

Copyright
by
Kenneth William Latimer
2015

The Dissertation Committee for Kenneth William Latimer
certifies that this is the approved version of the following dissertation:

**Statistical approaches for unraveling the neural code in
the visual system**

Committee:

Alexander Huk, Supervisor

Jonathan Pillow, Co-Supervisor

Ila Fiete

Nicholas Priebe

Carlos Carvalho

**Statistical approaches for unraveling the neural code in
the visual system**

by

Kenneth William Latimer, B.S.C.S.

DISSERTATION

Presented to the Faculty of the Graduate School of
The University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2015

Dedicated to Mr. Spock

Acknowledgments

I am grateful for all the guidance I have received during the past five years by my two advisers, Jonathan Pillow and Alex Huk. I would also like to thank my committee members, Nicholas Priebe, Ila Fiete, and Carlos Carvalho, for the valuable feedback they have provided.

I wish to thank my fellow INS graduate students for all the help and entertainment they've provided during my time at UT, especially Jacob Yates, without whom I would never have thought to look at models of LIP, Leor Katz and Ben Scholl. I thank Memming Park for guiding me through much of my time in the Pillow lab.

I thank E.J. Chichilnisky and Fred Rieke for providing helpful input and unique datasets for the retinal ganglion cell project and Miriam Meister for providing the LIP recordings. And I also extend my thanks to Jamie Roitman and Michael Shadlen for graciously making their response-time task recordings in LIP publicly available.

I thank my family for all their love and encouragement. And finally Heather and our two wonderful cats, Dax the Cat and Wanton Kitty, for their constant love and support.

Statistical approaches for unraveling the neural code in the visual system

Publication No. _____

Kenneth William Latimer, Ph.D.
The University of Texas at Austin, 2015

Supervisor: Alexander Huk
Co-Supervisor: Jonathan Pillow

Here we consider the neural coding problem at two levels of the macaque visual system. First, we analyze single neurons recorded in the lateral intraparietal (LIP) cortex while a monkey performed a perceptual decision-making task. We relate the single-trial responses in LIP to stochastic decision-making processes with latent dynamical models. We compare models with latent spike rates governed by either continuous diffusion-to-bound dynamics or discrete “stepping” dynamics. In contrast to previous findings, roughly three-quarters of the choice-selective neurons we recorded are better described by the stepping model. Second, we introduce a biophysically inspired point process model that explicitly incorporates stimulus-induced changes in synaptic conductance in a dynamical model of neuronal membrane potential. We show that our model provides a tractable model of spike responses in macaque parasol retinal ganglion cells that is both more accurate and more interpretable than the popular

generalized linear model. Most importantly, we show that we can accurately infer intracellular synaptic conductances from extracellularly recorded spike trains.

Table of Contents

Acknowledgments	v
Abstract	vi
List of Tables	xii
List of Figures	xiii
Chapter 1. Introduction	1
1.1 Neural representations of perceptual decisions	2
1.2 Coding of visual stimuli in the retina: combining excitation and inhibition	4
Chapter 2. The dynamics of neurons in area LIP during perceptual decision making	7
2.1 Encoding of perceptual decisions in parietal cortex	8
2.1.1 Previous work	14
2.2 Spike train models of decision-making dynamics in LIP	20
2.2.1 Diffusion-to-bound (ramping) model	22
2.2.2 Discrete stepping model	25
2.3 Bayesian model comparison	29
2.4 Markov chain Monte Carlo methods	33
2.4.1 Ramping model	34
2.4.1.1 Sampling the latent state with a particle filter	34
2.4.1.2 Sampling the model parameters	40
2.4.2 Stepping model	45
2.4.2.1 Sampling the latent steps	45
2.4.2.2 Sampling the model parameters	46
2.4.3 Visualizing samples from the posterior distribution	49

2.4.4	Analyzing simulated spike trains	53
2.4.4.1	Simulations: evaluating sampler efficiency	53
2.4.4.2	Estimating parameters in simulations	56
2.4.4.3	Model comparison identifies true model in simulations	57
2.5	Results of model comparison of LIP responses	58
2.5.1	Cell selection	59
2.5.2	Model fits to 40 LIP neurons	60
2.5.3	Model comparison supports stepping model in majority of cells	60
2.5.3.1	Population spiking statistics better fit by stepping mode	62
2.5.4	Inferring steps on single trials	64
2.5.4.1	Decoding choice with steps	67
2.5.4.2	Estimating step times	71
2.5.5	Robustness of results	74
2.5.5.1	Response-time task: Roitman & Shadlen, 2002	74
2.5.5.2	Model comparison results are unaffected by grouping coherence levels	78
2.5.5.3	Model comparison results are consistent across start time of analysis	78
2.6	Comparison to previous methods	81
2.6.1	Variance of the Conditional Expectation	81
2.6.2	Bollimunta et al., 2012	85
2.7	Conclusions	88
Chapter 3. Inferring synaptic conductances from spike trains in retinal ganglion cells		90
3.1	Methods	92
3.1.1	Electrophysiology	92
3.1.1.1	Synaptic current recordings	92
3.1.1.2	Dynamic clamp recordings	93
3.1.1.3	RGC population recordings: full-field stimulus	93
3.1.1.4	RGC population recordings: spatio-temporal stimulus	94

3.1.2	Modeling methods	95
3.1.2.1	The conductance-based spiking model	95
3.1.2.2	Spike rate nonlinearity	99
3.1.2.3	Conductance LN cascade	100
3.1.2.4	Generalized linear models	101
3.1.2.5	Modeling responses to spatio-temporal stimuli	102
3.1.2.6	Evaluating model performance	103
3.2	Results	104
3.2.1	A conductance-based extension to the GLM	104
3.2.1.1	Interpreting the GLM as a conductance-based model	105
3.2.2	Approximating the input and output properties of ON parasol cells with intracellular recordings	108
3.2.2.1	Spike rate nonlinearity	108
3.2.2.2	Conductances as a linear-nonlinear function of the stimulus	109
3.2.3	Model fitting in simulations	112
3.2.4	Predicting conductance tuning from spikes in ON parasol cells	112
3.2.5	Characterization of RGC responses by inferring excitatory and inhibitory inputs	117
3.2.6	RGC responses across different temporal contrasts	121
3.2.7	Excitatory and inhibitory responses to spatial stimuli	123
3.3	Conclusion	128
Chapter 4. Discussion		130
4.1	Future Directions: Modeling decision representations in multi-cell recordings	131
4.1.1	Confirming single-cell results in a new dataset	132
4.1.2	Are step times across different neurons correlated?	133
4.1.3	Extending the stepping model to multiple neurons	134
4.2	Future directions: conductance-based modeling of RGCs	137
4.3	Conclusion	139
Appendices		141

Appendix A. Methods used to analyze LIP response	142
A.1 Computing PSTHs	142
A.2 Parameter estimates	145
A.2.1 Ramping model	145
A.2.2 Stepping model	146
A.3 Example LIP cells	147
A.4 Cell examples over different start times of analysis	150
Appendix B. Gradients of CBSM log likelihood	152
B.1 Notation for the log likelihood	152
B.2 Gradients for conductance filters	154
B.2.1 Second derivative for conductance filters	156
B.2.2 Mixed second derivative for conductance filters	157
B.3 Gradients for the spike history filter	159
B.3.1 Second derivative for the spike history filter	159
B.3.2 Mixed second derivative for the spike history and conductance filters	159
Bibliography	160
Vita	177

List of Tables

A.1	Parameter estimates for LIP population: ramping model . . .	145
A.2	Parameter estimates for LIP population: stepping model . . .	146

List of Figures

2.1	LIP responses during a moving-dot discrimination task	10
2.2	Models of LIP spike train dynamics	12
2.3	Graphical model: diffusion-to-bound	23
2.4	Graphical model: switching	26
2.5	Step time distributions	27
2.6	Example of the estimated posterior distribution	50
2.7	Samples of the latent firing rate: stepping model	51
2.8	Samples of the latent firing rate: ramping model	52
2.9	Autocorrelation of the MCMC output	54
2.10	Autocorrelation of the MCMC output on real data	55
2.11	Convergence of parameter estimates	56
2.12	Model comparison results on simulated data	57
2.13	Stepping model fits to LIP population	61
2.14	Ramping model fits to LIP population	61
2.15	Model comparison results favor stepping model	62
2.16	Population PSTH and PSTV compared to model fits	63
2.17	Model-based analysis of a single LIP neuron	65
2.18	Population PSTH aligned to step time	66
2.19	Step-aligned firing rate of a simulated ramping cell	67
2.20	Decoding choice with the stepping model	69
2.21	Estimated step times in LIP population	73
2.22	Model comparison results for cells in a response-time task . . .	75
2.23	Choice-sorted activity of LIP cells in a response-time task . . .	76
2.24	Population average PSTH in a response-time task	77
2.25	Grouping coherence levels does not effect results	79
2.26	Model comparison results consistent over start time of analysis	80
2.27	Estimating spike rate variance: VarCE	84

2.28	Estimating spike rate covariance: CorCE	86
3.1	The conductance-based spiking model	96
3.2	Firing rate nonlinearity	109
3.3	Cascade model of synaptic inputs	111
3.4	Convergence of CBSM parameter estimates	113
3.5	Inferring conductances from spikes in ON parasol cells	115
3.6	Summary of conductance estimates	116
3.7	Estimating RGC spike responses with the CBSM	118
3.8	Including inhibition improves estimates of RGC spike responses	119
3.9	The GLM estimates excitatory input	120
3.10	The CBSM exhibits contrast gain control	122
3.11	Generalization of CBSM across coherences	124
3.12	Estimating RGC spike responses with the CBSM	125
3.13	CBSM improves characterization of RGC responses to spatio-temporal stimuli	126
3.14	Center-surround interactions in the CBSM	127
4.1	Model comparison of second dataset	133
4.2	Correlation of step times	134
4.3	Modeling a simultaneous step in cell pairs	136
A.1	Variance explained of population PSTH and PSTV	144
A.2	Example LIP cells 1	147
A.3	Example LIP cells 2	148
A.4	Example LIP cells 3	149
A.5	Cell examples: analysis start time 1	150
A.6	Cell examples: analysis start time 2	151

Chapter 1

Introduction

Neurons transmit information by emitting a series of action potentials, also known as a “spike train.” Finding a functional description between information in the world and spike trains recorded in the brain, the *neural coding problem*, is a fundamental goal of systems neuroscience. Statistical neuroscience techniques aim to solve this problem by defining probability distributions over spike trains (mathematically formulated as point processes) which depend on variables of interest.

Much insight about the neural code has been gained by studying neurons in the visual system (e.g., Hubel & Wiesel 1962, Kastner & Baccus 2014, Ma & Jazayeri 2014). Activity in the early visual system is closely linked with a visual stimulus that can be carefully controlled by an experimenter – although responses are variable across multiple presentations of the same stimulus. By presenting many stimuli to many visually responsive neurons, statistical relationships between spikes and the visual world can be mapped. However, an animal does not passively view the world: sensory signals are used to guide actions. Neurons can also be recorded while an animal performs an action in response to a stimulus in controlled tasks (Glimcher 2003). Animal

behavior varies from trial-to-trial, and therefore the neural responses during a behavioral task that are linked to an animal's internal decision to act will also vary from trial-to-trial. Averaging spike train responses over many trials with respect to an external variable such as stimulus onset can therefore fail to reveal the neural representation of internally generated signals during behavior. Statistical techniques that interpret single-trial spike train responses are therefore critical for understanding how perceptual decisions are formed and represented in cortex (Park, Meister, Huk & Pillow 2014).

This dissertation presents statistical modeling approaches to link observed spike trains to unobserved dynamical processes that underlie the spike trains. Chapter 2 examines the representation and formation of a decision variable in the lateral intraparietal (LIP) cortex. Chapter 3 aims to extend cascade models of coding in the early visual system with a simplified biophysical model of the membrane dynamics.

1.1 Neural representations of perceptual decisions

Chapter 2 introduces methods for comparing hypotheses of the dynamics of neural representations underlying perceptual decision making to single neuron recordings. A decision is an inherently unobservable variable that is represented by an entire network of neurons. The study of decision making in systems neuroscience strives to link internally generated behaviors to neural activity in order to understand how decisions to execute specific motor actions in response to a stimulus are formed (Hanks, Kopec, Brunton, Duan,

Erlich & Brody 2015). Single neuron recordings can provide a window into the decision making process in order to constrain behavioral models of decision making (Ditterich 2006). However, careful characterization the nature of the dynamics of the neural representation of perceptual decisions is needed to relate single-neuron recordings to cognitive-level notions of decisions.

At the behavioral level, a noisy evidence accumulation process successfully describes perceptual decision making by accounting for accuracy and reaction times (Ratcliff & Rouder 1998, Brunton, Botvinick & Brody 2013). The actual implementation of the evidence accumulation process in the brain, however, could appear very different than an integrator, especially at the level of single cells. Perceptual decisions could instead be implemented by an attractor network in which activity probabilistically jumps between discrete states (Deco, Pérez-Sanagustín, de Lafuente & Romo 2007, Durstewitz & Deco 2008, Miller & Katz 2010). Networks with multiple stable attractor states are of particular interest for their biological plausibility and robustness as a mechanism for integration (Koulakov, Raghavachari, Kepecs & Lisman 2002).

The average firing rate of neurons in several brain areas, notably LIP, correlate with an accumulation process during decision-making tasks (Roitman & Shadlen 2002, Mazurek, Roitman, Ditterich & Shadlen 2003, Gold & Shadlen 2007, Kiani, Hanks & Shadlen 2008, Kiani & Shadlen 2009, Purcell, Heitz, Cohen, Schall, Logan & Palmeri 2010, Shadlen & Kiani 2013). Concluding that the average observed neural responses demonstrate a direct link between single

neurons and the behavioral level description of decision formation is alluring, but the accumulation-to-bound hypothesis aims to describe single-trial decision making and not the average of independent decisions. We therefore aimed to test statistically the hypothesis of whether the single-trial activity of LIP neurons could be explained by an accumulation process or a stochastic discrete state-changing (or “stepping”) process.

We accomplish this by modeling the decision variable on each trial as an unobserved (latent), stochastic variable that follows either an accumulation-to-bound process or a discrete stepping process. In the models, the spike rate of the neuron is controlled by the latent decision variable. We fit these two candidate models to 40 choice-selective neurons recorded in LIP during a perceptual decision-making task using Markov chain Monte Carlo methods. Under a Bayesian model comparison framework, we show that the responses in a majority of these cells are in fact best captured by the discrete stepping process.

1.2 Coding of visual stimuli in the retina: combining excitation and inhibition

The response properties of visual neurons are often described as a linear filtering of the stimulus that is passed through a static nonlinear function to produce a spike rate – the linear-nonlinear (LN) model (Chichilnisky 2001, Chander & Chichilnisky 2001, Paninski 2004). The high-dimensionality of the visual world and the stochasticity of neural responses, however, have made

cracking the neural code a challenging task, even in the retina. Although there has been much progress in unraveling the input-output relationships in the retina, current statistical methods neither have a clear relationship to biophysical mechanisms nor do they accurately predict responses across ranges of naturalistic stimuli. In Chapter 3, we introduce a biophysically inspired extension to the LN cascade paradigm for describing single-neuron responses in the primate retina. Retinal ganglion cells (RGCs) provide an ideal foothold to decipher the neural code because responses are closely tied to a controllable visual input, and the neurons providing input to the cell are isolated.

RGCs integrate excitatory and inhibitory synaptic inputs to encode the visual information that is sent to the brain (Roska, Molnar & Werblin 2006). We demonstrate that a particular widely used instantiation of the LN cascade model, the generalized linear point process model (GLM), can be interpreted as a simple conductance-based model of the membrane voltage in which excitation and inhibition have equal and opposite tunings. We extend the GLM by explicitly modeling the excitatory and inhibitory conductances, and allowing them to have independent tunings. In contrast to the modeling of decision variables modeled in Chapter 2, excitatory and inhibitory synaptic conductances are concrete quantities that can be measured and used to constrain the model. We therefore show using intracellular recordings that the model accurately estimates the excitatory and inhibitory tuning received by the cell from spike trains alone.

Neural responses adapt to shifts in stimulus statistics and local stimulus

statistics can vary drastically within a single natural scene (Fairhall, Lewen, Bialek & van Steveninck 2001, David, Vinje & Gallant 2004, Wark, Lundstrom & Fairhall 2007, Hong, Lundstrom & Fairhall 2008, Ozuysal & Baccus 2012),. We show that the conductance-based modeling approach can exhibit adaptive coding across a range of temporal contrasts by having a stimulus-dependent gain. Additionally, we find that the integration of excitatory and inhibitory synaptic inputs can alter the spatial selectivity of simulated neurons compared to linear-nonlinear models. While the classification of RGC responses using a single linear filter has shed much light on the coding properties of retina under different conditions, considering both the excitatory and inhibitory stimulus tunings that a cell receives provides a more robust and accurate description of the neural code in the retina.

Chapter 2

The dynamics of neurons in area LIP during perceptual decision making

Neurons in the macaque lateral intraparietal (LIP) area exhibit firing rates that appear to ramp upwards or downwards during decision-making. These ramps are commonly assumed to reflect the gradual accumulation of evidence towards a decision threshold. However, the ramping in trial-averaged responses could instead arise from instantaneous jumps at different times on different trials. We examine single-trial responses in LIP using statistical methods for fitting and comparing latent dynamical spike train models. We compare models with latent spike rates governed by either continuous diffusion-to-bound dynamics or discrete “stepping” dynamics. Roughly three-quarters of the choice-selective neurons recorded in LIP are better described by the stepping model. Moreover, the inferred steps carry more information about the animal’s choice than spike counts.

This work has been published in *Science* (Latimer, Yates, Meister, Huk & Pillow in press) and was in part presented at the Computational and Systems Neuroscience (COSYNE) 2015 meeting. The statistical inference methods are described in Latimer, Huk & Pillow (2015).

2.1 Encoding of perceptual decisions in parietal cortex

A fundamental challenge in neuroscience is to understand how decisions are computed and represented in neural circuits. One popular approach to this problem is to record from single neurons in brain regions that lie between primary sensory and motor regions while an animal performs a perceptual decision-making task. Typical tasks require the animal to integrate noisy sensory evidence over time in order to make a binary decision about the stimulus. Such experiments have the tacit goal of characterizing the dynamics governing the transformation of sensory information into a representation of the decision. However, recorded spike trains do not reveal these dynamics directly; they represent noisy, incomplete emissions that reflect the underlying dynamics only indirectly.

This dissociation between observed spike trains and the unobserved dynamics governing neural population activity has posed a key challenge for using neural measurements to gain insight into how the brain computes decisions. Recording decision-related neural activity has certainly shed much light upon what parts of the brain are involved in forms of decision making and what sorts of roles each area plays. But without direct access to the dynamics underlying single-trial decision formation, most analyses of decision-related neural data rely on estimating spike rates by averaging over trials (and often over many independently recorded neurons as well). Although the central tendency is of course a reasonable starting point in data analysis, sole reliance on the mean can obscure single-trial dynamics when substantial stochastic components are

present. For example, as discussed in depth in this chapter, the average of a set of step functions — when the steps occur at different times on different trials — will yield an average that ramps continuously, masking the presence of discrete dynamics. Although the majority of averaging and regression-based analyses used in the field are straightforward to conceptualize and easy to apply to data, they provide limited insight into the dynamics that may govern how individual decisions are made.

This chapter presents statistical methods to compare hypotheses of the dynamics the single-trial dynamics governing spike trains in the LIP of macaques performing a well-studied motion-discrimination task (Fig. 2.1) (Newsome & Pare 1988, Kiani et al. 2008, Meister, Hennig & Huk 2013). The modeling approach formalizes two latent dynamical models of spike trains that have previously been proposed to govern decision-making behavior: (1) diffusion-to-bound (or “ramping”); and (2) discrete switching (or “stepping”). In the ramping model, also known as “diffusion-to-bound”, the spike rate evolves according to a Gaussian random walk with linear drift (Fig. 2.2A). The slope of drift depends on the strength of sensory evidence, and each trial’s trajectory continues until hitting an absorbing upper bound. Alternatively, in the stepping model, the latent spike rate jumps instantaneously from an initial “undecided” state to one of two discrete decision states during the trial (Fig. 2.2B). The probability of stepping up or stepping down and the timing of the step are determined by the strength of sensory evidence.

In the diffusion-to-bound model, the spike rate is a Gaussian random

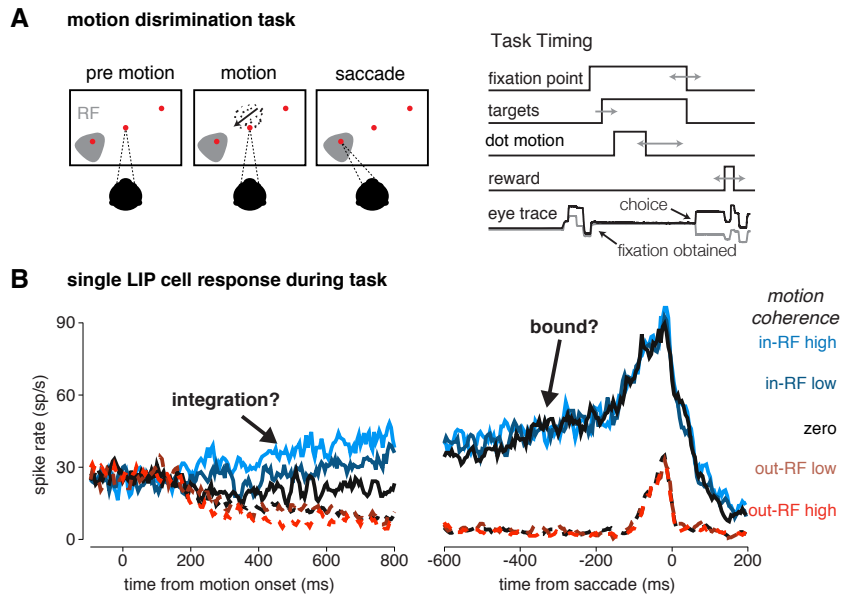


Figure 2.1: **(A)** Schematic of moving-dot direction-discrimination task. The monkey fixates while viewing a motion stimulus of variable motion strength and duration (500-1000 ms). The monkey must determine the direction of the moving dot stimulus, selecting from two possible choices. The monkey indicates its choice by making a saccade to one of two choice targets after receiving a go-signal 500 ms after motion offset. The difficulty of the task is controlled by varying the coherence of the noisy dot-motion stimulus. One choice target is in the response field of the neuron under study (RF; shaded patch on left); the other is outside it. **(B)** Average response of an LIP neuron during a motion-direction discrimination task. (left) The spike rates are aligned to the onset of the motion stimulus, and the traces are sorted by the coherence (strength and direction) of motion. The firing rate begins to show a stimulus-dependent ramping beginning at ~ 200 ms after motion onset (Churchland et al. 2008, de Lafuente et al. 2015). (right) The spike rates are aligned to the time of the saccade. Before a saccade into the cell’s RF, the firing rates converge to a common across coherence levels, as if the firing rates are achieving a common decision bound or threshold.

walk plus a constant linear drift, which stops once an upper bound is reached. In one popular instantiation of this model, the spike rate represents integrated

sensory evidence, and the bound corresponds to an evidentiary threshold at which the decision is made. Ramping responses have been observed in a variety of brain areas, such as the LIP and the frontal eye field (FEF), during decision-making, and have been widely interpreted as the neural implementation of evidence accumulation for forming decisions (Roitman & Shadlen 2002, Mazurek et al. 2003, Gold & Shadlen 2007, Kiani et al. 2008, Kiani & Shadlen 2009, Purcell et al. 2010, Shadlen & Kiani 2013). However, ramping can only be observed by averaging together responses from many trials (and often, many neurons), which obscures the dynamics governing responses on single trials. Although diffusion-to-bound models have been widely used to quantify behavior in decision-making tasks (Smith & Ratcliff 2004), they have not to our knowledge been formally incorporated into explicit point process models of spike train data. Thus, despite the appealing analogies between psychological accumulation and neural activity, it remains possible that the neural computations underlying decisions may be quite different than the idealized accumulation process that describes behavior. In particular, a discrete “stepping” process (Durstewitz & Deco 2008, Miller & Katz 2010), in which the spike rate jumps stochastically from one rate to another at some time during each trial, can also create the appearance of ramping (Goldman 2015, Churchland, Kiani, Chaudhuri, Wang, Pouget & Shadlen 2011). Although decision-making at the behavioral level is well described as an accumulation process (Ratcliff & Rouder 1998, Brunton et al. 2013), whether the brain computes decisions via a direct neural correlate (ramping) or a discrete implementation (stepping)

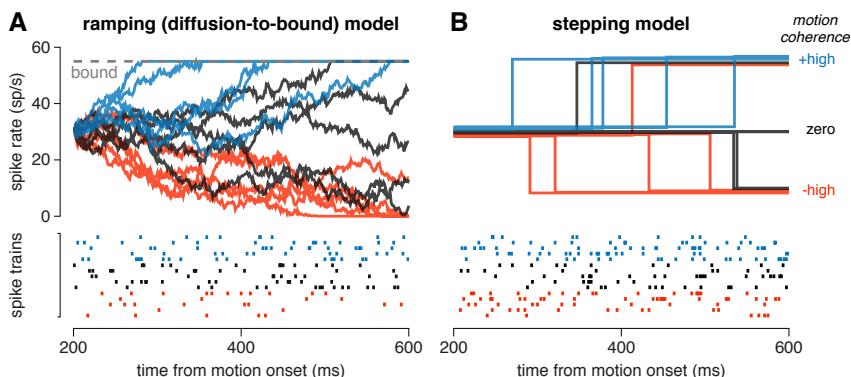


Figure 2.2: Two proposed models for the dynamics underlying LIP responses during motion evidence accumulation. Spike trains are modeled during a window beginning 200 ms after motion onset until 200 ms after motion offset (300 ms before the monkey receives the go signal). **(A)** Ramping (diffusion-to-bound) model. Spike rate trajectories (solid traces) were sampled from a diffusion-to-bound process for each of three motion coherences (strong positive, zero, and strong negative). The model parameters included an initial spike rate, a slope for each coherence, noise variance, and an upper bound. The model did not include a lower bound, consistent with the competing integrator (race) model of LIP (Shadlen & Kiani 2013). Spike trains (below) obey an inhomogeneous Poisson process for each spike rate trajectory. **(B)** Discrete stepping model. Spike rate trajectories (above) begin at an initial rate and jump “up” or “down” at a random time during each trial, and spike trains (below) are once again Poisson given the latent rate. The step times take a negative binomial distribution, which can resemble the time-to-bound distribution under a diffusion model. Parameters include the spike rates for the three discrete states and two parameters governing the distribution over step timing and direction for each motion coherence. Both models were fit using the spike trains and coherences for each neuron, without access to the animal’s choices.

remains a central, unresolved question in neuroscience.

An alternative to the diffusion-based model of decision-making is that populations of neurons jump between discrete states associated with different

firing rates (Miller & Katz 2010). If steps between low firing and high firing states occur at different times on each trial, this discrete “stepping model” can equally well produce an average spike rate profile that ramps. This class of models is attractive for decision-making because of the obvious mapping between states and decisions, and because it can be robustly implemented in neurally plausible models (Durstewitz & Deco 2008). Although the discrete transitions in this model are inherently distinct from the continuous ramping in the diffusion-to-bound model, the two models can resemble one another when the data are viewed through the conventional lens of trial-based averaging. Here, a restricted discrete-state model is examined in which the neuron starts in an initial state and is allowed to make a single jump to one of two choice states during a trial.

The following sections present latent variable models with either diffusion or switching dynamics, show how these models can be fit to spike train data, and describe Bayesian model selection methods for comparing and evaluating model fits to data. Both of the proposed models incorporate nonlinear dynamics, and there are no simple closed-form estimators for the model parameters. We therefore use Markov chain Monte Carlo (MCMC) methods to sample from the posterior distribution of the model parameters given the data. These MCMC methods can be tractably applied to the data from an entire experimental session. The samples, which act as an approximation of the full posterior, can be used to compute Bayesian estimates of model parameters, while also incorporating uncertainty about the parameters into conclusions

about model performance. Numerical experiments demonstrate that these methods reliably determine the true dynamics of spiking data simulated from each of the two models. We then use the models to determine quantitatively which family of dynamics best describes spike responses observed in a set of 40 choice-selective neurons recorded in area LIP.

2.1.1 Previous work

Models based on the “noisy accumulation of evidence” have been widely employed in the psychology literature on normative accounts of decision making (Wald 1973, Link 1975, Ratcliff & Rouder 1998). In the basic drift-diffusion model, a latent “particle” is initialized at zero and then moves with a linear drift plus diffusion noise. Once the particle reaches an upper or lower bound, each of which corresponds to a possible choice, the decision is made. Models of this class robustly capture subjects’ choice distributions and reaction times (Ditterich 2006). Previous work has demonstrated how to fit behavioral data with such models (Wiecki, Sofer & Frank 2013), including non-trivial stimulus-dependent effects on the diffusion process (Brunton et al. 2013). However, these models were originally used as a tool to study the representation of a decision at a cognitive level. More recently, this model has been applied to neural recordings taken while a subject performs a decision-making task (Mazurek et al. 2003, Shadlen & Kiani 2013). The responses of single neurons in LIP (and several other brain regions) appear to encode a diffusion-to-bound process, as their averaged spike rates look like a direct

neural correlate of diffusion-to-bound (Huk & Meister 2012).

Many initial claims about decision-related diffusion processes in single-neuron recordings in LIP were derived from average spike count responses that, when aligned to the onset of a noisy moving-dot stimulus, appear to ramp with a slope dependent on the strength of the motion signal. The slope of the average ramping also changes along with shifts in the stimulus strength (Huk & Shadlen 2005). The activity of LIP neurons appears to accumulate up to a stimulus-independent upper bound, or threshold, before the monkey chooses to saccade into the cell’s RF (Fig. 2.1B, right) (Roitman & Shadlen 2002, Shadlen & Kiani 2013). However, the firing rate does not appear to reach a lower bound when the monkey makes an out-RF decision. This is consistent with a competing accumulator model: LIP contains an independent integrator for each of the possible targets and the first accumulator to reach the threshold triggers the decision. The competing accumulators model extends the “noisy accumulator of evidence” model to tasks with more than two choices (Churchland et al. 2008). More recently, average ramping activity has been shown in LIP during tasks where the evidence for a decision is presented more abstractly by a series of shapes (Yang & Shadlen 2007, Kira, Yang & Shadlen 2015). In these studies, the monkey viewed a sequence of shapes. In each trial, the sequence shapes were drawn independently from a distribution that depended on the correct choice for the trial. Therefore, some shapes would provide evidence that the in-RF target was the correct choice, while other shapes would indicate that the out-RF target was more likely. The average single-unit ac-

tivity recorded in LIP during these tasks appeared to ramp proportionally to the log-likelihood ratio of the evidence that a choice into the cell's RF would result in a reward.

However, these conclusions are based on responses averaged over many trials and many neurons. The mean response provides only a limited window into the neuron's dynamics, because average responses do not reveal the trial-to-trial variability expected in a noisy accumulation process. Recent work has proposed to use the response mean and variance to reveal signatures of an underlying evidence integration process (Churchland et al. 2011). Although this is clearly superior to methods that use only on the mean, section 2.6.1 shows that this approach has limited ability to discriminate between continuous and discrete latent dynamics in a number of settings.

One recent study focused on model comparison for spike trains recorded in macaque LIP (Bollimunta, Totten & Ditterich 2012) (see Section 2.6.2). The authors concluded that single-neuron activity in this area reflected a drift-diffusion process rather than a discrete-stepping process. This is an important first step in applying this class of models to decision-making, but their initial efforts were affected by several limitations. Computational constraints forced the authors to examine restricted versions of the models and only a small percentage of a cell's spike trains at a time. The latent variable models we present in this chapter were formulated to match the models proposed elsewhere in the decision-making literature, and provide a computationally tractable method for fitting to a larger set of spike trains.

After defining potential models to describe spike train data, one must fit the models and compare the fits to the data. Fitting the diffusion-to-bound model to simple behavioral data is not trivial, and adding a spike train observation to this process adds to the complexity. The expectation-maximization algorithm (EM) is a commonly used tool for maximum likelihood inference for a variety of latent dynamical models of spike train data (e.g. Escola, Fontanini, Katz & Paninski 2011, Buesing, Macke & Sahani 2012). These models define the spiking process as a Poisson process where the intensity function (i.e., spike rate) is also a stochastic process. Typically, EM cannot be directly applied to the models of interest where the spiking is Poisson and the intensity follows a Gaussian process, and approximations must be made. Model fits obtained by EM algorithms are useful for studying data under a single model. However, EM provides only a point estimate of the model parameters, and particular choices of approximations could affect the quality of the fit. The goal in Bayesian model comparison is to incorporate uncertainty in parameter estimates, and model comparison ideally employs exact methods which do not favor one particular model due to the inaccuracies in approximate model fitting.

Yuan, Girolami & Niranjana (2012) applied exact Bayesian techniques to the problem of estimating latent dynamics underlying a spike rate. They assumed that the spike train could be modeled as a log Gaussian Cox process. MCMC methods were used to fit the model. These methods are typically much slower than EM, but the clever sampling techniques they proposed resulted in

efficient sampling. This previous use of MCMC methods for a spiking model fit a single model; for the application of these methods to perceptual decision making, this chapter presents Bayesian tools to compare multiple possible models that contain different dynamics.

Bromberg-Martin, Matsumoto & Hikosaka (2010) performed model comparison of the activity of dopamine neurons in a similar spirit to the models presented in this chapter. The authors compared two models of the spike rate for the observed neurons: (1) a deterministic, linearly ramping rate and (2) a two-state Hidden Markov model (HMM) that includes a baseline and a phasic response state. These models could be fit efficiently by maximum likelihood methods, and the authors compared the model fits using the maximum likelihood values. HMMs have been widely used to model discrete state transitions, like the stepping model, in neural data. They can be tractably fit to data using standard EM algorithms (Seidemann, Meilijson, Abeles, Bergman & Vaadia 1996, Miller & Katz 2010, Escola et al. 2011). One concern with applying an HMM to the neural data is that the HMM formulation forces state transitions to follow an exponential distribution. In order to compare a discrete transition process to a diffusion-to-bound in the context of decision making, the distribution decision times given by the discrete model should be comparable to the bound-hit times defined by the diffusion model, which obey an inverse-Gaussian distribution (Fig. 2.5). Hidden semi-Markov Models (HSMMs) allow for non-exponential switch times, and can therefore more closely capture transitions between bursting and non-bursting activity in neu-

rons (Tokdar, Xi, Kelly & Kass 2010).

A diffusion-to-bound model has previously been applied to spike train data recorded in the early visual system to determine the stimulus selectivity of single neurons (Pillow, Paninski, Uzzell, Simoncelli & Chichilnisky 2005). The membrane voltage is determined through linear “integrate-and-fire” dynamics, where the membrane voltage follows a (leaky) drift-diffusion process where the slope is a linear function of the visual stimulus. The neuron emits a spike when the membrane voltage reaches the upper bound. The integrate-and-fire process in this case is distinct from the diffusion-to-bound model for decision-making presented in this chapter. The integrate-and-fire model emits a spike every time the upper bound is hit. In the decision-making case, the bound-hit time is inherently unknown (it is a convenient mathematical abstraction instead of a measurable event) as it is a cognitive variable represented within a network. The spike rate in LIP is hypothesized to represent the decision process, and therefore spikes are still emitted to encode the amount of evidence for a decision while the integrator is below the bound. Therefore, both the model fitting procedure used by Pillow et al. (2005) and the interpretation of the diffusion-to-bound process are distinct from the ramping model examined here.

2.2 Spike train models of decision-making dynamics in LIP

Here we present the two models of LIP spike train responses during decision-making. The models define an observed spike train as a Poisson process with rate determined by an unobserved (latent) noise process on each trial, yielding a doubly stochastic model. Both models are fit to each cell independently.

The models do not include an urgency signal, as introduced by (Churchland et al. 2008). The urgency signal is defined as the residuals between the data and an unbiased diffusion model. Typically, the urgency signal manifests as a positive (sometimes linear) drive in the spike rate towards the decision threshold. This forces decisions to occur faster, perhaps in an attempt to increase reward rates in a reaction-time task. The LIP responses we examined did not exhibit an urgency signal and in practice characterized as the upward deflection of responses for 0% coherence trials not conditioned on choice. The latter is flat in the dataset examined here (e.g., Fig. 2.16).

We use MCMC methods to sample from the posterior distribution of model parameters given the observed spike trains. These methods provide samples from the posterior distribution over the model parameters by alternately sampling the model parameters and the latent variables for every trial. The resulting samples of the model parameters approximate the posterior distribution of parameters given the data, marginalized over the latent variables. This enables the inclusion of uncertainty in the parameter estimates,

and to avoid any approximations that be required to fit these non-linear, non-Gaussian models with deterministic algorithms such as expectation-maximization. The model fits use 90% of the trials for each neuron (selecting 90% from each coherence group), and hold out 10% for computing cross-validation statistics such as predictive log-likelihood. The results of our analysis is consistent with the results obtained from DIC and Bayes Factors, but much more costly to compute because they required 10 folds of fitting and validation, and we did not include them here.

For computational convenience, the models consider spike trains in discrete time using bins of length Δ_t ($\Delta_t = 10$ ms bins here). However, simulations from the models produce spikes at a 0.2 ms resolution, assuming a homogeneous spike rate within 10 ms bins. Here we provide some notation for the models. We denote the entire set of spike trains (for the single neuron being modeled) as \mathbf{y} , and the spike counts at time bin t in trial j as $y_{j,t}$, with trials numbered 1 to N . Some parameters depend on the stimulus coherence in the trial. The models treat the stimulus coherence level as categorical, rather than assuming a functional form for the stimulus-dependent parameters. The total number of categories is denoted C (here, $C = 5$). The coherence for trial j is $c(j)$. Some model parameters depend on stimulus coherence, and a subscript (for example, p_c) denotes coherence dependence. The latent-variable analysis method allows trials to be of varying length: the length of trial j in number of discrete bins is denoted T_j .

2.2.1 Diffusion-to-bound (ramping) model

The spike rate in the ramping model follows a diffusion-to-bound process. The parameters of the model are $\Theta = \{\beta_{1:C}, x_0, \omega^2, \gamma\}$. The β and ω^2 parameters are the drift and diffusion terms (respectively) for the drift-diffusion process. The diffusion process starts at x_0 , and the bound height is determined by γ . The drift-diffusion process $x_{j,1:T_j}$ determines the spike rate for trial j . The full model can be described:

$$x_{j,1} = x_0 + \epsilon_{j,0} \quad (2.1)$$

$$x_{j,t+1} = x_{j,t} + \beta_{c(j)} + \epsilon_{j,t} \quad (2.2)$$

$$\epsilon_{j,t} \sim \mathcal{N}(0, \omega^2) \quad (2.3)$$

$$\tau_j = \begin{cases} \inf_t x_{j,t} \geq 1 & : \text{ if there exists } x_{j,1:T_j} \geq 1 \\ \infty & : \text{ otherwise} \end{cases} \quad (2.4)$$

$$y_{j,t} | t < \tau_j \sim \text{Poisson}(\log(1 + \exp(\gamma x_t)) \Delta_t) \quad (2.5)$$

$$y_{j,t} | t \geq \tau_j \sim \text{Poisson}(\log(1 + \exp(\gamma)) \Delta_t) \quad (2.6)$$

The graphical model is shown in Figure 2.3. The drift term β is the only coherence-dependent parameter (representing the strength of evidence in the stimulus). Spike rates are kept positive with the soft-rectification function $\log(1 + \exp(\gamma x_t))$. The spike rates observed in choice-selective LIP cells are sufficiently high to make the soft-rectification function close to a true linear rectifier.

In the parameterization given here, the latent diffusion process, $x_{j,1:T_j}$, does not stop at the bound, but the spike rate is held constant after the

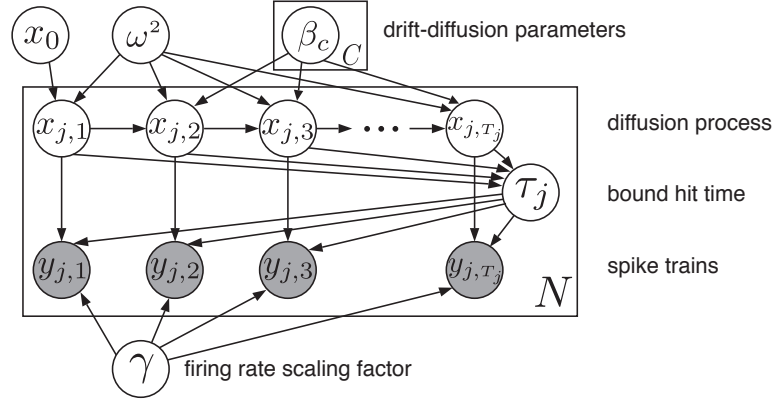


Figure 2.3: Graphical model for the diffusion-to-bound model. The box represents the N trials for a neuron. The diffusion-to-bound process begins each trial at the level x_0 and diffuses with noise w^2 and drift β . The diffusion process is bounded by setting τ_j equal to the first time t such that $x_{j,t} \geq 1$. The spike count at time t is drawn from a Poisson distribution where the rate is a bounded, scaled and linearly rectified version of $x_{j,t}$: the rate is $\log(1 + \exp(\gamma x_{j,t}))$ for $t < \tau_j$ and $\log(1 + \exp(\gamma))$ for $t \geq \tau_j$.

bound crossing time τ_j . This is equivalent to a model that stops the diffusion at bound hitting time, because spike rate is constant after this time in either case. Additionally, the bound-hitting time occurs when $x_{j,t}$ crosses a constant bound at 1; the bound height in terms of spike rate changes with the parameter γ . The transfer function makes bound in spikes per second equal to

$$\log(1 + \exp(\gamma)) \approx \gamma. \quad (2.7)$$

Therefore, the bound height is given by γ . This choice in parameterization not only simplifies model inference, but it makes the noise in the integration process independent of each neuron's firing rate *a priori*.

The prior distributions for the model parameters take the following

form

$$x_0 \sim \mathcal{N}(\mu_0, \sigma_0^2) \quad (2.8)$$

$$\beta \sim \mathcal{N}(\mu_\beta, \sigma_\beta^2) \quad (2.9)$$

$$\omega^2 \sim \text{Inv-Gamma}(\alpha_\omega, \beta_\omega) \quad (2.10)$$

$$\gamma \sim \text{Gamma}(\alpha_\gamma, \beta_\gamma) \quad (2.11)$$

We selected the following values for the priors

1. $\mu_0 = 0, \sigma_0 = 10$
2. $\mu_\beta = 0, \sigma_\beta = 0.1$
3. $\alpha_\omega = 0.02, \beta_\omega = 0.02$
4. $\alpha_\gamma = 2, \beta_\gamma = 0.05$

The prior over the maximum firing rate, γ , has mean 40 spikes/s with a standard deviation of 28.3, which covers the range of firing rates one expects to encounter in LIP. The typical values for the starting point of the diffusion process, x_0 , should lie between 0 and 1, because the diffusion process runs to a bound of 1. The broad Gaussian prior given for x_0 is nearly uniform over this region. The diffusion slope, β , corresponds to motion evidence (dot coherence), which was drawn for each trial from a distribution with mean 0%. We chose the standard deviation for the β prior by considering the range of realistic time-to-bound distributions: assuming $x_0 = 0$, then if $\beta > 0.1$ the

mean time-to-bound would be under 100 ms (extremely fast). Therefore, the standard deviation for the prior on β (0.1) places most of the prior probability mass on reasonable bound-hit times, without being too constricting. However, the specific choice of prior did not strongly affect results (not shown).

The MCMC section shows the prior parameter symbols instead of these specific values so that it is clear where the priors are placed.

2.2.2 Discrete stepping model

The stepping model allows the firing rate to occupy three discrete states: an initial state and two states that represent the possible decisions in a two-alternative forced choice task. The neuron maintains a single firing rate in each state, $(\alpha_{0,2})$, and the firing rates are constant across all trials. Transitions between states (or “steps”) occur instantaneously for simplicity. The model allows a single transition between states during a trial, although the model does not force a step to occur on every trial.

The step time for trial j is z_j (in discrete bins) and the state stepped to is d_j . If z_j is greater than the trial length, then no step occurs during the

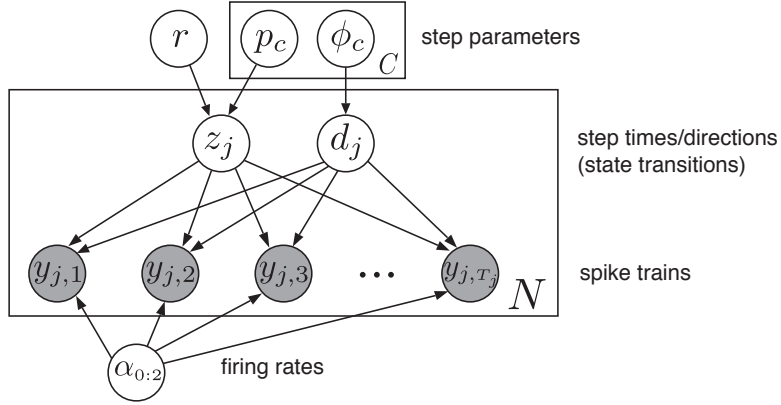


Figure 2.4: Graphical model for the discrete-switching hypothesis. The box represents the N trials for a neuron. The neuron begins each trial in the initial state with spike rate α_0 . Each trial has a step time, z_j , and a switch-to state, d_j . The observed spikes at each time t are drawn from a Poisson distribution with a rate depending on the state at t .

trial. The initial state is referred to as state 0.

$$z_j \sim \text{Negative Binomial}(p_{c(j)}, r) \quad (2.12)$$

$$P(d_j = 1) = \phi_{c(j)} \quad (2.13)$$

$$P(d_j = 2) = 1 - \phi_{c(j)} \quad (2.14)$$

$$y_{j,t} | t \leq z_j \sim \text{Poisson}(\alpha_0 \Delta_t) \quad (2.15)$$

$$y_{j,t} | t > z_j \sim \text{Poisson}(\alpha_{d_j} \Delta_t) \quad (2.16)$$

The graphical model is shown in Figure 2.4.

The negative binomial distribution is a discrete-time analogue of a gamma distribution. The distribution can be interpreted as the number of coin flips needed to get r heads, where $1 - p$ is the probability of a head (although r can take continuous values instead of just positive integers). If the

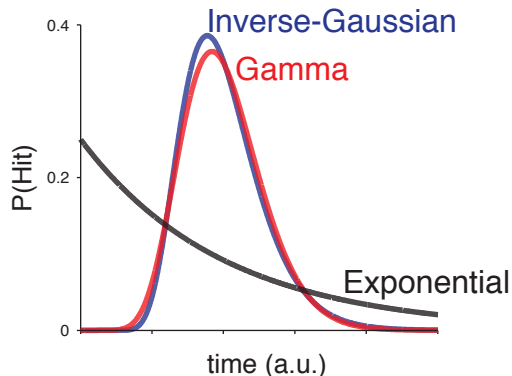


Figure 2.5: Comparison of the bound hit time distribution of a drift-diffusion process (inverse-Gaussian; blue curve) to a gamma distribution (red curve) and exponential distribution (black curve). The state transition times of a continuous time Hidden Markov model are exponentially distributed. The gamma distribution has the same mean and variance as the inverse-Gaussian, and the exponential shares only the mean.

negative binomial parameter r is set at 1, then this model becomes a more commonly used Hidden Markov model (HMM) with added restriction that the model cannot step out of states 1 or 2. We chose not to restrict the state transitions to be Markovian so that the step time distribution could be similarly shaped to the hit time distribution in the ramping model (Fig. 2.5). Additionally, trial lengths are finite, but z_j can take arbitrarily large values. Values larger than the trial length are interpreted to mean that no step took place on the trial.

The model parameters estimated are $\Psi = \{\alpha_{0:2}, r, \phi_{1:C}, p_{1:C}\}$.

The prior distributions for the model parameters take the following

form

$$\alpha_0 \sim \text{Gamma}(\alpha_\alpha, \beta_\alpha) \tag{2.17}$$

$$\alpha_1 \sim \text{Gamma}(\alpha_\alpha, \beta_\alpha) \tag{2.18}$$

$$p(\alpha_2|\alpha_1) \propto \mathbf{1}(\alpha_2 > \alpha_1)\text{Gamma}(\alpha_\alpha, \beta_\alpha) \tag{2.19}$$

$$p_c \sim \text{Beta}(\alpha_p, \beta_p) \tag{2.20}$$

$$\phi_c \sim \text{Beta}(\alpha_\phi, \beta_\phi) \tag{2.21}$$

$$r \sim \text{Gamma}(\alpha_r, \beta_r) \tag{2.22}$$

The truncated gamma prior enforced the inequality $\alpha_2 > \alpha_1$ in order to make state 2 always represent the up state and state 1 the down state (without this restriction, the model is unidentifiable because the state labels can be swapped). Otherwise, the prior distributions are independent.

We chose the following parameters

1. $\alpha_\alpha = 1, \beta_\alpha = 0.01$
2. $\alpha_p = 1, \beta_p = 1$
3. $\alpha_\phi = 1, \beta_\phi = 1$
4. $\alpha_r = 2, \beta_r = 1$

The distribution over firing rates is broadly tuned, with both a mean and standard deviation of 100 spikes/s. The beta distributions over ϕ and p are

uniform over the range of values $[0, 1]$. The prior on r peaks at 1, where the model becomes a HMM.

The MCMC section shows the prior parameter symbols instead of these specific values so that it is clear where the priors are placed.

2.3 Bayesian model comparison

The goal in modeling spike trains in this chapter is not just to provide estimates of the posterior distribution of model parameters, but to determine which model best describes actual neural data. Several metrics exist which can compare model fits using the output from the MCMC algorithms.

A classic tool in Bayesian model comparison is the Bayes factor. The quantity, denoted K , is the ratio of the marginal likelihoods ($P(\mathbf{y}|\mathcal{M})$) of the data given two models

$$K = \frac{p(\mathbf{y}|\mathcal{M}_d)}{p(\mathbf{y}|\mathcal{M}_s)} \quad (2.23)$$

$$p(\mathbf{y}|\mathcal{M}_d) = \frac{p(\mathbf{y}|\Theta, \mathcal{M}_d)p(\Theta|\mathcal{M}_d)}{p(\Theta|\mathbf{y}, \mathcal{M}_d)} \quad (2.24)$$

where \mathcal{M}_d and \mathcal{M}_s are the diffusion-to-bound and switching models respectively. The marginal likelihoods integrate over the parameter space in order to include the uncertainty in the parameter estimates. Chib & Jeliazkov (2001) and Chib (1995) describe how the output of the samplers can be used to calculate the marginal likelihood for each model using the samplers. Bayes factors are sensitive to the prior distribution, and the marginal likelihood is

only valid if proper prior distributions are used. Additionally, calculating the Bayes factors is often computationally intensive.

The posterior-predictive likelihood, PP , tests how well a model can predict new data (\mathbf{y}_{new}) given what the model can learn from a training data set (\mathbf{y}_{old}). This quantity can be simply estimated using the MCMC output.

$$PP = p(\mathbf{y}_{new}|\mathbf{y}_{old}, \mathcal{M}_d) = \int p(\mathbf{y}_{new}|\Theta, \mathcal{M}_d)p(\Theta|\mathbf{y}_{old}, \mathcal{M}_d)d\Theta \quad (2.25)$$

$$\approx \frac{1}{S} \sum_{s=1}^S p(\mathbf{y}_{new}|\Theta^{(s)}, \mathcal{M}_d) \quad (2.26)$$

This quantity is attractive, as it validates a model fit by using previously unseen data. Additionally, unlike the marginal likelihood, this quantity uses only a posterior distribution and is not heavily influenced by choice of prior. For the switching model, $p(\mathbf{y}_{new}|\Psi^{(s)}, \mathcal{M}_s)$ can be calculated in closed form. This quantity must be approximated for the diffusion model, which can be done by sampling

$$p(\mathbf{y}_{new}|\Theta^{(s)}, \mathcal{M}_d) = \int p(\mathbf{y}_{new}|\gamma^{(s)}, \mathbf{x}, \mathcal{M}_d)p(\mathbf{x}|\Theta^{(s)}, \mathcal{M}_d) \quad (2.27)$$

$$\mathbf{x}^{(m)} \sim p(\mathbf{x}|\Theta^{(s)}, \mathcal{M}_d) \quad (2.28)$$

$$p(\mathbf{y}_{new}|\Theta^{(s)}, \mathcal{M}_d) \approx \frac{1}{M} \sum_{m=1}^M p(\mathbf{y}_{new}|\gamma^{(s)}, \mathbf{x}^{(m)}, \mathcal{M}_d) \quad (2.29)$$

This process requires an M of several hundred to a few thousand to converge. However, this step is highly parallelizable.

Another model selection tool is the Deviance information criterion (DIC) (Spiegelhalter, Best, Carlin & van der Linde 2002). DIC constructs

a penalty for the log-likelihood of a model at the posterior mean, $\bar{\Theta}$ by examining the posterior distribution

$$DIC = D(\bar{\Theta}) + 2p_D \quad (2.30)$$

$$p_D = D(\bar{\Theta}) - D(\Theta) \quad (2.31)$$

$$D(\Theta) = -2\log(p(\mathbf{y}|\Theta)) + const \quad (2.32)$$

$$DIC = 2\log p(\mathbf{y}|\bar{\Theta}, \mathcal{M}_d) - 4\mathbb{E}_{\Theta|\mathbf{y}, \mathcal{M}_d} [\log p(\mathbf{y}|\Theta, \mathcal{M}_d)] \quad (2.33)$$

The expectation with respect to the posterior is approximated with the MCMC samples

$$\mathbb{E}_{\Theta|\mathbf{y}, \mathcal{M}_d} [\log p(\mathbf{y}|\Theta, \mathcal{M}_d)] \approx \frac{1}{S} \sum_{s=1}^S \log p(\mathbf{y}|\Theta^{(s)}, \mathcal{M}_d) \quad (2.34)$$

This quantity can be calculated for each model, and a lower DIC indicates more support for the model. The log likelihood becomes penalized when the support of the posterior distribution includes regions that do not describe the data well. This term is designed to act as an estimate of the effective number of parameters in the model. In fact, the term p_D converges exactly to the number of free parameters in some simple models. Because DIC provides the relative fit of each model to the data, the model comparison results relies only the difference in DIC values between the two models:

$$\Delta DIC = DIC_{ramping} - DIC_{stepping}. \quad (2.35)$$

Positive ΔDIC indicates support for the stepping model. Here, we interpret ΔDIC greater than 10 as strong support for the stepping model and a ΔDIC

less than -10 to indicate strong support for the ramping model (Burnham & Anderson 2002).

More well-known metrics, like the Akaike information criterion (AIC) the Bayesian information criterion (BIC), also use a penalized log-likelihood to compare model fits. These metrics are more difficult to apply to latent variable models: how does one include the latent variables in the count of model parameters? The latent spaces in the two models are starkly different: the stepping model's latent state can be described by two values (step time, and state stepped-to), while the state in the ramping model is much more complex.

We favor DIC for the purposes of comparing the latent-variable models discussed here, because it attempts to avoid these challenges by dynamically forming a penalty based on the shape of the posterior distribution. Additionally, DIC can be more widely applied to time series models than cross-validation. The models in this chapter divide the data into independent trials, but some models of interest to systems neuroscience consider one long stream of data (e.g., Turaga, Buesing, Packer, Dalglish, Pettit, Hausser & Macke 2013) that cannot be divided into independent chunks for cross-validation. DIC has been useful in these cases for comparing different latent dynamical models to financial time series data (Berg, Meyer & Yu 2004). Although DIC has not been widely applied in neuroscience, it has become a widely used model fitness criterion within many other domains (Spiegelhalter, Best, Carlin & van der Linde 2014). As with the posterior-predictive distribution, the expectation for

DIC is taken with respect to the posterior distribution. DIC is therefore less sensitive to the choice of prior than Bayes factors, and it permits the use of improper priors.

2.4 Markov chain Monte Carlo methods

The MCMC methods generate a sequence of samples from the posterior distribution of each model. We ran the MCMC algorithms for a total of 60 000 iterations and discarded the first 10 000 samples (the “burn in” period), to ensure that the Markov Chain had converged to its asymptotic distribution, the true posterior. The final result only took every 5th sample (a procedure known as “thinning” the chain) in order to reduce autocorrelation (or increase independence) between samples. Thus, we effectively obtained 10 000 samples from the posterior distribution of model parameters. These samples approximated the posterior distribution for performing model comparison analyses (Figs. 2.15, 2.12-2.6). All other analyses used the posterior mean (mean of the 10 000 samples) as a point estimate of the model parameters.

We implemented the sampling algorithms on a GPU using a combination of Matlab and CUDA. All sampling and analyses were performed on single desktop computer equipped with an Nvidia GTX Titan GPU and an Intel i7-4930K CPU (6 cores, 3.40 GHz). For the ramping (diffusion-to-bound) sampler running on a 500 trial dataset, the MCMC required 0.35 s to generate a sample. The stepping model sampler required 0.03 s per iteration. These times can be compared with an earlier CPU-only Matlab implementation of

the ramping MCMC algorithm, which required a prohibitively slow 17.6 s per sample. The use of both C/CUDA and a modern GPU produced an implementation suitable for running on a single desktop.

2.4.1 Ramping model

The sampler consists of two primary steps: (1) sampling the latent states given the previous value of the model parameters, and (2) sampling the model parameters given the newly sampled latent states. This gives a representation of the joint posterior of model parameters and latent states. By ignoring the latent state samples, we obtain an estimate of the posterior distribution over model parameters given the data alone, marginalizing over latent parameters. Sampling the latent states in each step results in a more efficient chain in which the model parameters are easily sampled using a fixed value of the latent states, especially because the data are divided into independent trials (Shephard & Pitt 1997, Yuan et al. 2012). We initialize the sampler by setting the bound height, γ , to the average spike rate in the final bin of all in-RF choice trials. The initial diffusion value x_0 is set to 0.1. The remaining parameters are set to the mode of the prior distribution.

2.4.1.1 Sampling the latent state with a particle filter

We obtain the s th sample of the latent state for all trials $\mathbf{x}^{(s)}$ conditioned on the previous sample of the parameters, $\Theta^{(s-1)}$, and the observed spikes \mathbf{y} . The latent state of each trial is independent of all other trials given

the model parameters and the data. Therefore, this section outlines the sampler for a single trial only, and the notation used here drops the subscript denoting trial number for simplicity of notation.

The posterior distribution over the latent states does not have a closed form. Instead, the posterior is decomposed using the Markovian structure of \mathbf{x} :

$$p(\mathbf{x}|\mathbf{y}, \Theta, \tau) = p(x_T|\Theta, y_{1:T}, \tau) \prod_{t=1}^{T-1} p(x_t|x_{t+1}, \Theta, y_{1:t}, \tau) \quad (2.36)$$

Using Bayes theorem, the right-hand side terms are computed as

$$p(x_t|x_{t+1}, \Theta, y_{1:t}, \tau) \propto p(x_{t+1}|x_t, \Theta, \tau)p(x_t|y_{1:t}, \Theta, \tau) \quad (2.37)$$

Therefore, if we can compute $p(x_t|y_{1:t}, \Theta, \tau)$ for $t = 1$ to T and $p(\tau|\Theta, y_{1:T})$, we can sample x_T from $p(x_T|\Theta, y_{1:T})$ and then work backwards, sampling x_t from $p(x_t|x_{t+1}, \Theta, y_{1:t}, \tau)$ for $t = \tau$ to 1 in order to obtain a sample from the complete posterior distribution, $p(\mathbf{x}, \tau|\mathbf{y}, \Theta)$.

We used a particle filter to approximate the distributions $p(x_t|y_{1:t}, \tau \geq t, \Theta^{(s-1)})$ for $t = 1$ to T (Gordon, Salmond & Smith 1993). In this algorithm, a set of M ($M = 200$ here) particles approximates the distribution of x_t for times $t = 1$ to T , considering only paths that have not crossed the bound by time t . At time t , particle k has position $\hat{x}_t^{(k)}$ and $w_t^{(k)}$ which form the distribution

$$p(x_t|y_{1:t}, \tau \geq t, \Theta^{(s-1)}) \approx \sum_{k=1}^M w_t^{(k)} \delta(x_t - \hat{x}_t^{(k)}) \quad (2.38)$$

where δ denotes the Dirac delta function. The weights must sum to 1 at each time ($\sum_{k=1}^M w_t^{(k)} = 1$). Additionally, we augmented the particle fil-

ter by tracking the distribution of the bound hit time, τ , relative to time t : $P(\tau < t|y_{1:t}, \Theta^{(s-1)})$ and $P(\tau = t|y_{1:t}, \Theta^{(s-1)})$, and $P(\tau > t|y_{1:t}, \Theta^{(s-1)})$. This formulation allowed the algorithm to account for the absorbing bound without needing to track each particle's history, which could result in a high percentage of degenerate particles. After obtaining these distributions over x_t and $\tau \geq t$ from time $t = 1$ to T , the sampler produced the values of $x_t^{(s)}$ working backwards from time $t = T$ to 1.

The particles were initialized by setting $\hat{x}_0^{(k)} = x_0$ and $w_0^{(k)} = \frac{1}{M}$ and $P(\tau > 0) = 1$. Particles are propagated through time using a sequential importance resampling (SIR) algorithm. Particle positions at time $t + 1$ are randomly sampled

$$\hat{x}_{t+1}^{(k)} \sim \pi(\hat{x}_{t+1}|\hat{x}_t^{(k)}, y_{t+1}, \Theta^{(s-1)}) \quad (2.39)$$

The particle weights are updated as

$$w_{t+1}^{(k)} \propto w_t^{(k)} \frac{p(y_{t+1}|\hat{x}_{t+1}^{(k)}, \Theta^{(s-1)})p(\hat{x}_{t+1}^{(k)}|\hat{x}_t^{(k)}, \Theta^{(s-1)})}{\pi(\hat{x}_{t+1}^{(k)}|\hat{x}_t^{(k)}, y_t, \Theta^{(s-1)})} \quad (2.40)$$

Because the particles tracked the distribution of x under the bound, the proposal distribution was a truncated Gaussian with mean and variance given by the drift-diffusion model

$$\pi(\hat{x}_{t+1}|\hat{x}_t^{(k)}, y_{t+1}, \Theta^{(s-1)}) \propto \mathbf{1}_{(-\infty, 1)}(\hat{x}_{t+1})\mathcal{N}(\hat{x}_{t+1}; \hat{x}_t + \beta^{(s-1)}, \omega^{2, (s-1)}). \quad (2.41)$$

The numerator terms are (from the model definition)

$$p(\hat{x}_{t+1}^{(k)}|\hat{x}_t^{(k)}, \Theta^{(s-1)}) = \mathcal{N}(\hat{x}_{t+1}^{(k)}; \hat{x}_t^{(k)} + \beta^{(s-1)}, \omega^{2, (s-1)}), \quad (2.42)$$

$$p(y_{t+1}|\hat{x}_{t+1}^{(k)}, \Theta^{(s-1)}) = \text{Poisson}\left(y_{t+1}; \log(1 + \exp(\gamma^{(s-1)}\hat{x}_{t+1}^{(k)}))\Delta_t\right) \quad (2.43)$$

The bound time, τ , is tracked through time. The model assumes the bound time was greater than 0 (i.e, $P(\tau > 0) = 1$). The bound-hit time probabilities are propagated through time, from $t = 1$ to T , with the following updates

$$P(\tau < t|y_{1:t}, \Theta^{(s-1)}) \propto p(y_t|\tau < t, \Theta^{(s-1)})P(\tau < t|y_{1:t-1}, \Theta^{(s-1)}) \quad (2.44)$$

$$P(\tau = t|y_{1:t}, \Theta^{(s-1)}) \propto p(y_t|\tau = t, \Theta^{(s-1)})P(\tau = t|y_{1:t-1}, \Theta^{(s-1)}) \quad (2.45)$$

$$P(\tau > t|y_{1:t}, \Theta^{(s-1)}) \propto p(y_t|\tau > t, y_{1:t-1}, \Theta^{(s-1)})P(\tau > t|y_{1:t-1}, \Theta^{(s-1)}) \quad (2.46)$$

and normalizing the probabilities so that

$$P(\tau > t|y_{1:t}, \Theta^{(s-1)}) + P(\tau = t|y_{1:t}, \Theta^{(s-1)}) + P(\tau < t|y_{1:t}, \Theta^{(s-1)}) = 1. \quad (2.47)$$

The particle distributions provide the updated probability distribution over τ relative to time t .

$$\begin{aligned} P(\tau = t|y_{1:t-1}, \Theta^{(s-1)}) &\approx P(\tau > t-1|y_{1:t-1}, \Theta^{(s-1)}) \\ &\quad \times \sum_{k=1}^M w_{t-1}^{(k)} (1 - \Phi(1; \hat{x}_t^{(k)} + \beta^{(s-1)}, \omega^{2,(s-1)})) \end{aligned} \quad (2.48)$$

$$p(y_t|\tau > t, y_{1:t-1}, \Theta^{(s-1)}) \approx \sum_{k=1}^M w_{t-1}^{*(k)} p(y_t|x_t = \hat{x}_t^{(k)}) \quad (2.49)$$

$$\text{where } w_{t+1}^{*(k)} = w_t^{(k)} \frac{p(\hat{x}_{t+1}^{(k)}|\hat{x}_t^{(k)}, \Theta^{(s-1)})}{\pi(\hat{x}_{t+1}^{(k)}|\hat{x}_t^{(k)}, y_t, \Theta^{(s-1)})} \quad (2.50)$$

$$\begin{aligned} P(\tau > t|y_{1:t-1}, \Theta^{(s-1)}) &\approx P(\tau > t-1|y_{1:t-1}, \Theta^{(s-1)}) \\ &\quad \times \left(\sum_{k=1}^M w_{t-1}^{(k)} \Phi(1; \hat{x}_t^{(k)} + \beta^{(s-1)}, \omega^{2,(s-1)}) \right) \end{aligned} \quad (2.51)$$

$\Phi(x; \mu, \sigma^2)$ denotes the normal cumulative density function with mean μ and variance σ^2 . The weights, $w_{t+1}^{*(k)}$ indicate the probability of particles $\hat{x}_{t+1}^{(k)}$ given only $y_{1:t}$ (without observing y_{t+1}).

After running the particle filter from $t = 1$ to T , the latent trajectory $x_{1:T}^{(s)}$ and $\tau^{(s)}$ is sampled. The sampler accomplished this by working backwards from time T to sample the value for $\tau^{(s)}$. Once $\tau^{(s)}$ was sampled, the sampler continues working backwards in time to establish the latent trajectory $x_{1:\tau}^{(s)}$.

With probability $P(\tau \leq T | y_{1:T})$ (calculated by the forward-pass), let $\tau^{(s)} \leq T$. Otherwise, let $\tau^{(s)} = \infty$, to signify that the diffusion process did not reach the bound on this trial. If instead $\tau^{(s)} \leq T$, then $\tau^{(s)}$ is sampled by working backwards from $t = T - 1$, then $t = T - 2$, and so on until an exact value for $\tau^{(s)}$ is found. The sampler works backwards setting $\tau^{(s)} \leq t$ with probability

$$P(\tau \leq t | \tau \leq t + 1, y_{1:t}, \Theta^{(s-1)}) = \frac{P(\tau \leq t | y_{1:t}, \Theta^{(s-1)})}{P(\tau \leq t | y_{1:t}, \Theta^{(s-1)}) + P(\tau = t + 1 | y_{1:t}, \Theta^{(s-1)})} \quad (2.52)$$

$$P(\tau = t + 1 | y_{1:t}, \Theta^{(s-1)}) \approx \sum_{k=1}^M w_t^{(k)} (1 - \Phi(1; \hat{x}_t^{(k)} + \beta^{(s-1)}, \omega^{(s-1)})) \times P(\tau > t | y_{1:t}, \Theta^{(s-1)}) \quad (2.53)$$

Otherwise, set $\tau^{(s)} = t + 1$.

If $\tau^{(s)} > T$, the sampler selects a value for $x_T^{(s)}$ from the particle set at

time T using the probability distribution

$$p(x_T|y_{1:T}, \tau^{(s)} > T, \Theta^{(s-1)}) \approx \sum_{k=1}^M w_T^{(k)} \delta(x_T - \hat{x}_T^{(k)}). \quad (2.54)$$

and then works backwards in time sampling $x_{1:T-1}^{(s)}$ (sampling first $x_{T-1}^{(s)}$, then $x_{T-2}^{(s)}$, and so on) as described below.

If instead $\tau^{(s)} < T$, the value of $x_{\tau-1}^{(s)}$ is set to one of the particles $\hat{x}_t^{(k)}$ where $t = \tau^{(s)} - 1$ by sampling from the distribution

$$\begin{aligned} p(x_t|\tau^{(s)}, y_{1:T}, \Theta^{(s-1)}) &\propto p(\tau^{(s)} = t + 1|x_t, \Theta^{(s-1)})p(x_t|y_{1:t}, \Theta^{(s-1)}) \\ &\approx \sum_{k=1}^M w_t^{(k)} \delta(x_t - \hat{x}_t^{(k)}) \left(1 - \Phi(1; \hat{x}_t^{(k)} + \beta^{(s-1)}, \omega^{2,(s-1)})\right). \end{aligned} \quad (2.55)$$

$$(2.56)$$

The remaining trajectory, $x_{1:\tau-2}^{(s)}$, is obtained by sampling backwards through time.

Backwards sampling again requires the particles. The value of $x_{t-1}^{(s)}$ given $x_t^{(s)}$ was sampled using the approximated distribution

$$\begin{aligned} p(x_{t-1}|x_t^{(s)}, y_{1:t-1}, \Theta^{(s-1)}, \tau > t - 1) &\propto p(x_t^{(s)}|x_{t-1}, \Theta^{(s-1)}) \\ &\times p(x_{1:t-1}|y_{1:t-1}, \Theta^{(s-1)}, \tau > t - 1) \end{aligned} \quad (2.57)$$

$$\begin{aligned} &\approx \sum_{k=1}^M \delta(x_{t-1} - \hat{x}_{t-1}^{(k)}) w_{t-1}^{(k)} \\ &\mathcal{N}(x_t^{(s)}; \hat{x}_{t-1}^{(k)} + \beta^{(s-1)}, \omega^{2,(s-1)}) \end{aligned} \quad (2.58)$$

After obtaining the latent values all the way back to $x_1^{(s)}$ and if $\tau^{(s)} \leq T$, the value of $x_{\tau^{(s)}}^{(s)}$ is sampled from the truncated normal distribution

$$p(x_\tau | x_{\tau-1}^{(s)}, \Theta^{(s-1)}) \propto \mathbf{1}_{[1, \infty)} \mathcal{N}(x_\tau; x_{\tau-1}^{(s)} + \beta, \omega^2) \quad (2.59)$$

where $\mathbf{1}$ is the indicator function

$$\mathbf{1}_{[1, \infty)}(x) = \begin{cases} 1 & \text{if } x \in [1, \infty) \\ 0 & \text{otherwise} \end{cases} \quad (2.60)$$

After $\tau^{(s)}$, the observations (spikes) no longer depend on the new values of the latent state. This independence also means that the sampler can drop $x_{\tau^{(s)}+1:T}^{(s)}$ from the model parameter sampling step (everything about the spike rate is given by $x_{1:\tau^{(s)}}^{(s)}$). Dropping these terms increases sampler efficiency.

2.4.1.2 Sampling the model parameters

With the s th sample of the latent states, the algorithm samples a new set of the model parameters. We first sample the parameters $x_0, \beta_{1:C}$, and ω^2 . Applying Bayes' rule to the posterior gives

$$p(x_0, \beta_{1:C}, \omega^2 | \mathbf{x}^{(s)}, \mathbf{y}) \propto p(\mathbf{x}^{(s)} | x_0, \beta_{1:C}, \omega^2) p(x_0, \beta_{1:C}, \omega^2) \quad (2.61)$$

in order to compute the distributions. The parameters $x_0, \beta_{1:C}$, and ω^2 conditioned on $\mathbf{x}^{(s)}$ are independent of the observations and γ . The model defines the latent paths as a simple linear-Gaussian process, and therefore the parameters can be sampled exactly using Gibbs steps. The model definition of $p(\mathbf{x}^{(s)} | x_0, \beta_{1:C}, \omega^2)$ states that the differences $(x_{j,t} - x_{j,t-1})$ are normally distributed with mean $\beta_{c(j)}$ and variance ω^2 .

The values of $\beta_{1:C}^{(s)}$ and $x_0^{(s)}$ can be sampled from independent Gaussian distributions given the previous diffusion variance term $\omega^{2,(s-1)}$.

$$\beta_c^{(s)} | \mathbf{x}^{(s)}, \mathbf{y}, \omega^{2,(s-1)} \sim \mathcal{N}(B \cdot A^{-1}, A^{-1}) \quad (2.62)$$

$$A = \frac{1}{\sigma_\beta^2} + \frac{1}{\omega^{2,(s-1)}} \sum_{j \in \{i:c(i)=c\}} ((T_j \wedge \tau_j^{(s)}) - 1) \quad (2.63)$$

$$B = \frac{\mu_\beta}{\sigma_\beta^2} + \frac{1}{\omega^{2,(s-1)}} \sum_{j \in \{i:c(i)=c\}} \sum_{t=2}^{T_j \wedge \tau_j^{(s)}} (x_{j,t}^{(s)} - x_{j,t-1}^{(s)}) \quad (2.64)$$

$$x_0^{(s)} | \mathbf{x}^{(s)}, \mathbf{y}, \omega^{2,(s-1)} \sim \mathcal{N}(D \cdot C^{-1}, C^{-1}) \quad (2.65)$$

$$C = \frac{1}{\sigma_0^2} + \frac{N}{\omega^{2,(s-1)}} \quad (2.66)$$

$$D = \frac{\mu_0}{\sigma_0^2} + \frac{1}{\omega^{2,(s-1)}} \sum_{j=1}^N x_{j,1}^{(s)} \quad (2.67)$$

$$T_j \wedge \tau_j^{(s)} = \min(T_j, \tau_j^{(s)}) \quad (2.68)$$

The term $T_j \wedge \tau_j^{(s)}$ signifies the bound crossing time (the “effective” length of the latent trajectory $x_{j,1:T_j}^{(s)}$). The spike rate is no longer dependent on the diffusion process once the process has crossed the bound (for $t \geq \tau_j^{(s)}$, $y_{j,t}$ is independent of the value of $x_{j,t}^{(s)}$). Therefore, the sampler only needs to consider the values of $x_{j,t}$ for $t \leq \tau_j$ in order to sample from the posterior (replacing $T_j \wedge \tau_j^{(s)}$ with T_j results in a correct, but slower, sampler).

The next step is to sample ω^2 given the newly generated samples of $\beta_{1:C}$ and x_0 .

$$\omega^{2,(s)} | \beta_{1:C}^{(s)}, x_0^{(s)}, \mathbf{x}^{(s)}, \mathbf{y} \sim \text{Inv-Gamma}(E, F) \quad (2.69)$$

$$E = \alpha_\omega + \frac{1}{2} \sum_{j=1}^N (T_j \wedge \tau_j^{(s)}) \quad (2.70)$$

$$F = \beta_\omega + \frac{1}{2} \sum_{j=1}^N \left[\left(x_{j,1}^{(s)} - x_0^{(s)} \right)^2 + \sum_{t=2}^{T_j \wedge \tau_j} \left(x_{j,t}^{(s)} - (x_{j,t-1}^{(s)} + \beta_{c(j)}^{(s)}) \right)^2 \right] \quad (2.71)$$

β_c depends only on trials of stimulus coherence c , while x_0 and ω^2 are coupled to all trials.

Even though the bound height parameter, γ , is independent of all other parameters given the latent states, there exists no single closed-form Gibbs step to sample $\gamma^{(s)}$. Instead, we generated samples via a Metropolis-Hastings (MH) step. The MH algorithm samples $\gamma^{(s)}$ with the following steps

- 1 Sample $\gamma^* \sim q(\gamma | \gamma^{(s-1)})$ where q is an arbitrary proposal distribution.
- 2 Sample $u \sim U([0, 1])$.
- 3 $\gamma^{(s)} = \begin{cases} \gamma^* & , \quad u < \alpha \\ \gamma^{(s-1)} & , \quad \text{otherwise} \end{cases}$ where $\alpha = \min \left(1, \frac{p(\gamma^* | \mathbf{y}, \mathbf{x}^{(s)}) q(\gamma^{(s-1)} | \gamma^*)}{p(\gamma^{(s-1)} | \mathbf{y}, \mathbf{x}^{(s)}) q(\gamma^* | \gamma^{(s-1)})} \right)$

A Langevin step was used for the proposal distribution (Roberts &

Stramer 2002)

$$q(\gamma^* | \gamma^{(s-1)}, \mathbf{y}, \mathbf{x}^{(s)}) = \mathcal{N} \left(\gamma^*; \gamma^{(s-1)} + \epsilon^2 \frac{1}{2} \mathbf{G}^{-1}(\gamma^{(s-1)}) \frac{d}{d\gamma} \mathcal{L}(\gamma^{(s-1)}), \epsilon^2 \mathbf{G}^{-1}(\gamma^{(s-1)}) \right) \quad (2.72)$$

$$\begin{aligned} \mathcal{L}(\gamma) &= \log p(\mathbf{y} | \mathbf{x}^{(s)}, \gamma) + \log p(\gamma) \\ &= \sum_{j=1}^N \sum_{t=1}^{T_j} \log p(y_{j,t} | x_{j,t}^{(s)}, \gamma) + \log p(\gamma) \end{aligned} \quad (2.73)$$

$$\mathbf{G}(\gamma) = -\mathbb{E}_{\mathbf{y} | \gamma, \mathbf{x}^{(s)}} \left[\frac{d^2}{d\gamma^2} \mathcal{L}(\gamma) \right] \quad (2.74)$$

This proposal uses the Fisher information plus the second derivative of the log prior, $\mathbf{G}(\gamma)$, to condition the step, as suggested by (Girolami & Calderhead 2011). The result of this conditioning made selecting an effective value for ϵ simple. We set ϵ to a small initial value (0.1) and slowly raised it during the burn-in period to a larger value of 1.

Labeling the soft-rectifying function

$$h(x, \gamma) = \log(1 + \exp(x\gamma)) \quad (2.75)$$

the derivative of the log likelihood is

$$\frac{d}{d\gamma} \mathcal{L}(\gamma) = \frac{d}{d\gamma} [\log p(\mathbf{y} | \mathbf{x}^{(s)}, \gamma) + \log p(\gamma)] \quad (2.76)$$

$$\begin{aligned} &= \frac{d}{d\gamma} \left[\sum_{j=1}^N \sum_{t=1}^{T_j} \left(-h(x_{j,t}^{(s)}, \gamma) \Delta_t + y_{j,t} \log(h(x_{j,t}^{(s)}, \gamma)) \right) \right. \\ &\quad \left. + (\alpha_\gamma - 1) \log(\gamma) - \gamma \beta_\gamma + \text{const} \right] \end{aligned} \quad (2.77)$$

$$= \sum_{j=1}^N \sum_{t=1}^{T_j} \left(h'(x_{j,t}^{(s)}, \gamma) \left(-\Delta_t + y_{j,t} \frac{1}{h(x_{j,t}^{(s)}, \gamma)} \right) \right) + \frac{\alpha_\gamma - 1}{\gamma} - \beta_\gamma \quad (2.78)$$

where

$$h'(x, \gamma) = \frac{d}{d\gamma} h(x, \gamma) = \frac{x}{1 + \exp(-x\gamma)} \quad (2.79)$$

The Fisher information combined with the prior Hessian is calculated

as

$$\begin{aligned} \mathbf{G}(\gamma) = & -\mathbb{E}_{\mathbf{y}|\gamma, \mathbf{x}^{(s)}} \left[\frac{d}{d\gamma} \left(\sum_{j=1}^N \sum_{t=1}^{T_j} \left(h'(x_{j,t}^{(s)}, \gamma^{(s-1)}) \left(-\Delta_t + y_{j,t} \frac{1}{h(x_{j,t}^{(s)}, \gamma)} \right) \right) \right) \right. \\ & \left. + \frac{\alpha_\gamma - 1}{\gamma} - \beta_\gamma \right] \end{aligned} \quad (2.80)$$

$$\begin{aligned} = & -\mathbb{E}_{\mathbf{y}|\gamma, \mathbf{x}^{(s)}} \left[\sum_{j=1}^N \sum_{t=1}^{T_j} \left(h''(x_{j,t}^{(s)}, \gamma) \left(-\Delta_t + y_{j,t} \frac{1}{h(x_{j,t}^{(s)}, \gamma)} \right) \right. \right. \\ & \left. \left. - y_{j,t} \left(\frac{h'(x_{j,t}^{(s)}, \gamma)}{h(x_{j,t}^{(s)}, \gamma)} \right)^2 \right) - \frac{\alpha_\gamma - 1}{(\gamma)^2} \right] \end{aligned} \quad (2.81)$$

$$\begin{aligned} = & -\sum_{j=1}^N \sum_{t=1}^{T_j} \left(h''(x_{j,t}^{(s)}, \gamma) \left(-\Delta_t + \mathbb{E}_{\mathbf{y}|\gamma, \mathbf{x}^{(s)}} [y_{j,t}] \frac{1}{h(x_{j,t}^{(s)}, \gamma)} \right) \right. \\ & \left. - \mathbb{E}_{\mathbf{y}|\gamma, \mathbf{x}^{(s)}} [y_{j,t}] \left(\frac{h'(x_{j,t}^{(s)}, \gamma)}{h(x_{j,t}^{(s)}, \gamma)} \right)^2 \right) + \frac{\alpha_\gamma - 1}{(\gamma)^2} \end{aligned} \quad (2.82)$$

Noting that $\mathbb{E}_{\mathbf{y}|\gamma, \mathbf{x}^{(s)}} [y_{j,t}] = h(x_{j,t}^{(s)}, \gamma) \Delta_t$

$$\mathbf{G}(\gamma) = \frac{\alpha_\gamma - 1}{(\gamma)^2} + \sum_{j=1}^N \sum_{t=1}^{T_j} \left(\Delta_t \frac{\left(h'(x_{j,t}^{(s)}, \gamma) \right)^2}{h(x_{j,t}^{(s)})} \right) \quad (2.83)$$

2.4.2 Stepping model

As with the ramping model, the sampler consists of two main steps: (1) sampling the latent stepping process given the parameters and (2) sampling the parameters given the new latent states. We initialize the chain by setting the rate parameters based on average spike rates: α_0 is set to the average spike rate over all trials in the first bin, α_2 is set to the average spike rate in the final bin of all in-RF choice trials, and α_1 is set to $\frac{1}{2}\alpha_2$. The initial values for all the ϕ and p parameters is 0.5 (the mean of the prior). The final parameter, r , is initially set to the mode of the prior distribution.

2.4.2.1 Sampling the latent steps

The s th sample of the latent states is obtained using the $(s - 1)$ th sample of the model parameters:

$$(\mathbf{z}, \mathbf{d})^{(s)} \sim \mathbf{z}, \mathbf{d} | \Psi^{(s-1)}, \mathbf{y} \quad (2.84)$$

With the model parameters fixed, each trial becomes independent. Unfortunately, there is no closed form distribution for this step. However, in this simple stepping case, truncating the distribution over z results in a finite discrete distribution that can be sampled. We set z to a maximum of 1500, which is 15 times longer than the longest trial and much longer than an interpretable step time in the decision-making task. Therefore, the truncated distribution is an extremely close approximation to the true distribution. The step time z_j can be greater than T_j , which is interpreted to mean no step occurred during

the trial. For $z_j = 1$ to 1500 and $d_j = 1$ to 2, the probability of the latent states is

$$\begin{aligned}
p(z_j, d_j | y_{j,1:T_j}, \Psi^{(s-1)}) &\propto p(y_{j,1:T_j} | z_j, d_j, \alpha^{(s-1)}) p(z_j | r^{(s-1)}, p_{c(j)}^{(s-1)}) p(d_j | \phi_{c(j)}^{(s-1)}) \\
&= p(z_j | r^{(s-1)}, p_{c(j)}^{(s-1)}) p(d_j | \phi_{c(j)}^{(s-1)}) \prod_{t=1}^{T_j} p(y_{j,t} | z_j, d_j, \alpha^{(s-1)})
\end{aligned} \tag{2.85}$$

The distributions on the right side are calculated using the model definition for the Poisson observation $p(y_{j,1:t} | z_j, d_j, \alpha^{(s-1)})$, the negative binomially distributed step time $p(z_j | r^{(s-1)}, p_{c(j)}^{(s-1)})$, and the Bernoulli state choice $p(d_j | \phi_{c(j)}^{(s-1)})$. Once all 3000 possible values are calculated, the joint distribution was normalized and sampled directly.

2.4.2.2 Sampling the model parameters

The s th sample of the latent states is used to draw the next value for the model parameters. This step consisted of two parts: the first part samples $\alpha_{0:2}$, p , and ϕ . The final step samples the parameter r with the new values of $p_{1:C}$ and z .

The parameters $\alpha_{0:2}$, p , and ϕ are all sampled independently with Gibbs' steps.

$$\phi_c^{(s)} | \mathbf{d}^{(s)} \sim \text{Beta}(\alpha_\phi + D_c, \beta_\phi + N - D_c)$$

$$D_c = \sum_{j \in \{i:c(i)=c\}} \mathbf{1}(d_j^{(s)} = 2) \quad (2.86)$$

$$p_c^{(s)} | \mathbf{z}^{(s)}, r^{(s-1)} \sim \text{Beta} \left(\alpha_p + \sum_{j \in \{i:c(i)=c\}} z_j^{(s)}, \beta_p + r^{(s-1)} \sum_{j \in \{i:c(i)=c\}} 1 \right) \quad (2.87)$$

$$\alpha_0^{(s)} | \mathbf{z}^{(s)}, \mathbf{d}^{(s)}, \mathbf{y} \sim \text{Gamma} \left(\alpha_\alpha + \sum_{j=1}^N \sum_{t=1}^{z_j^{(s)}} y_{j,t}, \beta_\alpha + \sum_{j=1}^N z_j^{(s)} \right) \quad (2.88)$$

$$\alpha_1^{(s)} | \mathbf{z}^{(s)}, \mathbf{d}^{(s)}, \mathbf{y} \sim \text{Gamma}(\alpha_1, \beta_1) \quad (2.89)$$

$$\alpha_2^{(s)} | \alpha_1^{(s)}, \mathbf{z}^{(s)}, \mathbf{d}^{(s)}, \mathbf{y} \sim \mathbf{1}(\alpha_2^{(s)} > \alpha_1^{(s)}) \text{Gamma}(\alpha_2, \beta_2) \quad (2.90)$$

$$\alpha_i = \alpha_\alpha + \sum_{j=1}^N \sum_{t=z_j^{(s)}+1}^{T_j} \mathbf{1}(d_j^{(s)} = i) y_{j,t} \quad (2.91)$$

$$\beta_i = \beta_\alpha + \sum_{j=1}^N \mathbf{1}(d_j^{(s)} = i) (T_j - z_j^{(s)}) \quad (2.92)$$

The truncated gamma distribution on $\alpha_2^{(s)}$ enforces that $\alpha_2^{(s)} > \alpha_1^{(s)}$ holds (see model prior section).

No simple closed form distribution exists for the posterior over the negative binomial parameter r . A Metropolis adjusted Langevin algorithm generates samples of r using the following proposals (see the ramping sampler

for a brief description of Metropolis-Hastings proposals):

$$q(r^*|r^{(s-1)}, \mathbf{y}, \mathbf{z}^{(s)}, p_{1:C}^{(s)}) = \mathcal{N}\left(r^*; r^{(s-1)} + \epsilon^2 \frac{1}{2} \frac{d}{dr} \mathcal{L}(r^{(s-1)}), \epsilon^2\right) \quad (2.93)$$

$$\begin{aligned} \mathcal{L}(r^{(s-1)}) &= \log p(\mathbf{z}^{(s)}|r^{(s-1)}, p_{1:C}^{(s)}) + \log p(r^{(s-1)}) \\ &= \sum_{j=1}^N \log p(z_j^{(s)}|r^{(s-1)}, p_{c(j)}^{(s)}) + \log p(r^{(s-1)}) \end{aligned} \quad (2.94)$$

We use a simple automatic procedure select a value for ϵ . Initially, ϵ is set to a small value (0.05), and during the burn-in period, the sampler raises or lowers ϵ (by multiplicative factors of 1.25 or 0.75) until the sampler stabilizes to an acceptance rate in the range of 30 – 70%.

The derivative can be calculated

$$\begin{aligned} \mathcal{L}(r^{(s-1)}) &= \sum_{j=1}^N \log p(z_j^{(s)}|r^{(s-1)}, p_{c(j)}^{(s)}) + \log p(r) \\ &= \sum_{j=1}^N \left[\log \Gamma(z_j^{(s)} + r^{(s-1)}) - \log \Gamma(z_j^{(s)} + 1) - \log \Gamma(r^{(s-1)}) - z_j^{(s)} \log(p_{c(j)}^{(s)}) \right. \\ &\quad \left. + r^{(s-1)} \log(1 - p_{c(j)}^{(s)}) \right] + (\alpha_r - 1) \log(r^{(s-1)}) - r^{(s-1)} \beta_r + \text{const} \end{aligned} \quad (2.95)$$

$$\frac{d}{dr} \mathcal{L}(r^{(s-1)}) = \sum_{j=1}^N \left[\psi(z_j^{(s)} + r^{(s-1)}) - \psi(r^{(s-1)}) + \log(1 - p_{c(j)}^{(s)}) \right] + \frac{\alpha_r - 1}{r^{(s-1)}} - \beta_r \quad (2.96)$$

where Γ and ψ are the gamma and digamma functions respectively. This sampling step uses the entire set of coherence-dependent p parameters.

2.4.3 Visualizing samples from the posterior distribution

To provide more intuition about what the MCMC methods produced, Figure 2.6 shows the posterior distributions for the model parameters estimated from the MCMC for one example LIP cell. These estimated posteriors are simply histograms of the samples from the Markov chain. For visualization purposes, these distributions are the marginal posterior distributions for each parameter given the set of spike trains. However, the samples came from the joint posterior. Figure 2.7 shows samples from the posterior distribution over the latent firing rates for the stepping model for 15 example trials from one cell. Figure 2.8 shows samples from the posterior distribution over the firing rates for the ramping model for the same 15 trials.

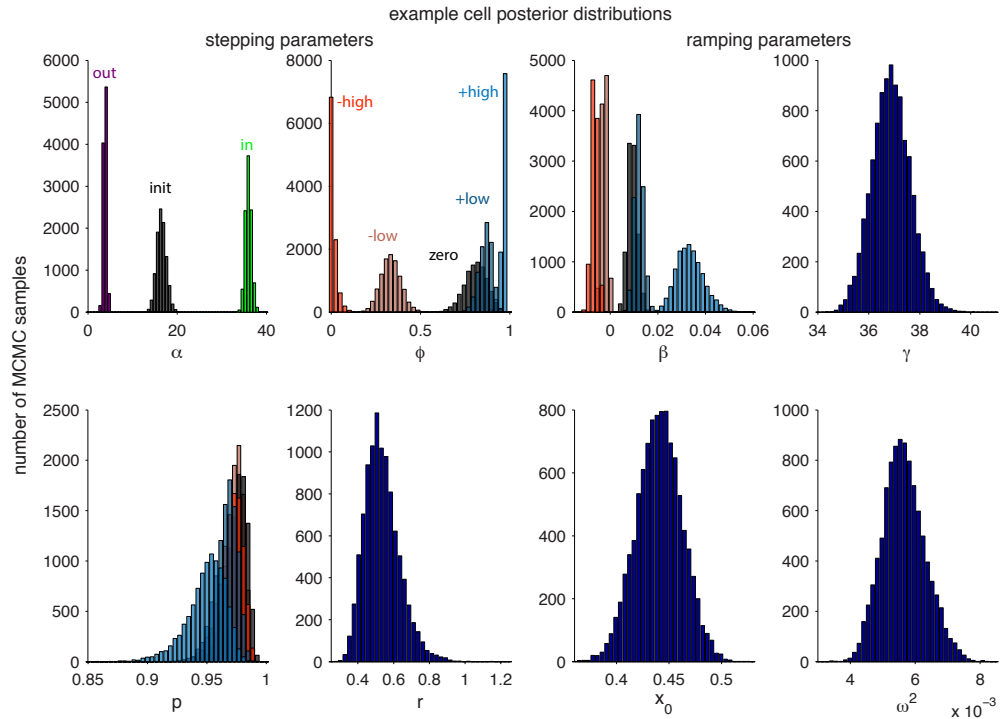


Figure 2.6: Example estimates of the posterior distributions for all the parameters for both models (stepping model left 2 columns, ramping model right 2 columns) for a single example LIP cell. For the ϕ , p , and β parameters, distributions for all 5 motion coherence levels are shown.

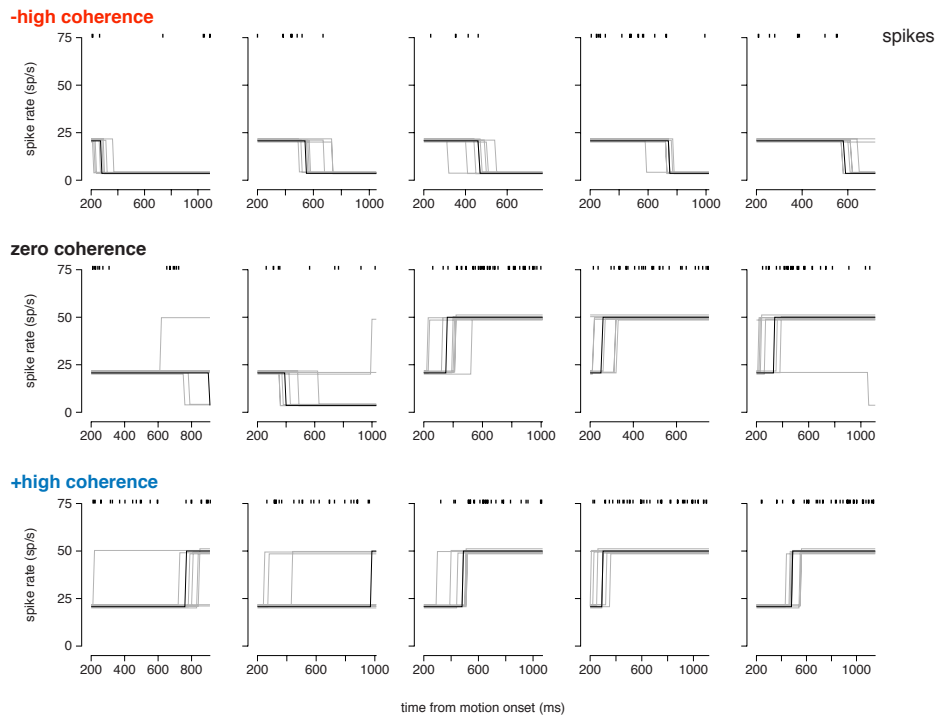


Figure 2.7: Each plot shows samples from the posterior distribution over the latent firing rates in the stepping model (gray traces) for five example trials at each of three motion coherence levels. The black trace shows the median step time estimate (Section 2.5.4.2) with the posterior mean spike rates for the states. Rasters above each plot indicate the spike times on each trial. (Top row) High negative coherence (strong out-RF motion) trials. (Middle row) 0% coherence trials. (Bottom row) high positive coherence (strong in-RF motion) trials. The trials shown are from cell 10.

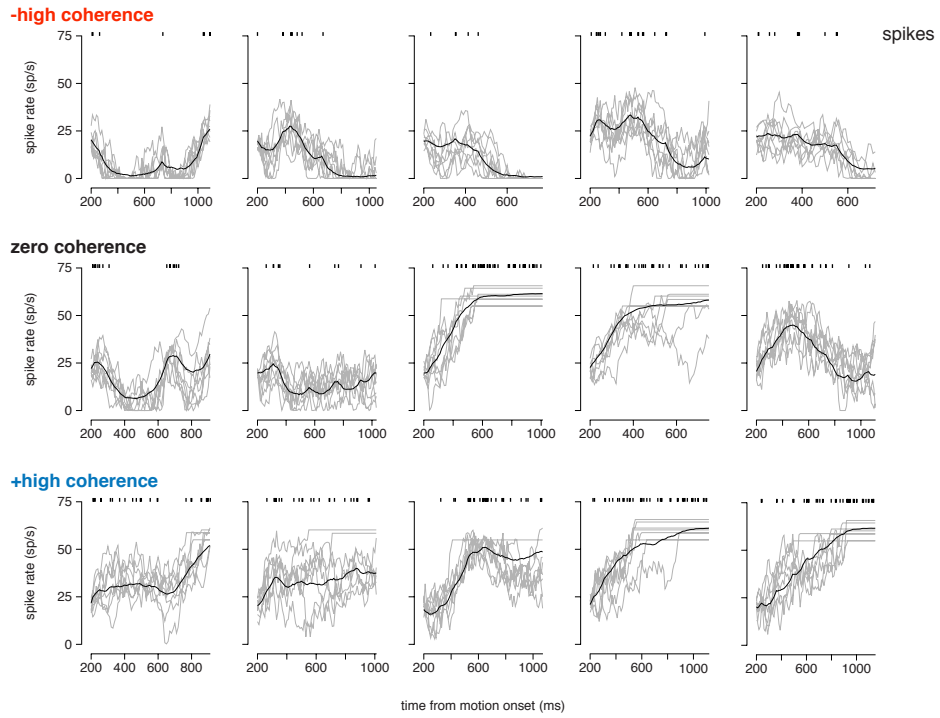


Figure 2.8: Each plot shows samples from the posterior distribution over the latent firing rates in the ramping model (gray traces) for the same trials as shown in Figure 2.7. The black trace shows the trial’s posterior mean firing rate. Rasters above each plot indicate the spike times on each trial. (Top row) High negative coherence (strong out-RF motion) trials. (Middle row) 0% coherence trials. (Bottom row) high positive coherence (strong in-RF motion) trials.

2.4.4 Analyzing simulated spike trains

To assess the ability of the MCMC methods to successfully fit and identify models, we simulated trials from the ramping and stepping models and applied the MCMC methods to the simulated data. We examined the ability of the sampler to recover true parameters and the DIC metric to identify the correct model.

2.4.4.1 Simulations: evaluating sampler efficiency

To evaluate how well the MCMC algorithms were mixing (i.e., how quickly the chain could start to produce effectively independent samples from the true posterior), we calculated the autocorrelations in the parameter samples from the chain. This is a basic visual tool that can indicate if the chain has reached the true posterior, and how independent the samples are (with lower autocorrelation meaning more independence). Testing this metric using all 50 000 samples from the MCMC output indicates that the ramping parameters ω^2 and γ , and the stepping model parameter r had a high autocorrelation. (Fig. 2.9, top). However, thinning the chain by taking only every 5th sample (described in Section 2.4), largely eliminated the autocorrelation (Fig. 2.9, bottom). The autocorrelation from the MCMC output of the fit to the cell shown in Figure 2.17 is shown in Figure 2.10 for a comparison.

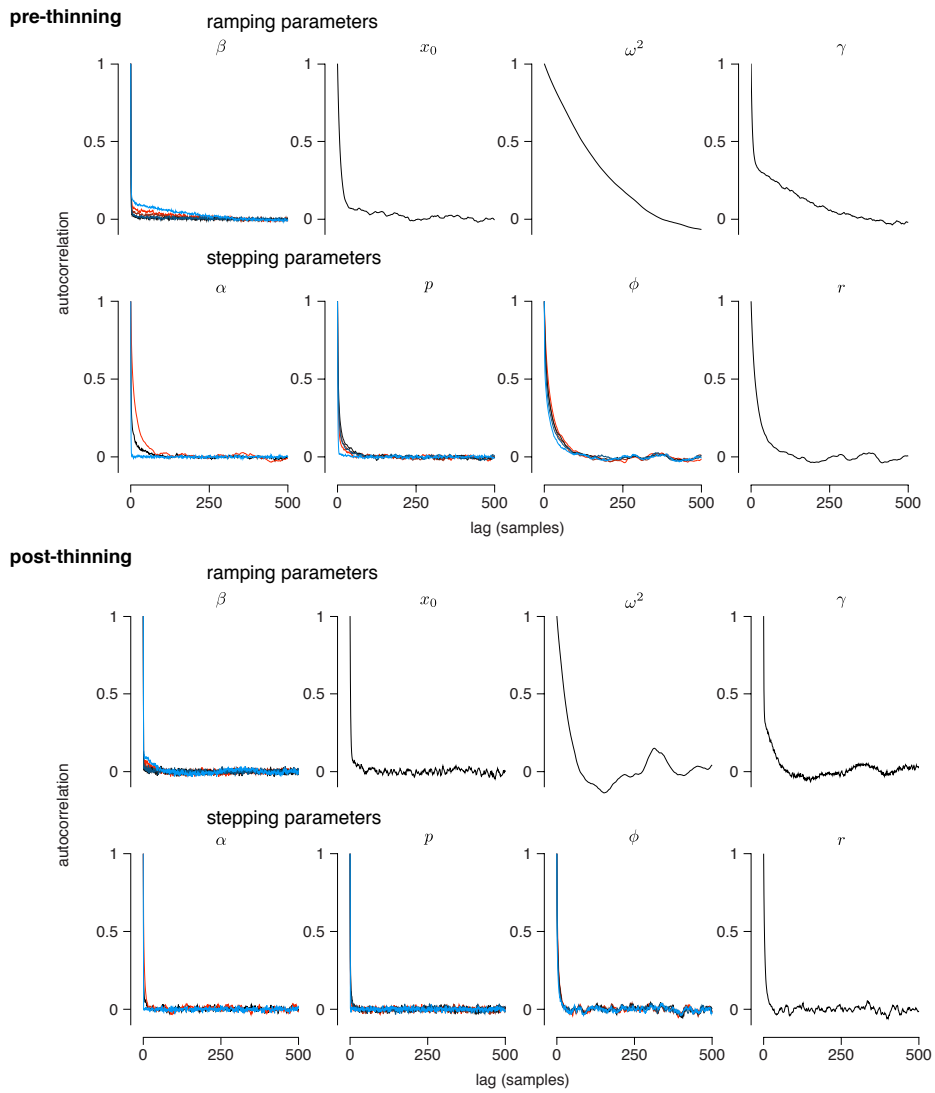


Figure 2.9: Autocorrelation plots of all the parameter samples from the MCMC output for all the parameters for both models for simulated data. The abscissa is in units of iterations of the MCMC algorithm. The top two rows show the autocorrelation of the samples from the original MCMC output. The bottom two rows show the autocorrelation after thinning the chain, taking only every 5th sample. The chain was run on datasets containing 500 trials, with 100 trials of each of the 5 possible motion coherence levels.

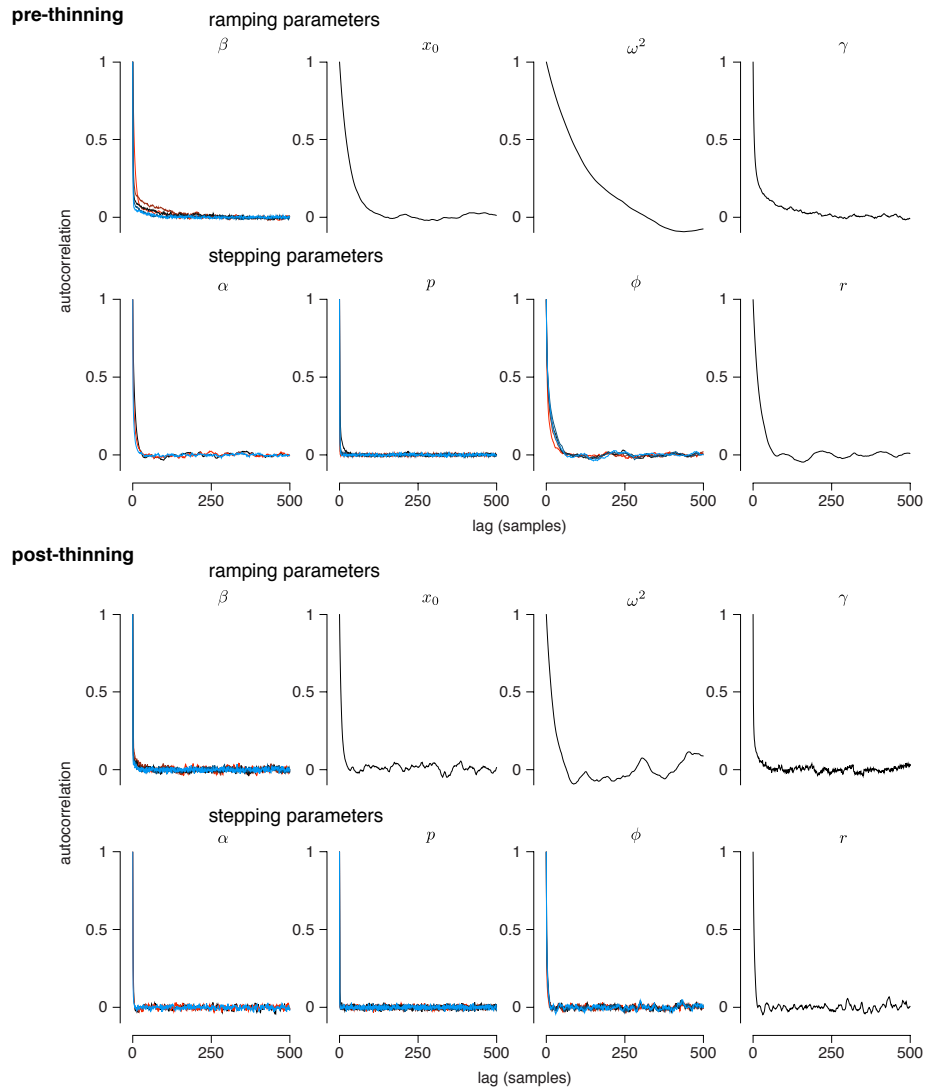


Figure 2.10: Autocorrelation plots of all the parameter samples from the MCMC for all the parameters for both models for the cell shown in Fig. 2.17. The result is similar to the simulation in Fig. 2.9

2.4.4.2 Estimating parameters in simulations

We confirmed that the samplers estimate the true parameters from simulated data (i.e., is the estimator consistent?). We ran the MCMC on various amounts of simulated data and compared the posterior mean estimates of the parameters to the true parameters. Figure 2.11 shows that the error tends towards zero with more data.

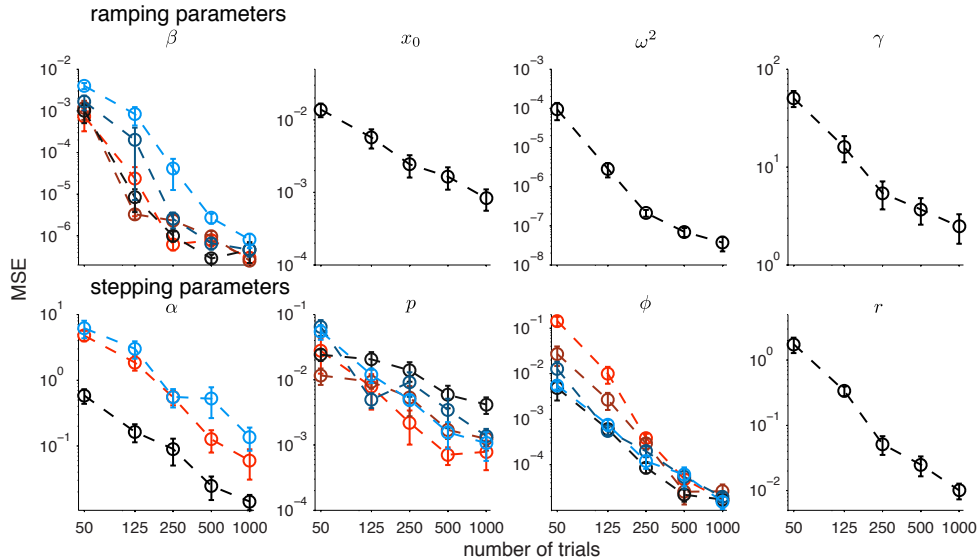


Figure 2.11: Mean-squared errors of the MCMC posterior mean estimate (log scaled) of the model parameters as a function of the amount of simulated data. Simulated datasets containing an equal number of trials of each of 5 possible motion coherence levels. For the 50 and 125 samples populations, 20 independent simulations and MCMC runs were used. 10 runs were used for the remaining sample sizes. Error bars show one standard error of the estimate of the MSE. The true parameters were chosen to be similar to the parameters estimated for a real LIP neuron: $\beta = \{-4.7e \times 10^{-3}, -2.4 \times 10^{-3}, -1.3 \times 10^{-3}, 6 \times 10^{-4}, 3.4 \times 10^{-3}\}$, $x_0 = 0.72$, $\omega^2 = 1.7 \times 10^{-3}$, $\gamma = 39.7$, $\alpha = \{4.1, 0.57, 41.0\}$, $\phi = \{0.10, 0.30, 0.71, 0.82, 0.98\}$, $p = \{0.990, 0.98, 0.975, 0.97\}$, $r = 1.05$

2.4.4.3 Model comparison identifies true model in simulations

Simulated model comparison results confirmed that the model comparison can identify the true model used to produce simulated spike trains, although several cells show a small (insignificant) DIC difference in these parameter regimes (Fig. 2.12). These simulations used parameters found in the model fits to actual data (see Section 2.5) to ensure these results held in parameter regimes relevant to the neurons in the real LIP dataset.

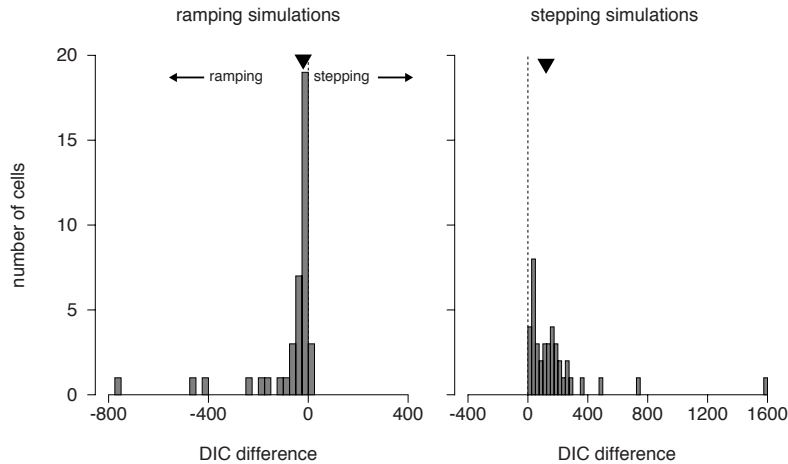


Figure 2.12: Model comparison with DIC (Eq. 2.35) on simulated data generated from the model fits to 40 LIP cells (Section 2.5). Each simulation contained the same number of trials as in the actual data (on an individual cell basis). For all 40 simulations of the stepping model (right), the model comparison showed strong evidence (DIC difference greater than 10) towards the correct model. For the ramping simulations (left), 31 simulations showed strong evidence supporting the ramping model. The model comparison showed that 3 ramping simulations could be explained better by stepping, but none of these simulations offered strong support for stepping (maximum DIC difference of 3, well below the threshold of 10 for strong support). Median DIC differences are given by the triangles: ramping simulations= -20.3 , and stepping simulations= 121.2 .

2.5 Results of model comparison of LIP responses

We applied Bayesian model comparison between the ramping and stepping model to a population of cells recorded in area LIP while the monkey performed the perceptual decision-making task shown in Figure 2.1A. The data shown here have been reported previously (Meister et al. 2013). The task and data collection are briefly summarized here. Single units were recorded while a monkey performed a moving-dot direction-discrimination task. The dot motion was displayed for random times, uniformly distributed from 500 to 1000 ms. The motion was directed towards one of the two choice targets: a target placed inside the cell’s response field (in-RF target), and one placed diametrically opposite (out-RF target). 500 ms after the motion stimulus ended, a go-signal was provided and the monkey made a saccade to one of two choice targets, indicating its choice of motion direction. Dot coherence (the strength of the motion stimulus) was varied across trials. Dot coherence was drawn from conventional values of 0, 3.2, 6.4, 12.8, 25.6, and 51.2%. Coherence levels for each trial were selected uniformly random and motion direction was sampled independently with a 50% chance of an in-RF direction. To simplify the analysis, we collapsed the coherences into 5 levels: zero=0%, positive/negative high = {25.6, 51.2}%, and positive/negative low = {3.2, 6.4, 12.8}% (positive values indicate motion towards the in-RF target and negative values towards the out-RF target). The original study included trials in which the choice targets were displayed during the entire trial and trials in which the targets were flashed briefly before motion onset. This study considered only trials for which the

targets were displayed throughout the trial.

2.5.1 Cell selection

In the original study (Meister et al. 2013), 80 spatially selective LIP cells were recorded from 2 adult, male rhesus macaques (*M. mulatta*; 14 from monkey J, 66 from monkey P). The full population included cells with a range of response styles, including cells selective primarily for motor response. The intent of this study was to include only cells with choice selectivity during the motion epoch. A d' analysis was used to quantify choice selectivity in the spike counts of the cells during the period 200-700 ms after dot motion onset (before the go signal was given on any trial). The d' value measures choice selectivity for a single cell as

$$d' = \frac{\mu_{in} - \mu_{out}}{\sqrt{\frac{1}{2} (\sigma_{in}^2 + \sigma_{out}^2)}} \quad (2.98)$$

where μ_{in} and μ_{out} are the mean spike counts on the in-RF and out-RF trials respectively. The variance of the spike counts on the in-RF and out-RF trials are σ_{in}^2 and σ_{out}^2 . We selected the top 50% of cells (40 cells, 6 from monkey J and 34 from monkey P). The d' of the 40 selected cells ranged from 0.397 to 1.661 with mean 0.819 and standard deviation 0.359.

For the statistical analyses, we used the spike trains for each trial starting 200ms after motion onset (assuming a $\tilde{200}$ ms latency of the decision-related activity to appear in LIP (Churchland et al. 2008)) until 200 ms after

motion offset. Therefore, the spike trains were 500-1000 ms long, uniformly distributed on this interval. The analyses produced similar results with earlier start times results (Sec. 2.5.5.3). The mean number of trials per cell was 385 (std=148.4, range=[96 750]). The mean number of spikes observed on each trial was 12.4 (std=11.7, range=[0 116]). This set of cells was qualitatively similar to those that have received focus in other studies, and exhibit response profiles (shown in Figures A.2- A.4) similar to those present in the well-studied Roitman & Shadlen dataset (Fig. 2.23).

2.5.2 Model fits to 40 LIP neurons

The posterior mean provided a specific estimate for the model parameters for a cell. Figures 2.13-2.14 plot the population summary of the parameters estimates for both models, and the exact values are provided in Tables A.2 and A.1 along with the model comparison results for each cell.

2.5.3 Model comparison supports stepping model in majority of cells

Using latent variable models, we formally addressed the ramping versus stepping hypothesis in LIP using statistical model comparison. Both models give a probability distribution over spike trains, and the model that better represents the data should place more probability mass over the observed spike trains. We compared the model fits using DIC (Section 2.3). The stepping model provided a superior account of LIP responses for 78% (31/40) of the cells compared to the ramping model (Fig. 2.15A). The superiority was supported

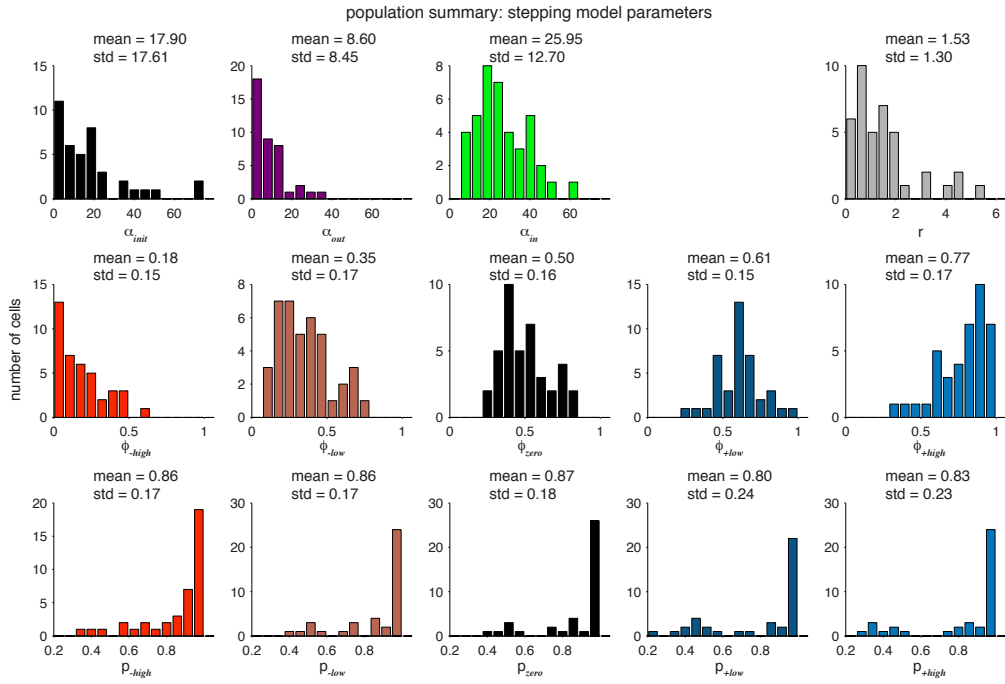


Figure 2.13: The stepping model parameters fit to all 40 cells.

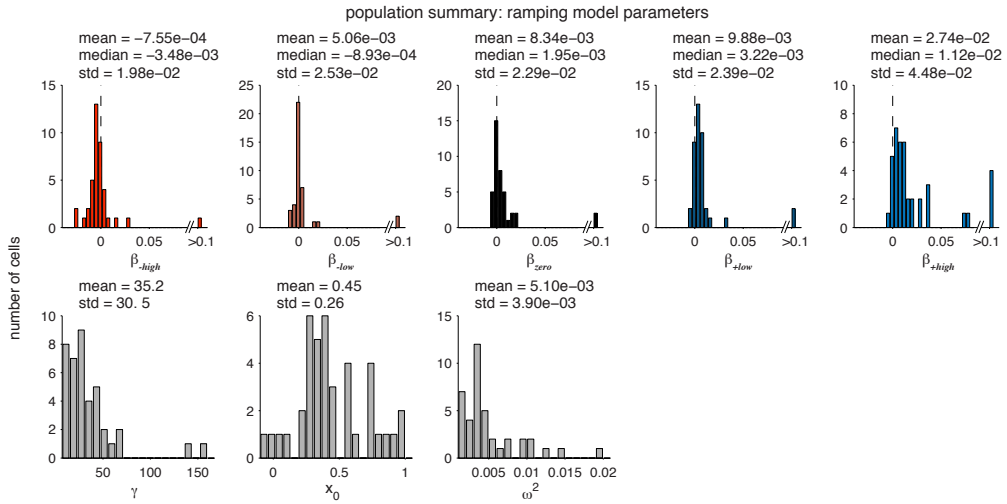


Figure 2.14: The ramping model parameters fit to all 40 cells.

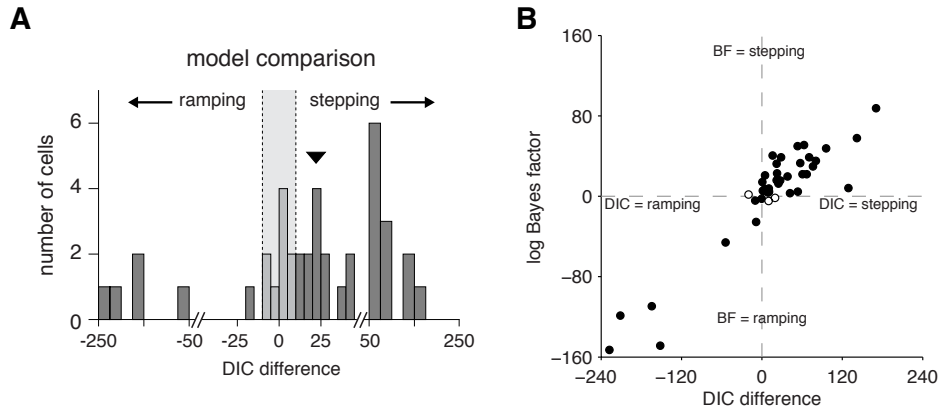


Figure 2.15: **(A)** Quantitative model comparison using DIC (Eq. 2.35) reveals a superior fit of the stepping model over the ramping model for the majority of cells (31 out of 40). A DIC difference greater than ± 10 (gray region) is commonly regarded as providing “strong” support for one model over the other (Burnham & Anderson 2002). We found substantially more cells with strong evidence for stepping over ramping (25 cells vs. 6 cells; median DIC difference = 22.1, sign test $p < 0.001$). **(B)** Comparison of the estimated DIC differences with the log Bayes factor model comparison metric for all 40 cells. These values are highly correlated ($r = 92.7, p < 10^{-8}$) and they provide the same model comparison result for all but 3 cells (open circles), all of which had weak (< 10) model assignments from the log Bayes factor.

not just by DIC but also by other model comparison metrics, such as Bayes factor (Fig. 2.15B).

2.5.3.1 Population spiking statistics better fit by stepping mode

We examined how well the two models account for the time-varying mean (or peri-stimulus time histogram, PSTH) and the variance (peri-stimulus time variance, PSTV) of neural responses. Figure 2.16 shows the comparison for the mean responses (top row) and variance (bottom row) for the data (left

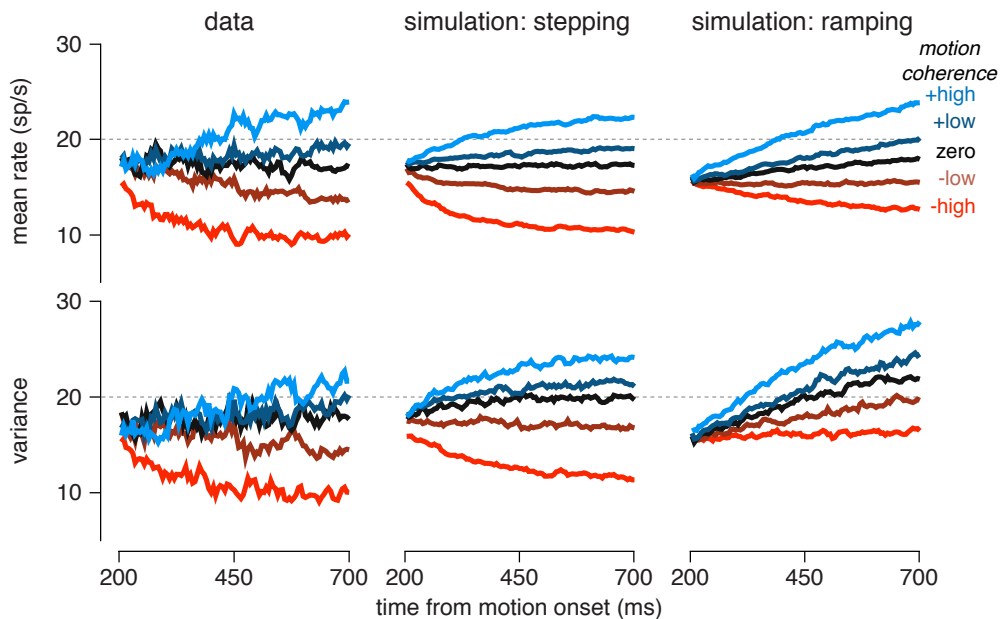


Figure 2.16: Comparison of model fits to average population activity, sorted by stimulus strength. Motion coherence and direction are indicated by color (blue, in-RF; red, out-RF). Average spike rate (top) and spike count variance (bottom) for the population aligned to motion onset. The data (left) and simulations from the stepping model (center) and the diffusion-to-bound model (right) fits to all 40 cells are shown. Spike rates and variances were calculated with a 25 ms sliding window (see Section A.1).

column), stepping model (middle column), and ramping model (right column). Although the models were fit to predict the spike responses on each trial, as opposed to these summary statistics, both models did an acceptable job of accounting for the mean response (fraction of variance in the PSTHs explained: stepping $R^2 = 0.94$, 95% credible interval:[0.90,0.94], ramping $R^2 = 0.78$ 95% CI: [0.71,0.79]). This is consistent with the long-standing difficulty in distinguishing between these two mechanisms. However, the stepping model provided a more accurate fit to the variance of neural responses (stepping

$R^2 = 0.40$, 95% CI:[0.09, 0.45], ramping $R^2 = -0.49$, 95% CI: [-0.86, -0.27]). In particular, the stepping model captured the decreasing variance observed in trials with strong negative motion much better than the ramping model. A similar result held for estimates of variance of the underlying spike rate (Fig. 2.27).

2.5.4 Inferring steps on single trials

Figure 2.17A shows the raster of spike trains from an example LIP neuron plotted in two different ways: first, aligned to the time of motion stimulus onset (left); and second, aligned to the step time inferred under the stepping model (right; see below in Section 2.5.4.2 for details about the step time estimate). The traditional raster and PSTH at left show that the average response ramps upward or downward depending on choice, as expected. The step-aligned raster at right, however, shows that these data are also consistent with discrete step-like transitions with variable timing across trials. Additional panels show the distribution of step times inferred under the model (Fig. 2.17B), and the dependence of step direction (“up” or “down”) on the motion signal (Fig. 2.17C). Discrete steps in the instantaneous spike rate could thus plausibly underlie the gradual ramping activity seen in stimulus-aligned and averaged LIP spike responses.

We applied the same analysis to the full set of LIP neurons and observed similar structure in step-aligned rasters (see Figs. A.2-A.4 for more examples). Figure 2.18A shows population-averaged PSTHs computed from stimulus-

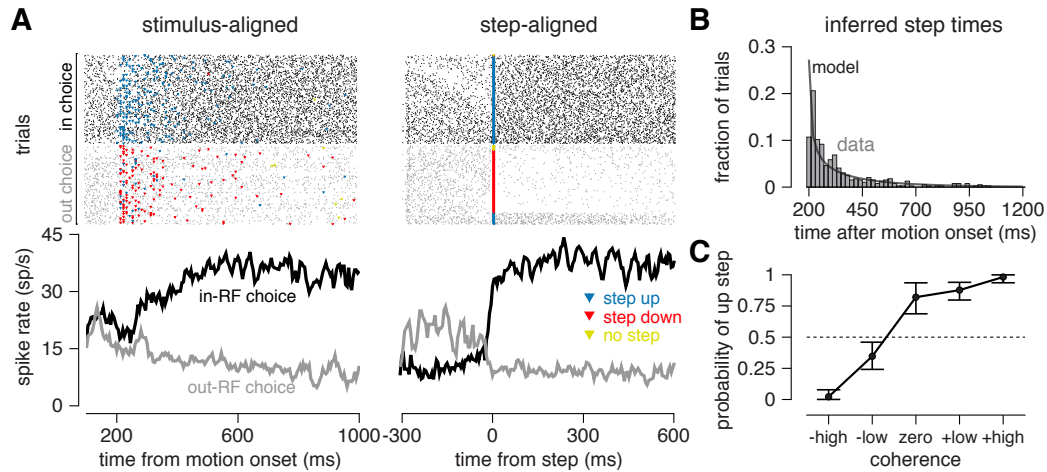


Figure 2.17: Model-based analysis of spike responses from an example LIP neuron. (A) Spike rasters sorted by the monkey’s choice in or out of the RF of the neuron under study (black=“in-RF”, gray=“out-RF”), and their associated averages (PSTHs, below). Left: Conventional stimulus-aligned rasters with each trial aligned to the time of motion onset exhibit commonly-observed ramping in the PSTH. Blue and red triangles indicate the inferred time of an “up” or “down” step on each trial under the fitted stepping model. Yellow triangles indicate that no step was found during the trial, and are placed at the end of the trial segment analyzed here (200 ms after motion offset). Right: The same spike trains aligned to the inferred step time for each trial. Note that estimated step direction of the neuron does not always match the animal’s decision on each trial. (B) The distribution of inferred step times shown in A (histogram), and the distribution over step times under the fitted parameters (black trace). (C) The probability of an “up” step, for each coherence level. Error bars indicate 95% credible intervals.

aligned and step-aligned responses, sorted by motion strength (Fig. 2.18B), or motion strength and step direction (Fig. 2.18C). Figures 2.18B and C plots show that spike rate is effectively constant when spike trains are aligned to the inferred step time on each trial. The gradient of step heights in the middle plot results from differential probabilities of stepping “up” or “down” as a function

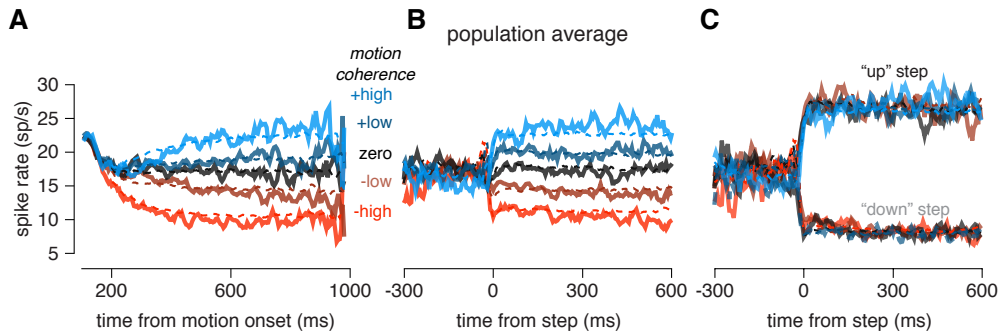


Figure 2.18: Population average PSTH sorted by motion coherence computed from spike trains: (A) aligned to motion onset and sorted by motion strength; (B) aligned to step times inferred under the stepping model and sorted by motion strength; (C) aligned to step times and sorted by both motion strength and inferred step direction. Simulated results from the stepping model (dashed lines) provide a close match to the real data under all types of alignment and conditioning.

of motion strength over trials. Figure 2.18C confirms that the firing rate, once conditioned on stepping “up” or “down”, is independent of motion strength. Furthermore, simulated spike responses, based on the fitted stepping models, resemble the real data under both kinds of alignment (dashed traces).

Although these analyses provided a visually compelling illustration of the plausibility of stepping dynamics in LIP, they could not by themselves definitively rule out the ramping model. We performed the same step-aligned analysis on a simulated diffusion-to-bound cell (Fig. 2.19). The step-aligned firing rate appeared to reveal discrete jump, even though the firing rate was truly a continuous diffusion process. Therefore, seeking to visualize a signature of a dynamical process in the mean spike rate of a cell aligned to an event - a “neural correlate” - can provide some evidence that a potential model is a

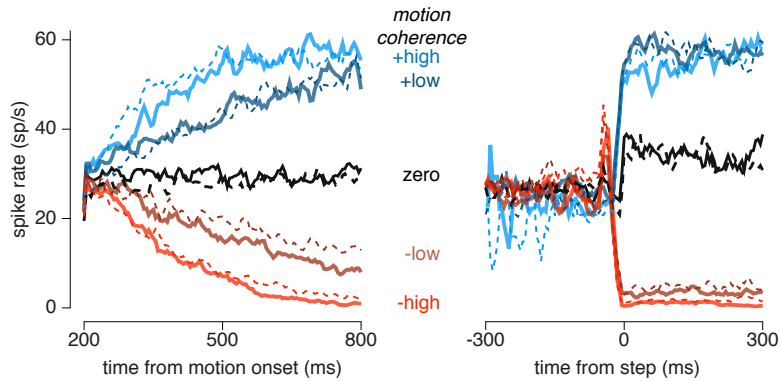


Figure 2.19: Coherence-sorted PSTH for a simulated ramping cell (solid traces) aligned to motion onset (left) and estimated step times (right). The output of the stepping model fit to the simulated ramping cell is plotted in the dashed traces. The step-aligned PSTH revealed an apparent step in the ramping simulation. Although the slope in the step-aligned PSTHs from the ramping simulation before and after the step appears to be nonzero, these ramping slopes were small enough that visual inspection should be deemed insufficient. However, the quantitative model comparison using DIC correctly identified the simulated responses as arising from the ramping model.

viable description of the data, but conclusive tests between ramping and stepping models require additional quantitative assays. Instead, Bayesian model comparison provides a stronger measure for hypothesis testing by quantifying which model best predicts the observed spike trains.

2.5.4.1 Decoding choice with steps

The stepping model provided an intuitive platform for decoding the monkey’s choice on single trials, as the posterior distribution over steps could be used for reading out decisions from the spikes on a single trial. We first quantified decoding performance using choice probability (CP), a popular met-

ric for quantifying the relationship between choice and a pair of spike counts. Aligned to motion onset, CP grew roughly linearly with time (Fig. 2.20A, left). However, the CP relative to the inferred step times (Fig. 2.20A, right) was consistent with an abrupt emergence of choice-related activity. We then compared classical CP with a model-based CP measure, which assumed that the direction of the neuron’s step predicted the animal’s choice. The model was fit to the spike trains without access to the animal’s choices. The model-based CP was on average greater than classical CP, indicating that the states estimated under the stepping model were more informative about the animal’s choice than raw spike counts (Fig. 2.20B).

Choice probability (CP) is a widely-used metric to quantify the relationship between spike count fluctuations and behavior. For a fixed stimulus, higher spike counts on single trials are often correlated with in-RF choices. CP is the probability that a spike count observed during an in-RF choice trial is greater than a spike count observed on an out-RF choice trial in response to the same stimulus condition. CP does not consider the percentage of in-RF or out-RF choices, only the distributions of spike counts observed for in-RF choice and out-RF choice trials. This way, one obtains a measure how the trial-to-trial fluctuations of a neuron correlate with choice, which discards any overall bias towards one choice.

CP is equivalent to decoding choice using a neuron-antineuron pair. The “anti-neuron” is a hypothetical neuron whose response distributions are equal to the recorded neuron, but with tuning to the targets reversed. A

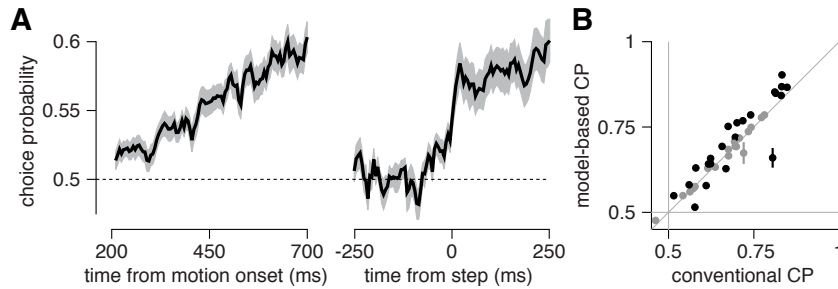


Figure 2.20: **(A)** Population average choice probability aligned to stimulus onset (left), and average CP aligned to estimated step times (right). Grey region indicates one standard error of the mean. CPs were calculated with a sliding 25 ms window. Conventional alignment suggests a ramp in choice selectivity, while the model-based alignment indicates a rapid transition. **(B)** Conventional choice probability based on spike counts using responses 200-700 ms after motion onset versus model-based choice probability using the probability of stepping to the up state by the end of the same period. Model-based CP is greater than conventional CP in the population (Wilcoxon signed rank test, $p < 0.05$). Stepping models were fit using 10 fold cross validation. Error bars show the standard error of CPs, as computed on each training data set. Black points indicate cells with significant differences between model-based and conventional CP (Student’s t-test, $p < 0.05$), and grey indicates not significant.

decoder would chose a target by selecting which of these two neurons gives the highest spike count on a single trial. An out-RF choice occurs when the antineuron gives the larger count, and the in-RF choice occurs when the antineuron has the lower count.

We calculated CP conditioned on stimulus coherence and direction and took the final CP as the average across conditions. CPs were calculated using dot coherences from 0 – 12.8%, including both directions of motion. The final CP value included only coherence levels with at least 3 in-RF choices

and 3 out-RF choices into the final CPs. We chose to average CPs across conditions instead of z -scoring each response within the stimulus categories and calculating a single CP on the pooled responses, because z -scoring can be biased by conditions with unbalanced choice selection (Kang & Maunsell 2012).

In addition to the classical, spike count-based CP, we calculated a model-based CP. The model fit was used to calculate the probability that the latent state had stepped up given a spike train. The model-based CP is the probability that an in-RF choice trial is more likely to have stepped into the up state than an out-RF choice trial (i.e., model-based CP was calculated by replacing the spike count observations in classical CP with the step probabilities).

The specific quantity used to compute the model-based CP trial j at time t was computed

$$P_{up}(j) = \frac{P(d_j = 1 | y_{j,1:t}, z_j < t, \bar{\Psi})}{P(d_j = 1 | y_{j,1:t}, z_j < t, \bar{\Psi}) + P(d_j = 2 | y_{j,1:t}, z_j < t, \bar{\Psi})} \quad (2.99)$$

where

$$P(d_j = 1 | y_{j,1:t}, z_j < t, \bar{\Psi}) = \sum_{z=0}^{t-1} P(z_j = z, d_j = 1 | y_{j,1:t}, z_j < t, \bar{\Psi}) \quad (2.100)$$

$$\propto \sum_{z=0}^{t-1} P(y_{j,1:t} | z_j = z, d_j = 1, \bar{\Psi}) P(z_j = z | \bar{\Psi}) P(d_j = 1 | \bar{\Psi}) \quad (2.101)$$

Here, the decoder assumes that a step has been made (i.e., it returns a “best guess” about the choice).

2.5.4.2 Estimating step times

Figures 2.17, 2.17, and A.2-A.4 visualize the LIP recordings using decoded steps. Step times and directions were estimated using a Bayesian decoder with the model fit parameters ($\bar{\Psi}$ = posterior mean of Ψ). For a trial j , the distribution over step times, marginalizing over step direction, is given

$$P(z_j = z | y_{j,1:T_j}, \bar{\Psi}) = \sum_{d \in \{1,2\}} P(z_j = z, d_j = d | y_{j,1:T_j}, \bar{\Psi}) P(d_j = d | \bar{\Psi}) \quad (2.102)$$

$$\propto P(y_{j,1:T_j} | z_j = z, d_j = d, \bar{\Psi}) P(z_j = z | \bar{\Psi}) P(d_j = d | \bar{\Psi}) \quad (2.103)$$

where the final terms $P(y_{j,1:T_j} | z_j = z, d_j = d, \bar{\Psi})$, $P(z_j = z | \bar{\Psi})$, and $P(d_j = d | \bar{\Psi})$ are all given by the model definition. For computational tractability, the possible step times was truncated to a maximum of 1500 time steps (15 times longer than the longest trial) and the distribution was normalized the distribution based on the truncation. We then estimated the median step time as

$$\hat{z}_j = \arg \min_{z \in \{0, \dots, 1500\}} \frac{1}{2} \geq \left(\sum_{x=1}^z P(z_j = x | y_{j,1:T_j}, \bar{\Psi}) \right) \quad (2.104)$$

The median time was selected instead of the mean because the distribution of step times tended to be highly skewed. A MAP estimator achieved similar results to the median.

If the step time occurred after the trial end, the decoder signaled that

no step occurred on that trial. Otherwise, step direction was decoded as

$$\hat{d}_j = \arg \max_{d \in \{1,2\}} P(d_j = d | y_{j,1:T_j}, \bar{\Psi}) \quad (2.105)$$

$$= \arg \max_{d \in \{1,2\}} \sum_{z=0}^{1500} P(d_j = d | z_j = z, y_{j,1:T_j}, \bar{\Psi}) P(z_j = z | \bar{\Psi}) \quad (2.106)$$

$$= \arg \max_{d \in \{1,2\}} \sum_{z=0}^{1500} P(y_{j,1:T_j} | d_j = d, z_j = z, \bar{\Psi}) P(d_j = d | \bar{\Psi}) P(z_j = z | \bar{\Psi}) \quad (2.107)$$

where the distributions $P(y_{j,1:T_j} | d_j = d, z_j = z, \bar{\Psi})$, $P(d_j = d | \bar{\Psi})$, and $P(z_j = z | \bar{\Psi})$ are all given by the model definition (product of independent Poissons, a Bernoulli distribution, and a negative binomial distribution respectively).

For all step aligned figures, spike rates on each trial were estimated by smoothing the spike trains with a centered boxcar filter (25 ms wide). The rates were then aligned to the step time on each trial and averaged. The step-aligned average included only the segment beginning 200 ms after motion onset until 200 ms after motion offset (300 ms from each trial. This avoided confounding step-related activity with motor effects or visual responses coupled to, for example, saccade target onset. For trials not containing a detectable step, we aligned the trial to the end of the analysis period (200 ms after motion offset).

Figure 2.21A shows the step times of 3 example cells. Figure 2.21C shows the average and variance of step times for each coherence, which was calculated by first taking the average (or standard deviation) of step times for each cell individually, then averaging across the population.

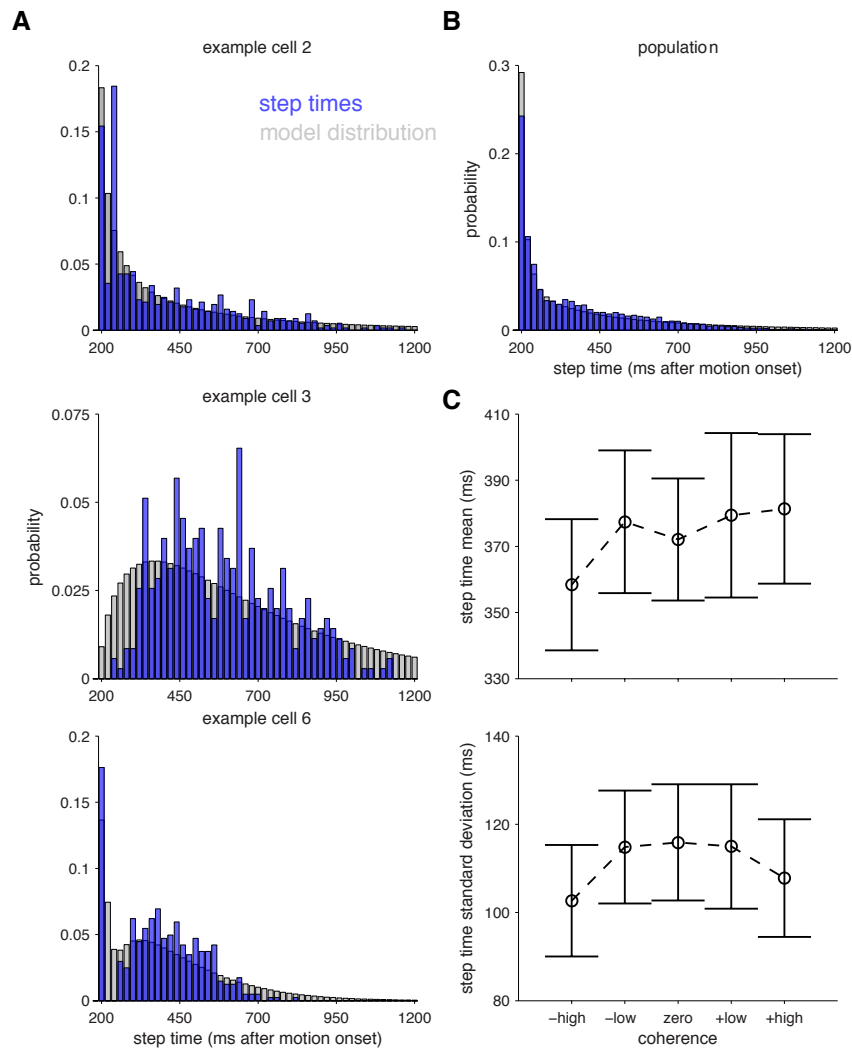


Figure 2.21: Model-estimated step times for the LIP cells. **(A)** The step times relative to motion onset estimated for 3 example cells (blue) across all motion coherences. The model fit step time distributions are shown in grey. **(B)** The step time distribution across the entire population. **(C)** The mean step times averaged across cells for each coherence level (top) and the standard deviation of step times averaged across cells (bottom). Only trials containing a detectable step step time were used in this plot.

2.5.5 Robustness of results

2.5.5.1 Response-time task: Roitman & Shadlen, 2002

Many perceptual decision-making studies have used a response-time task instead of the fixed-duration task in which the monkey must wait to receive a go signal before making a saccade used in this chapter. Although LIP is thought to reflect the same decision process during both tasks (Kiani et al. 2008), we applied the model comparison methods to a publicly available dataset from Roitman & Shadlen (2002) (downloaded from <https://www.shadlenlab.columbia.edu/>) to determine if the model comparison produced different results depending on the task. The task in this dataset was similar to the task presented earlier in this chapter, except the monkey was not given a “go” signal and instead viewed the dots until it chose to signal a decision with a saccade. The responses of these choice-selective LIP cells are summarized in Figure 7 of Roitman & Shadlen (2002).

Similar to the analysis of the data from Meister et al. (2013), we used spike trains starting from 200 ms after motion onset. Following Bollimunta et al. (2012), the model comparison included spikes occurring up until 50 ms before the saccade. Trials lasting less than 350 ms long, counting from motion onset until saccade, were discarded. Therefore, every trial we analyzed had at least 100 ms of data for the analysis. After weeding out short trials, we selected cells with at least 8 remaining trials per signed coherence level. 16 cells from this dataset met this criterion, and their motion-aligned and choice-separated PSTHs are shown in Fig. 2.23. The model comparison analysis indicated that

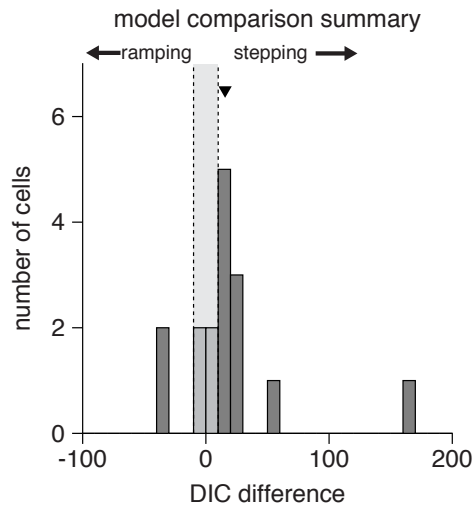


Figure 2.22: Model comparison results for 16 cells from Roitman & Shadlen (2002).

12 of these 16 cells were better fit by the stepping model than the ramping model (Fig. 2.22). Figure 2.24 shows the population average firing rate aligned to motion onset and inferred step times (similar to Figure 2.18). The similarity of this result to the model comparison using the larger dataset in the main paper suggests that the dynamics of LIP neurons are not fundamentally different in experimenter-controlled and reaction-time versions of the task, a conclusion that has been arrived at previously (Kiani et al. 2008).

The response-time paradigm, in contrast to the fixed-duration task, might allow for decision-related and saccade-related motor activity to overlap within a single spike train, even though we excluded a brief portion of the pre-saccadic activity from analysis. The models considered here were only intended to capture decision-related dynamics, so further analyses of LIP responses

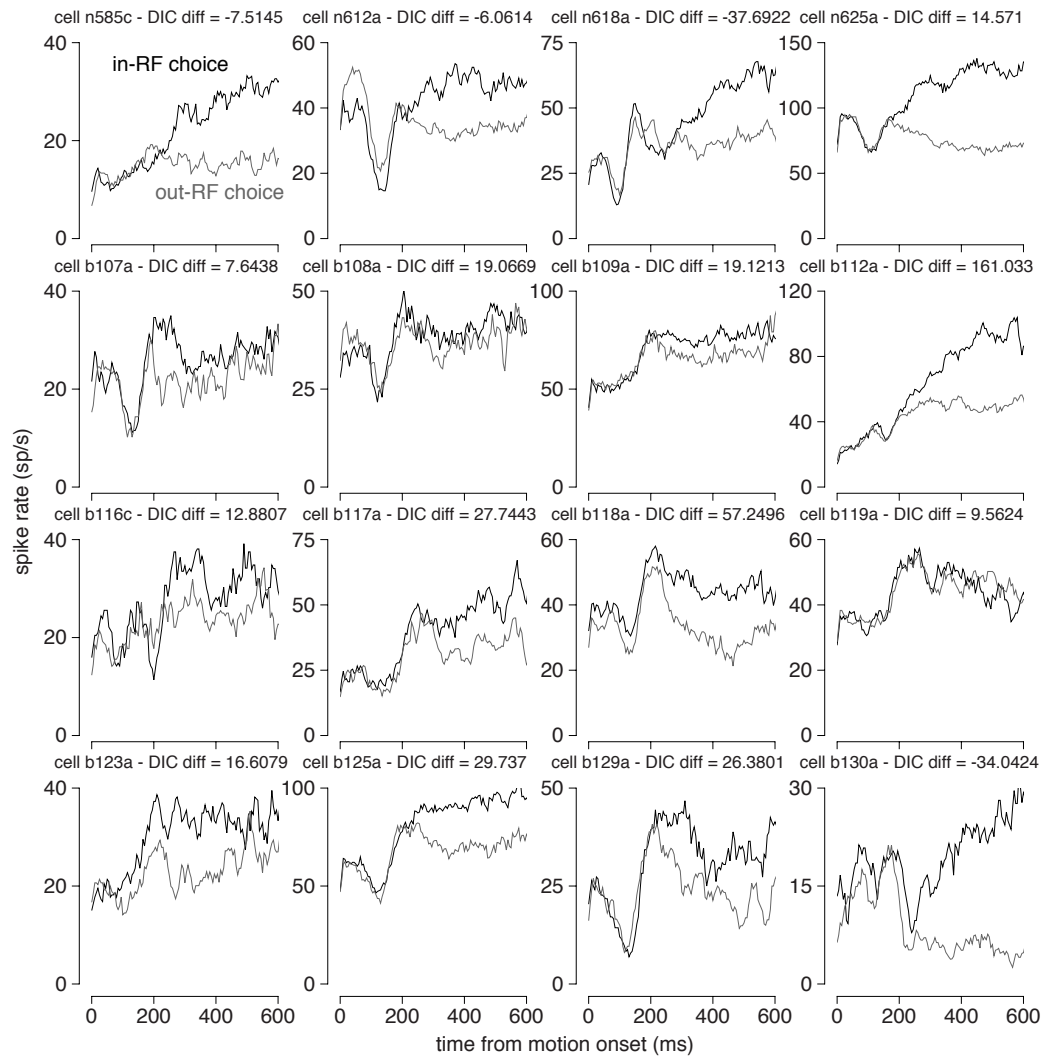


Figure 2.23: PSTHs aligned to motion onset from the cells analyzed in Figure 2.22 from Roitman & Shadlen (2002), sorted by the monkey's choice. DIC differences from the model comparison analyses are given for each cell (positive favoring the stepping model).

during the response time task will benefit from analyzing larger datasets with extended models that are explicitly designed to disentangle decision-related dynamics from pre-motor activity.

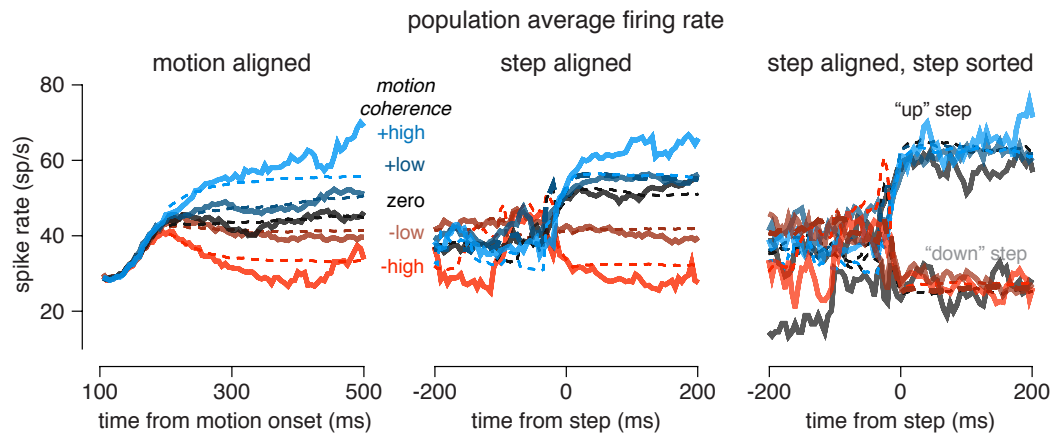


Figure 2.24: Population average PSTH sorted by motion coherence computed from spike trains for 16 cells from Roitman & Shadlen (2002): (left) aligned to motion onset and sorted by motion strength; (middle) aligned to step times inferred under the stepping model and sorted by motion strength; (right) aligned to step times and sorted by both motion strength and inferred step direction. Simulated results from the stepping model (dashed lines) provide a close match to the real data under all types of alignment and conditioning. These PSTHs were smoothed using a 50 ms sliding average, rather than a 25 ms window, because the number of trials in each condition was limited.

2.5.5.2 Model comparison results are unaffected by grouping coherence levels

We grouped stimulus coherence levels into 5 levels (\pm high, \pm low, and zero) in order to limit the number of model parameters. To ensure that this grouping did not unfairly disadvantage the ramping model, we fit the ramping model to the data using all coherence levels (11 levels) and compared to the stepping model fits with grouped coherences. In general, model comparison showed that the ramping model with grouped coherence levels performed better than including all coherences (Fig. 2.25). Only 6 cells showed slightly better performance on the all coherence levels. Of those cells, 2 were originally classified as “rampers.” For 3 of the remaining 4 “stepping” cells, the grouped-coherence stepping model still provided a better fit than the all-coherence ramping model. For the final cell, the all-coherence ramping model provided a slightly better fit than the grouped-coherence stepping model (DIC difference -1.27). However, the stepping model fit with all coherences to this cell better described the cell than the all-coherence ramping model (DIC difference 8.84).

2.5.5.3 Model comparison results are consistent across start time of analysis

The main analysis assumed that the motion integration signal in LIP began at approximately 200 ms (Churchland et al. 2011). However, the population PSTH in Figure 2.18A suggests that coherence sorting in the population may begin earlier than 200 ms. We therefore repeated the model comparison using spike trains beginning at 160 ms, 180 ms, and 220 ms after motion on-

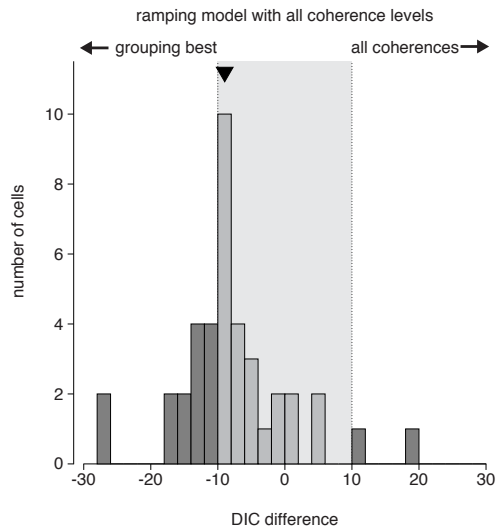


Figure 2.25: Model comparison between the ramping model fit using all coherences compared to the ramping model fit with grouped coherence levels. The median DIC difference is denoted by the black triangle.

set, but still ending 200 ms after motion offset. The model comparison results were similar across all 3 analysis windows (Fig. 2.26). The median DIC difference was comparable across time points: 17.3, 22.4, 22.1, 23.4, for the 160, 180, 200, and 220 ms start times respectively. In the 160 ms analysis, 30 cells were better fit by the stepping model (24 showed strong support) and 10 cells were better fit by the ramping model (7 showed strong support). In the 180 ms analysis, 29 cells were better fit by the stepping model (24 showed strong support) and 11 cells were better fit by the ramping model (7 showed strong support). In the 220 ms analysis, 31 cells were better fit by the stepping model (27 showed strong support) and 9 cells were better fit by the ramping model (8 showed strong support).

Only 3 out of the original 31 “stepping” cells were better fit by the

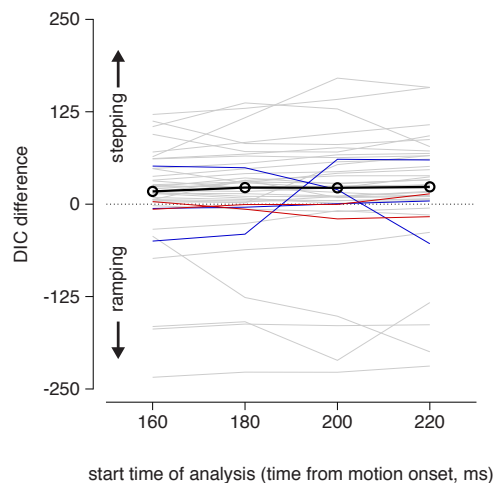


Figure 2.26: Model comparison between the ramping model fit using different start times for the analysis. The median ΔDIC (black) was significantly greater than 0 for all time points (sign test, $p < 0.004$ for each start time). 37 cells were consistently classified as a ramper or stepper across all times (grey traces). Two cells (red trace) were classified as a steppers by the 200 ms analyses, but showed slightly better support for stepping at the 160 ms, 180 ms, or 220 ms analyses. Three cells (blue traces) were classified as a stepper by the 200 ms analysis, but showed better support for ramping at one of the other start times.

ramping model at a different analysis start point (Fig. A.5). One of these cells showed weak support for stepping or ramping at all time points ($|\Delta DIC| < 10$). Two of the 9 original “ramping” cells were better fit by the stepping model at a different start point (Fig. A.6). One of these cells showed only small support for the stepping model (DIC difference = 3.5) in the 160 ms analysis. The other cell showed only weak support for ramping in the original analysis. These model comparison changes seen in these cells do not alter the overall result. These changes do not appear to be the result of an earlier onset of

integration, and are primarily the result of including responses to the onset of targets and dot motion that occur at the earlier portions of the trials (Meister et al. 2013, Park et al. 2014).

2.6 Comparison to previous methods

2.6.1 Variance of the Conditional Expectation

Churchland et al. (2011) introduced a method for analyzing the dynamics of spike trains based on the variance of the conditional expectation (or “VarCE”). This method employs the law-of-total-variance to divide the time-varying spike count variance (referred to above as the PSTV) into two components: one due to the point process or spiking variability, and another due to the variability of the underlying latent process. The Churchland et al. method assumes that the first component is proportional to the spike rate, which holds true for any inhomogeneous renewal process (including the conditionally Poisson stepping and ramping models considered here). The second component is the VarCE, which is the quantity of interest in this analysis, since it corresponds to the variability of the latent process that drives spiking. The VarCE is calculated within a single time bin as

$$VarCE = \sigma_N^2 - \phi\sigma_{N|x}^2 \quad (2.108)$$

where σ_N^2 is the total spike count variance (PSTV), and $\phi\sigma_{N|x}^2$ is an estimate of the point-process variance, obtained by multiplying the mean spike count $\sigma_{N|x}^2$ (i.e., the PSTH for that time bin) by a scale factor ϕ . This scale factor is estimated by setting it to the cell’s minimum observed Fano factor.

The basic intuition for this approach is that the shape of the VarCE over time should provide insight into the latent dynamics that underlie spiking. The VarCE of a continuous diffusion process should grow linearly, because noise accumulates linearly over time. A discrete stepping process, on the other hand, should have low VarCE at the beginning and end of a trial (assuming it always steps to the same final state), and high variance during the portion of the interval when steps are most likely.

Churchland et al. compared the shape of the VarCE curve estimated from neural data to that of a simulated stepping model, and concluded that LIP responses were inconsistent with stepping dynamics. However, the stepping simulations used to make this argument were restricted to “in”-RF choice, 0% coherence trials, and assumed that the mean response reflected a mixture of two response levels, one from early in the trial and one preceding the saccade. This produced a VarCE time course that was larger overall, and increased more steeply, than the VarCE from the data.

However, if a neuron’s latent state is allowed to step up, step down, or to not step on some trials (which is analogous to a drift-diffusion path that does not hit the bound, or wanders downward on some “in” trials), then a stepping model (like ours) can produce a more flexible range of VarCE timecourses. We therefore decided to explore whether VarCE could definitively distinguish between the ramping and stepping dynamics implemented in our models.

Figure 2.27A shows the VarCE calculated for 0% coherence trials for an example cell (black traces), along with the VarCE of spike trains simulated

from the two fitted models (red traces for ramping, blue traces for stepping). The ramping and stepping models produce nearly identical linear VarCE traces ($r^2 = 0.991$ and $r^2 = 0.992$ between a true linear ramp and the ramping and stepping model VarCEs respectively). Figure 2.27B shows an example cell for which the VarCE traces predicted by the two model fits show distinct nonlinear trends ($r^2 < 0.01$ and $r^2 = 0.83$ compared to a linear ramp for the ramping and stepping models respectively). However, estimates of VarCE are noisy for individual cells and the visual adequacy or superiority of either model is not particularly definitive. We therefore calculated the cell-averaged VarCE on the 0% coherence trials from the data and our model fits (Fig. 2.27C). We found that the VarCE of the stepping model fits (blue curve) matched the data more closely than ramping fits (red curve) (mean squared error ramping = 0.031, stepping = 0.026).

Churchland et. a. also presented an extension of the VarCE to estimate the correlation of the conditional expectation (CorCE), or the correlation of the latent spike rate over time. This process assumes that, given the latent spike rate, spike generation is independent across time in addition to the VarCE assumption that the variance of the spike count given a fixed rate is proportional to the mean. This assumption holds in the latent variable models presented in this chapter. The CorCE is calculated by taking the covariance matrix of

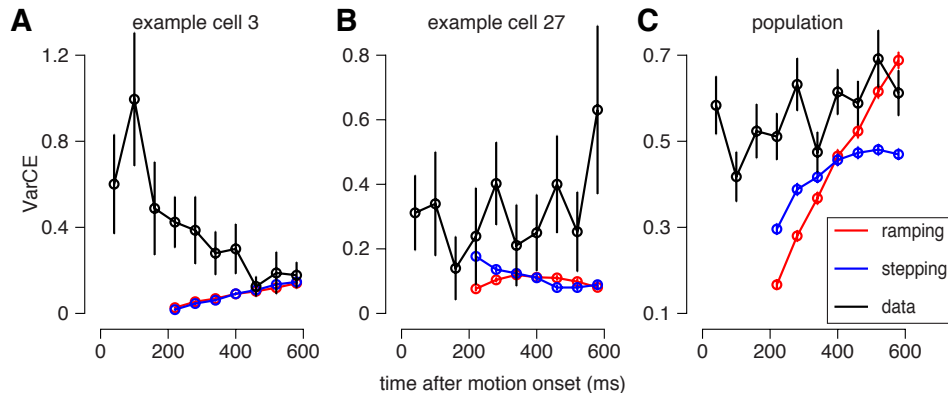


Figure 2.27: VarCE from Churchland et al. (2011) calculated from simulated spike trains and the data. The VarCE from the simulations was calculated in a 400 ms window beginning 200 ms after motion onset, and the VarCE from the data was calculated starting at motion onset. Spike count statistics were computed within a 60 ms time bin. (A) VarCE for the 0% coherence trials of an LIP cell (black). Simulations from ramping (red) and stepping fits to the cell (blue). Simulations included 50000 trials per cell— many more than can be collected in a real experiment. Cell parameters are given in Tables A.2- A.1. (B) Same as A for another cell. (C) VarCE for the LIP population (black) and from simulations of model fits to the population for 0% coherence trials.

spike counts over time, N_t , and then replacing the diagonal with the VarCE

$$\begin{pmatrix} VarCE_1 & \cdots & Cov[N_1, N_m] \\ \vdots & \ddots & \vdots \\ Cov[N_m, N_1] & \cdots & VarCE_m \end{pmatrix} = \begin{pmatrix} s_{\langle N_1 \rangle}^2 & \cdots & r_{1,m} \sqrt{s_{\langle N_1 \rangle}^2 s_{\langle N_m \rangle}^2} \\ \vdots & \ddots & \vdots \\ r_{1,m} \sqrt{s_{\langle N_1 \rangle}^2 s_{\langle N_m \rangle}^2} & \cdots & s_{\langle N_m \rangle}^2 \end{pmatrix} \quad (2.109)$$

The CorCE matrix is found by solving for the $r_{i,j}$.

Figure 2.28A shows the CorCE estimated for an example cell along with the CorCE's of the model fits to the data. The CorCE of the model fits were nearly identical, and both appeared similar to the predicted CorCE for the diffusion model given in Churchland et al. However, estimating the CorCE

matrix from a single cell is difficult, and therefore the estimate from the data is extremely noisy. It is therefore difficult to conclusively determine which model best resembled the CorCE calculated from the data. (The Frobenius norm between the model CorCE and the data CorCE was 0.82 and 0.80 for the ramping and stepping models respectively, compared to a Frobenius norm of 0.07 between the two model CorCEs.) Similarly, the CorCE calculated from the population (Fig. 2.28B) could not conclusively be differentiated between the CorCE calculated from the two models.

2.6.2 Bollimunta et al., 2012

Bollimunta et al. (2012) performed a statistical analysis of LIP spike trains similar in spirit to the one presented here. Specifically, they tested whether LIP spike trains were best fit by a ramping model or a discrete stepping model. Their analysis suggested that LIP spike trains were fit best by a ramping spike rate model, in contradiction to our results. However, their modeling and computational approaches differed strongly from our own.

There were two primary differences between our models and the models employed by Bollimunta et al. First, the Bollimunta ramping model defined a spike train as a Poisson process with a noiseless, linearly increasing rate, without any bound. This model is a severe simplification of the hypothesis of noisy accumulation to a decision bound. Second, Bollimunta et al. assumed in their stepping model that the distribution of possible step times was uniform, and that step direction was identical across the trials being analyzed. However,

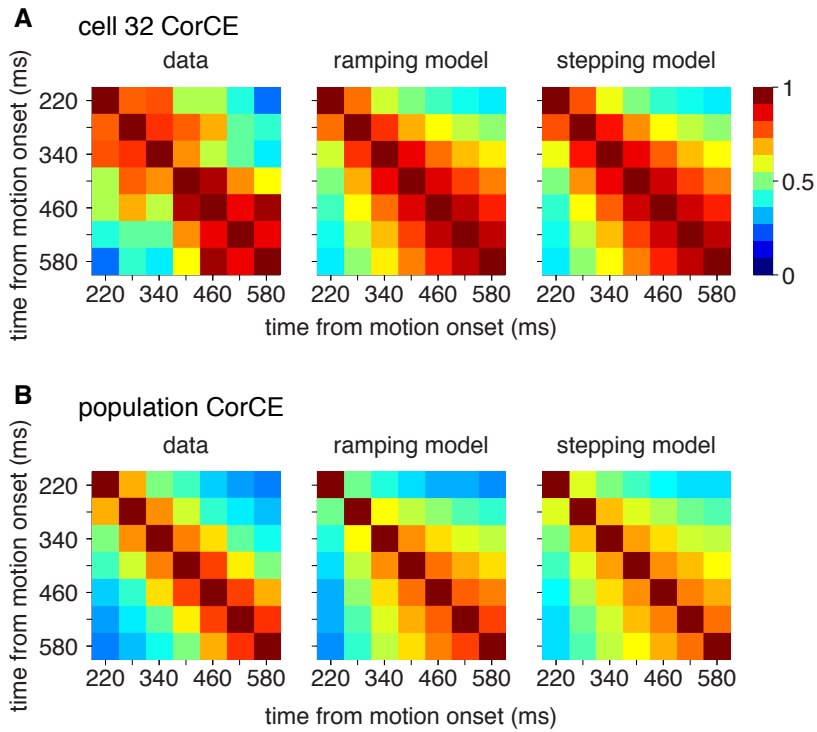


Figure 2.28: CorCE from Churchland et al. (2011) calculated from simulated spike trains and the data. The CorCE was calculated in a beginning 220 ms after motion onset, and spike count statistics were computed within a centered 60 ms time bin. **(A)** CorCE for all trials of an LIP cell (left). CorCE for simulations from ramping (middle) and stepping fits to the cell (right). Cell parameters are given in Tables A.2- A.1. **(C)** CorCE for the LIP population (left) and from simulations of model fits to the population for all trials.

step-time distributions, like reaction time distributions, are not expected to be uniform. Additionally, we found that the step-direction was not perfectly correlated with choice (Fig. 2A).

There are also issues of power. The computational costs of the model fitting techniques proposed by Bollimunta et al. made simultaneous analysis of more than 4 trials at a time infeasible. By only viewing a small portion of trials at a time, the model fits can produce drastically different parameter fits across trials. Bollimunta et al. used the Hannan-Quinn information criterion (HQIC) as a metric to compare models, similar to the use of DIC in this chapter. The distribution of HQIC values computed across model fits of different trials was tested for being significantly greater than 0, instead of computing a single HQIC for all the trial for a single cell. The magnitude of the median HQIC values reported by Bollimunta et al. are less than 0.01, which several orders of magnitude smaller than the model comparison values we reported. Additionally, Bollimunta et al. constrained their analysis to view only the 400ms preceding a saccade instead of the entire integration period. The MCMC methods we used to analyze a large number of trials, along with our definitions of both types of latent dynamics, increased the statistical power of our study.

These differences are substantial, and might explain the discrepancies in the findings reported here.

2.7 Conclusions

We have developed tractable, principled methods for fitting and comparing statistical models of single-neuron spike trains in which spike rates are governed by a latent stochastic process. We have applied these methods to determine the dynamics underlying single-neuron activity in area LIP. Although neurons in this area have been largely assumed to exhibit ramping dynamics, reflecting the temporal accumulation of evidence posited by models of decision-making, statistical model comparison supports an alternative hypothesis: LIP responses were better described by randomly timed, discrete steps between underlying states. In a supplementary analysis, we examined data from a response-time version of the dots task and found results consistent with the trend in the fixed duration version; this initial comparison will be strengthened by extending the models to account for overlapping decision and motor events, and application to larger datasets (Figs. 2.22-2.24) (Roitman & Shadlen 2002). In addition to accounting better for the dynamics of the mean firing rates, only the stepping model accounted accurately for the variance of neural responses. Finally, the estimation of single-trial step times provided a novel view of choice-related activity, revealing that choice-correlated fluctuations in response are also dominated by discrete step-like dynamics.

Although these results challenge the canonical perspective of LIP dynamics during decision-making, the approach facilitates new avenues of investigation. These analyses suggest that accumulation may be implemented by stochastic steps, but simultaneous recordings of multiple neurons will be

required to test whether population activity ramps or discretely transitions between states on single trials (Miller & Katz 2010); population-level ramping could still be implemented via step times that vary across neurons, even on the same trial. Fortunately, the statistical techniques reported here are scalable to simultaneously-recorded samples of multiple neurons, and newer recording techniques are starting to yield these multi-neuron datasets (Stevenson & Kording 2011, Bollimunta et al. 2012, Kiani, Cueva, Reppas & Newsome 2014, Kaufman, Churchland, Ryu & Shenoy 2015). It is also possible that single neurons with ramping dynamics implement evidence integration elsewhere in the brain, and that LIP neurons are post-decisional or pre-motor indicators of the binary result of this computation. More generally, we believe these techniques will have broad applicability for identifying and interpreting the latent factors governing multi-neuron spike responses, allowing for principled tests of the dynamics governing cognitive computations in many brain areas.

Chapter 3

Inferring synaptic conductances from spike trains in retinal ganglion cells

A popular approach to modeling sensory neurons describes neural responses in terms of a cascade of linear and nonlinear stages: a linear filter to describe stimulus integration, followed by a nonlinear function to convert the filter output to spike rate. In particular, the point process generalized linear model (GLM) has provided a useful and highly tractable tool for characterizing neural encoding in a variety of sensory, cognitive, and motor brain areas (Harris, Csicsvari, Hirase, Dragoi & Buzsaki 2003, Truccolo, Eden, Fellows, Donoghue & Brown 2005, Pillow, Shlens, Paninski, Sher, Litke & Chichilnisky 2008, Gerwinn, Macke & Bethge 2010, Stevenson, London, Oby, Sachs, Reimer, Englitz, David, Shamma, Blanche, Mizuseki, Zandvakili, Hatsopoulos, Miller & Kording 2012, Park et al. 2014). However, there is a substantial gap between descriptive statistical models like the GLM and more realistic, biophysically interpretable neural models. Cascade-type statistical models describe input to a neuron in terms of a set of linear (and sometimes nonlinear) filtering steps (Paninski 2004, Butts, Weng, Jin, Alonso & Paninski 2011, Vintch, Zaharia, Movshon & Simoncelli 2012, McFarland, Cui & Butts 2013, Park, Archer, Priebe & Pillow 2013, Theis, Chagas, Arnstein,

Schwarz & Bethge 2013). Real neurons, on the other hand, receive distinct excitatory and inhibitory synaptic inputs, which drive conductance changes that alter the nonlinear dynamics governing membrane potential. Previous work has shown that excitatory and inhibitory conductances in retina and other sensory areas can exhibit substantially different tuning (Roska et al. 2006, Trong & Rieke 2008, Poo & Isaacson 2009, Cafaro & Rieke 2013). Determining how the excitatory and inhibitory inputs combine to produce the spike train output of a neuron remains a challenge, because spikes and excitatory and inhibitory synaptic inputs cannot be measured simultaneously (Roska et al. 2006).

Here we introduce a quasi-biophysical interpretation of the generalized linear model. The resulting interpretation reveals that the GLM can be viewed in terms of a highly constrained conductance-based model. We expand on this interpretation to construct a more flexible and more plausible conductance-based spiking model (CBSM), which allows for independent excitatory and inhibitory synaptic inputs. The model can exhibit shunting as well as hyperpolarizing inhibition, and time-varying changes in both gain and membrane time constant.

We show that the stimulus-dependence of both excitatory and inhibitory conductances can be well described by a linear-nonlinear cascade, with the filter driving inhibition exhibiting opposite sign and a slight delay relative to the filter driving excitation. We find that the CBSM allows us to accurately infer excitatory and inhibitory synaptic conductances received by macaque retinal ganglion cells (RGCs) in response to visual stimuli from extracellu-

larly recorded spike trains. Additionally, the CBSM captures RGC responses more accurately than the standard GLM for both full-field and spatio-temporal stimuli.

This work has in part been published in the proceedings of the Neural Information Processing Systems (NIPS) 2014 meeting (Latimer, Chichilnisky, Rieke & Pillow 2014).

3.1 Methods

3.1.1 Electrophysiology

We analyzed four sets of parasol RGCs. All data were obtained from isolated, peripheral macaque monkey, *Macaca mulatta*, retina.

3.1.1.1 Synaptic current recordings

We analyzed the responses of 7 ON parasol cells previously described in Trong & Rieke (2008). Cell-attached records to obtain spike trains, and voltage clamp recordings were performed to measure excitatory and inhibitory currents in the same cells. The stimulus, delivered from an LED, consisted of a one dimensional, full-field white noise signal, filtered with a low pass filter with a 60Hz cutoff frequency, and sampled at a 0.1ms resolution. Either 30 or 40 trials were recorded from each cell, using 10 unique 6 second stimuli. After the spike trains were recorded, voltage clamp recordings were used to measure the excitatory and inhibitory currents to the same stimuli. For 5 of the cells, 2-4 trials were recorded for each of the 10 stimuli for the excitatory

and inhibitory currents. For the 2 remaining cells, 3-4 excitatory current trials were recorded for all 10 stimuli and 1-2 trials for the inhibitory current were obtained for 8 of the stimuli. Conductances were estimated by dividing the current by the approximate driving force (-70mV for the excitatory current, and 70mV for the inhibitory).

3.1.1.2 Dynamic clamp recordings

The membrane potentials of 2 ON parasol retinal ganglion cells were recorded during dynamic clamp experiments previously reported in Cafaro & Rieke (2013). The cells were current clamped and current was injected into the cells according to the equation

$$I(t) = g_e(t)(V(t - \Delta_t) - E_e) + g_i(t)(V(t - \Delta_t) - E_i) \quad (3.1)$$

where g_e and g_i were conductances recorded in RGCs in response to a light stimulus. The injected current at time t was computed using the previous measured voltage with offset $\Delta_t = 100\mu s$. The reversal potentials were $E_e = 0$ mV and $E_i = -90$ mV.

For the first cell, 18 repeat trials were recorded for a 19 s stimulation, and 24 repeat trials were obtained from the second cell.

3.1.1.3 RGC population recordings: full-field stimulus

We analyzed data from two experiments previously reported in Uzzell & Chichilnisky (2004) and Pillow et al. (2005). The first experiment included 9 simultaneously recorded parasol RGCs (5 ON and 4 OFF). The stimulus

consisted of a full-field binary noise stimulus with a root-mean-square contrast of 96%. The stimulus was displayed on a CRT monitor at a 120Hz refresh rate and the contrast of each frame was drawn independently. A 10 min stimulus was obtained for characterizing the cell responses, and a 5 min segment was used to obtain a cross-validated log-likelihood. Spike rates were obtained by recording 167 repeats of a 7.5 s stimulus.

In a second experiment, 8 cells (5 OFF parasol and 3 on parasol) were recorded in response to a full-field binary noise stimulus (120Hz) at 12%, 24%, and 48% contrast. An 8 min stimulus segment at each contrast level was used for model fitting, and cross-validated log-likelihoods were obtained using a novel 4 min segment at each contrast level.

3.1.1.4 RGC population recordings: spatio-temporal stimulus

We analyzed 11 ON and 16 OFF parasol RGCs which were previously reported in Pillow et al. (2005). The stimulus consisted of a spatio-temporal binary white noise pattern (i.e., a field of independent white and black pixels). The stimulus was 10 pixels by 10 pixels (pixel size of $120\mu\text{m} \times 120\mu\text{m}$ on the retina) displayed with a 120Hz refresh rate, and the contrasts of each pixel was drawn independently on each frame. The root-mean-square contrast of the stimulus was 96%.

A 10 min stimulus was obtained for characterizing the cell responses, and a 5 min segment was used to obtain a cross-validated log-likelihood. Spike rates were obtained by recording 600 repeats of a 10 s stimulus.

3.1.2 Modeling methods

3.1.2.1 The conductance-based spiking model

We introduce the CBSM which models the spike train response of a RGC to a visual stimulus as a Poisson process where the spike rate is a function of the membrane potential (Fig. 3.1). The membrane potential is approximated by considering a single-compartment neuron with linear membrane dynamics and conductance-based input:

$$\frac{dV}{dt} = -g_l(V - E_l) + g_e(t)(V - E_e) - g_i(t)(V - E_i). \quad (3.2)$$

The synaptic inputs take the form of linear-nonlinear functions of the stimulus, \mathbf{x} :

$$g_e(t) = f_e((\mathbf{k}_e * \mathbf{x})(t)), \quad g_i(t) = f_i((\mathbf{k}_i * \mathbf{x})(t)), \quad (3.3)$$

where f_e and f_i are nonlinear functions ensuring positivity of the conductances. The firing rate is a nonlinear function of the membrane voltage plus a GLM-like spike history (autoregressive) term to account for refractory periods or other spike-dependent behaviors:

$$r_t = R_{sp} \left(\frac{V(t) - V_T}{\Delta V} + (\mathbf{h}_{spk} * \mathbf{y}^{hist})(t) \right). \quad (3.4)$$

The voltage-to-spike rate nonlinearity, R_{sp} , follows the form proposed by Mensi, Naud & Gerstner (2011), where V_T is a soft spiking threshold and ΔV determines the steepness of the nonlinearity. Although spiking activity in real neurons influences both the membrane potential and the output nonlinearity (Johnston, Wu & Gray 1995, Badel, Lefort, Berger, Petersen, Gerstner

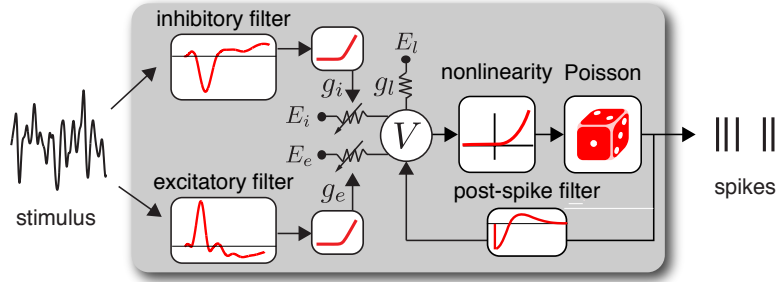


Figure 3.1: Schematic of conductance-based spiking encoding model.

& Richardson 2008), we do not include any spike-dependent currents in the CBSM’s membrane voltage dynamics.

For a set of spike times $s_{1:n_{sp}}$ in the interval $[0, T]$ and parameters Θ , the log-likelihood is

$$\log p(s_{1:n_{sp}} | \mathbf{x}_{1:T}, \Theta) = \sum_{i=1}^{n_{sp}} \log(r(s_i)) - \int_0^T r(t) dt. \quad (3.5)$$

This continuous time likelihood can be discretely approximated as the product of Bernoulli trials in bins of width Δ (Citi, Ba, Brown & Barbieri 2014)

$$\log p(y_{1:T} | \mathbf{x}_{1:T}, \Theta) = \sum_{t=1}^T y_t \log(1 - \exp(-r_t \Delta)) - (1 - y_t) r_t \Delta \quad (3.6)$$

where $y_i = 1$ if a spike occurred in the i th bin and 0 otherwise.

The membrane voltage (and firing rate) is computed by integrating the membrane dynamics equation (Eq. 3.2). In practice, we evaluate V along the same discrete lattice of points of width Δ ($t = 1, 2, 3, \dots T$) that we use to discretize the log-likelihood function. Assuming g_e and g_i remain constant within each bin, the voltage equation becomes a simple linear differential equation

with the solution

$$V(t + 1) = e^{-g_{tot}(t)\Delta} \left(V(t) - \frac{I_s(t)}{g_{tot}(t)} \right) + \frac{I_s(t)}{g_{tot}(t)} \quad (3.7)$$

$$V(1) = E_l \quad (3.8)$$

$$g_{tot}(t) = g_e(t) + g_i(t) + g_l \quad (3.9)$$

$$I_s(t) = g_e(t)E_e + g_i(t)E_i + g_lE_l. \quad (3.10)$$

The model parameters we fit were \mathbf{k}_e , \mathbf{k}_i , b_e , b_i , and \mathbf{h}_{spk} , which were selected using conjugate-gradient methods to maximize the log-likelihood. The derivatives of the log-likelihood with respect to these parameters are provided in Appendix B. The firing rate parameters were fixed at $V_T = -53mV$ and $\Delta V = 1.67mV$ (see Sections 3.1.2.2) and 3.2.2.1). The reversal potential and leak conductance parameters were kept fixed at $E_e = 0mV$, $g_l = 200$, $E_l = -60mV$, and $E_i = -80mV$. For modeling the cells in which we had access to intracellular recordings (Section 3.2.4), we set the time bin width to $\Delta = 0.1ms$ to match the sampling frequency of the synaptic current recordings. For the remaining cells, which were recorded in separate experiments, we set $\Delta = 0.083ms$ to equal 100 times the frame rate of the visual stimulus.

The stimulus filters spanned over 100 ms, or over 1000 time bins. Therefore, we restricted the excitation and inhibitory filters to a low dimensional basis to limit the total number of free parameters in the model. The basis consisted of 10 raised cosine ‘bumps’ (Pillow et al. 2005, Pillow et al. 2008) of the form

$$\mathbf{b}_j(t) = \begin{cases} \frac{1}{2} \cos(a \log[t + c] - \phi_j) & \text{for } a \log(t + c) \in [\phi_j - \pi, \phi_j + \pi] \\ 0 & \text{otherwise} \end{cases} \quad (3.11)$$

where t is in seconds. We set $c = 0.0004$ and $a = \phi_2 - \phi_1$. The ϕ_i were evenly spaced from $\phi_1 = \log(0.0 + c)$, $\phi_2 = \log(0.150 + c)$ so that the peaks of the filters spanned $0ms$ to $150ms$. The spike history filter was also represented in a low-dimensional basis. The refractory period was accounted for with 5 square basis functions of width $0.4ms$, spanning the period $0 - 2ms$ after a spike. The remaining spike history filter consisted of 7 raised cosine basis functions ($c = 0.0001$) with filter peaks spaced from $2ms$ to $90ms$.

The log-likelihood function for this model is not concave in the model parameters, which increases the importance of selecting a good initialization point compared to the GLM. We initialized the parameters by fitting a simplified model, which had only one conductance with a linear stimulus dependence, $g_{lin}(t) = \mathbf{k}_{lin}^\top \mathbf{x}_t$ (note that this allowed for unrealistic negative conductance values). We initialized this filter to 0, and then numerically maximized the log-likelihood for \mathbf{k}_{lin} . We then initialized the parameters for the complete model using $\mathbf{k}_e = c\mathbf{k}_{lin}$ and $\mathbf{k}_i = -c\mathbf{k}_{lin}$, where $0 < c \leq 1$, thereby exploiting a mapping between the GLM and the CBSM (Section 3.2.1.1).

When fitting the model to real spike trains, one conductance filter would occasionally become dominant early in the optimization process. This was likely due to the limited amount of data available for fitting, especially for the cells that were recorded intracellularly. The intracellular recordings clearly indicated that the cells received similarly scaled excitatory and inhibitory inputs. To alleviate this problem, we added a penalty term, ϕ , to the log-likelihood to

the L_2 norms of \mathbf{k}_e and \mathbf{k}_i :

$$\phi(\mathbf{k}_e, \mathbf{k}_i) = \lambda_e \|\mathbf{k}_e\|^2 + \lambda_i \|\mathbf{k}_i\|^2 \quad (3.12)$$

Thus we maximized

$$\mathcal{L}(\theta) = \log p(y_{1:T} | \mathbf{x}_{1:T}, \mathbf{k}_e, \mathbf{k}_i, b_e, b_i) - \phi(\mathbf{k}_e, \mathbf{k}_i) \quad (3.13)$$

All cells were fit using the same penalty weights: $\lambda_e = 1$ and $\lambda_i = 0.2$. We note that unlike the typical situation with cascade models that contain multiple filters, intracellular recordings can directly measure synaptic currents. Future work with this model could include more informative, data-driven priors on \mathbf{k}_e and \mathbf{k}_i .

In several analyses, we fit the CBSM without the inhibitory conductance, labeled as the CBSM_{exc}. All the fixed parameters used in the full CBSM were held at the same values in the CBSM_{exc}.

3.1.2.2 Spike rate nonlinearity

We used a spike-triggered analysis (De Boer & Kuyper 1968) on membrane voltage recordings to determine the spike rate nonlinearity, R_{sp} , as a function of voltage for the CBSM. The membrane potential and spikes were recorded in dynamic-clamp experiments over several repeats of simulated conductance traces for 2 cells. We computed the mean voltage recorded over all runs of the dynamic-clamp condition, which largely eliminated the action potential shapes from the voltage trace. Using the spike times from all the

repeats, we computed the probability of a spike given the mean voltage, \bar{V} :

$$p(Sp|\bar{V}) = \frac{p(\bar{V}|Sp)p(Sp)}{p(\bar{V})} \quad (3.14)$$

where $p(\bar{V}|Sp)$ is the spike-triggered distribution of the membrane potential.

We combined the spike times and voltage distributions for the two cells to compute a common spike rate function.

A least-squares fit approximated the nonlinearity with a soft-rectification function of the the form

$$R_{sp}(t) = a \log \left(1 + \exp \left(\frac{(V(t) - V_T)}{\Delta V} \right) \right) \quad (3.15)$$

where $a = 90$, $V_T = -53mV$ and $\Delta V = 1.67mV$.

3.1.2.3 Conductance LN cascade

The conductance nonlinearities, f_e and f_i were selected by fitting a linear-nonlinear cascade to measured conductances in response to a visual stimulus (Hunter & Korenberg 1986, Paninski, Vidne, DePasquale & Ferreira 2012, Park et al. 2013, Barreiro, Gjorgjieva, Rieke & Shea-Brown 2014). For a full-field stimulus \mathbf{x} , we modeled the mean conductance, $\bar{g}_e(t)$, as the LN cascade

$$\bar{g}_e(t) = a_e f_e((\mathbf{k}_e * \mathbf{x})(t) + b_e) + e_t \quad (3.16)$$

$$\epsilon_t \sim \mathcal{N}(0, \sigma^2) \quad (3.17)$$

where a_e and b_e are constants. We selected fixed functions for the nonlinearities f_e and f_i . Thus, we chose the \mathbf{k}_e , a_e , and b_e that minimized the squared error between the LN prediction and the measured excitatory conductance.

The soft-rectifying function was selected to model the conductance nonlinearities (see Sections 3.2.2.1-3.2.2.2):

$$f_e(s), f_i(s) = \log(1 + \exp(s)). \quad (3.18)$$

We chose to fix these nonlinearities to known functions rather than fitting with a more flexible empirical form (e.g., Ahrens, Paninski & Sahani 2008, McFarland et al. 2013). Fixing these nonlinearities to a simple, closed-form function allowed for fast and robust maximum likelihood parameter estimates while still providing an excellent description of the data.

3.1.2.4 Generalized linear models

For a baseline comparison to the CBSM, we also fit spike trains with a GLM. We used the same Bernoulli discretization of the point-process log-likelihood function for the GLM as we did with the CBSM:

$$\log p(y_{1:T} | \mathbf{x}_{1:T}, \mathbf{k}, b, \mathbf{h}) = \sum_{t=1}^T y_t \log(1 - \exp(-\lambda\Delta)) - (1 - y_t)\lambda\Delta \quad (3.19)$$

where the firing rate is

$$\lambda_t = f((\mathbf{k} * \mathbf{x})(t) + b + (\mathbf{h} * \mathbf{y}_{spk}^{hist})(t)). \quad (3.20)$$

The stimulus filter is \mathbf{k} and the spike history filter is \mathbf{y}_{spk}^{hist} . We used conjugate-gradient methods to find the maximum likelihood estimates for the parameters. We set $f(\cdot) = \exp(\cdot)$, which is the canonical inverse-link function for Poisson GLMs. We found that the soft-rectifying nonlinearity we used for the CBSM, $f(\cdot) = \log(1 + \exp(\cdot))$, did not capture RGC responses as well as the exponential function (results not shown).

3.1.2.5 Modeling responses to spatio-temporal stimuli

For spatio-temporal stimuli, the linear kernels for the CBSM and GLM (\mathbf{k} , \mathbf{k}_e , and \mathbf{k}_i) spanned both space and time. Although the stimulus we used was a 10x10 grid of pixels, the receptive field (RF) of the neurons did not span the complete grid. We therefore limited the spatial extent of the linear filters to a 5x5 grid of pixels, where the center pixel was the strongest point in the GLM stimulus filter.

The filters were represented as a matrix where the columns span the pixel space and the rows span the temporal dimension. The number of parameters was reduced by decomposing the spatio-temporal filters into a low-rank representation (Pillow et al. 2008). The filter at pixel x and time τ was

$$\mathbf{k}(x, \tau) = \sum_{j=1}^J \mathbf{k}_{s,j}(x) \mathbf{k}_{t,j}(\tau) \quad (3.21)$$

where $\mathbf{k}_{s,j}$ was a vector containing the spatial portion of the filter of length 25 (the number of pixels in the RF) and $\mathbf{k}_{t,j}$ represented the temporal portion of the filter. The temporal filters were projected into the same 10-dimensional basis as the temporal filters used to model the full-field stimuli and the spatial filters were represented in the natural pixel basis. For identifiability, we normalized the spatial filters and forced the sign of the center pixel of the spatial filters to be positive. We used rank 2 filters ($J = 2$) for the CBSM and GLM. Therefore, each filter contained 2×25 spatial and 2×10 temporal parameters for a total of 70 parameters. In the GLM, we found no significant

improvement using rank 3 filters. To fit these low-rank filters, we alternated between optimizing over the spatial and temporal components of the filters.

3.1.2.6 Evaluating model performance

We evaluated model predictions of spike rate by simulating 2500 trials from the model for a repeated stimulus. We computed the firing rate, or peri-stimulus time histogram (PSTH), by averaging the number of spikes observed in 1 ms bins and smoothing with a Gaussian filter with a standard deviation of 2 ms. The percent of variance in the PSTH explained by the model is

$$\% \text{ variance explained} = 100 \times 1 - \frac{\sum_{t=1}^T (PSTH_{data}(t) - PSTH_{model}(t))^2}{\sum_{t=1}^T (PSTH_{data}(t) - \overline{PSTH_{data}})^2} \quad (3.22)$$

where $\overline{PSTH_{data}}$ denotes the average value of the PSTH.

We evaluated single-trial spike train predictive performance by computing the log-likelihood of a spike train observed in response to a novel test stimulus. We computed the difference between the log (base-2) likelihood under the model and the log-likelihood under a homogeneous rate model (LL_h) that captured only the mean spike rate:

$$LL_h = n_{sp} * \log_2(\bar{\lambda}) + (T - n_{sp}) \log_2(1 - \bar{\lambda}) \quad (3.23)$$

$$\bar{\lambda} = \frac{n_{sp}}{T}. \quad (3.24)$$

where the test stimulus is of length T (in discrete bins) and contains n_{sp} spikes. We then divided by the number of spikes to obtain the predictive performance in units of bits-per-spike (Panzeri, Biella, Rolls, Skaggs & Treves

1996, Brenner, Strong, Koberle, Bialek & Steveninck 2000, Paninski, Shoham, Fellows, Hatsopoulos & Donoghue 2004)

$$\text{bits per spike} = \frac{LL_{model} - LL_h}{n_{sp}} \quad (3.25)$$

Intervals of bits-per-spike measures given here denote 1 standard deviation of the population mean.

3.2 Results

3.2.1 A conductance-based extension to the GLM

The GLM describes neural encoding in terms of a cascade of linear, nonlinear, and probabilistic spiking stages. A quasi-biological interpretation of GLM is known as “soft threshold” integrate-and-fire (Plesser & Gerstner 2000, Gerstner 2001, Paninski, Pillow & Lewi 2007, Mensi et al. 2011). This interpretation regards the linear filter output as a membrane potential, and the nonlinear stage as a “soft threshold” function that governs how the probability of spiking increases with membrane potential, specifically:

$$V_t = \int_0^t \mathbf{k}(t') \mathbf{x}(t - t') dt' = (\mathbf{k} * \mathbf{x})(t) \quad (3.26)$$

$$r_t = f_r(V_t + b) \quad (3.27)$$

$$y_t | r_t \sim \text{Pois}(r_t \Delta). \quad (3.28)$$

The linear filter \mathbf{k} maps the stimulus \mathbf{x} to the membrane potential V_t at time t , and a fixed nonlinear function f_r maps V_t to the conditional intensity (or spike rate) r_t . The spike count y_t is a Poisson random variable in a time bin of infinitesimal width Δ .

The stimulus vector \mathbf{x}_t can be augmented to include arbitrary covariates of the response such as the neuron’s own spike history or spikes from other neurons (Truccolo et al. 2005, Pillow et al. 2008). In such cases, the output does not form a Poisson process because spiking is history-dependent.

The nonlinearity f is fixed *a priori*. Therefore, the only parameters are the coefficients of the filter \mathbf{k} . The most common choice is exponential, $f(z) = \exp(z)$, corresponding to the canonical ‘log’ link function for Poisson GLMs. Prior work (Paninski 2004) has shown that if f grows at least linearly and at most exponentially, then the log-likelihood is jointly concave in model parameters θ . This ensures that the log-likelihood has no non-global maxima, and gradient ascent methods are guaranteed to find the maximum likelihood estimate.

3.2.1.1 Interpreting the GLM as a conductance-based model

A more biophysical interpretation of the GLM can be obtained by considering a single-compartment neuron with linear membrane dynamics and conductance-based input:

$$\begin{aligned} \frac{dV}{dt} &= -g_l V + g_e(t)(V - E_e) - g_i(t)(V - E_i) \\ &= -(g_l + g_e(t) + g_i(t))V + g_e(t)E_e + g_i(t)E_i \\ &= -g_{tot}(t)V + I_s(t), \end{aligned} \tag{3.29}$$

where (for simplicity) we have set the leak current reversal potential to zero. The “total conductance” at time t is $g_{tot}(t) = g_l + g_e(t) + g_i(t)$ and the “effective

input current” is $I_s(t) = g_e(t)E_e + g_i(t)E_i$.

Suppose that the stimulus affects the neuron via the synaptic conductances g_e and g_i . It is then natural to ask under which conditions, if any, the above model can correspond to a GLM. The definition of a GLM requires the solution $V(t)$ to be a linear (or affine) function of the stimulus. This arises if the two following conditions are met:

1. Total conductance g_{tot} is constant. Thus, for some constant c :

$$g_e(t) + g_i(t) = c. \quad (3.30)$$

2. The input I_s is linear in x . This holds if we set:

$$\begin{aligned} g_e(\mathbf{x}_t) &= (\mathbf{k}_e * \mathbf{x})(t) + b_e \\ g_i(\mathbf{x}_t) &= (\mathbf{k}_i * \mathbf{x})(t) + b_i. \end{aligned} \quad (3.31)$$

We can satisfy these two conditions by setting $\mathbf{k}_e = -\mathbf{k}_i$, so that the excitatory and inhibitory conductances are driven by equal and opposite linear projections of the stimulus. This allows us to rewrite the membrane equation (Eq. 3.29):

$$\begin{aligned} \frac{dV}{dt} &= -g_{tot}V + ((\mathbf{k}_e * \mathbf{x})(t) + b_e)E_e + ((\mathbf{k}_i * \mathbf{x})(t) + b_i)E_i \\ &= -g_{tot}V + (\mathbf{k}_{tot} * \mathbf{x})(t) + b_{tot}, \end{aligned} \quad (3.32)$$

where $g_{tot} = g_l + b_e + b_i$ is the (constant) total conductance, $\mathbf{k}_{tot} = \mathbf{k}_e E_e + \mathbf{k}_i E_i$, and $b_{tot} = b_e E_e + b_i E_i$. If we take the initial voltage V_0 to be b_{tot} , the equilibrium

voltage in the absence of a stimulus, then the solution to this differential equation is

$$\begin{aligned}
V_t &= \int_0^t e^{-g_{tot}(t-s)} ((\mathbf{k}_{tot} * \mathbf{x})(s)) ds + b_{tot} \\
&= \mathbf{k}_{leak} * ((\mathbf{k}_{tot} * \mathbf{x})(t)) + b_{tot} \\
&= (\mathbf{k}_{glm} * \mathbf{x})(t) + b_{tot},
\end{aligned} \tag{3.33}$$

where $\mathbf{k}_{leak} * (\mathbf{k}_{tot}^\top \mathbf{x}_t)$ denotes linear convolution of the exponential decay “leak” filter $\mathbf{k}_{leak}(t) = e^{-g_{tot}t}$ with the linearly projected stimulus train, and $\mathbf{k}_{glm} = \mathbf{k}_{tot} * \mathbf{k}_{leak}$ is the “true” GLM filter (from Eq. 3.26) that results from temporally convolving the conductance filter with the leak filter. Because the membrane potential is a linear (affine) function of the stimulus (as in Eq. 3.26), the model is clearly a GLM.

Thus, to summarize, the GLM can be equated with a synaptic conductance-based dynamical model in which the GLM filter \mathbf{k} results from a common linear filter driving excitatory and inhibitory synaptic conductances, blurred by convolution with an exponential leak filter determined by the total conductance.

From the above, it is easy to see how to create a more realistic conductance-based model of neural responses. We constructed the conductance-based spiking model (Section 3.1.2.1) which extends upon the GLM by allowing the stimulus tuning of excitation and inhibition to differ (i.e., allow $\mathbf{k}_e \neq -\mathbf{k}_i$), and includes a nonlinear relationship between \mathbf{x} and the conductances to preclude negative values (e.g., using a rectifying nonlinearity). As with the GLM, the only source of stochasticity on the CBSM is in the spiking mechanism: we

place no additional noise on the conductances or the voltage. This simplifying assumption allows us to perform efficient maximum likelihood inference using standard conjugate-gradient methods.

3.2.2 Approximating the input and output properties of ON parasol cells with intracellular recordings

Before applying the CBSM to characterize spiking responses in RGCs, we first analyzed intracellular recordings of ON parasol cells to validate and constrain the nonlinear components of the model.

3.2.2.1 Spike rate nonlinearity

We examined membrane voltages recorded from 2 ON parasol RGCs during a dynamic current clamp experiments to determine the spike rate as a function of membrane voltage (Fig. 3.2). The dynamic clamp recordings drove the RGCs with currents given by simulated conductances similar to those in the CBSM. We computed the firing rate as a function of the average membrane potential computed from several trials of the same simulated conductances. The firing rate function was closely approximated by the soft-rectification function

$$R_{sp} = a \log \left(1 + \exp \left(\frac{V - V_T}{\Delta V} \right) \right). \quad (3.34)$$

where where $a = 90sp/s$, $V_T = -53mV$ and $\Delta V = 1.67mV$.

Applications of the GLM have shown that an exponential nonlinearity captures the stimulus-to-spike rate relationship in RGC responses with greater accuracy than a soft-rectified function (Pillow et al. 2008). However,

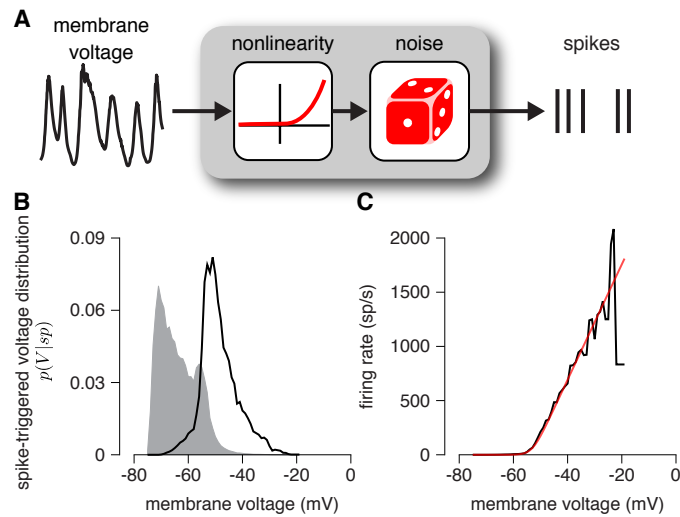


Figure 3.2: (A) Schematic of the transformation from voltage to firing rate of a neuron. The mean voltage recorded during repeated trials of current injection is passed through a nonlinear function to predict the spike rate. (B) The mean membrane voltage trace was computed by averaging voltages recorded over several repeats of the dynamic clamp condition. We computed the distribution of voltages observed in the mean trace (gray region) and compared this to the distribution conditioned on a spike occurring during one of the repeats (black trace) for two parasol RGCs. (C) The firing rate as a function of voltage (black trace) computed by applying Bayes' rule on the distributions shown in B. The firing rate function is closely approximated by a rectified linear function (red trace).

the voltage-to-firing rate function found here is much shallower than exponential, consistent with power-law relationships found in cortical neurons (Mechler & Ringach 2002, Priebe, Mechler, Carandini & Ferster 2004).

3.2.2.2 Conductances as a linear-nonlinear function of the stimulus

Ganglion cells receive excitatory synaptic input from bipolar cells (Gollisch & Meister 2010), and linear-nonlinear models can successfully characterize

the membrane potential response of bipolar cells to full-field visual stimuli (Rieke 2001). We therefore fit the measured conductances with a linear-nonlinear cascade model using a fixed nonlinear function (Fig. 3.3). The LN model with a soft rectified nonlinearity

$$f_e(\cdot) = \log(1 + \exp(\cdot)) \quad (3.35)$$

provided a close approximation to the observed conductances.

The LN model accounted for 76% of the variance of the mean excitatory conductance recorded in response to a novel stimulus. The LN model with the same fixed nonlinearity also accounted for 63% of the inhibitory conductances measured from the same cells.

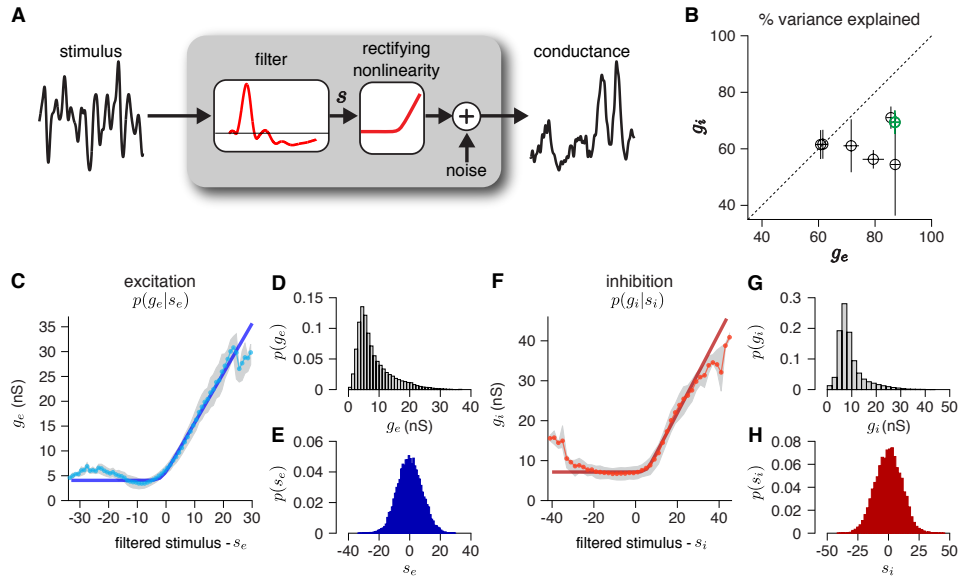


Figure 3.3: **(A)** We modeled the relationship between stimulus and synaptic conductances in a parasol RGC with an LN-cascade. The LN cascade assumed a fixed, rectifying nonlinearity. In order to fit filters directly to measured conductances, we assumed that the measured conductances contained additive Gaussian noise. Therefore, we fit the stimulus filter to predict the conductances by minimizing the squared error. (The CBSM, however, included the simplifying assumption that the conductances contain no noise.) **(B)** The percent variance explained for both excitatory and inhibitory conductances for 7 ON parasol RGCs, cross-validated using on a novel 6 s stimulus. The error bars indicate the standard deviation of the variance explained across all cross-validated stimuli. **(C)** The distribution of measured excitatory conductances given the filtered stimulus values for the example cell indicated in green in B. The gray region shows the middle 50-percentile of the distribution of observed excitatory conductance given the filtered stimulus value (s_e). The fixed soft-rectifying function (dark blue) closely matched the average conductance given the filtered stimulus value (light blue points). **(E)** The distribution of measured excitatory conductances for the example cell. **(D)** The distribution filtered stimulus values for the example cell given the excitatory stimulus filter. **(E)** The distribution inhibitory conductances conditioned on the filtered stimulus value for the same example cell. The fixed soft-rectifying function (dark red) provided a close approximation to the average inhibitory conductance given the filtered stimulus value (light red points). **(F)** The distribution of measured inhibitory conductances for the example cell. **(G)** The distribution filtered stimulus values for the example cell given the inhibitory filter.

3.2.3 Model fitting in simulations

To examine the performance of our numerical maximum likelihood estimation of the CBSM, we fit the parameters to simulated spike trains from the model with known parameters (Fig. 3.4). The first simulated cell qualitatively mimicked experimental RGC datasets, with input filters selected to reproduce the stimulus tuning of macaque ON parasol RGCs (excitation oppositely tuned and delayed compared to excitation, or “crossover” inhibition). The second simulated cell had similar excitatory tuning, but the inhibitory input had the same tuning as excitation with a short delay. The stimulus consisted of a one dimensional white noise signal, binned at a 0.1ms resolution, and filtered with a low pass filter with a 60Hz cutoff frequency. We validated our maximum likelihood fitting procedure by examining error in the fitted filters, and evaluating the log-likelihood on a 5-minute test set. With increasing amounts of training data, the parameter estimates converged to the true parameters for both simulated cells. Therefore, standard fast and non-global optimization algorithms can reliably fit the CBSM to spiking data, despite the fact that the model does not have the concavity guarantees of the standard GLM.

3.2.4 Predicting conductance tuning from spikes in ON parasol cells

We fit the excitatory and inhibitory conductance filters (and spike history filters) of the CBSM to spike trains recorded from 7 macaque on-parasol RGCs. The spike trains were obtained by cell attached recordings in response

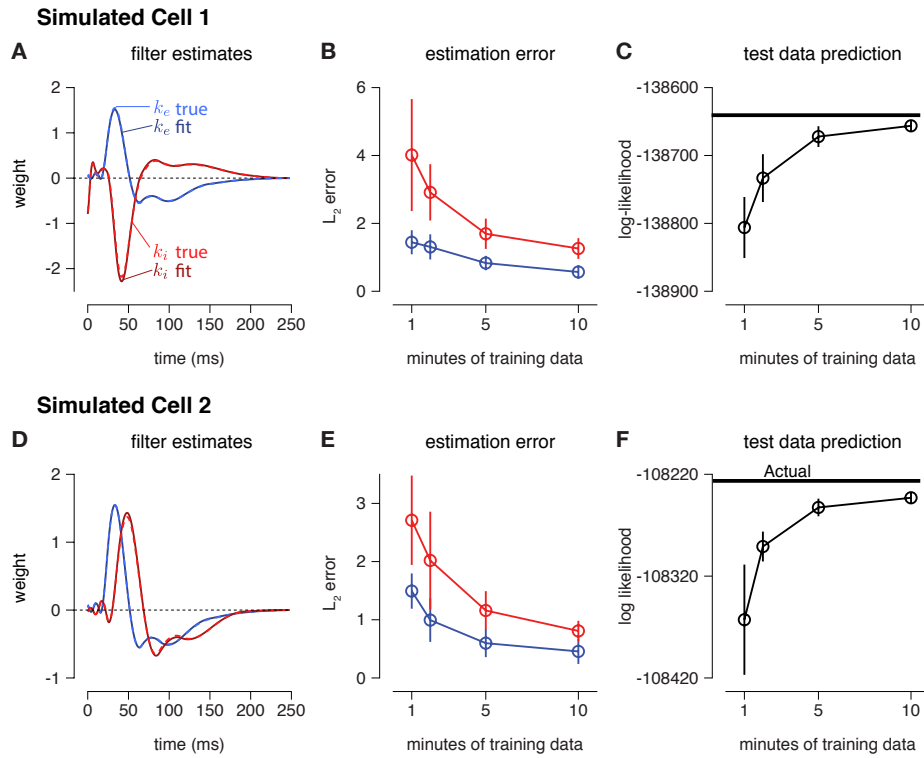


Figure 3.4: (A) Estimates (solid traces) of excitatory (blue) and inhibitory (red) filters from 10 minutes of simulated spike trains. (Dashed lines indicate true filters). The inhibitory filter was oppositely tuned and delayed compared to the excitatory filter. (B) The L_2 norm between the estimated input filters and the true filters as a function of the amount of training data. (C) The log-likelihood of the fit CBSM on the test data converged towards the log-likelihood of the true model. (D,E,F) Same as A-C for a simulation in which \mathbf{k}_i had the same tuning as \mathbf{k}_e with a delay.

to full-field, white noise stimuli (the same as in the simulations above). Either 30 or 40 trials were recorded from each cell, using 10 unique 6 second stimulus segments. After the spike trains were recorded, voltage clamp recordings were used to measure the excitatory and inhibitory conductances to the same 10 stimuli. We fit the model using the spike trains for 9 of the stimuli, and the

model fits were evaluated using the remaining 6 second stimulus. Thus, the models were effectively trained using 3 or 4 repeats of 54 seconds of full-field noise stimulus.

We compared the mean measured conductances in response to a test stimulus to the g_e and g_i predicted by the CBSM (Fig. 3.5). The CBSM correctly determined the stimulus tuning for excitation and inhibition for these cells: inhibition is oppositely tuned and slightly delayed from excitation. To visualize the measured conductances alongside the inferred conductances, we scaled the estimated conductances. Membrane voltage dynamics depend on the capacitance of the membrane, which we do not include because it introduces an arbitrary scaling factor that cannot be estimated from spikes alone. Therefore, for comparisons we chose a scaling factor for each cell to minimize the squared error between the predicted conductances and the measured conductances, while assuming a common scaling for the inhibitory and excitatory conductances. We similarly scaled the stimulus filters fit to the conductance measurements to match the average height of the CBSM filters.

The CBSM predicted the synaptic conductances showed strong correlation with the measured conductances for all 7 cells. Using only a few minutes of spiking data, the conductances inferred from the spike trains showed an average correlation of $r = 0.69$ for the excitatory input and $r = 0.68$ for the inhibitory input, compared to an average $r = 0.87$ and $r = 0.83$ for the excitatory and inhibitory conductances respectively from the LN fit directly to the conductances (Fig. 3.6).

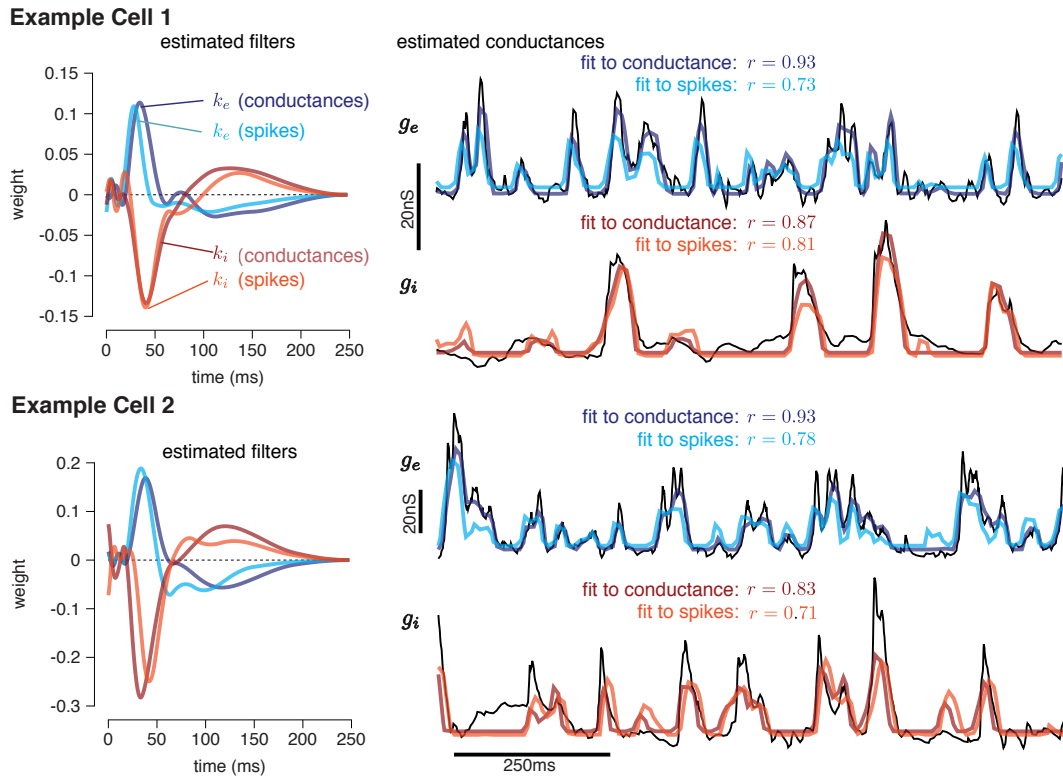


Figure 3.5: Two example ON parasol RGC responses to a full-field noise stimulus fit with the CBSM. The model parameters were fit to spike train data, and then used to predict excitatory and inhibitory synaptic currents recorded separately in response to novel stimuli. For comparison, we show predictions of an LN model fit directly to the conductance data. **Left:** Linear kernels for the excitatory (blue) and inhibitory (red) inputs estimated from the conductance-based model (light red, light blue) and estimated by fitting a linear-nonlinear model directly to the measured conductances (dark red, dark blue). The filters represent a combination of events that occur in the retinal circuitry in response to a visual stimulus, and are primarily shaped by the cone transduction process. **Right:** Conductances predicted by our model on a withheld test stimulus. Measured conductances (black) are compared to the predictions from the CBSM filters (fit to spiking data) and an LN model (fit to conductance data).

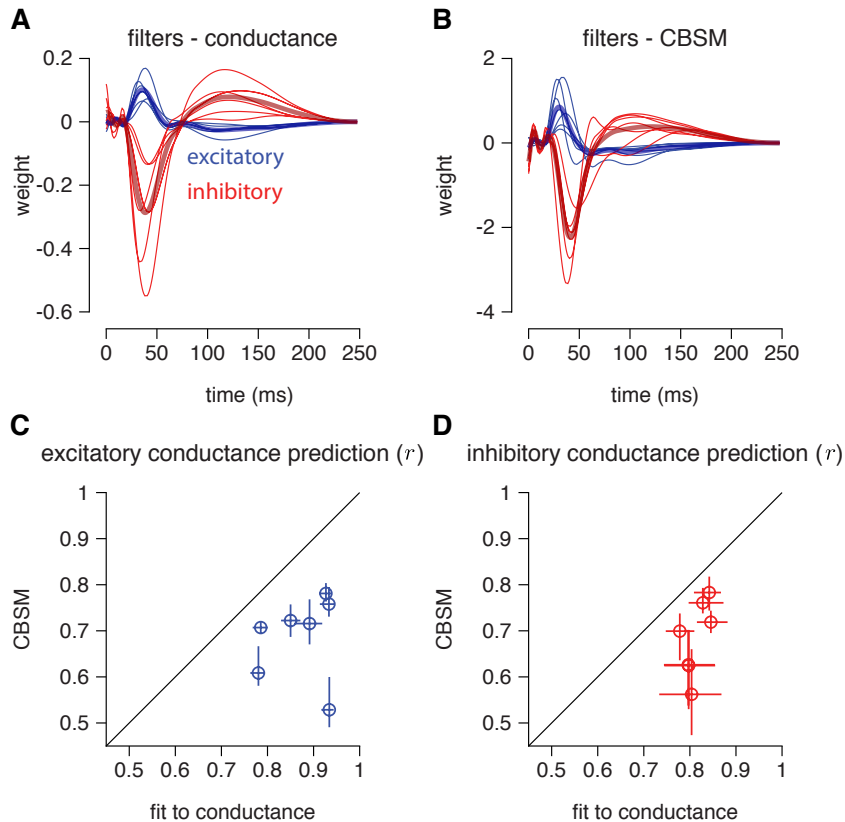


Figure 3.6: Summary of the CBSM fits to 7 ON parasol RGCs for which we had both spike train and conductance recordings. **(A)** The excitatory (blue) and inhibitory (red) filters estimated from voltage-clamp recordings. The thick traces show the mean filters. **(B)** The excitatory (blue) and inhibitory (red) filters estimated by the CBSM from spike trains. **(C)** The correlation coefficient (r) between the mean observed excitatory synaptic input to a novel 6 s stimulus and the conductance predicted by the LN cascade fit to the excitatory conductance (abscissa) and the CBSM prediction from spikes (ordinate) for each cell. Error bars denote the minimum and maximum of r values observed across all cross-validated stimulus **(D)** Same as C for the inhibitory conductance.

3.2.5 Characterization of RGC responses by inferring excitatory and inhibitory inputs

We fit the CBSM and GLM to a population of 9 extracellularly recorded macaque RGCs in response to a full-field binary noise stimulus (Pillow et al. 2005). The models were fit to 10 minutes of spike train recordings. We compared predicted spike rate of the CBSM to the GLM to a 5 s test stimulus for which we had repeated trials, and compared the predictive log-likelihood for a single a 5 minute recording with a novel test stimulus (Fig. 3.7). All cells showed an improved fit with the CBSM. The CBSM on average explained 86% of the variance of the PSTH, while the GLM on captured 77% of the spike rate. The CBSM also improved spike train prediction (0.34 ± 0.11 bits/sp over the GLM).

Previous experiments have indicated that inhibition only weakly modulates parasol cell responses to full-field stimuli (Cafaro & Rieke 2013). To test the effect inhibition in the model, we also fit the CBSM without any inhibitory synaptic conductance (CBSM_{exc}). The CBSM_{exc} still provided a superior prediction of the PSTH than the GLM (81% of the variance explained) and a larger cross-validated log-likelihood (mean improvement of 0.14 ± 0.10 bits/sp over the GLM. Fig. 3.8). Thus, inhibition only accounts for 48% of the CBSM's spike rate prediction performance over the GLM, and excitatory input alone can account for most of the RGC firing rate in response to a full-field, single-contrast noise stimulus. We compared the excitatory filters estimated by the CBSM_{exc} with the GLM filters and found that the filters are nearly identical

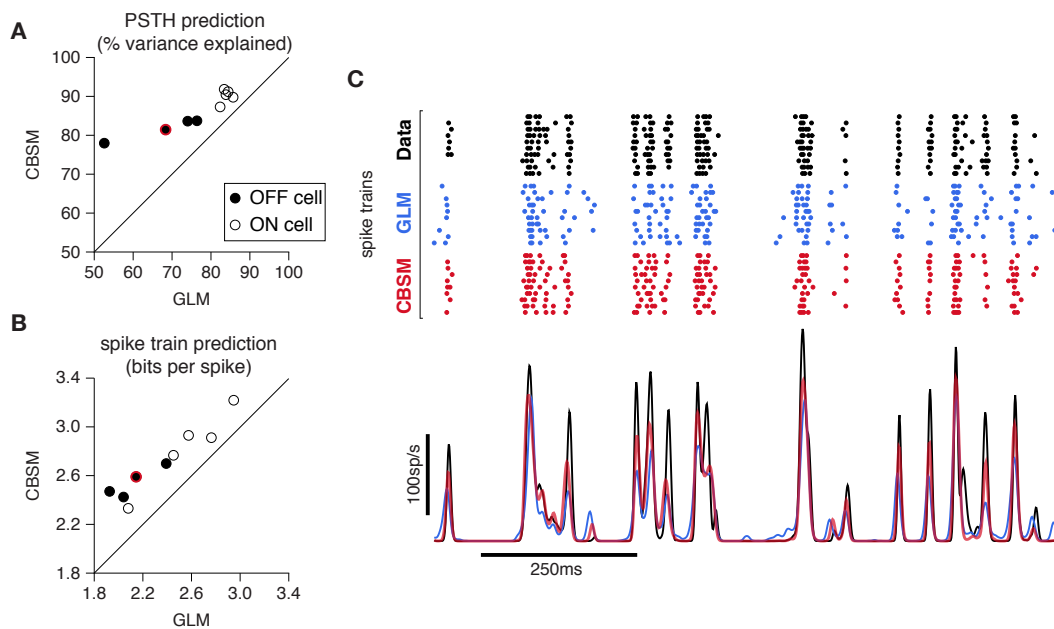


Figure 3.7: (A) Spike rate prediction performance for the population of 9 cells for 5 s test stimulus. The true rate (black) was estimated using 167 repeat trials. The red circle indicates the cell shown in C. (B) Log-likelihood of the CBSM compared to the GLM computed on a 5 min test stimulus. (C) (top) Raster of responses of an example OFF parasol RGC to repeats of a novel stimulus (black) and simulated responses from the GLM (blue) and the CBSM (red). (bottom) Spike rate (PSTH) of the RGC and the GLM (blue) and CBSM (red). The PSTHs were smoothed with a Gaussian kernel with a 2 ms standard deviation.

(Fig. 3.9). The GLM appears to account only for the excitatory input received by the cell. The CBSM improves characterization of the RGC responses both by including an inhibitory input, and by treating the excitatory input as a synaptic conductance in a simple biophysical model with data-constrained nonlinearities.

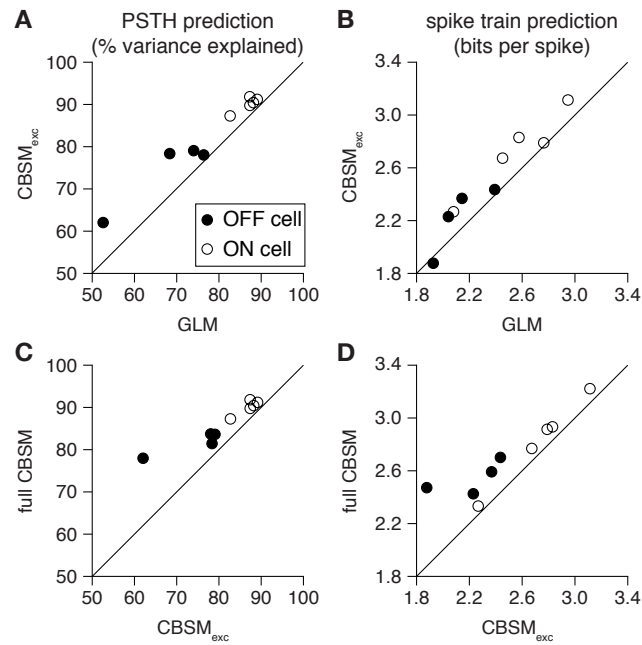


Figure 3.8: (A) Spike rate prediction performance and (B) cross-validated log-likelihood for the population of 9 cells for 7 s test stimulus for the GLM and the CBSM with only an excitatory input term ($CBSM_{exc}$). (C) The full CBSM with inhibition shows improved spike rate prediction and (D) cross-validated log-likelihood compared to the model without inhibition.

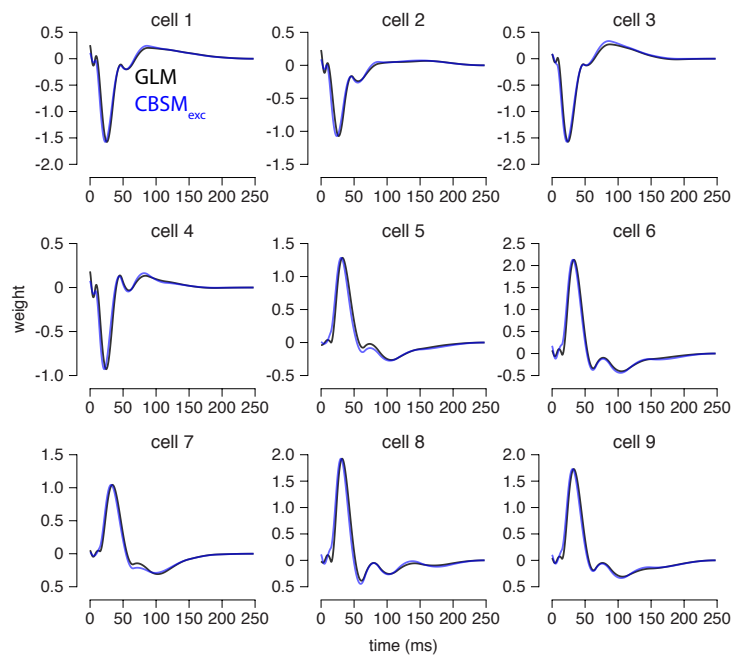


Figure 3.9: The GLM filters for 9 parasol RGCs (black) compared to the excitatory conductance filters estimated by the CBSM without an inhibitory input (blue). The GLM filters are shown scaled to match the height of the CBSM_{exc} filters.

3.2.6 RGC responses across different temporal contrasts

The characterization of adaptation to changes in stimulus statistics in sensory neurons has relied heavily on studying changes to single linear filter fit to a cell’s activity across different stimulus conditions (Chander & Chichilnisky 2001, Fairhall et al. 2001, Baccus & Meister 2002, Garvert & Gollisch 2013, Marava 2013). RGCs exhibit contrast gain control: as contrast (the variance of the luminance) increases, the sensitivity of the cell decreases. The GLM reveals this behavior by showing a contrast-dependent linear filter height (Fig. 3.10A).

We aimed to determine if the CBSM using a fixed pair of stimulus filters could predict gain scaling to temporal contrast. We fit 8 RGCs to responses to a binary, full-field stimulus at 12%, 24%, and 48% contrast. We then followed a procedure similar to an experimental paradigm to used to study contrast adaptation in real cells: we simulated from the CBSM fits using stimuli at all 3 contrast levels and we characterized the simulated fits with a GLM to obtain a linear filter (Fig. 3.10B). The CBSM simulations exhibited a similar scaling in linear filter in response to changes in contrast as the real cells. The change in filter heights of the CBSM simulations relative to the height of the 48% contrast filter correlated well with the filter height change seen in the GLM fits to the real cells ($r = 0.61, p < 0.05$; Fig. 3.10C). Therefore, much of the change in filter height over contrast levels can be reproduced by considering the intrinsic properties of a fixed model with excitation and inhibitory conductance inputs.

Classic LN models with single linear filter often fail to generalize over

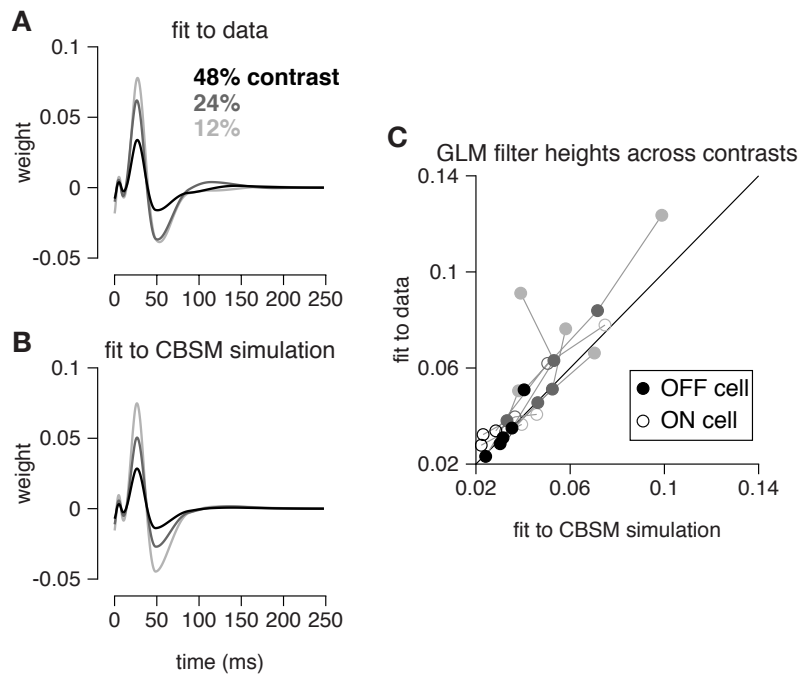


Figure 3.10: (A) GLM filters for an example ON cell fit to responses recorded at 12%, 24%, and 48% contrast. (B) GLM filters fit to spike trains simulated from the CBSM fit to the cell shown in A. The CBSM was fit to responses at all 3 contrast levels. (C) The filter heights (the absolute value of the peak of the filter) of the GLM fits to 8 cells at all three contrast levels (one point per contrast level per cell), compared to the GLM filters fit the CBSM simulations of those same cells.

stimulus conditions because of adaptation effects, and connecting the changes in a linear filter to specific biophysical mechanisms is difficult. Understanding how neural coding during natural stimuli and its biophysical implementation will require a single modeling framework that can explain RGC responses across stimulus conditions (Ozuysal & Baccus 2012, Clark, Benichou, Meister & Azeredo da Silveira 2013). To quantify the CBSM's ability to generalize

over coherences, we compared the model predictions to test stimuli at each contrast level with a model fit only to the test contrast level. We found that the CBSM showed greater ability to maintain predictive performance over the different contrast levels compared to the GLM (Fig. 3.11). At 12% contrast, the GLM fit to all contrasts lost an average 0.36 ± 0.41 bits/sp compared to GLM fit specifically to the 12% contrast stimulus, while the CBSM lost only 0.16 ± 0.2 bits/sp. At 24% contrast, the GLM lost 0.20bits/sp while CBSM only lost 0.07 ± 0.14 bits/sp. Finally, both models only lost 0.05 ± 0.08 bits/sp in the 48% contrast probe. Removing the inhibitory filter showed that the CBSM required inhibition to generalize over contrast levels (a loss of 0.95 ± 0.88 , 0.85 ± 1.20 , and 0.59 ± 0.65 bits/sp at the 12%, 24%, 48% probes respectively).

3.2.7 Excitatory and inhibitory responses to spatial stimuli

We next analyzed the responses of a population of 27 parasol cells in response to a spatio-temporal noise stimulus (Pillow et al. 2008) to explore how the excitatory and inhibitory inputs interact to shape the spatial selectivity of RGCs. The temporal response profiles of excitation and inhibition in the CBSM were qualitatively similar to those that we observed in the full-field stimulus condition (Fig. 3.12). However, the time-variant spatial profile of the filters could allow the synaptic inputs to have different temporal interactions compared to the full-field stimulus.

The CBSM captured the mean firing rate more closely than the GLM (average percent variance explained 83% and 79% for the CBSM and GLM

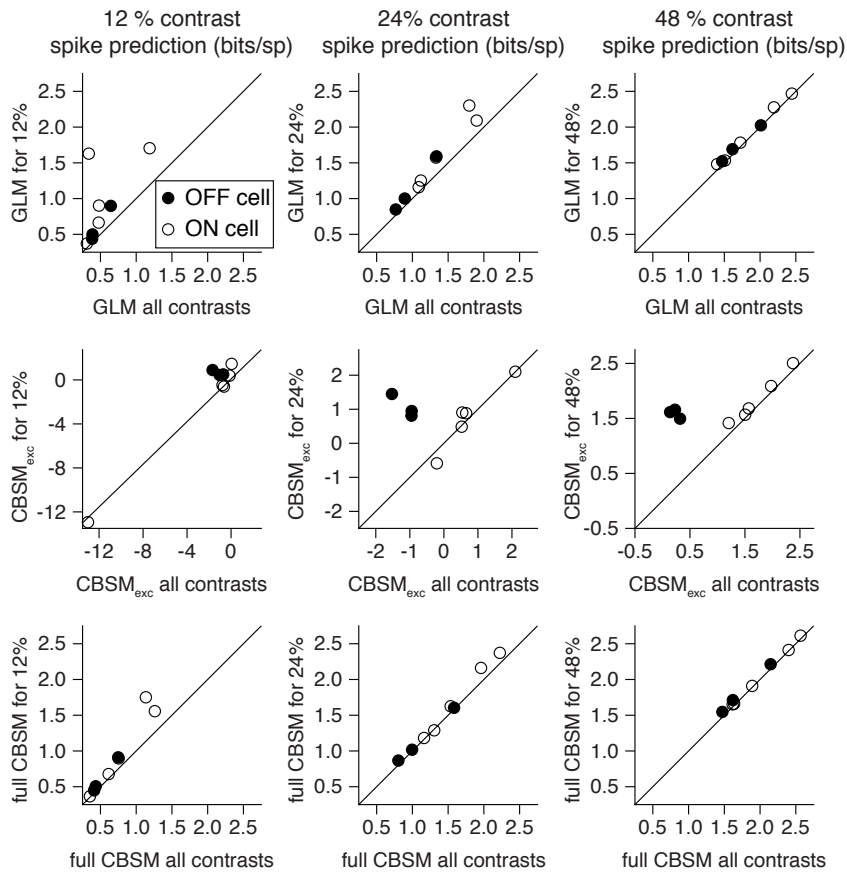


Figure 3.11: Spike train performance of the GLM (top row), CBSM with only excitatory input (middle row), and the CBSM (bottom row). The model is tested on a 4 minute stimulus at 12% (left column), 24% (middle column), and 48% (right column) contrast. The model fit to all 3 contrast levels (abscissa) is plotted against the same class of model trained only on the probe contrast level (ordinate).

respectively). The CBSM also predicted the single-trial responses with higher accuracy than the GLM (an average of 0.07 ± 0.04 bits/sp greater than the GLM). Removing the inhibitory conductance from the CBSM showed that the excitatory input alone provided a superior account of the mean firing

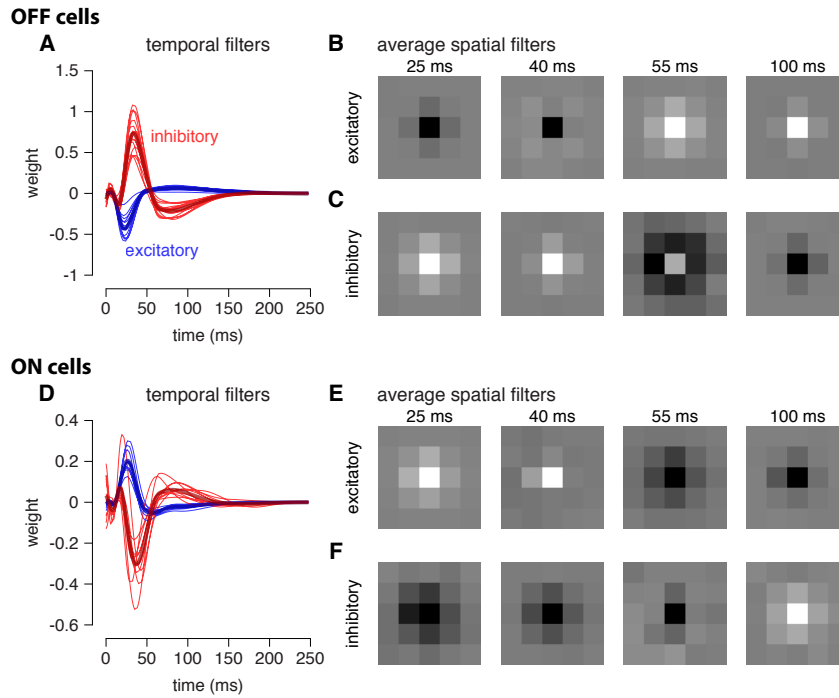


Figure 3.12: CBSM fits to a population of 27 RGCs. (A) Temporal profile of the excitatory (blue) and inhibitory (red) at the center pixel of the receptive field for 16 OFF parasol cells. The thick lines show the mean. (B) The mean spatial profiles of the excitatory and (C) inhibitory linear filters at four different time points for the OFF parasol cells. (D,E,F) same as A,B,C for 11 ON parasol cells.

rate (average percent variance explained 81%), but the single-trial spike train prediction was reduced by removing inhibition (an average of 0.02 ± 0.04 bits/sp greater than the GLM).

We simulated from the model using a binary, center-surround stimulus in which the center pixel was the opposite sign contrast of the outer pixels to explore how accounting for distinct excitatory and inhibitory inputs affects the characterization of spatial selectivity. Spatio-temporal noise and full-field noise

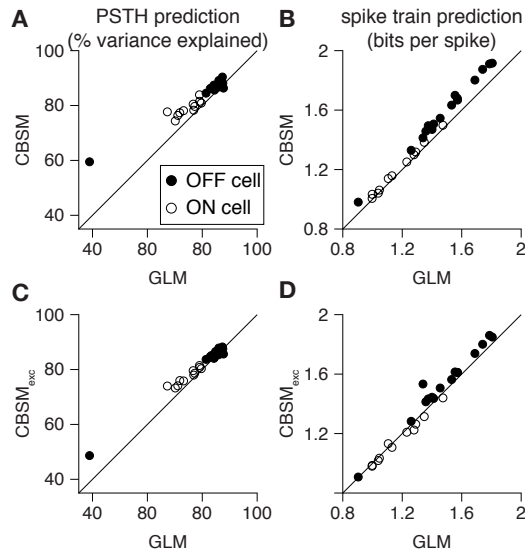


Figure 3.13: (A) Spike rate prediction performance of the CBSM compared to the GLM for the population of 27 cells for 8 s test stimulus. The true rate (black) was estimated using 600 repeat trials. (B) Log-likelihood of the CBSM compared to the GLM computed on a 5 min test stimulus. (C) Spike rate prediction performance of the CBSM_{exc} compared to the GLM. (D) Log-likelihood of the CBSM_{exc} compared to the GLM.

stimuli produced similar cross-correlation between the CBSM’s excitatory and inhibitory conductances (Fig. 3.14A,B). In response to these stimuli, the excitatory and inhibitory conductances showed a strong negative correlation with excitation preceding inhibition. The random opposing center-surround stimulus produced a distinct cross-correlation pattern with a larger positive peak at the positive lags. Additionally, we simulated responses to steps of the antagonistic center-surround stimulus. The CBSM and GLM showed similar onset responses, but the sustained responses of the CBSM simulations showed inhibition-dependent suppression for both ON and OFF cells (Fig. 3.14C).

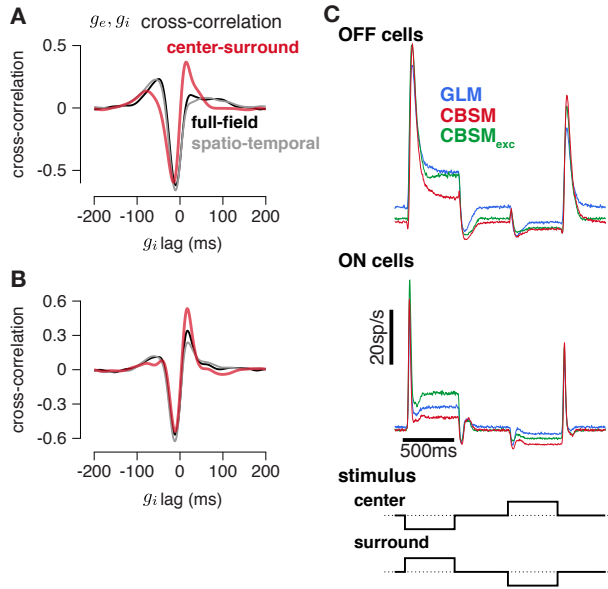


Figure 3.14: (A) The mean cross-correlation of the CBSM predicted excitatory and inhibitory conductances for the OFF cells in response to full-field noise (black), spatio-temporal noise (grey), and opposing center-surround noise (red). The spatio-temporal stimulus was the same binary noise stimulus used to fit the cells. The full-field stimulus consisted of binary noise at the same contrast and frame rate as the original spatio-temporal stimulus. In the opposing center-surround condition, the center pixel was of opposite contrasts to the surround pixels and the sign of the center pixel was selected randomly on each frame. The strong negative component showed that g_i is delayed and oppositely tuned compared to g_e . (B) Same as B for the ON cells. (C) Average firing rate of the GLM (blue), CBSM (red), and CBSM_{exc} (green) fits to 16 OFF cells (top) and 11 ON cells (middle) in response to opposing center-surround contrasts steps (bottom).

Thus, spatial correlations in the stimulus can alter the temporal structure of excitatory and inhibitory interactions in the CBSM that shape the model's response differently than the GLM.

3.3 Conclusion

The classic GLM serves as a valuable tool for describing the relationship between stimuli and spike responses. However, the GLM describes this map as a mathematically convenient linear-nonlinear cascade, which does not take account of the biophysical properties of neural processing (Herz, Gollisch, Machens & Jaeger 2006). This problem persists in abstract mathematical extensions to the traditional LN model that include quadratic terms (Schwartz, Chichilnisky & Simoncelli 2002, Agüera y Areas, Fairhall & Bialek 2001, Park & Pillow 2011), or additional linear-nonlinear processing steps (Ahrens et al. 2008, Butts et al. 2011, McFarland et al. 2013). Here we have shown that the GLM may be interpreted as a biophysically inspired, but highly constrained, synaptic conductance-based model. We proposed a more realistic model of the conductance, removing the artificial constraints present in the GLM interpretation, which results in a new, more accurate and more flexible conductance-based point process model for neural responses. Even without the benefit of a concave log-likelihood, tractable numerical optimization methods provide accurate estimates of model parameters.

The model also allows the excitatory and inhibitory conductances to be distinct functions of the sensory stimulus, as is expected in real neurons. We used intracellular measures to validate and constrain the nonlinear functions in the CBSM, in contrast to previous cascade modeling approaches which typically select these terms in order to maximize a likelihood function or based on computational convenience. We demonstrated that the CBSM not only

achieves improved performance as a phenomenological model of neural encoding compared to the GLM, the model accurately estimates the tuning of the excitatory and inhibitory synaptic inputs to RGCs purely from measured spike times. The CBSM has a stimulus-dependent time constant, which allows it change its gain as a function of stimulus statistics (e.g., contrast), an effect that cannot be captured by a standard GLM. Although the GLM is able to produce spike rate predictions to spatio-temporal noise stimuli that are almost as accurate as the CBSM, the interaction of the excitation and inhibition in the CBSM results in different predictions about the selectivity of RGCs to spatially structured, center-surround stimuli compared to the GLM. As we move towards more naturalistic stimulus conditions that have complex spatio-temporal statistics, we believe that the conductance-based approach will become a valuable tool for understanding the neural code in sensory systems.

Chapter 4

Discussion

The statistical tools presented in this dissertation targeted the neural coding problem at the level of single neurons in two stages of the primate visual system.

In Chapter 2, we applied latent variable models to compare candidate models of decision formation to single-trial spike trains recorded during a decision-making task. In stark contrast to previous findings, Bayesian model comparison revealed that decision-related activity in LIP neurons was better captured by a discrete stepping model than a continuous evidence accumulation process. Therefore the diffusion-to-bound model does not directly link the behavior of the monkey in the motion-discrimination task to single-neuron responses recorded in area LIP. These results demonstrate the limitations of “neural correlates” for studying decision making processes in the brain. While much of systems neuroscience has relied on viewing neural data in relation to external variables, applying modern statistical techniques enables us to examine neural data aligned to internal states to reveal new insights about the brain.

The modeling tools presented in Chapter 3 demonstrated that a more

accurate characterization of neural responses in the retina can be achieved by relating statistical models of neurons to the biological substrate. Using a simplified model of membrane dynamics and synaptic input, we found that we could recover the basic stimulus tuning of the excitatory and inhibitory input received by RGCs. Characterizing these two input sources separately described RGC responses more accurately than a single linear-nonlinear model. The conductance-based approach was more robust to changes in temporal contrast, and the model produced different predictions about spatial integration than the classic LN model approach.

4.1 Future Directions: Modeling decision representations in multi-cell recordings

The analyses presented in Chapter 2 examined what can be gleaned about the dynamics of the neural representation of perceptual decisions from single neurons recorded in LIP. To gain a true understanding of the dynamics of how a neural circuit integrates input will require observing and analyzing a larger portion of the neural circuit (Churchland, Cunningham, Kaufman, Ryu & Shenoy 2010, Mante, Sussillo, Shenoy & Newsome 2013, Kaufman, Churchland, Ryu & Shenoy 2014). Because decision-related activity is observed in multiple brain areas, including the medial intraparietal area (MIP) (de Lafuente et al. 2015), a complete study of perceptual decision-making must also examine the interactions between different brain regions. Collecting large-scale neural recordings, while challenging, is rapidly becoming more com-

mon (Stevenson & Kording 2011). However, many challenges for analyzing large neural datasets exist (Freeman, Vladimirov, Kawashima, Mu, Sofroniew, Bennett, Rosen, Yang, Looger & Ahrens 2014). Latent variable methods, like the models we described in Chapter 2 for single neurons, can provide a way to link the observed activity from multiple neurons, even neurons recorded from several regions, to a common variable. This section describes preliminary work modeling pairs of simultaneously recorded neurons.

4.1.1 Confirming single-cell results in a new dataset

We examined a set of 72 LIP neurons recorded while a monkey performed a motion discrimination task similar to the one described in Section 2.5 (Yates, Park, Cormack, Pillow & Huk 2013). The motion stimulus consisted of a set of drifting gabors instead of random-dot motion, and the motion stimulus was presented for 1.166 s on every trial. As with the earlier dataset, the top 50% (36 of the 72 cells) of choice-selective cells were chosen for further analysis. This selection included 19 pairs of simultaneously recorded neurons.

We applied the same model comparison analysis as before to the individual LIP cells in this new dataset. A majority of the 36 cells were better fit by the stepping model (Fig. 4.1), confirming the initial findings that the stepping model better captured the responses of LIP cells during the tentative integration period than the ramping model.

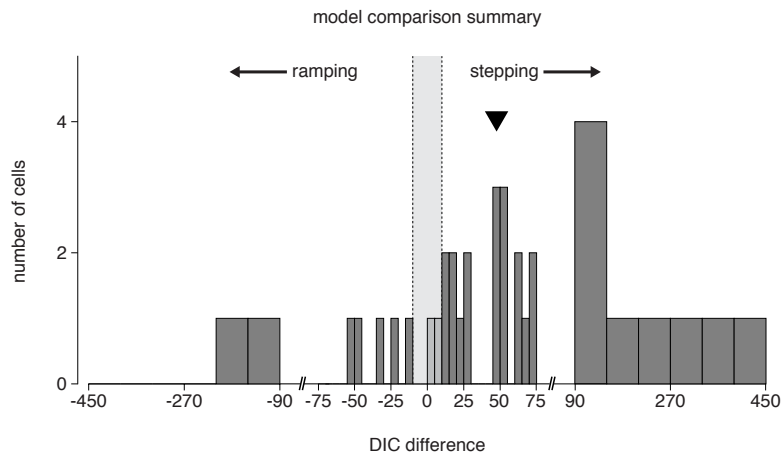


Figure 4.1: Model comparison results for 36 LIP neurons from a new dataset. The Δ DIC score shows that 29 of the 36 cells were fit better by the stepping model than the ramping model, and 27 of those cells showed strong support for stepping.

4.1.2 Are step times across different neurons correlated?

As a first step towards analyzing the representation of decisions in LIP populations, we examined the results of the stepping model fit to the individual cells. We then compared the step times inferred by the model across simultaneously recorded cell pairs (Fig. 4.2). Many cell pairs did not show significant correlations in step times across pairs of simultaneously recorded LIP neurons in this population, even though the majority of these cells were described better by the stepping model than the ramping model in the single-cell analysis. The mean correlation coefficient over all the pairs was centered near 0.

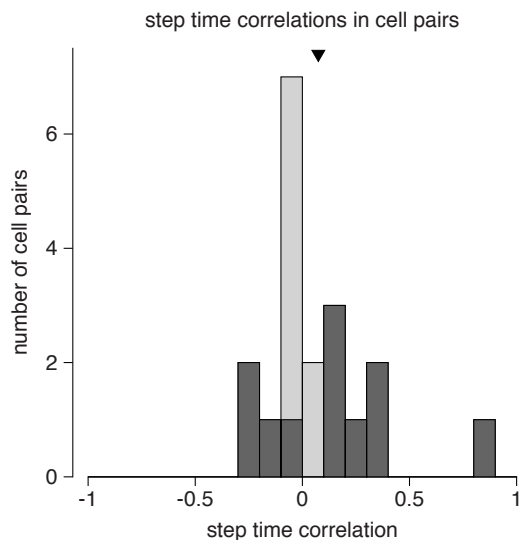


Figure 4.2: Correlation of the step times inferred by the stepping model between pairs of simultaneously recorded LIP cells. Cell pairs with a correlation coefficient that is statistically significant from zero ($p < 0.05$) are in dark gray. The mean of the correlations (black triangle) is slightly positive, but this is not statistically significant.

4.1.3 Extending the stepping model to multiple neurons

If the discrete state picked up by the stepping model truly reflected the actual decision signal, one might expect that all LIP cells with RFs containing the a choice target would step up or down together. To test this, we modeled the responses of the neuron pairs with a simultaneous step. The stepping model can easily be extended for multiple neurons by assuming that each observed cell is the product of a common latent process. We define the multi-

cell, simultaneous stepping model as

$$z_j \sim \text{Negative Binomial}(p_{c(j)}, r) \quad (4.1)$$

$$P(d_j = 1) = \phi_{c(j)} \quad (4.2)$$

$$P(d_j = 2) = 1 - \phi_{c(j)} \quad (4.3)$$

$$y_{k,j,t} | t \leq z_j \sim \text{Poisson}(\alpha_{k,0} \Delta_t) \quad (4.4)$$

$$y_{k,j,t} | t > z_j \sim \text{Poisson}(\alpha_{k,d_j} \Delta_t) \quad (4.5)$$

where the index k denotes the k -th neuron. This simple extension assumes that a step occurs simultaneously across all cells, without any temporal offsets. The primary difference between this model and the single-cell version in Section 2.2.2 is that there are multiple firing rates for each state, one for each neuron. This model is much more restrictive than more general HMM methods used to examine populations during decision-making tasks (e.g., Miller & Katz 2010, Bollimunta et al. 2012).

MCMC under this model proceeds almost identically to the MCMC for the single-cell model. There are two small differences. The first difference is that, to sample the latent variables, the likelihood term becomes a product of Poisson observations

$$p(y_{\cdot,j,t} | z_j, d_j, \alpha) = \prod_{k=1}^K p(y_{k,j,t} | z_j, d_j, \alpha_{k,\cdot}) \quad (4.6)$$

The second difference is that α values must be sampled for each neuron. This step is performed identically to the sampling step in the single neuron case,

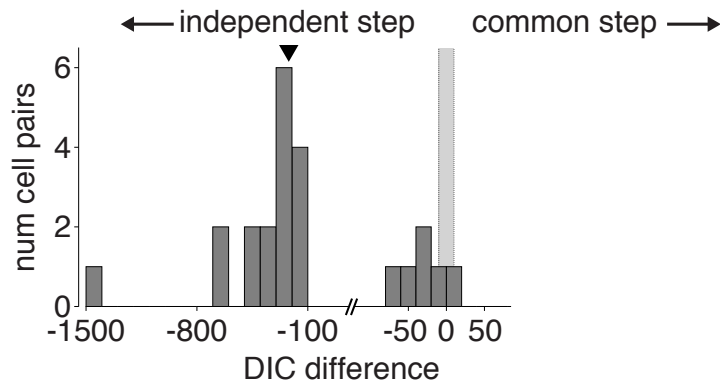


Figure 4.3: Model comparison results between a simultaneous stepping model and the single-cell stepping model which considers the cells independently. The model comparison strongly supports independent stepping.

but it is performed K times: once for each neuron (or once for each set of α_k , given $y_{k,\cdot,\cdot}$, d , and z).

For the 19 LIP cell pairs, we compared the simultaneous stepping to independent, single-cell stepping model fits. We found for these cell pairs, quantitative model comparison did not support the hypothesis of a common stepping process (Fig. 4.3). This is not surprising considering that the single-cell fits did not find strong evidence of correlated step times.

This preliminary analysis demonstrated that latent variable models can be applied to test hypotheses of representations of an evolving decision variable in a population of cells. The simple hypothesis test here shows that, while single LIP neurons were better described by a discrete-stepping model than a diffusion-to-bound process, the responses of pairs of LIP neurons were not well-described by a simultaneous state change. One possible interpretation is that

the steps seen in single cells act as “votes” for a particular choice. Therefore, the population could encode a more continuous decision-related variable even though the individual neurons appear to represent a discrete value. Future work, including analyses of more than two cells at a time, will be required to uncover the nature of the dynamics of the LIP population during decision-making.

4.2 Future directions: conductance-based modeling of RGCs

The CBSM took the GLM one step closer to a biophysical model of the neuron, but many properties of the biophysical neuron were not included in the model. Further developments to the CBSM can include additional sets of rectified, conductance inputs (McFarland et al. 2013). For example, the model could include multiple spatially distinct inputs to account for input from different bipolar or amacrine cells. Spatially selective rectification of inhibitory inputs helps determine RGC responses to spatial stimuli (Cafaro & Rieke 2013), and adaptation can occur in localized regions of a ganglion cell’s RF (Garvert & Gollisch 2013). Additionally, the CBSM could also be applied to explore the role of active conductances that depend spike history, such as an after hyper-polarization current, on the coding properties of sensory neurons. Spike-dependent conductances could also be included to within a population; although the analyses presented here focused on the coding properties of single neurons, many of the RGCs analyzed were recorded simultaneously (Pillow

et al. 2008).

Additionally, the spike generation of the CBSM, while constrained by dynamic clamp recordings, was a simple statistical approximation of the true spike generation process. The intrinsic properties of spike generation in real neurons can contribute to adaptive coding (Mease, Famulare, Gjorgjieva, Moody & Fairhall 2013). Including a more biophysically accurate spiking process could lead to a more experimentally applicable and accurate model.

The CBSM confined all noise in the model to the spike generation mechanism. Assuming deterministic conductances and voltages made the CBSM a tractable model that could be efficiently fit with standard conjugate-gradient methods. However, the synaptic conductances in real neurons are also variable from trial to trial, and the correlation of the noise between the two conductances can change the reliability of RGC spike trains (Cafaro & Rieke 2010). Latent variable models could extend the model to include variability in the conductances and voltage process (Paninski et al. 2012). The added flexibility would allow the model to match more closely neural data than the CBSM, but would also increase the difficulty of model fitting both in terms of the amount of data needed and the computational cost.

Finally, neural recordings analyzed here were exclusively driven by simple artificial noise stimuli. Solving the neural coding problem will ultimately require understanding the code during naturalistic stimulation, rather than experimentally convenient conditions (Carandini, Demb, Mante, Tol-

hurst, Dan, Olshausen, Gallant & Rust 2005, van Hateren, Rüttiger, Sun & Lee 2002, Butts, Weng, Jin, Yeh, Lesica, Alonso & Stanley 2007). Future work should aim to determine how excitatory and inhibitory synaptic inputs combine to shape RGC responses to natural stimuli.

4.3 Conclusion

The specific statistical tools we select to analyze data can transform our perception of the neural code. Ensuring that a statistical tool can truly capture the essence of a hypothesis we propose about neural coding is vital to construct a proper understanding of information processing in the brain, whether that hypothesis pertains to the integration of inputs that leads to spiking in single neurons or the network-level computation of a decision. Here we found that analyzing the neural representation of perceptual decisions in LIP with modern latent variable models leads to different conclusions than by viewing the data through the conventional lens of averaging spike trains over many trials. Unlike analyses of summary statistics like the PSTH, latent variable methods can relate neural activity directly to theories of perceptual decision-making because such decisions are single-trial events and not averages. Although the average firing rate of an LIP neuron appears as a correlate of a diffusion-to-bound process, modeling single-trial activity produced a different view of the data: choice-related activity in LIP on single trials is better explained by a discrete stepping process. In the retina, we developed a model relating visual stimuli to spike train responses in ganglion cells that was not only more accurate than

the traditional GLM, but the model provided a more general description of the data by accounting for gain scaling across temporal contrasts. This was accomplished not by constructing a purely phenomenological model, but by expanding an existing model (the GLM) towards a more realistic biophysical description of a neuron that integrates synaptic inputs.

As modern experimental methods advance to allow more and more neurons to be recorded simultaneously, statistical methods for analyzing neural must also advance in order to connect data to theory (Stevenson et al. 2012, Freeman et al. 2014). Simple statistical techniques, like computing the PSTH, have been invaluable for providing intuitive explanations of single-neuron responses, especially in early sensory regions. However, simply averaging the activity of single neurons may not be an effective method to analyze larger scale recordings, because even cells within the same region can exhibit heterogeneous properties (Meister et al. 2013) and correlations between cells can play an important role in the neural code (Pillow et al. 2008, Cohen & Kohn 2011). Additionally, essential aspects of neural responses recorded during complex tasks may be washed away by averaging (Goldman 2015). Advanced latent variable models will certainly play a central role in the future of systems neuroscience by tying together the activity of many brain regions in order to make sense of the dynamics in the whole population.

Appendices

Appendix A

Methods used to analyze LIP response

A.1 Computing PSTHs

The coherence-sorted spike rates (PSTH) and variances (PSTV) aligned to motion onset in Fig. 2.16 were computed for each cell using a 25 ms sliding window (boxcar filter) moved by 5 ms increments. At each time point, trials were only included in the average if the end of the window of the analysis (200 ms after motion offset) had not been reached. The population rates and variances were computed by averaging the rates and variances from all 40 cells.

I compared the model predictions of the coherence sorted PSTH/PSTV to the true population spike rate and variance during the interval 205-700 ms after motion onset (100 time points). I simulated 1000 spike trains per coherence level per cell using the posterior-mean parameters from each model to obtain the coherence-sorted model predicted rate (MR) and variance (MV). I then calculated the fraction of variance explained in the coherence-dependent PSTH (or PSTV):

$$R^2 = 1 - \frac{\sum_{c=1}^5 \sum_{t=205}^{700} (MR_{c,t} - PSTH_{c,t})^2}{\sum_{c=1}^5 \sum_{t=205}^{700} (\overline{PSTH} - PSTH_{c,t})^2} \quad (\text{A.1})$$

$$\overline{PSTH} = \frac{1}{5} \frac{1}{100} \sum_{c=1}^5 \sum_{t=205}^{700} PSTH_{c,t} \quad (\text{A.2})$$

where $PSTH_{c,t}$ is the PSTH at time t (in milliseconds) for coherence level c . The average spike rate over all time and coherence levels is \overline{PSTH} . The sums over t are in increments of 5 ms.

Credible intervals on the R^2 values included uncertainty in both the measured PSTH/PSTV as well as the model fit uncertainty. I obtained 1000 samples of the data PSTH (PSTV) by randomly drawing a set of trials with replacement for each cell and computing the population PSTH (PSTV) with those trials. Each sample used the number of trials per coherence as were actually observed for each cell. I obtained errors on the model PSTH (PSTV) by using the output from the MCMC. I simulated each model with each of the 10 000 parameters samples from the MCMC output. I calculated the R^2 for each of the 10 000 simulated PSTHs (PSTVs) against the 1000 bootstrapped data PSTHs (PSTVs), resulting 1000 000 R^2 values per model (Fig. A.1). The 95% credible interval was the 2.5 and 97.5 percentiles of the R^2 values.

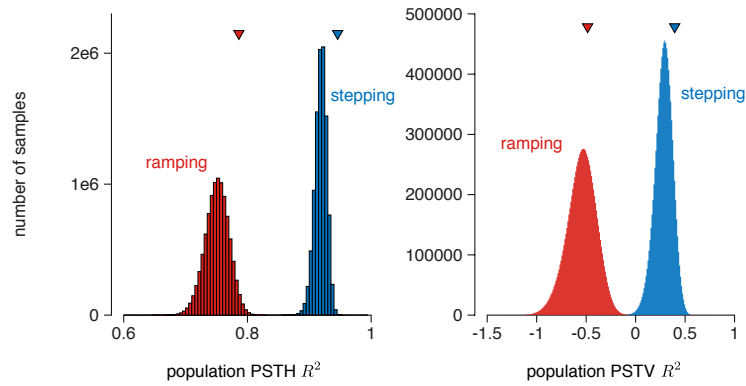


Figure A.1: Distribution of sampled R^2 values of the model predictions of the coherence-sorted PSTH (left) and PSTV (right). The distributions compare the ramping model PSTH/PSTV predictions (red) to the predictions from the stepping model (blue). Triangles indicate the R^2 values calculated for the PSTHs or PSTVs plotted in Fig. 4A, which were computed using the posterior mean parameters.

A.2 Parameter estimates

A.2.1 Ramping model

Cell	β_{-high}	β_{-low}	β_{zero}	β_{+low}	β_{+high}	x_0	ω^2	γ	DIC difference
1	-5.59e-03	-1.43e-03	1.02e-02	1.31e-02	3.45e-02	0.44	5.72e-03	37.0	141.8
2	-2.55e-02	-1.05e-02	-5.34e-03	8.48e-05	1.37e-02	0.36	1.09e-02	19.2	41.8
3	-1.47e-03	-3.21e-04	3.01e-05	3.28e-03	7.42e-03	0.40	3.05e-03	19.8	4.7
4*	-4.42e-03	-3.71e-03	-3.77e-03	-2.35e-03	3.94e-04	0.45	1.93e-03	138.6	-151.6
5	-8.72e-03	-6.58e-04	5.76e-04	9.20e-04	6.00e-03	0.38	3.80e-03	9.4	10.9
6	-8.47e-03	-7.17e-03	-3.56e-03	-2.16e-03	1.19e-01	0.76	1.56e-03	23.4	44.0
7*	-1.52e-02	-5.89e-04	4.69e-03	5.87e-03	1.21e-02	-0.05	1.01e-02	35.9	-164.4
8	-2.11e-03	1.73e-03	4.04e-03	6.12e-03	1.27e-02	0.19	6.75e-03	33.0	66.8
9	-3.45e-03	-1.97e-03	-7.70e-04	-9.58e-05	4.75e-04	0.36	1.71e-03	158.1	128.9
10*	-5.08e-03	-2.55e-03	1.36e-03	2.66e-03	4.82e-03	0.32	4.92e-03	62.0	-54.4
11	-1.30e-02	-6.06e-03	7.19e-03	1.49e-03	7.89e-02	0.93	5.18e-03	10.8	22.5
12	2.37e-03	1.67e-03	2.94e-03	5.53e-03	1.45e-02	0.29	2.87e-03	7.9	53.5
13	5.99e-03	4.03e-03	6.12e-03	7.78e-03	1.32e-02	0.30	3.16e-03	21.7	10.5
14	-2.79e-03	5.44e-04	6.52e-03	9.12e-03	2.21e-01	0.27	7.17e-03	11.7	70.6
15	-3.26e-03	-1.15e-03	-1.68e-03	-1.52e-03	-1.47e-03	0.38	2.67e-03	63.9	60.7
16	-2.02e-03	8.35e-04	1.17e-03	1.93e-03	4.38e-03	0.56	2.46e-03	23.4	28.4
17	3.01e-03	3.82e-03	4.97e-03	8.25e-03	1.03e-02	0.35	3.44e-03	70.4	96.0
18	1.08e-01	1.13e-01	1.09e-01	1.03e-01	1.17e-01	0.97	9.27e-03	39.5	27.5
19	-3.50e-03	-2.55e-03	5.31e-04	2.63e-03	3.51e-02	0.48	4.43e-03	14.5	22.1
20	-4.56e-03	-2.45e-03	1.51e-02	3.14e-03	2.73e-02	0.86	3.14e-03	26.0	0.5
21	2.77e-02	1.10e-01	9.63e-02	1.15e-01	1.10e-01	0.98	4.29e-03	29.5	170.5
22	-1.32e-05	2.67e-03	2.54e-03	4.09e-03	5.81e-03	0.29	4.31e-03	25.9	1.1
23*	-1.91e-03	-3.18e-04	1.36e-03	3.07e-03	5.32e-03	0.39	1.49e-03	46.4	-19.9
24	-1.07e-02	-6.20e-03	-3.44e-03	-5.04e-03	-3.94e-03	0.59	3.60e-03	25.5	38.4
25*	-2.36e-04	-6.52e-04	9.40e-04	3.17e-03	7.00e-03	0.40	3.45e-03	52.7	-211.3
26	2.43e-04	7.32e-04	3.59e-03	4.24e-03	4.86e-03	0.23	3.44e-03	39.8	19.6
27	3.77e-03	1.41e-02	1.72e-02	3.19e-02	2.67e-02	0.71	1.43e-02	20.5	9.9
28	-4.65e-03	-2.38e-03	-1.25e-03	5.69e-04	3.42e-03	0.72	1.69e-03	39.7	63.0
29	1.28e-03	2.89e-03	6.13e-03	9.51e-03	1.21e-02	0.09	4.00e-03	13.6	8.9
30	-8.49e-03	-2.40e-03	-1.66e-03	1.24e-03	6.76e-03	0.71	3.98e-03	26.5	57.3
31	-4.27e-03	2.47e-05	-1.11e-04	2.48e-03	4.74e-03	0.58	1.57e-03	32.2	76.4
32*	-4.55e-03	-2.68e-03	4.64e-03	8.97e-03	1.63e-02	0.03	9.99e-03	50.4	-227.4
33	-6.09e-03	-1.67e-03	-2.64e-03	-1.71e-05	1.04e-05	0.36	3.45e-03	26.6	80.7
34*	-1.23e-02	-2.32e-03	2.89e-04	2.51e-03	7.81e-04	0.55	3.56e-03	23.0	-9.9
35	1.57e-02	1.91e-02	2.12e-02	1.78e-02	3.76e-02	0.01	1.93e-02	11.0	24.6
36	1.89e-03	4.17e-03	7.94e-03	7.48e-03	1.92e-02	0.27	1.93e-03	28.8	4.7
37*	-2.44e-02	-9.12e-03	-3.99e-03	-3.45e-03	5.35e-03	0.65	7.49e-03	42.9	-8.9
38	-7.20e-03	-4.51e-03	4.67e-03	8.67e-03	2.11e-02	0.42	1.21e-02	11.6	22.0
39*	-4.72e-03	-2.48e-03	-1.43e-04	3.28e-03	7.49e-03	0.29	3.29e-03	18.1	-0.3
40	-1.76e-03	-1.13e-03	2.03e-02	1.11e-02	7.41e-02	0.81	2.40e-03	18.2	15.9

Table A.1: Posterior mean ramping model parameters for all cells, and the DIC differences from the model comparison (positive indicating support for the stepping model). Stars next to the cell number indicate those cells we identified as tentative rampers (negative DIC difference of any magnitude).

A.2.2 Stepping model

Cell	α_{init}	α_{out}	α_{in}	p_{-high}	p_{-low}	p_{zero}	p_{+low}	p_{+high}	ϕ_{-high}	ϕ_{-low}	ϕ_{zero}	ϕ_{+low}	ϕ_{+high}	r
1*	16.8	4.1	36.3	0.972	0.975	0.977	0.968	0.951	0.02	0.35	0.82	0.88	0.98	0.55
2*	7.7	0.7	19.9	0.815	0.957	0.977	0.982	0.987	0.01	0.20	0.36	0.48	0.94	0.79
3*	8.4	3.5	20.8	0.960	0.969	0.954	0.971	0.974	0.15	0.23	0.25	0.60	0.89	1.63
4	70.9	11.1	33.2	0.944	0.968	0.972	0.981	0.991	0.46	0.51	0.53	0.59	0.80	0.94
5*	3.4	0.8	9.5	0.953	0.979	0.944	0.980	0.957	0.03	0.39	0.50	0.52	0.78	1.42
6*	18.4	6.6	23.7	0.801	0.902	0.871	0.912	0.276	0.04	0.09	0.42	0.51	0.97	3.02
7	1.7	0.4	24.1	0.852	0.947	0.958	0.971	0.964	0.02	0.48	0.74	0.93	0.97	1.71
8*	1.3	10.9	29.2	0.994	0.990	0.990	0.987	0.970	0.07	0.26	0.57	0.59	0.62	0.38
9*	74.2	11.0	39.7	0.902	0.984	0.986	0.991	0.993	0.16	0.43	0.58	0.69	0.93	0.64
10	21.0	3.6	49.5	0.975	0.984	0.982	0.983	0.992	0.08	0.25	0.51	0.62	0.86	0.68
11*	11.8	3.3	11.0	0.829	0.737	0.406	0.374	0.360	0.37	0.72	0.84	0.84	0.92	2.13
12*	2.5	3.0	8.6	0.887	0.724	0.829	0.451	0.800	0.44	0.26	0.40	0.26	0.76	4.57
13*	5.1	12.0	21.9	0.986	0.985	0.975	0.972	0.924	0.58	0.29	0.37	0.44	0.72	0.86
14*	49.4	2.6	11.0	0.401	0.524	0.981	0.444	0.489	0.03	0.18	0.38	0.44	0.88	0.04
15*	38.9	5.4	16.9	0.972	0.989	0.990	0.985	0.977	0.45	0.67	0.54	0.67	0.80	0.38
16*	14.5	8.2	20.8	0.700	0.421	0.536	0.526	0.346	0.23	0.45	0.46	0.45	0.65	1.64
17*	11.6	31.5	60.9	0.978	0.957	0.951	0.964	0.957	0.34	0.31	0.38	0.69	0.72	0.54
18*	43.6	37.2	42.7	0.467	0.445	0.433	0.461	0.462	0.17	0.12	0.38	0.49	0.90	1.80
19*	11.2	3.6	13.3	0.545	0.695	0.762	0.559	0.436	0.15	0.20	0.34	0.48	0.81	1.87
20*	21.2	10.3	26.1	0.900	0.545	0.507	0.387	0.850	0.32	0.70	0.74	0.79	0.89	2.11
21*	35.6	24.3	27.4	0.361	0.757	0.779	0.878	0.908	0.17	0.39	0.69	0.64	0.64	5.42
22*	3.5	8.5	20.4	0.970	0.964	0.965	0.882	0.957	0.41	0.44	0.44	0.37	0.65	0.91
23	19.6	8.6	41.7	0.945	0.962	0.966	0.975	0.967	0.04	0.23	0.44	0.72	0.94	2.01
24*	21.3	1.7	10.4	0.965	0.967	0.953	0.962	0.982	0.15	0.48	0.75	0.64	0.58	0.73
25	8.6	23.6	44.9	0.996	0.993	0.991	0.988	0.985	0.23	0.14	0.37	0.43	0.63	0.27
26*	4.8	15.3	33.8	0.994	0.991	0.983	0.976	0.980	0.10	0.20	0.39	0.34	0.44	0.57
27*	3.1	13.8	21.7	0.994	0.963	0.972	0.727	0.967	0.38	0.68	0.75	0.69	0.83	0.12
28*	32.9	17.6	38.3	0.766	0.876	0.934	0.542	0.323	0.04	0.26	0.26	0.60	0.65	1.53
29*	1.6	7.1	14.2	0.986	0.985	0.974	0.969	0.953	0.09	0.40	0.40	0.61	0.36	1.44
30*	19.1	5.3	25.3	0.941	0.848	0.947	0.677	0.946	0.05	0.60	0.54	0.75	0.87	1.19
31*	19.9	10.9	29.4	0.871	0.859	0.882	0.449	0.850	0.06	0.41	0.40	0.59	0.80	0.84
32	4.1	0.6	41.0	0.988	0.983	0.982	0.975	0.972	0.10	0.30	0.71	0.82	0.98	1.05
33*	24.1	2.5	9.7	0.940	0.987	0.992	0.989	0.971	0.24	0.43	0.30	0.62	0.64	0.17
34	12.6	1.5	23.7	0.935	0.971	0.980	0.973	0.987	0.03	0.37	0.53	0.70	0.55	1.56
35*	2.5	7.2	18.8	0.910	0.925	0.862	0.923	0.739	0.27	0.57	0.31	0.58	0.31	4.21
36*	7.5	10.4	26.9	0.591	0.829	0.730	0.870	0.785	0.10	0.25	0.47	0.64	0.86	4.32
37	26.7	1.4	44.1	0.955	0.985	0.984	0.986	0.985	0.05	0.21	0.54	0.53	0.81	0.70
38*	5.8	1.7	10.8	0.646	0.538	0.569	0.325	0.882	0.14	0.33	0.59	0.71	0.96	3.39
39	5.8	1.1	16.8	0.979	0.987	0.987	0.990	0.983	0.10	0.16	0.33	0.62	0.82	0.74
40*	23.5	11.3	19.3	0.689	0.531	0.497	0.252	0.399	0.24	0.18	0.53	0.70	0.88	2.37

Table A.2: Posterior mean stepping model parameters for all cells. Stars next to the cell number indicate those cells we identified as tentative steppers (positive DIC difference of any magnitude).

A.3 Example LIP cells

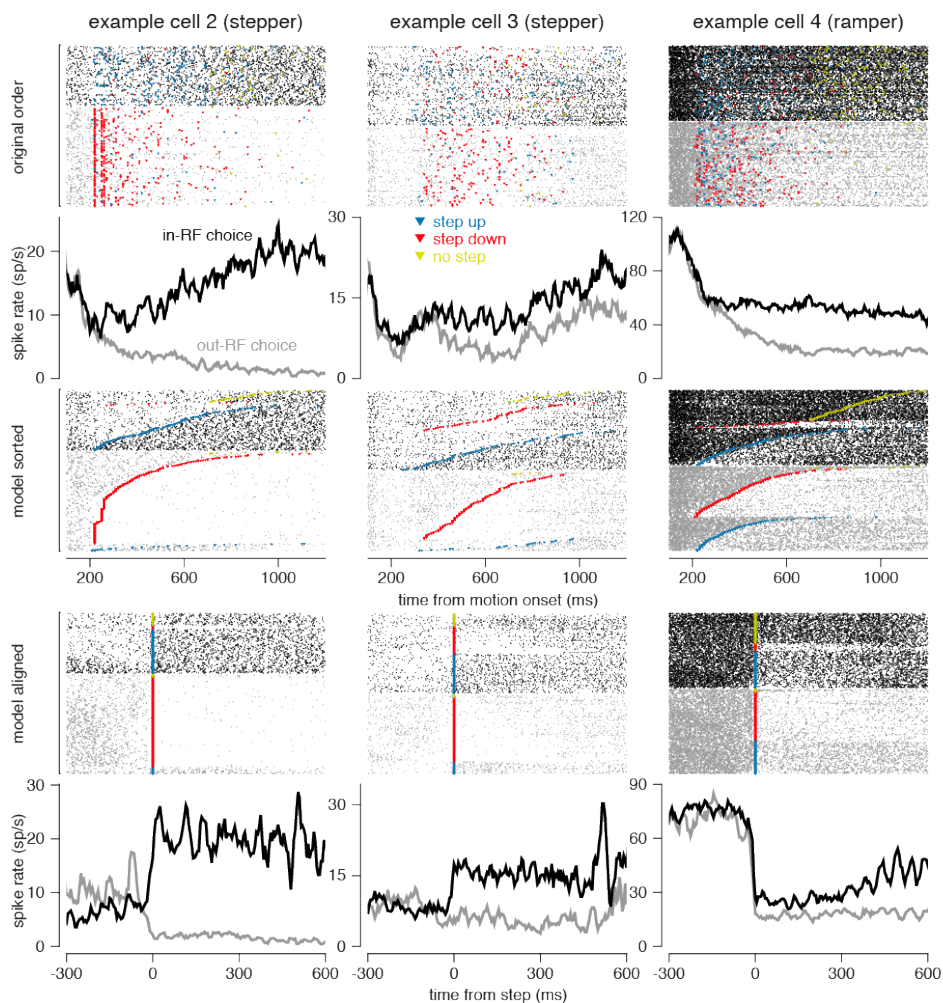


Figure A.2: Each column shows the responses of an example LIP cell with the same step-model analysis performed on the cell in Figure 2.17. The top row shows all the trials aligned to stimulus onset, sorted by choice, and ordered in the order the trials were collected. The 2nd row shows the average firing rate aligned to motion onset. The 3rd row is the same data as in row 1, although the trials have been ordered by step time. In the 4th row, the trials have been aligned to the step time. The spike rate aligned to step time is given in the bottom row. Cells that were fit better by the ramping model are labeled as “rampers” and cells better fit by the stepping model are labeled “steppers”.

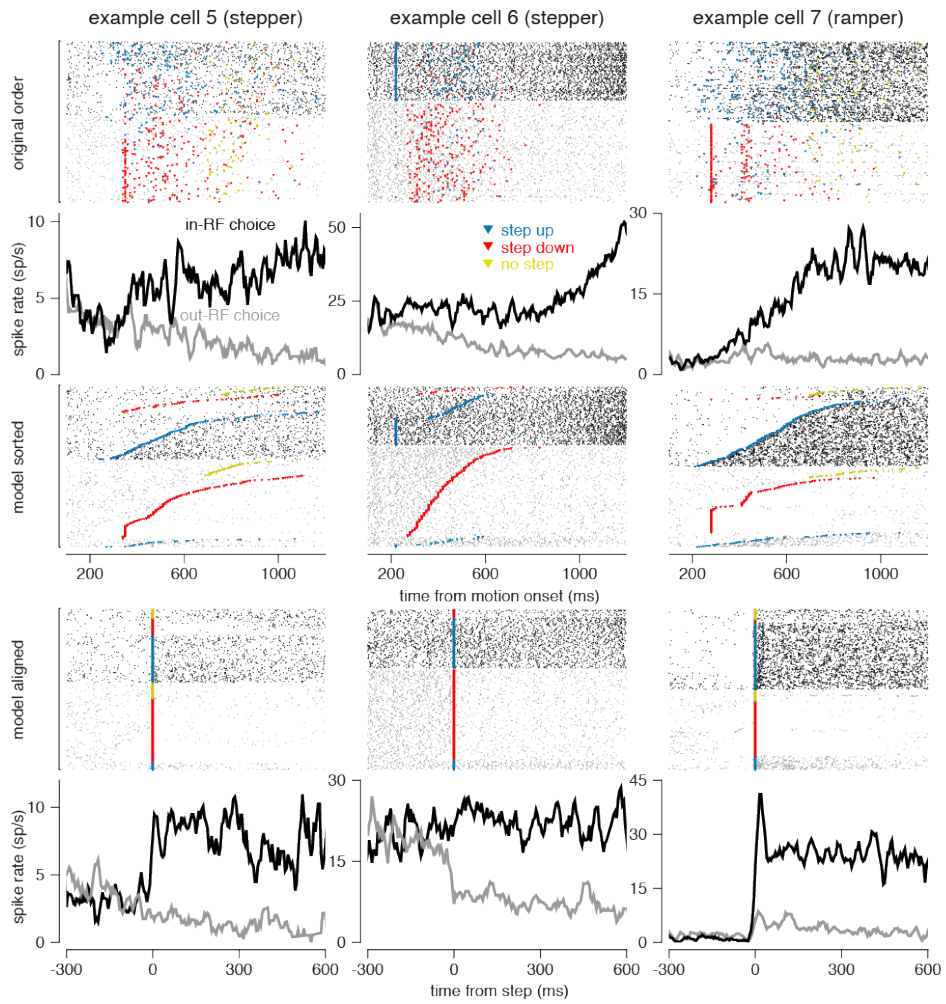


Figure A.3: Same as Figure A.2 for 3 more cells.

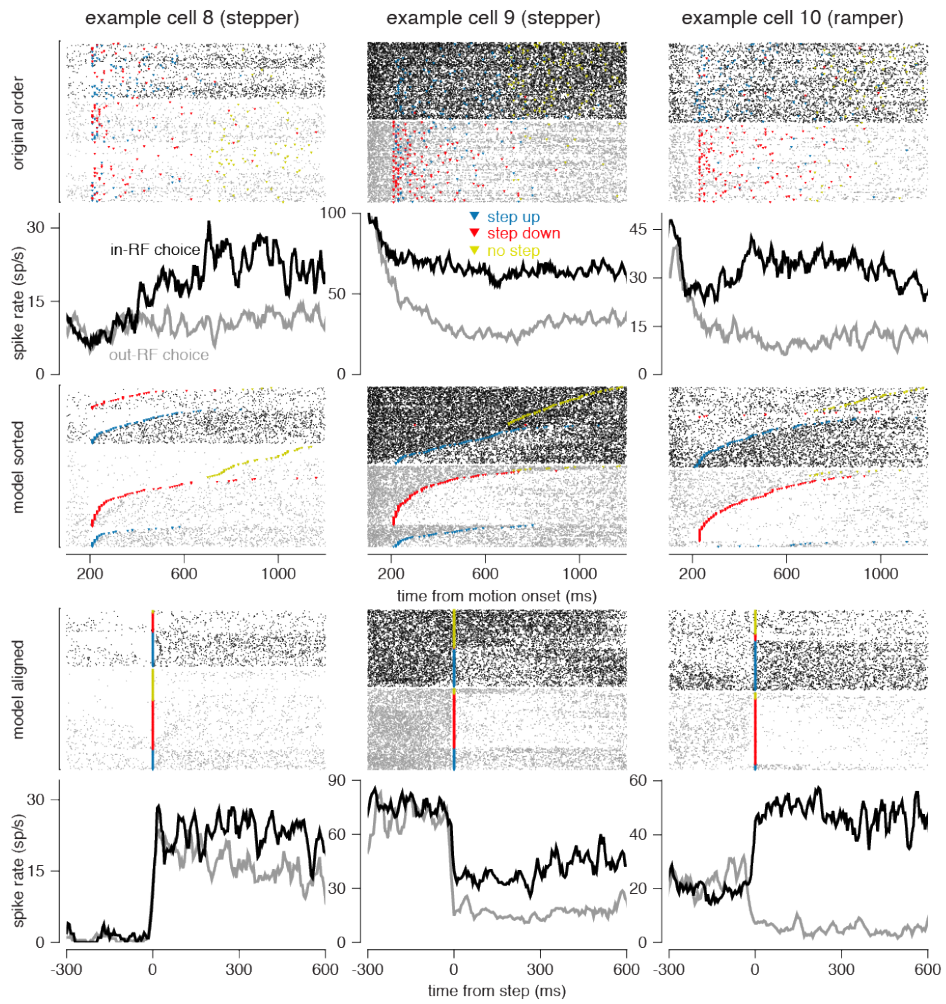


Figure A.4: Same as Figure A.2 for 3 more cells.

A.4 Cell examples over different start times of analysis

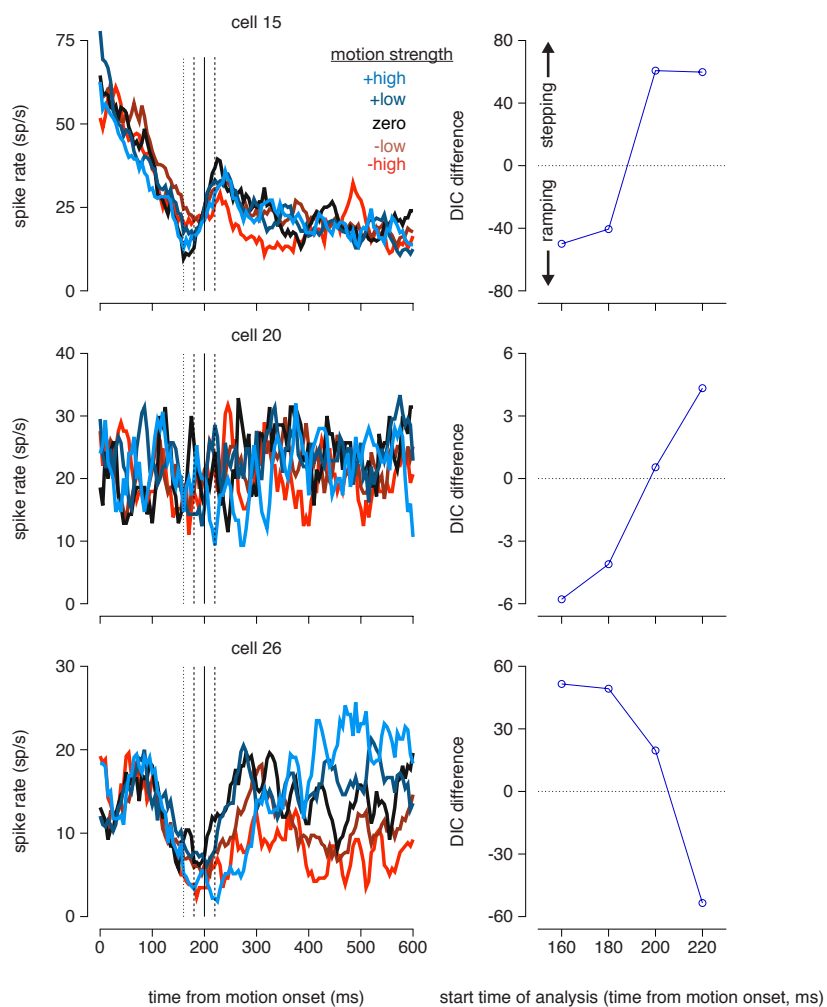


Figure A.5: The three rows show the coherence-sorted PSTH (left) of cells identified as steppers in the original analysis that were better fit by the ramping model at different start times of analysis. (right) The model comparison metric for each cell is given for all start times. The PSTHs were estimated using only a 25 ms sliding window, which makes the firing rate estimate very noisy for single cells.

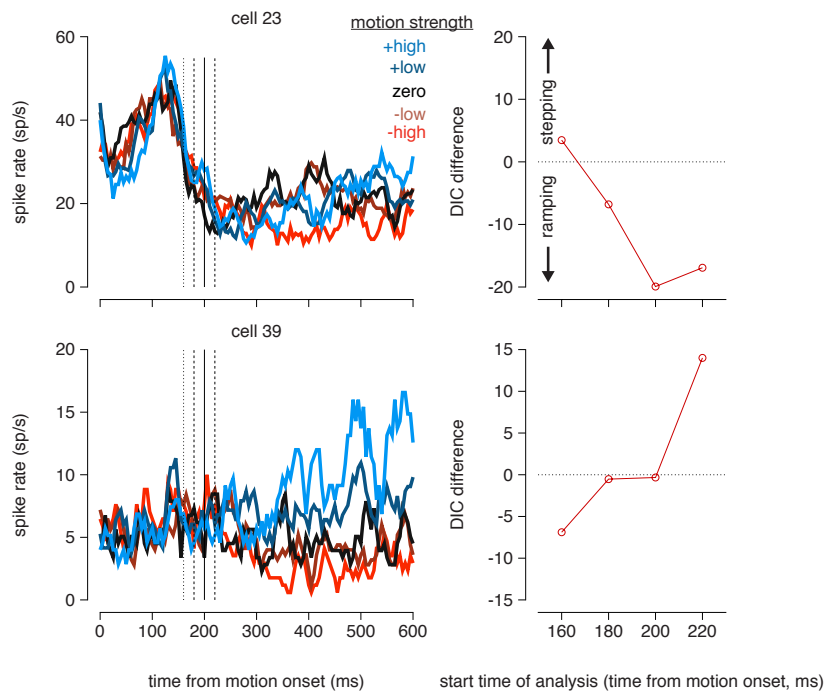


Figure A.6: The two rows show the coherence-sorted PSTH (left) of cells identified as rampers in the original analysis that were better fit bit the stepping model at different start times of analysis. (right) The model comparison metric for each cell is given for all start times.

Appendix B

Gradients of CBSM log likelihood

B.1 Notation for the log likelihood

The design matrix is denoted \mathbf{X} where the K columns of \mathbf{X} consist of the original stimulus convolved with the (temporal) basis vectors. The t th row of \mathbf{X} , denoted \mathbf{X}_t is the value of the stimulus at time t and $(\mathbf{k}_e * \mathbf{x})(t) = \mathbf{X}_t \mathbf{k}_e$. We assume that one of the columns of \mathbf{X} consists of all ones, so that the stimulus filters, \mathbf{k} , implicitly include the offset terms b_e and b_i . Similarly, the spike history convolved with the appropriate basis vectors is given by the matrix \mathbf{H} with J columns. The vector \mathbf{V} denotes the voltage at all discrete time bins and $V(t)$ is the voltage at time t .

The CBSM defines the membrane voltage according to the equations

$$g_e(t) = \log(1 + \exp(\mathbf{X}_t \mathbf{k}_e)) = f(\mathbf{X}_t \mathbf{x}_e) \quad (\text{B.1})$$

$$g_i(t) = \log(1 + \exp(\mathbf{X}_t \mathbf{k}_i)) = f(\mathbf{X}_t \mathbf{k}_i) \quad (\text{B.2})$$

$$\frac{dV(t)}{dt} = g_e(t)(E_e - V(t)) + g_i(t)(E_i - V(t)) + g_l(E_l - V(t)) \quad (\text{B.3})$$

$$= (g_e(t)E_e + g_i(t)E_i + g_l E_l) - V(t)(g_e(t) + g_i(t) + g_l) \quad (\text{B.4})$$

The voltage equation is solved along the bins $t = 1 \dots T$ of width Δ

$$c(t) = g_e(t) + g_i(t) + g_l \quad (\text{B.5})$$

$$b(t) = g_e(t)E_e + g_i(t)E_i + g_l E_l \quad (\text{B.6})$$

$$V(t+1) = \exp(-c(t)\Delta) \left(V(t) - \frac{b(t)}{c(t)} \right) + \frac{b(t)}{c(t)} \quad (\text{B.7})$$

$$V(t+1) - \exp(-c(t)\Delta)V(t) = \frac{b(t)}{c(t)}(1 - \exp(-c(t)\Delta)) \quad (\text{B.8})$$

$$\exp(c(t)\Delta)V(t+1) - V(t) = \frac{b(t)}{c(t)}(\exp(c(t)\Delta) - 1) \quad (\text{B.9})$$

$$M = \begin{bmatrix} \exp(c(1)\Delta) & 0 & \dots & 0 \\ -1 & \exp(c(2)\Delta) & 0 & \\ 0 & -1 & \exp(c(3)\Delta) & 0 \dots \\ & & \ddots & \\ 0 & & 0 & -1 \exp(c(T)\Delta) \end{bmatrix} \quad (\text{B.10})$$

$$\mathbf{d} = \begin{bmatrix} V_{\text{init}} + \frac{b(1)}{c(1)}(\exp(c(1)\Delta) - 1) \\ \frac{b(2)}{c(2)}(\exp(c(2)\Delta) - 1) \\ \vdots \\ \frac{b(T)}{c(T)}(\exp(c(T)\Delta) - 1) \end{bmatrix} \quad (\text{B.11})$$

The voltage at each time point is then

$$\mathbf{V} = M^{-1}\mathbf{d}. \quad (\text{B.12})$$

The firing rate at each time is the vector

$$\mathbf{r} = \begin{bmatrix} \alpha r(wV(1) + \mathbf{H}_t \mathbf{h}_{spk} + z) \\ \vdots \\ \alpha r(wV(T) + s(T) + z) \end{bmatrix}. \quad (\text{B.13})$$

The log-likelihood becomes

$$\mathcal{L}(\Theta) = \sum_{t=1}^T [y(t) \log(1 - \exp(-r(t)\Delta)) - \Delta r(t)(1 - y(t))] \quad (\text{B.14})$$

B.2 Gradients for conductance filters

Here we show the derivative of the log likelihood with respect to \mathbf{k}_e . The derivative for the inhibitory filter is computed analogously. The Hadamard product (element-wise multiplication) is denoted by \circ .

$$\frac{d}{d\Theta} \mathcal{L}(\Theta) = \frac{d}{d\Theta} \sum_{t=1}^T [y(t) \log(1 - \exp(-r(t)\Delta)) - \Delta r(t)(1 - y(t))] \quad (\text{B.15})$$

$$= \sum_{t=1}^T \left[y(t) \left(1 + \frac{1}{\exp(r(t)\Delta) - 1} \right) - 1 \right] \Delta \frac{dr(t)}{d\Theta} \quad (\text{B.16})$$

$$\frac{dr(t)}{d\mathbf{k}_e} = \frac{d}{d\mathbf{k}_e} \alpha \log(1 + \exp(wV(t) + s(t) + z)) \quad (\text{B.17})$$

$$= w\alpha \frac{1}{1 + \exp(-(wV(t) + s(t) + z))} \frac{dV(t)}{d\mathbf{k}_e} \quad (\text{B.18})$$

The derivative of the voltage is given

$$\begin{aligned} \mathbf{V}_e &= -M^{-1}M'M^{-1}\mathbf{d} + M^{-1}\mathbf{r}_e \\ &= -M^{-1}(M'\mathbf{V} - v\mathbf{d}_e) \\ &= -M^{-1}(\gamma_e) \end{aligned} \quad (\text{B.19})$$

$$\gamma_e = (\mathbf{m}_e \circ \mathbf{V} - \mathbf{d}_e) \quad (\text{B.20})$$

$$\mathbf{d}_e = \begin{bmatrix} \frac{b(1)}{c(1)}m_e(1) + p_e(1)(s(1) - 1) \\ \frac{b(2)}{c(2)}m_e(2) + p_e(2)(s(2) - 1) \\ \vdots \\ \frac{b(T)}{c(T)}m_e(T) + p_e(T)(s(T) - 1) \end{bmatrix} \quad (\text{B.21})$$

$$\mathbf{s} = \begin{bmatrix} \exp(c(1)\Delta) \\ \exp(c(2)\Delta) \\ \vdots \\ \exp(c(T)\Delta) \end{bmatrix} \quad (\text{B.22})$$

$$\mathbf{m}_e = \Delta \mathbf{c}_e \circ \mathbf{s} \quad (\text{B.23})$$

$$p_e(t) = \frac{b_e(t)c(t) - c_e(t)b(t)}{c(t)^2} \quad (\text{B.24})$$

$$c_e(t) = f'(\mathbf{X}_t \mathbf{k}_e) \quad (\text{B.25})$$

$$b_e(t) = E_e f'(\mathbf{X}_t \mathbf{k}_e) = E_e c_e(t) \quad (\text{B.26})$$

$$f'(x) = \frac{1}{1 + \exp(-x)} \quad (\text{B.27})$$

The derivative is then simply

$$\frac{d}{d\mathbf{k}_e} \mathcal{L}(\Theta) = \Delta w \alpha \mathbf{X}^\top (\mathbf{q} \circ \mathbf{a} \circ \mathbf{V}_e) \quad (\text{B.28})$$

$$\mathbf{q} = \begin{bmatrix} y(1) \left(\frac{1}{\exp(r(1)\Delta) - 1} + 1 \right) - 1 \\ y(2) \left(\frac{1}{\exp(h(2)\Delta) - 1} + 1 \right) - 1 \\ \vdots \\ y(T) \left(\frac{1}{\exp(h(T)\Delta) - 1} + 1 \right) - 1 \end{bmatrix} \quad (\text{B.29})$$

$$\mathbf{a} = \begin{bmatrix} \frac{1}{1 + \exp(-(wV(1) + s(1) + z))} \\ \vdots \\ \frac{1}{1 + \exp(-(wV(T) + s(T) + z))} \end{bmatrix} \quad (\text{B.30})$$

B.2.1 Second derivative for conductance filters

Here we show the second derivative of the log likelihood with respect to \mathbf{k}_e . We denote a column vector containing K ones as $\mathbf{1}_K$.

$$\begin{aligned} \frac{d^2}{d\mathbf{k}_e^2} \mathcal{L}(\Theta) = & w\alpha\Delta [([\mathbf{q} \odot (\mathbf{a} \odot \mathbf{V}_{ee} + \alpha\mathbf{a}_e \odot \mathbf{V}_e \odot \mathbf{V}_e) \\ & + w\alpha\mathbf{q}_e \odot (\mathbf{a} \odot \mathbf{V}_e) \odot (\mathbf{a} \odot \mathbf{V}_e)] \mathbf{1}_K^\top) \odot \mathbf{X}]^\top \mathbf{X} \end{aligned} \quad (\text{B.31})$$

$$\mathbf{q}_e = \begin{bmatrix} y(1) \frac{-\Delta}{\exp(r(1)\Delta) - 2 + \exp(-r(1)\Delta)} \\ y(2) \frac{-\Delta}{\exp(r(2)\Delta) - 2 + \exp(-r(2)\Delta)} \\ \vdots \\ y(T) \frac{-\Delta}{\exp(r(T)\Delta) - 2 + \exp(-r(T)\Delta)} \end{bmatrix} \quad (\text{B.32})$$

$$\mathbf{a}_e = \begin{bmatrix} \frac{1}{2 + \exp(-(wV(1) + s(1) + z)) + \exp(\alpha V(1) + s(1) + z)} \\ \vdots \\ \frac{1}{2 + \exp(-(wV(T) + s(T) + z)) + \exp(\alpha V(T) + s(T) + z)} \end{bmatrix} \quad (\text{B.33})$$

$$\mathbf{V}_{ee} = M^{-1} M' M^{-1} \gamma_e - M^{-1} \gamma_{ee} \quad (\text{B.34})$$

$$= M^{-1} (\mathbf{m}_e \odot (M^{-1} \gamma_e) - \gamma_{ee}) \quad (\text{B.35})$$

$$= M^{-1} (-\mathbf{m}_e \odot V_e - \gamma_{ee}) \quad (\text{B.36})$$

$$\gamma_{ee} = (\mathbf{m}_{ee} \odot V + m_e \odot \mathbf{V}_e - \mathbf{d}_{ee}) \quad (\text{B.37})$$

$$\mathbf{m}_{ee} = \Delta^2(\mathbf{c}_{ee} \circ \mathbf{s} + \mathbf{c}_e \circ \mathbf{c}_e \odot \mathbf{s}) \quad (\text{B.38})$$

$$d_{ee}(t) = p_e(t)m_e(t) + \frac{b(t)}{c(t)}m_{ee}(t) + p_{ee}(t)(s(t) - 1) + p_e(t)m_e(t) \quad (\text{B.39})$$

$$p_{ee}(t) = \frac{c(t)^2 (b_e e(t)c(t) - c_e e(t)b(t)) - 2c(t)c_e(t) (b_e(t)c(t) - c_e(t)b(t))}{c(t)^4} \quad (\text{B.40})$$

$$c_{ee}(t) = f''(\mathbf{X}_t \mathbf{k}_e) \quad (\text{B.41})$$

$$b_{ee}(t) = E_e c_e e t = E_e f''(\mathbf{X}_t \mathbf{k}_e) \quad (\text{B.42})$$

$$f''(x) = \frac{\exp(-x)}{(1 + \exp(-x))^2} \quad (\text{B.43})$$

$$\frac{d^2}{d^2\Theta} \mathcal{L}(\Theta) = \frac{d}{d\Theta} \sum_{t=1}^T \left[y(t) \left(1 + \frac{1}{\exp(r(t)\Delta) - 1} \right) - 1 \right] \Delta \frac{dr(t)}{d\Theta} \quad (\text{B.44})$$

$$\begin{aligned} &= \Delta \sum_{t=1}^T \left[y(t) \left(1 + \frac{1}{\exp(r(t)\Delta) - 1} \right) - 1 \right] \frac{d^2 r(t)}{d\Theta^2} \\ &\quad + \left[y(t) \frac{-\Delta}{\exp(r(t)\Delta) - 2 + \exp(-r(t)\Delta)} \right] \left(\frac{dr(t)}{d\Theta} \right)^2 \end{aligned} \quad (\text{B.45})$$

$$\frac{d^2 r(t)}{d\mathbf{k}_e^2} = \frac{dr(t)}{d\mathbf{k}_e} \left(w\alpha \frac{1}{1 + \exp(-(wV(t) + s(t) + z))} \frac{dV(t)}{d\mathbf{k}_e} \right) \quad (\text{B.46})$$

$$= w\alpha \left[\frac{d^2 V_t}{dk_e^2} a(t) + w a_e(t) \left(\frac{dV(t)}{d\mathbf{k}_e} \right)^2 \right] \quad (\text{B.47})$$

B.2.2 Mixed second derivative for conductance filters

Here we show the second derivative of the log likelihood with respect to \mathbf{k}_e and \mathbf{k}_i .

$$\begin{aligned} \frac{d^2}{d\mathbf{k}_i d\mathbf{k}_e} \mathcal{L}(\Theta) &= \Delta \sum_{t=1}^T \left[y(t) \left(1 + \frac{1}{\exp(r(t)\Delta) - 1} \right) - 1 \right] \frac{dr(t)}{d\mathbf{k}_i d\mathbf{k}_e} \\ &\quad + \left[y(t) \frac{-\Delta}{\exp(r(t)\Delta) - 2 + \exp(-r(t)\Delta)} \right] \frac{dr(t)}{d\mathbf{k}_e} \frac{dr(t)}{d\mathbf{k}_i} \end{aligned} \quad (\text{B.48})$$

$$\begin{aligned} &= \Delta w \alpha \left[([\mathbf{q} \circ \mathbf{a} \circ \mathbf{V}_{ei} \right. \\ &\quad \left. + w(\mathbf{q} \odot \mathbf{a}_e + w\mathbf{q}_e \odot \mathbf{a} \circ \mathbf{a}) \circ \mathbf{V}_e \circ \mathbf{V}_i] \mathbf{1}_K^\top) \circ \mathbf{X} \right]^\top \mathbf{X} \end{aligned} \quad (\text{B.49})$$

$$\mathbf{V}_{ei} = -M^{-1}(\gamma_{ei}) - M^{-1}\mathbf{m}_i \circ \mathbf{V}_e \quad (\text{B.50})$$

$$\gamma_{ei} = (\mathbf{m}_{ei} \circ \mathbf{V} + \mathbf{m}_e \circ \mathbf{V}_i - \mathbf{d}_{ei}) \quad (\text{B.51})$$

$$\mathbf{m}_{ei} = \Delta^2(\mathbf{c}_e \circ \mathbf{c}_i \circ \mathbf{s}) \quad (\text{B.52})$$

$$(\text{B.53})$$

$$d_{ei}(t) = \frac{b(t)}{c(t)} m_{ei}(t) + p_i(t) m_e(t) + p_e(t) m_i(t) + p_{ei}(t)(s(t) - 1) \quad (\text{B.54})$$

$$\begin{aligned} p_{ei}(t) &= \frac{c^2(t)(b_e(t)c_i(t) - b_i(t)c_e(t)) + 2c(t)c_i(t)(b_e(t)c(t) - b(t)c_e(t))}{c^4(t)} \\ &= \frac{c_e(t)c_i(t)(-c(t) * (E_e + E_i) + 2b(t))}{c(t)^3} \end{aligned} \quad (\text{B.55})$$

B.3 Gradients for the spike history filter

The derivative of the log likelihood with respect to the spike history filter is

$$\frac{d}{d\mathbf{h}_{spk}} \mathcal{L}(\Theta) = \Delta \alpha \mathbf{H}^\top (\mathbf{q} \circ \mathbf{a}). \quad (\text{B.56})$$

B.3.1 Second derivative for the spike history filter

The second derivative of the log likelihood with respect to the spike history filter is

$$\frac{d^2}{d\mathbf{h}_{spk}^2} \mathcal{L}(\Theta) = \alpha \Delta \left[\left([\mathbf{q} \circ \mathbf{a}_e + w \mathbf{q}_e \circ \mathbf{a} \circ \mathbf{a}] \mathbf{1}_J^\top \right) \circ \mathbf{H} \right]^\top \mathbf{H}. \quad (\text{B.57})$$

B.3.2 Mixed second derivative for the spike history and conductance filters

Here we show the second derivative of the log likelihood with respect to \mathbf{k}_e and \mathbf{h}_{spk} . The mixed second derivative with respect to \mathbf{k}_i and \mathbf{h}_{spk} is analogous.

$$\frac{d^2}{d\mathbf{h}_{spk} d\mathbf{k}_e} \mathcal{L}(\Theta) = w \alpha \Delta \left[\left([\mathbf{q} \circ d\mathbf{a}_e \circ \mathbf{V}_e + w \mathbf{q}_e \circ \mathbf{a} \circ \mathbf{a} \circ \mathbf{V}_e] \mathbf{1}_K^\top \right) \odot \mathbf{X} \right]^\top \mathbf{H} \quad (\text{B.58})$$

Bibliography

- Agüera y Areas, B., Fairhall, A. L. & Bialek, W. (2001), What can a single neuron compute?, *in* ‘Advances in Neural Information Processing Systems 13: Proceedings of the 2000 Conference’, Vol. 13, MIT Press, p. 75.
- Ahrens, M. B., Paninski, L. & Sahani, M. (2008), ‘Inferring input nonlinearities in neural encoding models.’, *Network: Computation in Neural Systems* **19**(1), 35–67.
- Baccus, S. A. & Meister, M. (2002), ‘Fast and slow contrast adaptation in retinal circuitry’, *Neuron* **36**(5), 909–919.
- Badel, L., Lefort, S., Berger, T. K., Petersen, C. C., Gerstner, W. & Richardson, M. J. (2008), ‘Extracting non-linear integrate-and-fire models from experimental data using dynamic i–v curves’, *Biological cybernetics* **99**(4–5), 361–370.
- Barreiro, A. K., Gjorgjieva, J., Rieke, F. & Shea-Brown, E. (2014), ‘When do microcircuits produce beyond-pairwise correlations?’, *Frontiers in computational neuroscience* **8**.
- Berg, A., Meyer, R. & Yu, J. (2004), ‘Deviance information criterion for comparing stochastic volatility models’, *Journal of Business & Economic Statistics* **22**(1), 107–120.

- Bollimunta, A., Totten, D. & Ditterich, J. (2012), ‘Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex.’, *The Journal of Neuroscience* **32**(37), 12684–701.
- Brenner, N., Strong, S. P., Koberle, R., Bialek, W. & Steveninck, R. R. (2000), ‘Synergy in a neural code’, *Neural Computation* **12**(7), 1531–1552.
- Bromberg-Martin, E. S., Matsumoto, M. & Hikosaka, O. (2010), ‘Distinct tonic and phasic anticipatory activity in lateral habenula and dopamine neurons’, *Neuron* **67**(1), 144–155.
- Brunton, B. W., Botvinick, M. M. & Brody, C. D. (2013), ‘Rats and humans can optimally accumulate evidence for decision-making.’, *Science* **340**(6128), 95–8.
- Buesing, L., Macke, J. H. & Sahani, M. (2012), ‘Learning stable, regularised latent models of neural population dynamics’, *Network* **23**(1-2), 24–47.
- Burnham, K. P. & Anderson, D. R. (2002), *Model selection and multimodel inference: a practical information-theoretic approach*, Springer Science & Business Media.
- Butts, D. A., Weng, C., Jin, J., Alonso, J. & Paninski, L. (2011), ‘Temporal precision in the visual pathway through the interplay of excitation and stimulus-driven suppression.’, *J Neurosci* **31**(31), 11313–11327.

- Butts, D. A., Weng, C., Jin, J., Yeh, C.-I., Lesica, N. A., Alonso, J.-M. & Stanley, G. B. (2007), ‘Temporal precision in the neural code and the timescales of natural vision’, *Nature* **449**(7158), 92–95.
- Cafaro, J. & Rieke, F. (2010), ‘Noise correlations improve response fidelity and stimulus encoding’, *Nature* **468**(7326), 964–967.
- Cafaro, J. & Rieke, F. (2013), ‘Regulation of spatial selectivity by crossover inhibition’, *The Journal of Neuroscience* **33**(15), 6310–6320.
- Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L. & Rust, N. C. (2005), ‘Do we know what the early visual system does?’, *The Journal of Neuroscience* **25**(46), 10577–10597.
- Chander, D. & Chichilnisky, E. J. (2001), ‘Adaptation to Temporal Contrast in Primate and Salamander Retina’, *The Journal of Neuroscience* **21**(24), 9904–16.
- Chib, S. (1995), ‘Marginal Likelihood from the Gibbs Output’, *Journal of the American Statistical Association* **90**(432), 1313–1321.
- Chib, S. & Jeliazkov, I. (2001), ‘Marginal Likelihood From the Metropolis-Hastings Output’, *Journal of the American Statistical Association* **96**(453), 270–281.
- Chichilnisky, E. (2001), ‘A simple white noise analysis of neuronal light responses’, *Network: Computation in Neural Systems* **12**(2), 199–213.

- Churchland, A. K., Kiani, R., Chaudhuri, R., Wang, X.-J., Pouget, A. & Shadlen, M. N. (2011), ‘Variance as a Signature of Neural Computations during Decision Making’, *Neuron* **69**(4), 818–831.
- Churchland, A. K., Kiani, R. & Shadlen, M. N. (2008), ‘Decision-making with multiple alternatives’, *Nature Neuroscience* **11**(6), 693–702.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. (2010), ‘Cortical preparatory activity: representation of movement or first cog in a dynamical machine?’, *Neuron* **68**(3), 387–400.
- Citi, L., Ba, D., Brown, E. N. & Barbieri, R. (2014), ‘Likelihood methods for point processes with refractoriness’, *Neural Computation* **26**(2), 237–263.
- Clark, D. A., Benichou, R., Meister, M. & Azeredo da Silveira, R. (2013), ‘Dynamical adaptation in photoreceptors.’, *PLoS computational biology* **9**(11), e1003289.
- Cohen, M. R. & Kohn, A. (2011), ‘Measuring and interpreting neuronal correlations’, *Nature neuroscience* **14**(7), 811–819.
- David, S. V., Vinje, W. E. & Gallant, J. L. (2004), ‘Natural stimulus statistics alter the receptive field structure of v1 neurons.’, *The Journal of neuroscience : the official journal of the Society for Neuroscience* **24**(31), 6991–7006.
- De Boer, E. & Kuyper, P. (1968), ‘Triggered correlation’, *IEEE Transactions on Biomedical Engineering* **15**(3), 169–179.

- de Lafuente, V., Jazayeri, M. & Shadlen, M. N. (2015), ‘Representation of accumulating evidence for a decision in two parietal areas’, *The Journal of Neuroscience* **35**(10), 4306–4318.
- Deco, G., Pérez-Sanagustín, M., de Lafuente, V. & Romo, R. (2007), ‘Perceptual detection as a dynamical bistability phenomenon: a neurocomputational correlate of sensation’, *Proceedings of the National Academy of Sciences* **104**(50), 20073–20077.
- Ditterich, J. (2006), ‘Stochastic models of decisions about motion direction: behavior and physiology.’, *Neural Networks* **19**(8), 981–1012.
- Durstewitz, D. & Deco, G. (2008), ‘Computational significance of transient dynamics in cortical networks.’, *The European Journal of Neuroscience* **27**(1), 217–27.
- Escola, S., Fontanini, A., Katz, D. B. & Paninski, L. (2011), ‘Hidden Markov models for the stimulus-response relationships of multistate neural systems.’, *Neural Computation* **23**(5), 1071–132.
- Fairhall, A. L., Lewen, G. D., Bialek, W. & van Steveninck, R. R. d. R. (2001), ‘Efficiency and ambiguity in an adaptive neural code’, *Nature* **412**(6849), 787–792.
- Freeman, J., Vladimirov, N., Kawashima, T., Mu, Y., Sofroniew, N. J., Bennett, D. V., Rosen, J., Yang, C.-T., Looger, L. L. & Ahrens, M. B. (2014),

- ‘Mapping brain activity at scale with cluster computing’, *Nature methods* **11**(9), 941–950.
- Garvert, M. M. & Gollisch, T. (2013), ‘Local and global contrast adaptation in retinal ganglion cells’, *Neuron* **77**(5), 915–928.
- Gerstner, W. (2001), A framework for spiking neuron models: The spike response model, *in* F. Moss & S. Gielen, eds, ‘The Handbook of Biological Physics’, Vol. 4, pp. 469–516.
- Gerwinn, S., Macke, J. H. & Bethge, M. (2010), ‘Bayesian inference for generalized linear models for spiking neurons’, *Frontiers in Computational Neuroscience* .
- Girolami, M. & Calderhead, B. (2011), ‘Riemann manifold Langevin and Hamiltonian Monte Carlo methods’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(2), 123–214.
- Glimcher, P. W. (2003), ‘The neurobiology of visual-saccadic decision making’, *Annual review of neuroscience* **26**(1), 133–179.
- Gold, J. I. & Shadlen, M. N. (2007), ‘The neural basis of decision making’, *Annual Review of Neuroscience* **30**, 535–574.
- Goldman, M. S. (2015), Failure of averaging, *in* D. Jaeger & R. Jung, eds, ‘Encyclopedia of Computational Neuroscience’, Springer.

- Gollisch, T. & Meister, M. (2010), ‘Eye smarter than scientists believed: neural computations in circuits of the retina’, *Neuron* **65**(2), 150–164.
- Gordon, N., Salmond, D. & Smith, A. (1993), ‘Novel approach to nonlinear/non-Gaussian Bayesian state estimation’, *IEEE Proceedings F Radar and Signal Processing* **140**(2), 107.
- Hanks, T. D., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C. & Brody, C. D. (2015), ‘Distinct relationships of parietal and prefrontal cortices to evidence accumulation’, *Nature* .
- Harris, K., Csicsvari, J., Hirase, H., Dragoi, G. & Buzsaki, G. (2003), ‘Organization of cell assemblies in the hippocampus’, *Nature* **424**, 552–556.
- Herz, A. V., Gollisch, T., Machens, C. K. & Jaeger, D. (2006), ‘Modeling single-neuron dynamics and computations: a balance of detail and abstraction’, *Science* **314**(5796), 80–85.
- Hong, S., Lundstrom, B. N. & Fairhall, A. L. (2008), ‘Intrinsic gain modulation and adaptive neural coding.’, *PLoS computational biology* **4**(7), e1000119.
- Hubel, D. H. & Wiesel, T. N. (1962), ‘Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex’, *The Journal of physiology* **160**(1), 106.
- Huk, A. C. & Meister, M. L. R. (2012), ‘Neural correlates and neural computations in posterior parietal cortex during perceptual decision-making.’, *Frontiers in Integrative Neuroscience* **6**, 86.

- Huk, A. C. & Shadlen, M. N. (2005), ‘Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making’, *The Journal of Neuroscience* **25**(45), 10420–10436.
- Hunter, I. & Korenberg, M. (1986), ‘The identification of nonlinear biological systems: Wiener and hammerstein cascade models’, *Biological cybernetics* **55**(2-3), 135–144.
- Johnston, D., Wu, S. M.-S. & Gray, R. (1995), *Foundations of cellular neurophysiology*, MIT press Cambridge, MA.
- Kang, I. & Maunsell, J. H. (2012), ‘Potential confounds in estimating trial-to-trial correlations between neuronal response and behavior using choice probabilities’, *Journal of Neurophysiology* **108**(12), 3403–3415.
- Kastner, D. B. & Baccus, S. A. (2014), ‘Insights from the retina into the diverse and general computations of adaptation, detection, and prediction’, *Current opinion in neurobiology* **25**, 63–69.
- Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. (2014), ‘Cortical activity in the null space: permitting preparation without movement’, *Nature neuroscience* **17**(3), 440–448.
- Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. (2015), ‘Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex’, *eLife* **4**, e04677.

- Kiani, R., Cueva, C. J., Reppas, J. B. & Newsome, W. T. (2014), ‘Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials’, *Current Biology* **24**(13), 1542–1547.
- Kiani, R., Hanks, T. D. & Shadlen, M. N. (2008), ‘Bounded Integration in Parietal Cortex Underlies Decisions Even When Viewing Duration Is Dictated by the Environment’, *Journal of Neuroscience* **28**(12), 3017–3029.
- Kiani, R. & Shadlen, M. N. (2009), ‘Representation of confidence associated with a decision by neurons in the parietal cortex’, *Science* **324**(5928), 759–764.
- Kira, S., Yang, T. & Shadlen, M. N. (2015), ‘A neural implementation of walds sequential probability ratio test’, *Neuron* **85**(4), 861–873.
- Koulakov, A. A., Raghavachari, S., Kepecs, A. & Lisman, J. E. (2002), ‘Model for a robust neural integrator’, *Nature neuroscience* **5**(8), 775–782.
- Latimer, K. W., Chichilnisky, E., Rieke, F. & Pillow, J. W. (2014), Inferring synaptic conductances from spike trains with a biophysically inspired point process model, *in* ‘Advances in Neural Information Processing Systems’, pp. 954–962.
- Latimer, K. W., Huk, A. C. & Pillow, J. W. (2015), Bayesian inference for latent stepping and ramping models of spike train data, *in* Z. Chen, ed., ‘Advanced State Space Methods for Neural and Clinical Data’, Cambridge University Press.

- Latimer, K. W., Yates, J. L., Meister, M. L. R., Huk, A. C. & Pillow, J. W. (in press), ‘Single-trial spike trains in parietal cortex reveal discrete steps during decision-making’, *Science* .
- Link, S. W. (1975), ‘The relative judgment theory of two choice response time’, *Journal of Mathematical Psychology* **12**(1), 114–135.
- Ma, W. J. & Jazayeri, M. (2014), ‘Neural coding of uncertainty and probability’, *Annual review of neuroscience* **37**, 205–220.
- Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. (2013), ‘Context-dependent computation by recurrent dynamics in prefrontal cortex’, *Nature* **503**(7474), 78–84.
- Marava, M. (2013), Adaptation and sensory coding, *in* R. Q. Quiroga & S. Panzeri, eds, ‘Principles of Neural Coding’, CRC Press, p. 357.
- Mazurek, M. E., Roitman, J. D., Ditterich, J. & Shadlen, M. N. (2003), ‘A role for neural integrators in perceptual decision making’, *Cerebral Cortex* **13**(11), 1257–1269.
- McFarland, J. M., Cui, Y. & Butts, D. A. (2013), ‘Inferring nonlinear neuronal computation based on physiologically plausible inputs.’, *PLoS computational biology* **9**(7), e1003143.
- Mease, R. A., Famulare, M., Gjorgjieva, J., Moody, W. J. & Fairhall, A. L. (2013), ‘Emergence of adaptive computation by single neurons in the developing cortex’, *The Journal of Neuroscience* **33**(30), 12154–12170.

- Mechler, F. & Ringach, D. L. (2002), ‘On the classification of simple and complex cells’, *Vision research* **42**(8), 1017–1033.
- Meister, M. L. R., Hennig, J. A. & Huk, A. C. (2013), ‘Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making.’, *The Journal of Neuroscience* **33**(6), 2254–67.
- Mensi, S., Naud, R. & Gerstner, W. (2011), From stochastic nonlinear integrate-and-fire to generalized linear models., *in* ‘NIPS’, pp. 1377–1385.
- Miller, P. & Katz, D. B. (2010), ‘Stochastic transitions between neural states in taste processing and decision-making.’, *The Journal of Neuroscience* **30**(7), 2559–70.
- Newsome, W. T. & Pare, E. B. (1988), ‘A selective impairment of motion perception following lesions of the middle temporal visual area (MT)’, *The Journal of Neuroscience* **8**(6), 2201–2211.
- Ozuysal, Y. & Baccus, S. A. (2012), ‘Linking the computational structure of variance adaptation to biophysical mechanisms’, *Neuron* **73**(5), 1002–1015.
- Paninski, L. (2004), ‘Maximum likelihood estimation of cascade point-process neural encoding models’, *Network: Computation in Neural Systems* **15**, 243–262.

- Paninski, L., Pillow, J. W. & Lewi, J. (2007), ‘Statistical models for neural encoding, decoding, and optimal stimulus design.’, *Progress in brain research* **165**, 493–507.
- Paninski, L., Shoham, S., Fellows, M. R., Hatsopoulos, N. G. & Donoghue, J. P. (2004), ‘Superlinear population encoding of dynamic hand trajectory in primary motor cortex’, *The Journal of neuroscience* **24**(39), 8551–8561.
- Paninski, L., Vidne, M., DePasquale, B. & Ferreira, D. G. (2012), ‘Inferring synaptic inputs given a noisy voltage trace via sequential monte carlo methods’, *Journal of computational neuroscience* **33**(1), 1–19.
- Panzeri, S., Biella, G., Rolls, E. T., Skaggs, W. & Treves, A. (1996), ‘Speed, noise, information and the graded nature of neuronal responses’, *Network: Computation in Neural Systems* **7**(2), 365–370.
- Park, I. M., Archer, E. W., Priebe, N. & Pillow, J. W. (2013), Spectral methods for neural characterization using generalized quadratic models, *in* ‘Advances in Neural Information Processing Systems’, pp. 2454–2462.
- Park, I. M., Meister, M. L., Huk, A. C. & Pillow, J. W. (2014), ‘Encoding and decoding in parietal cortex during sensorimotor decision-making’, *Nature neuroscience* .
- Park, I. M. & Pillow, J. W. (2011), Bayesian spike-triggered covariance analysis, *in* ‘Advances in neural information processing systems’, pp. 1692–1700.

- Pillow, J. W., Paninski, L., Uzzell, V. J., Simoncelli, E. P. & Chichilnisky, E. J. (2005), 'Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model.', *The Journal of neuroscience* **25**(47), 11003–11013.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M. & Chichilnisky, E. J. Simoncelli, E. P. (2008), 'Spatio-temporal correlations and visual signaling in a complete neuronal population', *Nature* **454**, 995–999.
- Plesser, H. E. & Gerstner, W. (2000), 'Noise in integrate-and-fire neurons: from stochastic input to escape rates.', *Neural Computation* **12**(2), 367–384.
- Poo, C. & Isaacson, J. S. (2009), 'Odor representations in olfactory cortex: "sparse" coding, global inhibition, and oscillations.', *Neuron* **62**(6), 850–61.
- Priebe, N. J., Mechler, F., Carandini, M. & Ferster, D. (2004), 'The contribution of spike threshold to the dichotomy of cortical simple and complex cells', *Nature neuroscience* **7**(10), 1113–1122.
- Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D. & Palmeri, T. J. (2010), 'Neurally constrained modeling of perceptual decision making.', *Psychological review* **117**(4), 1113.
- Ratcliff, R. & Rouder, J. N. (1998), 'Modeling response times for two-choice decisions', *Psychological Science* **9**(5), 347–356.

- Rieke, F. (2001), ‘Temporal contrast adaptation in salamander bipolar cells’, *The Journal of Neuroscience* **21**(23), 9445–9454.
- Roberts, G. O. & Stramer, O. (2002), ‘Langevin Diffusions and Metropolis-Hastings Algorithms’, *Methodology And Computing In Applied Probability* **4**(4), 337–357.
- Roitman, J. D. & Shadlen, M. N. (2002), ‘Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task’, *The Journal of Neuroscience* **22**(21), 9475–9489.
- Roska, B., Molnar, A. & Werblin, F. S. (2006), ‘Parallel processing in retinal ganglion cells: how integration of space-time patterns of excitation and inhibition form the spiking output’, *Journal of Neurophysiology* **95**(6), 3810–3822.
- Schwartz, O., Chichilnisky, E. & Simoncelli, E. P. (2002), ‘Characterizing neural gain control using spike-triggered covariance’, *Advances in neural information processing systems* **1**, 269–276.
- Seidemann, E., Meilijson, I., Abeles, M., Bergman, H. & Vaadia, E. (1996), ‘Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task’, *Journal of Neuroscience* **16**(2), 752–768.
- Shadlen, M. N. & Kiani, R. (2013), ‘Decision making as a window on cognition.’, *Neuron* **80**(3), 791–806.

- Shephard, N. & Pitt, M. K. (1997), ‘Likelihood analysis of non-gaussian measurement time series’, *Biometrika* **84**(3), 653–667.
- Smith, P. L. & Ratcliff, R. (2004), ‘Psychology and neurobiology of simple decisions’, *Trends in neurosciences* **27**(3), 161–168.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P. & van der Linde, A. (2002), ‘Bayesian measures of model complexity and fit’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64**(4), 583–639.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P. & van der Linde, A. (2014), ‘The deviance information criterion: 12 years on’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(3), 485–493.
- Stevenson, I. H. & Kording, K. P. (2011), ‘How advances in neural recording affect data analysis’, *Nature neuroscience* **14**(2), 139–142.
- Stevenson, I. H., London, B. M., Oby, E. R., Sachs, N. A., Reimer, J., Englitz, B., David, S. V., Shamma, S. A., Blanche, T. J., Mizuseki, K., Zandvakili, A., Hatsopoulos, N. G., Miller, L. E. & Kording, K. P. (2012), ‘Functional connectivity and tuning curves in populations of simultaneously recorded neurons.’, *PLoS Comput Biol* **8**(11), e1002775.
- Theis, L., Chagas, A. M., Arnstein, D., Schwarz, C. & Bethge, M. (2013), ‘Beyond glms: A generative mixture modeling approach to neural system identification’, *PLoS Computational Biology* **9**(11), e1003356.

- Tokdar, S., Xi, P., Kelly, R. C. & Kass, R. E. (2010), ‘Detection of bursts in extracellular spike trains using hidden semi-Markov point process models.’, *Journal of Computational Neuroscience* **29**(1-2), 203–12.
- Trong, P. K. & Rieke, F. (2008), ‘Origin of correlated activity between parasol retinal ganglion cells.’, *Nature neuroscience* **11**(11), 1343–51.
- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P. & Brown, E. N. (2005), ‘A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects’, *J. Neurophysiol* **93**(2), 1074–1089.
- Turaga, S., Buesing, L., Packer, A. M., Dalglish, H., Pettit, N., Hausser, M. & Macke, J. (2013), Inferring neural population dynamics from multiple partial recordings of the same neural circuit, *in* C. Burges, L. Bottou, M. Welling, Z. Ghahramani & K. Weinberger, eds, ‘Advances in Neural Information Processing Systems 26’, Curran Associates, Inc., pp. 539–547.
- Uzzell, V. & Chichilnisky, E. (2004), ‘Precision of spike trains in primate retinal ganglion cells’, *Journal of Neurophysiology* **92**(2), 780–789.
- van Hateren, J. v., Rüttiger, L., Sun, H. & Lee, B. (2002), ‘Processing of natural temporal stimuli by macaque retinal ganglion cells’, *The Journal of neuroscience* **22**(22), 9945–9960.

- Vintch, B., Zaharia, A., Movshon, J. & Simoncelli, E. P. (2012), Efficient and direct estimation of a neural subunit model for sensory coding, *in* ‘Advances in neural information processing systems’, pp. 3104–3112.
- Wald, A. (1973), *Sequential analysis*, Courier Corporation.
- Wark, B., Lundstrom, B. N. & Fairhall, A. L. (2007), ‘Sensory adaptation.’, *Current Opinion in Neurobiology* **17**(4), 423–9.
- Wiecki, T. V., Sofer, I. & Frank, M. J. (2013), ‘HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python.’, *Frontiers in Neuroinformatics* **7**, 14.
- Yang, T. & Shadlen, M. N. (2007), ‘Probabilistic reasoning by neurons’, *Nature* **447**(7148), 1075–1080.
- Yates, J., Park, I., Cormack, L., Pillow, J. & Huk, A. (2013), Precise characterization of multiple lip neurons in relation to stimulus and behavior, Presented at COSYNE, Salt Lake City, USA.
- Yuan, K., Girolami, M. & Niranjana, M. (2012), ‘Markov chain Monte Carlo methods for state-space models with point process observations.’, *Neural computation* **24**(6), 1462–86.

Vita

Kenneth William Latimer was born at the United States Air Force Academy on 12 September, 1987. He moved to Boulder, Colorado in 2005 where he learned to brew beer while studying computer science at the University of Colorado, Boulder. After joining Professor Randy O'Reilly's lab as an undergraduate research assistant, he decided to pursue a graduate degree in Neuroscience and in the summer of 2010, he moved to Austin to join the Institute for Neuroscience at UT Austin. He plans to move to Seattle in the fall of 2015 to join Adrienne Fairhall's lab at the University of Washington as a postdoc.

Permanent address: 1928 Chapel Hills Dr.
Colorado Springs, Colorado 80920

This dissertation was typeset with L^AT_EX[†] by the author.

[†]L^AT_EX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's T_EX Program.