Neural computation in the context of upstream dynamics in the retina

Philip R. Mardoum

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Fred Rieke, Chair

Rachel Wong, Chair

Adrienne Fairhall

David Perkel

Program Authorized to Offer Degree:
Neuroscience

University of Washington

Abstract

Neural computation in the context of upstream dynamics
in the retina

Philip R. Mardoum

Co-Chairs of the Supervisory Committee:
Professor Fred Rieke
Department of Physiology & Biophysics

Professor Rachel Wong
Department of Biological Structure

To understand neural circuit function, one would like to understand individual neurons' computation in the context of their physiological inputs. But these inputs are themselves subject to complex dynamics, and we rarely have tools to experimentally control synaptic inputs under physiological conditions that preserve their temporal features. Primary sensory structures present an exception because primary receptor neurons can be controlled experimentally under physiological conditions. Here, I present work carried out in the retina, where signaling by photoreceptors (the primary receptor neurons in vision) has been characterized in detail and can be controlled with light, and computation in downstream circuitry can therefore be investigated in the context of physiological input from photoreceptors. I first present a hybrid biophysical-statistical model of retinal output that disentangles the computational contributions of photoreceptors from those of other circuit elements and successfully predicts retinal ganglion cell responses to stimuli with dynamically changing statistics. Second, I present an investigation of synaptic specializations that could mediate parallel processing of input from different photoreceptor types within individual post-synaptic neurons.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

Thanks first to my advisers, Fred Rieke and Rachel Wong. I've been profoundly privileged to work closely with them for these years and to participate in the fascinating research programs they have built. They are brilliant scientists and supportive mentors, and I won't forget the example set by their leadership. Thanks to Fred for being an incredible teacher. His ability to keep the big picture in view and to animate the beauty in neural function has been a constant source of fuel. Thanks to Rachel for her infectious excitement and joy in the day-to-day work of doing science. Rachel makes it impossible to forget that we're lucky to do this work.

While working on our zebrafish project, Takeshi Yoshimatsu was effectively my third adviser, and I'm fortunate to have had his mentorship in many areas that were completely new to me. Thanks also to Owen Lawrence, Ayana Hellevik, and Raunak Sinha for their extensive contributions as collaborators on that project.

To my current and past colleagues at UW, you all have been the best part of this experience and the reason I felt excited to come to work every day. I learned from you the kind of colleague I want to be. A special thanks to fellow graduate students Jacob Baudin, Florie D'Orazi, Clare Gamlin, Max Turner, and Ali Weber for their constant help with the many challenges I inevitably encountered on a daily basis. I'm so grateful also for mentorship from Felice Dunn, Juan Angueyra, Mike Manookin, Raunak Sinha, Mrinalini Hoon, Sid Kuo, Will Grimes, Luca Della Santina, and Adam Bleckert. Thanks to Shellee Cunnington, Mark Ca-

faro, and Mike Ahlquist for their technical support and design of lab processes and software that multiplied my productivity and morale. To Greg Newkirk for making sure none of us ever took ourselves too seriously. And to Todd Appleby, Braden Brinkman, Chris Chen, Julian Freedland, Gabrielle Gutierrez, Norianne Ingram, Haruhisa Okawa, Adree Songco-Aguas, Sachihiro Suzuki, Zhou Yu, Wan-Qing Yu, and Chi Zhang. Thinking on the list of everyone I've been able to work with during my time as a graduate student, I am sincerely astounded by my luck.

I'm grateful to members of my thesis committee Adrienne Fairhall, David Perkel, and David Gire. Their guidance helped me chart a course through these years' twists and turns.

Thanks to the Neuroscience Graduate Program and to my departments, Physiology and Biophysics and Biological Structure, and for funding and career support from the University of Washington Institute for Neuroengineering and the Computational Neuroscience Training Grant.

I want to thank three people who played a transformative role in shaping my scientific interests before my graduate work and inspiring me to pursue the trajectory I have followed. This journey started in the classroom of my high school science teacher David Johnson. If not for his teaching, I think I would not be a scientist. My former adviser, Nicho Hatsopoulos, took a chance giving me a project in his lab when I was an undergraduate with no experience, and I have spent all the years since chasing the joys in research that I found in his lab. For a year before joining the Neuroscience program, I was fortunate to be welcomed as a member of Adrienne Fairhall's lab, and this was one of the best years I've spent in research. I can't imagine another lab where I could have learned so much about so many topics in short time, and Adrienne introduced me to many of the questions that drive my

curiosity still today.

To my family, my friends, and my partner Ali, I can't imagine any of this without your support and love. I appreciate you more than anything.

Chapter 1
INTRODUCTION

Even as there have been fast advances in the study of neural computation, underlying biophysical mechanism, and neural circuit structure, there remain some fundamental obstacles to bringing these domains together in such a way as would constitute a wholistic understanding of function. We can often fairly efficiently characterize neurons' responses to exogenous stimuli delivered directly through scientific instruments, but it is far more difficult to characterize a neuron's responses to physiological input from other neurons. This is because we almost never have tools to control those synaptic inputs under physiological conditions that would preserve their temporal features. With obstacles to controlled experiments in the context of physiological signals, how can we understand the contributions of particular cell types and mechanisms within circuits which distribute computations across many loci? Approaching this task, we encounter the reality that in neural circuits the activity of each cell depends on the others, with no clear foothold with which to begin an investigation without complex upstream dependencies. In the thick of the central nervous system, even the idea of a neural circuit becomes over-strained as immense interconnectivity defies any convenient notion of functional compartments. Surely there is a long road ahead.

This dissertation leverages an exception to the rule that upstream dependencies are inescapable. In sensory neural systems, following connections upstream eventually leads you to the sensory periphery, the vanguard of neurons that interface directly with the out-

side world. Physical stimuli – like molecules we smell and taste, mechanical disturbances we feel or hear, and photons we see – are these primary receptor neurons' main inputs, not signals from other neurons. This provides a rare foothold: not only can we characterize the behavior of primary receptor neurons without the worry of upstream neural dependencies, we can also use physical stimuli to control the output that these neurons convey to downstream circuitry, all under physiological conditions that preserve temporal features of signaling. This dissertation's focus is on questions we can ask about that downstream circuitry. How can a neuron's computation be understood in the context of input sources that are themselves subject to complex dynamics? How can multiple computations performed within a circuit be located within the circuit and tied to particular mechanisms? How are downstream computations specialized to the particular signal and noise properties of their inputs?

In this work we never stray more than two synapses beyond the primary receptors, always retaining access to this foothold. Hopefully, research in primary sensory areas will continue to serve as a proving ground for tools that will guide deeper expeditions in the future.

The subject of this work is the retina, where the primary receptor neurons are the photoreceptors. The retina is part of the central nervous system and is formed from three layers: one of photoreceptors, one of interneurons like bipolar cells, and one of retinal ganglion cells whose axons form the optic nerve and connect with other visual areas of the brain. Each layer processes visual information, parsing features of the visual scene and conveying them along specialized channels formed by a diversity of neuron types in each layer. The retina's relatively simple three-layer structure, along with numerous experimental conveniences, have made it a popular target for research aiming to understand principles of neural circuitry. It is also a focus of vision research and an effort to develop technologies for medical interventions, called retinal prostheses, that may soon help treat

some forms of blindness.

There is great interest, then, in developing computer models of retinal function. For example, many models aim to predict patterns of action potentials, or spikes, generated by RGCs in response to light stimuli. Reasons to pursue good models are both scientific and practical. As in many scientific fields, a useful benchmark of our understanding of a system's function is to formalize our current understanding in a set of rules that instantiates our best approximation of the rules followed by the real system. Even flawed models are useful because their failure modes can help us identify particular gaps in our knowledge. More research is done, a new or updated model is proposed, and we try again.

Currently there are no models that can predict retinal output in realistic, or 'naturalistic,' conditions. But many models perform quite well in limited experimental conditions. Also, smaller models of certain individual components of the retinal circuit (for example, cone photoreceptors) are close to a full representation of the mechanisms needed to predict responses to naturalistic stimuli. We may be well on our way toward generalizable models that will predict retinal output across a range of visual conditions. Success here would be good news for those hoping to model more complicated systems such as the deeply embedded circuits of the brain, and would also complete a key pillar in the advances needed to develop retinal prostheses.

We studied computation in the second and third layers of retina in the context of the dynamics of input received from upstream sources. Proximity to primary receptor neurons presents an opportunity to understand the role of upstream mechanisms in a setting where they have been deeply characterized and can be controlled experimentally. Most studies of retinal function, however, do not fully leverage this proximity. Even in the retina, the most common strategy to study function is to characterize neural responses with respect to external stimuli, not with respect to physiological input received by the cell. Most models of retinal output, for example, are built using only RGC responses to visual stimuli.

These models do not attempt to explicitly represent circuit components or mechanisms. Thus, insofar as they capture features of RGC output, nothing can be said about which features originate in the RGC itself and which originate in upstream circuitry. This is not by itself a problem if one's only goal is to predict retinal output, but I will argue (in Chapter 2) that an explicit representation of multiple processing stages in the retina is a fruitful strategy even for that goal.

In the following sections I will offer some background and reasoning in a small number of areas relevant to justifying the research presented in subsequent chapters. This background will focus on neural adaptation, two model classes, and photoreceptors.

Neural adaptation

Neural adaptation is the term broadly applied to changes in a neuron's response properties that occur when some feature of the stimulus changes. Adaptation is considered conceptually distinct from the response itself, i.e. it is not simply a change in the response (as would likely happen with any time-varying input) but rather a change in the input-to-output mapping. For example, neurons appear to adjust their sensitivity to suit the intensity of input being received: as the intensity of input increases, the cell decreases its sensitivity to avoid saturation of its response range (Fairhall et al., 2001; Nagel and Doupe, 2006; Tobler et al., 2005). This can be an important adjustment to maintain sensitivity when the dynamic range of the neural response is small compared to the range of all possible inputs. The retina contains a dramatic example of the mismatch between the input and response ranges: we can see over a range of light levels that exceeds the dynamic range of neural output a billion fold (Dunn and Rieke, 2008). A useful frame has been to understand adaptation in terms of information theory, wherein an optimally efficient coding strategy requires a matching of the dynamic range of responses to the local statistics of the input signals (Laughlin, 1981; Wark et al., 2007; Brenner et al., 2000).

Frequent changes in the local input statistics, then, demand an adjustment of the input-output mapping to account for changes in the input distribution. More concretely, if the dynamic range of responses is very large relative to the range of inputs, sensory signals could be overwhelmed by noise in neural output and be rendered imperceptible. If, on the other hand, the dynamic range of responses is small relative to the inputs, the neural response will frequently be saturated, causing information loss.

The ability of neurons to adapt to local input statistics is therefore an essential biological innovation, but quite an inconvenience for modelers wishing to define the mapping from neural input to output, since this relationship is in constant flux.

Adaptation is a ubiquitous property of neural systems carried out by a diversity of mechanisms across a wide spectrum of timescales (Fairhall et al., 2001; Weber and Fairhall, 2019; Hosoya et al., 2005). Indeed, the distinction between adaptation and other neural computational business is purely conceptual, as it often shares the same mechanisms and timescales as behavior traditionally categorized as “responses" (if a neuron responds only to stimulus changes and not steady activation, is this an aspect of its response selectivity, or is it “fast-adapting?"). Depending on the particular system, though, certain mechanisms may be found whose role is particularly adaptational in nature, especially under constrained experimental conditions. In such cases adaptation becomes a useful framework with which to understand function.

In the retina, two well studied categories of adaptation to temporal stimulus characteristics are adaptation to average luminance or light level, called mean adaptation, and adaptation to the variance of temporal fluctuations, called contrast adaptation (Grimes et al., 2014; Baccus and Meister, 2002; Tikidji-Hamburyan et al., 2015; Dunn et al., 2007). These are usually studied in the context of random gaussian-distributed white noise stimuli which have an explicit mean and variance. Indeed, this categorization of adaptation stems from the choice of stimulus: if a circuit with many adaptational mechanisms is pre-

sented with a stimulus whose main properties are mean and variance, the circuit will likely adapt to mean and variance (Fairhall et al., 2001). Adaptation to other stimulus statistics has also been observed using more complex stimuli (Tkačik et al., 2014; Sharpee et al., 2006). Even so, in the retina a distinction between mean and contrast adaptation proves to be informative: some studies have tied particular mechanisms more to one adaptation type than the other. For example, cones adapt strongly to changes in mean but weakly to contrast (Smirnakis et al., 1997; Rieke, 2001; Dunn and Rieke, 2006). Contrast adaptation has been attributed to various circuit elements including cone bipolar cell output synapses and RGCs (Kim and Rieke, 2001; Manookin and Demb, 2006; Ozuysal and Baccus, 2012; Cui et al., 2016).

The best-studied aspect of adaptation in retina is changes in response gain, but there are other effects including changes in mean response level, response delay, and response feature selectivity such as temporal frequency preference (Smirnakis et al., 1997; Baccus and Meister, 2002; Chander and Chichilnisky, 2001). Changes in response gain in mean adaptation have been studied using brief flashes delivered on top of different baseline illumination. The flash response strength scales inversely with the baseline light level (Angueyra et al., 2021). At most daytime light levels this scaling is equal in proportion to the change in stimulus baseline, i.e. a doubling of the stimulus baseline causes a halving of the response gain. The constancy of this ratio over a broad range of light levels is called Weber adaptation (after the related Weber's law, which was originally established in psychophysical experiments (Fechner, 1966)). Contrast adaptation has been studied using noise stimuli with periodic switches in variance, or temporal contrast. Following an increase in contrast, responses are briefly much stronger but over time partially relax back down as sensitivity adapts to the new contrast level (Smirnakis et al., 1997; Shapley and Victor, 1978).

Attributes and origin of mean and contrast adaptation likely depend on particular stim-

ulus features such as the timescales being investigated, in addition to intrinsic factors like species. For example, the speed and extent of adaptation has been shown to depend on the frequency with which stimulus statistics change, as if adaptational mechanisms become primed to adapt more quickly when it is frequently needed (Fairhall et al., 2001; Wark et al., 2009) (these adjustments are a kind of adaptation too). Dynamics of adaptation are complex and may reflect a neural solution to an inference problem in which neurons must infer the local stimulus distribution quickly enough to avoid saturation, but not so quickly that the estimation of the local stimulus distribution is overly noisy (Wark et al., 2009).

Two classes of models

The two main classes of models in sensory neuroscience are dynamical models and statistical models. Dynamical models use differential equations to describe neural activity, often explicitly modeling biophysical phenomena of a cell such as molecular interactions and ion conductances though channel proteins (Izhikevich, 2018). Statistical models express correlative relationships between neural activity and other variables such as a stimulus, often using dimensionality reduction techniques to generate phenomenological components like linear filters and nonlinear response functions that bear no explicit connection to physiological mechanisms (Rieke et al., 1997). Both model types can take in an input signal and predict an output, but while dynamical models often represent the computational machinery of a cell, statistical models often represent neural 'coding' (Weber and Fairhall, 2019).

The greater complexity of dynamical models (in this text often also called biophysical models) makes them much more demanding to constrain with experiments, and they are therefore comparatively rare. Their behavior is often high dimensional and difficult to visualize, and indeed statistical models are sometimes fit to responses generated by dynami-

cal models in order to visualize emergent coding properties (Agüera y Arcas et al., 2003; Hong et al., 2007; Latimer and Fairhall, 2020). But by capturing fundamental dynamical mechanisms underlying neural responses, dynamical models allow for time dependence in responses that has not been matched by statistical models. A complete biophysical model of the retina and all its cell types would solve many problems, but the strategy of detailing the molecular goings on of individual cell types does not easily scale to neural circuits. Even in the retina, perhaps the best understood brain region, there are many cell types and very few of them have been modeled at the biophysical level.

Fortunately, differential equations describing phototransduction, the molecular process whereby light is converted into an electrical current, have been developed over many years (Pugh and Lamb, 1993; Rieke and Baylor, 1998b; van Hateren and Snippe, 2007; Angueyra et al., 2021). Recent advances have enhanced the ability of dynamical models to capture adaptational behaviors observed in cones at multiple timescales (Angueyra et al., 2021). In Chapter 2 we propose a hybrid model with dynamical and statistical components. A dynamical model of phototransduction in cones provides a deep mechanistic representation of the first layer of retinal computation that captures nonlinear and adaptational processes of cones. A multi-stage statistical model of a post-receptor circuit, constrained using multiple recording techniques to distinguish computations performed in different locations, provides a mapping from cone output to the output of a common RGC type.

Photoreceptor types and convergence

The two main classes of photoreceptor are rods and cones. Rods mediate vision at low light levels and are exquisitely sensitive. Phototransduction in rods is specialized to massively amplify any signal from absorbed photons (Rieke and Baylor, 1998a). This comes at the cost of temporal resolution as amplification continues to ramp up in a rod

impulse response well after the time a cone impulse response reaches its end. Partially because of rods' poor temporal resolution (and because of a relative paucity of photons in rod-mediated vision in the first place), it is quite difficult to catch a baseball by moonlight, even if in the same lighting conditions our rods might easily allow us to take a walk.

Cones mediate daytime vision and transmit visual information with much greater temporal fidelity. With the relative glut of available photons, the slow work of amplification is a much smaller priority. Cones are much less sensitive than rods. At the lowest light levels, cones are unable to distinguish the small number of photons from the many spontaneous isomerizations of their photopigment (these spontaneous isomerizations are much more frequent than in rods). Throughout most of the range of light levels in which cones are active, rod signaling is entirely saturated and is therefore just as debilitated as cone signaling at night.

Despite major differences between rods and cones and the temporal characteristics of their signals, rod and cone signals do not traverse the retina through their own private channels. Instead, rod and cone signals share the same downstream circuitry: they converge in multiple circuit locations and are ultimately transmitted by the same RGCs (Grimes et al., 2018). This multiplexing of rod- and cone-mediated signals has remarkable implications for the circuitry's ability to process information with very different signal and noise spectra. In addition, there is an interval of light levels (think of dusk) in which both rods and cones contribute to vision (Grimes et al., 2015). The computational mechanisms that support this impressive parallel processing remain poorly understood.

In Chapter 3, I discuss a study of bipolar cells that receive convergent input from rods and cones. In this microcircuit, individual synapses could be specialized to differentially process these two input populations, performing parallel processing within the dendrites of a single cell. We show morphological and physiological evidence for specialization of rod and cone synapses onto one of the studied bipolar cell types. This project ultimately

falls short, though, of tying these specializations to the computational demands of each input's particular temporal characteristics. I hope future studies will soon investigate similar questions with success.

Chapter 2

A HYBRID BIOPHYSICAL-STATISTICAL MODEL OF RETINAL OUTPUT

2.1 Introduction

Under natural viewing conditions, visual input is characterized by rapid transitions in the statistics of light stimuli sensed by visual neurons. Large eye movements (saccades) segment vision into a series of fixations in which light-sensitive neurons are exposed to regions of the visual scene with vastly different brightness, and micro-saccades explore local variations in luminance that produce temporal fluctuations of different magnitudes (Reinagel and Zador, 1999; Kuang et al., 2012). In a glance from a shadowed patch of grass to a cloudless sky, for example, the visual stimulus changes from one with low mean luminance and high temporal contrast to one with high mean luminance and low contrast (Frazor and Geisler, 2006; Simoncelli and Olshausen, 2001). External and self motion can also cause rapid transitions of this kind.

In order to understand natural vision and to enable design of visual prosthetic technologies, then, models of visual processing must capture neural responses to stimuli whose statistics rapidly and frequently change. But current models that predict the responses of retinal ganglion cells (RGCs, the cells mediating signaling to the brain) fall short of maintaining accurate predictions of neural responses through large changes in brightness and contrast. This is in large part because changes in stimulus statistics engage adaptational circuit mechanisms that alter the mapping between visual input and the neural representation.

Importantly, multiple adaptational mechanisms are active in different locations in the

retina. Despite this reality, the current dominant strategy for modeling RGC responses is to constrain models using only recordings from RGCs themselves. This strategy misses an opportunity to incorporate features and constraints based on extensive existing knowledge of nonlinear processing and adaptation in other circuit locations. For example, cones, the first stage of retinal processing, contribute the majority of the retina's adaptation to mean luminance, exhibiting strong shifts in gain in response to sustained changes in activation (Dunn et al., 2007). Nonetheless the first stage of nearly all RGC models is simply a linear filter. This approach comes at a cost: mathematically, steps in a sequence of nonlinear operations are not commutative, i.e. their order matters. In addition, current models may simply employ too few components to capture adaptation to both mean and contrast. Grouping circuit mechanisms into too few model components forces mechanisms to trade against each other and diminishes a model's interpretability.

We constructed a sequential multi-stage model of primate RGC responses that is fit using data collected from multiple cell classes (cones and RGCs) and multiple recording techniques (voltage-clamp recordings in cones and RGCs, plus extracellular and current-clamp recordings in RGCs). The multi-stage model is built from three sub-models: 1) a detailed biophysical model of phototransduction in cones, 2) a linear-nonlinear (LN) cascade model that maps cone responses to RGC excitatory drive, and 3) a spike generation model that incorporates response dependence on spike history. The cone model was developed previously (Angueyra et al., 2021) and is based on differential equations used to describe phototransduction (Pugh and Lamb, 1993; Rieke and Baylor, 1998b; van Hateren and Snippe, 2007). The RGC model was constrained separately using specialized injected current experiments. Finally, we recorded primate RGC responses to visual noise stimuli with frequent shifts in temporal mean and contrast. The multi-stage model recapitulates the majority of mean- and contrast-driven gain control observed in these recordings, maintaining response prediction accuracy across the tested range of stimulus

statistics. We find that mean- and contrast-driven gain control occur at different stages of the model, with mean adaptation contributed by the cone stage and contrast adaptation contributed during spike generation. Our recordings and modeling illustrate a promising approach of hybridizing biophysical and empirical model components to capture multiple coincident adaptational mechanisms that are engaged during sensing.

2.2 Results

Eye movements around a visual scene interact with spatial variations in luminance and contrast to produce rapid changes in temporal statistics of stimuli reaching the retina (Fig. 2.1 A) (Kuang et al., 2012). We recorded responses of On-parasol RGCs to a spatially uniform noise stimulus whose mean and temporal contrast were changed periodically (Fig. 2.1 B). In each short epoch, stimulus mean and contrast were randomly drawn from a small set of combinations. In a typical experiment, epochs were 500 ms long with mean luminances of 1200, 4000, or 12000 isomerizations (R^*) per cone per second and with temporal contrasts of 25 or 50% (i.e. the noise standard deviation was 25 or 50% of the mean). There was some variation between experiments in the particular stimulus conditions and epoch lengths tested (see Methods).

Results are presented in four sections in which On-parasol RGCs (hereafter simply called RGCs) serve as a testing ground for our modeling approach. We first show that the majority of mean-driven gain control (here called mean adaptation) occurs in circuitry upstream of ganglion cells while contrast-driven gain control (contrast adaptation) is performed by RGCs. In the following two sections, we demonstrate that a biophysical model of phototransduction in cones accounts for most or all mean adaptation observed in our recordings, and that a simple model of spike generation accounts for most or all observed contrast adaptation. Finally, we test a multi-stage model's ability to capture responses across conditions.

2.2.1 Measuring mean- and contrast-driven gain control in RGC excitatory currents and spike responses

To distinguish between adaptation occurring in retinal circuitry upstream of the RGC soma and adaptation occurring during spike generation in the RGC itself, we measured exci-

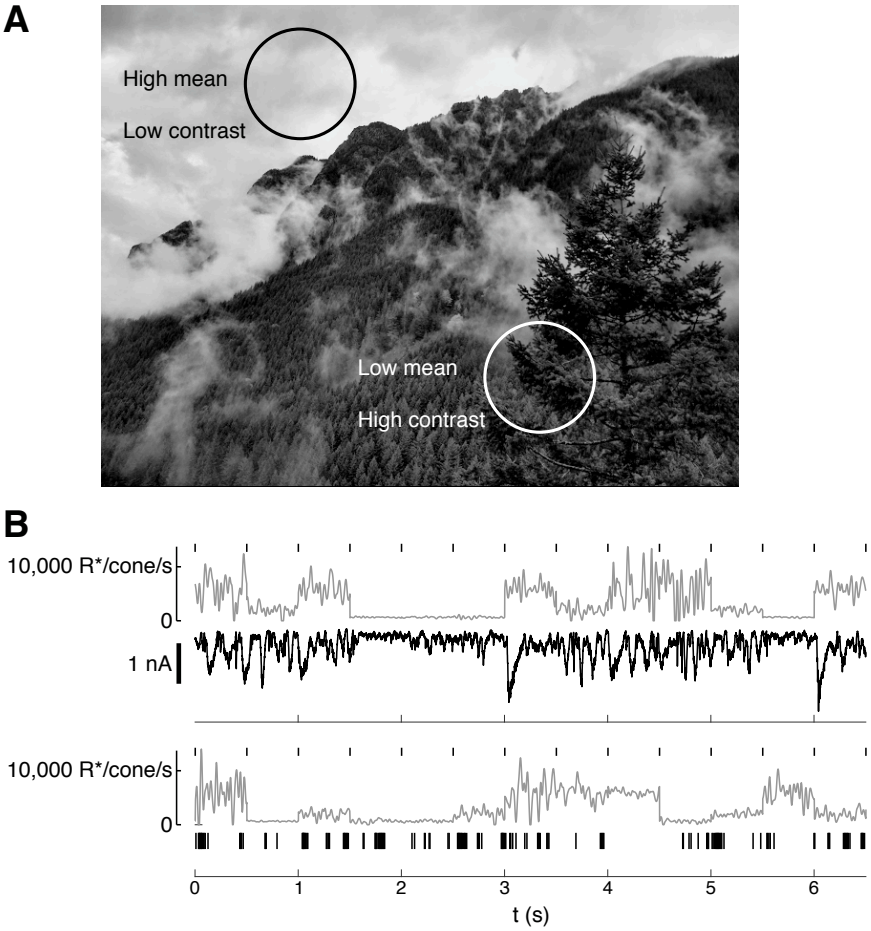

Figure 2.1: Dynamic noise stimuli engage adaptational mechanisms activated during eye movements. A: Spatial variations in luminance and contrast interact with eye movements to produce temporal variations in stimulus statistics. **B:** Top: Noise stimulus (gray) and excitatory current response (black) recorded in an On-parasol RGC. Bottom: Noise stimulus and spike response. Tick marks above stimulus: start of each epoch.

tatory currents and spike output in the same RGCs using whole-cell and cell-attached recording techniques respectively. To measure gain changes associated with adaptation to mean and contrast, we first pooled all epochs of each mean+contrast combination to create condition-specific collections which we used to characterize reponses in individual conditions (Fig. 2.2 A: two mean conditions, Fig. 2.3 A: two contrast conditions). In assembling these collections, the first 200 ms of each epoch was discarded to remove transient responses to the instantaneous change in stimulus condition that occurred between epochs.

Using epochs pooled at two different stimulus means (4000 and 8000 - 13,333 R^* / cone/ second) and the same contrast (50% contrast), we fit reverse correlation-based LN models (revcor LN models) to excitatory currents and spikes at each mean using standard white noise methods (Chichilnisky, 2001) (Fig. 2.2 A-B). To consolidate the representation of gain into the models' nonlinearities, the linear filters were normalized such that a noise signal passed through both filters would produce generator signals (filter output) with equal amplitude. An increase in stimulus mean corresponded with a decrease in the slope of the best-fit revcor LN models' response nonlinearities (Fig. 2.2 B, overlay). This effect is seen in the models based on excitatory current responses and in the models based on spike responses. The slope of the nonlinearity reflects the sensitivity of the response to changes in the generator signal. A slope decrease indicates a reduction in response gain to temporal fluctuations.

To quantify the extent to which the RGC adjusts its response gain when the stimulus mean changes, we measured the horizontal scaling of the response nonlinearities between means. We obtained the scaling factor α that, when applied to the x-axis of the nonlinearity corresponding to the lower mean intensity, best aligns that nonlinearity with the higher-mean nonlinearity (Fig. 2.2 B, bottom). This factor represents the change in slope (therefore change in gain) associated with the revcor LN models fit at the two

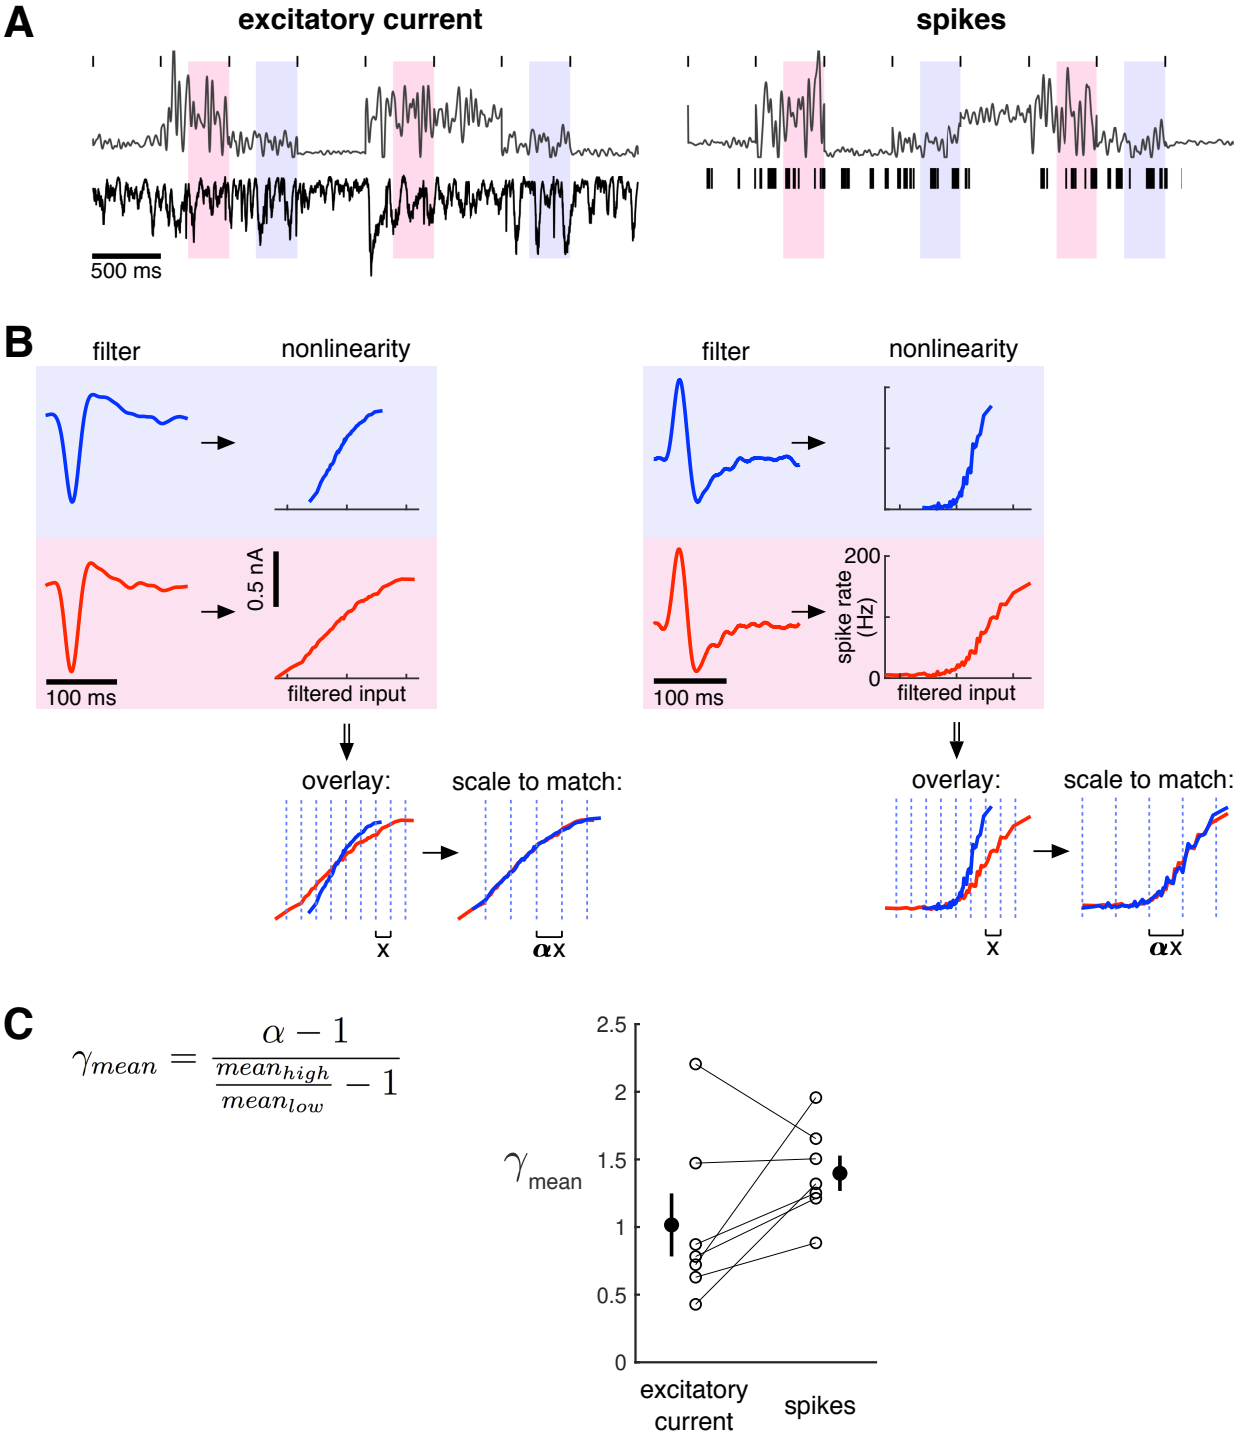
A
excitatory current
spikes
500 ms
B
filter
nonlinearity
filter
nonlinearity
0.5 nA
100 ms
filtered input
200
spike rate (Hz)
0
100 ms
filtered input
overlay:
scale to match:
overlay:
scale to match:
x
αx
x
αx
C
$\gamma_{mean} = \frac{\alpha - 1}{\frac{mean_{high}}{mean_{low}} - 1}$
γ_{mean}
2.5
2
1.5
1
0.5
0
excitatory current
spikes

Figure 2.2: *Previous page.* **Mean luminance-driven gain control measured in RGC responses. A:** Variable noise stimuli with excitatory current (left) and spike (right) responses. Shaded sections show example data segments pooled to fit models at each of two stimulus means. **B:** Top: Reverse correlation-based linear-nonlinear cascade (revcor LN) models fit to pooled data segments from each of two stimulus means. Bottom: Example of process to measure difference in gain between models fit at two means by quantifying the difference in x-axis scale. The scale difference factor α is stored for use computing γ in (C). **C:** Left: Formula for γ_{mean}, which is the fractional difference in gain relative to the fractional difference in stimulus mean. Right: γ_{mean} for 7 cells, computed using revcor LN models fit at two means on measured excitatory current and spike responses. Filled circles: mean $\pm$ SEM.

means. In all cells tested (n = 7), an expansion of the x-axis scale of the low-mean nonlinearity was required to align it with the higher-mean nonlinearity, meaning response gain was always decreased at the higher mean relative to the lower mean. To quantify the change in gain relative to the change in mean, we define an index γ_{mean} to describe the fractional change in gain ($\alpha - 1$) relative to the fractional change in the stimulus means:

$$\gamma_{mean} = \frac{\alpha - 1}{\frac{mean_{high}}{mean_{low}} - 1} \tag{2.1}$$

where $mean_{low}$ and $mean_{high}$ are the two stimulus means. For example, a gain change index γ of 0 indicates no scaling is required to match the two nonlinearities, i.e. the gain has not changed, while a γ of 1 indicates that the fractional increase in horizontal scale is comensurate with the fractional increase in mean or contrast (e.g. if the horizontal scale differs by a factor of 2 following a doubling of the mean luminance).

Fig. 2.2 C shows γ corresponding to a change in stimulus mean for excitatory currents and spikes. The average γ computed on excitatory input did not differ significantly from 1 (p = 0.95, one-sample t-test), which is consistent with previous reports that mean adaptation follows Weber's law at the light levels tested (Burkhardt, 1994). Similar γ values were

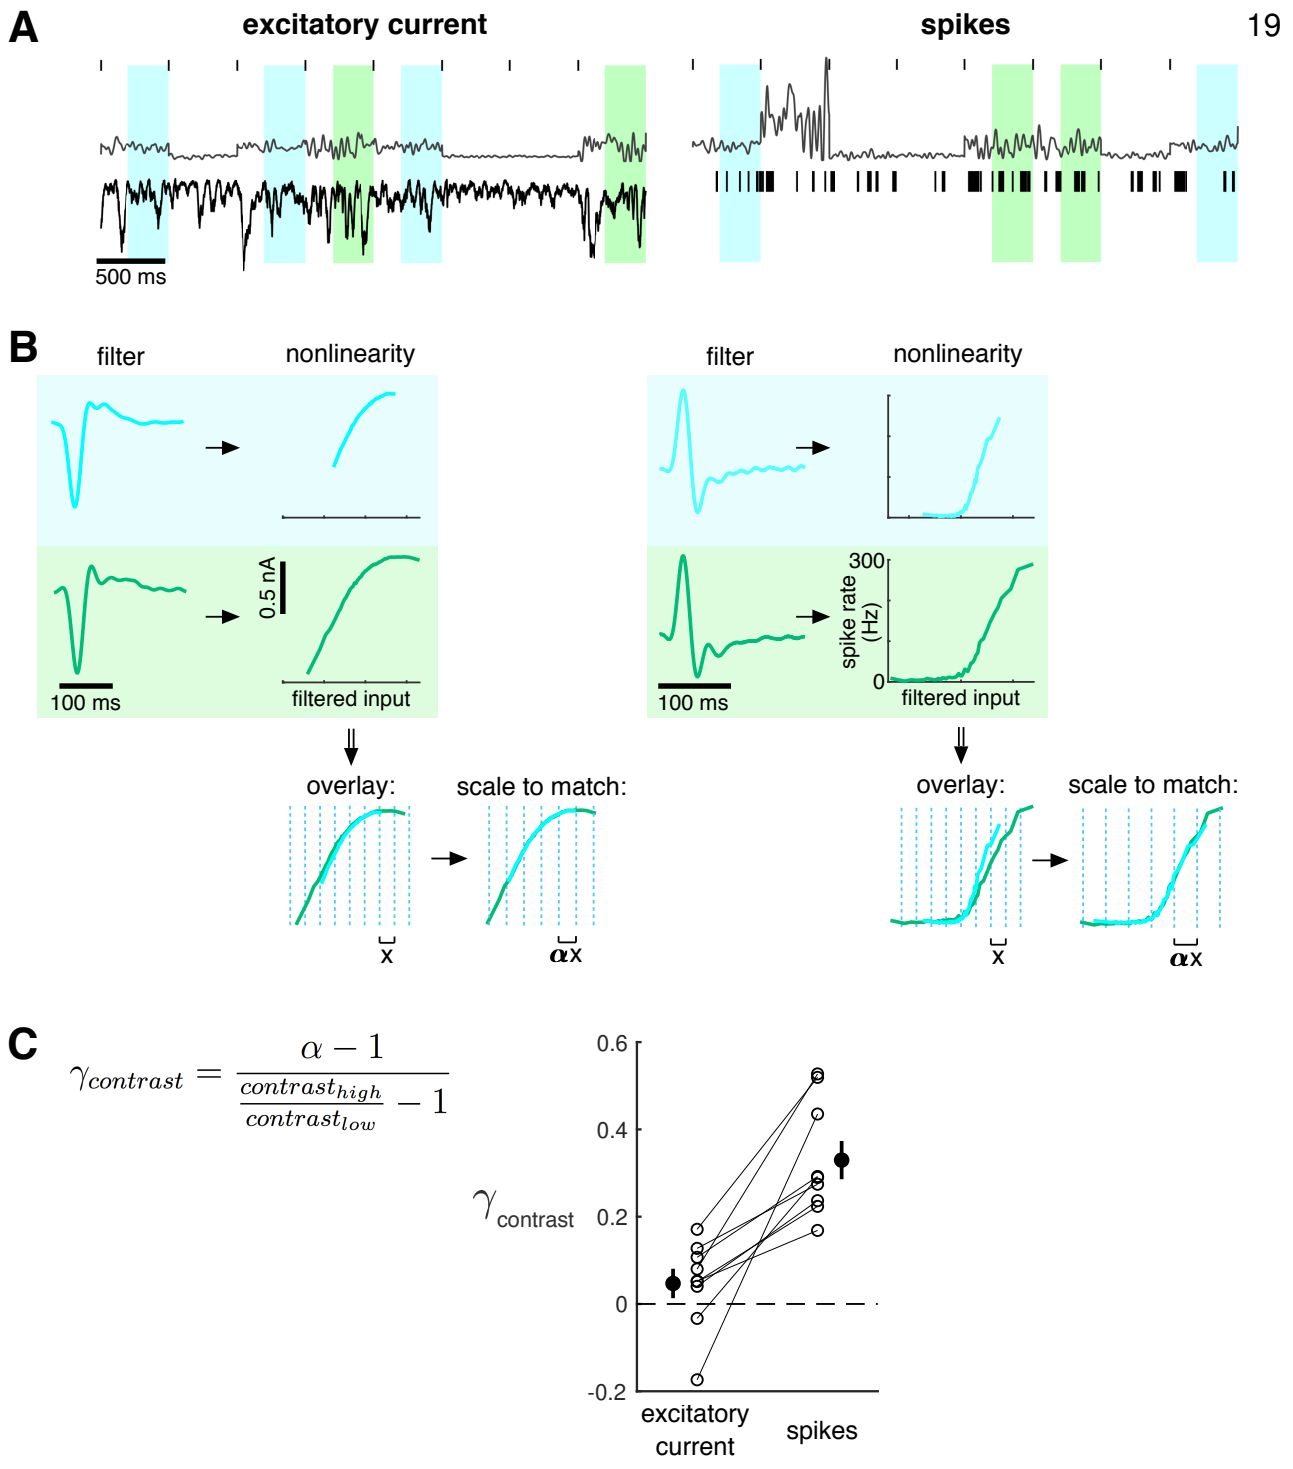

Figure 2.3: Temporal contrast-driven gain control measured in RGC responses. A: Variable noise stimuli and responses with shaded sections showing example data segments used to fit models at each of two stimulus contrasts. **B:** Revcor LN models fit to pooled data segments from each of two stimulus contrasts. Bottom: Example of process to measure difference in gain between models fit at two contrasts. **C:** $\gamma_{contrast}$ for 8 cells, computed using revcor LN models fit at two contrasts.

measured in spike output data (p = 0.13, paired-sample t-test), indicating that RGCs do not contribute a significant amount of mean adaptation on top of what is already present in the RGCs' excitatory input.

We made corresponding measuments to quantify response gain at different stimulus contrasts. Revcor LN models were fit using epochs pooled from two contrast conditions (20-25% and 50%) with the same mean (4000 R^*)(Fig. 2.3 A-B). Gain change factors α are measured to compute the γ index for contrast change:

$$\gamma_{contrast} = \frac{\alpha - 1}{\frac{contrast_{high}}{contrast_{low}} - 1} \tag{2.2}$$

In models fit using measured excitatory currents, there was little, if any, difference in slope between the response nonlinearities computed for the two contrast conditions, indicating an absence of contrast-driven gain control in excitatory inputs to recorded RGCs (Fig. 2.3 C). Accordingly, $\gamma_{contrast}$ measured in a collection of cells was not significantly different from 0 (p = 0.20). Models fit using measured spike responses, though, show significantly more gain control in RGC output (p = 8.5×10^{-4}), indicating that most or all contrast adaptation occurs during spike generation by the ganglion cell.

The absence of contrast adaptation in excitatory inputs to the RGC differs from results observed in other studies (Kim and Rieke, 2001; Manookin and Demb, 2006; Ozuysal and Baccus, 2012; Cui et al., 2016), and may stem from differences in primate retina (but see Chander and Chichilnisky (2001)) or between cell types, or from our use of comparatively short epochs. Epoch lengths have been shown to influence the speed and extent of adaptation (Wark et al., 2009), and, more generally, the dynamics of adaptation depend on stimulus history in ways that are poorly understood. Our epoch lengths were chosen to resemble the duration of visual fixations during natural eye movements around a scene, so our results represent mechanisms that are engaged in that context.

Under our experimental conditions, we find that mean and contrast adaptation occur in different circuit locations: most mean adaptation occurs upstream of ganglion cells while most contrast adaptation occurs in ganglion cells. In the next two sections, we subdivide our modeling approach along the same division. We first aim to predict excitatory input to RGCs given light stimuli and introduce a model that captures mean-driven gain control. Second, we aim to predict RGC spike output given excitatory input current and introduce a model that captures contrast-driven gain control. In the final section, all model components are combined in series to predict RGC spike output given light stimuli.

2.2.2 A biophysical cone model produces mean adaptation

Because we found that most mean adaptation and little or no contrast adaptation occurs upstream of the RGC, we sought to introduce mean adaptation into a predictive model of RGC excitatory input. Given that previous findings indicate that cones make significant contributions to mean adaptation (Dunn et al., 2007; Shapley and Enroth-Cugell, 1984), we built a multi-stage model of RGC excitatory input that explicitly includes a biophysical model of phototransduction in cones as a first stage. The cone model is presented in recent work (Angueyra et al., 2021), and is based on a set of differential equations (Pugh and Lamb, 1993; Rieke and Baylor, 1998b; van Hateren and Snippe, 2007) that represent the molecular interactions that comprise the phototransduction pathway. Angueyra et al. (2021) previously showed that the model accurately predicts cone responses to stimului designed to mimic the effect of large eye movements (saccades) that rapidly change the mean luminance (Fig. 2.4 C). Importantly, the cone model is used here without re-fitting or free parameters.

Given a light stimulus, the cone model simulates the photocurrent produced in phototransduction. This photocurrent can be passed to downstream model elements in place of a light stimulus. To predict excitatory current responses recorded in RGCs, we used a

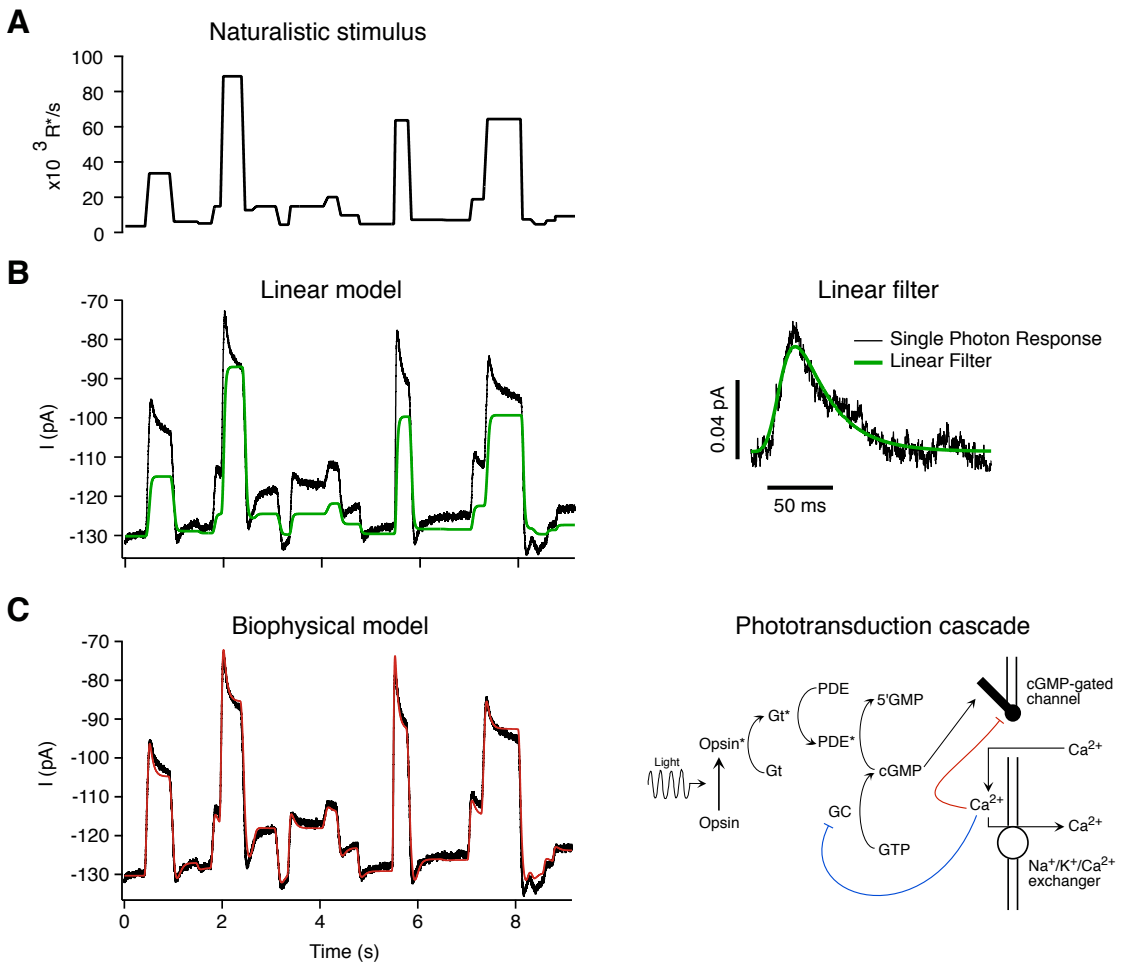

Figure 2.4: A biophysical cone model captures adaptational mechanisms in cone photocurrent responses. *Adapted from (Angueyra et al., 2021).* **A:** A stimulus with large changes in brightness mimics the effects of eye movements. **B:** Response predictions (left) of a linear model which is simply a linear filter (right) that captures the cone impulse response. **C:** Response predictions (left) of a biophysical model comprising differential equations that model the phototransduction cascade (schematized right).

multi-stage model comprising the biophysical cone model followed by an LN model. Unlike the light stimulus, the simulated cone photocurrent is not gaussian, so standard reverse correlation methods of fitting LN models are no longer valid (Sharpee, 2013). Instead, we use a parameterized LN model (linear filter and sigmoid nonlinearity, 9 free parameters total) fit to data using numerical optimization (see Methods).

We fit models with and without an initial cone phototransduction stage and compared downstream measures of mean-driven gain control. As an additional point of comparison, in one model we use a linear phototransduction stage that captures the linear kinetics of cone photocurrent but does not incorporate any nonlinear or adaptational aspects of cone responses (Fig. 2.4 B). This "linear cone" model is simply a linear filter – the impulse response of the full cone model. As a control, the LN model downstream of the linear cone is parameterized and optimized using the same method as is used downstream of the biophysical cone. We fit LN model components of each model type at each of two means and two contrasts. Fig. 2.5 A shows LN model components fit at each of two means (4000 and 8000 R^*/cone/second). Qualitatively, the slopes of the nonlinearities fit at the two means are more similar when input is first processed by the biophysical cone stage, as compared with model versions with no cone stage or a linear cone stage. Mean-driven gain control in phototransduction reduces amplitude differences in signals reaching downstream circuitry, and correspondingly the biophysical model reduces the need for gain adjustment in downstream model elements when the stimulus mean is changed.

γ_{mean} obtained in the linear cone + LN model case show relatively small deviations from those measured in the revcor LN model with no cone stage (Fig. 2.5 B, left; p = 3.6×10^{-2}) – both average $\sim$1 as expected if adhering to Weber's law. Inclusion of the biophysical cone model eliminates the majority of the x-axis scale difference between the nonlinearities fit at the two means as indicated by γ_{mean} falling close to 0 (p = 0.13, t-test against mean of 0). The full biophysical cone model, then, accounts for most or all of the

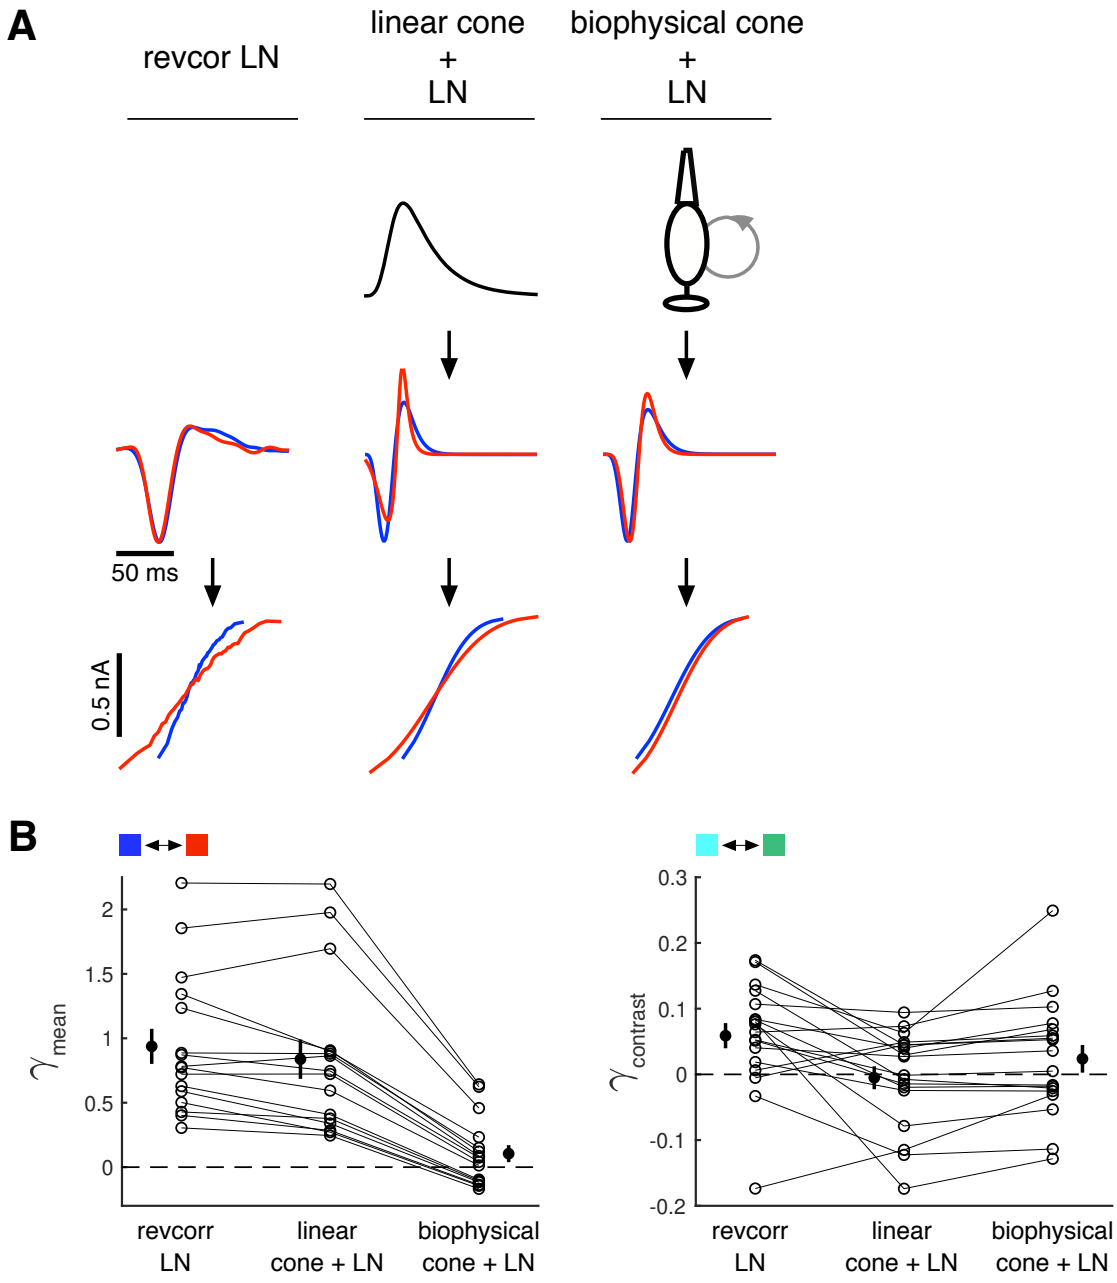

Figure 2.5: Biophysical cone model contributes mean-driven gain control to multi-stage model. A: Example LN model components in 3 model types, fit to responses at two means (4000 and 8000 R*/cone/second). **B:** γ_{mean} (left) and $\gamma_{contrast}$ (right) measured with nonlinearities from the same 3 model types as in (A).

mean adaptation observed in the recorded RGC excitatory inputs.

In the case of contrast adaptation, addition of the biophysical cone model has little effect on $\gamma_{contrast}$ relative to revcor LN models with no cone (p = 0.13) or LN models with a linear cone preprocessing stage (p = 1.7×10^{-2}) (Fig. 2.5 B, right). This is expected since we observed little or no contrast adaptation in RGC excitatory currents as reported above (Fig. 2.3 C).

Because of the high-fidelity representation of adaptational dynamics of phototransduction provided by the cone model, our findings illustrate that cones are the majority contributor of mean adaptation present in On-parasol RGC excitatory input. Taken together with our finding that most mean adaptation is already present in excitatory input with only minor (if any) contributons coming from spike generation, our results indicate that cones are also the majority contributor of mean adaptation seen in On-parasol spike output under the tested conditions. This is consistent with previous results in primate retina (Dunn et al., 2007).

2.2.3 A spike generation model produces contrast adaptation

In Fig. 2.3 we showed that most or all contrast adaptation occurs within the RGC in the tested conditions. One possible mechanistic origin of contrast adaptation in RGCs is spike history dependence in spike generation (Kim and Rieke, 2003). This refers to the fact that the dynamics of spike generation depend not only on the input received by the RGC but also on the history of activity in the cell as summarized by its spike history. Spike history dependence constitutes a time-dependent nonlinearity in RGCs' responses.

We sought to determine whether contrast adaptation could be captured using a simple model of spike generation with spike history dependence. We built a spike generation model that predicts spike times given an input current. The model comprises an LN model to predict subthreshold voltage, a spike threshold, and a post-spike filter (Fig. 2.6). The

post-spike filter introduces refractoriness by modulating the spike threshold via a 2 ms absolute refractory period and a single decaying exponential. Refractoriness in spiking neurons can produce adaptation by limiting the extent to which spike rate can scale with activation strength.

Like the cone model, the spike generation model was fit using a separate set of specialized experiments. To characterize the transformation from excitatory current to spikes in RGCs, we injected a noise current stimulus with variable mean and contrast (similar to the light stimulus described previously) and recorded RGC voltage responses. From these responses we used subthreshold voltage to fit the LN model components and used spike times to fit the baseline threshold and post-spike filter. Fig. 2.6 C illustrates isolation of subthreshold voltage with a low-pass filter. After fitting all parts of the RGC model with numerical optimization based on these injected current experiments, a single instantiation of the spike generation model was employed across all subsequent analyses (see Methods). Fig. 2.6 D shows the spike generation model's predicted spikes alongside recorded spikes from multiple trials of a repeated noise stimulus.

To test the spike generation model in the context of light stimuli while still isolating our analysis to RGC contributions, we used light-evoked excitatory currents recorded in RGCs as input to the spike generation model. Excitatory current has been shown to be the principal driver of On-parasol RGC spike responses to uniform stimuli, with inhibitory current having little impact (Cafaro and Rieke, 2013). After generating spike predictions using excitatory currents recorded in multiple stimulus conditions, we paired spike predictions with the associated light stimuli and fit revcor LN models for each stimulus condition (Fig. 2.7 A for contrast change example). These revcor LN models describe the relationship between the visual stimulus and the predicted spikes in each stimulus mean + contrast condition.

Fig. 2.7 B shows γ_{mean} and $\gamma_{contrast}$ derived from revcor LN models fit on predicted spikes compared to those fit on recorded excitatory currents and recorded spikes. If

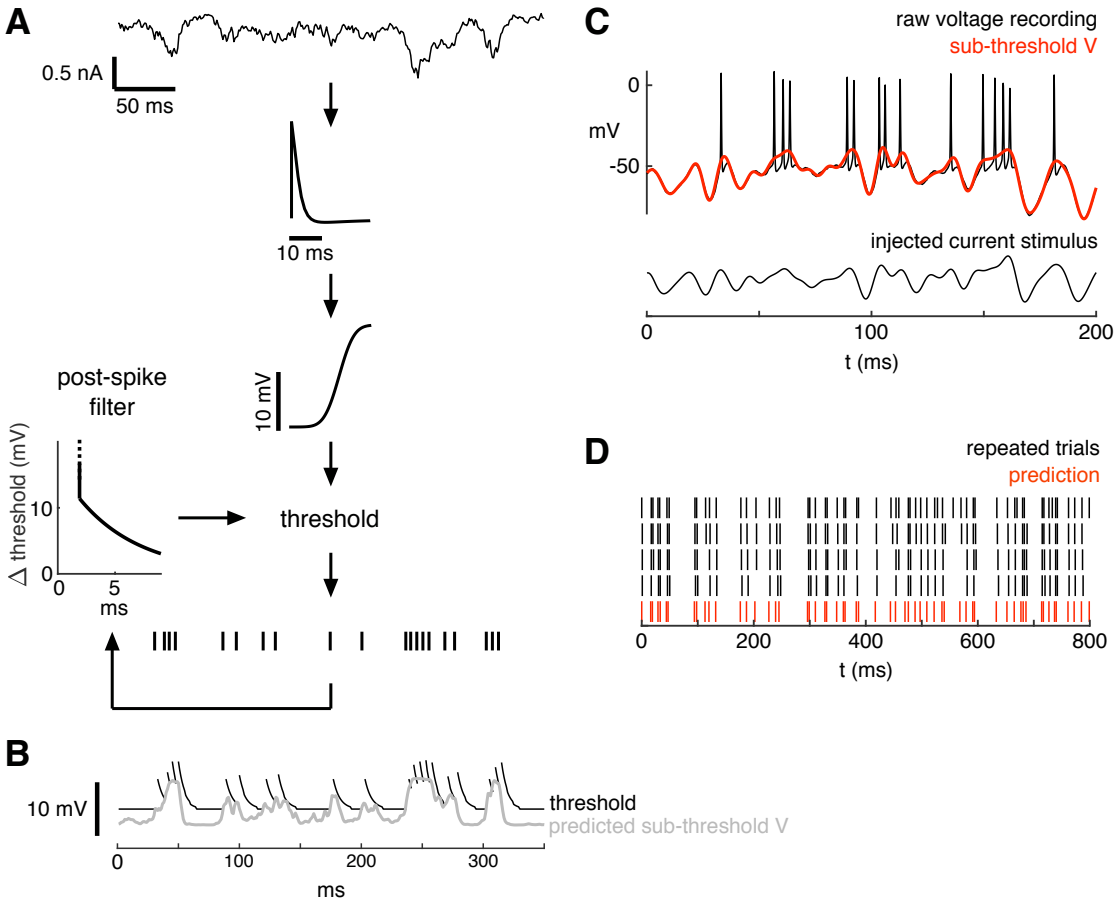

Figure 2.6: Spike generation model. A: Schematic of spike generation model. Input current is passed through a linear filter and nonlinearity, resulting in a predicted subthreshold voltage. Spike times are determined by a threshold that is modulated by a post-spike filter, which incorporates a 2 ms absolute refractory period followed by an exponential decay back to baseline threshold. **B:** Time series of threshold and predicted subthreshold voltage during model run shown in (A). Breaks in the threshold trace occur during absolute refractory periods during which the threshold is infinite. **C:** Example of injected current stimulus and response, and separation of spikes and subthreshold voltage by low-pass filtering. Subthreshold voltage (red) is used to fit the initial filter and nonlinearity in the spike generation model shown in (A). Spike times are used to fit the baseline threshold and post-spike filter shown in (A). **D:** Recorded spike times in response to a repeated injected noise stimulus segment, and the corresponding model prediction.

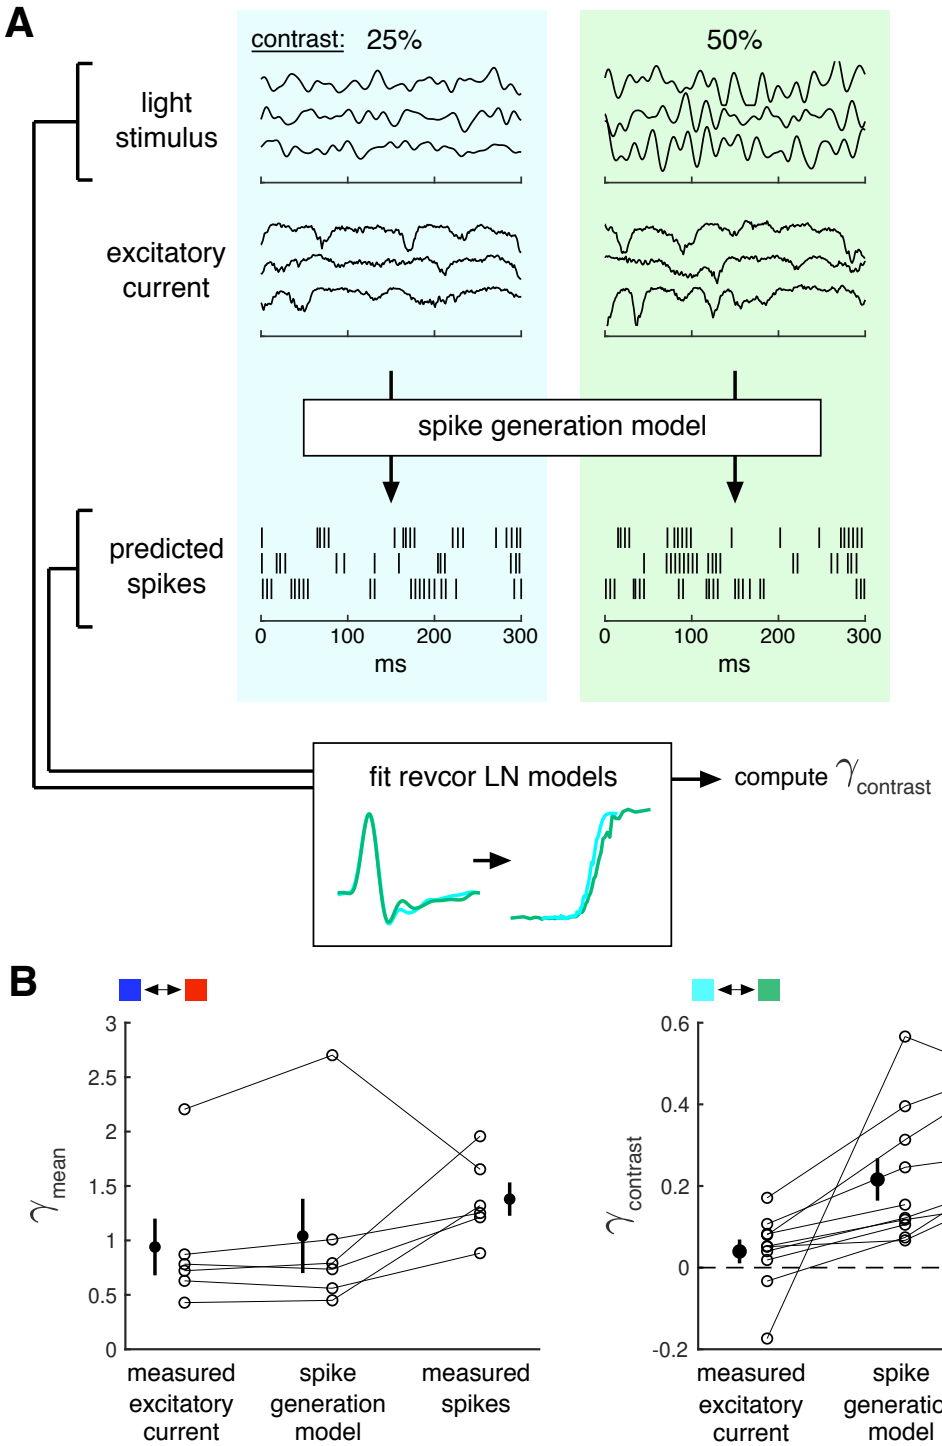

Figure 2.7: Spike generation model recapitulates contrast adaptation seen in recorded spikes. A: Method schematic. Spike generation model produces spike predictions given measured excitatory current at two stimulus contrasts. Spike predictions are reverse correlated with associated light stimuli to form LN models which are used to compute $\gamma_{contrast}$. **B:** γ_{mean} and $\gamma_{contrast}$ measured with nonlinearities derived from the spike generation model compared with those derived from measured excitatory current and spikes.

the spike generation model captures ganglion cell adaptational behavior, it is expected that γ values will be similar between the predicted and recorded spikes. With changes in temporal contrast, the RGC model produces contrast adaptation that is absent in the measured excitatory currents (Fig. 2.7 B, right; p = 2.6 $\times$ 10^{-2}). Further, RGC model spike responses recapitulate the majority of contrast adaptation observed in the measured spike responses, though measured spikes appear to show stronger signatures of adaptation (p = 4.2 $\times$ 10^{-2}). With changes in mean luminance (Fig. 2.7 B, left), γ_{mean} values are similar between predicted spikes and measured spikes (p = 0.33), as well as measured excitatory input (p = 0.28). This is expected given our result that spike generation does not seem to make significant contributions to mean adaptation, but is useful confirmation that the spike generation model does not deviate substantially from the mean adaptation behavior observed in measured spike responses.

Taken together, our recordings and modeling approach have so far shown that under the tested conditions mean adaptation stems principally from nonlinear processes of phototransduction, while contrast adaptation stems principally from history dependence in spike generation. A biophysical model of phototransduction is sufficient to capture most or all of the difference in RGC response gain at different mean light levels, while a simple model of spike generation that incorporates spike history dependence is sufficient to capture most of the difference in RGC response gain at different temporal contrasts.

2.2.4 Combining biophysical and statistical model components

We combined the cone model, an LN model, and the spike generation model in series (Fig. 2.8) and tested the resulting multi-stage model across a range of stimulus means and contrasts. We fit a single multi-stage model across all stimulus conditions to test the hypothesis that this model could dynamically produce mean- and contrast-driven gain control.

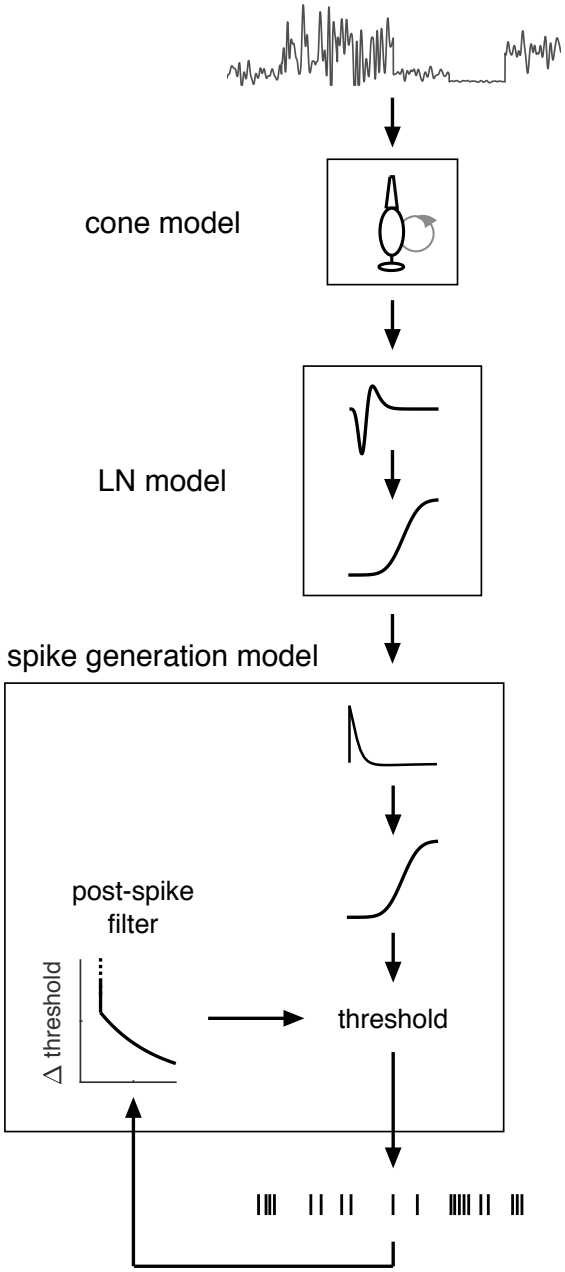

Figure 2.8: Multi-stage model. Sequential model incorporating the biophyisical cone model, an LN model, and the spike generation model.

Models were fit using excitatory currents and spike responses recorded in the same cells: the LN model (middle stage) was fit using only measured excitatory currents, and a single free scaling constant applied between the LN model and the spiking model was fit with measured spike output (see Methods). This scaling factor is needed in each cell to bring the dynamic range of predicted excitatory current into register with the effective input range of the spiking model, which is based on injected current experiments performed on a different set of cells. Altogether there are 10 free parameters in the multi-stage model: 9 to define the parameterized LN model (middle stage) and 1 to set the current scaling constant.

We fit multi-stage models on data collected from a subset of On-parasol RGCs (n=3) in which we collected additional trials using a repeated noise seed, allowing us to benchmark model performance against inter-trial variability in the recorded spike responses. In repeated noise seed experiments, noise mean and contrast were modulated from epoch to epoch as before – only the noise seed was held constant. Fig. 2.9 A shows measured (black) and predicted (red) spike responses to the repeated noise seed stimulus in each of 6 stimulus conditions. The multi-stage model is deterministic, so only one response is needed per condition to completely represent the model's response.

We used the Victor-Purpura spike distance metric (Victor and Purpura, 1996, 1997), a measure of similarity between two spike trains, to quantify the average spike distance between the predicted spike train and the measured spike trains, as well as the average spike distance between measured trains. We define "explainable spike distance" such that capturing 100% of explainable spike distance means the predicted train is on average as similar to the recorded trains as the recorded trains are to each other in repeated trials. Capturing 0% means the predicted train is on average only as similar to the recorded trains as is a random train with the same number of spikes. Fig. 2.9 B (left panel) shows the fraction of explainable spike distance captured by the full multi-stage model in each

of 6 stimulus conditions, plotted for 3 cells. For each cell, a single multi-stage model is fit simultaneously to data from all conditions (a non-repeating random noise seed stimulus was used to collect data for fitting). For 2 of 3 cells, the multi-stage model captures most explainable spike distance across all stimulus means and contrasts, averaging 84% and 86% captured across conditions. In the remaining cell, the multi-stage model performed similarly at 50% contrast conditions but failed to capture most explainable spike distance at 25% contrast, leading to an average of 51% captured. This may reflect cell-to-cell differences in the extent of contrast adaptation, or it may be a symptom of not capturing all contrast adaptation with our spike generation model. Or it could be that the health of this cell declined over the course of recording (we observed this cell's strength of response modulation was considerably less than the other two cells during the repeated noise experiment (data not shown)).

As a point of comparison, we fit multi-stage models separately in each individual condition (Fig. 2.9 B, right), which removes any need for the model to balance its performance across conditions. In this case, models from all three cells captured the majority of explainable spike distance in each condition, with an exception in which one model erroneously predicted zero spikes in one condition. Therefore, in the case of the cell whose cross-condition model performed poorly, the failure is associated with balancing performance across contrasts, not with challenges intrinsic to any individual condition. In the other two cells, the performance of the model fit across all conditions nearly matches the performance of models fit specifically to each condition.

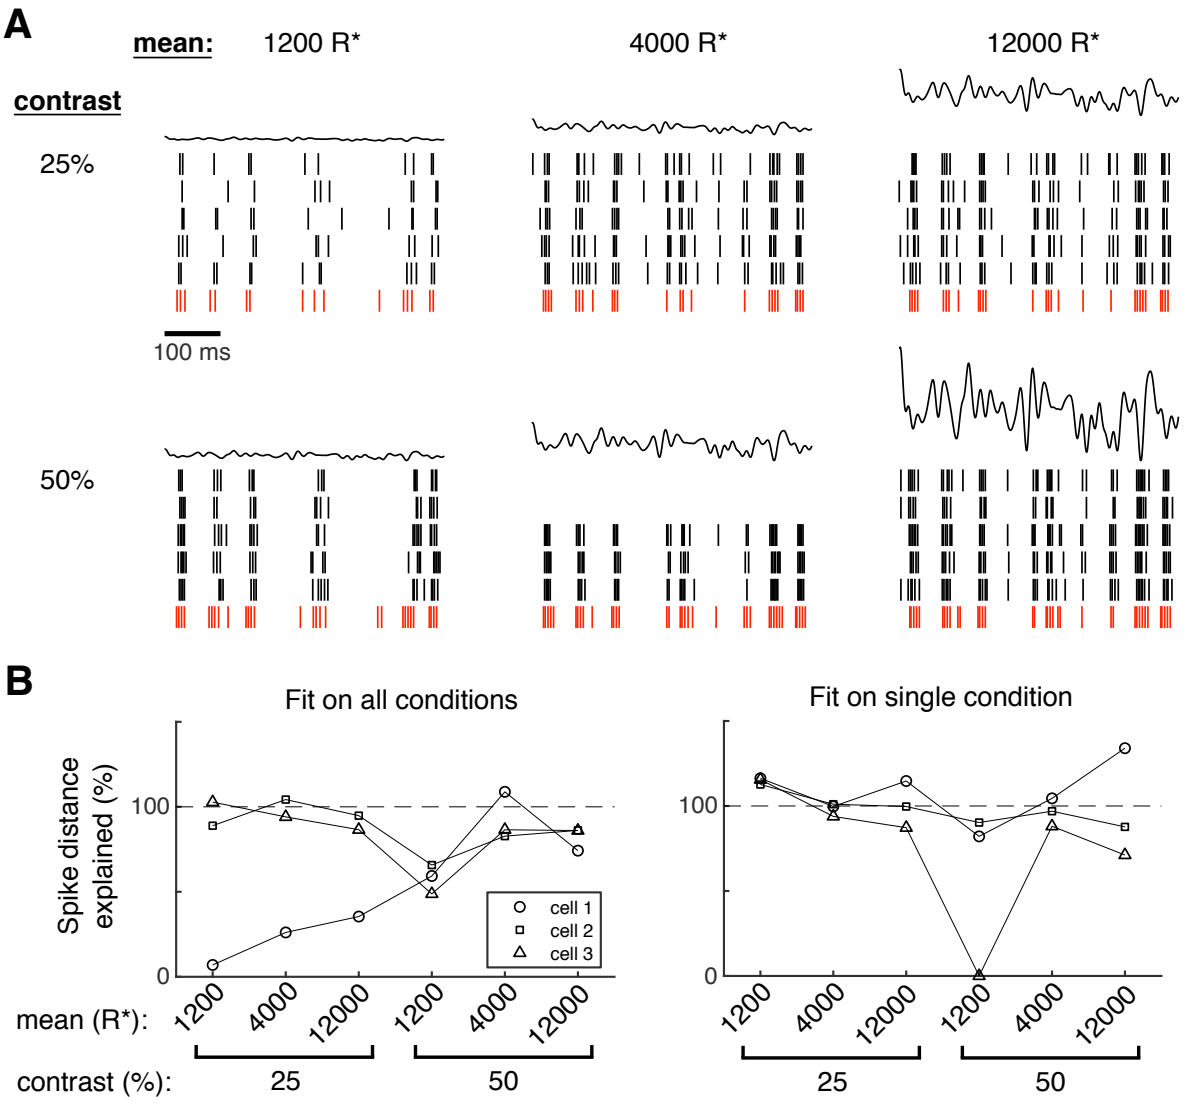

Figure 2.9: Multi-stage model performance across stimulus conditions. A: Example responses to repeated noise seed stimuli pooled for each stimulus condition (black) alongside spike response predictions from a single multi-stage model (see Fig. 2.8), which was fit on data from all conditions. Recordings and model are associated with cell 3 referenced in (B), left. **B:** Portion of explainable spike distance captured by multi-stage models fit simultaneously on all stimulus conditions (left) or separately on individual stimulus conditions (right). 0% spike distance explained connotes spike distances equivalent to a random spike train (with the same number of spikes), while 100% connotes spike distances as low as inter-trial variability in recorded responses to repeated stimuli.

2.3 Discussion

The statistics of natural scenes and of eye movements interact to produce large, rapid, and frequent changes in stimulus statistics that engage adaptational mechanisms at multiple sites in the retina. Relatively little is known about how these mechanisms together maintain visual sensitivity at timescales relevant to eye movements, even as much is known about adaptation at timescales of several seconds or longer. Our focus was on incorporating explicit representations of known adaptational mechanisms in multiple sites of a multi-stage predictive model of RGC responses. We chose a stimulus that enabled us to observe adaptation to both mean and contrast on timescales relevant to saccadic eye movements while retaining access to analytical tools that rely on gaussian-distributed white noise stimuli. In addition to measuring correlates of adaptation by fitting models separately in each of multiple conditions, we challenged our final multi-stage model to predict RGC responses across a range of stimulus means and contrasts. In two of three cells tested, the model successfully maintains prediction accuracy across the range of stimulus conditions tested. We consider only a single RGC type (albeit one of the most common types), and apart from the cone model we use deliberately simple model components that could readily be elaborated for different applications including other retinal circuits.

To our knowledge, there are no other existing models that predict RGC responses across variations in both mean and contrast, with the possible exception of deep learning models currently in development. At the very least, deep learning models will soon have this capability, but these models lack a mechanistic connection to retinal circuitry, making them difficult to interpret in terms of circuit function. They also have far more free parameters, and as a result require much more data to train. This demand for training data may be prohibitive for some applications.

One might worry that we too have started down a road to arbitrarily complex models

that are increasingly difficult to fit to data. But our full multi-stage model has only one more free parameter than the LN model at its center (see Methods). This is because both the cone model and the spike generation model are fit using separate specialized experiments, after which these models can be used off-the-shelf in other settings without re-fitting.

Stimulus design and timescales of adaptation

The stimulus chosen here differs from those used in most previous studies investigating adaptation in two important ways: 1) both mean and contrast are varied simultaneously and 2) there is a greater frequency of changes in the stimulus statistics. In certain regards, our results also differ from previous results in a manner that may be explained by differences in stimulus design. As in previous work on mean adaptation (Dunn et al., 2007), we find cones to be the primary site of mean adaptation in daylight vision. But unlike some previous studies on contrast adaptation (Kim and Rieke, 2001; Manookin and Demb, 2006; Ozuysal and Baccus, 2012; Cui et al., 2016), we do not observe significant contrast adaptation outside of the ganglion cells studied. Previous work has pointed to contrast adaptation in multiple circuit loci, most notably originating in bipolar cell output synapses and ganglion cells. Our recordings of excitatory input to ganglion cells showed little if any evidence of contrast adaptation when we fit revcor LN models at two contrasts (same mean).

Species or cell type differences may play a role in this discrepancy (these previous studies were conducted in non-primate retina), but the disparity may also be explained by a difference in stimulus switching period. In the studies in question, stimulus statistics were held constant for tens of seconds, compared to 0.5 to 4 s in our study. It may be the case, then, that the bipolar cell-mediated adaptation observed in those studies may occur too slowly to appear in our experiment. More generally, characteristics of adaptation

exhibit a long stimulus history dependence including dependencies on switching period (Fairhall et al., 2001; Wark et al., 2009). It may be difficult, then, to compare studies in which very different switching periods are used. Given that the dynamics of adaptation are subject to long-range stimulus dependencies, it will be important to develop experiments that apply increasingly naturalistic stimuli to understand adaptation during natural vision.

Relationship to existing adaptational models

No fully dynamical models of an entire retinal circuit currently exist as constraining models with such biophysical fidelity is especially demanding. The current generation of predictive models of retinal output therefore all rely in part on statistical components.

Several studies have offered extensions of the LN model framework that make these models to produce some form of adaptation. Ozuysal and Baccus (2012) introduced a linear-nonlinear-kinetic (LNK) model, which adds a dynamical stage that allows the model to capture adaptive response properties associated with changes in contrast. The kinetic module is amenable to mechanistic interpretation since the kinetic states represent states of synaptic transmission such as ion channel conformation and the pool of available neurotransmitter at the bipolar cell synaptic terminal. The LNK model does not include adaptation in cones, and the study does not include changes in stimulus mean. We also do not find synaptic mechanisms to play a significant role in near-saccade-timescale contrast adaptation in primate On-parasol RGCs (the LNK model study used a contrast switching period of 20 s, and investigated RGCs in salamander retina).

Other existing models forgo dynamical components but use additional statistical components to build in specific adaptational behaviors. Cui et al. (2016) showed that a divisive interaction between two LN branches, called divisive suppression (DivS), successfully captures contrast adaptation in an experiment with a 10 s contrast switching period in mouse retina.

Both the LNK and DivS models employ linear initial stages, and the studies introducing these models do not aim to capture mean adaptation. Since we and others have found that mean adaptation is carried out principally by cones, adding a dynamical cone model as an initial stage to the LNK and DivS models may allow them to predict responses across means as well as contrasts. LNK and DivS models offer more sophisticated representations of computation in the interneuron layer of retina than the simple LN model offered in this chapter, so incorporating these or similar models in a multi-stage model like the one we present in this work may lead to better performance than the version we present. Similar extensions of the middle stages are likely to be important, for example, in studies that emphasize the role of lateral interactions through horizontal and amacrine cells, which we have at least partially avoided here by using a spatially uniform stimulus. Clearly there are also some combinations of timescales, species, cell types, and other factors in which interneurons contribute to contrast adaptation, so the architecture of the middle model layer should be chosen accordingly. But the ubiquity of linear first stages in our current generation of predictive models has become limiting as we turn toward more naturalistic stimuli.

The simplicity of our spike generation model was helpful in attributing its contrast-adaptational behavior to a single component in the model (i.e. the post-spike filter), but other spiking models could be used, including a generalized integrate-and-fire model (GIF) (Mensi et al., 2016) or a generalized linear model (GLM) (Pillow et al., 2008).

Spatial and temporal naturalistic stimuli

In this work we use a spatially uniform stimulus to limit model complexity and lateral interactions in retinal circuitry. But lateral interactions and spatial nonlinearities are influential determinants of retinal output, are critical for a complete understanding of retinal function and mechanism, and provide a useful setting in which to study circuit motifs that are

present throughout the nervous system. Subunit models have been used to understand nonlinear spatial integration (Kuo et al., 2016; Turner and Rieke, 2016; Freeman et al., 2015; Maheswaranathan et al., 2018) and computations mediated by lateral interaction such as motion (Appleby and Manookin, 2020). Future studies using stimuli that are naturalistic in space as well as time should develop models that combine multiple nonlinear vertical stages (like those presented here) with subunit architectures that capture spatial effects.

We designed our stimulus to resemble naturalistic temporal stimuli in some key ways, but our stimulus is still far from what would be experienced by a cone during natural vision. Setting aside the fact that spatial stimuli are not considered, we sample only a small number of means and contrasts and switch them at a constant frequency. During natural vision, fixation lengths as well as mean and contrast within fixations all vary over a wide range, and stimulus parameters need not stay constant throughout a fixation. Future studies should use naturalistic stimuli derived from images or movies with eye-tracking data.

Selective partitioning of mechanisms within a model

Many statistical models of neural computation can be fit with little or no knowledge of neural mechanisms underlying the computation. The relative ease of fitting and interpreting these models has made them valuable tools no matter the extent of mechanistic detail understood in any particular system. On the other hand, models that explicitly incorporate known mechanisms may predict responses over a greater range of conditions by exhibiting emergent behaviors not yet achieved by existing phenomenological models. But mechanistic models are more difficult to create because they rely on sometimes very detailed characterization of mechanisms at play in the circuitry. We took a hybrid approach, starting with flexible phenomenological components and selectively incorporating explicit

representations of mechanisms known to make important contributions. This approach disentangles the contributions of those known mechanisms from all the rest. The result is both a more powerful predictive model and a more valuable tool to understand the mechanistic basis of function. By partitioning specific mechanisms within a model, one can more easily identify the roles of those mechanisms and, equally, the roles of remaining mechanisms that were not explicitly modeled. This can in turn guide further experimental work seeking to illuminate those least understood mechanisms within a circuit.

2.4 Methods

Tissue Preparation

Retinal tissue from macaque monkeys (*M. nemestrina*, *M. mulatta*, or *M. fascicularis*) was obtained via the Tissue Distribution Program at the Washington National Primate Research Center. All procedures were approved by the Institutional Animal Care and Use Committee at the University of Washington. After enucleation, the eye was hemisected and the vitreous humor was removed mechanically. Tissue was prepared as described previously (Turner and Rieke, 2016). Briefly, dark-adapted (>1 hour) retina in warm ($\sim$32° C), oxygenated Ames solution was placed photoreceptor side down on a poly-D-lysine-coated coverslip (BD biosciences) and transferred to the recording dish. Tissue was continuously perfused with oxygenated solution warmed to $\sim$32° C and visualized using infrared illumination.

Patch-clamp recordings

RGC excitatory current and spike recordings were performed as described previously (Turner and Rieke, 2016). Excitatory currents were measured using whole-cell voltage clamp. Spikes were collected separately using extracellular or cell-attached recording. Data were acquired using a Multiclamp 700B amplifier, low-pass filtered at 3 kHz, and digitized at 10 kHz. All stimulus presentation and data acquisition were controlled using MATLAB based, open-source Symphony Data Acquisition Software (https://github.com/symphony-das).

Injected current experiments were conducted in current clamp with 10 μM NBQX added to the superfusion and an internal solution containing (in mM): 123 K-aspartate, 10 KCl, 10 HEPES, 1 MgCl2, 1 CaCl2, 2 EGTA, 4 Mg-ATP, and 0.5 Tris-GTP.

Light stimulation

Light stimuli were presented using LEDs which delivered full-field illumination covering a $\sim$500 μm disk centered on the recorded cell. Stimulus protocols were generated using custom MATLAB-based extensions of Symphony Data Acquisition System and delivered at 10 kHz.

Data analysis and modeling

Data analysis was performed using custom software developed in MATLAB (Mathworks).

Linear-nonlinear (LN) cascade models

Reverse correlation LN models: Reverse correlation-based LN (revcor LN) models were fit using reverse correlation of measured or predicted responses with gaussian white noise stimuli in a standard procedure (Chichilnisky, 2001). In computing filters, the DC component is discarded (the mean is subtracted). In computing nonlinearities, the generator signal was binned such that an equal number of points were averaged in each bin.

Parameterized LN models: Revcor LN models are not appropriate when model input is non-gaussian (Sharpee, 2013). We therefore parameterized an LN model which could be optimized numerically. Filters were parameterized using the following equation (modified from (Angueyra and Rieke, 2013; Baylor et al., 1984)):

$$f(t) = \alpha \cdot \frac{\left(\frac{t}{|\tau_{rise}|}\right)^k}{1+\left(\frac{t}{|\tau_{rise}|}\right)^k} \cdot e^{-\left(\frac{t}{\tau_{decay}}\right)} \cdot cos\left(\frac{2\pi t}{\tau_{osc}} + \phi\right) \tag{2.3}$$

Nonlinearities were parameterized as a sigmoidal curve based on a cumulative Gaussian function:

$$N(x) = aC(bx + g) \tag{2.4}$$

where $C()$ is the cumulative normal distribution (*normcdf* in MATLAB), g is an offset along the horizontal axis, b determines the slope, and a is a scale factor which determines response magnitude (Chichilnisky, 2001).

Measuring gain change (γ) index

The gain change index γ is defined in Eqn. 2.1 (γ_{mean}) and Eqn. 2.2 ($\gamma_{contrast}$) and is used to quantify the difference in response gain in two stimulus conditions using LN models fit in each of the two conditions. We optimize the x-axis scaling factor α which applies a horizontal stretch to one nonlinearity that best matches it to the other nonlinearity according to a least squares objective function. This method allows us to measure the change in gain between two conditions without fitting a functional form (such as a sigmoid function) to each nonlinearity. γ is the ratio of the fractional change in gain to the fractional change in stimulus mean (for γ_{mean}) or contrast (for $\gamma_{contrast}$).

In an LN model the x-axis scaling of the nonlinearity is tied to the shape of the upstream filter since the filter determines the amplitude of the generator signal (the input to the nonlinearity). When comparing LN models across stimulus conditions, we normalized filters in the two LN models such that a noise signal passed through both filters would produce generator signals with equal amplitude. This consolidates gain-related model features in the nonlinearity, allowing for compact comparison of gain in two LN models without comparing filter shapes.

To compute the least squares objective function between two nonlinearities across arbitrary x-axis scalings, we use interpolation between points of each nonlinearity and sum squared differences within the range of overlap between the two nonlinearities. Free

parameters for vertical and horizontal offset are also introduced to ensure the best match between the slopes of the two nonlinearities.

Cone model

The cone model was developed previously (Angueyra et al., 2021) and is based on a set of differential equations used to describe phototransduction (Pugh and Lamb, 1993; Rieke and Baylor, 1998b; van Hateren and Snippe, 2007) with the addition of a second, slower Ca^{2+} feedback.

Because of the kinetics of phototransduction, cones convey a filtered, slower version of the light signal onto second-order cells, and likewise the cone model changes the kinetics of the signal available to the downstream LN model. Part of the effect of the cone model, then, can be captured by a linear filter, and this is in fact a good point of comparison to isolate the nonlinear and adaptational effects captured by the full cone model. Therefore alongside the biophysical cone model we tested a linear cone model which is simply a parameterized filter (Eqn. 2.3) fit to the biophysical cone model's impuse response (Fig. 2.4 B, right).

Placing the linear cone model in series with an LN model puts two linear filters in series. In the case of reverse correlation-based filters, two filters in series can be non-destructively combined into a single filter, but in the case of our parameteried filters this is not a guarantee. Including the linear cone model, then, is a useful control that introduces kinetic aspects of the full biophysical cone model without any of the nonlinear or adaptational contributions.

Spike generation model

The spike generation model comprises two parts: 1) a parameterized LN model that predicts subthreshold voltage (V) given an input current and 2) a spike generator composed

of a dynamic threshold modulated by a post-spike filter (Fig. 2.6). Both parts of the model were fit using current injection experiments performed on a separate group of On-parasol RGCs. In these experiments, we recorded subthreshold V and spikes while injecting a current stimulus. To fit the spike generation model, spike times were detected and stored before removing spikes from the V recording by low-pass filtering. The remaining signal is an approximation of subthreshold V, which was used to fit the LN model portion of the model (9 free parameters). With the LN model portion in hand, we generated subthreshold V predictions for all injected current stimuli and used this prediction, alongside stored spike times, to fit parameters of the spike generator. The spike generator's free parameters are the baseline threshold and two parameters which define a decaying exponential component of the post-spike filter, which is composed solely of this decaying exponential and an absolute refractory period. The absolute refractory component of the post-spike filter was set by hand to 2 ms. In the cells tested in injected current experiments (n = 6), the minimum inter-spike intervals were 1.8 - 2.1 ms, and we round to the nearest 1 ms since this is the time step size of our model. Functionally, then, there is only a single time step following a spike wherein an additional spike is disallowed.

We fit all components of the spike generation model in each of 6 cells. To choose a single spike generation model to use in other analyses, we tested each cell's best-fit model on data collected on all other cells and chose the model that performed the best overall. The same model performed best in terms of minimum average spike distance across data sets with and without normalization by spike count.

To use the spike generation model in the context of light-evoked responses, the model is placed downstream of recorded excitatory currents or excitatory currents predicted by other model components. To account for scale differences between light-evoked currents and the injected currents used to fit the spike generation model, a free scaling parameter was incorporated to scale the input to the spike generation model. This parameter was

fit for each use of the spike generation model using a numerical optimization minimizing Victor spike distance between measured and predicted spikes. In instances in which response gain was compared between conditions (i.e. in calculating γ_{mean} or $\gamma_{contrast}$), the same input scaling was used in both conditions (the parameter was fit once at the lower mean or contrast condition and then used in both conditions).

Optimization

Numerical optimization was performed using MATLAB's `fminsearch` function (Nelder-Mead method). In particular, we used John D'Errico's open-source modification, `fminsearchbnd`, to impose constraints disallowing impossible parameter regions. Because optimizations in this work are not guaranteed to have a unique solution, each optimization was initialized at 5 sets of initial conditions before choosing the best fit.

Chapter 3

SYNAPTIC SPECIALIZATION AT THE VISUAL SYSTEM'S FIRST SYNAPSE

3.1 Collaborators

This work was carried out in close collaboration with Takeshi Yoshimatsu, with additional scientific contributions from Owen Lawrence, Ayana Hellevik, and Stephen Neuhauss. I conducted physiology experiments while collaborators conducted imaging and morphological experiments. Study design and analysis of results in both domains were carried out jointly with Takeshi Yoshimatsu.

3.2 Introduction

It is common to conceive of neurons as the most basic units of computation in the nervous system, but in fact sub-cellular structures often function as important computational loci. Synapses are an important example: the site of signal transfer between cells presents unique opportunities to process and modify those signals, and different synapses onto a single cell can act independently in a kind of intra-neuron parallel processing (Asari and Meister, 2012; Baden et al., 2014). Synapse-specific specializations impart a functional richness to this parallel processing that remains largely unexplored. Synaptic specialization can diversify the influence a single cell has on multiple post-synaptic partners in a divergent circuit motif, or similarly diversify the computations performed on inputs from multiple pre-synaptic partners in a convergent circuit motif (Chabrol et al., 2015; Asari and Meister, 2012). Understanding how synaptic diversity contributes to the computational ca-

pacity of circuits, then, is an indispensable step toward understanding ubiquitous neural strategies like parallel processing and multiplexing of multiple signals within shared circuitry. More generally, understanding specialization in synaptic processing is fundamental to explaining the flow of information through neural circuits.

Research in this area is in very early stages, largely because it is usually impractical to isolate and characterize separate signals conveyed to different populations or classes of synapses, especially under physiological conditions. The retina, though, presents some exceptional opportunities. In the retina, signals from multiple photoreceptor types traverse the retinal circuit in parallel and are mixed and processed by many downstream cell types. Each photoreceptor type is functionally distinct and multiple types often converge onto single cells downstream (Li et al., 2012). Photoreceptors themselves are largely free from the influence of upstream neurons; as primary receptors, photoreceptors are driven mainly by light rather than potentially confounding input from other neurons. Importantly, each photoreceptor type exhibits unique kinetics and noise properties, meaning convergence of these types onto downstream cells also entails convergence of signals with different signal- and noise-spectral signatures.

The retina also offers some superlative experimental conveniences. Photoreceptors are readily accessed for recording in isolated retina, and their signaling can be controlled using light stimuli which preserve physiological signal properties like response kinetics and noise.

We investigated synapse specialization in zebrafish retinal bipolar cells (BCs) that receive convergent input from rod and cone photoreceptors. BCs, along with other cell classes, process photoreceptor signals and convey them to retinal ganglion cells (RGCs), the output cells of the retina. Each bipolar cell type has stereotypic connectivity, often receiving convergent input from two or more photoreceptor types and providing divergent output onto several RGC types (Hoon et al., 2014; Euler et al., 2014). Photoreceptor

synapses onto BCs are the site of a critical processing step in which relevant signal is separated from noise. Studies of the salamander rod-to-rod-BC synapse revealed that this synapse filters rod input to attenuate high temporal frequencies dominated by noise and to attenuate low frequencies that degrade temporal accuracy of the signal (Bialek and Owen, 1990; Rieke et al., 1991; Armstrong-Gold and Rieke, 2003). The filtering at the synapse is in fact precisely tailored to the input: it weights transmission of each temporal frequency component by its input signal-to-noise ratio (SNR), thereby matching the dynamics of the synapse to the input's signal and noise characteristics. This strategy is known as matched filtering. Because previous matched filtering studies only examined one photoreceptor type, it is not known whether retinal circuitry can implement this type of computation when parallel inputs with different signal and noise properties converge onto individual cells. Can synapses be specialized to process their particular input such that information is optimally separated from noise in each channel?

We use two transgenic zebrafish lines, *Vsx1:mCerulean* (Randlett et al., 2013) and *Vsx2:mCerulean* (newly generated), each with fluorescent labeling of a different BC type. We found both of these BC types to receive input from rods and long wavelength-sensitive cones (L cones). We studied the morphology and physiology of these two BC types to explore synapse specialization in cells receiving converging input from two functionally distinct streams.

3.3 Results

We first present two transgenic lines, each labeling a particular BC type, and show that both BC types contact rods and long wavelength-sensitive (L) cones. We then show that one of these BC types differentiates the anatomical structure of its contacts with rods and L-cones. Finally, using a combination of electrophysiological recording and microscopy, we show that in this BC population rod and cone input streams are mediated by different transmission mechanisms.

Transgenic lines label two molecularly distinct ON bipolar cell types both contacting rods and long wavelength-sensitive cones

We generated two transgenic zebrafish lines, Vsx1:mCerulean and Vsx2:mCerulean, in which distinct BC populations are brightly labeled (hereafter called Vsx1 and Vsx2 BCs) (Fig. 3.1). Both populations have large axonal boutons that stratify in the ON sublamina of the inner plexiform layer adjacent to the ganglion cell layer. Vsx1 but not Vsx2 BCs were immunoreactive for protein kinase C (PKC), a protein found in rod bipolar cells in mammals and previously described 'mixed' BCs (BCs contacting both rods and cones) in teleost fish (Connaughton, 2011; Caminos et al., 2000). The axonal bouton of the Vsx1 BC is relatively spherical compared to the flattened axonal ending of the Vsx2 BC.

We used light microscopy to determine the connectivity of Vsx1 and 2 BCs to upstream photoreceptor types. We confirmed that Vsx1 and Vsx2 BCs contact L-cones by crossing both transgenic lines with the trβ2:tdtomato line, in which L-cones express the red fluorescent protein tdTomato. We also immunostained rods (4C12 antibody) in these retinas. Dendrites of both BC types appear in L-cone pedicles and rod spherules (Fig. 3.2 A). Both BC types appear to contact L-cones and rods specifically, i.e. no contacts with other photoreceptor types were observed.

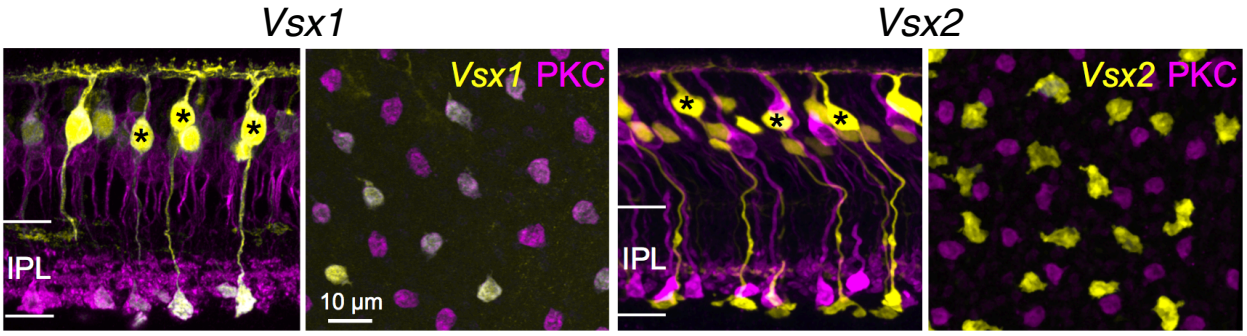

Figure 3.1: Identification of two bipolar cell types in transgenic Vsx-mCerulean lines. Confocal images of Vsx1- and Vsx2-mCerulean transgenic lines, immunostained for protein kinase C (PKC). Labeling in Vsx1-mCerulean but not Vsx2-mCerulean overlap with PKC staining, indicating that these lines label different bipolar cell types. *En face* views (second and fourth panels) are of axonal boutons in the inner plexiform layer (IPL, demarcated in the side views).

Anatomical synapse specialization

We noticed an input type-dependent structural synapse specialization in Vsx1 but not Vsx2 BCs. In fluorescence images, Vsx1 dendritic tips extending into cone pedicles appear to have a typical thin process, while dendritic tips extending into rod spherules have a halo-like appearance (Fig. 3.2 A, panels 1, 2). In contrast, Vsx2 BCs extend typical thin tips into both rod and cone terminals (Fig. 3.2 A, panels 3, 4).

We used serial block face scanning electron microscopy (SBFSEM) to more closely examine dendiric morphology. Vsx1 and 2 BCs were distinguished by the shape of their axonal boutons before tracing full reconstructions of one of each cell type (Fig. 3.2 B). Fig. 3.2 C shows details of each cell's dendritic tips extending into rod and cone terminals, as well as their orientation relative to synaptic "ribbons." Ribbons are a specialized structure in axon terminals of some cell types that act like a conveyor belt for vesicles of neurotransmitter, so the location of this structure reveals the site of neurotransmitter release. Rod spherules each have a single ribbon while cones have several. SBFSEM

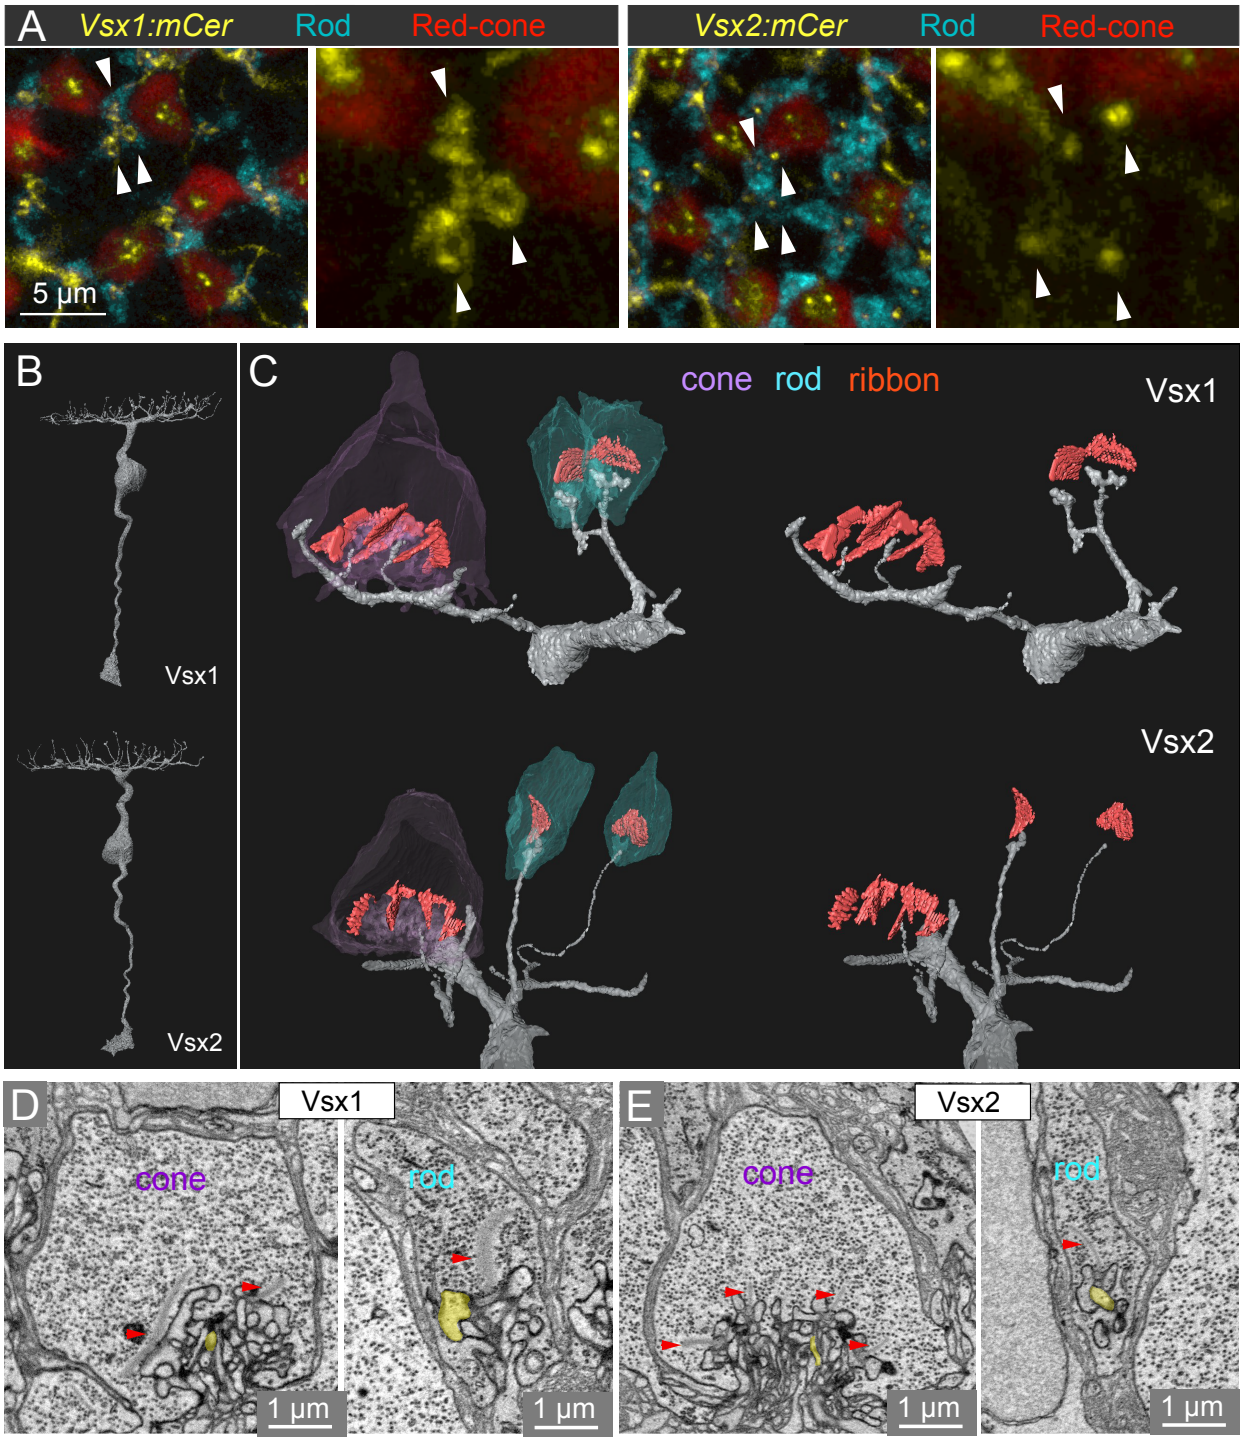
A
Vsx1:mCer
Rod
Red-cone
Vsx2:mCer
Rod
Red-cone
5 µm
B
Vsx1
Vsx2
C
cone
rod
ribbon
Vsx1
Vsx2
D
Vsx1
cone
rod
1 µm
1 µm
E
Vsx2
cone
rod
1 µm
1 µm

Figure 3.2: *Previous page.* **Vsx1 and Vsx2 connectivity with photoreceptors. A:** *En face* confocal images of Vsx1 and Vsx2 dendritic tips (yellow) in rod spherules (cyan, immunolabel with 4C12 antibody) and cone pedicles (trβ2:tdTomato transgenic label). White arrows mark examples of dendritic tips contacting rods, which differ in structure between Vsx1 and 2. **B:** Serial block face scanning electron microscopic (SBFSEM) reconstructions of Vsx1 and 2 BCs. **C:** Example reconstructed synaptic contacts with rods (cyan) and L-cones (purple) made by Vsx1 and 2 BCs. Reconstructions of synaptic ribbons are also shown. **D-E:** Example SBFSEM images of the synaptic contacts shown in (C). Yellow profiles: reconstructed BC dendrite sections, red arrows: synaptic ribbons.

reconstructions show that Vsx1 dendritic tips extending into rod spherules have a claw-like elaboration that casts a net around the site apposed to the rod synaptic ribbon, which contrasts with the cell's more typical thin tips that contact cones (Fig. 3.2 C, top). Vsx2 cells extend thin tips into both rod and cone terminals, with rod-contacting dendrites terminating in the middle of each rod spherule at a site directly apposed to the ribbon (Fig. 3.2 C, bottom).

Synapse architecture interacts with the dynamics of neurotransmitter release, reuptake, and diffusion to influence the nature of the signal received by post-synaptic neurons (DeVries et al., 2006). Therefore the input-specific diversification of dendritic tips we observe in Vsx1 BCs could differentially shape the signal and noise spectra and the magnitude of signals conveyed by the two presynaptic partner types.

Specialization in receptor expression

Previous experiments in fish have provided functional evidence that ON responses (responses excited by an increment in light) in BCs are mediated by two separate mechanisms (Grant and Dowling, 1995; Connaughton, 2011; Nawy and Copenhagen, 1987). Some results suggest that rod transmission is mediated by second-messenger-based

signaling via metabotropic glutamate receptor 6 (mGluR6, a G protein-coupled receptor), while cone transmission is primarily mediated by a chloride conductance gated by excitatory amino acid transporters (EAATs) and also partially by mGluR6 (Wong et al., 2005a,b, 2004). It is not known, however, whether receptor mechanisms are employed differentially at rod and cone synapses within the dendrites of individual BCs.

We immunostained mGluR6 and used super-resolution imaging to discern expression at the level of individual dendritic processes. Fig. 3.3 shows staining for mGluR6 as well as the fluorescence signal from Vsx1 and Vsx2 BC dendrites within the synaptic terminals of rods and L-cones. Overlap of the two signals (white) viewed from multiple perspectives (e.g. top and side views shown) is interpreted as colocalization of dendrites and mGluR6. Within the dendritic arbors of individual Vsx1 cells, the mGluR6 signal colocalizes with the Vsx1's rod-contacting dendrites but not its cone-contacting dendrites (Fig. 3.3, left side). No colocalization with Vsx2 in either synapse type is observed (Fig. 3.3, right side).

In the case of Vsx1 cells, differential expression of this receptor at rod and cone synapses represents a significant functional specialization. Much like anatomical specializations mentioned in the previous section, the particular receptor types mediating neurotransmission can influence many aspects of transmission including kinetics and sensitivity.

Physiology

Physiological experiments are ultimately needed to show functional synapse specialization. Functional effects of synapse specialization could take many forms (Euler et al., 2014; Thoreson, 2007; Attwell, 1986). We were particularly interested in testing the hypothesis that signals are filtered differently at rod and cone synapses. Many advantages of this circuit strategy can be imagined, especially because rod and cone signals differ in several critical ways. Across species, rods are sensitive at light levels 100 or more

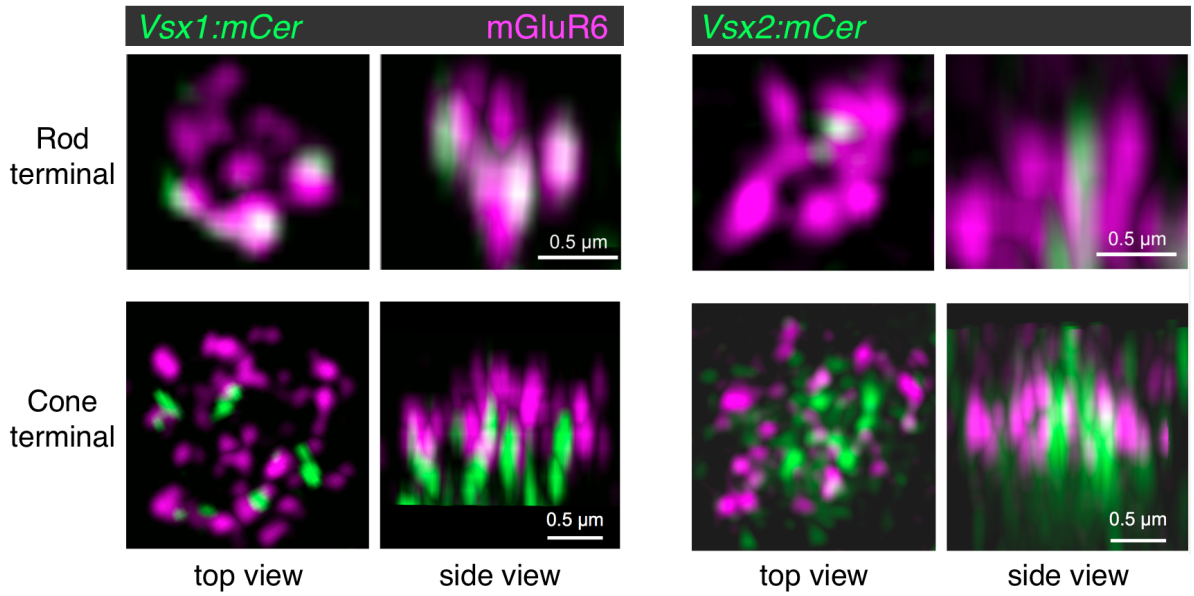

Figure 3.3: mGluR6 expression in dendritic tips contacting rods and cones. Super-resolution images of dendritic processes in rod and cone terminals in Vsx1:mCerulean and Vsx2:mCerulean lines with immunolabeling of mGluR6. Transgenic labels for Vsx1 and 2 dendrites (green) are shown along with the mGluR6 marker (magenta). Colocalization of the two signals (white) indicates the presence of mGluR6 in a dendrite of the transgenically labeled cell.

times dimmer than the threshold for cone sensitivity, their light responses have much slower kinetics (Schnapf and Copenhagen, 1982; Copenhagen et al., 1983; Armstrong-Gold and Rieke, 2003; Wong et al., 2005a) and greater amplitude (when scaled by stimulus strength), and they produce noise at different temporal frequencies (Rieke and Baylor, 2000; Wong et al., 2005b). Previous studies found filtering at a rod-to-BC synapse to be precisely tuned to maximize signal-to-noise ratio in transmission (i.e. the synapse performs "matched filtering") (Bialek and Owen, 1990; Rieke et al., 1991; Armstrong-Gold and Rieke, 2003). We hypothesize this kind of tuning can occur differentially and in parallel in different populations of a cell's synapses when multiple input types converge.

Testing this hypothesis requires characterizing signal and noise spectra of rod and cone signals before and after transmission through their respective synapses, as this reveals how each channel is filtered. With this goal, we recorded rod- and cone- mediated signals in Vsx1 bipolar cells. Our recordings bolster the finding from our imaging studies that rod and cone input is conveyed through different receptor mechanisms in the Vsx1 cell type. However, cone-mediated input did not appear to be conveyed through the expected pathways, and was sometimes entirely absent. Indeed, the cone-mediated input we observed in Vsx1 BCs may not even be dendritic. The absence of clear dendritic input from cones precludes a clear link between the cone-mediated signals we observed and the images of cone-to-Vsx1 BC synapses we presented above. We were therefore unable to connect the synaptic specializations we observed in imaging to a difference in signal filtering. The question of parallel matched filtering of distinct input channels, then, remains unanswered. Nonetheless we did obtain physiological evidence that supports our conclusions from the imaging studies detailed above. In particular, we find that rod and cone input to Vsx1 BCs is indeed transmitted via different receptor mechanisms.

Having found differential expression of the receptor mGluR6 at rod and cone contacts with Vsx1 BCs, we sought to collect functional evidence for a difference in receptor mech-

anisms mediating rod and cone signaling in this BC type. We looked for signatures of transmission through mGluR6 and EAATs since previous work established these two receptor mechanism classes as the primary mediators of ON transmission in teleost fish (Grant and Dowling, 1995; Connaughton, 2011; Nawy and Copenhagen, 1987; Wong et al., 2005a,b, 2004). Both mechanisms are sign-inverting, meaning these mechanisms transform a decrease in neurotransmitter (photoreceptors decrease neurotransmitter release in response to light increments) into a depolarizing (excitatory) post-synaptic current. But the two mechanims differ in one way that is critical to our experiment design: in response to an increment in light, the mGluR6 mechanism causes cation channels to open (Koike et al., 2010), while EAATs cause chloride (Cl^-) channels to close. Cations and Cl^- have different reversal potentials, which means during recording that these inputs can be selectively nulled during whole-cell voltage clamp by moving the holding potential to the reversal potential of cations or Cl^-. In the case of our whole-cell recordings, the cation and Cl^- concentrations of the internal and external solutions set the reversal potentials to +10 mV and -60 mV respectively. Clamping the cell at +10 mV, then, nulls current through the mGluR6 mechanism while clamping at -60 mV nulls the EAAT mechanism.

In some recordings we also used the mGluR6 agonist APB to selectively disrupt mGluR6 signaling. APB could be applied to the solution superfusing the retina at any time during the experiment and took effect after only a few seconds. Therefore we could record before and after addition of APB to solution, allowing us to observe which response components are mediated by mGluR6. All recordings presented here were performed with the addition of pharmacological blockers of classical inhibitory neurotransmission (gabazine, strychnine, and TPMPA), enabling us to selectively record excitatory post-synaptic currents.

In conjuntion with selective recording of mGluR6- and EAAT-mediated currents using different holding potentials, we performed selective activation of rods and cones using dif-

ferent light stimuli. Stimuli were designed to differentiate between rod- and cone-mediated signaling so input from these channels could be isolated. Rods mediate vision at low light levels and are more sensitive overall than cones, and they are especially sensitive to blue light. Cones mediate vision at daytime light levels (and are therefore less sensitive overall than rods) and L-cones are especially sensitive to red light.

We selectively stimulated rods by delivering dim blue flashes that are detectable by rods but not cones. Rod-mediated flash responses (Fig. 3.4 A, thick blue line) are stronger when the cell is held at -60 mV compared with +10 mV, indicating that rod signaling is not mediated by a Cl^- conductance. After bath application of APB (Fig. 3.4, thin blue lines), the blue flash response is eliminated at both holding potentials, indicating that rod signaling to Vsx1 BCs is mediated exclusively by mGluR6.

Activation of cones without any activation of rods was not possible given rods' exquisite sensitivity – a red flash strong enough to activate cones also activated rods despite rods' relative insensitivity to red light. To differentiate cone-mediated responses, we selected a red flash intensity that matches the rod activation strength between the red and blue flashes, thereby allowing us in subsequent analyses to subtract the blue flash response from the red flash response to obtain the cone-mediated component. At -60 mV, the red flash response looks almost identical to the blue flash response, indicating that most or all of this input originates from rods (Fig. 3.4 A). At +10 mV there is a component of the red flash response that is absent in the blue flash response, pointing to this component being cone mediated. This component's absence at the -60 mV clamp and appearance at the +10 mV clamp is consistent with a Cl^--mediated mechanism. Application of APB (thin red lines) eliminates the response at -60 mV as before (again consistent with this input coming from rods), but enhances the putative cone-driven response seen at +10 mV. The cone input is therfore not conveyed by mGluR6.

We repeated this experiment in 6 Vsx1 cells. In Fig. 3.4 B, blue flash responses are

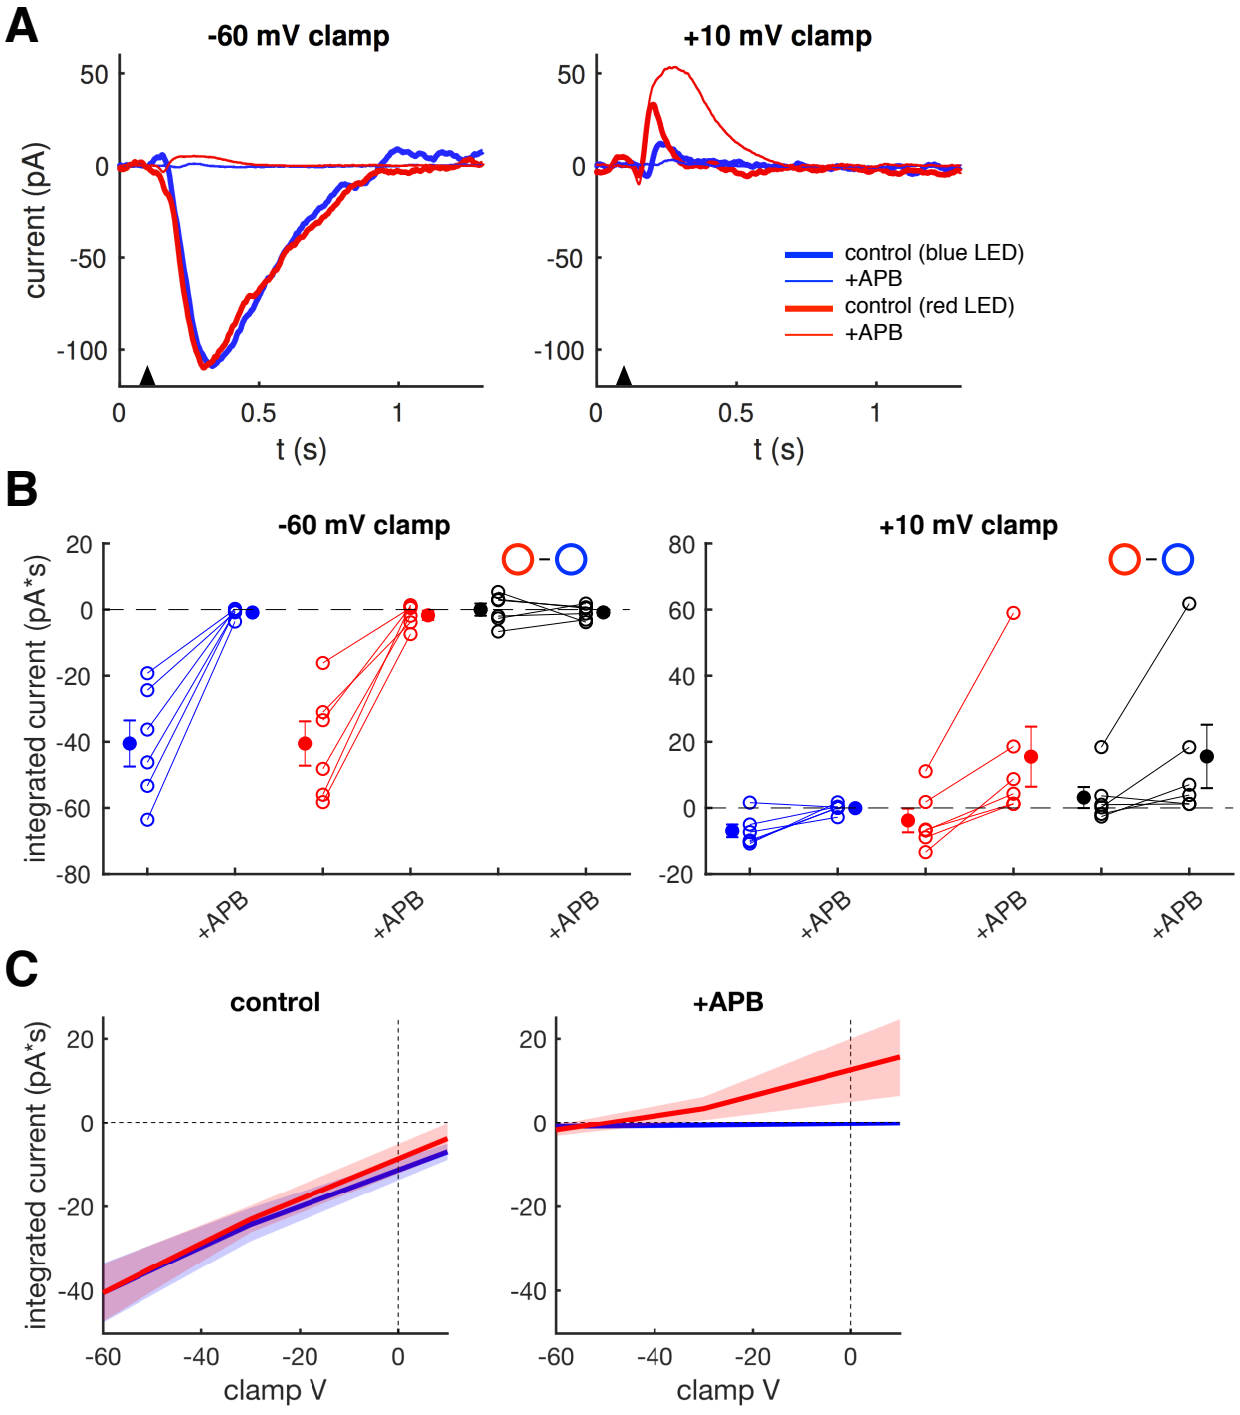
A
-60 mV clamp
+10 mV clamp
current (pA)
t (s)
control (blue LED)
+APB
control (red LED)
+APB
B
-60 mV clamp
+10 mV clamp
integrated current (pA*s)
+APB
C
control
+APB
integrated current (pA*s)
clamp V

Figure 3.4: *Previous page.* **Rod and cone inputs to Vsx1 BCs act through different mechanisms. A:** Excitatory current responses of an example cell to 10 ms flashes (avg. over $\sim$5 trials) from blue and red LEDs (response traces color coded by LED) at holding potentials of -60 mV (left) and +10 mV (right). Responses are shown before (thick lines) and after (thin lines) application of the mGluR6 blocker APB. Black triangles: time of flash. **B:** Summary across cells (n=6) of experiment shown in (A). Integrals of current response averages are shown for the same LEDs and holding potentials. Putative rod input is isolated with the blue flash (blue points), while putative cone input is found by subtracting the blue flash response from the red (black points). Solid circles: means, error bars: standard error of the mean (SEM). **C:** Current vs voltage curves summarizing integrated currents collected at -60 mV, -30 mV, and +10 mV. Shading: $\pm$ SEM.

subtracted from red flash responses to isolate cone contributions. Given these results, we argue that all rod input to Vsx1 BCs is through mGluR6 and that there is no mGluR6-mediated input from cones. This is consistent with our morphological studies which found colocalization of mGluR6 with Vsx1 dendrites only at rod contacts. Population results reveal that the APB-resistant cone-driven response is present in some but not all recorded Vsx1 cells (remaining cells did not have cone-driven responses with or without APB).

Application of APB shifts the reversal potential of the dominant current from $\sim$+10 mV to $\sim$-60 mV (Fig. 3.4 C). APB removes the cation current associated with mGluR6, and the only remaining input is red flash responses with a reversal potential around -60 mV, which is consistent with a Cl^- current.

To be consistent with the EAAT-mediated cone responses described in previous studies, these responses must be depolarizing (excitatory). In response to light, photoreceptors decrease glutamate release, which decreases glutamate transport through EAATs and reduces the associated Cl^- conductance. It is important to determine, then, whether the current we observe is excitatory, as expected of a dendritic EAAT-mediated input. This depends on the position of the cell's resting potential relative to the Cl^- reversal potential.

In experiments described above, the reversal potential of Cl^- is set by the Cl^- concentration in the internal solution of the micropipette used for whole-cell patch recording, which may differ from the concentration maintained by the cell when unperturbed. In a pilot experiment to measure the intrinsic resting potential and Cl^- reversal potentials, we used a perforated patch technique which enabled us to obtain current-clamp recordings without flooding the cell with internal solution. In a Vsx1 cell, we applied APB to isolate the APB-insensitive response component and presented red flashes across a range of holding currents (Fig. 3.5). The baseline voltage across the cell membrane and the integrated voltage responses to red flashes were acquired at each holding current. The resting potential of the cell (voltage across the cell membrane when no external currents are applied) is $\sim$-50 mV (Fig. 3.5 B), and the reversal potential for the APB-insensitive current is $\sim$-60 mV (Fig. 3.5 D). At the resting potential of this cell, the APB-insensitive response is hyperpolarizing, meaning it is inhibitory rather than excitatory and therefore is inconsistent with the proposed depolarizing EAAT mechanism.

It is unclear, then, what the origin of this current is. Cl^- currents also mediate inhibition, but our experiments were performed with a cocktail of pharmacological blockers of inhibition in solution. There are also multiple types of EAATs and they can be found in multiple sites. Indeed, axonal EAAT currents have been described in mouse retina (Ichinose and Lukasiewicz, 2012), and this type of input would be consistent with the currents we measured since ON responses originating in the axon are expected to be hyperpolarizing, unlike responses originating from direct photoreceptor input in the dendrites.

In another pilot experiment, we isolated the APB-insensitive current as before and then applied the pharmacological EAAT blocker TBOA to solution. TBOA eliminates the current (Fig. 3.6, suggesting the current may be mediated by EAATs. However, due to diverse duties of EAATs in the retina, blockade of EAATs by TBOA may in turn disrupt other types of signaling and other function. We have not ultimately made a confident

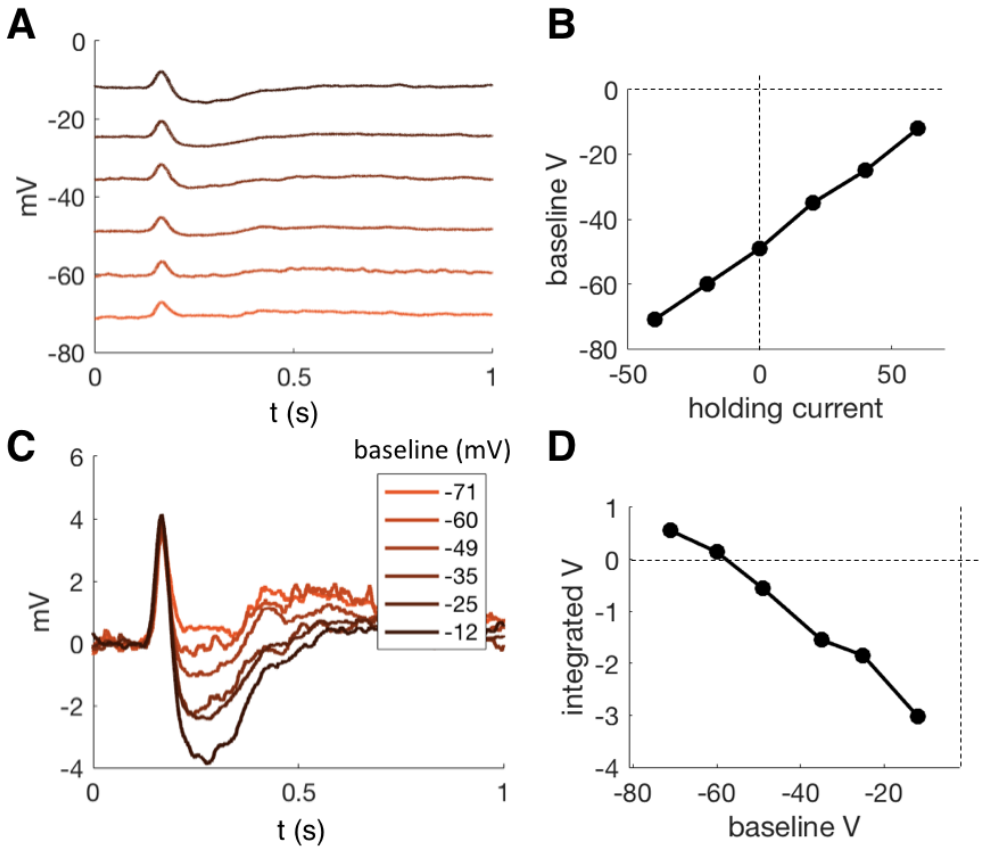

Figure 3.5: Pilot experiment: APB-insensitive input to Vsx1 BCs is hyperpolarizing. A: Responses to 10 ms red flashes across a range of holding currents (holding currents indicated in (B)). **B:** Baseline voltage of responses across holding potentials. The baseline voltage at 0 holding current is the cell's resting voltage. **C:** Baseline-subtracted responses from (A). The first positive peak is input through gap junctions (as evidenced by the fact that it does not change with baseline voltage) and is excluded from subsequent analyses. **D:** Integrated voltage response (baseline subtracted as in (C)). The 0-crossing point of integrated voltage is the input's reversal potential.

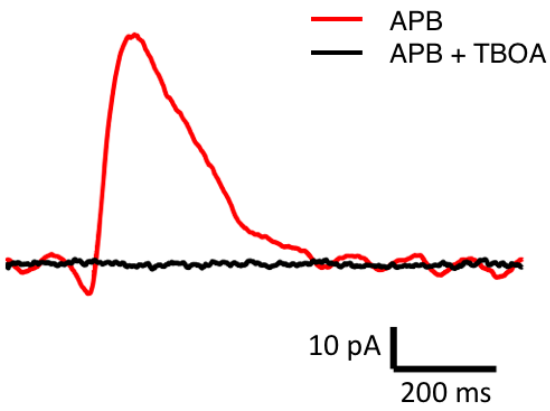

Figure 3.6: Pilot experiment: TBOA eliminates APB-insensitive current to Vsx1 BCs. Current responses in presence of APB, before and after application of TBOA, which disrupts EAATs.

determination, then, of the APB-insensitive current's origin.

3.4 Discussion

Across the nervous system, many circuits must integrate information from multiple channels, but little is known about how this can be achieved when the properties of those converging channels differ in important ways. Understanding how information is processed independently at sub-cellular structures like synapses will help to uncover how multiple signaling channels can be processed in parallel in the same circuitry. Matching of synaptic dynamics to the temporal characteristics of input signals, then, is an issue of general importance in the nervous system. These issues have largely evaded direct study due to the difficulty of characterizing the input signals under physiological conditions. We sought to explore morphological and functional features of synapse specialization in retinal bipolar cells that receive convergent input from rods and cones. This convergence is relevant to early processing in our own visual systems, as the multiplexing of rod and cone signals within the same circuitry is a feature of retinal architecture across species including humans. In Vsx1 BCs, we observed input-specific specializations in dendritic morphology, post-synaptic receptor expression, and physiological responses. Our study falls short of tying these specializations to particular computational needs of rod and cone channels, since the mechanism of cone-to-Vsx1 signaling remains unclear.

Targeting of molecular components to synapses with particular pre- and post-synaptic partner populations has been observed before in various circuits (Craig and Boudin, 2001; Rubio and Wenthold, 1997; Brandstätter et al., 1996). Future studies should investigate the connection between molecular and computational synaptic specializations and their role in parallel processing.

3.4.1 Rod-cone differences in synaptic transmission

Signals are shaped at synapses by numerous pre- and post-synaptic mechanisms involved in neurotransmitter release kinetics, diffusion of neurotransmitter in the synaptic cleft, and kinetics of post-synaptic receptors and associated channels (Attwell, 1986; Thoreson, 2007). Pre- and post-synaptic mechanisms have both been implicated in differentiating kinetics of transmission from rod and cone photoreceptors (Ashmore and Copenhagen, 1980; Rabl et al., 2005). Some studies have already tested the hypothesis that rod and cone synapses onto the same post-synaptic cell can exhibit different transfer kinetics. Schnapf and Copenhagen (1982) recorded rods, cones, and horizonal cells receiving convergent input from both photoreceptor types. By deconvolving pre- and post-synaptic responses, they calculated the impulse response of synapses associated with each input channel and found the kinetics of synaptic transfer to be approximately 10 times faster in cone synapses (Schnapf and Copenhagen, 1982; Copenhagen et al., 1983). This factor is roughly matched to differences in response kinetics in rods and cones themselves. Cadetti et al. (2005) compared synaptic transmission from a rod and cone to the same BC using paired simultaneous recording from pre- and post-synaptic cells. Pre-synaptic electrical stimulation evoked excitatory post-synaptic currents in the BC that were much more transient in the cone stimulation case than in the rod stimulation case.

We aimed to place kinetics differences in rod and cone transmission in the context of signal and noise spectra in rod and cone signals. This would enable us to determine whether differences in kinetics of synaptic transmission serve the signal processing function known as matched filtering (matched filtering was observed in rod-to-BC synapses previously (Bialek and Owen, 1990; Rieke et al., 1991)). We also saw potential to link kinetics differences to different post-synaptic receptor mechanisms, since previous studies showed differences in the receptors mediating rod and cone signal transmission in teleost fish (Wong et al., 2005a,b, 2004; Grant and Dowling, 1995; Connaughton, 2011; Nawy

and Copenhagen, 1987). Concretely tying kinetics differences to computation and mechanistic underpinnings would be a valuable contribution to this unfinished story, and would be a novel demonstration of parallel processing within single cells. Such a contribution is within reach to those venturing to study rod and cone signal convergence in any of many species. To continue this particular inquiry into Vsx1 BCs, though, would require a deeper investigation into the path of cone input to this cell.

Path of cone input to Vsx1

Previous work in teleost fish retina found EAATs and, secondarily, mGluR6, to mediate cone signaling to 'mixed' BCs receiving both rod and cone input (Wong et al., 2005b). Our imaging and physiology indicate no role of mGluR6 in cone signaling to Vsx1. The expected dendritic input through EAATs would be excitatory in response to a flash, but the cone-driven responses we observed were inhibitory. Therefore our observations in this cell type are not consistent with either of the pathways found previously in mixed-input BCs. It is possible that axonal EAATs (rather than dendritic) could mediate the cone input we recorded, and indeed EAATs have been found on BC axon terminals in mouse retina (Ichinose and Lukasiewicz, 2012). Axonal EAATs are sensitive to glutamate spillover from nearby synapses, which would be consistent with the polarity of the currents we observed. Experiments in a retinal slice preparation could confirm whether the recorded current is entering through the dendrites or the axon. One could puff-apply glutamate with a micropipette at the axon and the dendrites while recording from the cell, and the relative strength of the glutamate response in each case would reveal the location of input. One may also be able to sever the connection between the axon terminal and the cell body and record cone-driven light responses in the cell body. There will be no cone-driven response in the axotomized cell if the response relies on axonal EAATs.

3.5 Methods

3.5.1 Animals

Retinal tissue was obtained from adult zebrafish (*Danio rerio*) which were 3 - 18 months old in imaging experiments and 3 - 6 months old on physiology experiments. All procedures were conducted in accordance with University of Washington Institutional Animal Care and Use Committee guidelines. The following previously published transgenic lines were used: *Tg(vsx1:mCerulean)* (Q19) (Randlett et al., 2013) and *Tg(trβ2:Tomato)* (Suzuki et al., 2013). The *Tg(vsx2:Cerulean)* line was generated by injecting pTol2CG2-vsx2-memCerulean-pA plasmids into single-cell stage eggs. Injected fish were out-crossed with wild-type fish to screen for founders. Positive progenies were raised to establish transgenic lines.

3.5.2 Immunohistochemistry

Animals were euthanised by icing and retinas were isolated in oxygenated 0.1 M PBS (pH 7.4) at room temperature. Four relieving cuts were made to allow the retina to lie flat on filter paper (Millipore, HABP13). Retinas were then fixed in 4% paraformaldehyde in PBS for 15 - 20 min. After fixation, retinas were rinsed in PBS. For immunolabeling of retinal sections, eyecups without four relieving cuts were fixed, rinsed in PBS, embedded in 4% agarose in PBS, and sectioned (100 μm thick) using a microtome (Leica Microsystems).

Immunostaining was performed as follows: retinas were blocked in PBS containing 5% normal donkey serum and 0.5% Triton X-100 for 1–4 hr at 4°C. Tissue was incubated in primary antibody in blocking solution for 4–5 days at 4°C. After three washes in 0.5% Triton X-100 PBS, samples were incubated in secondary antibody solution for 1 day at 4°C. Samples were washed three times in 0.5% Triton X-100 PBS, mounted in 0.7% agarose, and coverslipped in Vectashield (Vector Labs).

The primary antibodies used were: rabbit polyclonal anti-mGluR6b (1:2000; Gift from Stephan Neuhauss), mouse monoclonal 4C12 antibody (1:50; Gift from James Fadool), rabbit polyclonal anti-PKCalpha (1:500; Sigma P4334). The secondary antibodies used were Alexa-568 anti-rabbit (1:500; Invitrogen), Alexa-568 anti-chicken (1:500; Invitrogen), Dy-light 649 anti-mouse (1:500; Jackson Immunoresearch laboratories).

3.5.3 Confocal image acquisition

Image stacks were acquired on a confocal microscope (Leica TCS SP8) using a 1.35 numerical aperture (NA) 63x (HC PL APO oil CS2, Leica), 40x (C PL APO CS2, Leica), and 20x (HC PL APO Dry CS2, Leica) objective lens. Typical voxel size was 150 nm and 0.25 μm in xy and z, respectively. Contrast, brightness and pseudo-color were adjusted for display in Fiji software (NIH). 3D image reconstructions were digitally sliced using the Amira software (FEI) slice functions. All measurements were made in Fiji. Image brightness, contrast, and hue were further adjusted in Adobe Photoshop.

3.5.4 Electron microscopy

Isolated retinas were fixed in 4% glutaraldehyde in 0.1M sodium cacodylate buffer (pH7.4) overnight. Tissue was then washed 3 x 5 min (all washes) in the same buffer and incubated in a solution containing 1.5% potassium ferrocyanide and 2% osmium tetroxide (OsO4) in 0.1M cacodylate buffer (0.66% lead in 0.03M aspartic acid, pH 5.5) for 1 hour. After washing, the tissue was placed in a freshly made thiocarbohydrazide solution (0.1g TCH in 10 mL double-distilled H_2O heated to 600°C for 1 h) for 20 min at room temperature. After another rinse, the tissue was incubated in 2% OsO4 for 30 min at room temperature. The samples were rinsed again and stained en bloc in 1% uranyl acetate overnight at 40°C, and washed and stained with Walton's lead aspartate for 30 min. After a final wash, retina pieces were dehydrated in a graded ice-cold alcohol series, and placed

in propylene oxide at room temperature for 10 min. Finally, the samples were embedded in Durcupan resin. The block was then trimmed and mounted in the serial-blockface scanning electron-microscope (GATAN/Zeiss, 3View). Serial sections were cut at 50 nm thickness, and imaged at an x-y resolution of 5.0 - 5.5 nm. 2 x 3 tiles, each about 40 x 40 μm were obtained with an overlap of about 10%. Image stacks were concatenated and aligned using TrackEM2 software (NIH). Neurons were traced using the tracing tools in TrackEM2.

3.5.5 Electrophysiology

Fish (3-6 months old) were dark adapted for at least 2 hours, and the retinas were isolated under infrared light following procedures approved by the Administrative Panel on Laboratory Animal Care at the University of Washington. Retinas were continuously perfused ($\sim$8 mL/min) with oxygenated (95% O_2, 5% CO_2) bicarbonate-buffered Ames solution (Sigma) maintained at 25-28°C.

We patched Vsx1 cells at their axon terminals, which terminate directly adjacent to the ganglion cell layer (Fig. 3.7 A). This allowed us to conduct experiments in a flat mount preparation (with photoreceptors down). We confirmed that recording via the axon terminal allowed sufficient electrical access to dendritic currents in the following experiment: after patching Vsx1 BCs at the axon terminal, we delivered rod-targeting flashes and confirmed that light-evoked currents could be quickly cancelled by switching the holding voltage (delivered at the axon terminal) to the reversal potential of the response (Fig. 3.7 B). To access BC terminals for recording, small groups of GCs were suctioned off the top of the retina to expose the inner plexiform layer. Terminals of vsx1 BCs could be targeted for recording using only infrared illumination. Whole-cell voltage-clamp recordings were obtained using patch pipettes filled with a Cs^+-based internal solution.

Data analysis was performed using custom software developed in MATLAB (Math-

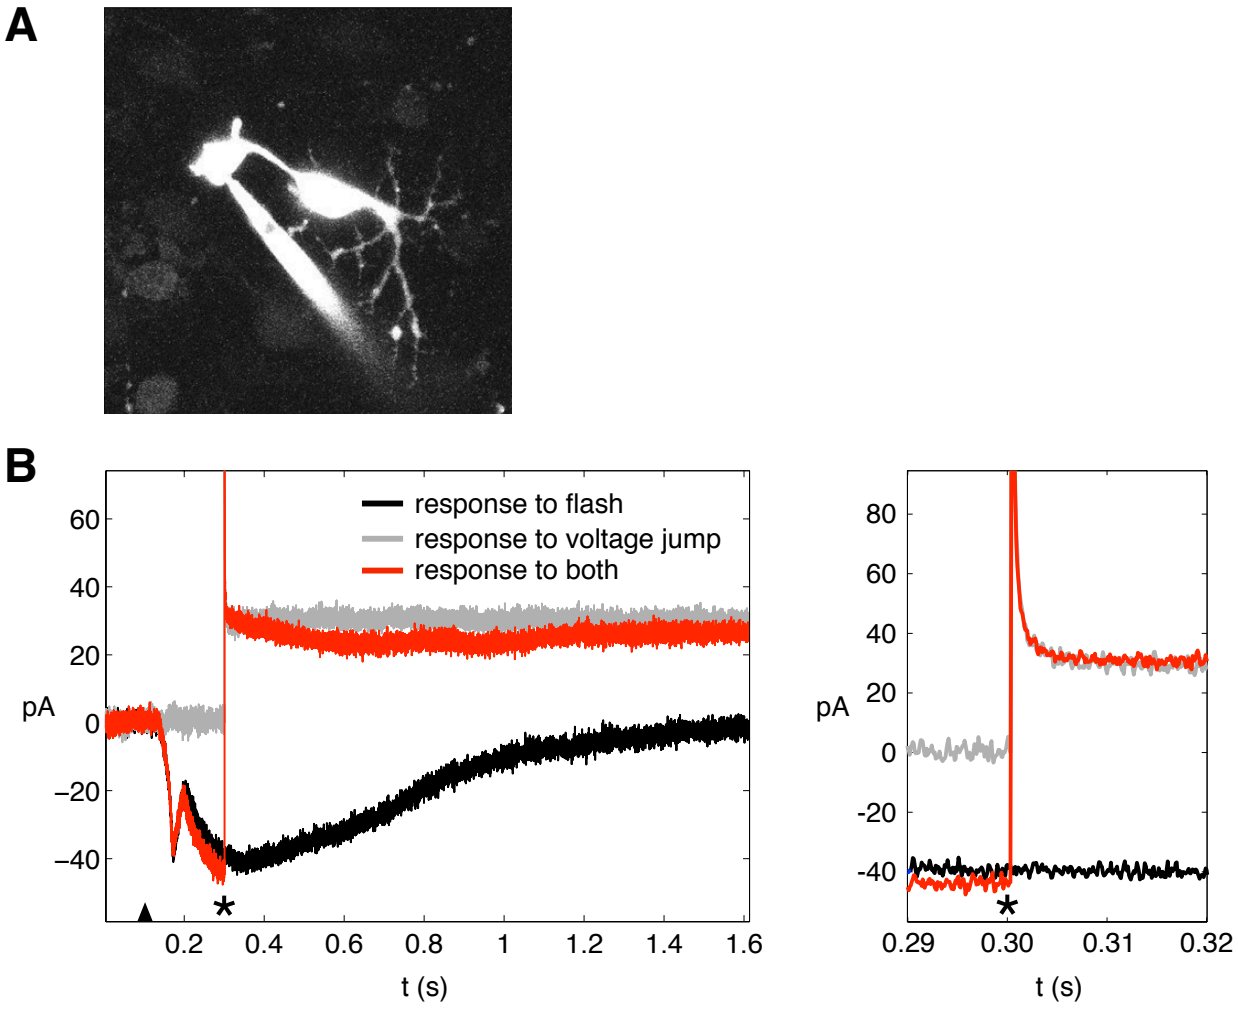

Figure 3.7: Patching the axon terminal of Vsx1. A: 2-photon image of a Vsx1 BC (flat mount with axon terminal on top) filled with alexa 594 after whole-cell patch recording. **B:** Test confirming sufficient electrical access to dendrites with patch recording at axon terminal. Dendritic input (the flash response) is quickly canceled by switching the holding voltage to the reversal potential of the response. The canceled light response (red) is compared with the uninterrupted flash response (black) and the response to the voltage switch alone (gray). Left: full responses, right: detail around time of voltage switch. Black triangle: time of flash, black asterisk: time of voltage switch.

works).

3.5.6 Light stimulation

Light from blue or red light-emitting diodes (LEDs, peak output = 470 nm and 640 nm respectively) was delivered to the recording chamber via fiber optic cable positioned beneath the microscope's condenser lens. The light uniformly illuminated a circular area through an aperture $\sim$0.5 mm in diameter centered on the recorded cell. Protocols for light stimulation were designed to either activate rods only (using the blue LED) or both rods and cones (using the red LED).

3.5.7 Pharmacology

We blocked inhibitory synaptic transmission in all experiments by adding the $GABA_A$ receptor antagonist gabazine (20 μM), the $GABA_C$ receptor antagonist TPMPA (50 μM), and the glycine receptor antagonist strychnine (3 μM) to the superfusion solution. In experiments in which mGluR6-mediated input was blocked, the mGluR6 receptor agonist APB (10 μM) was also added to the perfusion.

Chapter 4
CONCLUSION AND FUTURE DIRECTIONS

In Chapter 1, I developed a rationale for studying neural computation within the context of the particular dynamics of relevant upstream input sources. In subsequent chapters, I presented two projects carried out in the retina that hinge on opportunities to explore circuit function by leveraging deep existing mechanistic and descriptive knowledge of the retina's input layer. First, in Chapter 2 I showed that incorporating a dynamical model of phototransduction in a multi-stage predictive model of retinal output disentangles the roles of multiple circuit mechanisms in adapting to changes in stimulus statistics. Second, in Chapter 3 I detailed an effort to observe how synaptic specializations mediate parallel processing of rod and cone signals in circuitry shared by these functionally distinct input channels. Findings and limitations of these studies are discussed in each chapter. Here, I will focus on what I believe would be productive extensions for future work.

Capturing responses to stimulus transitions

Adaptation is often studied by fitting LN models to data collected in each of multiple stimulus conditions, such as before and after a change in stimulus statistics. But adaptation takes place over some timescale or timescales, meaning a decision must always be made about what time intervals relative to a switch in stimulus should be used for model fitting. LN models are static – nothing about the linear filter or static nonlinearity can exhibit time-dependent nonlinearities like adaptation – so they are most easily interpretable when response properties are sufficiently stable throughout the fitted data. In the work

presented in Chapter 2, we excluded the first 200 ms of the recording collected after stimulus switches because response transients were observed during that interval. Likewise when we tested our multi-stage model's ability to predict responses across conditions we looked only at predictions after the 200 ms buffer interval, meaning we only quantified the model's performance across relatively steady response states in each condition. But the response transients occuring immediately after stimulus changes, e.g. immediately after a saccade in natural vision, are of course essential aspects of retinal output, along with the dynamics of adaptation during this period.

Extensions of the work presented here should include the full timecourse of a dynamic stimulus with transitions between stimulus conditions. Performance can best be measured using multiple trials presenting the same stimulus so the prevalence of noise can be taken into account. We used repeats of the same noise seed in our study but did not repeat the same sequence of transitions between stimulus statistics. Future experiments should add this stimulus protocol modification to allow for more quantitative evaluation of models' performance on the full stimulus with transitions included.

Some improvements to our model and fitting procedure will increase the chance of successfully predicting responses during stimulus transitions. First, the middle section of our multi-stage model (between the cone stage and the spike generation stage) is simply an LN model with a sigmoid nonlinearity, but a sigmoid nonlinearity may be ill-suited to predicting responses to stimuli with changes in statistics. Sigmoid nonlinearities are best suited to gaussian inputs (indeed, a cumulative normal distribution is a sigmoid function), but stimuli with transitions between different statistics are non-gaussian. In preliminary testing of our multi-stage model with the full dynamic stimulus (including transitions between stimulus statistics) the sigmoid nonlinearity is revealed as a weak point because a sigmoid shape does not effectively describe the sampling we obtained of the mapping from generator signal to response (data not shown). Choosing a different shape or pa-

rameterization for the static nonlinearity is unlikely to be productive, however, since the underlying challenge is that the true nonlinearity used by the cell is not static. Rather, nonlinear processing by the cell adjusts as stimulus statistics change. A better strategy for augmenting our multi-stage model for use on a dynamic stimulus would therefore be to incorporate a mechanism to flexibly produce changes in the input-output mapping. Extensions of the LN model framework that use multiple interacting branches (Cui et al., 2016) or adaptive sub-modules (Ozuysal and Baccus, 2012) may provide sufficient flexibility.

The best model architecture will likely depend on which RGC type's responses are being modeled. This fact in itself may present a fruitful opportunity to explore differences in circuit mechanisms that give rise to response behaviors in different RGC types. Disentangling the contributions of phototransduction and spike generation from those of circuit computation should make it easier to attribute a particular architecture or mechanism to the circuit, and one may find that the best architecture depends on the cell type or circuit in question.

Naturalistic stimuli

Extensions of our modeling study should push toward naturalistic stimuli. Using a database of natural images and eye tracking trajectories such as DOVES (van der Linde et al., 2009) (images from van Hateren and van der Schaaf (1998)), one can generate naturalistic spatiotemporal stimuli. To create a purely temporal naturalistic stimulus, one can average over some amount of space around the point of fixation (or another coordinate relative to the point of fixation) at each time point to generate a time series of mean intensities around that point (Turner and Rieke, 2016). Even this temporal stimulus would introduce many more features of natural inputs than are included in our dynamic noise stimulus while still restricting the experiment to the time domain. Using a spatiotemporal stimulus will activate nonlinear spatial interactions which will require spatial models such as subunit

models to capture. Spatial processing places further demands not only on lateral model structure, but also on the vertical ordering of model components. For example, placing a nonlinearity prior to a point of convergence of signals (such as a summation point) will have a different effect than placing it after the convergence – the latter will act on signals after they are already combined (Turner and Rieke, 2016; Turner et al., 2018; Field and Rieke, 2002). By bringing these demands, spatial stimuli will help distinguish between good and bad model architectures.

Noise

I have argued that representing too many cirucit mechanisms using too few model components risks conflating multiple phenomena that do not truly vary in concert, causing different mechanisms to trade against each other in the model and diminishing the model's interpretability. This is also true of noise in circuit models. Multiple sources of noise affect retinal output, and the location of each noise source determines its particular effect (Brinkman et al., 2016; Weber et al., 2020). Adding noise sources to statistical models re-shapes other model components, and the nature of this reshaping is also dependent on noise location. Our model is entirely deterministic, meaning noise is not considered. Incorporating an appropriate number of noise sources in appropriate locations within a multi-stage model like ours may re-shape best-fit model components and potentially help constrain model architecture. This will ultimately help to produce a more faithful representation of computations carried out in the circuit.

Decoding

Models that predict RGC responses to visual stimuli can be said to capture the neural encoding of those stimuli, but one would also like to solve the converse problem of decoding: given the neural response, what can be predicted about the stimulus that evoked it? De-

coding can be valuable in understanding how sensory information is represented in neural signals, and may also be useful for neural prosthesis design applications. Similar to our encoding model, decoding may benefit from multiple stages that separate contributions from specific circuit elements. Many techniques used in decoding are similar to those used in encoding models, so there is reason to be optimistic about mapping our approach onto decoding. For example, it may be possible to perform the inverse of each of the major stages of our encoding model (which are all deterministic), first by predicting excitatory input to a RGC using its recorded spike response, then predicting the cone response from the RGC excitatory input, and finally predicting the stimulus from the predicted cone response.

Inverting the forward model is not always straightforward – for example, the particular differential equations comprising the biophysical model of phototransduction in cones cannot be completely inverted analytically. Different prediction strategies could be employed in such cases, including artificial neural networks which can capture time-dependent nonlinear relationships.

Linear (Parthasarathy et al., 2017) and nonlinear (Brackbill et al., 2020) methods including neural networks (Kim et al., 2020) have already been used in decoding applications, including efforts to decode spatial stimuli like natural images. But decoding tools may be more effective or easier to train (i.e. require fewer parameters and less training data) if embedded within a multi-stage decoder that subdivides the decoding problem along the lines of specific neural mechanisms or processing stages.

REFERENCES

Agüera y Arcas, B., Fairhall, A. L., and Bialek, W. (2003). Computation in a single neuron: Hodgkin and Huxley revisited. *Neural computation*, 15(8):1715–49.

Angueyra, J. M., Baudin, J., Schwartz, G. W., and Rieke, F. (2021). Multiple time scales of adaptation allow cones to encode the inputs created by visual exploration of natural scenes. *bioRxiv*.

Angueyra, J. M. and Rieke, F. (2013). Origin and effect of phototransduction noise in primate cone photoreceptors. *Nature Neuroscience*, 16(11):1692–1700.

Appleby, T. R. and Manookin, M. B. (2020). Selectivity to approaching motion in retinal inputs to the dorsal visual pathway. *eLife*, 9.

Armstrong-Gold, C. and Rieke, F. (2003). Bandpass filtering at the rod to second-order cell synapse in salamander (Ambystoma tigrinum) retina. *The Journal of neuroscience*, 23(9):3796–3806.

Asari, H. and Meister, M. (2012). Divergence of visual channels in the inner retina. *Nature Neuroscience*, 15(11):1581–1589.

Ashmore, J. F. and Copenhagen, D. R. (1980). Different postsynaptic events in two types of retinal bipolar cell. *Nature*, 288(5786).

Attwell, D. (1986). The Sharpey-Schafer lecture. Ion channels and signal processing in the outer retina. *Quarterly journal of experimental physiology (Cambridge, England)*, 71(4):497–536.

Baccus, S. a. and Meister, M. (2002). Fast and slow contrast adaptation in retinal circuitry. *Neuron*, 36(5):909–919.

Baden, T., Nikolaev, A., Esposti, F., Dreosti, E., Odermatt, B., and Lagnado, L. (2014). A Synaptic Mechanism for Temporal Filtering of Visual Signals. *PLoS Biology*, 12(10).

Baylor, D. A., Nunn, B. J., and Schnapf, J. L. (1984). The photocurrent, noise and spectral sensitivity of rods of the monkey Macaca fascicularis. *The Journal of Physiology*, 357(1).

Bialek, W. and Owen, W. G. (1990). Temporal filtering in retinal bipolar cells. Elements of an optimal computation? *Biophysical journal*, 58(5):1227–1233.

Brackbill, N., Rhoades, C., Kling, A., Shah, N. P., Sher, A., Litke, A. M., and Chichilnisky, E. J. (2020). Reconstruction of natural images from responses of primate retinal ganglion cells. *eLife*, 9.

Brandstätter, J. H., Koulen, P., Kuhn, R., Van Der Putten, H., and Wässle, H. (1996). Compartmental localization of a metabotropic glutamate receptor (mGluR7): Two different active sites at a retinal synapse. *Journal of Neuroscience*, 16(15).

Brenner, N., Bialek, W., and de Ruyter van Steveninck, R. (2000). Adaptive rescaling maximizes information transmission. *Neuron*, 26(3):695–702.

Brinkman, B. A., Weber, A. I., Rieke, F., and Shea-Brown, E. (2016). How Do Efficient Coding Strategies Depend on Origins of Noise in Neural Circuits? *PLoS Computational Biology*, 12(10).

Burkhardt, D. A. (1994). Light adaptation and photopigment bleaching in cone photoreceptors in situ in the retina of the turtle. *Journal of Neuroscience*, 14(3 I).

Cadetti, L., Tranchina, D., and Thoreson, W. B. (2005). A comparison of release kinetics and glutamate receptor properties in shaping rod-cone differences in EPSC kinetics in the salamander retina. *Journal of Physiology*, 569(3).

Cafaro, J. and Rieke, F. (2013). Regulation of spatial selectivity by crossover inhibition. *Journal of Neuroscience*, 33(15).

Caminos, E., Velasco, A., Jarrín, M., Lillo, C., Jimeno, D., Aijón, J., and Lara, J. M. (2000). A comparative study of protein kinase C-like immunoreactive cells in the retina. *Brain, Behavior and Evolution*, 56(6).

Chabrol, F. P., Arenz, A., Wiechert, M. T., Margrie, T. W., and DiGregorio, D. A. (2015). Synaptic diversity enables temporal coding of coincident multisensory inputs in single neurons. *Nature neuroscience*, 18(5):718–27.

Chander, D. and Chichilnisky, E. J. (2001). Adaptation to temporal contrast in primate and salamander retina. *J Neurosci*, 21(24):9904–9916.

Chichilnisky, E. J. (2001). A simple white noise analysis of neuronal light. 12:199–213.

Connaughton, V. (2011). Bipolar cells in the zebrafish retina. *Visual Neuroscience*, 28(01):77–93.

Copenhagen, D. R., Ashmore, J. F., and Schnapf, J. K. (1983). Kinetics of synaptic trans-

mission from photoreceptors to horizontal and bipolar cells in turtle retina. *Vision Research*, 23(4):363–369.

Craig, A. M. and Boudin, H. (2001). Molecular heterogeneity of central synapses: Afferent and target regulation.

Cui, Y., Wang, Y. V., Park, S. J. H., Demb, J. B., and Butts, D. A. (2016). Divisive suppression explains high-precision firing and contrast adaptation in retinal ganglion cells. *eLife*.

DeVries, S. H., Li, W., and Saszik, S. (2006). Parallel Processing in Two Transmitter Microenvironments at the Cone Photoreceptor Synapse. *Neuron*, 50(5).

Dunn, F. A., Lankheet, M. J., and Rieke, F. (2007). Light adaptation in cone vision involves switching between receptor and post-receptor sites. *Nature*, 449(7162):603–606.

Dunn, F. a. and Rieke, F. (2006). The impact of photoreceptor noise on retinal gain controls. *Current Opinion in Neurobiology*, 16:363–370.

Dunn, F. a. and Rieke, F. (2008). Single-photon absorptions evoke synaptic depression in the retina to extend the operational range of rod vision. *Neuron*, 57(6):894–904.

Euler, T., Haverkamp, S., Schubert, T., and Baden, T. (2014). Retinal bipolar cells: elementary building blocks of vision. *Nature reviews. Neuroscience*, 15(8):507–19.

Fairhall, A. L., Lewen, G. D., Bialek, W., and de Ruyter Van Steveninck, R. R. (2001). Efficiency and ambiguity in an adaptive neural code. *Nature*, 412(6849):787–92.

Fechner, G. T. (1966). Elements of psychophysics. Vol. I. [Originally published 1860]. *Readings in the history of psychology*.

Field, G. D. and Rieke, F. (2002). Nonlinear signal transfer from mouse rods to bipolar cells and implications for visual sensitivity. *Neuron*, 34:773–785.

Frazor, R. A. and Geisler, W. S. (2006). Local luminance and contrast in natural images. *Vision Research*.

Freeman, J., Field, G. D., Li, P. H., Greschner, M., Gunning, D. E., Mathieson, K., Sher, A., Litke, A. M., Paninski, L., Simoncelli, E. P., and Chichilnisky, E. J. (2015). Mapping nonlinear receptive field structure in primate retina at single cone resolution. *eLife*, 4(OCTOBER2015).

Grant, G. B. and Dowling, J. E. (1995). A glutamate-activated chloride current in cone-

driven ON bipolar cells of the white perch retina. *The Journal of neuroscience*, 15(5 Pt 2):3852–3862.

Grimes, W. N., Graves, L. R., Summers, M. T., and Rieke, F. (2015). A simple retinal mechanism contributes to perceptual interactions between rod-and cone-mediated responses in primates. *eLife*, 4(JUNE2015).

Grimes, W. N., Schwartz, G. W., and Rieke, F. (2014). The synaptic and circuit mechanisms underlying a change in spatial encoding in the retina. *Neuron*, 82(2).

Grimes, W. N., Songco-Aguas, A., and Rieke, F. (2018). Parallel processing of rod and cone signals: Retinal function and human perception.

Hong, S., Agüera y Arcas, B., and Fairhall, A. L. (2007). Single neuron computation: from dynamical system to feature detector. *Neural computation*, 19:3133–3172.

Hoon, M., Okawa, H., Della Santina, L., and Wong, R. O. (2014). Functional architecture of the retina: Development and disease. *Progress in Retinal and Eye Research*, 42:44–84.

Hosoya, T., Baccus, S. A., and Meister, M. (2005). Dynamic predictive coding by the retina. *Nature*, 436(7047).

Ichinose, T. and Lukasiewicz, P. D. (2012). The mode of retinal presynaptic inhibition switches with light intensity. *Journal of Neuroscience*, 32(13).

Izhikevich, E. M. (2018). *Dynamical Systems in Neuroscience*.

Kim, K. J. and Rieke, F. (2001). Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. *Journal of Neuroscience*, 21(1).

Kim, K. J. and Rieke, F. (2003). Slow Na+ inactivation and variance adaptation in salamander retinal ganglion cells. *Journal of Neuroscience*, 23(4).

Kim, Y. J., Brackbill, N., Batty, E., Lee, J. H., Mitelut, C., Tong, W., Chichilnisky, E. J., and Paninski, L. (2020). Nonlinear decoding of natural images from large-scale primate retinal ganglion recordings.

Koike, C., Numata, T., Ueda, H., Mori, Y., and Furukawa, T. (2010). TRPM1: A vertebrate TRP channel responsible for retinal ON bipolar function. *Cell Calcium*, 48(2-3):95–101.

Kuang, X., Poletti, M., Victor, J. D., and Rucci, M. (2012). Temporal encoding of spatial information during active visual fixation. *Current Biology*, 22(6).

Kuo, S. P., Schwartz, G. W., and Rieke, F. (2016). Nonlinear Spatiotemporal Integration by Electrical and Chemical Synapses in the Retina. *Neuron*, 90(2).

Latimer, K. W. and Fairhall, A. L. (2020). Capturing Multiple Timescales of Adaptation to Second-Order Statistics With Generalized Linear Models: Gain Scaling and Fractional Differentiation. *Frontiers in Systems Neuroscience*, 14.

Laughlin, S. (1981). A simple coding procedure enhances a neuron's information capacity.

Li, Y. N., Tsujimura, T., Kawamura, S., and Dowling, J. E. (2012). Bipolar cell-photoreceptor connectivity in the zebrafish (Danio rerio) retina. *The Journal of Comparative Neurology*, 520(16):3786–3802.

Maheswaranathan, N., Kastner, D. B., Baccus, S. A., and Ganguli, S. (2018). Inferring hidden structure in multilayered neural circuits. *PLoS Computational Biology*, 14(8).

Manookin, M. B. and Demb, J. B. (2006). Presynaptic Mechanism for Slow Contrast Adaptation in Mammalian Retinal Ganglion Cells. *Neuron*, 50(3).

Mensi, S., Hagens, O., Gerstner, W., and Pozzorini, C. (2016). Enhanced Sensitivity to Rapid Input Fluctuations by Nonlinear Threshold Dynamics in Neocortical Pyramidal Neurons. *PLoS Computational Biology*, 12(2).

Nagel, K. I. and Doupe, A. J. (2006). Temporal processing and adaptation in the songbird auditory forebrain. *Neuron*, 51(6):845–59.

Nawy, S. and Copenhagen, D. R. (1987). Multiple classes of glutamate receptor on depolarizing bipolar cells in retina. *Nature*, 325(6099):56–58.

Ozuysal, Y. and Baccus, S. A. (2012). Linking the Computational Structure of Variance Adaptation to Biophysical Mechanisms. *Neuron*.

Parthasarathy, N., Batty, E., Falcon, W., Rutten, T., Rajpal, M., Chichilnisky, E. J., and Paninski, L. (2017). Neural networks for efficient Bayesian decoding of natural images from retinal neurons. In *Advances in Neural Information Processing Systems*, volume 2017-December.

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., and Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207).

Pugh, E. N. and Lamb, T. D. (1993). Amplification and kinetics of the activation steps in phototransduction.

Rabl, K., Cadetti, L., and Thoreson, W. B. (2005). Kinetics of exocytosis is faster in cones than in rods. *Journal of Neuroscience*, 25(18).

Randlett, O., MacDonald, R. B., Yoshimatsu, T., Almeida, A. D., Suzuki, S. C., Wong, R. O., and Harris, W. A. (2013). Cellular Requirements for Building a Retinal Neuropil. *Cell Reports*, 3(2).

Reinagel, P. and Zador, A. M. (1999). Natural scene statistics at the centre of gaze. *Network: Computation in Neural Systems*, 10(4).

Rieke, F. (2001). Temporal contrast adaptation in salamander bipolar cells. *Journal of Neuroscience*, 21(23).

Rieke, F. and Baylor, D. (1998a). Single-photon detection by rod cells of the retina. *Reviews of Modern Physics*, 70(3):1027–1036.

Rieke, F. and Baylor, D. A. (1998b). Origin of reproducibility in the responses of retinal rods to single photons. *Biophysical Journal*, 75(4).

Rieke, F. and Baylor, D. a. (2000). Origin and functional impact of dark noise in retinal cones. *Neuron*, 26(1):181–186.

Rieke, F., Owen, W. G., and Bialek, W. (1991). Optimal Filtering in the Salamander Retina. *Advances in Neural Information Processing Systems 3*, pages 377–383.

Rieke, F., Warland, D., De Ruyter Van Steveninck, R., and Bialek, W. (1997). Spikes: Exploring the Neural Code.

Rubio, M. E. and Wenthold, R. J. (1997). Glutamate receptors are selectively targeted to postsynaptic sites in neurons. *Neuron*, 18(6).

Schnapf, J. L. and Copenhagen, D. R. (1982). Differences in the kinetics of rod and cone synaptic transmission. *Nature*, 296(5860):862–864.

Shapley, R. and Enroth-Cugell, C. (1984). Chapter 9 Visual adaptation and retinal gain controls.

Shapley, R. M. and Victor, J. D. (1978). The effect of contrast on the transfer properties of cat retinal ganglion cells. *The Journal of Physiology*, 285(1).

Sharpee, T. O. (2013). Computational identification of receptive fields.

Sharpee, T. O., Sugihara, H., Kurgansky, A. V., Rebrik, S. P., Stryker, M. P., and Miller,

K. D. (2006). Adaptive filtering enhances information transmission in visual cortex. *Nature*, 439(7079).

Simoncelli, E. P. and Olshausen, B. A. (2001). Natural image statistics and neural representation.

Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W., and Meister, M. (1997). Adaptation of retinal processing to image contrast and spatial scale. *Nature*, 386(6620).

Suzuki, S. C., Bleckert, A., Williams, P. R., Takechi, M., Kawamura, S., and Wong, R. O. (2013). Cone photoreceptor types in zebrafish are generated by symmetric terminal divisions of dedicated precursors. *Proceedings of the National Academy of Sciences of the United States of America*, 110(37).

Thoreson, W. B. (2007). Kinetics of synaptic transmission at ribbon synapses of rods and cones. *Molecular Neurobiology*, 36:205–223.

Tikidji-Hamburyan, A., Reinhard, K., Seitter, H., Hovhannisyan, A., Procyk, C. A., Allen, A. E., Schenk, M., Lucas, R. J., and Münch, T. A. (2015). Retinal output changes qualitatively with every change in ambient illuminance. *Nature Neuroscience*, 18(1).

Tkačik, G., Ghosh, A., Schneidman, E., and Segev, R. (2014). Adaptation to changes in higher-order stimulus statistics in the salamander retina. *PLoS ONE*, 9(1).

Tobler, P. N., Fiorillo, C. D., and Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science*, 307(5715).

Turner, M. H. and Rieke, F. (2016). Synaptic Rectification Controls Nonlinear Spatial Integration of Natural Visual Inputs. *Neuron*, 90(6).

Turner, M. H., Schwartz, G. W., and Rieke, F. (2018). Receptive field center-surround interactions mediate context-dependent spatial contrast encoding in the retina. *eLife*, 7.

van der Linde, I., Rajashekar, U., Bovik, A. C., and Cormack, L. K. (2009). DOVES: A database of visual eye movements. *Spatial Vision*, 22(2).

van Hateren, J. H. and Snippe, H. P. (2007). Simulating human cones from mid-mesopic up to high-photopic luminances. *Journal of Vision*, 7(4).

van Hateren, J. H. and van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society B: Biological Sciences*, 265(1394).

Victor, J. D. and Purpura, K. P. (1996). Nature and precision of temporal coding in visual cortex: A metric- space analysis. *Journal of Neurophysiology*, 76(2).

Victor, J. D. and Purpura, K. P. (1997). Metric-space analysis of spike trains: Theory, algorithms and application. *Network: Computation in Neural Systems*, 8(2).

Wark, B., Fairhall, A., and Rieke, F. (2009). Timescales of inference in visual adaptation. *Neuron*, 61(5):750–61.

Wark, B., Lundstrom, B. N., and Fairhall, A. (2007). Sensory adaptation. *Current opinion in neurobiology*, 17(4):423–9.

Weber, A. I. and Fairhall, A. L. (2019). The role of adaptation in neural coding.

Weber, A. I., Shea-Brown, E., and Rieke, F. (2020). Identification of multiple noise sources improves estimation of neural responses across stimulus conditions. *bioRxiv*.

Wong, K. Y., Adolph, A. R., and Dowling, J. E. (2005a). Retinal Bipolar Cell Input Mechanisms in Giant Danio. I. Electroretinographic Analysis. *Journal of Neurophysiology*, 93(1):84–93.

Wong, K. Y., Cohen, E. D., and Dowling, J. E. (2005b). Retinal Bipolar Cell Input Mechanisms in Giant Danio. II. Patch-Clamp Analysis of ON Bipolar Cells. *Journal of Neurophysiology*, 93(1):94–107.

Wong, K. Y., Gray, J., Hayward, C. J. C., Adolph, A. R., and Dowling, J. E. (2004). Glutamatergic mechanisms in the outer retina of larval zebrafish: analysis of electroretinogram b- and d-waves using a novel preparation. *Zebrafish*, 1(2):121–31.