Quantitative Modeling of Cone Signal Combination in Macaque Primary Visual Cortex

J Patrick Weller

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

Doctor of Philosophy

University of Washington

2018

Reading Committee:
Gregory Horwitz, Chair
Fred Rieke
Anitha Pasupathy

Program Authorized to Offer Degree:

Physiology and Biophysics

University of Washington

Abstract

Quantitative Modeling of Cone Signal Combination in Macaque Primary Visual Cortex

J Patrick Weller

Chair of the Supervisory Committee:
Gregory D Horwitz, PhD
Physiology and Biophysics

Vision at moderate to high light levels begins with the activity in three types of cone photoreceptors located in back of the retina: long-wavelength-sensitive (L-) cones, medium-wavelength-sensitive (M-) cones, and short-wavelength-sensitive (S-) cones. One central goal in the field of color neurophysiology is to understand how modulations in cone activity are transformed as they propagate through the visual system. My graduate work has been dedicated to understanding how cone signals are processed in the primary visual cortex (area V1). In the first chapter, I lay out some of the typical considerations for designing an effective electrophysiology experiment in the context of color vision, such as the spatiotemporal structure of the stimulus, color spaces and stimulus distributions, as well as the quantitative models under consideration. In the second chapter, I discuss some of the pitfalls and advantages of different stimulus distributions and analysis techniques under the linear-nonlinear cascade model, a standard model in the field.

In the third chapter, I present a novel single-cell electrophysiology experiment probing the effect of L- and M-cone modulations on the responses of single neurons in V1. In this chapter, I also present data from a new quantitative model for describing neural responses to L- and M-cone modulations. In the final chapter, I consider how we might use the results of these experiments to build strong priors for fitting an appropriate model of L- and M-cone processing in V1.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF EQUATIONS

LIST OF TABLES

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor, Gregory Horwitz, for taking a chance on me. When I entered graduate school, I did not have a strong quantitative or technical background, but Greg saw something in me. He has been so incredibly generous with his time and vast expertise: teaching me how to code, carefully explaining models of the visual system, walking me through derivations of countless statistical analyses, and shaping my ideas with a careful and critical eye. More than this, he has taught me by his own example that some of the truest signs of expertise are honesty about what we know and humility about what we do not. He has taught me more about being a thoughtful scientist than I could list here, and for that I am forever grateful.

I would also like to thank my wonderful committee members, Fred Rieke, Anitha Pasupathy, Adrienne Fairhall, and Jay Neitz, for providing invaluable feedback on the direction of my research. The work presented here, and so much more that didn't make it into my thesis, was advanced immeasurably by their guidance and thoughtfulness.

I would like to thank my Horwitz lab family, both past and present: Charlie Hass, Amy Ni, Yasmine El-Shamayleh, Abhishek De, Zack Lindbloom-Brown, Emily Gelfand, and Elise Grover. Their insight, support, and encouragement shaped my research, challenged me to consider new perspectives, and has ultimately made me a better scientist. Each of these people helped to make my papers, posters, and talks both accessible and enlightening to color neurophysiologists and non-scientists alike.

I would also like to thank my cohort and extended scientific family: Sweta Agrawal, Matt Brodsky, Dina Popovkina, Stephanie Seeman, Florie D'Orazi, Wucheng Tao, Tim Oleskiw, Pamela Baker, Max Turner, Eric Thomas, Atom Lesiak, and Dean Pospisil. Each of these people

were so incredibly generous with their time and expertise, especially in my first few years of graduate school when I was struggling the most to learn all sorts of new skills. Without them, I would never have made it past the first year of graduate school. They taught me that success is never achieved in isolation, but by surrounding oneself with a generous and supportive community.

I would like to thank my family for supporting me on this long and, at times, uncertain journey. I could never have earned this degree without their confidence and encouragement.

Although they will never read this, I would be remiss if I did not thank all of the monkeys that participated in my research through the years: Sedna, Apollo, Freya, Kali, Nut, and Maui. Working with these amazing animals has been one of the greatest privileges of my graduate school career. To those who feel conflicted about animal research, I would encourage you to befriend a primate physiologist. I hope you will see, as I did, that these animals are deeply loved and cared for by the people who choose to work with them. I love these monkeys like family, and I will miss them dearly.

Lastly and most importantly, I would like to thank my wife, Leah, for her boundless love and support. My time in graduate school has not always been easy—any success that I have had is due to her unwavering encouragement and confidence in my abilities. I cannot express how thankful I am to have her as my partner.

DEDICATION

For Leah, my amazing wife and constant cornerstone, whose love and support has carried me through graduate school. Leah, I love you more with every passing day, and I cannot thank you enough for being my amazing partner on this journey. I can't wait to live our next journey by your side.

Chapter 1. INTRODUCTION

The ability to perceive light has been an enormous benefit across the animal kingdom, evolving separately many times and in many species. In trichromatic humans and old-world primates, vision at moderate to high light levels is built on the activity of only three varieties of cone photoreceptors located in the retina: the long-wavelength-sensitive (L-), medium wavelength-sensitive (M-), and short-wavelength-sensitive (S-) cone photoreceptors. The transformation of these photoreceptor signals into a rich and informative sensory experience is a complicated process, as evidenced by the high percentage of our brains dedicated solely to this sensory modality. Understanding exactly how the brain achieves this incredible transformation is the goal of visual neuroscience. In my graduate work, I attempted to advance toward this goal by combining quantitative modeling with single-cell electrophysiology.

Broadly, the visual experience can be segregated into two kinds of information: changes in brightness and changes in color. Changes in brightness are produced when the activity of all three cone classes increases or decreases together, or by non-opponent cone signals. Changes in color are produced by differential changes between the cone classes, or by cone-opponent signals. In the brain, these two types of information are processed in tandem as they propagate through the visual hierarchy. Independently, brightness and color may reveal different aspects of the visual world. To illustrate this point, let us consider a painting by Monet. In **Figure 1.1**A, we see a grayscaled version of this painting which has been stripped of cone-opponent modulations. From this picture, the only distinctive images are two or three boats floating on a sea of gray swirls. In **Figure 1.1**B, we see a colored image that drives all of the cone-opponent modulations that were stripped from **Figure 1.1**A. From this image, the titular sunrise is clearly visible, but the boats are not. When

these images are viewed together, however, driving both cone-opponent and non-opponent signals, a much broader context is revealed: not only the boats and the sunrise, but an industrial harbor complete with cranes and smokestacks (**Figure 1.1**C).

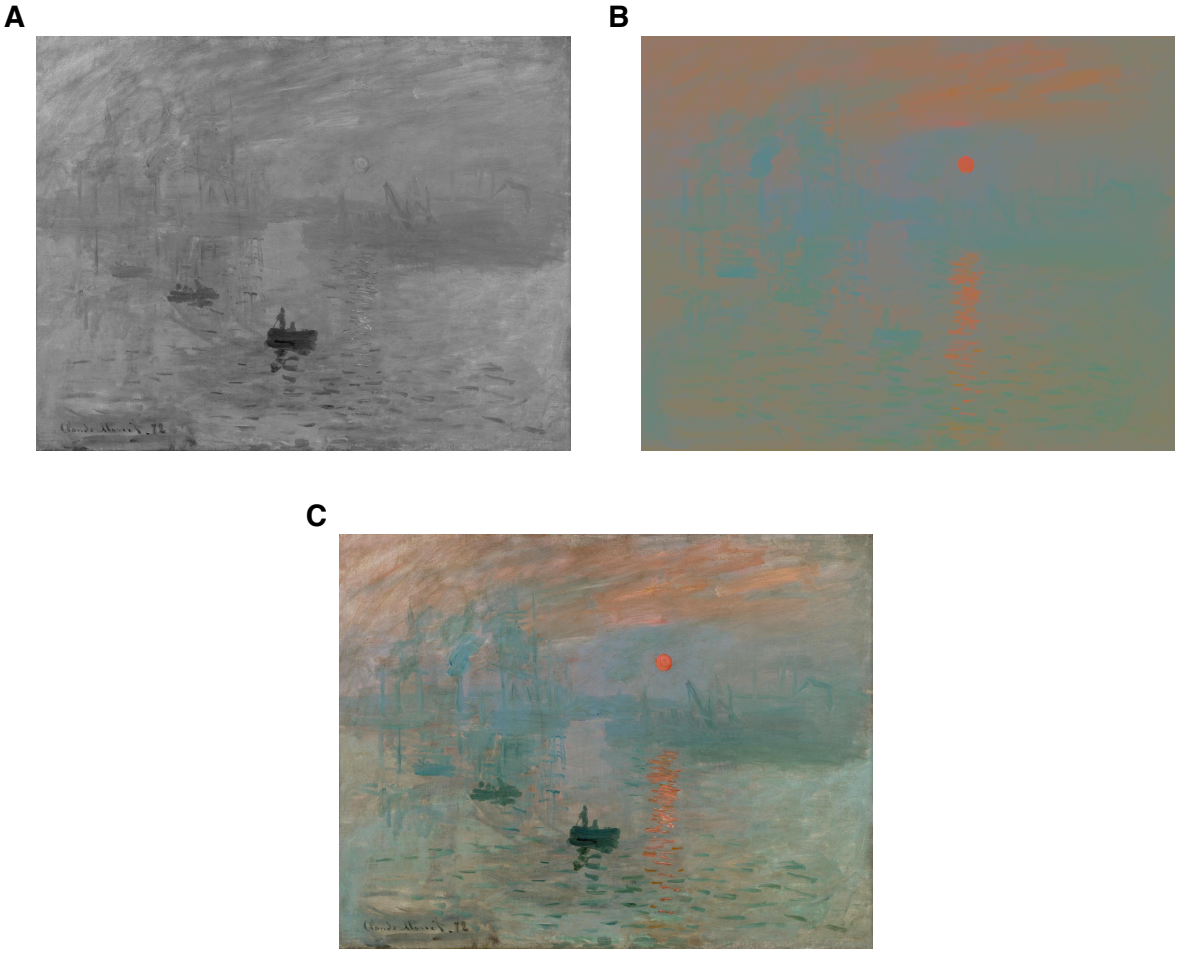

Figure 1.1 Claude Monet's *Impression, Sunrise*, 1872.
(A) Modified version of Monet's painting containing only changes in brightness. (B) Modified version of Monet's painting containing only changes in color. (C) Original paining, superposition of (A) and (B).

Although the propagation of photoreceptor activity is well understood in the pre-cortical visual system, the processing of these signals in the primary visual cortex (area V1) remains largely qualitative. We know that some V1 neurons respond to color but not to luminance, others to luminance but not to color, and still others to both color and luminance (Lennie, Krauskopf et al.

1990, Johnson, Hawken et al. 2001, Johnson, Hawken et al. 2004). A more comprehensive understanding of luminance and chromatic processing in V1 derives from rigorously tested quantitative models that map changes in photoreceptor activity to spiking responses in V1 neurons. It was the goal of my doctoral research to develop such a quantitative model.

In the paragraphs that follow, I will discuss some of the potential shortcomings in previous studies of color and luminance processing in V1—including spatiotemporal stimulus structures, stimulus distributions and color spaces, and the considered quantitative models—and how I have attempted to overcome these shortcomings in my own work.

1.1 SPATIOTEMPORAL STIMULUS STRUCTURE

A neuron's receptive field (RF) describes its sensitivity to changes in light over space and time. The spatiotemporal structure of a V1 RF is typically characterized by a central region, called the classical RF, which is sensitive to increments and/or decrements of cone-signals in some spatial and temporal arrangement, and a surrounding region, called the extraclassical surround, which provides a context for the activity in the central region—that is, how the light in the classical RF resembles or differs from light in the surrounding regions. Neurons with different sensitivities to cone signals in these regions will respond to different types of stimuli. For example, a V1 neuron may respond to a particular spatial arrangement of lights within the classical RF, such as a black bar next to a white bar, but not when their arrangement is reversed; another V1 neuron may respond to a small colored patch in the center of its RF, but not to a full-field stimulus of the same color. The mechanisms underlying these different responses can be binned into two categories: influences from within the classical RF and influences from beyond the classical RF.

Many studies have probed neural responses using stimuli that excite the whole classical RF simultaneously with both increments and decrements of light, such as drifting sinewave gratings (Lennie, Krauskopf et al. 1990, De Valois, Cottaris et al. 2000, Johnson, Hawken et al. 2004, Solomon and Lennie 2005, Horwitz and Hass 2012) or a spatiotemprochromatic white noise (Rovamo and Virsu 1984, Horwitz, Chichilnisky et al. 2005, Horwitz, Chichilnisky et al. 2007). Building quantitative models from the responses to these multicolored stimuli can be difficult because responses to light increments cannot be distinguished from responses to light decrements, which is important for determining which neurons contribute to hue discrimination. For example, a neuron that responds to both red and green stimuli can contribute to contrast detection but not to hue discrimination.

Many of these same studies have probed neural responses using stimuli that extend well beyond the classical RF into the extraclassical surround (Thorell, De Valois et al. 1984, Deluca, Balsevich et al. 1987, Lapp, Tyler et al. 1987, Pointer, Dean et al. 1989, Lennie, Krauskopf et al. 1990, Cottaris and De Valois 1998, De Valois, Cottaris et al. 2000, Hanazawa, Komatsu et al. 2000, Johnson, Hawken et al. 2004, Solomon and Lennie 2005, Webb, Dhruv et al. 2005). Such stimuli excite a wide network of V1 neurons that integrate signals across distant regions of visual space via lateral connections and feedback. Building quantitative models from the responses to these large stimuli is difficult because stimulation of a neuron's extraclassical surround has been observed to affect its color-tuning in complicated ways (Solomon, Peirce et al. 2004).

In my experiment, I attempted to stimulate each V1 neuron in as simple a regime as possible by avoiding complicating interactions from within the classical RF and from beyond. To do so, I stimulated each V1 neuron with uniformly colored flashes (to measure responses to light increments and decrements independently) that were tailored to fall within its classical RF

(avoiding the potentially complicating extraclassical surround) and, when possible, to individual subunits in isolation (e.g. ON- or OFF-) within the classical RF (avoiding potentially complicating subunit interactions).

1.2 COLOR SPACES

The color space in which stimulus-response pairs are analyzed can have a profound effect on our understanding of the relationship between stimuli and responses. This fact has been demonstrated in the evolution of models in the field of psychophysics. Early investigators of contrast detection and color discrimination used stimuli that were closely tied to the physical properties of light, such as monochromatic (single-wavelength) lights (Deluca, Balsevich et al. 1987, Tyler, Cowell et al. 1987). Although single-wavelength lights were intended to be relatively simple stimuli, the models of contrast detection derived from these stimuli were not straightforward. A great deal of progress was achieved when stimuli were specified in terms of the effects on the activity of photoreceptors. In cone excitation space, contrast detection could be described using a few linear mechanisms (Deluca, Balsevich et al. 1987, Lapp, Tyler et al. 1987, Roth, Unanue et al. 1987, Cole, Hine et al. 1993, Sankeralli and Mullen 1996, Giulianini and Eskew 1998). Further progress was achieved using stimuli that were specified in relative changes in photoreceptor activity, or cone contrast. In cone contrast space, contrast detection could be described under a variety of conditions using only a few linear mechanisms with fixed weights (Virsu, Rovamo et al. 1982, Leblanc, Yamamoto et al. 1987, Tyler 1987, Kalloniatis and Harwerth 1991).

Advances in V1 color neurophysiology have also been accompanied by the use of different color spaces. Early experiments in V1 were also conducted using monochromatic lights (Thorell, De Valois et al. 1984, Deluca, Balsevich et al. 1987, Cottaris and De Valois 1998). Later

experiments used stimuli that were specified according to their effect on three linear mechanisms derived from the psychophysical experiments above, e.g. Derrington-Krauskopf-Lennie color space, which defines stimuli based on their presumed effect on three presumed psychophysical detection mechanisms (Pointer, Dean et al. 1989, Lennie, Krauskopf et al. 1990, Cottaris and De Valois 1998, De Valois, Cottaris et al. 2000). More recent work has used stimulus distributions that are drawn from color spaces that are nonlinearly related to cone modulations, e.g. the CIE-xy chromaticity space (Hanazawa, Komatsu et al. 2000), but such color spaces are not clearly advantageous for characterizing the responses of neurons so early in the visual system as V1.

The goal of my thesis work was to develop a set of quantitative models that map changes in L- and M-cone activity to neural responses in V1. Therefore, it was necessary to specify stimuli in a color space more closely tied to photoreceptor activity than to perception. In particular, I represented stimuli in cone contrast space because this space had been so fruitful for developing quantitative models of psychophysical contrast detection and provides an empirically accurate approximation of the cone signals in V1, where neurons adapt to the background light level.

1.3 STIMULUS DISTRIBUTIONS

The selection of an appropriate color space is important when building quantitative models of color vision, but it must be accompanied with an appropriate stimulus distribution within that space. Psychophysical models of color vision were advanced by dense sampling of different color spaces, which allowed experimenters to rigorously test and compare different quantitative models of luminance and chromatic contrast detection (Krauskopf, Williams et al. 1982, Pointer and Hess 1989, Cole, Hine et al. 1993, Sankeralli and Mullen 1996). Most electrophysiology experiments in V1, on the other hand, typically use fewer than ten different stimuli to characterize color tuning

(Lennie, Krauskopf et al. 1990, Johnson, Hawken et al. 2001, Johnson, Hawken et al. 2004). This sparsity of sampled stimuli has hampered the development and refinement of appropriately specified quantitative models that can accurately map photoreceptor activity to neural responses.

A fundamental problem in the sampling of color spaces is that they are more than 1-dimensional, and it can be difficult to adequately sample enough color directions and stimulus contrasts. For example, let us consider a neuron that we know only receives input from L-cones. To uncover this neuron's sensitivity to L-cone modulation, we want to sample the relevant color space uniformly. This task is relatively straightforward in this 1-dimensional stimulus space, where uniform sampling means equally spaced L-cone contrast levels (**Figure 1.2**A). Instead, let us consider a neuron that receives input from both the L- and S-cones. Uncovering the color tuning of this neuron is more difficult, as uniform sampling in this 2-dimensional stimulus space is less clear. One possible uniform sampling scheme is a cartesian grid, with equally spaced L- and S-cone contrast levels (**Figure 1.2**B). Another possible uniform sampling scheme is a radial grid, with stimuli at equal distances from the origin along equally spaced color directions (**Figure 1.2**C). This problem of uniform sampling is compounded when we consider that L-cones outnumber S-cones by 5:1 (Tyler, Qualls et al. 1987, Roorda, Metha et al. 2001). If sensitivity to cone modulations scales with the population of each cone-type, then it would be necessary to sample S-cone modulations differently than L-cone modulations. A simple solution is to linearly scale the S-cone contrast levels of all stimuli in the distribution (**Figure 1.2**D-E), but this biases the distribution of color directions, which become tightly packed around the S-cone axis. An alternative is to keep color directions fixed and scale distance from the origin (**Figure 1.2**F), but this sacrifices the number of high contrast stimuli in the distribution. To make matters worse, the problem of uniform sampling is increased exponentially in a full 3-dimensional color space.

Uncovering the response properties of a neuron is difficult, especially when the selection of an appropriate stimulus distribution is not straightforward, as in the case of color. By probing neural responses with stimuli that are too sparse or inappropriately spaced, it can be difficult to uncover a meaningful estimation of color tuning. In Chapter II of my thesis, I discuss the estimation of color tuning using various stimulus distributions and analysis techniques. In Chapter III, I discuss the importance of measuring neuronal responses using a broad stimulus distribution, especially in absence of prior information of color tuning and contrast sensitivity. In this work, I probed the responses of V1 neurons using a stimulus distribution with a broad assortment of color directions and contrast levels, which was necessary for the development, refinement, and testing of quantitative models of cone signal combination.

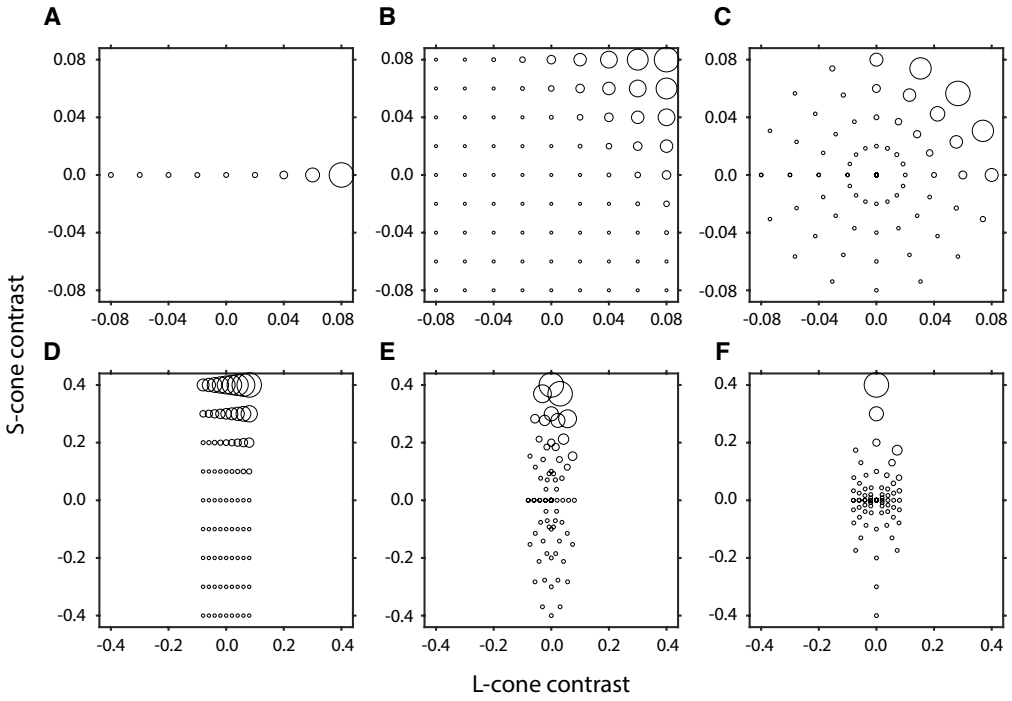

Figure 1.2 Stimulus distributions in the LS plane of cone contrast space.
Uniform distribution for L-cone-only stimulation. (B) Cartesian uniform sampling of LS plane. (C) Radially uniform sampling of LS plane. (D) Linearly scaled cartesian sampling of LS plane. (E) Linearly scaled radial sampling of LS plane. (F) Nonlinearly scaled sampling of LS plane.

1.4 QUANTITATIVE MODELS

Quantitative models play an important role in the effort to understand the visual system. They reflect our underlying assumptions about the role of each stage in the visual processing stream, and allow experimenters to summarize, predict, and compare responses between individual neurons or across neural populations. Without models, we would be forced to characterize neuronal tuning (i.e. preferences in orientation, spatial frequency, color) with an endless catalogue of responses to every stimulus imaginable. Instead, neuronal tuning (including color tuning) can be characterized using quantitative models that predict the responses of a neuron to a broad array of stimuli by sampling its responses to a small subset. Models also allow us to survey neuronal tuning across a population by comparing fitted parameter values, which are especially useful when the parameters have an intuitive interpretation (i.e. preferred orientation, preferred spatial frequency, or preferred color direction).

The development of quantitative models is advanced by comparing response predictions to actual neuronal responses and iteratively refining the model parameterization to achieve more accurate predictions. Here, I will here introduce a historically important model—the linear-nonlinear cascade (LN) model—that has shaped the field's understanding of V1. I will discuss the various ways in which this model has been used to describe neuronal responses in V1, the shortcomings in these fitting techniques, and some of the difficulties in recognizing when this model may not be appropriate.

The LN model, which has a long history in both color physiology and psychophysics, has profoundly shaped our understanding of cone signal representation in V1. As the name suggests, this model is characterized by a linear combination of cone signals followed by a (potentially) nonlinear contrast response function:

$$R = f\left([L\ M\ S]\begin{bmatrix}A\\B\\C\end{bmatrix}\right) \tag{1.1}$$

in which cone signals [L M S] are weighted by fixed amounts [A B C] and summed. The weighted sum is passed through a (potentially) nonlinear function *f*, which may capture realistic biological features such as spike-rate saturation, and is typically parameterized as an exponential, full- or half-rectification, or a Naka-Rushton. To model the variability in real neural responses, the output of the contrast response function, R, is passed through a noise function.

The linear portion of the LN model reduces the three-dimensional cone signals [L M S] to a single scalar by weighting and summing; this scalar weighted sum is called the generator signal (Chichilnisky, 2001). The weights [A B C] specify how strongly each cone-type drives the response of the neuron, and their sign indicates whether the neuron is sensitive to increments or decrements in the activity of that cone-type. The ratio of these weights specifies a direction through color space in which the neuron is most sensitive to modulation. This direction is called the neuron's preferred direction and is a common summary of color tuning in V1 (Johnson, Hawken et al. 2004).

Estimating the linear weights, however, can be difficult because neither the linear portion nor the nonlinear portion of the LN model may be observed in isolation: parameter estimates are obtained from neural responses, and neural responses are a product of both portions of the model (that is, if the model is accurate). Therefore, despite attempts to do so (Lennie, Krauskopf et al. 1990, Johnson, Hawken et al. 2004, Solomon and Lennie 2005), we generally cannot use standard linear regression to estimate the linear weights of the LN model in the domain of color.

Other strategies for estimating a neuron's preferred color direction also tend to be unreliable. One method is to estimate the preferred direction by recording responses to an assortment of stimuli and using one that evoked the greatest response from the neuron as a proxy

for the neuron's preferred color direction (Stoughton and Conway 2008, Bohon, Hermann et al. 2016). Such estimates are unlikely accurate because the potential preferred directions are restricted to the sampled color directions. More importantly, estimates obtained in this manner may be confounded by the various intensities of the considered stimuli. For example, the gamut of most display devices is extensive in some color directions and constricted in others; using high contrast stimuli on such devices, estimates of preferred color direction will tend toward color directions that can achieve higher contrasts.

As expressed above, one great advantage of using a quantitative model is that we can estimate a neuron's responses to a broad set of stimuli by measuring its responses to a few. The flip side to this utility, however, is that identifying violations of the presumed model is difficult from the responses to a small set of stimuli. Further, if the LN model is not an appropriate description of all V1 neuronal responses, then fitting this model to the responses of such neurons may produce misleading patterns in the data, or even obscure genuine structure in the tuning of neurons for which this model *is* appropriate. In my thesis work, I attempted to identify violations of the classical LN model by recording the responses of V1 neurons to a broad array of stimuli that modulated L- and M-cones in a multitude of ratios and intensities. Further, I considered broader classes of models that could provide more accurate descriptions of the observed responses.

1.5 OUTLINE

In the second chapter of this thesis, I discuss some difficulties in estimating preferred directions that can arise due to the stimulus distributions and the color spaces in which stimulus-response pairs are analyzed. I discuss several techniques for estimating preferred directions and their associated shortcomings and advantages. Additionally, I compare the accuracy of the

preferred directions estimated by these analysis techniques using both modeled and real neuronal responses. In the third chapter, I present data from a new set of experiments in which I measured the responses of V1 neurons to uniformly colored flashes presented within the classical RF. I used these responses to develop a family of quantitative models that accurately describe the breadth of color tuning that I observed in V1 and discuss the patterns that I observed across the population.

Chapter 2. MEASUREMENTS OF NEURONAL COLOR TUNING: PROCEDURES, PITFALLS, AND ALTERNATIVES

2.1 INTRODUCTION

Measuring the color tuning of visual neurons is important for understanding the neural basis of vision, but it is challenging because of the inherently three-dimensional nature of color. Color tuning cannot be represented by a one-dimensional curve, and measuring three-dimensional tuning curves is difficult. One approach to addressing this challenge is to analyze neuronal color tuning data through the lens of mathematical models that make assumptions about the shapes of tuning curves. In this paper, we discuss the linear-nonlinear cascade model as a platform for measuring neuronal color tuning. We compare fitting this model by three techniques: two using response-weighted averaging and one using numerical optimization. We highlight the advantages and disadvantages of each technique and emphasize the effects of the stimulus distribution on color tuning measurements.

In humans, apes, and Old World monkeys, color vision is mediated by the long (L), medium (M), and short (S) wavelength-sensitive cone photoreceptors. Signals from these three cone classes are combined in the retina, sent to the brain, and propagated through a complex processing hierarchy of recurrently connected visual areas. At each stage of this hierarchy, signals from the preceding stages are mixed to create new signals. Signals in the early stages are closely related to the physical properties of light, whereas signals in the later stages are more closely related to perception (Komatsu 1998, Gegenfurtner 2003, Solomon and Lennie 2007, Conway 2009, Conway, Chatterjee et al. 2010, Bohon, Hermann et al. 2016). Neuronal color tuning measurements can help reveal how and where these transformations occur.

2.1.1 *A Problem in Color Neurophysiology*

Color tuning measurements made in different laboratories are not always comparable. Different laboratories typically use different stimuli, and comparisons are based on data summaries, such as distributions of preferred colors or cone weights. In an ideal world, these summaries would not depend on an experimenter's choice of stimuli, but in reality they usually do. For example, differences in spatiotemporal stimulus parameters affect color tuning (Derrington and Lennie 1984, Thorell, De Valois et al. 1984, Cottaris and De Valois 1998, Solomon, Peirce et al. 2004, Conway and Livingstone 2006). These effects are important but beyond the scope of this article; we focus on how the spectra of lights presented in neurophysiology experiments affect measurements of color tuning.

To illustrate the type of problem we are considering, we analyze an example data set. We probed individual neurons in the primary visual cortex (V1) of a macaque with two types of white noise stimuli. In the *phosphor noise* stimulus, the three display primaries modulated independently (Figure 2.1A). In the *cone noise* stimulus, they modulated in ratios selected to stimulate the three cone types independently (Figure 2.1B) (see section 2.2.1 for methodological details of the experiment). The average phosphor noise stimulus that preceded a spike from a single example neuron (Figure 2.1C) appears different from the average cone noise stimulus that preceded a spike (Figure 2.1D); but do these two images reflect the same color tuning? We will return to this example neuron twice more as we present three techniques to estimate color tuning and discuss the transformation of estimates between color spaces. To begin, we discuss the model upon which the analysis techniques are based: the linear-nonlinear (LN) cascade.

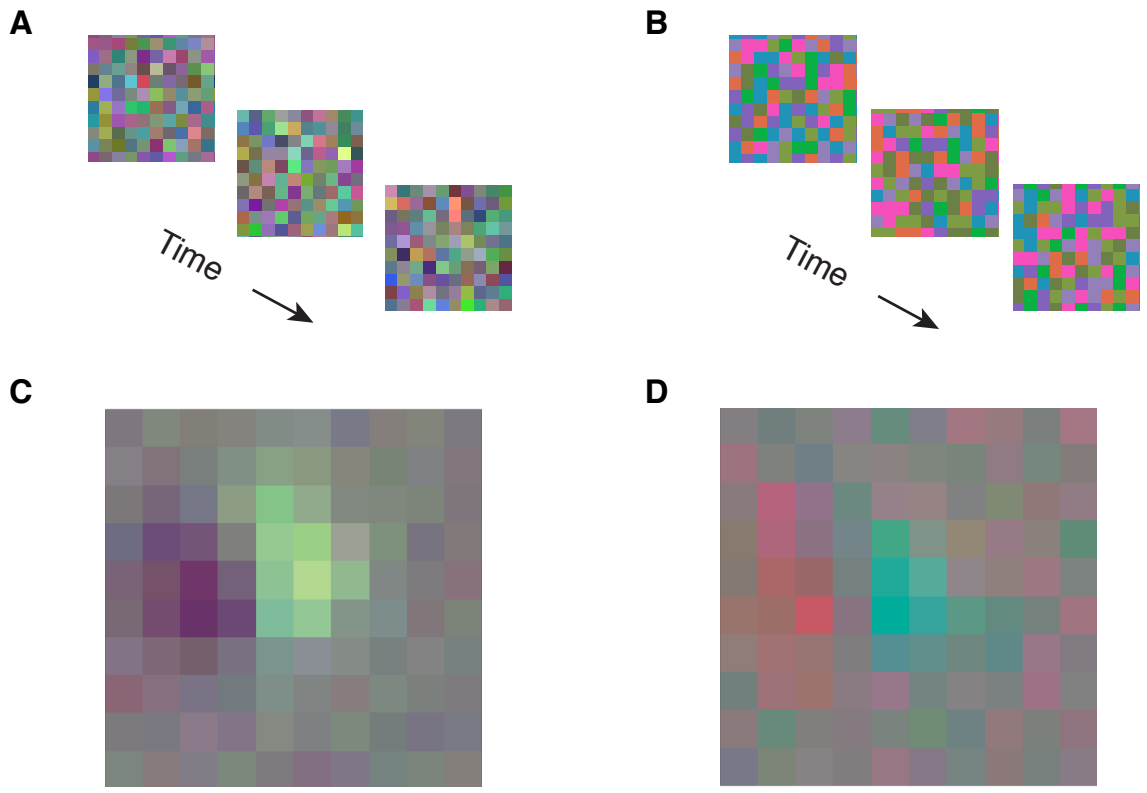

Figure 2.1. The responses of individual V1 neurons from an awake fixating monkey were probed with stimuli from two different distributions.
(A) Stimuli from the phosphor noise distribution. The color of each square was determined by the independent modulation of the three phosphors in a CRT monitor. (B) Stimuli from the cone noise distribution. Each colored square modulated the activity of the three photoreceptor-types independently. (C) The average stimulus that preceded a spike during the presentation of the phosphor noise. (D) The average stimulus that preceded a spike from the same neuron during the presentation of the cone noise.

2.1.2 *The Linear-Nonlinear Cascade Model*

Models have a central role in color neurophysiology. Without them, color tuning measurements are simply a collection of numerical tables that map stimuli to responses—tables that do not predict responses to untested stimuli and that can never be sufficiently large to characterize tuning fully. In this paper, we focus specifically on cone signal combination under the LN model. This class of model has a long history in the field of color psychophysics, specifically

in the areas of contrast detection (Krauskopf, Williams et al. 1982, Pointer and Hess 1989, Cole, Hine et al. 1993, Sankeralli and Mullen 1996), color discrimination (Ingling and Tsou 1977, Luke, Begun et al. 1989, Pointer 1989, Eskew, McLellan et al. 1999), and appearance judgments (Hurvich and Jameson 1957, May and Tyler 1987, Tyler 1987). Some of the issues we discuss regarding color spaces and transformations between them can be found in Cowell, Tyler et al. (1987), Dodd, Losco et al. (1987), Tyler (1987), Tyler, Cowell et al. (1987), Tyler, Qualls et al. (1987), Brainard (1996), Sankeralli and Mullen (2001). Here, we distill from this body of work a few key elements that are particularly useful for color neurophysiology. We apply these techniques to neurophysiological data, compare their accuracy for measuring neuronal color tuning, and provide intuitions into their respective strengths and weaknesses.

The LN model provides a simple but powerful description of neural responses. As the name suggests, this model has both a linear and a nonlinear component. The linear component specifies how stimulus elements are weighted and summed:

$$[v_1\ v_2\ v_3]\begin{bmatrix}A\\B\\C\end{bmatrix} \tag{2.1}$$

where the vector [v_1 v_2 v_3] represents three stimulus elements (e.g. signals from the three cone types), and the weighting vector [A B C] describes how the elements are combined. The sign and magnitude of each element in the weighting vector indicates how the corresponding element in the stimulus vector contributes to or detracts from the neuron's response. Input to the LN model can be represented by any number and kind of stimulus elements (e.g. cone excitations, DKL mechanism modulations, CIE chromaticity coordinates). The nonlinear component transforms the weighted sum into neural responses:

$$R = f\left([v_1\ v_2\ v_3] \begin{bmatrix} A \\ B \\ C \end{bmatrix} \right) \tag{2.2}$$

where *R* typically represents a spike count or spike rate. The function *f* captures nonlinear response features like spiking thresholds, spike-rate saturation, and rectification.

The LN model simplifies the problem of measuring color tuning, fundamentally, by reducing multiple stimulus elements to a single scalar through weighting and summing. The weighting vector describes this simplification by specifying a neuron's sensitivity to each element. The direction of the weighting vector though color space—the neuron's preferred direction—is a useful summary of color tuning. In contrast, the magnitude of the weighting vector is less informative. Any uniform scaling of the weights can be compensated exactly by scaling of the domain of *f*, so a family of LN models exists that make identical response predictions despite each having a different *f* and correspondingly different weights. For this reason, weighting vectors are often normalized so that their absolute values sum to one (Derrington, Krauskopf et al. 1984, May and Tyler 1987, Lennie, Krauskopf et al. 1990, Johnson, Hawken et al. 2004, Conway and Livingstone 2006, Horwitz, Chichilnisky et al. 2007).

2.2 Materials and Methods

2.2.1 *Electrophysiological Methods*

We recorded from 35 V1 neurons in an awake fixating rhesus monkey to interleaved white noise modulation from two different stimulus distributions. All experiments were done in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals and the Code of Ethics of the World Medical Association (Declaration of Helsinki).

The white noise stimulus comprised a 10 x 10 grid of 0.1° squares whose color changed randomly on every screen refresh (75 Hz). The color of each square during a given trial was drawn

from one of two distributions. The *phosphor noise* stimulus modulated each of the three display primaries independently (Figure 2.1A), creating a Gaussian distribution of stimuli that was RS in phosphor space. The *cone noise* stimulus modulated each cone-type independently between a high and low value (Figure 2.1B). Each stimulus was equidistant from the origin in *proportional cone-contrast space*, in which cone contrast is scaled according to the relative proportion of each cone-type in the retina, such that:

$$Lcc * Lcp \approx Mcc * Mcp \approx Scc * Scp \tag{2.3}$$

where *Lcc*, *Mcc*, *Scc* is the maximum L-, M- and S-cone contrast (0.09, 0.09, 0.4), and *Lcp*, *Mcp*, and *Scp* are the approximate proportions of L-, M-, and S-cones in the macaque retina (0.45, 0.45, 0.1). Both stimulus distributions had equal mean luminance and chromaticity, which was identical to the background.

For each neuron tested, two response-weighted averages (RWAs) were computed: one using the responses from cone noise trials and one using the responses from phosphor noise trials. Responses in these experiments were represented as the number of spikes occurring on each stimulus frame, measured after an estimated response latency.

To estimate a single preferred color direction from each RWA, we used linear regression to find a 3-element vector that describes the consensus color tuning across stimulus pixels and time-delays (Horwitz, Chichilnisky et al. 2005). Thus, the preferred color direction of each neuron was represented by a pair of 3-element vectors: one describing the weights in phosphor intensity space and the other describing the weights in proportional cone-contrast space. These weight estimates were then transformed to cone-contrast space by matrix multiplication, as per Table 2.1.

2.2.2 *Simulation Methods*

Modeled responses were generated according to an LN model (eq. 2.2) in which the nonlinear function *f* was a Naka-Rushton function:

$$\lambda = U * \frac{contrast^N}{{c_{50}}^N + contrast^N} + bl \tag{2.4}$$

where U is the peak response (the upper asymptote), bl is the baseline response (the lower asymptote), N is an exponent, *contrast* is the dot product between the stimuli [L M] and a unit-length weighting vector $\begin{bmatrix} A \\ B \end{bmatrix}$, and c_{50} is the contrast at which the response reaches half maximum. Noise was added to these modeled responses by passing the output of the Naka-Rushton function through a Poisson random number generator:

$$R = Pois(\lambda) \tag{2.5}$$

Each of 33 simulated neurons was tuned to a unique direction in the LM plane and was probed with three stimulus distributions, each comprising 64 unique stimuli. Each unique stimulus was presented 5 times. The upper asymptote (U), baseline (bl), and exponent (N) were fixed across datasets ($U = 50$, $bl = 0$, $N = 3$), and the c_{50} was defined for each neuron to be in the middle of the range of tested stimulus contrasts. For each simulated neuron and stimulus distribution, 100 datasets were generated and analyzed to estimate the neuron's preferred color direction.

To estimate preferred directions using maximum likelihood, each dataset was fitted with a Naka-Rushton function in which U, c_{50}, bl, N, and the weighting vector $\begin{bmatrix} A \\ B \end{bmatrix}$ were free to vary. The fitted values for each dataset were those that maximized the Poisson likelihood of the parameters given the responses R:

$$\hat{\theta} = argmax\big(\mathcal{L}(R_1 \dots R_n|\lambda_1 \dots \lambda_n)\big) = argmax\left(\prod_{i=1}^{n} \frac{{\lambda_i}^{R_i} * e^{-\lambda_i}}{R_i!}\right) \quad (2.6)$$

in which n is the total number of stimuli, R_i is the response to the i^{th} stimulus, λ_i is the predicted response to the i^{th} stimulus, and $\hat{\theta}$ is a vector of the best fitting parameter values (comprising U, bl, N, c_{50}, and $[A\ B]$). For practical reasons, we minimized the equivalent log-likelihood function:

$$\hat{\theta} = argmin\big(-\mathcal{L}(R_1 \dots R_n|\lambda_1 \dots \lambda_n)\big) = argmin\left(-\sum_{i=1}^{n} R_i * (\lg(\lambda_i) - \lambda_i)\right) \quad (2.7)$$

2.3 RESULTS

2.3.1 *Technique 1: Response Weighted Averaging*

Response weighted averaging is a data analysis technique that can provide valuable insight into how neurons represent visual stimuli. The response-weighted average stimulus is described mathematically as:

$$RWA = \frac{1}{n}\sum_{i=1}^{n} R_i[v_1\ v_2\ v_3]_i \quad (2.8)$$

where n is the total number of tested stimuli, $[v_1\ v_2\ v_3]_i$ is the i^{th} stimulus, and R_i is the response to the i^{th} stimulus. The elements of the RWA, like those of the weighting vector, reflect the influence of each visual signal on the response of the neuron. The average stimuli in Figure 2.1C and Figure 2.1D are examples of RWAs.

The RWA can be used to estimate a neuron's preferred color direction (Leblanc, Yamamoto et al. 1987, Cottaris and De Valois 1998, Sun, Smithson et al. 2006, Horwitz, Chichilnisky et al. 2007), but its accuracy depends on the stimulus distribution. This estimate is biased if the stimulus distribution is asymmetric (e.g. distended or unequally sampled). To visualize this bias, consider a hypothetical neuron that receives exclusively L-cone input. The stimuli that evoke the largest responses from such a neuron will strongly modulate the L-cones. If

such a neuron were probed with a stimulus distribution in which L- and M-cone signals are positively correlated (Figure 2.2A), then the stimuli that evoke the largest responses (and therefore dominate the RWA) will also strongly modulate the M-cones. The RWA (star) reflects this correlation, and, consequently, does not align to the L-cone axis, which is the neuron's preferred color direction (arrow).

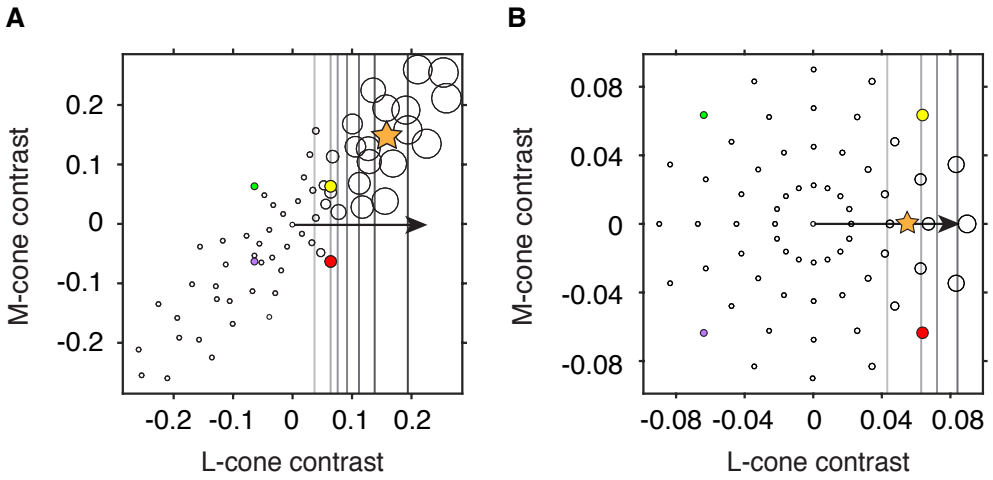

Figure 2.2. Responses of a hypothetical neuron that receives only L-cone input.
For simplicity, only L- and M-cone modulations are simulated in this example. The position of each circle represents a stimulus in the LM plane. The size of each circle represents the magnitude of the corresponding response. Contour lines, from light to dark colors, represent the increasing responses of the neuron to stimuli of progressively higher contrast. The arrow indicates the preferred color direction of the neuron, and the orange star represents the response-weighted average stimulus (RWA). (A) A distended stimulus distribution in which L- and M-cone modulations are positively correlated. The RWA does not align with the neuron's preferred direction. (B) A radially symmetric stimulus distribution of L- and M-cone modulations. The RWA aligns with the preferred direction. For reference, the purple, green, yellow, and red points represent identical lights in each distribution.

The RWA is an unbiased estimator of the preferred direction when the stimulus distribution is radially symmetric (RS) (Figure 2.2B) (Chichilnisky 2001). Intuitively, this is because there exist whole families of stimuli that give rise to the same weighted sum, and therefore, the same response. Under the LN model, all such stimuli describe contours through color space that are

straight and perpendicular to the preferred direction (colored lines). When stimuli that drive the same response are distributed symmetrically about the preferred direction (arrow), their average (star) aligns with the preferred direction. In most experiments, the preferred direction of a neuron is unknown a priori, so the stimulus distribution must be symmetric in all directions to guarantee an unbiased estimate[1].

Radial symmetry of a stimulus distribution depends on how the axes of the color space are defined. Stretching one axis of the color space distends the stimulus distribution along that axis, biasing the RWA in the direction of the stretch. This problem might appear formidable, since no single set of stimuli nor representation thereof is universally accepted as the best for measuring color tuning, and no distribution of stimuli is RS in every color space. Fortunately, an estimate of a neuron's preferred color direction, obtained in any color space by any means, can be transformed to any linearly related color space in a straightforward way. To demonstrate this procedure and to provide a practical application, we discuss how stimuli and weights are transformed between color spaces below.

2.3.1.1 The Representation and Transformation of Lights and Weights

The preferred direction of a neuron can be estimated in one color space (e.g. where the stimulus distribution is RS) and transformed into any linearly related space (e.g. where the distribution is not RS). This transformation can be demonstrated using the two color spaces introduced in Figure 2.1: a phosphor space and a cone space. In both spaces, the origin [0 0 0]

[1] In practice, stimulus distributions used in experiments are usually discrete and therefore can only approximate true radial symmetry. In this case, the denser the sampling in the stimulus space, the closer this approximation can become. Continuously varying stimuli can also be used to create a radially symmetric (RS) distribution, as in Sun, H. et al (2006). Stimuli that are presented in rapid succession or in close proximity, such that they are effectively averaged together in the visual system, create signals that approach a Gaussian distribution, which can always be made RS with a linear transformation.

represents the background of the display, which is also the average of the stimulus distribution. Changes in phosphor intensity relative to this light level are represented with positive and negative values, as are changes in cone excitation. Each light is represented by a single point in each space and can be transformed between them via matrix multiplication:

$$[L\ M\ S] = [R\ G\ B] * M \tag{2.9}$$

where $[R\ G\ B]$ is an nx3 matrix of stimuli represented in phosphor space, $[L\ M\ S]$ is this same collection of stimuli represented in cone space, and M is a 3x3 transformation matrix. Transforming the representation of lights alters their distribution. For example, a stimulus distribution that is RS in phosphor space (Figure 2.3A, gray points) is distended in cone space (Figure 2.3B, gray points) and vice versa (colored squares).

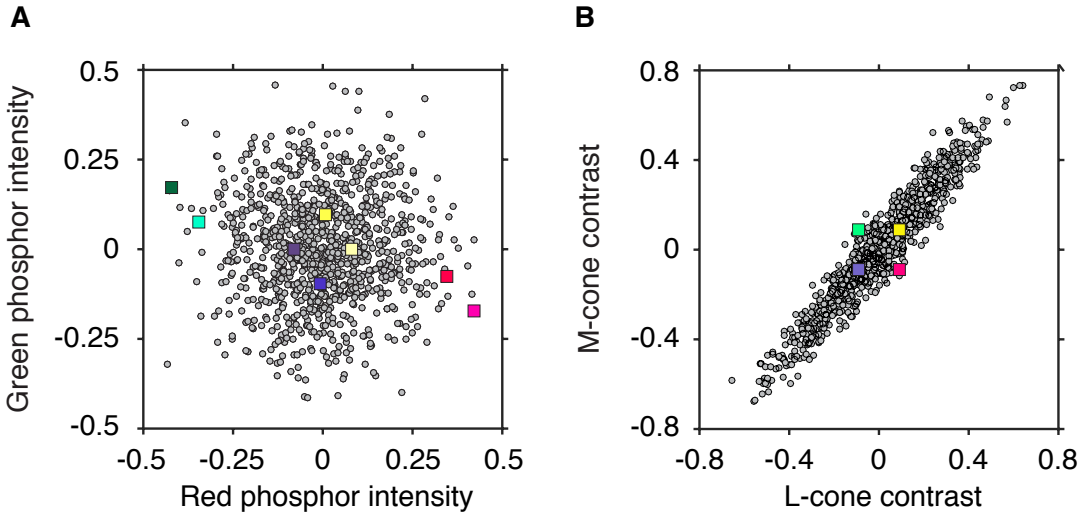

Figure 2.3 Two stimulus distributions in two color spaces.
Phosphor noise stimuli (gray circles) and cone noise stimuli (colored squares) represented in two color spaces. (A) Phosphor and cone noise stimulus distributions projected onto the red-green plane of phosphor space. The phosphor noise distribution is approximately radially symmetric (RS) in this space, and the cone noise distribution is distended. (B) The same two distributions projected onto the LM plane of cone space. In this color space, the cone noise distribution is approximately RS, and the phosphor noise distribution is distended. Only four colored squares are visible because +S- and –S-cone stimuli have an equal projection onto this plane.

Real neural responses do not depend on how an experimenter chooses to represent visual stimuli. For the responses of an LN model neuron to have this property, the weights must be transformed between color spaces such that a given stimulus produces the same weighed sum irrespective of its representation. To achieve this, the transformation of the stimuli must be counterbalanced by the transformation of the weights. If the matrix M transforms lights from phosphor space to cone space (eq. 2.9), then the weights are transformed by its inverse transpose (M^{-T}):

$$[l\ ms\ ] = [r\ g\ b] * M^{-T} \tag{2.10}$$

where $[r\ g\ b]$ and $[l\ m\ s]$ are the weights in phosphor and cone space, respectively, that reflect the same color tuning. Combining equations 2.9 and 2.10 clarifies the relationship:

$$R = [L\ M\ S] * \begin{bmatrix} l \\ m \\ s \end{bmatrix} = [R\ G\ B] * M * M^{-1} * \begin{bmatrix} r \\ g \\ b \end{bmatrix} \tag{2.11}$$

where $M^{-1} * \begin{bmatrix} r \\ g \\ b \end{bmatrix}$ is simply a rearrangement of terms in equation (2.10).

The general principle demonstrated here is that any invertible transformation of the stimuli must be accompanied by a compensatory transformation of the weights, so that weighted sums are invariant to changes in the stimulus representation. Using this principle, we can tabulate the transformations of stimuli and weights between linearly related color spaces (Table 2.1):

Table 2.1. Transformations of stimuli and weights between linearly related color spaces

	Space 1→Space 2 (e.g. phosphor space to cone space)	Space 2→Space 1 (e.g. cone space to phosphor space)
Stimuli $[v_1\ v_2\ v_3]$	M	M^{-1}
Weights $[A\ B\ C]$	M^{-T}	M^{T}

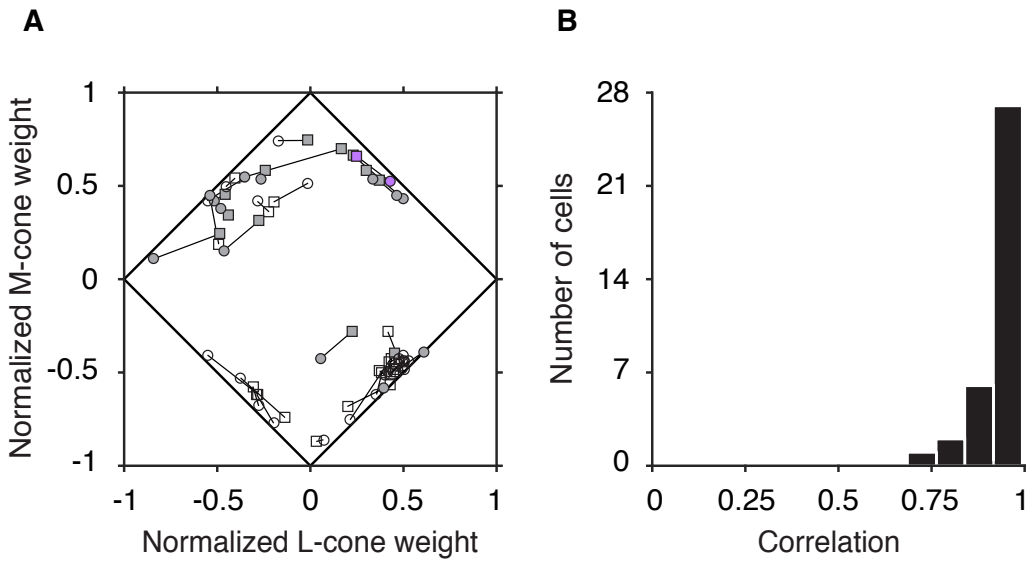

Figure 2.4 Cone weight estimates from two distributions.
(A) Normalized cone weight estimates from the population of neurons recorded in this experiment (n = 35). Two cone weight estimates (connected with black lines) were obtained from each neuron: one from the responses to the phosphor noise (circles), and one from the responses to the cone noise (squares). Shaded symbols indicate positive S-cone weights, and unshaded symbols indicate negative S-cone weights. The example neuron from Figure 1 is highlighted in purple. (B) Correlation coefficients between the normalized cone weight estimates obtained from the two distributions. Cone weight estimates derived from the phosphor noise agreed closely with those derived from the cone noise for the majority of tested neurons.

2.3.1.2 An Experimental Example

We return to the example neuron from Figure 2.1, which was probed with two stimulus distributions, each of which was RS in the color space in which it was constructed (Figure 2.3). We generate two estimates of the preferred color direction: one in phosphor space, and one in cone space. These two color spaces are linearly related, so the two estimates can be transformed from one space to the other, or into a third via equation 2.5. By convention, we represent both estimates in cone contrast space. Both estimates indicate similar color tuning (Figure 2.4A, purple symbols). Similar results were obtained for a larger population of V1 neurons, as captured by the proximity

of each neuron's normalized cone weight estimates (Figure 2.4A, gray symbols) and the high correlations between them (Figure 2.4B).

2.3.2 *Technique 2: Linear Regression*

To estimate the preferred color direction of a neuron whose responses have been probed with a non-RS stimulus distribution, the experimenter must turn to alternative techniques. One special case offers an easy solution: if the stimuli can be linearly transformed into a new color space in which their distribution is RS, the preferred color direction can be estimated in the new space using the RWA, then transformed to the original space using Table 2.1 (multiplication with the inverse-transpose of the stimulus transformation matrix).

This technique is not generally applicable, however, because most non-RS stimulus distributions cannot be made RS via linear transformation. Nevertheless, there is always an approximation that comes close: multiplying the stimuli with a whitening matrix. This multiplication transforms the stimuli so that they have variance of "1" in all directions (their *distribution* need not be the same in all directions, which is why the whitened distribution is not necessarily RS). The preferred color direction of a neuron can be estimated with the RWA of the whitened distribution, then transformed to the original color space. This solution does not necessarily provide an unbiased estimate of a neuron's preferred color direction, but it can provide a substantial improvement over the RWA of the non-whitened stimulus distribution. This technique is identical to linear regression.

2.3.3 *Technique 3: Numerical Optimization*

An even more general approach to the problem of estimating the preferred color direction is to use numerical optimization to iteratively adjust the weights of the LN model to minimize an

appropriate measure of error (e.g. the sum of squared differences between the actual and predicted responses). Numerical optimization is more computationally intensive than linear regression or calculating the RWA. Nevertheless, it can be used to accurately estimate the weighting vector, as well as any additional parameters of the model, largely irrespective of the stimulus distribution. Additional parameters may include those that govern the shape of *f* (the nonlinear component), or weights that represent input from non-stimulus sources (e.g. response history or the firing of other neurons; Pillow, Shlens et al. 2008). Through numerical optimization, all of these parameters can be fit simultaneously, and thus potentially more accurately than if each were estimated sequentially, as when the RWA is used to estimate the preferred color direction and other procedures are used to estimate *f*.

Generally, the goal of this procedure is to describe the data with a model that maximizes (or alternatively minimizes) an *objective function*. The objective function describes the relationship between the parameters of the model (e.g. the weights) and a number that represents how well each set of parameter values describes the data. The objective function may also include penalties for parameter values that are unrealistic or are unlikely to generalize well to new data.

The objective function must be tailored to the experiment, and although many circumstances warrant a particular set of functions, there is often no best choice. One set of objective functions that is well grounded in statistical theory—likelihood functions—represent the probability of observing a set of neuronal responses across all possible choices of model parameter values. The parameter values that maximize the likelihood function identify the model under which the observed responses are most likely to occur. This maximum likelihood estimate can be written:

$$\begin{bmatrix} A \\ B \\ C \end{bmatrix}, \hat{\theta} = argmax\left(\mathcal{L}\left(f\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}, \theta\right)\middle| R\right)\right) \tag{2.12}$$

where [a b c] is a weighting vector, θ is a vector of any additional parameters that f might have, R is a vector of all the observed responses, $\mathcal{L}$ is the likelihood, or the conditional probability of responses R given the candidate weights [a b c] and parameters values θ. [A B C], $\hat{\theta}$ are the parameter values that maximize this likelihood.

The parameter values that maximize likelihood may not be calculable analytically, but may be found using numerical optimization. Although many strategies exist for finding local maxima of the likelihood function, no strategy is guaranteed to produce a singular set of parameter values that maximizes the function globally. Instead, the fitting algorithm must search iteratively through many combinations of parameter values. Because the number of potential parameter combinations is infinite in most scenarios, not every possible combination can be tested; therefore, the possibility usually remains that a better set exists. Despite this shortcoming, an experimenter may confidently fit the data using techniques that avoid local maxima in the likelihood function if the number of parameters is reasonably low. Under restricted conditions, the likelihood function can be proven to have only a single maximum, facilitating fitting models with many parameters (Paninski, Pillow et al. 2007).

2.3.4 *Comparing Estimation Techniques*

To compare the accuracy of preferred direction estimates obtained by maximum likelihood with those obtained by response-weighted averaging, we performed two analyses. First, we used both techniques on the example neuron from Figure 2.1. Second, we used both techniques to estimate the preferred color directions of simulated neurons (see Section 2.2.2 for simulation details).

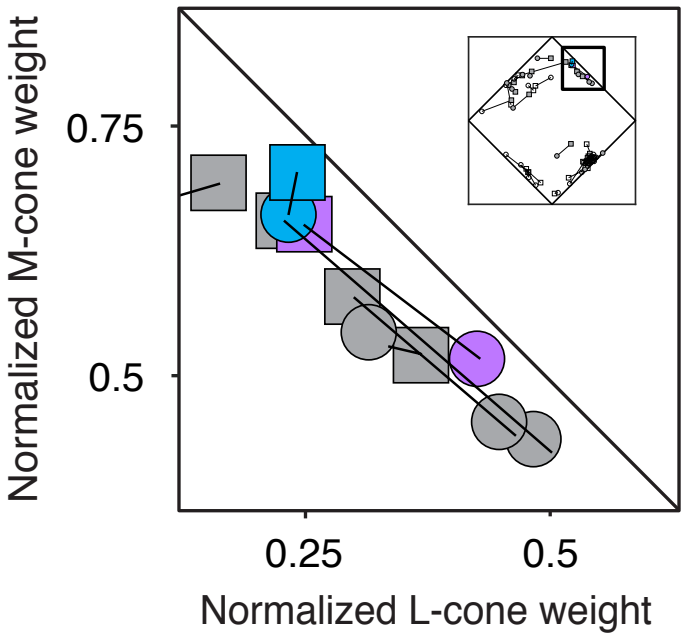

Figure 2.5 Cone weight estimates on two stimulus distributions using two techniques. Normalized cone weight estimates for the example neuron from Figure 1 estimated via the response-weighted average (purple) and maximum likelihood (blue). Normalized cone weight estimates derived from the phosphor noise (squares) agreed closely with those derived from the cone noise (circles) using both techniques.

The example dataset consisted of spike times from a single neuron stimulated with phosphor noise and cone noise. To estimate cone weights, we represented every stimulus as cone contrasts relative to the background and fitted the data with an LN model using numerical optimization (maximum likelihood). Fitted cone weights were similar whether computed from the phosphor noise data (Figure 2.5, circles) or from the cone noise data (Figure 2.5, squares), confirming that maximum likelihood is relatively insensitive to the distribution of stimuli used in the experiment.

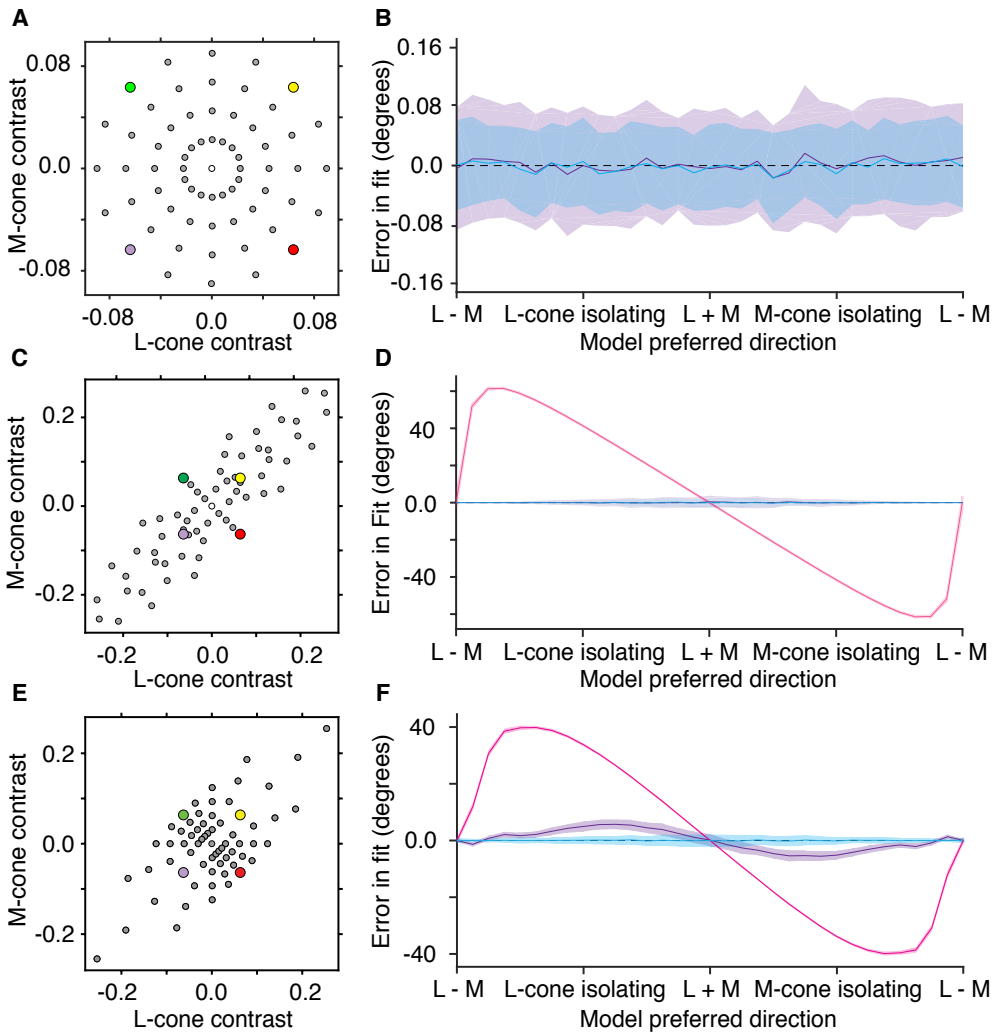

Figure 2.6 Estimates of preferred color direction obtained from the responses of simulated neurons.

Preferred color directions were estimated with three techniques and three stimulus distributions. Gray circles represent stimuli that were presented to the simulated neuron. The purple, green, yellow, and red points represent identical lights in each distribution. For simplicity, only L- and M-cone modulations were simulated. (A) A radially symmetric stimulus distribution. (B) The mean and standard deviation of the error between the true and estimated preferred directions. Estimation via the response-weighted average stimulus (RWA) (purple) and maximum likelihood (blue) were both unbiased. (C) A distended stimulus distribution that was RS when whitened. (D) Estimation via the RWA of the non-whitened distribution (pink) was biased, but estimation via linear regression (purple) and maximum likelihood (blue) were unbiased. (E) A distended stimulus distribution that could not be made RS by linear transformation. (F) Estimation via the RWA (pink) and linear regression (purple) were both biased, but estimation via maximum likelihood (blue) remained unbiased.

In our second analysis, we compared estimates of the preferred directions of simulated neurons using the RWA, linear regression, and maximum likelihood. Each technique was applied to three different stimulus distributions: an RS distribution (**Figure 2.6**A, left), a distended distribution that could be made RS by whitening (**Figure 2.6**B, left), and a distended distribution that could not be made RS by whitening (**Figure 2.6**C, left). For simplicity, only L- and M-cone modulations were simulated. We compared the true preferred direction with the estimated preferred direction obtained with each technique. The RWA estimate was unbiased only when the distribution was RS (Figure 2.6, pink curves). Linear regression was additionally unbiased when whitening the distribution produced radial symmetry (Figure 2.6, purple curves). Maximum likelihood was unbiased in all three cases (Figure 2.6, blue curves).

2.4 DISCUSSION

We described three techniques for measuring preferred color directions: two based on the RWA, and one based on numerical maximization of likelihood. Below, we summarize their advantages and disadvantages. Then, we discuss incorporating cone adaptation into the LN model. Finally, we outline a general strategy for extending the LN model.

The RWA has the advantages that it is easy to calculate and can be used to estimate a neuron's preferred color direction without explicit specification of the parametric form of f (e.g. Naka-Rushton or exponential) or the error model (e.g. Poisson distributed spike counts). Numerical optimization methods have the advantages that they can be used with any reasonable stimulus distribution, a variety of objective functions, and are easily generalizable to models that are more complex than the LN model. The RWA and the maximum likelihood estimate converge if f is strictly monotonic and the stimulus distribution is RS.

2.4.1 *Accounting for Adaptation with Visual Signals and Weights*

The techniques we described are for estimating the parameters of the LN model. They do not guide the selection of stimulus representation. Some stimulus representations can facilitate response descriptions under a narrow range of conditions, others under a broader range. For example, the red-green detection mechanism can be modeled as LN under constant illumination conditions:

$$R = [L\ M\ S]\begin{bmatrix} A \\ B \\ C \end{bmatrix} \tag{2.13}$$

where *L*, *M*, and *S* represent cone excitation differences, or the changes in photoisomerizations caused by a stimulus against a static background (Deluca, Balsevich et al. 1987, Lapp, Tyler et al. 1987, Roth, Unanue et al. 1987, Cole, Hine et al. 1993, Sankeralli and Mullen 1996, Giulianini and Eskew 1998), and *R* represents the response of the red-green detection mechanism.

If the background illumination is changed, however, the cones adapt, thresholds shift, and new weights are needed to describe them. To account for this adaptation, the background illumination must be included in the model. One possibility is to incorporate the background into the weights, yielding a dynamic weighting vector that changes with every background:

$$R = [L\ M\ S]\begin{bmatrix} \frac{A'}{L_0} \\ \frac{B'}{M_0} \\ \frac{C'}{S_0} \end{bmatrix} \tag{2.14}$$

where L_0, M_0, and S_0 represent cone excitations due to the background, and A' = A * L_0, etc. Alternatively, the background can be incorporated into the representation of the stimulus:

$$R = [L\ M\ S]\begin{bmatrix}\frac{A}{L_0}\\ \frac{B}{M_0}\\ \frac{C}{S_0}\end{bmatrix} = \begin{bmatrix}\frac{L}{L_0} & \frac{M}{M_0} & \frac{S}{S_0}\end{bmatrix}\begin{bmatrix}A\\ B\\ C\end{bmatrix} \quad (2.15)$$

Scaling the stimulus by the background changes the representation from cone excitation differences to cone contrasts. In this color space, a single static weighting vector describes detection thresholds under a wide array of illumination conditions (Virsu, Rovamo et al. 1982, Leblanc, Yamamoto et al. 1987, Tyler 1987, Kalloniatis and Harwerth 1991).

Cone contrast is a relatively simplistic way of modeling cone adaptation (for more complex alternatives, see Tyler, Qualls et al. (1987) and Angueyra and Rieke (2013)). This example is not meant to advocate for any particular color space, but to show that cone adaptation can be incorporated into the weights or into the representation of the stimulus, creating two isomorphic LN models. Both provide equivalent descriptions, but accounting for cone adaptation in the stimulus representation isolates receptoral from post-receptoral processes. This example underscores that the relationship between stimuli and responses may appear complex under one representation, but simple under another.

2.4.2 *Beyond the LN Model*

The LN model does not adequately describe neural responses under all conditions or in all color spaces. Some neurons are poorly described by the LN model (Thorell, De Valois et al. 1984, Komatsu, Ideura et al. 1992, Hanazawa, Komatsu et al. 2000, Solomon and Lennie 2005, Bushnell, Harding et al. 2011, Horwitz and Hass 2012). For such neurons, weight estimates are not meaningful. However, the LN model may provide a useful skeleton on which to build richer classes

of models that describe the responses of these neurons more accurately (Roth, Unanue et al. 1987, Rust, Mante et al. 2006, Horwitz and Hass 2012).

Improvements to the LN model will likely spur new advances in color neurophysiology. Finding a class of model that describes color tuning more accurately while being simple enough to work with is not trivial, but complementary approaches may provide leverage on this problem. First, patterns of residuals can be analyzed to find systematic deviations from the predictions of the LN model, and the model can be extended to eliminate these patterns. Second, known neuronal nonlinearities such as contrast gain control and contrast energy calculations can be incorporated into the model. Third, early stages of the visual system can be modeled with greater precision, thereby more accurately constraining how downstream neurons process color signals.

We described stimuli in terms of their effects on the cones, but the techniques we described generalize beyond these descriptions. Deep in the visual system, for example, neurons are poorly described by weighted sums of cone signals. One intriguing possibility is that such neurons may perform LN-like operations but on inputs that are abstract quantities (Cowell, Tyler et al. 1987). Finding stimulus representations that are combined quasi-linearly by neurons but are only distantly related to cone excitations may be a fruitful approach for understanding color tuning in higher-level visual areas.

Chapter 3. PROCESSING OF LUMINANCE AND RED-GREEN CRHOMATRICITY BY MACAQUE V1 NEURONS

3.1 INTRODUCTION

In trichromatic humans and Old World monkeys, luminance and red-green color vision are mediated by the same two photoreceptor classes: the long-wavelength-sensitive (L-) and medium-wavelength-sensitive (M-) cones. Luminance is defined as a weighted sum of L- and M-cone signals (Sharpe 1974, Lennie, Pokorny et al. 1993) and red-green chromaticity is closely related to their difference (Sharpe 1974, Koenderink and van Doorn 1978, Wuerger, Atkinson et al. 2005). Much is known about the processing of opponent (L+M) and non-opponent (L-M) signals in the retina and lateral geniculate nucleus (for reviews, see Dacey 1999, Weller and Horwitz 2017). However, relatively little is known about how these signals are processed once they reach the primary visual cortex (area V1). This uncertainty remains a major barrier to understanding the role of V1 in luminance and color vision.

Our current understanding of color tuning in V1 is largely qualitative: some V1 neurons respond to changes in luminance but not red-green chromaticity, others respond to changes in red-green chromaticity but not luminance, and yet others respond to both (Thorell, De Valois et al. 1984, Lennie, Krauskopf et al. 1990, De Valois, Cottaris et al. 2000, Johnson, Hawken et al. 2001). These observations constrain how L- and M-cone signals are combined at the level of V1 but are insufficient to build and rigorously test a quantitative model of cone signal combination. In the absence of such a model, cone signal combination is often approximated as a weighted sum, which is nonlinearly related to spiking responses. This simple linear-nonlinear cascade (LN) model has been useful for describing L- and M-cone signal combination in the pre-cortical visual system, but

fails to accurately describe the responses of many V1 neurons (Thorell, De Valois et al. 1984, Pointer, Dean et al. 1989, Hanazawa, Komatsu et al. 2000, Horwitz and Hass 2012).

To reveal how signals from the L- and M-cones are combined at the level of V1, we probed individual neurons with punctate flashes designed to modulate the L- and M-cones exclusively. The flashes were tailored to fall inside the receptive field (RF) of each recorded neuron, and when possible, inside individual subunits (e.g. ON or OFF) within the classical RF. This stimulus mitigated potentially confounding effects of stimulating functionally distinct regions of the RF with simultaneous contrast increments and decrements, as with a grating, and it reduced nonlinearities produced by integrating signals across (Pointer and Osur 1989, Conway and Livingstone 2006) and beyond (Hess and Pointer 1989, Pointer, Dean et al. 1989, Solomon, Peirce et al. 2004) the classical RF. Limiting stimulation to the L- and M-cones also reduced the influence of pre-retinal filters (Cottaris 2003, Lindbloom-Brown, Tait et al. 2014) and of S-cone-dominated pathways which show pronounced nonlinearities (Dennis, Howarth et al. 1989, Giulianini and Eskew 2007, Horwitz, Chichilnisky et al. 2007).

We found a population of neurons that responded to an equally weighted sum of L- and M-cone signals, rendering them sensitive to luminance. Another population of neurons responded to an equal but oppositely weighted sum of L- and M-cone signals, rendering them sensitive to red-green chromaticity. A third population responded to both luminance and chromaticity, but the color tuning of these neurons could not be described by a single weighted sum of cone signals. Instead, the responses of these neurons were better described by an energy-like combination of two orthogonally weighted sums of cone signals. These two weighted sums were consistently the same weighted sums as in the color-only and luminance-only neurons. Color-only neurons had greater

contrast sensitivity than luminance-only neurons, and color-luminance neurons were more sensitive to chromatic than to luminance modulations.

3.2 METHODS AND MATERIALS

3.2.1 *Experimental setup*

Two monkeys (*Macaca mulatta, one male and one female*) participated in these experiments. Each monkey was surgically implanted with a headpost and a recording chamber (Crist Instruments) over area V1. All procedures conformed to the guidelines provided by the US National Institutes of Health and the University of Washington Animal Care and Use Committee.

During experiments, monkeys sat in a primate chair 1 m from a cathode ray tube monitor in a dark room. Neural signals were recorded with extracellular tungsten microelectrodes of 1−2 MΩ (Frederick Haer) and digitized at 40 kHz. Spikes were isolated online based on waveform timing and amplitude criteria. Event timing, stimulus presentation, and data collection were controlled by three desktop computers (see Horwitz and Hass 2012 for details).

Emission spectra and voltage-intensity relationships of each monitor phosphor were measured with a PR-650 SpectraColorimeter (PhotoResearch). The depth of each color channel was increased from 8 to 14 bits using a Bits++ video signal processor (Cambridge Research) at the expense of spatial resolution; each pixel was twice as wide as it was tall. Background luminance was ($x = 0.3$, $y = 0.3$, $Y = 65$ cd m^{-2}). The monitor refreshed at 75 Hz.

The monkeys were trained to fixate a centrally located dot measuring 0.2 x 0.2 degrees of visual angle (DVA) and maintain their gaze within a 1 x 1 DVA fixation window. Eye position was monitored throughout the experiment using a digital eye-tracking system (SMI iView X Hi-Speed Primate, SensoMotoric Instruments). Successful fixation was rewarded with juice or water. Fixation breaks caused the stimulus to disappear, and only completed trials were analyzed.

3.2.2 *Experimental Design and Statistical Analysis*

Each experimental session was divided into two phases. During the first phase, the experimenter identified the position and shape of a recorded V1 neuron's classical RF. During the second phase, flashes of uniformly colored light were targeted to the classical RF.

3.2.2.1 Phase 1: White noise receptive field characterization

During the first phase of the experiment, the neuron was stimulated with spatiotemporochromatic white noise. The white noise stimulus was a grid of squares each subtending 0.075–0.1 DVA (Figure 3.1A). On every screen refresh, each square was randomly assigned one of four lights designed to modulate the L- and M-cones in or out of phase, producing luminance and chromatic signals, respectively. The rapidity of the stimulus update caused lights at the same location on successive frames to be integrated in the cones, thereby having nearly the same effect as the average of the two lights. This integration effectively increased the number of lights sampled (Chichilnisky 2001).

Cone excitations were calculated using the Stockman-MacLeod-Johnson 10-degree cone fundamentals (Stockman, MacLeod et al. 1993), and cone modulations were calculated using Weber contrast (Koenderink, Bouman et al. 1978). The L- and M-cone contrasts of the four lights were [0.64, 0.64] (ON-luminance), [-0.64, -0.64] (OFF-luminance), [0.06, -0.06] (red), and [-0.06, 0.06] (green).

Neuronal responses to the white noise stimulus were analyzed online by spike-triggered averaging. In this analysis, the four lights that composed the white noise stimulus were weighted by recent spiking activity and averaged separately (Figure 3.1B). This procedure allowed us to localize the RFs of both simple and complex cells. Simple cells responded to a single stimulus polarity (e.g. bright or dark) at each location in the RF, whereas complex cells responded to

opposite stimulus polarities (e.g. bright and dark) at the same locations. Averaging all of the stimuli together would have obscured the RFs of complex cells. RF regions were identified by the experimenter as sets of adjacent squares in the spike-triggered averages that indicated the same spectral sensitivity (Figure 3.1C). The constraint of adjacency was violated for one neuron.

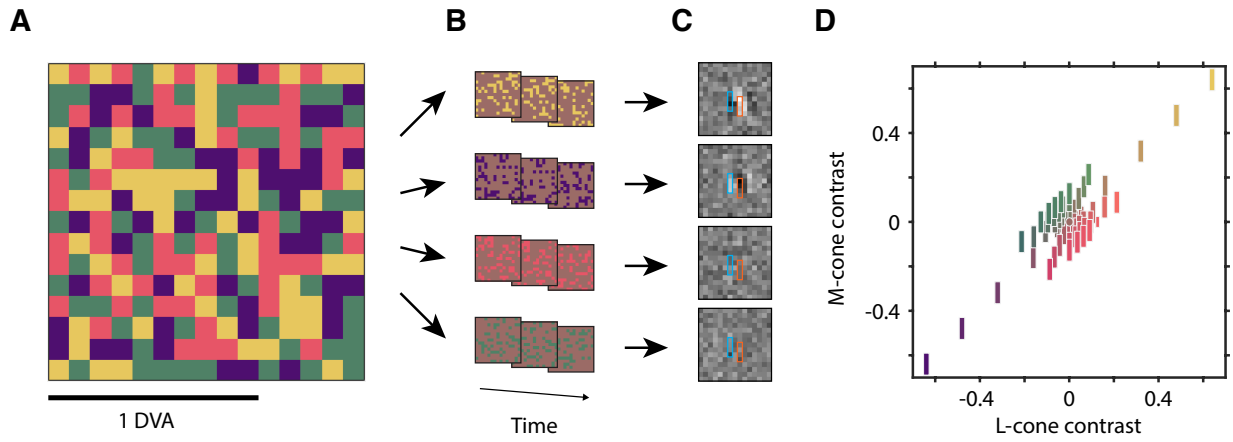

Figure 3.1 White noise and flashed stimuli.
A: Spatiotemprochromatic white noise. The color of each square was randomly reassigned on each frame refresh (every 13.3 ms) from a set of four: L+M (yellow), -L-M (purple), L-M (red), and M-L (green). B: Spike-triggered averaging. All of the stimulus frames that preceded each spike were separated by color. Frames sharing the same color were averaged together. C: Spike-triggered average stimuli for an example cell at 40-53 ms delay. Panels are separated by stimulus color. Dark and bright squares indicate the absence or presence, respectively, of each stimulus-type just before a spike. Receptive field subunit selections made by the experimenter are outlined in orange and cyan. D: Stimulus distribution in the LM plane. Each rectangle represents a stimulus in the LM plane. Stimulus shape matches the selected region in Figure 1C, and color approximates stimulus appearance. The distention of the stimulus distribution is due to the limited monitor gamut in the L-M direction.

3.2.2.2 Phase 2: Measurement of responses to uniformly colored flashes in the LM plane

In the second phase of the experiment, we measured responses to flashes of light targeted to individual RF regions. Stimuli were presented for 200 ms and were followed by a post-stimulus period of 500 ms during which the monkey maintained fixation. Flashes were uniformly colored

and designed to modulate the L- and M-cones over a wide range (Figure 3.1D). The maximum luminance and chromatic contrasts of the flashes were matched to those used in the white noise stimulus. For most neurons, each flash was presented five times within each block of trials. For two neurons, each flash was presented three times, and for one neuron each flash was presented 10 times. After each block, new combinations of L- and M-cone contrast were selected for testing, and this procedure was repeated for as long as the neuron remained well isolated. For 103 neurons, only a single RF region was targeted for stimulation. For 15 neurons (10 from monkey N and 5 from monkey M), two regions were targeted in interleaved trials. For the purposes of population analyses, only the more sensitive region was included for these 15 neurons.

We quantified neuronal responses by counting spikes within a 200 ms window that was customized for each neuron via the following procedure. A time-varying average firing rate was computed across stimulus conditions with respect to stimulus onset (1 ms bins, smoothed with a 1.7 ms standard deviation Gaussian filter). A threshold was defined as the highest firing rate achieved during the 200 ms prior to stimulus onset. The spike counting window for each neuron began when the average firing rate exceeded the threshold by 10% (mean latency ± standard deviation across all neurons = 65 ± 27 ms).

3.2.2.3 Gabor stimuli

To relate our results to those from previous studies, we probed fourteen neurons (8 from Monkey N and 6 from Monkey M) additionally with drifting Gabor patches in interleaved trials. Each Gabor patch subtended 2.4 DVA (σ = 0.4 DVA) and drifted at three Hz with a spatial frequency and orientation based on the subunit structure identified in the STA. Contrast increased linearly over a half-cycle, remained constant for one complete cycle, and decreased linearly over a half-cycle (full duration = 667 ms). The maximum contrast of each Gabor patch was matched to

those of the punctate flashes. All fourteen of these neurons were modulated (response - baseline) at least 10 spikes/sec by a Gabor patch.

A Gabor patch consists of both increments and decrements in light with respect to the background, whereas the flashes were unipolar. To compare responses across these two stimulus classes, we symmetrized flash responses by averaging responses to pairs of stimuli that had the same absolute contrast magnitude but opposite sign.

3.2.2.4 LN model

Under the LN model, L- and M-cone contrasts are weighted and summed to create a generator signal that underlies stochastic spiking (Figure 2A) (Chichilnisky 2001):

$$generator\ signal\ =\ [L\ M]\begin{bmatrix}A\\B\end{bmatrix} \tag{3.1}$$

where the generator signal can be expressed geometrically as the projection of cone contrasts, [L M], onto a unit-vector, $\begin{bmatrix}A\\B\end{bmatrix}$, that points in the neuron's preferred (color) direction.

We assume that the generator signal is related to the expected spike count, μ, via a Naka-Rushton function:

$$\mu = R_{max}\frac{max(generator\ signal, 0)^N}{max(generator\ signal, 0)^N + \ {c_{50}}^N} + bl \tag{3.2}$$

where R_{max} is the maximum response of the neuron (upper asymptote), *bl* is the baseline response in the absence of a stimulus (lower asymptote), *N* is an exponent that describes the slope of the contrast-response function, and c_{50} is the stimulus contrast at which the response reaches half-maximum.

As expected, some V1 neurons responded to flashes of only one contrast polarity at a single location in the RF (e.g. simple cells) and others responded to flashes of opposite polarities (e.g. complex cells). To describe the responses of both types of neurons, we used a pair of Naka-Rushton

functions that met tail-to-tail at the origin of the LM plane (which represents zero contrast). This model allowed us to describe responses to contrast increments and decrements along the preferred axis $\begin{bmatrix} A \\ B \end{bmatrix}$ independently:

$$generator\ signal_{LN} = max\left([L\ M]\begin{bmatrix} A \\ B \end{bmatrix}, 0\right) + u * max\left([L\ M]\begin{bmatrix} -A \\ -B \end{bmatrix}, 0\right) \quad (3.3)$$

The value of u was constrained to fall between 0 and 1, where 0 corresponds to half-wave rectification (Figure 3.2B, *left panel*), and 1 corresponds to full-wave rectification (Figure 3.2B, *right panel*).

3.2.2.5 LNLN model

Some neurons were poorly fit by the LN model but well fit by an extension of it that falls into the broad class of linear-nonlinear-linear-nonlinear (LNLN) cascade models. The particular subset of LNLN models we considered had the same basic architecture as the LN model: L- and M-cone contrast signals combine to produce a scalar-valued generator signal that is passed through a Naka-Rushton function to produce a spike count response. Under the LNLN model, however, the generator signal is produced by adding two orthogonally weighted sums of cone signals which are independently squared (Figure 3.2C):

$$generator\ signal_{LNLN} = \sqrt{\begin{matrix} \left(max\left([L\ M]\begin{bmatrix} A \\ B \end{bmatrix}, 0\right) + u * \max\left([L\ M]\begin{bmatrix} -A \\ -B \end{bmatrix}, 0\right)\right)^2 \\ + v * \left([L\ M]\begin{bmatrix} B \\ -A \end{bmatrix}\right)^2 \end{matrix}} \quad (3.4)$$

The neuron's sensitivity along the $\begin{bmatrix} A \\ B \end{bmatrix}$ axis is represented by the first term, just as in the LN model (Eq. 3.3). Unlike the LN model, the LNLN model includes a second term that represents the neuron's sensitivity along the orthogonal axis, $\begin{bmatrix} B \\ -A \end{bmatrix}$. Sensitivity along this axis is constrained to be identical for contrast increments and decrements. We added this constraint because a model

that allowed asymmetric sensitivity along this axis was overly flexible and unnecessary for describing the observed responses. For these reasons, we refer to the $\begin{bmatrix} A \\ B \end{bmatrix}$ axis as the 'asymmetric axis' and the $\begin{bmatrix} B \\ -A \end{bmatrix}$ axis as the 'symmetric axis'.

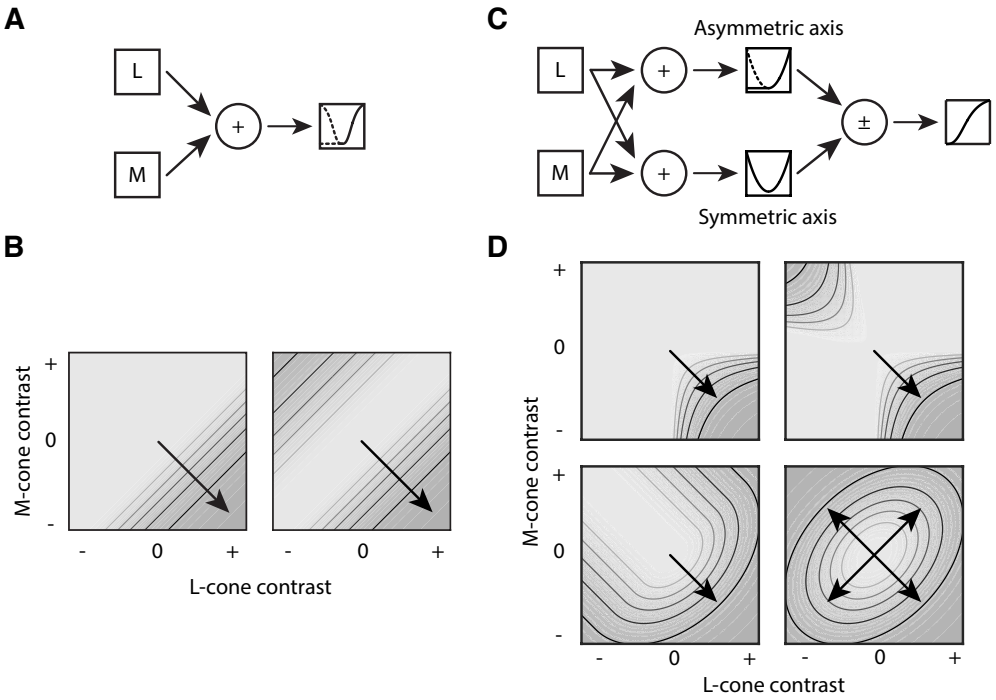

Figure 3.2 LN and LNLN models for L- and M-cone signal combination.
A: Under the LN model, responses are generated by passing a weighted sum of L- and M-cone signals through a (optionally two-sided) Naka-Rushton function. B: Response contours in the LM plane described by the LN model. Weak to strong responses are represented by light to dark shading, respectively. Contours are always straight and parallel to each other, and perpendicular to the neuron's preferred color direction (arrow). *Left*: half-wave rectified responses to cone modulations. *Right:* full-wave rectified responses to cone modulations. C: Under the LNLN model, responses are generated via an energy-like combination of two orthogonally weighted sums of L- and M-cone signals. One of these is squared or half-squared (asymmetric axis), and the other is squared (symmetric axis). The square root of these summed or subtracted quantities is then passed through a Naka-Rushton function. D: Response contours in the LM plane described by the LNLN model. Shading as in Figure 2B. Responses along the asymmetric axis may be half-rectified (*left*) or fully-rectified (*right*). Subtracting the two weighted sums results in narrower-than-linear color tuning (*top*), whereas adding them results in broader-than-linear color tuning (*bottom*).

The scalar v reflects the neuron's relative sensitivity between the symmetric and asymmetric axes. If $v < 0$, the neuron is narrowly tuned and responds to a smaller set of stimuli than predicted by the LN model. Narrowly tuned neurons have response contours that curve away from the origin (Figure 3.2D, *upper panels*). If $v > 0$, the neuron is broadly tuned and responds to a larger set of stimuli than predicted by the LN model. Broadly tuned neurons have response contours that curve toward the origin (Figure 3.2D, *lower panels*). The LNLN model is identical to the LN model if $v = 0$. Note that the classification of color tuning as linear, narrower-than-linear, or broader-than-linear depends entirely on the combination of cone signals into a generator signal (Eq. 3.3 and Eq. 3.4) and not on the contrast response function (Eq. 3.2), which transforms the generator signal into spiking responses.

3.2.2.6 Error model

Spike counts of many of the neurons we recorded had super-Poisson variability. We therefore modeled spike counts as having a negative binomial distribution (Koenderink, Bouman et al. 1978, Koenderink, Bouman et al. 1978). The negative binomial distribution flexibly couples the variance of responses to their mean:

$$\sigma^2 = \mu + \kappa * \mu^2 \tag{3.5}$$

where μ is the mean predicted response, σ^2 is the variance of the predicted response, and κ is the parameter that controls the relationship between mean and variance. The Poisson distribution is a limiting case of the negative binomial distribution for which $\kappa = 0$.

The parameters of the LN and LNLN models were fit by maximum likelihood. The log of the likelihood function, which was maximized numerically using the fmincon function (MATLAB, MathWorks), can be written:

$$\mathcal{L} = \sum_{i=1}^{n} \left(\log\left(\Gamma\left(\frac{{\mu_i}^2}{\sigma_i^2 - \mu_i} + R_i \right) \right) + \log(R_i!) + \log\big(\Gamma(R_i)\big) \ldots \right.$$
$$\left. - \log\left(\frac{\sigma_i^2 - \mu_i}{\sigma_i^2} R_i \right) - \left(\log\left(1 - \frac{\sigma_i^2 - \mu_i}{\sigma_i^2} \right) \frac{{\mu_i}^2}{\sigma_i^2 - \mu_i} \right) \right) \tag{3.6}$$

where $\mathcal{L}$ is the log-likelihood of the parameters given the observed responses, n is the total number of tested stimuli, R_i is the observed response to the i^{th} stimulus, μ_i is the predicted mean response to the i^{th} stimulus, σ_i^2 is the predicted variance of the responses to the i^{th} stimulus, and Γ is the gamma function.

3.2.2.7 Model comparisons

The LNLN model always fit the data at least as well as the less flexible LN model. To quantify the improvement in the fit of the LNLN model over the LN model, we calculated normalized log-likelihoods under each model (Koenderink, Bouman et al. 1978):

$$norm\mathcal{L}_{LN} = \frac{\mathcal{L}_{LN} - \mathcal{L}_{LB}}{\mathcal{L}_{UB} - \mathcal{L}_{LB}} \qquad norm\mathcal{L}_{LNLN} = \frac{\mathcal{L}_{LNLN} - \mathcal{L}_{LB}}{\mathcal{L}_{UB} - \mathcal{L}_{LB}} \tag{3.7}$$

where $\mathcal{L}_{LN}$ is the log-likelihood of the fitted LN model, $\mathcal{L}_{LNLN}$ is the log-likelihood of the fitted LNLN model, $\mathcal{L}_{LB}$ is a theoretical lower bound of the log-likelihood values, and $\mathcal{L}_{UB}$ is a theoretical upper bound. $\mathcal{L}_{LB}$ was determined using a model in which the predicted response to each stimulus was equal to the overall mean response observed across all stimuli. $\mathcal{L}_{UB}$ was determined using a model in which the predicted response to each stimulus was equal to the observed mean response to each stimulus. The resulting normalized log-likelihood values, as well as the difference between them, fall between 0 and 1.

We classified neurons into two categories by setting a bound on the normalized log-likelihood improvement necessary to justify the larger LNLN model. In the following analyses, this threshold was set to 0.08 on the basis that this threshold segregated the bulk of the distribution

from its long tail. The largest value of the normalized log-likelihood difference that we observed was 0.26; the theoretical upper bound was high because we presented many stimuli with few repeats. Note that a criterion based on a statistical hypothesis test would have biased neurons to be classified as better described by the LNLN model because the LN model is a mathematical idealization that is rejected with progressively higher probability as the number of stimulus presentations increases.

We also used the normalized log-likelihood difference to distinguish half-rectified responses from fully-rectified responses (Eq. 3.3 and Eq. 3.4). We fit the data with the more-constrained model ($u = 0$) unless the less-constrained model ($0 < u \leq 1$) increased the normalized log-likelihood by at least 0.08.

3.2.2.8 Comparing contrast sensitivity across neurons

We asked whether luminance- and chromaticity-tuned neurons differed in contrast sensitivity (c_{50} in Eq. 3.2). To control for potentially confounding effects of stimulus size, we performed the following analysis, which was inspired by the partial F-test. We fit c_{50} values across a pool of neurons as a function of stimulus size using a cubic spline and recorded the sum of squared error (SSE). Next, we divided the pool into luminance- and chromatic-tuned groups, fit each with a cubic spline, and recorded their in-group SSEs. Then, we computed a test statistic, which is the ratio of SSEs under both models:

$$F = \frac{SSE_{all}}{SSE_{lum} + SSE_{chrom}} \tag{3.8}$$

This ratio is large if luminance- and chromaticity-tuned neurons have different relationships between stimulus size and c_{50}, and it is small if both groups share the same relationship. To determine if the observed value of this ratio was larger than expected by chance, we shuffled the identities of the neurons between the luminance- and chromaticity-tuned groups, fit data from the

shuffled "luminance" and "chromatic" groups, and recalculated the ratio 10,000 times. The p-value is the proportion of ratios that exceeded the one calculated from the unshuffled data.

3.2.3 *Neuron inclusion criteria*

Receptive fields of the 118 neurons included in our study were located 1.5 – 6.2 DVA from the fovea (mean eccentricity = 4.3 DVA). To exclude noisily or weakly responding neurons from the analysis, we omitted any that had $\kappa > 4$ (Eq. 3.5) or failed to fire at least 10 spikes/sec to at least one of the stimuli. Two neurons had high baseline firing rates and were strongly suppressed by some stimuli. The model does not incorporate suppression below baseline, so these neurons were omitted from analysis.

3.3 RESULTS

The goal of this study was to develop quantitative models that map modulations in L- and M-cone activity to spiking responses in V1 neurons. We recorded responses of 118 V1 neurons (68 from monkey N and 50 from monkey M) to uniformly colored flashes that probed L- and M-cone signal processing inside the RF. Spike counts in response to the flashes were fit with two models. The first, the LN model, is a classical description of neural responses under which spike counts depend on a single weighted sum of L- and M-cone signals. The second, the LNLN model, is a broader class of models under which responses depend on a nonlinear combination of two orthogonally weighted sums of L- and M-cone signals.

3.3.1 *Luminance- and chromaticity-tuned neurons*

Most neurons were well described by the LN model, as exemplified by the neuron in Figure 3.3. During white noise stimulation, the example neuron was excited by green, suppressed by red,

and unaffected by luminance, as visible from the spike-triggered averages (STA) (Figure 3.3A). To test whether this neuron combined signals from the L- and M-cones as a weighted sum, we probed its RF with uniformly colored flashes that modulated L- and M-cones over a wide range. The neuron responded exclusively to stimuli that had greater M-cone contrast than L-cone contrast (color directions > 45° and < -135° in Figure 3.3B). These stimuli lie in the upper-left of the LM plane (Figure 3.3C).

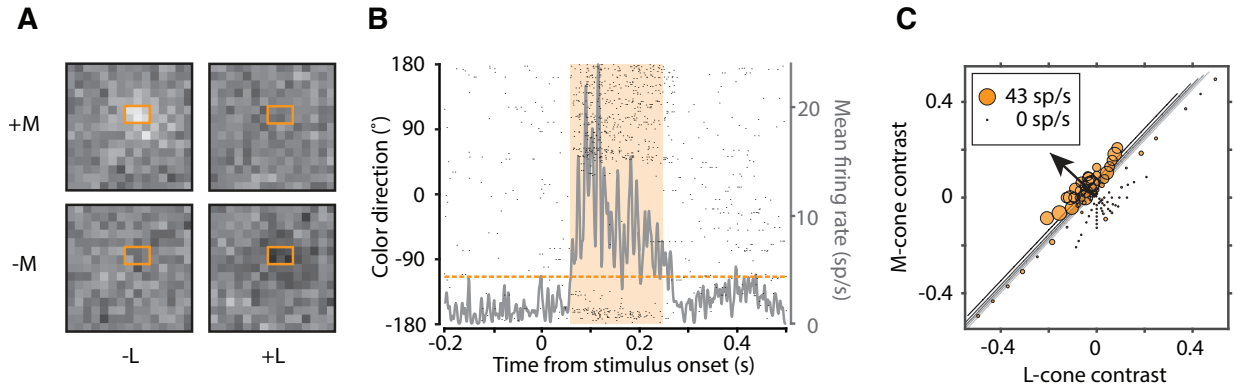

Figure 3.3 Example LN neuron.
A: Spike-triggered averages (STAs) of the four color channels. STAs are gray scaled, with dark to light shading representing absence or presence of a stimulus color in a given location prior to the neuron's response. The region selected by the experimenter for stimulation with flashes is outlined in orange. B: Peri-stimulus time histogram (gray) shows firing rate responses to flashed stimuli (right axis). Rasters show responses on individual trials, organized by color direction (left axis). Time 0 is stimulus onset. Spikes were counted in a 200 ms window (orange shading), which begins when the smoothed average firing rate exceeds the maximum observed baseline firing rate (dotted line) by 10%. C: Responses to flashed stimuli and LN model fit. The position of each orange circle indicates L- and M-cone contrasts, and circle size represents the corresponding mean response. Fitted response contours are gray scaled, with light to dark representing small to large responses, respectively. The minimum and maximum observed responses are indicated in the legend. The arrow represents the neuron's preferred color direction from the model fit.

We fit the responses of this neuron with the LN model and the LNLN model and compared normalized log-likelihoods. The normalized log-likelihood of these responses under the LN model was 0.76. Under the more flexible LNLN model, the normalized log-likelihood increased to 0.77.

This modest difference indicates that the LNLN model described this neuron's responses nearly as well as the LN, demonstrating that its added complexity was unwarranted.

Across the 118 neurons that we studied, differences in normalized log-likelihood ranged from 0 to 0.26. Small values indicate that the added complexity of the LNLN model improved the fit only slightly, whereas larger values indicate that the LNLN model improved the fit substantially. Based on the empirical distribution of normalized log-likelihood differences, we chose an improvement threshold of 0.08 to separate the bulk of the distribution from its long tail (Figure 3.4). By this criterion, 78 of the 118 neurons in our data set were well described by the LN model. Below, we analyze the tuning of these 78 neurons and focus on patterns in the parameter values from the LN model fits. We return to the remaining 40 neurons in the section on *Color-luminance neurons*.

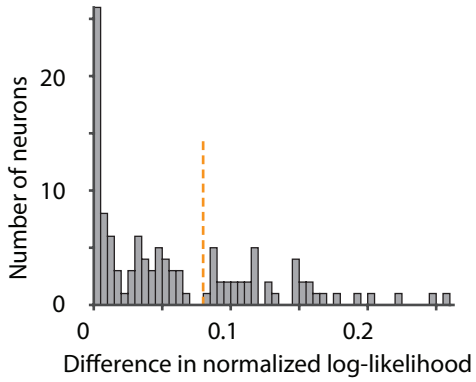

Figure 3.4 Histogram of normalized log-likelihood differences for 118 neurons.
Dotted orange line indicates the threshold (0.08) used to segregate neurons well described by the LN model from neurons well described by the LNLN model.

3.3.1.1 Preferred directions

The sensitivity of an LN neuron to L- and M-cone modulations can be summarized by a single set of cone weights (Eq. 3.3). These cone weights can also be expressed as a vector in the

LM plane pointing in the neuron's preferred direction. The estimated preferred direction of the example neuron in Figure 3.3 was 137º, which is close to the theoretical M-L direction (135°).

Across the 78 neurons that were well described by the LN model, the distribution of preferred directions had four modes: two cone-opponent and two non-opponent (Figure 3.5A). We define chromaticity- and luminance-tuned neurons as those with preferred directions within 22.5° of these opponent and non-opponent modes, respectively. No neuron falls into more than one category, and although these criteria account for only half of the potential preferred directions in the LM plane, they subsume 77% (60/78) of the neurons we studied. This organization, which has not been previously described using unbiased techniques, implies that chromatic- and luminance-tuned are separate populations and not points on a continuum.

We considered the possibility that our stimulus distribution, which was elongated along the luminance axis and constricted along the chromatic axis, biased the distribution of preferred color directions that we observed, which were oriented along the same axes. If the pattern of preferred directions is related to the elongation in the stimulus distribution, then eliminating the elongation should also eliminate the pattern. To test this possibility, we probed an additional 26 neurons in Monkey N with stimuli from a distribution that was approximately radially symmetric. All but one of these neurons were well-described by the LN model (difference in normalized log likelihood < 0.08). Although the stimulus distribution was virtually uniform in every direction, the distribution of preferred directions was not (Figure 3.5B). Instead, 19 neurons were tuned for red, six were tuned for green, and none were tuned for luminance. The absence of luminance-tuned neurons is likely due to the relative insensitivity of luminance-tuned neurons, as detailed in the following section on *Contrast sensitivity*. This control experiment shows that elongation of the

stimulus distribution does not account for structure in the distribution of preferred color directions (Figure 3.5A).

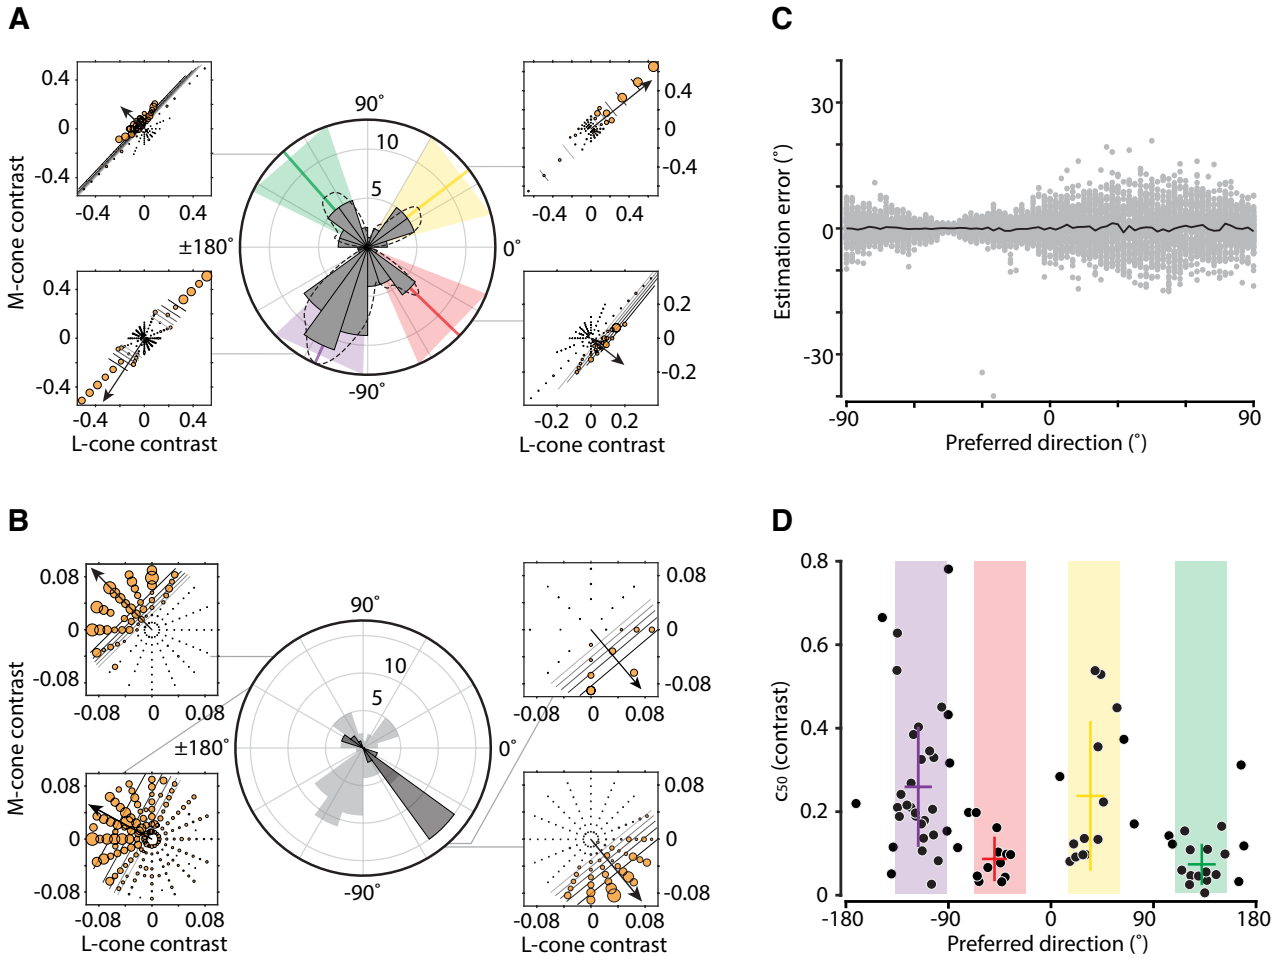

Figure 3.5 Preferred directions and contrast sensitivity of LN neurons.
A: Preferred directions of LN neurons in the LM plane are shown as a polar histogram and as a smoothed density estimate (black dotted line). Luminance and chromatic modes (thick colored lines) were identified as peaks in the smoothed density estimate. Wedges subtending 45° were centered on each mode to classify neurons into four groups. *Insets*: Data from representative neurons in each wedge. Conventions are as in Figure 3. B: Histogram of preferred directions of neurons probed with a radially symmetric stimulus distribution (dark gray). The histogram from Figure 5A is superimposed for comparison (light gray). C: Accuracy of preferred direction estimates. Errors in individual preferred direction estimates (gray circles) and averages (black curve) were calculated from simulated neuronal responses. D: c_{50} values in the preferred directions of all LN neurons (black circles). Colored boxes represent the wedges defined in panel A. Colored crosses indicate the mean and standard deviation of c_{50} values and preferred directions within each wedge.

We also considered the possibility that the predominance of chromaticity- and luminance-tuned neurons was a consequence of our fitting procedure. To test this possibility, we conducted a simulation in which we used our fitting procedure to estimate the preferred directions of model neurons with pre-determined preferred directions. The preferred directions of simulated neurons ranged from -90 to +90° in the LM plane at 2.5° intervals, and the c_{50} value was set to the middle of the range of tested stimulus contrasts (Eq. 3.1). Each simulated neuron was probed in 80 simulated experiments, and responses were fit identically to the real data. The resulting preferred direction estimates were essentially unbiased throughout the LM plane (Figure 3.5C), indicating that the observed patterns in the preferred directions did not arise from bias in the fitting procedure (see also Rovamo, Leinonen et al. 1984).

In summary, we found that neurons well-described by the LN model had stereotyped preferred directions, rendering these neurons chromaticity-tuned (sensitive to approximately equal but opposite modulations of L- and M-cones) or luminance-tuned (sensitive to approximately equal, non-opponent modulations of L- and M-cones). This structure in the distribution of preferred color directions does not arise from our stimulus distribution or fitting procedure. Next, we discuss differences between luminance- and chromaticity-tuned neurons in contrast sensitivity.

3.3.1.2 Contrast sensitivity

A preferred color direction describes a neuron's relative sensitivity to L- and M-cone modulations. Absolute contrast sensitivity can be quantified as the contrast at half-maximum response (c_{50} in eq. 3.2). Analysis of c_{50} values in preferred color directions showed that chromaticity-tuned neurons were more sensitive than luminance-tuned neurons (average c_{50} values: 0.08 and 0.25 for chromaticity- and luminance-tuned neurons, respectively; Mann-Whitney U-test, $p = 6.6 \times 10^{-7}$; Figure 3.5D). No significant difference in c_{50} was found between

ON- and OFF-luminance neurons (p = 0.32) or between red- and green-preferring neurons (p = 0.77). Other parameters of the Naka-Rushton function did not differ across the four groups of neurons defined in Figure 3.5A (Kruskal-Wallis test on R_{max}, mean across groups = 16 spikes/sec, p = 0.50; bl, mean across groups = 0.6 spikes/sec, p = 0.40; and *N*, mean across groups = 3.6, p = 0.11).

One possible explanation for the difference in contrast sensitivity between luminance- and chromaticity-tuned neurons is differences in RF eccentricity. Visual sensitivity decreases with eccentricity and, if the RFs of the luminance-tuned population were more eccentric than the chromaticity-tuned population, this might explain their lower sensitivity. However, RF eccentricity was similar between luminance-tuned (mean = 4.5 DVA) and chromaticity-tuned (mean = 4.2 DVA) groups (Mann-Whitney U-test, p = 0.1).

Another possible explanation for the greater contrast sensitivity of chromaticity-tuned neurons is that their RFs may be larger. Because the stimulus was tailored to each neuron, neurons with larger RFs were stimulated with larger stimuli. Large stimuli activate a greater number of cones than small stimuli and therefore might drive stronger responses in V1, even at low contrasts. Consistent with this explanation, neurons probed with large stimuli tended to have low c_{50} values (data not shown). Moreover, chromaticity-tuned neurons were probed with larger stimuli on average (mean = 0.49 DVA^2) than luminance-tuned neurons (mean = 0.34 DVA^2, Mann-Whitney test, p=0.001). These observations are consistent with the idea that chromaticity-tuned neurons are more sensitive because they have larger RFs. However, the relationship between stimulus size and c_{50} differed between luminance- and chromaticity-tuned populations (permutation test, $p < 0.01$, see Methods). We conclude that contrast sensitivity increases with stimulus size, but that this

relationship is different for luminance- and chromaticity-tuned populations and therefore does not account for the difference in contrast sensitivity between them.

3.3.2 *Color-luminance neurons*

Although most of the V1 neurons that we studied were well-described by the LN model, a substantial minority was not. A prominent feature of this minority is that they lacked a null axis—an axis through color space in which stimulus modulation does not drive a spiking response. Another prominent feature of these neurons is that they responded to both cone-opponent and non-opponent modulations. Qualitatively, we observed two varieties of these neurons: those which were responsive in every color direction, and those that responded to a broader-than-linear collection of stimuli but not in every color direction. We begin by describing neurons in the first category.

3.3.2.1 Omni-directional contrast detectors

Forty of the 118 neurons we studied were better fit by the LNLN model than the LN model based on our normalized log-likelihood criterion. Ten of these 40 neurons responded to modulations in every color direction tested, as exemplified by the neuron in Figure 3.6. Spike-triggered averaging for this neuron revealed a modest sensitivity to luminance modulation and no sensitivity to chromatic modulation (Figure 3.6A). Nevertheless, when the RF center was probed with punctate flashes, the neuron responded robustly to modulations in every color direction (Figure 3.6B & 3.6C). The responses of this neuron were poorly explained by the LN model (normalized log-likelihood = 0.62) but were well explained by the LNLN model (normalized log-likelihood = 0.88). This result shows that this neuron does not combine cone signals as a weighted sum.

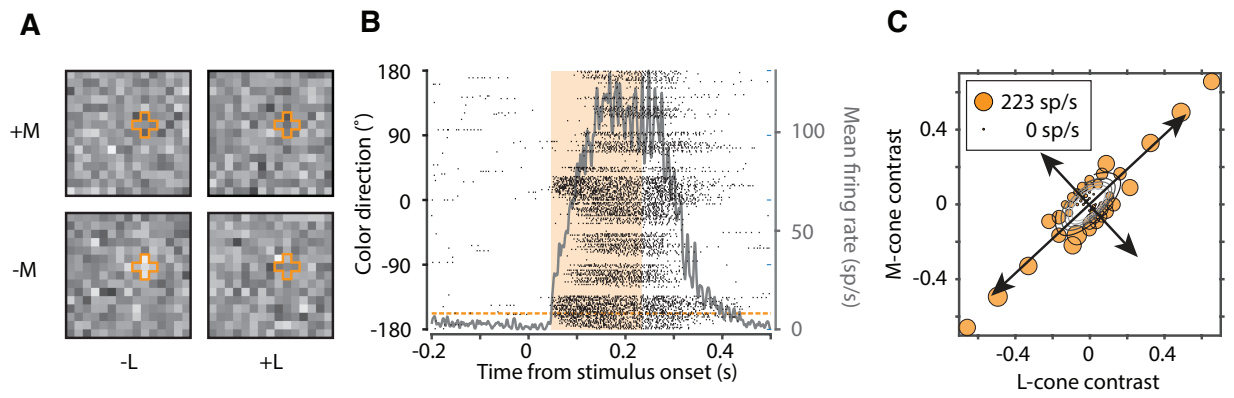

Figure 3.6 Example omni-directional LNLN neuron.
Conventions are as in Figure 3.3. Arrows in (C) show the orientation of the axes from the model fit.

For the 10 neurons that responded to modulations in all color directions, as in Figure 3.6, a single preferred color direction carries little information. On the other hand, the orientation of the symmetric ($\begin{bmatrix} A \\ B \end{bmatrix}$) and asymmetric ($\begin{bmatrix} B \\ -A \end{bmatrix}$) axes from the LNLN model fit (Eq. 3.4) is informative because contrast sensitivity in these two directions specify contrast sensitivity in all directions. The distinction between these two axes is nearly arbitrary for these neurons because they respond to modulation in all color directions, so we represent the orientation of both axes using a single 90° wedge.

The distribution of axes was tightly clustered with a mode at ~45° (Figure 3.7A). This result shows that L+M and L-M axes provide a natural basis for describing the responses of neurons that are broadly tuned for L- and M-cone contrast modulations. We verified that the observed structure in orientations of these axes was not due to the stimulus distribution or fitting procedure through simulation (data not shown).

These color-luminance-tuned neurons, like the population of neurons that were well described by the LN model, were more sensitive to chromatic modulations than luminance

modulations (average c_{50} values 0.13 vs. 0.22, Wilcoxon test, $p = 0.011$) (Figure 3.7B). This comparison was conducted within-neuron, thus stimuli were matched for size and retinal eccentricity in this analysis.

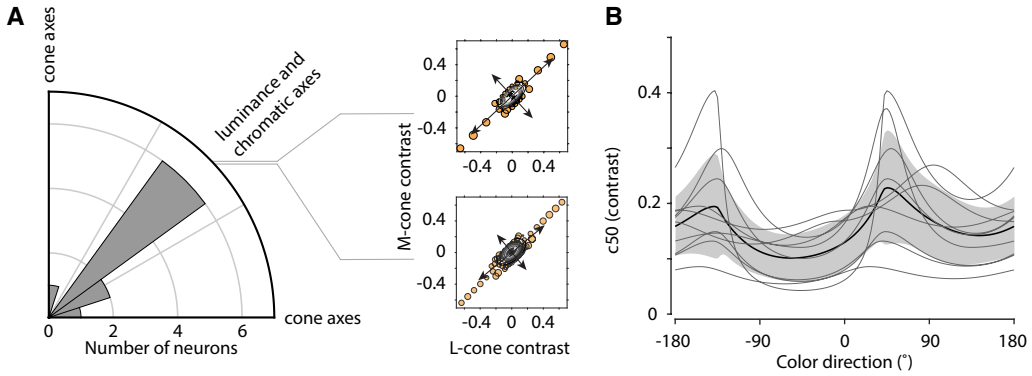

Figure 3.7 Population of omni-directional LNLN neurons.
A: Principal axes of omni-directional neurons. *Left*: Orientations of axes from model fits are shown as a polar histogram. *Right*: Data from two example neurons. Conventions as in Figure 3. B: Contrast sensitivity as a function of color direction. Average c_{50} values (thick black line), standard deviation (gray shading), and traces for individual neurons (thin black lines) are shown.

3.3.2.2 Multi-directional contrast detectors

The remaining 30 of the 40 neurons that were well fit by the LNLN model responded to a broader-than-linear set of stimuli but not to modulations in every color direction, as shown by an example neuron (Figure 3.8). On average, this neuron was excited by OFF-luminance modulation in the white noise and suppressed by green modulation, broadly consistent with sensitivity to M-cone decrements (Figure 3.8A). On the other hand, the neuron's responses to flashes revealed sensitivity in almost every color direction (Figure 3.8B). This color tuning is broader than predicted by the LN model, since it responds to green and both polarities of luminance, but not quite as broad as that shown in Figure 3.6, which additionally responds to red.

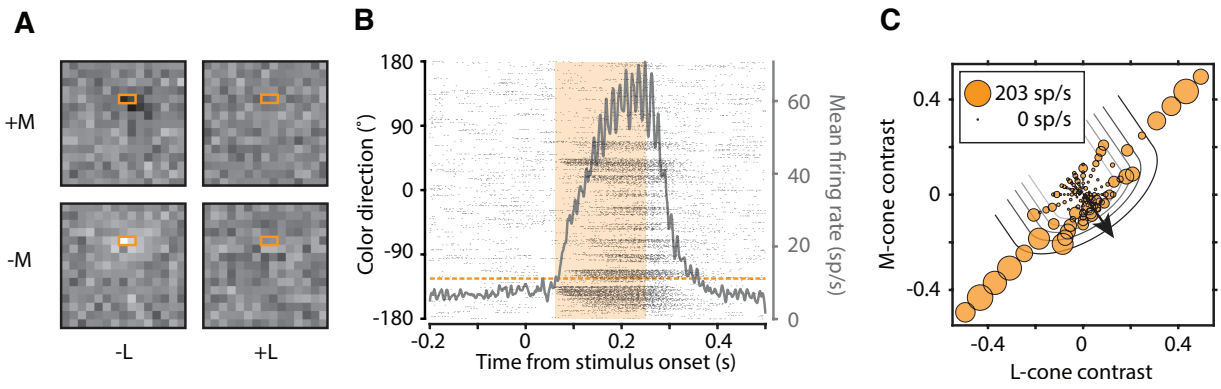

Figure 3.8 Example multi-directional LNLN neuron.
Conventions are as in Figure 3. Arrow in (C) shows the orientation of the 'U' from the model fit.

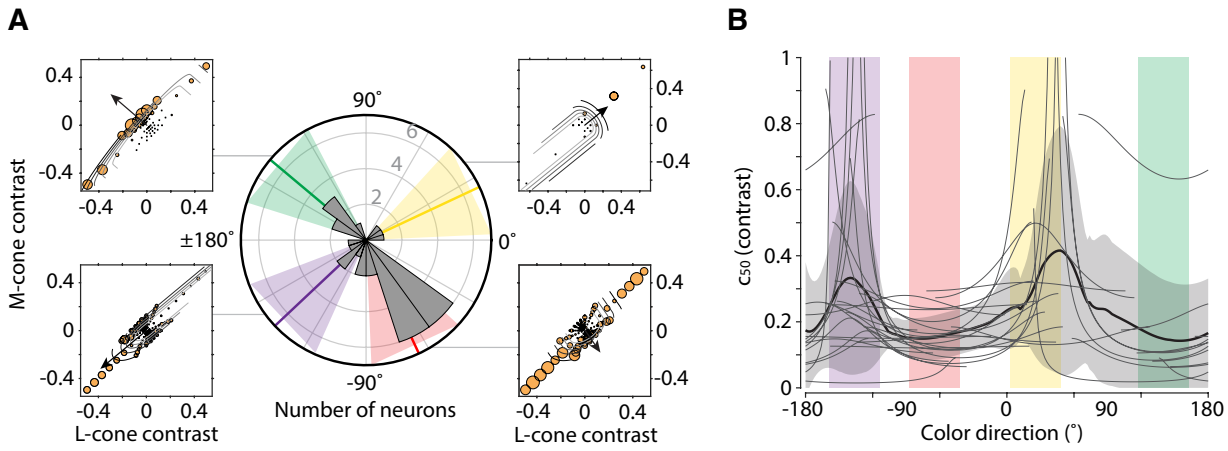

Figure 3.9 Population of multi-directional LNLN neurons.
A: Principal axes of multi-directional LNLN neurons. *Left*: Orientations of axes from model fits are shown as a polar histogram (dark gray) and as a smoothed density estimate (black dotted line). Luminance and chromatic modes (thick colored lines) were identified as peaks in the smoothed density estimate. Wedges subtending 45° were centered on each mode to classify neurons into four groups. *Insets*: Data from representative neurons within each wedge. Conventions are as in Figure 3. B: Contrast sensitivity as a function of color direction. Harmonic mean c_{50} values (thick black line), standard deviation (gray shading), and traces for individual neurons (thin black lines) are shown.

The absence of green responses from this neuron motivates a special case of our already-constrained LNLN model in which u in Eq. 3.4 was constrained to be 0, resulting in response

contours with a 'U' shape (Figure 3.8C). The color-tuning of neurons fit by this variant of the model was summarized by the orientation of this 'U' in the LM plane (Figure 3.9A). Most of the neurons in this category were sensitive to all modulations except for red or green (n=17). Neurons that were sensitive to all modulations except ON- or OFF-luminance were rarer (n=7). Like the other types of neurons that we studied, neurons in this category were more sensitive to chromatic modulation than to luminance modulation (c_{50} values for each neuron calculated for color directions within 90° of [A B] (see Eq. 4) 0.16 vs. 0.43, Wilcoxon test, $p = 4.1 \times 10^{-5}$) (Figure 3.9B).

3.3.3 *Stimulation of multiple RF regions*

In the initial phase of this experiment, we used spatiotemprochromatic white noise to identify the structure of V1 RFs. For 103 neurons, we stimulated only one region inside the classical RF. For 15 neurons whose STAs revealed different spectral sensitivity in different parts of the RF, we stimulated a pair of RF regions in interleaved trials and fit each set of responses separately. The LN model provided a good fit to responses produced by stimulation of both regions for 9 of these 15 neurons (difference in normalized log-likelihood < 0.08 for both sets of responses). For these 9 neurons, the preferred directions of the two RF regions were approximately opposite (e.g. Figure 3.10A). One neuron had a pair of RF regions that were each sensitive to one polarity of chromatic contrast and both polarities of luminance contrast (Figure 3.10B). This neuron was both spatially- and chromatically-opponent, and therefore double-opponent. Other double-opponent cells were weakly sensitive to luminance contrast, if at all (Figure 3.10C).

Five of these 15 neurons, including the example in Figure 3.10C, were better fit by the LN model when stimulation was directed to one part of the RF and by the LNLN model when stimulation was directed to another part. For these neurons, the improvement in the fit under the LNLN model was subtle. The fact that a single neuron could appear better fit by one model or the

other, depending on which part of the RF was stimulated, is likely a reflection of biological variability combined with a hard threshold for categorization.

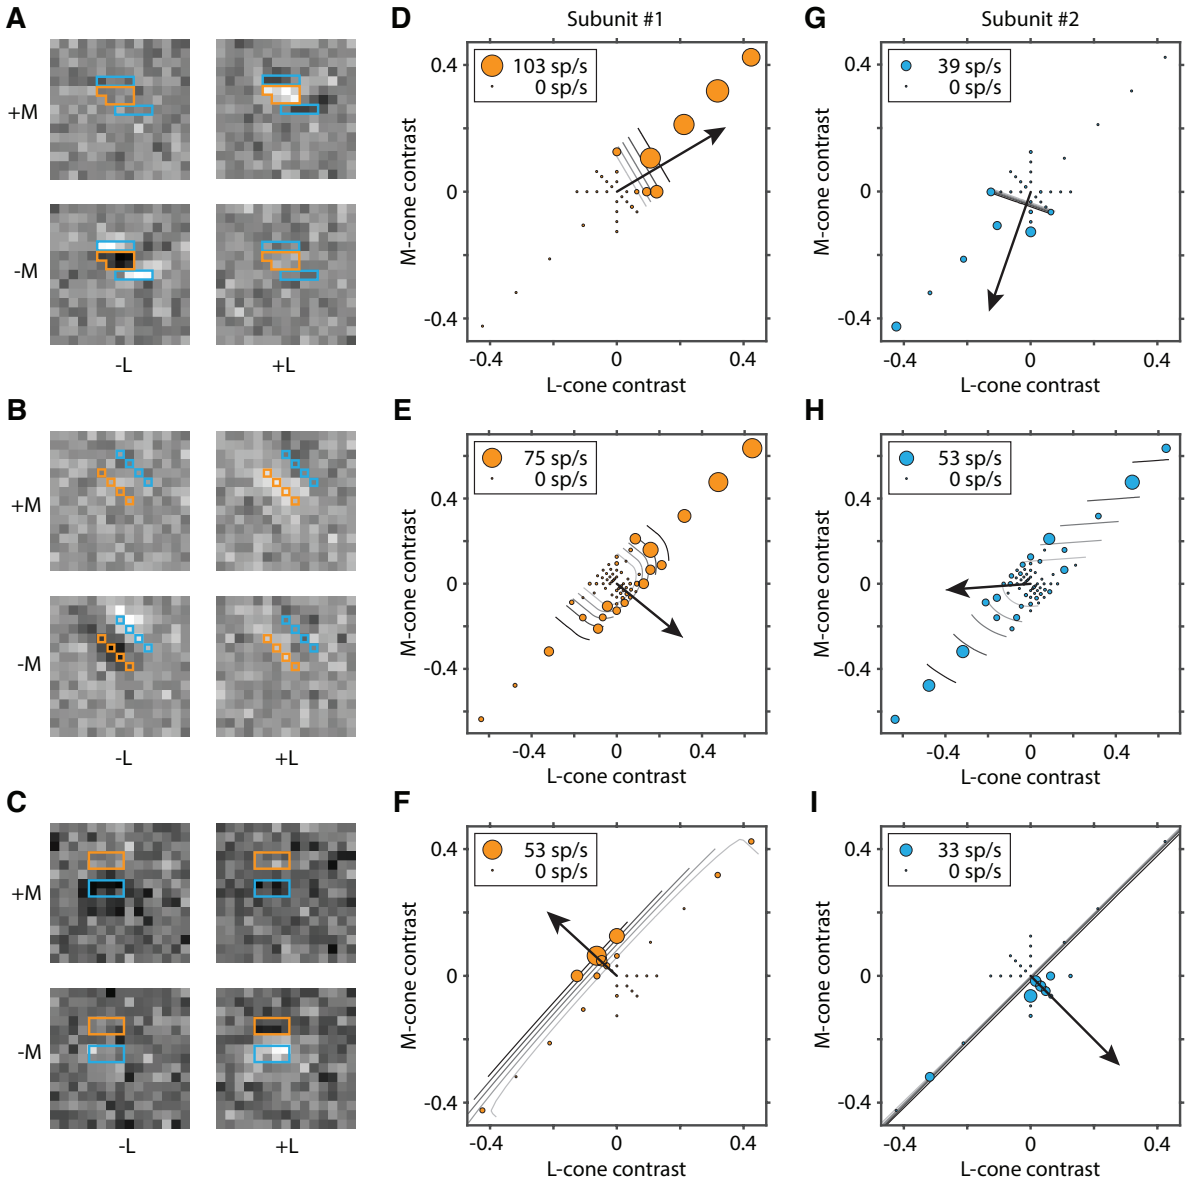

Figure 3.10 Tuning across receptive field (RF) locations for three example neurons. A–C: spike-triggered averages and flashed stimulus shape and location in the RF (blue and orange outlines). *D–F:* Responses and model fits to the flashes at the RF position that evoked the stronger responses. *G–I:* Responses and model fits to flashes at the RF position that produced weaker responses. Stimulus-response pairs and model fit conventions as in Figure 3. Preferred directions from LN fits are indicated with single arrows. For data sets fit by the LNLN model, the single arrow indicates the direction of the [A B] axis.

3.3.4 *Responses to Gabors and flashes are strongly correlated*

To determine whether our results were specific to the punctate flash stimuli we used, we stimulated 14 neurons with punctate flashes and drifting Gabor patches in interleaved trials. The correlation between responses to flashes and Gabors was > 0.5 ($p < 0.005$) for all neurons, indicating similar tuning (Figure 3.11A). For neurons that responded strongly to both sets of stimuli, the correlation was higher; 11 of 14 neurons responded > 18 spikes/sec for both stimuli, and these had $r > 0.8$ (e.g. Figure 3.11B). The three neurons with lower correlations responded weakly to one or both of the two stimuli (e.g. Figure 3.11C). We conclude that the spatial structure of the stimuli we used are unlikely to account for the tuning for L- and M-cone modulations we observed.

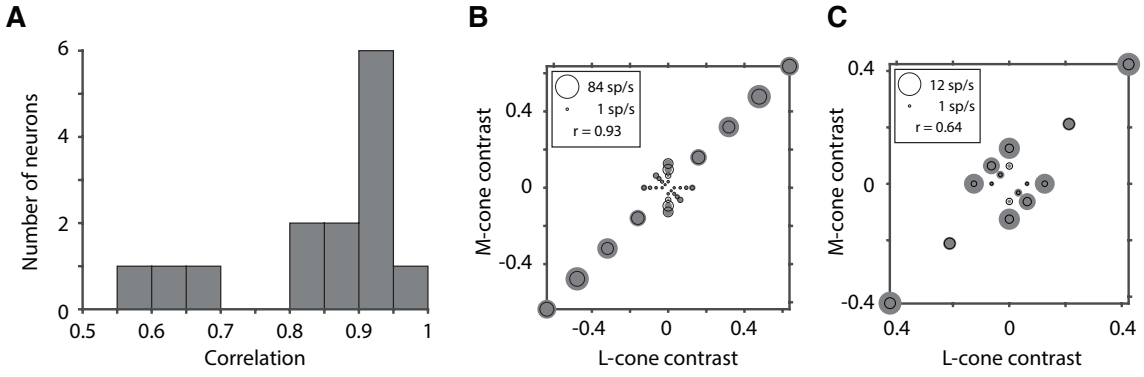

Figure 3.11 Responses to flashes and Gabors.
A: Histogram of correlation coefficients between responses to flashes and Gabors for 14 neurons. B and C: Data sets from two example neurons. Conventions are as in Figure 3 but model fits are not shown. Responses to Gabor patches (filled circles) and symmetrized responses to flashes (open circles) are shown at positions in the LM plane corresponding to their peak cone contrasts.

3.4 DISCUSSION

To investigate luminance and color processing in primate V1, we presented uniformly colored flashes inside the classical RFs of isolated V1 neurons. We found that some neurons responded to a single weighted sum of cone signals. These neurons were predominantly tuned for luminance or chromatic modulations but rarely for combined luminance and chromatic modulations. Neurons that were sensitive to both luminance and chromatic modulations responded to a nonlinear combination of two orthogonally weighted sums of cone signals. These two weighted sums were consistently the same weighted sums as in the first group.

This organization for L- and M-cone signals has not been previously observed in V1, and three aspects of our study helped to reveal it. First, our study is the first to manipulate L- and M-cone signals systematically and over a broad range to probe the responses of V1 neurons. Similar approaches propelled discoveries in retinal physiology (Pointer, Gauger et al. 1989, Lee, Martin et al. 1993) and psychophysics (Virsu, Rovamo et al. 1982, Pointer and Osur 1989, Metha, Vingrys et al. 1994, Gegenfurtner and Hawken 1995, Stromeyer, Kronauer et al. 1995, Giulianini and Eskew 1998). Second, we presented many stimuli, each repeated a few times as opposed to a presenting a few stimuli many times. This approach was necessary to test the LN model, estimate its parameters with precision when it fit the data well, and extend it when it failed. Third, we considered a broader class of models for describing color tuning in V1 than had been considered previously. This allowed us to analyze fitted parameter values for each neuron from an appropriately specified model. Fitted parameter values were consistent across neurons and comported with psychophysical results (the cardinality of L-M and L+M directions and differences in contrast sensitivity in these directions), which gives us confidence that the models described in

this paper are useful for understanding luminance and color processing in V1. Below, we discuss the implications of our results for neurophysiology, psychophysics, and V1 circuitry.

3.4.1 *Implications for neurophysiology*

A fundamental problem in visual neurophysiology is that, in the absence of prior information, measuring a tuning curve generally requires the presentation of many stimuli, as demonstrated by this study. This problem becomes exponentially harder as the dimensionality of the stimulus space increases. The results of our study mitigate these data requirements for the purposes of measuring tuning in the LM plane. The tuning of most V1 neurons throughout the LM plane was largely predictable from responses to luminance and chromatic modulations alone. These color directions are therefore preferable to cone-isolating stimuli for characterizing the color tuning of V1 neurons. The fact that contrast sensitivity was systematically higher for chromatic than luminance stimuli provides a guide to the contrasts needed for efficient neuronal characterization in the LM plane.

3.4.2 *V1 neurons are broadly tuned for L- and M-cone modulations*

Although we found none, V1 neurons with narrower-than-linear color tuning have been reported previously (Lennie, Krauskopf et al. 1990, Cottaris and De Valois 1998, De Valois, Cottaris et al. 2000, Hanazawa, Komatsu et al. 2000, Horwitz and Hass 2012). The apparent discrepancy between these studies and ours is due in part to the definition of "narrower-than-linear". Most previous studies have not distinguished nonlinear cone signal combination from nonlinear contrast response functions; in our study, we were able to distinguish these (see also Pointer and Pointer 1989). We identified many neurons with nonlinear contrast-response functions, which would likely have been classified as having narrower-than-linear tuning in some previous

studies, but no neurons that combined L- and M-cone signals nonlinearly to produce narrower-than-linear color tuning.

The absence of narrower-than-linear color tuning in our data may also be related to the size of the stimulus we used: most previous studies used stimuli that impinged on the extraclassical surround. Stimulation of the surround reduces neuronal sensitivity, potentially narrowing tuning by reducing responses in stimulus directions to which the surround is sensitive (Ts'o and Gilbert 1988, Wachtler, Albright et al. 2001, Solomon, Peirce et al. 2004). Although we did not observe narrower-than-linear color tuning in response to the drifting Gabors, we tested very few neurons with this stimulus (n=14). An additional possibility is that narrower-than-linear color tuning only appears in conjunction with S-cone modulation, which is strongly suppressive to some V1 neurons (Horwitz and Hass 2012).

3.4.3 *Functional significance of broadly tuned neurons*

Some neurons that we recorded responded to modulations in every direction through the LM plane. Understanding how these neurons contribute to vision will require a careful analysis of their spatiotemporal tuning. Most neurons in the present study were probed with stimuli of a single spatiotemporal structure, so our data provide little ground for speculation on this point. However, several previous studies have reported neurons that respond to both luminance and chromatic modulations, and these neurons tended to be orientation- and spatial frequency-tuned (Thorell, De Valois et al. 1984, Johnson, Hawken et al. 2004, Conway and Livingstone 2006). A likely possibility is that these neurons pool signals across cone-opponent and non-opponent channels to achieve a goal for which the distinction between luminance and chromaticity is not useful, for example, the identification of edges irrespective of the light spectra on either side of the boundary. An important future direction is to relate the responses of broadly tuned V1 neurons to light edges

produced by different physical phenomena (e.g. shadows, highlights, changes in material), which can be distinguished on the basis of nonlinear combinations of luminance and chromaticity gradients (Rovamo, Virsu et al. 1982).

Some broadly tuned neurons were unresponsive to modulations in a narrow range of color directions. Most of these neurons responded to both polarities of luminance (ON and OFF) and either red or green. These neurons may contribute to color-luminance interactions observed psychophysically (Virsu, Rovamo et al. 1982, Rovamo, Taskinen et al. 1984). Horwitz, Chichilnisky et al. (2005) identified a similar class of color-luminance neurons using spatiotemporal white noise, and they suggested that these neurons may facilitate blue/yellow signals in the presence of luminance contrast. However, we did not find any neurons for which luminance modulations clearly facilitated responses to chromatic modulations. Instead, luminance modulations were sufficient to drive responses by themselves, inconsistent with facilitation. One possibility is that facilitation only occurs in conjunction with S-cone signals, which were not modulated in our study.

3.4.4 *Implications for detection psychophysics*

A new result from this study is that luminance sensitivity is lower than chromatic sensitivity in V1 neurons. This observation is consistent with the fact that chromatic flashes are more easily seen than luminance flashes by humans (Virsu, Rovamo et al. 1982, Chaparro, Stromeyer et al. 1993) and monkeys (unpublished observations). It is also consistent with those of Koenderink and van Doorn (1978), who found that blood oxygen level dependent responses in V1 are more closely related to detection thresholds than cone contrast.

Psychophysical detection thresholds in the LM plane are well-approximated as pairs of straight lines orthogonal to the luminance and red-green chromatic mechanism directions (Virsu,

Rovamo et al. 1982, Giulianini and Eskew 1998). These properties are consistent with detection being mediated by luminance- and chromaticity-tuned V1 neurons exclusively. On the other hand, some color-luminance neurons are sufficiently sensitive that they may also contribute to detection (for a related result see Hass and Horwitz 2013).

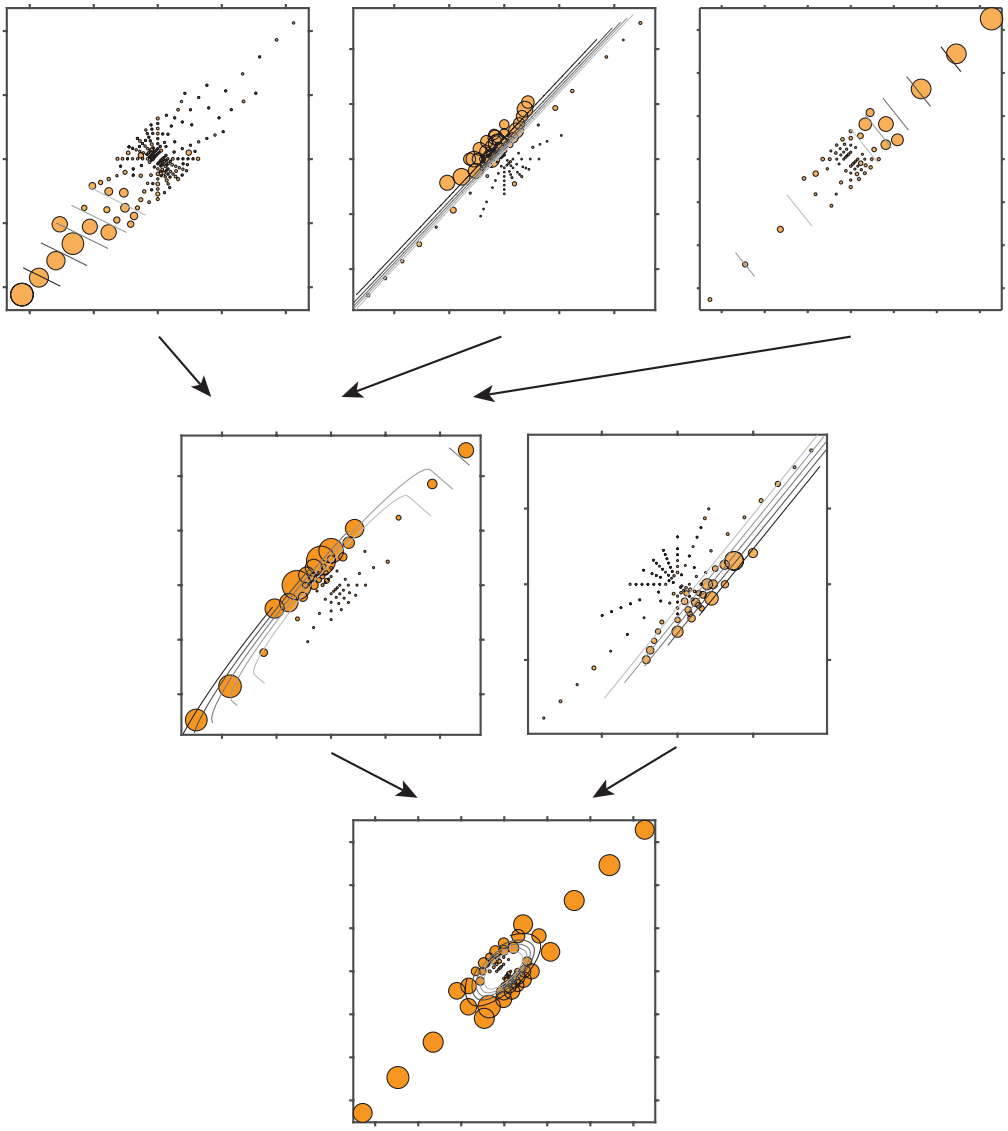

Figure 3.12 Proposed hierarchical relationship.
The responses of multi-directional LNLN neurons may arise from the convergence of three LN neurons. The responses of omni-directional LNLN neurons may likewise arise from the convergence of a multi-directional LNLN neuron and an LN neuron.

3.4.5 *Hierarchy of LN and LNLN neurons*

Color-luminance neurons may achieve their broad tuning via convergent input from luminance- and chromaticity-tuned V1 neurons. Support for this idea comes from the fact that luminance- and chromaticity-tuned neurons combine L- and M-cone signals approximately linearly and have nonlinear contrast-response functions. Adding the responses of such neurons would be expected to produce the nonlinear color-luminance tuning that we observed. The fact that neurons described by the LN model were tuned dominantly for luminance and red-green chromaticity is consistent with the orientations of axes from the LNLN model fits. One possibility is that the neurons that respond to modulations in all but one direction in the LM plane are an intermediate step in the construction of neurons that respond in all directions (Figure 3.12).

A hierarchical relationship between neurons well fit by the LN and LNLN models is consistent with the hierarchical relationship presumed to exist between simple and complex cells. V1 simple cells combine cone signals linearly across space, and many of the neurons we recorded appeared to be simple cells. We also recorded from presumed complex cells that responded to ON- and OFF-luminance modulations throughout their RFs (as assessed from their STAs), although we did not perform the F1/F0 measurement needed for a decisive classification. Note that these neurons were well fit by the LN model despite their full-wave response rectification due to their single preferred color direction. Empirically, only cells that responded to both luminance and chromatic modulations were better fit by the LNLN model. These color-luminance-tuned cells may combine signals from a population of luminance- and chromaticity-tuned simple cells to eliminate information about spatial phase, contrast polarity, and color direction. A critical role for double-opponent neurons, many of which appear to be essentially cone-opponent simple cells, may therefore be to serve as the building blocks for color-luminance-tuned complex cells.

Chapter 4. CONCLUSION

In the second chapter of my thesis, I discussed the LN model of cone signal combination, techniques for estimating a neuron's preferred direction using its responses to stimuli with various distributions, and the transformation of those stimulus distributions and estimated linear mechanisms between linearly related color spaces. In the third chapter, I presented my work developing quantitative models that describe the spiking responses of V1 neurons to modulation in the L- and M-cone photoreceptors. In order to accomplish this goal, I recorded the responses of V1 neurons to uniformly colored flashes presented within the classical RF. These flashes were designed to modulate the L- and M-cones over a broad range of color directions and contrast levels. I found that most V1 neurons combined signals from the L- and M-cones approximately linearly under the conditions of my experiment. The vast majority of these neurons were tuned for equal modulations of L- and M-cones, rendering them sensitive to luminance but not red-green chromaticity, or equal and opposite modulations of L- and M-cones, rendering them sensitive to red-green chromaticity and not luminance. Conversely, I found that a substantial minority of V1 neurons responded to both luminance and color, but that these neurons were poorly described by a linear combination of cone signals. Instead, these color-luminance neurons could be described by a combination of two linear mechanisms, and that these mechanisms were also tuned for luminance and chromaticity. Across the V1 population, I observed a higher sensitivity to chromatic modulation than to luminance modulation. I concluded this chapter with the suggestion that these neurons are hierarchically related, with color-only and luminance-only neurons converging to form color-luminance neurons.

The overarching goal of my graduate work was to advance our understanding of color tuning in V1 neurons. To achieve this goal, I developed quantitative models to connect the activity

in cone photoreceptors to spiking responses in V1. In the Introduction, I argued that building and a rigorously testing quantitative models of color tuning in V1 requires a large amount of data, and that it is difficult to collect a sufficiently broad sampling of 3-dimensional color space under the conditions of a single-cell electrophysiology experiment in an awake, behaving monkey. Therefore, in order to sample in a sufficiently broad manner, we reduced the dimensionality of the full 3-dimensional color space by modulating only the L- and M-cone photoreceptors and leaving the activity of the S-cones constant. However, further advancement toward understanding color tuning in V1 requires building more comprehensive quantitative models that include contributions from the S-cones and predict V1 neuronal responses to arbitrary combinations of cone signals.

Building quantitative models that are rich enough to describe color tuning in a full 3-dimensional color space is beyond the scope of my thesis. However, my thesis work can help to advance future research toward this goal by taking advantage of the structure in the color tuning that I observed in the LM plane. By approaching the task of sampling a fully 3-dimensional color space with prior information, it may be possible to reduce the number of stimuli needed to estimate color tuning, making such estimation possible under the constraints of an electrophysiology experiment in an awake, behaving monkey.

In particular, I suggest that color-tuning in the LM plane can be estimated efficiently using a small number of stimuli allowing experimenters more time to vary other stimulus parameters, such as S-cone contrast, to build more comprehensive models of color-tuning. To test this hypothesis, I used the recorded responses of V1 neurons from Chapter III of my thesis. For each neuron, I subsampled the observed stimulus-response pairs to generate response predictions. Generally, I found that different responses were predicted by using different distributions of

subsampled stimuli. I assessed the accuracy of different subsampling regimes by comparing the responses to all of the stimuli to the response predictions generated by the subset.

In the analysis that follows, I compared four potential methods for estimating color-tuning in the LM plane. These methods employed different subsampling regimes and assumptions about color tuning. Here, I describe each of the four methods that I used to estimate color-tuning, the subsampling regime employed by each, and associated assumptions.

The first method is based on a classical technique for estimating cone weights by fitting a linear model to neuronal responses (Lennie, Krauskopf et al. 1990, Johnson, Hawken et al. 2004, Solomon, Peirce et al. 2004). This method is typically applied by presenting a small number of stimuli in the neuron's RF—such as an L-, an M-, and an S-cone-isolating stimulus—of arbitrary contrast (usually determined by the display gamut), measuring the neuron's responses to each stimulus, and fitting a linear model to the response magnitudes to estimate the weight associated with cone-type. This method is built on two assumptions: first, a linear model of cone signal combination is an appropriate description of color tuning for all V1 neurons, and second, the contrast-response functions of V1 neurons are approximately linear. Our expectation is that this method will not be particularly accurate because neither of these assumptions are well founded. As demonstrated in Chapter III of my thesis, a linear model is a poor description of many V1 neurons, and contrast-response functions are not typically linear (average exponent (N) value = 3.6). To test the accuracy of this method, I used the single L- and M-cone isolating stimuli that generated the greatest contrast (Figure 4.1A, orange circles). The responses to these stimuli were fit with a linear model of cone signal combination:

$$R = max\left([L\ M]\begin{bmatrix}A\\B\end{bmatrix}, 0\right) + u * max\left([L\ M]\begin{bmatrix}-A\\-B\end{bmatrix}, 0\right) \quad (4.1)$$

where R is the response of the neuron, [L M] is the L- and M-cone contrast of the stimuli, $\begin{bmatrix} A \\ B \end{bmatrix}$ are the weights associated with each cone-type, and u represents the neuron's differential sensitivity to each stimulus polarity along the axis $\begin{bmatrix} A \\ B \end{bmatrix}$. Response estimates were obtained using maximum likelihood, as described in Chapter III.

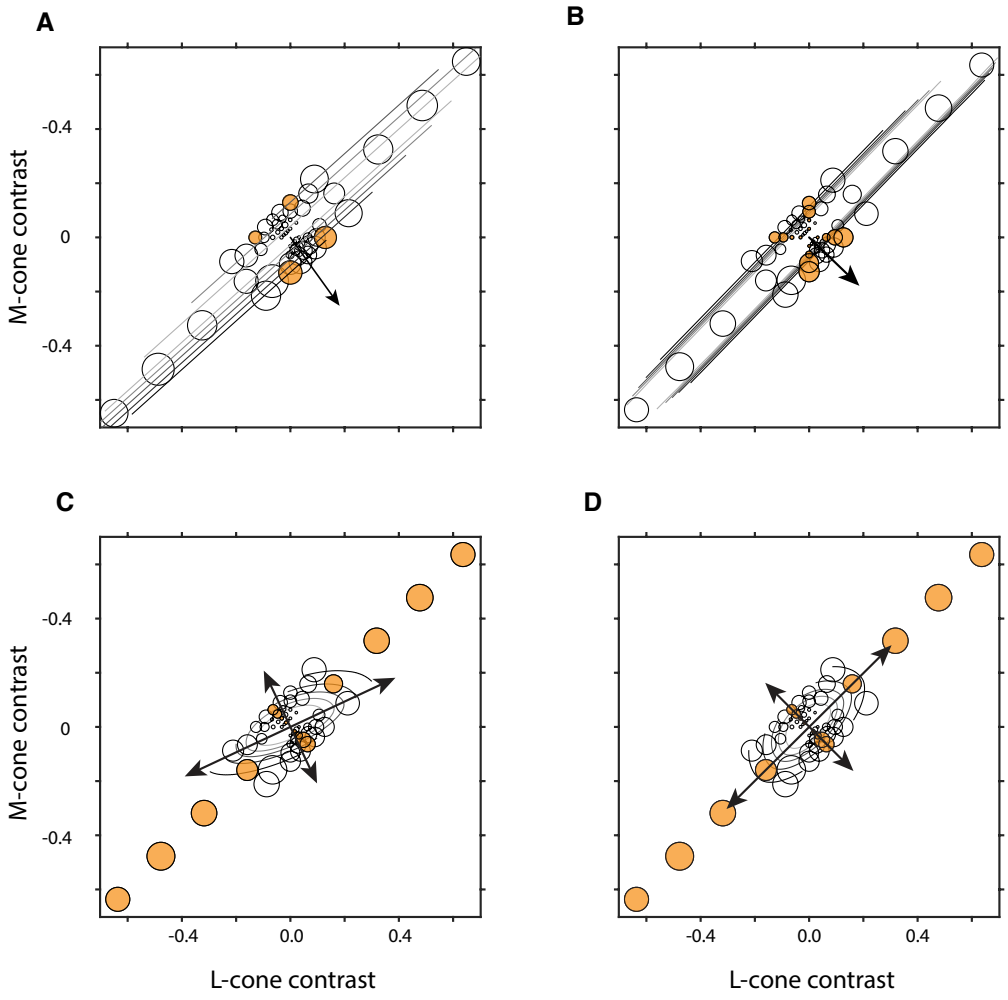

Figure 4.1 Four methods for estimating color tuning in the LM plane.
Each method is used to fit the responses of an example neuron. (A) A linear model based on single cone-isolating stimuli. (B) The LNLN model is freely fit to cone-isolating stimuli of various intensities. (C) The LNLN model is freely fit to cardinal luminance and chromatic stimuli of various intensities. (D) The LNLN model is fit to cardinal luminance and chromatic stimuli of various intensities, and the principal axes are constrained to the luminance and chromatic directions.

The second method estimates color tuning using neuronal responses to various contrasts of cone-isolating stimuli (Figure 4.1B, orange circles). Using a multitude of stimulus contrasts allows us to estimate both the linear cone weights and the nonlinear contrast-response function. Like the previous method, this second method assumes that the most useful subset of stimuli for estimating color tuning are cone-isolating. Unlike the previous method, however, this method does not assume that contrast-response functions are linear. To test the accuracy of this method, the responses to all L- and M-cone isolating stimuli were fit with the LNLN model (eq. 3.4) using maximum likelihood, as described in Chapter III. Because the LN model is a special case of the LNLN model, this broader model is able to accurately describe both varieties of neurons.

The third method estimates color tuning using neuronal responses to various contrasts of stimuli along the cardinal luminance and chromatic axes. This third method is identical to the second and shares the same assumptions, with the exception that color tuning is more accurately estimated using the responses to luminance-only and chromatic-only stimuli than the responses to cone-isolating stimuli (Figure 4.1C, orange circles). To test the accuracy of this method, I fit the observed responses to luminance-only and chromatic-only modulations with the LNLN model using maximum likelihood, as described in Chapter III.

The fourth method estimates color tuning using the same set of responses as the third method. Under the previous three models, the estimated cone weights were freely fit to the data. Under the fourth method, however, preferred directions are constrained to align along the luminance and chromatic axes. This model is built on the assumption that neurons are sensitive to an energy-like combination of luminance and chromatic modulations under the LNLN model and

that the color tuning of a V1 neuron may be quickly ascertained from its responses to luminance-only and chromatic-only modulations (Figure 4.1D, orange circles).

Each method was used to estimate the responses of all 118 neurons in the population. The accuracy of each method was assessed by comparing these response predictions to all of the actual observed responses (Figure 4.1). The quality of each response prediction was quantified as the sum of squared error (SSE) between all of the observed responses and the predictions. Predictions generated by methods 2-4 were generated using approximately the same number of stimulus-response pairs.

I found that distributions of SSE values for each of the four methods across all 118 neurons were unlikely to come from the same distribution (Kruskal-Wallis H test, $p = 4.33 \times 10^{-4}$). Here, I rank each of these methods according to the average SSE values across the population. The first method, which predicted responses using only four cone-isolating stimuli (Figure 4.1A), generated the highest average SSE values (mean SSE = 1.4×10^4). The relatively poor performance of this method is not unexpected, however, as it is based on the fewest observed responses and only considers a linear model when generating response predictions. The second method, which predicted responses based on all of the observed cone-isolating stimuli (Figure 4.1B), generated the second highest average SSE values (mean SSE = 7.8×10^3). The distribution of SSE values generated using method 2 was substantially lower than those generated using the method 1 (paired t-test, p-value = 0.16, Figure 4.2A). This improvement in performance is likely due to the facts that response predictions were generated using a greater number of observed responses and fit with the more flexible LNLN model. The third method, which predicted responses based on a multitude of cardinal luminance and chromatic modulations (Figure 4.1C), generated the third highest average SSE values (mean SSE = 4.2×10^3). The distribution of SSE value generated using method

3 was significantly lower than those generated using the method 2 (paired t-test, p-value = 1.4×10^{-4}, Figure 4.2B). The fourth method, which predicted responses using an LNLN model whose asymmetric and symmetric axes were aligned to the luminance and chromatic directions (Figure 4.1D), generated the lowest SSE values (mean SSE = 3.8×10^{3}). The distribution of SSE values generated using method 4 was significantly lower than those generated using the method 3 (paired t-test, p-value = 0.055, Figure 4.2C). Of the four considered methods, this one performed the best, despite being based on the same amount of data as the second and third methods and being constrained by greater assumptions (fixed axes vs fit axes).

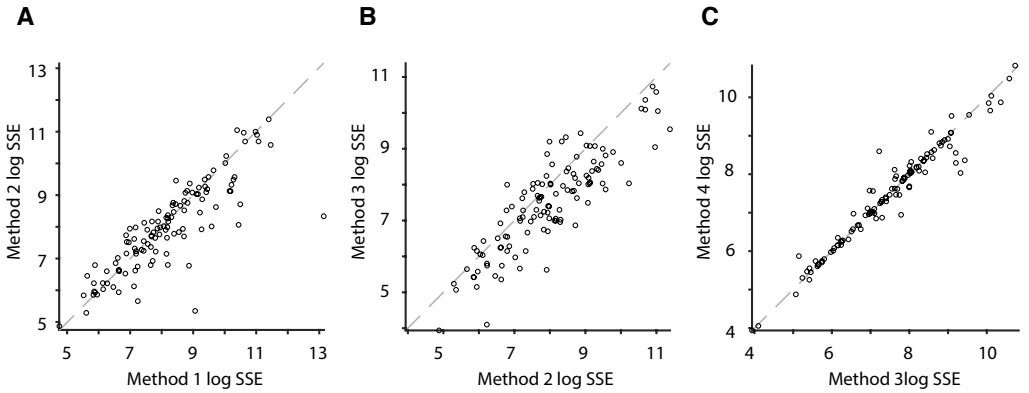

Figure 4.2 Comparing color tuning estimation methods.
Log SSE for all 118 neurons using methods 1 and 2. Dotted line is unity line. (B) Log SSE for all 118 neurons using methods 2 and 3. (C) Log SSE for all 118 neurons using methods 3 and 4.

In summary, I compared four methods for subsampling the LM plane and fitting neuronal responses. The most successful method sampled neuronal responses to luminance and chromatic modulations, and fit an LNLN model to the observed responses whose axes were aligned to the cardinal luminance and chromatic directions. This method was based on prior information that was ascertained from the experiment in Chapter III: that the color tuning in V1 was well described as a function of one or two weighted sums of L- and M-cone modulations, and that these weighted

sums were consistently aligned along the luminance and chromatic axes. Using this method, researcher can assess the color tuning of neurons in the LM plane quickly and efficiently by sampling the responses to a small number of stimuli along the luminance and chromatic axes and fitting their responses with a model that assumes that the principal axes are oriented in these same directions. This economy gives experimenters more time to present other stimuli, such as those that modulate the S-cones. For example, a future experiment might quickly estimate color tuning in the LM plane, then use the remainder of experimental time to estimate color tuning in another 2-dimensional slices of color space, such as the SM or the SL planes. Building a comprehensive 3-dimensional quantitative model of cone signal combination in V1 will be challenging, but we have demonstrated that a high density but lower-dimensional approach can help to advance this work by exposing structure that would likely be missed from sparser, higher-dimensional sampling. This strategy will undoubtedly reveal additional patterns in color tuning that will shed light on the mechanisms of color vision.

BIBLIOGRAPHY

Angueyra, J. M. and F. Rieke (2013). "Origin and effect of phototransduction noise in primate cone photoreceptors." Nat Neurosci **16**(11): 1692-1700.

Bohon, K. S., K. L. Hermann, T. Hansen and B. R. Conway (2016). "Representation of Perceptual Color Space in Macaque Posterior Inferior Temporal Cortex (the V4 Complex)." eNeuro **3**(4).

Brainard, D. H. (1996). Cone contrast and opponent modulation color spaces. Human Color Vision, 2nd edition. K. a. Boynton. Washington, DC, Optical Society of America.

Bushnell, B. N., P. J. Harding, Y. Kosai, W. Bair and A. Pasupathy (2011). "Equiluminance cells in visual cortical area v4." J Neurosci **31**(35): 12398-12412.

Chaparro, A., C. F. Stromeyer, E. P. Huang, R. E. Kronauer and R. T. Eskew (1993). "Color Is What the Eye Sees Best." Nature **361**(6410): 348-350.

Chichilnisky, E. J. (2001). "A simple white noise analysis of neuronal light responses." Network **12**(2): 199-213.

Cole, G. R., T. Hine and W. McIlhagga (1993). "Detection mechanisms in L-, M-, and S-cone contrast space." J Opt Soc Am A **10**(1): 38-51.

Conway, B. R. (2009). "Color vision, cones, and color-coding in the cortex." Neuroscientist **15**(3): 274-290.

Conway, B. R., S. Chatterjee, G. D. Field, G. D. Horwitz, E. N. Johnson, K. Koida and K. Mancuso (2010). "Advances in Color Science: From Retina to Behavior." Journal of Neuroscience **30**(45): 14955-14963.

Conway, B. R. and M. S. Livingstone (2006). "Spatial and temporal properties of cone signals in alert macaque primary visual cortex." J Neurosci **26**(42): 10826-10846.

Cottaris, N. P. (2003). "Artifacts in spatiochromatic stimuli due to variations in preretinal absorption and axial chromatic aberration: implications for color physiology." J Opt Soc Am A Opt Image Sci Vis **20**(9): 1694-1713.

Cottaris, N. P. and R. L. De Valois (1998). "Temporal dynamics of chromatic tuning in macaque primary visual cortex." Nature **395**(6705): 896-900.

Cowell, R. L., R. D. Tyler, K. D. Clinkenbeard and C. G. MacAllister (1987). "Collection and evaluation of equine peritoneal and pleural effusions." Vet Clin North Am Equine Pract **3**(3): 543-561.

Dacey, D. M. (1999). "Primate retina: cell types, circuits and color opponency." Prog Retin Eye Res **18**(6): 737-763.

De Valois, R. L., N. P. Cottaris, S. D. Elfar, L. E. Mahon and J. A. Wilson (2000). "Some transformations of color information from lateral geniculate nucleus to striate cortex." Proc Natl Acad Sci U S A **97**(9): 4997-5002.

Deluca, V., J. Balsevich, R. T. Tyler and W. G. Kurz (1987). "Characterization of a novel N-methyltransferase (NMT) from Catharanthus roseus plants : Detection of NMT and other enzymes of the indole alkaloid biosynthetic pathway in different cell suspension culture systems." Plant Cell Rep **6**(6): 458-461.

Dennis, M. J., N. Howarth, P. E. Key, M. Pointer and R. C. Massey (1989). "Investigation of ethyl carbamate levels in some fermented foods and alcoholic beverages." Food Addit Contam **6**(3): 383-389.

Derrington, A. M., J. Krauskopf and P. Lennie (1984). "Chromatic mechanisms in lateral geniculate nucleus of macaque." J Physiol **357**: 241-265.

Derrington, A. M. and P. Lennie (1984). "Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque." J Physiol **357**: 219-240.

Dodd, D. E., P. E. Losco, C. M. Troup, I. M. Pritts and T. R. Tyler (1987). "Hyalin droplet nephrosis in male Fischer-344 rats following inhalation of diisobutyl ketone." Toxicol Ind Health **3**(4): 443-457.

Eskew, R. T., J. S. McLellan and F. Giulianini (1999). Chromatic detection and discrimination. Color vision: From genes to perception**:** 345-368.

Gegenfurtner, K. R. (2003). "Cortical mechanisms of colour vision." Nat Rev Neurosci **4**(7): 563-572.

Gegenfurtner, K. R. and M. J. Hawken (1995). "Temporal and chromatic properties of motion mechanisms." Vision Res **35**(11): 1547-1563.

Giulianini, F. and R. T. Eskew, Jr. (1998). "Chromatic masking in the (delta L/L, delta M/M) plane of cone-contrast space reveals only two detection mechanisms." Vision Res **38**(24): 3913-3926.

Giulianini, F. and R. T. Eskew, Jr. (2007). "Theory of chromatic noise masking applied to testing linearity of S-cone detection mechanisms." J Opt Soc Am A Opt Image Sci Vis **24**(9): 2604-2621.

Hanazawa, A., H. Komatsu and I. Murakami (2000). "Neural selectivity for hue and saturation of colour in the primary visual cortex of the monkey." Eur J Neurosci **12**(5): 1753-1763.

Hass, C. A. and G. D. Horwitz (2013). "V1 mechanisms underlying chromatic contrast detection." Journal of Neurophysiology **109**(10): 2483-2494.

Hess, R. F. and J. S. Pointer (1989). "Spatial and temporal contrast sensitivity in hemianopia. A comparative study of the sighted and blind hemifields." Brain **112 (Pt 4)**: 871-894.

Horwitz, G. D., E. J. Chichilnisky and T. D. Albright (2005). "Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1." Journal of Neurophysiology **93**(4): 2263-2278.

Horwitz, G. D., E. J. Chichilnisky and T. D. Albright (2007). "Cone inputs to simple and complex cells in V1 of awake macaque." Journal of Neurophysiology **97**(4): 3070-3081.

Horwitz, G. D. and C. A. Hass (2012). "Nonlinear analysis of macaque V1 color tuning reveals cardinal directions for cortical color processing." Nat Neurosci **15**(6): 913-919.

Hurvich, L. M. and D. Jameson (1957). "An opponent-process theory of color vision." Psychol Rev **64, Part 1**(6): 384-404.

Ingling, C. R., Jr. and B. H. P. Tsou (1977). "Orthogonal combination of the three visual channels." Vision Res **17**(9): 1075-1082.

Johnson, E. N., M. J. Hawken and R. Shapley (2001). "The spatial transformation of color in the primary visual cortex of the macaque monkey." Nat Neurosci **4**(4): 409-416.

Johnson, E. N., M. J. Hawken and R. Shapley (2004). "Cone inputs in macaque primary visual cortex." Journal of Neurophysiology **91**(6): 2501-2514.

Kalloniatis, M. and R. S. Harwerth (1991). "Effects of chromatic adaptation on opponent interactions in monkey increment-threshold spectral-sensitivity functions." J Opt Soc Am A **8**(11): 1818-1831.

Koenderink, J. J., M. A. Bouman, A. E. Bueno de Mesquita and S. Slappendel (1978). "Perimetry of contrast detection thresholds of moving spatial sine patterns. II. The far peripheral visual field (eccentricity 0 degrees-50 degrees)." J Opt Soc Am **68**(6): 850-854.

Koenderink, J. J., M. A. Bouman, A. E. Bueno de Mesquita and S. Slappendel (1978). "Perimetry of contrast detection thresholds of moving spatial sine wave patterns. I. The near peripheral visual field (eccentricity 0 degrees-8 degrees)." J Opt Soc Am **68**(6): 845-849.

Koenderink, J. J., M. A. Bouman, A. E. Bueno de Mesquita and S. Slappendel (1978). "Perimetry of contrast detection thresholds of moving spatial sine wave patterns. III. The target extent as a sensitivity controlling parameter." J Opt Soc Am **68**(6): 854-860.

Koenderink, J. J., M. A. Bouman, A. E. Bueno de Mesquita and S. Slappendel (1978). "Perimetry of contrast detection thresholds of moving spatial sine wave patterns. IV. The influence of the mean retinal illuminance." J Opt Soc Am **68**(6): 860-865.

Koenderink, J. J. and A. J. van Doorn (1978). "Detectability of power fluctuations of temporal visual noise." Vision Res **18**(2): 191-195.

Koenderink, J. J. and A. J. van Doorn (1978). "Visual detection of spatial contrast; influence of location in the visual field, target extent and illuminance level." Biol Cybern **30**(3): 157-167.

Komatsu, H. (1998). "Mechanisms of central color vision." Current Opinion in Neurobiology **8**(4): 503-508.

Komatsu, H., Y. Ideura, S. Kaji and S. Yamane (1992). "Color Selectivity of Neurons in the Inferior Temporal Cortex of the Awake Macaque Monkey." Journal of Neuroscience **12**(2): 408-424.

Krauskopf, J., D. R. Williams and D. W. Heeley (1982). "Cardinal directions of color space." Vision Res **22**(9): 1123-1131.

Lapp, C. A., J. M. Tyler, M. E. Stachura and Y. S. Lee (1987). "Deleterious effects of fungizone on growth hormone and prolactin secretion by cultured GH3 cells." In Vitro Cell Dev Biol **23**(12): 837-840.

Leblanc, R., Y. L. Yamamoto, J. L. Tyler, M. Diksic and A. Hakim (1987). "Borderzone ischemia." Ann Neurol **22**(6): 707-713.

Lee, B. B., P. R. Martin, A. Valberg and J. Kremers (1993). "Physiological mechanisms underlying psychophysical sensitivity to combined luminance and chromatic modulation." J Opt Soc Am A **10**(6): 1403-1412.

Lennie, P., J. Krauskopf and G. Sclar (1990). "Chromatic mechanisms in striate cortex of macaque." J Neurosci **10**(2): 649-669.

Lennie, P., J. Pokorny and V. C. Smith (1993). "Luminance." J Opt Soc Am A **10**(6): 1283-1293.

Lindbloom-Brown, Z., L. J. Tait and G. D. Horwitz (2014). "Spectral sensitivity differences between rhesus monkeys and humans: Implications for neurophysiology." Journal of Neurophysiology.

Luke, R. D., J. W. Begun and D. D. Pointer (1989). "Quasi firms: strategic interorganizational forms in the health care industry." Acad Manage Rev **14**(1): 9-19.

May, W. S. and G. Tyler (1987). "Phosphorylation of the surface transferrin receptor stimulates receptor internalization in HL60 leukemic cells." J Biol Chem **262**(34): 16710-16718.

Metha, A. B., A. J. Vingrys and D. R. Badcock (1994). "Detection and discrimination of moving stimuli: the effects of color, luminance, and eccentricity." J Opt Soc Am A Opt Image Sci Vis **11**(6): 1697-1709.

Paninski, L., J. Pillow and J. Lewi (2007). "Statistical models for neural encoding, decoding, and optimal stimulus design." Prog Brain Res **165**: 493-507.

Pillow, J. W., J. Shlens, L. Paninski, A. Sher, A. M. Litke, E. J. Chichilnisky and E. P. Simoncelli (2008). "Spatio-temporal correlations and visual signalling in a complete neuronal population." Nature **454**(7207): 995-999.

Pointer, D. D. (1989). "Graduate education for the working professional." Healthc Exec **4**(1): 33-35.

Pointer, D. D., L. J. Dean and A. J. DeLellis (1989). "The MCV/VCU executive master's program: an experiment in the decentralization of health administration education." J Health Adm Educ **7**(2): 339-356.

Pointer, D. D. and T. K. Pointer (1989). "Case-based prospective price reimbursement." Nurs Manage **20**(4): 30-32, 34.

Pointer, J. E., G. E. Gauger, P. H. Kwok, B. Simon, M. E. Smith, S. W. Van Meter and B. K. Walker (1989). "Safety of axial traction." Ann Emerg Med **18**(4): 428-429.

Pointer, J. E. and M. A. Osur (1989). "Defibrillator study: just the first step." J Emerg Med **7**(2): 203.

Pointer, J. E. and M. A. Osur (1989). "Effect of standing orders on field times." Ann Emerg Med **18**(10): 1119-1121.

Pointer, J. S. and R. F. Hess (1989). "The contrast sensitivity gradient across the human visual field: with emphasis on the low spatial frequency range." Vision Res **29**(9): 1133-1151.

Roorda, A., A. B. Metha, P. Lennie and D. R. Williams (2001). "Packing arrangement of the three cone classes in primate retina." Vision Res **41**(10-11): 1291-1306.

Roth, K. A., R. A. Unanue, J. Leykam and A. N. Tyler (1987). "Isolation and characterization of beta-endorphin-(1-9) from human and rat pituitaries." Regul Pept **19**(5-6): 335-344.

Rovamo, J., L. Leinonen, P. Laurinen and V. Virsu (1984). "Temporal integration and contrast sensitivity in foveal and peripheral vision." Perception **13**(6): 665-674.

Rovamo, J. and V. Virsu (1984). "Isotropy of cortical magnification and topography of striate cortex." Vision Res **24**(3): 283-286.

Rovamo, J., V. Virsu, P. Laurinen and L. Hyvarinen (1982). "Resolution of gratings oriented along and across meridians in peripheral vision." Invest Ophthalmol Vis Sci **23**(5): 666-670.

Rovamo, L., M. R. Taskinen, T. Kuusi, E. A. Nikkila, C. Ehnholm and K. O. Raivio (1984). "Postheparin plasma lipase activities and plasma lipoproteins in newborn infants." Pediatr Res **18**(7): 642-647.

Rust, N. C., V. Mante, E. P. Simoncelli and J. A. Movshon (2006). "How MT cells analyze the motion of visual patterns." Nat Neurosci **9**(11): 1421-1431.

Sankeralli, M. J. and K. T. Mullen (1996). "Estimation of the L, M-, and S-cone weights of the postreceptoral detection mechanisms." Journal of the Optical Society of America a-Optics Image Science and Vision **13**(5): 906-915.

Sankeralli, M. J. and K. T. Mullen (2001). "Assumptions concerning orthogonality in threshold-sealed versus cone-contrast colour spaces." Vision Research **41**(1): 53-55.

Sharpe, C. J. (1974). "Letter: Propranolol in the treatment of migraine." Br Med J **3**(5929): 522.

Sharpe, C. R. (1974). "Letter: The contrast sensitivity of the peripheral visual field to drifting sinusoidal gratings." Vision Res **14**(9): 905-906.

Solomon, S. G. and P. Lennie (2005). "Chromatic gain controls in visual cortical neurons." J Neurosci **25**(19): 4779-4792.

Solomon, S. G. and P. Lennie (2007). "The machinery of colour vision." Nat Rev Neurosci **8**(4): 276-286.

Solomon, S. G., J. W. Peirce, N. T. Dhruv and P. Lennie (2004). "Profound contrast adaptation early in the visual pathway." Neuron **42**(1): 155-162.

Solomon, S. G., J. W. Peirce and P. Lennie (2004). "The impact of suppressive surrounds on chromatic properties of cortical neurons." J Neurosci **24**(1): 148-160.

Stockman, A., D. I. MacLeod and N. E. Johnson (1993). "Spectral sensitivities of the human cones." J Opt Soc Am A Opt Image Sci Vis **10**(12): 2491-2521.

Stoughton, C. M. and B. R. Conway (2008). "Neural basis for unique hues." Curr Biol **18**(16): R698-699.

Stromeyer, C. F., 3rd, R. E. Kronauer, A. Ryu, A. Chaparro and R. T. Eskew, Jr. (1995). "Contributions of human long-wave and middle-wave cones to motion detection." J Physiol **485 (Pt 1)**: 221-243.

Sun, H., H. E. Smithson, Q. Zaidi and B. B. Lee (2006). "Specificity of cone inputs to macaque retinal ganglion cells." Journal of Neurophysiology **95**(2): 837-849.

Thorell, L. G., R. L. De Valois and D. G. Albrecht (1984). "Spatial mapping of monkey V1 cells with pure color and luminance stimuli." Vision Res **24**(7): 751-769.

Ts'o, D. Y. and C. D. Gilbert (1988). "The organization of chromatic and spatial interactions in the primate striate cortex." J Neurosci **8**(5): 1712-1727.

Tyler, J. (1987). "Give us a break." Nurs Times **83**(50): 32-35.

Tyler, R. D., R. L. Cowell, K. D. Clinkenbeard and C. G. MacAllister (1987). "Hematologic values in horses and interpretation of hematologic data." Vet Clin North Am Equine Pract **3**(3): 461-484.

Tyler, R. D., C. W. Qualls, Jr., R. D. Heald, R. L. Cowell and K. D. Clinkenbeard (1987). "Renal concentrating ability in dehydrated hyponatremic dogs." J Am Vet Med Assoc **191**(9): 1095-1100.

Virsu, V., J. Rovamo, P. Laurinen and R. Nasanen (1982). "Temporal contrast sensitivity and cortical magnification." Vision Res **22**(9): 1211-1217.

Wachtler, T., T. D. Albright and T. J. Sejnowski (2001). "Nonlocal interactions in color perception: nonlinear processing of chromatic signals from remote inducers." Vision Res **41**(12): 1535-1546.

Webb, B. S., N. T. Dhruv, S. G. Solomon, C. Tailby and P. Lennie (2005). "Early and late mechanisms of surround suppression in striate cortex of macaque." J Neurosci **25**(50): 11666-11675.

Weller, J. P. and G. D. Horwitz (2017). "Measurements of neuronal color tuning: Procedures, pitfalls, and alternatives." Vision Res.

Wuerger, S. M., P. Atkinson and S. Cropper (2005). "The cone inputs to the unique-hue mechanisms." Vision Res **45**(25-26): 3210-3223.

VITA

J Patrick Weller earned his Bachelor of the Arts in Liberal Arts at St. John's College in 2008. Afterword, he worked and studied as a PostBac Intramural Research Training Award (IRTA) recipient at the National Institute of Mental Health in the laboratory of Dr. Leslie Ungerleider.