Layered image representations and the computation of surface lightness

Barton L. Anderson

School of Psychology,
University of New South Wales and University of Sydney,
Sydney, Australia

Jonathan Winawer

Department of Psychology, Stanford University,
Stanford, CA, USA

A fundamental goal of research in the perception of surfaces is to understand the nature of the computations and representations underlying lightness perception. A significant challenge posed to the visual system is recovering surface lightness from the multiple physical causes that contribute to image luminance. One view asserts that the visual system decomposes the image into estimates of illumination, lightness, and transparency, generating layered image representations. More recent views have questioned the need to posit layered representations to explain lightness perception. Here, a number of demonstrations and experiments involving the perception of transparency are presented that reveal a critical role played by layered image representations in the computation of surface lightness. We provide new evidence demonstrating that the contrast relationships along contours can play a decisive role in determining whether images are decomposed into multiple layers, and that the constraints that regulate how this decomposition occurs can have a dramatic influence on perceived lightness.

Keywords: lightness/brightness perception, perceptual organization, depth

Introduction

The past few decades has witnessed significant growth in experimental work in lightness perception, but there is little theoretical consensus about the underlying computations, representations, or mechanisms that support this dimension of perceptual experience. One view asserts that the visual system engages some form of source decomposition, wherein the image is separated (or *scised*) into contributions of reflectance, illumination, 3D shape, and intervening transparent surfaces (or media). Theories of this kind can be traced to the writings of Helmholtz (1866) and found modern expression in Barrow and Tenenbaum's (1978) intrinsic image analysis. In such models, the various causes of luminance variations are represented as separate "layers" that capture the contribution of each cause to the luminance distributions in the eyes. The main challenges facing scission models is to specify the information used to decompose the image into its underlying sources; to articulate the nature of the computations used to extract the properties of the different layers; and to determine how this information is (or is not) represented.

Other theoretical approaches to lightness perception question whether layered image representations play any direct role in the perception of surface lightness. For example, framework models stress the importance of grouping and segmentation processes in lightness computations but do not explicitly decompose the image into a set of causal layers (Adelson, 1999; Gilchrist et al., 1999). In Gilchrist et al.'s (1999) model, the image is decomposed into distinct regions, which may be contiguous ("like a country is divided into provinces"), nested, or intersecting. Each of these regions is considered a framework in which the highest luminance serves as the standard (or "anchor") that is assumed to be white, and other lightness values are computed relative to this standard. If a target region is in multiple overlapping frameworks, its lightness value is anchored in each, and the multiple values are combined in a weighted average. The role of depth in framework models is to determine which regions belong to a common framework and is just one of the grouping principles used to determine how the image is segmented into separate groups to which anchoring principles are subsequently applied. One of the critical differences between framework models and layer models is that there is no representation of the illuminant or transparent surfaces in a framework model; rather, luminance anchoring within and between frameworks is the sole means by which lightness values are computed. Each point in the image can thus be represented by only a single lightness value.

A related scheme is embodied in Adelson's (1999) concept of "adaptive windows." Adelson suggested that the perceived lightness of a region is determined through statistical estimation by comparing its luminance to other luminance values within a "soft-edged" adaptive

window. A critical component of the proposal is that the windows can change size and shape to avoid mixing "atmospheres." Using a variant of Metelli's equations, Adelson defined an image region's atmosphere via an atmospheric transfer function (ATF). This function contains one additive and one multiplicative term to describe the cumulative effects that the illuminant and intervening transparent surfaces or media have on the reflectance of underlying surfaces. According to this approach, various perceptual grouping mechanisms influence perceived lightness by providing cues about the presence of atmospheric boundaries, i.e., regions where the additive and multiplicative terms in the ATF differ in an image, thus influencing the size and shape of the window within which luminance values are compared to infer lightness values. Thus, although Adelson's model utilizes similar image properties as scission models (such as contour junctions) to infer atmosphere boundaries, the image is not decomposed into layers. Layers therefore do not play a causal role in determining perceived lightness in Adelson's model, and hence the nature of the underlying representations appears similar to those embodied in Gilchrist's model.

A common theme that underlies layer and framework models is that they both attempt to derive an estimate of surface reflectance from images. It should be noted, however, that there are also models that attempt to account for a variety of lightness phenomena without deriving explicit estimates of surface reflectance. These include spatial filtering models of brightness (e.g., Blakeslee & McCourt, 2004; Dakin & Bex, 2003; Kingdom & Moulden, 1992) and models of brightness filling-in (Rudd & Zemach, 2005, 2007). As with framework models, in their current form, such models can only transform image luminance into a single output value (which can be considered either a brightness or lightness value, depending on the model and context). Thus, such models cannot provide any understanding of the simultaneous experience of the illuminant and surface reflectance, or of the perception of two lightness values in the same visual direction, as occurs in the perception of transparency.

Although framework, filter, and layer models can explain the perception of surface lightness (or brightness) in a variety of displays, there are a number of phenomena that are difficult to explain with the frameworks, filtering, or filling-in approach. One of the oldest is the Mach (1886/1959) card. In this demonstration, a folded gray card is more strongly illuminated on one of its sides and is viewed from an angle such that it appears as a convex "roof." When the 3D shape of the card is perceived correctly, the card appears to be uniform in lightness, and the luminance difference on the two sides of the card is correctly perceived as an illumination difference. However, when viewed monocularly, the depth of the surfaces in this display is ambiguous and can take on an alternative appearance as a concave corner. In this percept, the two sides of the card appear to differ in reflectance, and the

shadowed side appears to have a significantly lower reflectance than the directly illuminated side. Mach argued that this result demonstrates that the visual system uses information about the direction of illumination to compute surface lightness, a conjecture supported by a series of studies by Beck (1965).

Another phenomenon that does not seem readily explained without layered image representations is Hering's outline shadow demonstration. In this illusion, a shadow is cast onto a uniformly colored surface (such as a white piece of paper) and is perceived appropriately, i.e., the decrease in luminance within the shadow is perceived as a change in illumination. However, if the boundary of the shadow is obscured by a thin black outline (say, with a black marker), the shadowed region now appears as a dark stain (i.e., a change in reflectance). The transformation in perceived lightness in Hering's demonstration appears to be a clear example of a phenomenon that arises as a failure to decompose the shadowed region appropriately, causing the decrease in luminance within the shadow to be attributed to reflectance rather than illumination.

Note that for both demonstrations, even if a frameworks theory, filtering theory, or filling-in model could explain why a shadowed region is seen as having the same surface lightness as its neighboring surface in direct illumination, it could not explain why the shadow is also simultaneously experienced. Indeed, a great shortcoming of all such models is that they cannot explain how both the surface lightness and the illuminant are simultaneously experienced in the same visual direction, or why objects that appear to be partly in one framework and partly in another appear to have to have uniform surface lightness (but for a recent attempt to address this in the context of a frameworks model, see Zdravković, Economou, & Gilchrist, 2006).

One common explanation of the Hering shadow demonstration is that the black outline eliminates the shadow penumbra, causing the edge to be reclassified as a reflectance boundary. A problem with this explanation is that many illumination edges can be sharp yet nonetheless are perceived veridically as illumination boundaries. Some striking examples were recently presented by Kennedy and Bai (2000) and may shed new insights into the cause of Hering's illusion. Kennedy and Bai applied Hering's outline technique to two-tone "mooney" faces, i.e., images in which the only information about 3D shape arises from interpreting the dark regions in the image as shadows. In these images, all of the candidate shadow boundaries were sharp—i.e., devoid of penumbras—yet the visual system correctly classified the dark regions of the image as shadows that arose from a 3D object. However, when the image regions corresponding to the shadows are surrounded by thin black outlines (i.e., outlines darker than the shadowed regions), the perception of 3D structure in these images is abolished. Critically, this loss in perceived 3D structure only occurs when the outline is darker than the enclosed shadow; if the outline

is lighter than the shadow (but still clearly visible), the 3D shape could still be readily perceived. Kennedy and Bai suggested that the black outline destroys the perception of a shadow in both their and Hering's demonstration because the contrast polarity of the transition from outside to inside the shadow is incompatible with the physics of shadows. Kennedy & Bai (2000) previously showed that the perception of 3D structure from shadows only occurs if the contrast polarity of the border separating the shadowed region from the surrounding regions has the correct sign, i.e., if the shadow is darker than the surrounding regions. When a black outline encloses a shadowed region, the contrast polarity from the surround to the shadow is inconsistent with an illumination change (i.e., it is a dark-light transition, rather than a light-dark transition). The inappropriate polarity relationship of the shadow boundary causes the shadowed region to appear as a change in reflectance rather than as a shadow, which in turn has a large impact on the perceptual organization of both Hering's and Kennedy and Bai's displays. In Hering's demonstration, this manipulation leads to a transformation in perceived surface lightness; in Kennedy and Bai's displays, it destroys the perception of 3D structure. These phenomena reveal the intimate link between the computation of reflectance, illumination, and 3D shape and hence the need for a theoretical framework that embodies computations that underlie these properties.

Some of the most phenomenally explicit evidence for the role of scission in perceived lightness occurs in conditions of transparency. Although early investigations focused on the geometric and photometric conditions that induce percepts of transparency (e.g., Gerbino, Stultiens, Troost, & de Weert, 1990; Metelli, 1974a, 1974b; Metelli, Da Pos, & Cavedon, 1985), recent work has shown that such decomposition can also have a large impact on the perceived lightness of a target (Anderson, 1997, 1999; Anderson & Winawer, 2005; cf. Adelson, 1993). The physics of transparency is deeply related to the physics of illumination, so there is good reason to believe that the computations underlying the perception of transparency are related to those involved in recovering surface lightness. There are two ways that scission can induce transformations in perceived lightness in conditions of transparency: by changing the apparent lightness of the near (transparent) layer (Anderson, 1997, 2003b; Taya, Ehrenstein, & Cavonius, 1995; Tse, 2005) or by transforming the lightness of the underlying layer (Anderson, 1999, 2003a, 2003b; Anderson & Winawer, 2005; Kingdom, Blakeslee, & McCourt, 1997; cf. Adelson, 1993). In both cases, the perceived lightness of the layers are inversely related: If scission causes the near layer to appear light, the more distant layer is perceived to be darker; if scission causes the near layer to be dark, the more distant layer is perceived to be lighter.

One theoretical advantage of using transparency displays to evaluate whether layered image representations

impact the perception of lightness is that such images elicit explicit percepts of layers; that is, the viewer perceives distinct surfaces along a line of sight separated in depth. Thus, transparency displays provide a powerful method to determine whether a layered image representation is being generated, which can then be used to assess whether the predicted shifts in lightness that should accompany such decompositions do or do not occur.

In addition to its relationship to illumination, the computation of transparency is also intimately related to the computation of occlusion relationships. The concept of layered image representations has been used to describe the decomposition of images into occluding and occluded surfaces (Kanizsa, 1979; Koffka, 1935),[1] as well as the separation of luminance into the types of maps suggested by intrinsic image analysis (i.e., surface orientation, reflectance, and illumination). The computation of transparency is deeply related to both usages. The transmittance of a transparent layer refers to the proportion of light reflected from an underlying surface that is not occluded by the transparent surface (or media). From the perspective of image generation, the difference between occlusion and transparency is therefore simply a matter of degree. Indeed, this link between occlusion and transparency is explicit in Metelli's episcotister model of transparency, where the transmittance of the transparent surface is determined by the size of the missing sector—the holes—in an otherwise opaque, occluding disc. Thus, in Metelli's model, occlusion is simply a limiting case in which the size of the missing sector goes to zero; complete transparency occurs when the area of the episcotister goes to zero.

In this paper, we provide further evidence demonstrating that scission *qua* transparency can cause dramatic transformations in perceived lightness and articulate the intimate relationship between occlusion, transparency, and lightness perception. In what follows, we present new phenomena demonstrating that the decomposition of images into layered representations can depend critically on the contrast relationships that occur along contours and that these relationships can have a dramatic role on the perception of depth, lightness, and opacity. The experiments and the demonstrations presented herein focus on the computations that underlie the formation of transparent and occluding surfaces, where the perception of layers is explicit. In what follows, we use the term "scission" to refer to the explicit decomposition into a transparent and underlying layer. It should be noted that intrinsic image analysis asserts that scission is ubiquitous and is used to recover the 3D structure and reflectance properties of all scenes. Here, we restrict our experimental and theoretical efforts to understanding the conditions and theoretical principles needed to initiate the decomposition of an image into a transparent layer and underlying surface and the consequences of this particular form of decomposition on perceived lightness. Our aims are as follows: first, to demonstrate that the decomposition of an

image into transparent layers can have powerful effects on the perception of surface lightness; and second, to characterize the image properties that are important for decomposition into layers in these displays. To the extent that the perception of transparency strongly and directly modulates lightness judgments, it will be critical for any complete theory of lightness perception to account for layered representations. We will consider the broader issue of scission *qua* intrinsic image analysis in the discussion. We present evidence first in the form of new displays, articulating the associated perceptual phenomena and the theory needed to understand them, followed by a series of psychophysical experiments that parametrically assess the percepts obtained with these displays. Some of these results have been presented previously in abstract form (Anderson & Winawer, 2006).

Phenomena and theory

To explore the relationship between transparency, lightness, and occlusion, we begin with some very simple displays that evoke percepts of transparency and articulate some basic principles of perceptual organization that determine the perceived depth, lightness, and opacity of the surfaces in such images. We will then progressively increase the complexity of our displays to reveal the generality and the robustness of these forms of image decomposition and the principles of perceptual organization that underlie them.

Consider the two images in Figure 1. The luminance values within the series of circular contours are identical in the left- and right-hand images; the only difference between these images is the luminance values in the surround. Note, however, that there is a very significant difference in the way the visual system computes the perceived lightness of the circular targets and the transparency relationships in the two images. In the left image, the discs appear nearly black, becoming progressively obscured by a milky transparent overlay on the upper right half of the image. In the right image, the discs appear nearly white, becoming progressively obscured by a dark transparent overlay on the bottom left of the image. The geometric factors that contribute to the transparency percepts evoked by these images include the continuity of the diagonal contours defining the near transparent layer and the continuity of the contours of the underlying discs. These geometric properties are identical in the two halves of Figure 1, so they cannot account for the perceived differences between the two figures. Rather, the difference in perceived lightness of the discs, the perception of transparency, and the relative depth of the surfaces in the image all arise from the differences in the photometric relationships between the circular figures and their surrounds. The problem, then, is to explain how
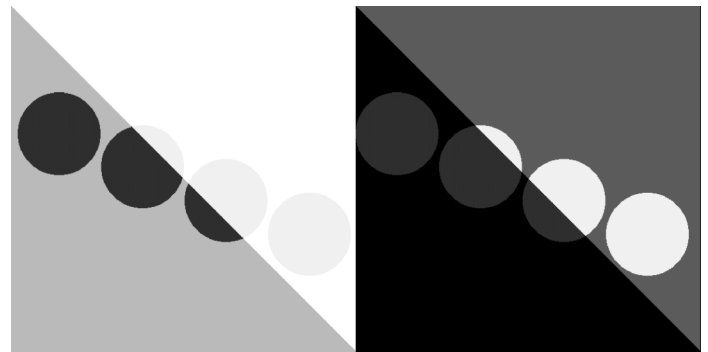

Figure 1. Simple transparency. The target discs in the left display and the right display are identical. The left circles appear as black discs partially obscured by a white transparent surface in the upper right, whereas the right discs appear as white discs partially obscured by a dark transparent surface in the lower left. The perceived lightness of the discs are determined by the regions that appear to be in plain view—the dark regions in the left display and the light regions in the right display. The other regions of the targets are scised—the luminance in light regions in the left figure are attributed to a combination of a light transparent layer and a more distant opaque black layer; the dark regions in the right figure appear to arise from a dark transparent layer and a more distant opaque white layer. Click the image to see the moving version.

changing the luminance values of the surround leads to such a dramatic transformation in the perceived lightness and transparency relationships in these images.

Any explanation of the percepts elicited by these images must account for the shift in perceived lightness of the discs; why transparency is perceived; the perceived lightness of and/or transmittance of the multiple surfaces; and why (and which) portions of these images appear in plain view. In what follows, we will argue that these computations are inherently coupled and depend critically on the local contrast variations that occur along the boundaries of partially occluded figures. (Here, we use the term "partially occluded" to include the partial obstruction of a distant surface by a transparent layer.)

Contrast and transparency

There are a number of physical constraints on the pattern of image contrasts formed by transparent surfaces and changes in illumination that provide information about the presence or absence of transparent surfaces. A growing body of evidence has shown that these constraints are used by the visual system to determine when transparency is and is not present. In conditions of overlay, transparent layers can only decrease or leave unchanged the contrast of underlying contours or textures; they cannot magnify the magnitudes or invert the polarity

relationships of the contrasts generated by underlying surfaces. If no contrast reduction occurs in a region of transparency, but there is a change in overall luminance level, the transformation of a transparent layer is purely multiplicative and therefore equivalent to a change in illumination. There are thus two photometric cues that can be used to infer the presence of transparent layers and/or illumination changes: the preservation of contrast polarity of contours, textures, or other surface markings; and the reduction (or preservation) of contrast magnitudes that accompany changes in overall luminance magnitudes. Several reports have shown that the polarity and magnitude of contrast plays a critical role in determining whether transparency is perceived (Adelson & Anandan, 1990; Anderson, 1997, 2003a, 2003b; Anderson & Julesz, 1995; Beck & Ivry, 1988; Kanizsa, 1979; Metelli, 1974a, 1974b; Singh & Anderson, 2002).

In Figure 1, these dimensions have been manipulated to induce percepts of transparency. In the left display, both the circular contours and the diagonal contours preserve contrast polarity, which implies that either contour is consistent with being overlaid by a transparent layer. In this and all of the other transparency displays presented herein, the contour junctions always preserve polarity across both sets of contours, which means that the region containing the transparent layer and the depth ordering of the layers in the transparent region are ambiguous with respect to the polarity constraint. Other constraints are therefore needed to resolve this ambiguity. These include figural constraints (such as those that bias smaller, convex regions to appear as figures) and constraints on the magnitude of contrast across the contour junctions. In Figure 1, figural constraints bias the circular regions to appear as discs (rather than holes), and the strong reduction in contrast along a portion of the circular contours provides a depth cue that the low contrast segment of the contour is being overlaid by a low-transmittance transparent layer (top right for the display on the left and bottom left for the display on the right). The bias to interpret large reductions in contrast along a contour as the presence of a transparent overlay makes intuitive sense and highlights the close relationship between transparency and occlusion: In the limit where the contrast of the low contrast segment goes to zero, the display becomes a simple occlusion display, and the X-junctions are transformed into T-junctions that signal occlusion (see Figure 2). Note that the perceived depth of the discs can also be reversed in the X-junction displays by varying the relative contrast magnitudes of the diagonal and circular contours. If the contrast across the diagonal is low relative to the contrast across the circular contours, the discs appear as transparent layers that lie on top of a two-toned background. Indeed, in these luminance regimes, the display is essentially identical to those studied extensively by Metelli.

Although models of transparency have focused on articulating the conditions that lead to the perception of

Figure 2. Simple occlusion. The display is identical to that in Figure 1, except that the contrast in the target regions is increased; the light parts of the discs are equal in luminance to the light half of the left surround and the dark parts of the disc are equal in luminance to the dark half of the right surround. This causes the discs to appear partly occluded instead of behind a transparent surface. Click the image to see the moving version.

transparency, there is a related problem that is not usually explicitly addressed; namely, specifying how the visual system determines whether a scene (or a region of a scene) is in plain view. There is a growing body of data demonstrating that the visual system uses contrast relationships to compute both the presence of transparent surfaces as well as their quantitative properties (such as transmittance).[2] But how does the visual system determine that a given image region is actually a portion of a scene in plain view rather than a higher contrast region viewed through a partially transmissive medium? Anderson (1999, 2003a) proposed that the visual system employs a *transmittance anchoring principle* (TAP) to determine when a transparent surface is present. Intuitively, the TAP states that the visual system treats the highest contrast image regions as regions in plain view and only infers the presence of transparent surfaces if there are spatial or spatio-temporal (Anderson, Singh, & Meng, 2006) perturbations in the contrast magnitude along contours, surfaces, or textures. More specifically, Anderson et al. (2006) showed that the visual system uses the region of highest contrast as a *transmittance anchor,* i.e., as the normalization factor used to compute the transmittance of transparent surfaces in lower contrast regions of the image. Note that the TAP only applies to partially occluded contours since contrast variations along an occluding contour are consistent with surfaces in plain view behind an occluding edge. Thus, there are a number of aspects of the contrast variations in Figure 1 that contribute to the percepts of transparency and shifts in perceived lightness of the circular targets. First, the geometric continuity of the diagonal and circular contours defines the figural unity of both the transparent layer and the discs. These geometric properties are identical in the two halves of Figure 1, so the difference in the perceived lightness, transparency, and depth in these images must be
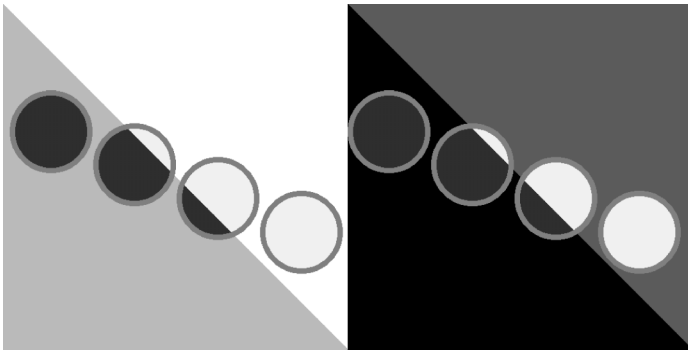
Figure 3. Simple transparency with rings. The display is identical to that in Figure 1, except that the targets have been outlined with thin gray rings. The rings introduce alternations in contrast polarity along the inner contour of the discs, such that the discs are sometimes lighter and sometimes darker than the surrounding gray ring. This interferes with the photometric cues to transparency and the geometric continuity with the surround, greatly diminishing the difference in perceived lightness between the targets in the left and right displays. Click the image to see the moving version.

due to the photometric differences in the displays. The preservation of contrast polarity along the continuous contours in this figure is consistent with either contour arising from a transparent surface or illumination change. However, the strong reduction in contrast magnitude of the circular contour provides information that these contours are obscured by a light (left image) or dark (right image) transparent surface. Finally, the TAP states that the highest contrast portions of the contour segments are perceived in plain view. This corresponds to contours belonging to the dark portions of the circles in the left image, and the light portions of the circles in the right image, giving rise to the striking difference in perceived lightness of the discs in the two images. This shift in perceived lightness arises from a reversal in the portions of the discs that are seen in plain view. In the limiting case, this principle reduces to conditions of occlusion, and portions of the circular figures are no longer visible (see Figure 2).

If this account is correct, it should be possible to destroy the percepts of transparency and the shifts in perceived lightness in these figures by reversing the contrast polarity along the far contour (the circular borders). One simple way to accomplish this goal is to cover the circular boundary with a gray ring, i.e., a luminance value that falls between the two colors within the circular boundary (see Figure 3). This manipulation causes the polarity along the ring and the texture within its interior to reverse (being light/dark along some portions, and dark/light along others). As can be seen in Figure 3, this manipulation blocks the percept of transparency in these images. Unlike the discs in Figure 1, the discs in Figure 3 are not seen as having a uniform surface lightness. This largely

eliminates the difference in perceived lightness of the regions enclosed by the circular contours in the left and right images, although a small brightness/lightness difference remains in both the light and the dark portions of the discs. It is currently unclear whether such residual lightness/brightness effects arise from a scission mechanism or whether it signifies a different process. We will discuss the possible causes and contributions to this residual simultaneous contrast lightness/brightness effect in these displays, and for those that follow, in the General discussion section. We will proceed by showing that the same principles needed to explain the percept of layers in Figures 1 and 2 can be used to predict when scission *qua* transparency occurs in more complex images, which in turn determines when transformations in perceived lightness are (and are not) experienced.

Texture variations

The images described in the preceding section used displays similar to Metelli's episcotister displays (e.g., Metelli, 1974a, 1974b). Note, however, that his displays could not have generated the particular luminance values in the discs in Figure 1 since the discs were always the transparent layers in his model. The same transformations in perceived lightness observed in Figure 1 can be observed in a large variety of images. Figures 4 and 5 provide examples of displays composed of four luminance values each that also invoke percepts of transparency and shifts in perceived lightness. In these figures, the textures within the circular boundaries are identical, and the geometric patterns outside the boundaries are also identical. The only difference is that the two luminance values outside the circular boundaries are higher on the left
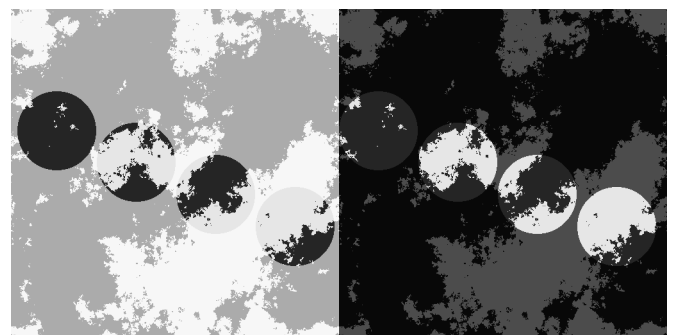

Figure 4. Thresholded clouds. The luminance values in this display are identical to those in the simple transparency display of Figure 1, but the geometric pattern is now random (thresholded noise with a $1/f^2$ power spectrum, with random phase and orientation). Although the discs on the left and the right are physically identical, the circles on the left are perceived as black discs seen through partially transparent white clouds, while the circles on the right are perceived as white discs viewed through dark clouds, similar to the effect in Figure 1. Click the image to see the moving version.
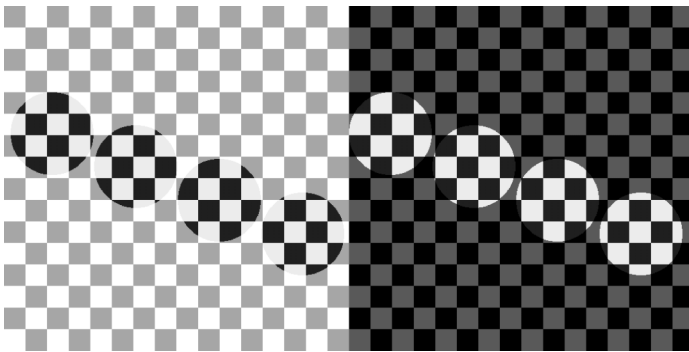
Figure 5. Checkerboard pattern (adapted from Anderson, 2007). The luminance values are identical to those in Figures 1 and 4. The discs on the left and right are identical, though on the left the circles appear as white discs seen through dark checkers and on the right, the circles appear as dark discs seen through white checkers. Click the image to see the moving version.

image than on the right. This creates a uniform polarity light/dark boundary between the background and discs (respectively) on the left and a uniform polarity dark/light boundary on the right. As in the previous display, the highest contrast regions along the circular contours appear in plain view in both figures. These are the black(ish)/mid-gray boundaries in the left image and the white(ish)/mid-gray boundaries in the right image. This causes the disk on the left to appear (nearly) black and the disk on the right to appear (nearly) white. The same effects can be observed when the lower contrast segments of the contours are reduced to zero, and the figures appear as a simple occlusion display.

As with our previous display, the difference in perceived lightness between the left and right images can be largely eliminated by covering the circular boundaries with thin gray rings (Figures 6 and 7), although a small lightness/brightness difference of the textures remains (which we discuss in the general discussion below). This
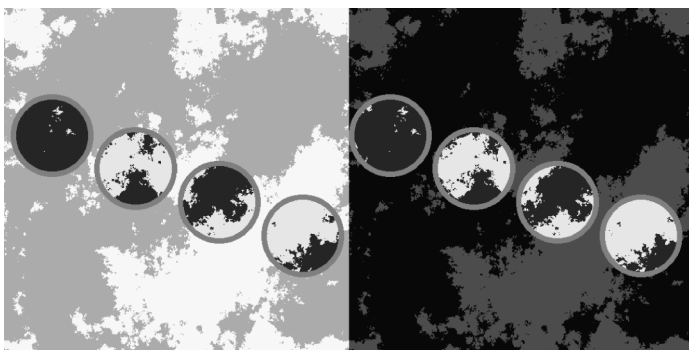

Figure 6. Thresholded clouds with rings. The gray rings, which cause an inconsistency in the contrast polarity along the contours of the discs, disrupts the percept of transparency and hence diminishes the perceived difference in lightness between the circles in the left and right displays. Click the image to see the moving version.
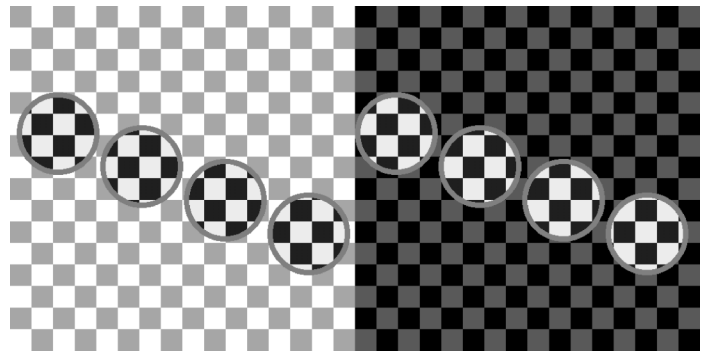

Figure 7. Checkerboard with rings. The gray rings break the percept of transparency. Click the image to see the moving version.

implies that these boundaries play a critical role in the segmentation processes that cause the images to appear layered, which in turn induces shifts in apparent lightness of the figures bounded by these contours.

The same effects can be observed in textures that contain a continuous distribution of luminance values, but the percept of transparency that accompanies these images is more complex (see Figure 8; from Anderson & Winawer, 2005). The textures in this image were created by constructing a noise pattern with a power spectrum of $1/f^4$ and summing the different frequency components with random phases and orientations. This texture pattern was used as a "seed" image that was used to create the target and background patterns. The target and the background regions were spatially identical to the seed image, but the range of luminance values was altered such that the targets spanned the full luminance range (from "black"
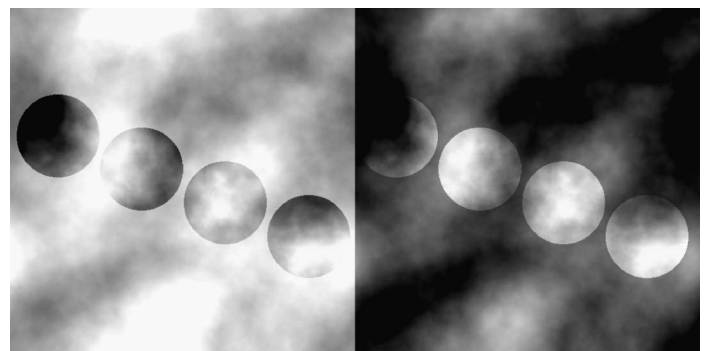

Figure 8. Continuous clouds (adapted from Anderson & Winawer, 2005). The circular targets on the left and right are identical but appear as black discs seen through white haze on the left and as white discs seen through dark haze on the right. The effect is similar to that in the previous figures, except that the luminance values span the full gray-scale range. The darkest regions of the targets in the left display and the lightest regions in the right display appear to be in plain view. The more the luminance departs from these peaks (lighter regions on the left and dark regions on the right), the greater the apparent opacity of the intervening layer. Click the image to see the moving version.
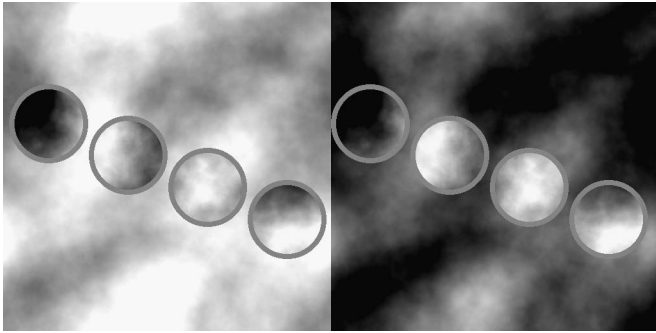
Figure 9. Continuous clouds with rings. The gray rings surrounding the targets destroy the perception of transparency. The perceived difference in lightness is now greatly reduced in the left versus right displays when compared to Figure 8. Click the image to see the moving version.

to "white") and the surround distributions were compressed and shifted upward (for the light surround) or downward (for the dark surround). For an appropriate choice of surround luminances, this created a uniform polarity light/dark boundary between the background and the discs (respectively) on the left and a uniform polarity dark/light boundary on the right. The discs on the left appear as (nearly) black discs visible through light mist, whereas the discs on the right appear as (nearly) white discs visible through black smoke. Note, however, that whereas the previous displays generated percepts of homogeneous (or "balanced") transparency, the continuous luminance variations in these displays evoke a percept of inhomogeneous transparency, i.e., the transparent layer appears to vary in transmittance. The transmittance of the transparent layer appears to be 100% where the contrast along the disc contour is greatest (i.e., these regions appear in plain view, as predicted by the TAP), and the perceived transmittance appears lowest in the regions of the image where the contrast of the target/background borders is lowest. In the limit where the contrast of portions of the contours goes to zero, the mist or the smoke appears completely opaque.

As with the previous displays, the apparent difference between the two sets of targets can be almost completely eliminated by covering the circular boundaries with thin gray rings (see Figure 9).

Geometric and photometric contributions to transparency

The preceding demonstrations show that the way that images are decomposed into layers can have a dramatic impact on the perception of lightness. These demonstrations also reveal that the computations underlying this decomposition depend critically on the boundaries between the targets and their surrounds. Although we introduced the rings to determine the effect of polarity

constraints along the contours, the rings also introduced a gap between the texture within the target discs and the texture in the surround. In other words, the texture contours that cross the circular target-surround border are missing in the gap created by the narrow, untextured ring region. Thus, it is unclear whether the disruptive effect of the rings on scission was due solely to photometric constraints, or whether the geometric break in continuity between the targets and the surround caused by the rings also contributed to the ring's impact.

One way to assess the relative contribution of the geometric and photometric effects of the rings is to introduce rings that retain the photometric conditions favorable to scission, while introducing a gap between the texture within the targets and their surrounds. This can be accomplished by surrounding the targets with white rings on the light background (generating a light/dark boundary from the surround/target) or dark rings on the dark background (generating a dark/light boundary from the surround/target). Examples are presented in Figures 10 and 11.

When viewed monocular (i.e., only one half of either figure), observers report that this manipulation gives rise to one of two percepts. In one organization, the rings appear as apertures in front of the entire display; the surrounding region appears as a window that reveals a new, higher contrast texture within the target region. The percept of scission of the targets does not occur in this configuration, and the difference in lightness between the targets on the two surrounds is largely abolished. In the alternative percept, the rings appear to surround a
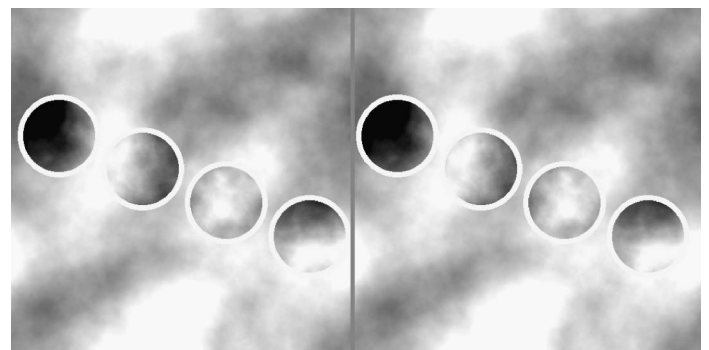

Figure 10. Equivalent to the left side of Figure 9, except that the boundaries of the discs have been surrounded by white rings instead of gray. The targets in the monocular displays are bistable and can appear as either opaque black discs surrounded by white rings, or the rings can appear in front of the surround texture as an aperture that reveals a region containing a higher contrast texture. A percept of transparency can be induced with binocular disparity that places the rings behind the surround (cross fuse the left and right of the displays). The disparity of the boundaries cause the edges to appear behind the texture, which induces scission and a percept of opaque black discs surrounded by white rings. Click the image to see the moving version.
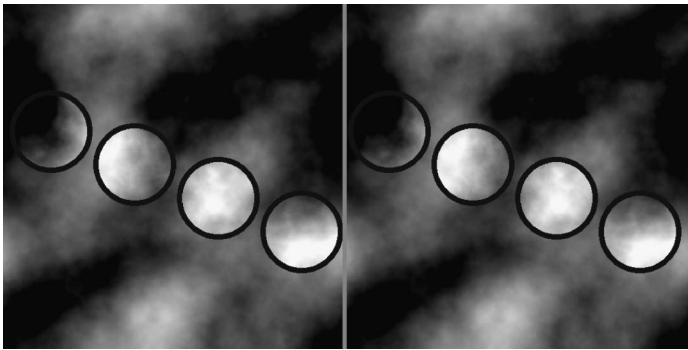
Figure 11. Equivalent to the right side of Figure 9, except that the discs have been surrounded by black rings instead of gray. The monocular percepts are bistable. However, when disparity places the rings behind the texture, a percept of homogeneous white discs surrounded by black rings is evoked. The figure should be viewed by cross-fusing the two sides of the display. Click the image to see the moving version.
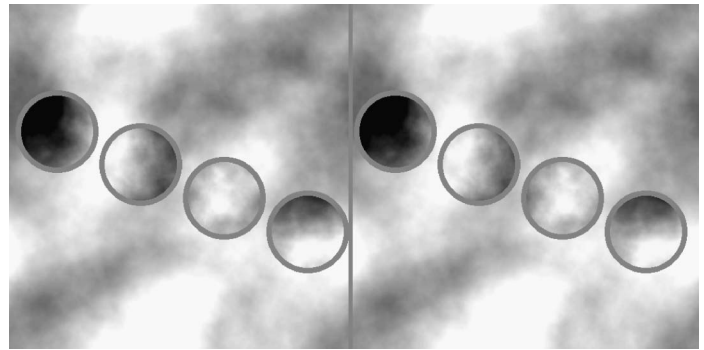

Figure 12. Continuous clouds with gray rings in stereo. Equivalent to the left side of Figure 9, except that stereoscopic depth has been added. Just as in the monocular case, in stereo the gray rings also eliminate the effects of scission, such that the targets no longer appear as uniform opaque black discs, but may appear black in some places and white in others, depending on the contrast polarity along the circular contour. The figure should be viewed by cross-fusing the two sides of the display. Click the image to see the moving version.

homogenously colored disc that lies in a more distant depth plane behind a cloudy transparent layer. In particular, when the surround is light and the rings are white (Figure 10), a white ring appears to surround a homogenous black disc; when the surround is dark and the rings are black (Figure 11), a black ring appears to surround a homogeneous white disc.

Each of the two possible depths associated with the rings can be enforced, and the bistability abolished, by adding a binocular disparity to the ring that places it either in front of or behind the texture of the surround (respectively). When Figure 10 is viewed stereoscopically, the targets appear as opaque black discs surrounded by white rings behind white clouds. In Figure 11, the targets appear as opaque white discs surrounded by black rings. However, when the same stereoscopic manipulation is applied to gray rings, the far depth organization does not produce a coherent lightness percept (Figure 12). Rather, the far disc appears to change in lightness depending on the local contrast relationship of the texture relative to the ring in the far plane.

Moreover, if the luminance values of the rings are reversed such that the rings on the light surround are black and on the dark surround are white, neither of these percepts is evoked, and most observers do not report any stable percept of transparency or any clear lightness difference between the two sides of the displays (Figure 13). In this arrangement, the contrast polarity along the boundaries of the rings and targets is opposite in sign to the polarity along the targets and surround in the absence of the rings. For example, in Figure 13 on the left, the rings are darker than the target, but the surrounding region is lighter.

Strictly speaking, the polarity constraints alone are consistent with transparency in these displays (Figure 13) since the polarity along the contour retains its sign along

its entire length. However, the luminance range is higher within the targets than in the surrounds. Thus, when black rings are placed on white surrounds, the polarity constraint is consistent with the target being a white disc surrounded by a black ring on a white surround (or an empty black ring floating on a white background). However, because the dark regions within the discs are actually darker than any luminance in the surround, there is no sensible interpretation by which a white underlying surface can cause a dark transparent layer to appear darker over the region of the (white) targets. Thus, the instability of these figures most likely arise from a violation of a contrast magnitude constraint.

The principles underlying these transformations in perceived lightness and depth have been discussed in detail elsewhere (Anderson, 1999, 2003a). The important
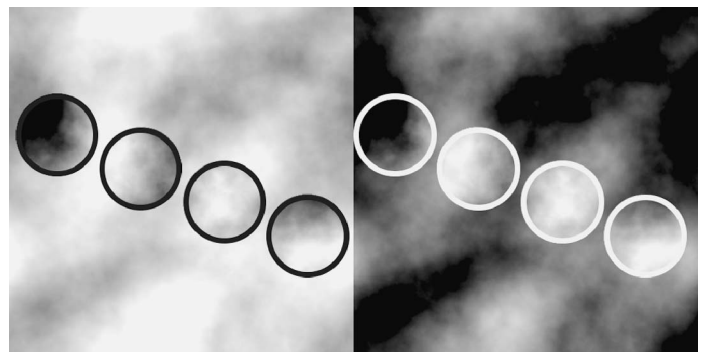

Figure 13. When black rings are placed around the targets on a white surround, or white rings around the targets on a black surround, no stable percept of transparency is reported. Click the image to see the moving version.

issue for the present work is that the geometric transformations induced by the rings are essentially identical in all of these displays; only the photometric properties (the relative luminance of the ring, surround, and discs) have been altered. Thus, these phenomena reveal the critical role of the photometric properties of the ring in modulating the occurrence (or nonoccurrence) of scission in these displays.

Experiments

We performed four sets of experiments to quantify and to document the phenomenological descriptions of the displays presented above. In all of the experiments, three types of textures were used: the simple transparency display (Figure 1); the thresholded texture display (Figure 4); and the continuous texture display (Figure 8). In the first experiment, subjects performed a matching task where they adjusted the luminance of a test patch to match the perceived lightness of the target (the luminance ranges within the target were varied from trial to trial). For targets on the continuous texture displays, this experiment replicated our previous results (Anderson & Winawer, 2005). For the thresholded and the simple transparency displays, this experiment extended the lightness matches to new stimuli. In the second experiment, observers rated how compelling the target disc appeared as a uniform, opaque disc for displays containing different colored rings (black, white, gray, or none). In our third experiment, we manipulated the stereoscopic depth of the rings, and observers were required to report whether the target regions appeared black, white, or not uniform/not opaque. In our fourth experiment, the thickness of the rings was varied, and the impact on the perception of scission and surface lightness was measured. Each of these experiments is described in detail below.

Experiment 1: Lightness matches
Subjects

Six naive subjects were recruited at Stanford University and participated in the experiment for course credit or cash payment. All subjects gave informed written consent.

Stimuli

Three types of textured stimuli were used for all of the targets and surrounding regions: a simple transparency display (Figure 1); a thresholded texture display (Figure 4); and a continuous texture display (Figure 8). The simple transparency and the thresholded texture displays contained the same luminance values and hence only differed in their geometric distribution of these

luminance values. As in Figures 1, 4, and 8, the displays consisted of a circular target (3.8° in diameter) made from a high contrast seed image on a dark or light surround composed of the same geometric pattern as the target. The target moved continuously back and forth across the surround at a rate of 1 Hz.

The luminance range of the high contrast seed images used to generate the targets was varied from trial to trial. For both the simple transparency and the thresholded textures, the two luminance values for the targets were 3 and 137 cd/m^2; for the continuous texture they ranged from 2 to 153 cd/m^2. For all three textures, additional targets were generated by reducing the luminance range in 4 approximately equal steps, so that there were 5 targets made from each texture. The minimal and the maximal pixel values for each target are indicated by the thin black lines in Figure 14.

The luminance ranges of the surround textures were as follows. For the light surrounds made from the transparency and thresholded textures, the luminance values were 72 and 153 cd/m^2; for the dark surrounds they were 1.6 and 11 cd/m^2. For the light version of the continuous texture displays, the luminance values ranged from 23 to 153 cd/m^2 (median of 103 cd/m^2) and for the dark versions from 1.6 to 95 cd/m^2 (median of 7 cd/m^2).

All displays were presented on a uniform gray background (44 cd/m^2) in a dark testing room using a 17-in. Apple Studio Display LCD monitor. The full gray scale luminance range of the monitor was measured every few days over several months with a Minolta chromometer (LS-110) to ensure that the appearance of the stimuli did not change over time. Subjects viewed the displays from a distance of 50 cm.

Methods

The stimuli were presented one at a time to subjects in random order. A static test circle equal in size to the moving target was displayed on a white noise pattern, a 6° × 6° square, and subjects adjusted the luminance of this test circle by sliding a mouse vertically over the screen. The initial luminance of the test region was randomized on every trial and the white noise pattern behind the test patch changed on every trial. The task was to adjust the luminance of the test circle such that it appeared to have the same surface lightness as the moving target. Subjects were instructed to click a button for "no match possible" if the moving target did not look like an opaque disc composed of a uniform color. This occurred on only 3.7% of trials. There was a 500-ms gray screen between trials and no time limit on the matches.

Each of 30 target-surround combinations (5 target contrast values × 3 texture types × 2 luminance values for the surrounds) was viewed and matched once in random order in a practice block and then once in each of two subsequent experimental blocks.
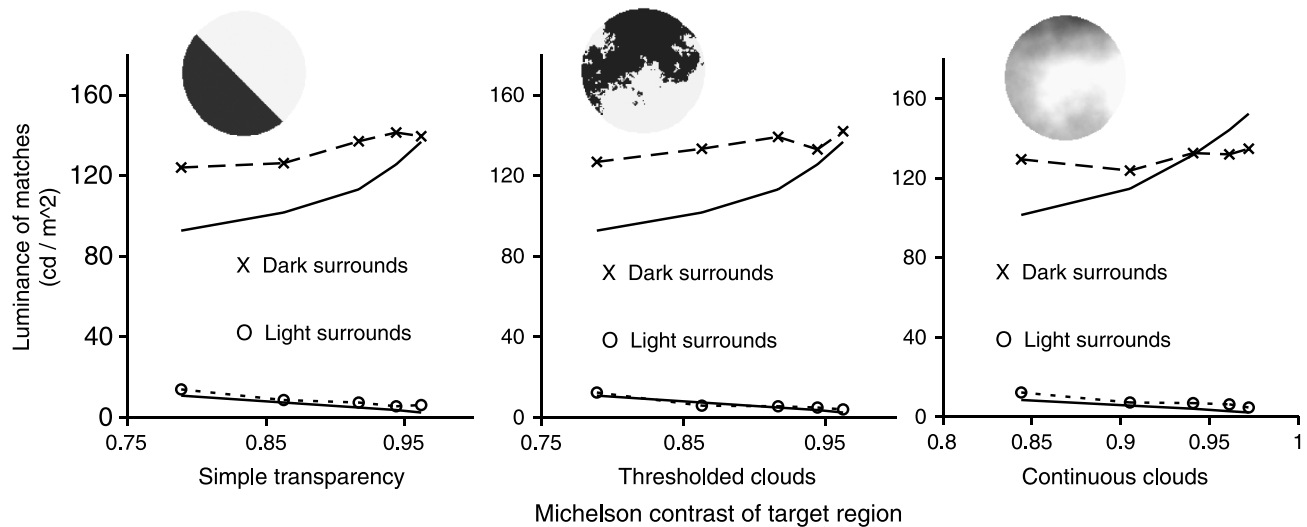
Figure 14. Experiment 1: Lightness matching. Subjects adjusted the luminance of a disc on a white noise background to match the apparent lightness of targets in either the simple transparency displays (Figure 1), the thresholded texture displays (Figure 4), or the continuous displays (Figure 8). The values of the matches are indicated on the y-axis. The contrast of the target regions (x-axis) was manipulated in each texture type. When the discs were shown on the light surrounds, the matches (open circles) indicated a nearly black percept on each texture type, whereas when the same discs were on dark surrounds, the matches (x's) were nearly white. The solid lines indicate the minimum and the maximum luminance value in the target regions. The general pattern in each display is that the matches on dark surrounds are close to the maximal value and on light surrounds close to the minimal value.

Results

The results from this experiment are presented in Figure 14. The circular insets indicate the type of target/surround used for each image. The most salient aspect of the results is that the identical target patches looked white when on dark surrounds and black when on light surrounds for each of the three textures tested (Figure 14). The difference in luminance of the lightness matches for the same target when on a dark versus a light surround, even for the targets with the lowest contrast, was greater than half of the luminance range available on the monitor. Across different target contrasts, the lightness matches were close to the minimum target luminance for targets on light surrounds and close to the maximal target luminance on dark surrounds, consistent with our previous report that assessed only the continuous cloud version of the stimulus (Anderson & Winawer, 2005).

Discussion

The results from Experiment 1 show that the decomposition of the image into layers can strongly affect which regions are perceived to be in plain view, which in turn transforms the perceived lightness of the entire discs. For targets that appear dark, the luminance settings chosen by observers in making their lightness judgments correspond almost perfectly to the minimum luminance value within the target region texture. For targets that appear light, the settings chosen by observers are somewhat higher than the luminance values of the lightest regions within the targets, a trend most apparent in the simple transparency and

thresholded clouds stimuli. This trend is consistent with a bias for the highest luminance to appear white (Gilchrist et al., 1999). The exact cause of this departure is not known and will not be evaluated in detail herein since it constitutes a small fraction of the difference in the lightness observed for targets on light and dark surrounds. It may arise from the misperception of the matching pattern's lightness or brightness induced by the random dot background on which it was superimposed, or it may arise from an asymmetric induction of dark backgrounds on light image regions. Our current method does not allow us to establish a "ground truth" that can allow us to resolve these possibilities. We will therefore focus on assessing the impact that the target/surround boundaries have on the perception of transparency and the lightness transformations that such decompositions induce.

Experiment 2: The effect of different colored rings on the perception of transparency

The purpose of our second experiment was to assess the impact of surrounding the circular target region with a thin ring. The primary intent of this manipulation was to assess the role of the surround/target boundary in inducing the perception of transparency and the accompanying changes in the perceived uniformity of the lightness of the targets in the image. More specifically, we were interested in determining the impact of surrounding the target disc with rings of different luminances that either preserved

contrast polarity along the boundary, or that did not preserve contrast polarity.

Subjects

Four naive subjects were recruited at Stanford University and participated in the experiment for course credit or cash payment. All subjects gave informed written consent.

Stimuli

The textured stimuli used in Experiment 1 were also used in Experiment 2, but only the highest contrast target of each texture type was used. The targets were either surrounded by a thin ring (8 pixels or 0.24 deg of visual angle) that was gray (44 cd/m^2), black (1.8 cd/m^2), or white (146 cd/m^2); or, as in Experiment 1, the targets were viewed with no ring. As in Experiment 1, the target moved continuously back and forth across the surround at a rate of 1 Hz. The experiment thus consisted of 24 different stimuli: targets made from 3 texture types, each on either a light or dark surround, each with one of 4 ring conditions. The testing room, the monitor, and the viewing distance were as in Experiment 1.

Methods

Subjects viewed a single display and were asked to give a rating from 1 to 7 on the extent to which the target region appeared as a solid, uniform disc surrounded by a colored ring. Subjects were instructed that a score of 7 meant the moving disc was perceived as being definitely opaque and uniform in color, and that a score of 1 meant the disc was perceived as definitely transparent or non-uniform in color. Scores between 1 and 7 were to be used when there was not complete certainty about the opacity or uniformity of the color. Subjects viewed and rated the 24 stimuli once each in random order in a practice block, and then once in each of two blocks in random order during the experiment.

Results

As expected, when the test stimuli were not surrounded by rings, scission ratings were consistently high (Figure 15, red bars). On the light surrounds, the mean scission ratings were 6.3 (.4), 6.6 (.1), and 6.3 (.4) for the transparency, the thresholded, and the continuous textures, respectively (standard error of the mean in parenthesis). For the dark
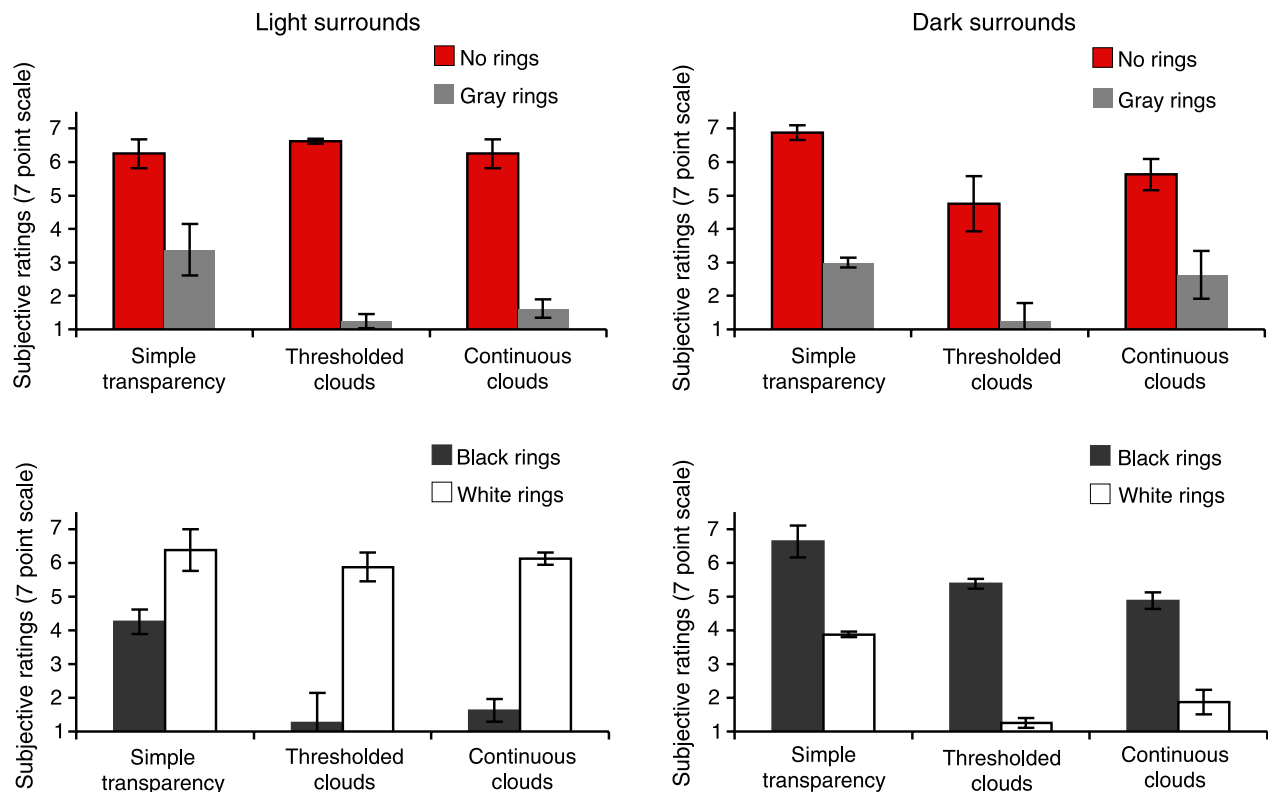

Figure 15. Experiment 2: Effect of rings on quality of scission. Subjects rated on a 1–7 scale the clarity that the target disc appeared to be a uniformly colored, opaque disc for each of three texture types (simple transparency, thresholded textures, and continuous textures). When the discs were surrounded by rings and transparency was experienced, the discs appear as uniformly colored discs surrounded by a uniformly colored ring. When there were no rings around the targets (red bars), ratings were consistently high. Gray rings significantly reduced the percept of transparency and the extent to which the discs appeared uniform in all displays (gray bars). Dark rings interfered with transparency on light surrounds (lower left) but not dark surrounds (lower right). Light rings had the opposite effect, reducing the quality of transparency for the dark but not light surrounds.

surrounds, the mean ratings were 6.9 (.2), 4.8 (.8), and 5.6 (.5). These results are consistent with the first experiment in which subjects saw the target as being uniform and opaque in almost all displays.

The presence of the gray rings severely impaired the percept of transparency (Figure 15, gray bars). For each subject and for each of the 6 background conditions (3 texture types, each with dark and light surrounds), the ratings were lower with gray rings than with no rings, except for two cases in which the ratings were identical for both displays (individual subject data not shown). Thus, the fluctuations in contrast polarity along the target boundary greatly impaired the decomposition of the target region into a layered representation.

When the circular targets were enclosed by white or black rings, a very different pattern was observed. On light surrounds, white rings did not impair scission but black rings did. The scission ratings for the targets with white rings were about the same as ratings with no rings. In contrast, the ratings with black rings were about the same as the ratings with gray rings. The opposite pattern was observed with the dark surrounds (although the target regions were identical): the black rings did not affect ratings of scission but the white rings did.

Discussion

The results of Experiment 2 revealed the importance of the photometric relationships along the target/surround boundaries in determining the clarity with which the image is decomposed into layers. In general, the gray rings disrupted the perception of transparency for all target/surround combinations for all texture types. The gray ring causes the polarity between the target region and the surrounding ring to reverse, thus blocking the perception of transparency and the accompanying percept of the discs as having uniform surface lightness. In contradistinction, the black rings primarily disrupted the perception of layers only on white surrounds, and the white rings primarily disrupted the perception of layers only on dark surrounds. This pattern of results reveals that the ring's disruptive effect on transparency cannot be entirely attributed to the preservation of contrast polarity along the ring's contour, as the dark and light rings both preserve polarity on both the dark and light surrounds.

There are two factors that may contribute to the specificity of the disruption induced by black and white rings. First, the contrast between the rings and the surround is greater on the displays for which they disrupted the perception of transparency (black rings on white surrounds and white rings on black surrounds). The stronger boundaries—i.e., more salient gaps—may interfere with the formation of a coherent transparent surface between the target region and the surround. Second, the contrast polarity between the ring, the surround, and the target requires the opposite pattern of lightness to be

experienced within the target region than is experienced in the absence of the rings. For example, when the rings are black on a light surround, the polarity relationships are the same on both sides of the ring's inner and outer boundaries. Thus, if any percept of transparency is experienced, the polarity relationships require that the black ring should appear behind dark clouds, and that the lightness of the far layer within and outside the ring boundary should be white (as constrained by the TAP). Although these polarity relationships are perfectly valid and supportive of transparency, the particular way in which these stimuli are constructed violates magnitude constraints on transparency.

In sum, there are two possible reasons for the selective disruptive effects of the black and white rings on the perception of transparency: the strength of their contrast relative to their surrounds; and the "accidental" contrast changes that would have to be accounted for to perceive transparency in the displays most disrupted by these rings. In an attempt to disentangle these possible factors, in Experiment 3 we manipulated the stereoscopic depth of the ring. Our informal observations revealed that the disparity manipulation was equally effective in placing the ring in a separate depth plane for all luminance values, allowing us to determine whether the clarity of the ring per se led to the particular pattern of results observed in Experiment 2.

Experiment 3: Stereoscopic depth and perceived surface lightness

Subjects

Six naive subjects were recruited at Stanford University and participated in the experiment for course credit or cash payment. All subjects gave informed written consent.

Stimuli

Experiment 3 used the same stimuli as Experiment 2, except that the binocular disparity of the target regions was manipulated. All 24 stimuli from Experiment 2 (3 target textures × 2 luminance surrounds × 4 ring conditions) were displayed in each of 3 conditions, such that the there was 0.3° of near disparity (with the target in front of the surround), 0.3° of far disparity, or no disparity. Disparity was introduced by shifting the position of the circular boundary of the target regions in the left and right eye views. Note that the texture within the boundary of the targets was not shifted in this method; only the circular boundary of the target was affected. In stimuli containing rings, the rings were also shifted in the two eye's views, but the texture inside the rings was not. Thus, the only features stereoscopically manipulated in the images were the boundaries of the targets or the rings surrounding the targets.

The monitor, the viewing distance, and the testing room were the same as Experiments 1 and 2. However, to create binocular disparity, two half-images were presented on two sides of the monitor and viewed through a mirror stereoscope. Unlike Experiments 1 and 2, the targets were displayed as 4 static objects instead of as a single moving target.

Methods

For each stimulus presentation, subjects were first asked whether the four discs appeared to be opaque. If the response was "yes," observers were required to report whether the targets appeared white, black, or "non-uniform"/"can't tell." Subjects were instructed to base their answer on the appearance of the disc and not the ring (so that a disc that appeared white inside a black ring would receive a response of "white"). If the response to the first question was "no," the experiment advanced to the next trial. Each of the 72 stimuli were viewed and judged once each in a practice block and then once in each of two experimental blocks.

Results

First, when the targets were placed stereoscopically behind the surrounds without rings (Figure 16; "far" condition, red lines), judgments of targets on the light

surrounds were very different from judgments on the dark surround, as expected. On the light surrounds, the targets were seen as black, and on the dark surrounds the same targets were seen as white. The frequency of these ratings was high, ranging from 86% to 100% across the 6 types of displays. In contrast, when binocular disparity placed the rings in front of the surrounds ("near" condition), targets were usually judged to be transparent; the judgments of black (on the light surrounds) or white (on the dark surrounds) fell to between 7% and 43%, far lower than in any of the conditions in which the target was behind the surround. With no disparity, the judgments fell between these values in all conditions and were always closer to the judgments with far disparity.

The pattern of results was strikingly similar when rings of the appropriate luminance for transparency were placed around the targets: black rings on dark surrounds (Figure 16, black lines versus red lines, right panel) and white rings on the light surround (dashed lines versus red lines, left panel). Again, the targets were seen as black (surrounded by white rings) when on light surrounds, and white (surrounded by black rings) when on dark surrounds, as long as the rings were stereoscopically behind the surround or at the same depth as the surround.

The gray rings severely disrupted the perception of transparency at all disparities (Figure 16). Even when the targets were behind the surround, the stereoscopic condition leading to the strongest percepts of transparency
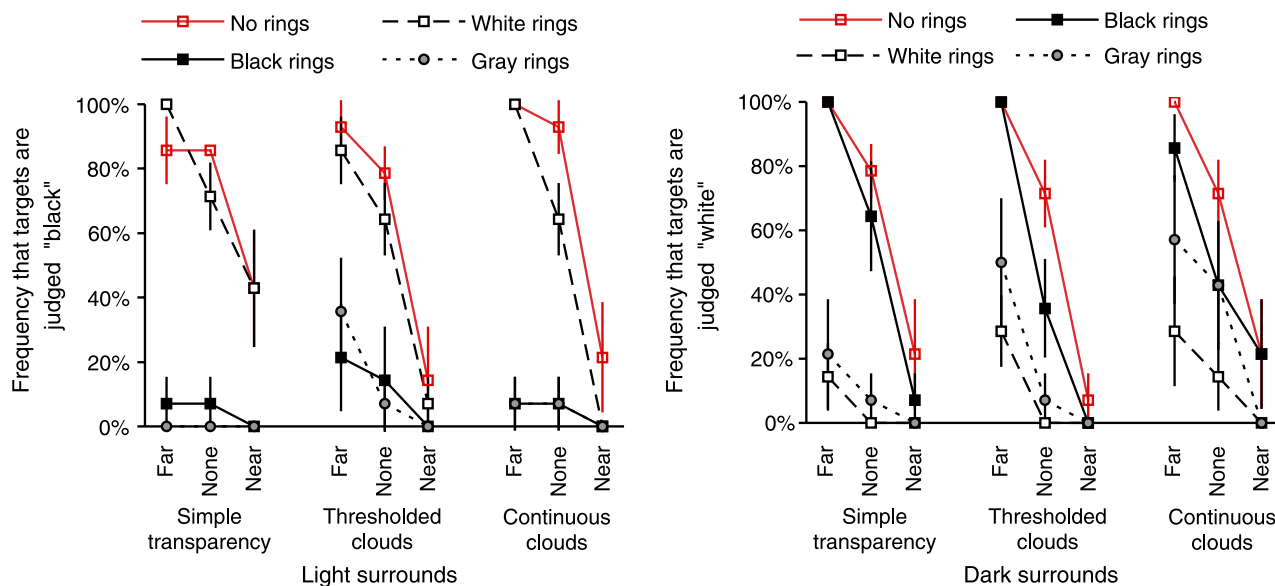

Figure 16. Experiment 3: Effect of binocular disparity and ring polarity on perception of targets. Subjects judged whether targets looked opaque and, if so, whether they appeared black, white, or "non-uniform/could not tell." The y-axis shows the frequency of judgments consistent with the expected transparency effects: judgments that the targets looked white when on dark surrounds or black when on light surrounds. When the rings were placed stereoscopically behind the surround ("far" condition), the targets with no rings (red lines) appeared as expected on nearly 100% of the trials. Placing the rings in front ("near") led to poorer scission, with the "no disparity" condition in between. Black rings on the dark surrounds and white rings on the light surrounds produced similar results to the no-rings condition. Gray rings as well as white rings on a dark surround and black rings on a light surround produced poor percepts of transparency.

with no rings, there was no coherent percept of a surface viewed through a transparent layer. Targets were most often judged either as transparent or as having a non-uniform color. In those cases when the targets were judged as white or black, the valence of the judgment was frequently opposite to judgments made with no ring.

When the rings were white or black but opposite in polarity to the surrounding region (white rings on a dark texture or vice versa), the percept of transparency was also largely eliminated. The pattern of responses in these cases was highly similar to the pattern with gray rings and quite different from the pattern with no rings, as is evident in Figure 16.

Discussion

The results of Experiment 3 reveal the importance of photometric relationships along boundaries in inducing percepts of transparency. It is important to note that only the boundaries of the target region were manipulated stereoscopically; any perception of a coherently colored disc within the interior of the ring was the result of perceptual processes (i.e., no explicit disparity was assigned to anything within the boundaries of the ring). As in Experiment 2, gray rings had a negative impact on perceiving the target as a coherent opaque disc for all ring disparities and backgrounds. In the stereo condition, the underlying cause of this effect can be directly experienced: The region inside the gray ring appears to change color, depending on the polarity relationships it has with the texture within its boundaries. One of us has recently argued that such effects arise from an inviolable constraint on how depth is assigned to local contrast relationships (see Anderson, 1999, 2003a, 2007). The core insight shaping this constraint is that near occluding and transparent surfaces can only modulate the strength, not the sign, of an underlying contrast signal. Thus, if the polarity between the ring and the interior of the target reverses, the target must appear to vary in its perceived lightness, leading to an incoherent percept of the target region. This is what observers report.

In contradistinction, black and white rings only had a substantial detrimental effect when they were placed on light and dark surrounds (respectively). This cannot be understood with a polarity constraint, as the polarity of the ring is constant on both surrounds. Rather, this incoherence results from a violation in the range of luminances present in the target region. Indeed, if the contrast within the target and the surround is the same, then a percept of transparency is readily experienced (Figure 17). Thus, these results reveal that the primary cause of the disruptive effects of the rings was due to their photometric relationships to the target and the surround not because they disrupted the continuity between the surround and the target region by introducing a gap. To further explore this question, we conducted a fourth experiment that
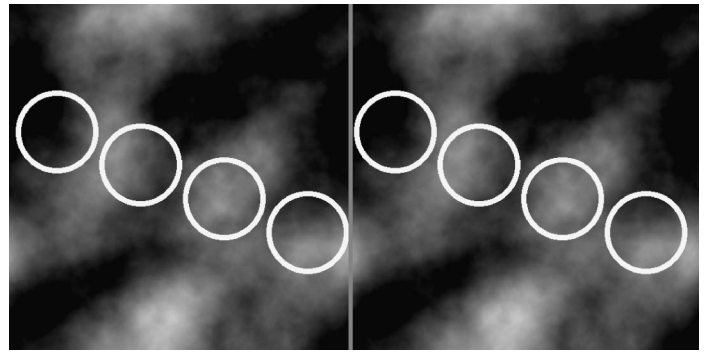

Figure 17. Continuous clouds with equal luminance range inside and outside white rings. Equivalent to the right side of Figure 12, except that the luminance range inside the rings has been lowered to be equated with the range in the surround. When viewed stereoscopically, the target regions can segment into near light clouds and far, uniform dark discs (or white hollow rings floating on a homogenously colored black background). The figure should be viewed by cross-fusing the two sides of the display. Click the image to see the moving version.

systematically varied the size of the gap between the target and the surround by varying ring thickness.

Experiment 4: The effect of ring thickness on the perception of surface lightness

Subjects

Five naive subjects were recruited at Stanford University and participated in the experiment for course credit or cash payment. All subjects gave informed written consent.

Stimuli

The stimuli in Experiment 4 were identical to those in Experiment 3, except for two differences. First, there was no binocular disparity in any of the stimuli. Second, the thickness of the rings varied from 0 (hence identical to the no-ring/no-disparity condition in Experiment 3), to one pixel, to 8 pixels (equal to the ring thickness in Experiments 2 and 3), to 20 pixels. As in Experiment 3, the stimuli were viewed through a mirror stereoscope (even though there was no disparity) and contained 4 static target discs.

Methods

The procedure was identical to that in Experiment 3.

Results

When the rings around the stimuli preserved contrast polarity and were consistent with the background luminance (white rings on a light surround and black rings on a dark surround), the expected pattern was observed: The
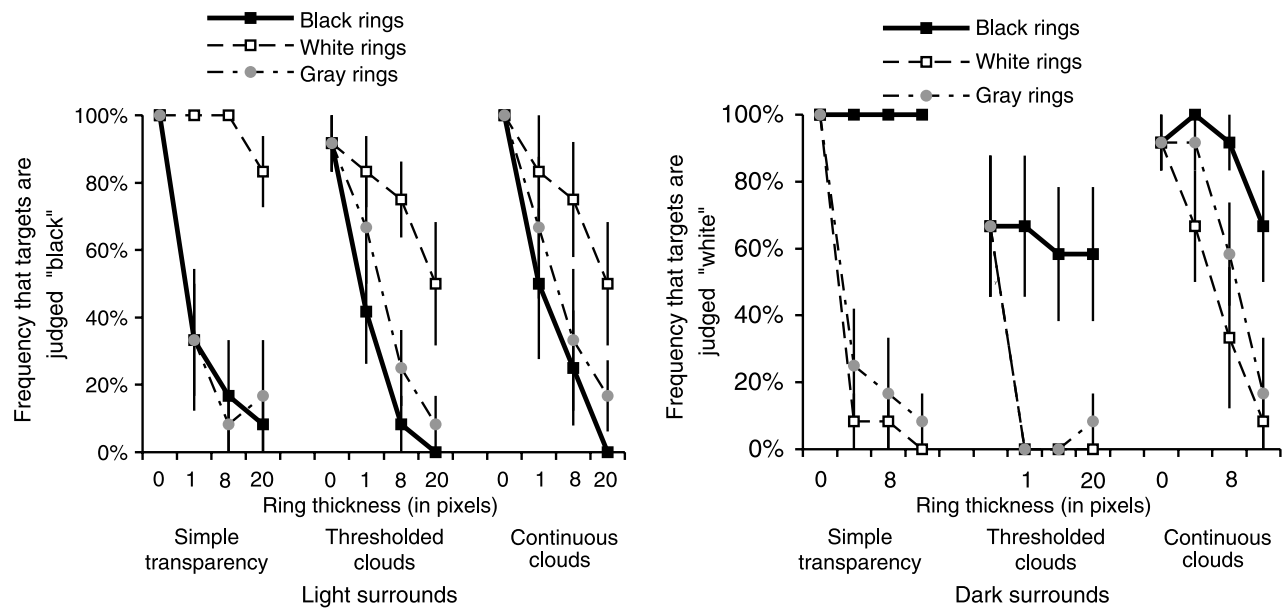
Figure 18. Experiment 4: Effect of ring thickness and ring polarity on perception of targets. For targets with black rings on dark surrounds (upper panel) or white rings on light surrounds (lower panel), the targets were usually seen as opaque discs of the expected lightness (white on dark surrounds and black on light surrounds). If rings were of opposite valence to the surround (black on light surrounds or white on dark surrounds) or if the rings were gray, the effect of scission was much poorer and the targets were less likely to be seen as opaque discs. Regardless of the luminance of the rings, scission decreased as the thickness of the rings increased.

targets appeared dark on the light surround and white on the dark surround (Figure 18). This effect weakened slightly as the rings increased in thickness. But even the thickest rings led to a high likelihood of seeing the target as solid and of the expected color, ranging from 50% for the thresholded texture on the light surround to 100% for the transparent texture on the dark surround.

In contradistinction, the gray rings strongly disrupted the perception of transparency. Rings of even one pixel in thickness (0.03°) reduced the likelihood of seeing the targets as black (on light surrounds) or as white (on dark surrounds). Averaging across the three textures, judgments of white for targets on dark surrounds fell from 89.9% to 56.2% and for judgments of black on light surrounds from 89.2% to 38.9%, when comparing gray rings to rings consistent with the surround luminance. The difference is starker when the rings were thicker (e.g., with 8 pixel rings or 0.24°): "black" responses on light surrounds fell from 83.3% to 22.1%, and "white" responses on dark surrounds fell from 83.3% to 25.1%. As before, rings of the opposite polarity (white rings on dark surrounds and black rings on light surrounds) acted similarly to gray rings, greatly reducing the likelihood that targets were seen as white on dark surrounds or black on light surrounds, even when the rings were one pixel thick.

Discussion

The results of Experiment 4 highlight the powerful role played by photometric contrasts along the target/surround

contour. Even very fine rings (~0.03°) significantly disrupted the perception of transparency, thereby blocking the large surround-dependent differences in perceived lightness of the targets. As in previous experiments, this disruptive effect was observed when the rings were gray or when the rings were inconsistent with the surround luminance, but not when the rings were consistent with the surround luminance. Increasing the thickness of the rings, thereby accentuating the geometric break between target and surround texture elements, also tended to decrease the likelihood of scission, suggesting a role for geometric constraints as well as photometric constraints. However, the fact that the ring luminance affected the likelihood of scission at each ring thickness, from 0.03° to 0.6°, highlights the particular importance of photometric relationships in inducing luminance decomposition.

General discussion

The preceding experiments and demonstrations provide further evidence that the decomposition of images into layered image representations can play a pivotal role in determining perceived surface lightness. Consistent with our previous findings (Anderson, 1997, 1999, 2003a, 2003b; Anderson & Winawer, 2005), these results show that the geometric and the photometric relationships along luminance discontinuities are critical in determining when

(and if) scission *qua* transparency is initiated, which in turn determines how luminance is (or is not) partitioned between different layers. By simply surrounding our target patches with thin outlines, we were able to disrupt the computations that initiate scission qua transparency. Although our data show that both geometric factors (the geometric continuity of the targets and the surrounds) and photometric factors (the consistent polarity relationships of the borders separating the targets from the surrounds) play a role in determining when this form of decomposition occurs, our stereoscopic experiments show that the polarity relationships between the targets and the rings are particularly critical. When the polarity relationships along the rings change sign (gray rings), no coherent pattern of scission is observed. However, when the polarity relationships between the rings and the targets have a uniform sign, and the magnitude constraints are satisfied, transparency is experienced. In these contexts, the contrast polarity between the border of the ring and the interior of the disc determine the perceived lightness polarity of the disc; the polarity of the ring and the surround determine the lightness polarity of the surround; and the TAP specifies which portions of the image are in plain view.

The phenomena reported here and previously (Anderson, 1999; Anderson & Winawer, 2005) reveal that scission can play a powerful role in determining the perceived lightness of surfaces in images that evoke percepts of transparency. One of the benefits of using the textured targets is that they evoke phenomenologically explicit percepts of layers when scission occurs. The percept is explicit because luminance edges in the targets caused by the transparent layer are seen, yet at the same time the targets are perceived as having uniform surface lightness. When uniform targets are used—as they have been in the vast majority of lightness studies (e.g., Adelson, 1999; Gilchrist et al., 1999; Todorović, 1997, 2006)—it is less perceptually obvious whether a region is decomposed into layers and hence whether such decompositions play a causal role in transformations in perceived lightness. This ambiguity has led some to question whether layered image representations play a critical role in lightness illusions, even those that evoke vivid percepts of transparency. For example, Adelson (1993) initially suggested that transparency might play a critical role in a number of lightness effects he constructed but subsequently downplayed the role of scission in favor of a frameworks-based model in which the image is divided into separate 2D regions via "adaptive windows" (Adelson, 1999). The results presented here reveal that scission *can* play a dramatic role in determining perceived lightness, but this does not imply that all lightness effects are the consequence of decomposing images into layered representations. We previously noted that the lightness transformations in textured images similar to those presented here are intimately related to occlusion computations (Anderson, 2003a, 2003b; Anderson & Winawer,

2005), whereas the vast majority of displays used to study surface lightness have no (obvious) relationship with occlusion computations. This leaves open the possibility that the processes involved in segmenting textured images into layers are different than those typically used to compute image lightness. On the other hand, if the computations of surface lightness are inherently layered, then the phenomena described herein may provide general insights into the computations that underlie the segmentation of images into causal layers of illumination, transparency, and occlusion.

One way to pursue the generality of these effects is to determine whether the phenomena reported here depend critically on the ability to define a particular image region as a portion of the scene in plain view. Note that in the paired images presented above, identical target regions are perceived as either black or white surfaces depending on their surrounds. However, for each surround, different regions within each target are perceived in plain view, which in turn determines the perceived lightness of the discs. For example, in Figure 1, the regions of the discs below the diagonal appear in plain view in the left image, but in the right image, the portions of the discs above the diagonal appear in plain view. Similar reversals occurred in the other textured images as well. It could therefore be argued that these effects are not lightness effects *per se*, as different image regions determine the perceived lightness of the layers in the two images. In contrast, most previous lightness studies have utilized targets containing a single luminance value, and if scission is involved in a target's perceived lightness, it must occur at each position in the target region (and more generally throughout the image).

One question, then, is whether it is possible to construct variants of our displays that exhibit the same form of lightness transformations as those presented in our other textured images but in which the decomposition of the target into layers occurs at all points in the target. One means of accomplishing this goal is to construct dynamic displays in which the luminance of each pixel is a constantly varying function of time. Consider the dynamic random-dot variant of this display depicted in Figure 19.
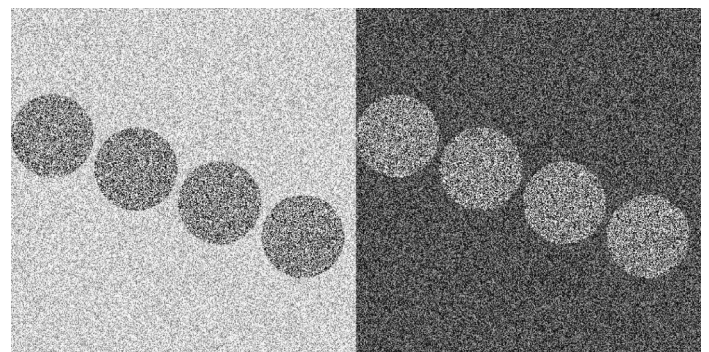

Figure 19. Scission with white noise textures. Click the image to see the moving version.
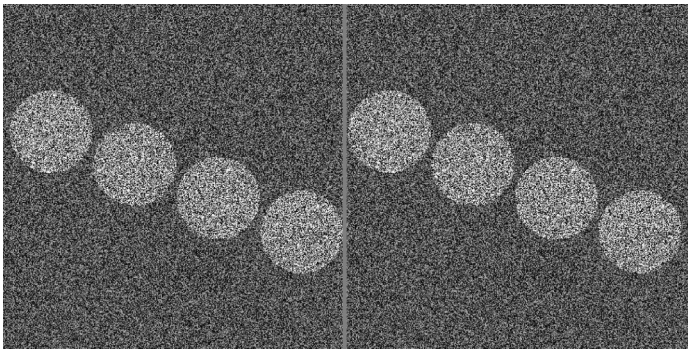
Figure 20. Scission with white noise textures in stereo. Placing the discs behind the surround can enhance the percept of the discs as opaque. The figure should be viewed by cross-fusing the two sides of the display. Click the image to see the moving version.

In this image, the central target region spans from black to white, whereas the surround values span half of the luminance range (from mid-gray to either black or white). In the moving version of the display, the luminance of each pixel is a randomly varying function of time, and hence the mean luminance of any pixel in the target region is mid-gray. Importantly, the same manipulation of the surround we used in our previous displays can induce scission of the target regions into a homogenous white or black patch overlaid with dynamic noise. Indeed, all of the critical characteristics of our previous displays are also observed with these dynamic noise displays: Placing the discs behind the surround enhances the effect of transparency (Figure 20), even when the targets are surrounded by rings of the appropriate luminance (Figure 21), but not if the luminance of the rings is inconsistent with magnitude constraints on transparency (Figure 22). The decomposition of images into layers in textured images therefore does not require the presence of stable image
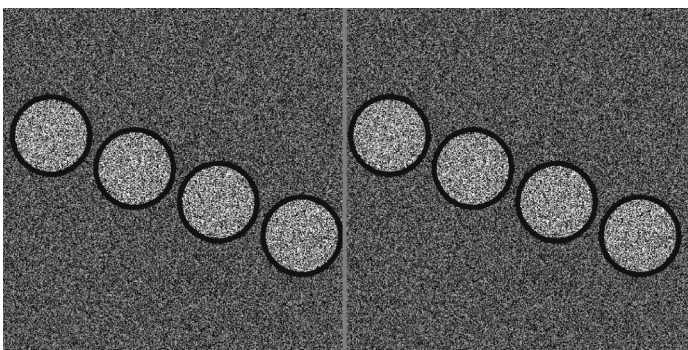

Figure 21. Scission with white noise textures in stereo surrounded by black rings. As for other texture patterns, scission can still be observed if targets are surrounded by rings of the appropriate contrast polarity. The figure should be viewed by cross-fusing the two sides of the display. Click the image to see the moving version.
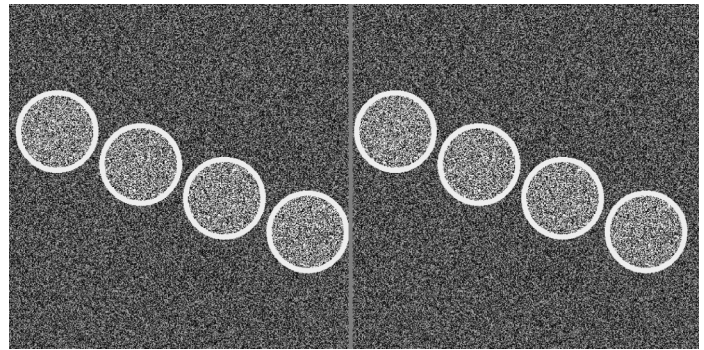

Figure 22. Lack of scission with white noise textures in stereo surrounded by white rings. As for other texture patterns, scission is not observed if targets are surrounded by rings of the inappropriate luminance values (e.g., white rings on a dark surround), even when the targets are placed behind the surround stereoscopically. The figure should be viewed by cross-fusing the two sides of the display. Click the image to see the moving version.

regions that can be assigned as regions in plain view, and hence explanations of these phenomena on the basis of traditional occlusion cues (such as border ownership, T-junctions, etc.) are untenable. (Note that other percepts are also possible, particularly in the static versions of the stereo displays. The interocular shifts used to generate a binocular disparity between the boundaries of the target regions and the noise texture also produce features that are unmatched in the left and right eye's views, which can lead to percepts of an opaque occluding surface, particularly when the near plane is fixated. These percepts have been discussed at length previously (Anderson, 1999, 2003a). Here we simply note that these alternative percepts are less stable when observers verge on the depth plane containing the target disc.)

Although the demonstrations presented in Figures 19–22 reveal that textures can be decomposed into layers at each point in a target, this does not imply that such processes underlie all context induced lightness phenomena. Nonetheless, it should be noted that in their current form, framework models do not provide a computational substrate capable of representing layered percepts and hence do not possess the representational complexity needed to assign multiple lightness values to the distinct layers that are experienced in conditions of transparency. The same problem occurs for low-level models that attempt to explain perceived lightness as a form of contrast induction (Blakeslee & McCourt, 2004; Rudd & Zemach, 2005, 2007). There is nothing in such models that could generate a layered representation of either a homogeneous or textured image region. The output of such models is a 2D map of transformed luminance values, and hence such models can only transform the relative intensities of one 2D image into another 2D image. However, such models might explain the residual lightness/brightness effects observed in all of our images.

For example, in Figure 1, the light portions of the disc in the left image appear darker than the light portions of the disc on the right, whereas the dark portions of the disc appear darker in the left image and lighter in the right. Such effects are well known, but their cause remains a source of debate, as there are a number of different models that have attempted to explain such phenomena. We will not attempt an exhaustive survey or discussion of these models here; we simply consider whether such effects implicate a separate process than that underlying the effects that have been the focus of our study.

One reason to suspect that these residual effects represent a different process is that they are much weaker than the lightness effects we have measured herein. When the textured targets are replaced with gray targets and the surrounds are unchanged, the perceived difference in the targets on the two surrounds is only 11% as large as the strongest lightness effects we observe using our continuous texture targets (Anderson & Winawer, 2005). On the other hand, such simultaneous contrast effects can also be explained with scission models: If the lightness of the surround is "discounted" from the target (i.e., attributed to a separate layer), then targets on light surrounds should appear darker, and dark surrounds lighter. In a model of this kind, the strength of the induction would be determined by the *extent* to which scission occurred, i.e., the extent to which the surround color is discounted. Indeed, some recent work provides striking evidence for the role of scission in chromatic variants of simultaneous contrast displays (Ekroll, Faul, & Niederée, 2004). Additional support for the importance of layered image decompositions in traditional lightness/brightness displays has also been reported in other displays, most notably White's effect (Anderson, 2003b). Textural variants of White's effect exhibit the same kind of decomposition as the displays described herein, and they exhibit the same dependencies on the spatial parameters of White's display as the more traditional untextured targets. This suggests that a common scission mechanism may be responsible for both forms of the illusion. Indeed, no other model of the textured variant of White's effect has yet to be offered, and existing lightness and brightness models do not appear to possess the requisite representational complexity to handle these forms of the illusion (for the same reasons that they fail to account for the effects reported herein). Thus, it is currently unclear whether the residual lightness/brightness effects observed in the displays herein represent a separate process, or whether they arise from a common scission mechanism that is only partially disrupted by our ring manipulations.

It should be noted that a number of authors have challenged whether scission and/or transparency plays a role in a number of lightness phenomena, especially a number developed by Adelson (1993, 1999). For example, Logvinenko (1999) showed that striking lightness transformations can be observed in Adelson's wall of blocks demonstration that contain luminance gradients that do not support a perception of transparency. More recently, Todorović (2006) presented a host of demonstrations that reveal the efficacy of luminance gradients in inducing changes in perceived lightness of otherwise identical stimuli. These results suggest that if scission plays a role in these phenomena, it is unlikely that such processes employ highly sophisticated processes to infer the illuminant. Indeed, it has been shown that observers are extremely insensitive to global inconsistencies in the distribution of shadows in a scene (Ostrovsky, Cavanagh, & Sinha, 2005). Thus, the persistence of lightness illusions in conditions that do not contain a globally consistent illuminant (or a highly probable illuminant) does not provide conclusive evidence that layered image decompositions are not responsible for these effects. Rather, they simply provide information about the kinds of information that may be used to induce scission if scission is occurring in these images (for a related discussion, see Todorović, 2006, pp. 244–245).

Bressan (2001, 2006) has also questioned the role of scission as an explanatory construct in lightness perception. She created a variant of the snake illusion (which she dubbed the "shredded snake") in which the X-junctions were replaced with polarity preserving T-junctions that are also consistent with transparency (cf. Anderson, 1997). She found that the lightness effect was weakened by this manipulation and actually attributed this weakening to a diminished contribution of scission in this version of the display. However, she constructed an "anti-snake" variant of Adelson's illusion in which the target on the lighter surround appears to be overlaid by a dark transparent layer, and the target on the darker surround appears to be overlaid by a light transparent layer. Bressan argued that a scission thesis should predict that the target on the light background should look lighter than the one on the dark surround (because it appears to be overlaid by a dark filter), but she found that the two targets were now perceived as identical. Despite her argument to the contrary, this result is also compatible with a scission model. If the local contrast effects (whatever mechanism is their cause) is taken as baseline (as they were for all other variants of the snake illusion), then the change in direction relative to the local contrast effect is exactly what would be predicted on the basis of scission. Since there is currently no way to determine the precise weights of the local (simultaneous contrast) and remote effects (the X-junctions favoring transparency) on the basis of either theory or data, it is impossible to predict whether the effect should reverse, be nulled, or simply diminish; all that can be predicted is the direction of the change, which is perfectly consistent with a scission model.

The demonstrations presented herein reveal that the perceived lightness of a region in conditions of transparency is constrained by how the visual system determines what portions of an image are considered to appear in plain view. Although such considerations might appear to be restricted to the domain of transparency and

Figure 23. Similarities and differences between the effects of perceived shadows and perceived transparency on lightness. The upper displays appear to be divided by a shadow boundary, whereas the lower displays appear to be divided by a transparent layer. In the case of shadows, the luminance of the target seems to be determined by the region in higher luminance (upper right), but in the case of transparency, the luminance appears to be determined by the region in higher contrast with the surround (lower left region in the lower left display and upper right region in the lower right display).

occlusion computations, it is also possible to demonstrate that related computational issues arise when a surface falls in two different illuminants. Consider several variants of Figure 1 depicted in Figure 23. The discs in all 4 images are composed of the same luminance range. In the upper displays, the diagonal contour that runs throughout the image has been blurred, yielding an impression of an illumination boundary separating the two halves of the image. Note that the discs in the upper left image are perceived to be light gray, whereas in the lower left image the discs are perceived to be dark gray. The perceived lightness of the discs in the transparency displays (lower panels) is determined by the highest contrast segments along the circular contours separating the targets and their surrounds (lower left portion of the lower left image, and upper right portion of the lower right image) but in the upper displays, the lightness of the discs is determined by the surfaces that appear in the higher illumination (upper right portion of both of the upper displays). This suggests that there is an asymmetry in the role of different illumination levels in determining perceived lightness (see Gilchrist & Zdravkovic, 2002; Singh & Anderson,

2006). This asymmetry is also reflected in transparency judgments. Whereas the perceived transition points between transparency and non-transparency of darkening filters are well predicted by their physical (or perceived) contrast relationships relative to their surrounds, the transition from transparency to non-transparency of lightening filters exhibits more complex and varied behavior across observers (Singh & Anderson, 2002, 2006). The important point for the present discussion is that the role of border ownership in determining perceived lightness is not restricted to cases of occlusion and transparency; it also arises in the interpretation of illumination boundaries and the lightness of the surfaces such boundaries cross.

The demonstrations and the experiments presented in the current work have focused on the role of geometric and photometric constraints in determining when scission occurs and its consequence on perceived lightness. It should be stressed, however, that we do not view the theoretical principles such as the TAP as inviolable constraints on the perception of transparency or the geometric and the photometric properties explored herein as the sole determinants of when transparency occurs. Rather, we view these properties as probabilistic constraints that work in concert with other sources of information in determining when scission qua transparency is and is not experienced. Our informal observations revealed that the perception of scission is enhanced when displays contain multiple, nearly contiguous targets. We presume that the repeated instances of the target (either in space or in space–time, as in our moving sequences), each consistent with the interpretation of transparency, strengthen the evidence for the existence of multiple layers, thereby favoring such a percept. More generally, any image features that signal the presence of transparency or occlusion may have similar effects on perceived surface lightness.

In sum, we have shown a variety of phenomena that reveal the close link between occlusion, transparency, and perceived lightness. These phenomena, as well as some classic effects such as the Mach card and Hering's outline shadow demonstration, provide strong evidence for the role of layered image computations in the perception of surface lightness. Further research is needed to determine the relative contribution of the computations that lead to layered image representations in determining perceived lightness.

Acknowledgments

Commercial relationships: none.
Corresponding author: Barton L. Anderson.
Email: barta@psych.usyd.edu.au.
Address: School of Psychology, Brennan MacCallum Building (A18), The University of Sydney, NSW 2006, Australia.

Footnotes

[1]For example, Kanizsa (1979) stated: "'Double presence' can also occur in the absence of transparency—for example, in the figure-ground phenomenon. The ground 'continues under' or 'passes behind' the figure. It is 'present,' even though invisible, behind the figure. The same region of stimulation is, thus, twice present in experience; this is not simply due to the characteristics of the regions itself but, rather, to factors that lie outside it (or, more strictly speaking, factors that belong to the *relation* between the shared region and those that border on it)" (p. 163).

[2]Although growing data demonstrate the importance of image contrast in determining when transparency is perceived, there remains a significant problem in constructing a definition of stimulus contrast that captures perceived contrast. It has been shown that ratios of Michelson contrast provide excellent fits of opacity matches in displays for which Michelson contrast maps well onto perceived contrast (Singh & Anderson, 2002). However, such measures do not generalize to transparency judgments in displays for which Michelson contrast does not provide a good measure of perceived contrast (Robilotto, Khang, & Zaidi, 2002). Indeed, in these studies, no existing definition of stimulus contrast captured perceived contrast, yet nonetheless, transparency judgments were well accounted for by observers perceived contrast of the targets in these images (see also Anderson et al., 2006).

References

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

Adelson, E. H. (1999). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., pp. 339–351). Cambridge, MA: MIT Press.

Adelson, E. H., & Anandan, P. (1990). Ordinal characteristics of transparency. In *Proceedings of the AAAI-90 Workshop on Qualitative Vision* (pp. 77–81). Boston: AAAI Press.

Anderson, B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: The role of contour junctions. *Perception, 26,* 419–453. [PubMed]

Anderson, B. L. (1999). Stereoscopic surface perception. *Neuron, 24,* 919–928. [PubMed] [Article]

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

Anderson, B. L. (2003b). Perceptual organization in White's illusion. *Perception, 32,* 269–284. [PubMed]

Anderson, B. L. (2007). Filling-in models of completion: Rejoinder to Kellman, Garrigan, Shipley, and Keane (2007) and Albert (2007). *Psychological Review, 114,* 509–527. [PubMed]

Anderson, B. L., & Julesz, B. (1995). A theoretical-analysis of illusory contour formation in stereopsis. *Psychological Review, 102,* 705–743.

Anderson, B. L., Singh, M., & Meng, J. (2006). The perceived transmittance of inhomogeneous surfaces and media. *Vision Research, 46,* 1982–1995. [PubMed]

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

Anderson, B. L., & Winawer, J. (2006). Scission and the perception of lightness [Abstract]. *Journal of Vision, 6*(6):392, 392a, http://journalofvision.org/6/6/392/, doi:10.1167/6.6.392.

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). New York: Academic Press.

Beck, J. (1965). Apparent spatial position and the perception of lightness. *Journal of Experimental Psychology, 69,* 170–179. [PubMed]

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

Blakeslee, B., & McCourt, M. E. (2004). A unified theory of brightness contrast and assimilation incorporating oriented multiscale spatial filtering and contrast normalization. *Vision Research, 44,* 2483–2503. [PubMed]

Bressan, P. (2001). Explaining lightness illusions. *Perception, 30,* 1031–1046. [PubMed]

Bressan, P. (2006). The place of white in a world of grays: A double-anchoring theory of lightness perception. *Psychological Review, 113,* 526–553. [PubMed]

Dakin, S. C., & Bex, P. J. (2003). Natural image statistics mediate brightness 'filling-in.' *Proceedings of the Royal Society B: Biological Sciences, 270,* 2341–2348. [PubMed] [Article]

Ekroll, V., Faul, F., & Niederée, R. (2004). The peculiar nature of simultaneous colour contrast in uniform surrounds. *Vision Research, 44,* 1765–1786. [PubMed]

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. [PubMed]

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

Gilchrist, A. L., & Zdravkovic, S. (2002). Highest luminance defines illumination level as well as lightness [Abstract]. *Journal of Vision, 2*(7):553, 553a, http://journalofvision.org/2/7/553/, doi:10.1167/2.7.553.

Kanizsa, G. (1979). *Organization in vision*. New York: Praeger.

Kennedy, J. M., & Bai, J. (2000). Cavanagh and Leclerc shape-from-shadow pictures: Do line versions fail because of the polarity of the regions or the contour? *Perception, 29,* 399–407. [PubMed]

Kingdom, F., & Moulden, B. (1992). A multi-channel approach to brightness coding. *Vision Research, 32,* 1565–1582. [PubMed]

Kingdom, F. A., Blakeslee, B., & McCourt, M. E. (1997). Brightness with and without perceived transparency: When does it make a difference? *Perception, 26,* 493–506. [PubMed]

Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt, Brace, & World.

Logvinenko, A. D. (1999). Lightness induction revisited. *Perception, 28,* 803–816. [PubMed]

Mach, E. (1959). *The analysis of sensations*. New York: Dover. (Original work published 1886)

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. [PubMed]

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

Ostrovsky, Y., Cavanagh, P., & Sinha, P. (2005). Perceiving illumination inconsistencies in scenes. *Perception, 34,* 1301–1314. [PubMed]

Robilotto, R., Khang, B. G., & Zaidi, Q. (2002). Sensory and physical determinants of perceived achromatic transparency. *Journal of Vision, 2*(5):3, 388–403, http://journalofvision.org/2/5/3/, doi:10.1167/2.5.3. [PubMed] [Article]

Rudd, M. E., & Zemach, I. K. (2005). The highest luminance rule in achromatic color perception: Some counterexamples and an alternative theory. *Journal of Vision, 5*(11):5, 983–1003, http://journalofvision.org/5/11/5/, doi:10.1167/5.11.5. [PubMed] [Article]

Rudd, M. E., & Zemach, I. K. (2007). Contrast polarity and edge integration in achromatic color perception. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 24,* 2134–2156. [PubMed]

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

Singh, M., & Anderson, B. L. (2006). Photometric determinants of perceived transparency. *Vision Research, 46,* 879–894. [PubMed]

Taya, R., Ehrenstein, W. H., & Cavonius, C. R. (1995). Varying the strength of the Munker-White effect by stereoscopic viewing. *Perception, 24,* 685–694. [PubMed]

Todorović, D. (1997). Lightness and junctions. *Perception, 26,* 379–394. [PubMed]

Todorović, D. (2006). Lightness, illumination, and gradients. *Spatial Vision, 19,* 219–261. [PubMed]

Tse, P. U. (2005). Voluntary attention modulates the brightness of overlapping transparent surfaces. *Vision Research, 45,* 1095–1098. [PubMed]

von Helmholtz, H. (1866). *Helmholtz's treatise on physiological optics*. New York: Optical Society of America.

Zdravković, S., Economou, E., & Gilchrist, A. (2006). Lightness of an object under two illumination levels. *Perception, 35,* 1185–1201. [PubMed]