

A simple principled approach for modeling and understanding uniform color metrics

KEVIN A. G. SMET,¹ MICHAEL A. WEBSTER,^{2,*} AND LORNE A. WHITEHEAD³

¹Light and Lighting Laboratory, KU Leuven, Ghent, Belgium

²Department of Psychology, University of Nevada, Reno, Nevada 89557, USA

³Department of Physics and Astronomy, University of British Columbia, Vancouver, Canada

*Corresponding author: mwebster@unr.edu

Received 5 October 2015; revised 20 January 2016; accepted 1 February 2016; posted 2 February 2016 (Doc. ID 251232); published 1 March 2016

An important goal in characterizing human color vision is to order color percepts in a way that captures their similarities and differences. This has resulted in the continuing evolution of “uniform color spaces,” in which the distances within the space represent the perceptual differences between the stimuli. While these metrics are now very successful in predicting how color percepts are scaled, they do so in largely empirical, ad hoc ways, with limited reference to actual mechanisms of color vision. In this article our aim is to instead begin with general and plausible assumptions about color coding, and then develop a model of color appearance that explicitly incorporates them. We show that many of the features of empirically defined color order systems (those of Munsell, Pantone, NCS, and others) as well as many of the basic phenomena of color perception, emerge naturally from fairly simple principles of color information encoding in the visual system and how it can be optimized for the spectral characteristics of the environment. © 2016 Optical Society of America

OCIS codes: (330.0330) Vision, color, and visual optics; (010.1690) Color; (330.1720) Color vision; (330.5310) Vision - photoreceptors.

<http://dx.doi.org/10.1364/JOSAA.33.00A319>

1. INTRODUCTION

The field of color vision encompasses a diverse range of questions and has spawned a number of subdisciplines. One involves colorimetry—“the branch of color science concerned... with specifying numerically the color of a physically defined visual stimulus” [1] in terms of *color matches* (basic colorimetry) and *color appearance correlates* (advanced colorimetry). The latter has involved developing a number of formal algorithms for specifying the basic structure of color appearance. Such algorithms have obvious practical value for gauging the perceptual consequences of colors in different conditions or applications, e.g., when images are rendered on different devices. However, they have largely been designed only by describing empirical measurements of color discrimination or similarity ratings, and not by asking what causes color appearances to be as they are. That is, while the values are useful from an engineering perspective, they are based on a nested series of multiparameter functions, for which the parameters have been adjusted to make the overall calculation fit known data, but wherein a number of the mathematical subfeatures lack a clear rationale or plausible neural mechanisms.

A second line of color research has focused on understanding the actual mechanisms of color coding. This has provided

deep insights into how information about the spectral characteristics of light is represented and transformed along the visual pathway, and the neural substrate of these mechanisms. This approach has also helped to elucidate computational principles that likely guided the evolutionary development of color vision. However, these approaches have generally not aimed to generate formal quantitative predictions for color metrics.

In the present work, our aim is to bridge the conceptual gap between these two important goals in color science—one focused on understanding the mechanisms and design principles underlying the neural encoding of color information, and the other focused on developing systems for quantifying and predicting the structure and characteristics of color appearance. Specifically, our aim is to illustrate how a quantitative model of color appearance—with predictive power approaching typical uniform color metrics—can be derived from reasonable and general assumptions about color coding, rather than purely empirical data fitting. Our model is thus in contrast to the many color metrics for which predictive performance is often the main goal at the cost of clarity and transparency of possible underlying explanatory physiological, neural, or cognitive mechanisms related to human color perception.

With these thoughts in mind, we begin by discussing the specific behavior we hope to explain—the perceptual organization of surface colors (specifically, colors as perceived within a uniform flat neutral background or context, as opposed to isolated or “aperture” colors). These are usually described by three perceptual attributes: hue (e.g., red versus green), chroma (pure versus diluted), and lightness (light versus dark). The fact that this representation has three key attributes follows plausibly (though not necessarily) from the fact that, as the color normal human eye scans the visible environment, light is sensed by three different types of cone photoreceptors in the retina. It is also widely assumed that these subjective attributes of color arise from combining cone signals by subtraction (opponency) or addition (non-opponency). This two-stage model (of an initial representation based on the three cone types, followed by combining the cones signals within color-opponent mechanisms) explains, in general terms, both the basic color matching characteristics of color vision and also the basic phenomena of color appearance.

However, at a finer level, the characteristics of color appearance remain complex. A wide variety of techniques and studies have been used to describe the relationships involved. Often these approaches arrange surface colors in terms of their perceived similarities and differences. Thus two shades of blue fall closer together than a blue and yellow. Ideally, such systems provide what subjects perceive to be an approximately “uniform color space,” in which the distances between any two points represent approximately equivalent perceptual differences. Here we will focus on the Munsell Color System [2] which, while not the most accurate, is possibly the most well-known perceptual color ordering system, with other examples being the line element models of Helmholtz and Schrödinger [3] and numerous improvements over the years, leading to what many feel is the most accurate prediction tool today, the CIECAM02 system [4,5] of the International Commission on Illumination (CIE).

The Munsell Color System [6] was developed through careful studies of classification of colored surfaces carried out by Albert H. Munsell in the early part of the 20th century, based on the reported observations of numerous human subjects. This system organizes color samples into a three-dimensional configuration, as depicted [7] in Fig. 1.

The system is organized into levels of equal perceived lightness (the vertical, axial direction), quantified by a number called

value (V) notionally ranging from $V = 0$ (black) to $V = 10$ (white). It is also organized into planar groupings of equal hue, in which radial distance from the center represents increasing intensity of color, represented by a number called chroma (C), which ranges from zero (gray) to a maximum value that depends on the hue (H) and value. The planes of equal hue are organized in equal perceived spacing in the circumferential direction, and can be defined by a labeling scheme or simply the angle, to identify a specific hue. As with many other color appearance systems, the Munsell chart was organized to have the distance between two samples correspond on average to subjects’ perception of the “size” of the appearance difference between them. Finally, note that the appearance of the samples depends critically on the background, and is normally relative to a neutral gray, which is the context we assume here.

People with normal color vision find such classification systems intuitively clear and meaningful and therefore may not realize they present a fascinating and important characteristic of human color vision—the fact that the arrangement appears sensible and appropriate for essentially everyone with normal color vision. But, what actually *are* hue, value, and chroma? More generally, what is the connection between the H , V , C coordinates of the Munsell system and the photoreceptor responses that can be calculated from the spectral power distributions (SPDs) of the reflected light that reaches the eye from each of the various surfaces in a scene (recognizing that these reflected light SPDs are the products of the incident illumination SPD and the spectral reflectance functions of the various surfaces)? One might have expected the answer to this question to be simple, but it is not! Actually, understanding this connection has been a goal of modern color metrics for some time, one that has previously been achieved through empirical approaches [4,8,9] that are significantly more complex than the method described here, and are much less clearly connected to the underlying physiology of color vision.

The challenge of this problem of connecting the H , V , C values to cone stimulations [10] is partially illustrated in Fig. 2. The figure’s rightmost panel shows a set of stimuli within the Munsell system, chosen to have equal chroma and lightness and spanning the circle of hues in equal intervals. As mentioned previously, this is a system in which the color differences portrayed are perceived by observers to be approximately perceptually uniform. However, the figure’s middle panel shows these stimuli plotted instead in terms of their corresponding cone photon absorption rates, using a variant of a standard diagram that plots the relative excitations of the cones to different spectra of the same luminance [11]. In terms of the cone absorption rates, the previously uniform set of stimuli is now strongly distorted—stretched along some axes and compressed along others. Moreover, the distortion pattern itself depends, in a complicated way, upon the SPD of the background. The challenge, then, is to identify the transformations that convert the cone photon absorption rates into perceptually meaningful (uniform) differences. As we have noted, in developing uniform color spaces, these connections have been predominantly empirical and rather complex. Thus, while they successfully describe color appearance, they provide little insight into the underlying processes or evolutionary factors driving the visual




Fig. 1. Munsell Color System.

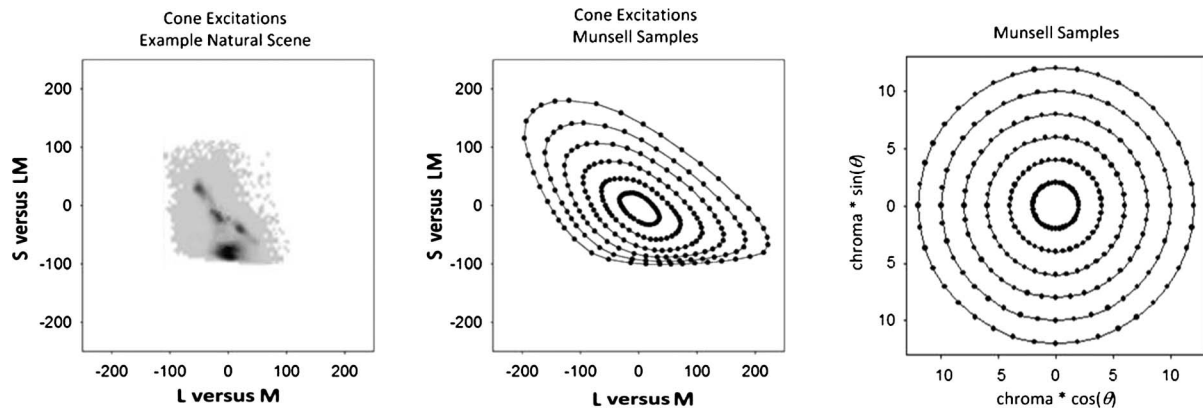


Fig. 2. Relative cone excitations for a set of natural outdoor scenes (left) and a set of uniform Munsell samples (middle), as shown by their Munsell coordinates (right).

representation of color, because there is little correspondence between the mathematical operations and the actual physiology, nor is there any explanation for how those calculations could be accurately emulated in the nervous system.

As a specific example, consider one of the most sophisticated instantiations of the family of CIE uniform spaces: CIECAM02 [12], which estimates the color differences with a basic set of 14 equation groups involving 39 constant numerical parameters. As with most other models, these include non-linear calculations resulting from the need to “compress” responses or “adapt” them according to the reference background, but without considering the actual basis or rationale for these adjustments. Moreover, many of the calculations in CIECAM02 involve complex polynomial expressions without any explanation for how such relationships could be emulated physiologically. Further, the many constant numerical parameters were selected to produce a reasonable match to observational data, but without any explanation for how such “tuning” might actually occur within the vision system. Thus, while CIECAM02 and its predecessors are very useful engineering tools, they provide little insight into color appearance (in the same way that the CIE 1931 chromaticity diagram provides an industry standard for color specification yet completely obscures the underlying basis of color matching and trichromacy).

Here, we show that a more explanatory description of color appearance is afforded by starting with some simple assumptions about sensory coding. Again, this is an area that also has an extensive, though more recent, history. A general principle emerging from this approach is that it would have been evolutionarily advantageous for sensory systems to evolve so as to optimize the efficiency of representation of the characteristics of an organism’s environment. Much of this work has drawn from information theory and arguments for coding efficiency, to identify how the limited dynamic range of neural responses should be allocated to carry the most information about the stimulus [13–16]. This includes adjusting individual neuron response functions based upon the frequency distribution of past stimulation values in order to equalize the information content associated with each response level [17,18], and adjusting the responses across neurons so that they are similar but independent (which maximizes the amount of useful information

carried by each neuron) [19,20]. Resource limitations may be an especially critical factor at early stages of the visual system, where the raw image must be compressed into a small number of channels, each restricted to a limited range of levels. As one example, the optic nerve consists of between 770,000 and 1,700,000 nerve fibers [21], transmitting somewhat less than 10,000,000 bits of information per second [22], but this must be used to carry the information arising from over 100,000,000 photoreceptors [23]. Clearly, then, this encoding must be optimized to preserve information about the image [24].

Coding efficiency also requires that the system be optimized for the specific characteristics of the stimulus environment, and be adaptable to changes in the environment [25]. The left panel of Fig. 2 shows the distribution of cone excitations from a sample of a natural outdoor setting [10]. Note that this environmental distribution is somewhat similar to the cone excitations of the Munsell samples, which appear perceptually uniform. This suggests that part of the mapping of the cones to color sensations is very much a mapping from the world, and indeed one could directly try to predict perceptually uniform color spaces from the non-uniform color statistics of natural images [10,26].

However, our goal is to illustrate how this mapping could be realized by feasible steps in color coding. For this reason, we focus on plausible simple sequential computations for implementing the transformation. We attempt to explain how each step can feasibly adjust its performance parameters in response to the statistical distribution of recent signals from the preceding stage. Admittedly, many of the neural processes underlying color perception remain poorly understood, and many of the mathematical transformations we present do not have uniquely defined neural correlates or might be instantiated in very different ways. Thus our model is not intended to imply the actual physiological basis for color metrics. Instead, the primary focus of the paper is to demonstrate what steps we might expect from fairly simple and general principles, and how close those steps come to emulating the content of empirically defined color metrics.

2. KEY COMPONENTS OF THE MODEL

With this background, we will begin with a brief overview of a model consisting of five sequential functional stages. Each

stage receives input information and produces a transformation that is input to the next stage. The input for stage one is light entering the eye. The output for stage five is the perception of surface color. Within each stage there is a simple, mathematically describable, transform function that contains one or more numerical parameters that “tune” its performance, and—where possible—we point to the potential mechanisms or principles underlying why this transformation may have evolved and how the parameters may be automatically adjusted over time. One type of “plausible tuning algorithm” would be to slowly adjust a parameter over time to achieve and maintain a uniform statistical distribution of recent output values.

Figure 3 depicts these stages as a flow chart. (We emphasize that these stages are not exhaustive and, for example, do not include many known contextual effects or how color interacts with other attributes such as form or motion. However they are intended to characterize the properties of color appearance as embodied by most color metrics.) Importantly, the figure shows two corresponding tracks. The left track describes putative processing stages and potential neural sites. The right track describes the mathematical models for these stages presented in this paper. Each stage is a function, with operating parameters, and a method for optimizing those parameters. Before exploring those mathematical relationships in detail, we will first briefly summarize the transformations at each of the five stages.

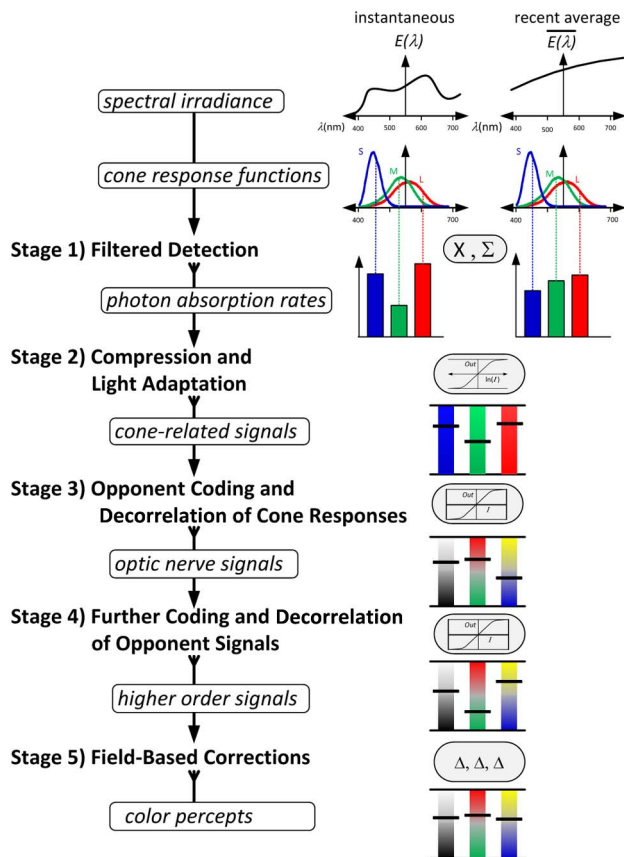


Fig. 3. Stages of information processing in the model.

- Stage 1. Filtered Detection

- Modeled process.

The input to Stage 1 is the time history of the SPD of the light incident at a given small region of the retina. This is sampled by the trichromatic cone mosaic, with output corresponding to values representing the cone photon absorption rates in the three classes of cones. For modeling surface colors we consider two of these outputs. The first is the instantaneous rate corresponding to the stimulus, and the second is a recent time average rate, representing the presumed response to the background due to sampling different regions with saccades.

- Mathematical representation.

The corresponding input is the SPD of light reflecting to the eye from a test object and also the average SPD of the scene. The filtering and detection of the incident light is simulated by multiplying each SPD with the corneal spectral sensitivities (cone fundamentals) of each of the three cone types and integrating them to determine the test and background “cone photon absorption rates.”

- Stage 2. Compression and Light Adaptation

- Modeled process.

The model-transformed cone signals are assumed to follow a compressive nonlinearity such that responses have greater sensitivity around the mean cone photon absorption rate and smoothly asymptote at very low or high levels. The response is adapted for the recent mean by a gain control that adjusts to the recent history of stimulation.

- Mathematical representation.

The proposed mathematical form for the response is the application of a sigmoidal transform function to a logarithmic response function. This requires two parameters: (1) a gain parameter that must be adjusted as a compromise between too low a gain, which reduces the signal-to-noise ratio, and too high a gain, which increases information loss through output saturation; and (2) an offset parameter that reduces the difference between the degree of output clipping at the low end and at the high end, thus optimizing information transfer.

- Stage 3. Opponent Coding and Decorrelation of Cone Responses

- Modeled process.

Following conventional color models, the model transformed cone signals are combined within three post-receptoral mechanisms. These help to remove the redundancies in the information carried by the different cones by recoding into channels that respond to the sum or differences of the model transformed cone signals. The summing channel is “nonopponent” and forms a mechanism sensitive to achromatic variations. The other two are “opponent” and take the difference of the model-transformed cone signals to represent information about the chromatic properties of the stimulus. Within each, the signals are gain controlled to match the range of inputs so that, over time, the statistical

distribution of output values is similar across the three mechanisms.

- Mathematical representation.

In the model shown here, there are three numerical coefficients that affect the calculation of each of the two color difference signals, and they are assumed to be tuned to the stimulus history. The proposed method is to slowly adjust those parameters to ensure that, over long times, (1) achromatic stimuli produce neutral color signals, (2) the average of the output values is also neutral, and (3) the r.m.s. output signal, relative to the neutral point, has a nominally specified value (e.g., 1), for optimal information transfer.

- Stage 4. Decorrelation of Opponent Signals

- Modeled process.

The previous stage decodes the model-transformed cone signals into three mechanisms with sensitivities corresponding to the “cardinal” luminance and chromatic dimensions that have been postulated for early post-receptor color coding. Stage 4 involves a subsequent rotation and scaling of the model chromatic signals, and is included because the responses of the cardinal mechanisms are often correlated for natural color distributions. Specifically, this stage is included to adjust to the “blue–yellow” bias in natural scenes, as illustrated in Fig. 1. We note that Stages 3 and 4 could be implemented with a single transformation. However, there is evidence to suggest that these adjustments occur at different stages in the visual system [27]. Moreover, if they were combined into a single transformation, there would be no simple mechanism for adaptively tuning those parameters over time to achieve and maintain accurate color perception.

- Mathematical representation.

The equations at Stage 4 apply a rotation and rescaling of the mechanism sensitivities. As in Stage 3, the control parameters are determined in a simple way to optimize the statistical distribution of the output values.

- Stage 5. Field-Based Corrections

- Modeled process.

To successfully predict appearance, color appearance models must include additional “chromatic adaptation” factors that normalize the model values to allow gray and white surfaces to be correctly perceived as uncolored under a range of differently colored illumination conditions. A wide variety of mechanisms contribute to this “color-constancy” correction.

- Mathematical representation.

The mathematical treatment of this stage is simply the global application of an additive correction to the two color signals, to ensure that a neutral gray surface appears approximately colorless. Additionally, to express the resultant color appearance information in terms of the scaling used within the Munsell system, three arbitrary scaling parameters are applied to express the model output in terms of the Munsell ranges of value and chroma, and with the correct zero point for hue.

In the following, we elaborate and provide quantitative details for each of these stages.

A. Stage 1: Filtered Detection

Stage 1 involves the absorption of visible light photons by the cone opsin molecules in the light-sensitive cone cells in the retina. The intensity of light incident on the eye is generally a function of the viewing direction, \hat{r} , and the wavelength λ . It is called the spectral radiance distribution (and more commonly the spectral power distribution or SPD), and is expressed as $L_{e\lambda}(\lambda, \hat{r})$, in units of $\text{W}/\text{sr} \cdot \text{m}^3$. (Although these are the correct units, it is more common outside of physics literature to use the units $\text{W}/\text{nm} \cdot \text{sr} \cdot \text{m}^2$.)

There are three different kinds of cone cells in the retina of the normal human trichromat, which differ in the photopigments they contain and, thus, in their relative spectral sensitivity as a function of wavelength. Because of absorptions in the lens, ocular media, macular pigment, and the retina itself, the pigment absorption spectra are typically less relevant than the corneal spectral sensitivities derived from psychophysical color matching functions. These functions, often referred to as cone fundamentals, are typically labeled $L(\lambda)$, $M(\lambda)$, and $S(\lambda)$ (based on the words long, medium, and short, which describe the relative wavelength of peak absorption). The model described in this paper uses the CIE 10° 2006 cone fundamentals [28]. The equal-area-normalized cone fundamentals are depicted in Fig. 4. For the purposes of this paper, the normalization is arbitrary because it does not affect the final calculation.

Using these three relative response functions, we then define three-cone photon absorption rates as follows:

$$L(\hat{r}) = (683 \text{ lm}/\text{W}) \int L(\lambda)L_{e\lambda}(\lambda, \hat{r})d\lambda / \int L(\lambda)d\lambda, \quad (1a)$$

$$M(\hat{r}) = (683 \text{ lm}/\text{W}) \int M(\lambda)L_{e\lambda}(\lambda, \hat{r})d\lambda / \int M(\lambda)d\lambda, \quad (1b)$$

$$S(\hat{r}) = (683 \text{ lm}/\text{W}) \int S(\lambda)L_{e\lambda}(\lambda, \hat{r})d\lambda / \int S(\lambda)d\lambda. \quad (1c)$$

(Note that in each of these equations, the cone fundamental spectral sensitivity functions appear in both the numerator and the denominator, so the normalization factors for them




Fig. 4. Spectral sensitivity functions of the cones.

cancel.) The units of $L(\hat{r})$, $M(\hat{r})$, and $S(\hat{r})$ are cd/m^2 . According to this measure, an SPD with uniform spectral radiance distribution and a photometric luminance L_v of $1 \text{ cd}/\text{m}^2$ will have that same numerical value, 1, for each of $L(\hat{r})$, $M(\hat{r})$, and $S(\hat{r})$. For a general, non-uniform spectral radiance distribution $L_{e,\lambda}(\lambda)$, the values $L(\hat{r})$, $M(\hat{r})$, and $S(\hat{r})$ will usually differ from one another, providing the basic information available for color and also fundamentally limiting that information. Specifically, the output of each cone is a function of its rate of photon absorption—no information is provided about the wavelength distribution of the absorbed photons, a fundamental constraint known as the “principle of univariance” [29]. Thus color information can be obtained only by comparing the outputs of the three cone types.

B. Stage 2: Compression and Light Adaptation

While photon absorption is an approximately linear process (at least up to an intensity that causes photopigment bleaching to become significant), it is important to realize that, as a whole, the visual system is highly nonlinear. For example, it is well known that in vision, as well as most areas of perception, the ratio of intensities of two stimuli is a better predictor of the perceived difference between them than is their intensity difference [30]. Nonlinearities of this kind are important given the large intensity range over which color vision must operate. People can see color over roughly a million-fold range of luminance values ranging from that of moonlit snow to sunlight shimmering on a lake. Over this entire range, it is desirable to be sensitive to small relative changes of intensity. This has the advantage that a given object can appear much the same when viewed under very different lighting conditions (because when the light intensity changes, it is the ratios of reflected light signals that remain constant).

A further important role of nonlinearity in neural processing relates to the fact that neurons have a quite limited ratio of dynamic range to noise. Typically, a graph of a neuron’s output as a function of stimulus level has a “sigmoidal” shape, as sketched roughly in Fig. 5. There may be a central region that is fairly linear, but the response then approaches asymptotes at both input value extremes.

A recognized rationale for this sigmoidal shape is that it provides maximum useful information when the frequency distribution of input values has the form of the ubiquitous normal (or Gaussian) distribution [31]. In this case, which we will use here for simplicity, the ideal mathematical form shown in Fig. 5




Fig. 5. Schematic model of how a neuron output signal depends on the net effect of its input signals.

is known as the error function, $\text{erf}(x)$, which is twice the cumulative probability of a Gaussian distribution, where x is measured from its center in units of standard deviation. Its lower limit, as x approaches $-\infty$, is -1 ; its upper limit, as x approaches ∞ , is 1 , and for small values of x , the function is approximately $(2/\sqrt{\pi})x$. Many consider the error function to be a very simple, meaningful, and useful sigmoidal function. (However, as with many other stages of the model we note that the assumption of a Gaussian distribution, while common, is an idealization that may not be perfectly realized in real systems.)

A further issue is that the world itself varies, and thus neural responses must adapt to adjust to these changes. For example, the narrow dynamic range of the system cannot represent the enormous range of light levels we encounter during the course of a day, and like a camera, must instead be continuously adjusted or centered on the much more limited range of contrasts that occur at a given moment [32]. Similarly for color vision, there is an ongoing need to adjust the operating parameters at each stage, guided by the frequency distribution of fairly recent previous signals, to optimize the transfer of information. This tuning adjusts not only to the mean stimulus but also to its variance or contrast, and could potentially also adjust to other moments of the stimulus distributions and thus adjust the shape of the transfer functions. Moreover, these adjustments can occur over many time scales, with some integrating over long intervals (of days, months, or years), while, in other cases, the time constants may range from seconds to milliseconds [33–37].

For the purpose of including these adaptation effects, we must now consider the initial encoding of color in more detail. To do so, we need to define certain key components that determine the reference frame or context for judging a color. Recall that Eqs. (1a)–(1c) defined the color information for a given SPD as $L(\hat{r})$, $M(\hat{r})$, and $S(\hat{r})$, with the dependence on \hat{r} indicating that, in a typical environment, these vary with the viewing direction. We must now be more specific about viewing directions. For the rest of this paper we will use three different triads of color information, as follows:

- L_t , M_t , and S_t will denote the $L(\hat{r})$, $M(\hat{r})$, and $S(\hat{r})$ values for a test surface in the direction \hat{r}_t ;
- L_f , M_f , and S_f will denote the $L(\hat{r})$, $M(\hat{r})$, and $S(\hat{r})$ values averaged over the current visual field; and
- L_o , M_o , and S_o will denote the $L(\hat{r})$, $M(\hat{r})$, and $S(\hat{r})$ values of a very long-term average over many scenes/fields.

Overall, the situation can be summarized approximately as follows: the apparent color of a “test” surface depends on the L_t , M_t , and S_t values for the light reflecting from it, and also depends on L_f , M_f , and S_f , and likely also on L_o , M_o , and S_o . Thus, we seek to understand the process by which these nine values produce the observed color experience denoted by the three values, H , V , and C , of the Munsell system.

Next, we model the nonlinearity in the cone response. The model response is depicted in Fig. 6. In this plot, the vertical axis has arbitrary scaling. In many measurements of cone responses to light, this form of response is observed, although it is usually described by a different function, known as the Naka–Rushton [29] equation, of the form $2I^n/(1 + I^n) - 1$. The equation we use here is essentially equivalent numerically,




Fig. 6. Model values (arbitrary scaling) representing a cone output versus logarithm of intensity, based on a truncated logarithm function of the form $\text{Output} = \text{erf}(0.26 \ln(I/(100 \text{ cd/m}^2 B)^{-.5}))$, where I is the intensity and B is the logarithmically averaged recent intensity, both in cd/m^2 .

but the form we use was selected because it is conceptually simpler, and because it represents the sequential application of two simple functional relationships that have a clear computational utility as described in this paper. In contrast, there is no known physiological mechanism for emulating the fractional powers within the Naka–Rushton equation. The response function shown in Fig. 6 is approximately log-linear in a central portion of its range, which corresponds to about 2 orders of magnitude of luminance. The range is centered by adaptation according to the average recently experienced (i.e., “ambient or background”) light level. For example, the dashed curve in Fig. 6 shows the response when the ambient luminance level is dim (1 cd/m^2), and, conversely, the dotted curve shows the opposite situation in response to much brighter outdoor daylight ($10,000 \text{ cd/m}^2$).

To quantitatively model the response, we’ll need to introduce two constants:

- C_c , a cone contrast factor, which determines the steepness of the slope of the cone response function; and
- C_f , a cone feedback factor, which is the coefficient for the automatic gain control based on a cone’s time-averaged signal.

For the three cones, we will denote the resultant model transformed cone signals as

$$L' = \text{erf}\left(C_c \left(\ln\left(\frac{L_t}{L_o}\right) + C_f \ln\left(\frac{L_f}{L_o}\right)\right)\right), \quad (2a)$$

$$M' = \text{erf}\left(C_c \left(\ln\left(\frac{M_t}{M_o}\right) + C_f \ln\left(\frac{M_f}{M_o}\right)\right)\right), \quad (2b)$$

$$S' = \text{erf}\left(C_c \left(\ln\left(\frac{S_t}{S_o}\right) + C_f \ln\left(\frac{S_f}{S_o}\right)\right)\right). \quad (2c)$$

In this model, the constants C_c , C_f , and L_o , M_o , S_o could potentially be determined by evolution and could be fixed for any individual. On the other hand, it is possible that L_o , M_o , and S_o could be determined by the very long term exposure of individuals to light. As depicted by the values shown in Fig. 11,

we have set the L_o , M_o , and S_o to be the values that would be obtained for a gray 20% reflective surface illuminated by mid-morning or mid-afternoon sunny daylight, with a luminance of 5000 cd/m^2 and a color temperature of 5500 K. The results depend only weakly on this choice and were selected as a reasonable proxy for an average viewing condition.

It should be noted that Eqs. (2a)–(2c) contain division operations, yet there is no known simple way for neurons to directly carry out a division operation, and it would seem even more improbable within a single cone cell. However, it does appear that cone cells have the ability to emulate the calculation of Eqs. (2a)–(2c). Potentially, this could involve a combination of an intrinsic logarithmic response with various time-averaged feedback mechanisms [38]. (Some might find the idea of physically emulating the mathematical logarithm function daunting. However, chemical systems can do this. For example, in the common pH meter, a simple cell produces an electrical potential difference directly related to pH, which is proportional to the logarithm of the concentration of H^+ ions. In general, exponential relationships are commonplace in chemical dynamics, and so, therefore, are their inverse, logarithmic relationships.)

If the frequency distribution of the logarithm of intensity is the common Gaussian distribution, then applying the error function to the logarithm of intensity, as depicted in Fig. 6, will convert to a uniform frequency distribution of output values. For this reason, we have used the error function applied to the logarithm as the model for the sigmoidal cone response function. For subsequent processing signals, for which a logarithmic transform is not relevant, we could simply use the error function to represent a sigmoidal input/output response. In reality, at least with surface samples and the light intensities studied in this paper, there is little or no computational benefit from using the sigmoidal form in the remaining stages.

C. Stage 3: Opponent Coding and Decorrelation of Cone Responses

To summarize, L' , M' , and S' are the model-transformed cone signals in response to the current color sample and its background. This recoding is a tremendous improvement over the raw L , M , and S excitations for the light itself, thanks to the automatic gain control and logarithmic compression. Nevertheless, these signals remain poorly suited for efficiently representing color.

The difficulty is shown in Fig. 7, where pairs of the values for L' , M' , and S' are plotted for 99 colored surfaces that were selected from a large data base of real samples, by means of a Monte Carlo procedure guided by criteria assuring a set of samples that are very uniformly distributed from several perspectives [39]. As can be seen, there is considerable correlation among these three signals, especially between the L' and M' signals [40]. This arises from the fact that there is considerable overlap of the spectral sensitivities of the cones. It is well known that it is computationally inefficient to transport highly correlated signals [16]—and thus it would be evolutionarily advantageous to develop a simple arrangement for combining cone responses in a way that reduces this problem, providing the conversion itself is simple. This is thought to be a major




Fig. 7. These graphs plots of pairs of L' , M' , and S' for 100 randomly selected points in color space, showing they are highly correlated.

rationale for converting the model-transformed cone signals into a color-opponent representation [16].

Thus, in Stage 3, the L' , M' , and S' independent model-transformed cone signals are combined to produce three different output signals [17]. In the actual retina, this is initiated within the various neural connections from the receptors to horizontal and bipolar cells in the retina, where emulation of the linear combination processing is quite feasible because neurons can have both excitatory and inhibitory inputs. A wide variety of distinct cell types with different spectral sensitivities have been identified in the primate retina, and how these might contribute to various aspects of color perception remains uncertain [41,42], but the general transformation of the independent cone signals into non-opponent or opponent mechanisms is well established.

For our model, the first step of Stage 3 is to recode the model-transformed cone signals into channels that separately carry the achromatic and chromatic content of the stimulus as follows:

$$\lambda^* = 0.5L' + 0.5M', \quad (3a)$$

$$\alpha = L' - M', \quad (3b)$$

$$\beta = 0.5L' + 0.5M' - S'. \quad (3c)$$

Note that partial decorrelation occurs because the channels now respond to the sums or differences of the model-transformed cone signals (effectively realigning the axes of the mechanisms in Fig. 7 so that they lie along the positive or negative diagonals). The first value, λ^* , is a signal related to the lightness of the stimulus. The value α corresponds to a reddish versus blue–green opponent color axis (red is positive, blue–green negative) and is based on comparing the L' and M' model-transformed cone signals. The value β relates to a yellow–green versus purple opponent color axis and is based on comparing L' and M' model-transformed cone signals to S' model-transformed cone signals. (Here these terms are simply a notation for labeling the content described mathematically in the relevant equations. They do not represent precisely defined color dimensions at this stage. Furthermore, the sign for α and for β is arbitrary, and here we have simply used the most common convention. The visual system in fact includes cell

types encoding both polarities, e.g., in “on” and “off” cells for luminance. The polarity of the β channel is such that “yellow–green” is positive to be consistent with the polarity of many color spaces. However, we should note that this is opposite to the polarity depicted in the diagram shown in Fig. 1, and is also at odds with the preponderance of S-cone “on” cells found in the retina [42].)

The specific combinations shown may be considered, by some, to be simple and appropriate for another reason: most mammals have only two cone types—an S and a single longer wave cone type. Among mammals three-cone vision is largely exclusive to old world primates, and arose from an evolutionarily recent (40 Mya) duplication of the long-wave pigment gene. For this reason, the two opponent dimensions have been described as the ancient (S versus LM) and modern (L versus M) subsystems of human color vision [43].

The second step of Stage 3 again involves the overall shape and scaling of each mechanism’s response. The opponent calculation for both α and β requires considerably more gain than the λ^* channel, since, as noted, the difference signals they measure are much smaller because of the overlapping $L(\lambda)$, $M(\lambda)$, and (to a lesser extent) $S(\lambda)$ sensitivity functions. Further, we again assume that it may be desirable for the signal transformation at this point to have its output functional form adjusted parametrically to yield a statistically uniform distribution of output values. We have not found a sigmoidal transfer function to be essential at this point, but we have found some advantage for fitting the Munsell system in allowing a different gain for positive outputs than for negative outputs, as depicted in Eqs. (3d) and (3e), although the potential mechanisms for achieving this affect and adjusting its performance are uncertain:

$$\text{if } \alpha \geq 0, \quad \alpha' = c_{g\alpha 1p} \alpha, \quad \text{if } \alpha < 0, \quad \alpha' = c_{g\alpha 1n} \alpha, \quad (3d)$$

$$\text{if } \beta > 0 \quad \beta' = c_{g\beta 1p} \beta, \quad \text{if } \beta < 0, \quad \beta' = c_{g\beta 1n} \beta. \quad (3e)$$

In this amplification stage, the expectation is that the four constants represent parameters that would adapt over time in response to stimuli [10], so that on average the mean value for both α' and β' would be 0 and their standard deviation would be an arbitrary constant of 1. In our model, we carry out that calculation by evaluating α' and β' for the previously




Fig. 8. Decorrelation of the modeled cone signals by recoding into opponent and non-opponent responses.

mentioned set of 99 uniformly distributed reflectance spectra and adjusting the constants in Eqs. (3d) and (3e) as just described.

Figure 8 shows the same stimulus set as that for Fig. 7, but now replotted in terms of the partially decorrelated and scaled opponent mechanisms β' versus α' , β' versus λ^* , and α' versus λ^* .

As noted above, for ordinary surfaces at typical light levels, it does not appear necessary to incorporate a sigmoidal character in the output of Eqs. (3d) and (3e). However, for highly saturated brightly illuminated stimuli, it is possible that some color distortion might arise from saturation. For intense narrowband LEDs and lasers, the distortion of hues could become apparent. Hue distortions of this type are observed in what is called the Bezold–Brücke effect [44], and it is possible that a nonlinearity could be feasibly incorporated to help predict this effect.

D. Stage 4. Decorrelation of Opponent Signals

As previously mentioned, the initial opponent stage significantly reduces the correlation between the model-transformed cone signals, but the opponent signals themselves may still have remaining correlations depending on the choice of axes. In fact, the left-hand plot in Fig. 8 reveals a positive correlation between the α' and β' responses, because the signals vary more along the positive diagonal (a yellowish–bluish axis that is also prevalent in many natural scenes but reflects covarying signals in the L versus M and S versus LM cardinal mechanisms [45]). This residual correlation could be removed by an additional subtractive stage, which applies a differential gain in the $\pm 45^\circ$ directions, as depicted in Eqs. (4a) and (4b), and as anticipated in previous observational studies [46,47]:

$$\alpha'' = c_{g\alpha 2}(\alpha' + \beta'), \tag{4a}$$

$$\beta'' = c_{g\beta 2}(\alpha' - \beta'). \tag{4b}$$

Unlike Stage 3, we found little value in a different gain for positive input values versus negative ones, and so there are just two constants that are to be determined, as in Stage 3, by adjustment to achieve a standard deviation of 1 for the distribution of α'' and β'' values found for the previously mentioned set of 99 uniformly distributed reflectance samples. It should also be noted that this is another case wherein the underlying

neural mechanisms may take a very different form. There is considerable evidence for—and also some against—the idea that at cortical levels the cardinal mechanisms are transformed into multiple “high-order” mechanisms tuned to different directions in color space [27,48,49]. The advantage of these multiple mechanisms is not immediately obvious from the perspective of efficient coding (since they are redundant with the information already carried by the three cardinal axes). However, this “multiple channels” code could allow the visual system to use a population code for hue, in the same way that it appears to represent other dimensions such as spatial orientation. In this case, the implementation would not involve decorrelation but rather adjustments of the relative gains within each mechanism. In either case, the critical point is that the mechanisms are adjusted so that the distribution of responses is made approximately spherical [49].

As an aside, it is important to ask whether the additional parameters in Stage 4 are actually adding value by significantly reducing the r.m.s. difference between the model-predicted H, V, C values and the actual Munsell values. This is an important consideration because, generally, when striving to reduce r.m.s. error between experimental data and a model fit, merely adding extra parameters, even if they have no meaning, will often slightly improve the fit if the parameters are carefully selected by means of an optimization algorithm. In this case, there are two reasons to believe these added parameters do have meaning. First, the added parameters were *not* adjusted by an optimization process to reduce the fitting error, so the fact that the fit was significantly improved strongly suggests the added parameters are an improvement to the model. Second, there are standard statistical tests to determine the fit improvement that would be expected purely from adding free, but meaningless, parameters to a model. To check this, we used the well-known Akaike [50] approach to statistical model selection, which provides what is called the Akaike Information Criterion (AIC), which depends on the number of data points and the fit improvement obtained by adding new parameters. We found these additional parameters significantly increase the AIC score, indicating they likely provide improved understanding. That is, the small increase in the overall number of parameters (effectively from five to eight), reduces the residual square error of the fit to less than 50% of its original value based on the five parameters. This improvement is far greater than one would expect from added parameters that did not significantly refine the model.




Fig. 9. Improved output distribution after secondary encoding of the color signals (Stage 4).

As evidence of the beneficial overall effect of this higher-order decorrelation, the same stimulus set as in Fig. 8 is again replotted as pairs of values β'' versus α'' , β'' versus λ^* , and α'' versus λ^* in Fig. 9, which shows that the distribution is fairly uniform.

E. Stage 5: Field-Based Corrections

There are necessarily many additional subtle and elaborate processes that adjust color percepts according to sophisticated inferences about the stimulus, but they are likely not required for the primary goal of this paper—to compare the output of the model with the organization of color experience implied by the Munsell system.

To make this comparison, we first need to scale the output values to match the arbitrary scaling range of the Munsell color system. This consists of an arbitrary rotation of angle θ , and arbitrary linear scaling to give values the correct range for the Munsell system equivalents of λ^* , α'' and β'' , here termed λ_{int}^* , α_{int} , and β_{int} (with the subscript standing for intermediate.) The scaling equations are as follows:

$$\lambda_{int}^* = 14.0(\lambda^* + 1), \tag{5a}$$

$$\alpha_{int} = 4.99(\cos(65.8^\circ)\alpha'' - \sin(65.8^\circ)\beta''), \tag{5b}$$

$$\beta_{int} = 4.99(\sin(65.8^\circ)\alpha'' + \cos(65.8^\circ)\beta''). \tag{5c}$$

Note that in Eq. (5a), the value 1 simply corresponds to the fact that the lowest possible value for λ^* , based on the error function, is -1 , which corresponds to ideal black, for which the Munsell value is zero. The last step is the problem of color constancy—separating the color of the surface from the color of the lighting. A certain amount of chromatic adaptation has already occurred as a result of the time-averaged adaptation that was introduced in Stage 2. (Indeed, if the coefficient C_f were to be set to -1 , all color shift arising from the illuminant color would, over time, be suppressed as this adaptation took place.) However, it is observed that this does not happen—color adaptation is not total, so the observer is able to know, weakly, the color and intensity of the ambient illumination. Another observation is that people are able to discount an illuminant immediately, in the sense that a piece of paper looks white, almost instantly, after the illuminant color has shifted [51], whereas adaptation is typically defined as a process that takes time. This instantaneous constancy requires spatial comparisons (e.g., the paper may not look white if it is seen in isolation), and can be quite sophisticated, taking into account the relative changes across surfaces and spatial variations in the intensity and color of the illuminant [52–54].




Fig. 10. Top row consists of seven separate plots of sets of Munsell samples for successively increasing value from 3, at the left, to 9, at the right. For each plot, the polar angle represents the Munsell hue, depicting 10 principal Munsell hues, and the radial distance represents the chroma, in increments of 2 Munsell units. By definition, this is a regular geometrical arrangement. The bottom row shows the corresponding plots for the output values generated by Eqs. (5d)–(5f). In converting from Cartesian to polar coordinates, the hue angle θ is determined by $\tan(\theta) = (\beta_{out}/\alpha_{out})$ and the chroma is equal to $\sqrt{(\alpha_{out})^2 + (\beta_{out})^2}$. If the fit were perfect, the patterns of the bottom row plots would perfectly match those of the top plots. The colors of the points shown are approximately the Munsell color corresponding to each plotted point.

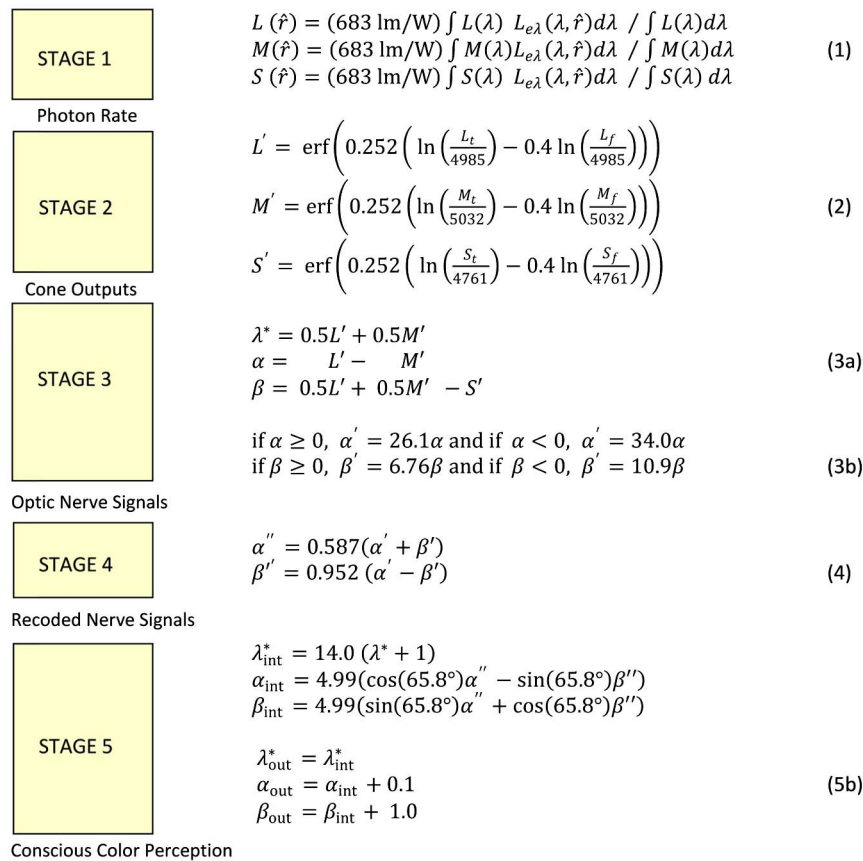


Fig. 11. Concise compilation of the equations of the five stages of the model, with the constants used to produce the results in Fig. 10. The only physiologically relevant parameters that were adjusted for that purpose are those in Eq. (2). The others arise from the statistical approach described here. For purposes of the data shown in Fig. 10, the Munsell samples were assumed to be illuminated with CIE Illuminant C and with uniform illuminance that causes a white surface to have a luminance of 400 cd/m².

For the purpose of comparing to the Munsell color set, it is sufficient to use a very simple white-balance procedure, in which the color shift (caused by changing the illuminant), for a neutral gray 20% reflective surface is evaluated. These evaluated shifts for the α and β channels are, respectively, labeled $\Delta\alpha_c$, and $\Delta\beta_c$. Using these shifts, the Cartesian output values for comparison to the Munsell color set are as follows:

$$\text{Munsell } V: \lambda_{\text{out}}^* = \lambda_{\text{int}}^* \quad (5d)$$

$$\text{Munsell } C \cos(H): \alpha_{\text{out}} = \alpha_{\text{int}} - \Delta\alpha_c \quad (5e)$$

$$\text{Munsell } C \sin(H): \beta_{\text{out}} = \beta_{\text{int}} - \Delta\beta_c \quad (5f)$$

Figure 10 shows that these five mathematical stages do a rather good job of modeling the perceptual experiences portrayed in the Munsell Color System. The top row depicts the arrangement of Munsell colors, for each of the 10 principal Munsell hues, with seven plots, one for each of the Munsell values ranging from 3 to 9. These are by definition placed in the regular arrays shown. The bottom row depicts the output of the model presented here. Ideally, the patterns in the bottom row would perfectly match those in the top and the fit is actually fairly good: the RMS error between the Munsell values

and the model values is about 1/2 unit in each direction, which is likely similar to the experimental error that is intrinsic to the Munsell system.

Figure 11 concisely summarizes the equations from the five stages described above, including the numerical values used for the parameters in the model.

Another characteristic of this model is that it also demonstrates several well-known observations that have long been attributed to nonlinear aspects of the color vision system. They include the Hunt effect, whereby surfaces appear less colorful at low light levels, the Abney effect, in which dilution of monochromatic light by white light produces hue shifts, and several others, which will be interesting areas for further study with this model.

3. CONCLUSIONS

The purpose of this paper has been to illustrate how the seemingly complex and unexplained characteristics of a popular approximately uniform color appearance system can be approached from fairly straightforward and plausible ideas about color coding in the visual system. We have done this by means of a simple computational approach that fulfills the standard scientific modeling objectives of simplicity, accuracy, integrability,

and correspondence, as well as plausibility from the perspectives of evolution, organismal development, and sensory adaptation. The quantitative expressions for each of the five stages are intentionally simple and undoubtedly imperfect, but nevertheless they describe the observable phenomena of color appearance metrics with reasonable accuracy. It is hoped the model's simplicity will be valuable pedagogically and perhaps stimulate new avenues of color perception research. As well, some may find it more satisfying to use the simple equations of this model for calculating anticipated color perception experiences. The results will not differ significantly from those of already available (but much more complex and less intuitive) calculation tools, but we hope the newfound simplicity will result in deeper understanding, and, perhaps, new insights.

Funding. National Eye Institute (NEI) (EY-10834); Natural Sciences and Engineering Research Council of Canada (NSERC) (EQPEQ 423559-12); Postdoctoral Fellowship of the Research Foundation Flanders (FWO) (12B4916N).

Acknowledgment. We thank Marty Banks, Jim Larimer, and Ed Pugh for thoughtful conversations that helped lead to this work.

REFERENCES

- G. Wyszecki, "Current developments in colorimetry," *Proc. AIC Color* **73**, 21–51 (1973).
- R. G. Kuehni and A. Schwarz, *Color Ordered: a Survey of Color Systems from Antiquity to the Present* (Oxford University, 2007).
- E. Schrödinger, "Grundlinien einer Theorie der Farbenmetrik im Tagessehen [Outline of a theory of colour measurement for daylight vision]," *Ann. Phys.* **368**, 397–426; 481–520 (1920).
- M. D. Fairchild, *Color Appearance Models*, 3rd ed. (Wiley-IS&T, 2013).
- M. D. Fairchild, "Colour appearance model," *Color Res. Appl.* **31**, 320–330 (2006).
- <http://munsell.com/>; retrieved September 27, 2015.
- <http://www.codeproject.com/Articles/7751/Use-Direct-D-To-Fly-Through-the-Munsell-Color-So>; retrieved May 31, 2015.
- R. G. D'Andrade and A. K. Romney, "A quantitative model for transforming reflectance spectra into the Munsell color space using cone sensitivity functions and opponent process weights," *Proc. Natl. Acad. Sci. USA* **100**, 6281–6286 (2003).
- P. Centore, "An open-source inversion algorithm for the Munsell renotation," *Color Res. Appl.* **37**, 455–464 (2012).
- K. C. McDermott and M. A. Webster, "Uniform color spaces and natural image statistics," *J. Opt. Soc. Am. A* **29**, A182–A187 (2012).
- D. I. MacLeod and R. M. Boynton, "Chromaticity diagram showing cone excitation by stimuli of equal luminance," *J. Opt. Soc. Am.* **69**, 1183–1186 (1979).
- M. R. Luo, G. Cui, and C. Li, "Uniform colour spaces based on CIECAM02 colour appearance model," *Color Res. Appl.* **31**, 320–330 (2006).
- H. B. Barlow, "Possible principles underlying the transformations of sensory messages," in *Sensory Communication*, W. A. Rosenblith, ed. (MIT, 1961), pp. 217–234.
- T. van der Twer and D. I. A. MacLeod, "Optimal nonlinear codes for the perception of natural colours," *Network* **12**, 395–407 (2001).
- C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication* (Urbana University of Illinois, 1949).
- Z. P. Li, *Understanding Vision: Theory, Models, and Data* (Oxford University, 2014).
- S. Laughlin, "A simple coding procedure enhances a neuron's information capacity," *Z. Naturforsch. C* **36**, 910–912 (1981).
- D. I. A. MacLeod and T. van der Twer, "Optimal nonlinear codes for the perception of natural colours," *Network* **12**, 395–407 (2001).
- G. Buchsbaum and A. Gottschalk, "Trichromacy, opponent colours coding and optimum colour information transmission in the retina," *Proc. R. Soc. B* **220**, 89–113 (1983).
- J. J. Atick, "Could information theory provide an ecological theory of sensory processing," *Network* **3**, 213–251 (1992).
- J. B. Jonas, A. M. Schmidt, J. A. Müller-Bergh, U. M. Schlötzer-Schrehardt, and G. O. Naumann, "Human optic nerve fiber count and optic disc size," *Invest. Ophthalmol. Vis. Sci.* **33**, 2012–2018 (1992).
- K. Kristin, J. McLean, R. Segev, M. A. Freed, M. J. Berry II, V. Balasubramanian, and P. Sterling, "How much the eye tells the brain," *Curr. Biol.* **16**, 1428–1434 (2006).
- C. A. Curcio, K. R. Sloan, R. E. Kalina, and A. E. Hendrickson, "Human photoreceptor topography," *J. Comp. Neurol.* **292**, 497–523 (1990).
- M. V. Srinivasan, S. B. Laughlin, and A. Dubs, "Predictive coding: a fresh view of inhibition in the retina," *Proc. R. Soc. B* **216**, 427–459 (1982).
- E. P. Simoncelli and B. A. Olshausen, "Natural image statistics and neural representation," *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
- F. Long, Z. Yang, and D. Purves, "Spectral statistics in natural scenes predict hue, saturation, and brightness," *Proc. Natl. Acad. Sci. USA* **103**, 6013–6018 (2006).
- J. Krauskopf, D. R. Williams, M. B. Mandler, and A. M. Brown, "Higher order color mechanisms," *Vis. Res.* **26**, 23–32 (1986).
- CIE, *Fundamental Chromaticity Diagram with Physiological Axes—Part I* (CIE, 2006).
- K. L. Naka and W. A. H. Rushton, "S-potentials from colour units in the retina of fish (Cyprinidae)," *J. Physiol.* **185**, 536–555 (1966).
- F. Ratliff, *On Mach's Contributions to the Analysis of Sensations* (Springer, 1970).
- S. B. Laughlin, "A simple coding procedure enhances a neuron's information capacity," *Z. Naturforsch.* **36**, 910–912 (1981).
- F. Rieke and M. E. Rudd, "The challenges natural images pose for visual adaptation," *Neuron* **64**, 605–616 (2009).
- J. Neitz, J. Carroll, Y. Yamauchi, M. Neitz, and D. R. Williams, "Color perception is mediated by a plastic neural mechanism that is adjustable in adults," *Neuron* **35**, 783–792 (2002).
- T. Jarsky, M. Cembrowski, S. M. Logan, W. L. Kath, H. Riecke, J. B. Demb, and J. H. Singer, "A synaptic mechanism for retinal adaptation to luminance and contrast," *J. Neurosci.* **31**, 11003–11015 (2011).
- B. Wark, A. Fairhall, and F. Rieke, "Timescales of inference in visual adaptation," *Neuron* **61**, 750–761 (2009).
- S. A. Baccus and M. Meister, "Fast and slow contrast adaptation in retinal circuitry," *Neuron* **36**, 909–919 (2002).
- M. A. Webster, "Visual adaptation," *Annu. Rev. Vis. Sci.* (to be published).
- D. Tranchina and C. S. Peskin, "Light adaptation in the turtle retina: embedding a parametric family of linear models in a single nonlinear model," *Vis. Neurosci.* **1**, 339–348 (1988).
- D. Aurelien, P. T. Fini, K. W. Houser, Y. Ohno, M. P. Royer, K. Smet, M. Wei, and L. Whitehead, "Development of the IES method for evaluating the color rendition of light sources," *Opt. Express* **23**, 15888–15906 (2015).
- D. L. Ruderman, T. W. Cronin, and C. C. Chiao, "Statistics of cone responses to natural images: implications for visual coding," *J. Opt. Soc. Am. A* **15**, 2036–2045 (1998).
- B. B. Lee, "Color coding in the primate visual pathway: a historical view," *J. Opt. Soc. Am. A* **31**, A103–A112 (2014).
- D. M. Dacey and B. B. Lee, "The 'blue-on' opponent pathway in primate retina originates from a distinct bistratified ganglion cell type," *Nature* **367**, 731–735 (1994).
- J. D. Mollon, "'Tho' she kneel'd in that place where they grew...' The uses and origins of primate colour vision," *J. Exp. Biol.* **146**, 21–38 (1989).

44. G. Wyszecki and W. S. Stiles, "Abney and Bezold-Brücke effects," in *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd ed. (Wiley, 2000), Section 5.9, p. 420.
45. M. A. Webster, Y. Mizokami, and S. M. Webster, "Seasonal variations in the color statistics of natural images," *Network* **18**, 213–233 (2007).
46. M. A. Webster and J. D. Mollon, "The influence of contrast adaptation on color appearance," *Vis. Res.* **34**, 1993–2020 (1994).
47. J. J. Atick, A. Li, and A. N. Redlich, "What does post-adaptation color appearance reveal about cortical color representation?" *Vis. Res.* **33**, 123–129 (1993).
48. R. T. Eskew, "Higher order color mechanisms: a critical review," *Vis. Res.* **49**, 2686–2704 (2009).
49. G. D. Horwitz and C. A. Hass, "Nonlinear analysis of macaque V1 color tuning reveals cardinal directions for cortical color processing," *Nat. Neurosci.* **15**, 913–919 (2012).
50. H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Autom. Control* **19**, 716–723 (1974).
51. D. H. Foster and S. M. C. Nascimento, "Relational colour constancy from invariant cone-excitation ratios," *Proc. Royal Soc. B* **257**, 115–121 (1994).
52. F. A. Kingdom, "Lightness, brightness and transparency: a quarter century of new ideas, captivating demonstrations and unrelenting controversy," *Vis. Res.* **51**, 652–673 (2011).
53. B. Anderson, "The perceptual organization of depth, lightness, color, and opacity," in *The New Visual Neurosciences*, J. S. Werner and L. M. Chalupa, eds. (MIT, 2014), pp. 653–664.
54. D. H. Brainard and A. Radonjić, "Color constancy," in *The New Visual Neurosciences*, J. S. Werner and L. M. Chalupa, eds. (MIT, 2014), pp. 545–556.