To be published in Journal of the Optical Society of America A:

Title: A spectral theory of color perception
Authors: James Clark and Sandra Skaff
Accepted: 29 September 2009
Posted: 30 September 2009
Doc. ID: 107314

A spectral theory of color perception

James J. Clark[1,*] and Sandra Skaff[1,2]

[1] *Centre for Intelligent Machines, McGill University, 3480 University Street, Montreal, Quebec, Canada*

[2] *Currently with Xerox Research Centre Europe, 6, chemin de Maupertuis, F - 38240 Meylan, France*

[*] *Corresponding author: clark@cim.mcgill.ca*

This paper adopts the philosophical stance that colors are real and can be identified with spectral models based on the photoreceptor signals. A statistical setting represents spectral profiles as probability density functions. This permits the use of analytic tools from the field of Information Geometry to determine a new kind of color space and structure deriving therefrom. In particular, the metric of the color space is shown to be the Fisher Information matrix. A maximum entropy technique for spectral modeling is proposed that takes into account measurement noise. Theoretical predictions provided by our approach are compared with empirical colorfulness and color similarity data.

OCIS codes: 330.1720,330.1690,330.4060.

1. INTRODUCTION: A PHYSICAL THEORY OF COLOR PERCEPTION

A long-standing philosophical question is whether color is a physical quantity or is purely a mental construct [1]. In other words, is color "out in the world" or "in the brain"? On one side of the debate stand the *Objectivists* or *Realists* who hold that colors have objective essences and are properties of physical objects. On the other side stand the proponents of the *Illusion* theory of color which holds that colors are virtual properties and our perceptions of color are illusory, in that physical objects are not colored, they just appear to be.

In this paper we propose taking the philosophy of color realism at face value and using it to develop a scientific theory of color perception. Our approach is based on the ideas of a prominent group of color realist philosophers who identify color with surface spectral reflectances. Influential representatives of this community are Byrne and Hilbert [2] and Tye [3]. We modify their ideas slightly to posit that the physiological, and thence psychological, quantity known as color is the particular *model* of spectra provided by the sensory apparatus of the human body. We claim then, that all statements about color can be derived from the physical laws underlying spectra and of the structure of the human sensory apparatus.

Noted philosopher of color Christopher Hardin commented in [4] that:

> *"It is a curious sociological fact that many philosophers, but very few visual scientists, are color realists."*

He suggests that the reason for this fact is that philosophers are generally loath to "banish colors from the physical world" for fear that they "will take up refuge in what Gilbert Ryle once called the dust-bin of the mind, never to be dislodged." [4]. He does not, however, speculate as to why visual scientists generally take the opposite view, that color is in the

mind. This may be a result of the intense focus of vision scientists on physiological and psychological studies of the color vision system. With such a strong focus on the brain it is perhaps natural that the bias of color scientists is to consider color as an internal quantity. This bias is not universal, however. Recently cognitive scientists have been considering perceptual systems in terms of the interaction of the observer with its environment. Such considerations lead to treating perception as a process of inducing sensori-motor relationships that reveal invariant structures of the world, independent of the specific details of the perceptual apparatus (see O'Regan and Noë [5] for a definitive statement of this viewpoint). In the case of color perception the relevant invariant physical quantities are the spectral energy distribution of lights and the spectral reflection functions of surfaces.

The closest that the current science of color comes to a focus on spectral modeling in understanding color vision processes is in the study of color constancy. Color constancy is the ability of the human visual system to perceive a surface as a given color independently of the color of the light illuminating the surface. While there are many successful color constancy theories and associated algorithms that do not consider the spectra of the light, the most flexible and robust methods are based on spectral models [6]. Most of these methods follow the seminal paper of Maloney and Wandell [7] who suggested utilizing a *linear model* of the spectrum of the illuminant and of the surface reflectance spectrum. This is not surprising, since the confounding of the illuminant and surface occurs at the level of the spectra, and not at the neural level. If one considers the role of color perception to be that of surface spectral modeling, then the problem of color constancy disappears, to be replaced by the fundamental problem of finding the invariant surface spectrum. Thus discounting of the illuminant spectrum becomes part of the determination of the surface spectrum rather than

part of a color constancy or color correction process.

In this paper our goal is to break with the mainstream internal theories of color that are based solely on photoreceptor channel signals. We describe a spectral modeling based theory of color vision that qualitatively explains various aspects of empirical color appearance phenomena. The approach we take is to begin with a statistical representation of the spectra of light. This allows the application of the mathematical framework known as *information geometry* to deduce the structure of color space, which in turn answers the question of *why* the empirical color appearance data takes the form it does.

In this paper we focus our efforts on the modeling of the spectra of lights, and do not consider here modeling the spectral reflectance functions of surfaces.

2. STATISTICAL SPECTRAL MODELING

In order to develop physically meaningful descriptions of color properties, one must identify these with physical quantities. As mentioned in the introduction, we propose associating with every spectrum a particular color perception. Thus, to determine the structure of color space we must determine the structure of the space of spectra. To do this we formulate the details of the theory in a statistical setting, and associate spectra with probability density functions. In this way, we will form an association between color space and the space of probability densities, for which there exists a rich mathematical structure.

We begin by considering a stream of photons, detected one at a time by a mechanism that can measure the wavelength of each photon. Each photon emitted by the source has a certain probability, $\Pr(\lambda_1 < \lambda < \lambda_2)$, that its wavelength, λ, lies in the range (λ_1, λ_2). This quantity obeys the laws of probability derived from Cox's postulates [8]. That is, taking A to be the

condition where $(\lambda_1 < \lambda < \lambda_2)$ then since a photon must have some particular wavelength $\Pr(A) = 1$ represents certain truth while $\Pr(A) = 0$ represents certain falsehood. Also $\Pr(A) + \Pr(A^C) = 1$ where A^C is the complement of A, and $Pr(A, B) = Pr(A)Pr(B|A) = Pr(B)P(A|B)$. If $\lambda_2 - \lambda_1 = \mathrm{d}\lambda$ becomes vanishingly small, then $\Pr(\lambda_1 < \lambda < \lambda_1 + \mathrm{d}\lambda) = p(\lambda_1)\mathrm{d}\lambda$. Thus $p(\lambda)$ is seen to be a probability density function. In this paper this probability density will be refereed to as the *spectral profile* of the photon source. Note that the spectral profile carries no information about the *intensity* of the photon source. The intensity of the photon source must be determined or given separately. The product of the spectral profile and the intensity will be referred to as the *spectrum*. The separation of the source intensity and the spectral profile is a key aspect of the theory and will impact the computational structure.

Let us denote the space of all spectral profiles (or probability density functions) by $\mathcal{S}$. In the theory being proposed in this paper, each element of $\mathcal{S}$ corresponds to a different possible spectral profile and hence to a distinct color perception. Only some of these possibilities can be obtained by the human visual system, however. The data that the human visual system has to work with to model the spectral profile of incoming light is the collection of signals provided by the retinal photoreceptors. These photoreceptor signals divide the space $\mathcal{S}$ into a group of equivalence classes, called *metamer sets* [9]. The spectral profiles in a given metamer set all produce the same photoreceptor signals. The metamer sets are exclusive (i.e. they do not have any spectral profiles in common) and cover the whole of $\mathcal{S}$. Therefore the metamer sets form a partition of $\mathcal{S}$.

2.A. Spectral Modeling as Metamer Selection - the Maximum Entropy Approach

Human color vision is marked by a singleness of the percept. When humans see a colored patch, they report perceiving a single color, rather than a whole set of colors. This suggests some mechanism for selecting one element from the metamer set. We refer to this selection process as *spectral modeling*.

One could argue that there is no need to select a specific element of the metamer set, and instead work with the more general equivalence classes defined by the metamers. If one considers only a single observer then the (equivalence) classes of metamers are well-defined by the characteristics of the photosensors. Inasmuch as these equivalence classes form a partition of the stimulus space (here taken to be the space of spectral power distributions) one could conceivably associate colors with these classes. This approach was taken by Krantz [10] in developing the geometry of color from very general principles.

Difficulties arise, however, when dealing with two or more observers, each having different sensing systems. For two observers one is faced with the problem that there are now two different definitions of the equivalence classes of metamers. How do we now define the notion of color? One approach is to say that each observer has their own, private, notion of color, defined by its own specific set of metameric equivalence classes. This amounts to the "illusion" theory of color mentioned in the introduction of the paper, where color is considered as purely a mental concept, and can be different from observer to observer.

We are proposing to take a different approach, that of "color realism", in which colors are associated with physical properties, and as such should be the same for all observers. This means that the notion of color can not be based on the metamer equivalence classes

(since these change from observer to observer) but rather on the underlying space of spectral energy distributions, which being a physical space outside of the observers is the same for all observers.

There is actually experimental data which supports the "realist" view, suggesting that different observers can perceive the same colors in spite of having quite different metameric equivalence classes. This support comes from reports of people that have unilateral color deficient vision, e.g. that have red or green-deficiency in one eye yet have normal color vision in the other. For example, the unilateral dichromat reported by Graham and Hsia [11] perceives yellow or blue in his dichromatic eye much in the same way as in his trichromatic eye. Color matching by this observer showed that monochromatic blue (or yellow) stimuli, with a wavelength of 470nm (resp. 570nm), looked similar in both eyes. Similarly, the unilateral dichromat reported by MacLeod and Lennie [12] perceived stimuli with a wavelength of 610nm (orange) to be the same in both eyes. These studies imply that observers with differing sensory systems, and hence differing metamer sets, can nonetheless perceive similar colors. This not to say that all stimuli are perceived in the same way by different observers, only that some may be - those that induce the same spectral model. For example, in the case of the dichromat and trichromat they may both perceive the color of light of 570nm wavelength in the same way (both see yellow) but will differ in their perception of light of 650nm wavelength (the trichomat perceiving red and the dichromat perceiving a dark yellow).

Note that we are not saying that different observers viewing the same physical surface will always have the same color percept. In general, the colors perceived by different observers will be different, reflecting differences (usually minor for normal color vision) in their metamer sets. Our example of the unilateral monochromat is intended to show that *some*

color percepts can be the same between observers (.e.g yellows and blues) in spite of differences in the metamer set structures. This is because different metamer sets can contain the same maximum entropy spectrum.

The realist view can therefore handle the individual differences in the wavelengths chosen as "unique hues" [13]. In this case each observer would have a particular stimulus wavelength that they would deem to correspond to a unique hue, and this stimulus would result in a specific spectral model and hence, in our view, to a particular color perception. What must be kept in mind, however, is that the spectral model obtained may, and probably will, differ from observer to observer, implying that the color ascribed to what an observer calls a unique hue will vary between observers. Thus, my shade of unique green will be different that yours.

Proceeding with the thesis that colors are associated with spectral models, we will now examine the geometry of the space of spectral models. We must first quantify the process for obtaining the model. We propose that the spectral model produced in response to light falling on the photoreceptors is determined by two factors - the spectral sensitivity profiles of the photoreceptors, $\boldsymbol{f}(\lambda)$, which define the partitioning of $\mathcal{S}$ into the metamer sets, and the particular selection process used which determines the specific element of the metamer set that is chosen as the spectral model. This is illustrated in figure 1. In the figure three different photoreceptors define metamer sets. In each of these sets, one particular element is singled out by the spectral modeling approach. The locus of all such distinguished elements over all metamer sets will be referred to as the spectral color manifold, and captures all spectral models that can be generated by the perceptual system.

There have been many ways proposed in the literature to select a single element of the metamer set. Metamer sets are generally difficult to compute, and to represent, so for reasons

of computational simplicity many approaches first identify a subset of the metamer set, and then a specific element is picked from this subset. One common approach to defining such metamer subsets is to use linear models [7] to describe spectra. The linear model approach uses the linear superposition of a small set of basis functions to approximate the spectral profile, with the weights of the basis functions chosen so as to make the resulting spectral profile be in the metamer set. Note that, in general, the linear model approach can only describe a subset of the metamers, so not all metamers can be obtained in these approaches. For example, monochromatic (impulse) spectra cannot be modeled, unless a very large number of basis functions are used. Morovic and Finlayson [9] describe three different ways of singling out a single member of a linear model metamer set: the centroid or mean metamer (in which the weight vectors in the set are assumed to be equally probable), the most likely metamer given an assumption that the weight vectors are normally distributed, and the metamer with the smoothest spectrum.

In this paper we perform spectral modeling by selecting the element of the metamer set that has the maximum entropy [14]. This approach is motivated by the writings of Jaynes [15] who held that, given a set of constraints on a probability distribution, the best choice of the distribution out of all those consistent with the constraints is the one with maximum entropy. The principal justification of this view is the so-called *entropy concentration* theorem [16,17]. In the context of spectral profiles the entropy concentration theorem states that the number of different ways in N detected photons can be placed into n discrete wavelength bins is $W(p_1, p_2...p_n) = N!/[(Np_1)!...(Np_n)!]$, where p_i is the proportion of photons placed into bin i. For large N the Stirling approximation gives $W \approx \exp(NH(p_1, p_2, ..., p_n))$, where H is the

entropy of the bin distribution.

$$H(p_1, p_2, ...p_n) = -\sum_{i=1}^{n} p_i \log(p_i) \tag{1}$$

It can be seen that higher entropy distributions have a much larger number of possible ways to fill the bins than do low entropy distributions. The entropy concentration theorem is based only on a counting of possibilities, however, and it does not consider the probability of these possibilities. It may be, for physical reasons, that some of the high entropy distributions are unlikely. But even so, the combinatorics is working against the lower entropy possibilities even if they may have higher probability of occurring, as there are just fewer of them. The entropy concentration theorem suggests that the photon bin counts consistent with the measured photoreceptor signals that will most frequently be encountered are those with the maximum entropy of all possibilities. Thus, in the remainder of this paper, we will select as our spectral model the metameric spectral profile which has the maximum entropy over all metamers.

3. THE INFORMATION GEOMETRY OF SPECTRAL PROFILES

Formulating spectral models in terms of probability density functions permits the application of the methods of what is known as *Information Geometry* [18]. Information geometry is based on the foundational work of Fisher [19] investigating the geometry of probability spaces. Information geometry considers such spaces as differentiable manifolds. This allows the use of tools from the field of differential geometry, in particular the concepts of Riemannian metrics and affine connections. In this paper, we will use these tools to determine the structure of color space, such as the similarity or distance between different colors.

In keeping with the position that color perceptions are synonymous with spectral profiles, the manifold defined by all possible maximum entropy spectral models will be referred to as the *spectral color manifold*. The spectral color manifold is a very restricted sub-manifold of the infinite-dimensional spectral manifold $\mathcal{S}$. The form and characteristics of the spectral color space depend on the photoreceptor sensitivity functions and on the spectral modeling technique. If either of these change, so too will the spectral color manifold. But the underlying space $\mathcal{S}$ is not changed - changing of the photoreceptors merely changes the subset of spectra (or colors) that can be perceived.

To begin the study of the structure of the spectral color manifold we first examine the measurement process. Measurements are assumed to be obtained by a linear projection of the spectral profile onto the photoreceptor sensitivity functions:

$$\boldsymbol{r} = \beta \int_{\Lambda} \boldsymbol{f}(\lambda) p(\lambda) \mathrm{d}\lambda \tag{2}$$

where $p(\lambda) \in \mathcal{S}$ is the spectral profile of the incident light, and $\boldsymbol{f}(\lambda)$ is the vector of projection kernels, corresponding to the spectral sensitivity profiles of the photoreceptors. $\Lambda = [\lambda_{\min}, \lambda_{\max}]$ is the support of the photoreceptor spectral sensitivities, which is the wavelength interval over which at least one of the photoreceptors has a non-zero sensitivity (i.e. the visible wavelength range). β is the intensity of the light incident on the photoreceptors over wavelengths in the support of the photoreceptor spectral sensitivities,

Spectral components outside the visible range (by definition) lie in the null space of the projection operation defined by the measurement equation. Any variation in the spectral power in wavelengths outside the visible range will cause no change in the measurements. Thus, in the absence of any other sources of information or additional constraints, nothing

can be said about the spectral power outside the visible range. When we look at a rainbow and see darkness at either side of a band, our vision system is not telling us that there is zero energy in the infrared and ultraviolet - on the contrary, it is telling us that there is zero energy in the visible wavelengths in these areas.

For the time being, we will assume that the intensity β is known, and focus on a normalized version of the measurement vector, $\boldsymbol{\eta} = \boldsymbol{r}/\beta$:

$$\boldsymbol{\eta} = \int_{\Lambda} \boldsymbol{f}(\lambda) p(\lambda) \mathrm{d}\lambda \tag{3}$$

It is a well-known result [15] that the probability density $\hat{p}(\lambda)$ having the maximum entropy consistent with the measurements defined by the equation above is an element of an exponential family of densities with the following form:

$$\hat{p}(\lambda; \theta) = \exp(\boldsymbol{\theta} \cdot \boldsymbol{f}(\lambda) - \psi(\boldsymbol{\theta})) \tag{4}$$

where $\boldsymbol{\theta}$ is the vector of natural parameters. The quantity $\psi(\boldsymbol{\theta})$ normalizes the model spectral profile so that it integrates to 1, as required for it to be a probability density:

$$\int_{\Lambda} \hat{p}(\lambda; \theta) \mathrm{d}\lambda = \int_{\Lambda} \exp(\boldsymbol{\theta} \cdot \boldsymbol{f}(\lambda) \exp\left(-\psi(\boldsymbol{\theta})\right) \mathrm{d}\lambda = 1 \tag{5}$$

Since $\exp\left(-\psi(\boldsymbol{\theta})\right)$ does not depend on λ it can be taken outside of the integral:

$$\exp\left(-\psi(\boldsymbol{\theta})\right) \int_{\Lambda} \exp(\boldsymbol{\theta} \cdot \boldsymbol{f}(\lambda) \mathrm{d}\lambda = 1 \tag{6}$$

which then yields the following expression for $\psi(\boldsymbol{\theta})$:

$$\psi(\boldsymbol{\theta}) = \log \int_{\Lambda} \exp(\boldsymbol{\theta} \cdot \boldsymbol{f}(\lambda)) \mathrm{d}\lambda \tag{7}$$

The natural parameters $\boldsymbol{\theta}$ can serve as a set of coordinates for the exponential family of distributions. The expectations of the $\boldsymbol{f}(\lambda)$ with respect to $\hat{p}(\lambda; \boldsymbol{\theta})$ yield another set of

coordinates:

$$\hat{\boldsymbol{\eta}} = E_{\hat{p}}\{\boldsymbol{f}(\lambda)\} = \int_{\Lambda} \boldsymbol{f}(\lambda)\hat{p}(\lambda;\theta)\mathrm{d}\lambda \tag{8}$$

If we take the $\boldsymbol{f}(\lambda)$ to be the photoreceptor spectral sensitivities, it is immediately evident that, since the true spectral profile $p(\lambda)$ and the model spectral profile $\hat{p}(\lambda;\boldsymbol{\theta})$ are metamers, the expectations in equation (8) are equal to the normalized responses of the photoreceptors. Thus we will not distinguish between $\boldsymbol{\eta}$ and $\hat{\boldsymbol{\eta}}$ in what follows.

Suppose we are given two different spectral profiles, $p_P(\lambda), p_Q(\lambda) \in \mathcal{S}$. These give rise to normalized measurements $\boldsymbol{\eta}_P$ and $\boldsymbol{\eta}_Q$ which in turn give rise to model spectral profiles $\hat{p}_P(\lambda)$ and $\hat{p}_Q(\lambda)$ which lie in the spectral color manifold. Because of the assumed linearity of the measurement process a linear combination of the two spectral profiles $ap_P(\lambda) + (1-a)p_Q(\lambda)$, $a \in [0,1]$, result in a measurement vector that is the linear combination of the individual measurements, $\boldsymbol{\eta}_{P+Q} = (1-a)\boldsymbol{\eta}_P + a\boldsymbol{\eta}_Q$. As this measurement arises from a physically realizable spectral profile, there exists a corresponding model spectral profile (i.e. the maximum entropy element of the metamer set associated with $\boldsymbol{\eta}_{P+Q}$) which lies on the spectral color manifold. The line formed in the spectral color manifold as a goes from $0 \rightarrow 1$ is referred to as the *m-geodesic*. This is a geodesic in the sense of being the curve which minimizes a certain distance measure, the canonical divergence (to be described later). As this line lies completely in the manifold, the manifold is considered to be, in the parlance of Amari [18], *m-flat* (where "m" stands for "mixture"). In addition, given the $\boldsymbol{\theta}$ coordinates of any two model spectral profiles, $\hat{p}_P$ and $\hat{p}_Q$, that lie on the spectral color manifold, a linear combination of these two coordinates, $(1-a)\boldsymbol{\theta}_P + a\boldsymbol{\theta}_Q$, defines a model spectral profile which also lies on the manifold. In this case the resulting spectral color manifold is said to be *e-flat*

(where "e" stands for "exponential"). The line formed in the spectral color manifold as a goes from $0 \rightarrow 1$ is referred to as the *e-geodesic*. The spectral color manifold is therefore said to be *dually-flat*, as it is both m-flat and e-flat. Note that this flatness is in the sense of the e- and m-geodesic, and does not mean that the manifold is Euclidean.

As pointed out by Amari [18] the dually-flat nature of the spectral color manifold implies a strong link between the $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ coordinates. In particular they are related through a *Legendre* transformation between two potential functions ψ and ϕ:

$$\phi(\boldsymbol{\eta}) = \max_{\boldsymbol{\theta}}[\boldsymbol{\theta} \cdot \boldsymbol{\eta} - \psi(\boldsymbol{\theta})] \tag{9}$$

The transformation is its own inverse, so that:

$$\psi(\boldsymbol{\theta}) = \max_{\boldsymbol{\eta}}[\boldsymbol{\theta} \cdot \boldsymbol{\eta} - \phi(\boldsymbol{\eta})] \tag{10}$$

and the two transformations are related through the following identity:

$$\psi(\boldsymbol{\theta}) + \phi(\boldsymbol{\eta}) = \boldsymbol{\theta} \cdot \boldsymbol{\eta} \tag{11}$$

The coordinates $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ are therefore related by:

$$\frac{\partial \psi(\boldsymbol{\theta})}{\partial \theta_i} = \eta_i \tag{12}$$

$$\frac{\partial \phi(\boldsymbol{\eta})}{\partial \eta_i} = \theta_i \tag{13}$$

The potential $\psi(\boldsymbol{\theta})$ can be identified with the normalizing term in equation (4). An interpretation of the $\phi(\boldsymbol{\eta})$ potential can be obtained by considering the entropy of the model spectral profile, $\hat{p}$. This is given by the negative of the expectation of the logarithm of $\hat{p}(\lambda;\theta)$:

$$H(\hat{p}) = -E\{\log(\hat{p}(\lambda;\theta))\} = -\int_{\Lambda} \hat{p}(\lambda;\boldsymbol{\eta}) \log(\hat{p}(\lambda;\boldsymbol{\eta}))\mathrm{d}\lambda \tag{14}$$

If we take $\hat{p}$ to be a member of an exponential family we can substitute the expression for $\hat{p}$ from equation (4) into the above equation and get:

$$H(\hat{p}) = -E\{\boldsymbol{\theta} \cdot \boldsymbol{f} - \psi(\boldsymbol{\theta})\} = -(\boldsymbol{\theta} \cdot \boldsymbol{\eta} - \psi(\boldsymbol{\theta})) = -\phi(\boldsymbol{\eta}) \tag{15}$$

Thus the potential $\phi(\boldsymbol{\eta})$ is equal to the negative of the entropy of the model spectral profile associated with a normalized measurement vector $\boldsymbol{\eta}$.

Note that, depending on the form of $\boldsymbol{f}$, there will be points in the $\boldsymbol{\eta}$ space that admit no corresponding element of the exponential family. A trivial example of this is that the elements of $\boldsymbol{\eta}$ cannot be negative, due to the positivity of spectra $p(\lambda)$ and the photoreceptor sensitivities $\boldsymbol{f}(\lambda)$. Also, it can be seen that the maximum possible value of an individual component η_i of $\boldsymbol{\eta}$ is induced by a spectrum $p(\lambda) = f_i(\lambda)/\int_\Lambda f_i(\lambda)\mathrm{d}\lambda$. For any measurement with η_i higher than this there is no metameric spectral model which gives this measurement. We will refer to the set of points $\boldsymbol{\eta}$ that admit spectral models as the *admissible normalized measurement space*, $\mathcal{M}$. In $\boldsymbol{\theta}$ space all points have corresponding elements of the exponential family. $\mathcal{M}$ is therefore the image of the entire $\boldsymbol{\theta}$ space.

3.A. The Fisher Information Metric

The information-geometric treatment of the manifold of spectral models permits specification of a metric and therefore a distance measure. Distance measures are of great importance in color science as they provide a way to judge the relative similarity between colors, and to quantify discriminibility of colors.

There are an infinite number of metrics that could be chosen, and each could predict different color similarity/discriminibility phenomena. There is, however, one particular metric that has unique properties which make it more appropriate than any other, and that is the

so-called Fisher metric (sometimes referred to as the Fisher-Rao metric). As Maybank [21] explains, it is desirable that a measure of distance between any two points in a probability space (i.e. the space of spectral power distributions in our application) be invariant under reparametrizations or coordinate transformations (e.g. in wavelength). Amari and Nagaoka [18] showed that this invariance property only holds for the Fisher metric (up to a scale factor). An even stronger statement was proven by Cencov [22] who showed that the Fisher metric is the only metric on the space of discrete probability distributions that is compatible with all Markov morphisms (or stochastic maps), in the sense that application of a stochastic map to two points does not increase the distance between the points. A stochastic map in this context is a linear map from the space of probability distributions into itself, or a subset of itself. Examples include permutation of wavelengths, and reducing the dimensionality of the probability space (e.g. by having wider histogram bins). In the context of color vision, this compatibility means that a stochastic map can not increase the discriminibility between two colors, as defined by the metric. Cencov's result shows that the Fisher metric is the only metric that has this property.

The Fisher metric (or the Fisher information matrix) has the following form:

$$g_{ij}(\boldsymbol{\theta}) = \int_{\Lambda} \hat{p}(\lambda;\boldsymbol{\theta}) \left(\frac{\partial}{\partial\theta_i} \log \hat{p}(\lambda;\boldsymbol{\theta}) \frac{\partial}{\partial\theta_j} \log \hat{p}(\lambda;\boldsymbol{\theta}) \right) \mathrm{d}\lambda \tag{16}$$

Rao [20] showed that the Fisher information matrix is a Riemannian metric on the space of probabilities. The matrix $G = [g_{ij}]$ formed by the elements of the metric is positive definite, as is its inverse G^{-1}. This means that $\boldsymbol{v}^T G \boldsymbol{v} > 0$ for any non-zero vector $\boldsymbol{v}$. For exponential families the Fisher information metric can be written as

$$g_{ij}(\boldsymbol{\theta}) = \frac{\partial\eta_j}{\partial\theta_i} = \int_{\Lambda} \hat{p}(\lambda;\boldsymbol{\theta})(f_i(\lambda) - \eta_i)(f_j(\lambda) - \eta_j)\mathrm{d}\lambda = \int_{\Lambda} \hat{p}(\lambda;\boldsymbol{\theta}) f_i(\lambda) f_j(\lambda)\mathrm{d}\lambda - \eta_i\eta_j \tag{17}$$

The integral in the rightmost equation above can be thought of as the response to the model spectrum of a synthetic photoreceptor with a spectral sensitivity profile of $f_{ij}(\lambda) = f_i(\lambda)f_j(\lambda)$. Let us denote this response by η_{ij}. Then the metric is given by:

$$g_{ij}(\boldsymbol{\theta}) = \eta_{ij}(\boldsymbol{\theta}) - \eta_i(\boldsymbol{\theta})\eta_j(\boldsymbol{\theta}) \tag{18}$$

The metric allows us to quantify the structure of the manifold i.e. the relative distances between points on the manifold. The *line element*, or the square of the arc length between two points separated by an infinitesimal distance, is given by:

$$\mathrm{ds}^2 = \mathrm{d}\boldsymbol{\theta}^T G(\boldsymbol{\theta})\mathrm{d}\boldsymbol{\theta} = \mathrm{d}\boldsymbol{\theta} \cdot \mathrm{d}\boldsymbol{\eta} \tag{19}$$

For a dually-flat space the metric relative to the dual parameters ($\boldsymbol{\eta}$) is the inverse of the metric relative to the primary parameters ($\boldsymbol{\theta}$) so that

$$\mathrm{ds}^2 = \mathrm{d}\boldsymbol{\eta}^T G^{-1}(\boldsymbol{\eta})\mathrm{d}\boldsymbol{\eta} = \mathrm{d}\boldsymbol{\eta} \cdot \mathrm{d}\boldsymbol{\theta} \tag{20}$$

The Riemannian distance between two points, P and Q, on the manifold is the length of the geodesic joining these two points:

$$L(P,Q) = \int_{s(P)}^{s(Q)} \mathrm{d}s = \int_0^1 \left(\frac{\mathrm{d}\boldsymbol{\theta}}{dt}^T G \frac{\mathrm{d}\boldsymbol{\theta}}{dt}\right)^{1/2} \mathrm{dt} \tag{21}$$

where t is a parameter along the geodesic. In general computation of the Riemannian distance is difficult due to the required integration, and to the difficulty in determining the geodesic curve. The integration in most cases needs to be done numerically, which becomes problematic for large θ values, where the elements of G become small. Instead, we work with a pseudo-distance that is much simpler to compute. This is the so-called *canonical divergence* [18]:

$$D(P||Q) = \psi(\boldsymbol{\theta}_P) + \phi(\boldsymbol{\eta}_Q) - \boldsymbol{\theta}_P \cdot \boldsymbol{\eta}_Q \tag{22}$$

For small differences in the $\boldsymbol{\theta}$ coordinates of the two points, the canonical divergence is a second order approximation to one half of the square arc length distance:

$$D(P||P+\Delta\boldsymbol{\theta}) = \Delta\boldsymbol{\theta}^T G \Delta\boldsymbol{\theta}/2 + o(||\Delta\boldsymbol{\theta}||^2) \approx \mathrm{d}s^2/2 \tag{23}$$

where $\Delta\theta_i$ is the difference in the θ_i coordinate between points Q and P.

The canonical divergence is related to the *Kullback-Leibler* divergence, which is a commonly used measure of the difference between probability distributions [23]:

$$D(P||Q) = D_{KL}(Q||P) = \int_\Lambda \hat{p}_Q(\lambda) \log\left(\frac{\hat{p}_Q(\lambda)}{\hat{p}_P(\lambda)}\right) \mathrm{d}\lambda \tag{24}$$

From equation (11) we see that the canonical divergence can be rewritten as:

$$D(P||Q) = \phi(Q) - \phi(P) + \boldsymbol{\theta}_P \cdot (\boldsymbol{\eta}_P - \boldsymbol{\eta}_Q) \tag{25}$$

From this equation and equation (15) relating ϕ to the entropy, we can see that the canonical divergence can be expressed in terms of the entropy as follows:

$$D(P||Q) = H(P) - H(Q) + \boldsymbol{\theta}_P \cdot (\boldsymbol{\eta}_P - \boldsymbol{\eta}_Q) \tag{26}$$

An important special case is that of the canonical divergence between an arbitrary spectral model Q and the origin in $\boldsymbol{\theta}$ coordinates:

$$D_0(Q) = D(0||Q) = H(0) - H(Q) \tag{27}$$

The origin in $\boldsymbol{\theta}$ coordinates corresponds to the uniform, or "white", spectral profile, and is the profile with the highest possible entropy. Thus the origin has a special physical significance.

4. SPECTRAL MODELING IN THE PRESENCE OF NOISE

The measurement provided by the photosensors will, in practice, be corrupted by noise. If we assume an additive noise process the measurement equation becomes

$$\boldsymbol{r} = \beta \int_{\Lambda} \boldsymbol{f}(\lambda) p(\lambda) \mathrm{d}\lambda + \boldsymbol{\nu} \tag{28}$$

The normalized measurement $\boldsymbol{\eta}$ is given by

$$\boldsymbol{\eta} = \int_{\Lambda} \boldsymbol{f}(\lambda) p(\lambda) \mathrm{d}\lambda + \boldsymbol{\nu}/\beta \tag{29}$$

The expectation coordinates $\hat{\boldsymbol{\eta}}$ associated with a model spectrum are:

$$\hat{\boldsymbol{\eta}} = \int_{\Lambda} \boldsymbol{f}(\lambda) \hat{p}(\lambda) \mathrm{d}\lambda \tag{30}$$

Note that $\hat{\boldsymbol{\eta}}$ is no longer equal to $\boldsymbol{\eta}$ in general, so we will now have to distinguish these whenever there is non-zero measurement noise.

The metamer set is broadened in the presence of noise, since there is a greater number of possible spectra that could give rise to the observed measurement. As the noise level increases, so does the volume of the metamer set.

We can still apply our maximum entropy spectral modeling approach. Let us assume that the noise is zero-mean and bounded such that

$$\boldsymbol{\nu}^T Q \boldsymbol{\nu} \leq 1 \tag{31}$$

where Q is a positive definite matrix. This constraint implies that the noise is bounded within an ellipsoid whose principal axes are given by the eigenvectors of Q, and with lengths given by the eigenvalues of Q. Let $\hat{\boldsymbol{\eta}}$ be the estimate of the normalized measurement that would be

obtained if there was no noise (we will informally refer to this as the "noise-free normalized measurement"). The actual measurement is related to the noise-free measurement by

$$\beta(\boldsymbol{\eta} - \hat{\boldsymbol{\eta}}) = \boldsymbol{\nu} \tag{32}$$

Thus the bounded noise constraint can be expressed in terms of the measurements as follows:

$$(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta})^T Q(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}) \leq 1/\beta^2 \tag{33}$$

Now we can formulate the spectral modeling problem as finding the spectrum which has maximum entropy subject to the constraint of equation (33). Let the set of $\hat{\boldsymbol{\eta}}$ in the space of admissible measurements, $\mathcal{M}$, that satisfy this constraint be referred to as $\mathcal{Z}(\boldsymbol{\eta})$. Then, for each $\hat{\boldsymbol{\eta}} \in \mathcal{Z}(\boldsymbol{\eta})$, we can find a maximum entropy spectral model $\hat{\boldsymbol{\theta}}$ consistent with $\hat{\boldsymbol{\eta}}$. We can then find the $\hat{\boldsymbol{\eta}}$ which has the corresponding spectral model with highest entropy over all $\mathcal{Z}(\boldsymbol{\eta})$.

$$\hat{\boldsymbol{\eta}} = \arg \max_{\hat{\boldsymbol{\eta}} \in \mathcal{Z}(\boldsymbol{\eta})} H(\hat{\boldsymbol{\theta}}(\hat{\boldsymbol{\eta}})) \tag{34}$$

where $H(\hat{\boldsymbol{\theta}}(\hat{\boldsymbol{\eta}}))$ is the entropy of the maximum entropy model associated with the estimate of the noise-free normalized measurement $\hat{\boldsymbol{\eta}}$.

Let us define $\boldsymbol{\eta}_0$ to be the normalized measurement associated with the uniform, or white, spectral model:

$$\boldsymbol{\eta}_0 = \frac{1}{\Lambda} \int_\Lambda \boldsymbol{f}(\lambda) \mathrm{d}\lambda \tag{35}$$

This spectrum is the one with globally maximum entropy. A simple test of whether $(\boldsymbol{\eta} - \boldsymbol{\eta}_0)^T Q(\boldsymbol{\eta} - \boldsymbol{\eta}_0) \leq 1/\beta^2$ suffices to determine whether $\boldsymbol{\eta}_0 \in \mathcal{Z}(\boldsymbol{\eta})$. If this test returns true then $\hat{\boldsymbol{\eta}} = \boldsymbol{\eta}_0$. If this is not the case, then $\hat{\boldsymbol{\eta}}$ must lie on the boundary of $\mathcal{Z}(\boldsymbol{\eta})$.

The solution detailed above assumes that the noise is bounded. Noise produced by practical noise sources can be unbounded, however (such as in the case of Gaussian noise). For such

unbounded noise processes this method will always yield $\boldsymbol{\eta}_0$ as the solution, since there will always be a possible (however improbable) noise value which, when added to $\boldsymbol{\eta}_0$, will result in the observed measurement. This is clearly unacceptable as it means that every surface would be perceived as a shade of gray. To avoid this problem one can instead employ an Bayesian method where the solution is that which maximizes a posterior probability incorporating a likelihood measure capturing the consistency of the model with the given measurement $\boldsymbol{\eta}$ and a prior based on the entropy of the model. In this case the estimate of the noise-free normalized measurement can be obtained as:

$$\hat{\boldsymbol{\eta}} = \arg\min_{\hat{\boldsymbol{\eta}}} \left\{ (\hat{\boldsymbol{\eta}} - \boldsymbol{\eta})^T A (\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}) - \gamma H(\hat{\boldsymbol{\theta}}(\hat{\boldsymbol{\eta}})) \right\} \tag{36}$$

where A is positive definite and γ is a constant trading off adherence to the prior as compared to consistency with the measurement. If $\gamma = 0$ then the estimate of the noise-free normalized measurement will simply be the given measurement. As γ gets larger the estimate will move towards $\boldsymbol{\eta}_0$. This approach is identical to maximum entropy methods used in Quantum Physics to estimate spectra [24].

As pointed out by Jaynes [16] the Bayesian method is equivalent to the bounded noise constrained optimization method. To see this, consider the solution, $\hat{\boldsymbol{\eta}}_1$, of the Bayesian method in a particular case for some $0 < \gamma < \infty$. One can always find a scalar q_1 such that $(\hat{\boldsymbol{\eta}}_1 - \boldsymbol{\eta})^T [q_1 A](\hat{\boldsymbol{\eta}}_1 - \boldsymbol{\eta}) = 1/\beta^2$. The data consistency term $(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta})^T A (\hat{\boldsymbol{\eta}} - \boldsymbol{\eta})$ is constant $(= 1/q_1\beta^2)$ for all points $\hat{\boldsymbol{\eta}}$ for which $(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta})^T [q_1 A](\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}) = 1/\beta^2$. Thus the solution $\hat{\boldsymbol{\eta}}_1$ to equation (36) has the maximum entropy for all such points (otherwise the optimization would have found a point with higher entropy). This means that $\hat{\boldsymbol{\eta}}_1$ is the same as the solution of the constrained minimization problem (eq. 34) where we interpret the noise

as being bounded with $Q = q_1 A$. We can consider $q_1 A$ as setting an *effective noise level*. Note that the effective noise level depends on the value of the trade-off factor γ that is used. Tikhonov and Arsenin [25], in their presentation of the *regularization* approach to solution of ill-posed problems described the converse procedure, whereby they show that the unconstrained optimization solution is equivalent to the constrained optimization solution. In their treatment γ is a Lagrange multiplier whose value is determined so as to make the solution satisfy the constraint $(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta})^T [q_1 A](\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}) = 1/\beta^2$.

Higher effective noise levels imply larger values of γ, and lead to solutions with higher entropy. In the limit as the noise level goes to infinity, the model spectrum goes to the uniform, or white, spectrum. For any non-zero noise level the entropy of the model spectrum will be finite so that one will never obtain an impulse monochromatic spectrum as a solution. This also means that the magnitude of $\boldsymbol{\theta}$ will be bounded, with lower bounds for higher noise levels.

5. DETERMINATION OF STIMULUS INTENSITY

In the preceding text we had left the intensity value β unspecified (or took it to be equal to 1). Now we will discuss how the value of β can be determined. Recall that $\boldsymbol{\eta} = \boldsymbol{r}/\beta$ where $\mathbf{r}$ is the measurement vector and β is the stimulus intensity. It should be noted that β is the intensity of the incident light only over the visible range (i.e. over the range of wavelengths for which the photoreceptors have non-zero sensitivity). The algorithm does not consider at all the intensity of the light outside of the visible range.

A photoreceptor with a constant sensitivity profile $f_0(\lambda) = 1$ can be seen to yield the value

of β:

$$r_0 = \beta \int_\Lambda f_0(\lambda) p(\lambda) \mathrm{d}\lambda = \beta \int_\Lambda p(\lambda) \mathrm{d}\lambda = \beta \tag{37}$$

The human visual system, at least, does not appear to be equipped with a photoreceptor with such a sensitivity profile, however, so we are left with the question of how the intensity can be determined. One possible approach to determination of the intensity is to extend the maximum entropy paradigm. To see how this might work, consider that a given measurement vector $\boldsymbol{r}$ defines a metamer set (possibly expanded by measurement noise) in $\mathcal{S}$. The maximum entropy spectral modeling approach then selects a particular element of this metamer set. But the particular element that is chosen depends on the value of the intensity, β. We could, therefore, determine which of these spectral models has the highest entropy.

For each value of β finding the maximum entropy solution amounts to finding the point $\hat{\boldsymbol{\eta}}$ on the surface of an ellipsoid with axes defined by the eigenvectors and eigenvalues of $\beta^2 Q$, centered at the normalized measurement $\boldsymbol{r}/\beta$. As β varies, these ellipsoids will change in size, with larger ellipsoids corresponding to smaller values of β, and will shift their centers to different points along the m-geodesic linking the origin and $\boldsymbol{r}$. As long as $\boldsymbol{r}^T Q \boldsymbol{r} > 1$ these ellipsoids will trace out a truncated solid cone with elliptical cross-section, with center-line given by the *m-geodesic* passing through the points $\boldsymbol{\eta} = \boldsymbol{r}$ (corresponding to $\beta = 1$) and $\boldsymbol{\eta} = \mathbf{0}$ (corresponding to $\beta = \infty$), and with the apex at the origin. Let us call this truncated cone the *intensity cone*. Each spectral model defined by points in the intensity cone has an associated entropy value, and the maximum entropy spectral modeling approach determines the model from the $\hat{\boldsymbol{\eta}}$ in the cone which has the maximum entropy. Figure 2 shows graphically the intensity cone intersecting with the set of allowable normalized measurements.

Note that if $\boldsymbol{r}^T Q \boldsymbol{r} \leq 1$ then the ellipsoids for different β values will all be contained within each other and each will contain the origin. This means that in this case the cones open up and the solution space becomes the entire admissible space $\mathcal{M}$ and the solution is therefore the white solution $\boldsymbol{\eta}_0$ (since this has maximum entropy over $\mathcal{M}$). It is interesting to speculate that the reason that colors are perceived as shades of gray in low light situations is not because the color channels are "turned off" but rather that the *signal-to-noise ratio* $||\boldsymbol{r}||/||\boldsymbol{\nu}||$ is so low that the spectral model is always chosen to be the white spectrum.

The corresponding Bayesian, unconstrained, solution is now given by the following modification of equation (36):

$$\{\hat{\boldsymbol{\theta}}, \beta\} = \arg\min_{\boldsymbol{\theta},\beta} \left\{ \left(\hat{\boldsymbol{\eta}}(\boldsymbol{\theta}) - \frac{\boldsymbol{r}}{\beta} \right)^T A \left(\hat{\boldsymbol{\eta}}(\boldsymbol{\theta}) - \frac{\boldsymbol{r}}{\beta} \right) - \gamma H(\hat{\boldsymbol{\theta}}) \right\} \tag{38}$$

Note that here we perform the search in the $\boldsymbol{\theta}$ space. The normalized measurement $\hat{\boldsymbol{\eta}}$ corresponding to the resulting $\hat{\boldsymbol{\theta}}$ is given by equation (8). The maximum entropy spectral model $\hat{p}(\lambda)$ is given by equation (4).

5.A. Implementation

Unlike most empirical color appearance models, except for a few special cases, our approach does not provide analytic formulae for the model parameters in terms of the photoreceptor signals. Instead, numerical solution processes must be employed to compute the model parameters. In this section we outline the process for computing the model.

- Given the measurement vector $\boldsymbol{r}$, solve the minimization problem of equation (38), which provides the spectral model coordinates $\hat{\boldsymbol{\theta}}$ and an estimate of the intensity β.
  - In carrying out the minimization the quantity $\hat{\boldsymbol{\eta}}$ is computed from the candi-

date model $\hat{\boldsymbol{\theta}}$ using equation (8). It is assumed that the photoreceptor sensitivity functions $\boldsymbol{f}(\lambda)$ are given.

– The entropy $H(\hat{\boldsymbol{\theta}})$ is computed using equation (14).

- Once the spectral model parameter vector $\hat{\boldsymbol{\theta}}$ is found, the canonical divergence can be computed using equation (26) (or in the special case where one point is the origin, equation (27)).

In all of the computations described in subsequent sections we take $A = I$, the identity matrix, implying isotropic noise. The Matlab function *fminsearch* is used to do the minimization. Because of the presence of local minima it is necessary to run the *fminsearch* function many times with different initial starting points. It is found that using 50 different starting points randomly selected was sufficient to handle the most difficult cases (cases in which the entropy was low - e.g. nearly monochromatic spectra).

The photoreceptor sensitivity profiles used in the examples are the L,M,S cone 2 degree data on the human retina obtained from Stockman *et al* [26].

6. COLORFULNESS AND SATURATION

Now that we have the spectral modeling theory defined and the information geometric tools at hand, we can evaluate how well the maximum entropy spectral modeling theory accounts for the empirical color appearance data. In this section we will first investigate the ability of the approach to predict the color appearance values of colorfulness and saturation.

Modern color appearance models have relied on nonlinear transformations of the cone responses to specify the coordinates of perceptual color space (e.g. Luv, Lab, and CIECAM02

[27, 28]). Unfortunately, the form of these nonlinear transformations have been determined in a mainly empirical fashion. There is, as yet, no derivation from first physical principles that produces the empirical color spaces most closely describing human color vision. In this section we will show that the maximum entropy spectral modeling view of color can predict the color appearance values of colorfulness quite well without the need for any empirical curve fitting.

We begin by considering how the entropy of the model spectra varies in $\boldsymbol{\theta}$ space. Along an e-geodesic emanating from the origin passing through any other point $\boldsymbol{\theta}_0$, we have that $\boldsymbol{\theta} = t\boldsymbol{\theta}_0$. The entropy along the geodesic is

$$H(t) = \psi(t\boldsymbol{\theta}_0) - t\boldsymbol{\theta}_0 \cdot \boldsymbol{\eta}(t) \tag{39}$$

The derivative of the entropy with respect to t is given by:

$$\frac{\mathrm{d}H(t)}{\mathrm{d}t} = \frac{\partial\psi(t\boldsymbol{\theta}_0)}{\partial t} - \frac{\partial(t\boldsymbol{\theta}_0)}{\partial t} \cdot \boldsymbol{\eta}(t) - (t\boldsymbol{\theta}_0) \cdot \frac{\partial\boldsymbol{\eta}(t)}{\partial t} \tag{40}$$

$$\frac{\mathrm{d}H(t)}{\mathrm{d}t} = \frac{\partial\psi(t\boldsymbol{\theta}_0)}{\partial(t\boldsymbol{\theta}_0)} \cdot \frac{\partial(t\boldsymbol{\theta}_0)}{\partial t} - \frac{\partial(t\boldsymbol{\theta}_0)}{\partial t} \cdot \boldsymbol{\eta}(t) - (t\boldsymbol{\theta}_0)^T \frac{\partial\boldsymbol{\eta}(t)}{\partial\boldsymbol{\theta}(t)}\boldsymbol{\theta}_0 = -t\boldsymbol{\theta}_0^T G(t)\boldsymbol{\theta}_0 \tag{41}$$

where $G(t)$ is the Fisher information matrix (metric tensor). Since the matrix $G(t)$ is positive-definite, we have that $\mathrm{d}H(t)/\mathrm{d}t < 0$ for all $t > 0$. Thus the entropy is strictly decreasing along an e-geodesic emanating from the origin, and the canonical divergence with the origin is strictly increasing. The iso-entropy surfaces (or the iso-divergence surfaces) are nested, with surfaces of higher entropy completely enclosed by surfaces of lower entropy. Thus either the entropy or the canonical divergence can be used as a radial coordinate in $\boldsymbol{\theta}$ space.

From a physical perspective the square root of the canonical divergence between the model spectral profile and the white spectral profile is a measure of the distance between these two

profiles. Thus the divergence can be considered to be a measure of the *purity* of the stimulus color. In color science color purity is quantitatively captured by three different values - *colorfulness, chroma*, and *saturation* [29]. Colorfulness refers to the perceived amount of a hue in the color of a stimulus. Chroma is the colorfulness judged relative to the brightness of a similarly illuminated white surface. Saturation is the colorfulness of a stimulus relative to the brightness of a stimulus. The colorfulness of a stimulus source depends on its brightness; darker surfaces appear less colorful than bright surfaces. Saturation, on the other hand, is invariant to the brightness of the surface. The CIE definition of saturation is as a function of the ratio of the colorfulness to the brightness. In the CIECAM02 appearance model [28] this function is a square root. Hunt [29] suggests that in the case of viewing surfaces under constant illumination chroma is the appropriate appearance quantity.

We saw earlier that as the stimulus intensity β falls the spectral model moves closer to the white point. This implies that the canonical divergence distance $D_0^{1/2}$ for the spectral model decreases as the intensity decreases. Thus we could hypothesize that this distance corresponds to the *colorfulness* of the stimulus. That is, the canonical divergence distance defines the following colorfulness measure:

$$C_D(\boldsymbol{\theta}) = D_0^{1/2}(\boldsymbol{\theta}) = [H(0) - H(\boldsymbol{\theta})]^{1/2} \tag{42}$$

It should be reiterated here that, in our theory, the reason for the decrease in colorfulness as stimulus intensity decreases has nothing to do with the color channels shutting off or becoming inoperative. Instead, what is happening is that the noise (assuming that the noise level is constant, independent of intensity) is becoming large relative to the signal being estimated. This expands the size of the metamer set, bringing in elements that are closer

to the white point. Eventually this set will expand sufficiently so as to include the white point, at which stage all colors will appear as shades of gray. Lights that have a yellowish hue have spectra that closest to the white point and will therefore become perceived gray first as intensity reduces, followed by the other hues.

6.A. Comparisons of the Theory with Experimental Data

Any theory of color vision must be in at least qualitative agreement with experimental observations. There exists experimental data on color appearance, specifically colorfulness, that could potentially be used to test our model predictions. The two most extensive experimental appearance datasets are from the LUTCHI [30] and the more recent UCL [31] studies.

However, it must be understood that our model as it stands is limited in that it does NOT consider receptor nonlinearities, chromatic and luminance adaptation, nor the effects of surround. It can be thought of as modeling a color appearance computation taking place downstream of any adaptation processes. The LUTCHI and UCL datasets were (purposefully) compiled under a wide range of luminance and surround conditions. Therefore the effects of receptor nonlinearities and adaptation are significant and must be taken into account when doing any model comparisons. As our model at this point does not take these effects into account its predictions cannot be directly compared with the LUTCHI and UCL data.

In the absence of the ability to compare our model predictions to the LUTCHI and UCL data, we employ the next-best thing, and that is to compare our predictions to the predictions of another model - one that is generally accepted as providing good matching to human data. The idea is that the other model will stand in as a proxy for the human data. We choose

to use the CIECAM02 [28] model as it is considered by some as one of the best empirical human appearance models.

We computed the maximum entropy spectral model (using the process outlined in section 5.A) for the matte *Munsell* patches and compared our theoretical colorfulness measure to that provided by the empirical CIECAM02 appearance model. The Munsell Book of Color [32] is a collection of color patches which uniformly occupy a large gamut in perceptual space. Parkkinen *et al* [33] measured the spectral reflectance of the patches in the Munsell book. These reflectances can be considered as spectral power distributions through the artifice of assuming that the surfaces are illuminated by an equal-energy (EE) illuminant. We use this collection of measured spectra as a set of light power spectra representative of those experienced by human observers in every-day life.

Figure 3 shows the spectral models produced by our procedure for a representative sample from the Parkkinen Munsell spectra database. The models are seen to be rather smooth, but fairly close approximations to the actual spectra.

Figure 4 shows the canonical divergence distance measure C_D of equation (42) plotted against the empirical chroma measure provided by the CIECAM02 color appearance model for all of the 1250 Munsell patch spectra in the Parkkinen database. The canonical divergences and intensity estimates were determined from the Munsell spectra using the optimization process described in section 5.A. The CIECAM02 chroma values were derived from CIE Tristimulus values obtained by projecting the Munsell spectra onto the CIE standard observer sensitivity curves. The equations described in [28] were used to compute the chroma values, under the assumption of equi-energy illumination, with complete adaptation to the white point. It was also assumed that the background has the same luminance as the patch.

Under these conditions the colorfulness and the chroma are equivalent. A γ value of 0.0005 was used in computing the maximum entropy spectral model. Note the very linear relationship between the theoretical distance values and the empirical chroma values.

The intensity estimates provided by the maximum entropy approach for the Munsell patch spectra are shown in figure 5, plotted against the true intensities. It can be seen that the relation is very linear, with the intensity estimate being higher than the true intensity by a factor of about 1.1. This over-estimation does not imply a failing of the algorithm. Rather, it is an expected by-product of the spectral modeling approach, which selects a metameric spectral profile different than the true Munsell patch spectrum. The increased intensity is needed for the model spectra to match the given photoreceptor measurements.

There are some psychophysical observations related to perceived saturation with which we can do a direct comparison with our theory. For example, it is well known that some pure monochromatic stimuli (very narrow band spectra) appear more saturated than others. Blue and red appear more saturated than yellow. To see if our theory can replicate this observation, we computed the canonical divergence colorfulness measure as a function of monochromatic stimulus wavelength. The results are shown in figure 6. Four curves are shown, two theoretical curves giving the canonical divergence colorfulness corresponding to two different γ values representing different effective noise levels, and two experimental curves taken from the study of Uchikawa *et al* [34]. It is evident that increasing γ (equivalent to reducing intensity or increasing the noise level) increases the entropy of the model spectrum and therefore reduces the saturation measure. It is also clear from the figure that certain wavelengths appear more saturated than others. In particular the saturation is highest around 420, 540, and 670 nm (or violet, green, and red) and lowest around 480 and 600 nm (cyan and yellow). These

theoretical results are quite similar to the experimental data reported by Uchikawa *et al*, as well as that of Kulp and Fuld ([35] - figures 8a, 9a, 10a, 11a). It must be pointed out that, although the shape of our theoretical curves are similar to those of the experimental data, to get the shapes to line up requires a shift in wavelength of about 20nm (in figure 6 the experimental curves have been shifted to longer wavelengths by 20nm).

6.B. An Analytic Expression for the Colorfulness of Monochromatic Stimuli as a Function of Wavelength

We can derive an analytical expression for the canonical divergence in the noise-free case for very narrow monochromatic pulses. This analysis will give us some insight as to why the saturation vs. wavelength curve has the shape it does. In the monochromatic case when the noise level is very low the magnitude of the $\boldsymbol{\theta}$ vector will be very large. Let λ_0 be the wavelength of the monochromatic impulse. The canonical divergence is then:

$$D(0|\lambda_0) = H(0) - H(\boldsymbol{\theta}(\lambda_0)) \tag{43}$$

This can be rewritten as

$$D(0||\lambda_0) = \psi(0) - \psi(\boldsymbol{\theta}(\lambda_0)) + \boldsymbol{\theta}(\lambda_0) \cdot \boldsymbol{\eta} \tag{44}$$

Since the stimulus is assumed to be an impulse at $\lambda = \lambda_0$ we have that $\boldsymbol{\eta} = \boldsymbol{f}(\lambda_0)$. Thus

$$D(0||\lambda_0) = \psi(0) - \psi(\boldsymbol{\theta}(\lambda_0)) + \boldsymbol{\theta}(\lambda_0) \cdot \boldsymbol{f}(\lambda_0) \tag{45}$$

From equation (7) we have the functional definition of ψ:

$$\psi(\boldsymbol{\theta}(\lambda_0)) = \log \int_\Lambda \exp(\boldsymbol{\theta}(\lambda_0) \cdot \boldsymbol{f}) \mathrm{d}\lambda \tag{46}$$

If $\boldsymbol{\theta}$ has a large enough magnitude, we can approximate the integral in the above equation using a *saddle-point approximation*, wherein the integrand is expanded in a Taylors series about the value of λ where $\boldsymbol{\theta}(\lambda_0) \cdot \boldsymbol{f}(\lambda)$ is maximum. If we call this value λ_m then the saddle-point approximation gives:

$$\psi(\boldsymbol{\theta}(\lambda_0)) = \boldsymbol{\theta}(\lambda_0) \cdot \boldsymbol{f}(\lambda_m) + 1/2 \log \left(\frac{-2\pi}{\boldsymbol{\theta}(\lambda_0) \cdot \frac{\mathrm{d}^2 \boldsymbol{f}}{\mathrm{d}\lambda^2}|_{\lambda=\lambda_m}} \right) \tag{47}$$

If the spectral model approaches a monochromatic impulse then the location of the maximum of $\boldsymbol{\theta}(\lambda_0) \cdot \boldsymbol{f}(\lambda)$ will approach $\lambda_m \rightarrow \lambda_0$. Thus we can write:

$$D(0||\lambda_0) \approx \psi(0) - 1/2 \log \left(\frac{-2\pi}{\boldsymbol{\theta}(\lambda_0) \cdot \frac{\mathrm{d}^2 \boldsymbol{f}}{\mathrm{d}\lambda^2}|_{\lambda=\lambda_0}} \right) \tag{48}$$

For large $\boldsymbol{\theta}$ the canonical divergence measure is dominated by the term in the denominator of the log function in the saddle-point approximation of $\psi(\boldsymbol{\theta})$, so that we have:

$$D(0||\lambda_0) \approx 1/2 \log \left(-\boldsymbol{\theta}(\lambda_0) \cdot \frac{\mathrm{d}^2 \boldsymbol{f}}{\mathrm{d}\lambda^2}|_{\lambda=\lambda_0} \right) \tag{49}$$

We can freely perform linear transformations of the $\boldsymbol{f}$ vectors and the $\boldsymbol{\theta}$ coordinates without changing the structure of the spectral color space and without altering distance relationships. We can obtain opponent channel/luminance channel versions of the $\boldsymbol{f}$ sensitivities by the transformations $f_1'(\lambda) = f_1(\lambda) - f_2(\lambda)$, $f_2' = f_3 - (f_1 + f_2)/2$, and $f_3' = (f_1 + f_2)/2$ (corresponding to the (R-G), (B-Y), and (R+G) channels). We can immediately determine the relative saturation levels of the "unique" hues corresponding to the peak (maximum and minimum) sensitivities of these opponent and luminance channels. For example, if $\boldsymbol{\theta}' = ||\boldsymbol{\theta}'||(1, 0, 0)$ where $||\boldsymbol{\theta}'||$ is very large, the corresponding spectral model is a sharp peak located at the positive peak of the R-G channel sensitivity curve (i.e. red). The canonical divergence is seen to be high when the magnitude of the second derivative of the relevant channel sensitivity

curve is high. The blue (S-cone), green (M-cone), and red (L-cone) peaks are narrower than the yellow (L+M) peak and hence have higher second derivative values. Thus the colorfulness for the blue, green, and red spectral models will be higher than that of the yellow model.

7. THE STRUCTURE OF THE SPECTRAL COLOR MANIFOLD

The previous section looked at the distance between a point on the spectral color manifold and the white point. In this section we look at the local structure of the spectral color manifold, expressed by the incremental distance between neighboring points on the manifold.

7.A. Line Elements

Most readers will have observed that when looking at a rainbow, or at the spectral band produced by passing white light through a prism, it is apparent that different colors take up different amounts of space in the band. The yellowish region is relatively thin compared to the greenish or reddish regions. Color scientists have often tackled the problem of understanding the non-uniformity of the CIE chromaticity diagram in terms of the just noticeable differences (JND) in color. The JND is the amount of wavelength shift that is needed for an observer to detect a change in the color of a stimulus. The starting point for most attempts to elucidate the structure of perceptual color space is the specification of a *line element* for the space. The line element is the incremental distance between neighboring points in the space as a function of incremental changes in the coordinates of the space. Vos [36,37] provides excellent reviews of the various line elements that have been put forward. The first specification of a line element was made by Helmholtz [38], as the sum of the squares of the small changes in

the photoreceptor signals weighted by the JNDs of these signals:

$$\mathrm{d}s_H^2 = \left(\frac{\mathrm{d}L}{JND(L)}\right)^2 + \left(\frac{\mathrm{d}M}{JND(M)}\right)^2 + \left(\frac{\mathrm{d}S}{JND(S)}\right)^2 \tag{50}$$

where L, M, S indicate the long, medium, and short wavelength cone photoreceptor signals respectively. He further assumed that the photoreceptor JNDs obeyed Weber's law and were proportional to the photoreceptor signal level (e.g. $JND(L) = \sigma L$). This assumption means that the Helmholtz line element is actually the Euclidean line element, with coordinates $\{\log(L), \log(M), \log(S)\}$. The Helmholtz line element has a serious problem in that making the log space isophotes to be lines orthogonal to the isohue lines resulted in a luminance that did not match the observed result.

Schrodinger [39] proposed, without theoretical justification, a modified version of the Helmholtz line element:

$$\mathrm{d}s_S^2 = \frac{1}{(L+M+S)}\left[\left(\frac{\mathrm{d}L}{\sqrt{L}}\right)^2 + \left(\frac{\mathrm{d}M}{\sqrt{M}}\right)^2 + \left(\frac{\mathrm{d}S}{\sqrt{S}}\right)^2\right] \tag{51}$$

This line element can be seen to be Euclidean in the coordinates $\{\sqrt{L}, \sqrt{M}, \sqrt{S}\}$. In this space the isophote and isohue curves are orthogonal resulting in a luminance very close to the observed. This line element found little use, however, due to its empirical choice of the square root function. Later, Bouman and Walraven [40] derived effectively the same line element (without the (L+M+S) term in the denominator) by applying photon signal detection theory to set $JND(X) = \sqrt{X}$.

Stiles [41] proposed a modification of the Helmholtz line element that consisted of a relative weighting of the L,M, and S components in the ratio of 1:2:16. This weighting was motivated by the observed difference in the relative Weber fractions σ of the cone classes. This slight

but important change greatly improved the matching of the predicted luminance function with the observed function.

Vos and Walraven [42] modified the Bouma-Walraven line element using an extended photon noise model, giving a polynomial functional form for the JND. This allowed a better modeling of the intensity variation of the JND, and also provided a connection to the Stiles line element at normal illuminant levels.

7.B. The Maximum Entropy Spectral Modeling Line Element

In our spectral modeling approach the differential structure of the spectral color manifold is captured by the line element associated with the Fisher metric:

$$\mathrm{d}s^2 = \mathrm{d}\boldsymbol{\eta}^T G^{-1} \mathrm{d}\boldsymbol{\eta} \tag{52}$$

The line elements described in the previous section were all expressed in terms of the photoreceptor signals L, M, S. Taking our measurement as $\beta\boldsymbol{\eta} = \boldsymbol{r} = (L, M, S)$ we can write the Fisher line element as:

$$\mathrm{d}s^2 = \frac{1}{\beta^2}([g^{-1}]_{11}\mathrm{d}L^2 + [g^{-1}]_{22}\mathrm{d}M^2 + [g^{-1}]_{33}\mathrm{d}S^2 + 2[g^{-1}]_{12}\mathrm{d}L\mathrm{d}M + 2[g^{-1}]_{23}\mathrm{d}M\mathrm{d}S + 2[g^{-1}]_{13}\mathrm{d}S\mathrm{d}L) \tag{53}$$

where the $[g^{-1}]_{ij}$ are the elements of the inverse metric G^{-1}. Unlike the previously defined line elements this equation does not assume that the metric is diagonal. In fact the metric provided by the Fisher Information matrix is, in general, not diagonal. Thus there are interactions between the photoreceptor signals. Since the metric matrix is positive definite we can find a coordinate transformation $\boldsymbol{\eta}' = U^T\boldsymbol{\eta}$ which diagonalizes the metric, where U is the matrix of the normalized eigenvectors of G^{-1}, removing the interactions between the

transformed coordinate values.

The form of the maximum entropy spectral manifold line element is rather opaque, and the relation, if any, to the previously proposed line elements is unclear. We can, however, compute the line element for the special case of monochromatic impulse spectra and use this to examine the wavelength discrimination predicted by our approach in comparison to empirical observations. To do this we take the spectral profile of the light incident on the photoreceptors to have the form of an impulse $p(\lambda) = \delta(\lambda - \lambda_0)$. From equation (20) we have that

$$\mathrm{d}s = \left(\frac{\mathrm{d}\boldsymbol{\eta}^T}{\mathrm{d}\lambda_0} G^{-1}(\boldsymbol{\eta}) \frac{\mathrm{d}\boldsymbol{\eta}}{\mathrm{d}\lambda_0} \right)^{1/2} \mathrm{d}\lambda_0 \tag{54}$$

or

$$\frac{\mathrm{d}\lambda_0}{\mathrm{d}s} = \left(\frac{\mathrm{d}\boldsymbol{\eta}^T}{\mathrm{d}\lambda_0} G^{-1}(\boldsymbol{\eta}) \frac{\mathrm{d}\boldsymbol{\eta}}{\mathrm{d}\lambda_0} \right)^{-1/2} \tag{55}$$

where λ_0 is the wavelength of the monochromatic stimulus. It can be seen that, since the stimulus spectral profile is an impulse,

$$\frac{\mathrm{d}\boldsymbol{\eta}}{\mathrm{d}\lambda_0} = \left. \frac{\mathrm{d}\boldsymbol{f}(\lambda)}{\mathrm{d}\lambda} \right|_{\lambda=\lambda_0} \tag{56}$$

which only depends on the shape of the photoreceptor sensitivity curves. Combining this equation with equation (55) yields:

$$\frac{\mathrm{d}\lambda_0}{\mathrm{d}s} = \left(\left. \frac{\mathrm{d}\boldsymbol{f}^T(\lambda)}{\mathrm{d}\lambda} \right|_{\lambda=\lambda_0} G^{-1}(\boldsymbol{\eta}) \left. \frac{\mathrm{d}\boldsymbol{f}(\lambda)}{\mathrm{d}\lambda} \right|_{\lambda=\lambda_0} \right)^{-1/2} \tag{57}$$

The quantity $\mathrm{d}\lambda_0/\mathrm{d}s$ can be taken to be proportional to the JND, $\Delta\lambda$, in wavelength, under the assumption that the minimum perceptible distance Δs is constant over the spectral color manifold. The variation of $\mathrm{d}\lambda_0/\mathrm{d}s$ with stimulus wavelength, computed using equation (57), is shown in figure 7 for three different intensity values (signal-to-noise ratios), computed

using a γ value of 0.0005. The metric is computed using equation (18), where the η_i are taken from the normalized measurements and the η_{ij} are computed by projecting the model spectrum associated with the measurements onto the synthetic photoreceptor sensitivity curves $f_i(\lambda)f_j(\lambda)$.

Compare the theoretical curves of figure 7 to the experimental wavelength discrimination data from McCree [43] for two different light levels, shown replotted in figure 7. The theoretical model provides a wavelength discrimination curve which lies somewhere between the low and high light intensity curves of McCree. In both our theoretical and the experimental curves there is a peak around 520-560 nm (green) and two minima near 470-500 nm (cyan) and 580 nm (yellow). The higher intensity curve of McCree shows significant shifting of these peaks and valleys as compared with the lower intensity curve. The locations of the peaks and valleys of the theoretical curve fall somewhere in between the empirical low and high intensity curves. The theory predicts an additional minimum around 670 nm, a wavelength which the McCree experiments did not cover. The higher intensity experimental data exhibits a significant peak around 460 nm. This latter peak is the so-called König-Dieterici anomaly [37] and may be due to nonlinear photoreceptor or opponent channel behavior, such as saturation. The papers of Vos [36, 37] provide good explanations of the effects and modeling of the photoreceptor response nonlinearities. We have not incorporated these nonlinear effects in our results, in order to focus on the intensity-related nonlinearities inherent in the maximum entropy spectral modeling approach. A more complete model will include the photoreceptor nonlinearities as well as a more detailed noise model.

8. CONCLUSIONS

In this paper we have presented a new theory of color perception, one that is based on "color realism", the philosophical view that that color is spectra. Taking this view we identify the space of color perceptions with the space of spectra. To facilitate the analysis of color space we take the further step of associating spectral profiles with probability density functions. This chain of development results in an identification of the space of color perceptions with the space of probability density functions. We then propose that a particular color perceiving system, such as the human eye and brain, performs *spectral modeling*, which is the selection of a distinct element of the metameric spectra consistent with the spectral measurements provided by the photoreceptors. The spectral modeling process defines a sub-manifold in the space of spectra, equivalent to a statistical model in the space of probability densities. This sub-manifold determines the range of color perceptions that are possible for a specific system. The structure of the spectral color manifold is set by the Fisher information metric associated with the specific statistical model implied by the photoreceptor sensitivity curves and the spectral modeling approach.

In this paper we propose a maximum entropy approach to spectral modeling. Other spectral modeling approaches are possible, such as selecting the average metamer, and may result in similar characteristics. We extend the maximum entropy method to handle situations in which the measurements are noisy, and also use it to provide an estimate for the intensity of the spectrum. The principal advantages of the maximum entropy spectral modeling technique are its wide applicability, as many physical processes generate high entropy stimuli, and the mathematical convenience of the resulting exponential models of the spectra. The

latter means that many quantities such as distances are straightforward to compute using the mathematical machinery provided by Information Geometry. One of the main results is the formula for distances between points on the spectral color manifold. This provides a theoretical basis for the computation of the *colorfulness* of a stimuli as the distance between the spectral model corresponding to the measurement induced by the stimulus and the uniform, or white, spectral model. This distance measure is shown in the paper to give theoretical predictions of colorfulness that qualitatively match empirical observations. In addition, the theory provides an explicit formula for the line element of the spectral color manifold. This provides theoretical predictions of wavelength discrimination curves which also are in agreement with empirical observations. In our theory the presence of measurement noise causes the metamer set to expand. In combination with the maximum entropy formalism, this means that, as the signal-to-noise level decreases, the resulting spectral model moves closer to the white point. This provides an explanation of why colorfulness decreases when stimulus intensity decreases.

It should be noted that the mathematical formalism presented in this paper is general in that it can be used for any set of photoreceptors, whether it be the three wavelength selective channels of the human trichromat, the two channels of the human dichromat, the six channels of birds, or even the sixteen channels of the mantis shrimp [44]. Our approach implies that the color perceptions obtained in these diverse systems can be directly compared. This is, in fact, a byproduct of the physical underpinnings of our theory - assuming that perception of color arises from a physical quantity, namely spectra, different sensory systems provide different models of this single physical quantity. It is possible that different systems may produce the same model in certain situations, in which case we could truly say that these

two observers are perceiving the same color. For example, the blue or yellow perceived by a dichromat could be said to be the same as the blue or yellow perceived by a trichromat. Such a statement is not possible, or is even meaningless under current theories of color perception as there is no common ground to serve as a basis for comparison.

In conclusion, this paper shows that a mathematically and computationally sound theory of color perception can be constructed from the color realist premise that color is identified with spectra, and demonstrates that the predictions of the theory are not at odds with empirical studies.

9. ACKNOWLEDGEMENTS

This work was supported by the Natural Sciences and Engineering Research Council of Canada through a Discovery Grant awarded to the first author. The authors would like to thank the anonymous reviewers for comments which strengthened the paper. They would also like to thank Tal Arbel, Mike Langer, Erik Myin, and Kevin O'Regan for discussions on color vision.

References

1. B. Maund, "Color", in *The Stanford Encyclopedia of Philosophy* (Winter 2008 Edition), Edward N. Zalta (ed.),
   URL = http://plato.stanford.edu/archives/win2008/entries/color/.
2. A. Byrne and D. Hilbert, "Color realism and color science", Behav. Brain Sci., Vol. 26, pp. 3-64 (2003).
3. M. Tye, "The puzzle of true blue", Analysis, Vol. 66, pp. 173-178 (2006).

4. C.L. Hardin, "A spectral reflectance doth not a color make", J. Phil., Vol. 100, No. 4, pp. 191-202 (2003).

5. J.K. O'Regan and A. Noë, "A sensorimotor account of vision and visual consciousness", Behav. Brain Sci., Vol. 24, No. 5, pp. 939-1011 (2001).

6. D.H. Brainard and W.T. Freeman, "Bayesian color constancy", J. Opt. Soc. Am. A Vol. 14, pp. 1393-1411 (1997).

7. L.T. Maloney and B.A. Wandell, "Color constancy: a method for recovering surface spectral reflectance", J. Opt. Soc. Am. A Vol. 3, pp. 29-33 (1986).

8. R.T. Cox, *The Algebra of Probable Inference*, (Johns Hopkins University Press, 1961).

9. P. Morovic, and G.D. Finlayson, "Metamer-set-based approach to estimating surface reflectance from camera RGB", J. Opt. Soc. Am. A Vol. 23, pp. 1814-1822 (2006).

10. D.H. Krantz, "Color Measurement and Color Theory: I. Representation Theorem for Grassmann Structures", J. Math. Psych., Vol. 12, pp. 283-303 (1975)

11. C.H. Graham and Y. Hsia, "Studies of color blindness: a unilaterally dichromatic subject", Proc. Natl. Acad. Sci. USA., Vol. 45(1), pp. 9699 (1959).

12. D.I.A. MacLeod and P. Lennie, "Red-green blindness confined to one eye", Vis. Res., Vol. 16, No. 7, pp. 691-702 (1976).

13. M.A. Webster, E. Miyahara, G. Malkoc, and V.E. Raker, "Variations in normal color vision. II. Unique hues," J. Opt. Soc. Am. A Vol. 17, pp. 1545-1555 (2000).

14. J.J. Clark and S. Skaff, "Maximum entropy models of surface reflectance spectra", Proceedings of the Instrumentation and Measurement Technology Conference, Vol. 2, pp. 1557-1560 (2005).

15. E.T. Jaynes, "Prior probabilities", IEEE Trans. on Systems Science and Cybernetics, Vol. SSC-4, pp. 227-241 (1968).

16. E.T. Jaynes, "On the rationale of maximum-entropy methods", Proceedings of the IEEE, Vol. 70, No. 9, pp. 939-952 (1982).

17. E.T. Jaynes, "Concentration of distributions at entropy maxima", in *E.T. Jaynes: Papers on Probability, Statistics and Statistical Physics*, R. D. Rosenkrantz (ed.), (D. Reidel, 1979), pp. 315-334.

18. S.-I. Amari and H. Nagaoka, *Methods of information geometry*, Transactions of mathematical monographs; v. 191, (American Mathematical Society, 2000).

19. R.A. Fisher, "Theory of statistical estimation", Proc. Cambridge Philos. Soc. Vol. 22, pp. 700-725 (1925).

20. C.R. Rao, "Information and accuracy attainable in the estimation of statistical parameters", Bull. Calcutta Math. Soc., Vol. 37, pp. 81-91 (1945).

21. S.J. Maybank, "The Fisher-Rao metric", Mathematics Today, Vol. 44, No. 6, pp. 255-257 (2008).

22. N.N. Cencov, *Statistical Decisions Rules and Optimal Inference*, Translations of Mathematical Monographs; v. 53, (American Mathematical Society, 1982).

23. S. Kullback and R.A. Leibler, "On information and sufficiency", Ann. Math. Stat., Vol. 22, No. 1, pp. 79-86 (1951).

24. N. Silver, D.S. Sivia, and J.E. Gubernatis, "Maximum-entropy method for analytic continuation of quantum Monte Carlo data", Phys. Rev. B, Vol. 41, pp. 2380-2389 (1990)

25. A.N. Tikhonov and V.Y. Arsenin, *Solutions of Ill-posed Problems*, (Winston & Sons,

1977).

26. A. Stockman, L.T. Sharpe, and C.C. Fach, "The spectral sensitivity of the human short-wavelength cones", Vision Res. Vol. 39, pp. 2901-2927 (1999).

27. D.B. Judd, and G. Wyszecki, *Color in Business, Science, and Industry.* 3rd ed., (John Wiley and Sons, 1975).

28. N. Moroney, M.D. Fairchild, R.W.G. Hunt, C.J. Li, M.R. Luo, and T. Newman, "The CIECAM02 color appearance model," Proceedings of the 10th IS&T/SID conference, pp. 23-27 (2002).

29. R.W.G. Hunt, "Saturation, superfluous or superior?" Proceedings of the 9th IS&T/SID conference (2001).

30. M.R. Luo, A.A. Clarke, P.A. Rhodes, A. Schappo, S.A.R. Scrivner, and C.J. Tait, "Quantifying colour appearance. Part I. LUTCHI colour appearance data", Color Res. Appl., Vol. 16, pp. 166-180 (1991).

31. M.H. Kim, T. Weyrich, and J. Kautz, "Modeling Human Color Perception under Extended Luminance Levels", ACM Transactions on Graphics (Proceedings SIGGRAPH 2009) 28(3), (2009), to appear.

32. A.H. Munsell, *Munsell Book of Color: Matte Finish Collection*, (Munsell Color, 1979).

33. J.P.S. Parkkinen, J. Hallikainen and T. Jaaskelainen, "Characteristic spectra of Munsell colors", J. Opt. Soc. Am. A , Vol. 6, No. 2, pp. 318-322 (1989).

34. K. Uchikawa, H. Uchikawa and P.K. Kaiser, "Luminance and saturation of equality bright colors", Color Res. Appl., Vol. 9, pp. 5-14 (1984).

35. T.D. Kulp and K. Fuld, "The prediction of hue and saturation for non-spectral lights",

Vision Res. , Vol. 35, No. 21, pp. 2967-2983 (1995).

36. J.J. Vos, "Line elements and physiological models of color vision," Color Res. Appl., Vol. 4, pp. 208-216 (1979).

37. J.J. Vos, "From lower to higher colour metrics: a historical account", Clinical & Experimental Optometry, Vol. 86, No. 6, pp. 348-360 (2006).

38. H. Helmholtz, *Handbuch der physiologischen Optik.* Voss, Hamburg, second edition (1896).

39. E. Schrodinger, "Grundlinien einer Theorie der Farbenmetrik im Tagessehen", Ann. Phys. Vol. 368, No. 22, pp. 481-584 (1920).

40. M.A. Bouman and P.L. Walraven, "Quantum theory of colour discrimination of dichromats", Vision Res. Vol. 2, pp. 177-187 (1962).

41. W.S. Stiles, "A modified Helmholtz line element in brightness-colour space". Proceedings of the Physical Society of London, pp. 58-41 (1946).

42. J.J. Vos, P.L. Walraven, "An analytical description of the line element in the zone fluctuation model of color vision, I and II", Vision Res. Vol. 12, pp. 1327-1344 and 1345-1365 (1972).

43. K.J. McCree, "Small field tritanopia and the effects of voluntary fixation", J. Modern Opt., Vol. 7, No. 4, pp. 317-323 (1960).

44. T.W. Cronin, and J. Marshall, "Parallel processing and image analysis in the eyes of mantis shrimps", Biological Bulletin, Vol. 200, No. 2, pp. 177-183 (2001).

List of Figure Captions

Fig. 1. Diagram illustrating the spectral modeling process. The photoreceptor signals define a metamer set, which is the set of all spectral profiles that can result in these signals. One element of the metamer set is distinguished by the spectral modeling approach and is taken as the spectral model. The space traced out by these distinguished elements of the metamer sets is called the spectral color manifold. It is the space of spectral models that are possible given the particular photoreceptor characteristics and spectral modeling approach.

Fig. 2. Diagram illustrating the determination of the normalized measurement estimate given noisy measurements. Given a photoreceptor measurement vector $\boldsymbol{r}$ a normalized measurement is obtained by dividing by the intensity β. The estimate of the noise-free normalized measurement will be located on the boundary of the ellipsoid centered on the normalized measurement with size given by the noise level scaled by the intensity. The intensity is chosen so as to maximize the entropy of the spectral model associated with the estimate of the noise-free normalized measurement.

Fig. 3. Maximum entropy spectral models compared with true spectra for four different Munsell patches. Computed using $\gamma = 0.0005$.

Fig. 4. Comparison between the CIECAM02 empirical chroma (colorfulness) measure and the canonical divergence colorfulness measure of equation (42). The dashed line represents the best linear fit to the data. The canonical divergence was computed using $\gamma = 0.0005$.

Fig. 5. Comparison between the actual Munsell patch intensity and the intensity estimated using the Maximum Entropy approach. The canonical divergence was computed using $\gamma = 0.0005$.

Fig. 6. The square root of the canonical divergence distance measure D_0 for monochromatic stimuli as a function of wavelength. Shown are curves (solid lines) for two different γ values corresponding to two different effective signal to noise ratios (larger γ corresponds to lower intensity). The upper solid line corresponds to $\gamma = 0.00056$ and the lower solid curve to $\gamma = 0.56$. For comparison are shown two curves derived from the experimental from Table I of Uchikawa et al (1984) (subject HU, purity 0.3 (dotted line) and 0.7 (dashed line). The data from Uchikawa et al has been shifted by 20nm to align with the theoretical curve.

Fig. 7. The quantity $d\lambda/ds$ for the maximum entropy spectral model as a function of monochromatic stimulus wavelength, computed using $\gamma = 0.0005$ and $\beta = 1$. Also shown are two curves replotted from McCree's [43] figure 4 showing empirical wavelength discrimination data for two different light levels. The scale for the theoretical curve is arbitrary, set for ease of comparison. It is evident in both the empirical and theoretical curves that wavelength discrimination is greatest in the yellow and cyan regions of the spectrum.

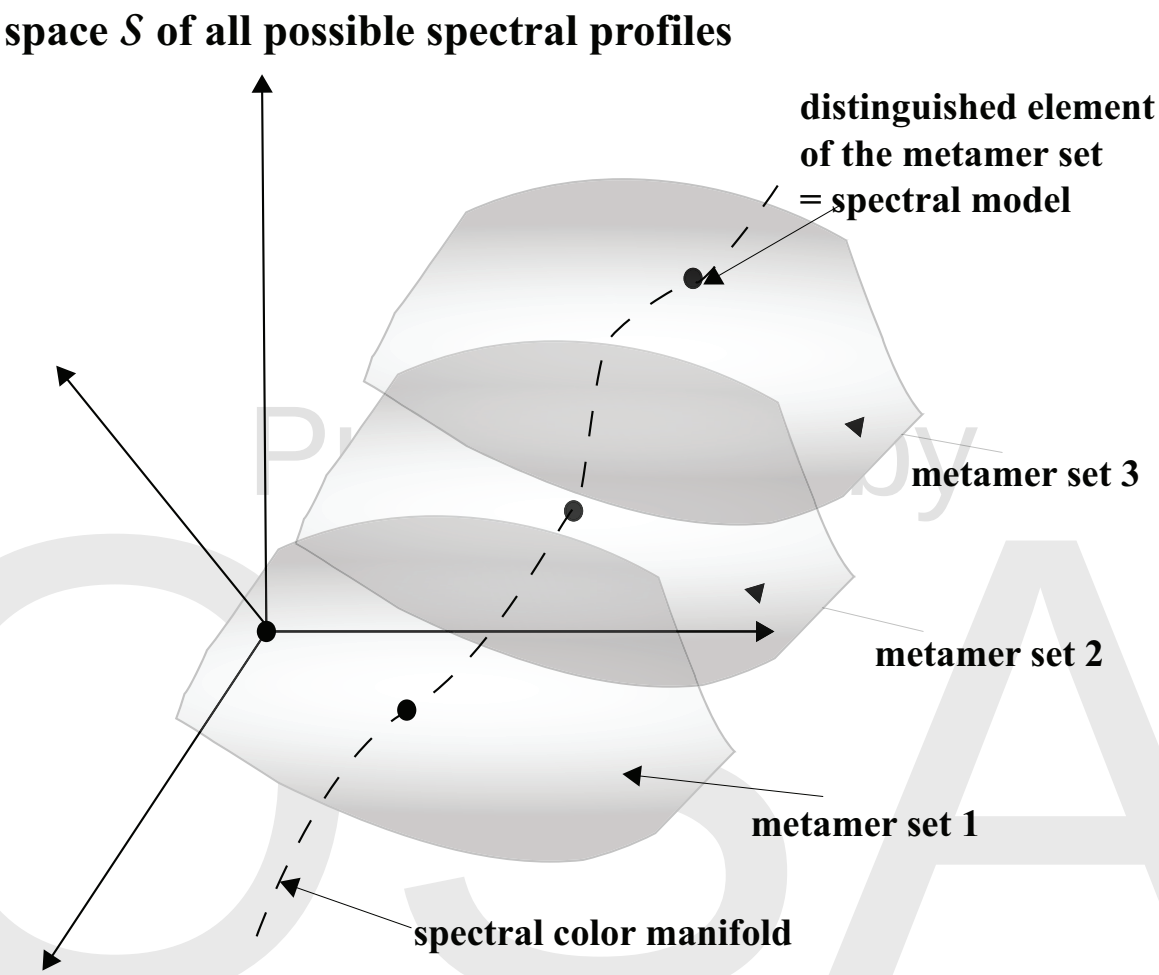

Fig. 1. Diagram illustrating the spectral modeling process. The photoreceptor signals define a metamer set, which is the set of all spectral profiles that can result in these signals. One element of the metamer set is distinguished by the spectral modeling approach and is taken as the spectral model. The space traced out by these distinguished elements of the metamer sets is called the spectral color manifold. It is the space of spectral models that are possible given the particular photoreceptor characteristics and spectral modeling approach. spectral_manifold_diagram.eps

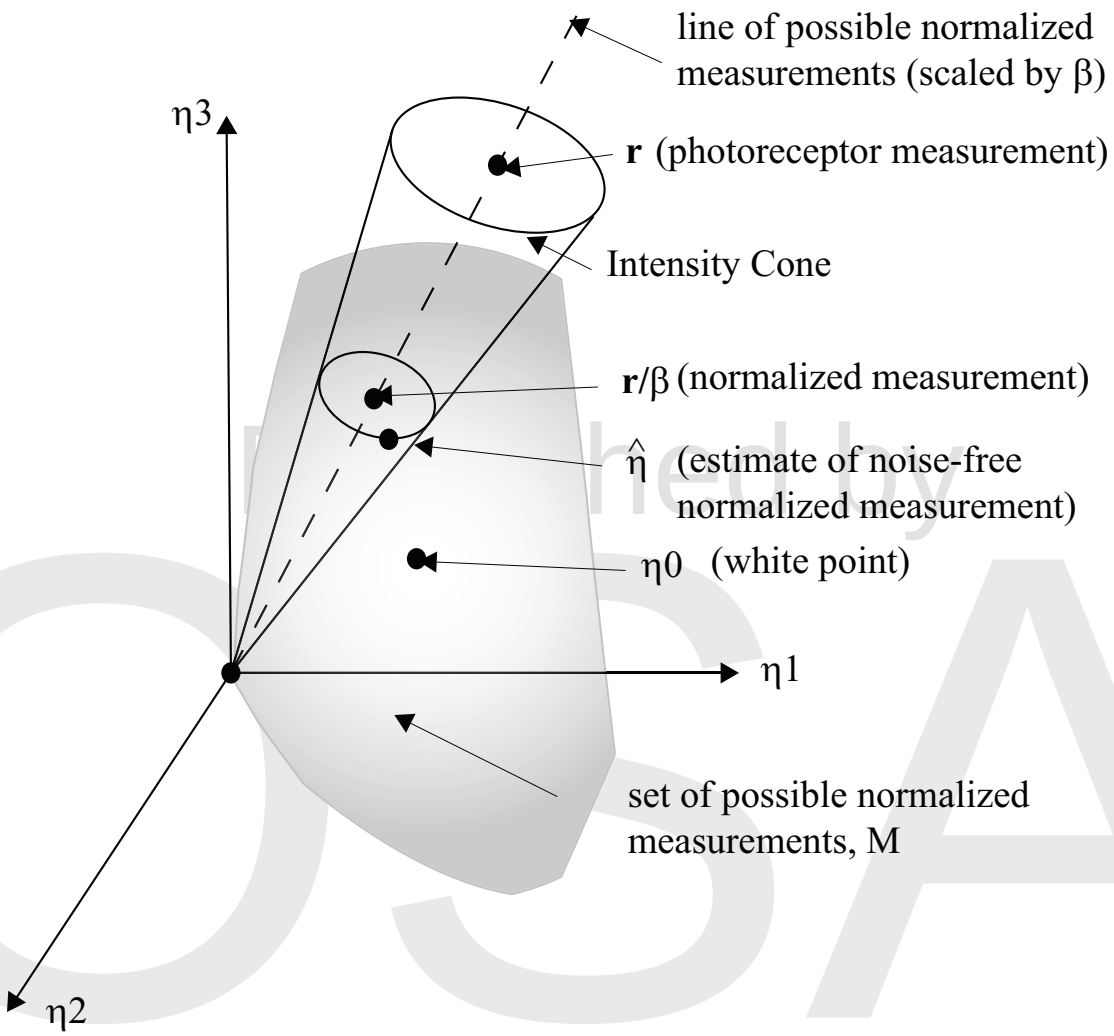

Fig. 2. Diagram illustrating the determination of the normalized measurement estimate given noisy measurements. Given a photoreceptor measurement vector $\boldsymbol{r}$ a normalized measurement is obtained by dividing by the intensity β. The estimate of the noise-free normalized measurement will be located on the boundary of the ellipsoid centered on the normalized measurement with size given by the noise level scaled by the intensity. The intensity is chosen so as to maximize the entropy of the spectral model associated with the estimate of the noise-free normalized measurement. noisy_maxent_diagram.eps

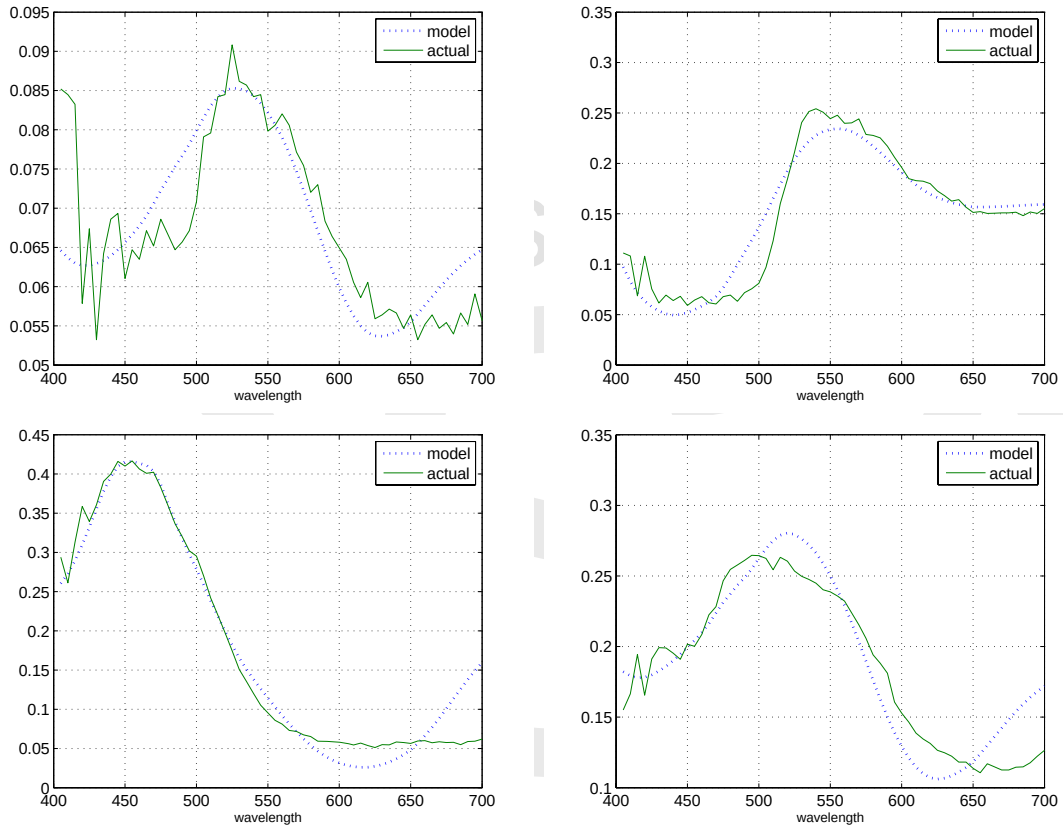

Fig. 3. Maximum entropy spectral models compared with true spectra for four different Munsell patches. Computed using $\gamma = 0.0005$. munsell_model_patch100.eps

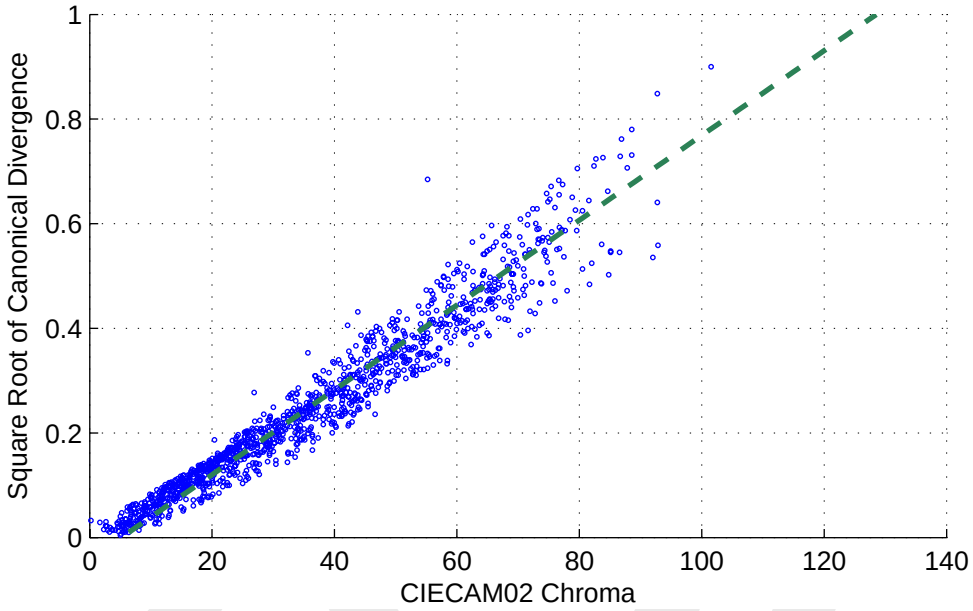

Fig. 4. Comparison between the CIECAM02 empirical chroma (colorfulness) measure and the canonical divergence colorfulness measure of equation (42). The dashed line represents the best linear fit to the data. The canonical divergence was computed using $\gamma = 0.0005$. munsell_ciecam02_divergence_comparison.eps

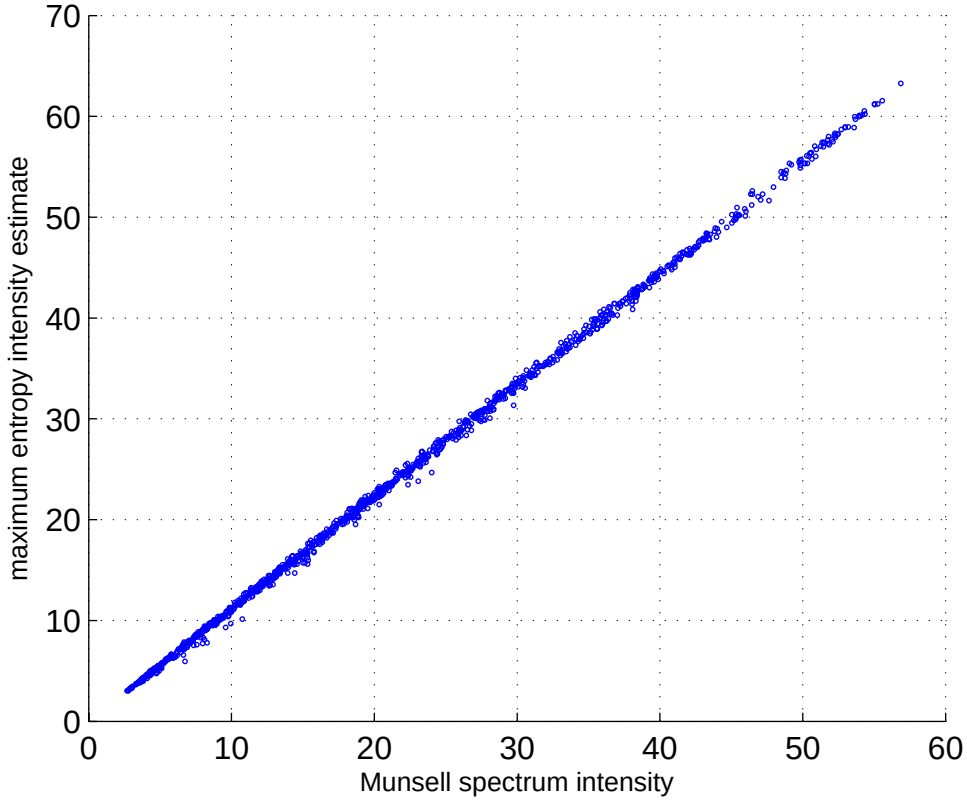

Fig. 5. Comparison between the actual Munsell patch intensity and the intensity estimated using the Maximum Entropy approach. The canonical divergence was computed using $\gamma = 0.0005$. munsell_intensity_estimate.eps

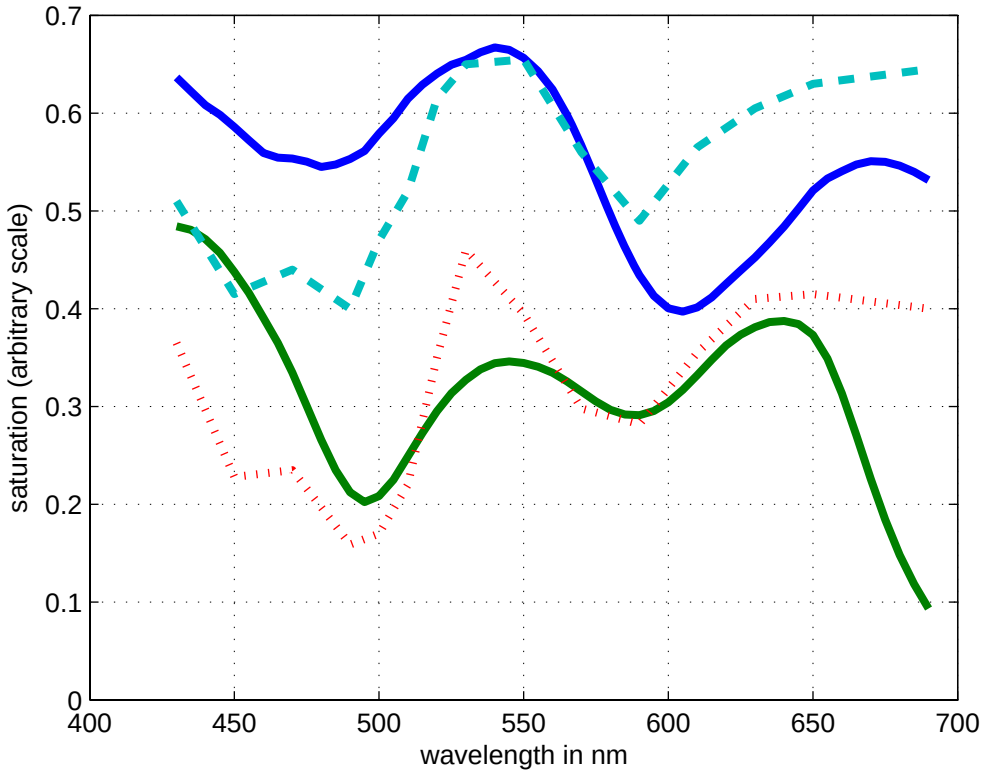

Fig. 6. The square root of the canonical divergence distance measure D_0 for monochromatic stimuli as a function of wavelength. Shown are curves (solid lines) for two different γ values corresponding to two different effective signal to noise ratios (larger γ corresponds to lower intensity). The upper solid line corresponds to $\gamma = 0.00056$ and the lower solid curve to $\gamma = 0.56$. For comparison are shown two curves derived from the experimental from Table I of Uchikawa et al (1984) (subject HU, purity 0.3 (dotted line) and 0.7 (dashed line). The data from Uchikawa et al has been shifted by 20nm to align with the theoretical curve. divergence_saturation_uchikawa.eps

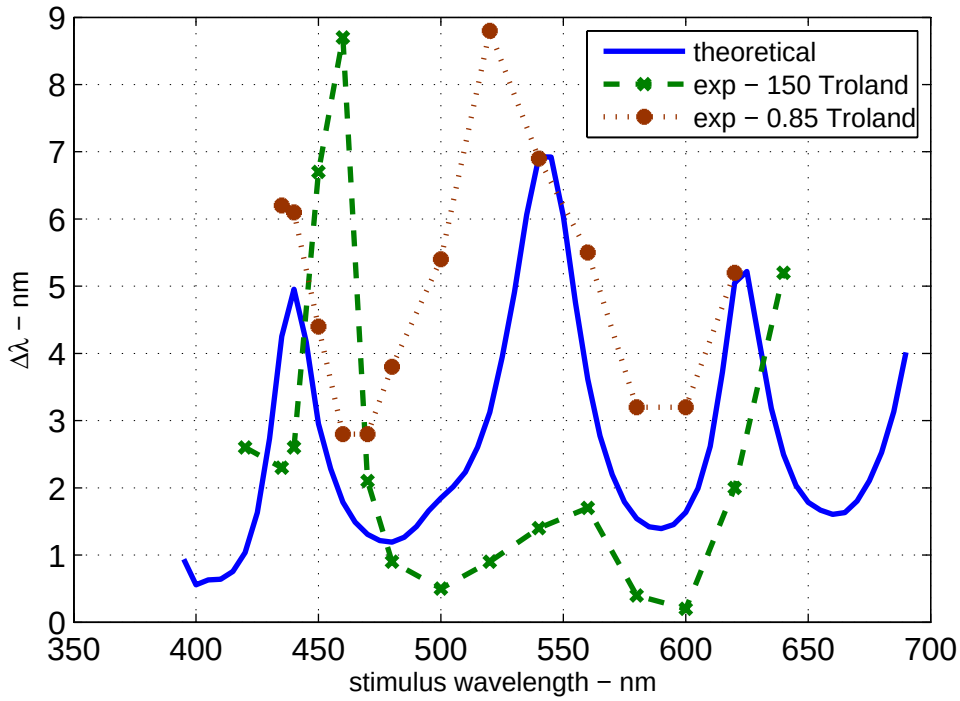

Fig. 7. The quantity $\mathrm{d}\lambda/\mathrm{ds}$ for the maximum entropy spectral model as a function of monochromatic stimulus wavelength, computed using $\gamma = 0.0005$ and $\beta = 1$. Also shown are two curves replotted from McCree's [43] figure 4 showing empirical wavelength discrimination data for two different light levels. The scale for the theoretical curve is arbitrary, set for ease of comparison. It is evident in both the empirical and theoretical curves that wavelength discrimination is greatest in the yellow and cyan regions of the spectrum. line_element_comparison.eps