Vision Research 48 (2008) 2090–2105

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Perceiving light versus material

Frederick A.A. Kingdom

McGill Vision Research, 687 Pine Av. W. Rm. H4-14, Montreal, Que., Canada H3A 1A1

ARTICLE INFO

Article history:
Received 10 January 2008
Received in revised form 25 March 2008

Keywords:
Intrinsic images
Layer perception
Reflectance
Albedo
Pigment
Shadows
Shading
Transparency
Specular reflections
Junctions
Neon-colour spreading

ABSTRACT

Humans rarely confuse variations in light intensity, such as shadows, shading, light sources and specular reflections, from variations in material properties, such as albedo or pigment. This review explores the cues, or regularities in the visual world that evidence suggests vision exploits to discriminate light from material. These cues include luminance relations, figural relations, 3D-shape, depth, colour, texture, and motion. On the basis of an examination of the cues together with the behavioural evidence that they are used by vision, I propose a set of heuristics that may guide vision in the task of distinguishing between light and material. I argue that while there is evidence for the use of these heuristics, little is known about their relative importance and the manner in which they are combined in naturalistic situations where there are multiple cues as to what is light and what is material. Finally, I discuss two theoretical frameworks, the generic view principle and Bayesian estimation, that are beginning to help us understand the visual processes involved in distinguishing between light and material.

1. Introduction

1.1. Light versus material

The visual world not only consists of materials but patterns of light, which include shadows, shading, specular reflections, spotlights and light sources. Distinguishing light from material is not only important for detecting what is material: the pattern of light itself provides information about the type of material, its shape and the object to which it belongs. How do we distinguish light from material? The problem may be considered one of decomposing an image into 'layers', or what Barrow and Tenenbaum (1978) have termed 'intrinsic images', in which each layer, or intrinsic image, has a different physical origin. Layer decomposition is a classic problem of inverse optics: given a piece of image data such as a change in luminance, which has many possible causes, what is the cause? This review aims to summarize what we know about how the visual system decomposes the visual input into light and material. The focus will be on the regularities, or cues in the natural visual environment that evidence suggests vision exploits for this purpose, cues that include luminance relationships, surface geometry, figural arrangements, 3D-shape, depth, colour, texture and motion. The review will not deal with how we distinguish different types of material (coloured versus achromatic, rough versus smooth etc.), except with regard to distinguishing reflectance surfaces from transparency, for reasons given below. Nor will the review consider how the pattern of light is used to determine object properties (e.g. lightness, colour, shape, texture), except in so far as these things tell us about the cues used by vision for discriminating light from material in the first place. For example, while studies of the lightness, or perceived shade of grey of stimulus patterns have revealed many cues used by vision to distinguish light from material, the significance of these studies for lightness perception is not discussed except in passing. Similarly, while studies of shape-from-shading have revealed how luminance, color and texture are exploited by vision for discriminating light from material, the review will not deal with how vision actually derives shape from shading. And while the review describes some of the cues that are likely used for identifying specular reflections, the strengths and limitations of specular reflections for identifying material properties will not be discussed.

1.2. The issue

Fig. 1 exemplifies the problem and suggests likely cues to the solution. The luminance image $L(x,y)$ of the cylinder can be decomposed into two layers, one reflectance $R(x,y)$, the other non-uniform illumination, $I(x,y)$. This follows from the relation $L(x,y) = I(x,y)$

E-mail address: fred.kingdom@mcgill.ca

0042-6989/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2008.03.020

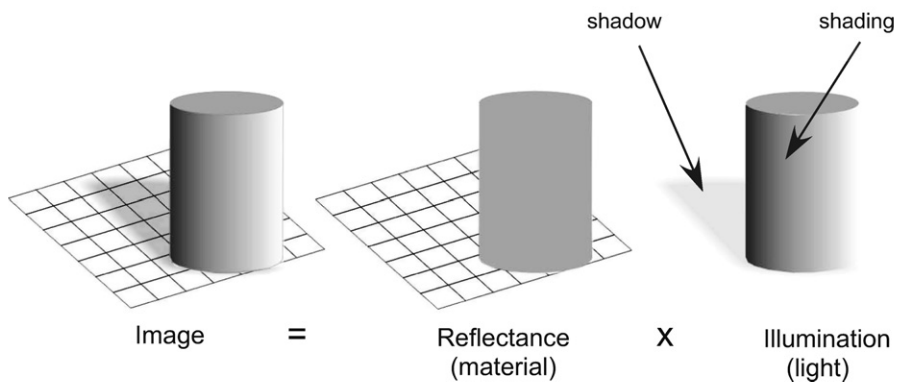

Fig. 1. An image decomposed into reflectance and illumination layers.

$R(x,y)$. It is impossible to determine from the luminance of any one pixel whether it is part of a light surface that is shaded or a dark surface brightly illuminated. Pixel luminance is inherently ambiguous because an infinite number of combinations of the two unknowns, $R(x,y)$ and $I(x,y)$ can produce $L(x,y)$. Only by examining the *relations* among pixels combined with *prior knowledge* about the circumstances in which reflectance changes and non-uniform illumination occur can one estimate the real-world properties $R(x,y)$ and $L(x,y)$. In Fig. 1 the cylinder appears to be shaded, and as a result looks rounded not flat, a consequence of 'shape-from-shading'. The main cue that the luminance changes on the cylinder are shading is that they are gradual, though shading is also consistent with the cylindrical shape suggested by the curved contour at the top of the cylinder (however this is insufficient on its own to elicit an impression of a cylinder, as can be seen from the reflectance image). The task of distinguishing reflectance from illumination is therefore accomplished by comparing points across the image and using prior knowledge of the spatial properties of both shading and shape.

1.3. Types of layer

Fig. 1 illustrates what is meant by reflectance, shadow and shading. When appropriate, intensive reflectance or the shade-of-grey of a surface will be termed *albedo*, and the spectral reflectance of a surface *pigment*. The perceptual correlates of intensive and spectral reflectance are lightness and colour. The term brightness will be used to denote perceived luminance. A *shadow* results from occlusion of a light source, while *shading* (sometimes termed an attached shadow) results from a change in the angle of the surface normal with respect to the direction of illumination. Other types of non-uniform illumination discussed here include: *spotlights* (Fig. 2a), which are local brightly illuminated regions; *self-luminosity*, i.e. light sources (Fig. 4); *specular reflections* (Fig. 5), which are sometimes called highlights. The last of these, specular reflections, are seen on shiny or metallic surfaces in non-diffuse illumination, and occur because shiny surfaces do not reflect light equally in all directions (termed 'non-Lambertian').

Another type of layer considered here is *transparency* (figures throughout text). Although transparencies are materials rather than light, they not only behave similarly to light, but are the basis of some of the most compelling demonstrations of the cues involved in light-versus-material perception (e.g. Adelson, 1993; Adelson, 2000; Logvinenko, Adelson, Ross, & Somers 2005). The physical properties of transparency include: transmittance, defined as the proportion of light that passes through the medium; reflectance, defined as the proportion of light reflected by the medium; scatter, defined as the amount of dispersion of light within the medium. A standard pair of dark glasses transmits a proportion of the incident light but reflects almost no light, and hence acts similarly to a shadow by dividing the background luminance by a constant factor, thus reducing the light level while leaving the contrast of objects viewed through them unchanged (Fig. 2b). Transparencies that reflect as well as transmit light are referred to as having a reflective or additive component, and reduce the contrast of objects seen through them. They include media such as frosted glass, grease-proof paper, milky water and fog, and are characterized by their cloudy appearance (Fig. 2c). Finally, a close relative of transparency is the delicate, illusory phenomenon of *neon-colour spreading* (Anderson, 1997; Bressan, Mingolla, Spillmann, & Watanabe, 1997; Hoffman, 1998; van Tuijl, 1975), whose appearance resembles both self-luminosity and reflective transparency (Figs. 6b and 10e).

In summary this review deals with the cues that underpin the decomposition of visual stimulation into different perceptual categories or 'layers', corresponding to reflectance (albedo and pigment), shadows, shading, spotlights, light sources, specular reflections, transparency and neon-colour spreading.

1.4. Brief history

The use of prior knowledge to help distinguish light from material is an example of what Helmholtz famously termed "unconscious inference" (Helmholtz, 1866/1962). One of Helmholtz's most celebrated examples of unconscious inference was simultaneous colour contrast, the phenomenon in which, to take just one example, a grey patch surrounded by green appears tinged with pink- the complementary colour of green. Helmholtz offered more than one explanation for simultaneous colour contrast (Turner, 1994), but they each had in common the idea that the visual system attributed to the illuminant the colour of the patch surround, and it was this that led to the misperception. Although Helmholtz's explanation(s) for simultaneous colour contrast continue to be disputed (e.g. see critique by James (1890/1981), discussed by Kingdom, 1997), they highlighted the idea that the visual system is able to decompose the image into illumination and pigment. Helmholtz's contemporary and adversary Hering also drew attention to the importance of layer perception in his famous shadows experiment (Hering, 1874/1964). Hering showed that if one drew a thick black line to cover the penumbra of a cast shadow, the appearance of the shadow changed from one of dark illumination to different shade of grey [see recent applications of this manipulation to face perception in Cavanagh and Leclerc (1989) and Cavanagh and Kennedy (2000)].

In more recent times Gilchrist & colleagues' theory of "edge classification", formulated in the context of lightness perception (Gilchrist, 1979; Gilchrist, Delman, & Jacobsen, 1983; reviewed by Gilchrist, 2006), and Metelli's (1974) theory of "scission", formulated in the context of transparency perception, have been seminal in revealing the importance of layer perception in vision. Using the display illustrated in Fig. 2a, Gilchrist (1979) and Gilchrist et al. (1983) showed that the lightness of a patch depended on whether

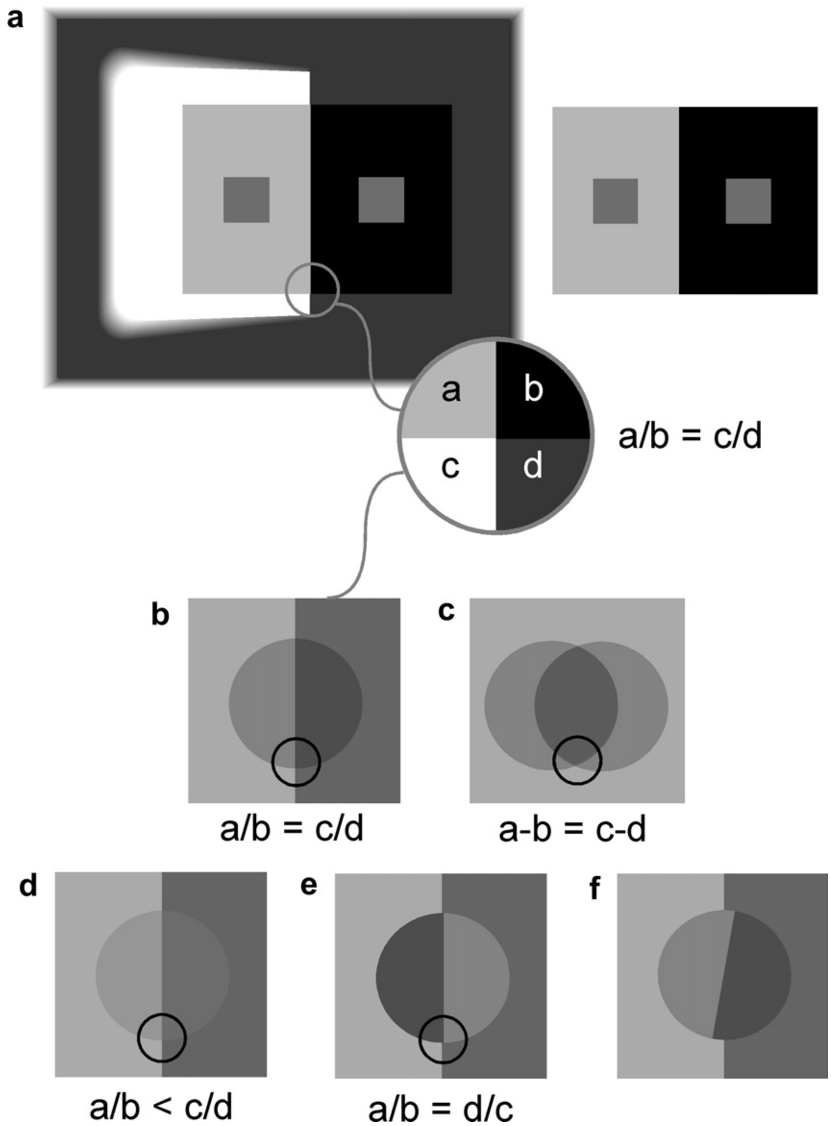

Fig. 2. (a) Gilchrist's (1979) experiment comparing the lightness of two test patches whose surrounds differed in either in illumination (left) or albedo (right). The inset shows the luminance arrangement at the at the illumination edge (an X-junction), which satisfies the ratio invariance rule. (b) A non-reflective transparency, or shadow, also satisfies ratio-invariance. (c) Two intersecting shadows show difference-invariance; (d) transparency with a reflective component, resulting in reduced contrast inside the transparent region, but of the same sign. (e) The transparent region has been rotated by 180 deg, which preserves the X-junctions but reverses the sign of the contrast of the transparent region, eliminating the impression of transparency. In (f) the transparency in b has been rotated by just 10 deg, and the lack of a well-defined X-junction eliminates the impression of transparency. Top figure without inset taken from Gilchrist (2006) and reproduced with permission.

its surround appeared to differ from the background in illumination or albedo, in spite of the fact that the luminance contrast between patch and surround was the same in both configurations[1] (the actual displays used real papers and lights, and the effects reported are larger than can be seen in the figure). Later, Gilchrist (1988) suggested that two cues were critical to classifying the surround as differing in illumination rather than albedo—X-junctions and luminance-ratio-invariance—cues that also figure in Metelli's work and which will be discussed later. Gilchrist's and Metelli's studies laid the foundation for many subsequent demonstrations showing how perceived transparency (e.g. Adelson, 1993; Anderson & Winawer, 2005; Logvinenko et al., 2005) and perceived shading, or shadows (Adelson, 2000; Adelson & Pentland, 1996; Logvinenko, 1999; see also earlier studies by Arend & Spehar, 1993a, 1993b) can dramatically influence the lightness and/or brightness of regions that appear part of the material layer. Moreover Metelli's studies established transparency perception as a distinct sub-discipline within vision science (e.g. Albert, 2006; Anderson, Singh, & Meng, 2006; Beck & Ivry, 1988; Gerbino, 1994, chap. 5; Gerbino, Stultiens, Troost, & de Weert, 1990; Gurnsey, Kingdom, & Schofield, submitted for publication; Kanizsa, 1979; Kasrai & Kingdom, 2001, 2002; Masin, 1997, 2006; Robilotto, Khang, & Zaidi, 2002; Robilotto & Zaidi, 2004; Singh & Anderson, 2002a, 2002b).

This brief summary of some of the pertinent historical moments sets the scene for an examination of the physical cues evidence suggests vision exploits to discriminate light from material. The examination will result in a list of heuristics, or 'rough-and-ready' rules used by vision to designate regions as either reflectance, shadow, shading, spotlight, light-source, specular reflection or transparency. It must be born in mind at the outset however that none of the heuristics should be regarded as a hard-and-fast rule. Each is a clue to whether a luminance change is more or less likely

[1] Although the present author has criticized the interpretation of this experiment with regard to *lightness* perception (Kingdom, 2003a), its historical significance for our understanding of layer perception is unquestionable.

to be light than material given its context. The heuristics are probabilistic (although no actual probabilities are provided) meaning that any one heuristic can in principle be overruled by any other if there is sufficient evidence to favour the alternative interpretation. In natural scenes multiple cues for distinguishing light from material co-exist and their relative weightings and the nature of their interactions are only beginning to be understood. Therefore at the very least one should append the following suffix to every heuristic listed below: "……provided no cues to the contrary are present".

2. Cues and heuristics

2.1. Luminance relations

2.1.1. Invariances at intersecting edges

When an illumination edge intersects a reflectance edge to form an X-junction, the ratio of luminances either side of the junction is the same (Gilchrist, 1988, 2006, p184). This is because the effect of illumination on a surface is multiplicative, changing its mean luminance but not its luminance ratios since a/b is unchanged when both a and b are multiplied by the same factor. Since a luminance ratio is a measure of contrast, the ratio invariance cue can also be termed the contrast invariance cue. Ratio invariance is illustrated for the spotlight in Gilchrist's display (Fig. 2a), and for the shadow or non-reflective transparency in Fig. 2b. Although ratio invariance is evidence for the presence of an illumination edge, it does not tell us which edge is the illumination edge, since if $a/b = c/d$, also $a/c = b/d$. The ambiguity is observable in the inset in Fig. 2a which shows the X-junction removed from its context. Ratio invariance is also consistent with two, intersecting, non-reflective transparencies.

Other types of edge violate ratio-invariance but preserve the sign, or polarity of contrast across the X-junction. Difference-invariance, i.e. $a - b = c - d$, occurs when two illumination edges intersect (Gilchrist, 2006, p. 184–185), as in the overlapping shadows example in Fig. 2c. Difference-invariance results from the fact that the intensities of superimposed illuminations add not multiply. Intersecting illumination edges are very common in nature, for example in the dappled illumination in foliage. Difference-invariance represents a special case of the more general situation in which the magnitude but not sign of contrast changes across the junction. The general situation is characterized by the fact that if $a - b = c - d$ and $a + b > c + d$, then $a/b < c/d$. In this situation, a reduction in contrast but not sign at an X-junction is a signature for a transparency with a reflective, or additive component, as In Fig. 2d.

On the behavioural side, the impressions of transparency in Fig. 2b–e are certainly compelling, and the importance of preserving the sign of contrast across the X-junction can be seen in Fig. 2e, where the polarity is reversed and the impression of transparency eliminated. Moreover, we are sensitive to the different physical properties of transparency: the cloudy appearance of the central region in Fig. 2d is a signature for a transparency with a reflective component. Singh and Anderson (2002a) have argued that the combination of contrast-sign-invariance and contrast reduction in Fig. 2d is a general trigger for unambiguous perceptual scission, unambiguous in the sense that the lower contrast region is always perceived as a foreground transparency and the higher contrast region the background. Compare for example Fig. 2b and d. Prolonged fixation of Fig. 2b may produce perceptual reversals in the ordering of layers, whereas in Fig. 2d one tends not to experience layer reversals.

There are many unanswered questions concerning the luminance relations at X-junctions. For example, can we distinguish overlapping non-reflective transparencies from overlapping shadows using X-junction information alone, given that one layer combination exhibits ratio-, the other difference-invariance? Since small violations from ratio invariance are probably imperceptible (Kasrai & Kingdom, 2001; Kingdom, Beauce, & Hunter, 2004) it would seem unlikely. Another question concerns how the two principle physical properties of transparency, transmittance and reflectance, are encoded by vision. This complex and controversial issue is beyond the scope of this review, but the interested reader may consult recent articles by Masin, Tommasi, and Da Pos (2007), Albert (in press), Anderson and Singh (in press) and Gurnsey et al. (submitted for publication).

Heuristic 1a. If two intersecting edges show luminance ratio invariance, one edge is likely to be an illumination edge or non-reflective transparency, the other a reflectance edge.

Heuristic 1b. If two intersecting edges show luminance difference invariance, both edges are likely to be illumination edges.

Heuristic 1c. If the contrast but not sign of an edge changes as it intersects another edge, but difference-invariance is not shown, the lower contrast side of the edge is likely to be a transparency with a reflective component.

Heuristic 1d. If the contrast sign changes across two intersecting edges, the edges are both likely to be reflectance edges.

2.1.2. Invariances in textures

Luminance-ratio (or -contrast) invariance applies also to textured surfaces. Just as a shadow reduces the mean but not contrast of a luminance edge, so too will it reduce the mean but not contrast of luminance variations in a texture. Schofield, Heese, Rock, and Georgeson (2006) investigated whether humans were sensitive to luminance-contrast invariance in textures in the context of shape-from shading. Fig. 3 demonstrates the main finding. The figure comprises two sinusoidal luminance gratings in a plaid configuration superimposed on a noise texture. In the right oblique component the local luminance amplitude of the texture is varied in-phase with the luminance grating, such that luminance contrast is constant. This component is consistent with shading of a uniform textured surface and is indeed seen as undulating in depth. In the left oblique component the luminance amplitude of the noise texture is varied in anti-phase with the luminance grating, and hence luminance contrast varies. In this case the modulation is not consistent with shading, and indeed observers perceive the anti-phase combination as relatively flat. Together with the results of other experiments, Schofield et al. (2006) conclude that local luminance amplitude can be a key to disambiguate the origin of luminance variations in an image (shading versus albedo), but that it is not effective in all circumstances.

Heuristic 2. If the luminance contrast of a texture is the same on either side of a luminance edge, the edge is likely to be an illumination edge or non-reflective transparency. On the other hand if texture luminance contrast is different on either side of the luminance edge, the lower contrast region is likely to be a reflective transparency or a texture with different luminance contrast.

2.1.3. Penumbra and gradients

Hering's demonstration described above, in which a shadow changed in appearance to albedo when its penumbra was covered with a black line may have less to do with removing the penumbra as with reversing the polarity at the shadow's edge (Cavanagh & Leclerc, 1989; Kennedy & Bai, 2000; Cavanagh & Kennedy, 2000). Indeed many shadows have sharp edges, and still appear as shadows. Irrespective of the precise reason for the loss in shadow appearance in Hering's demonstration however, numerous studies since Hering have suggested that gradual luminance changes are signatures for non-uniform illumination (Agostini & Galmonte, 2002; Hoffman, 1998; Kanizsa, 1979; Land & McCann, 1971; MacLeod, 1947; Schirillo & Shevell, 1997; Zavagno, 1999; Zavagno & Caputo, 2001; Correani, Scott-Samuel, & Leonard, 2006; reviewed by Gilchrist, 2006). An elegant demonstration of the importance of

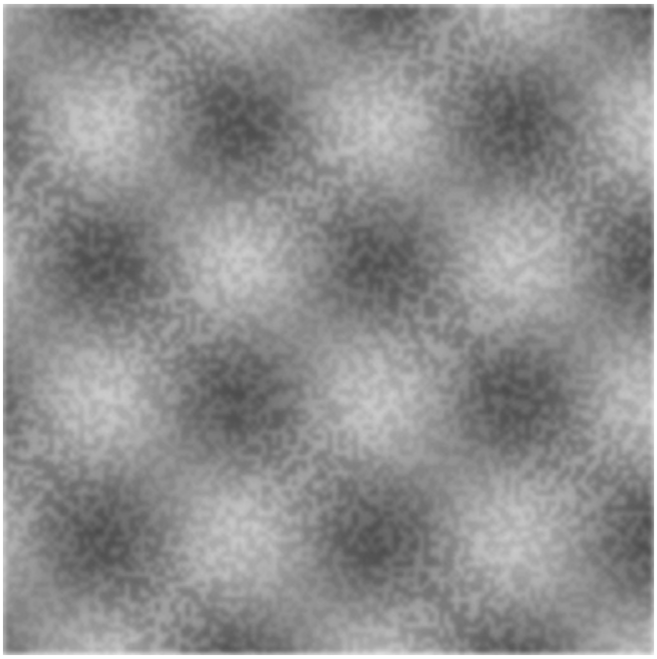

Fig. 3. A plaid consisting of two, orthogonal, sinusoidal luminance gratings is added to a noise texture. The noise texture itself is modulated in amplitude along both component orientations. In the left oblique component the luminance and amplitude modulations are in anti-phase, such that the peaks of luminance align with the troughs in amplitude, and no depth corrugations are perceived. In the right oblique component the luminance and amplitude modulations are in-phase, such that the peaks of luminance align with the peaks in amplitude, and depth corrugations are perceived. From Schofield et al. (2006), reproduced with permission.

penumbra for distinguishing shadows from shading can be seen in Elder, Trithart, Pintilie, and MacLean's (2004) crescent patterns in Fig. 4a. Switching the penumbra from the inside to the outside of a crescent changes the percept from shading to shadow. This seems to reveal one or other of two in-built assumptions, the first that the sharp edges of the crescent are occlusion boundaries (Anderson, personal communication), or second that objects tend to be convex-shaped.

Land and McCann (1971) based their famous Retinex algorithm, which was aimed at recovering the reflectance map of a Mondrian-like pattern that was obliquely illuminated, on the idea that gradual luminance gradients are caused by illumination. The algorithm removed the gradual luminance gradients associated with shading by (a) computing the 1D (first derivative) of the pattern's luminance map, then (b) imposing a threshold on the 1D map and then (c) re-integrating the 1D map to produce an image identical to the original minus the gradual luminance gradient. In principle one could go one step further and recover the shading image by subtracting the reflectance map from the original. The Retinex would presumably predict the appearance of Kingdom's (2003a) figure shown in Fig. 4b (which was inspired by Adelson's checkerboard-shadow illusion: http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html). One hardly notices the gradual luminance gradient in the figure, because the gradient is both gradual and low contrast. The important point however is that the lightnesses of the diamonds within each row appear very similar even though their luminances are very different. This is the 'correct' percept if the gradual luminance gradient is indeed shading[2]. The oft-noted shortcoming of the Retinex however is that it fails for patterns with relatively sharp illumination gradients, as in Fig. 2, since these would be identified as reflectance changes. Numerous studies attest to lightness being accurately perceived even in the context of sharp illumination changes (Arend & Spehar, 1993a; Arend & Spehar, 1993b; summarized by Arend, 1994), suggesting that sharp as well as gradual changes of luminance can be correctly identified as illumination.

Fig. 4c shows how a suitable arrangement of luminance gradients provides a compelling impression of a light source, one that glows and is 'brighter-than-white' (Agostini & Galmonte, 2002; Gori & Stubbs, 2006; Zavagno, 1999; Zavagno & Caputo, 2001; Correani et al., 2006). Thus as with Elder et al.s (2004) crescent figures, gradients are not only useful for discriminating light from material, but different types of light. The above considerations lead to the following heuristic:

Heuristic 2. Gradual luminance gradients are likely to be illumination changes, whereas sharp luminance gradients can be either changes in illumination, transparency or albedo.

2.1.4. Contrast magnitude

Although luminance gradients are probably the most effective cues for eliciting impressions of self-luminosity, such impressions can occur in the absence of gradients. Gilchrist (2006) argues that the conditions which produce impressions of self-luminosity in the absence of gradients are complex and not yet fully understood. Bonato and Gilchrist (1994); see also Gilchrist, 2006, p. 246–265) found that a region appeared self-luminous if the ratio of its luminance to that of a region appearing white was about 2.2. Fig. 4d suggests that at least under some circumstances small regions of high contrast, especially if they are 'outliers', appear either self-luminous or as spotlights. In this case self-luminosity could in part be detected by a mechanism that encoded the variance in contrasts within a display and identified as a light source any contrast that exceeded some threshold multiple of the contrast variance. If so this leads to the following heuristics.

Heuristic 3a. Regions appearing significantly brighter than white are likely to be self-luminous or spotlights.

Heuristic 3b. Positive outlier contrasts are likely to be self-luminous or spotlights.

Illumination edges appear lower in contrast than equal-in-contrast reflectance edges (Logvinenko, 2005), and the hues of grey regions illuminated by coloured light appear less saturated than if the colours are perceived as pigment (Jakobsson, Bergström, Gustaffson, & Fedorovskaya, 1997). Although it does not follow from these observations that low contrast regions are more likely perceived as illumination, Ekroll, Faul, and Niederee (2004) have put forward an interesting hypothesis consistent with the idea. Although their hypothesis has been framed within the context of chromatic contrast, it is also relevant to achromatic contrast, so is dealt with here. Ekroll et al. were struck by the observation that the hues of low contrast patches on uniform surrounds often appear *more* saturated than might be expected. To account for this they suggest that low contrast patches on uniform surrounds trigger a process of perceptual scission that results in a two-layered percept: a background layer the colour of the surround and a foreground, transparent layer whose colour is a function of the contrast between patch and surround. For a coloured patch on a grey background the two potential layers are grey background and saturated-hue foreground. At high contrasts the two layers are fused, and so the hue is desaturated because of the contribution of the grey component. However at low contrasts the components split, and the foreground transparent layer, which carries the colour, dominates, giving the patch its saturated hue appearance. Although the impression of transparency in low-contrast patches on uniform surrounds is not particularly compelling, Ekroll

[2] The more remarked-upon corollary of this type of figure is that diamonds **a** and **b**, which have the same luminance, appear very different in brightness. However this would only be remarkable if the visual system had access to absolute luminance. In this type of figure the dominant feature of the diamonds is lightness, and this is perceived near-veridically.

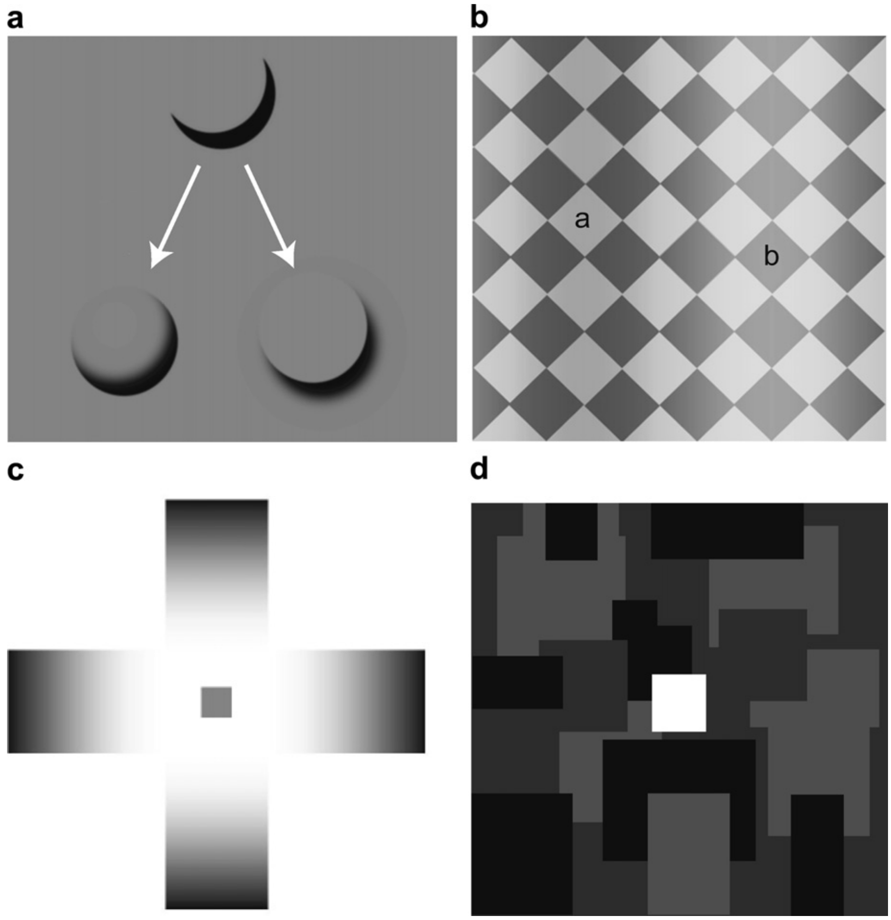

Fig. 4. (a) A crescent-shaped luminance profile can be perceived either as shading (bottom left) or shadow (bottom right) depending on which side is blurred; from Elder et al. (2004), reproduced with permission. (b) A single cycle of a vertical sinusoidal luminance grating is added to a checkerboard. The lightnesses of the diamonds in any row appear very similar, suggesting that the visual system identifies the sinusoidal luminance variation as shading not albedo. An incidental fact is that diamonds a and b have the same luminance. From Kingdom (2003a), reproduced with permission. (c) Four luminance ramps that converge towards the luminance of the background give a strong impression of a light source; from Agostini and Galmonte (2002), reproduced with permission. (d) Patches whose contrasts are positive outliers are sometimes perceived as spotlights or self-luminous.

et al.'s hypothesis offers an interesting alternative to the conventional view that changes in perceived saturation with contrast are caused by changes in contrast gain (e.g. Brown & MacLeod, 1997). Their scission hypothesis may also partly explain the impression of diffuse glow in the light-source figure in Fig. 4c. Although impressions of transparency experienced for low contrast reflectance patches on uniform surrounds are illusory, they may reflect an in-built assumption that isolated low-contrast patches are likely to be transparencies or illumination. In turn this might reflect the fact that in nature low contrast regions are more likely to be illumination than albedo or pigment. If so, an adventurous heuristic is:

Heuristic 4. Regions of low contrast on uniform surrounds are more likely to be caused by illumination or transparency than albedo or pigment.

2.1.5. Contrast polarity

The local luminance contrasts in textures not only have magnitude but sign (or polarity) depending on whether they are increments or decrements. The relative proportion of increments and decrements in a texture is reflected in the skew of the pixel luminance histogram: a predominance of decrements results in negative skew, and of increments positive skew (Kingdom, Hayes, & Field, 2001). Recently, Motoyoshi, Nishida, and Adelson (2005) and Motoyoshi, Nishida, Sharan, and Adelson (2007) showed that artificially skewing the luminance pixel histogram of a natural textured object can shift its appearance either to matte when skewed negatively or to glossy when skewed positively, as shown in Fig. 5. The shifts are consistent with a change in appearance of the surface markings from shading to specular reflection. Motoyoshi et al. (2007) suggest that the perceived differences in material resulting from differences in skew could result from the relative activities of neural filters with On- and Off-centre receptive fields, which are preferentially sensitive, respectively, to increments and decrements.

Heuristic 5. Negative skew in the luminance histogram of textured surfaces is evidence of shading, whereas positive skew is evidence of specular reflection.

2.2. Figural relations

Figural relations, as used here, refer to the arrangement and orientation of surface contours.

2.2.1. Junctions

Junctions form an important class of figural arrangement, and X-junctions, which have already been described, are widely believed to be signatures for non-uniform illumination and transparency (Adelson, 1993; Gilchrist, 1988). Eliminating the X-junction from Fig. 2b by rotating the central region by only 10 deg, as in

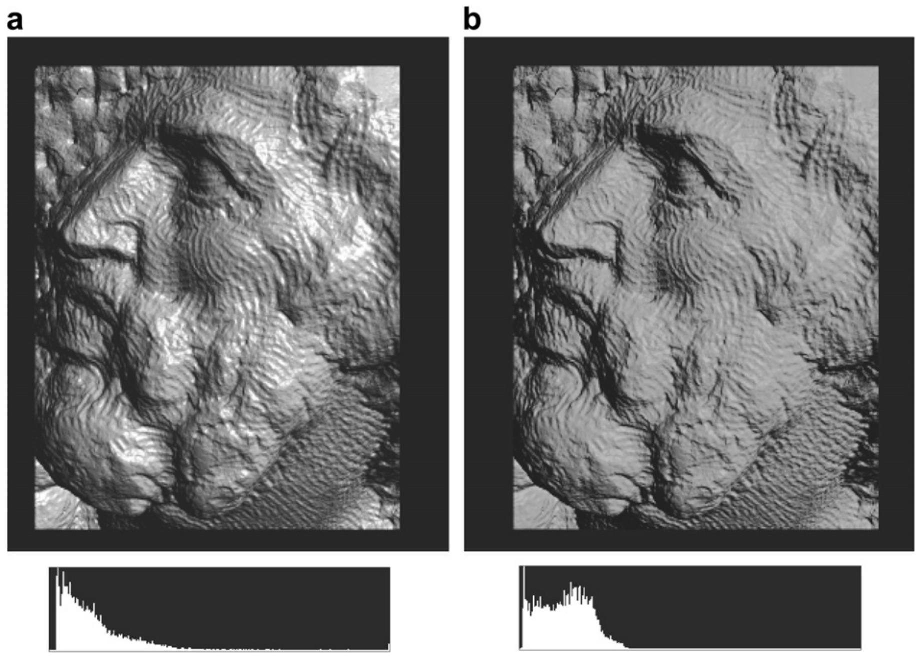

Fig. 5. (a) Skewing the pixel histogram of a textured surface towards the right creates the impression of specular reflections on a shiny metallic surface, whereas (b) skewing towards the left creates the impression of a dull matte surface with shading. From Motoyoshi et al. (2007), reproduced with permission.

Fig. 2e, destroys the impression of transparency, as does occluding the X-junction (Kasrai & Kingdom, 2002). Two other types of junction implicated in layer perception are T- and I-junctions (Anderson, 1997, 2001; Hoffman, 1998; Kanizsa, 1979). Fig. 6a, based on Kanizsa (1979), shows that T-junctions can trigger the impression of transparency, while Fig. 6b demonstrates that I-junctions (Anderson, 1997) can trigger the achromatic version of neon-colour spreading (Fig. 10e shows the more compelling chromatic version). An important constraint on the effectiveness with which T- and I-junctions generate impressions of transparency and neon-colour spreading is that they are multiply arranged in a manner consistent with the presence of a simple-shaped transparent overlay. X-junctions on the other hand are sufficient on their own to produce impressions of illumination/transparency, though the impressions are enhanced by multiple arrangements of the junction (see Fig. 2).

2.2.2. Straightness

If either border at an X-junction changes orientation at the junction, the impression of illumination/transparency is reduced, as Fig. 7a and b from Kanizsa (1979) attests. The change in percept reflects knowledge that in the natural world, illumination, transparency and reflectance edges do not by and large coincide, and therefore only rarely will a change in orientation of one type of edge coincide with the intersection of another. So when such an event does occur, we attribute it to different cause, namely a reflectance change. In Fig. 7c, Logvinenko et al. (2005) extend this theme and show that introducing curvature to an edge can change the impression from illumination/transparency to albedo. Besides the change in percept, Logvinenko et al. showed that the brightness difference between the equal-in-luminance spots placed either side of the edge was greater in the straight compared to curved edge condition. This would be expected if the brightness of the patches in the straight edge condition were strongly influenced by their lightnesses 'as seen through' the illumination/transparency. It would be wrong however to conclude from Fig. 7c that straight edges are in and of themselves perceived as illumination edges; the presence of X-junctions with the appropriate luminance relations in the figure strongly suggest illumination/transparency. Hence one can say:

Heuristic 6. Curvy edges are more likely to be changes in reflectance than illumination.

2.3. 3D-cues

2.3.1. 3D-shape

Changes in luminance that accompany changes in surface orientation tend to be perceived as shading (Fig. 8a; a versus b), whereas luminance changes on a flat surface tend to be perceived as changes in reflectance (Fig. 8a; b versus c)[3]. Knill and Kersten's (1991) demonstration in Fig. 8b neatly illustrates how the introduction of perceived surface curvature can shift the impression of a luminance discontinuity away from reflectance (left) towards shading (right).

Two other studies that reveal the importance of surface geometry for material-versus-light perception employ chromatic displays. Jakobsson et al's. (1997; see also Bergström, 2004) AMBEGUJAS demonstration (the acronym is a contraction of the authors' names) shows the effect in terms of perceived saturation. Subjects were shown a figure similar to the one shown in Fig. 10a-i. The top and bottom colours were produced by illuminating an achromatic stripe pattern with different coloured lights. The pattern appeared to alternate between the 'tile' and 'roof' percepts illustrated in Figs. 10a-ii and 10a-iii. During the tile percept, the hues looked saturated, whereas during the roof percept they appeared washed out. Jakobsson et al. suggest that the differences in perceived saturation dependent on perceived shape were caused by the fact that during the tile percept the two colours were perceived as different pigments, whereas during the roof percept they were perceived as different illuminations.

[3] As with the checkerboard-shading figure in Fig. 4b, the more remarked-upon percept in this type of figure is that **b** and **c**, which have the same luminance, appear markedly different in brightness and lightness. However as argued earlier, if the figure were a real object, the perceived lightness relations **a=b>c** are correct.

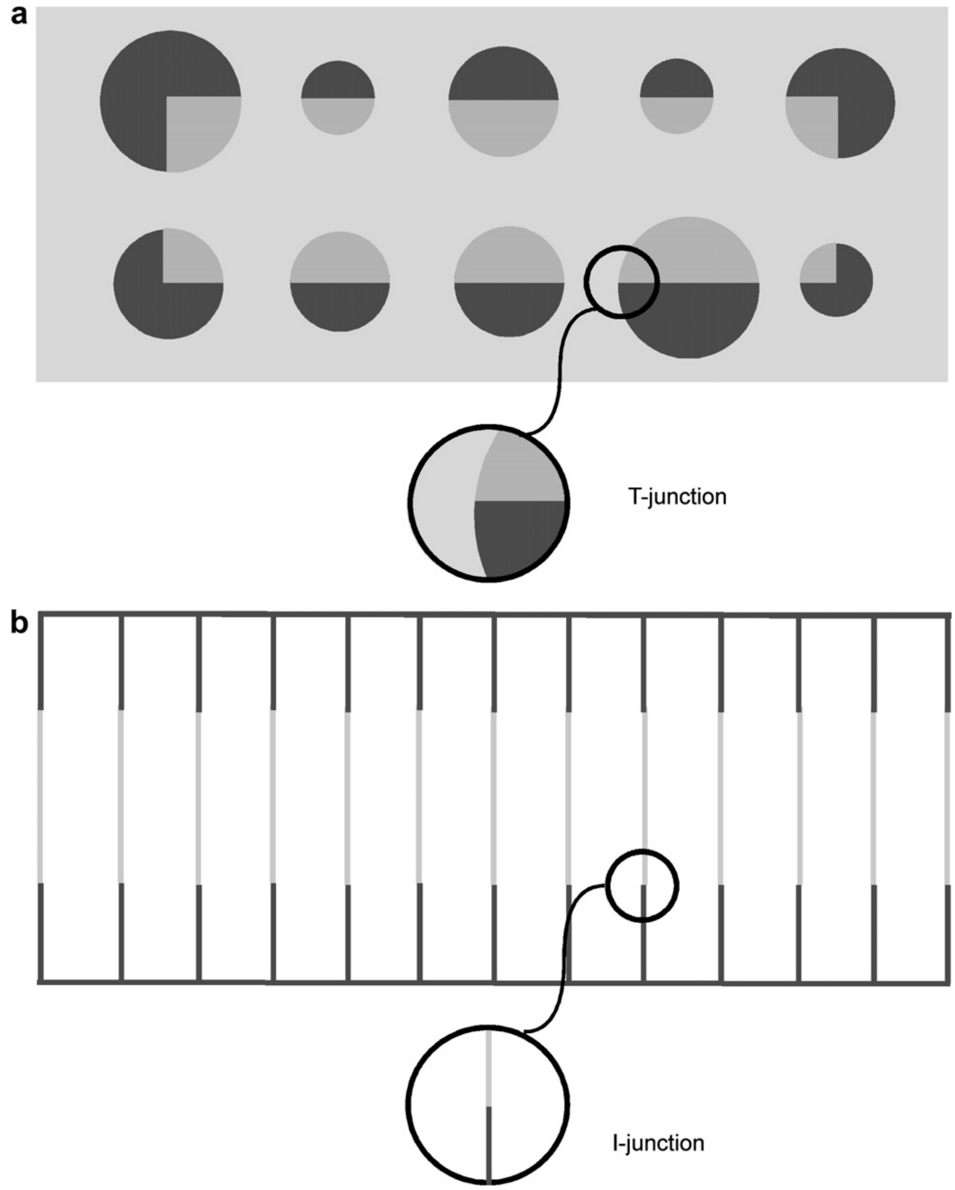

Fig. 6. (a) T-junctions can elicit an impression of transparency. Based on a similar figure in Kanizsa (1979). (b) I-junctions elicit an impression of achromatic neon-colour spreading.

Bloj, Kersten, and Hurlbert (1999) painted the inside of a V-shaped card magenta on one half and white on the other. Some of the light from the magenta half was reflected onto the white half, and as a result subjects perceived a pinkish glow on the white side. However when the perceived shape of the card was inverted with a pseudoscope, so that the card appeared convex (like so <, viewed from the left), the white half of the card took on a deep magenta color. Bloj et al. attributed the shift in colour percept caused by the change in perceived shape to the visual system's knowledge of physical inter-reflection. The reflected light from the magenta half that fell on the white half tended to be discounted in the normal view because the geometry was consistent with inter-reflection between surfaces. However in the inverted-shape, convex view, the arrangement was not consistent with inter-reflection and the light reflected from the magenta half was attributed to the surface of the white half, which therefore appeared magenta.

The reader may have noticed a contradiction between the results of AMBEGUJAS and Bloj et al. In the AMBEGUJAS experiment, the roof percept allegedly favoured the interpretation that the two colours were different illuminants, whereas in Bloj et al.'s demonstration the roof percept allegedly favoured the different pigment interpretation. It remains unclear as to the reason for these different percepts. Both studies nevertheless reveal the importance of surface geometry for the perception of illumination versus pigment.

Heuristic 7. A luminance change accompanied by a change in surface orientation is likely to be shading, whereas one unaccompanied by a change in surface orientation is likely to be change in reflectance.

Heuristic 8. The colours of concave-shaped surfaces are likely to be in part due to inter-reflected light.

2.3.2. Depth cues

The importance of depth relations for the perception of light versus material was revealed in a classic experiment by Gilchrist illustrated in Fig. 9 (Gilchrist, 1977; 1980; reviewed by Rock, 1977, and Gilchrist, 2006; see also experiments by Schirillo, Reeves, & Arend, 1990). When observers viewed the arrangement monocularly, as shown on the right, they perceived the two protruding patches as

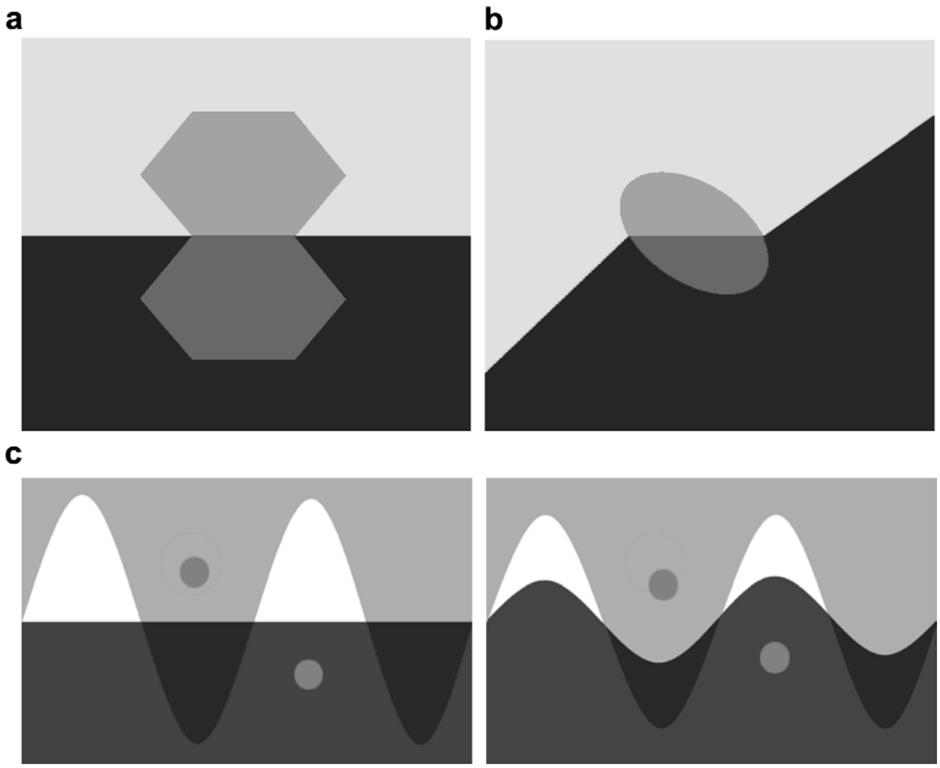

Fig. 7. If either (a) the border of the transparency, or (b) the border of the background changes orientation at the X-junction, the impression of transparency is lost; figures based on Kanizsa (1979). (c) A straight edge with X-junctions (left) elicits a strong impression of a transparency/shadow border, but if the edge is curved (right) it appears to be a pigment change; from Logvinenko et al. (2005), reproduced with permission.

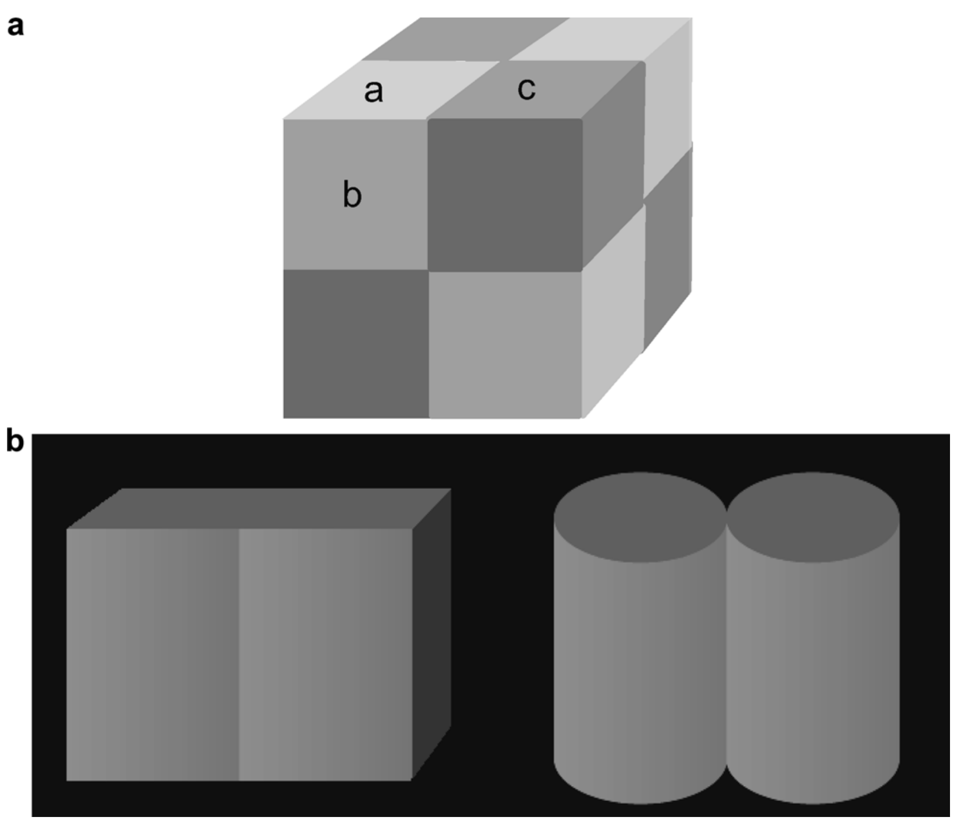

Fig. 8. Surface geometry can affect the perception of shading versus pigment. (a) A change in luminance is perceived as shading if accompanied by a change in surface orientation (a–b), but a change in pigment if unaccompanied by a change in surface orientation (a–c). An incidental fact is that the luminances b and c are the same. (b) Surface curvature can switch the percept of a brightness change from pigment to shading; from Knill and Kersten (1991), reproduced with permission.

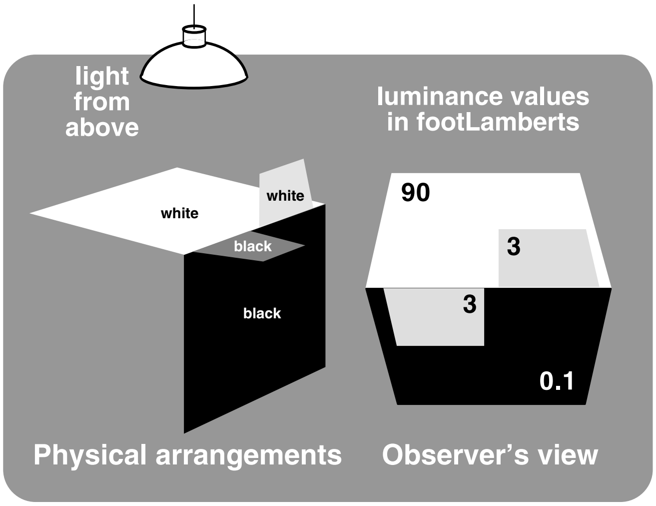

Fig. 9. Gilchrist's (1980) stimulus used to establish the coplanarity principle. See text for details. Reproduced from Gilchrist (2006) with permission.

being coplanar with their retinal backgrounds. Under these circumstances the lower patch appeared almost white while the upper patch appeared almost black. With binocular viewing however, observers saw the patches protruding with the upper patch coplanar with the black background and the lower patch coplanar with the white background. Now the patches appeared similar in lightness. What seems to be happening in this display is that subjects perceived the lightnesses of the patches as if they are illuminated by the same light as that illuminating the region with which they are coplanar. This observation led Gilchrist to formulate the 'coplanar ratio principle', which states that luminance changes between retinally-adjacent but non-coplanar papers play little or no role in lightness perception. It was Rock (1977) however who grasped the significance of Gilchrist's experiment for the lightness versus material issue, arguing that it revealed an in-built assumption that luminance differences within a depth plane are likely to be reflectance differences, whereas luminance differences between depth planes are likely to be illumination differences. The importance of this assumption is obvious when one considers the many situations in which other cues to non-uniform illumination, such as X-junctions and penumbra, are occluded by closer objects.

The impression of transparency is enhanced when the transparent layer is perceived to lie in front of the background (Anderson et al., 2006; Gurnsey et al., submitted for publication; Kingdom, Blakeslee, & McCourt, 1997). A categorical change in transparency perception dependent on relative depth can be seen in Nakayama and Shimojo's (1992) stereo display shown in Fig. 9f. When fused stereoscopically, the red bars appear projected in front in one stereo-image and behind in the other. In the former condition one perceives a delicate red-tinted transparency floating in front of the white cross, whereas in the latter condition one perceives a red form occluded by the black square. The wider theoretical significance of this experiment will be discussed later, but for now it suffices as a powerful demonstration of the potential effect of depth ordering on transparency versus pigment perception.

Heuristic 9. Luminance changes within a depth plane are likely to be reflectance changes, whereas luminance changes between depth planes are likely to be caused by different illumination.

2.4. Colour

2.4.1. Colour is material

In natural scenes most chromatic variations are changes in material, such as from pigment, whereas luminance variations are either changes in material or non-uniform illumination. These relationships can be observed in the photograph in Fig. 10c. The grass-pavement border is a change in both luminance and colour, whereas the shadow border is primarily a change in luminance. This suggests a simple rule: luminance edges that are co-aligned with chromatic edges are material changes, whereas luminance edges that are non-co-aligned with chromatic edges are illumination changes (Rubin & Richards, 1982). Although the potential usefulness of this rule for layer identification has been questioned on the grounds that shadows and shading are often coloured (Gilchrist, 2006), as in the bluing of deep shadows on sunny days (Churma, 1994), it has nevertheless proved effective in computer-vision algorithms designed to separate images into their reflectance and illumination layers (Finlayson, Hordley, & Drew, 2002a; Finlayson, Hordley, & Drew, 2002b; Olmos & Kingdom, 2004; Tappen, Freeman, & Adelson, 2003).

But what of perception? Notwithstanding that colour changes can under many circumstances be perceived to be caused by illumination (e.g. see previous section), Fig. 10b is evidence for the colour-is-material assumption. The building appears to be painted two-thirds orange, yet the orange part is in fact sunset illumination shining on a grey surface. Fig. 10d shows another photograph in which a shadow crosses a pigment boundary, forming an X-junction. In black-and-white (left), the edges of the junction are perceptually ambiguous when removed from their context (see inset), whereas the colour version (right) is less ambiguous.

The most powerful evidence in support of the colour-is-material assumption is the 'colour-shading effect' (Kingdom, 2003b; Kingdom, Rangwala, & Hammamji, 2005; Kingdom, Wong, Yoonessi, & Malkoc, 2006), illustrated in Fig. 11a (taken from Shevell & Kingdom, 2008). The luminance grating in Fig. 11a-i appears almost flat, yet when combined with the orthogonally oriented chromatic grating in Fig. 11a-ii produces a "plaid", as in Fig. 11a-iii. The plaid appears corrugated in depth, an example of shape-from-shading. In the plaid the changes in luminance are not aligned to the changes in colour, and this promotes the perceptual interpretation that the luminance variations are due to shading not pigment, characteristic of a corrugated surface illuminated obliquely. If the colour-is-material assumption is true, then one would predict that if the luminance variations in the plaid were now accompanied by corresponding colour variations, then the percept of shading and thus depth should be lost. This is exactly what happens. Adding a second colour grating in alignment with the luminance grating, as in Fig. 11a-iv, strongly reduces the impression of corrugated depth.

2.4.2. Colour transparency

The colour-is-material assumption also impacts transparency perception. Using stimuli similar to those illustrated in Fig. 11b-i–iii, Kingdom et al. (2004) found that "good" shadows/transparencies, i.e. those that exhibited luminance ratio invariance (see Section 2.1.1) were more easily identified when presented on chromatically-variegated (Fig. 11b-ii) compared to achromatically-variegated (Fig. 11b-i) backgrounds (see also Kingdom & Kasrai, 2006). In addition, along with other studies (Ripamonti & Westland, 2003), they found that the introduction of random colour changes across the transparency border (Fig. 11b-iii) disrupted the identification of "good" shadows/transparencies. The key to the negative impact of colour contrast on transparency perception however is that the colour changes are random. If the colour changes are consistent along the length of the transparency border, that is, always toward a particular color—blue in the example in Fig. 11b-iv—one retains a strong impression of transparency, but this time of a colored transparency such as color-tinted acetate (D'Zmura, Colantoni, Knoblauch, & Laget 1997; Fulvio, Singh, & Maloney, 2006; Khang & Zaidi 2002; Ripamonti & Westland 2003). Chromatic consistency at I-

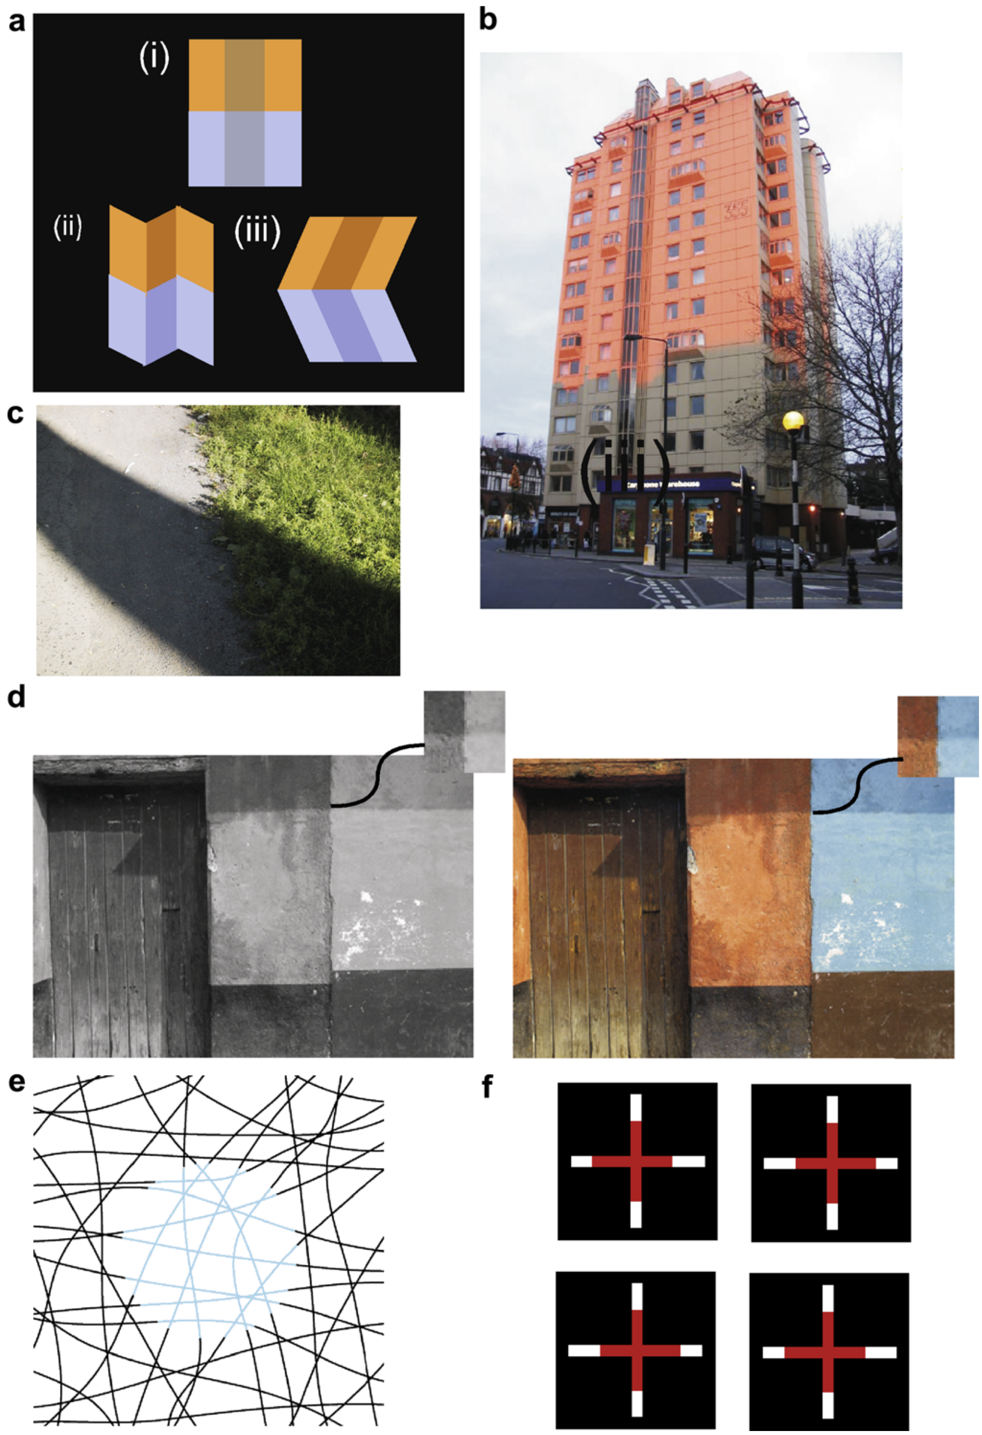

Fig. 10. (a) AMBEGUJAS stimulus based on Jakobsson et al. (1997), reproduced with permission. See text for details. (b) The orange colour of the tower block looks like paint but is in fact sunset illumination on a grey surface; photo courtesy of B. Micklethwait. (c) Shadow cast across a grass-pavement border reveals the spatial relations between colour and luminance in natural scenes. (d) Black-and-white (left) and colour (right) photograph of a wall with cast shadow. The insets showing the X-junction reveal how colour can disambiguate which border is pigment and which illumination. Taken from Kingdom et al. (2004); original colour photo courtesy of J. Marvullo and taken from *Color vision: A photographer's guide*. New York: Watson-Guptill Publications, 1989, p58. (e) Neon-colour spreading, courtesy of H. Knuchel, from http://www.blelb.ch/. (f) When free-fused, one of the stereo-pairs reveals a red disk occluded by a black square, the other a red transparent film floating in front of a white cross; based on Nakayama and Shimojo (1992), reproduced with permission.

junctions is also a critical factor in neon-colour spreading, as in Fig. 10e. The compelling impression one can obtain of a colored shadow or transparency is a counter-example to the colour-is-material assumption, but the percept occurs only in circumstances where there is already strong evidence for shadow or transparency, for example from X-junctions, as in Fig. 11b.

Heuristic 10. Luminance changes that are aligned with chromatic changes are perceived as pigment, whereas luminance changes that are non-aligned with chromatic changes are perceived as illumination.

Heuristic 11. Providing there are sufficient cues to the presence of a shadow or transparency, chromatic changes that are consistent

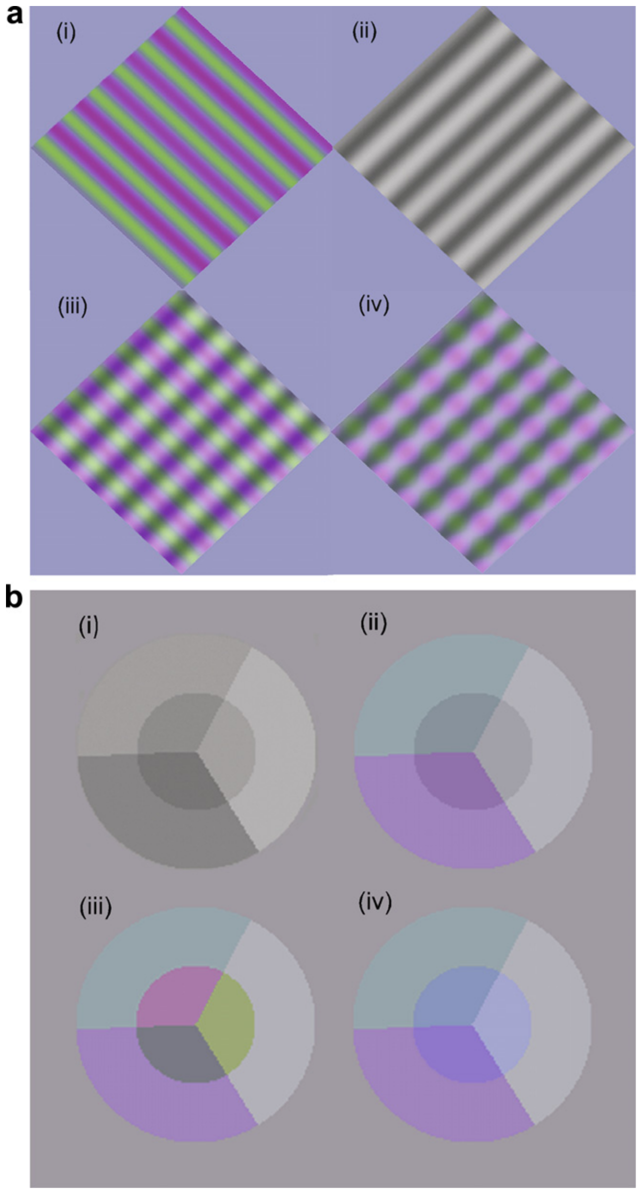

Fig. 11. (a) colour can promote or inhibit shape-from-shading. (i) Right oblique luminance grating; (ii) left oblique chromatic grating; (iii) Combining (i) and (ii) gives the impression of a depth-corrugated surface; (iv) adding a second chromatic grating aligned with the luminance grating in (ii) inhibits the impression of depth. (b) An achromatic non-reflective transparency/shadow on (i) a luminance-only background and (ii) a same-luminance but added-colour background. In (iii) the introduction of random colours changes across the transparency border reducs the impression of transparency. In (iv) the colour changes across the transparency border are systematic, e.g. all towards blue, giving rise to an impression of a colored transparency. Both figures based on Shevell and Kingdom (2008), reproduced with permission.

along the shadow/transparency border are evidence for coloured shadows or transparencies.

2.5. Texture

2.5.1. Orientation-defined textures

One might expect texture to operate in an analogous way to colour in facilitating the discrimination of light from material. The supposition in this case is that luminance changes that are aligned with texture changes are likely to be changes in reflectance, whereas luminance changes that are non-aligned with texture changes are likely to be changes in illumination or transparency, as in Fig. 10c. Although this may be true for some types of texture variation, for example variations in the size of texture markings, the situation is more complex for other types of texture variation, particularly those defined by orientation. Natural-scene textures that vary in local orientation arise in the retinal image primarily from changes in surface orientation (Knill, 2001; Li & Zaidi, 2000), and are exploited by vision for surface shape analysis (Li & Zaidi, 2000). Since shading often accompanies a change in surface orientation, it follows that a luminance change that is co-aligned with an orientation-defined texture boundary is not necessarily a change in reflectance; on the contrary it will more likely be shading. It should therefore come as no surprise that shading and orientation-defined texture variations combine synergistically to enhance surface shape perception (Kingdom et al., 2006; Mamassian & Landy, 2001). Thus one might expect that the visual system interprets as shading not only luminance changes that are non-aligned with texture changes, but ones that are aligned provided they are consistent with changes in surface orientation. Fig. 12c however suggests that *in the absence of other cues to shading* (such as non-sharp edges) the visual system is not so nuanced. Whereas the square-wave luminance grating in Fig. 12b, which is non-aligned with the orientation-defined texture grating, appears to be a shadow or transparency, the aligned-with-texture luminance gratings in Fig. 12a and c *both* appear more like changes in albedo, in spite of the fact that the texture variation in the latter figure is consistent with a change in surface orientation. One way to test this more thoroughly would be to compare the strengths of the perceived depth corrugations with and without the luminance grating, i.e. Fig. 12c compared to 12d. If the luminance grating in Fig. 12c is perceived as shading, the perceived corrugations should be larger in this stimulus.

Heuristic 12. Luminance changes unaccompanied by a change in texture are likely to be changes in illumination.

2.6. Motion

The cues described above for colour and texture potentially apply also to motion. A luminance pattern that moves coherently with colour, texture, shape etc. is likely to be a reflectance pattern, whereas if it moves incoherently with these dimensions it is likely to be a pattern of illumination or transparency. This principle has been used to extract the reflectance map of a scene as the common component when the scene is photographed under different lighting conditions (Weiss, 2001).

A stunning example of the importance of motion cues for layer perception is Hartung & Kersten's (2002; http://vision.psych.umn.edu/users/kersten/kersten-lab/demos/MatteOrShiny.html) 'Shiny versus Matte' demonstration. When the pattern of specular reflection on the surface of a simulated chrome teapot is made to move incoherently with the teapot's rotation, the impression of specularity is enhanced. However, if the pattern is moved coherently with the teapot's rotation, the appearance of the teapot changes dramatically from a shiny surface replete with specular reflections, to a matte surface painted in shades of yellow and brown.

Shadows can also be identified by incoherent motion. Jacobsen and Werner (2004) presented subjects with cubes whose surfaces were drawn with two patterns of luminance: stripes or leaves. As the cube was moved, one of the patterns was moved coherently with the cube's motion whereas the other was not. Whichever pattern was moved coherently was perceived as a reflectance pattern, whereas the incoherently moving pattern was invariably perceived as a shadow pattern.

Another cue to specular reflection is self-motion. When an observer moves, so too do the observable specular reflections. This does not happen with other types of non-uniform illumination

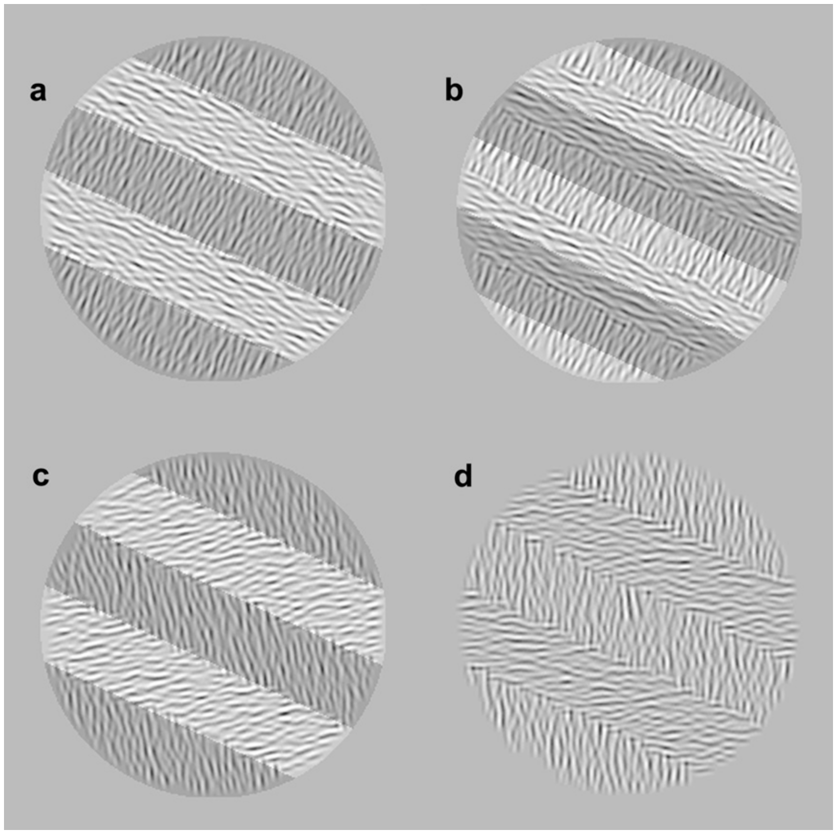

Fig. 12. Orientation-defined texture cues to layer perception. In (a) the square-wave orientation and luminance changes are co-aligned, and the latter appears as pigment. Shifting the relative phase of the luminance grating by 90 deg in (b) makes it appear as a shadow or transparency. In (c) the texture variation is consistent with a change in surface orientation. The co-aligned luminance grating, although physically consistent with shading, does not appear to be shading. If it were shading one would expect the perceived depth corrugations to be greater in (c) than in (d), where the luminance grating has been removed.

such as shading and shadows. The present author is not aware of any behavioural studies that have explored the efficacy of self-motion as a cue to specular reflection, but it is likely to prove particularly important precisely because it is so unique to this form of non-uniform illumination.

Heuristic 13. A luminance pattern that moves incoherently with an object's motion is likely to be a pattern of illumination, whereas if it moves coherently with an object's motion is likely to be a reflectance pattern.

Heuristic 14. If a luminance pattern is perceived to move as the observer moves it is likely to be specular reflection, otherwise it is likely to be some other form of non-uniform illumination or a reflectance pattern.

This completes the description of the physical cues available and the likely heuristics employed by the visual system for discriminating light from material. The following section discusses certain theoretical issues that arise from the above analysis.

3. Theoretical considerations

3.1. Importance of layer perception

By and large we do not wander around the world pondering whether luminance discontinuities look more like changes in reflectance or more like changes in illumination (well some of us admittedly do!). It is the behavioural consequences of light versus material identification that are, after all, what is important. Our ability to discriminate light from material is important for our perception of lightness, motion direction, depth, material composition, surface shape and object form. Some investigators regard layer perception as a critical stage of visual processing, lying between local feature extraction and fully-fledged object recognition, with behavioural consequences that extend far beyond the range of phenomena described here (Albert, 2001; Anderson, 1997, 2001, 2003; Nakayama, He, & Shimojo, 1995; Nakayama & Shimojo, 1992, 1996). On the other hand, one should not get too carried away. Artists often paint shadows and shading in the wrong positions and with the wrong properties, yet one tends not to notice (Cavanagh, 2005; Jacobsen & Werner, 2004). At the very least this suggests that global scene organization is not an important constraint on whether a region is perceived as shadow or reflectance. It also demonstrates that we still manage in spite of certain layer misperceptions.

3.2. A bag of tricks?

The heuristics outlined above constitute a sort of shopping list of rules used by vision for discriminating light from material, and might be thought of as a "bag of tricks" (Ramachandran, 1990), each pulled out of the bag as and when necessary. However the list is neither exhaustive nor sufficiently quantified to enable any predictions to be made except for the simplest of laboratory stimuli. It cannot be stressed enough that in the natural visual world multiple cues as to what is material and what is light co-exist, and it is the relative importance of each and the manner in which they are combined that is crucial. For example, the heuristic that colour variations are material is immediately discarded when watching cartoon films with brightly coloured shadows, because other cues such as motion and X-junctions render the shadows so effectively. For the same reason most natural shadows that do not have straight edges are still nevertheless perceived as shadows, and regions painted in subtle shades are often correctly perceived as paint, rather than as transparencies or changes in illumination. Readers of this review will doubtless have their own examples of heuristic failures, but should always bare in mind that the failures occur probably because one or more other heuristics are operating to promote the alternative interpretation.

The heuristics documented here embody knowledge of the regularities in the visual world, and these regularities add to those identified as useful to vision by numerous investigators beginning with Gibson (1979), and that include Marr (1982), Hoffman (1998) and most recently Purves and Lotto (2003). However the heuristics themselves say nothing about the mechanisms that implement them. Marr (1982) famously identified three levels at which visual processing could be analyzed: 1. identification of the regularities in the image exploited by vision for a particular computation; 2. the nature of the input representation and the algorithm that performs the computation; 3. the physiological hardware that implements the algorithm. The heuristics here embody the first, and to some extent the second of Marr's levels of analysis. If the goal of the computation is to identify whether a boundary is a shadow or albedo, ratio-invariance is one of the regularities apparently exploited by vision for this purpose. And one can easily envisage an algorithm along the following lines: find X-junction, compute contrasts at all junction edges, determine if ratio invariance exists, label junction as containing an illumination border etc. But with regard to the third, implementation stage in Marr's analysis, almost nothing is known. To this author's knowledge, there is no compelling evidence of neurons that either detect X-junctions or are sensitive to ratio-invariance in them. With some notable exceptions, such as Motoyoshi et al's (2007) suggestion that the relative activities of On and Off cells might underpin the impression of specular reflection, the neural mechanisms that implement the heuristics described above are neither known nor have even been guessed at. So if little or nothing is known about the underlying physiology, are there at least some general computational principles that we can point to, or should we just be content with regarding the heuristics as a bag of tricks?

3.3. Generic view principle

One mechanism that has been suggested to underpin some aspects of layer perception is the "generic view principle" (Hoffman, 1998; Nakayama & Shimojo, 1992, 1996, chap. 10; Freeman, 1994; Albert, 2001; Nakayama et al., 1995). This states that the favoured perceptual interpretation of an image is the one that is most invariant to small changes in viewpoint. For example consider Fig. 10f. In the stereo-pair in which the red bars project forward one perceives a reddish film floating in front of a white cross, rather than the alternative interpretation of a red cross with wings that are V-shaped in depth. According to the generic view principle, the reddish film interpretation is the one most valid across all viewing angles. The V-shaped wings interpretation would only be valid in the very rare or 'accidental' circumstance in which the cross was viewed precisely from the front, because only then would the cross completely occlude the white background. The generic view principle has been used to explain the conditions under which many illusory surface phenomena occur, such as neon-colour spreading (Hoffman, 1998) and specular reflection (Freeman, 1996), and it will be interesting to see to what extent it is applicable to the other layer percepts described here.

3.4. Bayesian estimation

An approach with potentially wider applicability than the generic view principle is Bayesian estimation (Kersten, Mamassian, & Yuille, 2004; Knill & Richards, 1996). Bayesian estimation offers the possibility of testing whether prior knowledge of the regularities in the world is used optimally for making perceptual decisions. The basic idea is that perceptual decisions are based on the computation of an *a posteriori* probability that a particular scene S is present given the image I, or $p(S|I)$. $p(S|I)$ can be calculated as proportional to the product of the likelihood of obtaining the image given the scene $p(I|S)$ and the prior probability of the scene $p(S)$. The generic view principle described above can in principle be incorporated within a Bayesian framework (Freeman, 1994; Nakayama & Shimojo 1996), though it has been argued that the Bayesian approach is problematic for situations such as Fig. 10f, because the prior probability $p(S)$ of a red-floating-transparency is presumably very low, yet our impression of the figure is nevertheless compelling (Albert, 2001; Nakayama & Shimojo, 1992, 1996; Nakayama et al., 1995).

Mamassian (personal communication) has suggested that some of the heuristics described in this review could be considered within the "explaining away" Bayesian framework described in Kersten et al. (2004). The general idea is that if there are two alternative hypotheses about a particular scene, one can "explain away" one of the two hypotheses if there is auxiliary evidence favouring the other. Fig. 13 illustrates how the idea applies to Knill and Kersten's (1991) demonstration of the effect of surface curvature on the interpretation of a luminance edge. According to Kersten et al. (2004), the two critical scene attributes in the figure are reflectance and shape, and each has two interpretations: 'discontinuous' versus 'uniform' for the reflectance attribute, 'flat' versus 'curved' for the shape attribute. The two interpretations are coupled, i.e. discontinuous reflectance with flat shape (left), uniform reflectance with curved shape (right). The Bayesian equation that captures "explaining away" in Fig. 13 is:

$$p(S_1, S_2|I_1, I_2) \propto p(I_2|S_2)p(I_1|S_1, S_2)p(S_1)p(S_2)$$

where the posterior probability $p(S_1, S_2 \mid I_1, I_2)$ is the probability of choosing one over the other of the two possible coupled scene interpretations. $p(S_1)$ and $p(S_2)$, are the prior probabilities that scenes have uniform reflectance and curved shapes, respectively. The likelihood function $p(I_1|S_1, S_2)$ is the probability of obtaining a luminance change (termed shading in the figure) given uniform reflectance and a curved shape. The other likelihood function $p(I_2|S_2)$ concerns the auxiliary information, and is the probability that a curved contour would arise in the image given a curved

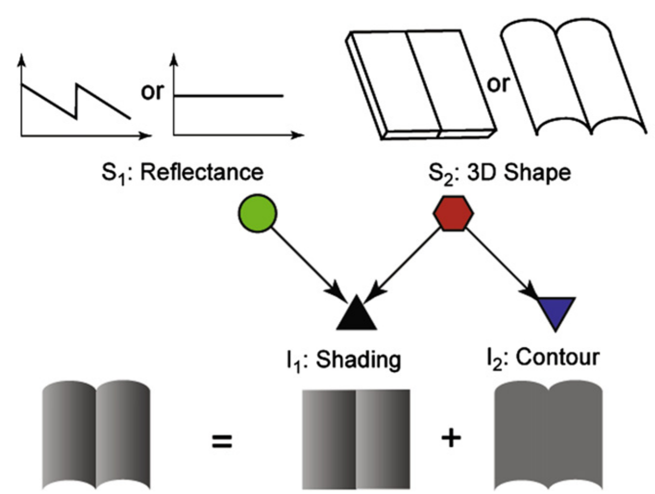

Fig. 13. Bayesian "explaining away" applied to Knill and Kersten's (1991) double-cylinder figure. See text for details. From Kersten et al. (2004), reproduced with permission.

shape. As this probability increases, the posterior probability that the scene is uniform in reflectance and curved in shape also increases. Mamassian suggests that this framework could be applied to some of the other phenomena described in this review, for example the colour-shading effect (Kingdom, 2003b), in which the introduction of non-aligned chromatic variations enhances shape-from-shading in a luminance pattern. Although there are good examples in the literature of how both prior probabilities and likelihood functions can be measured experimentally (e.g. Mamassian & Landy, 2001), it remains to be seen whether the Bayesian framework can be used to generate testable predictions as to the precise stimulus conditions that favour percepts of light rather than material.

4. Conclusion

Humans rarely confuse spatio-temporal variations in light from variations in albedo or pigment. This ability is achieved through the deployment of a range of heuristics that embody knowledge of the cues that exist as to what is light and what is material. However we currently have little understanding of how the heuristics are weighted and how they combine in situations where multiple cues co-exist. The generic view principle and Bayesian estimation offer a possible computational framework for understanding the perception of light versus material, and this will hopefully be complemented in the future by single-unit recordings and brain imaging studies that reveal the underlying physiology and anatomy.

Acknowledgements

This review was supported by Canadian Institute of Health Research Grant #11554 given to the author. Special thanks to Alan Gilchrist and Bart Anderson for providing the official reviews, and to Pascal Mamassian, Vejbörn Ekroll, Andrew Schofield and Curtis Baker for their comments and contributions during earlier drafts. Thanks also to all those who gave permission to publish the figures.

References

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

Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neuroscience* (2nd ed., pp. 339–351). Cambridge, MA: MIT Press.

Adelson, E. H., & Pentland, A. P. (1996). The perception of shading and reflectance. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian Inference*. Cambridge: Cambridge University Press. Ch. 11.

Agostini, T., & Galmonte, A. (2002). A new effect of luminance gradient on achromatic simultaneous contrast. *Psychonomic Bulletin & Review, 9*, 264–269.

Albert, M. K. (2001). Surface perception and the generic-view principle. *Trends in Cognitive Sciences, 5*, 197–203.

Albert, M. K. (2006). Lightness and perceptual transparency. *Perception, 35*, 433–443.

Albert, M. K. (in press). The role of contrast in the perception of achromatic transparency: Comment on Singh and Anderson (2002) and Anderson (2003). *Psychological Review*.

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

Anderson, B. L. (2001). Contrasting theories of White's illusion. *Perception, 30*, 1499–1501.

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

Anderson, B. L., Singh, M. (in press). Natural decompositions of perceived transparency: Reply to Albert. *Psychological Review*.

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

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

Arend, L. E. (1994). Surface colors, illumination, and surface geometry: Intrinsic-image models of human color perception. In E. Gilchrist (Ed.), *Lightness, brightness, and transparency* (pp. 159–213). Hillsdale: Erlbaum.

Arend, L. R., & Spehar, B. (1993a). Lightness, brightness and brightness contrast: 1. illuminance variation. *Perception & Psychophysics, 54*, 446–456.

Arend, L. R., & Spehar, B. (1993b). Lightness, brightness and brightness contrast: 2. Reflectance variation. *Perception & Psychophysics, 54*, 457–468.

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

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

Bergström, S. S. (2004). The AMBEGUJAS phenomenon and colour constancy. *Perception, 33*, 831–835.

Bloj, M. G., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature, 402*, 877–879.

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

Bressan, P., Mingolla, E., Spillmann, L., & Watanabe, T. (1997). Neon color spreading: a review. *Perception, 26*, 1353–1366.

Brown, R. O., & MacLeod, D. I. A. (1997). Color appearance depends on the variance of surround colors. *Current Biology, 7*, 844–849.

Cavanagh, P. (2005). The artist as neuroscientist. *Nature, 434*, 301–307.

Cavanagh, P., & Kennedy, J. M. (2000). Close encounters: Details veto depth from shadows. *Science (Letters), 287*, 2423–2424.

Cavanagh, P., & Leclerc, Y. (1989). Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 3–27.

Churma, M. E. (1994). Blue shadows: Physical, physiological, and psychological causes. *Applied Optics, 33*, 4719–4722.

Correani, A., Scott-Samuel, N. E., & Leonard, U. (2006). Luminosity-A perceptual "feature" of light-emitting objects? *Vision Research, 46*, 3915–3925.

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

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

Elder, J. H., Trithart, S., Pintilie, G., & MacLean, D. (2004). Rapid processing of cast and attached shadows. *Perception, 33*, 1319–1338.

Finlayson, G. D., Hordley, S. D., Drew, M. S. (2002a). Removing shadows from images. In *European conference on computer vision, ECCV'02* (Vol. 4). pp. 823–836.

Finlayson, G. D., Hordley, S. D., Drew, M. S. (2002b). Removing shadows from images using retinex. IS&T/SID tenth colour imagining conference. pp. 73–79.

Freeman, W. T. (1994). The generic viewpoint assumption in a framework for visual perception. *Nature, 368*, 542–545.

Freeman, W. T. (1996). The generic viewpoint assumption in a Bayesian framework. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian inference*. Cambridge: Cambridge University Press.

Fulvio, J. M., Singh, M., & Maloney, L. T. (2006). Combining achromatic and chromatic cues to transparency. *Journal of Vision, 6*, 760–776.

Gerbino, W. (1994). Achromatic transparency. In A. L. Gilchrist (Ed.), *Lightness, brightness and transparency* (pp. 15–255). Hove, UK, pp: Lawrence Erlbaum Associates.

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

Gibson, J. J. (1979). *The ecological approach to visual perception.* Houghton Mifflin, Boston.

Gilchrist, A. L. (1977). Perceived lightness depends on perceived spatial arrangement. *Science, 195*, 185–187.

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

Gilchrist, A. L. (1980). When does perceived lightness depend on perceived spatial arrangement. *Perception and Psychophysics, 28*, 527–538.

Gilchrist, A. (1988). Lightness contrast and failures of constancy: A common explanation. *Perception and Psychophysics, 43*, 415–424.

Gilchrist, A. (2006). *Seeing black and white*. Oxford University Press. pp. 180–187.

Gilchrist, A., Delman, S., & Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination. *Perception & Psychophysics, 33*, 425–436.

Gori, S., & Stubbs, D. A. (2006). A new set of illusions—The dynamic luminance-gradient illusion and the breathing light illusion. *Perception, 35*, 1573–1577.

Gurnsey, R., Kingdom, F. A. A. & Schofield, A. J. (submitted for publication). Matching neutral density filters placed over disparate backgrounds: brightness dominates and contrast is an afterthought. *Journal of Vision*.

Hartung, B., & Kersten, D. (2002). Distinguishing shiny from matte. *Journal of Vision, 2*(7), 551a. doi:10.1167/2.7.551. abstract.

Helmholtz, H. von. (1866/1962) *Treatise on physiological optics* (Vol. II). Trans. J.P.L. Southall, Dover Publications: New York. pp. 264–301. (Vol. II originally published in 1866).

Hering, E. (1874/1964) Outlines of a theory of light sense. Cambridge MA: Harvard University Press [originally published in 1874].

Hoffman, D. D. (1998). *Visual intelligence*. New York: Norton.

Jacobsen, J., & Werner, S. (2004). Why cast shadows are expendable: Insensitivity of human observers and the inherent ambiguity of cast shadows in pictorial art. *Perception, 33*, 1369–1383.

Jakobsson, T., Bergström, S. S., Gustaffson, K.-A., & Fedorovskaya, E. (1997). Ambiguities in colour constancy and shape-from-shading. *Perception, 26*, 531–541.

James, W. (1981) *The principles of psychology*. Cambridge, Massachusetts: Harvard University Press (Original work published in 1890).

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

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

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

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.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object Perception as Bayesian inference. *Annual Review of Psychology, 55*, 271–304.

Khang, B.-G., & Zaidi, Q. (2002). Accuracy of color scission for spectral transparencies. *Journal of Vision, 2*, 451–466.

Kingdom, F. A. A. (1997). Simultaneous contrast: The legacies of Hering and Helmholtz. *Perception, 26*, 673–677.

Kingdom, F. A. A. (2003a). Levels of brightness perception. In L. Harris & M. Jenkin (Eds.), *Levels of perception*. Springer-Verlag.

Kingdom, F. A. A. (2003b). Colour brings relief to human vision. *Nature Neuroscience, 6*, 641–644.

Kingdom, F. A. A., Beauce, C., & Hunter, L. (2004). Colour vision brings clarity to shadows. *Perception, 33*, 907–914.

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

Kingdom, F. A. A., Hayes, A., & Field, D. J. (2001). Sensitivity to contrast histogram differences in synthetic wavelet-textures. *Vision Research, 41*, 585–598.

Kingdom, F. A. A., & Kasrai, R. (2006). Colour unmasks dark targets in complex displays. *Vision Research, 46*, 814–822.

Kingdom, F. A. A., Rangwala, S., & Hammamji, K. (2005). Chromatic properties of the colour shading effect. *Vision Research, 45*, 1425–1437.

Kingdom, F. A. A., Wong, K., Yoonessi, A., & Malkoc (2006). Colour contrast influences perceived depth in combined shading and texture patterns. *Spatial Vision, 19*, 147–159.

Knill (2001). Contour into texture: Information content of surface contours and texture flow. *Journal of Optical Society of America A, 18*, 12–35.

Knill, D. C., & Kersten, D. (1991). Apparent surface curvature affects lightness perception. *Nature, 351*, 228–230.

Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*. London: Cambridge Univ. Press.

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

Li, A., & Zaidi, Q. (2000). Perception of three-dimensional shape from texture is based on patterns of local orientation. *Vision Research, 40*, 217–242.

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

Logvinenko, A. D. (2005). On achromatic colour appearance. In J. L. Nieves J. Hernanderez-Andres (Eds.), *AIC Colour 05. The 10th congress of the international colour association. 8–13 May 2005, Granada, Spain, Part 1(* pp. 639–642).

Logvinenko, A. D., Adelson, E. H., Ross, D. A., & Somers, D. (2005). Straightness as a cue for luminance edge classification. *Perception & Psychophysics, 67*(1), 120–128.

MacLeod, R. B. (1947) The effect of "artificial penumbrae" on the brightness of included areas. *Miscellanea Psychologica Albert Michotte* (pp. 138–154). Louvain: Institut Superieur de Philosophie.

Mamassian, P., & Landy, M. S. (2001). Interaction of prior visual constraints. *Vision Research, 41*, 2653–2668.

Marr, D. (1982). *Vision: A computational investigation into human representation and processing of visual information*. San Francisco: W.H. Freeman.

Masin, S. C. (1997). The luminance conditions of transparency. *Perception, 26*(1), 39–50.

Masin, S. C. (2006). Test of models of achromatic transparency. *Perception, 35*, 1611–1624.

Masin, S. C., Tommasi, M., & Da Pos, O. (2007). Test of the Singh–Anderson model of transparency. *Perceptual and Motor Skills, 104*, 1367–1374.

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

Motoyoshi, I., Nishida, S., & Adelson, E. H. (2005). Luminance re-mapping for the control of apparent material. In *ACM proceedings of the 2nd symposium on applied perception in graphics and visualization* (Vol. 95, p. 165).

Motoyoshi, I., Nishida, S., Sharan, L., & Adelson, E. H. (2007). Image statistics and the perception of surface qualities. *Nature, 447*, 206–209.

Nakayama, K., He, Z. J., & Shimojo, S. (1995). Intermediate-level vision (2nd ed.. In S. M. Kosslyn & D. N. Osherson (Eds.). *An invitation to cognitive science: Visual cognition* (Vol. 2, pp. 1–70). MIT Press.

Nakayama, K., & Shimojo, S. (1992). Experiencing and perceiving visual surfaces. *Science, 257*, 1357–1363.

Nakayama, K., & Shimojo, S. (1996). Experiencing and perceiving visual surfaces. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 391–407). Cambridge: Cambridge University Press.

Olmos, A., & Kingdom, F. A. A. (2004). A biologically-inspired algorithm for separating reflectance and shading images. *Perception, 33*, 1463–1473.

Purves, D., & Lotto, R. B. (2003). *Why we see what we do. An empirical theory of vision*. MA: Sinauer Associates, Sunderland.

Ramachandran, V. S. (1990). In C. Blakemore (Ed.), *The utilitarian theory of perception* (pp. 346–360). Cambridge: Cambridge Univ. Press.

Ripamonti, C., & Westland, S. (2003). Prediction of transparency perception based on cone-excitation ratios. *Journal of Optical Society of America A, 20*, 1673–1680.

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

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

Rock, I. (1977). In defense of unconscious inference. In W. Epstein (Ed.), *Stability and constancy in visual perception: Mechanisms and processes* (pp. 321–373). New York: Wiley.

Rubin, J. M., & Richards, W. A. (1982). Color vision and image intensities: When are changes material? *Biological Cybernatics, 45*, 215–226.

Schirillo, J., Reeves, A., & Arend, L. (1990). Perceived lightness, not brightness, of achromatic surfaces depends on perceived depth information. *Perception & Psychophysics, 48*, 82–90.

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

Schofield, A. J., Heese, G., Rock, P. B., & Georgeson, M. A. (2006). Local luminance amplitude modulates the interpretation of shape from shading in textured surfaces. *Vision Research, 46*, 3462–3482.

Shevell, S. K., & Kingdom, F. A. A. (2008). Color in complex scenes. *Annual Review of Psychology, 59*, 143–166.

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

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

Tappen, M. F., Freeman, W. T., & Adelson, E. H. (2003). Recovering intrinsic images from a single image. In *Advances in neural information processing systems 15 (NIPS)*. MIT Press.

Turner, R. S. (1994). *In the eye's mind: Vision and the Helmholtz–Hering controversy*. Princeton, New Jersey: Princeton University Press.

van Tuijl, H. F. J. M. (1975). A new visual illusion: Neonlike color spreading and complementary color induction between subjective contours. *Acta Psychologica, 39*, 441–445.

Weiss, Y. (2001) Deriving intrinsic images from image sequences. In *Proceedings of the eighth IEEE international conference on computer vision* (Vol. 2, pp. 68–75).

Zavagno, D. (1999). Some new luminance-gradient effects. *Perception, 28*, 835–838.

Zavagno, D., & Caputo, G. (2001). The glare effect and the perception of luminosity. *Perception, 30*, 209–222.