Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research 44 (2004) 1827–1842
Vision Research
www.elsevier.com/locate/visres

Lightness constancy through transparency: internal consistency in layered surface representations

Manish Singh *

Department of Psychology and Center for Cognitive Science, Rutgers University—New Brunswick, 152 Frelinghuysen Road, Piscataway, NJ 08854-8020, USA

Received 2 January 2003; received in revised form 5 February 2004

Abstract

Asymmetric lightness matching was employed to measure how the visual system assigns lightness to surface patches seen through partially-transmissive surfaces. Observers adjusted the luminance of a comparison patch seen through transparency, in order to match the lightness of a standard patch seen in plain view. Plots of matched-to-standard luminance were linear, and their slopes were consistent with Metelli's α. A control experiment confirmed that these matches were indeed transparency based. Consistent with recent results, however, when observers directly matched the transmittance of transparent surfaces, their matches deviated strongly and systematically from Metelli's α. Although the two sets of results appear to be contradictory, formal analysis reveals a deeper mutual consistency in the representation of the two layers. A ratio-of-contrasts model is shown to explain both the *success* of Metelli's model in predicting lightness through transparency, and its *failure* to predict perceived transmittance—and hence is seen to play the primary role in perceptual transparency.

Keywords: Lightness; Transparency; Surface segmentation; Layered representation; Image decomposition; Scission; Contrast; Transmittance; Opacity

1. Introduction

The intensity of light reaching the eyes from a given surface depends not only on its own intrinsic reflectance, but also on the visual context that it is placed in. A black surface in bright illumination and a white surface in dim illumination, for instance, can project identical luminance values onto the retinas. Nevertheless, they are usually perceived to be intrinsically dark and light surfaces, respectively. Although this ability, lightness constancy, has most often been studied in the context of surfaces viewed under different illumination conditions (Gilchrist, 1979; Gilchrist et al., 1999; Kraft, Maloney, & Brainard, 2002; Land & McCann, 1971; Maloney & Yang, 2003; Rutherford & Brainard, 2002; Schirillo & Shevell, 1997, 2002), a similar problem arises in cases where a surface is viewed through a partially-transmissive layer, such as a transparent filter, veil, or mesh screen, or—in the natural environment—haze, fog, murky water, or dense foliage. In order to compute the lightness of a surface seen through a partially-transmissive layer, the visual system must analyze image luminance into the separate contributions of the partially-transmissive layer and the underlying surface seen through it. The image decomposition implicit in such analysis is illustrated schematically in Fig. 1: the variation along the single dimension of luminance (i.e., the luminance profile in 1a) is decomposed into a representation of two 'layers' with distinct surface qualities. Indeed, it has been proposed that many lightness illusions can best be understood in terms of such layered surface representations (Anderson, 1997; Somers & Adelson, 1997).

It has long been known that, at least in sufficiently rich contexts, the visual system is able to 'discount' [1] the

* Tel.: +1-7324456823; fax: +1-7324456715.
E-mail address: manish@ruccs.rutgers.edu (M. Singh).

[1] 'Discount' is put in quotes to emphasize the fact that discounting does not mean discarding (see, e.g., Gilchrist & Jacobsen, 1983). This point is especially important in the context of perceptual transparency, because the visual system clearly has explicit representations both of the underlying surface, and of the transparent layer itself.

0042-6989/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2004.02.010

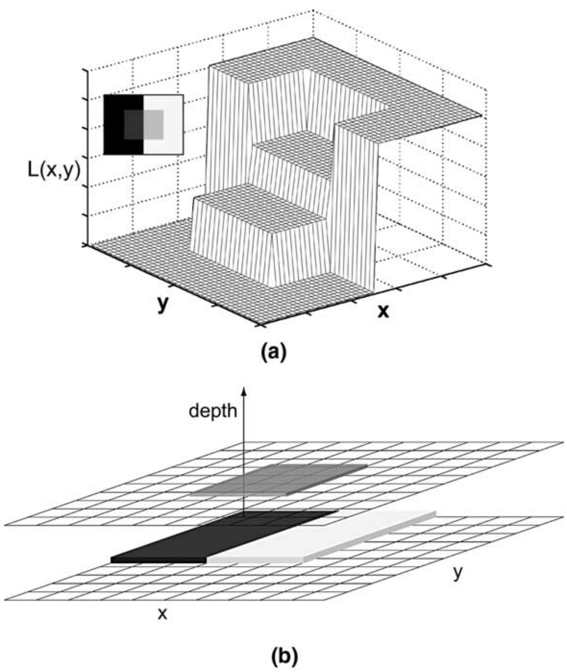

Fig. 1. The decomposition of image luminance into two separate layers in perceptual transparency. (a) The luminance profile of a simple grayscale image. (b) The layered surface representation generated by the visual system: a gray partially-transmissive surface and a bipartite opaque surface seen through it.

presence of *veiling luminance*—the introduction of a fixed luminance increment to a scene, for example, due to light reflected in a glass pane through which the scene is viewed (Gilchrist & Jacobsen, 1983). The current paper investigates lightness constancy in the more general context of partially-transmissive layers. The context of transparency is more general because of two (not altogether independent) reasons: (1) In addition to the purely additive component in veiling luminance, in transparency there is also an attenuating multiplicative component—which arises due to the partial transmittance of the transparent layer; and (2) whereas veiling luminance can, by definition, only increase the projected luminance, the presence of a transparent layer can either increase or decrease it—depending on whether the transparent layer is lighter or darker than the underlying surface. The current paper investigates how the visual system quantitatively assigns lightness to surface patches seen through transparency, using simulated transparent layers that either darken, lighten, or preserve overall mean luminance. Furthermore, it compares matches made on transparent layers with those made on the underlying surfaces seen through them. This direct comparison provides a more powerful means of investigating the visual system's internal representation of multiple layers in transparency. In particular, it allows one to test whether there is a quantitative consistency in the visual system's representation of the transparent layer and the underlying opaque surface. [2]

Two recent studies, suggesting opposing conclusions, are of direct relevance to the reported experiments. First, in the chromatic domain, results by D'Zmura, Rinner, and Gegenfurtner (2000) indicate that the functional form of Metelli's equations (see Section 2)—appropriately extended to color space—adequately captures the colors of surface patches seen through color filters (although observers underestimate the extent of color convergence produced by a color filter by a factor of almost two). However, in the context of achromatic transparency, Singh and Anderson (2002a) have recently demonstrated that the perceptual assignment of transmittance (sometimes referred to as *degree of transparency*) to transparent layers deviates strongly and systematically from the prediction of Metelli's model, thus arguing against its perceptual validity.

In the current paper, I investigate lightness constancy through achromatic transparency. In particular, I measure the extent to which Metelli's equations capture the lightness of surface patches seen through achromatic transparent layers. Moreover, I directly compare—and quantitatively reconcile—the results of matching the lightness of surface patches seen through transparent layers, with transmittance matches made on the transparent layers themselves. This comparison and reconciliation yield a more complete understanding of the visual computation and representation of multiple layers in transparency.

2. Metelli's model of transparency

Metelli (Metelli, 1970, 1974a, 1974b, 1985; Metelli, Da Pos, & Cavedon, 1985) proposed a model of transparency based on a physical setup involving an episcotister—a rotating disk with an open sector (or wedge) of relative area α (see Fig. 2a). When the rotation of the episcotister is sufficiently rapid, it is perceived as a homogeneous partially-transmissive layer (see Fig. 2b). If the surface of the episcotister has reflectance t, and an underlying surface region has reflectance a, the resulting 'color mixing' is described by

$$p = \alpha a + (1 - \alpha)t. \tag{1}$$

In other words, Metelli modeled transparency in terms of Talbot's equation of color mixing: the 'colors' [3] of the

[2] This issue is analogous to the *illuminant-estimation hypothesis* in color perception (Maloney & Yang, 2003; Yang & Maloney, 2001) according to which there is a mutual consistency in the visual system's representation of the illuminant and surface color (but see Rutherford & Brainard, 2002).

[3] Metelli used "color" to refer to "achromatic color," or the *reflectance* of an achromatic lambertian surface.

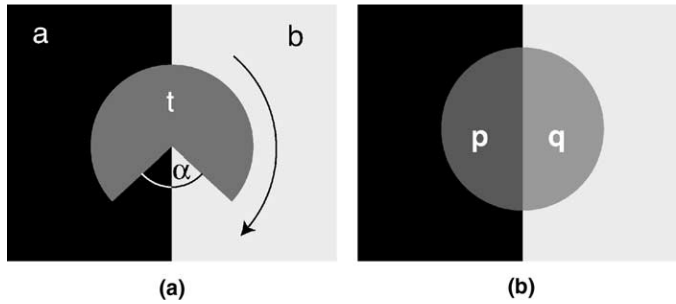

Fig. 2. Metelli's episcotister model of transparency. (a) A disk with an open sector (*episcotister*) is rotated in front of a bipartite background. (b) When the rotation is sufficiently rapid, it generates the percept of a homogeneous partially-transmissive surface. The resulting 'color mixing' can be described using Talbot's law (see Eqs. (1) and (2)).

transparent surface and an underlying opaque surface are mixed—their mixing proportions determined by the transmittance of the transparent layer (i.e., the proportion of light it lets through). Given only a and p in an image, it is clearly impossible to determine α and t uniquely. However, if there is one other surface region, with distinct reflectance b, which is also partly visible through the same episcotister, then one obtains an additional equation:

$$q = \alpha b + (1 - \alpha)t. \tag{2}$$

Eqs. (1) and (2) now yield the following solutions:

$$\alpha = \frac{p - q}{a - b}, \tag{3}$$

$$t = \frac{aq - bp}{a + q - b - p}. \tag{4}$$

Metelli argued that these equations describe both image generation and perception. According to this model, then, solutions (3) and (4) predict how observers will perceive the transmittance and lightness of a transparent layer. [4] It is noteworthy that the term α plays two distinct roles in Metelli's model: (i) α is the slope of the mapping from 'colors' in plain view (i.e., a and b) to those seen through transparency (i.e., p and q), and (ii) α is the perceived transmittance of the transparent layer. This dual role of α will assume importance for us later.

Although Metelli wrote his equations in terms of reflectance values, Gerbino, Stultiens, Troost, and de Weert (1990) have shown that the same equations also follow in the luminance domain, under the assumption of uniform illumination. The luminance version of the equations is more natural for perceptual theory, because the visual system is given luminance values, not reflectance values, as input. In what follows, the luminance version of Metelli's equations will be assumed—i.e., a, b, p, and q will be treated as luminance values. Moreover, although the equations were derived from a physical setup involving an episcotister, the same equations also follow if a veil, mesh, or dense foliage is used instead of an episcotister. In these cases, the partially-transmissive layer is naturally modeled as a surface containing a large number of holes that are too small to be resolved individually (Richards & Stevens, 1979)—and the 'color mixing' takes place spatially rather than temporally. Finally, despite their simplicity, Metelli's equations also provide a reasonable approximation to more complex cases involving physical filters and paint layers (Beck et al., 1984; Faul & Ekroll, 2002; Gerbino, 1994; Kubelka & Munk, 1931; Nakauchi, Silfsten, Parkkinen, & Ussui, 1999), as well as fog (Hagedorn & D'Zmura, 2000; Mahadev & Henry, 1999). [5]

D'Zmura et al. (1997) have extended Metelli's equations to the chromatic domain:

$$\begin{pmatrix} p_1 \\ p_2 \\ p_3 \end{pmatrix} = \alpha \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} + (1 - \alpha) \begin{pmatrix} t_1 \\ t_2 \\ t_3 \end{pmatrix}. \tag{5}$$

It is easily seen from this equation that the presence of an overlying transparent layer leads to a consistent convergence in color space: if $\vec{a}$, $\vec{b}$, $\vec{c}$, $\vec{d}$ are the colors of underlying surface patches seen in plain view, and $\vec{p}$, $\vec{q}$, $\vec{r}$, $\vec{s}$ are the colors of same patches seen through a homogeneous transparent layer, then the vectors $\vec{p} - \vec{a}$, $\vec{q} - \vec{b}$, $\vec{r} - \vec{c}$, and $\vec{s} - \vec{d}$ converge toward a common point in

[4] Although I do not delve into this issue in the current paper, it should be noted that Metelli's solutions for α and t also yield certain *qualitative constraints* for transparency (e.g., Metelli, 1974b). Such photometric constraints, plus geometric constraints involving junctions and contour continuity, have been developed in recent work to predict the perceptual decomposition of image luminance into a transparent and an underlying surface (Adelson & Anandan, 1990; Anderson, 1997, 2003; Beck, Prazdny, & Ivry, 1984; D'Zmura, Colantoni, Knoblauch, & Laget, 1997; Singh & Anderson, 2002a, Singh & Hoffman, 1998; Singh & Huang, 2003).

[5] The filter model converges to the episcotister model as the level of illumination gets increasingly higher (Gerbino, 1994), and the fog model approximates the Metelli model under the assumption that the fog extinction coefficient is independent of wavelength (Hagedorn & D'Zmura, 2000)—a reasonable assumption for naturally occurring clouds and fog (McClatchey, Fenn, Selby, Volz, & Garing, 1978).

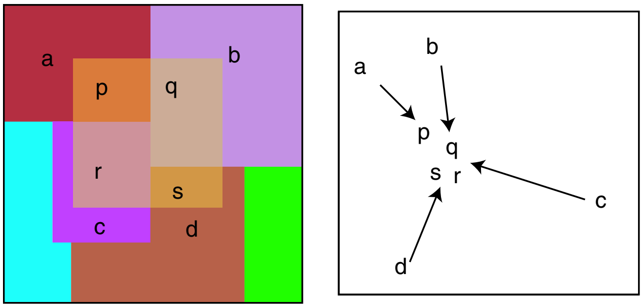

Fig. 3. Schematic illustration D'Zmura et al.'s (1997) convergence model of chromatic transparency. The presence of an overlying color filter leads to a consistent convergence (or, in the limit, a translation) in color space. The visual system can make use of such consistent convergence in the image to detect the presence of an overlying color filter.

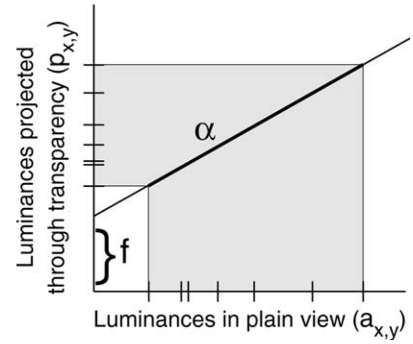

Fig. 4. The photometric transformation introduced by the presence of an intervening transparent layer. According to Metelli's model, this transformation is given by a linear map $p = \alpha a + f$, with $0 \leqslant \alpha \leqslant 1$ and $f \geqslant 0$. The compressive slope α results from the light-attenuating properties of the transparent layer, and the additive term f from its reflectance.

color space (see Fig. 3). The color $\vec{t}$ of the filter determines the common point toward which convergence occurs, and its transmittance α determines the extent of convergence toward this point. (The extent of convergence is inversely related to transmittance: zero convergence corresponds to the absence of any interposed layer, $\alpha = 1$, whereas complete convergence corresponds to the case of an opaque occluder, $\alpha = 0$. Indeed, the extent of convergence is given simply by $1 - \alpha$.)

The product $(1 - \alpha)t$ in Metelli's equations is often collapsed into a single additive term f (Adelson, 2000; Adelson & Anandan, 1990; D'Zmura et al., 2000; Gerbino et al., 1990):

$$p_{x,y} = \alpha a_{x,y} + f. \tag{6}$$

Here, the subscripts (x, y) index image position. [6] This form makes it explicit that the mapping from luminances $a_{x,y}$ projected directly from surfaces in plain view to luminances $p_{x,y}$ projected from the same surfaces through a transparent layer, is a linear function with slope α (see Fig. 4). Adelson (2000) similarly uses a linear mapping (the "atmospheric transfer function" or ATF) to capture the physical effects of not only transparency, but other transformations as well—such as those due to shadows, spotlights, and even contrast-enhancing transformations. [7] The constraints $0 \leqslant \alpha \leqslant 1$ and $f \geqslant 0$ ensure that one is in the domain of transparency. In the context of Metelli's setup (where the background surface projects only two distinct luminance values a and b), the additive term f is given by

$$f = \frac{aq - bp}{a - b}. \tag{7}$$

3. Lightness constancy through transparency

Since our goal is to investigate how the visual system quantitatively assigns lightness to surface patches seen through transparency—and whether this perceptual assignment is consistent with Metelli's equations—a natural way to pose our question is: to what extent does the linear form of Eq. (6) capture the lightness of surface patches seen through a transparent layer? To address this question, Experiment 1 uses stereoscopic transparency displays with textured backgrounds, such as the one shown in Fig. 5. In this stereoscopic display, a transparent surface is perceived to be floating in front of the left half of a textured background. In addition, the display contains two circular patches, located in between the depth planes of the textured background and the transparent layer (see the schematic of the depth layering in Fig. 6). One of these patches (the one on the right in Figs. 5 and 6) is seen as being in plain view, whereas the other is seen through the transparent layer. In the experiment, the patch in plain view acts as the *standard*, and its luminance is set to different values from trial to trial. Observers then adjust the luminance of the other *comparison* patch—seen through the transparent layer—in order to match the lightness of the standard patch. The matches thus obtained allow us to address two basic issues concerning the mapping from standard luminance values in plain view to adjusted luminance values through transparency:

[6] More generally, α and t can also be functions of image position (x, y), thus allowing the transparent layer to be inhomogeneous in transmittance and/or reflectance (see, e.g., Singh & Anderson, 2002a). For simplicity, I assume balanced, or homogeneous, transparency in the current discussion.

[7] Adelson's ATF is a mapping from *reflectance* values to luminance values projected through an atmosphere. Therefore, it is more appropriate to say that Eq. (6) describes the mapping from luminances under a "default atmosphere" (i.e., in plain view) to luminances through a transparent layer.

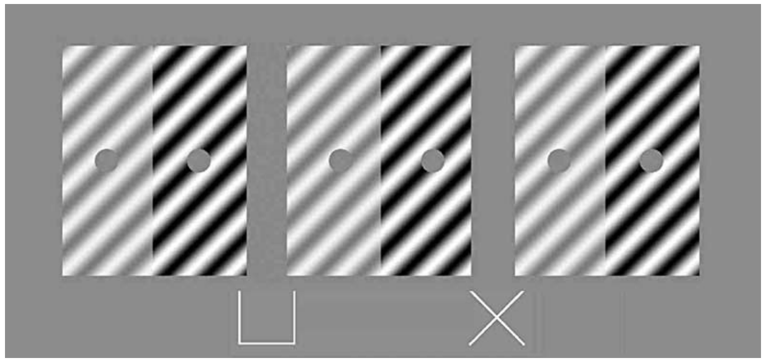

Fig. 5. The stereoscopic stimulus configuration used in Experiment 1. (Cross fusers should fuse the two right images, parallel fusers the two left images.) Observers adjusted the luminance of the comparison patch seen through a transparent layer in order to match the lightness of a standard patch seen in plain view.

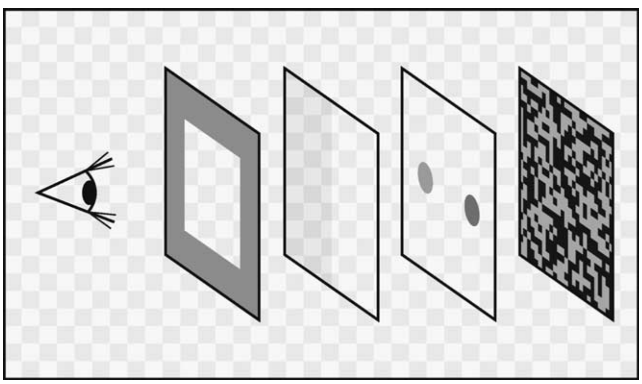

Fig. 6. A schematic of the depth stratification of surfaces seen in the stereoscopic stimuli used in Experiment 1.

(a) First, is this mapping linear, as predicted by Eq. (6)?
(b) Second, assuming linearity, is the pattern of slopes consistent with Metelli's solution for α (i.e., Eq. (3))?

As mentioned earlier, recent experiments by Singh and Anderson (2002a) demonstrated that the perceived transmittance of a transparent layer deviates systematically from Metelli's solution for α. In particular, whereas Metelli's α predicts that perceived transmittance should be independent of the mean luminance in the region of transparency (depending only on its luminance range) their results demonstrated that perceived transmittance decreases linearly with increasing mean luminance. Thus perceived transmittance scales with the contrast in the region of transparency (rather than its luminance range). This conclusion was also reached by a study by Robillotto, Khang, and Zaidi (2002). (A point of difference between the two studies, however, is that whereas for Singh and Anderson's (2002a) displays, perceived contrast was well captured by Michelson contrast, for the more complex and variegated texture displays used by Robillotto et al. (2002), neither Michelson contrast nor any other standard measure of contrast was found to capture perceived contrast.) Thus, perceived transmittance is consistent not with Metelli's α, but rather with

$$\alpha_c = \frac{p_{\text{contrast}}}{a_{\text{contrast}}}, \tag{8}$$

where p_{contrast} is the contrast of the region of transparency and a_{contrast} is the contrast in the region seen in plain view.

As we noted earlier, the term α plays two distinct roles in Metelli's model: (i) α is the slope of the mapping from luminances in plain view to luminances through transparency; and (ii) α is the transmittance of the transparent layer. If it is indeed true, perceptually speaking, that the slope of the mapping from luminances in plain view to luminances through transparency is quantitatively equal to the perceived transmittance of the transparent layer then, based on Singh and Anderson's transmittance-matching results, an alternative prediction would be that the slopes of the standard-to-matched luminance values are given by α_c rather than by α.

The simplest way to distinguish between slopes based on α and those based on α_c is to use a set of displays in which the region of transparency is defined by a common value of luminance range (i.e., amplitude $p_{\text{range}} = p_{\max} - p_{\min}$), but different values of mean luminance ($p_{\text{mean}} = \frac{p_{\max}+p_{\min}}{2}$); [8] see Fig. 7. In particular, the mean luminances can be chosen so that the simulated transparent layer either darkens, lightens, or preserves the mean luminance of the underlying surface. Given such a set of displays, slopes based on α are predicted to be constant across the three displays, whereas slopes based on α_c are predicted to decrease monotonically with increasing p_{mean}. Therefore, in order to test the validity of Eq. (6) (which has slope α), it is important not only to compare the specific values of the observed slopes with

[8] It should be noted that the textured backgrounds used in the current study have symmetric luminance distributions.

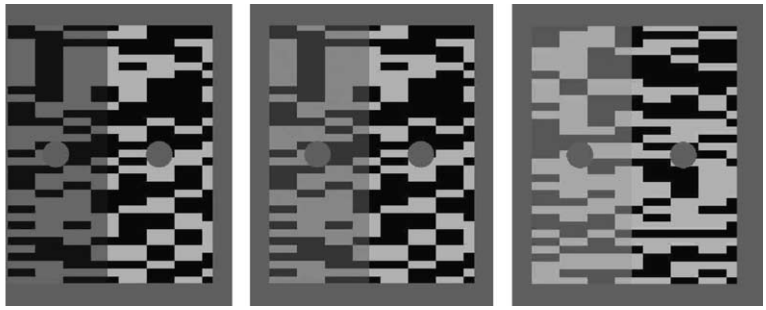

Fig. 7. The three simulated transparent layers used in Experiment 1. One decreased the mean luminance in the region of transparency (left), one preserved mean luminance (middle), and one increased mean luminance (right). In each case, however, the luminance range in the region of transparency was the same. These were presented stereoscopically, as in Fig. 5. The background contained either a random-dot texture with horizontally elongated elements (as shown here), or sinusoidal gratings (as shown in Fig. 5).

the slopes predicted by Metelli's equations, but also to test for this higher-order regularity in the pattern of slopes.

3.1. Experiment 1

3.1.1. Methods

Observers. Three observers, with normal or corrected-to-normal visual acuity, participated in the experiment. Two of the observers were naive to the purposes of the experiment; the third (O3) was the author.

Stimuli and apparatus. Each stimulus display consisted of a $6.3° \times 7.9°$ rectangular frame containing one of two texture patterns. The elements of the texture were given a far disparity of 12 min of arc relative to the rectangular frame, and half-occlusions were introduced at the left and right edges of the frame: texture elements adjacent to the left edge were present only in the right eye's image, and texture elements adjacent to the right edge were present only in the left eye's image ("Da Vinci stereopsis," see Nakayama & Shimojo, 1990). This generated the percept of a frontoparallel textured surface, viewed through a rectangular window.

The textured region was divided laterally by a vertical contrast border, so the left and right halves of the display could be given different values of mean luminance and luminance range. This vertical contrast border was given a near disparity of 7.25 min of arc relative to the texture elements. This generated the percept of an overlying transparent surface—on the side with lower contrast—which appeared to be floating at the depth of the contrast border. The high-contrast side of the display (which was seen to be in plain view) had a constant mean of 18.7 cd/m^2 and luminance range of 28.7 cd/m^2 (Michelson contrast = 0.768). The lower-contrast side (which was perceived to contain transparency) had a fixed value of luminance range of 14.3 cd/m^2, but could take one of three values of mean luminance: 12.9, 18.7, or 24.4 cd/m^2. (These corresponded to Michelson contrasts of 0.554, 0.384, and 0.294, respectively.) The ratio of luminance ranges for these three simulated transparent layers yielded a constant value of α (= 0.5); whereas the ratio of Michelson contrasts yielded three distinct values of α_c (namely, 0.72, 0.5, and 0.38, respectively).

The high- and low-contrast half of the display each contained a circular patch of homogeneous luminance. The two circular patches were given a near disparity of 4.9 min of arc relative to the background texture, and were thus perceived to be floating between the depth plane of the background texture and that of the transparent layer (see the schematic in Fig. 6). As mentioned above, the patch against the high-contrast background was assigned to be the standard. Its luminance was randomly set to one of six possible values: 2.2, 8.8, 15.4, 22.1, 28.7, and 35.3 cd/m^2 (three darker than the mean luminance of the background surface, and three lighter). The luminance of the comparison patch—seen through transparency—was to be adjusted by the observer. The transparent layer (and hence the comparison patch) was equally likely to appear on the left or the right half of the display. On any given trial, the assignment of standard and comparison was immediately apparent to the observers, as the comparison patch was always the one behind the transparent layer and, moreover, only the luminance of the comparison was under their control.

Two background texture patterns were used: sinusoidal gratings and a binary random-dot texture. The motivation behind using two different textured backgrounds was to ensure the robustness of the results—in particular, that the lightness matches are determined by the perceived surface properties of the transparent layers, and not by specific attributes of the luminance distributions (e.g., the continuous range of luminance values within the gratings versus the two discrete values in the random-dot texture). The sinusoidal gratings had a period of 1.01° of visual angle, and were oriented at +45°. Their phase was randomly set on every trial. The random-dot texture had rectangular elements, elongated in the horizontal direction ($0.85° \times 0.25°$). A new sample of the random-dot texture was generated for each trial. The presence of the vertical contrast border generated

X-junctions in each half image of the stereoscopic display, thereby reinforcing the percept of transparency.

The stimuli were presented on a linearized high-resolution 22″ (*Lacie Blue*) monitor. The monitor was calibrated so that screen luminance values (ranging from 0 to 44 cd/m^2) were linearly related to the 8-bit look-up table values. The stimuli were viewed through a mirror stereoscope, from an optical distance of 106 cm.

Procedure. Each observer performed adjustments in six separate sessions—one session for each of the three simulated transparent layers and the two background textures. On each trial, the luminance of the standard patch was randomly set to one of the six pre-determined values. Observers adjusted the luminance of the comparison patch in order to match the lightness of the standard. The task thus required observers to estimate what the comparison patch—seen through a transparent layer—would look like in plain view. Within each session, observers performed five experimental adjustments for each of the six preset values of standard luminance. These were preceded by six practice adjustments, one for each value of standard luminance.

3.1.2. Results

The data for the three observers are plotted in Fig. 8. No systematic differences were obtained across the two background textures, and the data are shown averaged over the two. Each data point thus corresponds to the mean of 10 adjustments by an observer. The three different curves correspond to matches made through the three different transparent layers, in order to match the lightness of the same set of standards. The middle curve corresponds to the display in which the region of transparency has the same mean luminance as the high-contrast background. The highest curve corresponds to the display in which the region of transparency has greater mean luminance (corresponding to a 'lightening' transparent layer), and the lowest curve corresponds to the display in which the region of transparency has lower mean luminance (corresponding to a 'darkening' transparent layer).

Two aspects of the data are prominent. First, the mapping from standard luminance values to the matched comparison values is linear (mean R^2-value of linear fits = 0.955; see Table 1). Thus, a linear form (i.e., Eq. (6)) predicts the perceived transformation in lightness produced by the presence of an overlying transparent layer. In particular, the linear slopes of these curves capture the way in which the presence of a transparent layer compresses the range of underlying luminance values, due to its partial transmittance. Second, within each observer's data, the slopes of the three curves are close to identical—with the pairwise differences to the baseline slope (i.e., the one obtained with same-luminance transparency display) not being

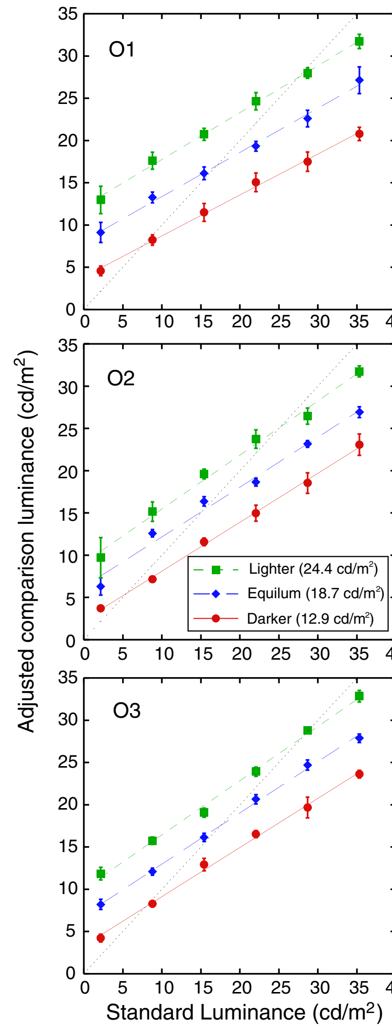

Fig. 8. The lightness-matching results of Experiment 1. The three curves correspond to matches made through the three different transparent layers in order to match the lightness of the same set of standard patches. The bottom curve corresponds to matches through the 'darkening' transparent layer, the top curve corresponds to matches through the 'lightening' transparent layer, and the middle curve to the one that preserves mean luminance. The data curves are linear and have essentially identical slopes—consistent with Metelli's α.

statistically significant. [9] The conclusion of essentially-identical slopes is also supported by the fact that a bivariate regression model that incorporated the 3 transparent layers as levels of a second variable yielded excellent fits (R^2-values of 0.953, 0.956, and 0.982 for observers O1, O2, and O3, respectively). Moreover, the small differences in slopes that do exist are in the

[9] In the case of the middle curve, where the region of transparency has the same mean luminance as the high-contrast region, α and α_c make the same prediction for the slope. Thus, the fitted slope in this case provides a baseline for each observer, against which the slopes of the other two data curves may be compared.

Table 1
Slopes and R^2 values for the linear fits to the lightness-through-transparency matches in Experiment 1

Predictions		Observer					
		O1		O2		O3	
α	α_c	$\hat{\alpha}$	R^2	$\hat{\alpha}$	R^2	$\hat{\alpha}$	R^2
0.5	0.72	0.484	0.937	0.579	0.960	0.581	0.972
0.5	0.5	0.523	0.931	0.591	0.965	0.606	0.984
0.5	0.38	0.555	0.939	0.639	0.925	0.642	0.982

Slope predictions based on α and α_c are shown on the left. The slopes of the matches through the three transparent layers are close-to-identical, and the pairwise differences to the baseline condition (middle row) are not statistically significant. (Moreover, the small differences that do exist are in the opposite direction from those predicted by α_c.) Thus the pattern of slopes is well captured by Metelli's α, but not by α_c.

opposite direction relative to the contrast-based prediction. As discussed above, slopes based on contrast (i.e., α_c) should exhibit a systematic *decrease* with increasing mean luminance in the region of transparency (i.e., the bottom-most curve should have the greatest slope, with the slope becoming successively shallower for the middle and top-most curves). Clearly, this prediction is not borne out in the data. Thus, the pattern of slopes is well captured by Metelli's α, but not by the contrast-based α_c (recall that Metelli's α is insensitive to the mean luminance in the region of transparency). Although there is some variability across observers in the specific values of the fitted slopes—with observers O2 and O3 overestimating it slightly (mean slopes 0.60 and 0.61, respectively), the important point is that the higher-order pattern of slopes is consistent with α (recall footnote 9).

These data also allow us to be more precise about the extent of lightness constancy through achromatic transparency. In the domain of color transparency, D'Zmura et al. (2000) reported that their observers underestimated the degree of convergence in color space due to the presence of an overlying color filter by almost 50%. (Recall that the degree of convergence is given by $1-\alpha$: zero convergence corresponds to the absence of an overlying transparent layer, i.e., $\alpha = 1$, and full convergence to the case of an opaque occluder, i.e., $\alpha = 0$; see Fig. 3.) The extent of lightness constancy in the current experiment is considerably better: one of the observers (O1) captures the degree of convergence almost perfectly, with only a slight underestimation of 4% ($\frac{1-\hat{\alpha}}{1-\alpha} = 0.96$), whereas the other two observers underestimate it by about 20% ($\frac{1-\hat{\alpha}}{1-\alpha} = 0.78$ and 0.8, respectively). One must be cautious, however, in attributing this difference in the extent of constancy solely to the difference in chromaticity in the two studies. D'Zmura et al.'s study used motion to reinforce the percept of an overlying transparent layer, whereas the current study used binocular disparity. It is thus also possible that binocular disparity is more effective in decomposing image luminance into two distinct layers than relative motion; and this factor too may be in part responsible for the improved constancy.

3.2. Experiment 2: Control for non-transparency-based factors

Can the pattern of lightness matches in Experiment 1 be explained by lightness factors that do not invoke the notion of transparency (i.e., that do not depend on a layered decomposition of image luminance into a background surface and an overlying transparent surface)? Two natural examples of such factors are simultaneous contrast and anchoring. For example, the results may be attributable to the fact that the surrounds of the standard and comparison patch have different mean luminances in two of the three transparency displays used. Similarly, if one considers the luminance distributions in the respective surrounds of the two patches, the highest luminances (i.e., a_{max} and p_{max}) of these two distributions are different. Thus, if the visual system anchors the highest luminances within each of these local "frameworks" to white (Gilchrist et al., 1999)—or uses separate "adaptive windows" in the two halves to perform the anchoring (Adelson, 2000)—this too might account for the pattern of matches.

In order to control for these lightness factors, the second experiment uses modified displays that have the same luminance distributions as those used in Experiment 1, but that suppress the percept of transparency. In particular, (a) the low-contrast and high-contrast half of each display are separated laterally, and (b) all disparities are set to zero (see Fig. 9). In these modified displays, then, the two circular patches are seen simply as lying against two different textured backgrounds, with no overlying transparent layer. The image factors influencing simultaneous contrast and local anchoring, on the other hand, are preserved. In particular, the mean luminance within the surround of each patch is preserved, and so is the highest luminance.

3.2.1. Methods

Observers. The same three observers participated as in Experiment 1.

Stimuli and apparatus. The stimuli were similar to Experiment 1, differing from them in only two ways. First, the high-contrast and low-contrast halves of each

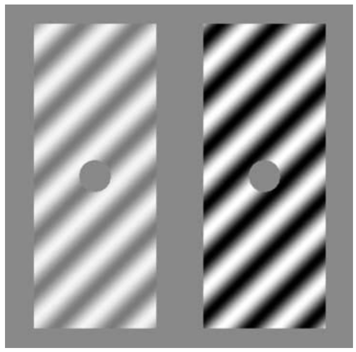

Fig. 9. The stimulus configuration used in the control experiment (Experiment 2). A lateral separation was introduced between the high- and low-contrast halves of the display, and all disparities were set to zero. Both of these manipulations were designed to suppress the percept of transparency, while preserving the distribution of luminance values around the standard and comparison patches.

display were laterally separated by 1.2 degrees of visual angle (see Fig. 9). Second, all disparities were set to zero. Both of these manipulations were designed to suppress the percept of an overlying transparent layer.

The same texture patterns, and the same set of values of mean luminance and luminance range were used within the two textured halves, as in Experiment 1. The high-contrast half had a fixed mean luminance of 18.7 cd/m^2 and luminance range of 28.7 cd/m^2 throughout. And the three versions of the lower-contrast half had the same luminance range of 14.3 cd/m^2, but different values of mean luminance, 12.9, 18.7, and 24.4 cd/m^2. Moreover, the same six values of luminance were used for the standard patch. As before, the lower-contrast half—and hence the comparison patch—was equally likely to appear on the left or right side of the display.

The apparatus was identical to Experiment 1. In order to maintain the same viewing conditions as Experiment 1, the stimuli were viewed through the mirror stereoscope even though they contained no binocular disparity.

Procedure. Each observer performed adjustments in three sessions—one session for each of the three values of mean luminance within the lower-contrast half. On each trial, the luminance of the standard patch was randomly set to one of the six pre-determined values. As in Experiment 1, observers adjusted the luminance of the comparison patch in order to match the lightness of the standard. (Unlike Experiment 1, however, the comparison patch no longer appeared to be seen through a transparent layer.) Within each session, observers performed six experimental adjustments for each setting of standard luminance. Three of these adjustments were made on displays containing the sinusoidal-gratings background, and three on displays with the random-dot texture. The experimental trials were preceded by 12 practice adjustments—one for each of the six values of standard luminance and the two textures.

3.2.2. Results

The data for the three observers are plotted in Fig. 10. As before, the three curves correspond to matches made against three different lower-contrast backgrounds, in order to match the lightness of the same set of standards.

The data are again linear (mean R^2-value of linear fits = 0.972; see Table 2). However, the magnitudes of the shifts from the identity function $y = x$ are quite small relative to those observed in Experiment 1 (compare with Fig. 8). In particular, the values of the fitted slopes (see Table 2) are now much closer to unity than in Experiment 1.

Thus, although factoring out transparency does leave some residual effects (the data curves do not lie perfectly on the identity function), the results clearly demonstrate that lightness factors such as simultaneous contrast and

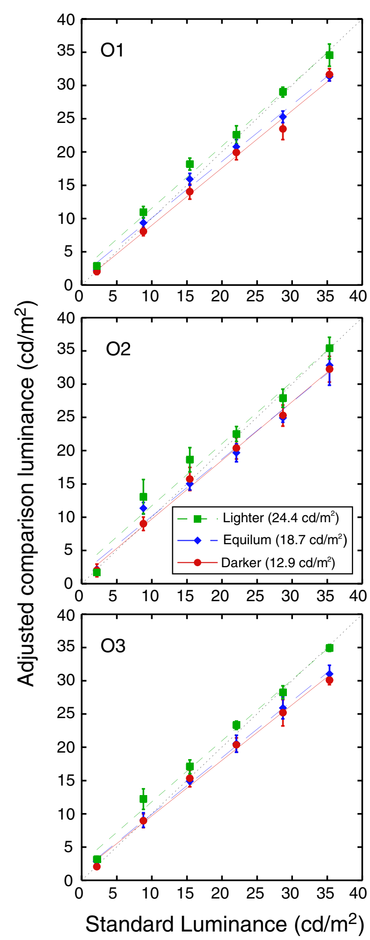

Fig. 10. The results of Experiment 2. The three curves correspond to lightness matches made when the background of the comparison display was darker, lighter, or had the same mean luminance as the background of the standard patch. Note that the data curves are now much closer to the identity function $y = x$ than in Experiment 1 (compare with Fig. 7), indicating that those results cannot simply be attributed to simultaneous contrast or anchoring within local frameworks.

Table 2
Slopes and R^2 values for the linear fits to the lightness-matching data in Experiment 2

Observer					
O1		O2		O3	
$\hat{\alpha}$	R^2	$\hat{\alpha}$	R^2	$\hat{\alpha}$	R^2
0.862	0.979	0.882	0.966	0.836	0.975
0.853	0.985	0.861	0.954	0.851	0.982
0.935	0.980	0.935	0.945	0.920	0.982

The slopes are now much closer to unity than in Experiment 1, indicating that the perceived transformation in lightness in that experiment cannot be attributed simply to simultaneous contrast or anchoring within local frameworks.

anchoring cannot, in themselves, explain the shifts obtained in Experiment 1. These shifts are thus indeed attributable to the visual system's attempting to 'discount' the contributions of an overlying transparent layer.

4. Directly matching the transparent layer

The results of Experiment 1 indicate that the linear form of Metelli's equations captures the perceived transformation of lightness values due to the presence of an overlying transparent layer. In particular, the mapping from standard luminance values in plain view to adjusted comparison values through transparency is linear, and the pattern of slopes is consistent with Metelli's α. These results appear to suggest, therefore, that Metelli's equations provide a perceptually valid model of transparency. Recall, however, that recent results on transmittance matching contradict this view (Singh & Anderson, 2002a, see also Robillotto et al., 2002). In particular, these results demonstrated that Metelli's α fails to capture perceived transmittance: Rather than being determined by the ratio of luminance differences (as in Metelli's solution for α, recall Eq. (3)), observers' transmittance matches are consistent with the ratio of contrasts—in other words, consistent with α_c (recall Eq. (8)), rather than α.

Applying this result to the three transparency displays used in Experiment 1 yields the following prediction: Even though lightness matches made through these three transparent layers yield close-to-identical slopes (consistent with Metelli's α), direct matches on the same three transparent layers should yield three very different transmittance matches. Specifically, transmittance matches on the darker layer should yield systematically higher settings of luminance range than those on the lighter layer. (Recall that the three transparent layers are defined by the same value of luminance range—and hence same α—but different Michelson contrasts—hence different α_c's.) Experiment 3 directly tests this prediction using a transmittance-matching method employed previously by Singh and Anderson (2002a).

4.1. Experiment 3

4.1.1. Methods

Observers. The same three observers participated as in the first two experiments.

Stimuli and apparatus. The stimuli consisted of (i) a standard transparency display with the same configuration as in Experiment 1 but without the two circular patches, and (ii) a comparison display located 1.3° below it. The standard display contained the same disparities and the same set of photometric values as in Experiment 1—thereby generating the same set of simulated transparent layers. In particular, the high-contrast side had a constant mean luminance of 18.7 cd/m^2 and luminance range of 28.7 cd/m^2. The lower-contrast side had a constant luminance range of 14.3 cd/m^2; but could take one of three possible values of mean luminance 12.9, 18.7, and 24.4 cd/m^2. As in Experiment 1, the ratio of luminance ranges for the three transparent layers yielded a constant value of $\alpha = 0.5$; whereas the ratio of Michelson contrasts yielded three distinct values of α_c, namely, 0.72, 0.5, and 0.38, respectively.

The comparison transparency display had a configuration used previously by Singh and Anderson (2002a) for transmittance matching. It consisted of a large circular disk (diameter = 6.05°) with a high-contrast sinusoidal grating, and a smaller lower-contrast disk (diameter = 2.4°) placed in its center (see Fig. 11). The period and orientation of the gratings were identical to Experiment 1. The lower-contrast disk was given a near disparity of 7.25 min of arc, and was perceived as a transparent disk floating in front of the high-contrast background. The high-contrast grating had the same mean luminance (18.7 cd/m^2) and luminance range (28.7 cd/m^2) as the high-contrast half of the standard display. The low-contrast center had the same mean luminance (18.7 cd/m^2); but its luminance range was to be adjusted by the observers. The apparatus and viewing conditions were identical to the first two experiments.

Procedure. On each trial, the mean luminance of the lower-contrast side of the standard display was set to one of the three preset values. Observers adjusted the luminance range within the central disk of the comparison display in order to match the transmittance of the

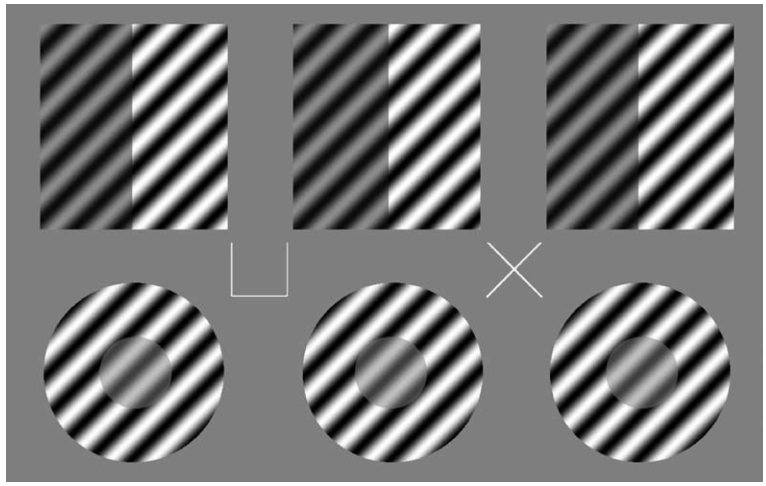

Fig. 11. The stereoscopic stimulus configuration used in Experiment 3 for transmittance matching. (Cross fusers should fuse the two right images, parallel fusers the two left images.) Observers adjusted the luminance range within the central disk of the comparison display (bottom) in order to match transmittance of the transparent layer in the standard display (top). The standard displays consisted of the same set of transparent layers used in Experiment 1, with either the grating or the random-dot texture in the background.

transparent layer in the standard display. Observers performed adjustments in two sessions. Within each session, they performed 10 experimental adjustments on each of the three standard transparent layers. Of these, five used the sinusoidal-grating background in the standard display, and five used the random-dot texture background. These were preceded by six practice adjustments.

4.1.2. Results

The data for the three observers are plotted in Fig. 12. It is clear from these graphs that observers perceive the three transparent layers as having very different transmittances—so that changing the mean luminance in the region of transparency (while preserving its luminance range) alters perceived transmittance. Thus, as in Singh and Anderson (2002a), Metelli's α—which is insensitive to mean luminance (see the dashed lines in Fig. 12)—fails to predict perceived transmittance. As in Experiment 1, the display in which the region of transparency has the same mean luminance as the background (the middle data point in each graph) acts as the baseline for each observer, because α and α_c make identical predictions for this display. Relative to their matches in this baseline case, observers consistently set a higher value of luminance range within the comparison display in order to match the transmittance of the darker transparent layer (which has higher contrast), and a lower value to match the transmittance of the lighter layer (which has lower contrast). Thus, although two of the observers overestimate the transmittance somewhat, the pattern of transmittance matches is consistent not with Metelli's α, but rather with the relative *contrast* in the region of transparency. Importantly, this occurs despite the fact that Metelli's α does capture the pattern of slopes of lightness matches through transparency. This point is examined more closely in Section 5.

5. Reconciling matches on the transparent and the underlying surface

The experimental results indicate that Metelli's equations capture the pattern of lightness matches through a transparent layer (Experiment 1)—even though they *fail* to capture the perceived transmittance of the transparent layer itself (Experiment 3; see also Singh & Anderson, 2002a). As discussed earlier, the term α plays two roles in Metelli's model: (i) α is the slope of the mapping from luminances projected from surface patches in plain view to luminances projected from the same patches through a transparent layer; and (ii) α is the transmittance of the transparent layer. The results demonstrate that, perceptually speaking, α plays role (i), but not role (ii). In other words, although the slope of the mapping from standard luminances in plain view to adjusted luminances through transparency is consistent with Metelli's α, the perceived transmittance of the transparent layer is not; rather, it is consistent with α_c. Does this entail a lack of internal consistency in the visual system's representation of the two layers in transparency? Or can these two facts be reconciled quantitatively?

In what follows, we will see that the two sets of results can in fact be reconciled in a mutually consistent manner. To motivate a reconciliation, consider first a visual context analogous to the lightness-matching context of Experiment 1—but involving an illumination change rather than transparency (e.g., half of the textured background is under shadow). Because illumination

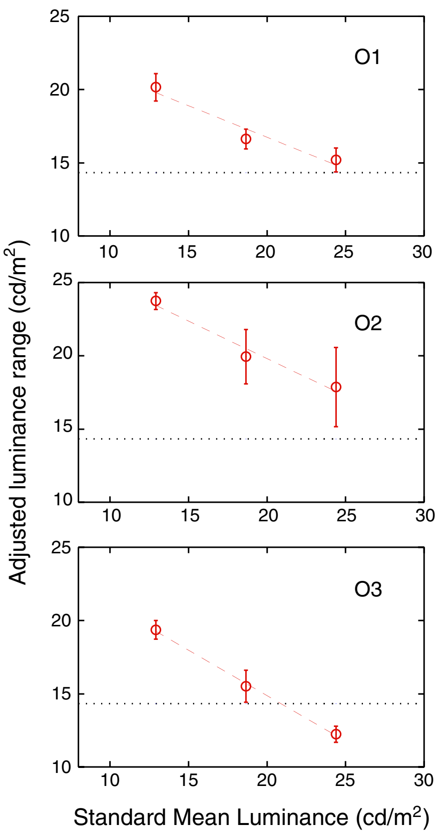

Fig. 12. The results of Experiment 3. Relative to the baseline condition in which α and α_c make the same prediction (middle data point), observers set a higher value of luminance range in the comparison display in order to match the transmittance of the darker filter (left data point), and a lower value in order to match the transmittance of the lighter filter (right data point). This pattern of results is thus consistent with α_c, but not with Metelli's α which predicts the same match in all three cases (see the dashed lines).

changes preserve contrast, a natural strategy in performing the lightness matching task is to equate the contrast of the comparison patch with its background with the contrast of the standard patch with its background. Let $a_{\max}$ and $a_{\min}$ be the peak and trough of the texture on the side under the default illumination, and $p_{\max}$ and $p_{\min}$ be the peak and trough of the texture under shadow. Then, if x is the luminance of the standard patch, and y is the adjusted luminance of the comparison, the lightness matches are defined by solutions to the following equations:

$$\frac{y-p_{\max}}{y+p_{\max}}=\frac{x-a_{\max}}{x+a_{\max}}, \tag{9}$$

$$\frac{y-p_{\min}}{y+p_{\min}}=\frac{x-a_{\min}}{x+a_{\min}}. \tag{10}$$

The case of lightness through transparency is similar, except that contrast is lowered, rather than preserved, in the region of transparency. Thus, rather than equating the contrasts of the comparison and standard patches with their respective backgrounds, the ratio of their values must be equated to the value of perceived transmittance, α_c (which essentially captures the extent of contrast reduction), i.e.:

$$\frac{y-p_{\max}}{y+p_{\max}}=\alpha_c\left(\frac{x-a_{\max}}{x+a_{\max}}\right), \tag{11}$$

$$\frac{y-p_{\min}}{y+p_{\min}}=\alpha_c\left(\frac{x-a_{\min}}{x+a_{\min}}\right). \tag{12}$$

Expanding the expression for α_c:

$$\frac{y-p_{\max}}{y+p_{\max}}=\left(\frac{\frac{p_{\max}-p_{\min}}{p_{\max}+p_{\min}}}{\frac{a_{\max}-a_{\min}}{a_{\max}+a_{\min}}}\right)\left(\frac{x-a_{\max}}{x+a_{\max}}\right), \tag{13}$$

$$\frac{y-p_{\min}}{y+p_{\min}}=\left(\frac{\frac{p_{\max}-p_{\min}}{p_{\max}+p_{\min}}}{\frac{a_{\max}-a_{\min}}{a_{\max}+a_{\min}}}\right)\left(\frac{x-a_{\min}}{x+a_{\min}}\right), \tag{14}$$

Simplifying these two equations yields:

$$(Dx+Sa_{\max})y=p_{\max}(Sx+Da_{\max}), \tag{15}$$

$$(Dx+Sa_{\min})y=p_{\min}(Sx+Da_{\min}) \tag{16}$$

where

$$D=\frac{a_{\max}-a_{\min}}{a_{\max}+a_{\min}}-\frac{p_{\max}-p_{\min}}{p_{\max}+p_{\min}}, \tag{17}$$

$$S=\frac{a_{\max}-a_{\min}}{a_{\max}+a_{\min}}+\frac{p_{\max}-p_{\min}}{p_{\max}+p_{\min}}. \tag{18}$$

are the difference and sum, respectively, of the Michelson contrasts. Subtracting Eq. (16) from Eq. (15), gives:

$$S(a_{\max}-a_{\min})y=S(p_{\max}-p_{\min})x+D(a_{\max}p_{\max}-a_{\min}p_{\min}) \tag{19}$$

i.e.,

$$\begin{aligned} y&=\left(\frac{p_{\max}-p_{\min}}{a_{\max}-a_{\min}}\right)x+\frac{D}{S}\left(\frac{a_{\max}p_{\max}-a_{\min}p_{\min}}{a_{\max}-a_{\min}}\right)\\ &=\left(\frac{p_{\max}-p_{\min}}{a_{\max}-a_{\min}}\right)x+\left(\frac{a_{\max}p_{\min}-a_{\min}p_{\max}}{a_{\max}-a_{\min}}\right)\\ &=\alpha x+f \end{aligned} \tag{20}$$

But this is precisely Metelli's equation! (Recall that $\alpha=\frac{p-q}{a-b}$ and $f=\frac{aq-bp}{a-b}$; Eqs. (3) and (7), respectively.) In other words, performing the lightness matching task by equating the ratio of Michelson contrasts yields Metelli's equations. This is somewhat surprising, given that—as we already know—equating the ratio of luminance ranges also yields Metelli's equations. That is, the two equations:

$$y-p_{\max}=\left(\frac{p_{\max}-p_{\min}}{a_{\max}-a_{\min}}\right)(x-a_{\max}) \tag{21}$$

$$y-p_{\min}=\left(\frac{p_{\max}-p_{\min}}{a_{\max}-a_{\min}}\right)(x-a_{\min}) \tag{22}$$

together also yield Metelli's equation $y = \alpha x + f$. Indeed, this is simply a consequence of the fact that the expression for α—the ratio of luminance ranges—was derived from Metelli's equations.

The outcome of this analysis is that, irrespective of whether the visual system performs the lightness-through-transparency matches by equating ratios of luminance differences (consistent with Metelli's α), or by equating ratios of contrasts (consistent with α_c), it is forced into the analytically identical solution. This, then, clarifies why the pattern of lightness matches through transparency is consistent with Metelli's equations (Experiment 1)—even though perceived transmittance is not predicted by Metelli's α (Experiment 3). *The work in both cases is really being done by contrast.* In the first case, performing the lightness-matching task by equating the ratio of Michelson contrasts yields matches that are analytically identical to those predicted by Metelli's equations (with slopes given by α); in the second case, it yields transmittance matches that deviate systematically from Metelli's α—being given instead by the ratio of contrasts, α_c. Thus, the success of Metelli's equations in predicting lightness through transparency, as well as their failure to predict perceived transmittance, can *both* be explained in terms of mechanisms whose main currency is contrast.

5.1. Contrast reduction and contrast normalization

The above analysis also helps resolve a question posed by Gilchrist and Jacobsen (1983) on veiling luminance—namely, how can contrast-based mechanisms account for lightness constancy under veiling luminance, given that veiling luminance lowers, rather than preserves, contrast? Since veiling luminance is a fixed luminance increment to a scene, or a portion thereof, the transformation from luminances in plain view to luminances under veiling luminance is given simply by:

$$p = a + f \tag{23}$$

This is essentially Eq. (6) with the value of α set to 1. As before, let x be the luminance of the standard patch in plain view, and y be the luminance of the comparison patch, now through veiling luminance. Then, as in the case of transparency, if the visual system equates the ratio of contrasts of the two patches against their respective backgrounds to the extent of contrast reduction α_c—as in Eqs. (11) and (12)—it would again be led to the correct solution. (Note that in the case of veiling luminance, the ratio of contrasts $\alpha_c = \frac{p_{\text{contrast}}}{a_{\text{contrast}}}$ simply equals the inverse ratio of the values mean luminance $\frac{a_{\text{mean}}}{p_{\text{mean}}}$.) Thus, adopting a strategy that explicitly factors in the extent of contrast reduction produced by the presence of veiling luminance yields the physically correct match.

Such 'factoring in' of the extent of contrast reduction may be readily achieved by known mechanisms of contrast normalization (e.g., Chubb, Sperling, & Solomon, 1989; D'Zmura & Singer, 1999). The role of such mechanisms in color constancy is of course not a new suggestion (see, e.g., Brown & MacLeod, 1992; D'Zmura et al., 2000; Hagedorn & D'Zmura, 2000). The current results and analysis, however, make two important points concerning the role of such mechanisms in lightness constancy. First, the large differences between the results of Experiment 1 and 2 demonstrate that the operation of these mechanisms is closely tied to the computation of a layered representation. In particular, the normalization effects implicit in the above lightness-matching analysis largely disappeared in Experiment 2 when the interpretation of transparency was suppressed (even though the contrasts of the respective surrounds were preserved). Second, previous proposals have not explicitly distinguished between different measures of 'contrast' relative to which normalization occurs, and their respective predictions for lightness or color constancy (e.g., luminance range versus Michelson contrast in the achromatic domain). This is most likely because these proposals were made in the chromatic domain, where the notion of contrast is substantially more complex. What is novel in the current analysis, therefore, is the result that analytically identical matches are predicted irrespective of whether the visual system performs the normalization using luminance differences (consistent with Metelli's solution for α), or using Michelson contrast (consistent with α_c). As we have seen, this result has important implications for quantitative models of transparency—in particular, for reconciling the failure of Metelli's equations in capturing perceived transmittance with its success in capturing lightness through transparency.

The analysis also sheds new light on recent results by Kasrai and Kingdom (2001) who investigated the precision and accuracy of perceived transparency using modified versions of Metelli's displays containing three background luminances and three luminances in the region of overlay. On each trial, the three background luminances and two of the overlay luminances were fixed. The observers' task was to adjust the third overlay luminance in order to generate a percept of transparency. Kasrai and Kingdom (2001) found that, although there was a relatively wide range of values that give rise to a percept of transparency, observers' adjustments were well predicted by Metelli's model. Moreover, a model based on the ratio of Michelson contrasts—motivated by Singh and Anderson's (2002a) transmittance-matching results—also predicted their results equally well. Although Kasrai and Kingdom did not comment further on the equivalent success of the two models, the analysis above makes it clear that it is a direct consequence of the fact that the two models in fact make identical predictions for their task. Even though their task required observers to generate a percept of homogeneous transparency—not to perform lightness-through-transparency matching—their predictions for the adjustable patch based on Metelli's model were nevertheless derived by equating ratios of luminance differences, whereas their predictions based on Singh and Anderson's model were derived by equating ratios of Michelson contrasts. As we have seen above, these two cases generate analytically identical predictions.

6. Discussion

The results of the lightness-through-transparency experiment (Experiment 1) indicate that when observers match the lightness of surface patches viewed through transparency, their matches are consistent with Metelli's equations and exhibit a high degree of constancy (96% for one observer; 78–80% for the other two). In particular, the perceptual mapping from standard luminances in plain view to the adjusted luminances through transparency is linear, and the pattern of slopes is consistent with Metelli's α. However, consistent with previous results (Singh & Anderson, 2002a) when observers directly match the transmittance of these transparent layers, their matches deviate systematically from Metelli's α, and are predicted instead by the ratio of contrasts α_c (Experiment 3). Taken together, these results falsify a basic assumption in Metelli's model, namely, that the slope of the linear mapping is quantitatively equal to the perceived transmittance of the transparent layer. They also emphasize the fact that, in order to fully test quantitative models of transparency, it is not sufficient to obtain matches on underlying surface alone, or on the transparent surface alone. The *combination* of the two sets of matches—on both the underlying and the transparent surface—yields a more complete picture than was available from looking at either set of matches in isolation.

Although the visual system's assignment of surface attributes to the two layers—the underlying opaque surface and the overlying transparent surface—initially appear to be contradictory (with the former, but not the latter, given by Metelli's model), formal analysis reveals a deeper mutual consistency. In particular, both sets of results can be explained by a model in which ratios of contrast are primary:

(a) The perceived transmittance of a transparent layer is computed using the ratio of contrasts, α_c.
(b) Lightness matching through transparency is performed by equating the ratio of contrasts of the standard and comparison patches against their respective backgrounds, to the value of perceived transmittance (which, effectively captures the extent of contrast reduction).

In particular, using the strategy in (b) above—which is a simple generalization of lightness matching under a change in illumination—yields predictions that are analytically identical to those of Metelli's equations. Apart from quantitatively reconciling the matching data on the underlying and overlying surfaces, the analysis also provides a principled reason for why Metelli's equations should provide a perceptually valid model of lightness through transparency. Researchers have sometimes expressed surprise (and/or skepticism) concerning why a rather simplistic physical model (namely, the episcotister setup) should yield perceptually valid equations. The analysis here provides an independent principled reason—based on known properties of visual mechanisms, rather than a physical model—for why this should be so.

One should bear in mind, though, that the above result yielding Metelli's equations for lightness through transparency was derived using a specific measure of contrast with respect to which normalization is performed, namely, Michelson contrast. Using Michelson contrast is certainly reasonable, given that it is a commonly used measure that is known to do a good job of capturing apparent contrast for a large class of textured displays. Moreover, in Singh and Anderson's (2002a) study, perceived transmittance was found to scale systematically with the Michelson contrast in the region of transparency. However, the physical determinants of apparent contrast constitute a complex and long-standing problem. In addition to Michelson contrast, various other measures of contrast have been proposed (e.g., RMS contrast, ratio contrast, Whittle contrast, King-Smith and Kulikowski contrast) and, unfortunately, no one measure is known to capture apparent contrast universally (see, e.g., Moulden, Kingdom, & Gatley, 1990; Peli, 1997). Indeed, in a recent study, Robillotto et al. (2002) found that although perceived transmittance scaled with the perceived contrast in the region of transparency, none of the standard measures of contrast could adequately capture perceived contrast for their displays (which contained complex and variegated textures). In such cases—assuming perfect internal consistency in the visual system's representation of the transparent and the underlying surface—an intriguing prediction is raised, namely, that lightness through transparency would be determined by Metelli's equations only to the extent that Michelson contrast predicts perceived transmittance of the transparent layer. In other words, if a different measure of contrast is appropriate, then an analysis completely analogous to the one above (Eqs. (11)–(20)) would yield alternate predictions for lightness-through-transparency matches—which may or may not correspond precisely to Metelli's equations. What the analysis above shows is that using either Michelson contrast or luminance differences predicts lightness matches that are, in fact, analytically identical to Metelli's equations. Finally, it should be noted that there are contexts—for example, when a partially-transmissive surface contains a light-scattering component that produces image blur (Singh & Anderson, 2002b), or a specular component—where the assignment of perceived transmittance becomes considerably more complicated. Lightness constancy through transparency for such cases, and its link to the perceived surface properties of the partially-transmissive layer, have yet to be systematically investigated.

7. Conclusions

When matching the lightness of surface patches seen through transparency, observers' matches are consistent with Metelli's equations: the mapping from standard luminance values (in plain view) to comparison luminance values (through transparency) is linear, and the pattern of slopes is consistent with Metelli's α. However, when matching the transmittance of the transparent layer itself, observers' matches deviate systematically from the predictions of Metelli's α, and are consistent instead with the ratio of contrasts α_c. Although these two sets of results appear to be contradictory, the analysis presented here demonstrates that both results can in fact be explained by a model based on the ratio-of-contrasts. This indicates that the ratio-of-contrasts model is primary as a perceptual model of transparency. The partial success of Metelli's model, on the other hand, may be epiphenomenal—being a by-product of the fact that its predictions for lightness through transparency happen to converge with those of the ratio-of-contrasts model.

Acknowledgements

This work was supported by NSF grant BCS-0216944. I am grateful to Jacob Feldman, Randy Gallistel, Eileen Kowler, Thomas Papathomas, and two anonymous reviewers for helpful comments on previous versions of the manuscript.

References

Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (pp. 339–351). Cambridge, MA: MIT Press.

Adelson, E. H., & Anandan, P. (1990). Ordinal characteristics of transparency. AAAI-90 Workshop on Qualitative vision, July 29, 1990, Boston, MA.

Anderson, B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images. *Perception, 26*, 419–453.

Anderson, B. L. (2003). The role of occlusion in the perception of depth, lightness, and opacity. *Psychological Review, 110*, 785–801.

Beck, J., Prazdny, K., & Ivry, R. (1984). The perception of transparency with achromatic colors. *Perception and Psychophysics, 35*, 407–422.

Brown, R. O., & MacLeod, D. I. A. (1992). *Saturation and color constancy. Advances in color vision technical digest*. Washington, DC: Optical Society of America, pp. 110–111.

Chubb, C., Sperling, G., & Solomon, J. (1989). Texture interactions determine perceived contrast. *Proceedings of the National Academy of Sciences, 86*, 9631–9635.

D'Zmura, M., Colantoni, P., Knoblauch, K., & Laget, B. (1997). Color transparency. *Perception, 26*, 471–492.

D'Zmura, M., Rinner, O., & Gegenfurtner, K. R. (2000). The colors seen behind transparent filters. *Perception, 29*, 911–926.

D'Zmura, M., & Singer, B. (1999). Contrast gain control. In L. T. Sharpe & K. R. Gegenfurtner (Eds.), *Color vision: From genes to perception*. Cambridge, UK: Cambridge University Press.

Faul, F., & Ekroll, V. (2002). Psychophysical model of chromatic perceptual transparency based on substractive color mixture. *Journal of the Optical Society of America A, 19*, 1084–1095.

Gerbino, W. (1994). Achromatic transparency. In A. L. Gilchrist (Ed.), *Lightness, brightness, and transparency* (pp. 215–255). Hillsdale, NJ: Erlbaum.

Gerbino, W., Stultiens, C. I., Troost, J. M., & de Weert, C. M. (1990). Transparent layer constancy. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 3–20.

Gilchrist, A. L. (1979). The perception of surface blacks and whites. *Scientific American, 240*, 112–123.

Gilchrist, A. L., & Jacobsen, A. (1983). Lightness constancy through a veiling luminance. *Journal of Experimental Psychology: Human Perception and Performance, 9*, 936–944.

Gilchrist, A. L., Kossyfidis, C., Agostini, T., Li, X., Bonato, F., Cataliotti, R., & Spehar, B. (1999). An anchoring theory of lightness perception. *Psychological Review, 106*, 795–834.

Hagedorn, J., & D'Zmura, M. (2000). Color appearance of surfaces viewed through fog. *Perception, 29*, 1169–1184.

Kasrai, R., & Kingdom, F. (2001). Precision, accuracy, and range of perceived achromatic transparency. *Journal of the Optical Society of America A, 18*, 1–11.

Kraft, J. M., Maloney, S. I., & Brainard, D. H. (2002). Surface-illumination ambiguity and color constancy: Effects of scene complexity and depth cues. *Perception, 31*, 247–263.

Kubelka, P., & Munk, F. (1931). Ein beitrag zur optik der farbanstriche [An article on the optics of paint layers]. *Zeitschrift für Technische Physik, 12*, 112.

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America, 61*, 1–11.

Mahadev, S., & Henry, R. C. (1999). Application of a color-appearance model to vision through atmospheric haze. *Color Research and Application, 24*, 112–120.

Maloney, L. T., & Yang, J. N. (2003). The illumination estimation hypothesis and surface color perception. In R. Mausfeld & D. Heyer (Eds.), *Colour: Connecting the mind to the physical world* (pp. 335–358). Oxford: Oxford University Press.

McClatchey, R. A., Fenn, R. W., Selby, J. E. A., Volz, F. E., & Garing, J. S. (1978). Optical properties of the atmosphere. In W. G. Driscoll (Ed.), *Handbook of optics* (pp. 14.1–14.65). New York: McGraw-Hill.

Metelli, F. (1970). An algebraic development of the theory of perceptual transparency. *Ergonomics, 13*, 59–66.

Metelli, F. (1974a). Achromatic color conditions in the perception of transparency. In R. B. MacLeod & H. L. Pick (Eds.), *Perception: Essays in honor of J.J. Gibson*. Ithaca, NY: Cornell University Press.

Metelli, F. (1974b). The perception of transparency. *Scientific American, 230*, 90–98.

Metelli, F. (1985). Stimulation and perception of transparency. *Psychological Research, 47*, 185–202.

Metelli, F., Da Pos, O., & Cavedon, A. (1985). Balanced and unbalanced, complete and partial transparency. *Perception and Psychophysics, 38*, 354–366.

Moulden, B., Kingdom, F., & Gatley, L. F. (1990). The standard deviation of luminance as a metric for contrast in random-dot images. *Perception, 19*, 79–101.

Nakauchi, S., Silfsten, P., Parkkinen, J., & Ussui, S. (1999). Computational theory of color transparency: Recovery of spectral properties for overlapping surfaces. *Journal of the Optical Society of America A, 16*, 2612–2624.

Nakayama, K., & Shimojo, S. (1990). Da vinci stereopsis: Depth and subjective occluding contours from unpaired image points. *Vision Research, 30*, 1811–1825.

Peli, E. (1997). In search of a contrast metric: Matching the perceived contrast of Gabor patches at different phases and bandwidths. *Vision Research, 37*, 3217–3224.

Richards, W. A., & Stevens, K. (1979). Efficient computations and representations of visible surfaces. Final Report Contract No. AFOSR-79-0200, Massachusetts Institute of Technology.

Robillotto, R., Khang, B., & Zaidi, Q. (2002). Sensory and physical determinants of perceived achromatic transparency. *Journal of Vision, 2*, 388–403.

Rutherford, M. D., & Brainard, D. H. (2002). Lightness constancy: A direct test of the illumination-estimation hypothesis. *Psychological Science, 13*, 142–149.

Schirillo, J. A., & Shevell, S. K. (1997). An account of brightness in complex scenes based on inferred illumination. *Perception, 26*, 507–518.

Schirillo, J. A., & Shevell, S. K. (2002). Articulation: brightness, apparent illumination, and contrast ratios. *Perception, 31*, 161–169.

Singh, M., & Anderson, B. L. (2002a). Toward a perceptual theory of transparency. *Psychological Review, 109*, 492–519.

Singh, M., & Anderson, B. L. (2002b). Perceptual assignment of opacity to translucent surfaces: The role of image blur. *Perception, 31*, 531–552.

Singh, M., & Hoffman, D. D. (1998). Part boundaries alter the perception of transparency. *Psychological Science, 9*, 370–378.

Singh, M., & Huang, X. (2003). Computing layered surface representations: An algorithm for detecting and separating transparent overlays. *Proceedings of Computer Vision and Pattern Recognition, Vol. II*, 11–18.

Somers, D. C., & Adelson, E. H. (1997). Junctions, transparency, and brightness. *Investigative Opthalmology and Vision Science* (Suppl.), *38*(4), 2126.

Yang, J. N., & Maloney, L. T. (2001). Illuminant cues in surface color perception: Tests of three candidate cues. *Vision Research, 41*, 2560–2581.