Human lightness perception is guided by simple assumptions about reflectance and lighting

Richard F. Murray*
Centre for Vision Research and Department of Psychology, York University,
4700 Keele Street, LAS 0009, Toronto, Ontario, Canada, M3J 1P3

ABSTRACT

Lightness constancy is the remarkable ability of human observers to perceive surface reflectance accurately despite variations in illumination and context. Two successful approaches to understanding lightness perception that have developed along independent paths are anchoring theory and Bayesian theories. Anchoring theory is a set of rules that predict lightness percepts under a wide range of conditions. Some of these rules are counterintuitive and difficult to motivate, e.g., a rule that large surfaces tend to look lighter than small surfaces. Bayesian theories are formulated as probabilistic assumptions about lights and objects, and they model percepts as rational inferences from sensory data. Here I reconcile these two seemingly divergent approaches by showing that many rules of anchoring theory follow from simple probabilistic assumptions about lighting and reflectance. I describe a simple Bayesian model that makes maximum a posteriori interpretations of luminance images, and I show that this model predicts many of the phenomena described by anchoring theory, including anchoring to white, scale normalization, and rules governing glow. Thus anchoring theory can be formulated naturally in a Bayesian framework, and this approach shows that many seemingly idiosyncratic properties of human lightness perception are actually rational consequences of simple assumptions about lighting and reflectance.

Keywords: lightness perception, Bayesian modelling, scene statistics, anchoring theory, inverse optics

1. INTRODUCTION

Visual perception is a difficult computational problem in large part because images are ambiguous. One of the clearest examples of this is lightness perception. Human observers are able to estimate the surface reflectance of objects depicted in greyscale images, even though any such image could have been generated by many different combinations of lighting conditions, surface shape, and surface reflectance patterns. How is this possible?

Gilchrist and colleagues[1,2] have developed a successful and wide-ranging anchoring theory of lightness perception and constancy, that accounts for the overall accuracy and systematic errors of human lightness perception under many conditions. Anchoring theory is formulated as a series of rules that predict lightness percepts. Some of these rules seem counterintuitive and difficult to motivate, e.g., one rule states that larger surfaces look lighter than small surfaces, all other things being equal. Nevertheless, these rules predict lightness percepts accurately in many scenes, so they capture important facts about human lightness perception.

Ambiguity is one of the fundamental obstacles to accurate lightness perception, so we might expect Bayesian models to be useful, and some investigators have taken this approach. Adelson[3] has suggested that lightness perception should be seen as the problem of mapping image luminance to surface reflectance, relying on soft constraints provided by knowledge of the statistical distribution of lighting conditions and reflectances in the world. Barron and Malik[4,5] have used similar ideas to develop successful computer vision algorithms for recovering shape and reflectance from single images.

Anchoring theory and Bayesian approaches are based on very different ways of predicting lightness percepts, and they have mostly developed independently. Anchoring theory is not probabilistic, and it may seem unlikely that Bayesian theories can account for the large number of unusual but systematic errors in lightness perception documented by anchoring theory. Here I show that these two approaches actually have much in common, in that that many rules of anchoring theory follow from a Bayesian theory that relies on a few simple and realistic assumptions about statistical properties of lighting and reflectance.

* rfm@yorku.ca; phone 1 416 736 2100 ext. 23025; fax 1 416 736 5857; www.yorku.ca/rfm

Human Vision and Electronic Imaging XVIII, edited by Bernice E. Rogowitz, Thrasyvoulos N. Pappas, Huib de Ridder,
Proc. of SPIE-IS&T Electronic Imaging, SPIE Vol. 8651, 865106 • © 2013 SPIE-IS&T
CCC code: 0277-786X/13/$18 • doi: 10.1117/12.2001782

2. A BAYESIAN MODEL OF GREYSCALE IMAGES

I will describe a generative model of greyscale images. I use the term "scene" to mean a 3D arrangement of surfaces and lights, and "image" to mean the 2D retinal luminance pattern that the scene gives rise to. The model makes the following assumptions about scenes and images.

Assumption 1. Surfaces are greyscale and Lambertian, so they can be described by their reflectance, i.e., the proportion of incident light they reflect. The prior probability of reflectances, $p_{\mathbf{r}}(x)$, peaks at 0.6, falls off smoothly at lower reflectances, and falls off smoothly but more steeply at higher reflectances (Figure 1a). Specifically, the prior probability is an asymmetric normal distribution, smoothly clipped to zero at its lower end via multiplication with a normal cumulative distribution function:

$$p_{\mathbf{r}}(x) = \begin{cases} k \exp\left(-0.5\left(\frac{x-0.6}{0.3}\right)^2\right)\Phi(x, 0.05, 0.025) & x \leq 0.6 \\ k \exp\left(-0.5\left(\frac{x-0.6}{0.2}\right)^2\right) & x > 0.6 \end{cases} \quad (1)$$

Here $\Phi(x,\mu,\sigma)$ is the normal cumulative density function, and k is a normalization constant set to 1.036 to give the probability density a total area of 1.0. This reflectance prior extends above the physical upper limit of 1.0, and this upper range represents glowing surfaces.

The important qualitative properties of this prior are that (a) it drops quickly to zero at very low reflectances, (b) a large middle range of reflectances has moderate probabilities, and (c) very high reflectances and glowing surfaces are unlikely, but possible. I do not assign any importance to the parametric form or parameter values of the priors that I use in this paper. I choose specific priors mostly so that I can show the results of actual fits of the model to various stimuli, but I will also explain how the results follow from qualitative properties of the priors.

The model assumes that a scene is divided into n patches, and reflectances are assigned to patches randomly and independently. The model represents a scene as an $n \times 1$ vector of reflectances $\mathbf{r}$, sampled from an $n \times 1$ vector random variable $\mathbf{R}$. The prior probability of a full vector of reflectances $\mathbf{r}$ is the product of the prior probabilities of its components:

$$P(\mathbf{R} = \mathbf{r}) = \prod_{i=1}^{n} p_{\mathbf{r}}(r_i) \quad (2)$$

The limited available data[6] suggest that the distribution of reflectances in the real world is qualitatively similar to this prior.

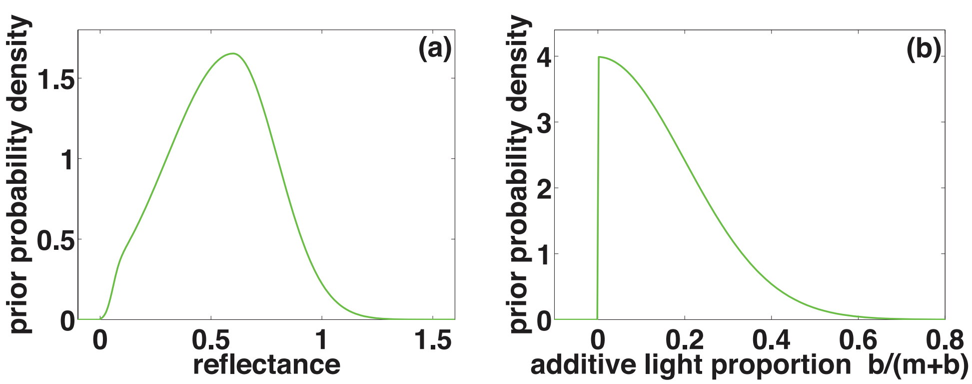

Figure 1. The model's (a) reflectance prior and (b) lighting prior.

Assumption 2. Lighting has a multiplicative component[3] m and an additive component b, and the image luminance l generated by a surface patch of reflectance r is $l = mr + b$. The multiplicative component corresponds to light that is reflected from the surface patch, and the additive component corresponds to veiling light from semi-transparent surfaces or volumes, such as mist, between the surface patch and the viewer. When convenient, I will denote the lighting parameters as λ instead of (m,b). According to this model, image luminance is an affine transformation of reflectance, and the probability density of image luminances l_i under lighting conditions $\lambda = (m,b)$ is a shifted and scaled version of the reflectance prior:

$$p_l(l_i;\lambda) = m^{-1} p_r((l_i - b)/m) \tag{3}$$

The probability density of the full luminance vector $\mathbf{l}$ is

$$P(\mathbf{L} = \mathbf{l} \mid \Lambda = \lambda) = \prod_{i=1}^{n} p_l(l_i;\lambda) \tag{4}$$

$$= m^{-n} \prod_{i=1}^{n} p_r((l_i - b)/m) \tag{5}$$

The model does not incorporate the fact that the luminance of a surface patch depends on the orientation of the patch relative to the light source, so it is a model of flat surfaces.

Assumption 3. The lighting conditions $\lambda = (m,b)$ are sampled from random variables $\Lambda = (M,B)$. The probability density of lighting conditions, p_λ, is a function of $b/(m+b)$ that peaks at zero, falls off smoothly at higher values, and is zero at negative values (Figure 1b):

$$p_\lambda(m,b) = \begin{cases} 0 & b/(m+b) < 0 \\ 2\phi(b/(m+b),0,0.2) & b/(m+b) \geq 0 \end{cases} \tag{6}$$

Here $\phi(x,\mu,\sigma)$ is the normal probability density function, and the factor of two in the second line is needed to make the total area under the prior equal to one, when the prior is considered as a function of $b/(m+b)$. According to this prior, the probability of a lighting condition depends only on the proportion of additive light b to total light $m+b$, and high proportions are unlikely. Arbitrarily large values of m are allowed (see endnote 1).

This is a highly simplified model of greyscale images. It has no notion of spatial adjacency, three-dimensional shape, grouping, material properties, or spatially extended lighting, all of which are important factors in lightness perception[1,2]. Nevertheless, even with these limitations it accounts for and unifies many of the empirical findings that support anchoring theory.

Using this model, we can estimate the reflectances depicted in an image. I will represent an image as an $n \times 1$ vector $\mathbf{l}$ of luminances, sampled from an $n \times 1$ vector random variable $\mathbf{L}$. According to the model, $\mathbf{l} = M\mathbf{r} + B$. In a lightness constancy task the goal is to recover the $n \times 1$ vector $\mathbf{r}$ of reflectances that correspond to the observed luminances $\mathbf{l}$. We could use Bayes' theorem to maximize $P(\mathbf{R} = \mathbf{r} \mid \mathbf{L} = \mathbf{l})$ over all possible $\mathbf{r}$, but this direct approach is inconvenient because $P(R = \mathbf{r} \mid L = \mathbf{l}) = 0$ for most values of $\mathbf{r}$: for fixed $\mathbf{l}$ and arbitrarily chosen $\mathbf{r}$, there is usually no lighting condition (m,b) that satisfies $\mathbf{l} = m\mathbf{r} + b$ for all vector components l_i and r_i. (This is simply because more than two points (r_i,l_i) overdetermine the line $l = mr + b$.) Instead, because the lighting parameters establish a one-to-one relationship between luminance and reflectance, we can make maximum a posteriori (MAP) estimates of the lighting parameters (m,b), and use these to transform the observed luminances into reflectances. Bayes' theorem shows that

$$P(\Lambda = \lambda \mid \mathbf{L} = \mathbf{l}) = \frac{P(\Lambda = \lambda)P(\mathbf{L} = \mathbf{l} \mid \Lambda = \lambda)}{P(\mathbf{L} = \mathbf{l})} \tag{7}$$

Equations (3) and (6) allow us to write this as

$$= \frac{p_\lambda(\lambda)\prod_{i=1}^{n} p_l(l_i;\lambda)}{P(\mathbf{L}=\mathbf{l})} \tag{8}$$

We can find the MAP lighting parameter estimates $\hat{\lambda} = (\hat{m}, \hat{b})$ numerically by maximizing equation (8), and the reflectance estimates are then $\hat{\mathbf{r}} = (\mathbf{l} - \hat{b}) / \hat{m}$.

Thus the search for reflectances is accomplished via a search through lighting conditions. In this sense, the model explicitly solves Gilchrist's[1,2] 'anchoring problem', the problem of finding a mapping from luminance to absolute reflectance. The model also supports Adelson's[3] view that an important part of lightness perception is estimating the 'atmospheric transfer function' that maps $\mathbf{r}$ to $\mathbf{l}$ in the scene being viewed.

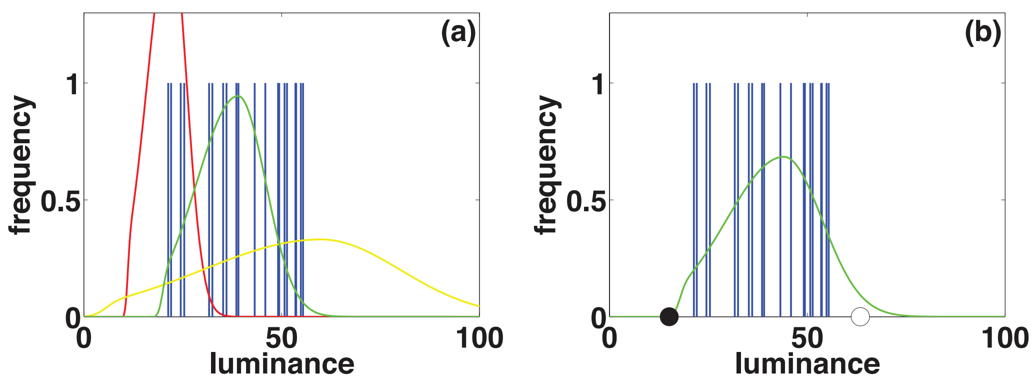

Figure 2. Fits of the shifted and scaled reflectance prior to the observed luminance distribution. Vertical blue lines represent observed luminance values, and solid lines represent the predicted distribution of luminance samples under various lighting conditions. (a) Predicted distribution too low (red line), too wide (yellow line), and approximately correct (green line). (b) Maximum a posteriori fit (green line), with circles indicating the luminance generated by surface patches with 3% reflectance (black dot) and 90% reflectance (white dot).

To understand the model's behaviour in the simulations that follow, it is helpful to think of its performance in terms of the observed image luminances $\mathbf{l}$. Equation (8) shows that the model chooses lighting parameters that maximize $p_\lambda(\lambda)\prod_{i=1}^{n} p_l(l_i;\lambda)$. First consider $\prod_{i=1}^{n} p_l(l_i;\lambda)$, the likelihood of the observed luminances given lighting conditions λ. Equation (3) shows that $p_l(l;\lambda)$ is the reflectance prior with its origin shifted to b, stretched by a factor of m, and scaled by m^{-1}. The model shifts and scales the reflectance prior so as to create a probability density function for luminances that maximizes the likelihood of the observed luminances (see endnote 2). Values of (m,b) that do not make the luminance density completely cover the range of observed luminance values (Figure 2, red line) do not maximize the likelihood, because they assign low probabilities to some observed luminances. Values of (m,b) that make the luminance density bracket the luminances too broadly (yellow line) assign low probabilities to all luminances, since the total area under the density must sum to one. The values of (m,b) that maximize the likelihood (green line) tend to match the luminance distribution $p_l(l;\lambda)$ to the range of observed luminance values. This fit is largely constrained by the facts that the reflectance prior does not extend below zero or much above one, so under lighting condition (m,b), the minimum possible luminance is around b, and the maximum possible luminance is around $m+b$. The model tends to maximize the likelihood by choosing b somewhere around the lowest luminance, and $m+b$ somewhere around the highest luminance. This is just a tendency, though, and we will see exceptions below.

Next consider the lighting prior, the factor $p_\lambda(\lambda)$ in equation (8). This lighting prior puts a cost on the shifting and stretching described in the previous paragraph. High values of $b/(m+b)$ are unlikely, so as the model fits the luminance distribution $p_l(l_i;\lambda)$ to the observed luminances, there is downward pressure on b. Furthermore, when $b=0$, the prior on m is uniform, and when $b>0$, the lighting prior assigns higher probabilities to higher values of m. Thus the lighting prior leaves the model free to choose arbitrarily high values of m, and when $b>0$ it creates pressure not to choose very low values.

3. ANCHORING THEORY

In this section I review several rules of anchoring theory, and show how they follow from the model I have described.

3.1 Anchoring to white

The anchor-to-white rule says that the image region with the highest luminance looks white (i.e., has perceived reflectance around 0.90), and that this highest-luminance region is an 'anchor' that determines the perceived reflectances of other regions in the image as well. Lower-luminance regions are assumed to have little effect on the anchoring point for white, and little effect on the perceived reflectance of other regions. Anchoring theory allows some exceptions to this rule, as discussed below, but assumes that it holds in most images.

The model described in Section 2 also anchors to white. Figure 3 shows the reflectances that the model assigns to 20 luminances evenly spaced from 5 cd/m^2 to 90 cd/m^2, and to 20 luminances evenly spaced from 5 cd/m^2 to 45 cd/m^2. In both cases the highest luminance is assigned a white reflectance. (Compare Gilchrist et al.'s[1] Figure 4, which shows similar results for human observers.) Naturally, there are many combinations of lighting and reflectance that could have generated the observed luminances in Figure 3. It might be, for instance, that all the luminances were generated under the same lighting conditions, but the reflectances in panel (a) were higher than those in panel (b). Nevertheless, in both scenes, anchoring theory and the model assign a white reflectance to the highest luminance.

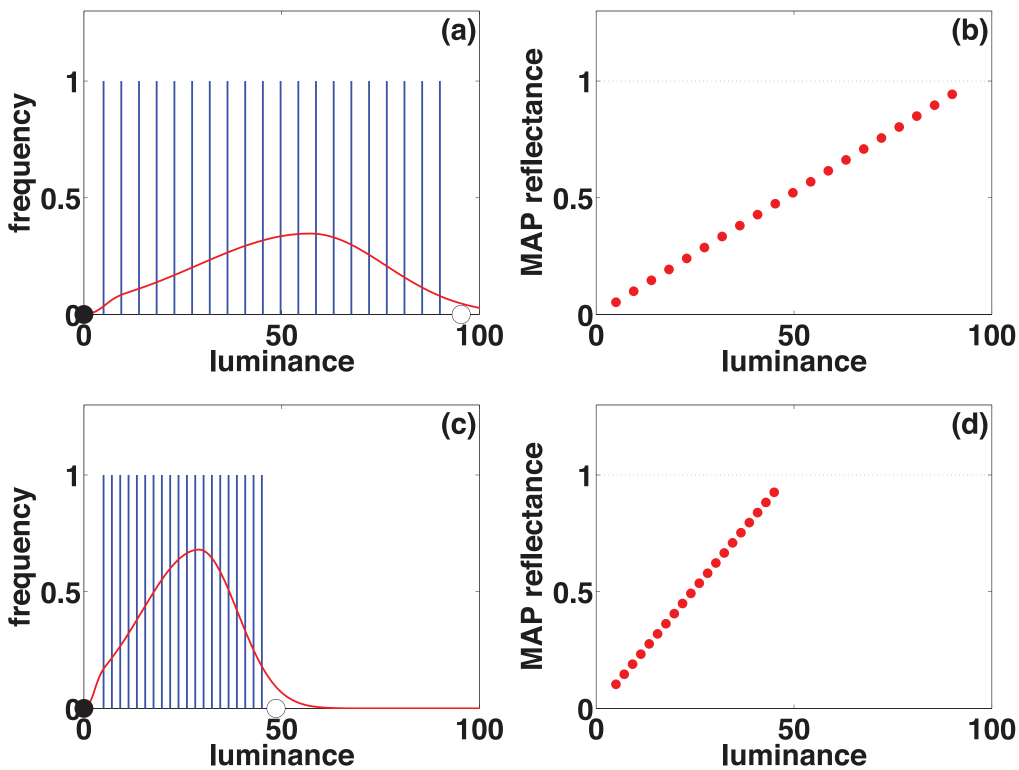

Figure 3. Anchoring to white. (a) Observed luminances spanning the range 5-90 cd/m^2 (vertical blue lines), and the theoretical luminance distribution (red line) that is matched to these observations in the course of making reflectance estimates. The black circle indicates the luminance that is assigned a reflectance of 0.0, and the white circle indicates the luminance that is assigned a reflectance of 1.0. (b) The reflectances assigned to the observed luminances. Panels (c) and (d) show corresponding results for observed luminances that span 5-45 cd/m^2.

The specific reflectance values shown in Figure 3 depend on the specific choice of priors in equations (1) and (6). Nevertheless, it should also be clear that the tendency to see the highest luminance as white depends only on qualitative properties of the priors. As explained earlier, the model tends to choose lighting parameters (m,b) that make the theoretical luminance distribution match the observed range of luminances (Figure 2). The lighting prior makes the model relatively free in choosing the multiplicative parameter m, so the upper end of the reflectance prior is usually mapped onto the upper end of the observed luminance range, and the highest luminance is seen as white.

Anchoring theory holds that although most images contain a patch that looks white, they do not necessarily contain a patch that looks black. The model shows similar behaviour, as shown in Figure 4. Here luminance ranges from 40 cd/m^2 to 90 cd/m^2, and the most luminous image patch is seen as white, but no image patch is seen as black (i.e., reflectance around 0.03). The lowest perceived reflectance is 0.16.

Why does the model usually see a white patch, but not necessarily a black patch? As the model searches for lighting parameters that account for the observed luminances, the lighting prior places much stronger constraints on the additive parameter b than on the multiplicative parameter m. If the model was free to choose any value of b, then the lower range of the reflectance prior would be mapped onto the lower end of the observed luminance range, and the lowest luminances would be seen as black. The additive parameter b is biased towards zero by the lighting prior, though, and so luminances well above zero are seen as grey, instead of black. This point is related to the phenomenon of scale normalization, which I discuss below (Section 3.2).

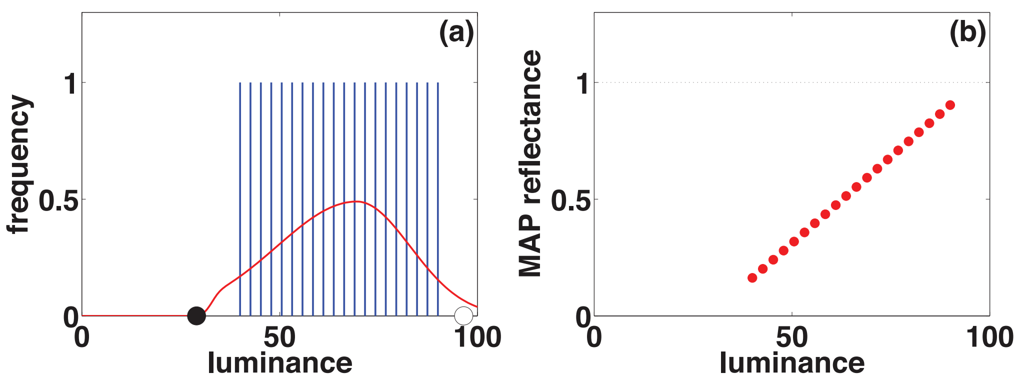

Figure 4. No anchoring to black. See caption for Figure 2 for explanation. Here we fit the model to a higher range of luminances (45-90 cd/m^2), and find that no luminance is assigned a reflectance of black.

Thus the model tends to anchor perceived white at the highest luminance in an image, but does not anchor perceived black at the lowest luminance. Furthermore, the highest luminance in the image has the largest influence on *all* the model's reflectance assignments. Suppose we have 19 luminances evenly spaced between 5 cd/m^2 and 90 cd/m^2, and a 20th luminance, 50 cd/m^2. How strongly does a perturbation of any of the first 19 luminances affect the reflectance assigned to the 20th ? Figure 5 shows dr_{20} / dl_i, the derivative of the reflectance r_{20} assigned to luminance $l_{20} = 50 \text{ cd/m}^2$, with respect to the other 19 luminances. Increasing l_i usually decreases r_{20}, but this effect is stronger at the lowest and highest luminances than at the middle luminances, and it is strongest of all at the highest luminance.

Again, this pattern of results follows from qualitative properties of the priors. The model tends to choose lighting parameters that match the theoretical luminance distribution to the observed luminances, and these lighting parameters determine the mapping from luminance to reflectance. When the highest luminance changes, the weak constraint on lighting parameter m means that the model is free to adjust this parameter to match the new lighting distribution, and all other reflectances are adjusted accordingly. When the lowest luminance changes, the stronger constraint on lighting parameter b means that the model is less free to adjust this parameter, so there is less of an effect on the assigned reflectances. When middle luminances change, the model is still under pressure to match the range of the theoretical luminance distribution to the range of the observed luminances, so there is even less of an effect on the chosen lighting parameters an the assigned luminances.

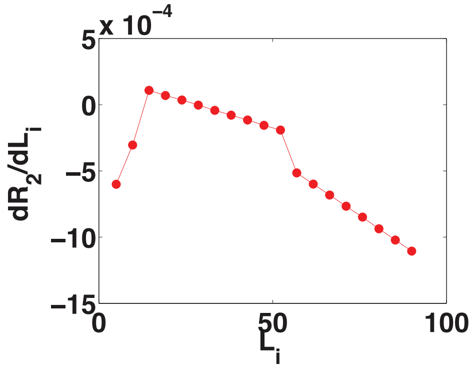

Figure 5. Influence function, showing dr_{20} / dl_i for l_i ranging from 5 to 95 cd/m^2, and $l_{20} = 50$ cd/m^2 .

Gilchrist counts the fact that human observers anchor their lightness percepts to white, and not to black or some intermediate level of grey, as one of the unexplained facts about anchoring theory. The model offers an explanation of why this happens. Interestingly, the model tends to anchor at white not because it believes white to be a particularly common reflectance, or because it has an explicit role for the highest luminance. Instead, anchoring to white emerges as a byproduct of finding a MAP explanation of the observed luminances (which tends to match the theoretical luminance distribution to the observed luminances), combined with a prior on lighting conditions that allows only small amounts of additive light.

3.2 Proportional reflectance

According to the anchor-to-white rule, the highest luminance in a scene looks white. Anchoring theory says that every other image patch is assigned its reflectance r_i proportionally, based on the ratio of its luminance l_i to the highest luminance l_{max} :

$$r_i = 0.90 \frac{l_i}{l_{\max}} \tag{9}$$

The model also tends to assign reflectances proportionally. The lighting model says that $\mathbf{l} = m\mathbf{r} + b$, and the lighting prior ensures that b is usually small. As a result, $\mathbf{l} \simeq m\mathbf{r}$, and reflectance is usually approximately proportional to luminance, as illustrated in Figure 3. This approximation fails when $b >> 0$, but as we will see in the next section, this also agrees with anchoring theory.

3.3 Scale normalization

According to the first two rules, the highest luminance in a scene looks white, and other regions are assigned reflectances proportionally based on their luminance. The scale normalization rule says that if the reflectances that result from the first two rules span a range that is larger or smaller than 30:1, then the range is adjusted towards 30:1 by lightening or darkening the lower reflectances. If the reflectance range that results from the first two rules is $R = r_{max} / r_{min}$, then scale normalization adjusts the range to:

$$R_{norm} = (1 + (0.56 * (\log(30) - \log(R))) * R \tag{10}$$

Equation (10) is based on fits to human lightness judgements.

The model also shows scale normalization. Figure 6 shows the model's predictions and anchoring theory's predictions (based on equation (10)), when luminances (and hence the perceived reflectances predicted by just the first two rules of anchoring theory) span a range of 30:1, 3:1, or 1.5:1. The model shows about the same amount of scale normalization as anchoring theory. Without scale normalization, perceived reflectance would be proportional to luminance, and all data points would lie on a line through the origin, as they mostly do in the condition with a 30:1 luminance range (red circles). When the luminance range is less than 30:1, the reflectance range expands, and the points lie on a line with a positive x-intercept.

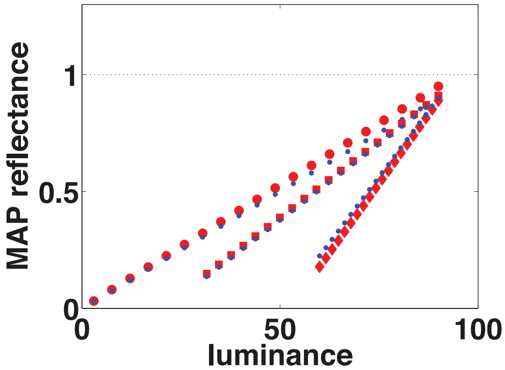

Figure 6. Scale normalization. The red circles represent 20 luminances evenly spaced between 3 and 90 cd/m^2, the squares span 32 to 90 cd/m^2, and the diamonds span 60 to 90 cd/m^2. The red data points show the reflectances that the model assigns to these luminances, and the small blue dots indicate anchoring theory's predictions. The model assigns reflectances that are very similar to those obtained from anchoring theory.

Scale normalization results from the model's tendency to choose lighting conditions that match the theoretical luminance distribution to the observed luminances. As shown in Figure 4a, when lowest luminance is well above zero and consequently the luminance range is small, the model chooses a high value of b, shifting the black point upwards, and luminances that would otherwise be seen as mid-grey look dark instead. Furthermore, the lighting prior puts a cost on additive light b, so the black point cannot be shifted upward with complete freedom, and scale normalization is incomplete (i.e., the reflectance range shifts towards 30:1, but does not reach that value), consistent with equation (10).

Anchoring theory also predicts that the reflectance range is compressed if the first two rules result in a range greater than 30:1. I will briefly mention that the model exhibits this behaviour as well, but for quite different reasons than those behind range expansion. Range compression results from the fact that the reflectance prior assigns a high cost to very low and very high reflectances, limiting the range of reflectances that can be assigned to even very widely spaced luminances. Allred, Lohnas, and Brainard[7] report a similar effect in a Bayesian lightness model.

3.4 Two rules for glow

Anchoring theory assumes that lightness is a psychological continuum that extends upwards beyond the physically possible reflectance range (0-1), and that high values on this continuum represent surfaces that appear to glow. Bonato and Gilchrist[8,9] demonstrated two properties of glow perception. First, large surfaces must have higher luminances than small surfaces in order to appear to glow. Second, if the luminance of a large glowing region increases, this creates a smaller increase in perceived glow than it would have created in a small surface, and furthermore it reduces the perceived lightness of other image regions.

The model accounts for these properties of glow. Figure 7a shows the model's reflectance assignments in an image that mostly spans a limited range of luminances, and has one patch with a much higher luminance. Here, instead of anchoring white to the highest luminance, the model shifts the theoretical luminance distribution mostly to match the main cluster of luminances, and the outlier patch is assigned a very high reflectance, above the white point, in the region that corresponds to percepts of glow.

When we increase the size of the high-luminance patch (Figure 7b), we add more samples in the far upper tail of the luminance density (assuming that the image is divided into patches of equal size), where the probabilities are low. Now the model shifts the theoretical luminance density upwards, to avoid assigning low-probability reflectances to a larger number of image patches. As a result, the high-luminance patch is not as far above the white point, so it appears to glow less. Thus, all other things being equal, large surface patches glow less than small surface patches.

When the small high-luminance patch is assigned an even higher luminance (Figure 7c), the white point rises only slightly, and the patch appears to glow more strongly. When the large high-luminance patch is assigned a higher luminance, though, the white point rises with it. Accordingly, the patch does not appear to glow much more than before, and furthermore the perceived reflectance of all other image patches goes down.

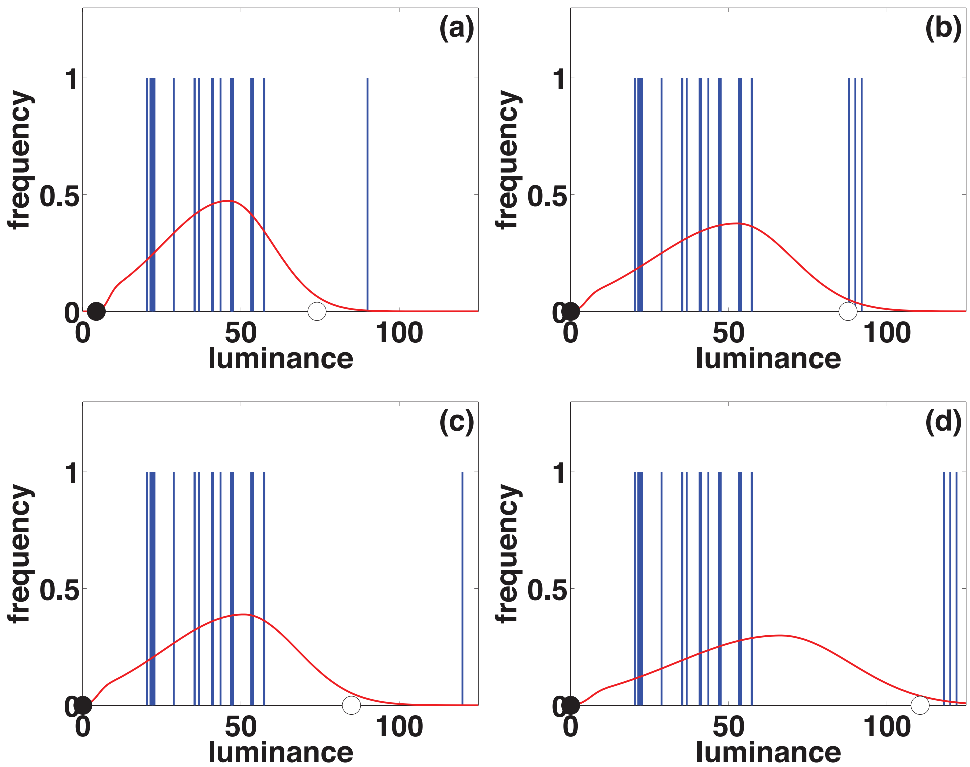

Figure 7. Glow perception. (a) A small high-luminance outlier lies well above the white point, and appears to glow (b) A larger high-luminance outlier does not lie as far above the white point, and so does not glow as strongly. (c) Increasing the luminance of the small high-luminance outlier increases its perceived glow (compare to panel (a)). (d) Increasing the luminance of the larger high-luminance outlier does not increase its perceived glow very much, but decreases the reflectances assigned to the other image patches (compare to panel (b)).

4. DISCUSSION

4.1 The model's contributions

If this model's predictions are so similar to those of anchoring theory, then what does it contribute to our understanding of lightness perception? I suggest that it makes the following contributions.

(a) The model shows that several unusual anchoring rules, that might seem simply to reflect failures of human lightness perception, actually follow from a rational strategy for solving the highly underconstrained problem of recovering surface reflectance. It may seem surprising that a model that aims at veridical perception produces so many systematic errors. Lightness perception is highly underconstrained, though, so in a Bayesian treatment of this problem the priors play an important role, and they create systematic biases under some conditions.

(b) The model draws connections between lightness phenomena that would otherwise seem unrelated, and makes novel experimental predictions. For example, the model implies that scale normalization is closely related to the perception of additive light, such as haze, and which suggests that manipulating cues to haze should affect the degree of scale normalization.

(c) The model formalizes ideas that already seem to be implicit in some of the anchoring rules. Scale normalization, for instance, may strike us as a reasonable property of lightness perception, if we think it unlikely for a rich scene to contain a very limited range of reflectances. It is also true, though, that the most probable sequence of samples from a random variable (such as the random variable described by the reflectance prior) is a sequence of identical values at the peak of the probability distribution function – in which case, is a limited range of reflectances unlikely after all? Such questions are difficult to resolve without a statistical understanding of the problem. In this example, the model shows that scale normalization results from treating lightness perception as a MAP estimation problem, which tends to match the theoretical luminance distribution to the observed luminance distribution, usually resulting in a wide range of perceived reflectances.

(d) The model simplifies our understanding of lightness perception. Instead of formulating lightness perception as a sequence of rules, some correcting the results of earlier rules, the model presents lightness perception as a single-stage optimization problem, based on simple assumptions about reflectance and lighting. This is preferable when we have no evidence that the intermediate results of anchoring theory (e.g., reflectance estimates before scale normalization) actually exist in the human visual system.

(e) The model suggests that the notion of choosing an anchor for white should be replaced by the notion of estimating the lighting conditions in the scene[3]. As anchoring theory has been extended to account for human lightness percepts more and more accurately, the notion of an anchor seems to have been progressively weakened. The anchor-to-white rule says that the highest luminance is an anchor that is perceived as white, and the proportional reflectance rule says that other luminances are perceived according to their ratio to the anchor. The area rule (not discussed in this paper; see endnote 3) revises this, so that large, low-luminance regions pull the anchor downwards, until it is lower than the highest luminance, and the highest luminance appears to glow. The scale normalization rule then removes the requirement that all other luminances are seen as proportional even to this modified anchor. In the end, the anchor becomes a hypothetical construct that appears nowhere in the image. From here it is just a small step to the model's arguably simpler approach of using all luminances in the scene simultaneously to estimate a single linear mapping from luminance to reflectance.

4.2 Future work

A natural next step in developing this model is to estimate priors that quantitatively explain human performance in specific lightness perception tasks. Here I have been mostly interested in showing that the model qualitatively predicts several phenomena in lightness perception. The various experiments demonstrating scale normalization, glow effects, and so on, were done under very different conditions (e.g., different lighting conditions), so there is little point in optimizing the reflectance and lighting priors to account for these previous results quantitatively. Allred et al.[7] have begun work on inferring observers' priors from their performance on lightness judgement tasks, using a model that is similar to the one I describe here.

Another natural direction for developing this model is to incorporate 3D shape. In 3D scenes, the luminance of a surface patch depends on it reflectance, and also on its orientation relative to the various light sources in the scene. This changes the lightness estimation problem considerably, e.g., the patch with the highest reflectance need not be the patch with the highest luminance. The model could be extended, for instance, to allow a scene to have a single Phong-like illuminant whose parameters need to be estimated (just as lighting parameters m and b are estimated in the current model), and to allow the orientation of each surface patch to be drawn from a statistical distribution on the sphere.

Finally, I have assumed that the model has noiseless measurements of the image luminances $\mathbf{l}$. If these measurements are contaminated with Gaussian noise, then every assignment of reflectances has a nonzero probability, and we must maximize $P(\mathbf{r} \mid \mathbf{l})$ over all possible reflectances $\mathbf{r}$, not just the subspace of reflectances that are an affine transform of the observed luminances Visual processing is permeated with noise, so incorporating noise into Bayesian models of lightness perception is an interesting avenue for future work.

ENDNOTES

1. This description does not completely specify the prior, because the integral of equation (6) over $\mathbb{R}^2$ is unbounded. We can limit the total volume under the prior by setting it to zero above some values of m and b that are so high they will never be encountered.

2. In this way, the model's approach to lightness perception is similar to the problem of observing n samples from a uniform distribution over an unknown interval $[a,b]$, and estimating the interval endpoints a and b. This is a classic example of a statistical problem in which maximum likelihood estimates are biased (the estimate of a is biased upwards, and b downwards), which raises the question of whether human lightness percepts are biased in the same way. (An

important difference is that the interval endpoint problem has no prior on *a* and *b*.)

3. Additional simulations, very similar to those in Section 3, show that the area rule also follows from the model.

ACKNOWLDGEMENTS

I thank David Brainard, Marianne Maertens, and Pascal Mamassian for helpful discussions.

REFERENCES

[1] Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, Xiaojun, Spehar, B., Annan, V., & Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review, 106,* 795-834.

[2] Gilchrist, A. (2006). *Seeing black and white.* New York: Oxford University Press.

[3] Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The New Cognitive Neurosciences* (2nd ed.) (pp. 339-351). Cambridge, MA: MIT Press.

[4] Barron, J. T., & Malik, J. (2011). High-frequency shape and albedo from shading using natural image statistics. *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2521-2528), June 21-23, 2011, Colorado Springs, CO.

[5] Barron, J. T., & Malik, J. (2012). Shape, albedo, and illumination from a single image of an unknown object. *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 334-341), June 16-21, 2012, Providence, RI.

[6] Attewell, D., & Baddeley, R. J. (2007). The distribution of reflectances within the visual environment. *Vision Research, 47,* 548-554.

[7] Allred, S. R., Lohnas, L. J., & Brainard, D. H. (2008). Bayesian model of the staircase Gelb effect [abstract]. *Vision Sciences Society Annual Meeting*, May 9-14, 2008, Naples, FL.

[8] Bonato, F., & Gilchrist, A. L. (1994). The perception of luminosity on different backgrounds and in different illuminations. *Perception, 23,* 991-1006.

[9] Bonato, F., & Gilchrist, A. L. (1999). Perceived area and the luminosity threshold. *Perception & Psychophysics, 61,* 786-797.