Available online at www.sciencedirect.com

SCIENCE DIRECT®

Vision Research 46 (2006) 879–894

Vision Research

www.elsevier.com/locate/visres

Photometric determinants of perceived transparency

Manish Singh [a,*], Barton L. Anderson [b]

[a] Department of Psychology and Center for Cognitive Science, Rutgers University, New Brunswick Campus, Piscataway, NJ 08854-8020, USA
[b] School of Psychology, University of New South Wales, Sydney, Australia

Received 31 July 2005; received in revised form 5 October 2005

Abstract

Photometric constraints for the perception of transparency were investigated using stereoscopic textured displays. A contrast discontinuity divided the textured displays into two lateral halves, with one (reference) half fixed. Observers adjusted the luminance range within the other (test) half in order to perform two tasks: (i) indicate the highest luminance range for which the test side is perceived to be transparent, and (ii) indicate the lowest luminance range for which the test side is seen as being in plain view. Settings were obtained for multiple values of test mean luminance, in order to map out the perceptual locus of transition between transparency and non-transparency. The results revealed a systematic violation of Metelli's magnitude constraint in predicting the percept of transparency. Observer settings were approximated instead by a constraint based on perceived contrast (which matched Michelson contrast for the textures used). The results also revealed large asymmetries between darkening and lightening transparency. When the test was darker than the reference, settings were highly consistent across observers and closely followed the Michelson-contrast prediction. When the test was lighter, however, there was greater variability across observers, with two observers exhibiting shifts toward Metelli's magnitude constraint. Moreover, each observer's setting reliability was significantly worse for lightening transparency than darkening transparency. These results suggest that (polarity-preserving) darkening serves as an additional cue to perceptual transparency.

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

1. Introduction

In the course of image projection, contributions of multiple scene variables are collapsed onto the rendered image. Any estimation of scene structure, therefore, must involve disentangling (whether explicitly or implicitly) the contributions of different scene variables. Computing 3D shape from image shading, for example, requires that the pattern of image intensity be analyzed in terms of the respective contributions of illumination, surface reflectance, and local surface orientation.

A particularly compelling case of image decomposition occurs in perceptual transparency. An example is shown in Fig. 1A, where the uniformly colored image region **P** is seen as containing two surfaces layered in depth—a black surface in the background and a mid-gray partially transmissive surface in front. Similarly, region **Q** is seen as containing a light-colored surface seen through a transparent mid-gray surface. The perceptual vividness of this decomposition underscores the fact that, unlike other forms of image decomposition, the perception of transparency involves a decomposition into two distinct *surfaces* or *material layers*—rather than, say, a surface and a pattern of shading.

A basic question concerning perceptual transparency is how the visual system is able to compute a representation of two surfaces along the same line of sight—one seen *through* the other—from a 2D array that varies along a single dimension, namely, image intensity. In articulating the constraints used to initiate the decomposition of image intensity into two distinct surfaces, researchers have consistently pointed to two classes of image properties: *photometric properties*, involving relative intensity values (Adelson, 2000; Adelson & Anandan, 1990; Anderson,

* Corresponding author. Tel.: +1 732 445 6714; fax: +1 732 445 6715.
E-mail address: manish@ruccs.rutgers.edu (M. Singh).

0042-6989/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2005.10.022

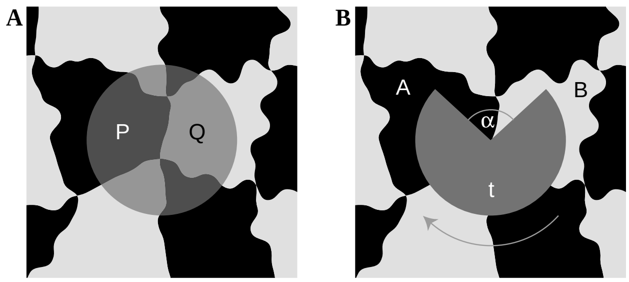

Fig. 1. (A) An example of perceptual transparency. The uniformly colored image region labeled **P** is perceived as containing two distinct surfaces: a black surface seen through a mid-gray transparent layer. Similarly, region **Q** is seen as a light-colored surface seen through transparency. (B) Metelli's episcotister model of transparency.

1997, 1999; Beck, Prazdny, & Ivry, 1984; Gerbino, Stultiens, Troost, & de Weert, 1990; Kasrai & Kingdom, 2001; Masin, 1997; Metelli, 1974; Metelli, Da Pos, & Cavedon, 1985; Robilotto, Khang, & Zaidi, 2002; Singh, 2004; Singh & Anderson, 2002), and *geometric properties*, involving spatial and configural factors (Adelson, 2000; Anderson, 1997; Beck & Ivry, 1988; Kanizsa, 1979; Kasrai & Kingdom, 2002; Metelli, 1974; Singh & Hoffman, 1998). In the current study, we investigate the role of photometric properties in initiating the perceptual decomposition that accompanies the percept of transparency, given a fixed geometric context.

1.1. Generative models

Metelli proposed a well-known model of transparency (Metelli, 1970, 1974, 1985; Metelli et al., 1985) using Talbot's equations of color mixing: The achromatic 'colors' of a partially transmissive (transparent) surface and an underlying (opaque) surface are 'mixed'—with the mixing proportions determined by the transmittance of the transparent surface (i.e., the proportion of incident light it transmits). Metelli modeled the transparent surface as a rotating *episcotister*—an opaque disk with an open sector (similar to a rotating fan, but with a single wide blade; see Fig. 1B). With sufficiently rapid rotations, the episcotister appears as a homogeneous partially transmissive surface (as in Fig. 1A)—its transmittance given by the relative size of the open sector.

Thus, if a and b are the reflectances of two underlying surface regions, t is the reflectance of the episcotister surface, and α is the relative size of its open sector, then the resulting 'color mixing' is described by Talbot's equations:

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

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

Since α and t are by construction identical—being the transmittance and reflectance of the same, homogeneous, transparent layer—the two equations are easily solved for α and t:

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

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

Although Metelli expressed his equations in terms of reflectance values, Gerbino et al. (1990) have shown that the same equations also follow in the luminance domain, under the assumption that the transparent layer and the background surface are illuminated equally.[1] Equations in terms of luminance values are more natural for perceptual theory since luminance values, not reflectances, constitute inputs to the visual system. Henceforth, we will assume the luminance formulation of Metelli's equations, i.e., we will treat a, b, p, q as luminance values.

Despite the simplicity of Metelli's equations, they have been shown to provide reasonable approximations to physical models of filters—which involve multiple reflections within the filter as well as between the filter and the background (Beck et al., 1984; Gerbino, 1994; Faul & Ekroll, 2002; Nakauchi, Silfsten, Parkkinen, & Ussui, 1999; Robilotto et al., 2002). The equations for the filter model converge to Eq. (5), as the illuminant strength gets increasingly higher (see Gerbino, 1994). Similarly, the fog model approximates the Metelli Eq. (4) under the assumption that the fog extinction coefficient is constant across wavelength (Hagedorn & D'Zmura, 2000)—a reasonable assumption for naturally occuring clouds and fog (McClatchey, Fenn, Selby, Volz, & Garing, 1978). Metelli's equations have also been extended to the chromatic domain, to define a convergene model of color transparency (D'Zmura, Colantoni, Knoblauch, & Laget, 1997).

1.2. Qualitative constraints for transparency

Based on solution (2) for α, Metelli derived two "qualitative constraints" for predicting the percept of transparency in the regions **P** and **Q** in his four-color displays (Metelli, 1970, 1974, 1985; Metelli et al., 1985). First, since α cannot be negative, the difference $p-q$ must have the same sign as $a-b$; so that, if region **A** is darker than **B**, then region **P** must be darker than **Q**. In other words, the **P/Q** border must have the same contrast polarity as the **A/B** border (the *polarity constraint*). Second, since α cannot exceed 1 (being the *proportion* of light transmitted), the magnitude of the difference $p-q$ must not exceed that of $a-b$ (the *magnitude constraint*).

- **C1.** Polarity constraint: $\text{sign}(p-q) = \text{sign}(a-b)$.
- **C2.** Magnitude constraint: $|p-q| \leqslant |a-b|$.

[1] Moreover, these equations apply equally if a veil or mesh is used instead of an episcotister—i.e., a surface with a large number of 'holes' that are too small to be resolved individually. In this case, the 'color mixing' takes place spatially, rather than temporally, with α corresponding to the areal density of the holes (Richards & Stevens, 1979).

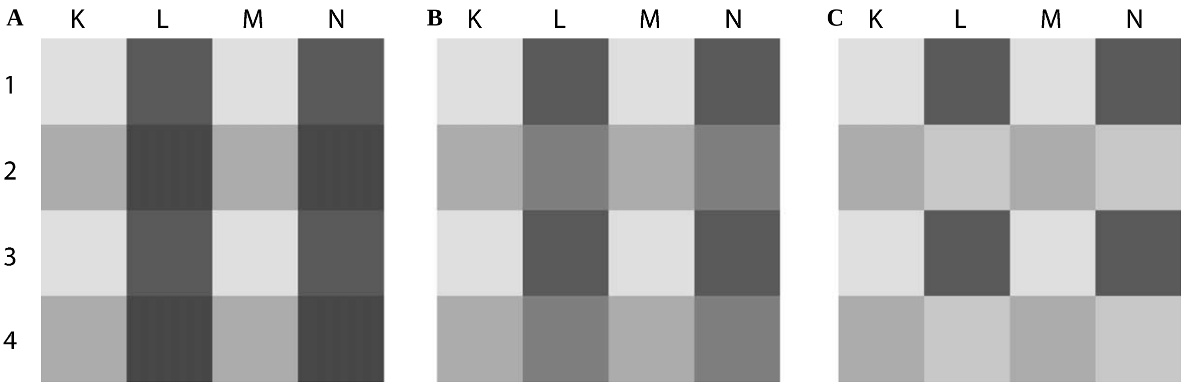

Fig. 2. Junction classification and transparency. (A) Contrast polarity is preserved across both horizontal and vertical contours; this display exhibits bistable depth layering. (B) Polarity is preserved across the horizontal, but not the vertical contours; this display exhibits a unique depth layering. (C) Polarity reverses across both sets of contours; this display is not perceived as containing a transparent overlap.

The role of the polarity constraint in perceptual transparency has been uncontroversial. Indeed, contrast polarity has been shown to provide a powerful constraint on image interpretation, and been used to devise striking illusions of lightness and transparency (e.g. Adelson, 1993, 2000; Anderson, 1997, 1999; Anderson & Winawer, 2005). Simple schemes for the classification of X-junctions, based on the polarity relationships across two intersecting edges, go a long way in predicting percepts of transparency and ordinal depth relationships between distinct layers (Adelson & Anandan, 1990; Beck & Ivry, 1988).[2] In Fig. 2, for instance, changing the common color of the four cells *L*2, *L*4, *N*2, and *N*4 changes the polarity relations at the X-junctions—thereby altering the percept of transparency. The display in A exhibits bistable transparency (either columns *L* and *N*, or rows 2 and 4, can be seen as transparent and in front), the display in B exhibits a unique transparency interpretation (only rows 2 and 4 can be seen as transparent), whereas the display in C is not perceived as containing a transparent overlap.

Although contrast-polarity relations at junctions can predict which contours—e.g., the vertical or horizontal contours in Fig. 2—will be perceived as boundaries of a transparent layer, they do not specify which *side* of a contour will be perceived as transparent. In Figs. 2A and B, for instance, the polarity relations do not in themselves explain why rows 2 and 4—but not rows 1 and 3—are seen as transparent. As this example makes it clear, a basic problem in the computation of transparency is determining which image regions correspond to surface patches seen in plain view, i.e., *without* any overlying transparency. A useful way to articulate the problem of transparency is thus in terms of an anchoring framework (Anderson, 1999, 2003), where the anchoring is performed on the dimension of transmittance or opacity (rather than lightness, e.g. Gilchrist et al., 1999). According to the transmittance-anchoring principle, the visual system anchors regions of highest contrast in an image to full transmittance (i.e., $\alpha = 1$), and decomposes regions of lower contrast into multiple layers. This principle thus predicts that the same image region may be perceived either as containing a transparent overlay, or as seen in plain view, depending on the contrasts of adjacent regions. Recent work also shows that transmittance anchoring has both a spatial and a temporal component, such that the visual system anchors the highest contrast within a spatio-temporal sequence to full transmittance (Anderson, Singh, & Meng, 2005).

The anchoring principle provides an elegant framework for inferring percepts of transparency and depth layering. However it does not settle the issue of what the appropriate measure of contrast magnitude is for the computation of perceptual transparency. Indeed, the issue of competing contrast measures does not arise in the context where the anchoring rule was originally articulated, i.e., transparency displays containing T-junctions (see Anderson, 1997, 1999). The various measures of contrast magnitude in this case all make the same prediction concerning which side has "lower contrast." As we will see below, this is not true in general.

As noted above, Metelli (1974) originally articulated his magnitude constraint in terms of reflectance differences: the magnitude of $p - q$ must be no greater than that of $a - b$. Since then, the notion of contrast magnitude relevant for transparency has been variously defined by researchers in terms of lightness differences (Beck et al., 1984), luminance differences (Gerbino et al., 1990), Michelson contrast (Singh & Anderson, 2002), and perceived contrast (Robilotto et al., 2002). Indeed, there is currently no general consensus on what the appropriate notion of contrast magnitude is for initiating percepts of transparency.[3]

[2] Polarity constraints at X-junctions provide strong, albeit entirely *local*, constraints on image interpretation. Computing and segmenting a transparent layer requires, in addition, that these local constraints be propagated and integrated along the bounding contour of a candidate transparent layer in order to check for mutual consistency (see Singh & Huang, 2003 for an implementation of such an integration scheme for the automatic segregation of transparent layers).

[3] This situation is further complicated by the fact that there is no known contrast measure that can universally capture *apparent contrast*—in both simple and complex textures.

(The latter two measures of contrast were, in fact, proposed based on experiments involving the matching of material properties of transparent layers, not on whether a layered surface structure involving transparency is perceived.) In the reported experiments, we map out the space of mean luminance and luminance range values that generate a percept of transparency, and compare these with observers' contrast matches, in order to better understand the photometric properties that the visual system uses in initiating a layered decomposition in transparency.

2. Assessing the validity of the magnitude constraint

Two considerations call into question the validity of Metelli's magnitude constraint **C**2. First, Metelli's derivation of this constraint was based on solution (3) for α (in particular, on the restriction $\alpha \leqslant 1$). As a result, the perceptual validity of this constraint depends critically on that of solution (3). In recent work, however, we have demonstrated that Metelli's solution fails to predict perceived transmittance (Singh & Anderson, 2002). Given fixed background luminances (e.g., a and b in Metelli's four-color displays), Metelli's solution predicts that perceived transmittance should be *independent* of the mean luminance within the region of transparency—depending only on the luminance difference or range $(p - q)$ therein. Contrary to this prediction, we found a systematic dependence of perceived transmittance on both luminance range and mean luminance within the region of transparency. For the textured backgrounds used in our study (sinusoidal gratings), this dependence was found to be captured by the Michelson contrast $\frac{L_{\max}-L_{\min}}{L_{\max}+L_{\min}}$ within the region of transparency. As a result, lightening filters appear more opaque than they should according to Metelli's α (i.e., when mean luminance in transparency region > mean luminance in background region), whereas darkening filters appear more transmissive than they should. A failure of Metelli's α was also found in Robilotto et al.'s (2002) study (see also Robilotto & Zaidi, 2004). However, all standard measures of contrast, including Michelson contrast, failed to predict perceived transmittance in their displays containing complex and variegated textures. Despite this, perceived transmittance was shown to be consistent with observers' matches of apparent contrast—and hence inconsistent with Metelli's α.

The second consideration that calls into question the validity of Metelli's magnitude constraint is that it was based entirely on a restriction on the solution for α. It did not make use of the solution for t, even though t is of course also restricted—in particular, t cannot be negative. However, Metelli and others (e.g., Beck et al., 1984) have argued that no simple, perceptually meaningful, prediction can be derived from restrictions on solution (4). We have previously shown, however, that this is not true (Singh & Anderson, 2002). Specifically, we demonstrated using the luminance formulation of Metelli's equations, that the constraint $t \geqslant 0$ (along with the restriction that $\alpha \leqslant 1$) implies that the region of transparency must have a lower Michelson contrast (i.e., $\frac{p-q}{p+q} \leqslant \frac{a-b}{a+b}$).[4] Thus, somewhat paradoxically, while Metelli's equations do not predict that perceived transmittance scales with Michelson contrast, they do predict that the region of transparency must have lower Michelson contrast relative to adjoining image regions.

The two considerations above suggest a natural alternative to Metelli's magnitude constraint, namely: in order for an image region to decompose perceptually into two layers, it must have lower Michelson contrast than adjoining image regions.

- **C3.** Michelson-contrast constraint: $\left|\frac{p-q}{p+q}\right| \leqslant \left|\frac{a-b}{a+b}\right|$.

For displays involving X-junctions, it is readily seen that the magnitude constraint **C2** and the Michelson-contrast constraint **C3** can make opposing predictions concerning which image region should undergo a perceptual decomposition into multiple layers. Specifically, if p, q, a, and b are the four luminance values at a polarity-preserving X-junction, it is straightforward to assign values such that $(p - q) \leqslant (a + b)$ but $\frac{p-q}{p+q} \geqslant \frac{a-b}{a+b}$. (A simple example is provided by setting $p = 12$, $q = 8$, $a = 48$, $b = 36$.) In this case, Metelli's magnitude constraint predicts that the **P/Q** region should be perceived as containing an overlying transparent layer, whereas the Michelson-contrast constraint predicts that the **A/B** region should decompose into two separate layers.

Finally, based on the results of Robilotto et al.'s (2002) study, we consider an additional constraint, based on perceived contrast:

- **C4.** Perceived-contrast constraint: $\mathrm{pcont}(p/q) < \mathrm{pcont}(a/b)$,

where pcont is short for perceived contrast. In many instances, constraint **C4** will coincide with constraint **C3**; however, as Robilotto et al.'s study showed, this need not always be the case—especially for complex textures. According to the more general constraint **C4**, image decomposition in perceptual transparency is determined by a lowering in *perceived contrast*—independently of how this might be defined for a given texture.

In the current experiments, we investigate the photometric relations that generate a percept of transparency, and compare these against the predictions based on luminance range, Michelson contrast, and perceived contrast.

3. Experiments

The class of stimulus displays we use is shown in Fig. 3. As in our previous experiments, we use stereoscopic displays with textured backgrounds—rather than a bipartite

[4] Indeed, one need only note that the numerator in the solution for t can be re-written as: $aq - bp = \frac{1}{2}(a+b)(p+q)[\frac{a-b}{a+b} - \frac{p-q}{p+q}]$, and that $(a+b)$ and $(p+q)$ cannot be negative.

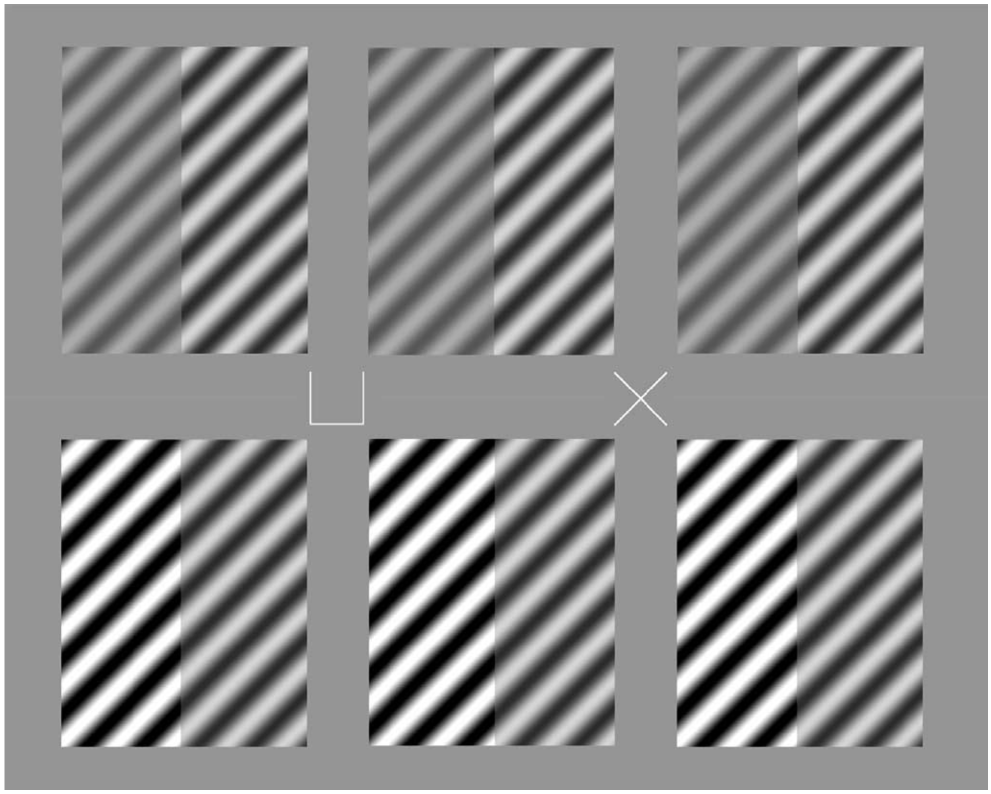

Fig. 3. Stereoscopic displays used in Experiment 1. The top and bottom displays contain the same binocular disparity, and differ only in the contrast within their left halves. However, their depth percepts are very different: In the top display, the contrast border is seen as bounding a transparent layer on the left, and its near depth propagates to the entire transparent surface on the left. In the bottom display, the near depth propagates to the entire transparent surface on the right.

one—because the perception of separate layers in depth, and the assignment of surface attributes to these layers, is more robust in such displays (e.g. Singh & Anderson, 2002). However, unlike our previous study, where we employed displays with a center-surround configuration, here the contrast discontinuity runs vertically down the midline, dividing the display into two lateral halves. This configuration is designed to ensure that the two textured regions have a symmetric status, insofar as configural factors are concerned. Since our goal is to study the role of photometric constraints in perceptual transparency, it is important to minimize biases from other sources. With a center-surround configuration, for instance, there is a potentially confounding figural bias, namely, the small enclosed central disk is more likely to be perceived as "figure," rather than "ground," which may interact in ill-understood ways with photometric constraints for transparency.

Consider the two stereoscopic displays in Fig. 3. They are identical as far as figural relations are concerned and, moreover, they contain precisely the same disparity information: First, the texture elements have been given far disparity relative to the rectangular frame, so that a frontoparallel textured plane appears to be viewed through a large rectangular window. Second, the locus of contrast discontinuity between the two halves has been given near disparity relative to the texture elements, which places this border in a depth plane between the underlying textured surface and the rectangular window. The resulting depth stratification is shown schematically in Fig. 4.

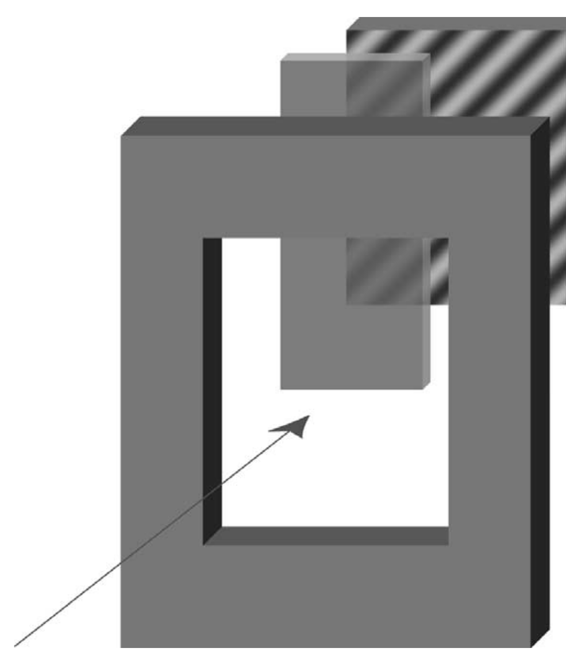

Fig. 4. Schematic depiction of the depth stratification of surfaces seen in the stereoscopic displays in Fig. 3.

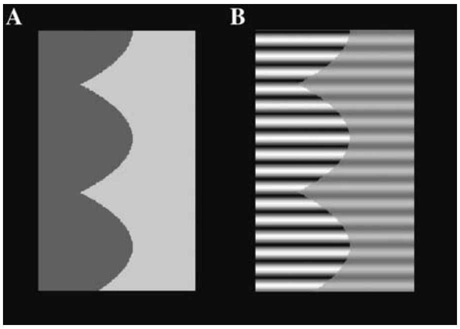

Fig. 5. Transparency as an instance of border ownership, or figure-ground assignment. (A) The luminance discontinuity is perceived as bounding a mid-gray surface of the left, with the light-gray surface perceived as extending behind it. (B) The contrast discontinuity is perceived as bounding a transparent layer on the right, with the striped background surface extending behind it.

It is important to note that although the disparity information places the vertical contrast discontinuity closer in depth relative to the textured background, it does not dictate which *side* of this boundary, if any, will be perceived as containing a surface at that depth. In the top panel of Fig. 3, the contrast discontinuity is perceived as being "owned" by a transparent surface that overlies the left half of the display—with its near depth propagating across this entire transparent surface. In the bottom panel, with precisely the same disparities and modifying only the contrast of the left half, the near depth of the contrast discontinuity now spreads to the entire right half of the display—or, more precisely, to the transparent surface perceived to overlie the right half of the display. This example demonstrates that the perceptual construction of transparency is naturally considered as an instance of figure-ground assignment (or border ownership): Whereas in a standard figure-ground display, a luminance discontinuity is perceived as being owned by one side or another (Fig. 5A), here a contrast discontinuity may be owned by a transparent surface on one side or the other (see Fig. 5B). Moreover, the perceptual difference between the top and bottom panels in Fig. 3 demonstrates that the percept of transparency can be altered by modifying only the relative contrast magnitudes—while preserving all configural and stereoscopic information. This fact provides us with an appropriate experimental handle on the investigation of photometric constraints in perceptual transparency.

3.1. Plotting the prediction space

A natural way to conceive of the problem at hand is in terms of the following question: If the luminance range (i.e., amplitude, $L_{\max} - L_{\min}$) of the textured region in the left half of the display is gradually increased (see Fig. 6) symmetrically about its mean, at what point along this spectrum does a switch in perceptual transparency occur? In other words, at what point does the percept switch from "left side is transparent" to "left side is seen in plain view?" Referring to the right side in this display as the *reference* side (which is fixed), and the left side the *test* side, this question can then be posed for different values of mean luminance on the test side. The loci at which a switch in the percept of transparency occurs—from the test side appearing transparent to appearing as seen in plain view—will naturally inform us about the photometric constraints being used to initiate percepts of transparency.

In slightly more general form, the question is simply: Given a reference side with fixed mean luminance and luminance range, what combinations of mean luminance and luminance range within the test side will generate a percept of transparency therein? Since the textures in the two halves of each display are generated using the same algorithm—differing only in mean luminance and luminance range—each texture patch can be quantified uniquely by its mean (L) and its range (R). The reference side has a fixed mean, L_0, and luminance range, R_0, and is thus represented by a fixed point in (L, R)-space (see Fig. 7). Relative to this reference point, we can now ask: What subset of the (L, R)-space can the test side reside in, and appear transparent?

Two very different prediction subsets follow from Metelli's magnitude constraint **C**2 and the Michelson-contrast constraint **C**3. The horizontal line passing through (L_0, R_0) in Fig. 7A corresponds to the locus of points for which the test side has the same luminance range as the reference side (i.e., $R = R_0$). Thus, according to Metelli's magnitude constraint, the shaded region below this line is the

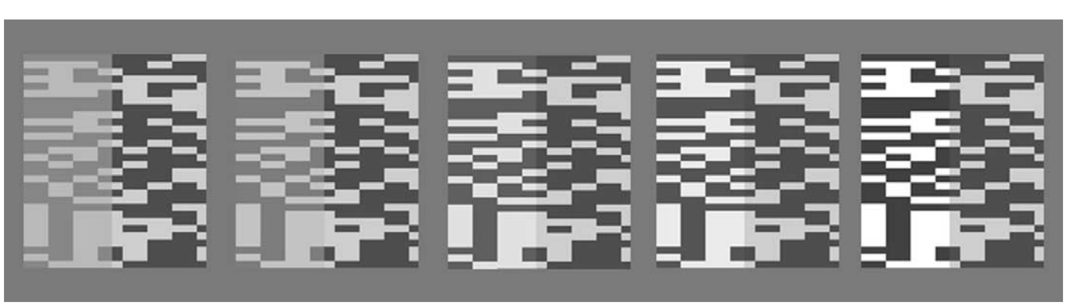

Fig. 6. A sequence of displays in which the right half is fixed, and the luminance range within the left half is gradually increased (symmetrically about its mean). The question addressed in Experiment 1 is: At what point along such a sequence does a perceptual switch occur from the left side being perceived as transparent to it being seen in plain view?

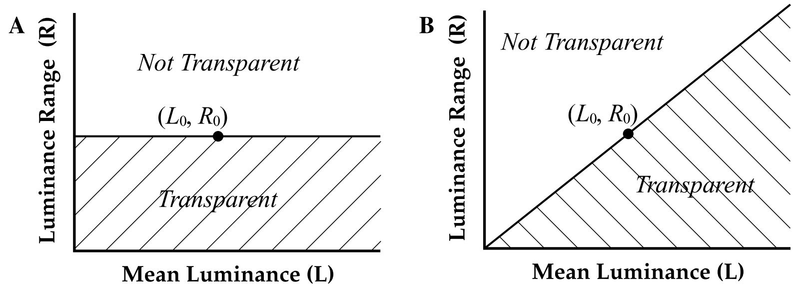

Fig. 7. Plotting the predictions of constraints **C2** and **C3** in initiating percepts of transparency. The reference side has fixed mean luminance L_0 and fixed luminance range R_0. The test side can take any value in (L, R)-space. (A) Shows the subset of the (L, R)-space that is predicted to generate a percept of transparency on the test side, according to Metelli's magnitude constraint **C2**. (B) Shows the corresponding transparent subset predicted by the Michelson-contrast constraint **C3**. Any other measure of contrast will similarly partition the (L, R) space into two.

set of (L, R) values for which the test side should appear transparent. On the other hand, the oblique line passing through (L_0, R_0) and the origin in Fig. 7B is the locus of points for which the test side has the same Michelson contrast as the reference side (since this line is defined by $\frac{R}{L} = \frac{R_0}{L_0}$). Thus, the Michelson-contrast constraint predicts that the shaded region below this oblique line is the set of (L, R) values for which the test side should appear transparent.[5]

Our experiments test the perceptual validity of these predictions. When the test and reference have the same mean luminance ($L = L_0$), note that the predicted transition point between "transparent" and "not transparent" is the same under both hypotheses. Thus, the two hypotheses can be experimentally distinguished only when the mean luminance of the test side differs sufficiently from that of the reference. Moreover, the case of equal mean luminances is also less interesting given that the contrast discontinuity simply disappears when the two sides have roughly the same contrast; whereas this is not the case when the mean luminances are unequal (e.g., see the sequence in Fig. 6). In Experiment 1, we measure the perceptual transition between "transparent" and "not transparent" on the test side, for multiple values of test mean luminance—both darker and lighter than the reference. In Experiment 2, we obtain observers' contrast matches on the same textures, in order to compare the transparency results against perceived contrast.

3.2. Experiment 1

To estimate the locus of switch in the percept of transparency, we adopt the strategy of approaching it from both directions. One of the two halves of the display—the reference—is fixed (with mean luminance L_0, and luminance range R_0). For different values of mean luminance L on the other—test—side, observers adjust its luminance range R to perform two different tasks. In one task, they indicate the highest luminance range for which the test side appears transparent. In the other task, they indicate the lowest luminance range for which the test side appears to be seen in plain view (i.e., with no overlying transparency). The mean of the settings across these two tasks provides an estimate of the locus of switch in the percept of transparency.

3.2.1. Methods

Observers. Three experienced observers, with normal or corrected-to-normal visual acuity, participated in the experiment. Two of the observers were naïve to the purposes of the experiment; the third was author MS (observer O2).

Stimuli and apparatus. Stimuli were stereoscopic displays, each consisting of a 4.31 deg $\times$ 5.38 deg rectangular frame containing one of three texture patterns. The texture elements were given a far disparity of 16.5 min of arc relative to the rectangular frame, and half-occlusions were introduced at its lateral edges ("da Vinci stereopsis," Nakayama & Shimojo, 1990). Texture elements adjacent to the left edge of the frame were present in the right eye's image only, and texture elements adjacent to the right edge were present in the left eye's image only. This resulted in the textured region being perceived as a frontoparallel surface, lying in a depth plane behind that of the frame, and seen through it.

The textured region in each eye's image was divided by a vertical mid-line into two lateral halves that could be assigned different values of luminance range ($L_{max} - L_{min}$) and mean luminance. This vertical contrast discontinuity was given a near disparity of 9.9 min of arc relative to the texture elements. Depending on the contrast relationships across the two sides, the contrast discontinuity was perceived as belonging to a transparent surface on either the left or the right side—floating at the depth of the contrast discontinuity, in front of the textured background (see the schematic depiction in Fig. 4).

The reference half of the display had a fixed mean luminance of 39.6 cd/m^2, and a fixed luminance range of 45.3 cd/m^2 (Michelson contrast $= 0.57$). The reference was

[5] It should be noted that although we highlight two specific predictions here (based on constraints C2 and C3), the framework is entirely general: Any other conceivable measure of contrast will similarly divide the (L, R)-space into two sub-regions, and thus predict its own locus of transition between transparency and non-transparency.

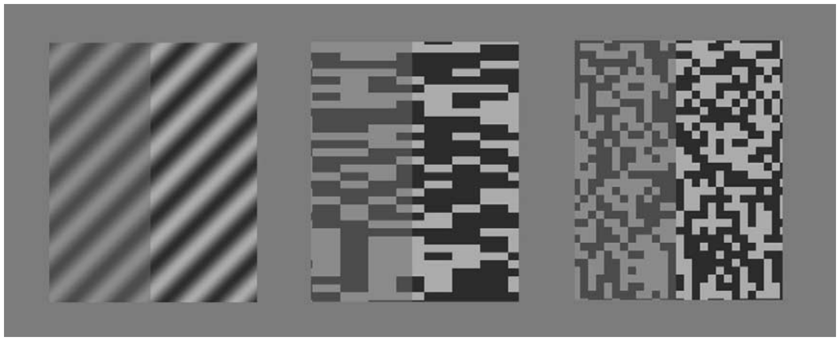

Fig. 8. Transparency displays showing the three texture patterns used in Experiment 1: sinusoidal gratings and two versions of random-dot textures.

equally likely to appear on the left or the right side of the display. From trial to trial, the mean luminance of the other, test, half of the display was clamped at one of 7 possible values: 24.8, 29.7, 34.7, 39.6, 44.6, 49.5, and 54.5 cd/m^2 (three darker than the reference, and three lighter). The luminance range within the test side was to be adjusted by the observers. On any given trial, the assignment of test and reference was immediately apparent to the observers, because only the luminance range within the test side was under their control.

Three different texture patterns were used: sinusoidal gratings and two versions of random-dot textures. The sinusoidal gratings had a period of 0.69° of visual angle, and were oriented at 45° (left panel of Fig. 8). The first version of the random-dot texture had rectangular elements, elongated in the horizontal direction (29.7 min × 9.9 min; middle panel of Fig. 8). The presence of the contrast discontinuity generated polarity-preserving X-junctions in each eye's image. The second version of random-dot texture had square dots, with sides measuring 9.9 min. The presence of the contrast discontinuity introduced no X-junctions in the monocular images for this texture (since the dot size equaled the disparity given to the contrast discontinuity; see right panel of Fig. 8).

The stimuli were presented on a linearized Radius PressView 17SR monitor. The monitor was calibrated so that screen luminance values (ranging from 0.58 to 90.14 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 99 cm.

Procedure. Each observer participated in 6 experimental sessions—one for each combination of the 2 experimental tasks and 3 textured backgrounds. On each trial, the mean luminance of the test side was randomly set to one of 7 predetermined values. Observers adjusted the luminance range within the test side (which varied symmetrically about the preset mean value) in order to perform one of two tasks. In one—the "highest-transparent"—task, the luminance range within the test was initially set to a low value (randomly drawn from the lower-most quarter of the scale), and observers set the *highest* luminance range for which the test side was perceived to contain an overlying transparent surface. In the other—"lowest not-transparent"—task, the luminance range within the test side was initially set to a high value (randomly drawn from the upper-most quarter of the scale), and observers set the *lowest* luminance range for which the test side was perceived to be in plain view, i.e., with no intervening transparent layer. Within each session, observers performed 35 experimental adjustments (5 adjustments for each of the 7 values of mean luminance), preceded by 7 practice adjustments (one for each value of mean luminance).

3.2.2. Results

The data for the three observers are plotted in Fig. 9A. No systematic differences were obtained across the three textured backgrounds; the data are thus shown averaged across the three textures. Each data point in Fig. 9A thus corresponds to the mean of 15 settings by an observer. The lower data curve in each plot represents an observer's settings in the "highest-transparent" task—i.e., the highest luminance range for which the observer perceived the test side to be transparent. The upper curve represents an observer's settings in the "lowest-not-transparent" task—i.e., the lowest luminance range for which the observer perceived the test side to *not* be transparent (i.e., seen in plain view). The region in between thus corresponds to a transition zone, where neither side has a clear preferential status in terms of a transparency percept. An estimate of the locus of switch in the percept of transparency may thus be obtained by taking the mean of the settings in the two tasks. These estimated transition loci for the 3 observers are displayed in the graphs in Fig. 9B.

Two characteristics of these data are salient: First, the data exhibit a systematic failure of Metelli's magnitude constraint in predicting a percept of transparency. As shown in Fig. 7A, the magnitude constraint predicts that the locus of transition between transparency and non-transparency should be flat—exhibiting no dependence on mean luminance. This prediction is clearly not borne out; rather observers' settings exhibit a systematic increase with the mean luminance on the test side. The deviations from the predictions of the magnitude constraint are particularly large in the left half of each plot (i.e., all points to the left of the reference point (L_0, R_0)).

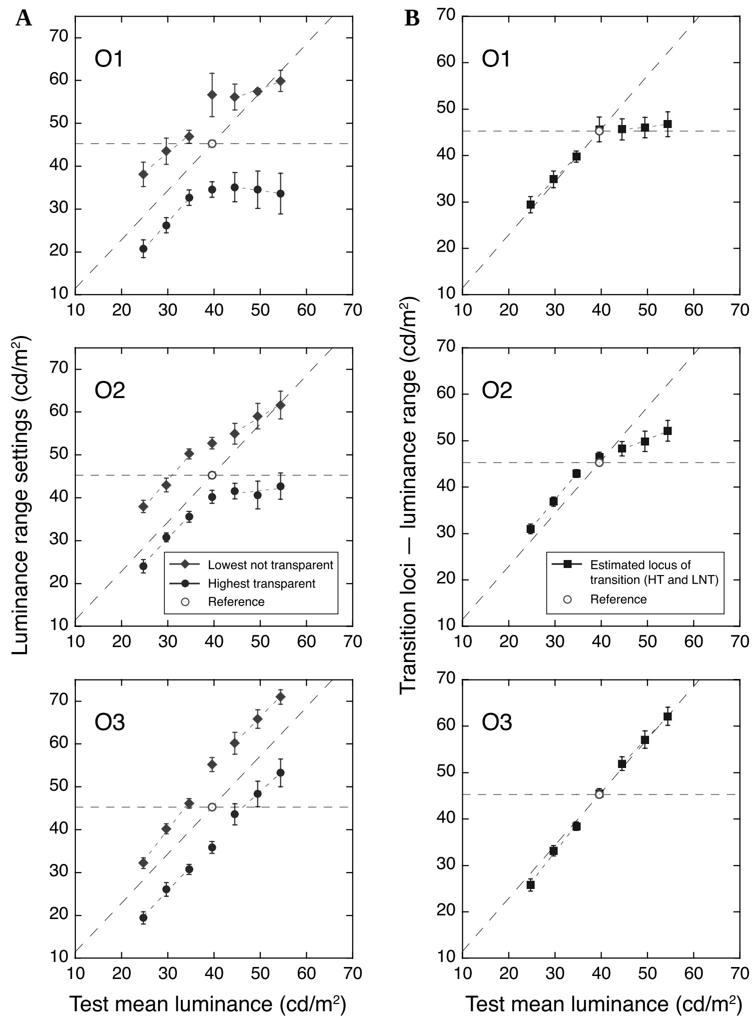

Fig. 9. Results of Experiment 1: (A) Observer settings of luminance range for the *highest-transparent* and the *lowest-not-transparent* tasks. The error bars depict 95% confidence intervals. (B) Estimated transition loci between transparency and non-transparency based on the average of the settings on the two tasks. Observers' settings consistently follow the Michelson-contrast constraint for darkening transparency (data points to the left of reference), but are more variable for lightening transparency (data points to the right of reference).

This half corresponds to the case of 'darkening' transparent layers—when the test side has a lower mean luminance than the reference. In this regime, the test side consistently fails to be perceived as transparent when its Michelson contrast exceeds that of the reference side—despite the fact that its luminance ranges is still lower than that of the reference (thus satisfying Metelli's magnitude constraint).

Second, a marked asymmetry is evident between tests that are darker than the reference (all points to the left of (L_0, R_0)) and tests that are lighter than the reference (all points to the right of (L_0, R_0)). When the test side is darker than the reference, the settings of the three observers agree closely with one another, and they closely follow the Michelson-contrast prediction (recall Fig. 7B). On the other hand, when the test side is lighter than the reference, there is greater variability across observers: Whereas observer O3's settings continue to follow the Michelson-contrast prediction, the settings of observers O1 and O2 are strongly biased toward the horizontal transition locus predicted by Metelli's magnitude constraint. (These biases are robust, and were obtained in each individual experimental session for observers O1 and O2.)

In addition to the greater variability *across* observers, there is also greater variability *within* each observer's settings in the regime where the test region is lighter than the reference ('lightening transparency'). For each observer, and for each of the two experimental tasks, setting variance was collapsed over the 3 test luminances that are darker than the reference ("darker variance," s_{d}^2), and similarly over the 3 test luminances that are lighter than the reference ("lighter variance," s_{l}^2). In 5 of the 6 cases (3 observers × 2 experimental tasks), setting variability was higher in the lighter-transparency condition than in the darker-transparency condition; in the sixth case it was almost identical. To test whether these differences were statistically reliable, we performed a test of the alternative hypothesis that $\sigma_{\mathrm{l}}^2 > \sigma_{\mathrm{d}}^2$, against the null that $\sigma_{\mathrm{l}}^2 = \sigma_{\mathrm{d}}^2$, using the F statistic:

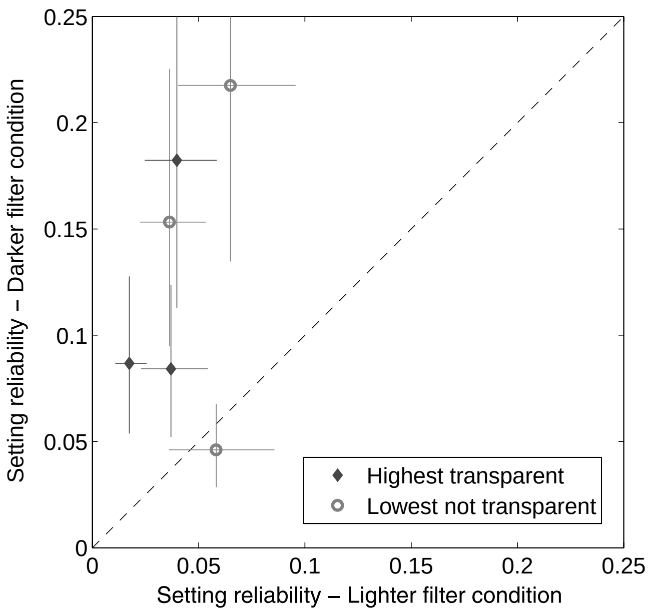

Fig. 10. Reliability (1/variance) of observer settings plotted in the 'darkening' transparency cases versus the 'lightening' transparency cases. Most points lie above the identity function, thereby indicating that setting reliability is greater when the test region is darker than the reference. Error bars denote 95% confidence intervals.

$$F(42,42) = \frac{s_l^2}{s_d^2}. \qquad (5)$$

The tests revealed that the increase in setting variance in the lighter-transparency condition was statistically significant (at the .05 level) for all 5 cases where $s_l^2 > s_d^2$. (As expected, the difference was not statistically significant for the sixth case.) Fig. 10 depicts this comparison in terms of setting reliability (=1/variance)—a commonly used measure of the effectiveness of a visual cue in determining a perceptual estimate (e.g. Backus & Banks, 1999; Ernst & Bülthoff, 2004; Maloney, 2002). (See Fulvio, Singh, & Maloney, 2005, for a similar analysis in investigating achromatic and chromatic contributions to perceptual transparency.) Observers' settings are consistently more reliable in the darkening-transparency condition than in the lightening-transparency condition: most points in Fig. 10 lie above the oblique line depicting the identity function.

3.2.3. Discussion

Experiment 1 reveals two basic results. First, the perceptual construction of transparency is not predicted by Metelli's magnitude constraint. Rather, the perceptual transition between transparency and non-transparency is better approximated by the Michelson-contrast constraint. Second, there is a marked asymmetry between transparent layers that darken versus lighten the underlying surface. In particular, settings in the two tasks (highest-transparent and lowest-not-transparent) were both more consistent across observers, as well as more reliable within individual observers, in the 'darkening-transparency' case. Thus the perceptual transition between transparency and non-transparency is sharply defined for 'darkening' transparency, but rather imprecise and ill-defined for 'lightening' transparency. This asymmetry in precision is consistent with observers' reports that there was a wide range of settings in the lighter-filter conditions for which the percept of transparency (i.e., whether the left or right side was transparent) was relatively ambiguous.

Even though Michelson contrast does a considerably better job of explaining observers' settings than Metelli's magnitude constraint, it does not entirely capture the transition between transparency and non-transparency—at least for cases involving lightening transparency. Whereas it perfectly predicts the transition loci for all observers in the regime of darkening transparency, in the regime of lightening transparency, two of the three observers exhibited large deviations from the Michelson-contrast constraint. Specifically, their settings were systematically biased toward the luminance-range prediction—the 'horizontal' prediction that follows from Metelli's magnitude constraint; recall Fig. 7A.

A natural question is whether these deviations from Michelson contrast can be attributed to the way in which observers scale perceived contrast in the textures used. For the more complex textures used in Robilotto et al.'s (2002) study, for instance, neither Michelson contrast nor any other standard contrast measure could capture observers' perceived contrast; however their apparent contrast nevertheless predicted the perceived transmittance of transparent filters. Similarly, in the current context, it may be the case that perceived-contrast constraint **C4** perfectly captures observers' percept of transparency, but that perceived contrast is simply not determined by Michelson contrast. Experiment 2 investigates this possibility by obtaining contrast matches on the textures used in Experiment 1.

3.3. Experiment 2

Experiment 2 measures how observers scale perceived contrast in the textures used in Experiment 1. We use display configurations similar to Experiment 1, but with the test and reference halves separated laterally, so the displays no longer generate the percept of transparency (see Fig. 11). Observers adjust the luminance range within the test texture patch in order to match the apparent contrast of the reference texture patch.

3.3.1. Methods

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

Stimuli. The stimuli differed from those used Experiment 1 in only two ways. First, the two (test and reference) halves of each transparency display were laterally separated (see Fig. 11). Second, the displays contained no disparity information. Both manipulations were designed to suppress the percept of transparency. The same set of photometric values

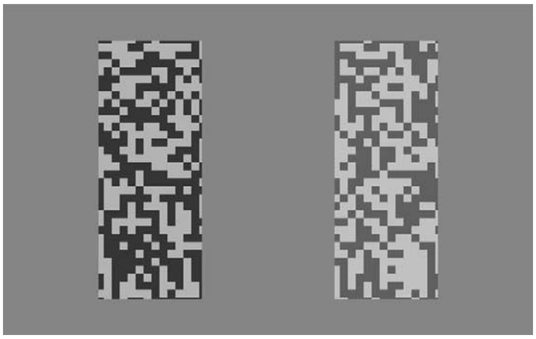

Fig. 11. Stimulus configuration used for contrast matching in Experiment 2. The mean luminance of the test patch (the one on the right in this display) was clamped to different values, and observers adjusted the luminance range within it in order to match the apparent contrast of the reference patch. This experiment used the same three texture patterns as in Experiment 1 (see Fig. 8).

were used, within the same set of textures, for the test and reference patches.

Procedure. Each observer participated in 3 sessions—one for each of the 3 textures. On each trial, the test patch was randomly set to one of the 7 values of mean luminance. Observers adjusted the luminance range within the test patch in order to match the perceived contrast of the reference. Within each session, observers performed 35 experimental adjustments (5 adjustments for each of the 7 values of mean luminance), preceded by 7 practice adjustments.

3.3.2. Results

The data for the three observers are plotted in Fig. 12. As in Experiment 1, no systematic differences were obtained across the three textures, and the data are shown averaged over them. Each data point thus corresponds to the mean of 15 adjustments for an observer.

Unlike the transparency settings for the 'lightening' filters is Experiment 1 (recall Fig. 9), the contrast data exhibit a high degree of consistency across observers. In particular, the data curves consistently follow the oblique prediction based on Michelson contrast—thereby indicating that for the texture patterns used, observers' matches of perceived contrast are well-captured by Michelson contrast. Consequently, the deviations obtained in Experiment 1 from the Michelson-contrast prediction—in two of the three observers, in the case of 'lightening' filters—cannot be attributed to how the observers scale perceived contrast in these displays. Indeed, these results show that the transparency setting deviate systematically from perceived contrast. Thus the deviations obtained in Experiment 1 appear to be specific to the context of perceptual transparency.[6]

[6] Observer O2's contrast matches exhibit a small deviation from the Michelson-contrast prediction; however this deviation is in the opposite direction, relative to the large deviation obtained in the transparency settings in Experiment 1; see Fig. 9B.

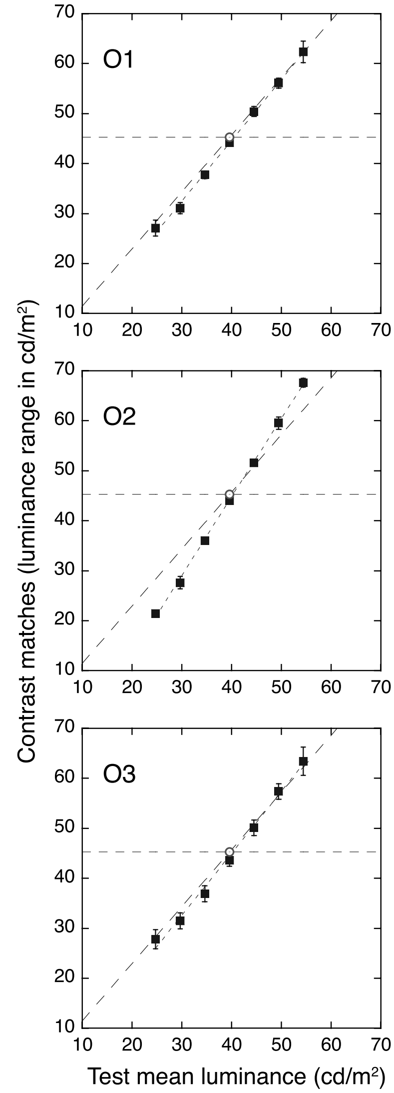

Fig. 12. Results of Experiment 2. Observers' contrast matches on the textures used in Experiment 1 follow the Michelson-contrast prediction, given by the oblique dashed line. (The error bars denote 95% confidence intervals.) Hence the deviations from the Michelson-contrast constraint observed in Experiment 1 (Fig. 9) cannot be attributed to the way in which the observers scale apparent contrast in the textures used.

4. General discussion

4.1. Patten of bias

The results of Experiment 1 demonstrate that Metelli's magnitude constraint fails to predict the perception of transparency. Specifically, the locus of transition between transparency and non-transparency (measured in terms of luminance range, $L_{max} - L_{min}$) increases systematically with the mean luminance on the test side—whereas the magnitude constraint predicts that it should be constant. This transition locus is better captured by the Michelson contrast on the test side, relative to the reference. These results thus roughly parallel previous results on matching the material properties of transparent surfaces, where the perceived transmittance of a transparent surface was found to deviate systematically from Metelli's formula for α, and given instead by the Michelson contrast (Singh & Anderson, 2002) or perceived contrast (Robilotto et al., 2002) in the region of transparency.

However, the results of Experiment 1 also exhibited deviations from the Michelson-contrast prediction. Specifically, in the regime of lightening transparency (i.e., when the test has higher mean luminance than the reference), two of the three observers exhibited a shift toward the luminance-range prediction (Fig. 9). These deviations are surprising, given that they constitute deviations from perceived contrast as well. (As the results of Experiment 2 showed, perceived contrast in the textures used here coincides with Michelson contrast.)

What might be the source of these deviations, in the regime of lightening transparency? One possibility is suggested by considering the derivation of the two constraints **C**2 and **C**3. Recall that the luminance-range constraint **C**2: $p - q \leqslant a - b$ was derived from the restriction $\alpha \leqslant 1$ on the solution for α (Gerbino et al., 1990; Metelli, 1974), whereas the Michelson-contrast constraint **C3** : $\frac{p-q}{p+q} \leqslant \frac{a-b}{a+b}$ was derived from the additional restriction $t \geqslant 0$ on the solution for t (Singh & Anderson, 2002). Because these two restrictions ($\alpha \leqslant 1$ and $t \geqslant 0$) apply simultaneously, however, it follows that both constraints **C**2 and **C**3 must be satisfied together.

The two plots on the left of Fig. 13 display the respective subsets of the (L, R)-space where constraints **C**2 and **C**3 are satisfied individually (recall Fig. 7). Thus the region of (L, R)-space where *C***2** and **C**3 are satisfied simultaneously is given by the intersection of these two prediction subsets. This intersection-of-constraints prediction is plotted on the right-hand side of Fig. 13. As is evident from this graph, the intersection hypothesis does indeed predict an asymmetry between darkening and lightening transparency. For darkening transparency, the locus of transition from transparency to non-transparency is predicted to be given by Michelson contrast; whereas for lightening transparent layers, the transition locus is predicted to be given by luminance range. Thus the combined operation of the two constraints **C**2 and **C**3 predicts the "flattening" of the data curves observed in 2 of 3 three observers.

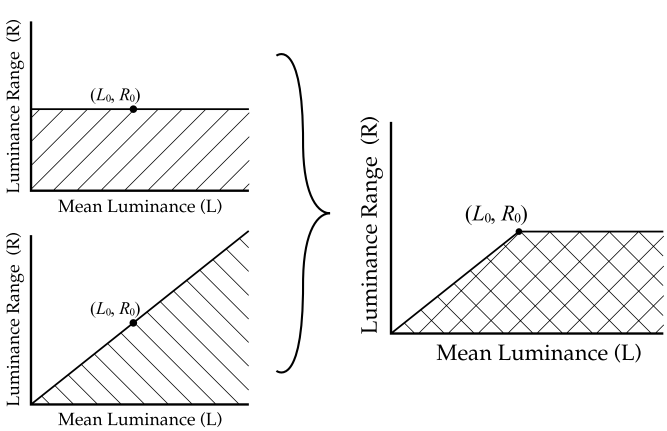

Fig. 13. Predictions of transparency based on the intersection of the two constraints shown in Fig. 7. The left side depicts the respective subsets of (L, R)-space that are predicted to generate a percept of transparency according to Metelli's magnitude constraint (top) and the Michelson-contrast constraint (bottom). However, because these two constraints must be satisfied simultaneously, the percept of transparency is predicted by the intersection of these two subsets, as shown on the right.

4.2. Asymmetry in setting variability

In addition to the asymmetry in bias, a large asymmetry between darkening and lightening filters was also observed with respect to setting variability. For darkening filters, settings were highly consistent across observers, as well as highly reliable within individual observers. For lightening filters, on the other hand, there was considerably greater variability both across observers, as well as within individual observers' settings. Thus, whereas the perceptual transition between transparency and non-transparency is sharply defined for darkening transparency, it is substantially more imprecise and ill-defined for lightening transparency. Although the intersection-of-constraints hypothesis considered above explains the pattern of bias seen in observers' settings, it does not make any specific prediction concerning setting reliability.

A natural hypothesis to consider is whether the light/dark asymmetry in variability is simply the result of Weber's law. After all, the test mean luminance is, by definition, higher for lightening transparency than darkening transparency. In order to test this hypothesis, we compared the pattern of increase in variance across the transparency experiment (Experiment 1) and the contrast experiment (Experiment 2). If the light/dark asymmetry observed in the transparency experiment were simply the result of Weber's law, one would predict a similar light/dark asymmetry in the contrast experiment as well. When F tests comparing the variances in the lighter versus darker test patches were performed for the contrast settings in Experiment 2, none of the observers exhibited a significantly higher variance for the lighter reference patches (although there was a trend in that direction). This differs markedly from the results of Experiment 1, where 5 of the 6 tests performed revealed a statistically reliable increase in variance in the lighter-transparency condition (see Fig. 11). Because the set of values of mean luminance and contrast used in the two experiments were identical, this difference indicates that the large light/dark asymmetry observed in Experiment 1 is specific to the context of transparency.

Visually comparing the plots of standard deviations against mean luminance in the contrast settings versus transparency settings also suggests the same conclusion (see Fig. 14). In the contrast settings, there was little evidence for a systematic increase in standard deviation, except at the very highest value of mean luminance. In the "highest-transparent" task of Experiment 1, by contrast, standard-deviations remained essentially constant within the darker-transarency regime; then they increased suddenly as the mean-luminance values enter the lightening-transparency regime. This differential pattern of increase again indi-

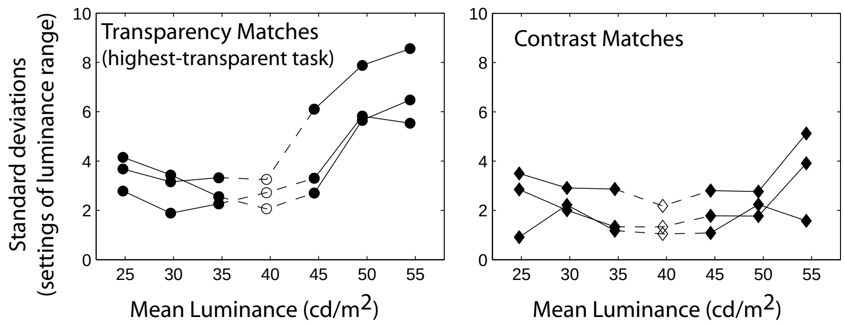

Fig. 14. A comparison of the increase in standard deviations of observers' settings in the *highest-transparent* task in Experiment 1 (left), and the contrast-matching task in Experiment 2 (right). This comparison suggests that the light/dark asymmetry obtained in Experiment 1 is not simply the result of a Weber-like behavior.

cates that the light/dark asymmetry in the transparency matches cannot be reduced to Weber's law.

4.3. Darkening as a separate cue

One way to understand these dark/light asymmetries is in terms of the statistics of natural images (see, e.g., Brunswik & Kamiya, 1953; Elder & Goldberg, 2002; Geisler, Perry, Super, & Gallogly, 2001, for the role of statistics of natural images in understanding visual processing). A ubiquitous feature of the natural environment is the presence of shadows. Although shadows are not *material* layers, they share a number of photometric properties with transparency. From a generative point of view, their presence preserves contrast polarity along underlying contrast borders, and their photometric influence is in fact consistent with that of a neutral-density filter, or a veil of zero reflectance.[7] From a perceptual point of view, shadows also involve computing multiple photometric causes by decomposing image luminance. The sheer frequency and ubiquity of shadows in the natural environment, and the critical need for visual organisms to interpret them correctly, has led many researchers to posit that general mechanisms of luminance decomposition (such as those responsible for the perception of transparency) evolved from visual mechanisms for computing shadows (Noest & van den Berg, 1993; Stoner & Albright, 1996; see also von Helmholtz, 1860/1924).

If visual mechanisms of luminance decomposition are indeed tuned to shadows, this would suggest the existence of asymmetries in how the visual system interprets shadows versus spotlights. Such asymmetries have indeed been demonstrated. The bottom display in Fig. 15 can be generated by two very different physical setups: a light-colored surface with a shadow covering its left half, or a dark-colored surface with a spotlight shining on its right half. Using backgrounds such as these, Gilchrist and Zdravkovic (2002) found that when observers are shown a dark-colored surface, half of which is illuminated with a spotlight, they perceive it instead as a light-colored surface, half covered by a shadow. Based on these results, Gilchrist and Zdravkovic (2002) posited that, in addition to anchoring the highest luminance to white, the visual system also uses the highest luminance in the image to define the default illumination level of the scene (see also Kozaki, 1973). Thus the brighter half of the display is perceived as a light-colored surface under default illumination, rather than a dark surface under spotlight.

From the point of view of the current article, Gilchrist and Zdravkovic's result also speaks to the light/dark asymmetry observed in our own data. Indeed, there are two ways in which the light/dark asymmetry observed in lightness perception may bear on the current results. Although the two are related, and each is capable of providing insight into the light/dark asymmetry observed in our transparency data, they embody different assumptions concerning the nature of the underlying representations in lightness perception and their relationship to transparency. We describe each in turn.

One interpretation of the light/dark asymmetry is that the visual system treats darkening as a cue to luminance decomposition, or scission. In this analysis, the concept of "default illumination" in lightness perception is identified with perceiving a surface "in plain view" in transparency. From this perspective, when viewing the bottom display in Fig. 15, the darker half is represented as *light-colored surface + shadow*, whereas the lighter patches on the right are perceived in plain view and as continuing "underneath" the shadow. Thus, unlike traditional intrinsic-image models which assert that all image regions are decomposed into illumination and reflectance maps (e.g. Barrow & Tenenbaum, 1978), this view asserts that only *shadowed* (or lower-illumination) regions are explicitly decomposed into these two sources. The regions in the default (highest) illumination are perceived in "plain view," in the sense that the lightness of surfaces under the default illumination are determined via some kind of anchoring and normalization principles that interpret the highest luminance as white (e.g. Gilchrist et al., 1999). Note that this view embodies a *hybrid* model of light-

[7] To a first-order approximation, the photometric transformation introduced by a shadow is given by $P = \alpha A$, which is seen to be a special instance of Metelli's equation with $T = 0$. Thus shadows can only darken underlying surfaces, whereas the presence of a potentially non-zero term T in transparency means that a transparent surface can either darken, lighten, or preserve the underlying luminance.

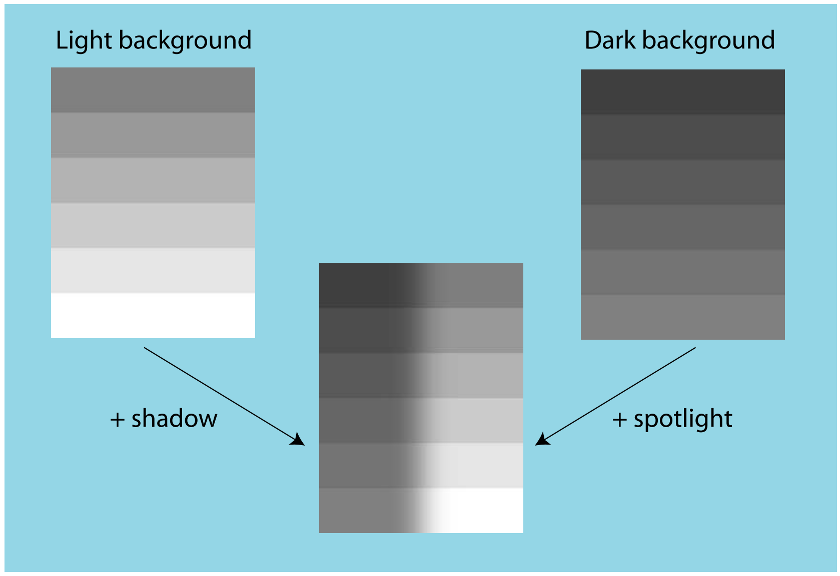

Fig. 15. Demonstrating the perceptual asymmetry in the interpretation of shadows versus spotlights. When the display at the bottom is generated by shining a spotlight on the right half of a dark-colored background, observers perceive it instead as a light-colored background, whose left half is covered by a shadow (after Gilchrist and Zdravkovic, 2002).

ness perception—asserting the role of scission in the interpretation of shadows, but relying on anchoring and normalization principles to determine the perceived lightness of surfaces in the default (or highest) illumination.

An alternative, but related, interpretation of the light/dark asymmetry focuses on the role of *border ownership* of illumination boundaries. The fact that observers treat the illumination boundary as arising from a shadow rather than a spotlight can be re-expressed as a statement about the interpreted ''ownership'' of illumination boundaries (such as the central graded contrast border in the bottom display in Fig. 15). Indeed, another way of stating Gilchrist and Zdravkovic (2002) result is that observers treat the illumination boundary as *arising* from a cast shadow, not a spotlight—i.e., from the presence of an object that has interrupted the default illumination. From this perspective, polarity-preserving darkening of surfaces and contours provides evidence for the presence of a cast shadow, and hence, serves as a cue to the border ownership of the contrast border. Note that in our experiments, the central dividing border was presented in stereoscopic depth, and observers made settings in the highest-transparent and lowest-not-transparent tasks. It is possible that observers based their settings simply on which side appeared to ''own'' the contour floating in depth.

Thus there are two interpretations of Gilchrist and Zdravkovic's result that may provide some understanding of the light/dark asymmetry observed in our experiments. According to the luminance-decomposition version of the hypothesis, only one side of the contour (the one treated as a perturbation from the default atmosphere) is explicitly given a layered representation; the other side (containing the default atmosphere) is represented as a single layer. Although this is a reasonable and efficient strategy for biological vision, it deviates from traditional models of intrinsic-image analysis (e.g. Barrow & Tenenbaum, 1978).[8] The border-ownership version of the hypothesis, on the other hand, is neutral with respect to whether or not a layered representation is generated at every point in the image, or only for perturbations from the default atmosphere. It asserts simply that the contrast border is more likely to be owned by the darker side—making no claims concerning the side that does not own the contrast border. Determining which of the two versions is ultimately more appropriate will require further empirical tests aimed specifically at distinguishing them.

From the perspective of the current paper, however, the important point is that the dark/light asymmetries observed in the context of spotlights and shadows, and the statistics of natural images, both suggest that the visual system may use polarity-preserving darkening as a cue to transparency. Indeed, treating darkening as an additional cue to transparency perfectly accounts for the light/dark asymmetries seen in our data—both in setting bias and in setting reliability (Experiment 1). When the test mean luminance is lower than the reference, the darkening cue on the test side is always consistent with transparency, and the locus of transition between transparency and non-transpar-

[8] In the context of transparency, it is evident that representing a surface in plain view as two layers in depth—the surface plus an overlying transparent layer that happens to have full transmittance (and does not contribute to image luminance)—would constitute an uneconomical representation, and a waste of cortical resources. A more parsimonious approach is to have a single-layer representation be the default (i.e., corresponding to image regions interpreted as seen in plain view—which generally constitute a large proportion of the image), and generate a multilayered representation only in those regions that deviate from this default ''atmosphere.''

ency is therefore determined entirely by perceived-contrast (=Michelson contrast for the textures used). This transition locus is sharply defined (setting reliability is high) because as soon as the test contrast goes below the reference contrast, the contrast and luminance cues both become consistent with perceptual transparency—thereby combining to yield a powerful signal for transparency. On the other hand, when the test mean luminance is higher than the reference, the darkening cue becomes inconsistent with transparency. Thus, in order to overcome the inconsistent signal from the darkening cue, observers must increase the strength of the contrast cue by setting the test contrast substantially lower than the reference contrast. (This was the case for two of our three observers.) Moreover, in the lightening-transparency case, varying the contrast on the test side can only manipulate the cues from (a) neither cue being consistent with transparency (when the test contrast exceeds the reference contrast) to (b) a situation involving cue conflict (when the test contrast is lower than the reference). As a result, the locus of transition between transparency and non-transparency is more imprecise and ill-defined (lower setting reliability within observers, and lower consistency across observers).

In sum, our data is well explained by treating darkening as a separate cue to transparency that interacts with the contrast-reduction cue to transparency. This cue may work either by introducing a border-ownership bias, or by introducing an additional cue to luminance scission. Although further research is needed to fully resolve how this additional cue is instantiated, the important point is that treating darkening as a separate cue to transparency provides a coherent account of the pattern of results obtained in our experiment—both in terms of the pattern in bias and the asymmetry in setting variability.

5. Conclusions

In measuring the locus of perceptual transition between transparency and non-transparency, we found a failure Metelli's magnitude constraint. The locus of transition was approximated instead by a constraint based on Michelson contrast (=perceived contrast in the textures used). In addition, however, settings also revealed a pronounced asymmetry between darkening and lightening transparency, both in terms of bias and reliability. For darkening transparency, the estimated transition locus closely and consistently followed the prediction of perceived contrast, with high setting reliability. For lightening transparency, however, the settings were more variable across observers (with two of the observers exhibiting deviations from the prediction of perceived contrast), as well as less reliable within observers. These asymmetries in bias and reliability can both be understood in terms of an interaction between the lowering-contrast cue to transparency, and a proposed darkening cue to transparency. Independently measuring the strength of the darkening cue, and investigating its interactions with other cues to perceptual transparency, remains an important topic for future research.

Acknowledgments

We thank two anonymous reviewers for their comments on a previous version of the manuscript. The authors were supported by NSF grant BCS-0216944 to M.S., and NIH grant EY0D11487-01 and ARC grant DP0559984 to B.A.

References

Adelson, E. H. (1993). Perceptual organization and the judgment of brightness. *Science, 262*, 2042–2044.

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. In *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. (1999). Stereoscopic surface perception. *Neuron, 26*, 919–928.

Anderson, B. L. (2003). The role of occlusion in the perception of depth, lightness, and opacity. *Psychological Review, 110*, 762–784.

Anderson, B. L., Singh, M., & Meng, J. (2005). The perceived transmittance of inhomogeneous surfaces and media. Vision Research, in press.

Anderson, B. L., & Winawer, J. (2005). Image segmentation and lightness perception. *Nature, 434*, 79–83.

Backus, B. T., & Banks, M. S. (1999). Estimator reliability and distance scaling in stereoscopic slant perception. *Perception, 28*, 217–242.

Barrow, H. G., & Tenenbaum, J. (1978). Recovering intrinsic scene characteristics from images. In A. R. Hanson & E. M. Riseman (Eds.), *Computer vision systems* (pp. 3–26). Orlando: Academic Press.

Beck, J., & Ivry, R. (1988). On the role of figural organization in perceptual transparency. *Perception and Psychophysics, 44*, 585–594.

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

Brunswik, E., & Kamiya, J. (1953). Ecological cue-validity of proximity and other Gestalt factors. *American Journal of Psychology, 66*, 20–32.

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

Elder, J. H., & Goldberg, R. M. (2002). Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision, 2*(4), 324–353.

Ernst, M. O., & Bülthoff, H. H. (2004). Merging the senses into a robust percept. *Trends in Cognitive Sciences, 8*, 162–169.

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.

Fulvio, J. M., Singh, M., & Maloney, L. T. (2005). Combining achromatic and chromatic cues to transparency. *Journal of Vision, 5*(8), 566a.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research, 41*(6), 711–724.

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., Kossyfidis, C., Agostini, T., Li, X., Bonato, F., Cataliotti, R., et al. (1999). An anchoring theory of lightness perception. *Psychological Review, 106*, 795–834.

Gilchrist, A. L., & Zdravkovic, S. (2002). Highest luminance defines illumination level as well as lightness. *Journal of Vision, 2*(7), 553a.

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

Kanizsa, G. (1979). *Organization in vision: Essays on Gestalt perception*. New York: Präger.

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

Kasrai, R., & Kingdom, F. (2002). Achromatic transparency and the role of local contours. *Perception, 31*, 775–790.

Kozaki, A. (1973). Perception of lightness and brightness of achromatic surface color and impression of illumination. *Japanese Psychological Research, 15*, 194–203.

Maloney, L. T. (2002). Statistical decision theory and biological vision. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world* (pp. 145–189). New York: Wiley.

Masin, S. (1997). The luminance conditions of transparency. *Perception, 26*, 39–50.

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. (1974). 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.

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.

Noest, A. J., & van den Berg, A. V. (1993). The role of early mechanisms in motion transparency and coherence. *Spatial Vision, 7*, 125–147.

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

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

Robilotto, R., & Zaidi, Q. (2004). Perceived transparency of neutral density filters across dissimilar backgrounds. *Journal of Vision, 4*, 183–195.

Singh, M. (2004). Lightness constancy through transparency: Internal consistency in layered surface representations. *Vision Research, 44*, 1827–1842.

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

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, 2*, 11–18.

Stoner, G. R., & Albright, T. D. (1996). The interpretation of visual motion: Evidence for surface segmentation mechanisms. *Vision Research, 36*, 1291–1310.

von Helmholtz, H. (1860/1924). *Physiological optics*. English translation by Southall, J.P.C. for the Optical Society of America from the 3rd German edition of 1909, *Handbuch der Physiologischen Optik*. Voss, Hamburg.