Measurements of neuronal color tuning: Procedures, pitfalls, and alternatives

J. Patrick Weller, Gregory D. Horwitz*

Department of Physiology & Biophysics, Washington National Primate Research Center, University of Washington, Seattle, WA 98195, United States

ARTICLE INFO

ABSTRACT

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 of likelihood. We highlight the advantages and disadvantages of each technique and emphasize the effects of the stimulus distribution on color tuning measurements.

1. Introduction

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 (Bohon, Hermann, Hansen, & Conway, 2016; Conway, 2009; Conway et al., 2010; Gegenfurtner, 2003; Komatsu, 1998; Solomon & Lennie, 2007). Neuronal color tuning measurements can help reveal how and where these transformations occur.

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 spatio-temporal stimulus parameters affect color tuning (Conway & Livingstone, 2006; Cottaris & DeVolois, 1998; Derrington, Krauskopf, & Lennie, 1984; Solomon, Peirce, & Lennie, 2004; Thorell, De Valois, & Albrecht, 1984). 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 (Fig. 1A). In the *cone noise* stimulus, they modulated in ratios selected to stimulate the three cone types independently (Fig. 1B) (see Appendix I for methodological details of the experiment). The average phosphor noise stimulus that preceded a spike from a single example neuron (Fig. 1C) appears different from the average cone noise stimulus that preceded a spike (Fig. 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.

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

---

* Corresponding author at: Department of Physiology & Biophysics, University of Washington, 1959 NE Pacific Street, HSB G-424, Box 357290, Seattle, WA 98195-7330, United States.
 E-mail address: ghorwitz@uw.edu (G.D. Horwitz).
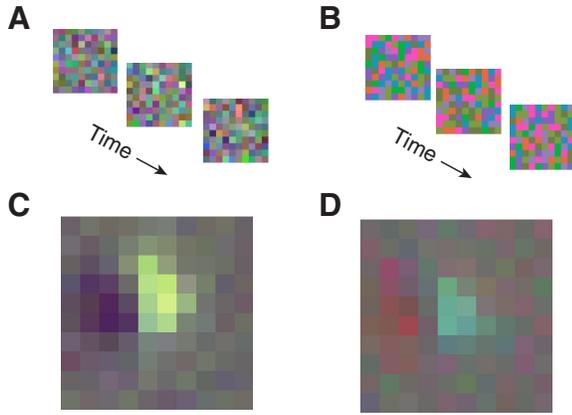
Fig. 1. The responses of an individual V1 neuron to stimuli from two distributions. (a) Stimuli from the phosphor noise distribution. The color of each square was determined by the independent modulation of the three phosphors in a CRT monitor. (b) Stimuli from the cone noise distribution. Each colored square modulated the activity of the three cone 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.

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 (Cole, Hine, & McIlhagga, 1993; D'Zmura & Knoblauch, 1998; Krauskopf, Williams, & Heeley, 1982; Sankeralli & Mullen, 1996), color discrimination (Eskew, McLellan, & Giulianini, 1999; Ingling & Tsou, 1977; Krauskopf & Gegenfurtner, 1992; Wandell, 1985), and appearance judgments (Hurvich & Jameson, 1957; Poirson & Wandell, 1993; Thornton & Pugh, 1983). Some of the issues we discuss regarding color spaces and transformations between them can be found in Knoblauch, 1995; Sankeralli & Mullen, 2001; D'Zmura & Knoblauch, 2001; Brainard, 1996; Brainard & Stockman, 2010; Stockman & Brainard, 2010; Hansen & Gegenfurtner, 2013. Here, we distill from this body of work a few key techniques 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:

$$[v_1 v_2 v_3] \begin{bmatrix} A \\ B \\ C \end{bmatrix} \tag{1}$$

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:

$$R = f\left([v_1 v_2 v_3] \begin{bmatrix} A \\ B \\ C \end{bmatrix}\right) \tag{2}$$

where R typically represents a spike count or spike rate. The function f captures nonlinear response features like 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 (Conway & Livingstone, 2006; Derrington et al., 1984; Horwitz, Chichilnisky, & Albright, 2007; Johnson, Hawken, & Shapley, 2004; Lennie, Krauskopf, & Sclar, 1990; Reid & Shapley, 2002).

2. 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 (RWA) is described mathematically as:

$$RWA = \frac{1}{n} \sum_{i=1}^{n} R_i [v_1 v_2 v_3]_i \tag{3}$$

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 Fig. 1C and D are examples of RWAs.

The RWA can be used to estimate a neuron's preferred color direction (Chichilnisky & Baylor, 1999; Cottaris & DeValois, 1998; Horwitz et al., 2007; Sun, Smithson, Zaidi, & Lee, 2006), 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 (Fig. 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).

The RWA is an unbiased estimator of the preferred direction when the stimulus distribution is radially symmetric (RS) (Fig. 2B) (Chichilnisky, 2001). Intuitively, this is because there exist families of stimuli that give rise to the same weighted sum, and therefore, the same response. Under the LN model, all such stimuli lie on 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,

---

[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 is. Continuously varying stimuli can also be used to create a radially symmetric (RS) distribution (Sun 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.
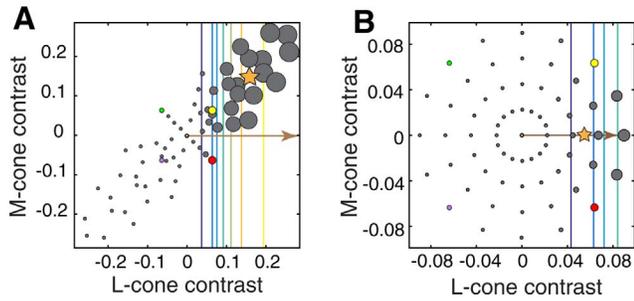
Fig. 2. Responses of a hypothetical neuron that receives exclusively L-cone input. For simplicity, only L- and M-cone modulations are simulated in this example. The position of each gray point represents a stimulus in the LM plane. The size of each point represents the magnitude of the corresponding response. Contour lines, from cool to warm colors, represent the increasing responses of the neuron to stimuli of progressively higher contrast. The brown 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.

Fig. 3. 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.

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.

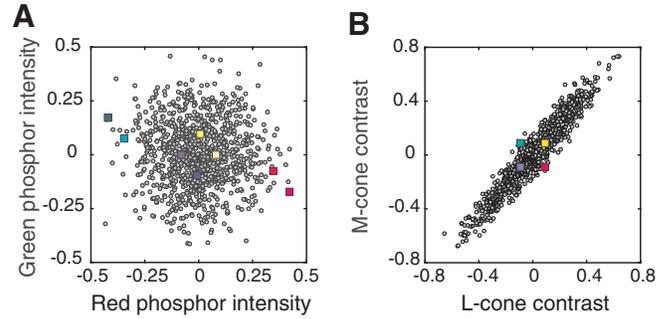### 2.1. The representation and transformation of lights and weights

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 Fig. 1: a phosphor space and a cone space. In both spaces, the origin $[0\,0\,0]$ 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:

$$[LMS] = [RGB] * M \tag{4}$$

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

Neural responses do not depend on how an experimenter choses 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. (4)), then the weights are transformed by its inverse transpose (M^{-T}):

$$[lms] = [rgb] * M^{-T} \tag{5}$$

where $[r\,g\,b]$ and $[l\,m\,s]$ are the weights in phosphor and cone space, respectively, that reflect the same color tuning. Combining Eqs. (4) and (5) clarifies the relationship:

$$R = [LMS] * \begin{bmatrix} l \\ m \\ s \end{bmatrix} = [RGB] * M * M^{-1} * \begin{bmatrix} r \\ g \\ b \end{bmatrix} \tag{6}$$

where $M^{-1} * \begin{bmatrix} r \\ g \\ b \end{bmatrix}$ is simply a rearrangement of terms in Eq. (5).

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 1).

2.2. An experimental example

We return to the example neuron from Fig. 1, which was probed with two stimulus distributions, each of which was RS in the color space in which it was constructed (Fig. 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 Eq. (5). By convention, we represent both estimates in cone contrast space. Both estimates indicate similar color tuning (Fig. 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 (Fig. 4A, gray symbols) and the high correlations between them (Fig. 4B).

3. 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 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 an approximation that comes close: multiplying the stimuli with a whitening matrix. This multiplication

Table 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 tophosphor space)
Stimuli $[v_1\ v_2\ v_3]$	M	M^{-1}
Weights $[A\ B\ C]$	M^{-T}	M^T
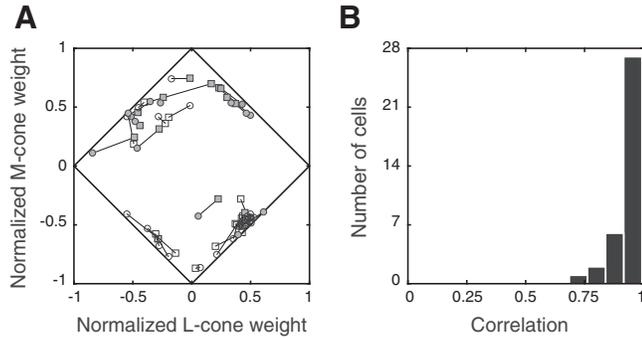
Fig. 4. (a) Normalized cone weight estimates from a population of V1 neurons (n = 35). Two cone weight estimates (connected with black lines) were obtained from each neuron: one from the responses to the phosphor noise (circles) and one from the responses to the cone noise (squares). Shaded symbols indicate positive S-cone weights, and unshaded symbols indicate negative S-cone weights. The example neuron from Fig. 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 most of the neurons tested.

Fig. 5. Normalized cone weight estimates for the example neuron from Fig. 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.

transforms the stimuli so that they have variance 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.

4. 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 et al., 2008). Through numerical optimization, all of these parameters can be estimated 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
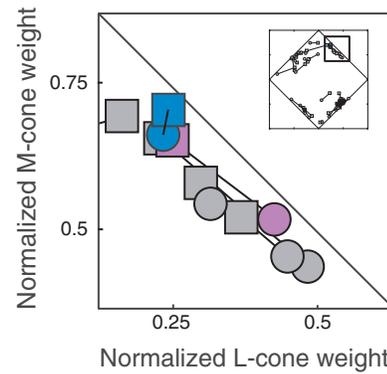
function may also include penalties for parameter values that are unrealistic or are unlikely to generalize well to new data.

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:

$$\begin{bmatrix} A \\ B \\ C \end{bmatrix}, \hat{\theta} = argmax\left(\mathscr{L}\left(f\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}, \theta \right) \middle|\ R \right) \right) \tag{7}$$

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 the observed responses, $\mathscr{L}$ is the likelihood, or the conditional probability of responses R given the candidate weights $[a\ b\ c]$ and parameters values θ. $[A\ B\ C]$ and $\hat{\theta}$ 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 a single maximum, facilitating fitting models with many parameters (Paninski, Pillow, & Lewi, 2007).

J.P. Weller, G.D. Horwitz
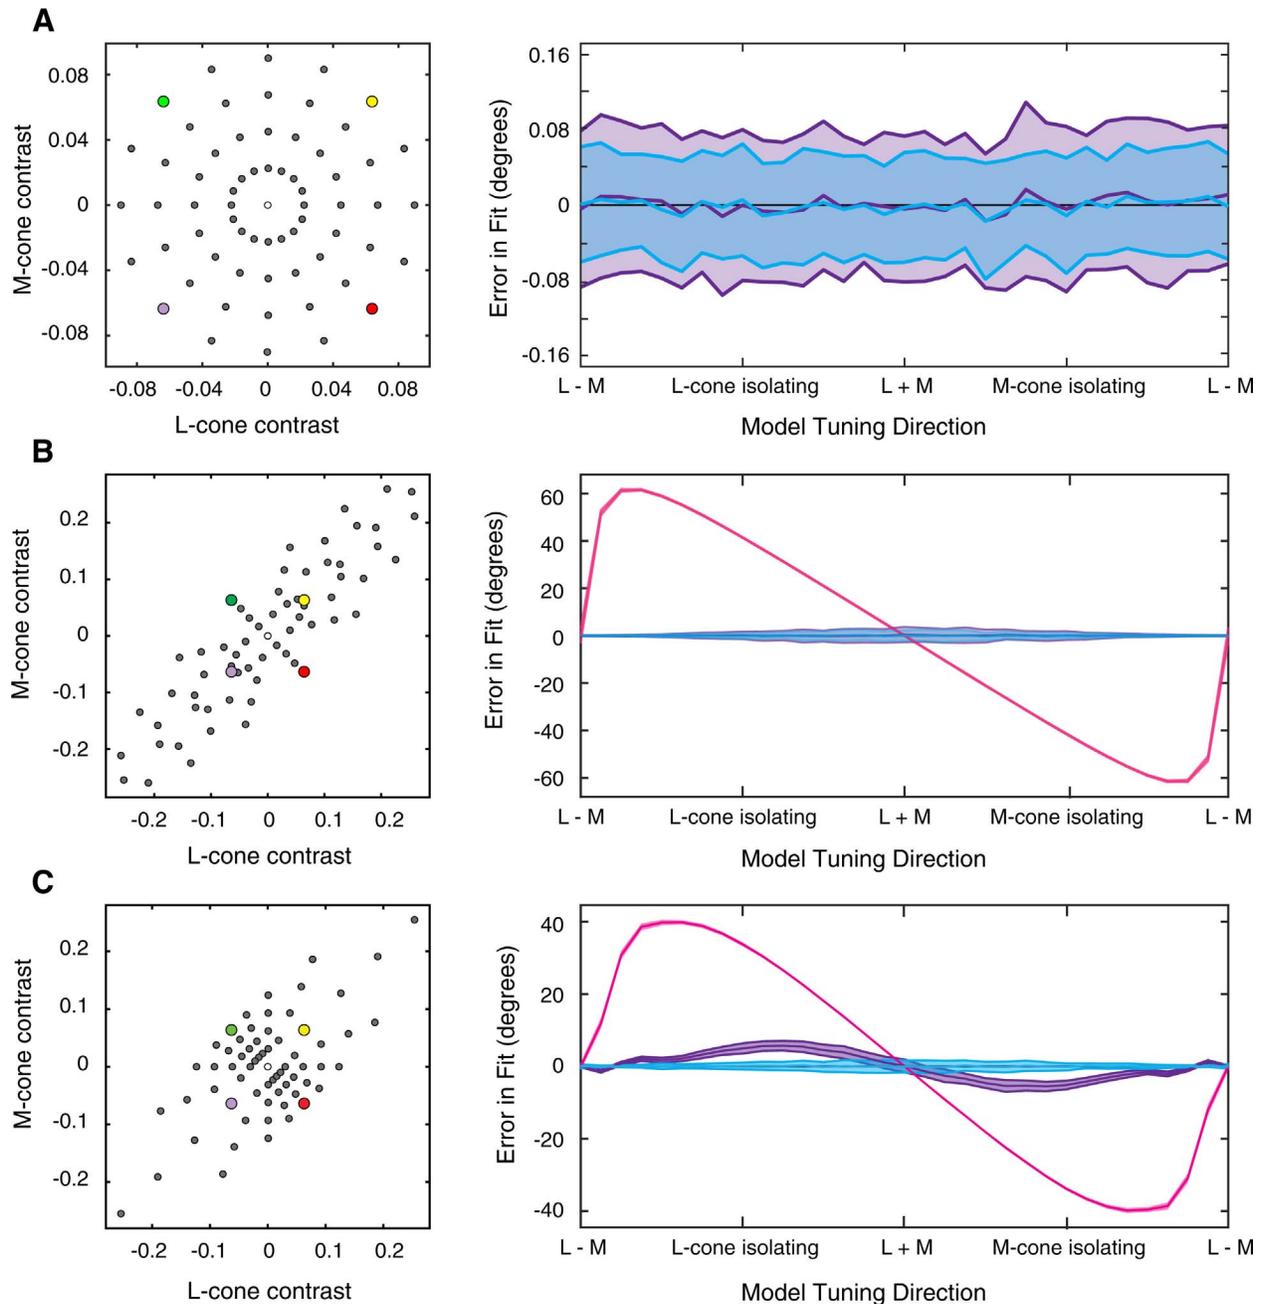
Fig. 6. Estimates of preferred color direction obtained from the responses of simulated neurons. Preferred color directions were estimated with three techniques and three stimulus distributions. Gray circles represent stimuli that were presented to the simulated neuron. The purple, green, yellow, and red points represent identical lights in each distribution. For simplicity, only L- and M-cone modulations were simulated. (a) Left: a radially symmetric stimulus distribution. Right: the mean and standard deviation of the error between the true and estimated preferred directions. Estimates obtained via the response-weighted average stimulus (RWA) (purple) and maximum likelihood (blue) were both unbiased. For this stimulus distribution, RWA and linear regression estimates are the same. (b) Left: a distended stimulus distribution that was RS when whitened. Right: estimation via the RWA of the non-whitened distribution (pink) was biased, but estimation via linear regression (purple) and maximum likelihood (blue) were not. (c) Left: a distended stimulus distribution that could not be made RS by linear transformation. Right: estimates obtained via RWA (pink) and linear regression (purple) were both biased, but maximum likelihood estimates (blue) were not.

4.1. Comparing estimation techniques

To compare the accuracy of preferred direction estimates obtained by numerical maximization of likelihood with those obtained by response-weighted averaging, we performed two analyses. First, we used both techniques on the example dataset from Fig. 1. Second, we used both techniques to estimate the preferred color directions of simulated neurons (see Appendix II for simulation details).

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 (Fig. 5, circles) or from the cone noise data (Fig. 5, squares), confirming that maximum likelihood estimates are relatively insensitive to the distribution of stimuli used in the experiment.

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 (Fig. 6A, left), a distended distribution that could be made RS by whitening (Fig. 6B, left), and a distended distribution that could not be made RS by whitening (Fig. 6C, 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 (Fig. 6, pink curves). Linear regression was additionally unbiased when whitening the distribution produced radial symmetry (Fig. 6, purple curves). Maximum likelihood was unbiased in all three cases (Fig. 6, blue curves).

5. 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.

5.1. Accounting for adaptation with visual signals or 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:

$$R = [LMS]\begin{bmatrix} A \\ B \\ C \end{bmatrix} \tag{8}$$

where L, M, and S represent cone excitation differences, or the changes in photoisomerizations caused by a stimulus against a static background (Boynton, Ikeda, & Stiles, 1964; Cole et al., 1993; Giulianini & Eskew, 1998; Kranda & King-Smith, 1979; Sankeralli & Mullen, 1996; Sperling & Harwerth, 1971), 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:

$$R = [LMS]\begin{bmatrix} \frac{A'}{L_0} \\ \frac{B'}{M_0} \\ \frac{C'}{S_0} \end{bmatrix} \tag{9}$$

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:

$$R = [LMS]\begin{bmatrix} \frac{A}{L_0} \\ \frac{B}{M_0} \\ \frac{C}{S_0} \end{bmatrix} = \begin{bmatrix} \frac{L}{L_0} & \frac{M}{M_0} & \frac{S}{S_0} \end{bmatrix}\begin{bmatrix} A \\ B \\ C \end{bmatrix} \tag{10}$$

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 (Chaparro, Stromeyer, Kronauer, & Eskew, 1994; Kalloniatis & Harwerth, 1991; Stromeyer, Cole, & Kronauer, 1985; Thornton & Pugh, 1983).

Cone contrast is a simple way of modeling cone adaptation (for more complex alternatives, see Stockman & Brainard, 2010 and Angueyra & 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.

5.2. Beyond the LN model

Some neurons are poorly described by the LN model (Bushnell, Harding, Kosai, Bair, & Pasupathy, 2011; Hanazawa, Komatsu, & Murakami, 2000; Komatsu, Ideura, Kaji, & Yamane, 1992; Solomon & Lennie, 2005; Thorell et al., 1984). 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 (Golden, Vilankar, Wu, & Field, 2016; Horwitz & Hass, 2012; Rust, Mante, Simoncelli, & Movshon, 2006).

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. 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, neuronal responses are poorly described as weighted sums of cone signals. One intriguing possibility is that such neurons may perform LN-like operations but on inputs that are abstract quantities (Pagan, Simoncelli, & Rust, 2016). 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.

Appendix I: Electrophysiological methods

We recorded from 35 V1 neurons in an awake fixating rhesus monkey. All experiments were done in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals.

The white noise stimulus comprised a 10×10 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 (Figs. 1A and 3A), 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 (Figs. 1B and 3B). 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:

$$Lcc*Lcp \approx Mcc*Mcp \approx Scc*Scp \tag{1A}$$

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 RWAs were computed: one using the responses from phosphor noise trials and one using the responses from cone 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 a singular value decomposition to find a 3-element vector that describes the consensus color tuning across stimulus pixels and time-delays (Horwitz & Albright, 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 1.

Appendix II: Simulation methods

Modeled responses were generated according to an LN model (Eq. (2)) in which the nonlinear function *f* was a Naka-Rushton function:

$$\lambda = U*\frac{contrast^N}{c_{50}^N + contrast^N} + bl \tag{2A}$$

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 [*A B*], 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:

$$R = Pois(\lambda) \tag{3A}$$

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

$$\hat{\theta} = argmax(\mathscr{L}(R_1...R_n|\lambda_1...\lambda_n)) = argmax\left(\prod_{i=1}^{n} \frac{\lambda_i^{R_i}*e^{-\lambda_i}}{R_i!}\right) \tag{4A}$$

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:

$$\hat{\theta} = argmin(-\mathscr{L}(R_1...R_n|\lambda_1...\lambda_n)) = argmin\left(-\sum_{i=1}^{n} R_i*(\log(\lambda_i)-\lambda_i)\right) \tag{5A}$$

References

Bohon, K. S., Hermann, K. L., Hansen, T., & Conway, B. R. (2016). Representation of perceptual color space in macaque posterior inferior temporal cortex (the V4 complex). *eNeuro, 4*(3).

Boynton, R. M., Ikeda, M., & Stiles, W. S. (1964). Interactions among chromatic mechanisms as inferred from positive and negative increment thresholds. *Vision Research, 4*(1–2), 87–117.

Brainard, D. H. (1996). Cone contrast and opponent modulation color spaces. In Kasier, & Boynton (Eds.). *Human color vision*(2nd ed.). Washington DC: Optical Society of America.

Brainard, D. H., & Stockman, A. (2010). Colorimetry. In M. Bass (Ed.). *OSA handbook of optics* (pp. 10.1–10.56). (3rd ed.). New York: McGraw-Hill.

Bushnell, B. N., Harding, P. J., Kosai, Y., Bair, W., & Pasupathy, A. (2011). Equiluminance cells in visual cortical area V4. *Journal of Neuroscience, 31*(35), 12398–12412.

Chaparro, A., Stromeyer, C. F., III, Kronauer, R. E., & Eskew, R. T., Jr. (1994). Separable red-green and luminance detectors for small flashes. *Vision Research, 34*(6), 751–762.

Chichilnisky, E. J. (2001). A simple white noise analysis of neuronal light responses.

Network, 12, 199–213.

Chichilnisky, E. J., & Baylor, D. A. (1999). Receptive-field microstructure of blue-yellow ganglion cells in primate retina. *Nature Neuroscience, 2*, 889–893.

Cole, G. R., Hine, T., & McIlhagga, W. (1993). Detection mechanisms in L-, M-, and S-cone contrast space. *Journal of the Optical Society of America A, 10*, 38–51.

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

Conway, B. R., Chatterjee, S., Field, G. D., Horwitz, G. D., Johnson, E. N., Koida, K., & Mancuso, K. (2010). Advances in color science: From retina to behavior. *Journal of Neuroscience, 30*, 14955–14963.

Conway, B. R., & Livingstone, M. S. (2006). Spatial and temporal properties of cone signals in alert macaque primary visual cortex. *Journal of Neuroscience, 26*(42), 10826–10846.

Cottaris, N. P., & DeValois, R. L. (1998). Temporal dynamics of chromatic tuning in macaque primary visual cortex. *Nature, 395*, 896–900.

D'Zmura, M. D., & Knoblauch, K. (1998). Spectral bandwidths for the detection color. *Vision Research, 38*, 3117–3128.

D'Zmura, M. D., & Knoblauch, K. (2001). Reply to letter to editor by MJ Sankeralli and KT

Mullen. *Vision Research, 41*, 1683–1684.

Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Neurophsiology, 357*, 241–265.

Eskew, R. T., Jr., McLellan, J. S., & Giulianini, F. (1999). Chromatic detection and discrimination. In K. Gegenfurtner, & L. T. Sharpe (Eds.). *Color vision: From genes to perception* (pp. 345–368). Cambridge: Cambridge University Press.

Gegenfurtner, K. R. (2003). Cortical mechanisms of colour vision. *Nature Reviews Neuroscience, 4*, 563–572.

Giulianini, F., & Eskew, R. T., Jr. (1998). Chromatic masking in the (ΔL/L, ΔM/M) plane of cone-contrast space reveals only two detection mechanisms. *Vision Research, 38*(24), 3913–3926.

Golden, J. W., Vilankar, K. P., Wu, M. C. K., & Field, D. J. (2016). Conjectures regarding the nonlinear geometry of visual neurons. *Vision Research, 120*, 74–92.

Hanazawa, A., Komatsu, H., & Murakami, I. (2000). Neural selectivity for hue and saturation of colour in the primary visual cortex of the monkey. *European Journal of Neuroscience, 12*, 1753–1763.

Hansen, T., & Gegenfurtner, K. R. (2013). Higher order color mechanisms: Evidence from noise-masking experiments in cone contrast space. *Journal of Vision, 13*, 26.

Horwitz, G. D., & Albright, T. D. (2005). Paucity of chromaic linear motion detectors in macaque V1. *Journal of Vision, 5*, 525–533.

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

Horwitz, G. D., & Hass, C. A. (2012). Nonlinear analysis of macaque V1 color tuning reveals cardinal directions for cortical color processing. *Nature Neuroscience, 15*, 913–919.

Hurvich, L. M., & Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review, 64*(Part 1(6)), 384–404.

Ingling, C. R., Jr., & Tsou, B. H. P. (1977). Orthogonal combinations of the three visual channels. *Vision Research, 17*, 1075–1082.

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

Kalloniatis, M., & Harwerth, R. S. (1991). Effects of chromatic adaptation on opponent interactions in monkey increment-threshold sectral-sensitivity functions. *Journal of the Optical Society of America A, 8*(11), 1818–1831.

Knoblauch, K. (1995). Dual bases in dichromatic color space. *Colour Vision Deficiencies XII*, 165–176.

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

Komatsu, H., Ideura, Y., Kaji, S., & Yamane, S. (1992). Color selectivity of neurons in the inferior temporal cortex of the awake macaque monkey. *Journal of Neuroscience, 12*, 408–424.

Kranda, K., & King-Smith, P. E. (1979). Detection of coloured stimuli by independent linear systems. *Vision Research, 19*(7), 733–745.

Krauskopf, J., & Gegenfurtner, K. (1992). Color discrimination and adaptation. *Vision Research, 32*, 2165–2175.

Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research, 22*, 1123–1131.

Lennie, P., Krauskopf, J., & Sclar, G. (1990). Chromatic mechanisms in striate cortex of macaque. *Journal of Neuroscience, 10*(2), 649–669.

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

Pagan, M., Simoncelli, E. P., & Rust, N. C. (2016). Neural quadratic discriminant analysis: Nonlinear decoding with V1-like computation. *Neural Computation, 28*(11), 2291–2319.

Paninski, L., Pillow, J., & Lewi, J. (2007). Statistical models for neural encoding, decoding, and optimal stimulus design. *Progress in Brain Research, 165*, 493–507.

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

Poirson, A. B., & Wandell, B. A. (1993). Appearance of colored patterns: Pattern-color separability. *Journal of the Optical Society of America A, 10*(12), 2458–2470.

Reid, R. C., & Shapley, R. M. (2002). Space and time maps of cone photoreceptor signals in macaque lateral geniculate nucleus. *Journal of Neuroscience, 22*(14), 6158–6175.

Rust, N. C., Mante, V., Simoncelli, E. P., & Movshon, J. A. (2006). How MT cells analyze the motion of visual patterns. *Nature Neuroscience, 9*, 1421–1431.

Sankeralli, M. J., & Mullen, K. T. (1996). Estimation of the L-, M-, and S-cone weights of the postreceptoral detection mechanisms. *Journal of the Optical Society of America A, 13*, 906–915.

Sankeralli, M. J., & Mullen, K. T. (2001). Assumptions concerning orthogonality in threshold-scaled versus cone-contrast colour spaces. *Vision Research, 41*, 53–55.

Solomon, S. G., & Lennie, P. (2005). Chromatic gain controls in visual cortical neurons. *Journal of Neuroscience, 25*, 4779–4792.

Solomon, S. G., & Lennie, P. (2007). The machinery of colour vision. *Nature Reviews Neuroscience, 8*, 276–286.

Solomon, S. G., Peirce, J. W., & Lennie, P. (2004). The impact of suppressive surrounds on chromatic properties of cortical neurons. *Journal of Neuroscience, 24*(1), 148–160.

Sperling, H. G., & Harwerth, R. S. (1971). Red-green cone interactions in the increment-threshold spectral sensitivity of primaries. *Science, 172*, 180–184.

Stockman, A., & Brainard, D. H. (2010). Color vision mechanisms. In M. Bass (Ed.). *In the OSA handbook of optics* (pp. 11.1–11.104). (3rd edition). New York: McGraw-Hill.

Stromeyer, C. F., Cole, G. R., & Kronauer, R. E. (1985). Second-site adaptation in the red-green chromatic pathways. *Vision Resarch, 25*(2), 219–237.

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

Thorell, L. G., De Valois, R. L., & Albrecht, D. G. (1984). Spatial mapping of monkey V1 cells with pure color and luminance stimuli. *Vision Research, 24*, 751–769.

Thornton, J. E., & Pugh, E. N., Jr. (1983). Red/green color opponency at detection threshold. *Science, 219*, 191–193.

Wandell, B. A. (1985). Color measurement and discrimination. *Journal of the Optcial Society of America A, 2*, 62–71.