
Colour vision brings clarity to shadows

Frederick A A Kingdom, Catherine Beauce, Lyndsay Hunter

McGill Vision Research Unit, Department of Ophthalmology, McGill University, 687 Pine Avenue West, Room H4-14, Montréal, Québec H3A 1A1, Canada; e-mail: fred@jiffy.vision.mcgill.ca

Received 11 December 2003, in revised form 27 March 2004; published online 30 June 2004

Abstract. We have revealed a new role for colour vision in visual scene analysis: colour vision facilitates shadow identification. Shadows are important features of the visual scene, providing information about the shape, depth, and movement of objects. To be useful for perception, however, shadows must be distinguished from other types of luminance variation, principally the variation in object reflectance. A potential cue for distinguishing shadows from reflectance variations is colour, since chromatic changes typically occur at object but not shadow boundaries. We tested whether colour cues were exploited by the visual system for shadow identification, by comparing the ability of human test subjects to identify simulated shadows on chromatically variegated versus achromatically variegated backgrounds with identical luminance compositions. Performance was superior with the chromatically variegated backgrounds. Furthermore, introducing random colour contrast across the shadow boundaries degraded their identification. These findings demonstrate that the visual system exploits inbuilt assumptions about the relationships between colour and luminance in the natural visual world.

1 Introduction

Shadows are ubiquitous features of natural scenes. Often termed ‘cast shadows’, they occur wherever objects occlude light. It is useful to distinguish shadow from shading, which is the variation in luminance due to a change in the angle of a surface with respect to the direction of illumination. Gilchrist and colleagues (Gilchrist 1979; Gilchrist et al 1983) were one of the first groups of researchers to demonstrate the importance of being able to distinguish shadows from changes in surface reflectance, a process they referred to as “edge classification”. They showed that the perceived reflectance, or lightness of a surface changed when its immediate surround was perceived to lie in shadow, as opposed to being surrounded by a lower reflectance. More recently, the importance of shadow perception to vision has been demonstrated in tasks such as object recognition (Cavanagh and Leclerc 1989) and motion perception (Knill et al 1996).

Shadows are primarily luminance-defined features; although sometimes tinged with colour (for example blue when formed in sunlight) shadows tend to have minimal colour contrast (Parraga et al 2002). Objects, on the other hand, typically vary in colour as well as lightness, that is in spectral as well as intensive reflectance. These relationships are illustrated in the photograph in figure 1, which shows a shadow falling across a grass/pavement border. The shadow is primarily a change in luminance (bright to dark) whereas the grass/pavement border is a change in both colour (green to grey) and luminance (dark to light). Colour is therefore a potential cue for helping disambiguate shadows from reflectance changes via the following rule: luminance variations that are accompanied by colour variations are variations in reflectance, whereas luminance variations that are unaccompanied by colour variations are variations in illumination (Rubin and Richards 1982; Cavanagh 1991; Mullen and Kingdom 1991; Tappen et al 2003). In the image-processing domain, this rule has recently been applied to help separate the shading and reflectance components of natural images (Tappen et al 2003; Olmos and Kingdom 2004), and there is evidence that the rule is exploited by the human visual system when determining the shape of a surface from the pattern of its shading (Kingdom 2003).

A hint that the visual system is sensitive to the aforementioned colour–luminance relationships in the context of shadow perception is revealed when one examines scenes where only a limited number of cues are available to help parse the image into reflectance and shadow. A good example is the photograph in figure 2. The insets show a shadow crossing a painted edge (the crossing of two luminance borders is often termed an ‘X junction’). When the inset is viewed in colour, most observers regard the chromatic border as a change in reflectance and the near-pure luminance border as a shadow. However, the situation is more ambiguous when the same inset is viewed in black-and-white.

Figure 2 appeals to subjective impression. Is there a more objective indicator of the modulatory impact of colour on shadow perception? We have attempted to answer this question with a psychophysical experiment that compares the ability of human test subjects to discern simulated shadows on chromatically variegated versus achromatically variegated backgrounds. Our stimuli are termed ‘6-luminance’ displays (Kasrai and Kingdom 2001), and examples are shown in figure 3. Each comprises three ‘background’ sectors and a central ‘shadow’. The shadows in our displays may also be regarded as simulated achromatic transparencies (or neutral density filters)—indeed some look more like transparencies than shadows. An achromatic transparency, although a material medium, is like a shadow in that, when overlaid onto a surface, it divides all the surface luminances by the same amount according to its transmissivity (an achromatic transparency should be distinguished from a translucent material which has an additive, as well as divisive luminance component). The difference between an achromatic transparency and a shadow is that shadows tend to have penumbra, or blurred edges (though see figure 2 for a counterexample). However, because the luminance relationships between the uniform parts of a shadow and its associated background is the same as for an achromatic transparency, we see no reason why the results of our study would be any different if our simulated shadows possessed penumbra. Therefore we will refer to the simulated overlays in figure 3 as shadows.

2 Method

2.1 Stimuli

2.1.1 *Generation.* All stimuli were generated with the VSG2/3F video-graphics card (Cambridge Research Systems) hosted by a Gateway 2000 P5 computer, and displayed on a BARCO Calibrator monitor.

2.1.2 *Stimulus backgrounds.* Each stimulus was 4.6 deg in diameter positioned in the middle of the screen, and consisted of three equal-area background sectors subtending 120 deg, with the leading edge of the first sector randomised in orientation, with a simulated shadow overlaying the centre. The three conditions in figures 3a to 3c are ‘achromatic’, ‘chromatic’, and ‘chromatic all-border’. For every stimulus, the luminances of the three background sectors were randomly drawn from a distribution that was the same across all three conditions (see below for details). This is a critical property of our stimuli. If the luminances of the background sectors were, on average, different for the different conditions, then any measured differences in shadow identification between conditions could not unequivocally be attributed to differences in their chromatic content. The colours of the background sectors in the chromatic and chromatic all-border conditions were drawn randomly from the gamut available on the monitor. In the chromatic condition, the colours were the same both inside and outside the shadow, whereas in the chromatic all-border condition, the colours on either side of the shadow border were randomly selected.

The hues and luminances of the background sectors were allocated as follows. For each stimulus, three image planes, referred to as R, G, and B, were created in the VSG's video memory. The R, G, and B planes were alternated in sequence at 180 Hz, and thus each plane was refreshed at 60 Hz. While the R plane was displayed, only the red monitor phosphor was activated; for the G plane only the green phosphor was activated; and for the B plane only the blue phosphor was activated. There were 256 linearly spaced intensity levels available for each R, G, and B plane, resulting in a total of 256^3 possible colours. Let each colour be described as a 'tristimulus' value, rgb , where r , g , and b are the intensities of the RGB colours expressed as proportions of 256 (ie 0–1). For the chromatic background, each rectangle was randomly allocated a tristimulus value, and hence hue and luminance, from the full range available. If the resulting tristimulus value for a given rectangle in the chromatic condition is given by $r_c g_c b_c$, the luminance L_c of the rectangle is:

$$L_c = r_c R_{\max} + g_c G_{\max} + b_c B_{\max} \quad (1)$$

where R_{\max} , G_{\max} , and B_{\max} are the maximum luminances of the R, G, and B planes (the maximum luminance was measured by setting the other two planes to zero luminance and presenting the frame-alternating sequence in the same way as for the stimulus). For the achromatic condition, the distribution of background luminances was the same as for the chromatic condition. To achieve this, we first generated random tristimulus values (r_c , g_c , b_c) as for the chromatic condition, and then calculated the achromatic tristimulus values r_a , g_a , b_a that gave the same luminance L_a but under the constraint that $r_a = g_a = b_a$. The formula was:

$$r_a = g_a = b_a = \frac{r_c R_{\max} + g_c G_{\max} + b_c B_{\max}}{R_{\max} + G_{\max} + B_{\max}}. \quad (2)$$

This method for equating the luminance distribution of the chromatic and achromatic background sectors is robust to any variations in spectral sensitivity between observers. Suppose for example that R_{\max} , as measured by our photometer, underestimated the sensitivity of a subject to the R plane by a factor of two. Doubling R_{\max} has no effect on the equality between L_c and L_a , nor therefore on the equality of shadow contrast between the achromatic and chromatic conditions.

2.1.3 Shadows. The simulated shadow in the centre was a circular patch 2.3 deg in diameter. There were two types of shadow: 'correct' and 'incorrect'. A correct shadow was one in which all three shadow sectors were allocated a contrast of half ($\times 0.5$) the luminance of their associated background sector, as in figures 3a to 3d. The notion of 'correctness' here is based on the assumption that in natural scenes the effect of a shadow is to divide the luminance of the shadowed regions by a more-or-less constant amount. An incorrect shadow was generated by setting two of the sectors to a contrast of $\times 0.5$, and the third, 'odd', sector to one of 11 contrasts chosen from an equal-interval range spanning $\times 0.025$ to $\times 0.975$. The range of odd-sector contrasts thus included contrasts that were both decremental as well as incremental with respect to the $\times 0.5$ shadow sectors, though of course all odd-shadow sectors were decremental with respect to their associated backgrounds. Thus, in the $\times 0.025$ odd-sector stimulus, one of the sectors appeared unnaturally dark, while in the $\times 0.975$ odd-sector stimulus, it appeared unnaturally bright. An example of an incorrect shadow whose odd-sector contrast was $\times 0.025$ is shown in figure 3d. For the chromatic all-border condition, the tristimulus values of the three shadow sectors were randomly generated anew before being multiplied by the appropriate t values.



(a)



(b)

Figure 1. Natural shadow in (a) colour, and (b) black-and-white.



(a)



(b)

Figure 2. Photograph in (a) black-and-white and (b) colour. Insets show the X junction towards the top-right of the figure. The colour photograph is taken from J Marvullo (1989) *Color Vision: A Photographer's Guide* (New York: Watson–Guptill Publications) page 58, reproduced with permission of the author.




Figure 3. (a–c) Main conditions and (d, e) example task for shadow-identification experiment. The conditions are (a) achromatic, (b) chromatic, and (c) chromatic all-border. (d) and (e) show an example forced-choice pair in the chromatic condition. In (d) one of the three shadow sectors is set to a higher negative contrast ($\times 0.025$) with respect to its background than are the other two sectors ($\times 0.5$) making the shadow ‘incorrect’. In all the other patterns the three shadow sectors have been set to the same contrast ($\times 0.5$), and are ‘correct’. Subjects were required to indicate on each trial which of the two shadows, (d) or (e), was the correct shadow.

2.2 Subjects

Six subjects were tested, the three authors (LH, CB, and FK) and three undergraduate volunteers (MY, KH, and HW). The volunteers were naïve as to the purpose of the experiment. All subjects had normal, or corrected-to-normal acuity, and normal colour vision.

2.3 Procedure

A two-interval forced-choice procedure with the method of constant stimuli was employed. On each trial, two stimuli were presented, one with a correct shadow and one with an incorrect shadow, and the subject was required to indicate, by a button press, which was the correct shadow. Feedback was given in the form of a tone for an incorrect decision. The stimuli were presented for 500 ms. There were 110 trials per session with the 11 values of ($\times 0.025 - \times 0.975$) presented 10 times each, in random order.

2.4 Data analysis

Weibull functions were fitted to the proportion of correct data by the formula:

$$1.0 - 0.5 \exp(-x/a)^b,$$

where x is the independent variable and a and b are free parameters determining, respectively, the threshold at the 81% correct level and the slope of the psychometric function.

3 Results

Figure 4a shows example psychometric functions for one subject, KH. Proportion of correct responses is plotted as a function of the absolute difference in contrast between the $\times 0.5$ and odd-shadow sectors, $|\Delta C|$, with data collapsed across both incremental and decremental odd sectors. As can be seen, as the odd-sector contrast increasingly differs from $\times 0.5$, the proportion of correct identifications increases, ie the task becomes easier. The smooth curves are Weibull-function fits to the data (see section 2), and as can be seen the curves do not overlap. Figure 4b shows threshold values of $|\Delta C|$ calculated from the Weibull fits at the 81% correct level for each subject and each condition. As the figure shows, the ordering of thresholds from lowest to highest was chromatic < achromatic < chromatic all-border, for all six subjects tested. Correlated-sample one-tailed t -tests show that the lower thresholds of the chromatic compared to the achromatic condition is significant ($t_5 = 0.056$, $p < 0.005$), and the higher thresholds of the chromatic all-border condition compared to the achromatic condition is significant ($t_5 = 0.03$, $p < 0.05$).

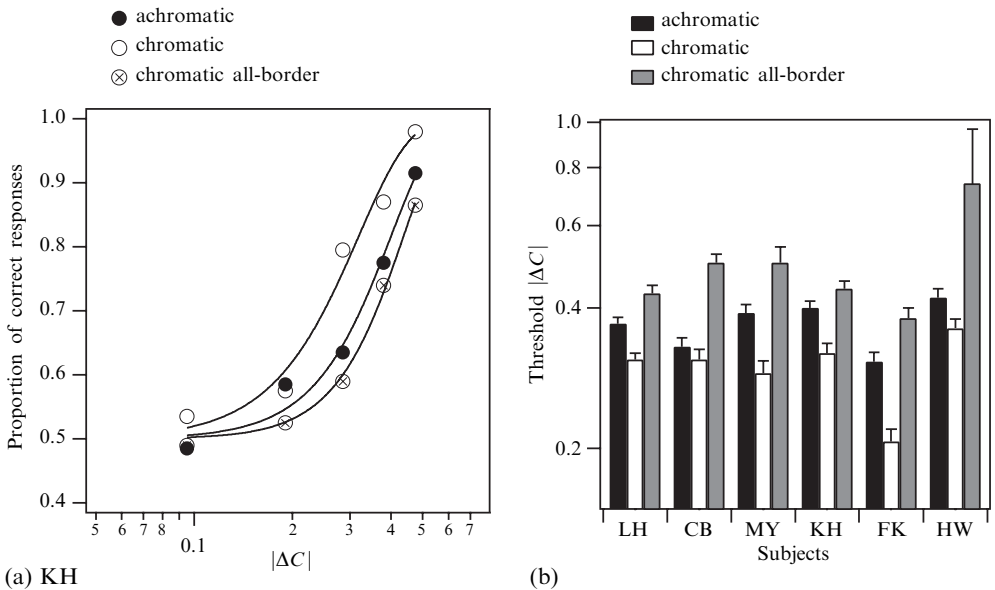


Figure 4. (a) Results and Weibull fits for one subject. The proportion of correct responses is plotted against the absolute difference $|\Delta C|$ between the $\times 0.5$ and odd-shadow sector contrasts, averaged across positive and negative values of ΔC . The three sets of data and fits are for the three main conditions. (b) Thresholds for six subjects obtained from the Weibull fits.

4 Discussion

We have found that chromatic variations that are nonaligned with shadow borders facilitate their identification, whereas chromatic variations that are aligned with shadow borders suppress their identification. There is nothing inevitable about these results. The chromatic variations in the two chromatic display conditions had no impact on their information content. Although in all conditions there was stochastic variability in the selection of the background luminances, and in the case of the chromatic displays also their colours, there was no stochastic variability in the luminance ratios of the shadow sectors to their respective backgrounds. And it was in the discernment of the pattern of these luminance ratios that the task was defined. Thus for an ideal observer, ie one with access to all the information in the display, performance would have been perfect on all trials in all conditions. Therefore any differences in performance between the conditions must have been solely due to limitations in visual processing.

In general, we do not confuse shadows with the objects from which they arise. Although it is conceivable that the shadows in figures 1 and 2 could have been painted by a clever artist, that is not our impression. The experimental findings of this study suggest that one of the cues used by the visual system to help identify shadows is colour. Although shadows are primarily luminance-defined features, their perception appears to be significantly impacted by colour, and in a manner consistent with the idea that the human visual system has inbuilt assumptions about the origin of colour–luminance relationships found in natural scenes. Specifically, chromatic variations, and luminance variations that are spatially aligned with them, are assumed to arise from surfaces, whereas pure or near-pure luminance variations are assumed to arise from inhomogenous illumination.

Cavanagh and Leclerc (1989), using a task in which subjects rated their ability to recognise objects such as faces and cups defined solely by shadows and/or shading, failed to find evidence that shape-from-shadows was impaired when colour contrast was introduced across shadow borders. It may be that form judgments based on shadows are less susceptible to the negative influence of aligned chromatic variations than are shadow judgments. This may be a property unique to shadows, since shape-from-shading has recently been shown to be highly susceptible to the negative impact of aligned chromatic variations (Kingdom 2003). Further research is needed to establish precisely the conditions under which form judgments based on illumination are affected by colour.

It has been suggested that it would make good sense for the visual system to suppress luminance borders in favour of chromatic ones, because chromatic borders are more reliable indicators of object boundaries (Switkes et al 1988). Evidence from the detection of sinusoidal luminance gratings in the presence of chromatic contrast masks has provided some support to this idea (Switkes et al 1988). The results of the present study, however, suggest that there are circumstances when the presence of colour contrast can facilitate the detection and identification of certain luminance features. In the case of shadows this makes good sense, since shadows, correctly identified, can be used by the visual system for a variety of form and motion tasks. This positive role of colour vision in the perception of shadows complements the well-attested benefits of colour vision in the analysis of other aspects of image structure, for example in the detection of fruit and young leaves in dense foliage (Mollon 1989; Sumner and Mollon 2000; Dominy and Lucas 2001). It seems unlikely, however, that colour vision would have evolved primarily to help animals discern shadows, and is more likely an ancillary benefit with regard to the primary advantage of making chromatic distinctions.

As we mentioned earlier, our simulated shadows can also be regarded as simulated achromatic transparencies, and thus our results are relevant to models of achromatic transparency perception (Metelli 1974; Beck et al 1984; Gerbino 1994; Kasrai and Kingdom 2001; Robiletto et al 2002; Singh and Anderson 2002). Our finding that the perception of achromatic transparency was improved when the colours were the same on either side of the transparency border, yet impaired when they were different, shows that colour should be included among the factors known to impact upon achromatic transparency perception. The impairment of transparency perception when random colour contrast was added across the borders of the overlay is in keeping with a recent study in which the impression of achromatic transparency was found to be optimal when the ratios of cone excitations between any two surfaces viewed through the transparency were preserved (Ripamonti and Westland 2003). In our chromatic all-border condition this cone-invariant ratio rule was violated. However, the cone-invariant ratio rule presumably does not predict the superior performance of the chromatic over achromatic conditions. The impairment of transparency perception with added random colour contrast is also in keeping with some recent models of chromatic transparency perception. D'Zmura et al (1997) found that for a simulated overlay to appear as a

uniform coloured transparency, systematic shifts in the colours of the overlaid regions were necessary, specifically those that corresponded to translations or convergences in colour space. In our chromatic all-border stimulus, the shifts in colour space from background to shadow sector were random, and on the D'Zmura et al model would therefore be expected to give rise to a poor impression of transparency. However, it is not clear how the superior performance of the chromatic over achromatic conditions could be explained by the D'Zmura et al model, since an achromatic transparency, whether on a chromatic or achromatic background, does not produce any shifts in colour.

Acknowledgment. This research was supported by Canadian Institute of Health Research grant #MOP-11554 given to FK.

References

- Beck J, Prazdny K, Ivry R, 1984 "The perception of transparency with achromatic colors" *Perception & Psychophysics* **35** 407–422
- Cavanagh P, 1991 "Vision at equiluminance", in *Limits of Vision* Eds J J Kulikowski, I J Murray, V Walsh, volume 5 of *Vision and Visual Dysfunction* Ed. J Cronly-Dillon (Boca Raton, FL: CRC Press) pp 234–250
- Cavanagh P, Leclerc Y, 1989 "Shape from shadows" *Journal of Experimental Psychology: Human Perception and Performance* **15** 3–27
- Dominy N J, Lucas P W, 2001 "Ecological importance of trichromatic vision to primates" *Nature* **410** 363–366
- D'Zmura M, Colantoni P, Knoblauch K, Laget B, 1997 "Color transparency" *Perception* **26** 471–492
- Gerbino W, 1994 "Achromatic transparency", in *Lightness, Brightness, and Transparency* Ed. A L Gilchrist (Hove, Sussex: Lawrence Erlbaum Associates) pp 215–255
- Gilchrist A L, 1979 "The perception of surface blacks and whites" *Scientific American* **240**(3) 112–123
- Gilchrist A L, 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
- Kasrai R, Kingdom F A A, 2001 "The precision, accuracy, and range of perceived achromatic transparency" *Journal of the Optical Society of America A* **18** 1–11
- Kingdom F A A, 2003 "Color brings relief to human vision" *Nature Neuroscience* **6** 641–644
- Knill D C, Kersten D, Mamassian P, 1996 "Implications of a Bayesian formulation of visual information for processing for psychophysics", in *Perception as Bayesian Inference* Eds D C Knill, W Richards (Cambridge: Cambridge University Press) pp 239–286
- Metelli F, 1974 "The perception of transparency" *Scientific American* **230**(4) 90–98
- Mollon J D, 1989 " 'Tho' she kneel'd in that place where they grew...': The uses and origins of primate colour vision" *Journal of Experimental Biology* **146** 21–38
- Mullen K T, Kingdom F A A, 1991 "Colour contrast in form perception", in *The Perception of Colour* Ed. P Gouras, volume 6 of *Vision and Visual Dysfunction* Ed. J Cronly-Dillon (London: Macmillan) pp 198–217
- Olmos A, Kingdom F A A, 2004 "A biologically inspired algorithm for the recovery of shading and reflectance images" *Perception* **33** in press
- Parraga C A, Troscianko T, Tolhurst D J, 2002 "Spatiochromatic properties of natural images and human vision" *Current Biology* **12** 483–487
- Ripamonti C, Westland S, 2003 "Prediction of transparency perception based upon cone-excitation ratios" *Journal of the Optical Society of America A* **20** 1673–1680
- Robilotto R, Khang B, Zaidi Q, 2002 "Sensory and physical determinants of perceived achromatic transparency" *Journal of Vision* **2**(5) 388–403
- Rubin J M, Richards W A, 1982 "Color vision and image intensities: When are changes material?" *Biological Cybernetics* **45** 215–226
- Singh M, Anderson B L, 2002 "Toward a perceptual theory of transparency" *Psychological Review* **109** 492–519
- Sumner P, Mollon J D, 2000 "Catarrhine photopigments are optimized for detecting targets against a foliage background" *Journal of Experimental Biology* **23** 1963–1986
- Switkes E, Bradley A, De Valois K K, 1988 "Contrast dependence and mechanisms of masking interactions among chromatic and luminance gratings" *Journal of the Optical Society of America A* **5** 1149–1162

ISSN 0301-0066 (print)

ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 33 2004

www.perceptionweb.com

Conditions of use. This article may be downloaded from the Perception website for personal research by members of subscribing organisations. Authors are entitled to distribute their own article (in printed form or by e-mail) to up to 50 people. This PDF may not be placed on any website (or other online distribution system) without permission of the publisher.