour control of average population firing rate via CLOC, with vs. without AMPA-receptor blockade via CNQX. (c) Excitatory post-synaptic current (EPSC) under control conditions, AMPA receptor blockade, AMPA receptor blockade with CLOC of firing rate. Maintaining firing with CLOC showed that it was changes in synaptic transmission, not reduced spiking, that led to homeostatic synaptic upscaling.

Concurrent with this work and using the same experimental preparation, Newman *et al* demonstrated the ability to hold firing rates of cultured neurons steady for hours, explored the effectiveness of continuously-modulated light delivery vs. control-modulated patterned stimulation (triangle, sinusoidal, square pulses), and demonstrated robustness to different types of pharmacological receptor blockade ([6]). Importantly, while most of the study was

in vitro, the authors conducted an initial set of anesthetized *in vivo* experiments in which they held the firing rate of single thalamic neurons fixed at multiple target firing rates.

1.5 Overview

The initial *in vivo* demonstration of CLOC by Newman *et al* ([6]) was provocative, but the design of the model-free (*i.e.*, proportional-integral) control scheme was *ad hoc*, only applied to the problem of holding firing rate steady over long periods, was limited to the anesthetized animal, and did not investigate the effects of stimulation on thalamic activity more generally. In this thesis, Chapter 2 first develops a design methodology for using the model-free control scheme of Newman *et al* ([6]) to entrain naturally relevant patterns of rate modulation such as observed in the rodent VPM during whisking, rather than holding firing rate constant. In order to generalize more gracefully to future multi-input/multi-output control problems, Chapter 3 applies state-space model-based control and estimation to the problem of manipulating thalamic firing rates. Importantly, we investigate the effectiveness of CLOC in the awake animal for the first time. Finally, in Chapter 4 we investigate the effect of CLOC on thalamic state more generally, analyze the robustness of control to a naturally-occurring disturbance (animal whisking), and look at its consequences for downstream cortical activity and sensory response characteristics. As part of this analysis, a broadly-applicable state-space model based notion of thalamic state is put forth, marrying neuroscientific and engineering notions of “state”.

CHAPTER 2

MODEL-FREE FEEDBACK OPTOGENETIC CONTROL FOR ENTRAINING TIME-VARYING PATTERNS OF NEURAL ACTIVITY

2.1 Introduction

The ability to control neuronal activity has deep clinical and scientific significance, ranging from treating movement disorders and epilepsy to understanding fundamental operating principles of single cells and intricately interconnected networks. Indeed, it was the use of a feedback controller—Cole’s voltage clamp—that enabled the pioneering studies of Hodgkin and Huxley and gave rise to our understanding of the ionic currents underlying the action potential [61]. Today, the neuroscience community is faced more acutely than ever with the task of dissecting the functions of neurons in the context of connected networks. To disentangle the roles of different cell types or structures under such conditions, systems neuroscience requires a set of tools for controlling neural activity at a meso-scale, between the extremes of stimulating single neurons and non-selectively manipulating large, diverse populations. Importantly, these control methodologies should also be robust to changes in ongoing activity in the brain which could otherwise be sources of unexplained experimental variability.

Over the last fifteen years, optogenetic stimulation has emerged as a tool for understanding neural circuit function. Unlike electrical stimulation, optogenetic manipulations have the ability to target the expression of opsins genetically (*e.g.*, [122, 113]) and/or anatomically (*e.g.*, single-cell [62] or retrogradely-labeled [124]). Given the flexibility of this technique and a maturing genetic toolbox, there is growing interest in the intersection between optogenetics and engineering control theory as a method for dissecting circuit function [5, 95, 136]. Notably, optogenetics lends itself particularly well to closed-loop control,

pairing electrophysiological recordings with optical stimulation. In contrast with simultaneous electrical stimulation and recording which is plagued by stimulation artifacts, there is comparatively less concern for such artifacts corrupting measurements when using optical stimulation.

Control theory has already been brought to bear on the problem of manipulating neural activity for the purposes of halting seizures ([139, 140, 57, 131, 141]), reducing oscillatory activity in models of Parkinson's ([142]), and artificially replicating LFP patterns naturally evoked in response to touch ([73]). To date, most so-called neurocontrol—whether by means of optical or electrical stimulation—has been conducted in an open-loop or event-triggered or on-off closed-loop fashion. In the latter case, stimulation is triggered by activity of interest (*e.g.*, [57, 131, 132, 143, 133]). In such applications, the stimulation that is delivered has been previously determined, through experimentation or using previously identified mathematical models. However, neuronal responses elicited by stimulation can vary across individuals, cells, and even over time. Rather than using feedback merely to trigger pre-determined stimulation, another strategy is to make continuous use of feedback to update stimulation in real-time. Recently, we demonstrated the first such use of closed-loop optogenetic control *in vivo* ([6]), where light intensity was adjusted to maintain a constant firing rate over time using a proportional-integral (PI) controller. In this model-free control strategy, output feedback was provided to the controller using a moving average filter to smooth spiking activity into an estimate of instantaneous firing rate. While the first of its kind, the study did not offer a design methodology, nor did it extend the method beyond static reference tracking. In many applications, the objective of neurocontrol may be to entrain patterns of activity such as embedded artificial surrogates for sensory responses ([132, 73]) or desired spike trains ([136, 95, 137, 92]). To this end, we have developed and demonstrated a strategy for tuning the elements of the control loop for eliciting desired patterns of temporal firing rate modulation.

Here, we demonstrate in the anesthetized rat that the simple model-free control scheme

used by Newman *et al* ([6]) can be effective for eliciting desired patterns of firing rate when appropriately designed, and that the use of feedback confers a reduction in trial-to-trial variability as compared to open-loop stimulation. Design of this control loop entails choices for the degree to which measured spikes are smoothed into an estimated firing rate as well as the proportional and integral feedback controller gains on minimizing instantaneous and integrated firing rate tracking error, respectively. We developed an approach for tuning the parameters of this system for eliciting sinusoidally modulated patterns of firing rate by optimizing feedback controller gains in a simulated control loop where the spiking neuronal system was approximated using the simplest of the widely-used spiking models, the linear-nonlinear-Poisson cascade model (Section 1.4.1). Finally, we demonstrate that this procedure can generalize to more complex, non-sinusoidal signals of interest.

2.2 Methods

We have developed a principled design strategy for closed-loop control of dynamic trajectories in neural firing patterns through the use of a PI controller and a moving average filter characterized by its exponential decay time constant serving to estimate the latent firing rate from measured spikes (“Firing rate (FR) filter”). We utilized an anesthetized rodent model where *in vivo* we recorded from and optically stimulated neurons in the ventral posteromedial (VPM) region of the thalamus that were transfected with a depolarizing light-sensitive ion channel. The experimental preparation is outlined in Figure 2.1(a), illustrating the single-unit thalamic recording, optical drive of opsin-expressing thalamic neurons, and the relationship of the thalamus to the whisker-driven afferent input from the periphery and the ascending and descending connections with the primary somatosensory cortex. The block diagram in Figure 2.1(b) illustrates the control framework, where the “neural system” (in this case, a single unit in the thalamus) emits measured spiking patterns (z), which are in turn utilized to estimate firing rate (\hat{y}). The difference in estimated firing rate and reference (or, desired) firing rate (r) is defined as the online estimate of tracking error (\hat{e}). A con-

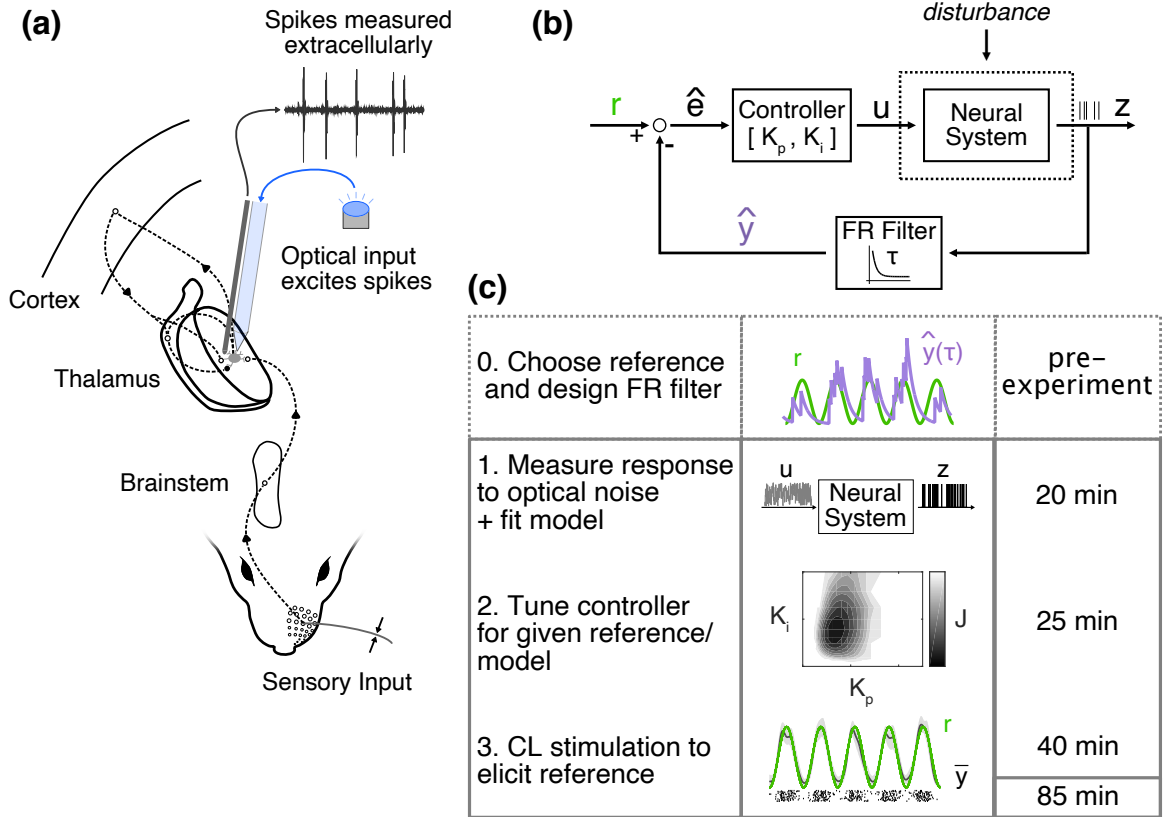


Figure 2.1: **Closed loop optogenetic control of firing rate** (a) Physical diagram. (b) System block flow diagram. (c) Procedure for closed-loop stimulation experiments. The firing rate filter used to estimate output feedback was designed for a given reference firing rate pattern previous to experiments. A model was fit to data recorded for system identification during the experiment. Using this model, controller gains were optimized in simulation. These parameters were then used for experimental closed-loop stimulation.

ventional PI controller operates on this error signal, yielding a light input (u) which drives optogenetic excitation of the thalamic neurons.

2.2.1 Experimental preparation

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 NIH. Experiments were carried out using female albino (Sprague-Dawley) rats. Expression of channelrhodopsin was targeted to excitatory neurons (rAAV5/CamKIIa-hChR2(H134R)-mcherry-WPRE-pA; UNC Vector Core, Chapel Hill, NC) in the ventro-posteromedial nucleus (VPM) of the thalamus by way of stereotactic injections (3 x 3 x 5.2 mm rostro-caudal

x medio-lateral x depth)[16]. A 1 μ L volume of virus was injected at a rate of 0.1 μ L/min. The animals were given buprenorphine for pain management (0.03 mg/kg). Animals were then monitored daily following injection surgery. Wound clips were removed at 10-13 days post-surgery. Animals were allowed to recover and opsins allowed to express > 3 weeks. While VPM of the thalamus was the anatomical target, any optogenetically-driven neuron in thalamus that exhibited well-isolated spiking activity was considered a candidate for this study.

On the day of the experiment, rats were anesthetized using a cocktail of fentanyl (5 μ g/kg), midazolam (2 mg/kg), and dexmedetomidine (150 μ g/kg) delivered intravenously through the tail vein ([16, 144]). Animal body temperature was maintained at approximately 37° C using a feedback-controlled heating pad. A 3 mm x 3 mm cranial window centered over the left hemisphere at 3 mm rostral and caudal of bregma was created and the *dura mater* carefully removed. Single units were isolated in thalamus using an optrode: 80 μ m, 2 M Ω tungsten electrode (FHC), coupled to a 200 μ m optic fiber (Thorlabs). Blue light was conducted from an LED (470 nm, ThorLabs) to the thalamus via the optic fiber. Prior to the experiment, an optical power meter was used to measure the light intensity emitted from the tip of the fiber when peak command voltage (5 V) was sent to the LED driver. We found a linear relationship between command voltage and light intensity. Therefore, we estimated the light inputs reported in this study by scaling the command voltage accordingly.

All signals were recorded using a Tucker Davis Technologies (TDT) RZ2 Bioprocessor. Extracellular voltage was recorded at 24.414 kHz and bandpass filtered from 500 to 5000 Hz. Single unit spikes were manually thresholded, their dimensionality reduced using principal component analysis, and clustered in this reduced space using K-means. Measured spikes were then smoothed into an online estimate of instantaneous firing rate, and a PI controller was used to modulate the amplitude of optical stimulation of the neuron (Figure 2.1(a-b)). In post-hoc analysis, we found that optical artifacts were minimal, especially when considering the 500 to 5000 Hz band used for spike thresholding/sorting

(see Appendix, Figures B.1-B.2). We also found that spikes used for feedback during control epochs did not significantly differ from the spikes that occurred spontaneously in inter-trial-intervals.

At the conclusion of each experiment, animals were sacrificed using an overdose of sodium pentobarbital.

2.2.2 Reference trajectories

Sinusoidal

To extend beyond the methodology laid out by [6] for maintaining a constant firing rate, we designed the control system to elicit sinusoidally-modulated rates of the form

$$y_t = \sigma_{\sin} \sin(2\pi f_{\text{mod}} \Delta t) + \mu_{\sin} ,$$

where σ_{\sin} , f_{mod} , μ_{\sin} , t , and Δ are the sinusoidal amplitude, modulation frequency, DC firing rate, time index, and sample period, respectively. Such a parametric time-varying reference is the ideal starting point for this design problem. Motivated by firing rates observed in the awake animal (see next section), all reference trajectories had mean firing rates of 20 spikes/s, which was found to be an appropriate average firing rate during awake rodent whisking (see Appendix A, [27]). To focus our investigation further, the sinusoidal reference trajectories used here were maximally modulated about the mean (*i.e.*, $\sigma_{\sin} = \mu_{\sin}$). The control system was tuned for 1, 5, or 10 Hz modulated patterns.

Non-sinusoidal trajectory

In order to test the results of the design procedure on a non-sinusoidal reference trajectory of interest, we used an example of rate modulation observed in single-unit data recorded in the VPM of an awake rat (data from [145, 16]; see Appendix A).

Rhythmic spiking activity possibly related to the animal moving its whiskers was iden-

tified in a subset of trials. To inspect for rhythmic spiking, spike trains were smoothed using a Gaussian window with standard deviation (SD) of 20 ms and autocorrelograms calculated for each trial. Putative “whisking” trials ($n = 3$) were identified by peaks in correlation located at lead/lag of 100 ms, corresponding to 10 Hz which is within the natural frequency range of whisking [146, 147, 148]. Spike trains from these trials were aligned such that their cross-correlations had a peak at zero-lag, were averaged across trials, and then smoothed with a Gaussian window of 20-ms SD, resulting in the reference trajectory shown in Figure 2.2(c-d) (see also Appendix, Figure A.1). The mean (*i.e.*, DC) firing rate of this signal was approximately 20 spikes/s, and 95% of the total power in this signal was between DC and 10 Hz. Note that approximately 50% of the power in this signal occurred at DC.

2.2.3 Firing rate estimator design

The closed-loop control system had two designed components: an estimator of output firing rate and the feedback controller. In this application, measurements ($z \in \mathbb{R}^1$) took the form of spiking activity. Given a raw measure of instantaneous firing rate, z/Δ (either Δ^{-1} or 0 if there is/is not a spike, where Δ is the sampling period), the firing rate filter yields an online estimate of the latent rate, $\hat{y} \in \mathbb{R}^1$.

Fixed-bandwidth smoothing

Fixed-bandwidth smoothing was used to estimate the firing rate online. An exponential window was chosen because it is causal and efficiently computed online as a first-order recursive filter. The firing rate estimator filter was characterized by a single parameter, its decay time constant (τ):

$$\hat{y}_t = a\hat{y}_{t-1} + (1 - a) \frac{z_t}{\Delta}$$

where,

$$a = \exp\left(-\frac{\Delta}{\tau}\right),$$

and t is the sample index.

Parametric sweeps for optimal sinusoidal rate estimation

In designing the bandwidth of the exponential filter, the goal was to choose a filter which provided an appropriate amount of smoothing such that the underlying firing rate could be accurately recovered from measured spikes. Consistent with previous work (*e.g.*, [149, 150, 151, 152]), the filter bandwidth was designed in simulation by minimizing the mean integrated squared error (MISE) between a ground truth firing rate pattern and the estimate:

$$\tau^* = \arg \min_{\tau} \left\langle \frac{1}{T} \sum_{t=1}^T (y_t - \hat{y}_t)^2 \right\rangle,$$

where $\langle \cdot \rangle$ denotes trial-averaging and \hat{y} is the rate estimated by filtering the spike trains generated according to Poisson statistics from the ground truth rate, $y \in \mathbb{R}^1$. Here, y was taken to be the reference firing rate pattern of interest, $r \in \mathbb{R}^1$. Depending on the properties of this pattern, the MISE-optimal estimator may smooth out the temporal modulation due to a scarcity of spikes within a period of oscillation. Koyama and Shinomoto [150] discussed such “divergent” solutions in the context of designing the optimal bin width for constructing a peristimulus time histogram (PSTH).

In the case of sinusoidal rates used here, the optimal filter time constant depended upon the expected number of spikes per period of the sinusoid, $(\mu_{\sin} N_{\text{trials}}) / f_{\text{mod}}$, and the degree of the modulation as defined by the ratio of amplitude to DC offset of the sinusoid, $\sigma_{\sin} / \mu_{\sin}$, where N_{trials} is the number of trials used for estimation. To develop a parametric expression for the filter design, we determined the optimal time constant for different conditions of a simulated rate-modulated Poisson process.

In simulation, four parameters were swept: the modulation frequency of the driving

sinusoid (f_{mod}), the mean firing rate (DC offset of the sinusoid, μ_{sin}), the amplitude of the sinusoid around the mean (σ_{sin}), and the number of trials used for estimation (N_{trials}). Note that at the conclusion of the design, we sought time constants that were optimized for single-trial estimation: *i.e.*, where $N_{\text{trials}} = 1$. For the rates used in this study, we focused on the specific cases where $\mu = 20$ spikes/s and $\sigma_{\text{sin}}/\mu_{\text{sin}} = 1$.

The MISE-optimal filters (τ^*) were calculated for a range of frequencies, mean rates, modulation intensities, and number of trials. In keeping with Koyama and Shinomoto [150], a power law relationship was fit to these data for τ^* to yield a tuning curve of estimated time constants, $\hat{\tau}$:

$$\hat{\tau} f_{\text{mod}} = b_{\tau} \zeta^{-a_{\tau}} ,$$

where [150] defined ζ as

$$\zeta \equiv \left(\frac{\mu_{\text{sin}}}{f_{\text{mod}}} \quad N_{\text{trials}} \right) \left(\frac{\sigma_{\text{sin}}}{\mu_{\text{sin}}} \right)^2 .$$

2.2.4 Linear-nonlinear Poisson model

Previously, an *ad hoc* approach to controller design demonstrated proof-of-principle feedback optogenetic control. Here, however, we sought a principled approach that generalized to dynamic reference trajectories. To finely tune the controller during experiments, we used a linear-nonlinear Poisson (LNP) model to approximate the response of the neural system to optical stimulation. The controller gains (K_p , K_i) were tuned around the LNP model in simulation. In this model structure, a linear system (denoted L) was cascaded with a static nonlinearity (N) to produce a latent firing rate which drives a Poisson spike generator (P) and emits spikes.

The optical input was filtered through a finite impulse response (FIR) stimulus filter, or kernel, \mathbf{k} :

$$x_t = \mathbf{k}^\top \mathbf{h}_t ,$$

where \mathbf{h}_t is the stimulus a vector of stimulus history over a time window up to time t , inclusive. The output of this filter, $x_t \in \mathbb{R}^1$, was then mapped through a static nonlinearity (*e.g.*, [153]):

$$y_t(\theta, x_t) = \alpha \log [1 + \exp (gx_t + d)] ,$$

where $\theta = [d, g, \alpha]^\top$ are parameters describing the static nonlinearity. d serves as a bias term reflecting neural firing that does not co-vary with stimulation (*i.e.*, spontaneous firing rate). g and α together set the effective DC gain of the model neuron's response to light stimulation as it approaches the asymptotically linear region of the curve, whereas their relative values set the knee of the nonlinearity.

During the experiment, an LNP model was fit to spiking data recorded in response to repeated presentations of a 5-second instantiation of optical uniform white noise. The range of this noise was titrated for each cell to avoid apparent depolarization block, but on average ranged between 0 and 8.5 mW/mm². The white noise stimulus was mean-subtracted and a light-to-spiking kernel, $\hat{\mathbf{k}}$, was estimated by reverse correlation (“rotated” or “whitened” spike-triggered averaging, *e.g.*, [85, 80, 154]):

$$\hat{\mathbf{k}} = (\mathbf{U}^\top \mathbf{U})^{-1} \mathbf{U}^\top \mathbf{z} ,$$

where \mathbf{U} is the stimulus design matrix in which the t^{th} column is a vector of stimulus history over a time window up to time t , $\mathbf{z} \in \mathbb{R}^T$ is the across-time vector spike signal (1 or 0 if there is/is not a spike measured). To avoid ambiguity between the static gain of the kernel and the scaling factor applied to the output of the kernel in the nonlinearity (g), the kernel was normalized by its static gain. The stimulus was filtered with this normalized

kernel, yielding x_t , and the remaining parameters $(\hat{\theta})$ were fit by maximum likelihood:

$$\hat{\theta} = \arg \max_{\theta} \sum_K \sum_T z_{tk} \log [y_{tk}(\theta, x_{tk})\Delta] - y_{tk}(\theta, x_{tk})\Delta ,$$

where t and k denote the time and trial index, respectively. Note that [93] fit LNP models by maximum likelihood, with the kernel parameters only *initialized* using the values predicted by whitened spike-triggered averaging (STA). In our application, we found that the kernels resulting from this approach were primarily scaled versions of the STA estimate, \hat{k} . Because parameter g accounts for this scaling, we found that it was unnecessary to re-estimate the kernel parameters in practice. This allowed for the model to be fit quickly and dependably during the experiment, where experimental viability is time-limited.

2.2.5 Controller design

We implemented a proportional-integral (PI) controller, which was defined by two parameters, K_p and K_i . Given a reference trajectory of interest, a choice for the firing rate filter time constant (Section 2.2.3), and a model, these controller gains were tuned in simulation during the experiment.

Control law

The PI controller was implemented in its parallel form:

$$\begin{aligned} \hat{e}_t &= r_t - \hat{y}_t \\ u_t &= K_p \hat{e}_t + K_i \sum_{i=0}^t \hat{e}_i \Delta , \end{aligned}$$

where $u \in \mathbb{R}^1$, K_p and K_i are the control input, proportional, and integral gains, respectively.

Note that due to the fact the light delivered can neither be negative nor exceed the limits

of the LED output, the variable u was bounded between $[0,1]$ and scaled to the dynamic range of the LED driver ($5 \text{ V} = 1 \text{ A}$). For all simulations, the control signal was hard rectified and bounded as was done in hardware to ensure actuation was subject to the same limitations.

Objective function

A common objective for controller design is to minimize the integrated squared tracking error (*e.g.*, in the context of PID control [155]). We found that tuning the controller to minimize the integrated squared error between the reference and online estimate of rate (*i.e.*, between r and \hat{y}) yielded undesired tracking behavior, especially at higher frequency references where the firing rate estimate is less accurate (see insets of Figure 2.3C, Section 2.3.2). Because filtered estimates are on average slightly lagged and attenuated in amplitude compared to ground truth, if stimulation were optimized around this estimate, the resulting neural firing would lead the target and be of larger amplitude. To avoid this behavior, we instead tuned the controller by minimizing the squared error between the reference (r) and the raw measure of instantaneous rate (z/Δ), which neither imposes lag nor amplitude attenuation on the estimate.

While minimizing this raw tracking error ameliorates the aforementioned problems, the highly punctate nature of the z/Δ estimate of rate and the large error incurred each time the neuron spikes ($r - \Delta^{-1}$) mean that simply minimizing the integrated square of the error, $e = r - z/\Delta$, yielded trivially low solutions for controller gains where the neuron never spikes. This motivated controller design in the frequency domain, where the effects of the spiking error can be down-weighted relative to important tracking criteria. Because we are willing to tolerate error that occurs at frequencies higher than those of interest for a control task, we avoided trivial low-gain solutions by weighting the squared error in the frequency domain according to the spectral content of the reference trajectory (w), thereby penalizing error at frequencies according to their significance for the intended control.

We calculated the frequency-domain amplitude of the tracking error, e , as well as that of the reference trajectory:

$$E[f] = |\mathcal{F}\{e\}|$$

$$R[f] = |\mathcal{F}\{r\}| ,$$

where $\mathcal{F}\{\cdot\}$ denotes the Fourier transform.

The spectral content of the reference was used to create frequency-dependent weights:

$$w[f] = \frac{R^2[f]}{\sum_{f=0}^{N_f} R^2[f]} ,$$

where N_f corresponds to the Nyquist frequency.

These weights were used to calculate a modified squared error metric where the frequency spectrum of the error was penalized as a function of importance for the control task:

$$J_{\text{fwt}}(K_p, K_i) = \left\langle \sum_{f=0}^{N_f} w[f] E^2[f] \right\rangle ,$$

where $\langle \cdot \rangle$ denotes the across-trial average.

Finally, controller gains were chosen to minimize J_{fwt} , using a numerical solver (*ga*, Mathworks, Inc.):

$$[K_p^*, K_i^*] = \arg \min_{K_p, K_i} J_{\text{fwt}}(K_p, K_i) .$$

2.2.6 Open-loop stimulus design

To assess the benefits of utilizing feedback with optogenetic stimulation, we designed open-loop stimuli for each control task (*i.e.*, sinusoids, and a more natural non-sinusoidal trajectory). In Figure 2.2(a)&(b), we compare pulsatile open-loop stimulation (most commonly

used stimulation mode) with closed-loop, continuous modulation of light amplitude. In this scenario, 5-ms pulses were generated by varying the frequency of a carrier wave according to the desired firing rate. To determine the appropriate pulse amplitude at the time of the experiment, a static logistic mapping was fit to the response to pulsatile inputs of varying amplitude, where light inputs were presented at the DC firing rate of the target trajectory (*i.e.*, 20 pulses/s, see Appendix A). The amplitude of stimulation was chosen such that approximately 1 spike/pulse was generated on average.

In all other cases of open-loop control, a static mapping from light intensity to firing rate was estimated by fitting a logistic curve to the steady-state firing rate in response to 1.5-second step inputs of light at various amplitudes. Open-loop control signals were designed for sinusoidal and non-sinusoidal firing rate trajectories by inverting this curve (*i.e.*, using this curve as a lookup table).

2.2.7 Disturbance

A load disturbance in the form of whisker stimulation was used to challenge both open-loop and closed-loop control strategies. Using a computer-controlled galvanometer motor ([145, 16, 67]) positioned approximately 10 mm from the animal's face, sensory white noise was applied to the thalamic neuron's principal whisker [16] at two seconds into the control epoch.

2.2.8 Offline firing rate estimation

All reported firing rates were estimated offline (\bar{y}) using an appropriate Gaussian window to smooth either single-trial (Figure 2.2(a-b)) or trial-averaged spike trains binned at 1-ms resolution. In the case of the slowly-modulated rates in Figure 2.2(a-b), the same procedure used for designing the filter time constant for online rate estimation was employed to estimate an MISE-optimal SD of a Gaussian window. Because accurate single-trial estimation is difficult at higher frequencies, for all other firing rate estimation reported here, a Gaus-

sian window was used to smooth a trial-averaged peri-stimulus time histogram (PSTH). The SD of this filter was chosen for each reference trajectory (*i.e.*, 5 Hz or non-sinusoidal) in the same way as before except now for multiple trials: the MISE-optimal width for recovering the reference rate from a PSTH of simulated Poisson spikes, averaging the same number of trials collected experimentally. Bands around these trial-averaged firing rate estimates are 95% confidence intervals for the smoothed PSTH, bootstrapped by sampling the trials with replacement.

2.2.9 Fano factor

Trial-to-trial variability in spike count was quantified using the Fano factor [156] calculated in a 250-ms sliding window:

$$FF = \frac{\text{var}[N_{250}]}{\langle N_{250} \rangle},$$

where N_{250} is the spike count per 250-ms window of time and $\langle N_{250} \rangle$ indicates the across-trial average.

2.3 Results

Here, we develop a design strategy for closed loop control that could be applied to a range of different of neural circuits. We have applied this design strategy to the problem of closing the loop around the spiking activity of a single neuron in the somatosensory thalamus of the rat *in vivo*, as illustrated in Figure 2.1(a). Specifically, the input to the “neural system” was light delivered by way of a fiber optic cable inserted deep into the brain, targetting the somatosensory thalamus. Blue light (470 nm) drove the depolarization of excitatory thalamic neurons expressing channelrhodopsin (ChR2). Spiking activity of a single thalamic neuron was measured using a tungsten extracellular recording electrode bundled to the fiber optic cable (often referred to as an “optrode”).

The block-diagram of the control system is shown in Figure 2.1(b). We applied light

input u to the “neural system”, or plant, whose activity was measured through spiking activity $z \in \mathbb{R}^1$. In this application, the control objective was to achieve time-varying trajectories in firing rate, and thus the feedback signal consisted of a filter’s estimate of instantaneous firing rate, “Firing rate (FR) filter” in Figure 2.1(b). Optical stimulation intensity, $u \in \mathbb{R}^1$, was varied continuously by a proportional-integral (PI) controller, which acted on the error signal between the reference/desired time-varying firing rate r and the online estimate of the instantaneous output firing rate, $\hat{y} \in \mathbb{R}^1$. Control was effected in the face of unobserved disturbances, which could take the form of uncontrolled inputs and/or changes in the dynamics of the system.

The design procedure utilized here required identification of a model for the neural system during the experiment; combined with testing the controller design, this places serious demands on the duration of an experiment, which is typically limited to approximately 2-3 hours. The timeline and task demands are outlined in Figure 2.1(c).

2.3.1 Examples of closed- vs. open-loop stimulation

Before unpacking the technical details of the system design in subsequent sections, we first present examples representing the basic abilities of the closed-loop framework vs. open-loop stimulation strategies to track slowly-modulated reference firing rates (here, 1 Hz), reject exogenous disturbances, and track complex, biologically relevant trajectories in firing rate.

In the example for controlling a 1 Hz sinusoidal firing rate provided in Figure 2.2(a), pulsatile input was used for open-loop stimulation (bottom). Note that for this instance of open-loop control, pulsatile inputs were used since this is the most common way to stimulate ChR2. Single-trial and trial-averaged estimates of firing rate are shown for closed-loop (top) and open-loop (bottom) cases. For the open-loop case, measured firing activity was not utilized in shaping the light input, but instead the amplitude of the light inputs were designed based on previous measurements of number of spikes elicited as function of pulse

amplitude (see Section 2.2.6). In each scenario, the trial-averaged input (*i.e.*, light intensity) is plotted below the corresponding firing rate. The control epochs begin at time zero. Both closed- and open-loop stimulation strategies achieved the target firing behavior *on average* in the undisturbed scenario (Figure 2.2A, thick red and black versus green). However, single-trial estimates of firing rate (thin red lines, smoothed with a 120-ms SD Gaussian window) reveal that open-loop stimulation resulted in more variable rate trajectories than closed-loop for this example.

A major benefit of a closed-loop system is its capacity to react to changes in ongoing activity and reject disturbances, as illustrated in Figure 2.2(b). For the same example thalamic neuron in Figure 2.2(a), we identified the whisker on the contralateral side of the animal's face to which the neuron responded most robustly, often referred to as the “principal whisker”. A whisker disturbance (see Section 2.2.7) begins at 2 seconds into the control epoch. During the application of the disturbance, closed-loop stimulation was able to adjust to maintain reasonable control of the 1 Hz trajectory (Figure 2.2(b), top). Conversely, in the case of open-loop stimulation (Figure 2.2(b), bottom), the firing rate was unsurprisingly increased well above the reference. It is also of interest to note that, while the same open-loop pulsatile stimulus was used in Figure 2.2(a-b), the effectiveness of stimulation was weaker in the initial 2 seconds of the control epoch before the onset of the disturbance. Given that time elapsed between these two recordings, this phenomenon speaks to apparent non-stationarity in the system for which closed-loop stimulation is able to compensate, even over relatively short timeframes.

We also challenged the control framework with a non-sinusoidal, reference trajectory derived from thalamic spiking measured in an awake animal (Section 2.2.2, Figure A.1), as shown for a different thalamic neuron in Figure 2.2(c). As described in the Section 2.2.2, we analyzed previously-recorded single unit firing activity in the VPM thalamus of the awake rat during active whisking, and used this to generate a more “naturalistic”, non-sinusoidal firing rate trajectory to track in these separate experiments. Note that open-

loop control performed qualitatively similarly to closed-loop control on average in this example, although closed-loop does provide modest improvements in tracking. However, in agreement with the sinusoidal trajectories above, closed-loop stimulation results in lower trial-to-trial variability, as shown by Fano factor (FF), where the variability of closed-loop controlled firing generally falls below open-loop.

These examples of control shown in Figure 2.2 depended upon the design and implementation of both the firing rate filter and controller elements of a closed-loop system, which we further detail below, before returning to further analyses of the control performance.

2.3.2 Firing rate estimator design

In order to perform feedback control of output firing rate, this underlying instantaneous rate function had to be estimated online, given measured spiking activity. As described in detail previously (Section 2.2.3), for simplicity of real-time implementation, instantaneous firing rate was estimated by a first-order moving average filter with a single parameter, the decay time constant τ . Figure 2.3(a) shows the basic concept for the design strategy, where a Poisson spike generator was driven with a sinusoidal rate function, y , and the latent rate was estimated by filtering the spike train. The goal of the design procedure was to choose the filter time constant that optimally recovered the underlying sinusoidal rate. Shown are the results for an illustrative example where a simulated spike train was filtered with one of three values of the decay time constant τ (actual rate in green and estimated rates in purple shown on the right; bottom: too slow, top: too fast, middle: MISE-optimal).

As described in detail in Section 2.2.3, for sinusoidal modulation the optimal filter for firing rate estimation depended upon the baseline (DC) offset of the sinusoid (μ_{sin}), the modulation frequency (f_{mod}), the amplitude (σ_{sin}), and the number of trials used for estimation (N_{trials}), which together influenced the optimal filter bandwidth for recovering sinusoidal rate modulation from recorded spikes. Plotted on the top panel of Figure 2.3(b)

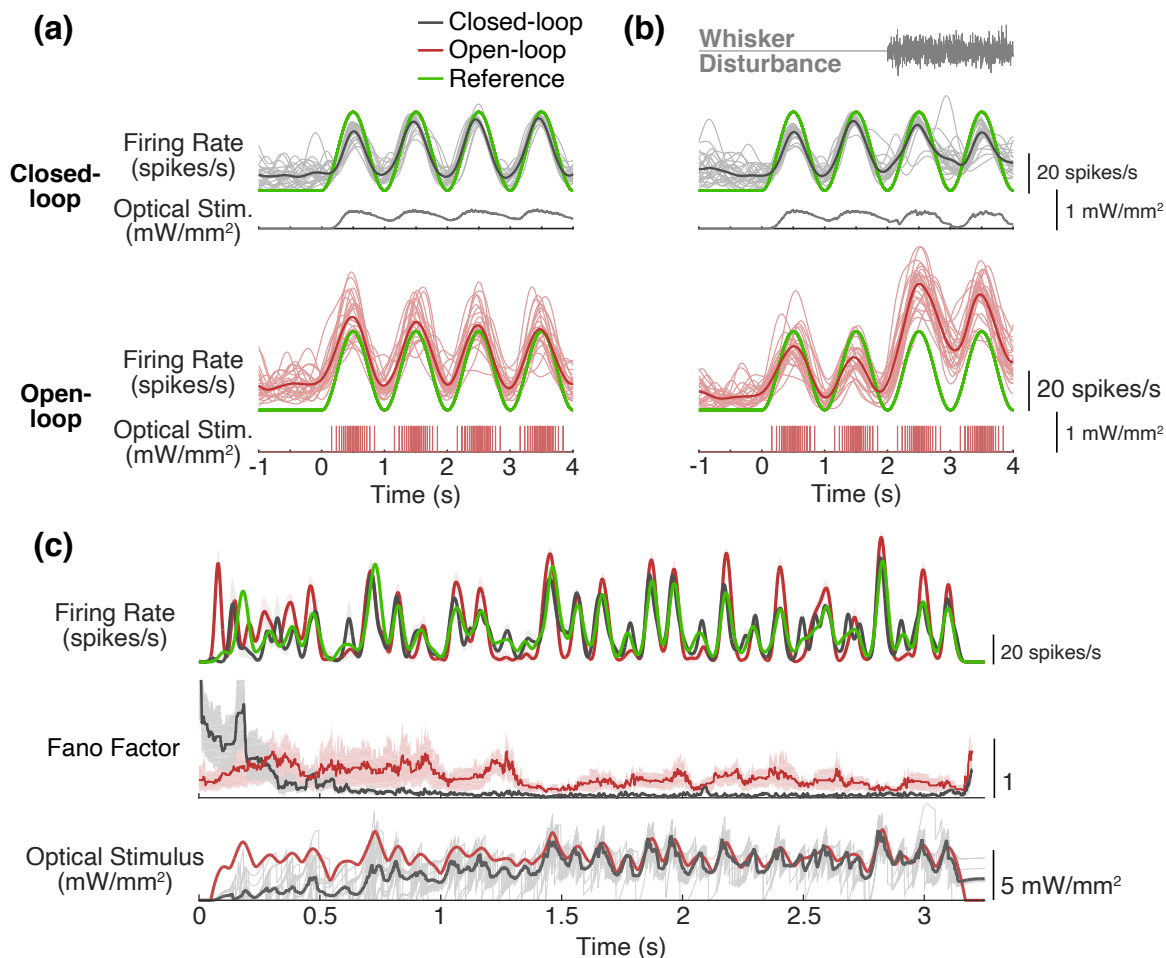


Figure 2.2: Closed- vs. open-loop optogenetic control of dynamic firing rate trajectories (a) Closed- and open-loop control of sinusoidally-modulated firing rate. Closed-loop (black) control and pulsatile open-loop (red) stimulation were used to elicit a 1 Hz sinusoidally modulated firing rate. Light lines correspond to single trial firing rates estimated by smoothing spike trains with a Gaussian window (120 ms SD); bold lines are the trial-averaged rate. Average control inputs (*i.e.*, light) are below the corresponding firing rate trajectory. (b) Closed- and open-loop control in presence of a disturbance. Control was challenged with a whisker disturbance at 2 seconds into the control epoch, as shown in gray (top). (c) Closed-loop and open-loop control of non-sinusoidal firing rate. Top, firing rates for closed-loop (black) vs. open-loop (red) control: average in bold, while fills represent 95% confidence intervals for smoothed PSTH. Middle, Fano factor calculated in 250-ms sliding window for closed- and open-loop control ($n = 25$ trials). Bottom, trial-averaged control inputs for closed-loop (black) or open-loop (red).

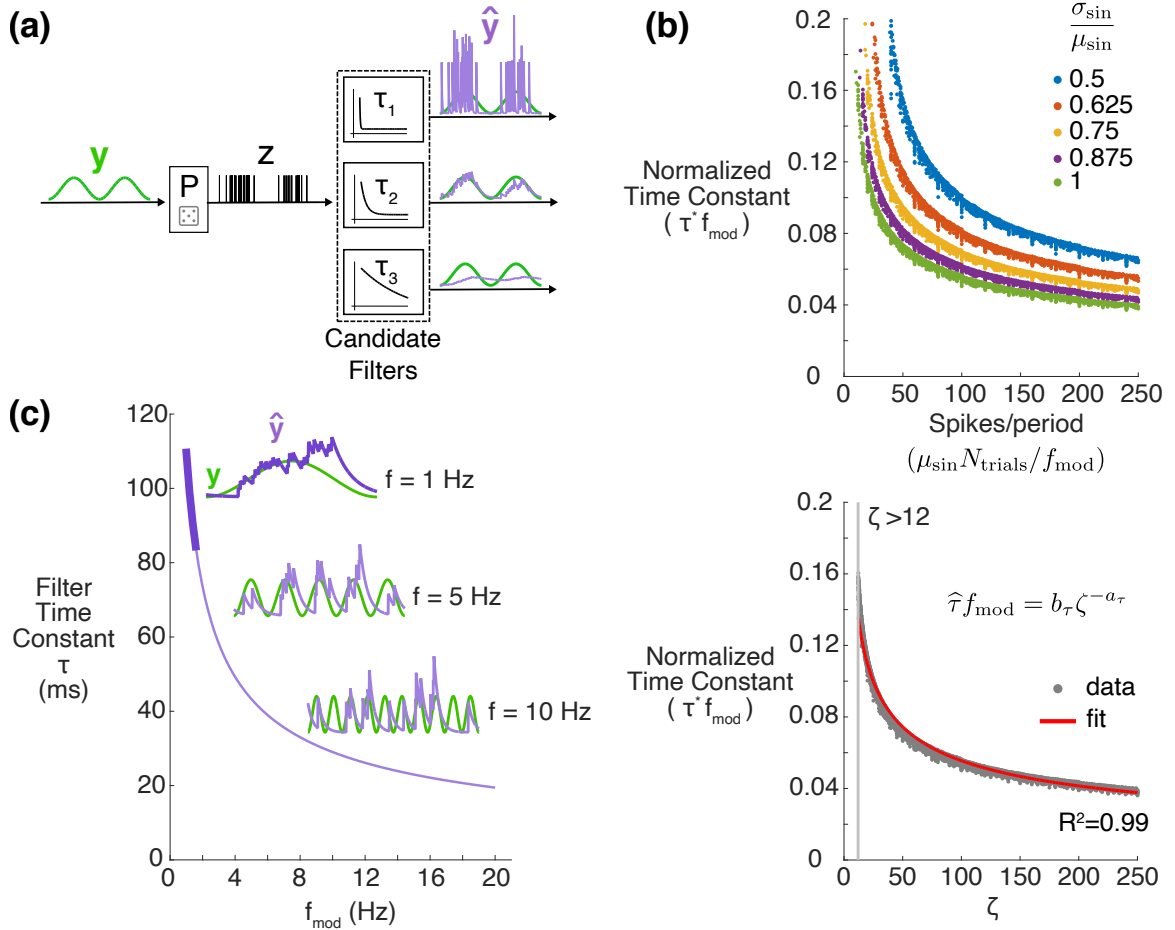


Figure 2.3: **Firing rate filter: Choosing filter bandwidth** (a) Conceptual diagram. A given sinusoidal firing rate (λ) drove a Poisson spike generator (P). The resulting spike train was multiplied by Δ^{-1} (not shown) ahead of filtering. Filters parameterized by a time constant, τ , yielded an estimate of the true rate. (b) Optimal time constant as a function of expected number of spikes per period and degree of modulation about mean (top) and as a function of ζ (bottom). For visualization, the optimal filter has been normalized by the frequency of each sinusoid. Fit: $a_{\tau} = 0.423$, $b_{\tau} = 0.389$. (c) Filter time constants designed for single-trial estimation where $\mu_{\text{sin}} = 20$ spikes/s, $\sigma_{\text{sin}} / \mu_{\text{sin}} = 1$. Bold purple, region of frequencies where the derived design equation was fit. Light purple, frequencies at which the the design equation was extrapolated ($\zeta < 12$). Insets, example single-trial estimates (purple) of the ground truth rate (green) at indicated modulation frequencies.

are the MISE-optimal values of the filter’s decay time constant, τ^* , as a function of the number of spikes/period of the sinusoid, for varying degrees of modulation (σ_{\sin}/μ_{\sin}). As expected, the MISE-optimal filter time constant decreased with increasing number of spikes per period and modulation intensity. We then used the quantity ζ defined by Koyama and Shinomoto [150] to reduce the dimensionality (see Section 2.2.3). The disparate curves in the top panel of Figure 2.3(b) then collapsed to a single curve when plotted as a function of ζ (Figure 2.3(b), bottom). We fit a power law relationship for the optimal time constant using least-squares, the result of which is shown in red in Figure 2.3(b) (bottom): $a_{\tau} = 0.423$, $b_{\tau} = 0.389$.

In the case of the sinusoidal firing rates utilized in the remainder of this study, where modulation amplitude was maximal (*i.e.*, $\sigma_{\sin}/\mu_{\sin} = 1$), the MISE-optimal filter masked temporal modulation in favor of capturing the DC firing rate if there were fewer than 12 spikes per period. More generally, this corresponds to a regime where $\zeta < 12$ in which there was a deficit in spiking data on a single-trial basis (denoted by the vertical gray line in Figure 2.3(b), bottom). As such, these data points were excluded from the regression. For the rates used in the remainder of this study where $\mu_{\sin} = 20$ spikes/s, this corresponded to frequencies greater than 2 Hz (Figure 2.3(c), light purple). By extrapolating the expression for filter time constant that was developed and fit in the regime where $\zeta > 12$, we arrived at choices for time constants that did not smooth out temporal modulation at higher frequencies despite this fundamental limit while still providing a reasonable degree of smoothing, as seen in Figure 2.3(c) (insets). From the insets which provide examples of single-trial estimation of a 1 Hz, 5 Hz, and 10 Hz sinusoidal rates from spikes, it is apparent that when the modulation frequency reached 10 Hz for this DC offset, spikes were too infrequent per period of oscillation to recover the sinusoid accurately on a single-trial basis.

The net result of this procedure was a relationship between the filter time constant and the parameters of the sinusoidal rate which was valid in a regime where there was sufficient

data to resolve the temporal modulation. When the rate estimation was data-impooverished, this relationship allowed us to estimate what the filter time constant would be if not for this fundamental limit. Therefore, while it is impossible to faithfully recover higher frequency modulation on a single-trial basis, this approach provides a principled choice of firing rate filter time constant that does not smooth out the rate modulation.

2.3.3 Controller design

Given the design for the firing rate filter time constant, we developed a PI controller design strategy for tracking sinusoidal trajectories of different frequencies (1, 5, 10 Hz). For principled design of the controller, we undertook a brief system identification step during the experiment to fit a model for the neural system, consisting of a cascade of a linear stimulus filter, static nonlinear function and Poisson spike generator (LNP model). Controller gains were then tuned in simulation using the LNP model fit to data in place of the neural system (Figure 2.4(a)), followed by implementation in the experiment.

To tune the controller, we chose gains which minimized the squared error between the reference rate and the measured spiking signal. This noisy error signal was weighted in the frequency-domain according to the importance of a frequency for the intended control (see Section 2.2.5). In the context of sinusoidal trajectories used here ($\sigma_{\sin} = \mu_{\sin}$), the frequency-weighted squared error objective function (Section 2.2.5) evenly penalized error at DC and the modulation frequency.

To illustrate the tuning procedure, we provide example simulations of the control loop using an LNP model fit to experimental data. The frequency-weighted error metric is shown as a function of the proportional controller gain (K_p) and the integral controller gain (K_i) for a 5 Hz sinusoidal trajectory (Figure 2.4(b)). The result of the minimization was that there were optimal controller gains (circle symbol) which yielded improvement over two suboptimal examples (x & triangle symbols). This is shown more explicitly with simulated examples (Figure 2.4(c), bottom), where desired (green) and achieved (black) firing rates

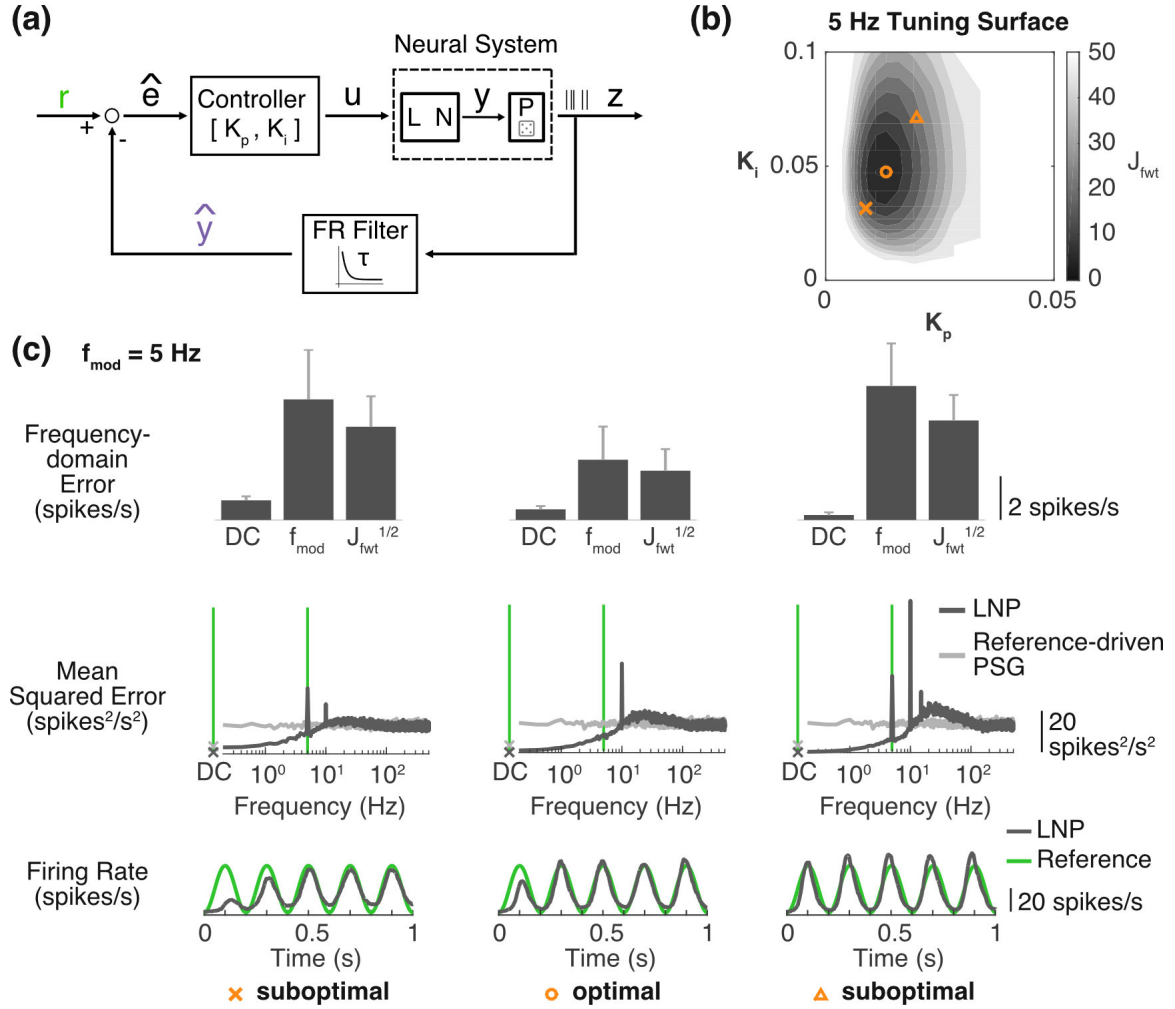


Figure 2.4: Controller Design: Tuning the controller around an LNP model neuron (a) Controller design through simulation. The closed-loop system was simulated with a model of the neural system for design purposes. (b) Example tuning surface for 5 Hz sinusoidal trajectory. In simulation, the controller was tasked with tracking a sinusoidal trajectory (here, 5 Hz), using the firing rate filter designed previously for the corresponding reference. The objective was to minimize the squared tracking error, weighted as a function of frequencies important for the control task. (c) Examples of optimal and suboptimal controller gains. Frequency-domain error (top row) corresponds to amplitude of error between the raw instantaneous rate (z/Δ) and the reference at DC and the modulation frequency (here, 5 Hz). For comparison, the square root of the frequency-weighted squared error (J_{fwt}) is also provided. Error bars correspond to +1 SD. Corresponding error spectra are provided (middle row), as compared to a simulated Poisson process modulated at the reference rate (light grey). Green lines highlight DC and f_{mod} . Finally, time-domain tracking is provided (bottom).

are plotted for the three sets of control parameters. In the middle row of Figure 2.4(c), the resulting error spectra for these three tunings are compared to the error expected for a Poisson spike generator (PSG) that has been driven at the reference rate (light grey). The spectrum for the reference-driven PSG shows the error that results purely as a function of random spiking, rather than off-target firing rate modulation. Because of the use of feedback, the spikes resulting from simulated closed-loop control are not truly random, leading to less error at low frequencies than would be expected in the Poisson case (black vs. grey error spectra). The peaks in error at $2f_{mod}$ (here, 10 Hz) in Figure 2.4(c) occur because the controlled neuron can spontaneously fire even when the reference rate is zero, resulting in errors at the troughs of each period. This in combination with error that occurs at peaks of the sinusoidal reference results in some power at double the modulation frequency. In comparison to the other two examples, the optimal tuning clearly reduces the error at both DC and f_{mod} . For the sinusoidal trajectory used here, the tuning procedure involved minimizing the combined error at DC and modulation frequency, illustrated in Figure 2.4(c) (top) as the square-root of the frequency-weighted squared error metric (J_{fwt}).

2.3.4 Model accuracy and closed-loop performance

While the closed-loop control system is model-free in implementation, the above results relied upon knowledge of a model relating the light input for optogenetic stimulation and the neuronal firing. During experiments, we therefore estimated a model that captures this relationship between the light input, $u \in \mathbb{R}^1$, and the measurement, $z \in \mathbb{R}^1$, (Figure 2.5(a)) for use in controller design. We utilized a band-limited “white-noise” optical input for stimulating the thalamic neuron and fit a simple linear-nonlinear-Poisson (LNP) cascade to capture these responses (see Section 2.2.4). An example LNP fit is shown in Figure 2.5(b). Note that the finite impulse response (FIR) stimulus filter, or kernel, is plotted time-reversed for visualization.

To understand how accurate these LNP models were, we tested their ability to pre-

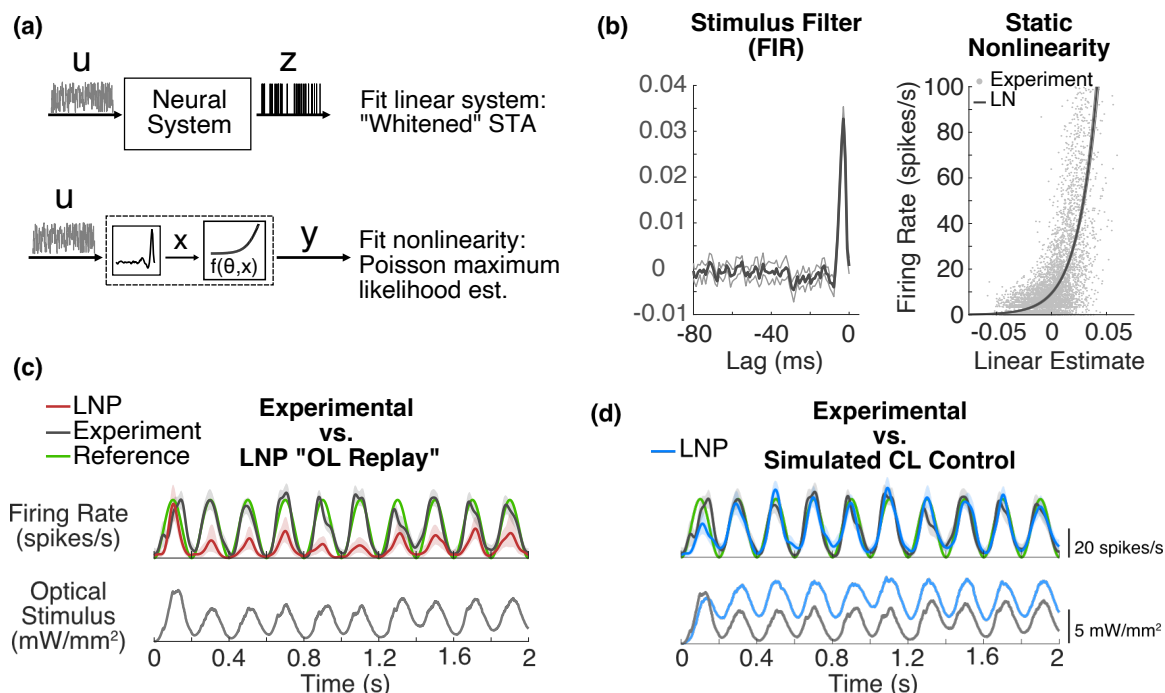


Figure 2.5: LNP Model Performance: Open-loop vs. closed-loop (a) Fitting the linear-nonlinear (LN) model components. (b) A typical LN model fit to training data. Left, the stimulus FIR filter estimated using whitened spike triggered averaging; error bars correspond to ± 1 SD for the lagged coefficients of the kernel when fit to 10 subsets of the full training dataset. Right, the static nonlinearity fit by Poisson maximum likelihood, given output of the linear stage and measured spikes. Gray points indicate the experimental firing rate (PSTH smoothed with 1-ms Gaussian) versus the kernel-filtered stimulus. (c) LNP prediction of response to open-loop 'replay' of stimulation used experimentally during a 5 Hz sinusoidal control task, using the same LNP model shown in (b). Top, firing rate predicted by the model (red) as compared to the experimental data for the same cell (black); bottom, experimental optical input. (d) LNP prediction of response to simulated closed-loop stimulation. Firing rate and light input predicted by the model (blue) as compared to the experimental data for the same cell (black); bottom, simulated (blue) and experimental (black) closed-loop stimulation.

dict the input-output relationship in the operating regime observed during a control epoch. Upon investigating how the LNP models responded to control inputs used during closed-loop stimulation experiments, we found that the models produced varying levels of success in predicting the observed firing activity of the thalamic neurons. Shown in Figure 2.5(c) is one such representative example where light inputs used experimentally for tracking a 5 Hz sinusoidal target trajectory were presented to the neuron’s corresponding LNP model. The response of the LNP model neuron (red) did not accurately predict the experimental rate (black). Given a wealth of literature suggesting difficulty in obtaining predictive models of neuronal firing *in vivo* beyond the sensory periphery (*e.g.*, [89, 157]), this finding is perhaps unsurprising.

However, because these models were used for tuning the controller gains in simulation, it is more important to assess their predictive capability in the context of closed-loop operation. The use of feedback for control conferred some robustness to model inaccuracy, and although the LNP model was generally not a good predictor of neuronal firing activity in open-loop, the model produced outputs that agreed with experimental results when simulated in the context of closed-loop control. Figure 2.5(d) provides an example where the same model neuron in Figure 2.5(c) was simulated in the closed-loop system parameterized by the firing rate filter time constant (τ) and controller gains used experimentally. In this case, the firing rate of the model neuron (blue) more closely matched the experimentally-realized firing (black).

2.3.5 Robustness of control to model inaccuracy

Given the observation that the LNP models used for controller tuning were generally poor predictors of experimental data in the open-loop sense, it is important to know to what extent the controller performance was robust to modeling error in this context. Therefore, we set out to determine more systematically through simulation how robust control performance was to inaccuracy in two identified LNP model parameters: gain and bias. Taking an

LNP fit to experimental data, the PI controller was tuned around it, as would be done during an experiment. We then tested these controller parameters on perturbed versions of the original LNP model, where the static gain and the bias term of the linear component of the LN model were changed systematically. To quantify the distance between the performance of the original and perturbed systems, the percentage change in tracking performance was calculated by comparing the error (J_{fwt}) for the perturbed models to that of control around the original LNP. Further, to assess the effect of model inaccuracy on “steady-state” tracking performance, J_{fwt} was calculated from 1 second onward for 5-second control epochs. To inspect for added benefit of closed-loop stimulation as opposed to open-loop, the tracking error was assessed for the perturbed models in both closed- and open-loop contexts. For the latter, light inputs used to control the original LNP were presented to the perturbed versions of the model in open-loop.

Decrements in control performance due to model inaccuracy are illustrated in Figure 2.6 for tracking a 5 Hz sinusoidal reference, using an example LNP fit to experimental data. The schematic in Figure 2.6(a) shows the two parameters that were varied: the static gain (g) and the bias term (d) of the LNP model (see Section 2.2.4). Altering g increases or decreases the apparent ‘sensitivity’ of the neuron to optical drive. The bias term reflects the baseline firing rate in the absence of optical drive. Contrary to Section 2.2.4, for this analysis, LNP models were fit with α equal to 1 to control for differences in the knee of the nonlinearity. Note also that in the analysis shown here, the shape of the kernel was retained. The parameters g and d were perturbed over a range extending from five times smaller to larger than the original values for which the controller was designed and the decrements in controller performance quantified (Figure 2.6(b)).

While there is interaction between the effects of the two parameters on control performance, changing the bias led to binary effects, as the performance quickly transitioned from optimal to very poor, moving from left to right in Figure 2.6(b). This observation is made plain when holding the static gain at the nominal value and fractionally changing

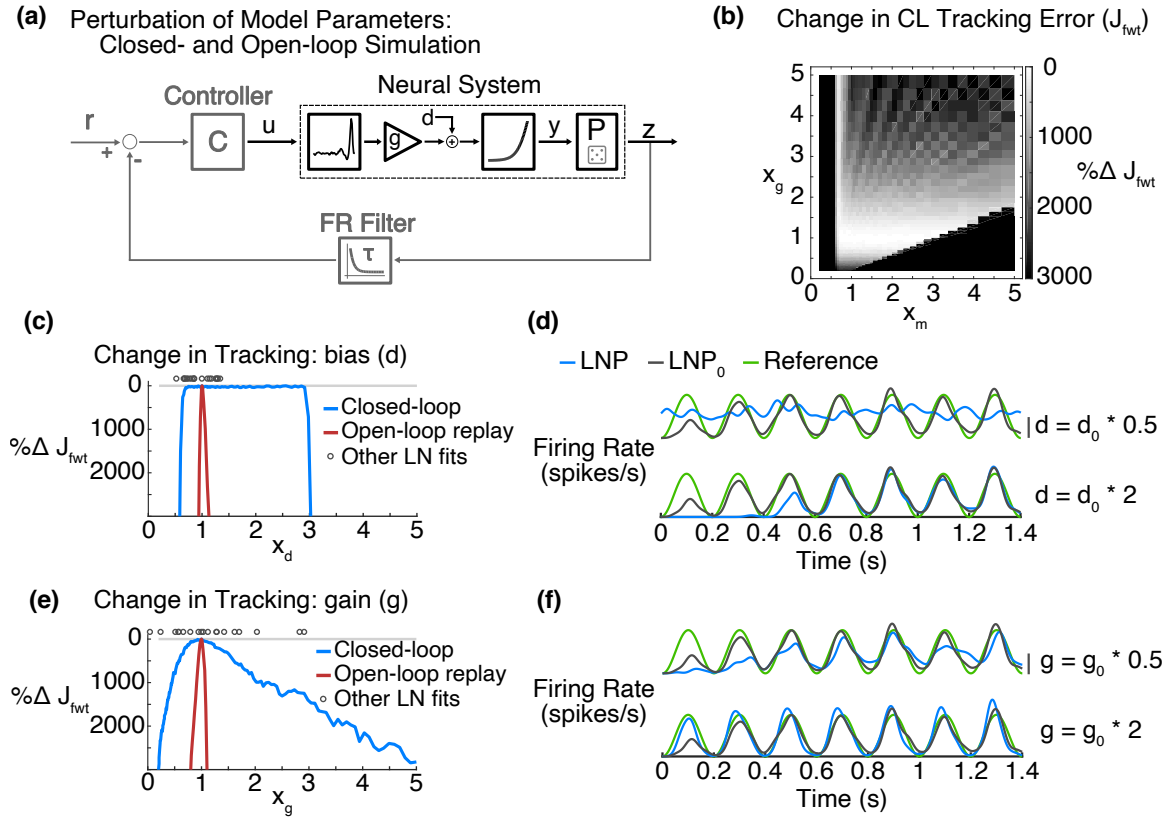


Figure 2.6: **Robustness of Control to Model Inaccuracy** (a) Model Perturbation and Simulation. The static gain (g) and the bias (d) of the linear component of the LNP model were systematically perturbed and simulated either in closed- or open-loop. (b) Grid Search Over Gain and Bias. Fractional changes in g and d ranged from five times smaller to five times greater than the original parameter value. Grayscale represents the percentage change in tracking performance (J_{fwt}) between that of the original model and each perturbed version. This tracking error was calculated from 1 second onward for 5-second control epochs. (c) Changing Bias. Holding all else constant, d was changed, and the tracking performance was quantified for closed-loop control vs. open-loop replay of the light traces used to stimulate the original neuron. Gray circles indicate the estimated biases of all recorded neurons relative to the model used for perturbation study ($n=20$). (d) Examples of Simulated Control When Bias Estimation Inaccurate. (Top) Outcome when the actual neuron (blue) was 2 times less biased than the model around which controller were tuned (black); (Bottom) the outcome when the actual neuron (blue) was 2 times more negatively biased than the model around which the controller was tuned (black). Scale bar indicates 20 spikes/s. (e) Changing Gain. Holding all else constant, g was changed and the tracking performance quantified for closed-loop control vs. open-loop replay of the light traces used to stimulate the original neuron. Gray circles indicate the estimated gains of all recorded neurons relative to the model used for perturbation study ($n=20$). (f) Examples of Simulated Control When Gain Estimation Inaccurate. (Top) Outcome when the actual neuron (blue) was 2 times less sensitive than the model around which controller were tuned (black); (Bottom) the outcome when the actual neuron (blue) was 2 times more sensitive than the model around which the controller was tuned (black). Scale bar indicates 20 spikes/s.

the bias (Figure 2.6(c)). Note that the identified bias term was negative. As the bias term approached zero (*i.e.*, x_d decreases), this increased the baseline firing rate of the model neuron. Conversely, as the bias became more negative than the original value (x_d increases), this mapped the output of the linear portion of the LNP model to the highly nonlinear portion of the static nonlinearity.

In Figure 2.6(c), the change in tracking performance for a 5 Hz sinusoid is shown for both closed-loop stimulation (blue) as well as open-loop replay of the stimulation pattern used to manipulate the original (*i.e.*, unperturbed) model (red). Except at extremes, closed-loop control was very robust to inaccurate estimation of the bias term. At these extremes, either the neuron's baseline firing rate exceeded the reference firing rate or the output of the linear component of the model was mapped to a regime of the nonlinearity where the LED or other light source was incapable of delivering inputs intense enough to raise the firing rate. In both of these scenarios, there would be little if anything a controller could do to salvage performance. Figure 2.6(d) provides simulated examples for when the neuron being controlled (blue) was two times less biased (top) or two times more negatively biased (bottom) than the neuron for which the controller was tuned (black). In the first case, the baseline firing rate of the neuron in blue was too high to be effectively controlled. Conversely, in the case where the bias of the neuron was more negative (bottom), it took more light input (not shown) and, therefore, more time for the system to reach “steady state” behavior. However, once the controlled system reached steady state, the tracking performance (blue) was identical to the that of original model neuron the controller was designed around (black).

Holding the bias term at the original value and instead perturbing the static gain, Figure 2.6(e) reveals that while the impact of inaccurate estimation of static gain on controller performance was more graded, closed-loop control of time-varying trajectories was less robust to inaccurate estimation of gain than to the bias term. As before, it was the case that closed-loop control (blue) offers a buffer against the effects of such model inaccuracy

as compared to open-loop stimulation (red). Figure 2.6(f) provides examples for when the controller was tuned around a model that was two times more sensitive (top) or less sensitive (bottom) to light than the neuron being controlled in blue. In the case where the neuron being controlled was less sensitive than the model around which the controller was designed (top), the controller achieved the correct DC firing rate, but was more weakly modulated than the target oscillatory activity. Conversely, where the neuron being controlled was more sensitive to optical drive, there were periodic overshoots of the reference. Unlike the case of the bias term, there are relatively simple actions that could be taken in the future to ameliorate errors due to inaccurate gain estimation, including online re-estimation.

2.3.6 Control summary: Tracking a sinusoidal trajectory

Figure 2.2 provided a single example of tracking at 1 Hz. Here, we expand on this by presenting results for eliciting a 5 Hz sinusoidal pattern in thalamic neurons recorded in separate animals. Figure 2.7(a) provides an experimental example for tracking a 5 Hz sinusoid at “steady-state” 3 seconds into a control epoch. In this example, controller performance reached the level predicted in simulation for the LNP model used for design (“neuron 3” in Figure 2.7(b), left vs. middle). Again, the tracking error was quantified as the frequency-weighted squared error (J_{fwt}) between the raw measurement of rate (z/Δ) and the reference (r), either for the LNP model used to design the controller (Figure 2.7(b), left) or for the experiment (middle, right). In the case of the sinusoidal trajectories used here, J_{fwt} reflects even penalty placed on error at DC and the modulation frequency. As for the perturbation analysis (Section 2.3.5), this metric was calculated from 1 second onward for 5-second control epochs. 95% confidence intervals for the metric were calculated for simulated Poisson firing at the reference rate and plotted as a light grey horizontal band (Figure 2.7(b)).

In Figure 2.7(b), each colored line corresponds to a different neuron recorded in a separate animal, and the black bars represent the population average in simulated closed-loop

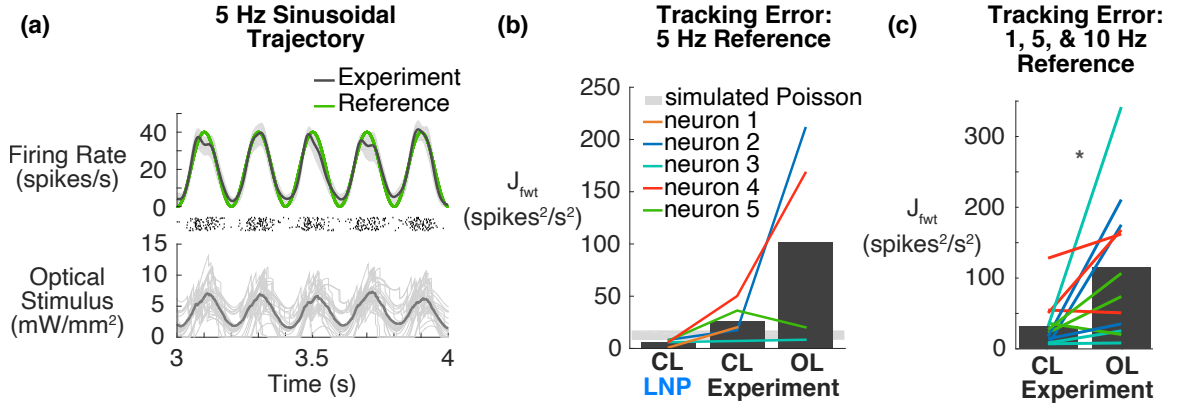


Figure 2.7: Sinusoidal tracking performance (a) Example experimental implementation (“neuron 3”) where the controller and firing rate filter were tuned for a trajectory modulated at 5 Hz. The third of a five-second control epoch is shown. (b) Population performance for tracking a 5 Hz sinusoidal trajectory (LNP prediction vs. experimental): average (black bar) and individual cells (colored symbols). 95% confidence intervals for this metric were calculated for simulated Poisson firing at the reference rate and plotted in light grey. Left, results of design procedure predicted by the LNP models fit and tuned around during the experiment. Middle, experimental closed-loop tracking performance. Right, experimental open-loop tracking performance. (c) Closed- vs. open-loop experimental performance on 1, 5, & 10 Hz sinusoidal control tasks. Closed-loop tracking error is significantly less than open-loop ($p < 0.05$, Wilcoxon signed rank test, $n = 12$ comparisons, 4 different cells).

(left), experimental closed-loop (middle), and experimental open-loop (right). Here, open-loop stimulation was a continuously modulated light signal, designed as previously detailed (Section 2.2.6). Simulated control around the LNP model used for tuning provides a lower bound for the control error (CL LNP, left). While the level of performance achieved in simulation was not achieved experimentally (Figure 2.7(b), middle), reasonable control was nonetheless achieved for all neurons in closed-loop. In contrast, open-loop control was not as robust, with the designed stimulation sometimes resulted in much worse performance than the closed-loop scenario (Figure 2.7(b), right).

The control loop was also tuned for and tested on 1 and 10 Hz sinusoidal trajectories. The tracking error of closed- vs. open-loop control is provided for all frequencies tested (*i.e.*, 1, 5, 10 Hz) in Figure 2.7c. When all sinusoidal tasks are considered, closed-loop control provides significantly lower tracking error ($p < 0.05$, Wilcoxon signed rank test).

2.3.7 Control summary: Tracking a non-sinusoidal trajectory

For the non-sinusoidal rate trajectory estimated from previously recorded spiking in the awake animal, 95% of the power in the reference signal was below 10 Hz. We found that of control systems designed for 1, 5, or 10 Hz sinusoidal trajectories with the same DC firing rate (20 spikes/s) as this non-sinusoidal signal, controller gains and filter time constants designed for the 10 Hz trajectory outperformed lower frequencies (not shown).

Figure 2.8(a) provides an example where the system was tuned for a 10 Hz sinusoidal trajectory and tasked with tracking the non-sinusoidal pattern of rate modulation (“neuron 5” in Figure 2.8(b), color code same as in Figure 2.7), as well as its corresponding error spectrum (Figure 2.8(a), bottom). For comparison, the error spectrum of simulated Poisson firing at the reference rate is plotted in light grey. After a period of approximately 750 ms, the neuron’s firing rate faithfully followed the reference. A controller tuned for a 10 Hz sinusoidal trajectory did approximately as well as within-experiment simulations predicted (Figure 2.8(b)). Again, controller performance was quantified using the frequency-weighted squared error. Note that for this non-sinusoidal reference trajectory, 50% of its power lies at DC and 45% of the remaining power ranges from DC to 10 Hz; therefore, half the penalty is placed on achieving the correct average firing rate, while the remaining half of the penalty is primarily exerted at and below 10 Hz. In all cases, open-loop stimulation proved less effective than closed-loop. Furthermore, as was first shown qualitatively for a single cell in Figure 2.2(c), the response to open-loop stimulation was also more variable than in the case of closed-loop stimulation. We observed this phenomenon across cells ($n=5$), where the spike-count variability, as measured by time-averaged Fano factor calculated in a 250-ms sliding window, was greater in open-loop than in closed-loop (Figure 2.8(c)).

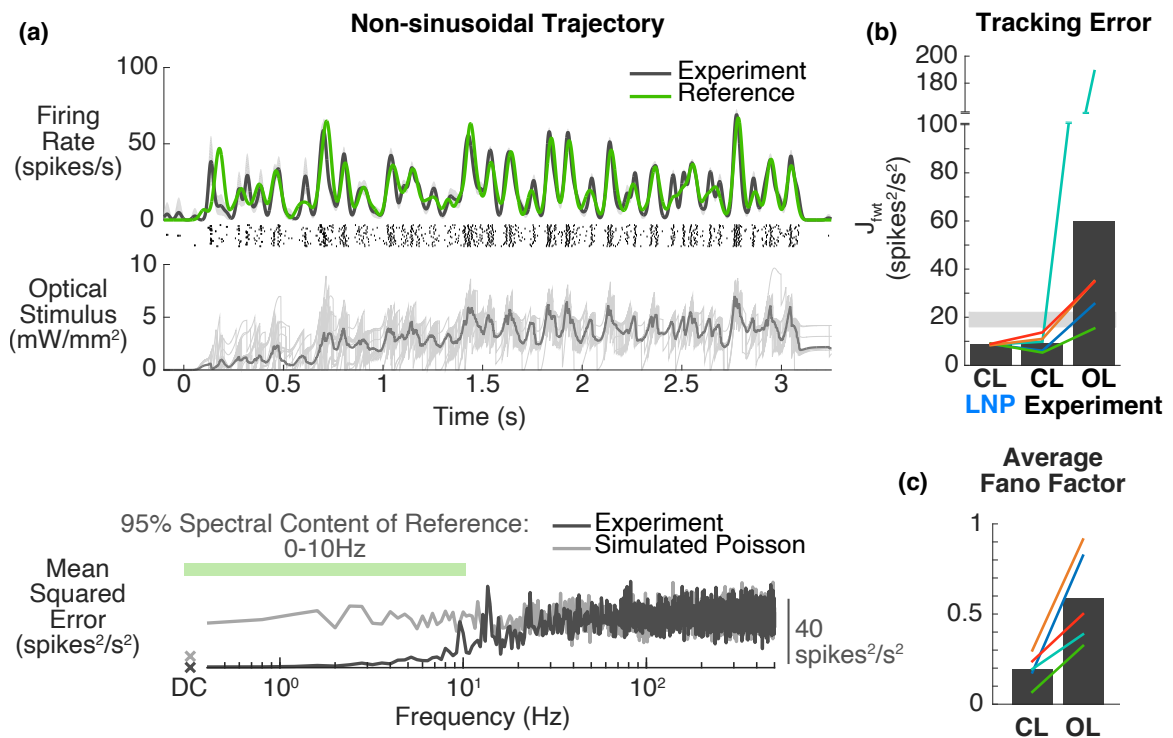


Figure 2.8: Non-sinusoidal tracking performance. (a) Example implementation (“neuron 5”) on the more naturalistic, non-sinusoidal trajectory. Controller and filter were tuned for a 10 Hz sinusoidally modulated trajectory. The corresponding error spectrum is also shown (bottom). For comparison, the error spectrum of simulated Poisson firing at the reference rate is plotted in light grey. (b) Population tracking error for non-sinusoidal trajectory. 95% confidence intervals for this metric are calculated for simulated Poisson firing at the reference rate and plotted in light grey. Left, simulated LNP performance predicted by design procedure. Middle, experimental closed-loop tracking performance. Right, experimental open-loop tracking performance. (c) Experimental across-trial variability in closed-loop vs. open-loop. Treating 750-ms onward as ‘steady-state’, time-averaged Fano factor calculated in a 250-ms sliding window for the closed- vs. open-loop control cases.

2.4 Discussion

The advent of new tools for measuring and manipulating activity within complex neural circuits opens up a wealth of possibilities for more interactive electrophysiological experiments, where feedback is used to inform stimulation continuously. Previously, we provided an initial demonstration for optogenetic control of neural activity with continuous use of feedback using a simple model-free control strategy ([6]). Specifically, the previous study was a proof-of-concept concerned with the demonstration of holding firing rates constant over long control epochs. Further, while holding firing rate steady is certainly of scientific interest for probing the functions of neural circuitry near steady-state, generalization of the closed-loop control approach to time-varying patterns of firing would open up a range of new applications and lines of investigation. Therefore, in the present study we have developed and demonstrated a design strategy for using closed-loop stimulation to track time-varying patterns of firing rate.

Ranging from intracellular current injection to electrical and optogenetic stimulation, there has been previous work using open- and closed-loop control strategies to manipulate neural firing patterns (*e.g.* [158, 159, 136, 95, 92, 137, 73]). Because spiking is often thought of as the information currency of the nervous system, such studies are usually concerned with achieving a target train of spike times. However, as a matter of practical concern, some jitter in the elicited spike times is tolerated, whether explicitly as in [92] or by means of a quadratic penalty between a spike time and the next desired spike in a train ([137]). For this reason, the problem of controlling spike times can be reconceptualized as a problem of controlling a time-varying pattern of firing rate, where the timescale of rate variation depends upon the tolerated jitter in spike timing. Therefore, while we have chosen to tackle the problem of manipulating instantaneous firing rate, the design methodology laid out here could in principle be applied to the problem of controlling spike times, where the allowable jitter sets the timescale over which desired spike times are smoothed

into a rate function that is fed to the control loop as a reference (*e.g.*, [158]).

However, we have presently tasked a control system with tracking not precise spike timing, but relatively slow patterns of rate modulation (1-10 Hz). Therefore, in this study we are not exercising control over spiking at timescales faster than approximately 100 ms. Notable rate modulation in this slower range of timescales occurs across neural systems, including in the context of active sensation [27, 160], hippocampal theta-phase precession [161], and movement [162]. The ability to insert such patterns of rate modulation would enable causal investigation into how instantaneous rate affects the function of such systems. Moreover, the basic framework developed here could be extended to finer timescales, opening up additional avenues of investigation that revolve around questions of precision timing.

The choice to use firing rate as the controlled output of the system necessitated designing a filter for estimating this quantity from noise-corrupted measurements in the form of spikes. In the interest of simplicity needed for widespread adoption of this technique, we took a model-free approach to estimation, without taking into account any dependence of the rate on inputs to the system. The obvious benefit of this approach is that it allows an experimenter to design this part of the control loop prior to the experiment, and it is simple to implement. However, given the point process nature of the measurements, this approach is not ideal for recovering time-varying rates, depending on the timescale of rate modulation. This is especially true of low firing rate regimes: *e.g.*, at modulation frequencies greater than approximately 2 Hz, with an average rate of 20 spikes/s (Figure 2.3), where estimates become increasingly poor.

Given the difficulty of accurate rate estimation at higher frequencies of rate modulation, we took the practical approach of designing the controller with this firing rate output filter in the loop. That is to say, we did not make the customary assumption that we could design the controller for the case where the latent state/controlled output (here, y) was known. Instead, a Poisson spike generator and filter were included in the simulated control loop when tuning

the PI controller. By minimizing the error between the measured spiking and the reference at frequencies of interest for control, we endeavored to safeguard performance of the final closed-loop system against deleterious effects of imperfect online rate estimation.

Another design choice made in this study was to directly modulate the amplitude of optical input to the system. While channelrhodopsin (ChR) is most often stimulated using pulsatile inputs, our choice to use feedback to modulate the amplitude of light was made on the basis that it requires the fewest design decisions for the experimenter. In contrast to continuous modulation of amplitude, manipulating pulsatile inputs would require a mapping between the control signal and pulse amplitude, width, and frequency (*e.g.*, [6]). That said, there is nothing preventing the details of the methodology laid out here to be applied to the modulation of a single pulse parameter such as amplitude or width. We expect that pulsatile and continuously-modulated modes of optical stimulation will have differential effects on higher-order aspects of activity such as local population synchrony [6]. While a robust way to stimulate ChR-expressing cells, pulsatile inputs may not suit all control applications, such as manipulating subthreshold neuron polarization. Another alternative to continuous modulation or modulation of pulsatile inputs would be to wait to update the stimulation intensity each time the neuron spikes, as in [159]. In either case, the current methodology may still be applied if the simulations used for controller tuning were altered to reflect the chosen implementation of the control signal.

In this study, we have demonstrated the successful use of a single-degree-of-freedom controller for tracking a time-varying trajectory of interest. In such a system, the designed controller is necessarily a compromise between feeding forward an optimally transformed version of the reference and making best use of feedback to attenuate the influence of disturbances. In this case, the goal of the design procedure was tracking a time-varying reference trajectory. The success of such a simple approach is noteworthy in that it indicates scaling of the reference trajectory is a reasonable first-order approach to stimulus design for the desired rate modulation, while the feedback signal ensured tracking at or near DC by

providing the appropriate bias to this input. The downside of the use of a single-degree-of-freedom controller and tuning for trajectory tracking is that there are no explicit constraints put on the response to load disturbances or other changes in the system. Satisfying requirements on both reference tracking and disturbance rejection independently begs for a two-degree-of-freedom control strategy [163].

Since the predominant manner in which stimulation is applied to neural systems is through open-loop control and event-triggered or on-off closed-loop control ([57, 131, 58, 132, 143, 5]), we have compared the effectiveness of our current approach to an open-loop strategy. The principal reason one might utilize feedback continuously would be to compensate for disturbances, whether these come in the form of unmeasured inputs to the system or changes in the underlying dynamics. Although not an explicit goal of controller tuning in this study, the use of feedback does, of course, grant some robustness to such disturbances. As shown previously for the static reference case ([6]) and again here in the case of a time-varying reference (Figure 2.2), the use of feedback enables effective control even in the face of uncontrolled input to the system in the form of sensory drive. However, a noteworthy limitation of the current single excitatory opsin approach is the inability to actively inhibit neural activity. As a consequence, disturbances which raise the firing rate above the desired rate at timescales of interest for control will not be effectively rejected. This scenario will necessitate a two-input approach, whether it be expression of both inhibitory and excitatory opsins or expression of opsins targeted to inhibitory neurons. Notably, this extension will require additional modeling for capturing effects of inhibitory inputs (whether direct or indirectly mediated through inhibitory neurons) as well as design of multi-input controllers. A state-space approach to modeling and control would generalize to such applications.

Related to the ability to reject disturbances, another difference between the results of closed- vs. open-loop stimulation highlighted in Figure 2.2 and again in Figure 2.8 is a reduction in across-trial variability provided by reactively updating stimulation in real-time.

Given that stimulation of neural systems is often plagued by response variability, this effect of closed-loop stimulation will certainly be a positive for many applications. However, it is also the case that there is naturally-occurring variability in spiking within a given cell and across a population [164, 165, 156]. There may be scenarios, such as artificial stimulation aimed at naturalistic perception, in which variability will be an important goal for control. Since adding response-variability will be an easier task than reducing it, the current approach represents a good initial step toward a more nuanced control of neural activity.

At least in an anesthetized preparation, we have shown that it is feasible to undergo a brief period of system identification followed by controller tuning in simulation within the timespan of an experiment. The fact that the LNP models used for this process were often poor predictors of the input-output relationship in the open-loop sense and yet closely predicted the performance of the closed-loop system highlights the need to assess the goodness of a model in the context of its intended use. If, for example, the use of a mathematical model is purely as a tool for design rather than for making mechanistic inferences, then more complicated models that may be more accurate but also more difficult to estimate may prove unnecessary. Indeed, while not tested here, a nonlinear spiking model may not have been necessary for effective control, given the robustness of closed-loop stimulation to the inaccuracies already present in the LNP model. It is possible a linear model would have been sufficient in this context.

Since the controller was designed around an identified model for the system, it is important to ask how sensitive, or conversely how robust, control system performance is to model inaccuracy. If the resulting design is sensitive to small deviations in estimated parameters, either the modeling/system identification approach or the control scheme should be modified accordingly. Here, we have found that even with often inaccurate LNP models, we achieved reasonable controller performance (Figures 2.7&2.8). Although closed-loop stimulation does relax the requirements for model accuracy, we have also shown that something

as simple as inaccurate estimation of static gain of the LNP model can have deleterious effects on control performance in the case of a time-varying reference signal. Importantly, changes in response properties such as gain and baseline firing rate often occur over time in neural systems in the contexts of sensory adaptation [86, 84, 166, 88], level of arousal [20], *etc.* It is noteworthy that changes in response gain observed in thalamic neurons in the context of sensory input (*e.g.* [86, 88]) are often on par or greater than those seen here in the context of optogenetic stimulation. Therefore, the question of controller robustness will be important moving forward.

2.5 Conclusions

Here, we have found that when appropriately tuned, the simple exponential firing rate filter and PI controller framework used by Newman *et al* [6] to maintain a constant firing rate can also be effective for eliciting time-varying firing rate trajectories. The strength of this approach lies in its relative simplicity, which should render it readily adoptable for the neuroscience community. With the availability of an open-source platform for real-time control of electrophysiology (RTXI, [167]), the design methodology laid out here should enable widespread application of closed-loop control to optogenetics experiments.

Of particular note is the fact that the simple models used for controller tuning could be poor predictors of the light-to-spiking transformation and still prove useful. This phenomenon highlights that it is not necessary in all contexts to fit the best possible mathematical model. Instead, modeling decisions should be made in light of the intended end use, where considerations such as expedience or mathematical tractability may be of great importance. Indeed, it is possible that in the context of this study, an even simpler model structure could have sufficed. Moreover, the finding that closed-loop optogenetic stimulation can reduce spike-count variability at the timescales investigated highlights the technique's utility for more reliable electrophysiological recordings and opens the door to control strategies that make variability explicit goals of the design process. Taken together,

the methodology developed in this study will lay the groundwork for more refined use of optogenetic stimulation and may enable a new class of experiments aimed at elucidating the functional roles of neural populations in networks.

CHAPTER 3

STATE-SPACE MODEL-BASED FEEDBACK CONTROL OF OPTOGENETICALLY DRIVEN NEURAL ACTIVITY

3.1 Introduction

Over the last two decades, there has been a rapid expansion of tools and technologies for recording the large-scale activity within and across brain structures at single neuron resolution ([63, 168]). In parallel, the development of optogenetics provided the ability to optically excite or inhibit neural activity in a cell-type specific manner ([3]). Together, these advances in the ability to “read” or “write” the neural code have led to a wide range of discoveries of the circuit mechanisms underlying sensory, motor, and cognitive processes ([169]). The integration of recording and optogenetic stimulation techniques, however, has received comparatively little attention until recently ([6, 138, 126, 132, 57, 131, 133, 134, 69]; for review [5]), and in most cases these closed-loop systems utilize event-triggered or on-off control rather than continuous feedback. While feedback control is the engineering cornerstone for the function of a wide range of complex technologies ranging from communication to flight, applying this perspective in the nervous system remains more theoretical ([135, 170, 158, 137, 141]) than experimental. In this study, we establish a general framework for continuously-modulated closed-loop optogenetic control of neuronal circuits, where optical actuation is determined in real-time by comparing measured neuronal spiking to target activity. This work opens up possibilities for investigation of poorly understood mechanisms of the underlying circuitry and for adaptively interacting with the circuit dynamics within and across brain regions that constantly change in response to the internal and external environments.

Electrical stimulation has been the gold standard for manipulating the activity of neu-

rons at fast time-scales, and remains the basis for clinical interventions like deep brain stimulation ([171]). However, this approach suffers from lack of specificity while also typically precluding simultaneous measurement of the activity of the neurons being stimulated. While not yet clinically viable, optogenetics offers an alternative approach that enables cell-type specificity, bi-directional actuation, the ability to simultaneously stimulate and obtain electrophysiological recordings, and a potentially lesser degree of unnatural synchronization of the local population ([72]). This presents an attractive toolbox for the development of continuous, feedback control strategies where stimulation is continuously modulated based on real-time measurements of the local neuronal activity. There has been a range of studies where previously-determined stimulation is triggered based on recorded activity in a reactive closed-loop fashion ([57, 131, 132, 133]). In addition to event-triggered control, a recent study has also used on-off closed-loop control to gate photostimulation when recorded neuronal activity was below target levels ([69]). Although these approaches to stimulation have proven effective for their uses, they are fundamentally different from the continuously-graded feedback control we describe here, where stimulation is updated on a moment-by-moment basis as a function of the current and past measured neural activity. In previous studies, we have developed and demonstrated strategies for closed-loop optogenetic control of spiking activity of neurons in a cultured network and single neurons *in vivo* in the anesthetized brain ([6, 126]). While laying the conceptual groundwork, these approaches do not scale well to neuronal populations and do not take advantage of more modern approaches in control theory. Additionally, these previous studies had not yet applied optogenetic control in the context of wakefulness.

In this study we bridge the gap between optogenetics and established paradigms of more modern control theory by utilizing state-space models to capture single- and multi-neuron responses to optogenetic stimulation and employing optimal control to design the control loop for driving desired neuronal activity. Specifically, precise manipulation of neurons was carried out via optical activation of the excitatory opsin channelrhodopsin-2

(ChR2) expressed in the somatosensory thalamus of the awake, head-fixed mouse. A feedback controller updated light intensity in real-time based on simultaneous electrophysiological measurements of the thalamic neurons being manipulated. A linear dynamical system model structure was used to approximate the light-to-spiking input-output relationship in both single-neuron as well as multi-neuron scenarios in cases where multiple neurons were measured simultaneously using multielectrode arrays. These linear state-space models were used in combination with linear quadratic optimal control to design feedback controller gains for the purpose of regulating thalamic firing around a desired target rate. The models were also used online for estimation of state feedback, using a parameter-adaptive Kalman filter for robustness to model-mismatch. The resulting controller-estimator feedback loop was deployed experimentally by way of a custom-written program running in real-time. This control scheme provided effective optogenetic control of firing rate in the awake brain, owing to the robustness to model inaccuracy granted by a parameter-adaptive Kalman filter that estimated a stochastically-varying process disturbance. Feedback control using this estimator resulted in very good firing rate tracking experimentally for the single neurons whose activity was used for feedback. By comparison, control was not as effective for other simultaneously-measured neurons not used for feedback. To investigate the generalizability and efficacy of these methods for future multi-output control scenarios, we demonstrate their application to multi-neuron feedback control of a population in simulation.

3.2 Methods

3.2.1 Animal preparation

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 NIH. Experiments were carried out using either C57BL/6J mice that were virally transduced to express channelrhodopsin-2 (ChR2) or by single-generation crosses of an

Ai32 mouse (Jax) with an NR133 cre-recombinase driver line (Jax) which grants better specificity of ChR2 expression in ventral posteromedial/posterolateral thalamus ([123]). In the case of viral transfection, ChR2 expression was targeted to excitatory neurons in the thalamus via stereotactic injection relative to bregma (approximately $2 \times 2 \times 3.25$ mm caudal \times lateral \times depth) using 0.5 μ L of virus (rAAV5/CamKIIa-hChR2(H134R)-mCherry-WPRE-pA; UNC Vector Core, Chapel Hill, NC) at a rate of 1 nL/s.

At least three weeks prior to recordings, a custom-made metal plate was affixed to the skull for head fixation and a recording chamber was made using dental cement while the animals were maintained under 1–2% isoflurane anesthesia ([172]). After allowing a week for recovery, mice were gradually habituated to head fixation over the course of at least five days before proceeding to electrophysiological recordings and optical stimulation. On the day of the first recording attempt, animals were again anesthetized under 1–2% isoflurane and a small craniotomy (1–2 mm in diameter) was centered at approximately 2×2 mm caudal and lateral of bregma. The animals were allowed to recover for a minimum of three hours before awake recording. At the time of recording, animals were headfixed and either a single electrode or an electrode array coupled to an optic fiber (Section 3.2.2) was advanced through this craniotomy to a depth between 3–4 mm for thalamic recording and stimulation. Between repeated recording attempts, this craniotomy was covered using a biocompatible silicone sealant (Kwik-Cast, WPI). Following termination of recordings, animals were deeply anesthetized (4–5% isoflurane) and sacrificed using a euthanasia cocktail.

3.2.2 Experimental setup

All optical stimuli were presented deep in the brain via a 200 or 100 μ m diameter optic fiber attached to a single tungsten electrode (FHC) or to a 32-channel NeuroNexus optoelectric probe in a 25 μ m-spaced “poly3” configuration (A1x32-Poly3-5mm-25s-177-OA32LP, NeuroNexus Technologies, Inc.), respectively. Command voltages were gener-

ated by a data acquisition device (National Instruments Corporation) in a dedicated computer running a custom-written RealTime eXperimental Interface (RTXI, [167]) program at 1 ms resolution. Command voltages were sent to a Thorlabs LED driver, which drove a Thorlabs M470F3 LED (470 nm wavelength blue light) connected to the 100-200 μm optical fiber. A commercially available data acquisition device and processor (Tucker Davis Technologies RZ2) measured extracellular electrophysiology. This system was used for single-channel PCA spike sorting, binning, and sending these binned spike counts at 2 ms resolution to the computer running RTXI over ethernet via UDP. The computer running RTXI for realtime control listened for datagrams over ethernet and linearly interpolated from 2 ms to the operating resolution of 1 ms. All told, the closed-loop processing loop was approximately 10 ms.

As mentioned above, the control and estimation algorithms were carried out in real-time at 1 ms resolution using a custom-written program. The program consisted of an RTXI “plugin” linked against a C++ dynamic library that was responsible for online estimation of state feedback (Section 3.2.5) and the generation of control signals (Section 3.2.6). This functionality was provided as part of a state-space controller C++ class. The RTXI plugin forwarded the reference, or target, firing rate, model parameters, and feedback controller gains to a state-space controller object, and the controller returned an updated control signal each time it was queried by RTXI. This control signal was then routed by RTXI to the LED driver via a DAC (see above). All linear algebra was carried out using the C++ library Armadillo ([173]).

3.2.3 Offline spike sorting

For online control applications, single-channel PCA-based spike sorting was carried out in real-time using a commercially available electrophysiology system (Tucker Davis Technologies RZ2). Beyond tetrode recordings, spike sorting from high-density electrode arrays requires a multi-step process that is not feasible within the timescale of experiments

with head-fixed awake animals. Kilosort2 ([174]) was used for all offline spike sorting, including single-channel recordings, in which case spatial whitening and common mode referencing steps were disabled. Initial sorting by Kilosort2 was then manually curated (additional merging/splitting of clusters) using the phy viewer. After manual curation, any clusters that met the following criteria were considered single units and used in this study: sub 1 ms ISI violations of $<0.5\%$, sub 2 ms ISI violations of $<2\%$, and mean waveform amplitude-to-standard-deviation ratio >4 .

3.2.4 Mathematical modeling

Linear and Gaussian state space models were used in designing feedback controller gains before experiments as well as during the experiment as part of the state feedback estimator. These models were fit offline to neuronal data before experimental application of optogenetic control and were fit at 1 ms time resolution, which was also the operating resolution of the RTX software used for real-time control and estimation during experiments. In addition to the single-unit quality selection criteria in Section 3.2.3, models were only fit to putative single units (called “neurons” hereafter) whose activity was significantly modulated by optical stimulation. Following Sahani and Linden ([175]), a neuron’s response was considered significantly modulated if the amount of “signal power” in the response was greater than one standard error above zero. Note that Sahani and Linden ([175]) define “power” as the variance in time. We will refer to “signal power” as “signal variance” in this study.

The underlying dynamics of neural activity were approximated as a linear dynamical system (LDS) in which a number of latent “state” variables, represented as the vector $\mathbf{x} \in \mathbb{R}^n$, evolve linearly in time:

$$\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{B}u_{t-1} + \mu_{t-1} + \mathbf{w}_{t-1} , \quad (3.1)$$

where $u_t \in \mathbb{R}^1$ is the optical stimulus at time t , $\mu_t \in \mathbb{R}^n$ is a process disturbance,

$\mathbf{w}_t \sim \mathcal{N}(0, \mathbf{Q})$ is Gaussian noise of covariance \mathbf{Q} , $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the state transition matrix, and $\mathbf{B} \in \mathbb{R}^{n \times 1}$ is the input vector (generally a matrix). Note that the disturbance, μ , was assumed to be zero during model fitting. However, for robustness in control applications, μ was allowed to be non-zero and to vary stochastically over time for the purpose of online state estimation (Section 3.2.5).

Gaussian linear dynamical system

Linear and Gaussian models were used for control system design and implementation because of the relative simplicity and ubiquity of linear control approaches. In this case, the output of an LDS $\mathbf{y} \in \mathbb{R}^p$ is modeled as a linear transformation of a latent state \mathbf{x} and is assumed to be corrupted by additive Gaussian noise before measurement in the form of binned spiking, $\mathbf{z} \in \mathbb{R}^p$:

$$\mathbf{y}_t = \mathbf{C}\mathbf{x}_t + \mathbf{d} , \quad (3.2)$$

$$\mathbf{z}_t = \mathbf{y}_t + \mathbf{v}_t , \quad (3.3)$$

where $\mathbf{d} (\in \mathbb{R}^p)$ is an output bias term that describes the baseline firing rates of the p outputs (here, neurons), and $\mathbf{v}_t \sim \mathcal{N}(0, \mathbf{R})$ is zero-mean Gaussian measurement noise of covariance $\mathbf{R} \in \mathbb{R}^{p \times p}$. As a system whose dynamics evolve linearly and whose observation statistics are assumed to be additive/Gaussian, this is termed a Gaussian LDS, or GLDS ([101]). The bias term \mathbf{d} was estimated as the average firing rate of each channel during spontaneous periods without optical stimulation, and GLDS models were fit relating u_t and $(\mathbf{z}_t - \mathbf{d})$ using subspace identification (N4SID algorithm, [176]).

Poisson linear dynamical system

While GLDS models were used for control and estimation, we evaluated their performance in capturing light-driven firing rate relative to an equivalent spiking model type. As it is

a more accurate statistical observation model for spike count data, we fit linear dynamical systems with Poisson observations, so-called Poisson LDS, or PLDS ([101, 97]). In this case, the underlying latent state(s) of the LDS is mapped to an output firing rate by a rectifying exponential nonlinearity and the measured spike counts are assumed to be drawn from a Poisson process driven at the given rate:

$$y_t^i = \exp(\gamma_i \mathbf{x}_t + d_i) , \quad (3.4)$$

$$z_t^i | y_t^i \sim \text{Poisson} [y_t^i] , \quad (3.5)$$

where y_t^i is the firing rate and z_t^i the measured spike counts of the i^{th} output at time t . For the purposes of this study, PLDS models were fit by first estimating a GLDS model. The row vectors γ_i that describe the log-linear contributions of each state to output firing rates were assumed to be scaled versions of the GLDS output matrix rows: *i.e.*, for the i^{th} output,

$$\gamma_i = g_i \mathbf{c}_i , \quad (3.6)$$

where \mathbf{c}_i is the corresponding row of the GLDS output matrix \mathbf{C} .

Note that at each time point the outputs are statistically independent conditioned on the state, allowing the output function parameters to be estimated in an output-by-output fashion. The resulting $2p$ -parameters of the PLDS output function were fit by maximizing the log-likelihood of the model one output at a time, given the predicted state sequence:

$$\theta_i^* = [g_i \quad d_i]^* = \arg \max_{\theta_i} \mathcal{L}_i(\theta_i) , \quad (3.7)$$

$$\theta_i^* = \arg \max_{\theta_i} \sum_{t=1}^T (z_t^i \log y_t^i | \mathbf{x}_t(\theta_i) - y_t^i | \mathbf{x}_t(\theta_i)) , \quad (3.8)$$

where θ_i^* are the parameters and \mathcal{L}_i the log-likelihood of the model for the i^{th} output, $(\cdot)^*$ denotes the result of the optimization, and

$$y_t^i | \mathbf{x}_t (\theta_i) = \exp (g_i \mathbf{c}_i \mathbf{x}_t + d_i) . \quad (3.9)$$

This optimization was carried out iteratively until parameter convergence for each output by analytically solving for d_i and numerically solving for g_i using Newton's method in a manner analogous to Smith *et al* ([97]).

Finite impulse response model

While state-space models were used in this study, finite impulse response (FIR) models were also fit in order to provide empirical estimates of the light-to-spiking responses that did not depend on choices such as number of latent states. Moreover, FIR models, often termed (“whitened”) spike-triggered average (STA) models, are widely used to characterize neuronal responses to stimuli ([82]), so they are a more familiar model type for much of the neuroscience community and provide a useful point of comparison for state-space models which are less frequently used in this context. Contrary to state-space models whose outputs share a set of dynamical states, in FIR models the optical stimulus (u) is related to the output firing rates (\mathbf{y}) of p neurons in the following manner:

$$\mathbf{y}_t = \mathbf{K} \mathbf{h}_t + \mathbf{d} , \quad (3.10)$$

where \mathbf{h}_t is a q -dimensional column vector of stimulus history up to time step t inclusive,

$$\mathbf{h}_t = \begin{bmatrix} u_t \\ u_{t-1} \\ \vdots \\ u_{t-q+1} \end{bmatrix} , \quad (3.11)$$

and \mathbf{k}_i is the impulse response of the i^{th} output, comprising the rows of \mathbf{K} . Note that this is effectively a convolution of a set of p FIR filters with the stimulus. The output \mathbf{y} is assumed to be corrupted by additive Gaussian noise before being observed/measured in the form of binned spiking, \mathbf{z} :

$$\mathbf{z}_t = \mathbf{y}_t + \mathbf{v}_t, \quad (3.12)$$

where \mathbf{v}_t is the measurement noise as described in the case of the GLDS model previously. This FIR model was fit by ordinary least squares linear regression between $(\mathbf{z}_t - \mathbf{d})$ and corresponding 100 ms stimulus histories, *i.e.*, $\mathbf{h}_t \in \mathbb{R}^{100}$ at $\Delta = 1$ ms sample period.

Optical stimulus for model fitting

While the approaches in this study can generalize to multi-input systems (*e.g.*, multiple light sources spread spatially or multiple wavelengths), only single-input systems are considered and tested here. As in Bolus *et al* ([126]), a repeated 5-second instantiation of 1 ms resolution uniform optical noise was used to stimulate spiking activity for model fitting. While the amplitude of this stimulus varied across experiments based on perceived neuronal sensitivity to light, the average range of this uniform-distributed noise was from 0 to 14.4 mW/mm², and the same pattern of noise was always presented. State-space models were fit using data from the first 2.5 seconds of each stimulus trial, while the remaining 2.5 seconds of stimulation were held out and used to assess model performance.

3.2.5 Estimator

GLDS models were used both offline for designing the control law and online for estimating state feedback. For online estimation, two variants of GLDS model-based state estimation are considered. The first is a standard implementation of the Kalman filter (Section 3.2.5, [108, 177]). Another variant of this approach that was used to achieve greater robustness to plant-model mismatch was to apply Kalman filtering to estimate a parameter-augmented state vector (Section 3.2.5), which we will refer to here as a parameter-adaptive Kalman

filter but has elsewhere been described as a proportional-integral (PI) Kalman filter ([178, 179]).

Kalman filtering

The Kalman filter proceeds by alternating between a one-step prediction of the state and updating this estimate when the corresponding measurement is available ([177]). The filter has two design parameters which are reflected in the GLDS model structure (Equations 3.1 & 3.3): the covariances of the process and measurement noise, or \mathbf{Q} and \mathbf{R} , respectively. The value for \mathbf{R} was taken from fits of the GLDS models to training data. In analyzing the performance of the Kalman filter on previously-collected spiking data, the fit matrix for \mathbf{Q} was rescaled to minimize the mean squared error (MSE) between the Kalman-filter-estimated firing rate and an output of a model-free estimation method: smoothing the spikes with a 25 ms Gaussian window.

At each time point, a one-step prediction of the estimated state mean ($\hat{\mathbf{x}}$), state covariance (\mathbf{P}), and output ($\hat{\mathbf{y}}$) were calculated:

$$\hat{\mathbf{x}}_{t|t-1} = \mathbf{A}\hat{\mathbf{x}}_{t-1|t-1} + \mathbf{B}u_{t-1} , \quad (3.13)$$

$$\mathbf{P}_{t|t-1} = \mathbf{A}\mathbf{P}_{t-1|t-1}\mathbf{A}^\top + \mathbf{Q} , \quad (3.14)$$

$$\hat{\mathbf{y}}_{t|t-1} = \mathbf{C}\hat{\mathbf{x}}_{t|t-1} + \mathbf{d} , \quad (3.15)$$

where $\widehat{(\cdot)}$ denotes estimates, $(\cdot)_{t|t-1}$ denotes the prediction at time t , given data up to time $t-1$, and $(\cdot)_{t|t}$ denotes filtered estimates. Recall that all model parameters were fit to optical noise-driven spiking activity, and note that μ was assumed to be zero unless adaptively re-estimated (Section 3.2.5). The one-step prediction was updated taking into account the

latest measurement as

$$\mathbf{K}_t^{\text{est}} = \mathbf{P}_{t|t-1} \mathbf{C}^\top (\mathbf{R} + \mathbf{C} \mathbf{P}_{t|t-1} \mathbf{C}^\top)^{-1} , \quad (3.16)$$

$$\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t^{\text{est}} (\mathbf{z}_t - \hat{\mathbf{y}}_{t|t-1}) , \quad (3.17)$$

$$\mathbf{P}_{t|t} = (\mathbf{I} - \mathbf{K}_t^{\text{est}} \mathbf{C}) \mathbf{P}_{t|t-1} , \quad (3.18)$$

$$\hat{\mathbf{y}}_{t|t} = \mathbf{C} \hat{\mathbf{x}}_{t|t} + \mathbf{d} , \quad (3.19)$$

where $\mathbf{K}_t^{\text{est}}$ is the Kalman filter gain and \mathbf{I} denotes an identity matrix.

Parameter-adaptive Kalman filtering

For robustness of state estimation to plant-model mismatch, the state and a model parameter were jointly re-estimated by the Kalman filter. Specifically, the mean of the process disturbance, μ , was assumed to vary stochastically over time as a random walk:

$$\mu_t = \mu_{t-1} + \mathbf{w}_{t-1}^\mu , \quad (3.20)$$

where $\mathbf{w}_t^\mu \sim \mathcal{N}(0, \mathbf{Q}_\mu)$ is noise disturbing the stochastic evolution of μ . The covariance of this process \mathbf{Q}_μ effectively sets the timescale of adaptive re-estimation of μ . To jointly estimate this disturbance, the state and model parameters were augmented as follows:

$$\mathbf{x}_t^{\text{aug}} = \begin{bmatrix} \mathbf{x}_t \\ \mu_t \end{bmatrix} , \quad (3.21)$$

$$\mathbf{A}^{\text{aug}} = \begin{bmatrix} \mathbf{A} & \mathbf{I} \\ 0 & \mathbf{I} \end{bmatrix}, \quad (3.22)$$

$$\mathbf{Q}^{\text{aug}} = \begin{bmatrix} \mathbf{Q} & 0 \\ 0 & \mathbf{Q}_{\mu} \end{bmatrix}, \quad (3.23)$$

$$\mathbf{B}^{\text{aug}} = \begin{bmatrix} \mathbf{B} \\ 0 \end{bmatrix}, \quad (3.24)$$

$$\mathbf{C}^{\text{aug}} = \begin{bmatrix} \mathbf{C} & 0 \end{bmatrix}. \quad (3.25)$$

In general, such joint parameter-state estimation would require the use of the extended Kalman filter (*e.g.*, [180]). However, in this case, the augmented dynamics and output equations remain linear with respect to the augmented state. Therefore, Kalman filtering was carried out on this augmented form of the state and GLDS model as detailed before in Section 3.2.5. For the purposes of this study, \mathbf{Q}^{aug} was assumed to be a diagonal matrix. In analyzing the performance of this adaptive Kalman filter on spiking data, the elements of \mathbf{Q}^{aug} were scaled to minimize the mean squared error between the Kalman-filter-estimated firing rate and the Gaussian smoothed estimate as before (Section 3.2.5).

3.2.6 Controller

While the state-space modeling and control framework can be readily used for trajectory tracking, the control objective in this study was holding the output neuronal firing to a fixed target, or reference, rate (r), corresponding to a nonzero-setpoint regulation problem ([181]), also described here as “clamping”.

Control setpoint

In order to use state feedback for the case where the target is an output, we first calculated the state and optical input that would be required to achieve the target firing rate, \mathbf{r} . Since this was a regulation problem, we calculated the state and input for achieving the target at steady-state. This steady-state setpoint $[\mathbf{y}^{*\top} \mathbf{x}^{*\top}]^\top$ was calculated using models fit to previously collected optical noise driven data. This problem was solved by linearly-constrained least-squares [182], where the objective was to minimize the 2-norm $\|\mathbf{y}^* - \mathbf{r}\|^2$, subject to the system being at steady-state $\mathbf{x}^* = \mathbf{A}\mathbf{x}^* + \mathbf{B}u^*$. The control signal required to achieve the target at steady state, u^* , was served as a nominal control signal, about which feedback controller gains modulated light intensity. For single-input/single-output (SISO) applications, there was a solution that resulted in zero-offset tracking (*i.e.*, $\mathbf{y}^* = \mathbf{r}$). However, for multi-output control where the responses to control are heterogeneous, the steady-state solutions do not result in zero-offset tracking, but rather the least-squares compromise across neurons.

Linear quadratic regulator design

Linear quadratic optimal control was used to design controller gains K^{ctrl} for non-zero-set-point regulation ([181]):

$$u_t = u^* - K^{\text{ctrl}} \begin{bmatrix} \mathbf{x}_t - \mathbf{x}^* \\ \sum_{i=1}^t (\mathbf{y}_i - \mathbf{y}^*) \Delta \end{bmatrix}, \quad (3.26)$$

where both instantaneous state error (top row) as well as integrated output error (bottom row) were used for feedback to ensure robustness of control. Δ is the sample period (1 ms). The controller gains were chosen to minimize a quadratic cost (J) placed on these tracking

errors and on deviations in the control ([181, 28]):

$$J(K^{\text{ctrl}}) = \sum_{t=1}^{\infty} \frac{1}{2} \begin{bmatrix} \mathbf{x}_t - \mathbf{x}^* \\ \sum_{i=1}^t (\mathbf{y}_i - \mathbf{y}^*) \Delta \end{bmatrix}^{\top} \mathbf{Q}^{\text{ctrl}} \begin{bmatrix} \mathbf{x}_t - \mathbf{x}^* \\ \sum_{i=1}^t (\mathbf{y}_i - \mathbf{y}^*) \Delta \end{bmatrix} + \quad (3.27)$$

$$+ \frac{1}{2} (u_t - u^*)^{\top} r^{\text{ctrl}} (u_t - u^*) ,$$

where \mathbf{Q}^{ctrl} is the weight placed on minimizing squared instantaneous state error and integrated output error,

$$\mathbf{Q}^{\text{ctrl}} = \begin{bmatrix} \mathbf{C}^{\top} \mathbf{C} & 0 \\ 0 & q^{\text{int}} \mathbf{I} \end{bmatrix} , \quad (3.28)$$

and r^{ctrl} is the weight placed on control deviations. Minimization of this quadratic cost function is linearly constrained by the error system dynamics

$$\begin{bmatrix} \mathbf{x}_t - \mathbf{x}^* \\ \sum_{i=1}^t (\mathbf{y}_i - \mathbf{y}^*) \Delta \end{bmatrix} = \begin{bmatrix} \mathbf{A} & 0 \\ \mathbf{C} \Delta & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{x}_{t-1} - \mathbf{x}^* \\ \sum_{i=1}^{t-1} (\mathbf{y}_i - \mathbf{y}^*) \Delta \end{bmatrix} + \begin{bmatrix} \mathbf{B} \\ 0 \end{bmatrix} (u_{t-1} - u^*) . \quad (3.29)$$

This optimization was carried out numerically by backward recursion of the discrete-time matrix Riccati equation until convergence ([181]) or calculated using the MATLAB function *dlqr()* (MathWorks). Generally, a stabilizing solution was not possible for multi-output control scenarios with integral action because of nonzero output error; however, the numerical solution for feedback controller gains still converged in practice.

Experimental SISO control

First-order GLDS models fit to previously collected spiking responses to optical noise were used offline for designing feedback controller gains, K^{ctrl} , (Section 3.2.6) and online for the parameter-adaptive Kalman filtering (Section 3.2.5). The diagonal elements of the assumed process noise covariance, \mathbf{Q}^{aug} , used in the parameter-adaptive Kalman filter ranged from 1×10^{-9} to 5×10^{-8} . For controller design, the quadratic weight chosen for integral error

(q^{int}) was 1×10^2 , while the weight placed on control deviation, r^{ctrl} , ranged from 1×10^{-4} to 1×10^{-3} . The online-sorted spiking data fed back to the controller was used to assess performance of the control scheme; however, in cases where a 32-channel electrode array was used for recording, offline-sorted population activity was inspected to understand the local effects of closing the loop around a given putative single neuron.

Simulated SIMO vs. SISO control

In addition to experimental validation in the SISO case, the state-space modeling, estimation, and control methods were also applied to a simulated multi-output control problem in which the objective was to push the outputs toward a common target firing rate. In this case, 5th-order models were used. When fitting GLDS models to SIMO datasets, we found that there was often great heterogeneity in input-output gain across outputs. Therefore, a two-output PLDS model was the simulated system being controlled, whose second output (“neuron 2”) was a gain-modulated version of the first (“neuron 1”), before exponentiation and spike generation. The dynamics and the first output channel of this PLDS came from a fit to an example SISO dataset. The log-linear gain of neuron 2 was swept between 0.1 and 3 times that of neuron 1. A multi-output controller and estimator were designed using a 2-output GLDS model fit to simulated PLDS data, where optical noise stimulated the PLDS in the case where the both neurons had the same gain. The neuron-averaged mean squared error performance of the SIMO control loop was compared to the SISO scenario when only neuron 1 data was fed back. For both SIMO and SISO control loops, the diagonal elements of the process noise covariance for the parameter-adaptive Kalman filter (\mathbf{Q}^{aug}) were all taken as 1×10^{-6} , while the weights placed on quadratic cost of integrated tracking error (q^{int}) versus control deviation (r^{ctrl}) were 1×10^2 and 1×10^{-3} , respectively.

3.2.7 Performance measures

Various measures of performance are used throughout this study to quantify goodness of fit for state-space models and the effectiveness of the estimators as well as the controller.

Model performance

The performance of GLDS and PLDS models were assessed using variance of the raw 1 ms binned PSTH explained in the held-out second half of each 5 second trial of optical noise stimulation. The variance explained was either taken as a proportion of the variance in the PSTH (pVE), or relative to the amount of “signal” or explainable variance in the PSTH (pSVE, [175]). These two metrics were computed for each SISO and SIMO dataset for 5th order PLDS models and 1st and 5th order GLDS models.

Estimator performance

Because the control objective in this study was to track a constant reference firing rate, it was important that the estimator achieve low bias; otherwise, the integral action of the controller cannot serve its ideal purpose to eliminate steady state tracking errors. Therefore, the performance metric considered here for the online estimator was the squared bias of the single-trial-estimated firing rate compared to the corresponding spiking responses to 5-second step inputs of light.

Control performance

To assess controller performance, the mean squared error (MSE) as well as squared bias between the achieved single-trial firing rate and the reference firing rate were calculated. Single-trial firing rate was taken as the online-sorted spike train fed back to the controller, smoothed offline with a 25 ms standard deviation Gaussian window. While MSE takes into account variance, we separately considered across-trial variability using the Fano factor ([183]) of spike counts in a 500 ms sliding window, a mean-normalized measure of

spike count variability. Finally, in cases where a 32-channel multielectrode array was used for recording local population activity, the degree of synchrony between simultaneously recorded neurons was quantified in a manner similar to Wang *et al* ([19]). Briefly, a cross-correlogram was constructed by binning the relative spike times of simultaneously recorded neuron pairs. To quantify degree of synchrony, the number of correlated events in a ± 7.5 ms window (N_{cc}) was normalized by the total number of spikes in a ± 50 ms window (N_{tot}):

$$\text{synchrony} = \frac{N_{cc}}{N_{tot}} . \quad (3.30)$$

Allowing 1-second for non-steady state performance, all four of these performance metrics were calculated in a 4-second period of time during closed-loop control. As a point of comparison, the same metrics were also calculated using 4-second periods of spontaneous data recorded between trials of closed-loop stimulation.

3.3 Results

In this study, we applied a model-based optimal control framework to the experimental control of neural activity in vivo using optogenetic stimulation. Specifically, we utilized the ventral posteromedial (VPM) region of the sensory thalamus in the vibrissa/whisker pathway of the awake mouse as an experimental model system, where single-unit electrophysiological recordings were obtained while optically stimulating light sensitive channels with an inserted optical fiber. The optimal control framework relies on a state-space representation of the optically-driven dynamics of neural activity. This model is used both for the offline design of the optimal controller and for the online estimation of state feedback. Although experimental results are presented from this specific pathway and brain region, the approach is directly applicable to others. Furthermore, while the methods used here generalize to trajectory tracking for multi-input and multi-output (MIMO) applications, we first focus on the single-input and single-output (SISO) case where the measured outputs

were single-unit spiking activity and the control objective was to track step commands (*i.e.*, clamp neural activity at a fixed target firing rate). In the context of these experiments, we were able to use a linear and Gaussian model to approximate light-driven spiking responses for the purposes of controlling firing rate; moreover, we found that a low order approximation of the neural dynamics was sufficient at least for the slow timescale control/estimation objectives studied here. In experiments where multi-electrode arrays were employed to record thalamic activity, we found that simultaneously-recorded neurons responded to optical stimulation with a high degree of diversity, motivating investigation of applicability of this control framework to multi-output scenarios. We applied this framework in a simulated single-input/multi-output (SIMO) scenario, where the “output” consisted of the activity of multiple simultaneously-recorded neurons, and the control objective was to force the population activity as close as possible to a common target firing rate. Feeding back multi-output population activity to the controller enhanced the robustness of the control scheme’s ability to drive the collective population activity to a desired target in the face of heterogeneity in sensitivity to light.

Figure 3.1(a) illustrates the control scheme that was implemented experimentally in the awake, head-fixed mouse, where an “optrode” consisting of an electrode attached to an optical fiber was inserted into the VPM. Given binned single-unit spiking activity, control and estimation was carried in realtime at 1 ms resolution using custom-written software (Section 3.2.2). We designed an estimator that generated an online estimate of the state of neural activity, and a feedback controller that maintained a target firing rate in the face of potential disturbances, such as reafferent sensory input (*i.e.*, whisker motion) and changing brain states.

To develop a control methodology that would generalize to MIMO applications, we applied a state-space model-based control and estimation scheme where the model is used not only in the design phase but as an online estimator for the control scheme (Figure 3.1(b)). The model structure utilized here was a linear dynamical system (LDS), where optical in-

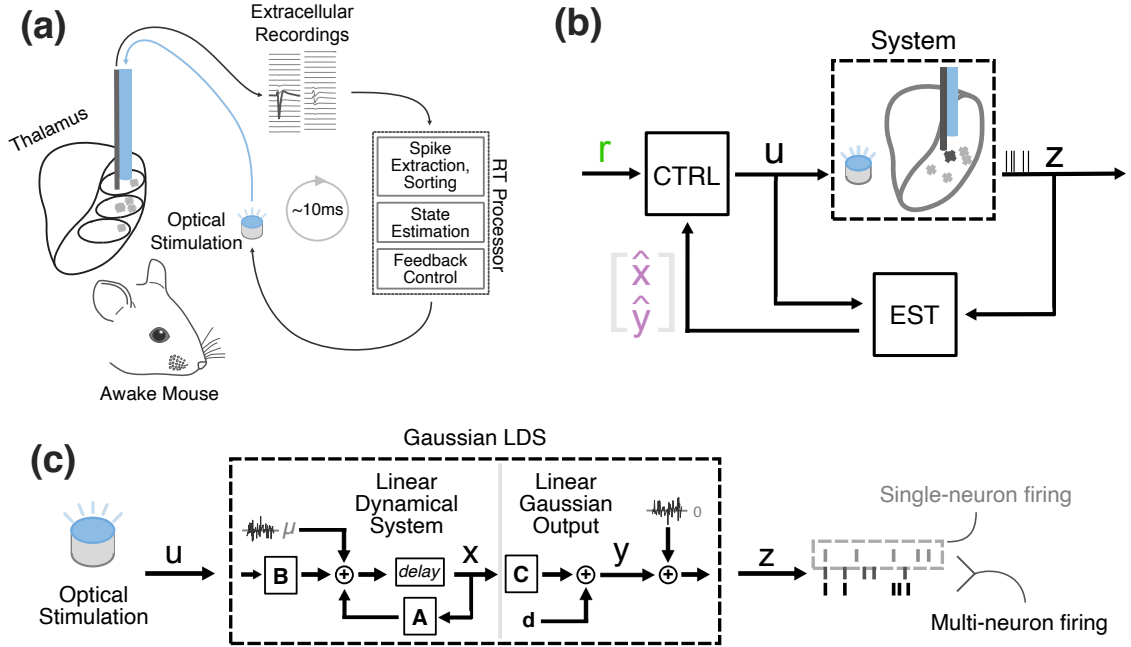


Figure 3.1: Closed-loop optogenetic control using state-space linear dynamical systems models. (a) Experimental Setup. (b) Control system block flow diagram. Spiking activity is fed back to a model-based estimator (“EST”), which provides online estimates of the underlying state of the system (\mathbf{x}) and the output (\mathbf{y}), which is firing rate in the current application. The controller (“CTRL”) uses a model to generate the system setpoint $[\mathbf{y}^{*\top} \mathbf{x}^{*\top} \mathbf{u}^{*\top}]^\top$ that corresponds to user-specified reference firing rate (\mathbf{r}). An updated control signal is generated using feedback controller gains and the error between this setpoint and the online estimates of the system state/output. The updated control signal is sent to an LED driver to modulate light intensity. (c) Structure of the Gaussian LDS Model. The GLDS used throughout the control loop consists of a linear dynamical system (LDS) describing the evolution of the state (\mathbf{x}) and a linear remapping of \mathbf{x} to the output firing rate and eventually measured spiking (\mathbf{z}). This model is used for single-neuron and multi-neuron estimation/control.

put(s) modulate the activity of latent state variables. More specifically, for the purposes of this study we employed a Gaussian linear dynamical system (GLDS), in which a linear combination of the states is observed after being corrupted by additive Gaussian noise (Figure 3.1(c)). Here, the output of the model was either single or multi-neuron firing rate, although in principle these same techniques could be applied to other neural signals of interest such as local field potential or voltage/calcium signals.

3.3.1 GLDS captures optical noise-driven responses

The control framework used here depends on a model of the underlying dynamics for both the design of the controller and online state estimation to execute the control strategy. As we have previously described in a simpler, classical control framework ([126]), feedback control is robust to a degree of model inaccuracy. Therefore, there is an application-specific balance to be struck between model complexity/fidelity and simplicity. Here, we first asked to what extent a GLDS model could predict the experimentally observed SISO firing rate modulation with optogenetic stimulation, as this would provide a relatively simple modeling framework that is attractive in terms of its widespread applicability and ease of implementation. Since the measurements were spike counts in 1 ms bins at relatively low firing rates, a Gaussian observation model is an obvious violation of these statistics. For comparison, we also fit an LDS model whose observation model is Poisson (PLDS), which has been utilized in a range of studies for describing the dynamics of spiking neurons (Figure 3.2(a)).

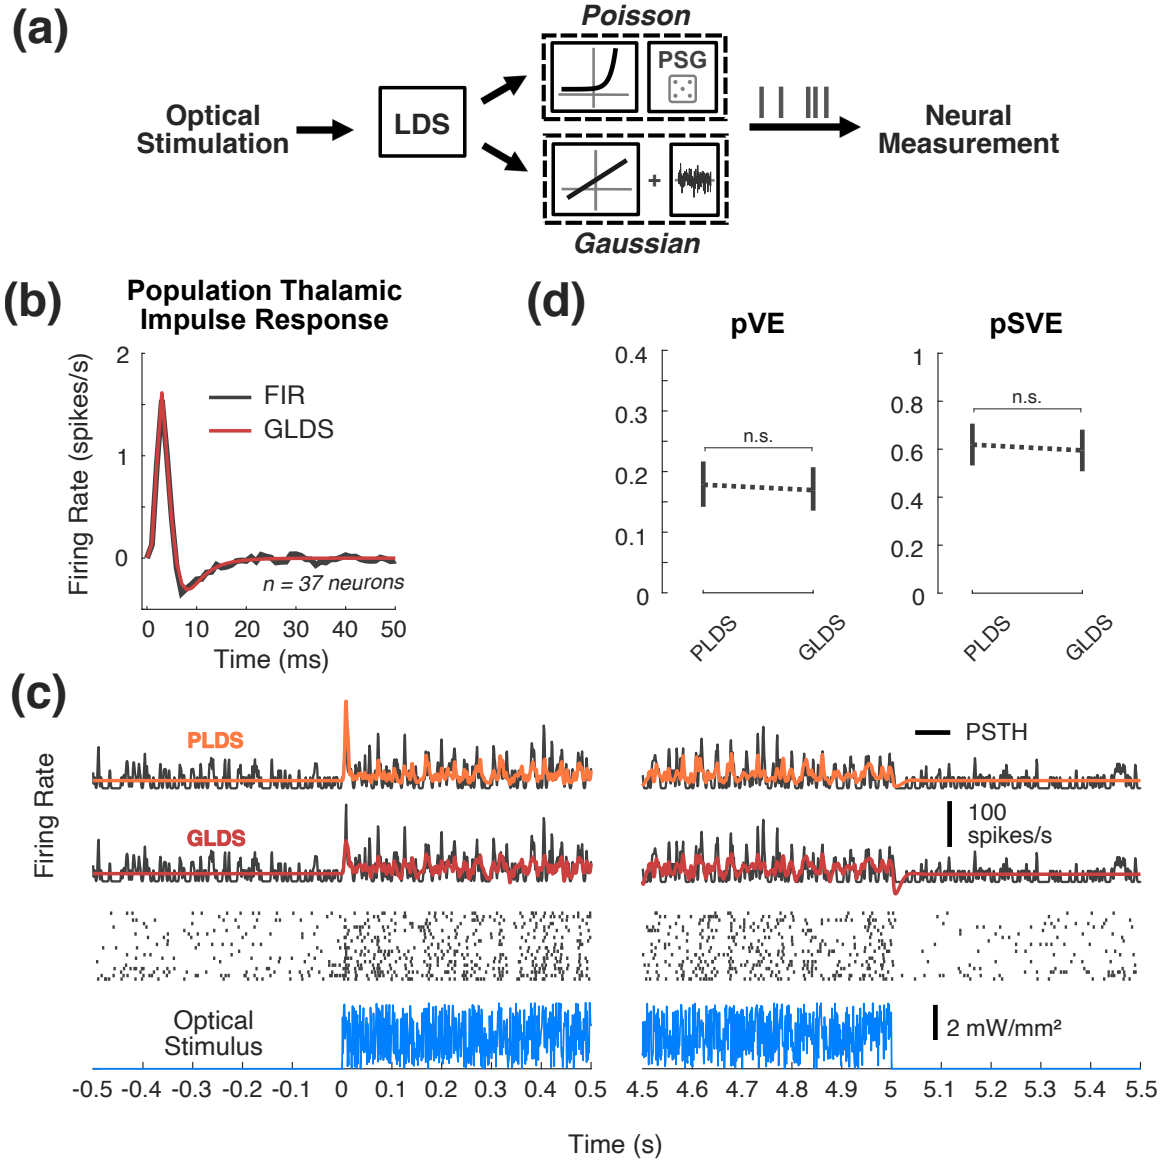


Figure 3.2: State-space models of SISO optogenetic responses. (a) SISO LDS model structure: Poisson (top) or Gaussian (bottom) output functions being considered. (b) Population impulse response. This impulse response was fit using pooled data from 37 neurons that were excited by optical noise. An FIR model fit to population data (black) is plotted alongside the impulse response from the 5th-order GLDS model fit to the same data (red). (c) Example Data and Model Fits. Top, the PSTH (black) was smoothed with a 1 ms standard deviation Gaussian window for visualization. The fit types include 5th-order PLDS (orange), 5th-order GLDS (red). Middle, the corresponding trial-by-trial spike raster. Bottom, repeated instantiation of uniform optical noise. (d) Proportion variance in PSTH explained (pVE) and signal variance explained (pSVE) by model response to noise. All models were trained on data from first half of each trial, while model performance metrics (pVE, pSVE) were calculated from the second half of each trial. Error bars represent bootstrapped 95% confidence intervals about the population mean ($n=48$ neurons, 17 recordings, 9 animals).

In fitting a state-space model, the order of the model (the dimensionality of the latent state vector) must be specified. To ascertain the appropriate order of these models, we pooled together noise-driven response data from 37 neurons that were all significantly excited by the optical stimulus, and fit GLDS models to this population. For comparison, we separately fit a finite impulse response (FIR) model to the same population dataset (see Section 3.2.4), as it is widely used in the neuroscience literature ([81, 82]). Models were fit from recorded responses to white-noise optical inputs (Section 3.2.4). Shown in Figure 3.2(b) are the impulse responses for both the GLDS (red) and FIR (black) models, as a head-to-head comparison. This can be interpreted as the model prediction of the instantaneous firing rate in response to a light impulse input at time zero. Prominent in both is an initial peak at approximately 3 ms reflecting a relatively short latency excitation, followed by a subsequent drop below baseline at 7-8 ms reflecting a post-excitatory inhibition. We found that a 4th to 5th order state-space model was sufficient for these data, striking a balance between goodness of fit and model complexity. Note that the above analysis was restricted to thalamic neurons that were found to be excited by the optical input, which excluded other thalamic neurons that exhibited more heterogeneous behaviour (i.e. a minority of recorded neurons were indirectly inhibited by the optical input, interestingly). To capture the full heterogeneity of the population, therefore, we fit 5th order PLDS and GLDS models to each single-output dataset individually ($n = 48$ neurons, 17 recordings in 9 mice). A representative example SISO dataset is shown in Figure 3.2(c), where the firing rate estimates for the 5th order PLDS (first row, orange) and 5th order GLDS (second row, red) are superimposed onto the corresponding PSTH (black) at white-noise onset and offset. Qualitatively, there is little gained in using a PLDS model instead of a GLDS for this example, aside from the non-negativity of the PLDS firing rate. Across the population of units, there is no significant difference between the performance of the Poisson vs. Gaussian models (Figure 3.2(d), $n = 48$ neurons, $p = 0.234$, Wilcoxon signed-rank test). Specifically, the left plot of Figure 3.2(d) shows the proportion of the variance in the raw

1 ms PSTH explained by 5th order PLDS and GLDS fits (pVE). Note that a relatively low proportion of the variance in the raw PSTHs was explained, due to levels of intrinsic noise in the observed responses at fine timescales. For this reason, we assessed the quality of the model using a metric that takes into account the fact that some of the observed variability is not explainable across trials ([175]), instead quantifying the amount of explainable, or “signal”, variance the model captures. The right panel of Figure 3.2(d) presents the proportion of the signal variance explained (pSVE), showing that the models captured approximately 60% of the explainable variance and that there was not a significant difference in the predictive capabilities between the GLDS and the PLDS models in this dataset. Therefore, with the exception of multi-output modeling where the same PLDS versus GLDS analysis was conducted for comparison, GLDS models are used for the remainder of this study in order to leverage linear controls approaches.

3.3.2 Parameter-adaptive Kalman filtering provides robust online estimation

These GLDS models are used online as part of the Kalman-filter-based estimator (Figure 3.3(a), grey box) which is used to provide state feedback to the controller. While the models performed relatively well in the case of uniform white-noise optical stimulation as shown in Figure 3.2, when challenged with step changes in input that are often utilized in control scenarios, non-zero-mean model mismatch is clearly revealed (Figure 3.3(b)). In this example the open-loop model predicted firing rate (OL Prediction, red) initially underestimates the experimentally-measured firing rate (PSTH, black) during the first second of stimulation and then consistently underestimates the firing rate at steady state. Model-based control and estimation schemes are particularly sensitive to such plant-model mismatch, as is apparent here when standard Kalman filtering used for online estimation is applied to these datasets for step changes in input. In this example in Figure 3.3(b) there is still an obvious bias in the average Kalman-filter estimated firing rate (KF Estimate, purple) when compared to the smoothed PSTH (PSTH, black). Moreover, because of the rapid time-

course of the fit neuronal dynamics (Figure 3.2(b)) and the spiking nature of the measurements, the single-trial KF estimates of firing rate which will be fed back to a controller are full of extreme transients each time a new spike is measured (Figure 3.3(c), purple trace). Online estimation of firing rate can be made more robust by assuming there is an unmeasured, non-zero-mean disturbance that varies stochastically (e.g., other exogenous inputs), augmenting the state with the mean(s) of this disturbance (μ), and jointly re-estimating this along with the state using Kalman filtering (Figure 3.3(d), see methods for details), which we refer to here as the parameter-adaptive Kalman filter, but has elsewhere been described as a proportional-integral Kalman filter ([178, 179]). As can be seen in the example in Figure 3.3(e), this adaptive Kalman filter produces an effectively unbiased estimate of the experimentally-observed PSTH in SISO applications (Figure 3.3(e), purple vs. black), and it is able to do so with a single-trial estimate of firing rate that is smoother than that achieved by the standard Kalman filter (Figure 3.3(f), *c.f.* Figure 3.3(c)). In this example, the parameter-adaptive Kalman filter approach accounts for apparent model mismatch by estimating a process disturbance μ that on average pushes the firing rate above the model prediction for the first second of optical stimulation and then pulls the estimated firing rate below that prediction at steady state (Figure 3.3(g)). The filtering approach works well in this illustrative example and at a population level, as it brings the estimation bias to near-zero levels compared to the standard Kalman filter (Figure 3.3(h), $p = 1.63 \times 10^{-9}$, Wilcoxon signed-rank test, $n = 48$ neurons, 17 experiments, 9 animals). At least in the context of estimating step responses, we see there is little benefit in using a 5th order versus 1st-order GLDS model for this SISO application (Figure 3.3(h), black, $p = 0.0830$, Wilcoxon signed-rank test). Importantly, the parameter adaptation provides enough robustness that even the population-average GLDS model in Figure 3.2(b) was able to estimate SISO firing nearly as well as models fit to each neuron individually (Figure 3.3(h), gray, $p = 0.0142$, Wilcoxon signed-rank test). Since the control objective in this study is to clamp firing rate at relatively long timescales, we therefore used a 1st-order Gaussian

approximation for the system. However, for fast timescale trajectory tracking problems, a higher-order model would almost certainly be warranted (see Discussion), and higher-order models are important even for long timescale control/estimation in multi-output scenarios (see Section 3.3.6).

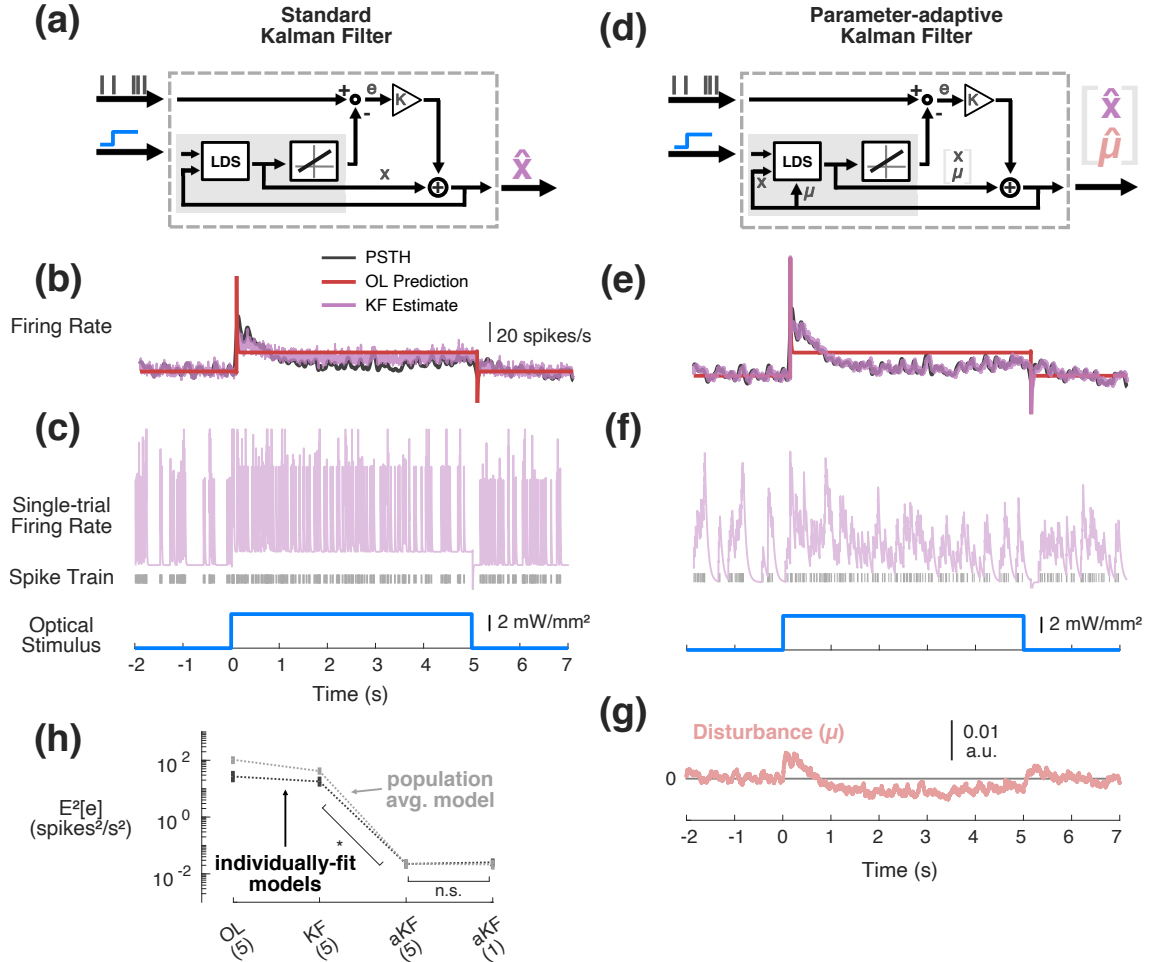


Figure 3.3: **Kalman filtering for online estimation in SISO applications.** (a) Standard Kalman filter. Prediction error (e) is used to correct the estimate of state at each time step. (b) Example open-loop (OL) prediction of neuronal response (red) to step input of light (blue) using 5th order GLDS fit to noise-driven data, compared to PSTH smoothed with 25 ms Gaussian window (black) and the trial-averaged response estimated using the standard implementation of the Kalman filter (5th-order GLDS) (purple). (c) Example single-trial Kalman filter estimate (purple) along with corresponding spike raster (grey). (d) Parameter-adaptive Kalman filter. In addition to estimating the state of the system, this approach jointly re-estimates a state disturbance (μ) at each time step. (e) Same as (b) but trial-averaged estimate of firing rate using the parameter-adaptive Kalman filter. (f) Same as (c) except single-trial estimate using parameter-adaptive Kalman filter. (g) Trial-averaged disturbance on the first state estimated using parameter-adaptive Kalman filter. (h) Population average squared-bias in estimation calculated between the single-trial spiking responses and the OL prediction of a 5th-order GLDS, the standard Kalman filter using the 5th-order GLDS, and the parameter-adaptive Kalman filter (aKF) using a 1st- or 5th-order GLDS. Black and grey data points correspond to error associated with using individually-fit models vs. a single population average fit model, respectively. Error bars represent bootstrapped 95% confidence intervals about the mean ($n=48$ neurons, 17 recordings, 9 animals).

3.3.3 State-space control performs well in SISO clamping applications

The control and estimation framework was tested experimentally in the awake head-fixed mouse in a SISO configuration, where spiking activity of a single neuron was fed back to a controller with a single channel of optical input. As previously noted, 1st-order GLDS models were sufficient and were therefore used for this particular application; however, higher order models would be merited or necessary in other scenarios. The robustness of the estimator (Figure 3.3), the use of feedback, and the slow timescale nature of the control objective allowed GLDS models fit to previously-collected noise response data to be used for experimental control and estimation, rather than fitting a model during an experiment, the timespan of which is limited in the context of awake, head-fixed recordings. The feedback controller was designed using output-weighted LQR ([181]), where the state of the system was augmented with the integrated output in order to find not only proportional feedback gains on the state, but integral feedback gains to minimize steady state tracking errors (Section 3.2.6). Additionally, since this particular application is a non-zero setpoint regulation problem, the steady-state set-point of the system $[\mathbf{y}^{*\top} \mathbf{x}^{*\top} u^{*\top}]^\top$ at the desired output firing rate (r) was calculated as described in Section 3.2.6.

Figure 3.4 illustrates the performance of the control framework for a typical single thalamic neuron and the summary performance across experiments. Figure 3.4(a) is an illustration of the control implementation, highlighting the feedback controller and the on-line estimator. In the case of the estimator, Parameter-adaptive Kalman filtering is being used to estimate not only the state of the system being controlled but also the uncontrolled disturbance (Figure 3.4(a), estimator block). On the other hand, the controller is operating on the error between the estimated state of the system and the desired steady-state set point as well as the integrated output error (Figure 3.4(a), controller block). In this example (Figure 3.4(b)), the baseline ongoing activity of the recorded neuron was approximately 5 spikes/s, and the controller was activated at time zero with a target firing rate of 20 spikes/s. Upon activating the controller, the neuron reached and remained at the target

firing rate (green), as reflected in the average firing rate (black). Importantly, the controller operated using online estimates of state and corresponding output firing rate provided by the estimator (Figure 3.4(b), purple). The firing rate of the online estimator (purple) also quickly reached the target (green) and remained there. As shown previously in Figure 3.3, the online estimate was on average unbiased, as it matched the offline estimate of the average firing (black, PSTH smoothed with 25 ms s.d. Gaussian). The controller achieved the target with well-below spontaneous levels of across-trial variability, quantified using the Fano factor (FF) that captures the spike count variance relative to the mean spike count (Figure 3.4(b), middle). In this particular example, the controller's use of feedback resulted in a gradual increase in light intensity that was needed to maintain the target level of spiking over the control epoch. It is possible that this is a signature of increased inhibitory feedback from the reciprocally connected reticular nucleus of the thalamus that the controller worked against to maintain target activity. Also note that this control signal varied substantially across individual trials (Figure 3.4(b), bottom, light blue), with significant individual trial variability serving to drive the firing rate tracking and quench the variability.

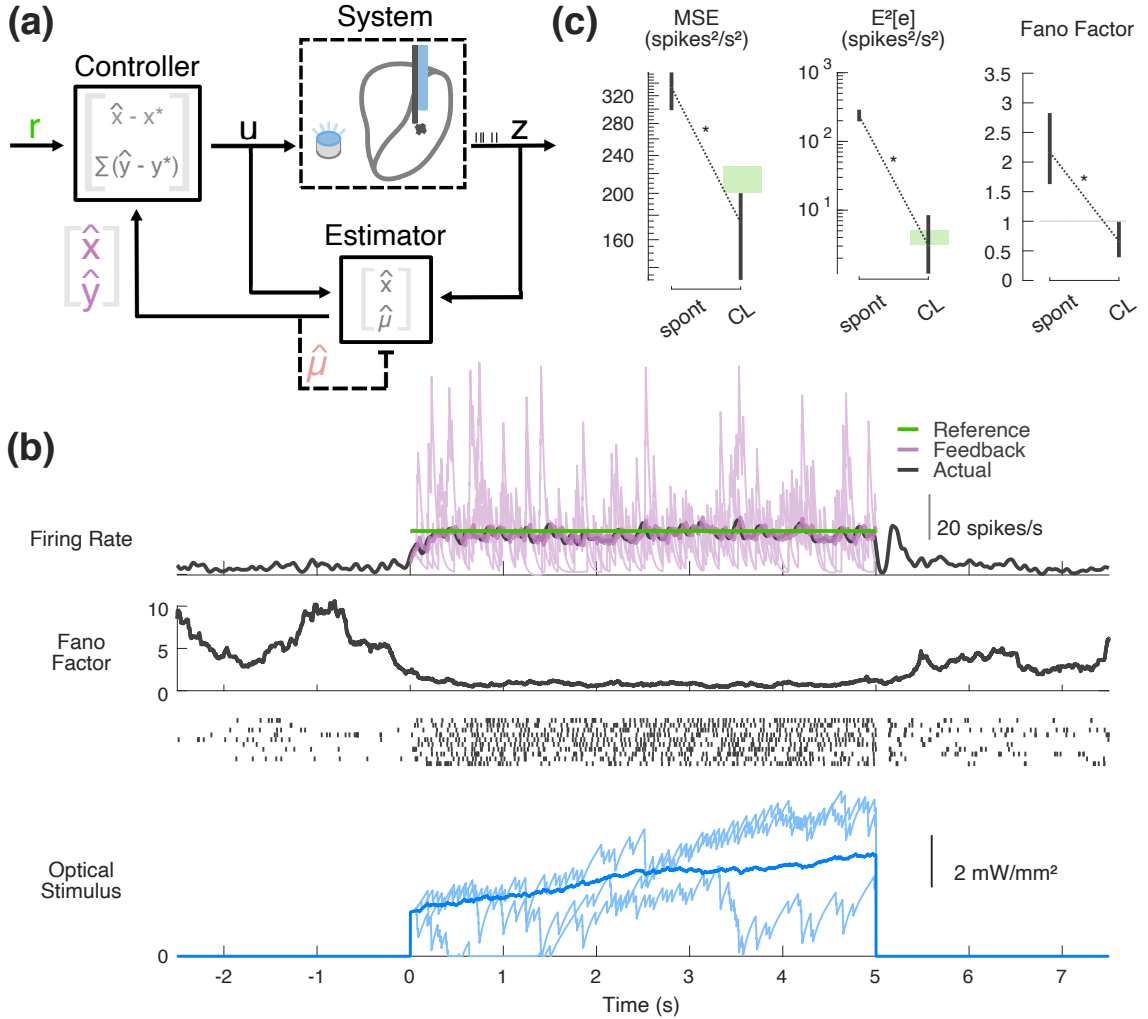


Figure 3.4: **Experimental SISO control and estimation.** (a) SISO control block flow diagram. Shown inside the controller and estimator blocks are the notions of state being used in each operation. (b) Example experimental SISO control. (top) Fed-back online estimate in purple (single trial in light purple, trial-averaged in bold), along with the corresponding trial-average offline estimate (25 ms s.d. Gaussian-smoothed PSTH); (middle) across-trial spike count variability (Fano factor in 500 ms sliding window) and corresponding example spike rasters from 10 randomly selected trials; (bottom) controller input. (c) Population controller performance. In spontaneous vs. closed-loop (CL) control conditions, mean squared error (left) and squared bias (middle) were calculated between the reference (20 spikes/s) and single-trial feedback spiking data smoothed with a 25 ms s.d. Gaussian window; average Fano factor was also calculated (right). For each trial, four seconds of spontaneous data were compared to four seconds of CL control data. The first second was ignored in order to obtain a measure of steady-state performance. Error bars represent bootstrapped 95% confidence intervals about the mean. Green bands represent 95% confidence band for the metrics calculated from simulated Poisson firing at the target rate.

Across experiments ($n = 11$ neurons, 11 experiments), the control framework performed well as quantified by the summary of performance metrics in Figure 3.4(c). For each of these metrics, the measure during closed-loop (CL) control is compared to that from the spontaneous period (spont) before the control was activated at time zero. The mean squared-error (MSE) between an offline estimate of the single-trial firing rate and the target (Figure 3.4(c) left) decreased significantly with activation of the control law as expected ($p = 0.00195$, Wilcoxon signed-rank test), and the MSE during closed-loop control was even below that of a Poisson spike generator driven at the target rate (green bar), consistent with the sub-Poisson variability as revealed by the Fano-factor in Figure 3.4(b). Because the MSE captures a combination of the variance and the bias, we separately computed the bias in the control (Figure 3.4(c) middle), substantially reduced with the activation of the control ($p = 0.000977$, Wilcoxon signed-rank test) and at the level expected for a Poisson spike generator driven at the target rate (green band). To further quantify the reduction in across-trial variability during the control, we computed the average Fano-factor in a 500 ms sliding window, exhibiting substantial reduction from supra-Poisson variability ($FF > 1$) in the spontaneous activity to sub-Poisson variability ($FF < 1$) during the control (Figure 3.4(c) right).

3.3.4 Multi-electrode recordings reveal effects of SISO control on simultaneously recorded neurons

Up to this point, the state-space control framework has been shown effective for tracking step commands in single-neuron scenarios. However, neural recording methodologies (electrophysiology and imaging) continue to scale in size (*e.g.*, larger numbers of channels for electrophysiology, or pixels for imaging) and one of the main benefits of using state-space models for control and estimation is the generalizability to such multi-output problems. While the preceding experimental demonstration was presented in the context of a single channel of light input and a single channel of neuronal output, in a subset of

experiments, we simultaneously recorded multiple nearby neurons in the thalamus of the awake, head-fixed mouse. This provides a window into the effect of the stimulation on the local population while a single neuron is used as an “antenna” around which the controller is operating, which we will refer to as the feedback (FB) neuron (Figure 3.5(a)). For the purposes of this analysis, we inspected simultaneously-recorded neurons that were excited by 5 ms square pulses of light with sub-10 ms latency. Figure 3.5(b) provides an example in which one neuron is being used for feedback (purple, top), while offline spike sorting reveals the activity of six other simultaneously recorded neurons, which we will refer to as non-FB neurons (black, trial-averaged firing rates; green shows control target). While nearby on this 25 μm spaced electrode array (Figure 3.5(b), right), these neurons nevertheless responded heterogeneously to the optical stimulation. In this particular example, the FB neuron was substantially more sensitive to light compared to the non-FB neurons, as evidenced by their modest response following the controller activation at time zero. While all increase their firing rates in response to the controller input, none are driven to or above the target firing rate of 20 spikes/sec in this example. This was not always the case, as in other experiments the non-FB neurons could be either more or less sensitive to the light input as compared to the FB neuron. This heterogeneity in sensitivity is likely due to a combination of factors ranging from differences in opsin expression levels to differences in intrinsic excitability of cells. Across experiments ($n = 8$ feedback neurons, 23 non-feedback neurons, 8 experiments), we calculated the average per-trial firing rate during the pre-control spontaneous (spont) versus control periods for the FB neuron and the non-FB neurons recorded simultaneously. As expected, for the FB neurons, the controller reliably pushed the firing rate to the 20 spikes/s target. In contrast, while the average firing rate of non-FB neurons was significantly elevated from spontaneous levels and toward the 20 spikes/s target ($p = 0.00781$, Wilcoxon signed-rank test), it did so with high variability as evidenced by very wide confidence intervals about the across-experiment average (12.7 to 25.1 spikes/s, Figure 3.5(c), left, grey) and made more plain by the fact that FF did not

change from its spontaneous levels in the non-FB neuron case (Figure 3.5(c), right, grey, $p = 0.844$, Wilcoxon signed rank test).

Beyond the firing rate of individual neurons within the population, it is important to determine what effect the optical stimulation has on the spike timing and synchronization across the population. Although we have previously shown that optical stimulation over some ranges results in a somewhat reduced synchronization relative to comparable electrical stimulation ([72]), it remains an important issue to quantify the effect in the context of the control scheme used here. We find that the use of continuously graded closed-loop stimulation did not significantly synchronize the recorded thalamic neurons when compared to commonly used pulsatile stimulation (Figure 3.5(d-e)). While stimulation elevated population firing over longer timescales, continuously graded closed-loop stimulation did so while allowing natural variability in spike timing across cells, whereas pulsatile inputs tended to coax cells to fire in concert. Spike cross-correlograms were calculated from relative spike times for each of 33 simultaneously-recorded pairs of neurons (see Section 3.2.7). The population correlogram shows no peak at or around zero-lag for the case of closed-loop control (Figure 3.5(d), black). In contrast, 5 ms square pulses of light delivered in open-loop to the same neurons caused clearly aligned spiking (Figure 3.5(d), red). The lack of a peak in the correlogram shows there was very little synchronization of recorded neurons during closed-loop control epochs compared to the results using pulsatile stimulation (Figure 3.5(e), $p = 5.39 \times 10^{-7}$, Wilcoxon signed-rank test), where synchrony was quantified as the number of temporally-aligned spikes in ± 7.5 ms window, relative to the total number of spikes in a ± 50 ms window. Note that these open-loop pulses were in general higher amplitude than the continuously modulated closed-loop stimulation, so it is not necessarily the case that pulsatile inputs would have such synchronizing effects at all stimulation intensities.

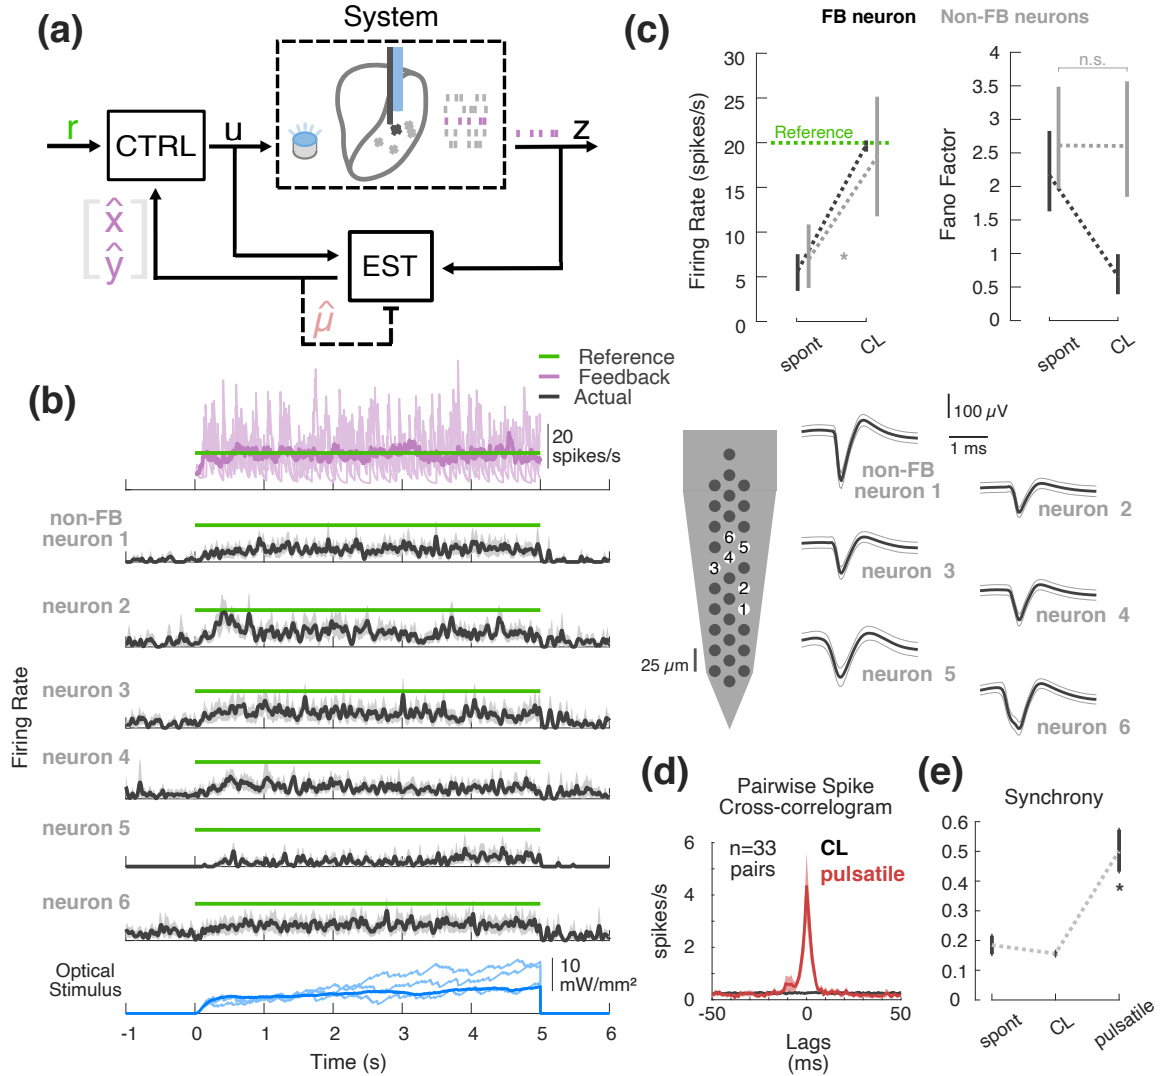


Figure 3.5: Effects of SISO control on local population. (a) SISO control block flow diagram with multi-output recordings. (b) Example experimental SISO control with simultaneous multi-output recordings. (top) Fed-back online firing rate estimate in purple (single trial in light purple, trial-averaged in bold) relative to reference (green); (middle) trial-averaged firing rate estimates for simultaneously recorded non-FB neurons (25 ms s.d. Gaussian-smoothed PSTH); (bottom) controller input. To the right are the waveforms of each neuron in this example (average waveform in black, ± 1 s.d. in grey). (c) Spontaneous vs. CL population average firing rate (left) and Fano factor (right) for the feedback neuron (black) as compared to the other non-feedback-neurons recorded simultaneously (grey). Error bars represent bootstrapped 95% confidence intervals about the mean. (d) Population spike cross-correlogram of simultaneously recorded pairs during optical stimulation. Bold black represents population mean in each 1 ms bin for CL, while fills represent 2 standard errors about the mean. For comparison, red represents population average spike cross-correlogram for response of the same cells to 5 ms square pulses of light presented in open-loop. (e) Population synchrony for spontaneous vs. closed-loop vs. pulsatile conditions. Synchrony was taken as the number of spikes occurring in the ± 7.5 ms bins, relative to the total number of spikes in the ± 50 ms window. Error bars represent 95% confidence intervals about the mean.

3.3.5 GLDS models generalize to multi-output datasets

In the previous section, we considered the effects of closed-loop optogenetic control on nearby neurons when the control was applied in a single FB neuron (*i.e.*, SISO) scenario. More generally, the goal of control may be bringing the firing rate of a neuronal population toward a common target, rather than a single neuron, in order to provide a more controlled and uniform input to downstream neurons. An open question is whether multi-output control would serve this goal better than the above single-neuron “antenna” approach. One of the strengths of the state-space control and estimation framework is that it is amenable to such multi-output applications.

To investigate the multi-output capabilities of this approach, we first demonstrate that these GLDS models can be used to capture the SIMO systems in cases where we recorded multiple neurons simultaneously. In such multi-output GLDS models. We found that the response of multiple neurons to optical noise could be represented by 5th order GLDS models due to the similarity in dynamics and coupling across the channels. In other words, a common state vector is mapped to individual outputs (Figure 3.6(a)). Figure 3.6(b) and (c) provide example results of the GLDS state-space modeling for an example set of four thalamic neurons recorded simultaneously. Figure 3.6(b) shows the impulse response of the GLDS model of the dynamics across these recorded neurons (red), superimposed on the corresponding FIR estimates (black), showing good correspondence as previously exhibited for the single neuron case in Figure 3.2(b). Figure 3.6(c) shows the model predictions of the responses to uniform white noise optical stimulation (red) as compared to average experimentally recorded trial averaged firing (black) for this same set of neurons. The neurons clearly responded heterogeneously to light in terms of overall gain, and the GLDS model captures this and the temporal characteristics of the response to optical noise well. On average ($n = 11$ experiments, 42 neurons), 5th order GLDS models predict population PSTHs approximately as well as in the previously shown SISO case (pSVE 60%, Figure 3.6(d)). As before (Figure 3.2), multi-output PLDS models were also fit to the same data and we

found no significant difference between the performance of the Poisson vs. Gaussian LDS models in explaining the PSTHs under these conditions ($p = 0.923$, Wilcoxon signed-rank test). As is clear in the example responses in Figure 3.6(c), across recordings we found there was often large (sometimes tenfold) heterogeneity in overall sensitivity to light as measured by the static input-output gain, even though the dynamics could be qualitatively similar. To explicitly characterize this heterogeneity, Figure 3.6(e) represents the static gain for each recorded neuron, calculated from the steady state input-output gain of the GLDS fits (circle represents mean, bars represent range). It should be noted that among the inclusion criteria for this study was that neurons must be significantly modulated by light (Section 3.2.3); however, this does not mean that all neurons expressed ChR2 and were being directly stimulated. Instead, they could be indirectly excited or even inhibited (*i.e.*, have negative gains) by optical stimulation of ChR2 expressed in other cells in the network, which could explain the sometimes large differences in neurons' sensitivity to light.

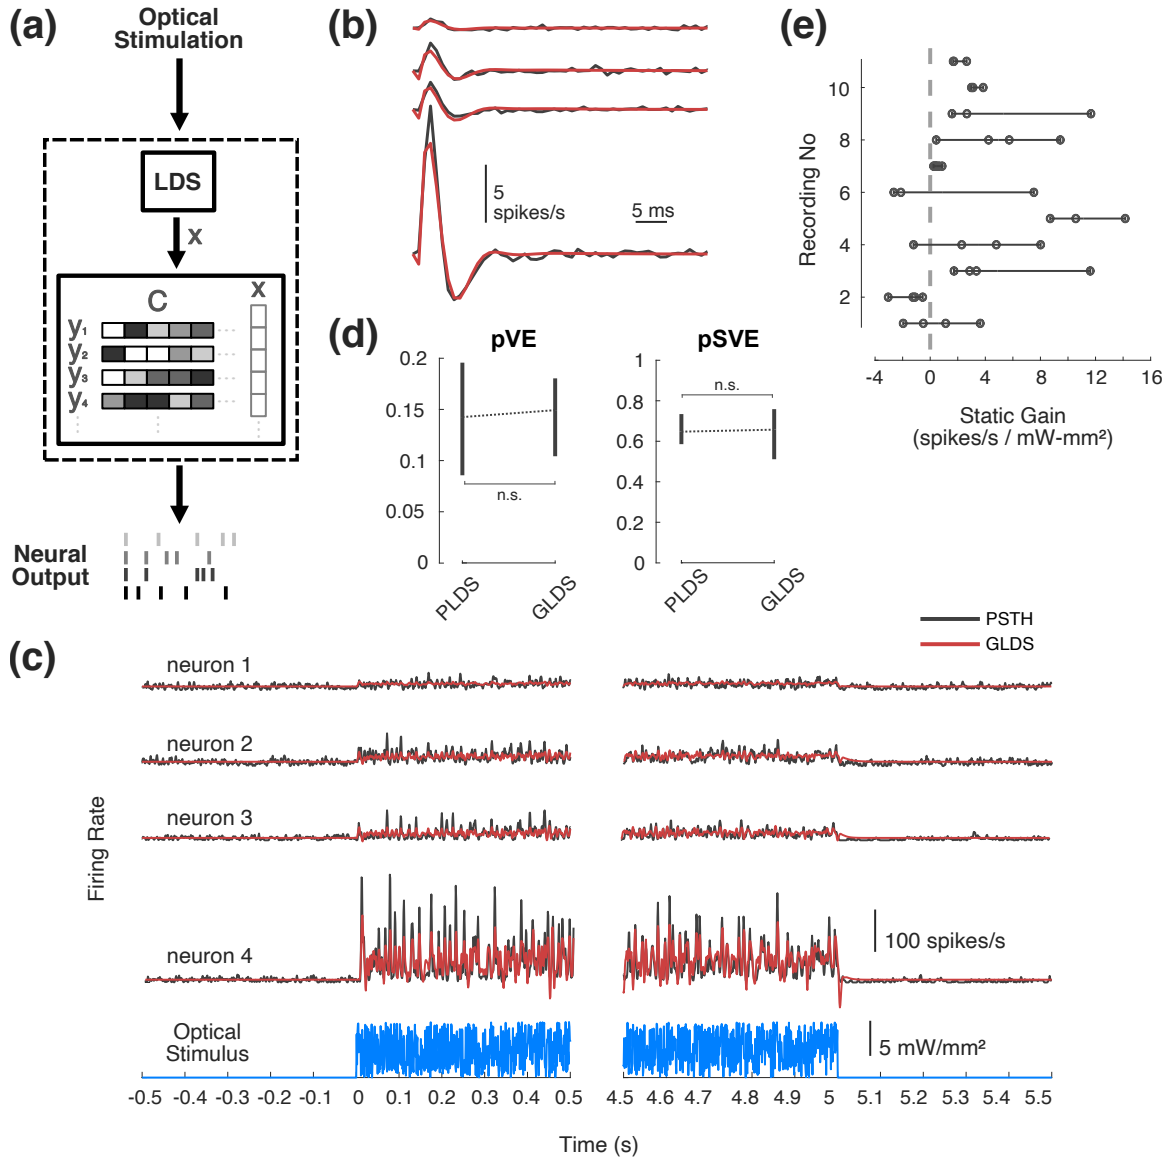


Figure 3.6: **State-space models of SIMO optogenetic responses.** (a) Multi-output GLDS model diagram. A common state is mapped to multiple outputs. (b) Example impulse responses from a single multi-output GLDS model (red) vs. multiple single-output FIR models (black). (c) (top) Example multi-output GLDS model response to optical noise (red) vs. PSTH (black). (bottom) Optical noise stimulus used to fit the model. (d) Population proportion variance in PSTH explained (pVE) and signal variance explained (pSVE) by model response to noise. All models were trained on data from first half of each trial, while model performance metrics (pVE, pSVE) were calculated from the second half of each trial. Error bars represent bootstrapped 95% confidence intervals about the population mean. (e) Range of static input-output gain across and within recordings.

3.3.6 Online estimation methods generalize to multi-output applications

The state-space models were used for both offline controller design as well as for the on-line estimation of feedback provided to the controller. Figure 3.7 details the performance of the parameter-adaptive Kalman filter estimator for the single-input, multi-output (SIMO) case. While a 1st-order GLDS Kalman filter performed sufficiently well as an estimator for the SISO application (Figure 3.3), a 1st-order approximation leads to substantial bias in the estimates in multi-output scenarios, as shown in the example 3-neuron recording of Figure 3.7(a) (trial averaged estimate in dark purple vs. PSTH in black). In this example, the adaptive Kalman filter overestimates the firing activity of neurons 1 and 3, while underestimating the firing of neuron 2. Moreover, the filter fails to capture the gradual decline in firing of neuron 3 over the course of the step response. This is to be expected, as there is only one state disturbance being estimated in the 1st-order case for multiple outputs that may be independently perturbed. For the same multi-output example, a higher order adaptive Kalman filter (5th-order, Figure 3.7(b)) achieves substantially lower estimation bias, albeit not unbiased like the SISO scenario (Figure 3.3). In this example, the average activity of neuron 1 is accurately estimated and the gradual decline in neuron 3 firing is captured; however, estimates for neuron 2 and neuron 3 activity remain biased. Across the population of recordings ($n = 11$ experiments, 42 neurons), the 1storder adaptive Kalman filter provides lower estimation bias than a standard Kalman filter using a 5thorder GLDS fit (Figure 3.7(c), $p = 0.00982$, Wilcoxon signed-rank test) as before in single-output application; however, an adaptive 5thorder filter provides an improvement in the multi-output scenario across recordings ($p = 1.65 \times 10^{-8}$, Wilcoxon signed-rank test). It is unsurprising that the higher order adaptive filter improved performance because there are more state disturbances being estimated. However, as is evident in Figure 3.7(b), it is important to note that while higher order models perform better, this form of parameter-adaptive Kalman filtering does not independently minimize the estimation error for each output neuron because the estimated process disturbances act on a set of common state variables.

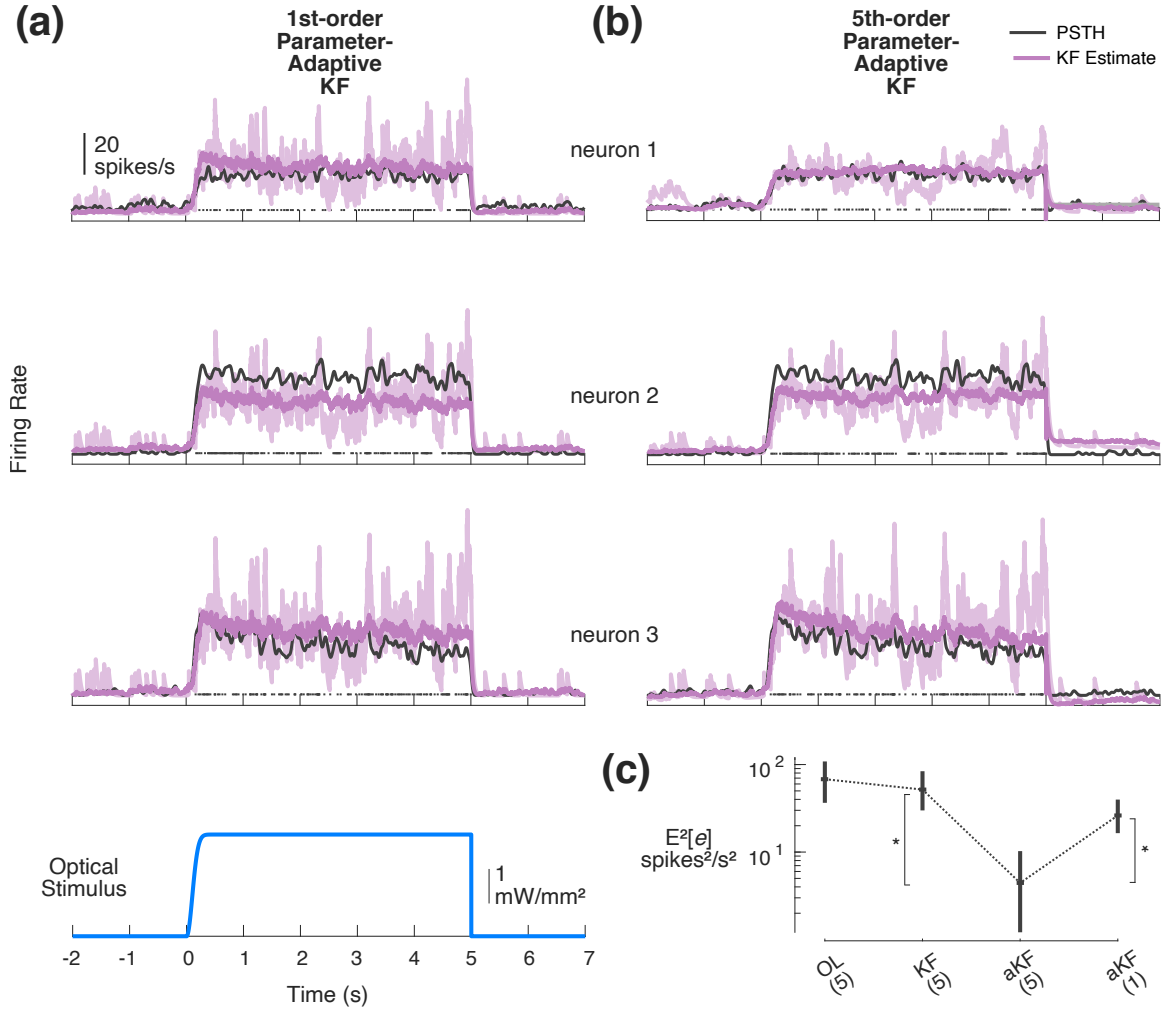


Figure 3.7: **Kalman filtering for online estimation in SIMO applications.** (a) Online estimation for example SIMO system using parameter-adaptive Kalman filter with a 1st-order GLDS. Shown for each of three simultaneously recorded single neurons: PSTH (black), trial-averaged estimate (bold purple), and example single trial (light purple). (b) Online estimation of firing rate using parameter-adaptive 5th-order GLDS (same data as in (a)). (c) Population summary squared-bias: Open-loop (OL) prediction of 5th order GLDS, standard Kalman filter (KF), and parameter-adaptive KF (akF) for 1st- and 5th-order models. Fills/error bars represent 95% confidence intervals about the mean.

3.3.7 Simulated SIMO control more robust to population heterogeneity than a SISO “antenna” approach

To demonstrate the generalizability of this approach to multi-output control problems and to explore the robustness of SIMO control to population heterogeneity in light sensitivity (Figures 3.5(b), 3.6), we simulated a 2-output system whose second output (neuron 2) ranged from much less sensitive than neuron 1 to much more sensitive (Figure 3.8(a)). To simulate spiking, an example PLDS model previously fit to SISO experimental data was chosen to most accurately represent the complexities in the data (*e.g.*, spiking). The output matrix of this model was augmented with a second row whose elements were gain-modulated versions of the first row (Figure 3.8(a), bottom inset). This log-linear gain term was swept from 0.1 to 3 times that of neuron 1. After fitting GLDS models to simulated SIMO datasets (below), this resulted in linear gain of neuron 2 ranging from 0.027 to 13.3 times that of neuron 1. In these simulations, the dynamics of the PLDS model neurons were held fixed. This resulted in a set of simulated datasets representing a range of similarity between the two output neurons.

Using previously described methods, a multi-output GLDS model was fit to simulated spiking responses to optical white noise in the case where neurons 1 and 2 had the same log-linear gain. A single-output model was fit to the responses of neuron 1. An estimator and controller were designed using these SIMO and SISO models and they were applied to control of the 2-output PLDS across a range of output gain disparities. In the SISO case, only the activity of neuron 1 was used for feedback control, while in the SIMO, spiking activity from both neurons was fed back to the estimator and controller. Note that in these simulations, the control system only had to overcome static differences in light sensitivity and not the apparent decrease in neuron sensitivity over time observed in the previous examples of experimental control. Examples where the log-linear gain of neuron 2 is 1.5-times that of neuron 1 are provided in Figure 3.8(b) and (c), showing the SISO and SIMO control results, respectively. Qualitatively, while SISO control of neuron 1 successfully

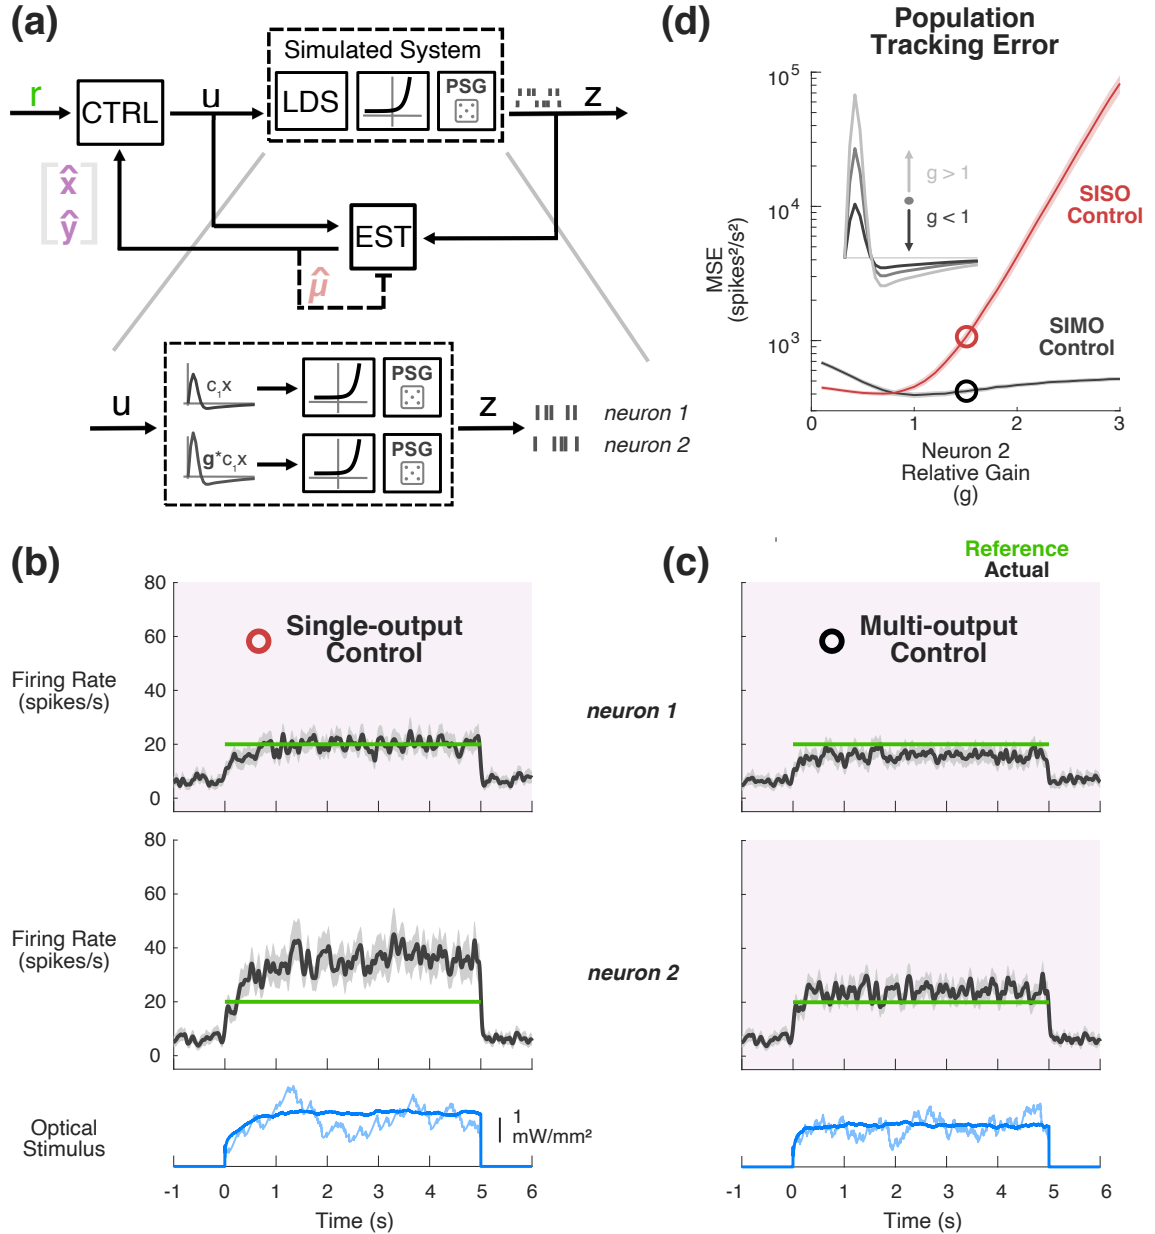


Figure 3.8: **Simulated SIMO control and estimation.** (a) Simulated control of a SIMO PLDS model system. A 2-output PLDS was simulated, whose second channel of output was a gain-modulated version of the first channel. (b) Example SISO Control. Only data from neuron 1 (top) was fed back during the simulated control. PLDS neuron 2 was 1.5 times more sensitive than neuron 1 before exponentiation in this example. Actual firing rate of each neuron (PSTH smoothed with 25 ms s.d. Gaussian) is shown in black as compared to reference (green). (c) Example of multi-output control. Both neurons' data are fed back to controller for online estimation and control. Signals are the same as in (b). (d) Simulated mean-squared tracking error for SISO control of neuron 1 (red) vs. SIMO control (black) as a function of the relative log-linear gain of neuron 2. Circles denote mean-squared error for examples in (b), (c).

clamps activity of that channel at the target 20 spikes/sec, neuron 2 is well above the target (Figure 3.8b). Conversely, the multi-output control case strikes a balance between the two, allowing neuron 1 to fall below the target and reigning in the above-target activity of neuron 2 (Figure 3.8(c)). As summarized in Figure 3.8(d), population tracking performance was quantified as the total MSE of neuron 1 and neuron 2 single-trial firing rates vs. a 20 spikes/s reference firing rate, both for single-output (red) and multi-output (black) control strategies. Highlighted with the open symbols are the performances of the examples given in Figure 3.8(b) and (c), where the value of the log-linear gain of neuron 2 was 1.5 times that of neuron 1. As expected, a multi-output control strategy is more robust to population heterogeneity, as the SISO control performance rapidly degrades when the ignored neuron 2 is increasingly sensitive. Note that this effect is not symmetric, as discrepancies in sensitivity are substantially less problematic when neurons in the population are less sensitive to the light input as compared to the feedback neuron (relative gain < 1). Taken together with previous multi-output modeling and estimation on experimental data, these simulations demonstrate that the techniques used first for SISO applications are readily applicable to multi-output problems and that such approaches could grant better control of population neural activity of interest.

3.4 Discussion

With the continued development of tools for precisely and selectively manipulating neuronal ensembles using multiple inputs ([184, 69]) and corresponding technologies for measuring large-scale neuronal activity ([63]), a framework for the integration of these technologies enables more intelligent interaction with neuronal circuits within and across brain regions (for review, see [5]). The state-space model structure is a natural choice for describing systems that involve a number of inputs and outputs (referred to as multi-input, multi-output or MIMO) ([185]). State-space models for system dynamics in combination with the framework of optimal control and estimation allows design and implementation of control loops to scale without cumbersome changes in methodology, as is evident in this study where the same model structure has been applied successfully in online estimation and control of activity in single- as well as multi-neuron systems.

In more modern control approaches, a model of the underlying dynamics to be controlled is used for both design and implementation of the control law. Across different pathways and circuits, numerous model types have been used to predict neural activity at fast time scales ([82, 90, 153]), some of which have been applied to ensembles of neurons ([101, 91, 99, 186, 187]). For the purpose of control applications, we modeled optically-driven responses of neurons in the sensory thalamus using a state-space dynamical systems representation, where any higher order dynamics are captured in sets of coupled, first-order difference equations. In contrast to more widely used phenomenological models for neuronal responses to stimuli such as the linear receptive field ([79, 80]), the linear-nonlinear-Poisson (LNP) model ([81, 87]), and the generalized linear model (GLM) ([90, 91]), state-space models describe what could be large-scale recordings as arising from some potentially small number of latent “states” that evolve dynamically in time as a function of themselves and covariates such as sensory or optical stimuli. We found that linear state-space (*i.e.*, GLDS) models could be used in the context of the control objective of maintain-

ing a steady firing rate in the face of ongoing activity changes during wakefulness in these early sensory neurons. This was not a given, as many applications of state-space models to neuronal spiking data have used nonlinear dynamical systems, or at least linear dynamical systems with Poisson observations ([135, 95, 137, 99, 101, 97, 188, 189, 190]). While the GLDS models used here clearly do not respect the statistics of the measured spiking data and while they can grossly mis-predict neuronal responses to optical stimulation (*e.g.*, Figure 3.3(a)), they do capture the basic dynamics and the robustness of the parameter-adaptive Kalman filter in combination with feedback control allowed the use of this relatively simple modeling framework and unlocks a wealth of other linear design and analysis methods developed over years of study. Similarly, a recent study that used state-space feedback control in the context of hardware-simulated manipulation of electrocorticography (ECoG) made a practical concession to use a linear model that was more amenable to commonly-used feedback control techniques ([191]). In contrast, we previously used a linear-nonlinear-Poisson (LNP) spiking model in to design a classical proportional-integral (PI) controller ([126]). While that simple control strategy proved quite effective even for tracking patterns of rate modulation, the controller was designed numerically around a simulated spiking system, owing to the multiple nonlinearities that precluded the use of such design tools as LQR used in the present study. Simply put, there is a natural trade-off between the complexity of models and complexity of the control design and implementation itself, especially as the dimensionality of problems scale. That said, a nonlinear model will likely be needed in some control applications and it is possible the use of a PLDS model would have improved control performance even in this application but at the expense of complexity. In cases where a Poisson model is necessary and/or beneficial, there are previously-developed methods for estimating the underlying state of a PLDS (“point process filter”, [98, 188]). One could leverage these nonlinear filtering techniques and design/implement feedback control in the log-linear state space as described here for linear systems.

Aside from the observed robustness to the linear approximation of nonlinear neuronal

activity, we also found in the context of SISO control problems that the parameter-adaptive Kalman filtering and feedback control granted enough robustness to model-mismatch that the GLDS did not have to be fit during the tight time constraints of awake, head-fixed recording sessions. Instead, data from previous experiments were used for the controller design and online estimation. Certainly, a large part of this success comes down to the fact that the control objective was relatively long-timescale firing rate regulation. In the context of trajectory tracking, we previously showed that, while closed-loop control grants some robustness to model mismatch, even things as simple as DC gain mis-estimation can lead to off-target activity when the target trajectory is time-varying ([126]). Indeed, in the present study we have observed a wide array of neuron sensitivity to light. Therefore, there are certainly scenarios in which the control objective would warrant better model fits, or at least adaptively re-estimating other parameters such as the input matrix (B) or output matrix (C) rather than attempting to capture all model mismatch with a linear disturbance as was implemented here. In general, this would call for nonlinear variants of the Kalman filter, such as the extended Kalman filter ([180, 98]).

In addition to the fact that the control method proved quite robust to model-mismatch in the sense of the statistics of measured data and the aforementioned long timescale biases in model predictions, we also found that we were able to effectively carry out the SISO control and estimation problems using a first order approximation of systems that appeared to be fourth or fifth order (Figure 3.2(b)). Since the control objective was to track a firing rate step command over relatively long timescales, this should be expected. After all, the dynamics of these systems tended to have died out after tens of milliseconds (Figure 3.2(b)). In applications where the objective is to entrain precisely-timed sequences of spiking activity rather than an overall firing rate (*e.g.*, [136, 137, 95, 92]), a higher-order model would be merited and more emphasis would need to be placed on stimulus design.

To this point, all references to the robustness of this control framework have pertained to activity of the putative single neuron which was used to adjust stimulation in real-time.

Across recordings, the feedback neuron's activity was maintained at the target firing rate with low error on average and, importantly, with low trial-to-trial variability. However, we found that the local population of neurons also excited by the optical stimulation did not exhibit this same lowered variability. It is worth noting that it is likely the case that pulsatile stimulation rather than the continuously-graded stimulation we used here would have had a less variable effect on the population. However, we showed that presenting 5 ms pulses of light synchronizes the population and would thus likely strongly impact downstream targets in a way that is unsuitable for many applications. Therefore, to reap the benefits of closed-loop control in neural circuits, feedback of population activity rather than a putative single neuron will likely be of great importance moving forward. Given the difficulty of online identification of individual neurons from raw electrophysiological recordings (i.e. spike-sorting) for dense multielectrode arrays such as the ones used here, the thresholded multi-unit activity often utilized in brain-machine interface applications may prove an effective alternative measure of population activity ([64]). Alternatively, it is conceivable in the case of chronic implants to sort and track single units across experiments ([192, 193]).

Aside from providing multi-output feedback to the controller, the addition of multiple light sources (*e.g.*, [184]) would afford some degree of population control spatially. The current preparation is highly underactuated in that there is a single light source being used to manipulate local activity, and there will in practice always be heterogeneities in responsiveness to light in space, whether it be due to varying distance from a common light source or differences in expression of opsins, *etc.* In addition to multiple spatially-distributed light sources, having the ability to simultaneously excite and inhibit neuronal activity using light of different wavelengths will also be key for robustness of optogenetic control moving forward. Note that in the present study, a single excitatory opsin (ChR2) was expressed in excitatory cells, meaning that the control is limited to pushing activity of those neurons toward higher firing rates. This effectively limits the control problem to one that maintains firing at an above-average desired level: here, 20 spikes/s which naturally occurs in this

pathway. Conversely, if inhibitory opsins were expressed (or excitatory opsins were expressed in inhibitory interneurons), the control objective would be limited to maintaining or pulling down spontaneous levels of activity. Therefore, the ability to effectively push as well as pull back on neuronal activity would greatly expand the utility of this approach. Importantly, while not tested here, the state-space control and estimation methods used in the present study should generalize to such multi-input control problems. However, it is likely the differing kinetics of excitatory and inhibitory opsins would necessitate higher order models.

Besides the utility in treatment of neurological disorders and diseases ([171, 142, 141, 140]), or in augmenting normal brain function, the precise, closed-loop control of neural circuits has the potential to significantly enhance our understanding of underlying mechanisms of basic brain function. After all, feedback control enabled the seminal work of Hodgkin and Huxley in uncovering the nature of the ionic currents that underlie the generation of a neuron's action potential, for which they won the Nobel Prize in 1963 ([61]). The key to this experimental work was the use of a feedback controller to “clamp” the trans-membrane voltage by injecting current to counter-act naturally occurring changes in ionic currents. This functional decoupling of constituent ionic and capacitive currents led to a quantitative description of the nonlinear dynamics of the action potential. Single-cell voltage clamp and dynamic clamp experiments ([194]) continue to be a powerful tool for scientific discovery, but continuously-graded feedback control of this sort has not been translated to the circuit-level, where the dynamics are complex and can adaptively change from moment to moment. Fundamentally different from lesioning or reversibly silencing brain regions, closed-loop optogenetic control has the opportunity to aid investigation of the mechanisms governing circuit level dynamics in a similar way voltage clamp did for the single neuron.

3.5 Conclusions

In this study, a state-space control and estimation framework has been developed and demonstrated to work well in the context of wakefulness, where there is spontaneous fluctuation in neuronal activity. Compared to Bolus *et al* ([126]), this updated approach is more naturally suited to the MIMO control problems that are important in studying complicated neural circuits. Notably, we were able to use a simple, linear approximation to this nonlinear system at least for long-timescale control objectives, such as maintaining an overall firing rate. The relative simplicity of these approaches achieved at the expense of modeling fidelity represents one of the chief strengths of the methodology, as linear control is well understood and is widely used. Moreover, while tested in the context of controlling spiking activity which is often statistically modeled as a point-process, the methods laid out here are immediately applicable to the control of continuous-valued neuronal signals of interest such as local field potential and voltage/calcium imaging. This demonstration of state-space models being used for single- and multi-output applications of optogenetic control opens the door to other established control strategies that use this modeling framework, such as model predictive control ([185]). As a whole, this work lays the foundation for future advances in manipulation and study of neuronal circuits using the integration of neuronal recordings and optogenetic stimulation.

CHAPTER 4

OPTOGENETIC CONTROL OF THALAMIC STATE AND ITS EFFECTS ON THALAMOCORTICAL ACTIVITY

4.1 Introduction

From sleep at one extreme to alert wakefulness at the other, the brain moves through a gamut of states that are accompanied by measurable changes in behavioral and neuronal activity. Thalamic and cortical activity is state-dependent, ranging from “Up” and “Down” cortical activity and thalamocortical spindle oscillations during slow wave sleep and under certain anesthetics to rapid, small variations in cortical subthreshold membrane potential and elevated thalamic firing rates during periods of active arousal (for reviews, [29, 30, 31, 32]). Previous work defines the state of the brain during wakefulness in discrete extremes such as quiet vs. active/aroused ([33, 34, 35, 27]), inattentive vs. alert [36], or passive vs. engaged in a task ([37]). Despite this discretization, wakefulness is likely a continuum ([29]), albeit with some well-studied extrema. On one end of the spectrum, when animals are quiet/inactive the neuronal activity of cortex as measured by membrane potential or extracellular modalities like local field potential (LFP) or electroencephalography is characterized by large amplitude, low frequency fluctuations coordinated across cells, whereas the thalamic activity is characterized by low firing rates and some degree of burst firing. At the other end of this spectrum, in an active state, the cortex is characterized by lower amplitude, high-frequency fluctuations in subthreshold potential, while the thalamus exhibits elevated tonic firing. Such state changes can also be behaviorally measured by factors such as pupil diameter ([36, 38, 39, 40]) and movement ([41, 33]).

While the nature of the effects are apparently different in the visual versus auditory and somatosensory pathways, these state changes in thalamus and cortex have been shown

to have consequences for sensation both in terms of neuronal responses as well as in the perception of sensory stimuli across modalities ([42, 41, 43, 44, 37, 38, 45, 46, 47, 48, 49]). These states are likely governed by changing tone of neuromodulators such as acetylcholine and norepinephrine in thalamus and cortex ([20, 51, 52, 38, 39]); however, it has been shown previously that direct manipulations of thalamic firing have been sufficient to control cortical state measures ([51, 35]). Therefore, it is as yet unclear what the relative contributions of direct neuromodulation of cortex versus thalamic state itself are to variability in cortical and, ultimately, perceptual responses to sensory stimuli. An experimental approach that could hold the state of one brain region steady in the face of changing neuromodulatory and behavioral activity would be an invaluable tool for answering such questions.

Rejecting such disturbances is one of the chief benefits of using feedback control systems ([28]). Moreover, optogenetic stimulation in combination with electrophysiology allows unprecedented ability to stimulate and record at fast timescales with less concern for stimulation artifacts seen with electrical stimulation, as well as the ability to do so with cell-type specificity (*e.g.*, [122, 113, 62, 124]). Closed-loop optogenetic control is a general approach to stimulation that utilizes continuous feedback of neuronal activity to modulate light intensity in order to achieve a target level of activity. To date, this form of optogenetic control has been demonstrated to work well in vitro ([6, 138]) as well as in the anesthetized animal ([6, 126]). Here, closed-loop optogenetic control has been applied to the problem of regulating thalamic activity in the vibrissa, or whisker, system of the rodent primary somatosensory pathway about an elevated level that is known to occur during an active state. The control was implemented as previously described (Chapter 3) using the spiking feedback of a putative single neuron. This study investigates the effect of optogenetic control on the state of the thalamus in the awake animal with a special emphasis placed on understanding the approach's ability to reject disturbances that would otherwise change this activity, such as the state changes that occur when animals move their

facial whiskers (*i.e.*, “whisk” [33, 27, 160]). We also explore the downstream effects on ongoing cortical activity, and consequences for sensory response properties in each region.

Given the importance in changing thalamic firing properties for downstream cortical activity and the lack of agreed-upon thresholds for discrete state delineations, a state-space dynamical systems notion of thalamic state was used to provide a continuous-valued, latent variable underlying coordinated changes in population thalamic firing, as measured by multi-unit spiking activity recorded extracellularly. While in general such state variables could be multidimensional, for the purposes of this work, the state was assumed to be unidimensional to capture the first-order characteristics of coordinated changes in firing activity. This variable was used to quantify the effects of closed-loop stimulation on thalamic activity globally, juxtaposed with activity seen under spontaneous and open-loop optical stimulation conditions. While applied specifically in the context of thalamic firing, such a mathematical approach to state definition and estimation could be fruitful for other systems as well.

4.2 Methods

4.2.1 Animal preparation

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 NIH. Experiments were carried out using either C57BL/6J mice that were virally transduced to express channelrhodopsin-2 (ChR2) or by single-generation crosses of an Ai32 mouse (Jax) with an NR133 cre-recombinase driver line (Jax) which grants better specificity of ChR2 expression in ventral posteromedial/posterolateral thalamus ([123]). In the case of viral transfection, ChR2 expression was targeted to excitatory neurons in the thalamus via stereotactic injection relative to bregma (approximately $2 \times 2 \times 3.25$ mm caudal \times lateral \times depth) using 0.5 μ L of virus (rAAV5/CamKIIa-hChR2(H134R)-mCherry-WPRE-pA; UNC Vector Core, Chapel Hill, NC) at a rate of 1 nL/s.

At least three weeks prior to recordings, a custom-made metal plate was affixed to the skull for head fixation and a recording chamber was made using dental cement while the animals were maintained under 1–2% isoflurane anesthesia ([62, 195, 172]). After allowing a week for recovery, mice were gradually habituated to head fixation over the course of at least five days before proceeding to electrophysiological recordings and optical stimulation. In cases where recordings were also carried out in primary somatosensory cortex (S1), non-invasive intrinsic optical imaging ([196, 197, 198, 199]) was used to localize the approximate location of the cortical column in S1 (or “barrel”) that corresponded to individually-stimulated whiskers on the animal’s face. Animals were lightly anesthetized for 1-3 hours using 0.8-1.5% isoflurane, their skulls carefully thinned over S1 to optimize optical translucency, and single whiskers repetitively stimulated for 4 s at 10 Hz and 500 degrees/s using a computer-controlled galvanometer (Cambridge Technology, Inc.). The cortex was illuminated with red light (625 nm, M625F2 Thorlabs), while the reflected light was recorded at 10 fps using a tandem “macroscope” ([196]) comprised of 50 and 105 mm focal length lenses (Nikon Nikkor) connected front-to-front and mounted on a camera (Basler AG acA1920). The location of the column in cortex corresponding to the stimulated whisker was functionally identified as the point in the frame that had the largest dip in reflectance ([199]).

On the day of the first recording attempt, animals were again anesthetized under 1-2% isoflurane and a small craniotomy (1-2 mm in diameter) was centered at approximately 2×2 mm caudal and lateral of bregma. In cases where recordings were also carried out in S1 cortex, a second craniotomy was made (approximately 0.5-1 mm diameter) centered at the point functionally identified using intrinsic optical imaging (see above). Animals were allowed to recover for a minimum of three hours before awake recording. At the time of recording, animals were headfixed and either a single electrode or an electrode array coupled to an optic fiber (Section 4.2.2) was advanced through this craniotomy to a depth between 3-4 mm for thalamic recording and stimulation. Between repeated record-

ing attempts, this craniotomy was covered using a biocompatible silicone sealant (Kwik-Cast, WPI). Following termination of recordings, animals were deeply anesthetized (4–5% isoflurane) and sacrificed using a euthanasia cocktail.

4.2.2 Experimental setup

All optical stimuli were presented deep in the brain via a 200 or 100 μm diameter optic fiber attached to a single tungsten electrode (FHC) or to a 32-channel Neuronexus optoelectric probe in a 25 μm -spaced “poly3” configuration (A1x32-Poly3-5mm-25s-177-OA32LP, NeuroNexus Technologies, Inc.), respectively. Additionally, in experiments where cortical recordings were also carried out, a linearly-space 32-channel Neuronexus electrode array (A1x32-5mm-25-177-A32) was inserted 1 mm into the cortical craniotomy (Section 4.2.1). Note that even recordings where thalamic and cortical activity was not topographically aligned (*i.e.*, not responsive to the same facial whisker) were included in this study. A computer-controlled galvanometer was used to administer angular single-whisker stimulation ([200, 16, 67]) (Cambridge Technology, Inc). Command voltages were generated by a data acquisition device (National Instruments Corporation) in a computer running a custom-written RealTime eXperimental Interface (RTXI, [167]) program at 1-ms resolution. Command voltages were sent to a Thorlabs LED driver, which drove a Thorlabs 470F3 LED (470 nm wavelength blue light) connected to the 100-200 μm optical fiber. A Tucker Davis Technologies RZ2 system measured extracellular electrophysiology. This system was used for single-channel PCA spike sorting, binning, and sending these binned spike counts at 2-ms resolution to the computer running RTXI over ethernet via UDP. The computer running RTXI for realtime control listened for datagrams over ethernet and linearly interpolated from 2-ms to the operating resolution of 1-ms. All told, the closed-loop processing loop was approximately 10 ms.

In addition to optical stimulation and electrophysiology, the animal’s facial whiskers were recorded from below using high-speed videography (Basler AG acA1920). The com-

puter running RTXI triggered frame acquisition at 100 fps with digital strobes. These frames and frame timestamps were buffered, compiled into video files, and saved on the same computer using a custom-written C/C++ program.

4.2.3 Quantification of whisker motion

The behavioral videography was used to assess when the animals were moving their whiskers (*i.e.*, whisking) and with what intensity. Offline, movement of one hemisphere of the animal's facial whiskers was quantified as the mean squared luminance at a given point in time in a region of interest (ROI) manually centered on the whiskers. In order to provide an estimate of gross movements of the whiskers, this signal was smoothed with a 100-ms rectangular window, was subtracted by 1 standard deviation of the signal in time, and thresholded such that values below zero (*i.e.*, below 1 s.d.) were set to zero. When this measure of whisker motion was above zero, the animals were said to be whisking; all other times were taken as non-whisking.

4.2.4 Thresholded multi-unit activity

In contrast to previous chapters, thresholded multi-unit data comprises most of the neuronal spiking activity considered in this study, as opposed to sorting this spikes into clusters belonging to putative single neurons (Section 4.2.5). Each channel of extracellular voltage recorded from single tungsten or an array of silicon electrodes was high-pass filtered at 250 Hz ([64]) using a fifth-order Butterworth filter. The negative threshold past which voltage crossings were taken as spikes of one or multiple neurons near the electrode (*i.e.*, multi-unit activity, hereon referred to as MUA) was chosen to be four times the standard deviation of the noise on a given channel, or $-4 \times \sigma_n$. Following [201, 202], the standard deviation of the noise on each channel was estimated as

$$\hat{\sigma}_n = \frac{\text{median}(|\mathbf{v}|)}{0.6745} ,$$

where v is the channel voltage time series. The number of standard deviations of the noise used for the negative threshold is typically between 3 and 5 ([202]), and [64] tried three different thresholds of $-3.5\times$, $-4\times$, and $-4.5\times$ for MUA, and showed little difference between them in the context of that study. Therefore, the middle $-4\times$ was used for the purposes of this study.

4.2.5 Offline spike sorting

While most analysis of spiking activity was confined to MUA, offline-sorted single unit activity was used in one analysis in this study. Kilosort2 ([174]) was used for all offline spike sorting, including single-channel recordings, in which case spatial whitening and common mode referencing steps were disabled. Initial sorting by Kilosort2 was then manually curated (additional merging/splitting of clusters) using the “phy” viewer. After manual curation, any clusters that met the following criteria were considered single units and used in this study: sub-1-ms ISI violations of $<0.5\%$, sub-2-ms ISI violations of $<2\%$, and mean waveform amplitude-to-standard-deviation ratio >4 . Additionally, only those units that exhibited significant above-baseline excitation within 10 ms from onset of 5 ms optical square pulses were considered in this study ($n=52$ units, 16 recordings, 10 animals).

4.2.6 Cortical local field potential

Cortical local field potential was used to quantify cortical state and sensory response characteristics. Raw extracellular voltages were recorded in S1 of cortex at the operating sampling rate of 24.414 kHz (Section 4.2.2) and downsampled by $48\times$ to a sampling rate of approximately 500 Hz (508.6 Hz) and bandpass-filtered between 0.5 and 500 Hz using a fifth-order Butterworth filter. All other signals (optical and sensory stimuli and whisker motion) were resampled to this same time resolution, and spikes (single- or multi-unit) were likewise binned at this resolution.

4.2.7 Thalamic state estimation

For the purposes of this study, a one-dimensional representation of changing population thalamic firing was taken as the state of the thalamus. In general, higher-dimensional state estimates may be used; however, the goal of this analysis was a single variable that best explained coordinated changes in across-channel firing activity. To this end, a first-order Gaussian linear dynamical system model (GLDS, see Section 3.2.4) was fit to multi-channel MUA spiking data, $(\mathbf{y} \in \mathbb{R}^p)$. The number of output channels p was at most 32; however, on the electrode arrays used, some channels were dysfunctional and so were omitted from analysis. As recorded covariates of thalamic firing, the two inputs to this multi-input/multi-output (MIMO) GLDS model, $(\mathbf{u} \in \mathbb{R}^2)$, were the optical control signal as well as the whisker motion signal (Section 4.2.2). The 1-dimensional thalamic state variable, $x \in \mathbb{R}^1$, was assumed to evolve as

$$x_t = ax_{t-1} + \mathbf{b}\mathbf{u}_{t-1} + w_{t-1} ,$$

where $w_t \sim \mathcal{N}(\mu, q)$ is process noise of mean μ and variance q disturbing the state dynamics. As before (Section 3.2.4) the state was linearly mapped to a p -output firing rate and measured in this case as 2-ms-binned multi-unit spiking \mathbf{z} :

$$\mathbf{y}_t = \mathbf{c}x_t + \mathbf{d} ,$$

$$\mathbf{z}_t = \mathbf{y}_t + \mathbf{v}_t ,$$

where \mathbf{d} is an output bias term corresponding to the baseline firing rate of each channel, $\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ is the assumed zero-mean Gaussian measurement noise of covariance \mathbf{R} . GLDS state-space models were fit using training data from periods of optical noise stimulation (Section 3.2.4) as well as spontaneous activity. The output bias term, \mathbf{d} , was taken as the average firing rate across all periods of optical and whisker movement silence in the

training data. Subspace identification ([176]) was used to fit an initial GLDS between \mathbf{u} and $(\mathbf{y} - \mathbf{d})$ to each dataset individually.

In fitting these models, care had to be taken to ensure a consistent meaning and scale for this thalamic state variable, x , across experiments. First of all, because of different lighting conditions from experiment to experiment, the amplitudes of the whisker motion signal recorded by videography were widely varied. To ensure consistency across experiments, the whisker motion signal described previously (Section 4.2.2) was normalized by the standard deviation across time in each training dataset. The same was done for the optical input, resulting in a two-dimensional input, \mathbf{u} , where both channels were of unit variance.

Most importantly, to ensure consistent notion of state across experiments, a single state transition coefficient, a , input vector, \mathbf{b} , and process noise variance, q , were fit to the entire population dataset ($n=18$ recordings, 11 animals). This was achieved by first fitting a GLDS model for each recording individually, and then performing a single step of the Expectation-Maximization (EM) algorithm detailed in [177] to solve for *maximum a priori* estimates of a and \mathbf{b} . Briefly, Kalman smoothing ([203, 177]) was performed with GLDS models fit individually to each dataset in order to obtain estimates for the state mean and covariance across all training data. These state and covariance estimates from individual experiments were pooled together, and the parameters a and \mathbf{b} were fit as detailed in the maximization step of [177]. With this single set of dynamics, the output vector, \mathbf{c} , was then refit for each experiment by iterating the EM algorithm, while holding the dynamics fixed. In this way, a single consistent set of LDS dynamics were used across experiments.

These GLDS models were used as the basis for thalamic state estimation throughout the study. In order to estimate thalamic state from optical and whisker motion covariates and measured MUA, parameter-adaptive Kalman filtering was performed as in Section 3.2.5, where the process disturbance μ is jointly re-estimated with the state of the system. In this application, the disturbance was assumed to vary as a random walk of variance $q_\mu =$

$$1 \times 10^{-4}.$$

4.2.8 Cortical LFP frequency ratio

The “state” of cortex is often said to be correlated to the frequency content of signals such as electroencephalography (EEG) and intracortical local field potential (LFP) ([32]), such that a so-called deactivated cortical state is characterized by higher amplitude, low-frequency activity and an activated state is characterized by lower-amplitude, high-frequency activity. A quantity that captures this spectral redistribution of signal power is the ratio of power in low frequency (1-10 Hz) to high frequency (30-90 Hz) bands ([32]). In order to estimate this ratio as a function of time, the power spectral density was estimated in a 1-second sliding window using the MATLAB function *spectrogram()* (MathWorks). The frequency bins corresponding to each of the aforementioned bands were integrated and the low to high frequency ratio was taken.

4.2.9 Closed- and open-loop optogenetic control

All optogenetic control was performed as detailed in Chapter 3, in a single-input/single-output (SISO) configuration where spiking activity of a putative single neuron was used for feedback. As before, the control objective in this study was holding the output neuronal firing to a fixed target, or reference, rate (r), corresponding to a nonzero-setpoint regulation problem ([181]), also described here as “clamping”. As in Sections 3.2.6 and 3.3.3, first order GLDS models for light-to-spiking were used as the basis for the design and implementation of optogenetic control and estimation in the feedback loop.

A steady-state solution was used for this control problem, and the system state-control setpoint was calculated using models fit to previously collected optical noise driven spiking data. This problem was solved by linearly-constrained least-squares ([182]) (Section 3.2.6). The resulting nominal control input was not only used during feedback control epochs but was also used as an open-loop (OL) control analogue for comparison to closed-loop (CL).

In closed-loop control epochs, the adaptive Kalman filter detailed in Section 3.2.5 was used to provide state feedback to the controller. The state and integrated-output feedback controller gains were designed offline using linear quadratic optimal control (See Section 3.2.6).

4.2.10 Dataset inclusion criteria

All told, the data in this study come from 18 recordings in 11 different animals. However, not every recording was appropriate for every analysis. For all analyses pertaining to the effects of closed- and open-loop optogenetic control on thalamic state, only experiments where the control objective was to maintain feedback neuron firing rate at 20 spikes/s were used (n=16 recordings, 10 animals). To assess the effects of optogenetic control of thalamic state on ongoing cortical activity, only those datasets in which there was significant cortical LFP and MUA spiking responses to optical pulse stimulation of thalamus were included; additionally, there were a number of datasets in which large low-frequency artifacts in the LFP possibly related to animal movement prevented the necessary frequency content analysis. After these were left out, there was a remaining 7 recordings in 4 animals for analysis of ongoing cortical activity. In analyzing the effects of optogenetic thalamic activation on thalamic sensory responses, the dataset was limited to those in which there was significant MUA excitation in response to whisker stimulation within a 30 ms window post stimulus (n=8 recordings, 7 animals). Likewise, in analyzing the effects of optogenetic activation on cortical sensory responses, the dataset was limited to those recordings in which there was significant cortical LFP and MUA response to whisker stimulation within a 30 ms window post stimulus (n=8 recordings, 6 animals).

4.3 Results

In this study, we examine the effects that optogenetic control has on thalamic activity, downstream cortical activity, as well as consequences for sensory responses in both areas.

Optical stimulation and extracellular electrophysiological recordings were conducted in the primary somatosensory thalamus of the awake mouse. Specifically, the target was the ventral posteromedial nucleus of the thalamus which receives topographically-aligned sensory input from the facial whiskers ([23]). Thresholded multi-unit activity (MUA) recorded across an electrode array was used to provide a large-scale measure of thalamic activity, rather than focusing on the activity of single neurons. The use of MUA has the added benefit of not requiring online spike sorting of recordings from dense electrode arrays, which is currently experimentally intractable. A one-dimensional dynamic state variable relating optical inputs and animals' whisker self-motion to measured MUA was used as a representation of thalamic state. In general, higher-dimensional state estimates may be used; however, the goal of this analysis was a single variable that best explained coordinated changes in across-channel firing activity. In addition to thalamic measurement and stimulation, extracellular electrophysiology was also recorded downstream of VPM in the whisker-recipient region of primary somatosensory cortex (S1) known as the barrel cortex ([204, 24]).

4.3.1 Optogenetic control and multi-unit thalamic activity

Closed-loop optogenetic control of thalamic activity was carried out using spiking activity of a putative single neuron (hereon termed the “feedback (FB) neuron”), as detailed in Section 3.2.6. Briefly, a reference (r), or target, firing rate was fed to the control loop consisting of a controller (“CTRL”) and a state estimator (“EST”) (Figure 4.1(a)). The controller consisted of feedback gains on instantaneous state error and integrated output tracking error as well as a feedforward model-based optical input for maintaining steady state firing at the reference rate (Section 3.2.6). The estimator that provided state feedback was the parameter-adaptive Kalman filter described previously (Section 3.2.5). This filter provides a robust estimate of output firing rates by adaptively re-estimating an unmeasured additive disturbance.

The feedback control system was tasked with clamping FB neuron firing rate at an elevated level of 20 spikes/s (Figure 4.1(b)), which can occur naturally in this pathway in the context of animal whisking (*e.g.*, [35, 27]). While the control system updated light intensity as a function of FB neuron activity, in most experiments, 32-channel electrode arrays were used to record population thalamic activity simultaneously. Extracellular voltages were high-pass filtered and thresholded (Section 4.2.4), giving spike times from single or multiple neurons/units in the vicinity of each electrode termed multi-unit activity (MUA, Figure 4.1(c)). As evident in the peristimulus time histogram (PSTH) of MUA in Figure 4.1(d), although activity of a putative single FB neuron was used to modulate the light intensity in a closed-loop fashion, the effects of the light are felt across a wide swath of the approximately 400 μm span between bottom and top electrode contacts on the array. The PSTH for each channel of MUA is arranged according to depth on the probe (left), such that lower channels were deeper in the brain and vice-versa. In this particular example, deeper channels were excited more than shallower channels, although this was not always the case.

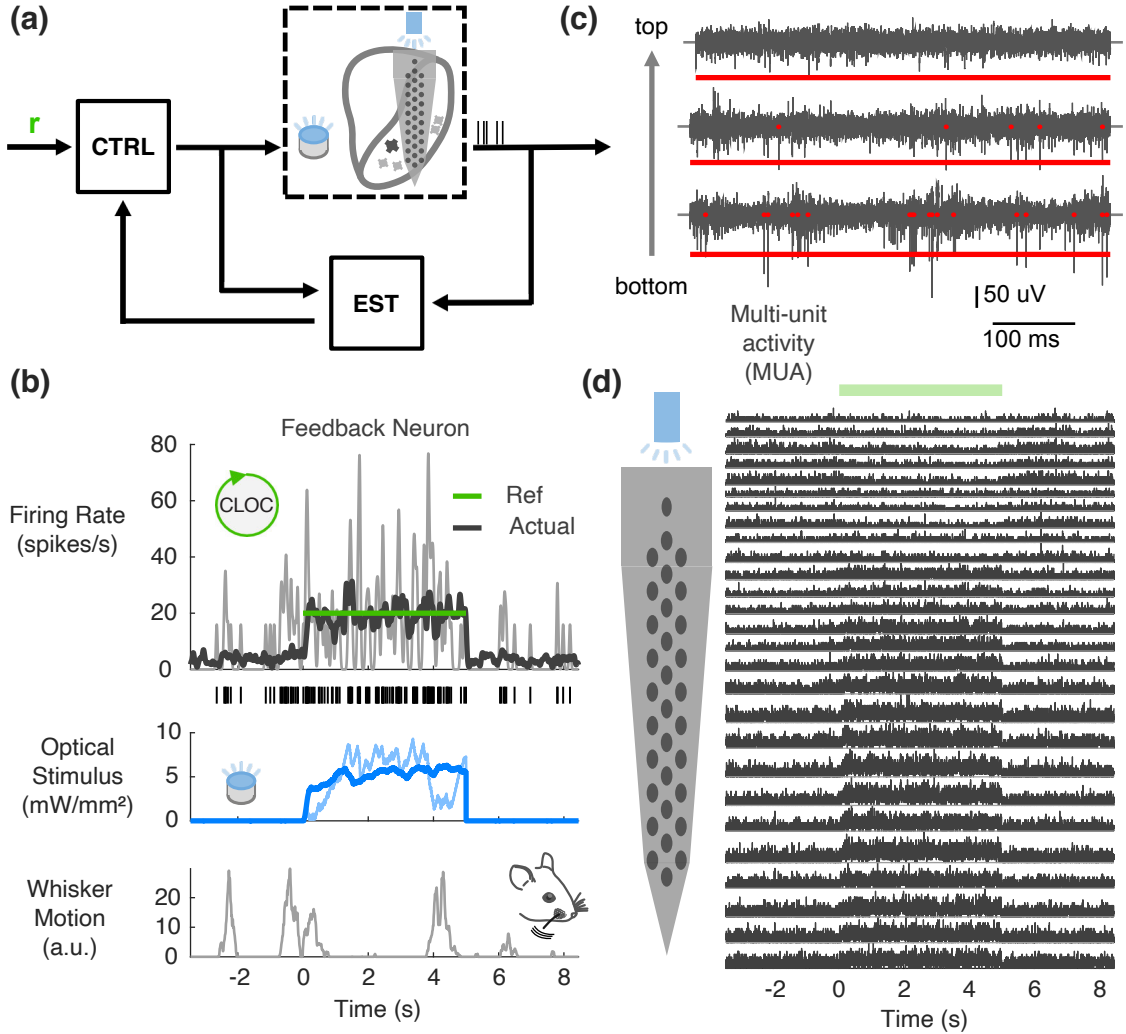


Figure 4.1: **Closed-loop optogenetic control and multi-unit activity.** (a) Block flow diagram for closed-loop optogenetic control. A single neuron is used for feedback while simultaneously recording a local population. (b) Example of closed-loop optogenetic control of feedback neuron. Top, trial-averaged (black) and example single-trial (grey) feedback neuron spiking smoothed with 25-ms standard deviation Gaussian window. Middle, optical control stimulus. Bottom, animal self-motion of whiskers. (c) Three example channels of thresholded multi-unit activity, from bottom to top of 32-channel electrode array. Thresholds and thresholded spike events shown in red. (d) Trial-averaged firing rate of MUA for recording in (b).

4.3.2 Thalamic state estimation

In order to capture coordinated changes in across-channel firing, a one-dimensional dynamic state variable, x , of a Gaussian linear dynamical system (GLDS) model relating optical inputs and animal whisker movement to measured MUA was used as a representation of thalamic state (Figure 4.2(a)). As mentioned previously, such state estimates may in general be multidimensional. Using training data consisting of spontaneous and optical noise trial conditions, a single set of dynamics (*i.e.*, parameters a , b , and q as detailed in Section 4.2.7) were fit across all recordings ($n=18$ recordings, 11 animals) to ensure a consistent, albeit abstract, meaning for this thalamic state variable (Section 4.2.7). Effectively, this state term describes the dynamics in the (anti-)correlated firing rate across the electrode array.

Estimation of this latent state variable was carried out using the adaptive Kalman filter detailed previously in the context of feedback estimation for control (Section 3.2.5). An example of the GLDS model-predicted and Kalman-filtered state estimate is provided in Figure 4.2(b). A single trial of optical step input, as well as animal whisker motion, is shown to have a predicted excitatory effect on the state variable (red). Given measured MUA, Kalman filtering provides data-driven corrections to this model estimate (purple). A raster plot of multi-unit spiking is also shown for the same trial, where each row is a channel on the electrode array arranged from top to bottom as before in Figure 4.1. The estimated thalamic state changes (purple) follow the apparent changes in spike density in the raster. Although difficult to see, same can be said of trial-averaged estimates of output firing rate on each channel, shown at the top of Figure 4.2(b).

The accuracy of the GLDS model state prediction and the filtered state estimates was assessed based on the estimation accuracy of the output firing rates compared to measured MUA. The squared across-time estimation bias was calculated between the model/filter output and MUA spiking on a single-trial basis in each of four trial conditions: spontaneous (Spont), optical noise (Noise), open-loop (OL), and closed-loop (CL) steps (Figure 4.2(c)).

As expected, the filtered estimate of state (purple) was a more accurate representation than the forward prediction of the model. In general, the filtered state estimate was only off by between 20-40 spikes²/s² (4.5-6.3 spikes/s in root-mean-squared error) on average. In the case of optical noise and open-loop step inputs, where identical optical inputs were repeated for multiple trials, variance explained in the raw 2-ms-binned PSTH was calculated (Figure 4.2(d)). This explained variance was either taken as a proportion of the variance in the PSTH (pVE, top), or relative to the amount of “signal” (*i.e.*, explainable) variance in the PSTH (pSVE [175], bottom). Again, as should be the case, the filtered estimates captured greater variance than the model prediction without correction. In terms of pVE, the filtered state estimates captured 30-35% of the variance in PSTHs recorded in response to noise and OL steps. In the case of optical noise stimulus conditions, this translated to capturing approximately 50% of the explainable signal variance. Because there was less explainable variation in MUA responses to simple step inputs of light, the filtered state estimates were actually able to capture 100% of this signal variance.

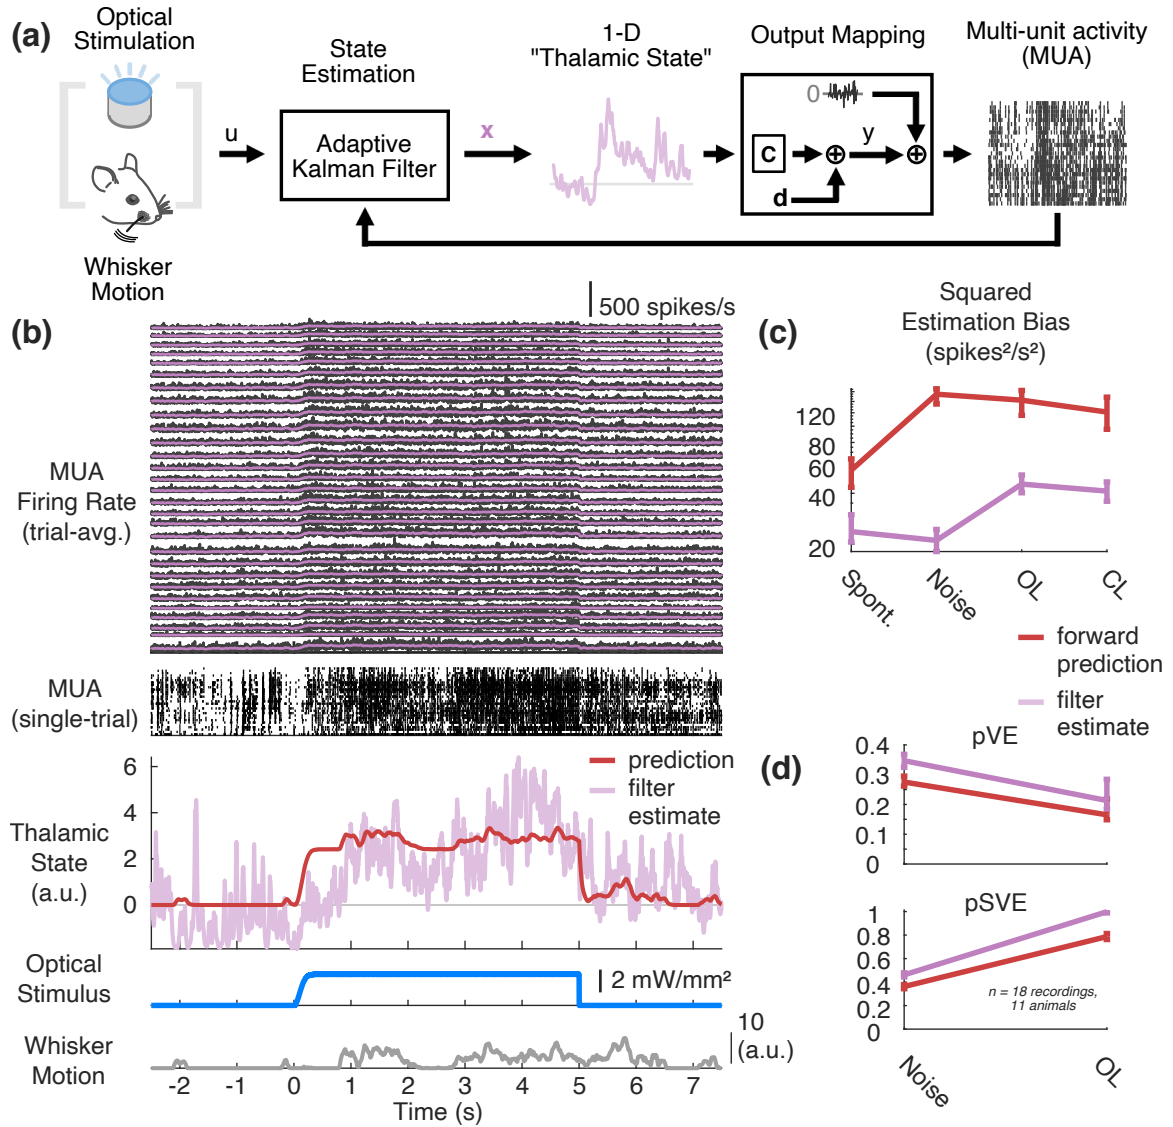


Figure 4.2: Estimating thalamic state. (a) Thalamic state estimation. Two-dimensional input (u) to a GLDS model comprised of the optical stimulus and whisker motion. Given this input and MUA spiking, an adaptive Kalman filter estimates a one-dimensional thalamic state, which is linearly mapped to MUA. (b) Example response to open-loop (OL) step input of light. Top, trial-averaged MUA firing rate (black) and output estimate (purple). Below is example single-trial MUA spiking. Middle, corresponding single-trial GLDS-predicted (red) and filtered (purple) state estimates. Bottom, corresponding single-trial optical and whisker inputs. (c) Population summary of squared output estimation bias for model-predicted (red) and filtered (purple) estimates for spontaneous, optical noise, closed-loop step, and open-loop step trial conditions. Error bars correspond to bootstrapped 95% confidence intervals about the mean. (d) Population summary of proportion of variance and signal variance explained for model-predicted (red) and filtered (purple) estimates for optical noise, and open-loop step trial conditions.

4.3.3 Effect of optogenetic control on thalamic state

While we have previously characterized the effect of closed-loop optogenetic control on the FB neuron (Section 3.3.3) and other simultaneously recorded neurons (Section 3.3.4), here we investigate the effect of optogenetic stimulation, both closed-loop (CL) and open-loop (OL), on the estimated state of the thalamus. Figure 4.3(a-b) provides an illustrative example of the effect of optogenetic control on thalamic state both in a single-trial (light purple) and on average (bold purple), as well as corresponding FB neuron activity around which the controller operated and the average MUA (Figure 4.3(b)), where output channels are arranged according to depth as before. The across-experiment distribution of thalamic state during closed-loop control versus a spontaneous baseline condition ($n=16$ recordings, 10 animals) is plotted in Figure 4.3(c). At baseline the thalamic state is distributed about zero, albeit skewed. During optogenetic stimulation (blue), the distribution shifts significantly rightward to a mean value of approximately 2 (Figure 4.3(c), $p = 4.38 \times 10^{-4}$, Wilcoxon signed-rank test between per-recording baseline and CL means). Despite the decrease in FB neuron variability observed previously (Section 3.3.3), closed-loop stimulation did not produce lower state variability. Instead, the thalamic state variability in panel (c) is very clearly elevated during control; in fact, it is approximately twice as great in periods of CL stimulation (1.41 a.u.^2) as compared to the spontaneous baseline condition (0.734 a.u.^2) ($p = 0.0061$, Wilcoxon signed-rank test). In addition, variability during OL control epochs such as shown in Figure 4.2(a) was not significantly greater than that observed during CL ($p = 0.179$, Wilcoxon signed-rank test). This begs the question what CL stimulation is achieving over OL in terms of thalamic activity more globally. We have seen previously that CL stimulation has the ability to compensate for experimentally-imposed disturbances (Chapter 2). An open question is to what extent CL optogenetic stimulation achieves this in the context of wakefulness where natural disturbances occur continuously.

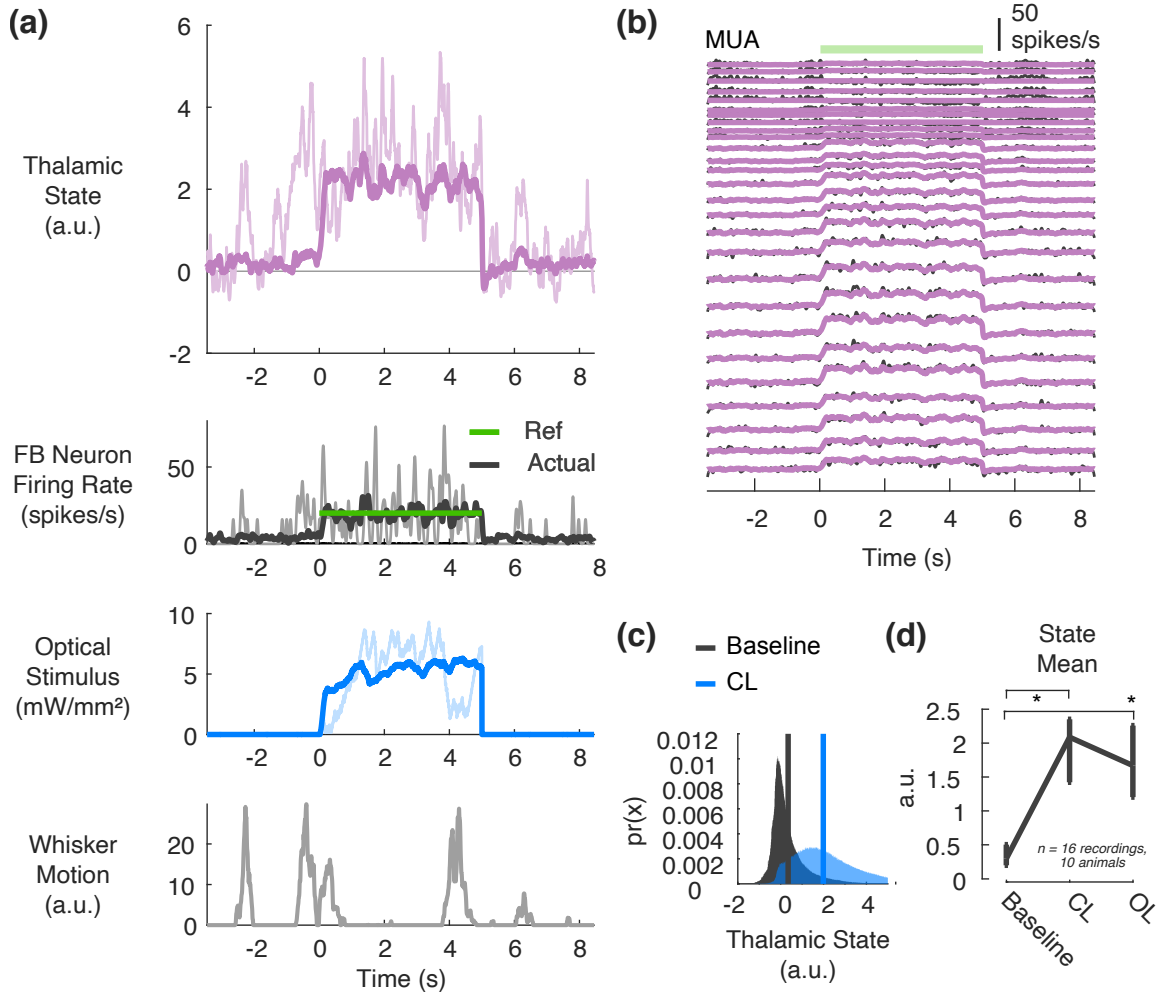


Figure 4.3: **Effect of optogenetic control on thalamic state.** (a) Example thalamic state during closed-loop optogenetic control of feedback neuron. Single-trial (light purple) and trial-averaged (bold purple) thalamic state estimates during feedback control of a putative single neuron (below). (b) Trial-averaged MUA firing rate for the same data in (a). Green bar signifies five-second control epoch. (c) Thalamic state distribution during baseline spontaneous (black) versus closed-loop step (blue) conditions. Vertical bars signify means. (d) Thalamic state means for spontaneous baseline vs. closed-loop ($p = 4.38 \times 10^{-4}$), and open-loop step ($p = 4.38 \times 10^{-4}$) input conditions. Error bars correspond to bootstrapped 95% confidence intervals about the mean.

4.3.4 Whisking as a natural thalamic state disturbance

So far, we have focused on the effects of optical stimulation on thalamic activity. However, under awake conditions, there are spontaneous (*i.e.*, uncontrolled) changes in neuronal activity, such as those that occur when animals move their facial whiskers (Section 1.3). Such behaviorally-correlated activity changes are apparent in the above example. Especially clear in periods of whisker motion, such as just before two seconds pre-control, at control onset, and at approximately 1 second after the control epoch, the single-trial activity of the FB neuron (Figure 4.3(a), grey, FB neuron spike train smoothed with 25-ms standard deviation Gaussian window) and the single-trial estimate of thalamic state (light purple) are correlated to the recorded whisker motion. Both the feedback neuron and the more global estimate of thalamic state are transiently excited as the animal “whisks”. Notably, the optical control signal on the same trial (light blue) sensed and compensated for this apparent disturbance due to whisking, as shown by lower light intensity at the beginning and towards the end of the 5-second control epoch, concurrent with whisking.

With this observation and the knowledge that firing activity in VPM thalamus is modulated by whisking ([35, 27, 160]), we investigated the extent at which this representation of thalamic state was modulated with whisking. Figure 4.4(a) provides an example of spontaneous single-unit activity (top), thalamic state (middle), and MUA (bottom raster) as it is modulated in time with whisking. Since previous studies have quantified average firing rates in whisking vs. non-whisking conditions, we first considered the population of all single units recorded ($n=52$) and segmented the spontaneous data into non-whisking or whisking at all times when the whisker motion signal was zero or greater than zero, respectively (see Section 4.2.3). During periods of whisking, the average single unit firing rate increased from 4.9 spikes/s to 8.0 spikes/s ($p = 1.40 \times 10^{-8}$, Wilcoxon signed-rank test), while previous studies have reported firing rates of 15-20 spikes/s during whisking in VPM ([35, 27, 160]). It is important to note that many, or at least some, of the recorded units in the present study were likely not VPM neurons because the inclusion criterion

used for these analyses was that units be significantly excited by optical pulses with short latency. In the case where opsin expression was achieved by viral injection, excitatory cells in other nearby nuclei could have been picked up on the recording electrodes. In the case of transgenic mice, cells in the VPL thalamus could likewise have been recorded. Since only a single whisker was probed in each experiment, it is impossible to know which units were not whisker-driven. For these reasons, the discrepancy in the degree of firing rate modulation with whisking observed here versus previous studies is not unexpected.

The distribution of spontaneous thalamic state was segmented according to whisker motion in the same way as the putative single-unit activity (Figure 4.4(c)). As with single unit activity, thalamic state was significantly activated during whisking ($p = 1.96 \times 10^{-4}$, Wilcoxon signed-rank test). Note that while the movement of the whiskers could be providing reafferent input to the thalamus, it has also been previously shown that thalamic firing rates are elevated during whisking even when the infraorbital nerve is severed ([35]). For this reason, the measured whisker motion is most accurately described as a correlate of an unmeasured thalamic state disturbance. For simplicity, we will here consider animal whisking itself to be a disturbance of thalamic state.

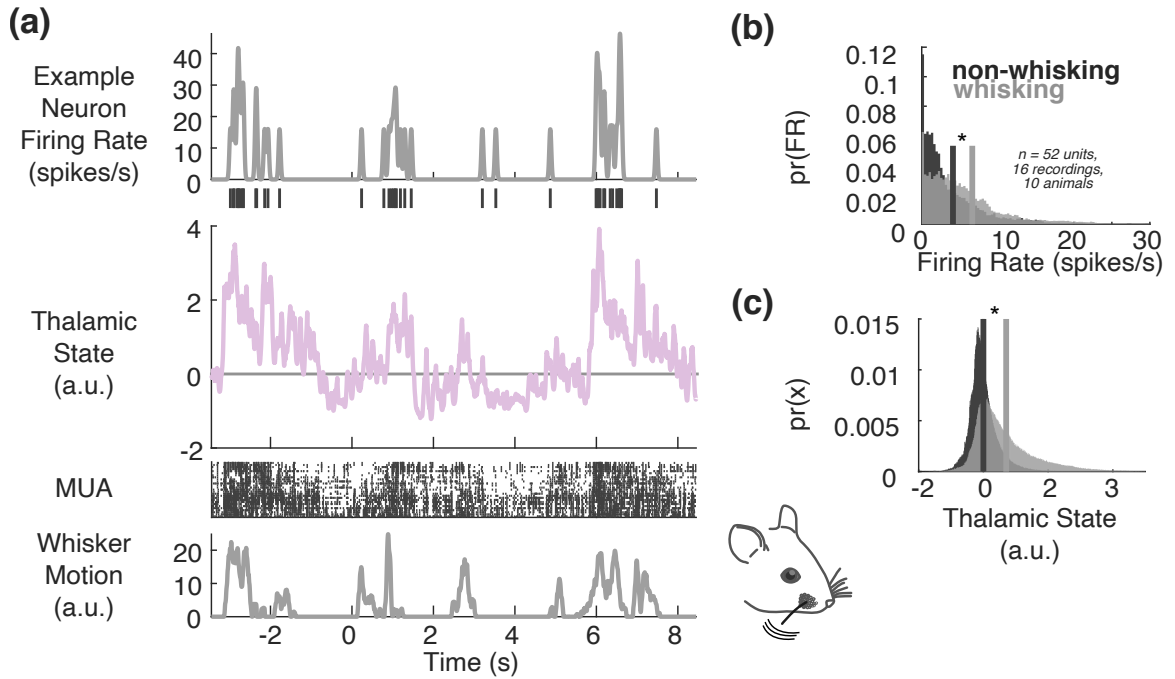


Figure 4.4: **Whisking as a natural thalamic state disturbance.** (a) Example effect of whisking on a putative single neuron (top, smoothed with 25-ms s.d. Gaussian window), thalamic state estimate (middle), and corresponding single-trial MUA raster. (b) Summary firing rate distribution for neurons in non-whisking vs. whisking conditions. Vertical lines indicate means ($n = 52$ neurons, $p = 1.40 \times 10^{-8}$). (c) Summary state distribution for non-whisking vs. whisking conditions. Vertical lights indicate means ($p = 1.96 \times 10^{-4}$).

4.3.5 Closed-loop optogenetic stimulation resists whisking thalamic state disturbance

Since one of the chief reasons to use feedback is to reject disturbances, we examined the extent to which closed-loop optogenetic control resisted thalamic state change during whisking. Figure 4.5(a) provides an example of a long (30-second) CL control epoch. In the single-trial case (light colored traces) during periods of whisking (*e.g.*, approximately 5, 12, 20, and 27 seconds into control epoch), there are compensatory dips in the optical control signal (light blue). To visualize this more clearly, 1-second snippets of thalamic state, optical inputs and whisker motion were collected around the onset of any whisker motion during control. The ensemble average of these snippets is plotted in Figure 4.5(b), where the whisking-onset-triggered average thalamic state changes very little while there is a subtle decrease in optical input from on average 1.5 mW/mm^2 to 1.4 mW/mm^2 . The population distribution of closed-loop optical stimulation during non-whisking (black) vs. whisking (grey) periods (Figure 4.5(c), $n=16$ recordings, 10 animals), shows that on average the light intensity needed to keep the FB neuron firing at 20 spikes/s went from 4.7 mW/mm^2 to 3.8 mW/mm^2 during whisking ($p = 2.3 \times 10^{-3}$, Wilcoxon signed-rank test).

While the approximately 20% decrease in population average light intensity with whisking shows that the feedback controller is resisting changes in FB neuron firing rate, the question remains to what extent this compensation effectively rejects the disturbance to thalamic state. To investigate this, distributions of thalamic state were constructed for periods of non-whisking vs. whisking in a baseline condition of spontaneous activity (Figure 4.5(d), top), CL optogenetic control epochs (middle), and OL control epochs (bottom). It is clear in comparing black and grey distributions that thalamic state is comparatively more invariant to the whisking disturbance in the CL case as compared to the other trial types. While significantly different ($p = 0.0084$, Wilcoxon signed-rank test), the means of the non-whisking vs. whisking CL distributions (vertical black and grey lines) are closer than in either the baseline or OL conditions. The sensitivity index (d') was used to quantify

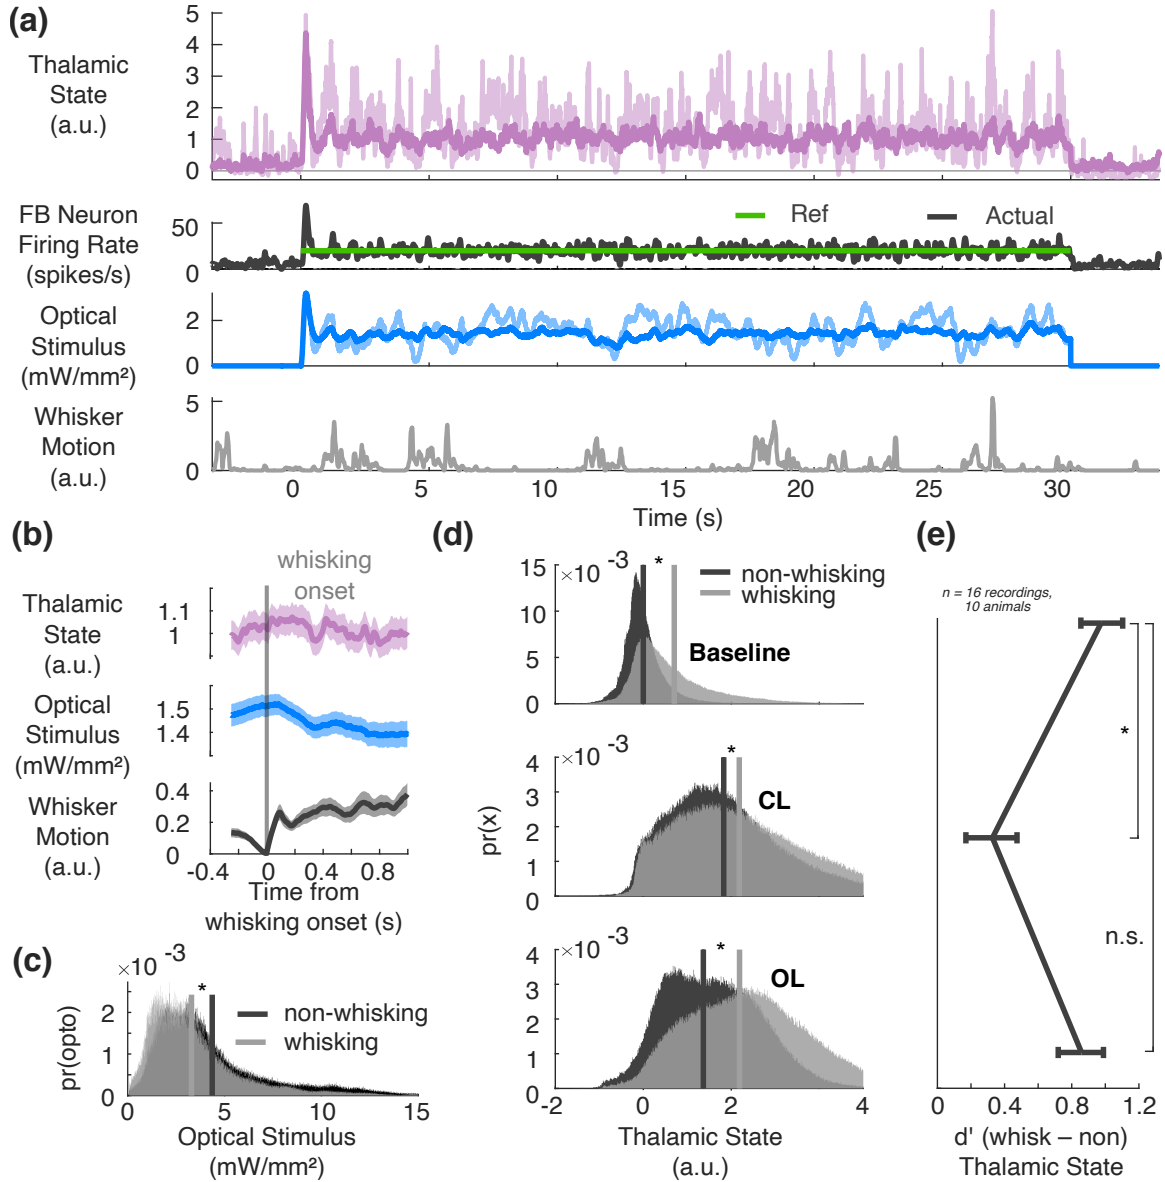


Figure 4.5: Closed-loop optogenetics resists whisking thalamic state perturbation. (a) Example 30-second control epoch. Single trial state, optical stimulus and whisker motion shown in light colors. Trial-averaged activity in bold. Feedback neuron firing has been smoothed with 25-ms s.d. Gaussian window for visualization. (b) Whisking onset-triggered average state and optical stimulus for example in (a). (c) Summary distribution of optical stimulus in non-whisking vs. whisking conditions (vertical bars indicate means, $p = 2.3 \times 10^{-3}$). (d) Non-whisking vs. whisking distributions (vertical bars indicate means) for spontaneous baseline (1.96×10^{-4}), closed-loop ($p = 0.0084$), and open-loop ($p = 4.38 \times 10^{-4}$) conditions. (e) Non-whisking vs. whisking selectivity index (d') for distributions in (d) in baseline vs. closed-loop ($p = 4.38 \times 10^{-4}$), and open-loop ($p = 0.109$) conditions. Error bars correspond to bootstrapped 95% confidence intervals about the mean.

the distance between non-whisking vs. whisking distributions in each trial type:

$$d' = \frac{\mu_{\text{whisk}} - \mu_{\text{non-whisk}}}{\sqrt{\frac{1}{2} (\sigma_{\text{whisk}}^2 + \sigma_{\text{non-whisk}}^2)}},$$

where μ_i and σ_i^2 denote the mean and variance of the state in whisking or non-whisking, respectively. Figure 4.5(e) shows that while the non-whisking and whisking distributions are still discernible in the CL case (*i.e.*, $d' > 0$, $p = 2.7 \times 10^{-3}$, Wilcoxon signed-rank test), there is significantly less difference between the two conditions than in the OL or spontaneous case ($p = 4.38 \times 10^{-4}$, Wilcoxon signed-rank test). Therefore, although the disturbance is not completely rejected, closed-loop optogenetic control does indeed resist thalamic state perturbation by whisking.

4.3.6 Optogenetic control of thalamus pushes cortex towards activated state

The cerebral cortex, like the thalamus, displays different population activity as a function of the state of the brain. During wakefulness, cortex is often said to be in a “deactivated” state when the local field potential (LFP) is characterized by higher amplitude, low-frequency fluctuations ([32]). Such a cortical state has also been termed “synchronized”, as it has been shown in such regimes that the activity of nearby cortical neurons are more highly correlated ([34]). Conversely, in an “activated” cortical state of wakefulness, the LFP is characterized by lower-amplitude, high-frequency activity. A continuous-valued quantity that captures this spectral redistribution of signal power is the ratio of power in low frequency (1-10 Hz) to high frequency (30-90 Hz) bands ([32]).

While there are many factors effecting the state of the cortex, it has been previously shown through direct thalamic manipulation that spiking activity in the thalamus has a profound effect on cortical state [51, 35]. And similar to the thalamus, it is known that the cortex displays characteristics of an activated state during periods of whisking ([34]) or locomotion ([41, 38]). Therefore, we investigated the extent to which optogenetic control of thalamic activity changed the state of the cortex, and how robust this change was to natural changes concomitant with whisking.

Example snippets of spontaneous cortical LFP and thalamic state and whisker motion are shown in Figure 4.6(a). Immediately preceding whisking onset, the thalamic state is low (hovering around zero, grey line) and the cortical LFP is characterized by higher amplitude, low-frequency activity. During whisking, the thalamic state is elevated and the cortical LFP is lower-amplitude and higher-frequency. To quantify the LFP spectral characteristics indicative of changing cortical state, the power spectral density was estimated in a 1-second sliding window, the frequency bins corresponding to the aforementioned low- and high-frequency bands (Figure 4.6(b), red and pink, respectively) were integrated and the low-to-high frequency ratio was taken. Figure 4.6(c) provides an example of a 30-second CL control epoch in which thalamic activity was elevated with direct optogenetic stimulation

and downstream cortical LFP was recorded simultaneously. In this case, the frequency ratio was reduced from approximately 9 before control onset (dotted line) to approximately 5 during control in thalamus. While not significant due to variability across datasets ($p = 0.0781$, Wilcoxon signed-rank test, $n = 7$ recordings, 4 animals), CL control of thalamic activity pushed the cortical frequency ratio from 15.2 on average at baseline to a value of 12.5. Therefore, while the effect was variable, optogenetic control of thalamus pushes the cortex towards a so-called “activated” state, consistent with [35].

As was done previously for thalamic state, to investigate the effectiveness of thalamic control in maintaining cortical state in the face of a perturbation, the cortical frequency ratio was segmented into distributions for when the animal was whisking (grey) vs. not (black) in each trial condition (Figure 4.6(d)). As before in thalamus, the non-whisking and whisking distributions are very close to each other in the CL control case, as compared to the baseline. In fact, the means of these two distributions (vertical lines) are not significantly different across recordings ($p = 0.375$). In contrast, there is a clear change in the baseline condition such that whisking takes the frequency ratio from a value of 22.5 to 12.5 ($p = 0.0156$, Wilcoxon signed-rank test). Sensitivity index, d' , was again used quantify the distances between the whisking/non-whisking distributions, and while the change between whisking and non-whisking conditions was less on average for CL stimulation than at baseline, this change was not significant due to variability across recordings ($p = 0.156$, Wilcoxon signed-rank test). In contrast to findings in thalamus (Section 4.3.5), the OL control of thalamus resulted in the same effects in cortex as CL. While speculative, it may be the case that the effects of thalamic firing and animal whisking on cortical activity are more binary in nature, such that continuous CL titration of optical stimulation has more subtle consequences for downstream cortical activity.

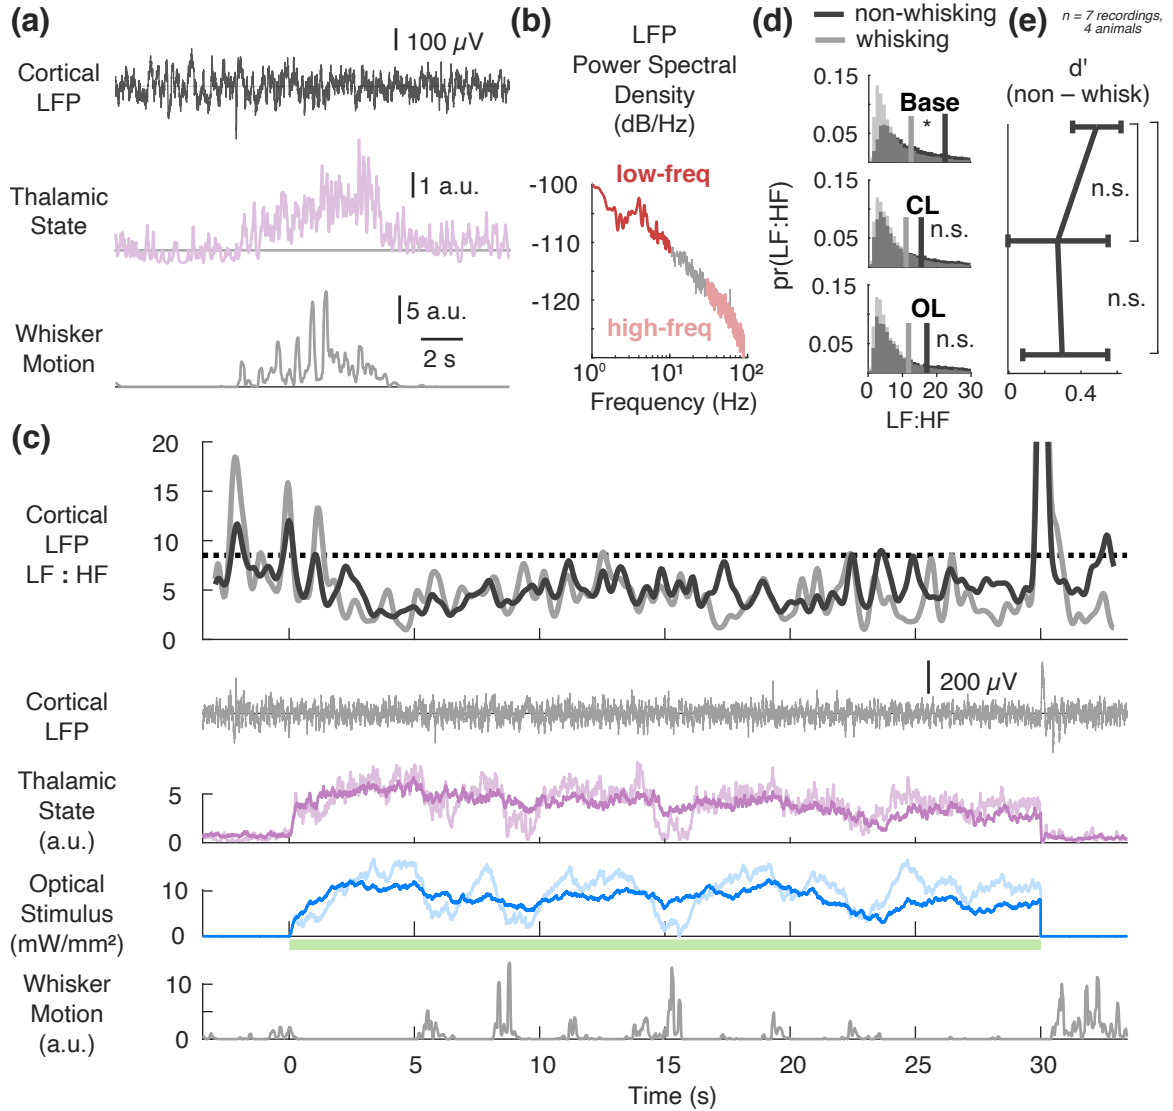


Figure 4.6: Effect of thalamic optogenetic stimulation on cortical activity. (a) Example of spontaneous cortical local field potential (LFP) and thalamic state during whisking. (b) Example power spectral density estimate for cortical LFP. Highlighted in red and pink are the low (1-10 Hz) and high (30-90 Hz) frequency ranges used to calculate low-to-high frequency ratio. (c) Example of cortical activity during closed-loop thalamic stimulation. Light colors signify single-trial estimates for each signal, whereas bold signify trial-average. Dotted line in the top plot corresponds to average low-to-high frequency ratio pre/post optogenetic control. (d) Summary non-whisking vs. whisking distributions (means, vertical bars) for cortical LFP low-to-high frequency ratio in baseline ($p = 0.0156$), closed-loop ($p = 0.375$), open-loop ($p = 0.0781$) conditions. (e) Non-whisking vs. whisking selectivity index (d') for distributions in (d) in baseline vs. closed-loop ($p = 0.156$), and open-loop ($p = 0.0781$) conditions. Error bars correspond to bootstrapped 95% confidence intervals about the mean.

4.3.7 Effect of thalamic state elevation on thalamic and cortical sensory responses

Given these effects of optogenetic control on thalamic and cortical states, the question remains what effect if any such state manipulations have on thalamic sensory responses. In the rodent whisker pathway, it has been previously put forward that an elevated thalamic state both activates the cortex ([35, 51]) and depresses the thalamocortical synapse, resulting in suppressed cortical responses to whisker stimulation ([20]). Such cortical response suppression is thought to sharpen cortical somatotopic receptive fields ([22, 50]) and mediate better discrimination between whiskers at the expense of simple sensory detection ([21]).

We first investigated the extent to which an elevated pre-stimulus thalamic state affected thalamic sensory responses. For this investigation, the state of the thalamus would ideally be clamped a longer timescales but let through sensory responses at a faster timescale. After all, if closed-loop stimulation actively rejected responses to sensory inputs, it would be impossible to investigate the effects of overall thalamic state on neuronal sensory representations. A computer-control galvanometer delivered one second of 10-Hz repetitive stimulation to a single whisker, where each punctate deflection consisted of a rapid exponential rise and fall (300 °/s average angular velocity; 99% rise time, 5 ms; 99% fall time, 5 ms). This sensory stimulation was delivered in a baseline condition without optical stimulation or embedded in the middle of a 5-second OL or CL control epoch, an example of which is shown in Figure 4.7(a). In the case of CL control, feedback was disabled for the 50 ms after each deflection to prevent compensatory changes in optical input. In this particular case, the whisker responses can be seen in the middle 5-10 channels of the electrode array (Figure 4.7(a), top, MUA PSTH smoothed with 5 ms s.d. Gaussian window). In each dataset ($n = 8$ recordings, 7 animals), thalamic MUA sensory responses were calculated for the channel which responded with the largest average magnitude to the first whisker deflection. Shown in Figure 4.7(b) is the population average peak response in the PSTH (*e.g.*, left) broken down by stimulus number in the 10 Hz train, where the response at base-

line is shown in black and the response during optogenetic state manipulation is shown in blue. For each recording, the responses were normalized to that of the first deflection at baseline. The peak of the PSTH response to the first stimulus adapted over the course of the 10 whisker deflections, such that the last response was approximately 80% the amplitude of the first at baseline; however, this change was not statistically significant ($p = 0.148$, Wilcoxon signed-rank test). Therefore, the responses to all 10 deflections were pooled together across datasets.

The resulting population-averaged response to whisker stimulation is shown in Figure 4.7(c), where each grey or blue line corresponds to the population average response to a single whisker deflection in the baseline or CL condition, respectively, and the bold lines indicate the overall response average. In the case of CL stimulation, it is clear there is a larger response to the sensory stimulus; however, this response is also riding on top of a heightened background of thalamic activity, shown both in the elevated pre-stimulus thalamic state (Figure 4.7(a)), as well as the return to this background post-stimulus (Figure 4.7(c), blue).

Since it is known that the thalamocortical synapse depresses with heightened thalamic activity ([20]), it is unclear whether from a cortical response or even perceptual perspective whether the absolute magnitude of the thalamic response is an appropriate measure of response strength. An alternative measure is to quantify the magnitude of the response, relative to the pre-stimulus background. Therefore, two measures were used to quantify single-trial responses to whisker deflections: the sum of spikes that occurred in a 30 ms window post stimulus (horizontal bar, Figure 4.7(c)), and this sum subtracted by the background activity in a 30 ms window immediately preceding the sensory stimuli. Both of these quantities were z-scored relative to the baseline condition on a recording-by-recording basis, such that the quantities in Figure 4.7(d-e) represent the number of standard deviations away from baseline the response was in each trial type. Consistent with previous findings in the urethane-anesthetized rat ([20]), an activated thalamic state achieved by either CL

or OL optogenetic control significantly increased the absolute response to sensory stimulation (Figure 4.7(d), CL condition $p = 7.58 \times 10^{-12}$). However, the response relative to background activity was actually smaller (Figure 4.7(e), CL condition $p = 0.001$).

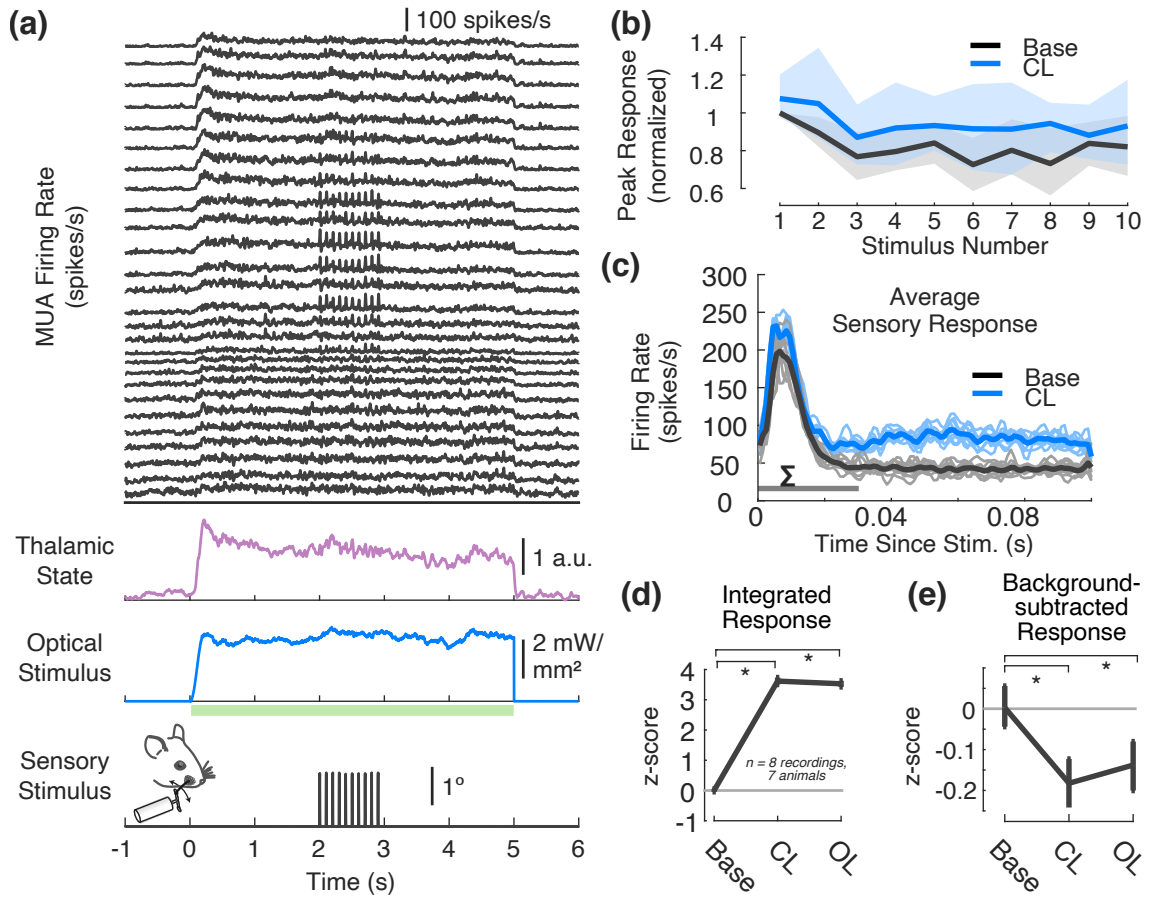


Figure 4.7: Effect of pre-stimulus thalamic activation on thalamic sensory responses. (a) Example MUA activity response to 10 Hz single-whisker deflections at 300 degrees/s velocity. All signals are trial-averaged. (b) Summary per-deflection average MUA firing rate response to whisker stimuli in baseline (black) vs. closed-loop optogenetic activation (blue) conditions, normalized to first response at baseline. Fills correspond to bootstrapped 95% confidence intervals about the mean. (c) Summary average MUA response to whisker deflection in baseline (black) vs. closed-loop optogenetic (blue) conditions. Light lines correspond to single-deflection average response, while bold indicates ensemble average. Horizontal bar indicates 30 ms window of time in which spikes were integrated for single-trial response measure (d-e). (d) Summary thalamic sensory response in baseline vs. closed-loop and open-loop optogenetic stimulation conditions (CL: $p = 7.5778 \times 10^{-12}$; OL: $p = 7.5767 \times 10^{-12}$; Wilcoxon signed-rank test). Spike counts in 30-ms post whisker deflection were z-scored according to baseline condition for that recording. Error bars correspond to bootstrapped 95% confidence intervals about the mean. (e) Summary pre-stimulus-subtracted sensory response in baseline vs. closed-loop and open-loop optogenetic stimulation conditions (CL: $p = 0.001$; OL: $p = 0.0031$; Wilcoxon signed-rank test). Spike counts in 30-ms post whisker deflection were subtracted by preceding spikes in the same 30-ms window.

Since the overall magnitude of thalamic sensory responses was greater while the background-subtracted response was smaller, the consequences of optogenetic activation of thalamus for cortical sensory processing are not simple to predict. On one hand, the number of spikes coming from thalamus to cortex is larger during optogenetic stimulation of thalamus; on the other, the increase in pre-sensory-stimulus firing rate leads to depression at the thalamocortical synapse, such that each impulse exerts less influence downstream. Previously, it has been reported in the urethane-anesthetized rat that thalamic activation by brainstem reticular formation stimulation or cholinergic agonist infusion leads to suppressed cortical responses downstream ([20, 51, 48]). Similarly, in the auditory pathway, sensory responses in primary cortex of the awake rat are suppressed when animals are engaged in a task and thalamic firing rate is likewise increased ([37]).

Cortical sensory responses were measured using LFP, such as in the example recording provided in Figure 4.8(a). For each recording, the channel that displayed the earliest significant depolarization in response to stimulation was taken as putative LIV. The response to each punctate deflection was quantified as the peak amplitude of the evoked response in a 30 ms. Similar to the analysis of thalamic spiking responses (above), these LFP response amplitudes were z-scored against the pooled responses in the baseline condition for each session ($n = 8$ recordings, 6 animals). Consistent with the trend in the thalamic responses, cortical responses adapt to repetitive stimuli (Figure 4.8(d), $p = 0.0078$, Wilcoxon signed-rank test of first vs. last response). As before, to quantify overall sensory response magnitudes, data for all whisker deflections were pooled together. The resulting population average evoked LFP response is shown in Figure 4.8. Note that this population average was carried out on raw LFP, without standardization, and only very minor differences between CL and baseline conditions survive the population averaging. Z-scored LFP response amplitudes (Figure 4.8(d)) reveal that there is a significant suppression of cortical responses as a result of CL optogenetic manipulation of thalamic state as compared to baseline ($p = 0.0223$, Wilcoxon signed-rank test). This suppression was expected due to

the well-known depression of the thalamocortical synapse with heightened thalamic activity ([20]). This decrease in efficacy of the TC synapse apparently outweighed the additional number of spikes observed in thalamus.

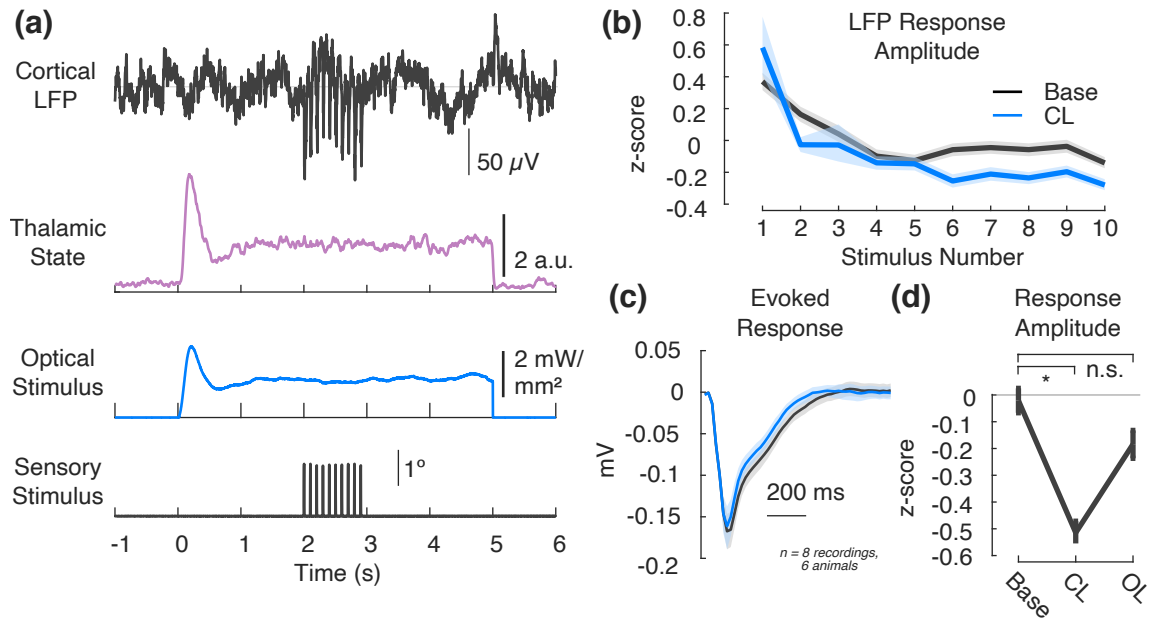


Figure 4.8: Effect of pre-stimulus thalamic activation on cortical sensory responses. (a) Example putative LIV cortical LFP response to 10 Hz single-whisker deflections at 300 degrees/s velocity. All signals are trial-averaged. (b) Summary per-deflection average LFP amplitude in response to whisker stimuli in baseline (black) vs. closed-loop optogenetic activation (blue) conditions. Values are z-scored against baseline condition responses for each recording, with all ten deflection responses pooled together. Fills correspond to bootstrapped 95% confidence intervals about the mean. (c) Summary average LFP response to whisker deflection in baseline (black) vs. closed-loop optogenetic (blue) conditions. (d) Population average LFP sensory response amplitude (CL: $p = 0.0019$; OL: $p = 0.2003$; Wilcoxon signed-rank test). Not shown, cortical multi-unit spiking response displays same trend (CL: $p = 0.0223$; OL: $p = 0.1053$; Wilcoxon signed-rank test). Values are z-scored against baseline condition responses for each recording. Error bars correspond to bootstrapped 95% confidence intervals about the mean.

4.4 Discussion

Unlike previous work that defines the state of the brain during wakefulness in discrete extremes such as quiet vs. active/whisking ([34, 35, 27]) or inattentive/passive vs. attentive/engaged ([36, 37]), the definition of thalamic state used in this study is a continuous-valued quantitative measure of changing population activity. The reason for this choice was multi-fold. First, by whatever measure, the state of the brain during wakefulness is a continuum ([29]), albeit with some well-studied extrema. Secondly, there are not agreed-up thresholds, for example, in thalamic firing rate that are indicative of an active vs. inactive, or quiet, state. And finally, unlike the periods of concerted “Up” and “Down” cortical activity that can occur in slow-wave sleep and under certain anesthetics ([46, 42, 205]), we saw no clear evidence for bi-/multi-modality in distributions of thalamic firing rate or the continuous state measure, or at least not until segmenting the data according to whisking behavior (Figure 4.4(b-c)). Given that the firing rate of thalamic neurons is known to be important both for changing the state of cortex as well as altering the sensory response characteristics ([35, 20, 51, 48]), the quantitative measure of large-scale, coordinated changes in firing rate used here would seem an ideal estimate of thalamic state.

Importantly, the notion of state developed in this study provided a way to investigate effects of optogenetic control on the thalamus more generally when closing the loop around a single feedback neuron. One of the most surprising results was that, despite a significant reduction in the variability in that feedback neuron’s activity both in the anesthetized rat (Section 2.3.7) and in the awake mouse (Section 3.3.3), this control actually yielded more variation in population activity than occurred under spontaneous, uncontrolled conditions. One might assume the cause of this heightened variability was simply the variation in optical stimulation in time, which arises by virtue of using feedback control. While this is certainly possible, the amount of thalamic state variability in the case of open-loop step inputs was not significantly different from the closed-loop case (Section 4.3.3). Since con-

trolling systems in a robust manner, with low variability is one of the chief reasons to use closed-loop control in general, this deserves further investigation. It is likely that closing the loop around this population state rather than a single neuron's activity would result in more effective population control. Since thresholded MUA was used as the measurement modality for the state estimator, this extension is imminently feasible, as it requires no sophisticated spike sorting techniques to distinguish action potentials of individual neurons online. Such activity has long been used in the motor field as the basis for decoders in brain-machine-interface applications ([103, 188, 206]). Moreover, the state-space estimation and control approaches used experimentally in the single feedback neuron case are immediately generalizable to such multi-output scenarios.

Among other things, this work demonstrates for the first time an example of closed-loop optogenetic control compensating for a naturally-occurring disturbance in the awake animal: in this case, whisking. While it cannot be said closed-loop stimulation completely rejected changes in thalamic state concomitant with whisking, thalamic activity changed significantly less during periods of closed-loop control than spontaneously or during open-loop control. It is also interesting that whisking minimally changed the state of the cortex when the thalamus was undergoing closed-loop stimulation. Since whisking is known to be accompanied by cholinergic inputs to both thalamus and cortex ([52]), one possible future application of this technique may be to hold thalamic state invariant to such perturbation and investigate non-thalamic origins of cortical state and sensory-response variability. However, this would require further investigation of the thalamic state variation finding discussed above.

In general, the effects of thalamic activation on cortex were more subtle and variable than expected. Given prior publications on the effects of direct thalamic activation either via cholinergic receptor agonist infusion ([51, 48]) or open-loop optogenetic stimulation ([35]), a much more robust change in cortical LFP was anticipated. In the case of [51, 48], the discrepancy could be due to the fact that these investigations were carried out under urethane

anesthesia, rather than in awake conditions. However, [35] used a very similar preparation: direct optogenetic stimulation of thalamus in the awake, head-fixed mouse. In that study, the authors used a ramp of blue light that went from 0 to approximately 16 mW/mm^2 (0.5 mW) over the course of five seconds. On average, the closed-loop controller used blue light of intensities lower than 5 mW/mm^2 (Figure 4.5(c)). The example dataset that showed the clearest effects of thalamic activation on cortical LFP used closer to an intensity of 10 mW/mm^2 (Figure 4.6(c)), so it is possible that the intensity of light used was disparate enough to explain the disparity. Also, [35] used a $200 \mu\text{m}$ optical fiber rather than the $100 \mu\text{m}$ used here. In that case, for a given light intensity, a larger area of tissue would be activated. All this said, it is also true that clear periods of “quiet wakefulness” were infrequent in this dataset, suggesting that animals were in general already in a heightened state of alertness in these experiments and that thalamic stimulation may have had little room for further activation of cortex.

4.5 Conclusions

This work presents a novel representation of thalamic state and shows that, even when using a single neuron for feedback, closed-loop optogenetic control effectively resists state changes concomitant with disturbances such as whisking in the awake mouse. Since whisking is known to be accompanied by cholinergic inputs to both thalamus and cortex, one possible future application of this technique may be to hold thalamic state invariant to such perturbation and investigate non-thalamic origins of cortical state and sensory-response variability. In addition to using the thalamic state defined here for feedback, future work could benefit from bidirectional control either through coexpression of excitatory and inhibitory opsins in thalamus or expression of a red-shifted excitatory opsin in the inhibitory reticular nucleus of the thalamus. Moreover, the utility of the state-space dynamical systems approach to neuronal state estimation need not be limited to thalamic activity. Such an approach to state definition and estimation could be applied to other systems to provide

continuous-valued, large-scale estimates of dynamic population activity in cases where there are not clearly defined discrete changes in state-related neuronal activity. Likewise, while developed and tested in the thalamus of the somatosensory system, closed-loop optogenetic control is a broadly applicable methodology that should enable decoupling activity of brain regions or cell types within a region and provide unprecedented ability to assess the relative roles of otherwise inextricably linked variables such as thalamic and cortical activity across states of wakefulness.

CHAPTER 5

CONCLUSIONS AND FUTURE DIRECTIONS

A common thread running throughout the work in this thesis is the integration of a number of engineering concepts and their application to problems in neuroscience. Chief among these include mathematical modeling and practical system identification for relating optical stimulation to neuronal activity, numerical or analytical optimization for design of feedback controllers, and estimation of latent variables from noisy neuronal data. Chapter 3 and the transition to state-space modeling for control as well as state estimation perhaps best embodies this intersection. While developed and tested in the context of control of thalamic spiking activity, the hope is that the work of Chapters 3-4 enables closed-loop optogenetic control to be more broadly applied in neuroscience. With this in mind, what follows is a discussion of the strengths as well as limitations of the current approach, what I see to be prudent next steps algorithmically and experimentally, as well as some ideas for future applications outside the realm of controlling thalamic firing rate.

5.1 Robustness of closed-loop optogenetic control

In the context of feedback control of neuronal firing rate, we have demonstrated robustness to sensory white noise disturbance under anesthetized conditions (Chapter 2) as well as robustness to model inaccuracy in controller design and implementation (Chapters 2 and 3). Over relatively long timescales, these advantages of feedback were conferred by the use of integral action for the control, but—just as importantly—by employing estimators that have zero DC estimation error. For the simple model-free control loop in Chapter 2, the moving average filter used to smooth spiking activity into an estimate of instantaneous firing rate has unit gain at DC. In this way, even if higher frequency content in the underlying firing rate were not accurately conveyed to the feedback controller, the DC component of this

feedback signal was always accurate. Similarly, in the case of the state estimator of Chapter 3, the random walk model for the adaptation of the process noise mean (Section 3.2.5) ensured near-zero estimation bias. In fact, this parameter-adaptive Kalman filtering has been termed a PI Kalman filter ([207, 178, 179]) because the random walk is a stochastic integration of an additive state correction term. Admittedly, this approach does not always ensure zero DC estimation error for multi-output systems (Section 3.3.6), but it still provides a vast improvement over a standard Kalman filter. Given that this estimation step proved so important for reaping the benefits of PI control in both model-free and model-based control applications, whatever future developments there are in control approaches for optogenetic stimulation, a similar degree of thought should be put into the quality of the feedback signal provided to the controller.

One of the chief reasons to use feedback control is the rejection of disturbances. As mentioned above, we demonstrated some degree of robustness to an artificial sensory disturbance in the anesthetized animal. However, it was not until later work (Chapters 3-4) that we were able to apply CLOC in the awake animal where natural disturbances occur continuously. One such disturbance was the thalamic state change concomitant with animal whisking. Throughout this work, feedback from single neurons was used to modulate optical stimulation whose effects were felt more generally. Given the variability in response to stimulation (see below), it was perhaps surprising to find that this stimulation approach was able to confer a degree of robustness to animal whisking when viewed through a large-scale measure of thalamic state developed in Chapter 4. While the single-neuron feedback approach resisted some changes in the state variable, moving forward, the MUA-based state estimate should be used for CLOC of thalamus. The same could be said of other potential applications of CLOC outside of the thalamus, such as controlling laminar activity in cortex.

5.2 Variability and feedback control of single-neuron vs. population activity

Going hand-in-hand with questions of robustness is the issue of variability. Perhaps the most surprising result from Chapters 3 and 4 was the fact that while closed-loop control reduced the trial-to-trial variability in spiking activity of the single neuron used for feedback, population activity was not less variable, whether analyzing other simultaneously-recorded single neurons (Chapter 3) or using a more wholistic measure of underlying thalamic state (Chapter 4). In part, this could be due to the fact that online spike sorting was imperfect, effectively allowing noise spikes from other nearby neurons to creep into the feedback signal. However, it is also the case that simultaneously-recorded neuronal populations exhibited heterogeneity in their response to light. While lower spiking variability may not be desirable in all control applications, I believe the utility of this approach over open-loop alternatives will be limited unless at least long-timescale population variability is reduced. Given heterogeneity in population responses, a prudent future direction would be to apply optogenetic control where a population state estimate such as that developed in Chapter 4 were used for feedback. Although not suitable for all applications, if the goal of a future study is to investigate how the activity in a given region (*e.g.*, thalamus) affects a downstream target (*e.g.*, cortex), invariant to other disturbances (*e.g.*, whisking), closing the loop around population activity would likely be more effective than doing the same around a single neuron.

5.3 Opportunities for dual-opsin bidirectional control

All optogenetic manipulation in this work utilized a single excitatory opsin. A dual-opsin approach to bidirectional manipulation presents an opportunity for finer control of activity and would serve to broaden the applicability of this technique. Perhaps the greatest limitation of the current approach is the inability to pull back on neuronal activity in addition to exciting it, analogous to a cruise control system that only controls throttle and not brakes.

Throughout this work, the non-scientific, practical reason for making the target for control an above-baseline firing rate was that ChR is a non-selective cation channel, and so opening these channels with light will depolarize cells. Choosing an elevated target firing rate that naturally occurs in this pathway allowed feedback control to display its utility by pulling back on the light intensity as needed during natural periods of rate elevation like whisking. Moreover, if firing rate went above a target for any substantial period of time, integral control would cease to be helpful, becoming instead a nuisance due to a phenomenon called integrator windup [208]. The PI controller would sense firing rate was too high and compensate by turning down the light intensity until bottoming out, at which point this uncontrolled error would continue to accumulate, making any future excitatory action of the controller delayed.

Adding simultaneous optogenetic inhibition to this loop would greatly improve the utility of the approach and should be prioritized over any algorithmic changes, since the state-space control approach in Chapter 3 will generalize to this two-input scenario. The most straight-forward way to achieve such bidirectional control would be to co-express ChR along with an inhibitory opsin like the chloride pump HR (see Section 1.4.2). Functional co-expression of these two opsins has been achieved *in vitro* (*e.g.*, [121, 6]). The two-opsin plasmid eNPAC2.0 ([118]) is the ideal candidate for this bidirectional CLOC, as it contains genes for both ChR and HR which should help to ensure there are stoichiometric equivalents of each opsin expressed (barring major differences in membrane trafficking). Usage of viral injection of cre-recombinase-dependent vector AAV-DIO-eNPAC2.0 in appropriate transgenic mice should allow cell type specific bidirectional control in, for example, VPM/VPL thalamus (NR133 mouse, [123]), parvalbumin-expressing interneurons ([122]), or LVI CT neurons ([54]). Notably, however, this construct has only been used in one published *in vivo* study, and it is not clear how effective it was bidirectionally ([119]). Another consideration is that even most red-shifted opsins have a degree of blue light sensitivity (see HR activation spectrum in yellow, Figure 1.5). Therefore, to minimize cross-talk between

stimulation of ChR and HR (or archaerhodopsin-based alternatives), it would be prudent to use blue light closer to 415 nm (rather than 470 nm used here) to stimulate ChR. Amber (*e.g.*, 590 nm) or even red light could be used to selectively stimulate HR. Before pursuing this opsin co-expression approach, however, the authors of the aforementioned *in vivo* study as well as the originating lab ([118]) should be consulted.

In my view, this co-expression of opsins in the cell type(s) undergoing control would be ideal in most applications because this way you are directly manipulating one population of cells at time. However, in cases where the cell type of interest for control is inhibited by other cells in the network, there is an alternative. A red-shifted ChR variant (*e.g.*, [116, 117]) could be expressed in inhibitory cell populations such as reticular nucleus of the thalamus (*e.g.*, [124, 209]) or inhibitory interneurons in cortex (*e.g.*, [122]). In this way, neurons could be excited directly using a blue-light sensitive ChR and inhibited via the local circuitry. In addition to perhaps being a more natural avenue for inhibition, one practical benefit of this is that some inhibitory opsins have been reported to display paradoxical excitatory effects at certain extremes ([210]).

5.4 Choice of mathematical model for optically-driven neural activity

Throughout this work, an attempt has been made to use the simplest models that are still useful for the intended application (*i.e.*, control and estimation of neuronal firing). The initial use of an LNP model (Chapter 2) was motivated primarily by the fact that it was the simplest commonly used spiking model that relates stimulation (usually sensory) to neural activity (Section 1.4.1). However, while the LNP and GLM are widely used in sensory neuroscience, they are built around feed-forward, non-dynamical filters and therefore are not amenable to most control methods which generally revolve around the use of transfer functions or state-space dynamical systems models. As a result, while the design methodology laid out in Chapter 2 respects the spiking activity of neurons, it did not leverage more commonly used control and estimation methods and was not designed to handle

future multi-input (*e.g.*, bidirectional control above) and multi-output (*e.g.*, multi-channel MUA) control problems. For this reason, when trying to develop an optogenetic control approach that would be more generally applicable—even beyond spiking activity—a state-space dynamical system approach was taken and the necessity of spiking nonlinearity was investigated (Chapter 3). At least in the context of holding firing rates steady at relatively long timescales, we found little was functionally gained by using a PLDS instead of the simpler GLDS, and what you gained was access to widely-used linear controls methods. Moreover, because other neuronal signals of interest like cortical LFP or voltage/calcium fluorescence imaging data are continuous-valued rather than spiking in nature, a GLDS model would actually be a more accurate statistical model in those cases. All these things together suggested ultimately that the GLDS model was an appropriate compromise between accuracy and utility in this context and was applied both in CLOC (Chapter 3) as well as in offline thalamic state estimation (Chapter 4).

This is not to say PLDS models are unnecessary for optogenetic control more generally. Under different experimental circumstances or control objectives, it may be the case that a PLDS model is a substantially better choice. Controlling spike timing at finer timescales is likely an example application where the Poisson LDS would provide substantially better results. Given that there are published methods developed for estimating the underlying state of these PLDS models (“stochastic-state point process filter”, [98, 109, 110]), one could simply leverage these nonlinear filtering techniques and design/implement feedback control in this log-linear state space as described for linear systems in Chapter 3.

In addition to this output nonlinearity for spike generation, the underlying dynamics of single cells and populations are clearly not linear, so in some applications even a PLDS would be an inappropriate modeling choice. One example of a clear nonlinearity that was encountered in the early days of this project is that of depolarization blockade. If the optical stimulation is intense and sustained enough that it opens too many ChR channels, the cell(s) can become depolarized to the point where they cannot repolarize for the generation of

subsequent action potentials ([211]). In this way, overly intense actuation of the system can effectively inhibit neuronal firing. This is a potentially-important behavior that cannot be captured by an LDS with a static input or output nonlinearity. However, in practice, such nonlinearities were not problematic if the target firing rate was in an achievable range. It is possible that using feedback control effectively maintained these nonlinear systems within a regime that could be more effectively capture using linear models than would otherwise have been the case.

5.5 Trajectory tracking vs. constant set point

In this work, we have applied model-free and model-based forms of proportional-integral control. In Chapter 2, a simple model-free PI controller was used to entrain patterns of temporal firing rate modulation, whereas in Chapter 3, model-based control was tasked with clamping the firing rate at a constant set point. This begs the question of how applicable the latter approach is to reference trajectory tracking problems. The controller was designed by solving an infinite horizon optimization problem (specifically, LQR), so it was not explicitly designed for tracking target patterns of activity. That said, the methods laid out in Chapter 3 would be directly applicable to tracking problems where the desired pattern of activity is slow compared to the dynamics of the system being controlled. For example, the average light-to-spiking impulse response for thalamic units tended to die out over tens of milliseconds, indicating that the methods laid out here for model-based control may not be completely applicable to control of patterns at that time scale or faster. In such cases, a finite horizon optimization of feedback controller gains and a nominal control input using a technique like iterative LQR ([212]) may prove beneficial or necessary.

5.6 Dynamical systems approach to brain state definition and estimation

A dynamical systems approach was taken in Chapter 4 to the problems of defining and estimating the “state” of the thalamus. In neuroscience literature, converging evidence

points to thalamic firing rate being at least correlated to changes in behavioral and cortical measures of brain state and to it being an important factor in determining the gating of sensory information flow to cortex and ultimately perception (see Section 1.3). Therefore, consistent with the engineering definition of state as being the minimum set of variables needed to explain future activity of a dynamical system, the state of the thalamus has been defined here as a one-dimensional variable that underlies coordinated changes in population thalamic firing, measured by MUA across an electrode array. Importantly, this approach is not equivalent to simply averaging activity across the channels of the array. In theory, at the extreme where half of the channels were suppressed and half were excited, a simple across-channel average of the activity may be destructive and obscure what were actually gross changes. In contrast, such anticorrelated changes would be represented in the output matrix (C) of the GLDS models used as the basis for this state estimation. This is far from a contrived scenario, as it has been shown in the visual pathway that exciting deep cortical layer VI CT cells to fire actively inhibits other cortical layers like LIV ([213, 214]).

While a one-dimensional state was used here, it is important to note that this need not be the case; in thalamus and certainly other regions like cortex, a higher number of latent variables may be needed to explain activity of interest. In motor systems especially, variations on this dynamical systems approach have been used to distill high dimensional array recordings down to a small, albeit greater than unidimensional, state vector (*e.g.*, [100, 104, 106, 64, 107]). Notably, however, this approach is rarely applied to sensory systems as a way to explain state-dependent variability in responses to stimuli. Therefore, such dynamical systems state definition and estimation may prove to be a fruitful intersection of engineering and neuroscience.

As was done for single-neuron feedback control applications in Chapter 3, a GLDS model and Kalman filtering was used for thalamic state estimation in Chapter 4. One of the somewhat surprising results noted above was that the state variation appeared to be greater at CLOC-elevated firing rates than under spontaneous conditions. Although speculative, it

is possible the use of a GLDS instead of a PLDS model could have resulted in this apparently elevated variability. While not truly Poisson, the variability in neuronal spiking tends to correlate positively with the firing rate ([165]). If a PLDS model with an exponential rectifying nonlinearity were used, it is possible that, in a compressed logarithmic state-space, the estimated state variation would not have been so grossly elevated at higher firing rates during CLOC versus spontaneous conditions. Therefore, this variability analysis may be an example case where a PLDS, rather than GLDS, would be warranted (see above).

Finally, the use of a dynamical system approach to defining state assumes this state variable is in fact continuous-valued, rather than discrete. While even experts on states of wakefulness will often say brain state varies along on a continuum (*e.g.*, [29]), neuroscientists still tend to bin states into discrete categories such as active vs. quiet or synchronized vs. desynchronized. In the context of the thalamic recordings analyzed here, there was not clear multi-modality in the distribution of firing rates over time. Therefore, in situations such as this where there is no agreed-upon threshold for transitions between discrete states, the dynamical systems approach seems appropriate. That said, in cases where there are clear transitions (*e.g.*, UP vs. DOWN states in cortex under anesthesia), the discrete state space equivalent to the dynamical system could be used: the hidden Markov model (HMM). In fact, HMMs have been used to date to capture coordinated On vs. Off spiking activity regimes in cortex ([215]). See also [216, 217] for other neuronal applications of HMMs.

5.7 Applicable measurement modalities for CLOC

Experimentally, spiking from single neurons were used as feedback for CLOC. However, as has been noted already, control around a larger-scale measure of activity such as thresholded MUA (Chapter 4) is likely the way to push forward with this technique. Using MUA also has an important practical advantage in that it lessens the need to have well-isolated single units and opens the window to using chronically implanted electrodes ([193, 63])

rather than acutely inserting probes each recording. This will be especially important for longitudinal studies involving animal behavior. Because of the generality of the approach to control and estimation laid out in Chapter 3, using MUA as the measurement modality would require no change in methodology.

Besides spiking activity, these approaches could also be used for feedback control of cortical LFP, calcium imaging signals, or voltage imaging signals. If anything, the use of GLDS models would be more appropriate in the case of these continuous-valued signals and, especially in the context of imaging, using low-dimensional state-space representations for feedback control as well as offline analysis would likely be beneficial.

5.8 Thalamic state and burst vs. tonic firing modes

Because of how rare they are during wakefulness, the role of low-threshold calcium-mediated bursts in the wakefulness is debated. In a recent study that used intracellular and juxtacellular recordings in VPM of awake, head-fixed mice, the authors reported that they observed no calcium-mediated bursting. Moreover, across putative single units recorded in the awake animal here (Chapters 3-4), on average less than 2% of spikes were members of bursts (data not shown). That said, when bursts do occur in response to sensory stimuli, it is thought they provide an “wakeup call” to cortical circuitry downstream ([14], see Section 1.3). By the very language chosen, this implies that the brain had to have been in a quiescent state for this bursting to occur (*e.g.*, “inattentive” period shown in Figure 1.3(a)). Therefore, it is possible that the role of thalamic burst firing in sensation is difficult to probe under experimental conditions because animals are likely in a state of alertness much of the time. One potential use of bidirectional CLOC (see above) could be to step the thalamus through different states ranging from very quiet/hyperpolarized to very active/depolarized, and apply sensory stimuli while recording in cortex. Because calcium-mediated bursting requires hyperpolarized membrane potential and is therefore anticorrelated to overall firing rate, the thalamic state definition developed here is ideal for this application.

5.9 Across-region feedback control

Instead of closing the loop locally in a single brain region as was done in this work, CLOC could also be used to entrain target activity in a brain region of interest (*e.g.*, cortex) by way of stimulation in an upstream region (*e.g.*, thalamus). For example, the control objective may be to make cortical activity invariant to unmeasured state disturbances using optogenetic thalamic stimulation, and subsequently analyze the thalamic activity required to maintain target cortical state in a way not dissimilar to how voltage clamp experiments are analyzed: inspecting the current injection necessary to keep membrane voltage fixed. While scientifically interesting by analogy to current-day and classical experiments using voltage clamp, I am not confident this would yield a clear result. First of all, while voltage clamp is able to maintain membrane potential of a single patched cell quite effectively, it is unlikely closed-loop optogenetic stimulation would ever be able to overcome all other modulatory forces that can underlie changes in the activity of a local population of cells, let alone a downstream population of cells. This would leave the results of such experiments difficult to interpret. That said, inability to completely regulate cortical state via thalamic stimulation, for example, is in itself a result of sorts, but it is always difficult to prove a negative.

5.10 Closing the loop around behavior

Another potentially interesting application of CLOC would be to use neuronal stimulation to achieve target behavioral performance in, for example, a sensory detection task. As discussed in Section 1.3.2 the state of the brain at the time of sensory input can have significant effects on the perception of a stimuli. In the context of audition and somatosensation, when the brain is in an active/aroused state, thalamic firing is elevated ([27]), cortical sensory responses to inputs are suppressed ([37, 20]) and sensory detection suffers ([38, 49]). In a set of experiments, one could causally probe the role of thalamic state in perceptual sensi-

tivity by combining optogenetic stimulation of thalamus and behavioral feedback with the goal of keeping animal detection performance (*e.g.*, correct positive or hit rate) invariant to the strength of a sensory stimulus. If this instantiation of CLOC were able to significantly flatten the psychometric curve of animal hit rate versus stimulus intensity, one could then analyze the thalamic state necessary to achieve such stimulus strength invariance. The expected result would be that for weaker sensory stimuli, a bidirectional controller would have hyperpolarized the thalamus, resulting in a larger cortical response downstream. Conversely, for stronger sensory inputs, one would expect the thalamus to have been pushed to an elevated state by depolarization.

5.11 Closed-loop optogenetic control for testing role of thalamic state in seizure susceptibility in absence epilepsy

Finally, CLOC may provide a useful technique to test the role of thalamic state in seizure susceptibility. In multiple rodent studies of absence epilepsy, there is a reported decrease in thalamic firing rate immediately before onset of the generalized seizures that occur in this disease (GAERS rat, [59]; Scn8a-mutant mice, [60]). This may be merely an epiphenomenon or it may be the case that keeping thalamic activity at a baseline level can effectively prevent such seizures from occurring. This is, of course, speculative, and absence epilepsies are often responsive to pharmacological treatment. However, this could provide a tractable model for future applications in other forms of generalized epilepsies in which the thalamocortical loop is known to play a major role.

Appendices

APPENDIX A
NON-SINUSOIDAL REFERENCE FIRING RATE: MEASURED THALAMIC
FIRING RATE MODULATION

Chapter 2 applies the model-free control scheme of Newman *et al* ([6]) to the problem of entraining patterns of rate modulation. In order to arrive at an example pattern of naturally-relevant rate modulation in the rodent somatosensory thalamus used in this study, data recorded in the VPM of awake rats from Waiblinger *et al* ([145, 16]) were analyzed. In example VPM neuron spike trains (Figure A.1(a)), there was evidence in a subset of trials for a tone in spiking autocorrelation at approximately 10 Hz (Figure B.1(b), green), which is in the frequency range in which rats move their whiskers. Such trials were taken as putative whisking trials. For one example VPM neuron, three such trials were identified, their spike trains were phase-aligned, averaged, and smoothed with a 20 ms s.d. Gaussian window, resulting in an estimate of naturally-occurring rate modulation (Figure A.1(d)). Analysis of the frequency content of this pattern showed that there was a DC firing rate of 20 spikes/s and that 95% of the power in this signal occurs between DC and 10 Hz (Figure A.1(e)).

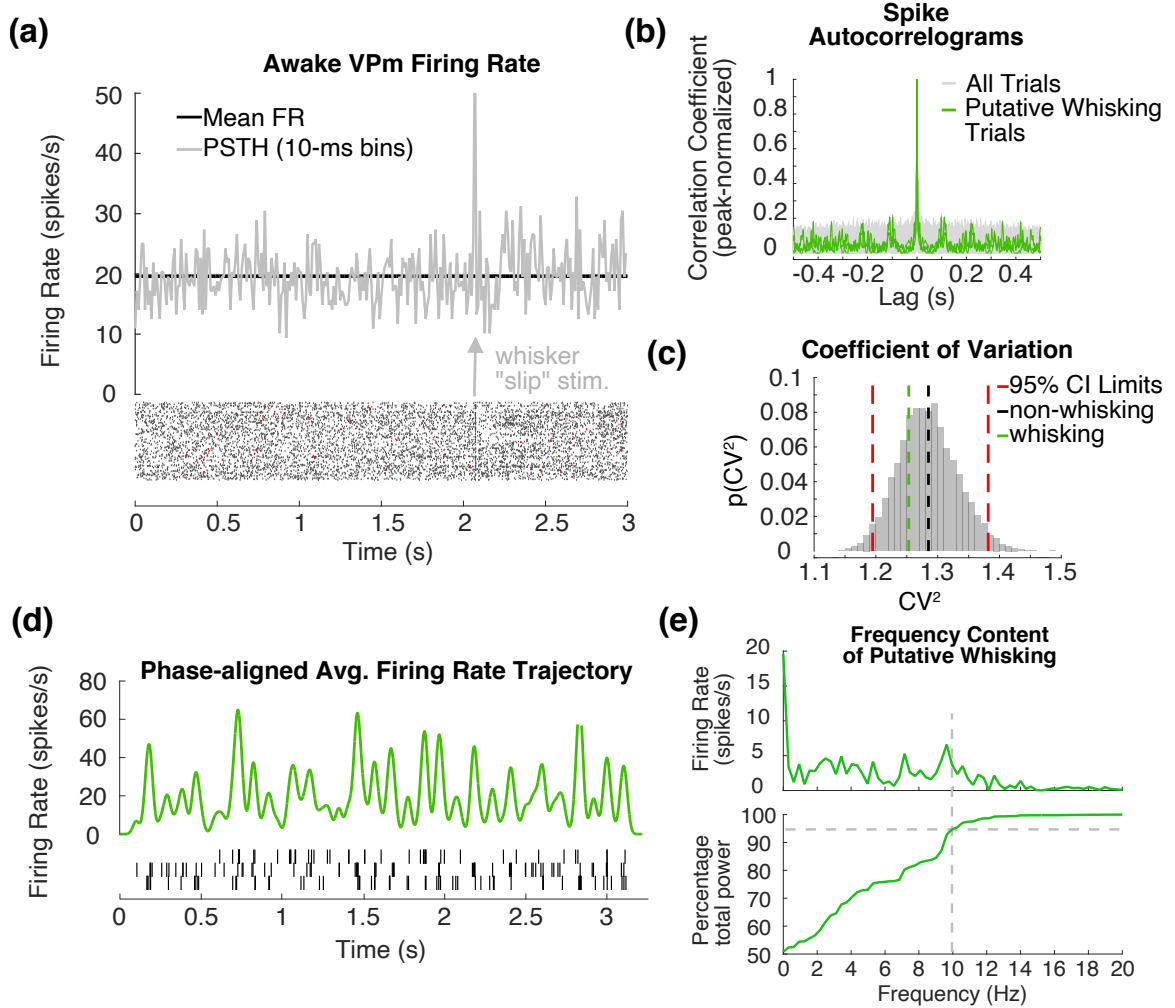


Figure A.1: **Firing rate modulation during putative whisking** (a) Firing of example single-unit recorded in awake rat (data from [145, 16]). Whisker stimulus occurs at arrow. (b) Spike autocorrelograms after smoothing trains with 20 ms s.d. Gaussian window. Putative whisking trials highlighted in green ($n=3$). (c) Bootstrapped distribution of squared coefficient of variation (CV^2) in interspike intervals. (d) Phase-aligned putative whisking spike trains. Firing rate trajectory estimated by smoothing spikes with 20 ms s.d. Gaussian window and averaging across phase-aligned trials. (e) Power spectrum for firing rate trajectory (top). Cumulative power spectrum (bottom). Grey lines indicate 95% of power.

APPENDIX B

INSPECTION OF ELECTRICAL RECORDINGS FOR OPTICAL STIMULATION ARTIFACTS

Feedback control relies on the ability to continuously recording and stimulate. Therefore, the possibility of optical stimulation artifact was was investigated first by pulsing blue light (470 nm) through a 200 μm diameter optic fiber, while recording from an affixed Tungsten electrode (FHC) in saline. Five-millisecond square pulses of light were strobed at 10 Hz. Results of this test in Figure B.1 reveal that while there are small artifacts at low frequencies, after the bandpass filtering that is done before spike thresholding and subsequent sorting, this artifact is negligible.

To further investigate the effects of light on recorded feedback, we analyzed the spike waveforms of two example feedback neurons used in Chapter 2 for any artifactual changes in spike waveform. For both large and small example recordings, extracellular spike waveforms during spontaneous activity versus during optical control epochs did not display any major differences (Figure B.2).

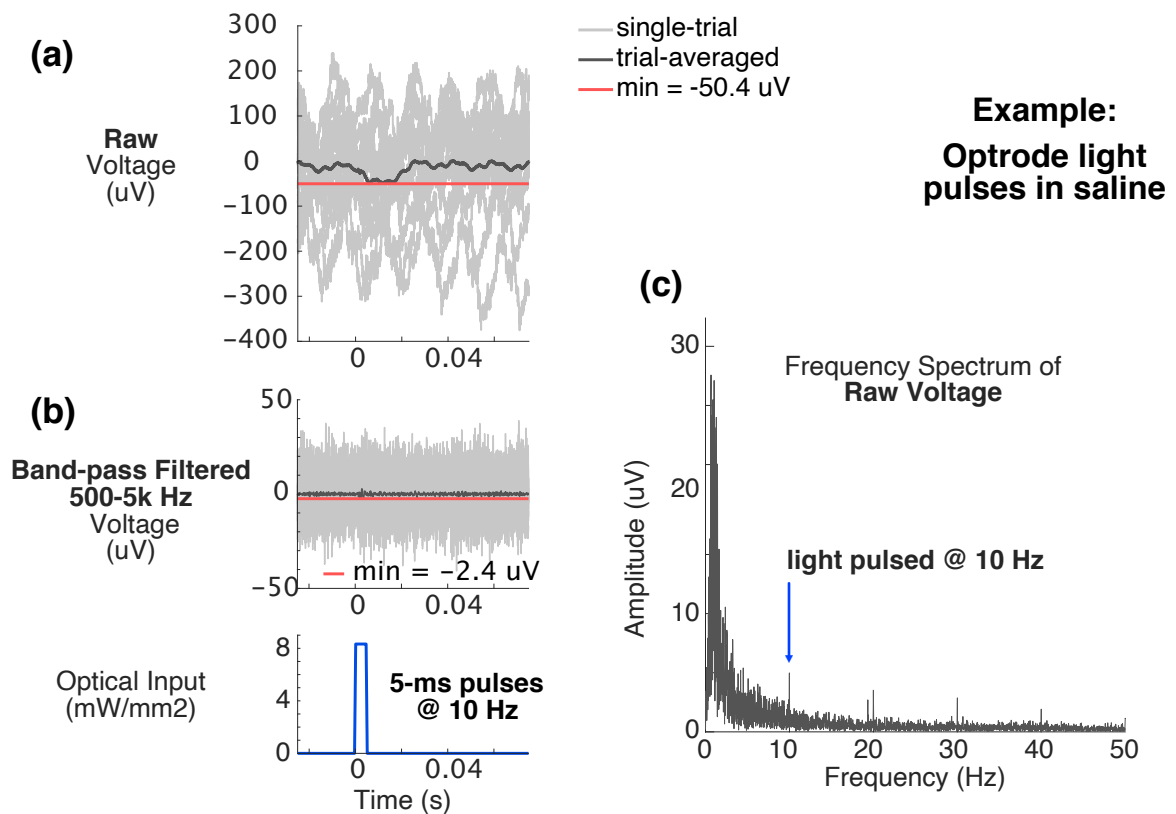


Figure B.1: **Inspection of electrical recordings for optical stimulation artifacts** (a) Raw voltage recordings for Tungsten optrode in saline during 10 Hz square optical pulses (bottom). Average shown in black. Minimum in red. (b) Bandpass filtered version of (a). (c) Power spectrum for raw recordings in (a). Arrow points to tone at 10 Hz due to optical pulsation.

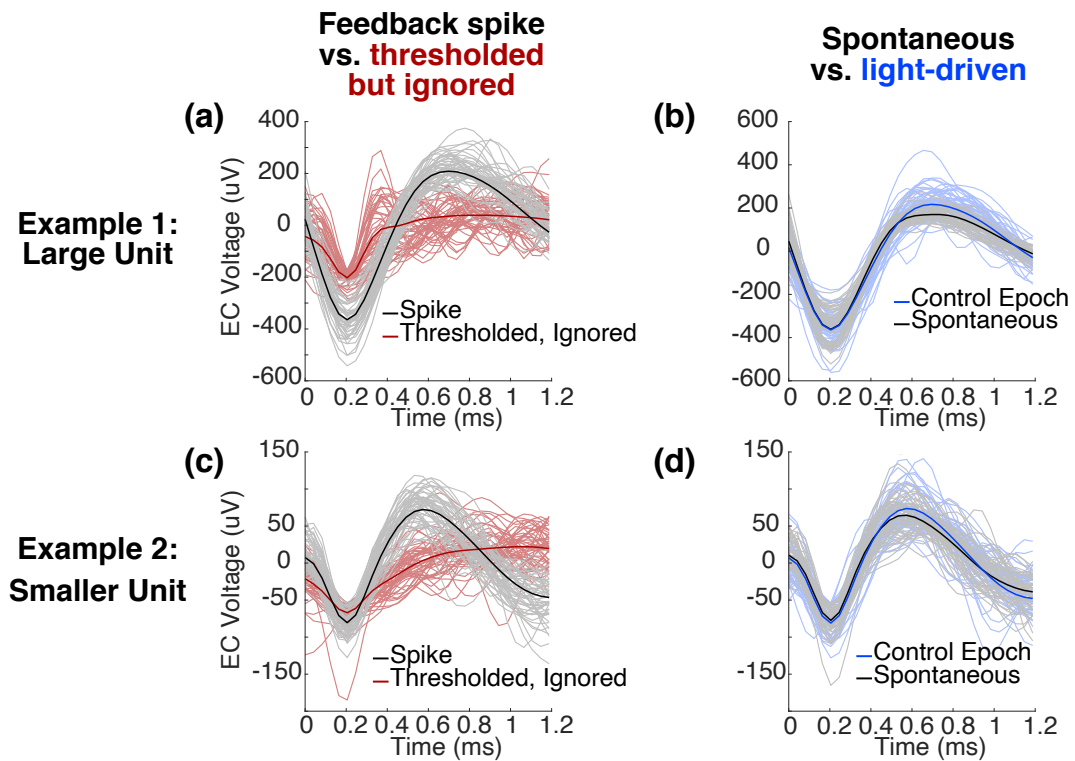


Figure B.2: **Inspection of spiking activity for optical stimulation artifacts** (a) Example sorted feedback unit spikes (black) vs. thresholded but ignored activity (red), recorded using Tungsten electrode in thalamus (Chapter 2). (b) Same unit as (a), spikes during optogenetic control epochs (blue) vs. spontaneous activity (grey). (c-d) Same as (a-b), but smaller unit.

APPENDIX C

TARGETING OPTOGENETICS TO VPM/VPL: AI32;NR133

In a subset of experiments in Chapters 3-4, a transgenic mouse line was used to restrict expression of ChR2 to VPM/VPL thalamus. A previous anatomical study of cre-recombinase drive mouse lines reported that a few transgenic lines displayed selectivity for VPM/VPL thalamus: one example being the NR133 line used here ([123]). A single-generation cross between this driver mouse and a cre-dependent ChR-expressing mouse (Ai32, Jax), allowed selective ChR expression. Figure C.1 provides an example coronal section of one such transgenic cross. The ChR2 is reported using EYFP, and cell nuclei are stained blue using DAPI.




Figure C.1: **Coronal section of Ai32;NR133 mouse.** Histological sample for single-generation Ai32;NR133 mouse cross (Chapters 3-4). ChR2 labeled using EYFP reporter. Nuclei labeled blue with DAPI. Manually overlaid with Paxinos mouse atlas.

REFERENCES

- [1] L. Galvani, “De viribus electricitatis in motu muscularis commentarius,” *De Bononiensi Scientiarum et Artium Instituto atque Academia Commentarii*, 1791.
- [2] H. Y. Zoghbi, “From anatomy to electrophysiology: clinical Lasker goes deep.,” *Cell*, vol. 158, no. 6, pp. 1225–1229, 2014.
- [3] L. Fenno, O. Yizhar, and K. Deisseroth, “The Development and Application of Optogenetics,” *Annual Review of Neuroscience*, vol. 34, no. 1, pp. 389–412, 2011.
- [4] C. Lee, A. Lavoie, J. Liu, S. X. Chen, and B. H. Liu, “Light Up the Brain: The Application of Optogenetics in Cell-Type Specific Dissection of Mouse Brain Circuits,” *Frontiers in Neural Circuits*, vol. 14, no. April, pp. 1–14, 2020.
- [5] L. Grosenick, J. H. Marshel, and K. Deisseroth, “Closed-Loop and Activity-Guided Optogenetic Control,” *Neuron*, vol. 86, no. 1, pp. 106–139, 2015.
- [6] J. P. Newman, M.-f. Fong, D. C. Millard, C. J. Whitmire, G. B. Stanley, and S. M. Potter, “Optogenetic feedback control of neural activity,” *eLife*, vol. 4, no. December, e07192, 2015.
- [7] M. B. Carpenter, *Core Text of Neuroanatomy*. Williams & Wilkins, 1985.
- [8] D. Pinault, “The thalamic reticular nucleus: structure, function and concept.,” *Brain Research Reviews*, vol. 46, pp. 1–31, 2004.
- [9] M. Steriade and I. Timofeev, “Neuronal plasticity in thalamocortical networks during sleep and waking oscillations,” *Neuron*, vol. 37, no. 4, pp. 563–576, 2003.
- [10] R. R. Llina and M. Steriade, “Bursting of Thalamic Neurons and States of Vigilance,” vol. 10016, pp. 3297–3308, 2006.
- [11] W. Guido and T. Weyand, “Burst responses in thalamic relay cells of the awake behaving cat,” *Journal of Neurophysiology*, vol. 74, no. 4, pp. 1782–1786, 1995.
- [12] E. J. Ramcharan, J. W. Gnadt, and S. M. Sherman, “Burst and tonic firing in thalamic cells of unanesthetized, behaving monkeys,” *Visual Neuroscience*, vol. 17, no. 1, pp. 55–62, 2000.

- [13] T. G. Weyand, M. Boudreaux, and W. Guido, “Burst and tonic response modes in thalamic neurons during sleep and wakefulness,” *Journal of Neurophysiology*, vol. 85, no. 3, pp. 1107–1118, 2001.
- [14] S. M. Sherman, “A wake-up call from the thalamus.,” *Nature Neuroscience*, vol. 4, no. 4, pp. 344–346, 2001.
- [15] H. a. Swadlow and A. G. Gusev, “The impact of ‘bursting’ thalamic impulses at a neocortical synapse.,” *Nature Neuroscience*, vol. 4, no. 4, pp. 402–8, 2001.
- [16] C. J. Whitmire, C. Waiblinger, C. Schwarz, and G. B. Stanley, “Information Coding through Adaptive Gating of Synchronized Thalamic Bursting,” *Cell Reports*, vol. 14, no. 4, pp. 795–807, 2016.
- [17] N. A. Lesica and G. B. Stanley, “Encoding of natural scene movies by tonic and burst spikes in the lateral geniculate nucleus.,” *The Journal of Neuroscience*, vol. 24, no. 47, pp. 10 731–40, 2004.
- [18] R. M. Bruno, “Cortex Is Driven by Weak but Synchronously Active Thalamocortical Synapses,” *Science*, vol. 312, no. 5780, pp. 1622–1627, 2006.
- [19] Q. Wang, R. M. Webber, and G. B. Stanley, “Thalamic synchrony and the adaptive gating of information flow to cortex,” *Nature Neuroscience*, vol. 13, no. 12, pp. 1534–1541, 2010.
- [20] M. A. Castro-Alamancos and E. Oldford, “Cortical sensory suppression during arousal is due to the activity-dependent depression of thalamocortical synapses,” *Journal of Physiology*, vol. 541, no. 1, pp. 319–331, 2002.
- [21] D. R. Ollerenshaw, H. J. Zheng, D. C. Millard, Q. Wang, and G. B. Stanley, “The Adaptive Trade-Off between Detection and Discrimination in Cortical Representations and Behavior,” *Neuron*, vol. 81, no. 5, pp. 1152–1164, 2014.
- [22] M. a. Castro-Alamancos, “Role of thalamocortical sensory suppression during arousal: focusing sensory inputs in neocortex,” *The Journal of Neuroscience*, vol. 22, no. 22, pp. 9651–9655, 2002.
- [23] M. E. Diamond, M. von Heimendahl, P. M. Knutsen, D. Kleinfeld, and E. Ahissar, “‘Where’ and ‘what’ in the whisker sensorimotor system.,” *Nature Reviews Neuroscience*, vol. 9, no. 8, pp. 601–12, 2008.
- [24] C. C. Petersen, “The functional organization of the barrel cortex.,” *Neuron*, vol. 56, no. 2, pp. 339–355, 2007.

- [25] D. Feldmeyer, M. Brecht, F. Helmchen, C. C. H. Petersen, J. F. a. Poulet, J. F. Staiger, H. J. Luhmann, and C. Schwarz, “Barrel cortex function.,” *Progress in Neurobiology*, vol. 103, pp. 3–27, 2013.
- [26] E. Ahlssar, R. Sosnik, and S. Haldarilu, “Transformation from temporal to rate coding in a somatosensory thalamocortical pathway,” *Nature*, vol. 406, no. 6793, pp. 302–306, 2000.
- [27] N. Urbain, P. A. Salin, P.-a. Libourel, C. C. H. Petersen, N. Urbain, P. A. Salin, P.-a. Libourel, J.-c. Comte, L. J. Gentet, and C. C. H. Petersen, “Whisking-Related Changes in Neuronal Firing and Membrane Potential Dynamics in the Somatosensory Thalamus of Awake Mice,” *Cell Reports*, vol. 13, pp. 1–10, 2015.
- [28] K. J. Åström and R. M. Murray, *Feedback systems: An introduction for scientists and engineers*. Princeton: Princeton University Press, 2008, ISBN: 9781400828739.
- [29] K. D. Harris and A. Thiele, “Cortical state and attention.,” *Nature Reviews Neuroscience*, vol. 12, no. 9, pp. 509–23, 2011.
- [30] E. Zagha and D. A. McCormick, “Neural control of brain state,” *Current Opinion in Neurobiology*, vol. 29, no. 1, pp. 178–186, 2014.
- [31] D. A. McCormick, M. J. McGinley, and D. B. Salkoff, “Brain state dependent activity in the cortex and thalamus,” *Current Opinion in Neurobiology*, vol. 31, pp. 133–140, 2015.
- [32] J. F. Poulet and S. Crochet, “The cortical states of wakefulness,” *Frontiers in Systems Neuroscience*, vol. 12, no. January, pp. 1–18, 2019.
- [33] S. Crochet and C. C. Petersen, “Correlating whisker behavior with membrane potential in barrel cortex of awake mice,” *Nature Neuroscience*, vol. 9, no. 5, pp. 608–610, 2006.
- [34] J. F. A. Poulet and C. C. H. Petersen, “Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice,” *Nature*, vol. 454, no. 7206, pp. 881–5, 2008.
- [35] J. F. Poulet, L. M. J. Fernandez, S. Crochet, and C. C. H. Petersen, “Thalamic control of cortical states.,” *Nature Neuroscience*, vol. 15, no. 3, pp. 370–2, 2012.
- [36] T. Bezdudnaya, M. Cano, Y. Bereshpolova, C. R. Stoelzel, J. M. Alonso, and H. A. Swadlow, “Thalamic burst mode and inattention in the awake LGNd,” *Neuron*, vol. 49, no. 3, pp. 421–432, 2006.

- [37] G. H. Otazu, L. H. Tai, Y. Yang, and A. M. Zador, “Engaging in an auditory task suppresses responses in auditory cortex,” *Nature Neuroscience*, vol. 12, no. 5, pp. 646–654, 2009.
- [38] M. J. McGinley, S. V. David, and D. A. McCormick, “Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection,” *Neuron*, vol. 87, no. 1, pp. 179–192, 2015.
- [39] J. Reimer, M. J. McGinley, Y. Liu, C. Rodenkirch, Q. Wang, D. A. McCormick, and A. S. Tolias, “Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex,” *Nature Communications*, vol. 7, no. May, pp. 1–7, 2016.
- [40] C. R. Lee and D. J. Margolis, “Pupil dynamics reflect behavioral choice and learning in a Go/NoGo tactile decision-making task in mice,” *Frontiers in Behavioral Neuroscience*, vol. 10, no. NOV, pp. 1–14, 2016.
- [41] C. Bennett, S. Arroyo, and S. Hestrin, “Subthreshold mechanisms underlying state-dependent modulation of visual responses,” *Neuron*, vol. 80, no. 2, pp. 350–357, 2013.
- [42] B. Haider, A. Duque, A. R. Hasenstaub, Y. Yu, and D. A. McCormick, “Enhancement of visual responsiveness by spontaneous local network activity in vivo,” *Journal of Neurophysiology*, vol. 97, no. 6, pp. 4186–4202, 2007.
- [43] P. O. Polack, J. Friedman, and P. Golshani, “Cellular mechanisms of brain state-dependent gain modulation in visual cortex,” *Nature Neuroscience*, vol. 16, no. 9, pp. 1331–1339, 2013.
- [44] M. L. Scholvinck, A. B. Saleem, A. Benucci, K. D. Harris, and M. Carandini, “Cortical State Determines Global Variability and Correlations in Visual Cortex,” *Journal of Neuroscience*, vol. 35, no. 1, pp. 170–178, 2015.
- [45] J. R. Aguilar and M. a. Castro-Alamancos, “Spatiotemporal gating of sensory inputs in thalamus during quiescent and activated states,” *Journal of Neuroscience*, vol. 25, no. 47, pp. 10 990–1002, 2005.
- [46] A. Hasenstaub, R. N. S. Sachdev, and D. A. McCormick, “State Changes Rapidly Modulate Cortical Neuronal Responsiveness,” *Journal of Neuroscience*, vol. 27, no. 36, pp. 9607–9622, 2007.
- [47] I. Ferezou, F. Haiss, L. J. Gentet, R. Aronoff, B. Weber, and C. C. H. Petersen, “Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice,” *Neuron*, vol. 56, no. 5, pp. 907–23, 2007.

- [48] A. Hirata and M. A. Castro-Alamancos, “Effects of cortical activation on sensory responses in barrel cortex,” *Journal of Neurophysiology*, vol. 105, no. 4, pp. 1495–1505, 2011.
- [49] D. R. Ollerenshaw, B. a. Bari, D. C. Millard, L. E. Orr, Q. Wang, and G. B. Stanley, “Detection of tactile inputs in the rat vibrissa pathway,” *Journal of Neurophysiology*, vol. 108, no. 2, pp. 479–90, 2012.
- [50] I. Ferezou, S. Bolea, and C. C. H. Petersen, “Visualizing the Cortical Representation of Whisker Touch: Voltage-Sensitive Dye Imaging in Freely Moving Mice,” *Neuron*, vol. 50, no. 4, pp. 617–629, 2006.
- [51] A. Hirata and M. A. Castro-alamancos, “Neocortex Network Activation and Deactivation States Controlled by the Thalamus,” *Journal of Neurophysiology*, vol. 103, pp. 1147–1157, 2010.
- [52] E. Eggermann, Y. Kremer, S. Crochet, and C. C. Petersen, “Cholinergic Signals in Mouse Barrel Cortex during Active Whisker Sensing,” *Cell Reports*, vol. 9, no. 5, pp. 1654–1660, 2014.
- [53] D. W. Godwin, J. W. Vaughan, and S. M. Sherman, “Metabotropic glutamate receptors switch visual response mode of lateral geniculate nucleus cells from burst to tonic,” *Journal of Neurophysiology*, vol. 76, no. 3, pp. 1800–1816, 1996.
- [54] R. a. Mease, P. Krieger, and A. Groh, “Cortical control of adaptation and sensory relay mode in the thalamus,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 111, no. 18, pp. 6798–803, 2014.
- [55] S. R. Crandall, S. J. Cruikshank, and B. W. Connors, “A Corticothalamic Switch: Controlling the Thalamus with Dynamic Synapses,” *Neuron*, vol. 86, no. 3, pp. 768–782, 2015.
- [56] J. T. Paz and J. R. Huguenard, “Microcircuits and their interactions in epilepsy: is the focus out of focus?” *Nature Neuroscience*, vol. 18, no. 3, pp. 351–359, 2015.
- [57] J. T. Paz, T. J. Davidson, E. S. Frechette, B. Delord, I. Parada, K. Peng, K. Deisseroth, and J. R. Huguenard, “Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury,” *Nature Neuroscience*, vol. 16, no. 1, pp. 64–70, 2013.
- [58] J. M. Sorokin, T. J. Davidson, E. Frechette, A. M. Abramian, K. Deisseroth, J. R. Huguenard, and J. T. Paz, “Bidirectional Control of Generalized Epilepsy Networks via Rapid Real-Time Switching of Firing Mode,” *Neuron*, vol. 93, no. 1, pp. 194–210, 2017.

- [59] C. McCafferty, F. David, M. Venzi, M. L. Lorincz, F. Delicata, Z. Atherton, G. Recchia, G. Orban, R. C. Lambert, G. Di Giovanni, N. Leresche, and V. Crunelli, "Cortical drive and thalamic feed-forward inhibition control thalamic output synchrony during absence seizures," *Nature Neuroscience*, vol. 21, no. 5, pp. 744–756, 2018.
- [60] C. D. Makinson, B. S. Tanaka, J. M. Sorokin, J. C. Wong, C. A. Christian, A. L. Goldin, A. Escayg, and J. R. Huguenard, "Regulation of Thalamic and Cortical Network Synchrony by Scn8a," *Neuron*, vol. 93, no. 5, 1165–1179.e6, 2017.
- [61] A. L. Hodgkin and A. F. Huxley, "A Quantitative Description of Membrane Current and its Application to Conduction and Excitation in Nerves," *J. Physiol.*, vol. 117, pp. 500–544, 1952.
- [62] A. Pala and C. C. Petersen, "In Vivo Measurement of Cell-Type-Specific Synaptic Connectivity and Synaptic Transmission in Layer 2/3 Mouse Barrel Cortex," *Neuron*, vol. 85, no. 1, pp. 68–75, 2015.
- [63] J. J. Jun, N. A. Steinmetz, J. H. Siegle, D. J. Denman, M. Bauza, B. Barbarits, A. K. Lee, C. A. Anastassiou, A. Andrei, Ç. Aydin, M. Barbic, T. J. Blanche, V. Bonin, J. Couto, B. Dutta, S. L. Gratiy, D. A. Gutnisky, M. Häusser, B. Karsh, P. Ledochowitsch, C. M. Lopez, C. Mitelut, S. Musa, M. Okun, M. Pachitariu, J. Putzeys, P. D. Rich, C. Rossant, W. L. Sun, K. Svoboda, M. Carandini, K. D. Harris, C. Koch, J. O’Keefe, and T. D. Harris, "Fully integrated silicon probes for high-density recording of neural activity," *Nature*, vol. 551, no. 7679, pp. 232–236, 2017.
- [64] E. M. Trautmann, S. D. Stavisky, S. Lahiri, K. C. Ames, M. T. Kaufman, D. J. O’Shea, S. Vyas, X. Sun, S. I. Ryu, S. Ganguli, and K. V. Shenoy, "Accurate Estimation of Neural Population Dynamics without Spike Sorting," *Neuron*, vol. 103, no. 2, 292–308.e4, 2019.
- [65] A. Arieli, A. Sterkin, A. Grinvald, and A. Aertsen, "Dynamics of Ongoing Activity : Explanation of the Large Variability in Evoked Cortical Responses," vol. 273, no. 5283, pp. 1868–1871, 1996.
- [66] H. Mutoh and T. Knöpfel, "Probing neuronal activities with genetically encoded optical indicators: from a historical to a forward-looking perspective.," *Pflügers Archiv : European Journal of Physiology*, vol. 465, no. 3, pp. 361–71, 2013.
- [67] P. Y. Borden, A. D. Ortiz, C. Waiblinger, A. J. Sederberg, A. E. Morrisette, C. R. Forest, D. Jaeger, and G. B. Stanley, "Genetically expressed voltage sensor ArcLight for imaging large scale cortical activity in the anesthetized and awake mouse," *Neurophotonics*, vol. 4, no. 3, p. 031 212, 2017.

- [68] T. W. Chen, N. Li, K. Daie, and K. Svoboda, “A Map of Anticipatory Activity in Mouse Motor Cortex,” *Neuron*, vol. 94, no. 4, 866–879.e4, 2017.
- [69] Z. Zhang, L. E. Russell, A. M. Packer, O. M. Gauld, and M. Häusser, “Closed-loop all-optical interrogation of neural circuits in vivo,” *Nature Methods*, vol. 15, no. 12, pp. 1037–1040, 2018.
- [70] H. Dana, Y. Sun, B. Mohar, B. K. Hulse, A. M. Kerlin, J. P. Hasseman, G. Tsegaye, A. Tsang, A. Wong, R. Patel, J. J. Macklin, Y. Chen, A. Konnerth, V. Jayaraman, L. L. Looger, E. R. Schreier, K. Svoboda, and D. S. Kim, “High-performance calcium sensors for imaging activity in neuronal populations and microcompartments,” *Nature Methods*, vol. 16, no. 7, pp. 649–657, 2019.
- [71] B. Rosin, M. Slovik, R. Mitelman, M. Rivlin-Etzion, S. N. Haber, Z. Israel, E. Vaadia, and H. Bergman, “Closed-loop deep brain stimulation is superior in ameliorating parkinsonism,” *Neuron*, vol. 72, no. 2, pp. 370–384, 2011.
- [72] D. C. Millard, C. J. Whitmire, C. A. Gollnick, C. J. Rozell, and G. B. Stanley, “Electrical and Optical Activation of Mesoscale Neural Circuits with Implications for Coding,” *Journal of Neuroscience*, vol. 35, no. 47, pp. 15 702–15 715, 2015.
- [73] J. S. Choi, A. J. Brockmeier, D. B. McNiel, L. M. Von Kraus, J. C. Príncipe, and J. T. Francis, “Eliciting naturalistic cortical responses with a sensory prosthesis via optimized microstimulation,” *Journal of Neural Engineering*, vol. 13, no. 5, 2016.
- [74] E. S. Boyden, F. Zhang, E. Bamberg, G. Nagel, and K. Deisseroth, “Millisecond-timescale, genetically targeted optical control of neural activity,” *Nature Neuroscience*, vol. 8, no. 9, pp. 1263–1268, 2005.
- [75] F. Zhang, L. P. Wang, M. Brauner, J. F. Liewald, K. Kay, N. Watzke, P. G. Wood, E. Bamberg, G. Nagel, A. Gottschalk, and K. Deisseroth, “Multimodal fast optical interrogation of neural circuitry,” *Nature*, vol. 446, no. 7136, pp. 633–639, 2007.
- [76] B. Y. Chow, X. Han, A. S. Dobry, X. Qian, A. S. Chuong, M. Li, M. A. Henninger, G. M. Belfort, Y. Lin, P. E. Monahan, and E. S. Boyden, “High-performance genetically targetable optical neural silencing by light-driven proton pumps,” *Nature*, vol. 463, no. 7277, pp. 98–102, 2010.
- [77] C. Enroth-Cugell and J. G. Robson, “The contrast sensitivity of retinal ganglion cells of the cat,” *The Journal of Physiology*, 1966.
- [78] J. A. Movshon, I. D. Thompson, and D. J. Tolhurst, “Spatial summation in the receptive fields of simple cells in the cat’s striate cortex,” *The Journal of Physiology*, 1978.

- [79] K. H. Britten and H. W. Heuer, “Spatial summation in the receptive fields of MT neurons,” *Journal of Neuroscience*, vol. 19, no. 12, pp. 5074–5084, 1999.
- [80] M. Park and J. W. Pillow, “Receptive field inference with localized priors,” *PLoS Computational Biology*, vol. 7, no. 10, 2011.
- [81] E. Chichilnisky, “A simple white noise analysis of neuronal light responses,” *Network: Computation in Neural Systems*, vol. 12, pp. 199–213, 2001.
- [82] O. Schwartz, J. W. Pillow, N. C. Rust, and E. P. Simoncelli, “Spike-triggered neural characterization,” *Journal of Vision*, vol. 6, no. 4, pp. 484–507, 2006.
- [83] M. Park and J. Pillow, “Bayesian Spike-Triggered Covariance Analysis,” *Advances in Neural Information Processing Systems*, pp. 1692–1700, 2011.
- [84] D. Chander and E. J. Chichilnisky, “Adaptation to temporal contrast in primate and salamander retina,” *Journal of Neuroscience*, 2001.
- [85] L. Paninski, “Convergence properties of three spike-triggered analysis techniques,” *Network: Computation in Neural Systems*, vol. 14, no. 3, pp. 437–464, 2003.
- [86] N. A. Lesica, J. Jin, C. Weng, C. I. Yeh, D. A. Butts, G. B. Stanley, and J. M. Alonso, “Adaptation to Stimulus Contrast and Correlations during Natural Visual Stimulation,” *Neuron*, vol. 55, no. 3, pp. 479–491, 2007.
- [87] L. Estebanez, S. E. Boustani, A. Destexhe, and D. E. Shulz, “Correlated input reveals coexisting coding schemes in a sensory cortex,” *Nature Neuroscience*, vol. 15, no. 12, pp. 1691–1699, 2012.
- [88] M. Maravall, A. Alenda, M. R. Bale, and R. S. Petersen, “Transformation of adaptation and gain rescaling along the whisker sensory pathway,” *PLoS ONE*, vol. 8, no. 12, 2013.
- [89] M. Carandini, J. B. Demb, V. Mante, D. J. Tolhurst, Y. Dan, B. A. Olshausen, J. L. Gallant, and N. C. Rust, “Do we know what the early visual system does?” *Journal of Neuroscience*, vol. 25, no. 46, pp. 10 577–10 597, 2005.
- [90] L. Paninski, J. Pillow, and J. Lewi, “Statistical models for neural encoding, decoding, and optimal stimulus design,” *Progress in Brain Research*, vol. 165, pp. 493–507, 2007.
- [91] J. W. Pillow, J. Shlens, L. Paninski, A. Sher, A. M. Litke, E. J. Chichilnisky, and E. P. Simoncelli, “Spatio-temporal correlations and visual signalling in a complete neuronal population,” *Nature*, vol. 454, no. 7207, pp. 995–999, 2008.

- [92] A. Nandi, M. Kafashan, and S. Ching, “Controlling Point Process Generalized Linear Models of Neural Spiking,” in *2016 American Control Conference*, 2016, pp. 5779–5784, ISBN: 9781467386814.
- [93] J. W. Pillow, “Prediction and Decoding of Retinal Ganglion Cell Responses with a Probabilistic Spiking Model,” *Journal of Neuroscience*, vol. 25, no. 47, pp. 11 003–11 013, 2005.
- [94] L. Paninski, “Maximum likelihood estimation of cascade point-process neural encoding models,” *Network: Computation in Neural Systems*, vol. 15, no. 4, pp. 243–262, 2004.
- [95] S. Ching and J. Ritt, “Control strategies for underactuated neural ensembles driven by optogenetic stimulation,” *Frontiers in Neural Circuits*, vol. 7, p. 54, 2013.
- [96] E. N. Brown, D. P. Nguyen, L. M. Frank, M. A. Wilson, and V. Solo, “An analysis of neural receptive field plasticity by point process adaptive filtering,” *Proceedings of the National Academy of Sciences*, vol. 98, no. 21, pp. 12 261–12 266, 2001.
- [97] A. C. Smith and E. N. Brown, “Estimating a State-Space Model from Point Process Observations,” *Neural Computation*, vol. 15, pp. 965–991, 2003.
- [98] U. T. Eden, L. M. Frank, R. Barbieri, V. Solo, and E. N. Brown, “Dynamic analysis of neural encoding by point process adaptive filtering,” *Neural Computation*, vol. 16, no. 5, pp. 971–98, 2004.
- [99] L. Paninski, Y. Ahmadian, D. G. Ferreira, S. Koyama, K. Rahnema Rad, M. Vidne, J. T. Vogelstein, and W. Wu, “A new look at state-space models for neural data,” *Journal of Computational Neuroscience*, vol. 29, no. 1-2, pp. 107–126, 2009.
- [100] B. M. Yu, J. P. Cunningham, G. Santhanam, S. I. Ryu, K. V. Shenoy, and M. Sahani, “Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity,” *Journal of Neurophysiology*, vol. 102, no. 3, p. 2008, 2009.
- [101] J. H. Macke, L. Buesing, J. P. Cunningham, B. M. Yu, K. V. Shenoy, and M. Sahani, “Empirical models of spiking in neural populations,” in *Advances in Neural Information Processing Systems 24*, 2011, pp. 1350–1358.
- [102] L. Buesing, J. H. Macke, and M. Sahani, “Spectral learning of linear dynamics from generalised-linear observations with application to neural population data,” *Advances in Neural Information Processing Systems (NIPS)*, pp. 1–9, 2012.
- [103] V. Gilja, P. Nuyujukian, C. A. Chestek, J. P. Cunningham, B. M. Yu, J. M. Fan, M. M. Churchland, M. T. Kaufman, J. C. Kao, S. I. Ryu, and K. V. Shenoy, “A

- high-performance neural prosthesis enabled by control algorithm design.,” *Nature Neuroscience*, vol. 15, no. 12, pp. 1752–7, 2012.
- [104] K. V. Shenoy, M. Sahani, and M. M. Churchland, “Cortical Control of Arm Movements: A Dynamical Systems Perspective,” *Annual Review of Neuroscience*, vol. 36, no. 1, pp. 337–359, 2013.
 - [105] M. M. Shanechi, G. W. Wornell, Z. M. Williams, and E. N. Brown, “Feedback-controlled parallel point process filter for estimation of goal-directed movements from neural signals,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 21, no. 1, pp. 129–140, 2013.
 - [106] C. Pandarinath, D. J. O’Shea, J. Collins, R. Jozefowicz, S. D. Stavisky, J. C. Kao, E. M. Trautmann, M. T. Kaufman, S. I. Ryu, L. R. Hochberg, J. M. Henderson, K. V. Shenoy, L. F. Abbott, and D. Sussillo, “Inferring single-trial neural population dynamics using sequential auto-encoders,” *Nature Methods*, vol. 15, no. 10, pp. 805–815, 2018.
 - [107] Z. Wei, H. Inagaki, N. Li, K. Svoboda, and S. Druckmann, “An orderly single-trial organization of population dynamics in premotor cortex predicts behavioral variability,” *Nature Communications*, vol. 10, no. 1, pp. 1–14, 2019.
 - [108] R. E. Kalman, “A New Approach to Linear Filtering and Prediction Problems,” *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
 - [109] L. Srinivasan, U. T. Eden, A. S. Willsky, and E. N. Brown, “A state-space analysis for reconstruction of goal-directed movements using neural signals,” *Neural Computation*, vol. 18, no. 10, pp. 2465–2494, 2006.
 - [110] M. M. Shanechi, G. W. Wornell, Z. Williams, and E. N. Brown, “A parallel point-process filter for estimation of goal-directed movements from neural signals,” in *IEEE*, 2010, pp. 521–524.
 - [111] F. Zhang, V. Gradinaru, A. R. Adamantidis, R. Durand, R. D. Airan, L. de Lecea, and K. Deisseroth, “Optogenetic interrogation of neural circuits: technology for probing mammalian brain structures.,” *Nature protocols*, vol. 5, no. 3, pp. 439–56, 2010.
 - [112] J. A. Cardin, M. Carlén, K. Meletis, U. Knoblich, F. Zhang, K. Deisseroth, L.-H. Tsai, and C. I. Moore, “Targeted optogenetic stimulation and recording of neurons in vivo using cell-type-specific expression of Channelrhodopsin-2.,” *Nature Protocols*, vol. 5, no. 2, pp. 247–54, 2010.
 - [113] L. Madisen, T. Mao, H. Koch, J. M. Zhuo, A. Berenyi, S. Fujisawa, Y. W. A. Hsu, A. J. Garcia, X. Gu, S. Zanella, J. Kidney, H. Gu, Y. Mao, B. M. Hooks, E. S.

- Boyden, G. Buzsáki, J. M. Ramirez, A. R. Jones, K. Svoboda, X. Han, E. E. Turner, and H. Zeng, “A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing,” *Nature Neuroscience*, vol. 15, no. 5, pp. 793–802, 2012.
- [114] A. M. Packer, B. Roska, and M. Häusser, “Targeting neurons and photons for optogenetics,” *Nature Neuroscience*, vol. 16, no. 7, pp. 805–15, 2013.
- [115] A. Galvan, W. R. Stauffer, L. Acker, Y. El-Shamayleh, K. I. Inoue, S. Ohayon, and M. C. Schmid, “Nonhuman primate optogenetics: Recent advances and future directions,” *Journal of Neuroscience*, vol. 37, no. 45, pp. 10 894–10 903, 2017.
- [116] O. Yizhar, L. E. Fenno, M. Prigge, F. Schneider, T. J. Davidson, D. J. Ogshea, V. S. Sohal, I. Goshen, J. Finkelstein, J. T. Paz, K. Stehfest, R. Fudim, C. Ramakrishnan, J. R. Huguenard, P. Hegemann, and K. Deisseroth, “Neocortical excitation/inhibition balance in information processing and social dysfunction,” *Nature*, vol. 477, no. 7363, pp. 171–178, 2011.
- [117] J. Y. Lin, P. M. Knutsen, A. Muller, D. Kleinfeld, and R. Y. Tsien, “ReaChR: a red-shifted variant of channelrhodopsin enables deep transcranial optogenetic excitation,” *Nature Neuroscience*, vol. 16, no. 10, pp. 1499–508, 2013.
- [118] V. Gradinaru, F. Zhang, C. Ramakrishnan, J. Mattis, R. Prakash, I. Diester, I. Goshen, K. R. Thompson, and K. Deisseroth, “Molecular and Cellular Approaches for Diversifying and Extending Optogenetics,” *Cell*, vol. 141, no. 1, pp. 154–165, 2010.
- [119] M. Carus-Cadavieco, M. Gorbati, L. Ye, F. Bender, S. Van Der Veldt, C. Kosse, C. Börgers, S. Y. Lee, C. Ramakrishnan, Y. Hu, N. Denisova, F. Ramm, E. Volitaki, D. Burdakov, K. Deisseroth, A. Ponomarenko, and T. Korotkova, “Gamma oscillations organize top-down signalling to hypothalamus and enable food seeking,” *Nature*, vol. 542, no. 7640, pp. 232–236, 2017.
- [120] G. Nagel, M. Brauner, J. F. Liewald, N. Adeishvili, E. Bamberg, and A. Gottschalk, “Light activation of Channelrhodopsin-2 in excitable cells of *Caenorhabditis elegans* triggers rapid behavioral responses,” *Current Biology*, vol. 15, no. 24, pp. 2279–2284, 2005.
- [121] V. Gradinaru, K. R. Thompson, F. Zhang, M. Mogri, K. Kay, M. B. Schneider, and K. Deisseroth, “Targeting and readout strategies for fast optical neural control in vitro and in vivo,” *The Journal of Neuroscience*, vol. 27, no. 52, pp. 14 231–8, 2007.
- [122] J. A. Cardin, M. Carlén, K. Meletis, U. Knoblich, F. Zhang, K. Deisseroth, L.-H. Tsai, and C. I. Moore, “Driving fast-spiking cells induces gamma rhythm and controls sensory responses,” *Nature*, vol. 459, no. 7247, pp. 663–7, 2009.

- [123] C. R. Gerfen, R. Paletzki, and N. Heintz, “GENSAT BAC cre-recombinase driver lines to study the functional organization of cerebral cortical and basal ganglia circuits,” *Neuron*, vol. 80, no. 6, pp. 1368–1383, 2013.
- [124] R. D. Wimmer, L. I. Schmitt, T. J. Davidson, M. Nakajima, K. Deisseroth, and M. M. Halassa, “Thalamic control of sensory selection in divided attention,” *Nature*, vol. 526, no. 7575, pp. 705–709, 2015.
- [125] A. Galvan, X. Hu, Y. Smith, and T. Wichmann, “Effects of optogenetic activation of corticothalamic terminals in the motor thalamus of awake monkeys,” *Journal of Neuroscience*, vol. 36, no. 12, pp. 3519–3530, 2016.
- [126] M. F. Bolus, A. A. Willats, C. J. Whitmire, C. J. Rozell, and G. B. Stanley, “Design strategies for dynamic closed-loop optogenetic neurocontrol in vivo,” *Journal of Neural Engineering*, vol. 15, no. 2, 2018.
- [127] E. Ronzitti, R. Conti, V. Zampini, D. Tanese, A. Foust, N. Klapoetke, E. Boyden, E. Papagiakoumou, and V. Emiliani, “Sub-millisecond optogenetic control of neuronal firing with two-photon holographic photoactivation of Chronos,” *The Journal of Neuroscience*, pp. 1246–17, 2017.
- [128] K. Kampasi, E. Stark, J. Seymour, K. Na, H. G. Winful, G. Buzsáki, K. D. Wise, and E. Yoon, “Fiberless multicolor neural optoelectrode for in vivo circuit analysis,” *Scientific Reports*, vol. 6, no. August, pp. 1–13, 2016.
- [129] K. Kampasi, D. F. English, J. Seymour, E. Stark, S. McKenzie, M. Vöröslakos, G. Buzsáki, K. D. Wise, and E. Yoon, “Dual color optogenetic control of neural populations using low-noise, multishank optoelectrodes,” *Microsystems and Nano-engineering*, vol. 4, no. 1, 2018.
- [130] A. E. Mendrela, K. Kim, D. English, S. McKenzie, J. P. Seymour, G. Buzsaki, and E. Yoon, “A High-Resolution Opto-Electrophysiology System with a Miniature Integrated Headstage,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 12, no. 5, pp. 1065–1075, 2018.
- [131] E. Krook-Magnuson, C. Armstrong, M. Oijala, and I. Soltesz, “On-demand optogenetic control of spontaneous seizures in temporal lobe epilepsy,” *Nature Communications*, vol. 4, p. 1376, 2013.
- [132] D. H. O’Connor, S. A. Hires, Z. V. Guo, N. Li, J. Yu, Q.-Q. Sun, D. Huber, and K. Svoboda, “Neural coding during active somatosensation revealed using illusory touch,” *Nature Neuroscience*, vol. 16, no. 7, pp. 958–65, 2013.
- [133] C. F. V. Latchoumane, H. V. V. Ngo, J. Born, and H. S. Shin, “Thalamic Spindles Promote Memory Formation during Sleep through Triple Phase-Locking of Cor-

- tical, Thalamic, and Hippocampal Rhythms,” *Neuron*, vol. 95, no. 2, 424–435.e6, 2017.
- [134] S. S. Srinivasan, B. E. Maimon, M. Diaz, H. Song, and H. M. Herr, “Closed-loop functional optogenetic stimulation,” *Nature Communications*, vol. 9, no. 1, pp. 1–10, 2018.
 - [135] G. Ullah and S. J. Schiff, “Tracking and control of neuronal Hodgkin-Huxley dynamics,” *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, vol. 79, no. 4, pp. 1–4, 2009.
 - [136] Y. Ahmadian, A. M. Packer, R. Yuste, and L. Paninski, “Designing optimal stimuli to control neuronal spike timing,” *Journal of Neurophysiology*, vol. 106, no. 2, pp. 1038–1053, 2011.
 - [137] A. Iolov, S. Ditlevsen, and A. Longtin, “Stochastic optimal control of single neuron spike trains,” *Journal of Neural Engineering*, vol. 11, no. 4, 2014.
 - [138] M. F. Fong, J. P. Newman, S. M. Potter, and P. Wenner, “Upward synaptic scaling is dependent on neurotransmission rather than spiking,” *Nature Communications*, vol. 6, pp. 1–11, 2015.
 - [139] A. Berenyi, M. Belluscio, D. Mao, and G. Buzsáki, “Closed-Loop Control of Epilepsy by Transcranial Electrical Stimulation,” *Science*, vol. 337, no. 6095, pp. 735–737, 2012.
 - [140] G. Kozák and A. Berényi, “Sustained efficacy of closed loop electrical stimulation for long-term treatment of absence epilepsy in rats,” *Scientific Reports*, vol. 7, no. 1, p. 6300, 2017.
 - [141] D. Ehrens, D. Sritharan, and S. V. Sarma, “Closed-loop control of a fragile network: Application to seizure-like dynamics of an epilepsy model,” *Frontiers in Neuroscience*, vol. 9, no. MAR, pp. 1–9, 2015.
 - [142] S. Santaniello, G. Fiengo, L. Glielmo, and W. M. Grill, “Closed-loop control of deep brain stimulation: a simulation study,” *IEEE Trans Neural Syst Rehabil Eng*, vol. 19, no. 1, pp. 15–24, 2011.
 - [143] E. Stark, L. Roux, R. Eichler, Y. Senzai, S. Royer, and G. Buzsáki, “Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations,” *Neuron*, vol. 83, no. 2, pp. 467–480, 2014.
 - [144] C. J. Whitmire, D. C. Millard, and G. B. Stanley, “Thalamic state control of cortical paired-pulse dynamics,” *Journal of Neurophysiology*, vol. 117, no. 1, pp. 163–177, 2017.

- [145] C. Waiblinger, C. J. Whitmire, A. Sederberg, G. B. Stanley, and C. Schwarz, “Primary Tactile Thalamus Spiking Reflects Cognitive Signals,” *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 38, no. 21, pp. 4870–4885, 2018.
- [146] D. Kleinfeld, E. Ahissar, and M. E. Diamond, “Active sensation: insights from the rodent vibrissa sensorimotor system,” *Current Opinion in Neurobiology*, vol. 16, no. 4, pp. 435–444, 2006.
- [147] D. N. Hill, R. Bermejo, H. P. Zeigler, and D. Kleinfeld, “Biomechanics of the vibrissa motor plant in rat: Rhythmic whisking consists of triphasic neuromuscular activity,” *Journal of Neuroscience*, vol. 28, no. 13, pp. 3438–3455, 2008.
- [148] C. C. Petersen, “Cortical Control of Whisker Movement,” *Annual Review of Neuroscience*, vol. 37, no. 1, pp. 183–203, 2014.
- [149] M. Nawrot, A. Aertsen, and S. Rotter, “Single-trial estimation of neuronal firing rates: From single-neuron spike trains to population activity,” *Journal of Neuroscience Methods*, vol. 94, pp. 81–92, 1999.
- [150] S. Koyama and S. Shinomoto, “Histogram bin width selection for time-dependent Poisson processes,” *Journal of Physics A: Mathematical and General*, vol. 37, no. 29, pp. 7255–7265, 2004.
- [151] H. Shimazaki and S. Shinomoto, “A method for selecting the bin size of a time histogram,” *Neural Computation*, vol. 19, no. 6, pp. 1503–1527, 2007.
- [152] H Shimazaki and S Shinomoto, “Kernel bandwidth optimization in spike rate estimation,” *Journal of Computational Neuroscience*, vol. 29, no. 1-2, pp. 171–182, 2010.
- [153] J. M. McFarland, Y. Cui, and D. A. Butts, “Inferring Nonlinear Neuronal Computation Based on Physiologically Plausible Inputs,” *PLoS Computational Biology*, vol. 9, no. 7, 2013.
- [154] G. B. Stanley, “Adaptive spatiotemporal receptive field estimation in the visual pathway,” *Neural Computation*, vol. 14, no. 12, pp. 2925–2946, 2002.
- [155] M. Zhuang and D. Atherton, “Automatic tuning of optimum PID controllers,” *IEE Proceedings D Control Theory and Applications*, vol. 140, no. 3, p. 216, 1993.
- [156] M. M. Churchland, B. M. Yu, J. P. Cunningham, L. P. Sugrue, M. R. Cohen, G. S. Corrado, W. T. Newsome, A. M. Clark, P. Hosseini, B. B. Scott, D. C. Bradley, M. A. Smith, A. Kohn, J. A. Movshon, K. M. Armstrong, T. Moore, S. W. Chang, L. H. Snyder, S. G. Lisberger, N. J. Priebe, I. M. Finn, D. Ferster, S. I. Ryu, G.

- Santhanam, M. Sahani, and K. V. Shenoy, “Stimulus onset quenches neural variability : a widespread cortical phenomenon,” *Nature Neuroscience*, vol. 13, no. 3, pp. 369–378, 2010.
- [157] B. A. Olshausen and D. J. Field, “How close are we to understanding V1?” *Neural Computation*, vol. 17, no. 8, pp. 1665–1699, 2005.
- [158] J. Liu, H. K. Khalil, and K. G. Oweiss, “Model-based analysis and control of a network of basal ganglia spiking neurons in the normal and Parkinsonian states,” *Journal of Neural Engineering*, vol. 8, no. 4, 2011.
- [159] O. Miranda-Domínguez, J. Gonía, and T. I. Netoff, “Firing rate control of a neuron using a linear proportional-integral controller,” *Journal of Neural Engineering*, vol. 7, no. 6, 2010.
- [160] D. A. Gutnisky, J. Yu, S. A. Hires, M.-S. To, M. R. Bale, K. Svoboda, and D. Golomb, “Mechanisms underlying a thalamocortical transformation during active tactile sensation,” *PLOS Computational Biology*, vol. 13, no. 6, pp. 1–31, 2017.
- [161] K. D. Harris, D. A. Henze, H. Hirase, X. Leinekugel, G. Dragoi, A. Czurkó, and G. Buzsáki, “Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells,” *Nature*, vol. 417, no. 6890, pp. 738–741, 2002.
- [162] M. M. Churchland, J. P. Cunningham, M. T. Kaufman, J. D. Foster, P. Nuyujukian, S. I. Ryu, K. V. Shenoy, and K. V. Shenoy, “Neural population dynamics during reaching,” *Nature*, vol. 487, no. 7405, pp. 51–56, 2012.
- [163] K. J. Åström and B. Wittenmark, *Adaptive Control (2nd Edition)*. 2013, ISBN: 9780201558661.
- [164] M. N. Shadlen and W. T. Newsome, “The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding,” *The Journal of Neuroscience*, vol. 18, no. 10, pp. 3870–3896, 1998.
- [165] R. L. T. Goris, J. A. Movshon, and E. P. Simoncelli, “Partitioning neuronal variability,” *Nature Neuroscience*, vol. 17, no. 6, pp. 858–865, 2014.
- [166] V. Mante, R. A. Frazor, V. Bonin, W. S. Geisler, and M. Carandini, “Independence of luminance and contrast in natural scenes and in the early visual system,” *Nature Neuroscience*, vol. 8, no. 12, pp. 1690–1697, 2005.
- [167] Y. A. Patel, A. George, A. D. Dorval, J. A. White, D. J. Christini, and R. J. Butera, “Hard real-time closed-loop electrophysiology with the Real-Time eXperiment Interface (RTXI),” *PLoS Computational Biology*, vol. 13, no. 5, pp. 1–22, 2017.

- [168] W. Yang and R. Yuste, “In vivo imaging of neural activity,” *Nature Methods*, vol. 14, no. 4, pp. 349–359, 2017.
- [169] G. B. Stanley, “Reading and writing the neural code,” *Nature Neuroscience*, vol. 16, no. 3, pp. 259–263, 2013.
- [170] S. J. Schiff, “Towards model-based control of parkinson’s disease,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1918, pp. 2269–2308, 2010.
- [171] F. Hell, C. Palleis, J. H. Mehrkens, T. Koeglsperger, and K. Bötzel, “Deep brain stimulation programming 2.0: Future perspectives for target identification and adaptive closed loop stimulation,” *Frontiers in Neurology*, vol. 10, pp. 1–11, 2019.
- [172] A. J. Sederberg, A. Pala, H. J. Zheng, B. J. He, and G. B. Stanley, “State-aware detection of sensory stimuli in the cortex of the awake mouse,” *PLoS Computational Biology*, vol. 15, no. 5, e1006716, 2019.
- [173] C. Sanderson and R. Curtin, “Armadillo: a template-based C++ library for linear algebra,” *The Journal of Open Source Software*, vol. 1, no. 2, p. 26, 2016.
- [174] M. Pachitariu, N. Steinmetz, S. Kadir, M. Carandini, and K. Harris, “Fast and accurate spike sorting of high-channel count probes with KiloSort,” in *Advances in Neural Information Processing Systems 30*, 2016, pp. 4455–4463.
- [175] M. Sahani and J. F. Linden, “How linear are auditory cortical responses?” In *Advances in Neural Information Processing Systems 15*, 2003, pp. 109–116.
- [176] P. Van Overschee and B. De Moor, *Subspace Identification for Linear Systems*. Boston/London/Dordrecht: Kluwer Academic Publishers, 1996, ISBN: 978-1-4613-0465-4.
- [177] Z. Ghahramani and G. E. Hinton, “Parameter Estimation for Linear Dynamical Systems,” Tech. Rep. CRG-TR-96-2, 1996, pp. 1–6.
- [178] B. Shafai, S. Beale, H. H. Niemann, and J. Stoustrup, “Proportional-Integral Observers for Discrete Time Systems,” in *Proceedings of the European Control Conference*, 1995, pp. 520–525.
- [179] S. P. Linder and B. Shafai, “Rejecting disturbances to flexible structures using PI Kalman filters,” in *Proceedings of the 1997 IEEE Conference on Control Applications*, 1997.

- [180] V. A. Bavdekar, R. B. Gopaluni, and S. L. Shah, *Evaluation of adaptive extended Kalman filter algorithms for state estimation in presence of model-plant mismatch*. IFAC, 2013, vol. 10, pp. 184–189, ISBN: 9783902823595.
- [181] R. F. Stengel, *Optimal Control and Estimation*. Courier Corporation, 1994, ISBN: 9780486682006.
- [182] S. Boyd and L. Vandenberghe, *Introduction to Applied Linear Algebra*. Cambridge: Cambridge University Press, 2018, ISBN: 9781316518960.
- [183] M. P. Nawrot, C. Boucsein, V. Rodriguez Molina, A. Riehle, A. Aertsen, and S. Rotter, “Measurement of variability dynamics in cortical spike trains,” *Journal of Neuroscience Methods*, vol. 169, no. 2, pp. 374–390, 2008.
- [184] F. Wu, E. Stark, P. C. Ku, K. D. Wise, G. Buzsáki, and E. Yoon, “Monolithically Integrated μ LEDs on Silicon Neural Probes for High-Resolution Optogenetic Studies in Behaving Animals,” *Neuron*, vol. 88, no. 6, pp. 1136–1148, 2015.
- [185] J. B. Rawlings, “Tutorial Overview of Model Predictive Control,” *IEEE Control Systems*, 2000.
- [186] E. Archer, U. Köster, J. Pillow, and J. H. Macke, “Low-dimensional models of neural population activity in sensory cortical circuits,” *Advances in Neural Information Processing Systems*, vol. 1, no. January, pp. 343–351, 2014.
- [187] D. M. Zoltowski and J. W. Pillow, “Scaling the Poisson GLM to massive neural datasets through polynomial approximations,” in *Advances in Neural Information Processing Systems 32*, 2018, pp. 3517–3527.
- [188] M. M. Shanechi, A. L. Orsborn, and J. M. Carmena, “Robust Brain-Machine Interface Design Using Optimal Feedback Control Modeling and Adaptive Point Process Filtering,” *PLoS Computational Biology*, vol. 12, no. 4, pp. 1–29, 2016.
- [189] K. Oweiss, *Statistical Signal Processing for Neuroscience and Neurotechnology*. Burlington: Elsevier, 2010, ISBN: 9780123750273.
- [190] Z. Chen and S. V. Sarma, *Dynamic Neuroscience*, Z. Chen and S. V. Sarma, Eds. Cham: Springer International Publishing AG, 2018, ISBN: 9783319719764.
- [191] Y. Yang, A. T. Connolly, and M. M. Shanechi, “A control-theoretic system identification framework and a real-time closed-loop clinical simulation testbed for electrical brain stimulation,” *Journal of Neural Engineering*, vol. 15, no. 6, 2018.
- [192] A. Eleryan, M. Vaidya, J. Southerland, I. S. Badreldin, K. Balasubramanian, A. H. Fagg, N. Hatsopoulos, and K. Oweiss, “Tracking single units in chronic, large scale,

- neural recordings for brain machine interface applications,” *Frontiers in Neuroengineering*, vol. 7, no. JUL, pp. 1–13, 2014.
- [193] M. Okun, A. Lak, M. Carandini, and K. D. Harris, “Long term recordings with immobile silicon probes in the mouse cortex,” *PLoS ONE*, vol. 11, no. 3, pp. 1–17, 2016.
 - [194] A. A. Prinz, L. F. Abbott, and E. Marder, “The dynamic clamp comes of age,” *Trends in Neurosciences*, vol. 27, no. 4, pp. 218–224, 2004.
 - [195] A. Pala and C. C. Petersen, “State-dependent cell-type-specific membrane potential dynamics and unitary synaptic inputs in awake mice,” *eLife*, vol. 7, pp. 1–13, 2018.
 - [196] E. H. Ratzlaff and A. Grinvald, “A tandem-lens epifluorescence macroscope: Hundred-fold brightness advantage for wide-field imaging,” *Journal of Neuroscience Methods*, 1991.
 - [197] A. Grinvald, R. D. Frostig, R. M. Siegel, and E. Bartfeld, “High-resolution optical imaging of functional brain architecture in the awake monkey,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 88, no. 24, pp. 11 559–11 563, 1991.
 - [198] V. A. Kalatsky and M. P. Stryker, “New paradigm for optical imaging: Temporally encoded maps of intrinsic signal,” *Neuron*, vol. 38, no. 4, pp. 529–545, 2003.
 - [199] S. Lefort, C. Tómm, J. C. Floyd Sarria, and C. C. Petersen, “The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex,” *Neuron*, vol. 61, no. 2, pp. 301–316, 2009.
 - [200] C. Waiblinger, D. Brugger, C. J. Whitmire, G. B. Stanley, and C. Schwarz, “Support for the slip hypothesis from whisker-related tactile perception of rats in a noisy environment,” *Frontiers in Integrative Neuroscience*, vol. 9, no. October, pp. 1–11, 2015.
 - [201] C. Rossant, S. N. Kadir, D. F. Goodman, J. Schulman, M. L. Hunter, A. B. Saleem, A. Grosmark, M. Belluscio, G. H. Denfield, A. S. Ecker, A. S. Tolias, S. Solomon, G. Buzski, M. Carandini, and K. D. Harris, “Spike sorting for large, dense electrode arrays,” *Nature Neuroscience*, vol. 19, no. 4, pp. 634–641, 2016.
 - [202] H. G. Rey, C. Pedreira, and R. Quiñan Quiroga, “Past, present and future of spike sorting techniques,” *Brain Research Bulletin*, vol. 119, pp. 106–117, 2015.
 - [203] R. H. Shumway and D. S. Stoffer, “An Approach to The Time Series Smoothing and Forecasting Using the EM Algorithm,” *Journal of Time Series Analysis*, vol. 3, no. 4, pp. 253–264, 1982.

- [204] T. A. Woolsey and H. Van der Loos, “The structural organization of layer IV in the somatosensory region (S I) of mouse cerebral cortex,” *Brain Research*, 1970.
- [205] B. Haider, M. Häusser, and M. Carandini, “Inhibition dominates sensory responses in the awake cortex,” *Nature*, vol. 493, no. 7430, pp. 97–102, 2013.
- [206] P. Nuyujukian, J. Albites Sanabria, J. Saab, C. Pandarinath, B. Jarosiewicz, C. H. Blabe, B. Franco, S. T. Mernoff, E. N. Eskandar, J. D. Simeral, L. R. Hochberg, K. V. Shenoy, and J. M. Henderson, “Cortical control of a tablet computer by people with paralysis,” *PLoS ONE*, vol. 13, no. 11, pp. 1–16, 2018.
- [207] S. R. Beale and B. Shafai, “Robust control system design with proportional integral observer,” in *Proceedings of the 27th IEEE Conference on Decision and Control*, 1988, 554–557 vol.1.
- [208] K. J. Astrom and L. Rundqwist, “Integrator windup and how to avoid it,” 1989.
- [209] A. Clemente-Perez, S. R. Makinson, B. Higashikubo, S. Brovarney, F. S. Cho, A. Urry, S. S. Holden, M. Wimer, C. Dávid, L. E. Fenno, L. Acsády, K. Deisseroth, and J. T. Paz, “Distinct Thalamic Reticular Cell Types Differentially Modulate Normal and Pathological Cortical Rhythms,” *Cell Reports*, vol. 19, no. 10, pp. 2130–2142, 2017.
- [210] J. V. Raimondo, L. Kay, T. J. Ellender, and C. J. Akerman, “Optogenetic silencing strategies differ in their effects on inhibitory synaptic transmission,” *Nature Neuroscience*, vol. 15, no. 8, pp. 1102–1104, 2012.
- [211] A. M. Herman, L. Huang, D. K. Murphey, I. Garcia, and B. R. Arenkiel, “Cell type-specific and time-dependent light exposure contribute to silencing in neurons expressing Channelrhodopsin-2,” *eLife*, vol. 3, pp. 1–18, 2014.
- [212] E. Todorov and Weiwei Li, “A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems,” in *Proceedings of the 2005, American Control Conference, 2005.*, 2005, pp. 300–306, ISBN: 0-7803-9098-9.
- [213] S. R. Olsen, D. S. Bortone, H. Adesnik, and M. Scanziani, “Gain control by layer six in cortical circuits of vision,” *Nature*, vol. 483, no. 7387, pp. 47–54, 2012.
- [214] D. S. Bortone, S. R. Olsen, and M. Scanziani, “Translaminar Inhibitory Cells Recruited by Layer 6 Corticothalamic Neurons Suppress Visual Cortex,” *Neuron*, vol. 82, pp. 474–485, 2014.

- [215] T. A. Engel, N. A. Steinmetz, M. A. Gieselmann, A. Thiele, T. Moore, and K. Boahen, "Selective modulation of cortical state during spatial attention," *Science*, vol. 354, no. 6316, pp. 1140–1144, 2016.
- [216] G. Radons, J. D. Becker, B. Dülfer, and J. Krüger, "Analysis, classification, and coding of multielectrode spike trains with hidden Markov models," *Biological Cybernetics*, vol. 71, no. 4, pp. 359–373, 1994.
- [217] A. C. Camproux, F. Saunier, G. Chouvet, J. C. Thalabard, and G. Thomas, "A hidden Markov model approach to neuron firing patterns," *Biophysical Journal*, vol. 71, no. 5, pp. 2404–2412, 1996.