Vision Research xxx (2010) xxx–xxx

Contents lists available at ScienceDirect

Vision Research

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

Review

Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy

Frederick A.A. Kingdom

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

ARTICLE INFO

Article history:
Received 18 March 2010
Received in revised form 3 September 2010
Available online xxxx

Keywords:
Lightness
Brightness
Transparency
Illumination
Layer decomposition
Intrinsic images

ABSTRACT

The past quarter century has witnessed considerable advances in our understanding of Lightness (perceived reflectance), Brightness (perceived luminance) and perceived Transparency (LBT). This review poses eight major conceptual questions that have engaged researchers during this period, and considers to what extent they have been answered. The questions concern 1. the relationship between lightness, brightness and perceived non-uniform illumination, 2. the brain site for lightness and brightness perception, 3 the effects of context on lightness and brightness, 4. the relationship between brightness and contrast for simple patch-background stimuli, 5. brightness "filling-in", 6. lightness anchoring, 7. the conditions for perceptual transparency, and 8. the perceptual representation of transparency. The discussion of progress on major conceptual questions inevitably requires an evaluation of which approaches to LBT are likely and which are unlikely to bear fruit in the long term, and which issues remain unresolved. It is concluded that the most promising developments in LBT are (a) models of brightness coding based on multi-scale filtering combined with contrast normalization, (b) the idea that the visual system decomposes the image into "layers" of reflectance, illumination and transparency, (c) that an understanding of image statistics is important to an understanding of lightness errors, (d) Whittle's log *W* metric for contrast-brightness, (e) the idea that "filling-in" is mediated by low spatial frequencies rather than neural spreading, and (f) that there exist multiple cues for identifying non-uniform illumination and transparency. Unresolved issues include how relative lightness values are anchored to produce absolute lightness values, and the perceptual representation of transparency. Bridging the gap between multi-scale filtering and layer decomposition approaches to LBT is a major task for future research.

1. Introduction

The past 25 years have seen the study of lightness, brightness and transparency (LBT) grow in leaps and bounds. New theories have emerged, and some established theories have found new expression. Other theories have failed to fulfil their promise and have withered away. No theory has yet fully flowered. Aided by the computer graphics revolution, we have witnessed an explosion of captivating demonstrations showing how much we may err in our judgements of brightness and lightness. And the controversies continue unabated. Divided into different camps, each with its own preferred stimuli, methodology and theory, the study of LBT is sometimes more reminiscent of the social sciences with its deep ideological divides than it is of the neurosciences.

This review poses major conceptual questions that engaged LBT researchers a quarter of a century ago and considers to what extent progress has been made towards answering them. In particular, the review will critically analyze the principal theoretical approaches to LBT that have been advanced in the recent period and suggest which are likely and which are unlikely to bear fruit in the long term. Of course what the major conceptual questions are, and which approaches will ultimately succeed, is largely a matter of opinion, but hopefully what transpires below will resonate not only with those who hold a particular viewpoint about LBT, but also with the interested, albeit uncommitted members of the wider vision community. The questions posed are as follows:

1. What is the relationship between brightness, lightness and perceived illumination?
2. Where is lightness and brightness encoded in the brain?
3. Is there a general theory for the effects of context on brightness and lightness?
4. What is the function relating brightness to contrast in simple patch-background displays?
5. How are brightness and lightness "filled-in"?
6. How are relative lightness values converted to absolute ones?
7. What are the cues for perceiving non-uniform illumination and transparency?
8. What are the dimensions of perceptual transparency and how are they encoded?

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

0042-6989/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2010.09.012

Before considering these questions in detail, the following introductory sections provide some rudimentary background.

1.1. Image decomposition

Any luminance image can be decomposed into 'layers', or 'intrinsic-images' (Barrow & Tenenbaum, 1978), and the reflectance, illumination and transparency layers are critical to vision. Surface reflectance, or 'albedo', is the proportion of light incident on a surface that is reflected from it. Spatio-temporal variations in surface albedo arise from changes in material composition, such as from different shades-of-grey of paint. Illumination is neither temporally nor spatially uniform. Temporally, the most dramatic changes in illumination occur in the ambient level as a result of the diurnal cycle. Spatially, the visual world is replete with various types of non-uniform illumination: shadows, shading, spotlights, inter-reflections, highlights and light sources, to name the main varieties. Shadows are caused by occlusion of the light source; shading (sometimes termed "attached" shadows) from changes in the angle of a surface with respect to the direction of illumination; spotlights by the projection of a light beam onto a surface; inter-reflections from light bouncing back and forth between surfaces; highlights, or specular reflections, from non-diffuse illumination of shiny or metallic surfaces. Example light sources are the sun, moon and artificial lighting.

Transparencies are media that we can see through, and vary along a number of dimensions. The principal two physical dimensions of transparency are transmittance, which is the proportion of light that passes through the medium, and reflectance, which is the proportion of incident light reflected from the medium. Transparent media have a dispersion, which is the degree of internal light scatter, and a specular reflection, the degree to which the medium acts like a mirror. A standard pair of dark glasses transmits a proportion of the incident light but does not reflect any light. Dark glasses hence act similarly to a shadow, in that they divide the background luminance by a constant factor, reducing the light level but leaving the contrast of objects viewed through them unchanged. Transparencies that reflect as well as transmit light are said to possess a reflective or additive component, since they reduce the contrast of objects seen through them. Reflective transparencies invariably appear 'milky', and include media such as frosted glass, grease-proof paper, milky water and fog. These example media also have a dispersion. Example simulated transparencies that vary in transmittance and reflectance are shown in Fig. 23.

Fig. 2. Perceptual layer decomposition enables us to distinguish among reflectance (the white walls) shading (underside of roof, inside of left wall) and shadows (on front-facing walls).

Fig. 1 provides a simple framework for describing the physical dimensions of the achromatic visual world. The focus of this review however is on the perceptual representation of these dimensions. Some of these dimensions can be seen in the photograph in Fig. 2, an image that will serve as a focal point for many of the issues discussed below.

1.2. Lightness

Lightness is perceived reflectance. If x and y are Cartesian coordinates in the two-dimensional image plane, the luminance $L(x, y)$, reflectance $R(x, y)$ and illumination $I(x, y)$ layers are related by the simple equation: $L(x, y) = I(x, y)\,R(x, y)$. This equation reveals that it is impossible to determine from the luminance of any one image point whether it is part of a light surface that is shaded or a dark surface that is brightly illuminated: pixel luminance is inherently ambiguous because an infinite combination of the two unknowns, $R(x, y)$ and $I(x, y)$ can produce the same $L(x, y)$. It is therefore axiomatic that only by examining the *relations* among image points can one determine the perceptual correlates of $R(x, y)$, i.e. lightness. The ability to recognize that surfaces under different illuminations

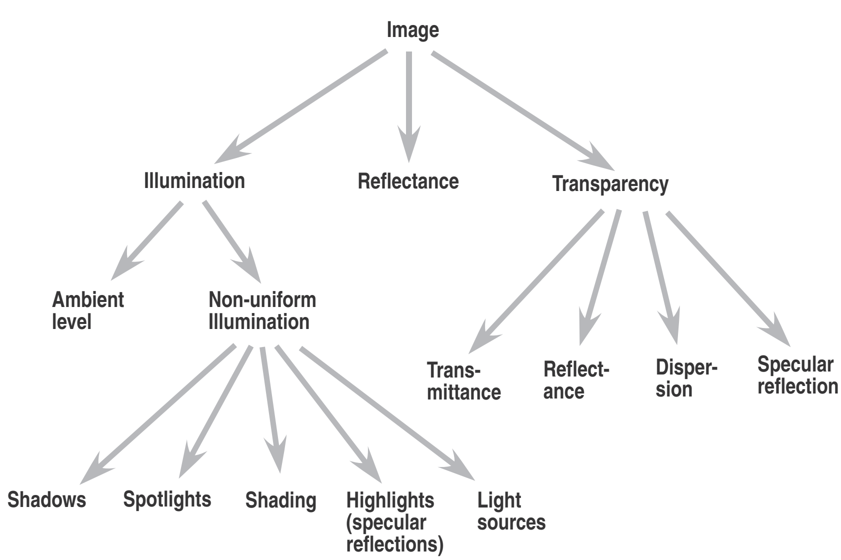

Fig. 1. The physical dimensions of achromatic experience.

Please cite this article in press as: Kingdom, F. A. A. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* (2010), doi:10.1016/j.visres.2010.09.012

have the same reflectance is termed 'lightness constancy'. Lightness constancy in natural scenes is generally believed to be quite good (though see recent findings to the contrary by Baddeley, Attewall, & Patel, 2010) – the white walls in Fig. 2 appear white even though their luminances vary considerably over space and time. It is virtually axiomatic that the visual system achieves lightness constancy in part by *comparing* the luminances of surfaces between different parts of the image. Although the details of the underlying computations are not fully understood, the basic strategy works because luminance ratios, or contrasts, between surfaces remain invariant with changes in illumination (Jacobsen & Gilchrist, 1988; Wallach, 1976). The notion that brightness and lightness involves the computation of *contrast* originated in Hering's (1874/1964) ideas on the role of reciprocal neural interactions for the perception of brightness, whose closest modern expression are the spatial-filtering models discussed in Section 4.

While the computation of contrast is in principle sufficient to generate a scale of *relative* lightness values (i.e. this surface is lighter than that one but only slightly darker than that other one), *absolute* lightness values (this surface is light-grey, that one is white) require that the contrasts be 'anchored' to some value. The main contendors for lightness anchors are the highest luminance, which is assigned white, and/or the mean luminance which is assigned mid-grey. Lightness anchoring is discussed in Section 7.

It is no exaggeration to say that the study of lightness perception over the past 25 years has been dominated by an exhaustive examination of its errors. This however needs to be put in context. In our everyday visual experience we rarely encounter lightness errors in the form of blatantly contradictory lightness percepts, as for example when a surface appears to change in lightness when viewed under a different illumination or against a different background. Good lightness constancy is presumably related to the fact that our natural visual world is sufficiently rich in contextual information to enable our visual system by-and-large to get it right (though again see Baddeley et al., 2010). It is in the terrain of the laboratory stimulus that we are confronted with the limitations of lightness perception, precisely because laboratory stimuli are deliberately fashioned to bring out those limitations. The simplest of these laboratory stimuli is the classic simultaneous contrast display shown in Fig. 3, in which two equal-in-luminance patches appear different in lightness when set against different luminance backgrounds. This stimulus, together with its innumerable variants, continues to play a big part in the scientific investigation of lightness (and brightness).

An even more striking error that also reveals the importance of contrast, but in addition the process of layer decomposition for lightness perception, is the well-known Gelb effect, popularized in recent times by Alan Gilchrist in numerous live-audience demonstrations (Gelb, 1929; a translation of Gelb's own description of the effect can be found in Gilchrist, 2006, p. 62). The schematic in Fig. 4 shows a smooth black disk lit from below by a powerful stage lamp. The light from the lamp falls on the wall behind the disk but well above it, such that an observer standing in front of

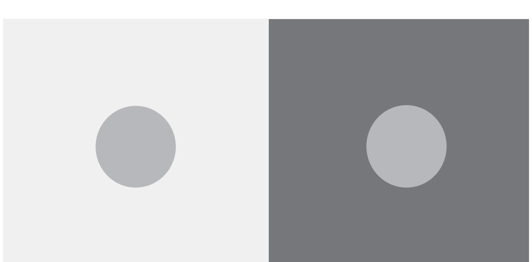

Fig. 3. Demonstration of simultaneous contrast.

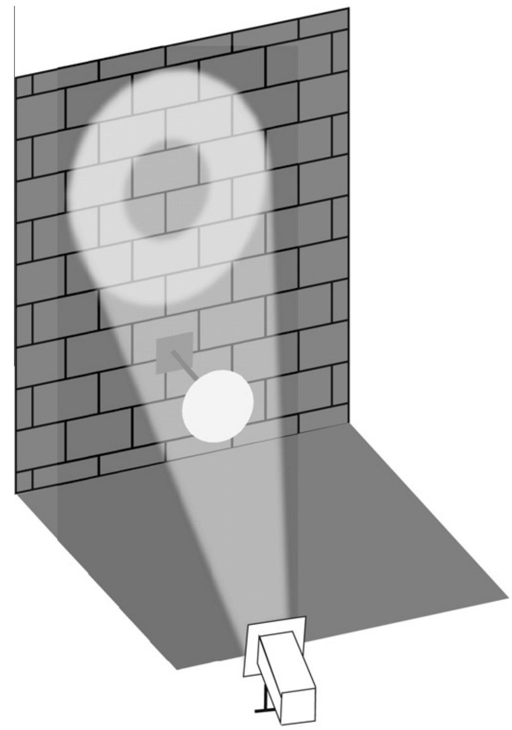

Fig. 4. Schematic of demonstration of Gelb Effect. The disk protruding from the wall is painted black, but appears light-grey or white when illuminated from the theatre lamp in the way shown.

the disk views a strongly lit disk set against an unlit background. Although the disk is painted black it appears light-grey, or even white if the lamp is powerful enough and the disk sufficiently smooth. Even when one realizes that the whiteness of the disk is an illusion caused by the light from the lamp, it is hard to shake the impression that it is…white. The critical physical property underpinning the illusion is that a discontinuity in illumination occurs precisely at the border of the disk with its background, such that the change in illumination is perfectly correlated with the change in reflectance. In the absence of the normal visual cues for decoupling illumination edges from reflectance edges, lightness is computed primarily from the contrast of the disk relative to its background. The role of contrast in producing the illusion is pressed home further when a large white card is placed directly behind the disk (and hence also illuminated by the stage lamp). Now the contrast reverses and the disk appears dark grey or black.

Another influential class of lightness error is assimilation. Assimilation is the opposite of contrast, in that lightness shifts *towards* rather than away from that of the immediate surround. One of the most celebrated and intensely researched phenomenon, which in its most popular form may be considered an example of assimilation, is White's Effect (White, 1979), shown in Fig. 5. The two sets of grey bars have the same luminance but differ markedly in lightness. The lightnesses of the test bars are shifted towards those of the flanking phases of the grating, which abut more with the test bars than the coaxial phases. However, White's Effect should not be regarded as a general case of assimilation; the same direction of illusion occurs with test bars that are horizontally elongated and as a result abut more with the coaxial than the flanking phases of the grating, thus making the effect an instance of contrast (Kingdom & Moulden, 1991). Indeed, the fact that White's effect can switch from assimilation to contrast depending on spatial configuration shows that 'contrast' and 'assimilation' are best regarded as operational short-hands for the direction of lightness errors not signatures for different underlying mechanisms.

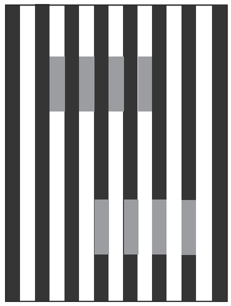

Fig. 5. White's Effect. The two sets of grey bars are the same luminance. Based on the original stimulus in White (1979).

1.3. Brightness

Brightness, the perceptual correlate of perceived luminance, could in principle be derived from the sensory measurement of a single image point. However, it has been understood for over two millenia that brightness, like lightness, depends on context, and specifically on contrast. The similarity between brightness and lightness under many circumstances is exemplified by the simultaneous contrast display in Fig. 3, where the percepts of brightness and lightness are synonymous. In fact the literature refers to Fig. 3 either as either simultaneous brightness contrast, simultaneous lightness contrast or plain simultaneous contrast.

Brightness and lightness assume very different qualia however when there are visible illumination variations in the scene, as in Fig. 2. The surfaces of the walls in the photograph appear uniformly white (a lightness judgement), yet are brighter in some places than others (a brightness judgement) due to the presence of shading and shadows.

As with lightness perception, our understanding of brightness perception over the past 25 years has developed primarily through an exhaustive examination of its errors. The relationship between brightness, lightness and perceived non-uniform illumination is discussed in Section 2, and again in Section 4.5.

1.4. Transparency

Research on perceptual transparency during the past quarter century has focused on two issues: the conditions that give rise to the impression of transparency and the perceptual dimensions of transparency. The first issue is discussed in Section 8, the second Section 9. Most research on the latter issue has focused on understanding the relationship between the perception of transparency and the physical dimensions of transmittance and reflectance. Transparent media have also figured prominently in studies of brightness and lightness perception; some of the most compelling demonstrations of brightness errors have been revealed when surfaces are overlaid by simulated transparencies, as discussed in Section 4.5.

Now to examine each of the eight questions posed in this review.

2. Question 1: What is the relationship between brightness, lightness and perceived illumination?

The study of LBT over the past 25 years has been profoundly influenced by the growing appreciation of the importance of the distinction between three percepts: brightness (perceived luminance), lightness (perceived reflectance) and perceived illumination. Yet in spite of the apparent ease with which these percepts can be separated in natural scenes such as the photograph in Fig. 2, the importance of the distinction between them has been slow in coming. Notwithstanding important early work by Katz (1935), and studies by Beck (1959), Arend and Goldstein (1987) and Schirillo, Reeves, and Arend (1990), it took two seminal studies by Arend and Spehar (1993a, 1993b)to definitively change the situation. Using Mondrian-like displays in which test patches were compared under visibly different illuminations, Arend and Spehar found that subjects made different judgements depending on whether they were required to compare lightness, brightness or perceived contrast, revealing unequivocally that there were circumstances in which these three percepts were separable dimensions of achromatic experience.

Experiments such as Arend and Spehar's, which revealed that lightness is perceptually distinct from brightness when illumination is visibly non-uniform, have been pivotal in persuading some researchers that brightness is the primary and lightness the secondary sensation (Blakeslee & McCourt, 2003; Blakeslee, Reetz, & McCourt, 2008). According to Blakeslee et al. (2008), when there is no visible illumination component, lightness follows directly from brightness; however with a visible illumination component, lightness comparisons across differently illuminated regions are inferential. The idea that lightness is either direct or inferential depending on whether there are visible illumination regions parallels the "phenomenal" versus "projective" distinction advanced by other authors for surface colour (i.e. chromatic) perception (Reeves, Amano, & Foster, 2008). Blakeslee et al. (2008) argue that lightness inferences may be easy or difficult depending on how much contextual information is present; a difficult case would be when two comparison patches are completely isolated within their own illumination regions. This view of the relationship between lightness, brightness and perceived non-uniform illumination is schematized in Fig. 6a.

Blakeslee and McCourt's arguments raise two important questions: 1. does brightness precede, or follow layer decomposition and 2. what does it mean to say that lightness judgements are in some circumstances inferential? Taking the first question, an immediate rejoinder is the mounting body of evidence that brightness *depends* on perceived illumination or on transparency. For example, shadows appear brighter than their equal-in-luminance reflectance counterparts (Logvinenko, 2005), light sources appear brighter than their equal-in-luminance reflectance counterparts (Agostini & Galmonte, 2002; Correani, Scott-Samuel, & Leonard, 2006; Gori & Stubbs, 2006; Zavagno & Caputo, 2001), and surfaces appearing to lie in shadow or behind dark transparencies appear brighter than if presented on equivalent reflectance backgrounds(Adelson, 1993, 1997; Kingdom, Blakeslee, & McCourt, 1997; Logvinenko, 1999). Although the interpretation of some of these findings is controversial (see Section 4.5), they nevertheless allude to the idea that brightness *follows* rather than precedes layer decomposition, as schematized in Fig. 6b (see also Gilchrist, 2006, pp. 338–339). This scheme also accords with recent ideas on the influence of layer decomposition on perceived chromaticity in chromatic-contrast displays (Ekroll, Faul, & Niederee, 2004). Remember however that for both schemes in Fig. 6 brightness and lightness are synonymous when there are no visible discontinuities in illumination.

Are lightness judgements across illumination boundaries inferential, as Blakeslee and McCourt suggest? It depends what one means by inferential. If one means that the visual system draws upon inbuilt knowledge of the spatio-temporal relations between reflectance and illumination in the natural visual world, then few would disagree. Lightness inferences in this sense are akin, for example, to the way perceived shape-from-shading is influenced by the lighting-comes-from-above assumption (Ramachandran, 1988); in other words lightness inferences are examples of what Helmholtz famously termed "unconscious inferences". However, if by inferential one means that lightness judgements across

Please cite this article in press as: Kingdom, F. A. A. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* (2010), doi:10.1016/j.visres.2010.09.012

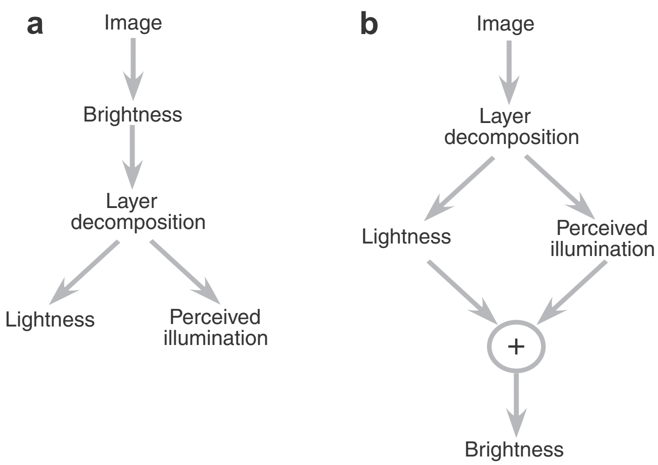

Fig. 6. Alternative schemes for the relationship between lightness, brightness and perceived illumination.

illumination boundaries are cognitive and effortful, then this would seem contrary to our everyday visual experience. The qualia of surface albedo are no less compelling in a world with shadows than in one without, as the photograph in Fig. 2 attests. Indeed, it would be odd for the visual system to have developed a smooth and effortless mechanism for estimating lightness in the context of variations in the ambient level of illumination, but a clunky and effortful one in the context of shadows and shading. Does it not in fact make sense for the visual system to prioritize those computations that are functionally important, such as surface lightness and non-uniform illumination, rather than brightness?

It is possible that both schemes in Fig. 6 describe what happens in vision, but under different circumstances. Perhaps the primacy of lightness over brightness is a characteristic of our visual experience with highly articulated scenes such as Fig. 2, whereas the reverse situation pertains to simple laboratory stimuli.

2.1. Summary

During the past quarter century numerous studies have attested to the importance of the distinction between three percepts – brightness, lightness and perceived illumination – as well as to the influence of perceived non-uniform illumination on brightness and lightness. However, the precise relationship between the three percepts and their relationship to the process of layer decomposition is still unresolved.

3. Question 2: Where is lightness and brightness encoded in the brain?

Prior to the period dealt with in this review, the prevailing wisdom was that lightness was encoded in the retina. For example, Cornsweet's (1970) influential book on vision championed the idea that reciprocal interactions among retinal neurons were responsible for simultaneous brightness contrast, and by implication lightness. During the past quarter-century however, opinion has shifted in favour of a cortical locus for lightness perception. One must bear in mind however that even if lightness values are 'read-out' at a cortical level, the retina plays a crucial role in lightness perception. Light adaptation, the process whereby the visual system adjusts to the local average light level is universally believed to be the result of rapid gain changes among *retinal* neurons. Thus it is the retina that normalizes luminance differences to the local light level, converting those differences into contrasts (or 'ratios'), a critical computation for achieving lightness constancy (Shapley & Enroth-Cugell, 1984; Walraven, Enroth-Cugell, Hood, MacLeod, & Schnapf, 1990). Second, the retina possesses neurones, such as ganglion cells, that encode those contrasts. A pre-cortical involvement in brightness-from-contrast processing is evidenced by the 'haploscopic superimposed-background display' shown in Fig. 7 (Whittle, 1994a). When the two patterns in the figure are free-fused, the two equal-in-luminance patches are seen on a common background yet appear very different in brightness, suggesting that the lightnesses of the patches are encoded before the point of binocular combination, i.e. pre-cortically. Using the haploscopic display method, Shevell, Holliday, and Whittle (1992) confirmed a pre-cortical involvement of local contrast in processing lightness perception, but went on to show that the influence of remote context on lightness occurred *after* the point of binocular combination, i.e. was cortical.

Besides the study of Shevell et al. (1992), the challenge to a *solely* pre-cortical locus for lightness perception has emerged from a number of quarters. One of these are the studies showing how depth relations influence lightness perception; for if depth perception is cortical and depth relations affect lightness, lightness perception must be at least partly cortical in origin. Gilchrist (1977, 1980) was the first to demonstrate that the lightness of a patch depended not only on the local surround (as seen retinally) but on the context that is specific to the patch's perceived depth plane. He employed a scaled-down model of two adjoining rooms, with the near room dimly lit and containing a black patch, and the far room brightly lit and containing a white patch. Subjects judged the lightness of a grey test patch that even though fixed in position, could be made to appear either in the near or far room depending on its interposition relationship with the white and black patches. Gilchrist found that when the test patch appeared to be coplanar with the near room (containing the black patch), subjects judged the test patch as white, but when it appeared to be coplanar with the far room (containing the white patch) the test patch appeared dark grey. This led Gilchrist to formulate his "coplanar ratio principle", which states that lightness is computed from luminance ratios *within the perceived depth plane* of the surface rather than in the retinal image.

Schirillo et al. (1990) repeated Gilchrist's experiment and replicated the main result, though they found that the effects were weaker than those reported by Gilchrist. Howe (2006) has argued from other evidence that it is the perceived illumination not local contrast within each depth plane that is critical, and Dalby, Saillant, and Wooten (1995) have shown that the coplanar ratio principle does not hold if the luminance range is small. Zaidi, Spehar, and Shy (1997) found that the effects predicted by the coplanar ratio hypothesis were reversible when the configuration induced strong illusory lightness differences in the opposite direction – they achieved the reversal using a three-dimensional version of White's Effect. However, Zaidi et al.'s finding is less evidence

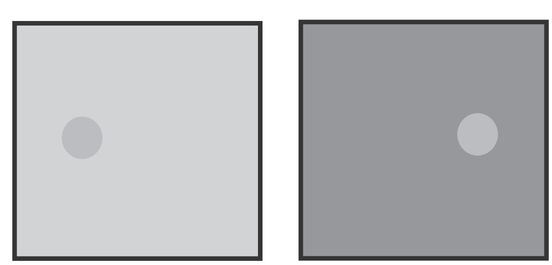

Fig. 7. In the haploscopic superimposed-background display, the two patterns are dichoptically superimposed such that the two circular patches are seen next to one another on the same background. This can be achieved here by free-fusion, and observers may notice that when free-fused the two equal-in-luminance patches appear more different in lightness than when the backgrounds are viewed side-by-side. Based on Whittle (1994a).

against the coplanar ratio principle as it is a demonstration that other contextual factors influence and in some circumstances counteract the effects of coplanar ratios. A recent summary of the research on the effects of depth and lightness can be found in Gilchrist (2008) pp. 120–122 and pp. 159–172.

The influence of depth relations on lightness has been further corroborated by studies showing that changes to the perceived three-dimensional structure of a stimulus may cause surfaces to switch in appearance from being different in reflectance to different in illumination (Buckley, Frisby, & Freeman, 1994; Knill & Kersten, 1991; Logvinenko & Menshikova, 1994). All the above studies are therefore consistent with a cortical locus for lightness perception.

What of physiology and brain-imaging? It has been argued that evidence for a cortical involvement in lightness perception is the presence of cortical cells or brain areas that respond to changes in surface lightness but that are unresponsive either to the borders of the surfaces or to changes in illumination. For example, Roe, Lu, and Hung (2005) reported cells in monkey V2 that responded to the purely illusory lightness variations in the Cornsweet illusion, and MacEvoy and Paradiso (2001) described cells in cat V1 whose responses were modulated by interactions from outside the classical receptive field in a manner that rendered them immune to changes in illumination. Human fMRI studies report strong responses to changes in surface lightness but not surface borders in retinotopic areas V1, V2 and V3 (Boyaci, Fang, Murray, & Kersten, 2007; Haynes, Lotto, & Rees, 2004).

When considering the evidence for "surface" as opposed to "border" neurons it is worth bearing in mind that neurons sensitive to low spatial frequencies and that have even-symmetric receptive fields would be expected to respond optimally when centred on a surface rather than on a surface's border, so the aforementioned studies are compatible with the multi-scale filtering accounts of lightness perception discussed later. However, studies showing evidence for surface-lightness-sensitive neurons that are nevertheless insensitive to illumination are not so easily compatible with such accounts.

What of the higher cortical areas? Zeki (1983a, 1983b) has reported that area V4 of macaque contains cells that respond to the perceived colour (i.e. chromaticity) of a surface rather than to its spectral (i.e. wavelength) composition, in other words that exhibit colour constancy (see also Schein & Desimone, 1990). Moreover, damage to human cortical area V4 (see Wandell & Wade, 2003, for a discussion of the relationship between macaque and human V4) can result in a failure of constancy in a colour-naming task (Zeki, Aglioti, McKeefry, & Berlucchi, 1999). One might reasonably assume that lightness constancy and colour constancy would be instantiated in the same brain region, but there does not seem to be strong evidence that V4 is involved in lightness constancy. Taken together however, the above evidence suggests that a number of brain areas in the ventral cortical stream are involved in lightness perception.

Is brightness also cortical, as implied by models based on multi-scale filtering combined with contrast normalization (see Section 4.4)? A cortical origin for brightness perception is consistent with demonstrations showing the influence of depth relations on brightness and not just lightness (Adelson, 1993, 2000). It is worth noting however that in Schirillo et al.s (1990) study, one of the only studies on depth relations in which subjects were explicitly required to make both lightness and brightness judgements, there was no effect of depth on brightness.

3.1. Summary

Lightness and brightness perception is a multi-stage process involving both retina and cortex. While the retina normalizes luminance variations to the local average light level and encodes rudimentary contrast information, neurons in the ventral cortical stream explicitly signal lightness and brightness.

4. Question 3. Is there a general theory for the effects of context on brightness and lightness?

This, of course, is the big question. The demonstrations of simultaneous contrast, the Gelb-effect, White's effect, and a multitude of other phenomena amassed during the previous quarter century have revealed the profound influence that context plays on lightness and brightness. However the devil lies with the details and it is with the details that the approaches reviewed in this section are concerned. We now critically examine the various models that have attempted to account for how the particular spatio-luminance structure of the image influences the brightness and lightness of test regions within. These models reveal more than anything else the fault-lines that demarcate the study of LBT. The models reviewed are: edge-integration, Gestalt-anchoring, spatial-filtering feature, multi-scale filtering with contrast normalization, intrinsic-image and empirical.

4.1. Edge-integration models

Land and McCann (1971) first championed the idea that lightness was computed by integrating local edge contrasts across the image. Their Retinex algorithm was designed to recover the lightnesses of Mondrian-like patterns that were subject to gradual illumination gradients such as from shading. Lightness computation was achieved by a four stage process: (1) the detection of edges via differentiation of the image; (2) thresholding to remove any gradual discontinuities; (3) integration to combine edge contrasts across space in order to establish a scale of relative lightness values; and (4) anchoring of relative lightness values to convert them to absolute ones. The Retinex, as well as the models it spawned (e.g. Hurlbert & Poggio, 1988; Land, 1986), would fail however with the photograph in Fig. 2, because any sharp-bordered shadows would be incorrectly identified as reflectance changes. The Retinex would also fail to predict the illusory lightness differences in the Gelb effect, simultaneous contrast and White's Effect, designed as it is to generate an approximately veridical representation of lightness.

Whittle (1994b), more than anyone else at the time, understood the contradiction between the role of contrast in producing errors such as simultaneous contrast and the role of edge integration (or some process analogous to it) in providing a veridical representation of lightness. He noticed that in the haploscopic superimposed-background display shown in Fig. 7, simultaneous contrast was considerably enhanced, reasoning that this was because the effects of edge integration that normally occurred when the backgrounds were seen side-by-side were being bypassed. Whittle proposed that are two types of constancy: Type I and Type II. Type I is constancy with respect to the ambient level, is achieved by computing contrast, and produces errors such as simultaneous contrast. Type II is constancy with respect to the changing background, and is achieved by integrating contrasts across the image. Type I and Type II constancy thus tend to work in opposite directions, with Type II serving to mitigate the errors resulting from Type I in order to provide a more veridical representation. In keeping with Whittle's ideas, the study by Shevell et al. (1992) mentioned earlier provided empirical support for the Type I versus Type II distinction by showing that 'contrast' processes were precortical whereas 'surround-integrative' processes were cortical.

More recently, Rudd and colleagues (Rudd, 2001, 2003; Rudd & Arrington, 2001; Rudd & Popa, 2007; Rudd & Zemach 2004, 2005,

2007) have developed the idea of edge-integration into a fully-fledged model of lightness perception, one capable of quantitatively accounting for both contrast, assimilation and edge-integrative phenomena. By doing so they have shifted the focus of the edge-integration approach away from its original goal of modelling veridical lightness perception to towards an entirely new role of modelling its errors. Critical to the development of Rudd and colleagues edge-integration models is the finding that with patch-ring-surround stimuli (Shapley & Reid, 1985; see Fig. 8), not only contrast but assimilation may be observed between the ring and test patch, especially when the ring is relatively thin. Rudd and colleagues argue that this assimilation is a consequence of a cortical contrast gain control mechanism acting between neighbouring edge detectors, with the gain of each edge detector being a positive function of the response magnitudes of neighbouring edge detectors, but a negative function of the distances between them. Thus if the spatial span of the gain control bridges the distance between the inner and outer edges of the ring (for example when the ring is relatively thin), assimilation is predicted.

Rudd et al. have successfully dove-tailed a well-established physiological mechanism, contrast gain control, to a model of edge-integration, and as a result have been able to quantitatively model assimilation, contrast and edge-integration data. However the edge-integration approach, as it stands, has profound limitations. The language of edge integration is primarily one of stimulus properties: "edges", "contrasts", "areas", "widths" etc. As such, theories of edge-integration are restricted to a Mondrian-world, and a simplified one of radially-symmetric stimuli at that. Why? The defining computation of all edge-integration models is that edge signals are integrated across space to generate a map of lightness values. For simple stimuli this is relatively straightforward, but with complex two-dimensional images the process is computationally expensive and, one cannot help feel, physiologically implausible. Rudd and colleagues are keenly aware of this limitation of the edge-integration approach and anticipate incorporating multi-scale filtering into their model (e.g. see discussion in Rudd & Zemach, 2005). This may make the approach more amenable to dealing with complex stimuli such as natural scenes, replete as they are with luminance gradients, textured surfaces and other forms of structural complexity.

4.2. Gestalt-anchoring models

'Anchoring' refers to the process that converts relative lightness values into absolute ones, by associating a particular lightness (e.g. mid-grey, white) to a particular luminance (mean, maximum). Anchoring as such is discussed in Section 7. "Anchoring theory" on the other hand refers to a theory of lightness perception in which multiple anchors exist within nested image frameworks, with all anchors contributing to the computation of lightness. Because the frameworks in anchoring theory are determined by Gestalt grouping principles, the models are termed here 'Gestalt-anchoring' models. The best-known is that of Gilchrist and colleagues (Gilchrist, 2006; Gilchrist et al., 1999). Bressan (2006a, 2006b, 2007) has advanced a version of Gestalt-anchoring model based on different anchoring principles and modified putative frameworks.

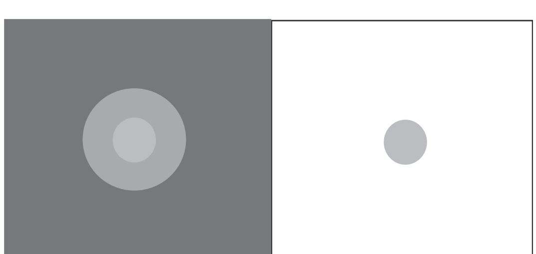

Fig. 8. Left: patch-ring-surround stimulus. Typically the subject adjusts the luminance of the patch on the right to match the central test patch on the left.

While the full details of these models are beyond the remit of this review, the principle of the best-known, that of Gilchrist and colleagues, can be grasped by considering how it applies to the simultaneous contrast display in Fig. 3. The idea is that the stimulus is divided into two local and one global perceptual framework, with the lightness of each test patch computed as a weighted average of two lightness values, one derived from the local, the other from the global framework. The global framework is comprised of the stimulus as a whole, whereas the local frameworks consist of the two surrounds each with their respective test patch. Within each framework the anchor is the highest luminance, which is assigned white, and all other regions within the framework are assigned greys according to their luminance ratio with respect to the white anchor. Within the global framework, the white background of the stimulus as a whole is the highest luminance and therefore assigned white, and the two test patches are assigned identical lower lightnesses, since their luminance ratios with respect to the background are the same. However, for the two local frameworks, the patch lightness assignments are different. For the patch on the dark surround, the patch is the highest luminance and therefore assigned white. For the patch on the light surround, the surround is the highest luminance and therefore assigned white, while the patch is assigned a value relative to it, i.e. mid-grey. The net patch lightnesses are the averages of the global framework values (both patches equal and mid-grey) and the local framework values (patch on dark surround white, patch on light surround mid-grey). The result is a higher lightness value for the patch on the dark compared to the light surround. Thus although the global framework has an important influence on the lightnesses of the patches, it is the two local frameworks that cause the difference in patch lightness and hence the illusion.

Although the model makes successful predictions with relatively simple stimuli, there are some notable failures. One is the prediction that if the patches in the simultaneous contrast display are both equal-in-luminance increments, they should appear equal in lightness (two equal-in-luminance decrements on the other hand are predicted to have different lightnesses). This prediction follows from the rule that the two patch lightnesses are computed only with respect to the white background of the stimulus; the local frameworks do not contribute to patch lightness because they are each of lower luminance than the patches they surround. Double increments however *do* differ in lightness (Blakeslee, Reetz, & McCourt, 2009; Bressan, 2006a, 2006b; Bressan & Actis-Grosso, 2001; Rudd & Zemach, 2005; Whittle, 1994b), as the top panel of Fig. 9 shows, though in general by not as much as double decrements (Whittle, 1994b) (bottom panel).

Bressan's (2006a, 2006b, 2007) 'double-anchoring' model successfully predicts double-increment lightness differences by assuming that there are two anchors per framework: highest-luminance-is-white and surround-luminance-is-white. One cannot help however but question the plausibility of the idea that, for example, a black surround can take on the contributory role of a white anchor.

A problem generic to *all* Gestalt-anchoring models is their failure to adequately deal with gradients in brightness/lightness, such as the form of simultaneous contrast known as grating-induction (Blakeslee & McCourt, 1997; McCourt, 1982), shown here in Fig. 10. The challenge of grating induction for Gestalt-anchoring models is in dividing the stimulus into the necessary frameworks in order to predict the continuous lightness variations along the test stripe. Although Bressan's double-anchoring model correctly

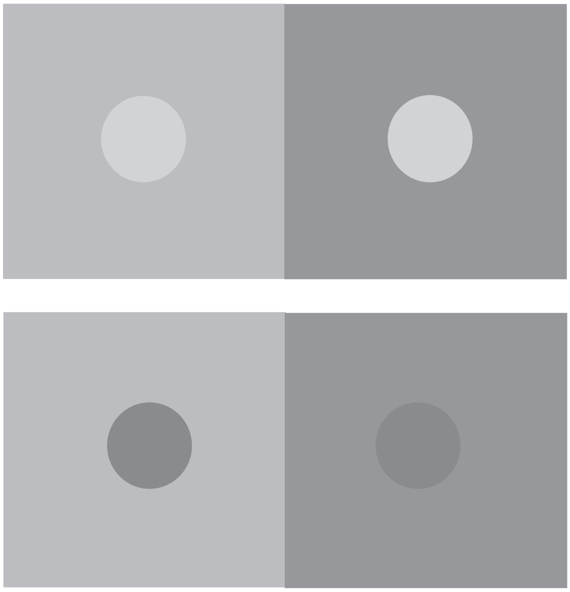

Fig. 9. Top: double increment; bottom: double-decrement simultaneous contrast displays. The two patches in each stimulus are equal-in-luminance.

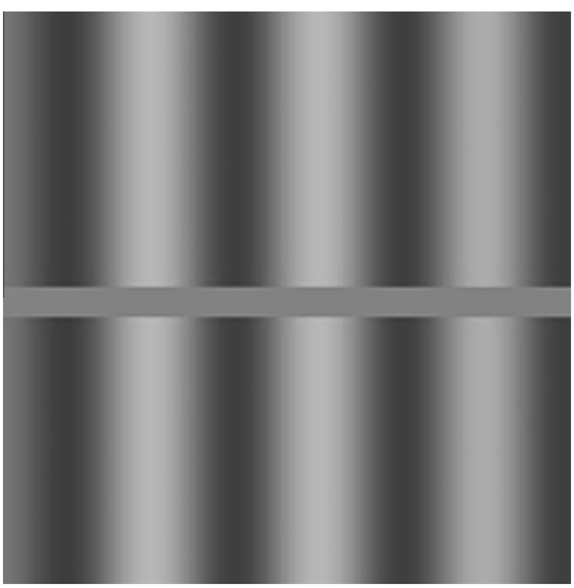

Fig. 10. Grating induction. The horizontal stripe is uniform in luminance but appears modulated in brightness out-of-phase with the inducer grating. Based on McCourt (1982).

predicts that with a square-wave grating inducer, the lightnesses of the test stripe are different depending on whether they fall on the dark or light phase of the grating (Bressan, 2007), it is hard to see how the model predicts the *continuous variation* in lightness along the test stripe, which for the sinusoidal inducer in Fig. 10 is itself near-sinusoidal. The model presumably requires a unique confluence of anchor(s) and framework(s) to predict each of the many different lightness values along the test stripe, and the specification of the confluences is inevitably ad hoc. Contrast this with the relative parsimony with which grating induction is explained by multi-scale filtering (Blakeslee & McCourt, 2005; Moulden & Kingdom, 1991). Note also that if the test stripe is placed into a completely separate "framework" from the inducing grating by presenting it in a different stereoscopic depth plane, grating induction is still observed (Kingdom, 2003b). In short, grating induction highlights a critical problem with Gestalt-anchoring models: specifying the frameworks (see also Howe et al.'s (2007) critique of Bressan's double-anchoring model, but also reply by Bressan (2007)). For anything other than the simplest of Mondrian worlds, specifying the frameworks is always going to be difficult, and for really complex stimuli such as natural scenes, it is hard to see how one could ever begin. Therefore, as with edge-integration models, Gestalt-anchoring models do not at present have the inherent flexibility for dealing with the relatively complex world of our everyday visual experience (as noted also by Corney & Lotto, 2007).

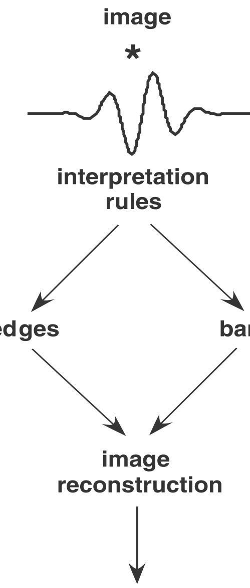

Fig. 11. Flow diagram of principle components of spatial-filtering feature models.

4.3. Spatial-filtering feature models

Although the idea that the clues to brightness errors lie in the responses of spatial bandpass filters has its roots in Hering, it was Marr's (1982) theory of the primal sketch that inspired a class of model that gave Hering's ideas one of its most modern expressions. I am referring to the class of model that posits that the responses of spatial filters are interrogated by a set of rules in order to create a map of the salient features in the image, specifically a symbolic description of the image in terms of 'edges' and 'bars'. This symbolic description is then used to construct an internal image by integrating the edge signals and combining them with the bar signals, and it is within this internal image that brightness errors are found. Fig. 11 illustrates the principle. Although individual models bases upon the principle differ in terms of the details of filtering and the degree to which they are concerned with brightness errors as opposed to spatial vision in general, all have something to say about illusory brightness phenomena. The three main protagonists are MIRAGE (Watt & Morgan, 1985), the Local Energy Model (Morrone & Burr, 1988) and MIDAAS (Kingdom & Moulden, 1992), the last of which is aimed specifically at modelling brightness phenomena.

Briefly, MIRAGE proposes that the outputs of centre-surround spatial filters at various spatial scales are (non-linearly) combined to produce a pattern of positive and negative response 'lobes'. The interpretation rules posit that any response lobe flanked either side by two opposite-polarity response lobes indicates the presence of a bar (whose width is proportional to the lobe's standard deviation), while an abutting pair of positive and negative response lobes indicates the presence of an edge (whose polarity is determined by the ordering of the two lobes and whose slope or 'blur' is proportional to the lobes' standard deviations). Although formulated primarily to account for data obtained from positional acuity and edge blur discrimination experiments, Watt and Morgan (1985) and Watt (1988) show that MIRAGE predicts illusory brightness phenomena such as Mach bands – these are the illusory bars at the 'foot' and 'knee' of a trapezoid – as well as the Chevreull illusion, the phenomenon in which a staircase in luminance appears as a triangular wave in brightness.

MIDAAS (Kingdom & Moulden, 1992) draws upon the same implementation rules as MIRAGE but applies them to the individual filter rather than combined filter responses, and then averages the symbolic edge-bar description across the different filter scales. One of the cases for this model variant is that it explains why a trapezoid can simultaneously appear trapezoidal *and* have Mach bands; MIRAGE predicts that a trapezoid appears *either* as a blurred edge *or* as a uniform region demarcated by two Mach bands (depending on the trapezoid's slope), but not both at the same time, unlike what is usually perceived.

The Local Energy model (Morrone & Burr, 1988) locates edges and bars as peaks in 'local energy', defined as the square root of the sums of squares of the responses of odd and even-symmetric spatial filters. The nature of the feature (edge or bar) is obtained by evaluating the relative responses of the even and odd detectors at the local energy location. If the peak in local energy coincides with that of the even-symmetric filter, the feature is a bar, whereas if it coincides with the peak of the odd-symmetric filter, the feature is an edge. The Local Energy model gives a good account of both the presence and magnitude of Mach bands (Morrone, Ross, Burr, & Owens, 1986; Ross, Morrone, & Burr, 1989), as well as other illusions such as the Craik–Cornsweet–O'Brien effect (Burr, 1987; and see Fig. 19).

Although spatial-filtering feature models are undoubtedly pertinent to some aspects of vision, they have arguably failed to provide a plausible and lasting account of brightness errors. There are a number of reasons for this. The first is that evidence has failed to accrue that the early visual system classifies image discontinuities into edges and bars, the defining and unifying theme of these models (see review of the literature in Huang, Kingdom, & Hess, 2006). Although Burr, Morrone, and Spinelli (1989) found apparently good evidence for dedicated edge and bar detectors in human vision, others have failed to replicate their results and have instead proposed that dedicated detectors only exist for coding increments and decrements, i.e. for the different polarities of even-symmetric stimuli, such as bars (Huang et al., 2006) (and see Section 5.2). The second problem with these models is that as with edge-integration models, edge signals must be integrated across space to generate the internal image, which as was noted earlier is for two-dimensional images computationally very expensive and probably physiologically implausible.

The Local Energy model runs into particular difficulties as a model of brightness perception because it predicts that a sinusoidal grating is featureless, since $(\sin^2 + \cos^2) = 1$ everywhere in the stimulus (Georgeson, 1994; Hesse & Georgeson, 2005; Kingdom & Moulden, 1992; Meese & Georgeson, 2005). A sine-wave grating is a visible modulation in brightness that would be expected to be captured by any model of brightness perception. Although it might be argued that the luminance non-linearities in early vision sufficiently distort a sinusoidal grating to make its features visible to a local energy transformation, this leads to the implausible prediction that if the non-linearities were nulled, for example by multiplying the sine-wave luminance profile by the appropriate function, the bars should disappear (Meese, 1999; Meese & Georgeson, 2005). Furthermore, at contrasts at which the distorting effects of a luminance non-linearity are probably negligible (<5%), the light and dark bars of sine-wave gratings are visible across a range of spatial frequencies.

The final reason why spatial-filtering feature models have failed in their bid to account for brightness phenomena is that they have been superseded by another class of spatial-filtering model that is more physiologically realistic, easier to implement in two dimensions and which accounts for a wider range of phenomena. It is to these models that we now turn.

4.4. Spatial-filtering models involving contrast normalization

The other class of filtering model anticipated by Hering involves the combination of multi-scale spatial filtering and contrast normalization. Unlike the models reviewed in the previous section, the models discussed here do not require rules to interpret the outputs of spatial filters in order to generate an 'output'. Rather, the filter responses themselves constitute the output. The best-known of model of this class is the ODOG (Oriented Difference-of-Gaussian) model of Blakeslee, McCourt and colleagues (Blakeslee & McCourt, 1999, 2001, 2003, 2004, 2005; Blakeslee, Pasieka, & McCourt, 2005). Dakin and Bex (2003) have argued for a similar approach but with a different model implementation, and have applied it to the well-known Craik–Cornsweet–O'Brien illusion (see Fig. 19). Their model is discussed again in Section 6.2.

Convolving an image with an array of narrowband, linear filters tuned to different spatial frequencies and orientations, followed by summation of outputs does not in itself predict brightness errors, because if the filters form a 'complete' set, the output is a more-or-less veridical copy of the original. In the Blakeslee and McCourt approach, two processes conspire to produce brightness errors. First, very low spatial frequencies are attenuated, and this accounts for a variety of contrast errors such as simultaneous contrast and grating induction (see also Shapiro, Knight, & Lu, 2008). The second is contrast normalization, which we saw earlier was a key feature of Rudd and colleagues' edge-integration model. In Blakeslee and McCourt's ODOG model, the contrast normalization stage equates the root-mean-square (RMS) output across the six orientation channels, each of which is a weighted linear sum of seven spatial-frequency channels. The contrast normalization stage is the key to assimilation phenomena such as White's Effect (Fig. 5). The most responsive filters to White's stimulus are the relatively high spatial frequency, vertically-oriented filters tuned to the inducer grating. Contrast normalization has the effect of reducing the responses of these filters relative to those sensitive to horizontal orientations. Since it is the horizontally-oriented filters that pool the luminances of the flanking bars with those of the test patches, the enhancement of their relative contribution is the cause of the illusion. In Dakin and Bex's (2003) model, the contrast normalization equates filter responses across spatial-frequency not orientation, and they show that it provides a good explanation for the Craik–Cornsweet–O'Brien illusion (Section 6.2).

The idea that brightness phenomena result from the attenuation of low spatial frequencies is not a new idea (e.g. see early review by Kingdom & Moulden, 1988), but the idea that contrast normalization is the key to assimilation and other brightness errors is arguably one of the most important developments in LBT research to have emerged in the past quarter century.

A particular strength of the ODOG model is that its parameters are fixed across all stimuli to which it is applied. Another strength is that it predicts quantitative data obtained using point-wise brightness matching of the stimulus (as in Fig. 12), enabling brightness gradients and not just uniform surfaces to be modelled.

"Contrast" theories of lightness and brightness perception, of which multi-scale filtering models are the most modern expression, have traditionally been criticized for failing to take into

Please cite this article in press as: Kingdom, F. A. A. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* (2010), doi:10.1016/j.visres.2010.09.012

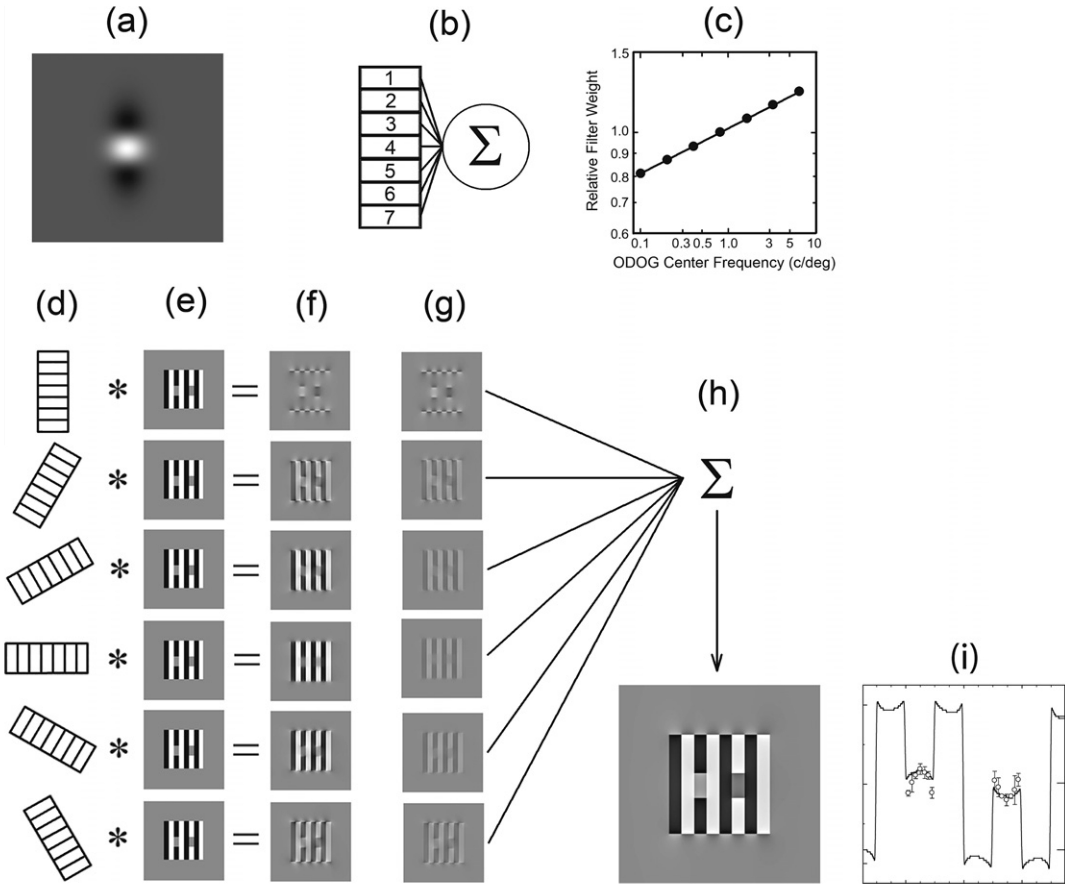

Fig. 12. ODOG model applied to White's Effect. (a) filter profile, (b) filters of different spatial scales are summed at a given orientation, (c) filter gains as a function of centre spatial frequency, (d) orientations of combined filters (e) stimulus, (f) result of convolving each stimulus with d, (g) contrast normalization equates RMSs of filter outputs (h) outputs are summed across filter orientation, (i) the continuous line is the predicted brightness of a horizontal cross-section of the stimulus, while circles are the results of point-wise brightness matching. Figure supplied by B. Blakeslee and M. McCourt, based on Fig. 1 in Blakeslee and McCourt (1999).

account the effects of remote context, the assumption being that they only deal with contrast in the immediate vicinity of the edge. The models of Blakeslee and McCourt and Dakin and Bex however should hopefully bury this red herring once and for all: in these models remote context makes its impact through the coarse-scale filters. Moreover, many of the phenomena traditionally explained by 'edge-integration' (e.g. Fig. 7) may similarly be accounted for by coarse-scale filtering.

Are there shortcomings to the current generation of models that combine multi-scale-filtering with contrast normalization? A general criticism is that they fail to deal with the effects of illumination/transparency on brightness perception. As we shall see however, there is at present no mechanistic model of such effects, so it is perhaps not surprising that models such as ODOG do not deal with them. Another shortcoming, though less problematic in the long run, is that at present multi-scale filtering models do not account for anisotropies in the brightnesses of increments and decrements (Corney & Lotto, 2007). However, as will be argued in Section 5.1, this could easily be remedied by applying a suitable non-linearity to the filter responses to embody the effects of local light adaptation.

One other concern with ODOG is choice of filter. The ODOG response profile does not match anything known physiologically. There is no reason however to suppose that the use of more physiologically-realistic filters such as ones modelled on cortical simple cells should not work equally well.

Finally, however appealing is the idea that contrast normalization is responsible for many brightness errors, there is at present little actual evidence for it. Experiments that manipulate the amount of contrast normalization, perhaps via adaptation or masking, in order to test whether the predicted changes in the magnitude and direction of brightness errors occur, would be welcome.

4.5. Intrinsic-image models

For the most part, intrinsic-image 'models' are not really models at all, but rather compilations of demonstrations showing the influence of non-uniform illumination and transparency on surface lightness and brightness (e.g. Adelson, 1993, 2000; Adelson & Pentland, 1996; Albert, 2006; Arend, 1994; Bergström, 1994, Chap. 6; Logvinenko, 1999; Logvinenko, Adelson, Ross, & Somers, 2005). Some protagonist however go further and suggest that even in displays containing no obvious regions of non-uniform illumination or transparency, the visual system assumes that some form of illumination/transparency exists with profound consequences for perception (Anderson, 1997, 2001; Ekroll et al., 2004). This latter viewpoint should perhaps be termed the 'strong' form of intrinsic-image model, to differentiate it from the 'weak' form described above which is concerned only with visible regions of illumination/transparency.

In situations where there are visible regions of illumination or transparency, for which both weak and strong forms of intrinsic-image model apply, there are two issues. The first issue concerns the cues used by vision to identify the presence of non-uniform illumination and transparency, such as X-junctions, three-dimensional shape, motion, occlusion, and colour. This issue is dealt with

Please cite this article in press as: Kingdom, F. A. A. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* (2010), doi:10.1016/j.visres.2010.09.012

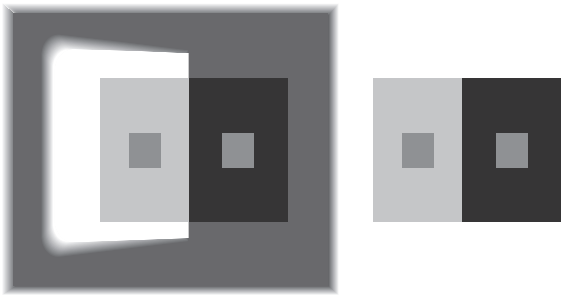

Fig. 13. Schematic of display used by Gilchrist et al. (1983). On the right is the standard simultaneous contrast display, while on the left one half of the display appears to be brightly illuminated. The luminances of the grey patches and their immediate surrounds are the same in both displays. Gilchrist used real rather than simulated illumination, and its effect on the magnitude of simultaneous contrast was reportedly much greater than seen here. The figure is taken from Fig. 6.21 in Gilchrist (2006) and supplied by the author.

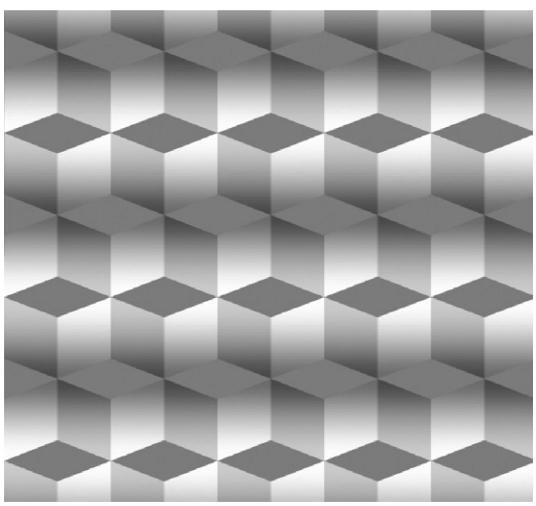

Fig. 14. Logvinenko's figure of a wall-of-blocks with shading. All rows of diamonds have the same luminance but alternating rows appear to differ dramatically in lightness and brightness. From Logvinenko (1999), supplied by the author.

in Section 8. The second issue concerns the effect of perceived illumination and transparency on lightness and brightness. Some of the contested ideas concerning the relationship between brightness, lightness and layer decomposition were discussed earlier in Section 2. Here we focus on the evidence that visible non-uniform illuminations and transparencies influences surface brightness and lightness.

The seminal studies by Gilchrist and colleagues (Gilchrist, Delman, & Jacobsen, 1983; see also Gilchrist, 1988) were unquestionably the watershed that precipitated the recent interest in intrinsic-image models, though the basic idea goes back to Helmholtz (see below). Gilchrist observed that the magnitude of simultaneous contrast was enhanced when the two grey patches, normally seen as surrounded by materials of different reflectance, were instead seen as lying in different illuminations, as illustrated in Fig. 13. The luminance contrasts between the patches and their immediate surrounds were kept the same under both configurations, so Gilchrist argued that the enhancement of simultaneous contrast could not be due to the effects of contrast but instead due to the perceived non-uniform illumination.

Capitalizing on the development of computer-graphics technology, a multitude of demonstrations soon followed. These showed how depictions of transparency (e.g. Adelson, 1993; Logvinenko et al., 2005), shading/shadows (Adelson, 2000; Adelson & Pentland, 1996; Knill & Kersten, 1991; Logvinenko, 1999) and figure-ground relationships (Anderson & Winawer, 2005) could profoundly influence brightness perception. One of the author's favourites is Logvinenko's (1999) wall-of-blocks-with-shading demonstration in Fig. 14; another celebrated example is Adelson's (2000) 'snake' figure in Fig. 15. The allure of these demonstrations is the sheer magnitude of their illusory brightness differences, which seem to far surpass those of standard simultaneous contrast displays. All appear to demonstrate that with depictions of non-uniform illumination or transparency, brightness shifts dramatically towards the 'true' lightnesses of the test regions 'beneath' the transparencies/shadows. So impressive are these brightness illusions that many in the vision community have come to eschew the importance of basic luminance relations for brightness perception in favour of the idea that intrinsic-image relations are all that are really important.

Impressive as these demonstrations are, caution must be exercised when interpreting them (Kingdom, 2003a; Todorovic, 2006). Consider the snake/anti-snake figure. The difference in brightness between the two rows of equiluminant diamonds is clearly much bigger in the snake than it is in the 'control', anti-snake figure. The argument goes that because the corresponding rows of diamonds in the two figures have the same luminance contrast with their immediate surrounds, the difference in illusion magnitude must be a result of the apparent transparency in the snake figure. The problem with this argument is that it is not the case that the luminance contrasts in the snake and anti-snake are equivalent, providing one accepts that luminance contrast is not just something that happens "at the edge", but involves the wider context. In the snake figure, the upper row of diamonds is surrounded by a larger area of black, and the lower row of diamonds by a larger area of white, than the corresponding diamonds

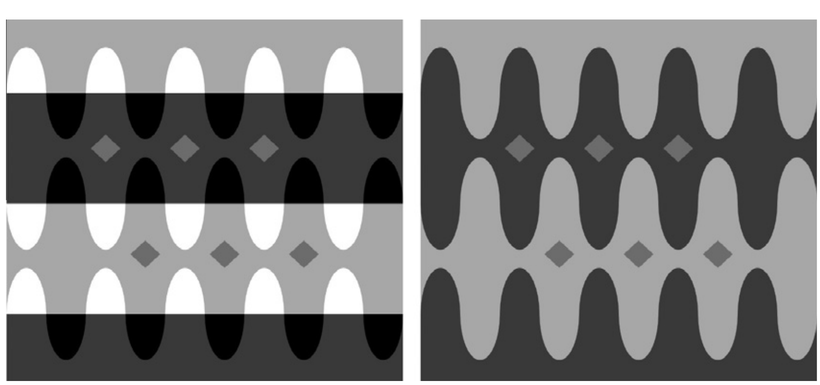

Fig. 15. Adelson's 'snake'(left) and 'anti-snake'(right) figures. The small diamonds in both figures all have the same luminance. The brightness difference between the upper and lower rows of diamonds is much bigger in the snake compared to the anti-snake. From Adelson (2000), supplied by the author.

in the anti-snake figure. Although it seems unlikely, it is hard to be certain that local luminance contrasts operating at multiple spatial scales are not the reason for the difference in size of illusion in the two figures (Kingdom, 2003a; Todorovic, 2006).

Few researchers have seriously attempted to tackle this problem. One exception is an experiment conducted by Kingdom et al. (1997). They used a stimulus designed to minimize the difference in luminance relations between the with- and without-transparency conditions. They found that the depicted transparency did have an impact on brightness, as in Adelson's and Logvinenko's figures, though in general the effects were quite small with the biggest effect being about a factor of two.

What of the strong form of intrinsic-image model? The roots of the strong form lie in (one of) Helmholtz's (1866/1962) explanations of simultaneous colour (i.e. chromatic) contrast (see Kingdom, 1997). Helmholtz opined that two identical-in-colour patches set against differently-coloured backgrounds appeared different in colour because it was assumed that they were bathed in differently-coloured illuminations; for if the patches are differently-illuminated yet identical in reflected colour, they must, by inference, have different surface colours. This explanation is underpinned by the idea that prior knowledge about the conditions in which non-uniform illumination normally occurs leads to assumptions about non-uniform illumination being present even if there are no strong cues.

The evidence in support of the strong form of intrinsic-image model includes those demonstrations described above that support the weak form of the model. For example, support comes from the observation that simultaneous contrast is increased when the two patch backgrounds appear as different illumination rather than reflectance regions (Gilchrist et al., 1983; Williams, McCoy, & Purves, 1998b). The argument goes that if simultaneous contrast is especially large when strong illumination cues are present, then inferred illumination must be the cause of simultaneous contrast even in the absence of those cues. The problem with this argument is that it imputes causality by association. Consider in this regard Anderson's (2001) account of White's Effect (Fig. 5). Anderson correctly notes that one obtains impressions of transparency and occlusion for the test bars placed on, respectively, the black and white phases of the grating, and concludes that the illusion must be a result of the effects of layer decomposition, or "scission". However, the illusion is pronounced for single test bars (e.g. Moulden and Kingdom, 1989) and occurs even when the T-junctions at the ends of the test bars (which are regarded as critical to the impressions of occlusion and transparency) are eliminated through the use of ellipsoid-shaped test bars (Yazdanbakhsh, Arabzadeh, Babadi, & Fazl, 2002). Moreover, provided the gratings are of relatively high spatial frequency, the same direction of illusion as White's Effect is observed in continous circular test bars embedded in circular gratings, i.e. also without T-junctions (Hong & Shevell, 2004a, 2004b). It is hard to convince oneself that in those versions of White's Effect and its relatives that do not elicit impressions of transparency and occlusion, transparency and occlusion are the cause.

A current limitation of all intrinsic-image models is that they offer no mechanistic explanation as to how layer decomposition is combined with luminance values to compute lightness.

Notwithstanding the above caveats, intrinsic-image models have provided a lush new terrain for exploring both the visual cues that facilitate layer decomposition and the impact of non-uniform illumination and transparency on our perception of brightness and lightness.

4.6. Empirical models

In their "empirical" approach to lightness perception, Purves and Lotto (2003) suggest that when organisms are confronted with the need to identify reflectances in the context of spatially non-uniform illumination, they estimate the most likely reflectance values based on the pattern of luminances observed together with their knowledge of image statistics learnt through goal-directed behaviour. Lightness illusions occur because in any given situation the most likely value of reflectance will often differ from its true value. For example, in the case of simultaneous contrast, Purves and Lotto argue that patches on dark backgrounds are more likely to be lying in shadow compared to patches on bright backgrounds, and hence that two equal-in-luminance patches, one on a dark the other on a bright background, will likely have a different reflectance, and that is how they are perceived (Purves & Lotto, 2003 – for details see Williams, McCoy, & Purves, 1998a). Other illusions such as Mach bands and the Craik–Cornsweet–O'Brien illusion (Fig. 19) are similarly explained: the illusory percepts match physical illumination patterns that often arise because of the way non-diffuse illumination is reflected from the surfaces of three-dimensional objects. Moreover, according to Nundy and Purves (2000), the scaling of brightness values (Section 5) is also explicable in terms of learned image statistics.

At first sight the empirical approach appears to echo that of intrinsic-image models: note the similarity between the empirical explanation of simultaneous contrast and that given by Helmholtz described above. However there is nominally at least an important difference. With intrinsic-image models the visual system decomposes the image into separate representations of illumination and reflectance prior to coding the lightnesses of the reflectance layer. In the empirical approach no process of layer-decomposition occurs. Rather, image statistics learnt through goal-directed behaviour are used to make lightness estimates. To illustrate the difference with the intrinsic-image approach and to demonstrate the operation of the empirical approach as it relates to lightness errors, Corney and Lotto (2007) trained an artificial back-propagation neural network to identify the reflectances of target surfaces in synthetic images. The synthetic images consisted of three-dimensional arrangements of multiple-sized reflectance patches subject to simulated non-uniform illumination. The only sense data available to the network were the luminances of the patches, so the network had to learn to identify surface reflectance in spite of the inherent ambiguity of patch luminance. Having learned to identify the target reflectances in the synthetic images to a criterion level of accuracy, the network was then given the task of estimating the reflectances of target patches in classic lightness-illusion displays such as simultaneous contrast, White's Effect and Mach bands. The network made very similar lightness errors to those reported by human observers.

Is the empirical approach really different from the mechanistic accounts of lightness errors discussed elsewhere in this section? It has been understood for some time that visual mechanisms have evolved to code the useful information in the visual environment world in an optimally efficient manner (Field, 1994; Geisler, 2008; Olshausen & Field, 2004). The pivotal idea in the empirical approach, namely that knowledge of image statistics learnt through goal-directed behaviour leads to predictable lightness errors, is not a far cry from the idea that visual mechanisms have evolved to encode the useful image statistics in the environments and that these mechanisms produce lightness errors. For example, the response of a cortical bandpass filter is designed to be largely invariant to the ambient level of illumination, so the output of such filters will, on average, more closely correlate with the pattern of image reflectances than with the pattern of image luminances. In other words cortical bandpass filters serve to reduce the range of lightness values from which the visual system has to choose. By the same token, a mechanism that discounts spatially-varying illumination via a process of layer decomposition also reduces the potential range of lightness choices. In short, the mechanisms

Please cite this article in press as: Kingdom, F. A. A. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* (2010), doi:10.1016/j.visres.2010.09.012

deployed by vision for coding lightness (bandpass spatial filtering, contrast gain control, layer decomposition, etc.) have been honed during evolution and/or development to maximize the probability of judging lightness correctly, but being imperfect they nevertheless make errors. The difference between the empirical and mechanistic approaches to lightness perception then seems to come down to the type of image statistics exploited by vision for judging lightness. For example the process of layer decomposition is believed to exploit relatively high-order statistical relations such as X-junctions and three-dimensional shape (Kingdom, 2008), whereas the neural networks employed in Corney and Lotto's study presumably captures relatively low-order image statistics. The exploitation of high-order image statistics is a defining property of our visual system, so we should not be surprised that these statistics are used by vision for judging lightness. Therefore, we will conclude that the main contribution of the empirical approach is in impressing upon the vision community the need to understand the role of image statistics in the visual system's computation of lightness.

4.7. Summary

Three promising theoretical developments of the past quarter century for understanding the role of context in lightness and brightness perception are 1. multi-scale filtering combined with contrast normalization, 2. the idea that the visual system performs intrinsic-image, or layer decomposition, and 3. that an understanding of image statistics is important to an understanding of lightness errors. Multi-scale filtering, when combined with contrast normalization, has the potential for dealing with complex stimuli such as images of natural scenes. An important refinement to this approach will be the incorporation of the effects of local light adaptation to model anisotropies in the brightness of increments and decrements (see Section 5.2). Intrinsic-image approaches have helped refine our understanding of the cues available to vision for segmenting the image into reflectance, illumination and transparency (see Section 8), and have demonstrated the importance of layer decomposition for brightness and lightness perception. However, intrinsic-image studies are at present compromised by inadequate controls for ruling out the effects of contrast. Finally, both multi-scale filtering and intrinsic-image approaches will continue to benefit from a deeper understanding of image statistics.

5. Question 4: What is the function relating brightness to contrast in simple patch-surround displays?

A quarter of a century ago there was no clear formulation of the relationship between brightness and contrast for the simple case of a patch on a uniform background. This situation has been remedied largely by Whittle (1986, 1992; summarized along with earlier relevant studies in Whittle, 1994a). Using the haploscopic superimposed-background display described in Section 3, Whittle measured the brightness of patches across a wide range of background luminances and increment contrasts, and, presciently, across the full range of decrement contrasts. The measurements were obtained from three tasks: brightness matching, brightness discrimination and brightness scaling. The results turned out to be closely related. Whittle coined the term "contrast-brightness" to capture the intimate relationship he observed between brightness and contrast (Whittle, 1994a).

5.1. Log W

Whittle showed that a simple metric of contrast captured the relationship between brightness and contrast for patch-on-background stimuli: log W. In log W, $W = \Delta L/(L_{min} + k)$, where ΔL is the difference in luminance between patch and background, L_{min} is the smaller of the patch and background luminances, and k is a constant that prevents W approaching infinity when L_{min} approaches zero. k can be regarded as a measure of the internal noise level when luminance is zero, but if L_{min} is not close to zero the constant can be safely omitted. Fig. 16 illustrates how W is calculated for both an incremental and decremental patch. Note first the difference between W and the two more commonly employed measures of contrast: Weber contrast and Michelson contrast. Weber contrast is defined as $\Delta L/L_b$, where L_b is the background luminance, and Michelson contrast as $(L_{max} - L_{min})/(L_{max} + L_{min})$. In the parameters of the figure, Michelson contrast translates to $\Delta L/(2L_b + \Delta L)$ for the increment and $\Delta L/(2L_b - \Delta L)$ for the decrement. Note that W is calculated differently for the increment and decrement, because L_{min} is the background luminance for the increment but the patch luminance for the decrement.

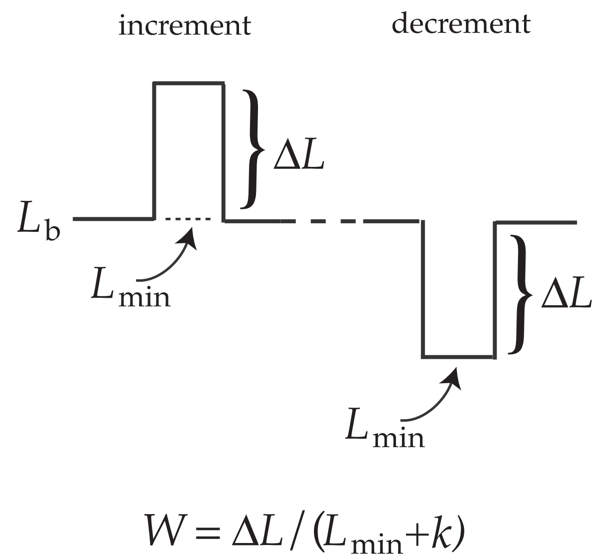

Fig. 16. Calculation of W for both an increment and a decrement.

Fig. 17 shows data from Whittle (1992) along with an illustration of the two types of measurement involved: brightness scaling and brightness discrimination. In the brightness scaling experiment, subjects set the luminances of a series of patches so that they appeared to be at equal brightness intervals. The data (Fig. 17a, closed squares) are the *differences* in ΔL, termed $\Delta^2 L$, between adjacent settings as a function of the luminance of the lower setting. The brightness discrimination data, shown as crosses, are taken from Whittle (1986). These were the just-noticeable-differences (JNDs) in patch luminance, with one of the patches serving as a baseline, or 'pedestal'. The $\Delta^2 L$ values in this case are therefore threshold luminance differences, and have been scaled upwards by a suitable factor to bring them into line with the brightness scaling $\Delta^2 L$s. For both tasks, $\Delta^2 L$ rises with pedestal luminance for increments, but is inverse-U-shaped for decrements. Fig. 17c shows that when the data are re-plotted in terms of log W, the increment and decrement data nearly superimpose and are linearized.

Why does log W capture the brightness behaviour of both increments and decrements? The likely reason is that it encapsulates two visual processes: local light adaptation and a compressive, specifically logarithmic, contrast non-linearity (Kingdom & Whittle, 1996; see also McIlhagga & Peterson, 2006). In all metrics of contrast, the light-adaptation level is embodied in the equation's denominator, and with the denominator set to L_{min}, W embodies the idea that a neural filter sensitive to patch contrast light-adapts to the lower of the two luminances falling within its receptive field.

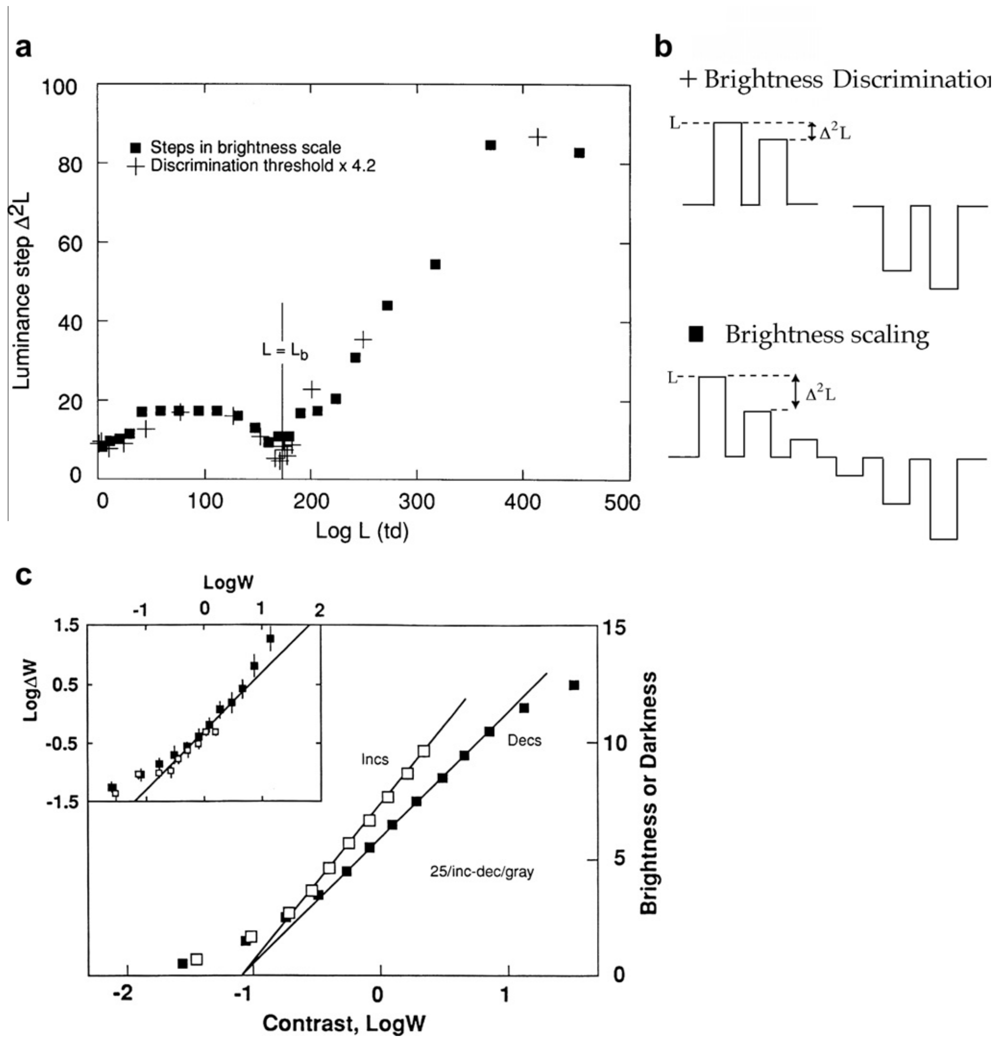

Fig. 17. (a) Data from two types of task, contrast discrimination (crosses) and brightness scaling (filled squares). (b) Illustration of the two types of task, from Kingdom, (2003a). (c) Data from the brightness scaling, and in the inset from the brightness discrimination task, re-plotted in terms of log *W*. Both graphs reproduced from Whittle (1992).

With regard to the logarithmic transformation of *W*, this presumably embodies the 'true' shape of the contrast transducer, and is notably different from the power exponent typically used to model contrast behaviour when using Michelson contrast (see Kingdom & Whittle, 1996).

Is log *W* useful in situations other than that of the haploscopic superimposed-background display? The brightness scaling data in Fig. 17 was measured for a number of patches presented alongside each other on the same background, and Kingdom and Whittle (1996) found that log *W* was a good model of grating contrast discrimination. In both these studies however, the test patches/gratings were on a common background, as one perceives in the haploscopic display. However for the conventional simultaneous contrast display, in which two *different* backgrounds are viewed side-by-side, log *W* ostensibly runs into difficulties (Whittle, 1994b). For example, for corresponding increment and decrement patches with the same luminance ratios, log *W* predicts that two equal-in-luminance increments on different backgrounds should appear no less different in brightness than two equal-in-luminance decrements on different backgrounds, whereas most of the data shows the brightness difference to be bigger for decrements (Whittle, 1994b). However, this may simply reveal the limitations of considering log *W* solely in terms of the simple patch-background arrangement. Log *W* may turn out to be useful complex displays if incorporated at the individual filter response level of a multi-scale transformation. Log *W* thus remains to this day an under-appreciated theoretical tool in vision research.

5.2. Increments versus decrements

Although log *W* is a good model for increment and decrement brightness perception, it does not follow that it is instantiated by a single mechanism. An enduring home-truth of the past quarter century is that increments and decrements are processed in different neural pathways. This should be seen against the notable failure to find consistent evidence that edges enjoy dedicated channels (e.g. Burr et al., 1989 versus Huang et al., 2006). It would appear that the visual system possesses specialized mechanisms for encoding the two luminance-contrast polarities of even-symmetric stimuli (e.g. bars, patches), but no specialized mechanisms for either contrast-polarity of odd-symmetric stimuli (e.g. edges). This leads to the idea that the discrimination of opposite polarity edges, and of edges from bars, is based on the spatial ordering of even-symmetric detector responses. Put another way, a black-white edge is a decrement abutting an increment and a white–black edge is an increment abutting a decrement. The fact that the brightnesses either side of an edge can extend to infinity merely alludes to the fact that increment and decrement channels operate at multiple spatial scales.

The evidence for separate channels for increments and decrements is both physiological––specifically the ''On'' and ''Off'' pathways of the mammalian visual system beginning in the retina (Schiller, 1982; and for a review see Fiorentini, Baumgartner, Magnusson, Schiller, & Thomas, 1990)––and psychophysical (reviewed by Huang et al., 2006). Phenomenologically, it is striking how difficult it is to find a luminance setting of an increment that matches the brightness of a decrement and vice versa: they simply never look the same. Increments invariably appear brighter than decrements whatever their luminance or contrast. A simple demonstration of the categorical nature of increment and decrement perception is shown in Fig. 18 (from Kingdom, 2003a). Fusion of the two stereo-halves reveals two fusible and one rivalrous stereo-pair. The difference in luminance between the top two increments, and between the bottom two decrements, is much greater than between the increment-decrement pair in the middle, yet only the top and bottom pair fuse to produce patches more-or-less midway in brightness between their monocular half-images. It should be noted however that Fig. 18 is equally consistent with the view that opposite edge-polarities, not increments versus decrements, cannot be fused.

5.3. Summary

Log *W* represents an important development in our understanding of brightness and lightness perception, one that will likely prove seminal in our understanding of a number of anisotropies in increment and decrement brightness. Increments and decrements, not edges and bars, are processed in separate neural channels.

6. Question 5: How are brightness and lightness ''filled-in?

A quarter of a century ago it was virtually axiomatic that the brightness and lightness of a uniform region was computed by a process that began with the detection of the region's edges and was then followed by a propagation, or spreading of neural activity to ''fill-in'' the region in between. This idea, along with its corollary that brightness and lightness values are determined by what happens ''at the edge'' (because the putative edge detectors are necessarily small-in-scale otherwise there would be no need for filling-in), continues to exercise a powerful hold on our thinking (e.g. Grossberg & Todorovic, 1988; Paradiso & Nakayama, 1991; Rossi & Paradiso, 1996). The term ''filling-in'' is not however synonymous with neural spreading. Filling-in has been used as a metaphor for both edge integration as well as the representation of uniform areas by low spatial frequencies (see later this section).

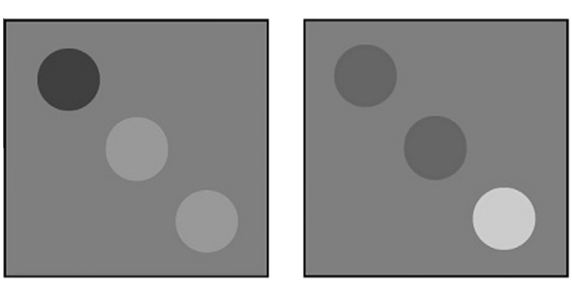

Fig. 18. Increment and decrement perception is categorical. When free-fused, the two decrements (top) and the two increments (bottom) easily fuse, but the increment-decrement pair in the middle whose luminance difference is actually smaller, is rivalrous. From Kingdom (2003a).

Fig. 19. Craik–Cornsweet–O'Brien illusion. The figure appears like a step-edge, but its luminance profile below shows otherwise.

However it is with the idea that filling-in involves neural spreading that is of issue here.

Interestingly, the primary test-bed of the idea of neural spreading, for both protagonists and adversaries alike, has been induced (or illusory) brightness phenomena, even though the principle of neural spreading applies to both real and illusory brightness. Although anecdotal evidence against the idea of neural spreading has been prevalent in the literature for some time (e.g. Kingdom, 2003a), two recent studies by Dakin and Bex (2003) and Blakeslee and McCourt (2008) have been instrumental in persuading a number of people, the present author included, that neural spreading, at least as it relates to brightness perception, is a myth. However, before reviewing these studies, recent evidence in support of the neural spreading hypothesis should be discussed.

6.1. Evidence for the neural spreading hypothesis

The main approach to evidencing neural spreading has been to demonstrate that it has a time course. In one study, Paradiso and Nakayama (1991) flashed a large white disk followed by a mask consisting of a white ring on a black background. The mask was found to reduce the brightness of the target within the region bounded by the ring, as if the ring prevented the brightness signal from propagating inwards from the target's edges. Moreover, as the white disk was increased in size relative to the ring, masking was obtained at longer inter-stimulus-intervals, consistent with the idea of a time course for the propagation. In another study, Paradiso and Hahn (1996) showed that steadily decreasing, or increasing the luminance of a disk led to a slight delay in the corresponding brightness changes in the disk's centre, again consistent with a time course for the propagation. Rossi and Paradiso (1996)used a similar approach with a square-wave grating. They temporally modulated alternate bars of the grating while holding the luminance of the other bars constant, and measured the temporal-frequency at which the induced brightness modulation disappeared, which they termed the ''critical flicker frequency'' or CFF. For narrow-bar gratings the induced CFF was around 5 Hz, decreasing to around 2 Hz as bar width was increased. Because an increase in bar width would be expected to *increase* the CFF if the brightness modulation came from a real luminance modulation, the fact that it decreased, and by an amount that was width-dependent, suggested that there was a slow filling-in process for the induced brightness modulations. Attempts to put a precise value on the time course of the putative neural spreading process in the above studies were hampered however by imprecision. Robinson and de Sa (2008) were more successful in this regard; they used a similar induction stimulus to Rossi and Paradiso but used static presentations and a noise mask, and found that brightness induction was still visible with presentations as brief as 58 ms.

Single unit recordings, some of which have produced evidence in support neural spreading, have also failed to come up with a consistent value for its time course. Hung, Ramsden, and Roe

(2007) based their estimates on spike temporal correlations, and concluded that filling-in within V1 travels between 1300 and 2400 deg/s, and as fast as 4000 deg/s when going from V1 to V2. The authors pointed out however that the spreading activity in V1 that they measured was unrelated to perceived brightness. In contrast, Huang and Paradiso (2008) found that many cells in V1 fired much earlier to a contrast border than to the interior of a large uniform region and calculated the propagation speed to be about 270 deg/s.

The evidence from fMRI studies is contradictory with regard to neural spreading, with a study by Pereverzeva and Murray (2008) in favour and one by Cornelissen, Wade, Vladusich, Dougherty, and Wandell (2006) against.

6.2. Evidence against the neural spreading hypothesis

Blakeslee and McCourt (2008) attempted to measure the time course of brightness induction in the grating-induction stimulus (Fig. 10) using a highly-sensitive method that afforded millisecond precision. They employed the quadrature motion technique to exploit the visual system's excellent motion sensitivity in order to leverage tiny temporal differences into conspicuous changes in motion direction. If two counterphasing gratings with phases in spatial and temporal quadrature (i.e. differing by one-quarter cycle in both space and time) are summed linearly, a moving grating results. The travelling wave moves leftwards when the temporal phase difference is minus one-quarter cycle, and rightward when the temporal phase difference is plus one-quarter cycle. When the temporal phase difference is zero the sum is a standing wave. By counterphasing the inducing grating of a grating-induction stimulus, Blakeslee and McCourt produced a counterphasing induced grating whose spatial phase was opposite to that of the inducing grating but whose temporal phase lagged the inducing grating by 180° plus some unknown quantity that depended on the time lag of brightness induction. They added to the induced grating a like-frequency luminance grating that was counterphasing in spatial quadrature to the inducing grating, but varying in temporal phase. They found that the temporal response of the induced brightness differences lagged by less than 1 ms, and furthermore was constant across wide variations in test field height. The fact that a method capable of measuring the time course of brightness induction with millisecond precision showed that it is virtually instantaneous must be seen as a serious blow to the idea of neural spreading.

The second major piece of evidence against neural spreading comes from a study by Dakin and Bex (2003). They took a new look at what is often heralded as the flagship of neural spreading: the Craik–Cornsweet–O'Brien (CCOB) illusion (see the varieties of this illusion in Todorovic, 1987), a version of which is shown here in Fig. 19. In the CCOB stimulus a step function in luminance is connected by sharp ramps to uniform, equal-in-luminance regions on either side. Yet the stimulus appears compellingly like a real step-edge (indeed one of the methods for constructing a CCOB pattern is to high-pass filter an actual step-edge). The proponents of the neural spreading idea suppose that neural signals propagate outwards from the edge to fill in the regions either side and deliver the illusion. Dakin and Bex (2003) however show in Fig. 20 that the CCOB is critically dependent on the presence of residual low-spatial-frequency information, in sine-phase, at the edge: if the low frequencies are phase-scrambled, the illusion disappears. They propose that contrast normalization, which in their model has the effect of equalizing the energy across log bands of spatial frequency, boosts the residual low-frequency information and it is this that produces the illusion. They also show, as did Burr (1987), that filling the CCOB stimulus with luminance noise, which would be expected to prevent any propagation of neural signals from the edges, fails to destroy the illusion.

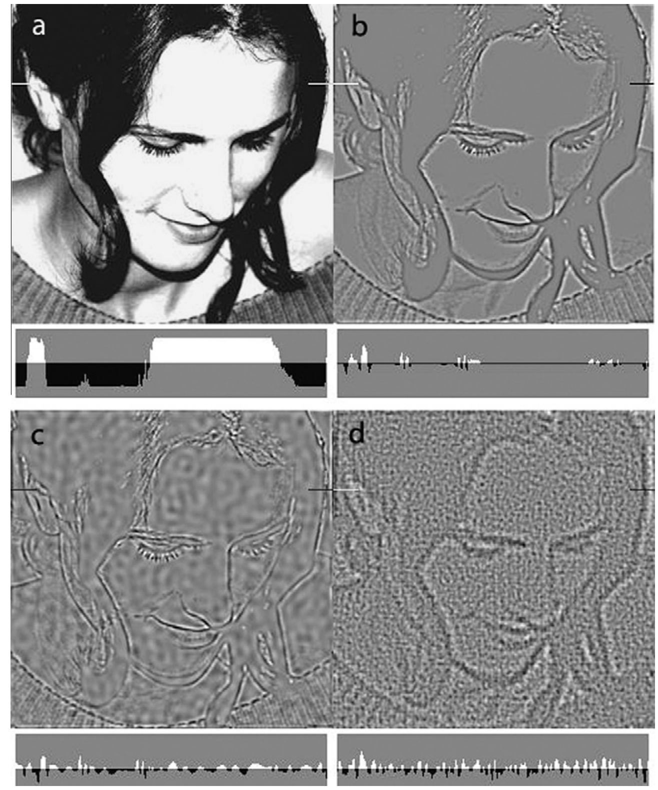

Fig. 20. (a) input image; (b) high-pass filtered gives the Craik–Cornsweet–O'Brien illusion; (c) if the low frequencies are completely removed rather than attenuated, the illusion disappears; (d) adding luminance noise to b. does not destroy the illusion. Below each figure is the cross-section luminance profile at the point marked by the short horizontal lines. From Dakin and Bex (2003), supplied by the authors.

A phenomenon related to the CCOB illusion is the Missing Fundamental (MF) illusion (Campbell, Howell, & Robson, 1971; Campbell, Howell, & Johnstone, 1978). A square-wave grating with its fundamental Fourier component removed appears indistinguishable from a square-wave provided its contrast is below that necessary for the fundamental to be visible if presented alone (Campbell, Howell, & Johnstone, 1978), a finding confirmed also for aperiodic stimuli (Burr, 1987; Kingdom, 1996). By definition linear filtering cannot restore a missing Fourier component no matter what the centre frequency, bandwidth or gain of the filters, so the Dakin and Bex (2003) model cannot predict the MF illusion (confirmed by Dakin, personal communication). Campbell et al.'s (1971) explanation for why the MF appeared like a square-wave is the so-called "default to square-wave" rule, which although arguably ad hoc, remains to this day unchallenged. Therefore, although Dakin and Bex's model fails to predict the MF illusion, it is arguably salvageable by inclusion of this rule.

The view that illusory filling-in phenomena such as the CCOB illusion are explained by the enhancement of low frequencies via contrast normalization, rather than neural spreading, has considerable appeal. Apart from fitting so well with our current understanding of cortical processing and being consistent with other evidence against neural filling-in, it has the additional merit of putting filling-in phenomena such as the CCOB illusion within the same theoretical framework as assimilation phenomena, which as we have seen are also explicable in terms of contrast normalization. It will be interesting to see if other filling-in phenomenon that have so far evaded explanation in terms of contrast normalization can be similarly explained. Two phenomena that spring to mind

Please cite this article in press as: Kingdom, F. A. A. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* (2010), doi:10.1016/j.visres.2010.09.012

are the watercolour illusion (Pinna, Brelstaff, & Spillman, 2001) and neon-colour-spreading (Anderson, 1997; Bressan, Mingolla, Spillmann, & Watanabe, 1997; Hoffman, 1998), both particularly striking in their chromatic versions.

6.3. Summary

The idea that brightness "filling-in" of uniform regions is mediated by neural spreading has been seriously challenged by two sets of findings: 1. That brightness induction is near-instantaneous and 2. That the Craik–Cornsweet–O'Brien illusion is dependent on the presence of residual low-frequency information and is not disrupted by the addition of luminance noise. "Filling-in" should at best therefore be considered as a metaphor for the representation of uniform regions by relatively low spatial frequencies. These low spatial frequencies may in some circumstances be enhanced by contrast normalization. Low spatial frequency enhancement by contrast normalization might explain a range of illusory phenomena that have hitherto attracted less parsimonious theoretical treatment.

7. Question 6: How are relative lightness values converted to absolute ones?

Gilchrist (2006) defines anchoring as "…a rule that identifies a specific value of lightness with some physical property of the image (highest luminance, average luminance, average area)". He argues that most evidence supports the idea that the lightness anchor is the highest luminance in the display, which is designated white, an idea suggested originally by Wallach (1976), incorporated into models of lightness constancy such as the Retinex (Land & McCann, 1971), and assumed to be the case by some other researchers (e.g. Horn, 1986; Marr, 1982). Historically, the other main contender for the lightness anchor is the average scene luminance, designated as mid-grey (Buchsbaum, 1980; Hurlbert, 1986; Hurlbert & Poggio, 1988). One model has posited that both the highest and the surround luminance are used as anchors (Bressan, 2006a, 2006b).

Li and Glichrist (1999), and more recently Gilchrist and Radonjic (2009), tested between the two main anchor contenders. Subjects viewed a dome that filled the entire field of view, with one half painted black the other mid-grey, creating an approximately 5:1 luminance range. Subjects reported that the mid-grey half looked white and the black half mid-grey. Using a more complex stimulus, Cataliotti and Gilchrist (1995) presented observers with a 15-patch Mondrian containing a restricted range of greys from black to mid-grey. The Mondrian was either presented in a spotlight within an otherwise dimly lit room or in a closed chamber. The highest luminance was seen as white, while no black surfaces were seen.

Although this evidence supports the idea that the highest luminance is white, the rule as it stands is untenable. First, as Gilchrist (2006) himself points out, surfaces can appear white even in rooms where the light source, for example a fluorescent light, is clearly visible and invariably the highest luminance. Second, the shadowed area in the photograph in Fig. 2 looks white, even though it is not the highest luminance in the image. These observations point inexorably to the alternative suggested by Rudd and Zemach (2005): white is determined not by the highest luminance but by the highest lightness. If this is true anchoring must follow rather than precede (or at least be independent of) any process of layer decomposition.

Bressan (2006a) has published a compelling figure shown here in Fig. 21 that at first sight also appears to contradict the highest luminance rule. The alternating rows of diamonds, as well as the background white of this paper, are all of the highest luminance, yet look strikingly different in lightness. However, readers may perceive the very bright diamonds in Fig. 21 as self-luminous, and if so the white appearance of the surrounding paper and intervening rows of diamonds *is* consistent with a 'highest' rule, but highest lightness not highest luminance.

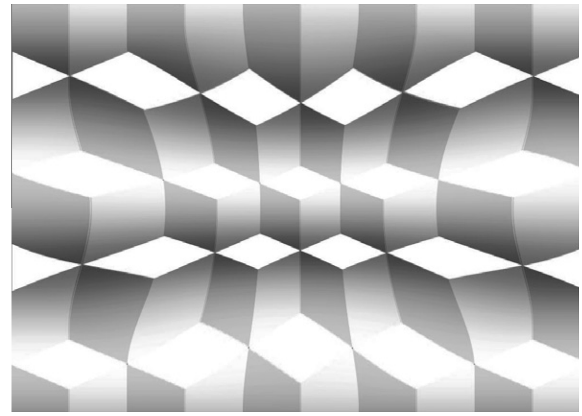

Fig. 21. Bressan's figure challenges the highest-luminance-is-white rule. The different rows of diamonds, as well as the white of the surrounding paper all have the same luminance, yet appear markedly different. From Bressan (2006a), supplied by the author.

Experiments on anchoring have invariably employed simple two-region or Mondrian-like displays, and thus we know little about how relative lightness values are anchored in complex stimuli such as natural scenes. It is possible for example that no one surface in the image serves as the anchor for the whole image, as Bressan (2006a) has suggested, whether encoded before or after layer decomposition, and that instead anchoring operates at multiple spatial scales. It will be interesting to see if the multi-scale filtering approach to brightness perception described earlier can be made to incorporate anchoring at different spatial scales to enable it to predict absolute lightness.

7.1. Summary

While some evidence favours the highest-luminance-is-white rather than average-luminance-is-grey anchoring rule, a plausible alternative is that it is the highest lightness not highest luminance that serves as the anchor, implying that anchoring is instantiated after the process of layer decomposition. At present however little is known about how anchoring operates in complex stimuli, such as natural scenes. Anchoring in complex stimuli may be instantiated at multiple spatial scales.

8. Question 7: What are the cues for perceiving non-uniform illumination and transparency?

During the past quarter-century a considerable effort has been made towards identifying the cues that enable observers to distinguish between reflectance surfaces, transparent media and various types of non-uniform illumination, such as shadows, shading, spotlights, light sources and specular reflections. The cues that have been identified as useful to vision for this purpose are wide-ranging and include photometric, chromatic, geometric, figural, motion, depth and shape cues. A catalogue of these cues and an examination of how they are used by vision has recently been provided by the present author in a separate review (Kingdom, 2008), and the interested reader is therefore directed to this source.

9. Question 8: What are the dimensions of perceptual transparency and how are they encoded?

Metelli's (1974) classic study on perceptual transparency launched a distinct sub-discipline within vision science that has proved fertile ground for both creativity and controversy (e.g., Albert, 2006, 2008; Anderson, Singh, & Meng, 2006; Gerbino, 1994, Chap. 5; Gerbino, Stultiens, Troost, & de Weert, 1990; Kasrai & Kingdom, 2001; Masin, 2006; Masin, Tommasi, & Da Pos, 2007; Robilotto, Khang, & Zaidi, 2002; Robilotto & Zaidi, 2004; Singh & Anderson, 2002a, 2002b, 2006). This section examines recent ideas on the perceptual dimensions of transparency, that is the dimensions along which human observers naturally decompose transparent media.

9.1. Physical transparency

In the introduction it was argued that the two critical physical properties of transparent media are transmittance, which is the proportion of light that passes through the medium, and reflectance, which is the proportion of light reflected from the medium. Fig. 22 shows a simulated transparency overlaying a bipartite background, resulting in four luminances: A, B, P and Q. Metelli's (1974) original equations for transmittance and reflectance were formulated in terms of the reflectances of the surfaces involved, but Gerbino (1994, Chap. 5) has re-cast them in terms of luminances, making them more amenable to analysis for transparencies simulated on CRT displays where luminance is the variable. Although Gerbino's formulations make some simplifying assumptions about the passage of light through the transparent medium to the surface beneath and then back to the observer, the transmittance t and reflectance r of the transparency can be approximated by:

$$t = \frac{P - Q}{A - B}, \quad (1)$$

$$r = \frac{AQ - BP}{A - B} \quad (2)$$

(Gerbino, 1994, Chap. 5; Kasrai & Kingdom, 2001) (an alternative formulation is provided by Beck, Prazny, & Ivry, 1984, and one based on a different physical model of transparency by Robilotto et al. (2002) and Robilotto and Zaidi (2004)). Eqs. (1) and (2) suffice to show that a minimum of four luminances are required to determine both t and r.

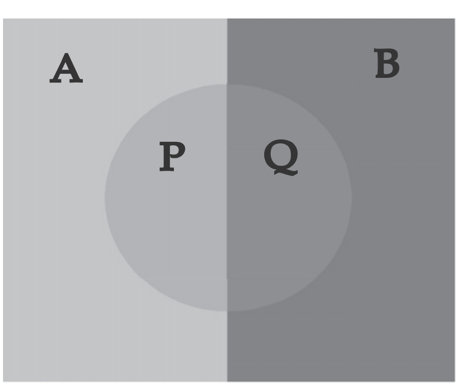

Fig. 22. Simulated transparency of a bipartite background. The luminances A, B, P and Q are sufficient to specify the transmittance and reflectance of the transparency.

9.2. Perceptual dimensions of transparency

The physical analysis of transparency in terms of t and r raises the question of whether the perceptual dimensions of transparency axe correlates of these two dimensions. The issue is controversial. Singh and Anderson (2002a, 2002b, 2006) and Anderson et al. (2006) have argued from transparency matching data that the two perceptual dimensions of transparency are not correlates of t and r, but instead opacity, or 'hiding power' (which they regard as the inverse of perceived transmittance) and lightness. On the other hand Albert (2006, 2008), based on his own transparency matching data, disputes this claim, and Petrini and Logvinenko (2006), using a multi-dimensional scaling (MDS) method, find evidence in support of t and r as the two perceptual dimensions. In Petrini and Logvinenko's study subjects were required to rate the perceived dissimilarity of pairs of transparencies that varied in both t and r, and found that the output configuration from the non-metric MDS was two-dimensional, with one dimension correlated with t, the other with r.

In motivating their case for opacity and lightness as the two perceptual dimensions of transparency, Singh and Anderson (2002a) make the important point that the perceptual representation of transparency is likely to be based on computations that make up the common currency of perception, for example contrast and lightness. They suggest that perceived transmittance (the inverse of opacity) is proportional to the ratio of two contrasts: the contrast of the transparency region and the contrast of the background. Later they modified this formulation to express the ratio in terms of perceived rather than physical contrast (Anderson et al., 2006), but for brevity we will refer to the model as the 'ratio-of-contrast' model. Interestingly, although Singh and Anderson (2002a) identify a ratio-of-contrast computation with perceived transmittance, the measure correlates better with r than t, since when a transparency takes on a reflective component (i.e. r becomes non-zero), which adds a constant luminance to the transparency region, its physical contrast decreases. In keeping with Singh and Anderson (2002a), Robiloto et al. (2002) and Robilotto and Zaidi (2004) found that subjects matched transparencies on the basis of their contrasts, even though the precise dimensions along which they were required to match were unspecified and they were instructed simply to match "transparency".

To obtain a flavour of how these conflicting views on the perceptual dimensions of transparency relate to the physics of transparency, consider the simulated transparencies in Fig. 23 along with their descriptions in terms of their qualitative differences in t and r. The condition in which the ratio-of-contrast model is most noticeably at odds with a model based on t is when $r = 0$, i.e. a transparency without a reflective component.[1] Two such transparencies are shown at the top of Fig. 23. When $r = 0$, the effect of transparency on luminance is purely multiplicative, so for a given background, varying t does not change the physical contrast of the transparency region. Moreover for all backgrounds when $r = 0$, the ratio-of-contrasts for all t is unity. Thus if the evidence were to show that observers perceive the different t but same $r = 0$ transparencies at the top of Fig. 23 as markedly different in perceived transmittance, the ratio-of-contrast model of perceived transmittance would clearly be wrong. However, if the evidence were to show that the two transparencies were not perceived as significantly different in transmittance, then this would be consistent with the ratio-of-contrast model.

[1] If contrast is calculated by Michelson contrast (Singh & Anderson, 2002a), the contrast of the transparency region is $(P - Q)/(P + Q)$, and that of the background $(A - B)/(A + B)$. The ratio of contrasts is then $[(P - Q)/(A - B)] * [(A + B)/(P + Q)]$, or $t(A + B)/(P + Q)$.

Please cite this article in press as: Kingdom, F. A. A. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* (2010), doi:10.1016/j.visres.2010.09.012

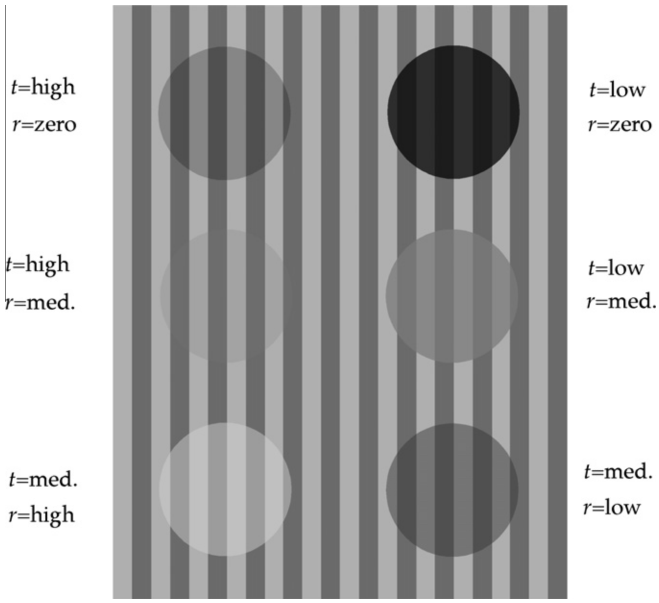

Fig. 23. Simulated transparencies varying in transmittance t and reflectance r. The qualitative descriptions are given next to each transparency. med. = medium. Note that on an uncalibrated monitor these reproductions may not be accurate.

For the middle and bottom pair of transparencies in Fig. 23, the relative contrasts of the transparency and background appear to correlate more closely with t and so the difference between the two models is less apparent. Therefore, whether the perceptual dimensions of transparency are opacity (defined by the ratio-of-contrasts) and lightness, as suggested by Singh and Anderson (2002a), or whether the perceptual dimensions of transparency are correlates of t and r (e.g. Petrini & Logvinenko, 2006), hinges to a large extent on whether or not observers perceive non-reflective transparencies (those with $r = 0$) with very different t as also very different in perceived transmittance.

One reason why the data obtained so far has failed to resolve this issue is because of limitations in current transparency-matching protocols. Psychophysical studies that have employed matching protocols to study transparency have invariably required subjects to adjust one of the parameters of a simulated transparency to match some pre-specified dimension (such as transparency, transmittance, lightness, reflectance) of another simulated transparency (Albert, 2008; Gerbino et al., 1990; Robilotto et al., 2004; Singh & Anderson, 2002a). For example, Singh and Anderson (2002a), using simulated transparencies placed over identical sine-wave gratings, fixed the mean luminances of the two transparencies at different levels and required subjects to adjust the luminance range of one transparency to match that of another in terms of perceived transmittance. Robilotto and Zaidi's (2004) subjects adjusted either the inner transmittance or reflectivity (with the other dimension fixed) of one transparency to match another in terms of perceived transparency, and Albert's (2008) subjects adjusted the transmittance of one transparency (with reflectance fixed) to match the perceived transmittance of another transparency. The potential problem with all these protocols is that as soon as one constrains the form of luminance relations within the transparency region that is being adjusted (e.g. constraining mean luminance, transmittance, reflectance, contrast etc.), subjects become locked into a parameter space for which the adjustable physical dimension may not correlate well with the subject's perceptual representation (interestingly, Anderson, Singh, and O'Vari (2008) highlight this problem with regard to Albert's (2008) method in their reply to Albert's critique of the ratio-of-contrast model). The result is that subjects' matches may end up being significantly biased. For example, suppose hypothetically that the perceptual dimensions of transparency are correlates of t and r so subjects perceive the two transparencies at the top of Fig. 23 as very different in perceived transmittance. Suppose now that the subjects are required to adjust the luminance range but not mean luminance of one transparency to match that of the other transparency in terms of perceived transmittance, as in Singh and Anderson's (2002a) experiments. With the mean luminances fixed and only the luminance range of the match transparency adjustable, any adjustments involve simultaneous changes to *both* t and r. If t and r are correlates of the natural perceptual dimensions of transparency, the two dimensions might be difficult to disentangle when forced to co-vary in this way and as a result the perceptual judgements may be biased.

An alternative approach that would avoid this problem is not to constrain the form of luminance relations in the transparency region during the adjustment procedure. In other words allow P and Q in Fig. 22 to take on *any* value and instruct subjects simply to make the transparencies "look like the same transparency". This could be achieved either by subjects adjusting P and Q, or by adjusting the mean *and* range of P and Q, or by adjusting t and r. The models described above all predict a unique value of P and Q for each match, so the method could in principle decide between different models. One arguable disadvantage of this method however is that it does not allow one to hold one of the putative dimensions of transparency constant while the other is varied. Another possible disadvantage is that it necessitates that the backgrounds of the comparison transparencies are different in order that their Ps and Qs will be set differently.

9.3. Summary

Important new ideas have emerged during the past 25 years as to the perceptual dimensions of transparency. However current transparency-matching protocols are too restricted to provide a clear picture as to what these dimensions are. This could be remedied by the use of matching protocols in which subjects are not constrained by the form of the luminance relations that they are allowed to adjust. Multi-dimensional scaling (MDS) is also a promising approach for determining the perceptual dimensions of transparency.

10. Conclusion: Is a unified account of LBT possible?

Ideally one would like to take any image and decompose it into separate representations, or 'maps', of brightness, lightness, in homogenous illumination and transparency, with the last two dimensions being further subdivided into shadows, spotlights, shading, highlights, light sources and the (two or more) perceptual dimensions of transparency. Will this ever be possible? It seems to be a tall order. The above examination of the current state of knowledge about LBT has revealed a multitude of impressive phenomena, models and partial theories, but not yet the beginnings of a general theory.

Nevertheless, there have been important new developments. Models based on multi-scale filtering combined with contrast normalization are particularly promising. Not only have they proven themselves capable of predicting quantitatively a range of brightness phenomena, but they possess the inherent flexibility to deal with stimuli much more complex than the stock-in-trade ones of the vision laboratory. Some of the shortcomings of these models, such as their failure to account for anisotropies in the perception of increments and decrements, can easily be remedied. A major outstanding problem however is in combining the multi-scale filtering approach with the other big success story of the last quarter century: the intrinsic-image, or layer decomposition approach. Combining the two approaches is a pre-requisite for being able to predict patterns of brightness, lightness, non-uniform illumination and transparency. The problem is a profound one however, because the languages of layer decomposition and filtering are so very different. Bridging the gap between the two will therefore be a major task for future research.

Acknowledgements

This review was supported by a Canadian Institute of Health (CIHR) research Grant No. 11554 given to the author. The author would like to thank the following persons for informal discussions during the writing of the manuscript and/or written comments/reviews on the first draft: Bart Anderson, Jason Bell, Barbara Blakeslee, Paola Bressan, Steve Dakin, Mark Georgeson, Alan Gilchrist, Rick Gursney, Mark McCourt, Beau Lotto, Steve Shevell and Michael Rudd. Special thanks also for the comments from an anonymous reviewer. Finally, thanks to all those who provided 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.

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

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

Albert, M. K. (2008). The role of contrast in the perception of achromatic transparency: Comment on Singh and Anderson (2002) and Anderson (2003). *Psychological Review, 115*, 1127–1143.

Anderson, B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: The role of contour junctions. *Perception, 26*(4), 419–453.

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

Anderson, B. L., Singh, M., & O'Vari, J. (2008). Natural decomposition of perceived transparency: Reply to Albert (2008). *Psychological Review, 115*, 1144–1151.

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

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

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

Arend, L. E., & Goldstein, R. (1987). Simultaneous constancy, lightness and brightness. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 4*, 2281–2285.

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.

Baddeley, R. J., Attewall, D., & Patel, S. (2010). Real world lightness constancy is bad (and why), but perfectly matched to language and memory. *Perception, 39*(Suppl.), 8.

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. (1959). Stimulus correlates for the judged illumination of a surface. *Journal of Experimental Psychology, 58*, 267–274.

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

Bergström, S. S. (1994). Color constancy: Arguments for a vector model for the perception and illumination, color and depth. In A. L. Gilchrist (Ed.), *Lightness, brightness, and transparency* (pp. 215–255). Hove, UK: Lawrence Erlbaum Associates.

Blakeslee, B., & McCourt, M. E. (1997). Similar mechanisms underlie simultaneous brightness contrast and grating induction. *Vision Research, 37*, 2849–2869.

Blakeslee, B., & McCourt, M. E. (1999). A multiscale spatial filtering account of the White effect, simultaneous brightness contrast and grating induction. *Vision Research, 39*, 4361–4377.

Blakeslee, B., & McCourt, M. E. (2001). A multiscale spatial filtering account of the Wertheimer–Benary effect and the corrugated Mondrian. *Vision Research, 41*, 2487–2502.

Blakeslee, B., & McCourt, M. E. (2003). A multiscale spatial filtering account of brightness perception. In L. R. Harris & Jenkin (Eds.), *L*. Springer-Verlag.

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.

Blakeslee, B., & McCourt, M. E. (2005). A multiscale spatial filtering account of grating induction and remote brightness induction effects: Reply to Logvinenko. *Perception, 34*, 793–802.

Blakeslee, B., & McCourt, M. E. (2008). Nearly instantaneous brightness induction. *Journal of Vision, 8*(2), 1–8. http://journalofvision.org/8/2/15/.

Blakeslee, B., Pasieka, W., & McCourt, M. E. (2005). Oriented multiscale spatial filtering and contrast normalization: A parsimonious model of brightness induction in a continuum of stimuli including White, Howe and simultaneous brightness contrast. *Vision Research, 45*, 607–615.

Blakeslee, B., Reetz, D., & McCourt, M. E. (2008). Coming to terms with lightness and brightness: Effects of stimulus configuration and instructions on brightness and lightness judgments. *Journal of Vision, 8*(11), 1–14.

Blakeslee, B., Reetz, D., & McCourt, M. E. (2009). Spatial filtering versus anchoring accounts of brightness/lightness perception in staircase and simultaneous brightness/lightness contrast stimuli. *Journal of Vision, 9*(3), 1–17.

Boyaci, H., Fang, F., Murray, S. O., & Kersten, D. (2007). Responses to lightness variations in early visual cortex. *Current Biology, 17*, 989–993.

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

Bressan, P. (2006b). In homogenous surrounds, conflicting frameworks, and the double-anchoring theory of lightness. *Psychonomic Bulletin and Review, 13*, 22–32.

Bressan, P. (2007). Dungeons, gratings, and black rooms: A defense of double-anchoring theory and a reply to Howe et al. (2007). *Psychological Review, 114*(4), 1111–1115.

Bressan, P., & Actis-Grosso, R. (2001). Simultaneous lightness contrast with double increments. *Perception, 30*, 889–897.

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

Buchsbaum, G. (1980). A spatial processor model for object color perception. *Journal of the Franklin Institute, 310*, 1–26.

Buckley, D., Frisby, J. P., & Freeman, J. (1994). Lightness perception can be affected by surface curvature from stereopsis. *Perception, 23*, 869–881.

Burr, D. C. (1987). Implications of the Craik-O'Brien illusion for brightness perception. *Vision Research, 27*, 1903–1913.

Burr, D. C., Morrone, M., & Spinelli, D. (1989). Evidence for edge and bar detectors in human vision. *Vision Research, 29*, 419.

Campbell, F. W., Howell, E. R., & Robson, J. G. (1971). The appearance of gratings with and without the fundamental Fourier component. *Journal of Physiology, London, 217*, 17–19.

Campbell, F. W., Howell, E. R., & Johnstone, J. R. (1978). A comparison of threshold and supra threshold appearance of gratings with components in the low and high spatial frequency range. *Journal of Physiology, London, 284*, 193–201.

Cataliotti, J., & Gilchrist, A. K. (1995). Local and global processes in lightness perception. *Perception & Psychophysics, 57*, 125–135.

Cornelissen, F. W., Wade, A. R., Vladusich, T., Dougherty, R. F., & Wandell, B. A. (2006). No functional magnetic resonance imaging evidence for brightness and color filling-in in early human visual cortex. *Journal of Neuroscience, 26*, 3633–3641.

Cornsweet, T. N. (1970). *Visual perception*. Academic Press (London) Ltd.

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

Corney, D., & Lotto, B. (2007). What are lightness illusions and why do we see them? *PLoS Computational Biology, 3*, 1790–1800.

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

Dalby, T. A., Saillant, M. L., & Wooten, B. R. (1995). The relation of lightness and stereoscopic depth in a simple viewing situation. *Perception & Psychophysics, 57*, 318–332.

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

Field, D. J. (1994). What is the goal of sensory coding. *Neural Computation, 6*, 559–601.

Fiorentini, A., Baumgartner, G., Magnusson, S., Schiller, P. H., & Thomas, J. P. (1990). The perception of brightness and darkness. In Lothar Spillman & John S. Werner (Eds.), *Visual perception: The neurophysiological foundations*. San Diego: Academic Press, Inc..

Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annual Review of Psychology, 59*, 167–192.

Gelb, A. (1929). Die "Farbenkonstanz" derSehdinge [The color of seen things]. In W. A. von Bethe (Ed.), *Handbuch der normalen und pathologischen physiologie* (Vol. 12, pp. 594–678).

Georgeson, M. A. (1994). From filters to features: Location, orientation, contrast and blur. In *Higher order processing in the visual system (Ciba foundation symposium 184)* (pp 147–165). Chichester: Wiley.

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

Please cite this article in press as: Kingdom, F. A. A. Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research* (2010), doi:10.1016/j.visres.2010.09.012

Gilchrist, A. L. (1977). Perceived lightness depends on perceived spatial arrangement. *Science, 195*, 185–187.
Gilchrist, A. L. (1980). When does perceived lightness depend on perceived spatial arrangment? *Perception & Psychophysics, 28*, 527–538.
Gilchrist, A. (1988). Lightness contrast and failures of lightness 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., & Radonjic, A. (2009). Anchoring of lightness values by relative luminance and relative area. *Journal of Vision, 9*(9), 1–10.
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.
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.
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.
Grossberg, S., & Todorovic, D. (1988). Neural dynamics of 1-D and 2-D brightness perception: A unified model of classical and recent phenomena. *Perception and Psychophysics, 43*, 241–277.
Haynes, J., Lotto, R. B., & Rees, G. (2004). Responses of human visual cortex to uniform surfaces. *Proceedings of the National Academy of Sciences of the United States of America, 101*, 4286–4291.
von Helmholtz, H. (1866/1962). *Treatise on physiological optics, Vol. II. Trans. J.P.L. Southall*. New York: Dover Publications [Vol. II originally published in 1866].
Hering, E. (1874/1964). *Outlines of a theory of the light sense. (L.M.H.D. Jameson, Trans.)*. Cambridge, MA: Harvard University Press.
Hesse, G. S., & Georgeson, M. A. (2005). Edges and Bars: Where do people see features in 1-D images? *Vision Research, 45*, 507–525.
Hong, S. W., & Shevell, S. K. (2004a). Brightness induction: Unequal spatial integration with increments and decrements. *Visual Neuroscience, 21*, 353–357.
Hong, S. W., & Shevell, S. K. (2004b). Brightness contrast and assimilation from patterned inducing backgrounds. *Vision Research, 44*, 35–43.
Hoffman, D. D. (1998). *Visual intelligence*. New York: Norton.
Horn, B. K. P. (1986). *Robot vision*. Cambridge, MA: MIT Press.
Howe, P. (2006). Testing the coplanar ratio hypothesis of lightness perception. *Perception, 2006*(35), 291–301.
Howe, P. D. L., Sagreiya, H., Curtis, D. L., Zheng, C., & Livingstone, M. S. (2007). The double-anchoring theory of lightness of lightness perception: A comment on Bressan (2006). *Psychological Review, 114*, 1105–1110.
Huang, P.-C., Kingdom, F. A. A., & Hess, R. F. (2006). Only two phase mechanisms, ±cosine, in human vision. *Vision Research, 46*, 2069–2081.
Huang, X., & Paradiso, M. A. (2008). VI response timing and surface filling-in. *Journal of Neurophysiology, 100*, 539–547.
Hung, C. P., Ramsden, R. M., & Roe, A. W. (2007). A functional circuitry for edge-induced brightness perception. *Nature Neuroscience, 10*, 1185–1190.
Hurlbert, A. (1986). Formal connections between lightness algorithms. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 3*, 1684–1693.
Hurlbert, A., & Poggio, T. (1988). Synthesizing a color algorithm from examples. *Science, 239*, 484–485.
Jacobsen, A., & Gilchrist, A. (1988). The ratio principle holds over a million-to-one range of illumination. *Perception & Psychophysics, 43*, 1–6.
Kasrai, R., & Kingdom, F. A. (2001). Precision, accuracy, and range of perceived achromatic transparency. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 18*, 1–11.
Katz, D. (1935). *The world of colour*. London: Kegan Paul, Trench, Trubner & Co..
Kingdom, F. A. A. (1996). Pattern discrimination with increment and decrement Craik–Cornsweet–O'Brien stimuli. *Spatial Vision, 10*, 285–297.
Kingdom, F. A. A. (2008). Perceiving light versus material. *Vision Research, 48*, 2090–2105.
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.), *Ch. in levels of perception*. Springer-Verlag.
Kingdom, F. A. A. (2003b). Comment on Gilchrist. In L. Harris & M. Jenkin (Eds.), *Levels of perception*. Springer-Verlag.
Kingdom, F., & Moulden, B. (1988). Border effects on brightness: A review of findings, models and issues. *Spatial Vision, 3*, 225–262.
Kingdom, F., & Moulden, B. (1991). White's effect and assimilation. *Vision Research, 31*, 151–159.
Kingdom, F., & Moulden, B. (1992). A multi-channel approach to brightness coding. *Vision Research, 32*, 1565–1582.
Kingdom, F. A. A., & Whittle, P. (1996). Contrast discrimination at high contrasts reveals the influence of local light adaptation on contrast processing. *Vision Research, 36*, 817–829.
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.
Knill, D. C., & Kersten, D. (1991). Apparent surface curvature affects lightness perception. *Nature, 351*, 228–230.
Land, E. H. (1986). An alternative technique for the computation of the designator in the retinex theory of color vision. *Proceedings of the National Academy of Sciences of the United States of America, 83*, 3078–3080.
Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America, 61*, 1–11.
Li, Z., & Glichrist, A. (1999). Relative area and relative luminance combine to anchor surface lightness values. *Perception & Psychophysics, 61*, 771–785.
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.
Logvinenko, A. D. (1999). Lightness induction revisited. *Perception, 28*, 803–816.
Logvinenko, A. D., & Menshikova, G. (1994). Trade-off between achromatic colour and perceived illumination as revealed by the use of pseudoscopic inversion of apparent depth. *Perception, 23*, 1007–1023.
MacEvoy, S. P., & Paradiso, M. A. (2001). Lightness constancy in primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America, 98*, 8827–8831.
Marr, D. (1982). *Vision: A computational investigation into human representation and processing of visual information*. San Francisco: W.H. Freeman.
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.
McCourt, M. E. (1982). A spatial frequency dependent grating induction effect. *Vision Research, 22*, 119–134.
McIlhagga, W., & Peterson, R. (2006). Sinusoid = light bar + dark bar? *Vision Research, 46*, 1934–1945.
Meese, T. S. (1999). A model of human pattern vision: Association fields for adaptive filters. *Spatial Vision, 12*, 363–394.
Meese, T. S., & Georgeson, M. A. (2005). *Carving up the patchwise transform: Towards a filter combination model for spatial vision*. *Advances in Psychology Research* (Vol. 34). New York: Nova Science Publishers. pp. 51–88.
Metelli, F. (1974). The perception of transparency. *Scientific American, 230*, 90–98.
Moulden, B., & Kingdom, F. (1991). The local border mechanism in brightness induction. *Vision Research, 31*, 1999–2008.
Morrone, M. C., Ross, J., Burr, D. C., & Owens, R. (1986). Mach bands depend on spatial phase. *Nature, 324*, 250–253.
Morrone, M. C., & Burr, D. C. (1988). Feature detection in human vision: A phase dependent energy model. *Proceedings of the Royal Society of London Series B – Biological Sciences, 235*, 221–245.
Moulden, B., & Kingdom, F. (1989). White's effect: A dual mechanism. *Vision Research, 29*, 1245–1256.
Nundy, S., & Purves, D. (2000). A probabilistic explanation of brightness scaling. *Proceedings of the National Academy of Science, 99*, 14482–14487.
Olshausen, B. A., & Field, D. J. (2004). Sparse coding of sensory inputs. *Current Opinion in Neurobiology, 14*, 481–487.
Paradiso, M. A., & Hahn, S. (1996). Filling-in percepts produced by luminance modulation. *Vision Research, 36*, 2657–2663.
Paradiso, M. A., & Nakayama, K. (1991). Brightness perception and filling-in. *Vision Research, 31*, 1221–1236.
Petrini, K., & Logvinenko, A. D. (2006). Multidimensional scaling (MDS) analysis of achromatic transparency. *Journal of Vision, 6*(6), 394a.
Pinna, B., Brelstaff, G., & Spillman, L. (2001). Surface color from boundaries: Anew 'watercolor' illusion. *Vision Research, 41*, 2669–2676.
Pereverzeva, M., & Murray, S. O. (2008). Neural activity in human V1 correlates with dynamic lightness induction. *Journal of Vision, 8*, 1–10.
Purves, D., & Lotto, R. B. (2003). *Why we see what we do. An empirical theory of vision*. Sunderland, MA: Sinauer Associates.
Ramachandran, V. S. (1988). Perceiving shape from shading. *Scientific American, 259*, 76–83.
Reeves, A. J., Amano, K., & Foster, D. H. (2008). Color constancy: Phenomenal or projective? *Perception & Psychophysics, 70*, 219–228.
Robilotto, R., & Zaidi, Q. (2004). Perceived transparency of neutral density filters across dissimilar backgrounds. *Journal of Vision, 4*, 183–195.
Robilotto, R., Khang, B.-G., & Zaidi, Q. (2002). Sensory and physical determinants of perceived achromatic transparency. *Journal of Vision, 2*, 388–403.
Robinson, A. E., & de Sa, V. R. (2008). Brief presentations reveal the temporal dynamics of brightness induction and White's Illusion. *Vision Research, 48*, 2370–2381.
Roe, A. W., Lu, H. D., & Hung, C. P. (2005). Cortical processing of a brightness illusion, Proc. *Proceedings of the National Academy of Sciences of the United States of America, 102*, 3869–3874.
Ross, J., Morrone, M. C., & Burr, D. C. (1989). The conditions under which Mach bands are visible. *Vision Research, 29*, 699–715.
Rossi, A. F., & Paradiso, M. A. (1996). Temporal limits of brightness induction and mechanisms of brightness perception. *Vision Research, 36*, 1391–1398.
Rudd, M. E. (2001). Lightness computation by a neural filling-in mechanism. *Proceedings of the Society of Photo-Optical Engineers, 4299*, 400–413.
Rudd, M. E (2003). Progress on a computational model of human achromatic color processing. *Proceedings of the Society of Photo-Optical Instrumentation Engineers, 170*, 181.
Rudd, M. E., & Arrington, K. F. (2001). Darkness filling-in: A neural model of darkness induction. *Vision Research, 41*, 3649–3662.
Rudd, M. E., & Popa, D. (2007). Stevens' brightness law, contrast gain control, and edge integration in achromatic color perception: A unified model. Journal of the Optical Society of America A: Optics, Image Science, and Vision, 24, 2766–2782. Errata. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 24*, 3335.

Rudd, M. E., & Zemach, I. K. (2004). Quantitative properties of achromatic color induction: An edge integration analysis. *Vision Research, 44*, 971–981.

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*, 983–1003.

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.

Schein, S. J., & Desimone, R. (1990). Spectral properties of V4 neurons in the macaque. *Journal of Neuroscience, 10*, 3369–3389.

Schiller, P. H. (1982). Central connections of the retinal ON and OFF pathways. *Nature, 297*, 580–583.

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

Shapiro, A., Knight, E., & Lu, Z.-L. (2008). Lightness, brightness and luminance: Spatial scale models of lightness illusions: Contrast, anchoring, and tunable filters. *Journal of Vision, 8*(6), 288a.

Shapley, R., & Enroth-Cugell, C. (1984). Visual adaptation and retinal gain controls. In N. N. Osbourne & G. J. Chader (Eds.), *Progress in retinal research* (pp. 263–346). Oxford: Pergamon Press.

Shapley, R., & Reid, R. C. (1985). Contrast and assimilation in the perception of brightness. *Proceedings of the National Academy of Sciences of the United States of America, 82*, 5983–5986.

Singh, M., & Anderson, B. L. (2002a). Toward a perceptual theory of transparency. *Psychological Review, 109*(3), 492–519.

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

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

Shevell, S. K., Holliday, I., & Whittle, P. (1992). Two separate neural mechanisms of brightness induction. *Vision Research, 32*, 2331–2340.

Todorovic, D. (1987). The Craik–O'Brien–Cornweet effect: New varieties and their theoretical implications. *Perception & Psychophysics, 42*, 545–560.

Todorovic, D. (2006). Lightness, illumination, and gradients. *Spatial Vision, 19*, 219–261.

Wallach, H. (1976). *On Perception*. New York: Quadrangle/The New York Times Book Co.

Walraven, J., Enroth-Cugell, C., Hood, D. C., MacLeod, D. I. A., & Schnapf, J. L. (1990). The control of visual sensitivity: Receptoral and postreceptoral processes. In L. Spillman & J. S. Werner (Eds.), *Visual perception: The neurophysiological foundations*. San Diego, California: Academic Press, Inc.

Wandell, B. A., & Wade, A. R. (2003). Functional imaging of the visual pathways. *Neurologic Clinics (North America), 21*, 417–443.

Watt, R. J., & Morgan, M. J. (1985). A theory of the primitive spatial code in human vision. *Vision Research, 25*, 1661–1674.

White, M. (1979). A new effect on perceived lightness. *Perception, 8*, 413–416.

Whittle, P. (1986). Increments and decrements: Luminance discrimination. *Vision Research, 26*, 1677–1691.

Whittle, P. (1992). Brightness, discriminability and the "crispening effect". *Vision Research, 32*, 1493–1507.

Whittle, P. (1994a). The psychophysics of contrast brightness. In A. Gilchrist (Ed.), *Lightness, brightness, and transparency* (pp. 35–110). Hillsdale: Erlbaum.

Whittle, P. (1994b). Contrast brightness and ordinary seeing. In A. Gilchrist (Ed.), *Lightness, brightness, and transparency* (pp. 111–157). Hillsdale: Erlbaum.

Williams, S. M., McCoy, A. N., & Purves, D. (1998a). An empirical explanation of brightness. *Proceedings of the National Academy of Sciences of the United States of America, 95*, 13296–13300.

Williams, S. M., McCoy, A. N., & Purves, D. (1998b). The influence of depicted illumination on perceived brightness. *Proceedings of the National Academy of Sciences of the United States of America, 95*, 13301–13306.

Yazdanbakhsh, A., Arabzadeh, E., Babadi, B., & Fazl, A. (2002). Munker–White-like illusions without T-junctions. *Perception, 31*, 711–715.

Zaidi, Q., Spehar, B., & Shy, M. (1997). Induced effects of backgrounds and foregrounds in three-dimensional configurations: The role of T junctions. *Perception, 26*, 395–408.

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

Zeki, S. (1983a). Colour coding in the cerebral cortex: The reaction of cells in monkey visual cortex to wavelengths and colours. *Neuroscience, 9*, 741–765.

Zeki, S. (1983b). The relationship between wavelength and color studied in single cells of monkey striate cortex. *Progress in Brain Research, 58*, 219–227.

Zeki, S., Aglioti, S., McKeefry, D., & Berlucchi, G. (1999). The neurological basis of conscious color perception in a blind patient. *Proceedings of the National Academy of Sciences of the United States of America, 96*, 14124–14129.