3

Levels of Brightness Perception

Frederick A. A. Kingdom

3.1 Introduction

Most vision scientists are comfortable with the idea that perception is a multi-level process. For example we take it for granted that the two best understood properties of human colour vision, trichromacy and colour-opponency, are underpinned by physiological mechanisms operating at different stages in the visual pathway. The observation that any colour can be matched by a suitable mixture of three primaries - the definition of trichromacy - is understood to be a consequence of having three cones rather than two or four. On the other hand, our inability to perceive reddish greens or bluish yellows, one of the hallmarks of colour-opponency, is understood to be a consequence of the particular way the three cones are combined post-receptorally. In other words trichromacy and colour-opponency have independently measurable behavioural consequences that reflect their different physiological origins.

In this chapter I will argue that a multi-level approach is also the right approach for brightness and lightness perception. No single process mechanism can, in my view, account for the many fascinating brightness/lightness phenomena that presently fill the pages of journals and textbooks alike. This admittedly unglamorous viewpoint is not, as one might expect, shared by all. Notably, Gilchrist and Economou (this volume) argue that all brightness/lightness phenomena can be explained within a single theoretical framework. Their approach, inspired by Gestalt psychology, rejects the very idea of "levels" in perception. Their viewpoint has come to the fore at the same time as a renaissance of interest in contextual effects on surface colour appearance (e.g., see the recent special editions of Perception, 1997, Vol. 26, Nos. 4 and 7). Today's emphasis is on configurational relationships, and these are believed to be the major, if not the sole determinant of the perceived pattern of brightness/lightness variations in the image.

I will argue that such contextual effects are best considered within a multi-level framework that includes both low-level contrast and mid-level configurational mechanisms. While the famous Gestalt maxim "the whole is greater than the sum of the parts" undoubtedly holds for brightness/lightness perception, it will be argued that the parts, when considered as levels or stages, nevertheless do exist and can be behaviourally identified.

It should be emphasised that his chapter is not a review of the rapidly expanding literature on lightness and brightness perception. Nor does it present a fully-fledged model of lightness/brightness perception. It is essentially a polemic, and only a handful of studies are described that are necessary to make the point. There are many excellent published studies that are highly relevant to the issues dealt with here that are not discussed, and I apologise in advance to anyone who feels their work should have been mentioned but was not.

3.2 Simultaneous brightness contrast

Simultaneous brightness contrast, or SBC, is the observation that a grey patch looks brighter on a dark compared to a bright surround. The phenomenon illustrated in Figure 3.1a has intrigued philosophers and scientists for over two millenia (see Wade, 1999, for a historical overview), and it is a sobering thought that even after all this time there still seems no consensus as to why such an apparently simple effect occurs.

When discussing SBC, I will be somewhat cavalier in my use of the terms ’lightness’, or perceived reflectance, and ’brightness’, or perceived luminance. A good discussion of the definitions, uses and misuses of these terms can be found in the accompanying chapter by Blakeslee and McCourt. Here, I will assume that for figures without an explicit illumination component, such as Figure 3.1a, it is immaterial whether one makes a relative brightness, or a relative lightness judgement. Indeed it would be equally valid to refer to Figure 3.1a as an example of simultaneous lightness, rather than brightness contrast. The situation is very different however with Figure 3.1b, where there is a pictorial impression of a change in illumination, or specifically a highlight. The distinction between brightness and lightness, as we shall see, becomes critical in any discussion of such stimuli.

Many are familiar with the controversy over SBC that began with Hering and Helmholtz in the 19th century. They disagreed as to whether SBC was based on peripheral sensory processes sensitive to contrast (Hering’s view), or central influences involving assumptions about the configuration of the display as a whole (Helmholtz’s view) (and see Kingdom, 1997). Up to thirty years ago, the dominant view was that contrast underlied SBC, a view sustained by Hurvich and Jameson (e.g., Hurvich and Jameson, 1966), whose own ideas were inspired by Hering (1874/1964). The undergraduate textbook explanation for SBC in Figure 3.2a, which is that SBC results from the operation of filters with centre-surround receptive fields (such as retinal ganglion cells), is the modern version of Hering’s explanation of SBC in terms of “lateral inhibition”. The idea that low-level filters sensitive to contrast underlies SBC has been the principle theme behind a new generation of brightness models whose other defining characteric is filtering at multiple spatial scales (e.g. Kingdom and Moulden, 1992; Blakeslee and

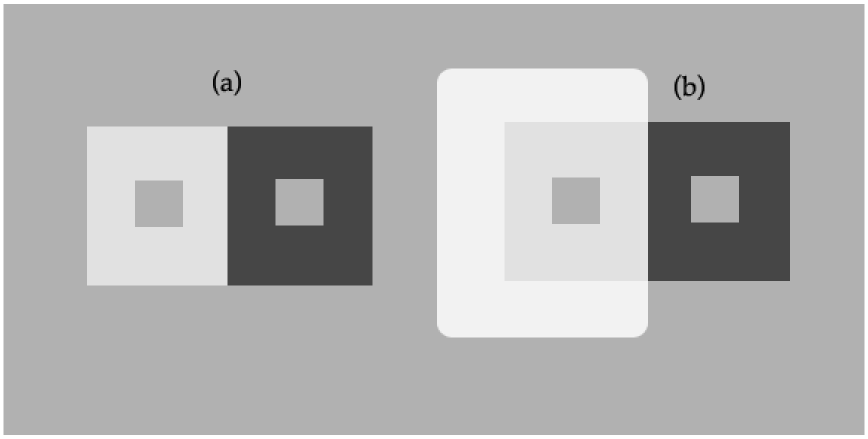

FIGURE 3.1. (a) Standard simultaneous brightness contrast (SBC) display. The two grey patches are equal in luminance, yet appear different in brightness. (b) Simulation of arrangement employed by Gilchrist (1979). One half of the display appears to be lit by a highlight rather than being of different reflectance. 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 SBC was reportedly much greater than seen here.

McCourt, 1999; 2001a,b).

During the last decade, the pendulum has swung away from the idea that errors such as SBC are due to contrast, and more towards the idea that they result from of mechanisms sensitive to the overall configuration of the display. I have argued previously (Kingdom, 1999) that the studies of Gilchrist and his colleagues in the 70's and 80's (e.g. Gilchrist, 1977; Gilchrist 1979; Gilchrist, Delman and Jacobsen, 1983; see also Gilchrist, 1988) were instrumental in precipitating this change in opinion. One of the stimuli that was central to establishing the new way of thinking is illustrated in Figure 3.1b. It must be made clear at the outset that Figure 1b is only an illustration; Gilchrist's original experiments used Munsell papers and real illuminations, and the effects were reportedly much greater than can be seen here. In Figure 3.1b the standard SBC display has been replaced by one consisting of a uniform background with one half illuminated by a bright light. The luminances of the test patches and their surrounds are however identical to the standard display. With Gilchrist's stimulus, subjects reported an enhancement of the lightness difference between the two patches in the part-highlighted display. Since the contrasts of the patches with their surrounds are the same under both configurations (we will return to a critical examination of this assumption later on), the enhancement of SBC can not, it seems be due to the effects of contrast. It must instead be due to the way subjects interpreted the display as a whole.

Gilchrist et al.'s experiments laid the foundation for many recent demonstrations on a similar theme (e.g. Knill and Kersten, 1991; Adelson, 1993; Anderson, 1997; Logvinenko, 1999). With the aid of modern computer graphics, pictorial

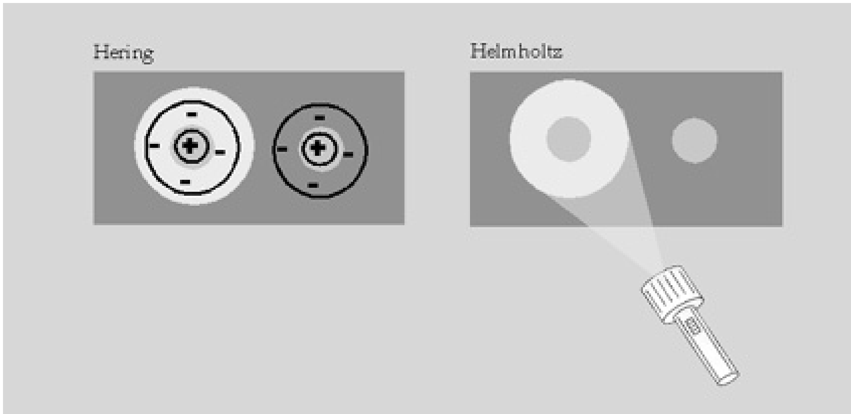

FIGURE 3.2. (a) Explanation for SBC in terms of centre-surround receptive-field filters, the modern version of Hering's explanation in terms of "lateral inhibition". (b) Explanation of SBC based on Helmholtz's veiling hypothesis. See text for details.

representations of complex three-dimensional patterns with vivid impressions of shadows, transparency and shading have replaced the conventional SBC display, and with impressive results. I have provided an example of my own in Figure 3.12, a figure inspired by Adelson's (unpublished) checkerboard-shadow illusion. The allure of these new demonstrations is the sheer magnitude of their illusory brightness differences, which far surpass that found with standard SBC displays. Yet there is a negative side. For some protagonists it has meant downplaying the value of not just simple forms of SBC such as the standard display, but more importantly their explanation in terms of contrast. After all, if such stunning illusions are apparently inexplicable in terms of contrast, is contrast really that important ? I will argue in the next section that it is. Moreover, demonstrations suggesting that contrast may be insufficient to account for certain brightness/lightness phenomena are best considered in terms of the multi-level framework advocated here, where contrast forms an essential component. Let us therefore now look to the evidence that contrast plays a central role in brightness/lightness perception.

3.3 Contrast-brightness and low-level filtering

3.3.1 A common transducer function for brightness discrimination and brightness scaling

My first piece of evidence comes from the work of Paul Whittle. Whittle's quantitative measurements of brightness obtained under a variety of task conditions have provided some of the best evidence for the role of contrast in brightness perception. A comprehensive exposition of Whittle's findings and their theoretical implications is provided in two book chapters, Whittle (1994a,b), and here the reader will obtain the full story of 'contrast brightness', the term Whittle used to capture the idea of the intimate relationship between contrast and brightness. I consider here a subset of Whittle's findings that for me provides the most succinct evidence for contrast-brightness. Figure 3.3 shows data taken from Whittle (1992) (see also Whittle, 1994a), along with my own illustration of the two types of measurement involved, namely brightness discrimination and brightness scaling. In the brightness discrimination task, subjects were required to detect a difference in the luminance of two patches, where one of the patches served as a baseline, or 'pedestal'. With this task, the term 'brightness discrimination' is synonymous with both 'luminance discrimination' and 'contrast discrimination', as it is a *threshold* task involving a comparison between two patches against the *same* background. Results for one background are shown as the crosses in Figure 3.3 (data originally from Whittle, 1986). For increments, the thresholds rise with pedestal luminance, whereas for decrements the function is inverse U-shaped. The different shapes of the increment and decrement functions is of interest in itself (e.g. see Whittle, 1986; Whittle, 1994a; Kingdom and Whittle, 1996), but for the present purpose one need only assume that both functions reflect the shape of the underlying transducer function for contrast. The second set of measurements in Figure 3.3, the closed squares, are from the brightness scaling experiment (original data from Whittle, 1992). In this task subjects were required to set the luminances of a series of patches so that they appeared to be equally different in brightness. As with the threshold task, one of the patches served as a pedestal. In Figure 3.3 the difference in luminance between adjacent pairs of patches is plotted as a function of pedestal luminance. When the brightness discrimination thresholds were scaled upwards by a suitable factor so that they could be compared directly to the brightness scaling data, the two functions almost perfectly superimposed. This strongly suggests that the underlying transducer function for the threshold brightness discrimination task is the same as that for the suprathreshold brightness scaling task. Given that the detection of threshold differences in brightness/luminance/contrast is universally believed to be mediated by bandpass filters in the visual cortex, Whittle's experiment provides powerful evidence that the same filters are also involved in signalling suprathreshold brightnesses.

3.3.2 Illusory gratings facilitate the detection of real gratings

Whittle's experiment demonstrated that a critical behavioural signature for contrast transduction could be revealed in data from a prototypical brightness task. A similar rationale lay behind an experiment I recently conducted in collaboration with Mark McCourt (McCourt and Kingdom, 1996), my second piece of evidence for the role of contrast in brightness perception. We used a form of SBC known as grating induction that was first demonstrated by McCourt (1982; see also Blakeslee and McCourt, 1997, for the evidence that grating induction is SBC). Figure 3.4 shows an example grating induction stimulus. An opposite-phase illusory sine-wave grating is observed in the uniform mid-grey stripe that runs horizontally through an 'inducer' sine-wave grating. Grating induction is a

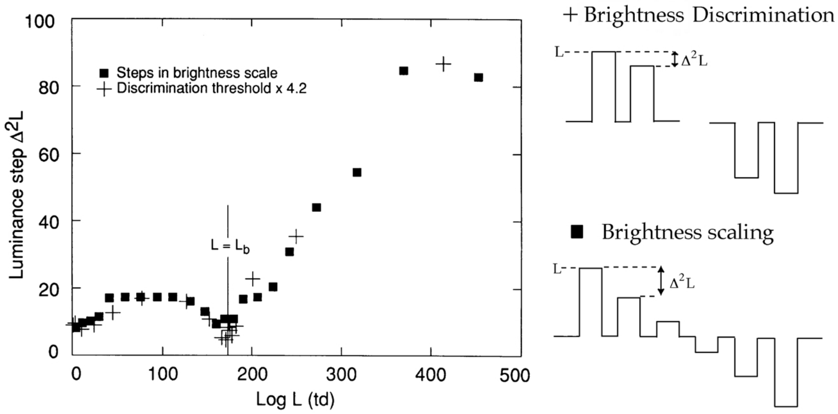

FIGURE 3.3. Left: data from two types of task, contrast discrimination (crosses) and brightness scaling (filled squares), from Whittle (1992). Right: my illustration of the two types of task. Left: reprinted from Whittle, P., Brightness discriminability and the “crispening effect”, *Vis. Res.*, 1493-1507, 1992, with permission from Elsevier Science.

useful tool for studying SBC as it lends itself easily to parametric manipulation, and under some circumstances can be quite compelling (e.g. see Figure 3.7). We reasoned that if the induced brightness variations in grating induction were signalled by the same mechanisms that detect real gratings - and here we come to the idea of a critical behavioural signature - an illusory grating should facilitate the detection of a real grating. Facilitation, as used here, means a reduction in the threshold for detecting a stimulus as a result of the presence of another stimulus. The best-known form this takes is the ‘dipper’ observed in the function relating contrast increment thresholds to pedestal contrast (Campbell and Kulikowski, 1966; Foley and Legge, 1981). When the test is added to a different type of stimulus, the pedestal is usually referred to as a mask.

Our experiment is illustrated in Figure 3.5. We first measured increment thresholds for real gratings whose spatial characteristics were the same as the illusory gratings that formed the main part of the study (Figure 3.5a). We then repeated the experiment this time using *illusory* rather than real pedestals (Figure 3.5b). By varying the contrast of the inducer we were able to vary the apparent contrast of the illusory pedestal. The real test grating was added in phase with the illusory pedestal (which at very low inducer contrasts was not visible) in one of the two forced-choice intervals, and subjects had to decide which interval contained the test. Finally we used a matching technique to find the contrast of a real grating that matched that of the illusory grating at each inducer contrast (Figure 3.5c). This allowed us to recast the contrast of the inducer in terms of ’equivalent’ real grating contrast.

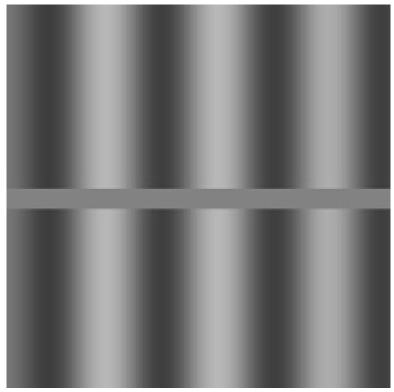

FIGURE 3.4. Grating induction stimulus, first described by McCourt (1982).

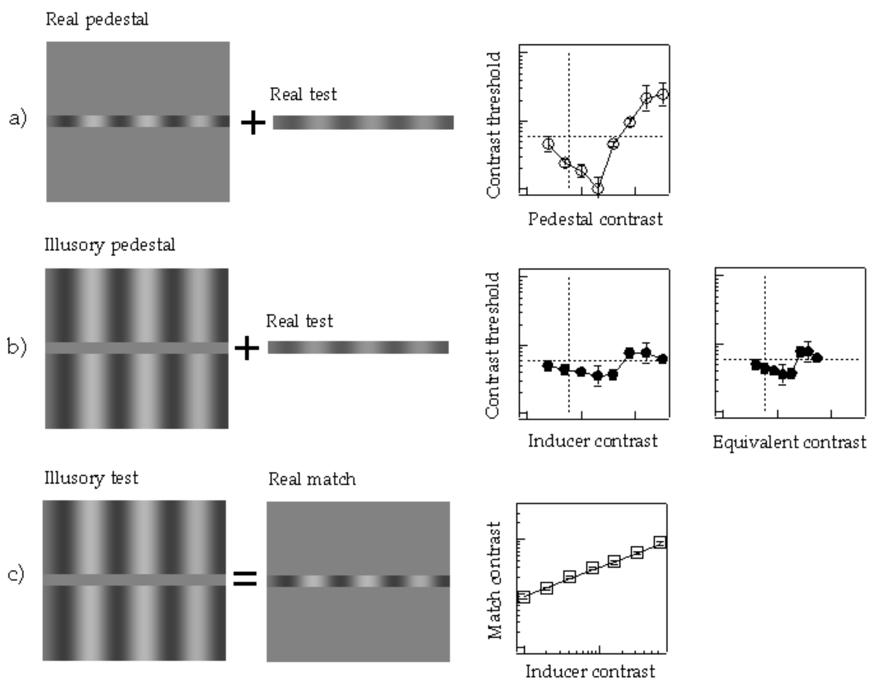

FIGURE 3.5. Method employed by McCourt and Kingdom. In (a) increment thresholds are measured as a function of pedestal contrast for a real grating gated into a narrow stripe. (b) Detection thresholds are measured for a real grating added in phase to an illusory grating, for various contrasts of inducer. (c) The apparent contrast of the illusory gratings was measured by matching them to real gratings. This allowed the contrast of the inducer to be re-cast in terms of ‘equivalent’ real grating contrast. Based on Figure 1 of McCourt, M.E. and Kingdom, F. A. A., Facilitation of luminance grating detection by induced gratings, *Vis. Res.*, 36: 2563-2573, 1996. Copyright Elsevier Science. Used with permission.

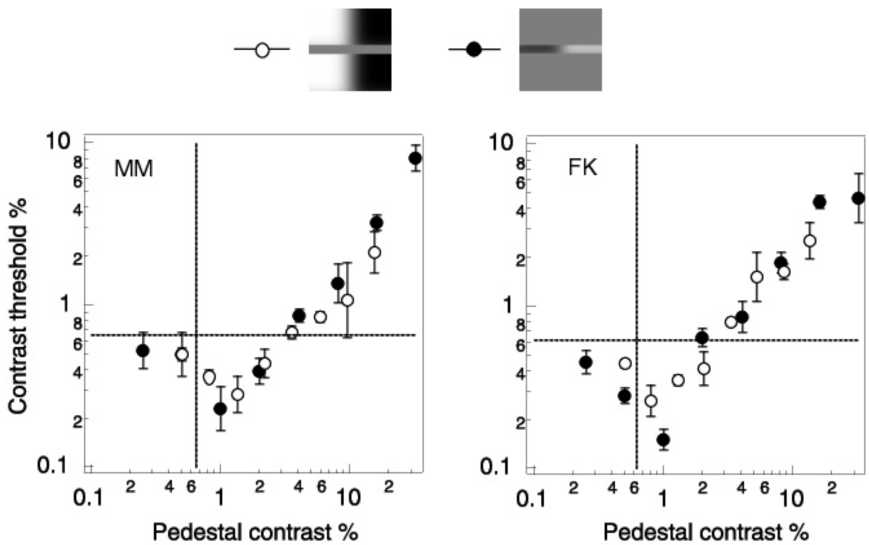

FIGURE 3.6. Results from McCourt and Kingdom using 0.0625 cpd gratings. Filled circles are increment thresholds for a real grating plotted against real grating pedestal contrast. Open circles are thresholds for detecting a real grating on an illusory grating pedestal. The contrast of the illusory pedestal is given as the equivalant contrast of a matched real grating. Note how the real and illusory grating two functions neatly superimpose. Data taken from Figure 3 of McCourt, M. E. and Kingdom, F. A. A., Facilitation of luminance grating detection by induced gratings, *Vis. Res.*, 36: 2563-2573, 1996. Copyright Elsevier Sciences. Used with permission.

Figure 3.6 shows results from the grating spatial frequency that produced one cycle of modulation across the display (0.0625 cpd). The data for the real and illusory grating pedestals almost perfectly superimpose when the contrast of the inducer is couched in terms of equivalent contrast. This shows that, at least for one set of spatial characteristics, an illusory grating acts as an almost perfect metamer of a real grating of the same apparent contrast, in terms of its ability to facilitate (and mask) the detection of a superimposed real grating. I see no alternative explanation for these results other than that illusory gratings are signalled by the same mechanisms that detect real gratings (and see Kingdom, McCourt and Blakeslee, 1997, for further evidence in support of this conclusion). Given the abundance of evidence that real gratings are detected by narrowband filters in the visual cortex, one is once again drawn irrevocably to the conclusion that the same filters are involved in signalling brightness variations, in this case illusory ones.

Besides this quantitative evidence, there are some simple demonstrations of grating induction that provide additional evidence for a central role for low-level contrast-sensitive filters. Two of my favourites are shown in Figures 3.7 and 3.8. Figure 3.7 shows two patterns, each appearing to consist of a low contrast, single cycle sine-wave grating in a narrow horizontal stripe on a uniform surround (based on a similar figure in Moulden and Kingdom, 1991). However, only one of the two

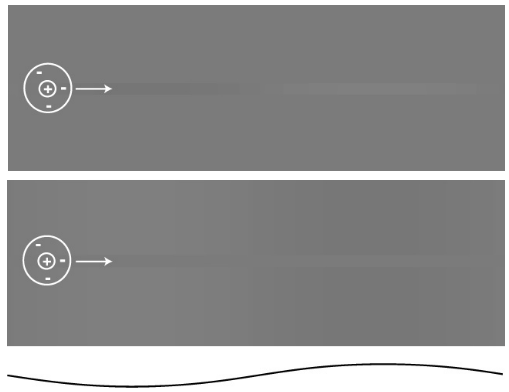

FIGURE 3.7. Top: a single cycle of a real grating runs along a narrow stripe in the middle of a uniform background. Bottom: a uniform stripe lies in the middle of a single-cycle inducer grating. The illusory grating in the bottom figure is more visible than the real grating that induces it. The appearance of both stimuli is most parsimoniously explained by the convolution response (shown below) of a bandpass filter whose receptive field centre is similar in diameter to the height of the stripe.

patterns physically accords with this description - the one at the top. In the bottom pattern it is the stripe that is uniform, and the surround that contains the sine-wave; hence the sine-wave in the stripe is illusory. It is hard to tell the two patterns apart. With scrutiny, the digital quantization of the low amplitude luminance gradients gives it away, but the metamerism of the two patterns is nevertheless striking. Also striking is that in the bottom figure the illusory grating is more visible than the surround grating that induces it.

It is easy to explain the appearance of both patterns in Figure 3.7 with filtering. At the bottom of the figure is shown the horizontal convolution response of a centre-surround filter, obtained when centred on either stripe. Because both the real (top) and illusory (bottom) gratings are gated into narrow stripes, the filter giving the biggest response is one whose centre diameter is approximately the same as the height of the stripe. The surround grating in the bottom pattern however will only weakly stimulate the same filter because its dominant spatial frequency lies almost outside the filter's passband.

In Figure 3.8a, stripes containing ramps in luminance alternate with uniform stripes (see also Moulden and Kingdom, 1991). It is hard to distringuish the stripes containing ramps from those that are uniform. In this instance the induced brightness variations are almost as salient as the inducing brightness variations. Once again, the filtering explanation suffices. A filter matched to the height of the stripe produces a response of opposite phase to the ramp and uniform stripes, but of more-or-less identical amplitude, in accord with the percept. Finally, Figure 3.8b shows that at higher contrasts the illusion begins to break down, in that one can easily distinguish the ramp from the uniform stripes. I will discuss the significance of this last demonstration in the following section.

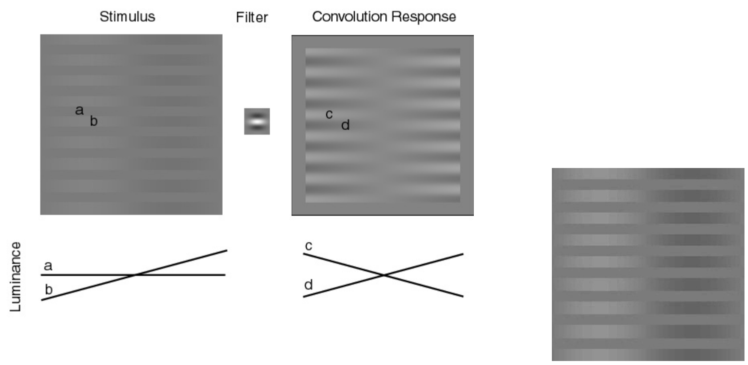

(a) (b)

FIGURE 3.8. (a) Ramp-induced brightness. The stimulus on the top left consists of horizontal uniform stripes (e.g. a) alternatinb with luminance stripes (e.g. b). It is difficult to tell which stripes are uniform and which ramps. On the top right is the convolution response of a centre-surround filter whose centre diameter is approximately the same as the height of the stripes. The response accords with the percept of the stimulus, as shown in the luminance profiles of c and d. (b) The ramps and uniform stripes become more discriminable at higher contrasts. Figure reprinted from Kingdom, F. A. A., Guest editorial: Comments on Lonvinenko "Lightness induction revisited", *Percept.*, 28: 929-934, 1999, with permission from Pion Ltd.

3.3.3 Increment and decrement perception is categorical

My final piece of evidence for a low-level contrast mechanism involves an examination of the differences between ’increments’ and ’decrements’. I refer here to the sign, or polarity of contrast of relatively small, usually closed regions in the image. When I began my research into brightness perception in the 1980s I was often struck, and frequently irritated, by just how difficult it was to find a luminance setting of an increment that matched the brightness of a decrement, and vice-versa. Somehow they never quite looked the same. More often than not, increments looked brighter than decrements whatever their luminance (see Whittle, 1994a), and sometimes they even seemed to take on a slightly different hue. These observations complement a substantial psychophysical and neurophysiological literature suggesting that increments and decrements are processed by different mechanisms, specifically the “On” and “Off” pathways of the mammalian visual system that begin at the retina (Schiller, 1982; and for a review see Fiorentini et al., 1990).

A simple demonstration of the categorical nature of increment and decrement perception is shown in Figure 3.9, which is based on an early finding by Whittle (1965). Fusion of the two stereo-halves reveals two fusable, and one rivalrous stereo-pair. The difference in luminance between the top two increments, and between the bottom two decrements, is 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. The categorical nature of increment and decrement perception, with its ready physiological substrate in the form of “On” and “Off” pathways, shows that our brightness perception is in part a result of low-level physiological processes.

The unique perceptual properties of increments and decrements also pose a special challenge for modellers. The output of brightness models (e.g. Kingdom and Moulden, 1992; Blakeslee and McCourt, 1999; 2001a,b) is a map of (relative) brightness values. For example with SBC a ’successful’ prediction is a lower value for the patch on the white background compared to the patch on the black background. Although this accords with our perception that one patch looks darker than the other, it does not capture the categorical nature of the difference.

3.4 Multiscale filtering and edged-based filling-in

So far I have omitted to discuss details of the filters involved, showing instead how in principle filtering is a valid and simple explanatory tool. In Figures 3.7 and 3.8, a single, linear, circularly-symmetric filter captured the qualitative appearance of the stimuli. This is, of course, a gross over-simplification. We know that contrast (and hence brightness) coding is a multi-scale process, involving cortical filters tuned to a range of scales and orientations. A full multi-scale (such as wavelet) transform of an image produces a veridical output, and if this is what the visual system performed, illusions such as SBC would not occur. One of the main rea-

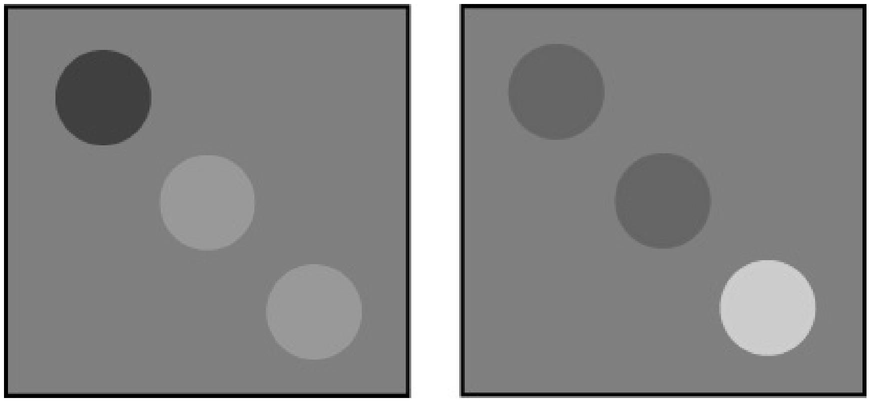

FIGURE 3.9. Increment and decrement perception is categorical. When free-fused, the two decrements (top) and two increments (bottom) easily fuse, but the increment-decrement pair in the middle is rivalrous. The difference in luminance between the increment pair, and also between the decrement pair, is actually bigger than between the increment-decrement pair.

sons why filtering results in brightness illusions is our relative insensitivity to low spatial frequencies, which is particularly marked at low contrasts. In Figure 3.8c, unlike its low contrast version in Figure 3.8a, the uniform and ramp stripes have different perceived amplitudes. This is almost certainly due to the increased involvement of filters tuned to relatively low spatial frequencies.

I suggest that all sizes of filters that are active contribute to our percept of brightness/lightnesss. Precisely how the outputs of filters at different scales (and orientations) are combined for brightness and lightness perception, and in particular what types of nonlinearities are involved, is not fully understood. Yet the most successful attempts at modelling brightness phenomena in terms of filtering have employed filters at multiple spatial scales (Kingdom and Moulden, 1991; and especially McCourt and Blakeslee, 1999, 2001a, b).

Given the abundance of evidence that the early stages of vision involve multiscale filtering, it is somewhat surprising that its importance for brightness/lightness perception has yet to be fully appreciated. I believe one reason for this is a wrong idea that has become entrenched over the years, namely that it is the contrast (or ratio) *in the immediate vicinity of the edge* that is critical to brightness/lightness perception. This idea follows from one of the most enduring themes in the recent history of this topic, namely that the visual system first locates edges, and then "fills-in" the gaps between them by some kind of spreading of neural activity (Ratliff, 1972; Gilchrist, 1979; Grossberg and Todorovic, 1988; see review by Kingdom and Moulden, 1988; see discussion by Blakeslee and McCourt, this volume). In this view, the luminance relationships between those parts of a stimulus that lie at a distance from the edge exercise little influence on brightness. An almost anecdotal but nevertheless striking demonstration of the importance of distal luminance relationships is illustrated in Figure 3.10. If brightness perception is critically dependent on the luminance relationships at the edge, then it must follow that blurring the edge should *at the very least* reduce the magnitude of any perceived brightness variations. Yet the opposite is found. As can be seen in Figure 3.10b, blurring the edges if anything increases the magnitude of SBC, and this has been confirmed quantitatively by McCourt and Blakeslee (1993) using the grating induction stimulus. On the other hand selective removal of the low spatial frequencies, which define the more distal luminance relationships, substantially reduces SBC, as shown in Figure 3.10c.

This is not to say that the luminance relationships at the edge play no role in brightness perception. The Craik-Cornsweet-O'Brien illusion (e.g. Cornsweet, 1970; Todorovic, 1987; see a weak version of the illusion in Figure 3.10c), in which an illusory brightness difference is observed on either side of a highpass-filtered (or equivalent) edge, suggests that an edge-based filling-in mechanism may contribute to brightness, perhaps even playing a crucial role in the perceived uniformity of physically uniform regions. The point being made here, and Figure 3.10 seems persuasive evidence, is that the role played by an edge-based filling-in mechanism is probably quite minor.

There are important ramifications to the idea that lightness/brightness perception is a multi-scale process. An often-heard refrain against contrast theories of brightness/lightness is that two patches with the same luminance and edge contrast can nevertheless appear very different in brightness/lightness (e.g. see the discussion of Figure 3.1; Gilchrist, 1979; Gilchrist et al., 1999). However, once we accept the idea that contrast-sensitive mechanisms operate at multiple spatial scales, we cannot reject an explanation base on contrast merely because of what happens at the edge. We must also consider the distal luminance relationships. Bearing in mind this caveat, let us now turn to a consideration of those brightness/lightness phenomena that appear to defy explanation in terms of contrast.

3.5 Helmholtz and the illumination-interpretive approach

In previous sections I considered the evidence for a contrast-sensitive mechanism based on multi-scale filtering. One purpose of such a mechanism is to achieve lightness constancy with respect to the ambient level of illumination (Whittle, 1994a,b). By tying lightness to contrast, lightness becomes invariant to changes in light level. There is a penalty however: errors such as SBC.

In this section I examine the component of brightness/lightness perception that is thought to be involved in discounting spatial, as opposed to ambient changes in illumination such as shadows, highlights, shading and transparency. Although the last of these, transparency, is a material property, its luminance relationships are identical to those of shadows. The distinction between lightness and brightness becomes very important when considering spatially varying illumination. Con-

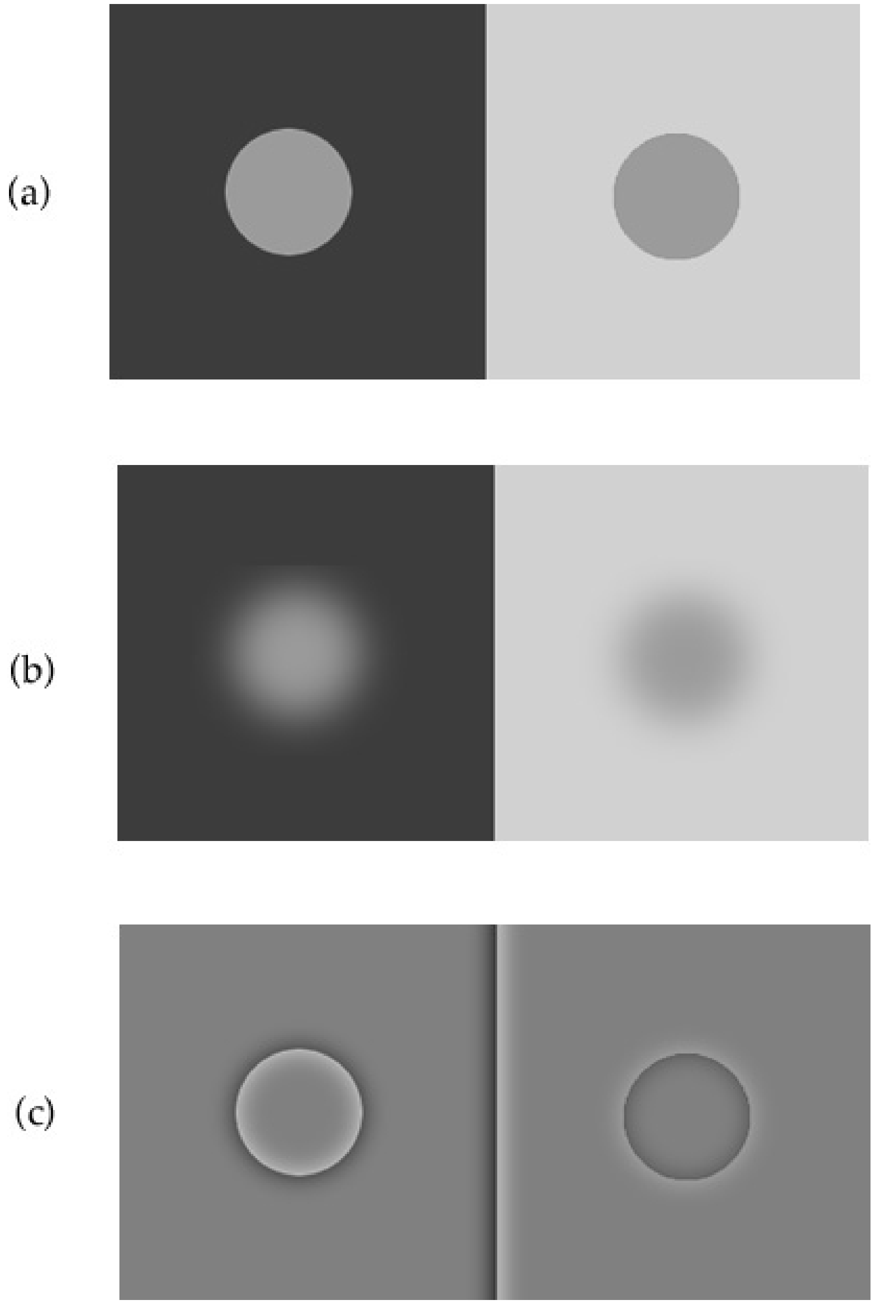

FIGURE 3.10. Simultaneous brightness contrast (top) is slightly enhanced when the stimulus is lowpass filtered (middle), but diminished when highpass filtered (bottom).

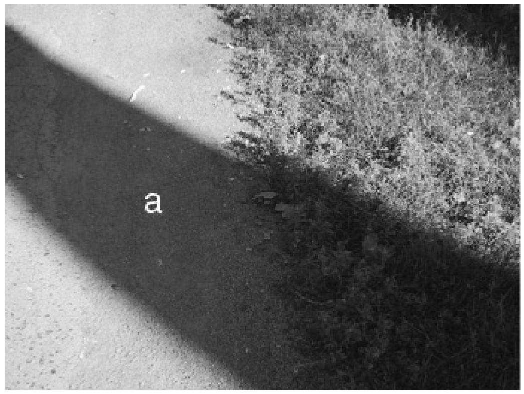

FIGURE 3.11. Natural shadow

sider for example the natural scene in Figure 3.11. Two judgements concerning the shadowed region at **a** can be made. On the one hand we observe that it is darker than its surround - a relative brightness judgement. On the other hand we infer that it is the same shade of grey as its surround - a relative lightness judgement. While it is conceivable that a clever artist might have painted the grass and road with dark paint to simulate the effect of a shadow, which would make our lightness judgement wrong, this is not our impression.

When discussing the effects of spatial variations in illumination on brightness/lightness, it is instructive to begin with Helmholtz. Helmholtz (1866/1962) mainly considered the chromatic version of SBC, simultaneous colour contrast (SCC), in which a grey patch appears tinted with the complementary colour of its surround. Helmholtz believed that all forms of SCC resulted from 'errors of judgement'. In some cases SCC occurred because of the mistaken assumption that the grey patch was covered by a transparent veil the colour of the surround, the eye compensating for the veil when estimating the colour of the patch (pp. 282-87). Helmholtz was influenced by an earlier experiment conducted by Heinrich Meyer in 1855. Meyer had shown that the red tinge seen in a grey patch on an intense green background became even redder when both were overlaid with a piece of transparent white paper, which had the effect of desaturating the green background. Helmholtz suggested that the overlay of transparent white paper helped create the illusion that the grey patch was being viewed through a green veil. However, because the eye received from the grey patch a composition of light normally associated with grey, an inference was made that the patch must be pinkish, as the effect of the green veil would be to absorb the long wavelengths associated with the pinkish tint. Thus according to Helmholtz, we have learned to 'correct' for the effects of intervening, transparent media, just as we have learned to 'correct' for the prevailing illumination in assessing the intrinsic lightness of objects. A Helmholtzian account of achromatic SBC would be based on an analogous argument, as illustrated in Figure 3.2. We assume that the patch on the bright background is more intensely illuminated than the one on the dark background. However, because the intensity of light reaching the eye is the same for both patches, an inference is made that the patch on the bright background must be of lower reflectance, and that is how it is perceived.

Helmholtz believed that other types of judgement error were also involved in SCC (e.g., see Helmholtz, pp. 274-278; also Turner, 1994, pp. 108-113, for a recent review), but it is interesting that it is his veiling hypothesis that William James seized upon when discussing of SCC in his classic work The Principles of Psychology. James criticized Helmholtz's veiling hypothesis because SCC occurred under conditions where it was quite implausible to suppose that the test regions were differently illuminated (James, 1890/1981, pp. 662-674). For example, James describes how a pinkish tinge can be seen in grey concentric rings that alternate with green concentric rings, yet one has no impression that any one part of the stimulus is differently illuminated from any other. James's argument is an important one because it suggests not only that there are other explanations for SCC besides the veiling hypothesis, but that one needs to have visible illumination borders before entertaining what I refer to here as an 'illumination- interpretative' explanation of SCC. Notwithstanding James's critique, it is the way Helmholtz's veiling hypothesis anticipated the remarkable series of demonstrations alluded to in the Introduction and now considered in more detail that makes his ideas so prescient.

Figure 3.12 is my own figure that was inspired by Adelson's (unpublished) checkerboard-shadow illusion. I multiplied a black-white checkerboard by a low amplitude, single cycle of near-sinusoidal luminance modulation, such that the luminance of the dark square at **a** in the bright shaded region is identical to the light square at **b** in the dark shaded region. In an important sense this figure is a brightness and not a lightness illusion. The checks **a** and **b** look different in brightness, yet have the same luminance. However, once we attribute the slowly varying luminance component of the figure to shading, we are correct to judge the lightnesses of **a** and **b** as different, even though physically on the page they are the same. The illusion appears to demonstrate our ability to parse the image into its illumination and reflectance components, or its 'intrinsic images' (Bergstrom, 1977; Barrow and Tenenbaum, 1978; Adelson and Pentland, 1996).

What is striking about Figure 3.12 is the way our brightness perception appears to be so dominated by our lightness perception. It is as if in discounting the shading we ceased to be aware of its presence altogether, and as a result conclude that **a** and **b** must be different in brightness and not just lightness. Our 'intrinsic image' processing seems to work well for lightness, but fails for brightness perception. One can legitimately argue that the goal of the system is lightness constancy, and thus brightness per se is unimportant. Be that as it may, observers often express incredulity when told that **a** and **b** have the same luminance (or told that they are the same shade of grey), which suggests that at the very least they feel they ought to be able to correctly judge their relative brightnesses.

The illusion in Figure 3.12 is strongly suggestive of the involvement of a Helmholtzian, lightness constancy mechanism that discounts spatially varying illumination, i.e. is 'illumination-interpretative'. But before jumping to this con-

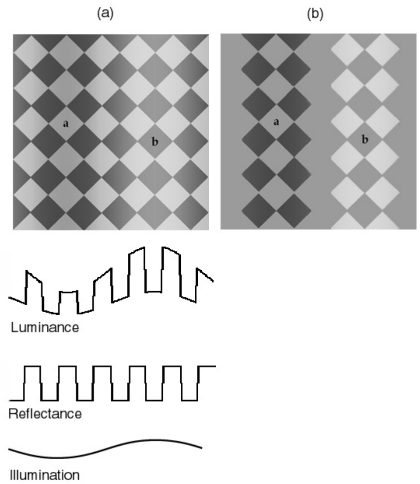

FIGURE 3.12. (a) Checkboard-shading illusion, similar to Adelson's (unpublished) checkerboard-shadow illusion. **a** and **b** are identical in luminance, as shown in the luminance profile. The image can be decomposed into its "intrinsic images": a reflectance and illumination layer, as illustrated below. In (b) the three columns of diamonds centred respectively on **a** and **b** have been placed on a background of the same luminance as **a** and **b**. The brighness difference between **a** and **b** is markedly reduced.

clusion, we must be careful. Lest we missed the fact, **a** and **b** are surrounded by different luminances. Is the brightness illusion really illumination-interpretative, or is it simply a result of contrast? To answer this question we must demonstrate that the pictorial representation of shading enhances the brightness illusion over and above that due to contrast, and for this we need a 'control' stimulus with the same pattern of luminance, but without the impression of shading. But here lies the rub. To remove the impression of shading we must change the arrangement of luminances. Can we be certain when doing this that we have not inadvertently altered contrast, bearing in mind what was said in the previous section about the importance of distal luminance relationships in contrast processing? In Gilchrist's classic experiment illustrated in Figure 3.1b, the highlight increases the area around the test square that is surrounded by a high luminance, and this alone might have caused the patch to appear darker than in the standard display (1a), irrespective of how the surround was interpreted. Consider also my attempt in Figure 3.12b to provide a suitable control. The three columns of diamonds centred on **a** and **b** have been placed on a background of uniform grey the same luminance as **a** and **b**. The impression of shading disappears, and the brightness difference between **a** and **b** is unquestionably reduced. Although it is therefore difficult to see how contrast alone could account for the dramatic reduction in the strength of the illusion, even when taking into account the distal luminance relationships in the figure, one can not be certain. The point being made here is not that illumination-interpretative processes do not influence brightness, on the

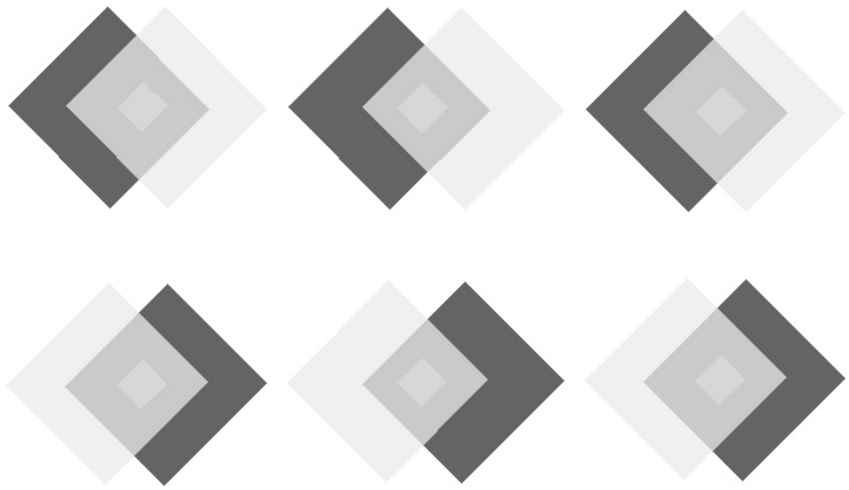

FIGURE 3.13. Effect of transparency on brightness. When free-fused one sees four figures: two consist of a bright transparency in front of a dark background, two of a dark transparency in front of a bright background. Equal-in-luminance test diamonds lie either on a transparency or on a background. For most observers the grey patch on the bright background behind the dark transparency looks brighter than the other test diamonds. Note that the pattern of luminances surrounding all test diamonds is near-identical in the monocular view.

contrary, but that one must be very careful before rejecting an explanation based on contrast.

Ideally, what one would like are two stimuli whose test regions are surrounded by near-identical patterns of luminance, but whose perceived pattern of illumination is nevertheless very different. Such a stimulus would then *isolate* the putative illumination-interpretative mechanism from the effects of contrast. I think Figure 3.13 goes some way towards achieving this. When free-fused, one sees four figures in stereoscopic depth, each consisting of a simulated transparency in front of a background, with equal-in-luminance test diamonds either on the transparency or on the background. The pattern of luminances surrounding the test diamonds is more-or-less identical, at least in the monocular view of all four figures, and importantly not just at the edges of the test diamonds. Most observers agree that the test diamond on the white background behind the dark transparency looks both brighter and lighter than the others. This is in keeping with the Helmholtzian idea that the lightness attributed to the test diamond is what it would be if the transparency was removed, with the added observation that the brightness of each test diamond is strongly influenced by its lightness.

In a recent experiment, Barbara Blakeslee, Mark McCourt and I measured the brightness of test patches perceived to lie either behind a simulated transparency, as in Figure 3.13, or on a reflectance background with a near-identical pattern of surround luminance (Kingdom, Blakeslee, and McCourt, 1997). We found that the perception of transparency did effect brightness in the expected direction, though in general the effects were quite small (the biggest effect we found was about a factor of two). Thus in spite of the concerns expressed above about the potentially confounding effects of contrast, our experiment confirmed the findings of Gilchrist (1979), Adelson (1993), Logvinenko (1999) and others, and provided additional evidence for an illumination-interpretative component of brightness and lightness perception.

3.6 Integration and anchoring

I have now argued that two mechanisms contribute to brightness/lightness perception, a low-level contrast-sensitive, and a mid-level illumination-interpretative mechanism. Other processes are presumably also involved, and a few remarks will be made about just two of them.

The first is integration. If contrast-sensitive mechanisms operate locally, some method of combining their signals across the image may be necessary, one purpose being to compare brightesses/lightnesses across a distance. Since contrast is a differencing operation, the putative mechanism, if it exists, is arguably analagous to mathematical integration (see Kingdom and Moulden, 1988; Gilchrist, 1994; Whittle, 1994b; Arend, 1994 for reviews). Whittle (1994b) has suggested that an important function of integration is to achieve lightness constancy with respect to the surround, so that surfaces viewed against different backgrounds do not appear to differ in lightness. Whittle refers to this type of lightness constancy as Type II, as distinct from Type I, which is constancy with respect to the ambient level. Whittle includes constancy with respect to spatially varying illumination, the illumination-interpretative constancy mechanism I described in the previous section, Type II. The putative integration stage would work in the opposite direction to contrast, serving to mitigate its effects and derive a more veridical representation (Whittle, 1994b). That SBC exists at all is testament to the fact that such a mechanism is, however, unable to fully override the effects of contrast. Traditionally, the integration stage has been assumed to operate on edge contrasts, or ratios, perhaps most famously in the Retinex model of Land and McCann (1971). However Land (1986) and Hurlbert and Poggio (1988) have suggested that lightness constancy with respect to spatial variations in illumination might be achieved directly via the use of filters with small receptive field centres and much larger receptive field surrounds, without need for an explicit integration stage. This raises the tantalising possibility that what appears to be integration might in fact be large-scale filtering.

The second process that deserves to be mentioned is ‘anchoring’, the term coined by Gilchrist et al. (1999) for the mechanism that turns relative lightness judgements into absolute ones. In mathematical terms, anchoring is traditionally associated with the restoration of the d.c. level. Gilchrist et al. have provided evidence that relative lightness values are anchored to the highest luminance in the display, which is ascribed white, an idea suggested by Wallach (1976) and incorporated into models of lightness constancy such as the Retinex (Land and McCann, 1971). Moreover, Gilchrist et al. suggest that anchoring is itself responsible for errors such as SBC, because it operates not only globally but also locally within different perceptual frameworks. This is an interesting idea worth pursuing. In the debate that accompanies this chapter I describe how the anchoring model predicts SBC in the standard display, and offer a critical appraisal of the model's plausibility when applied to other types of SBC, such as grating induction.

3.7 Conclusions

Much can be learnt about how we perceive brightness and lightness from the errors we make when doing so. Brightness and lightness perception involve a number of mechanisms operating at different levels of visual processing. One mechanism is low- level, and processes spatial variations in brightness via multiscale filtering. It serves to achieve lightness constancy with respect to the ambient level of illumination. However it comes at a cost: errors such as simultaneous brightness contrast. A second, mid-level mechanism aims to achieve lightness constancy with respect to spatially varying illumination such as shading, shadows, highlights and transparency. The cost in this case is an enhancement of errors in brightness judgement.

Acknowledgements

I would like to thank Beverley Mullings for her valuable help with many of the figures, and Alan Gilchrist for helpful comments regarding terminology

References

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

Adelson, E. H. and Pentland, A. P. (1996). Lightness Perception and lightness illusions. In M. Gazzaniga (Ed.), *The Cognitive Neurosciences (2nd ed.)*, Cambridge, MA: MIT Press.

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

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

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

Blakeslee, B. and McCourt, M. E. (1997). Similar mechanisms underlie simultaneous brightness contrast and grating induction. *Vis. Res.*, 37: 2849-2869.

Blakeslee, B. and McCourt, M. E. (1999). A multiscale spatial filtering account of the White effect, simultaneous brightness contrast and grating induction. *Vis. Res.*, 39: 4361-4377.

Blakeslee, B. and McCourt, M. E. (2001a). A multiscale spatial filtering account of the Wertheimer-Benary effect and the corrugated Mondrian. *Vis. Res.*, 41: 2487-2502.

Blakeslee, B. and McCourt (2001b). A multiscale spatial filtering account of brightness perception". In L. R. Harris and M. Jenkin (Eds.), *Levels of Perception*, Springer Verlag: New York, NY.

Bergstrom, S. S. (1977). Common and relative components of reflective light as information about the illumination, colour, and three-dimensional form of objects. *Scandinavia J. Psych.*, 18: 180-186.

Campbell, F. W. and Kulikowski, J. J. (1966). Orientation selectivity of the human visual system. *J. Physiol. Lond.*, 187: 437-445.

Cornsweet, T. (1970). *Visual Perception*. Academic Press: New York, NY.

Fiorentini, A., Baumgartner, G., Magnusson, S., Schiller, P. H. and Thomas, J. P. (1990). The Perception of Brightness and Darkness". In L. Spillman and J. S. Werner (Eds.), *Visual perception: The Neurophysiological Foundations*. Academic Press: San Diego, CA.

Foley, J. M. and Legge, G. E. (1981). Contrast detection and near-threshold discrimination in human vision. *Vis. Res.* , 21: 1041-1053.

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

Gilchrist, A. L. (1979). The perception of surface blacks and whites. *Sci. Am.*, 240: 112-123.

Gilchrist, A. L. (1988). Lightness contrast and failures of lightness constancy: a common explanation. *Percept. and Psychophys.*, 43: 415-424.

Gilchrist, A. L. (1994). Absolute versus relative theories of lightness perception. In A. L. Gilchrist (Ed.), *Lightness, Brightness, and Transparency*, pp. 1-33. Hillsdale: Erlbaum.

Gilchrist, A. L., Delman, S. and Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination. *Percept. and Psychophys.*, 33: 425-436.

Gilcrhist, A. L. and Economou, E. (2001). Dualistic versus monistic accounts of lightness perception. In L. R. Harris and M. Jenkin (Eds.), *Levels of Perception*, in press.

Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V. and Economou, E. (1999). An anchoring theory of lightness perception. *Psych. Rev.* 106: 795-834.

Grossberg, S. and Todorovic, D. (1988). Neural dynamics of 1-D and 2-D brightness perception: A unified model of classical and recent phenomena. *Percept. and Psychophys.*, 43: 241-277.

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

Hering, E. (1964). *Outlines of a Theory of the Light Sense*, L. M. Hurvich and D. Jameson, Trans. Harvard University Press: Cambridge, Massachusetts. (Original work published in 1874).

Hurlbert, A. and Poggio, T. (1988). Synthesizing a color Algorithm from examples. *Science*, 239: 484-485.

Hurvich, L. M. and Jameson, D. (1966). *The Perception of Brightness and Darkness*. Allyn and Bacon, Inc: Boston.

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

Kingdom, F. (1997). Simultaneous contrast: the legacies of Hering and Helmholtz. *Percept.*, 26: 673-677.

Kingdom, F. (1999). Commentary on Logvinenko “Lightness induction revisited”. *Percept.*, 28: 929-934.

Kingdom, F. McCourt, M. E. and Blakeslee, B. (1997). In defence of ”lateral inhibition” as the underlying cause of induced brightness phenomena. A reply to Spehar, Gilchrist and Arend. *Vis. Res.*, 37: 1039-1044.

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

Kingdom, F. and Moulden, B. (1988). Border effects on brightness: A review of findings, models and issues. *Spat. Vis.*, 3: 225-262.

Kingdom, F. and Moulden, B. (1992). A multi-channel approach to brightness coding. *Vis. Res.*, 32: 1565-1582.

Kingdom, F. and Whittle, P. (1996). Contrast discrimination at high contrasts reveals the influence of local light adaptation on contrast processing. *Vis. Res.*, 36: 817-829.

Knill, D. C. and 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. *Proc. Nat. Acad. Sci. USA*, 83: 3078-3080.

Land, E. H. and McCann, J. J. (1971). Lightness and retinex theory. *J. Opt. Soc. Am.*, 61: 1-11.

Logvinenko, A. D. (1999). Lightness induction revisited. *Percept.*, 28: 803-816.

McCourt, M. E. (1982). A spatial frequency dependent grating induction effect. *Vis, Res,*, 22: 119-134

McCourt, M. E. and Blakeslee, B. (1993). The effect of edge blur on grating induction magnitude. *Vis, Res.*, 33: 2499-2508.

McCourt, M. E. and Kingdom, F. A. A. (1996) ”Facilitation of luminance grating detection by induced gratings”. Vision Research, 36, 2563-2573.

Moulden, B. and Kingdom, F. (1991). The local border mechanism in brightness induction. *Vis. Res.*, 31: 1999-2008.

Ratliff, F. (1972) Contour and contrast. Scientific American, 226, 91-101.

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

Todorovic, D. (1987). The Craik-O’Brien-Cornweet effect: New varieties and their theoretical implications. *Percept. and Psychophys.*, 42: 545-560.

Turner, R. S. (1994). *In the Eye’s Mind Vision and the Helmholtz-Hering Controversy*. Princeton University press: Princeton, NJ.

Wade, N. J. (1996). Descriptions of visual phenomena from Aristotle to Wheatstone. *Percept.*, 25: 1137-1175.

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

Whittle, P. (1965). Binocular rivalry and the contrast at contours. *Quar. J. Exp. Psych.*, 17: 217-226.

Whittle, P. (1986). Increments and decrements: Luminance discrimination. *Vis. Res.*, 26: 1677-1691.

Whittle, P. (1993). Brightness, discriminability and the “crispening effect”. *Vis. Res.*, 32: 1493-1507.

Whittle, P. (1994a). The Psychophysics of Contrast Brightness. In Gilchrist (Ed.) *Lightness, Brightness, and Transparency*, pp. 35-110. Hillsdale: Erlbaum.

Whittle, P. (1994b). Contrast Brightness and Ordinary Seeing. In Gilchrist (Ed.) *Lightness, Brightness, and Transparency*, pp. 111-157. Hillsdale: Erlbaum.