

Chromatic properties of the colour-shading effect

Frederick A.A. Kingdom ^{*}, Sohil Rangwala, Karim Hammamji

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

Received 18 June 2004; received in revised form 19 November 2004

Abstract

The ‘colour-shading effect’ describes the phenomenon whereby chromatic variations affect the magnitude of perceived shape-from-shading in luminance patterns. A previous study showed that in mixed colour-plus-luminance sine-wave plaids, impressions of depth in the luminance component were enhanced by non-aligned chromatic components, and suppressed by aligned chromatic components [Nature Neuroscience 6 (2003) 641–644]. Here we examine the chromatic determinants of these effects. Colour contrast was defined along the cardinal axes of colour space in order to isolate the $L-M$ and $S-(L+M)$ post-receptoral chromatic mechanisms. We found no difference in the potency of $L-M$ -only and $S-(L+M)$ -only gratings, either for enhancing or suppressing perceived depth. Moreover, the magnitude of depth-suppression was no different for any combination of depth-enhancing and depth-suppressing cardinal directions. Finally we tested whether the visual system carried the assumption that natural shading is tinged with blue, by measuring perceived depth in a colour-plus-luminance grating that was made to appear either bright-yellow/dark-blue or bright-blue/dark-yellow. However there was no difference in the magnitude of depth-suppression between conditions, suggesting that the visual system does not make any assumption about the colour of natural shading. Taken together, the results suggest that while the colour-shading effect is highly sensitive to colour contrast, it is agnostic with respect to colour direction.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Colour; Depth; Shape-from-shading

1. Introduction

The role of colour vision in the analysis of image structure has been of considerable interest to vision scientists in recent decades (reviewed by Regan, 2000). The topic has been mainly studied using isoluminant (or equiluminant) stimuli, which ostensibly isolate the colour vision system and allow its spatio-temporal properties to be probed directly. However, much can be learnt about the role of colour vision in the analysis of image structure by studying how colour and luminance interact in stimuli that embody the spatio-temporal relationships that exist between these two dimensions in the natural visual world (Kingdom, 2003).

The ‘colour-shading effect’ (Kingdom, 2003) is a recent example of this approach. When a colour (used here to mean chromatic) grating is added to a differently-oriented luminance grating, an impression of a corrugated depth surface is triggered—a clear instance of perceived shape-from-shading—and termed here ‘depth-enhancement’. However when a second colour grating of the same orientation and spatial phase as the luminance grating is now added, the impression of depth is reduced or eliminated, termed here ‘depth-suppression’. Some of these effects may be seen in Fig. 1. Shape-from-shading has hitherto been studied almost entirely in the achromatic domain (Attick, Griffin, & Redlich, 1996; Lehky & Sejnowski, 1988; Ramachandran, 1988; Sun & Perona, 1997); the colour-shading effect demonstrates that shape-from-shading can be profoundly affected by colour contrast.

^{*} Corresponding author.

E-mail address: fred.kingdom@mcgill.ca (F.A.A. Kingdom).

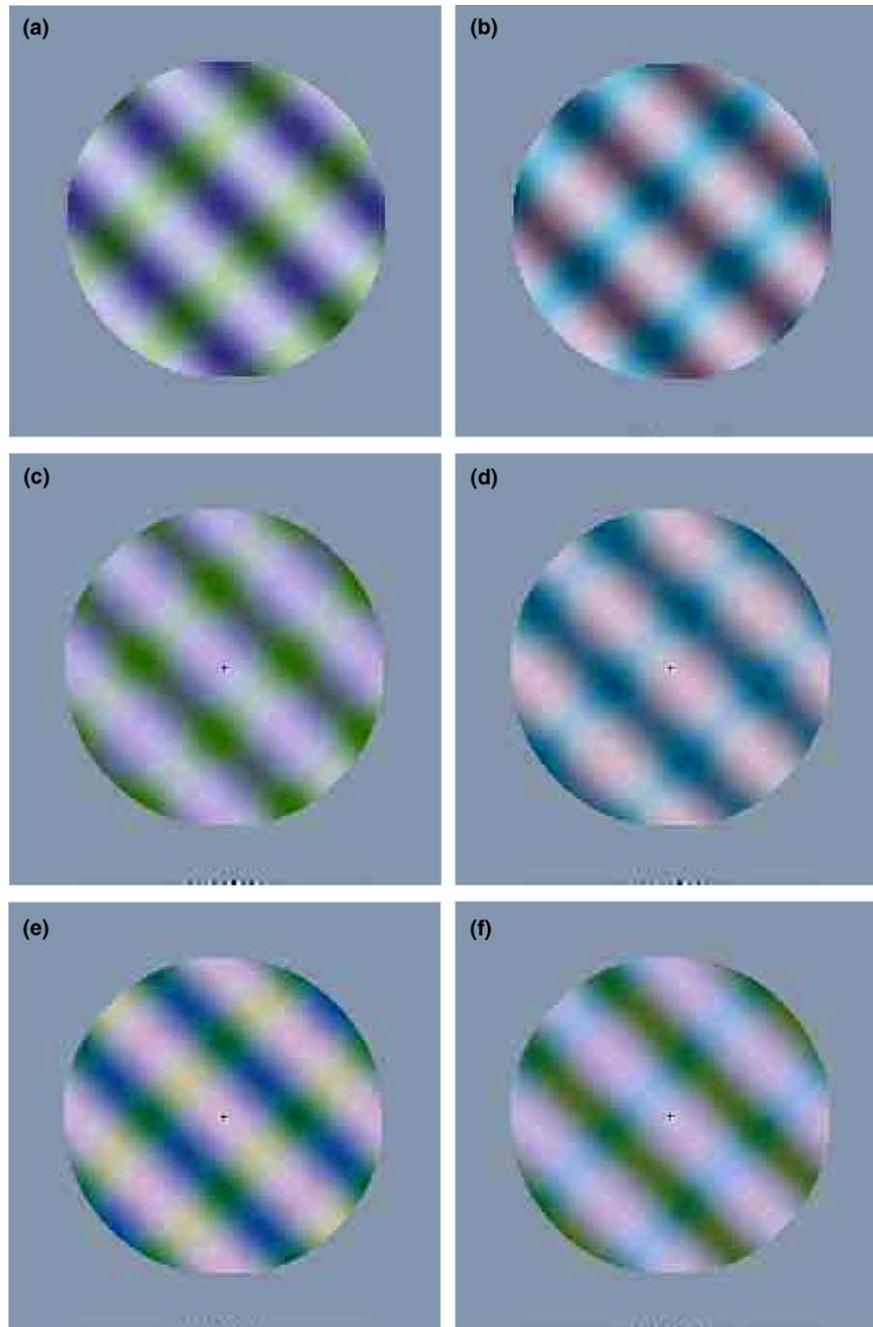


Fig. 1. Example plaids used in Experiments 1 and 2. (a) and (b) are two-component plaids used in Experiment 1. These consist of a right-oblique, 'depth-enhancing' chromatic grating together with a left-oblique, luminance 'shading' grating. In (a) the chromatic grating is defined along the S and in (b) the $L-M$ cardinal axes. Most observers perceive left-oblique depth corrugations in these plaids. (c)–(f), example plaids used in Experiment 2. These consist of the two components as in (a) and (b), plus a third 'depth-suppressing' chromatic grating that is added in phase to the luminance grating. Most observers report a reduction in depth in (c)–(f). The four combinations of depth-enhancing and depth-suppressing chromatic grating are: (c) S and S , (d) $L-M$ and $L-M$, (e) S and $L-M$ and (f) $L-M$ and S .

Kingdom (2003) argued that the depth-enhancing and depth-suppressing capabilities of colour contrast in plaids such as those in Fig. 1 revealed the influence of the visual systems' built-in assumptions about the relationship between colour and luminance in the natural visual world. The assumptions are that chromatic variations, and those luminance variations that are spa-

tially aligned with them, arise from changes in surface reflectance, whereas pure, or near-pure luminance variations arise from spatially non-uniform illumination, such as shading and shadows. From these assumptions it follows that the right-oblique colour gratings in Fig. 1a and b are interpreted as changes in spectral reflectance, i.e. as surfaces, and that the luminance-defined

left-oblique gratings are interpreted as shading. In Fig. 1c–f, where a second colour grating has been added in spatial alignment to the luminance grating, the interpretation of the luminance variations shifts from being that of shading towards being that of a surface, with an attendant reduction in perceived depth. Although the physical reality that underpins these assumptions has been appreciated by vision scientists for some time (Cavanagh, 1991; Mullen & Kingdom, 1991; Rubin & Richards, 1982), the colour-shading effect is, to our knowledge, the first evidence that these assumptions are built into the fabric of the human visual system.

The unique and positive role that colour vision appears to play in the perception of shape-from-shading is especially pertinent given that colour vision is traditionally considered the poor cousin of luminance vision in its capacity to analyse the third dimension, i.e. depth. For example, stereoscopic depth judgements of isoluminant stimuli are generally worse than those made with purely luminance-defined patterns (Gregory, 1977; Kingdom & Simmons, 2000; Livingstone & Hubel, 1987; Lu & Fender, 1972). The colour-shading effect demonstrates however that colour vision *in combination with* luminance vision can significantly impact depth perception. This positive role of colour vision in the perception of shape-from-shading complements other positive roles of colour vision in the analysis of image-structure, for example for detecting fruit and flowers in foliage (Domini & Lucas, 2001; Mollon, 1989; Sumner & Mollon, 2000), identifying shadows and transparency (Kingdom, Beauce, & Hunter, 2004) and memorising scenes (Gegenfurtner & Rieger, 2000).

A number of questions concerning the chromatic properties of the colour-shading effect naturally arise. First, are all colour directions equally effective at depth-enhancement, and are all colour directions equally effective at depth-suppression? At the post-receptor level the primate colour vision system divides into two colour-opponent pathways, one that differences the outputs of the *L* (long-wavelength-sensitive) and *M* (middle-wavelength-sensitive) cones—the '*L–M*' pathway—the other that differences the outputs of the *S* (short-wavelength-sensitive) from the sum of outputs of the *L* and *M* cones—the '*S–(L + M)*' pathway (DeValois, 1965; Derrington, Krauskopf, & Lennie, 1984; DeValois & DeValois, 1975; Krauskopf, Williams, & Heeley, 1982). It is reasonable to ask whether stimuli that selectively stimulate the two pathways are equally potent at driving the colour-shading effect. Mollon (2000) has argued that the *S–(L + M)* pathway is the more phylogenetically primordial of the two pathways. Might the older colour system have the bigger impact on perceived shape-from-shading, on the grounds that it has had more time to develop an intimate relationship with luminance vision?

Second, is the particular combination of depth-enhancing and depth-suppressing colour directions important? It is possible that the colour-shading effect is weaker when the depth-enhancing and depth-suppressing colour directions are the same, as in these circumstances the visual system might bind together both colour patterns into a single object, 'releasing' the luminance variations from being designated as changes in reflectance, and designating them instead as shading, even though they are spatially aligned with one of colour patterns. The result might be more depth-enhancement/less depth-suppression.

Ecological considerations lead to a third question about the chromatic properties of the colour-shading effect. In natural scenes, although shadows and shading are predominantly luminance-defined features, they are often tinged with colour, and in particular blue (Churma, 1994; Parraga, Troscianko, & Tolhurst, 2002). Bluing in shadows and shading occurs especially on sunny days; because shaded regions are bathed predominantly in blue skylight, whereas un-shaded regions are bathed in both blue skylight and yellow sunlight. If the visual system has knowledge that shading in natural scenes tends to be bluish, then it might refrain from making the assumption that blue-yellow variations are inevitably changes in surface reflectance, especially when of the bright-yellow/dark-blue variety. The question is thus. Is the magnitude of depth-suppression from blue-yellow gratings (which are not the same as the gratings employed to isolate the *S–(L + M)* pathway—see below) phase-dependent? In other words if a blue-yellow grating is added to a luminance grating that is perceived as shading, will it suppress perceived depth to a lesser extent when the blue falls on the dark, compared to the bright phase of the grating?

We have attempted to answer these questions using mixed colour-plus-luminance plaids in which we have manipulated the colour direction and saturation of both depth-enhancing and depth-suppressing colour contrasts. The results have enabled us to refine our knowledge of the chromatic properties of the colour-shading effect, and therefore of the assumptions made by the visual system concerning the relationships between colour and luminance in the visual world.

2. Methods

2.1. Stimuli—generation and display

The stimuli were generated by a VSG2/5 graphics card (Cambridge Research Systems) and displayed on a Sony Trinitron F500 flat-screen monitor. The R (red), G (green) and B (blue) gun outputs of the monitor were gamma-corrected after calibration with an Optical photometer (Cambridge Research Systems). The spectral emission functions of the R, G and B phosphors

were measured using a PR 640 spectral radiometer (Photo Research), with the monitor screen filled with red, green or blue at maximum luminance. The CIE coordinates of the monitors' phosphors were R: $x = 0.624$, $y = 0.341$; G: $x = 0.293$, $y = 0.609$; B: $x = 0.148$, $y = 0.075$. The stimuli were viewed through a custom-built, modified 8-mirror Wheatstone stereoscope. Viewing distance along the light path through the stereoscope was 105 cm.

2.2. Stimuli—component gratings

Component gratings can be seen in the example plaids in Figs. 1 and 2. All gratings were combinations of sinusoidal modulations of cone contrast, with cone contrast defined as $L_c = \Delta L/L_b$, $M_c = \Delta M/M_b$ and $S_c = \Delta S/S_b$ (Cole, Hine, & McIlhagga, 1993; Norlander & Koenderink, 1983; Sankeralli & Mullen, 1997; Stromeyer, Cole, & Kronauer, 1985). The denominator in

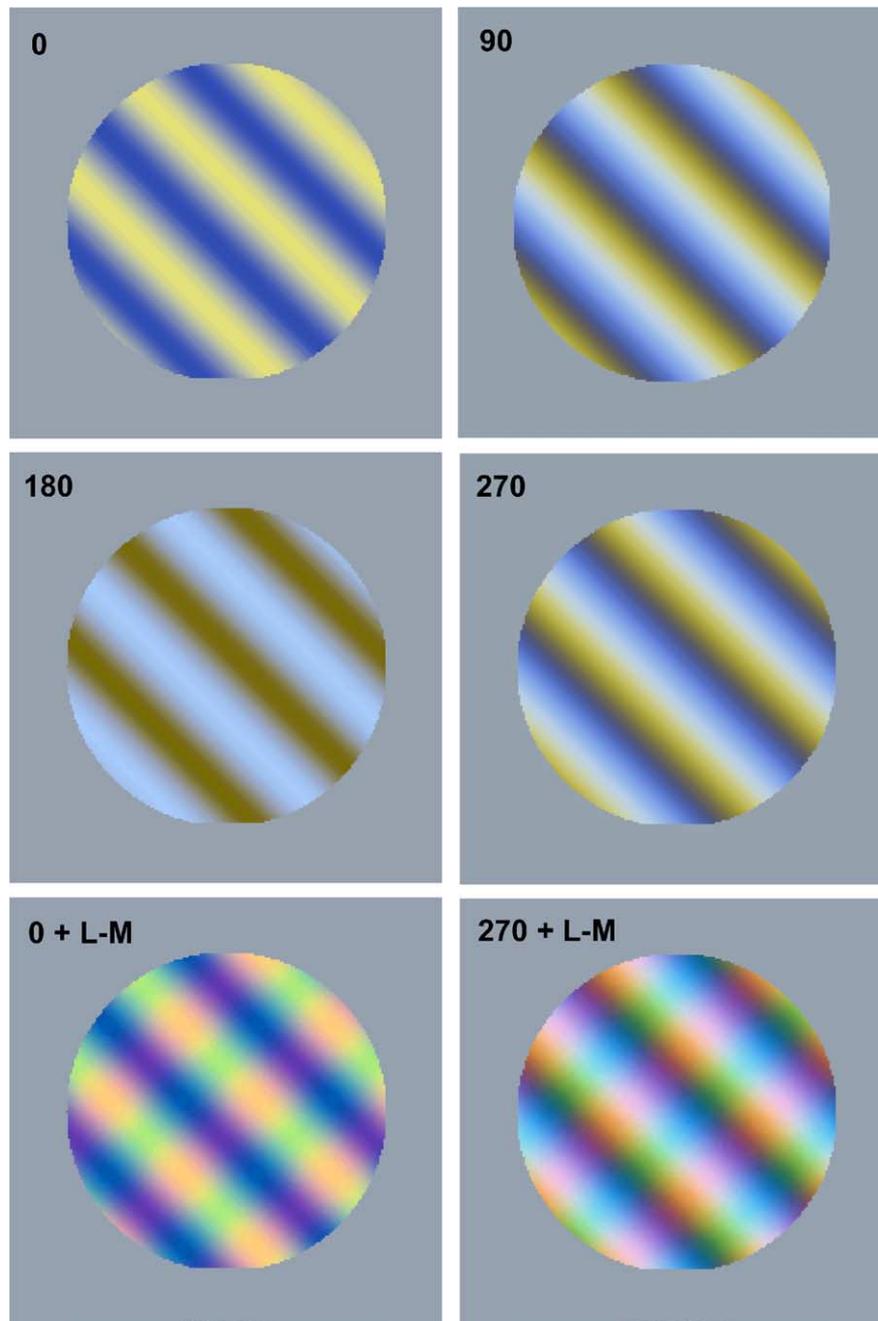


Fig. 2. Component gratings and plaids used in Experiment 3. The top and middle pair show the four phase relationships between the BY (blue-yellow) and the luminance shading grating. 0° and 180° are the two in-phase conditions, with $0 =$ bright-yellow/dark-blue and $180 =$ bright-blue/dark-yellow. 90° and 270° are the two out-of-phase conditions. In the bottom pair a right-oblique, depth-enhancing $L-M$ grating has been added to the 0° (left) and 270° (right) condition to produce the three-component plaids used in the actual experiment.

each term refers to the cone excitation produced by the grating's d.c., a mid-grey colour with CIE chromaticity $x = 0.282$ and $y = 0.311$, and luminance 40 cd/m^2 . The nominator in each cone contrast term represents the difference in cone excitation between the peak of the grating's modulation and the d.c. The resulting *LMS* cone excitations assigned to each pixel were converted to RGB phosphor intensities using the cone spectral sensitivity functions provided by Smith and Pokorny (1975) and the measured RGB spectral functions of the monitor.

Three of the four types of component gratings were defined along the cardinal axes of a modified version of the MacLeod–Boynton colour space (MacLeod & Boynton, 1979), illustrated in Fig. 3. The axes are termed *LUM*, *L–M* and *S*. The term cardinal implies that each grating uniquely stimulates one of the three post-receptor mechanisms (Cole et al., 1993; Derrington et al., 1984; Krauskopf et al., 1982; Norlander & Koenderink, 1983; Sankeralli & Mullen, 1997; Stromeyer et al., 1985). The relative cone contrast inputs to the three mechanisms have been estimated to be as follows: $kL_c + M_c$ for the luminance mechanism, $L_c - M_c$ for the mechanism that differences *L* and *M* cone-contrasts, and $S_c - (L_c + M_c)/2$ for the mechanism that differences *S* from the sum of *L* plus *M* cone-contrasts (Cole et al., 1993; Sankeralli & Mullen, 1996; Stromeyer et al., 1985). The parameter *k* determines the relative weightings of the *L* and *M* cone-contrast inputs to the luminance mechanism, and varies between observers. Once *k* was established for each subject (see below) the cone contrasts of the gratings necessary to make them orthogonal were:

$$LUM = L_c + M_c + S_c \tag{1a}$$

$$L-M = L_c - kM_c + S_c(1 - k)/2 \tag{1b}$$

$$S = S_c \tag{1c}$$

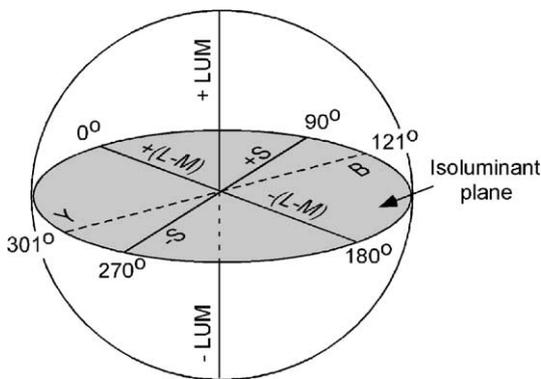


Fig. 3. The modified version of the MacLeod–Boynton colour space used to define the stimuli. See text for details.

Grating contrast was defined as follows: for *LUM*, the contrast assigned to each cone; for *L–M*, the difference in *L* and *M* cone contrasts; for *S*, the contrast assigned to the *S* cone.

In one of the experiments, a ‘blue-yellow’, or BY grating was employed, whose colours were modulated between, approximately, unique blue and unique yellow. The BY modulation was defined along an axis in the isoluminant plane that lay at an angle between the two principle axes, as shown in Fig. 3. The $L_c - M_c$, and S_c contrasts of the BY grating were calculated respectively as $C_{BY}\sin(\theta)$ and $rC_{BY}\cos(\theta)$ where C_{BY} was the contrast of the BY grating, and θ its vector direction in colour space. Note that in Fig. 3, + (*L–M*) which is reddish, is defined at 0° , and + *S*, which is violet, is defined at 90° . *r* is a scaling factor used to equate the perceived contrasts of the *L–M* and *S* components, and was calculated from their individual contrast detection thresholds (see below). Examples of a BY grating combined with *LUM* gratings of various relative phases can be seen in Fig. 2.

2.3. Stimuli—plaids

All plaids were constructed by combining the luminance and chromatic gratings additively. In order to minimize any on-screen interactions between the different plaid components, the components were stored on separate pages of the VSG's video memory and displayed in rapid alternation at 180 Hz. In Experiment 1 there were two plaid components, so the stimulus was seen at 90 Hz. In Experiments 2 and 3 there were three plaid components, so the stimulus was seen at 60 Hz. No subject reported any visible flicker in any of the experiments. The grating components had a spatial frequency of 0.75 cpd, and the plaids were presented in a circular hard-edged window of diameter 4° . The orientation of the depth-enhancing colour gratings was always $+45^\circ$ (right-oblique), and the orientation of the *LUM*-shading and depth-suppressing colour gratings was always -45° (left-oblique). The phases of the depth-enhancing and *LUM*-shading components were randomized for each stimulus presentation. The phases of the depth-suppressing chromatic gratings were specified relative to the phase of the *LUM*-shading component, and details will be provided with each experiment. The plaids were presented against a grey background, whose cone excitation levels were the same as those of the d.c. of the plaid.

2.4. Stimuli—matching stereo-grating

The matching stimulus was a random-element disparity grating containing left-oblique depth corrugations whose amplitude could be adjusted by the subject to match the depth corrugations perceived in the plaids.

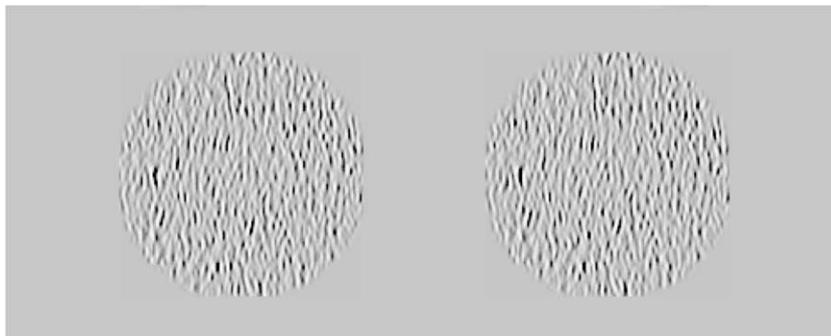


Fig. 4. Matching stimulus. When free-fused, one can see left-oblique depth corrugations. The amplitude of these corrugations was adjusted by the subject to match the perceived depth of the corrugations in the test plaids.

An example stereo-pair is shown in Fig. 4. Each half of the disparity grating consisted of 3000 gabors, whose positions were random but whose disparities were selected to produce 0.75 cpd left-oblique depth corrugations (i.e. similar to those in the plaid). The gabors were all odd-symmetric (phase = 90°), with a spatial frequency of 8.0 cpd, bandwidth 1.5 octaves and Michelson contrast 12%. When two gabors fell on top of one another their amplitudes, but not d.c. levels were added, and clipping prevented any underflow and overflow in the look-up-table.

In pilot experiments we discovered that when the disparity modulation in the matching grating was defined sinusoidally, its perceived depth corrugations had sharper peaks than those seen in the plaids, making it difficult to make a match. To render the corrugations in the matching stimulus more rounded, the disparity modulation was defined as the sum of three sinusoids. Each sinusoid was defined as $A\sin(fx + \rho)$, where a is amplitude, f spatial frequency, ρ phase and x the position along the axis orthogonal to the corrugation's orientation. The relative values of A for the three sinusoids were 1, 1/5 and 1/25, and the relative values of f 1, 3 and 5. The phase ρ of disparity modulation was randomised on each trial, and amplitude A was adjusted by the subject during each trial.

To produce the disparity modulation the gabors in the stereo-grating were selected from 100 templates that were pre-generated and stored in computer memory. Each template was a square patch containing a gabor that was horizontally offset from the middle of the patch, with sub-pixel accuracy, by an amount that determined its disparity. Before each trial, the positions of the gabor templates were randomly assigned, but during the adjustment phase, the positions were unchanged. During the adjustment phase the display was updated about five times a second. At each update, depending on the current setting of A , pairs of gabors with $+1/2$ and $-1/2$ the required disparity were selected from the templates and painted into each gabor position in the two stereo-half-pairs. The updating process itself was

invisible; the stereo-grating appeared as a static stimulus whose corrugations grew or receded in depth as A was adjusted.

3. Procedures

3.1. Cardinal settings

Because of inter-subject variation in the relative weightings of the L and M cones in the luminance mechanism, it was necessary to ensure that both the $L-M$ and BY gratings were isoluminant (as both received inputs from L and M cones). We used the criterion of minimum perceived motion. The contrast of the $L-M$ and BY gratings was set to 0.025, and the gratings were set to drift at about 1.0 Hz. By pressing a key on the CB3 response box (Cambridge Research Systems), subjects added (or subtracted) an equal amount to both the L and M cone contrasts until perceived motion was at a minimum. Each subject made between 10 and 15 settings. For the $L-M$ grating the average amount of luminance contrast added (or subtracted) was used to calculate the parameter k in Eq. (1b), where k is the ratio of L to M cone contrasts in the putative luminance mechanism. k was determined to be for subject SR = 1.10, MT = 1.67, HW = 1.14 and FK = 1.78. For the BY grating the average amount of luminance contrast added or subtracted was used to calculate the ratio of luminance to colour contrast needed to achieve isoluminance, and these ratios were for SR = -0.819 , MT = -0.171 and FK = -0.315 .

Although it has been established that S cones have a negligible input to the luminance mechanism across all subjects (e.g. Eskew, McLellan, & Giulianini, 1999; see also Mullen & Sankeralli, 1999) we checked that our S gratings were isoluminant in three of the four subjects tested (SR, MT and FK) using the same procedure as for the $L-M$ gratings. We found that the amount of luminance contrast added to produce minimum perceived motion was not significantly different from zero.

Individual variation in lens and macular pigment density means that stimuli calculated to modulate only *S* cones of the standard observer might additionally modulate the *L*- and *M*-cones of an individual observer. One way to check for *S*-cone isolation is to use the criterion of minimum visibility (Mullen & Kingdom, 2002), which exploits the fact that contrast sensitivity for *S* gratings is much lower than that for *L–M* gratings. Three of the four subjects (SR, MT and FK) adjusted the direction within the isoluminant plane of a 0.04 contrast grating until it appeared minimally visible. The settings for each subject were not significantly different from the direction that conformed to the calibrated *S*-cone axis (90°–270°).

3.2. BY colour direction

Using a grating with C_{BY} set to 0.025 and r to 5 (see above), subjects adjusted the direction θ of the BY grating until it appeared to modulate between “sky-blue” and “sunny-yellow”. Each subject made several adjustments, and the resulting mean values of θ in the blue direction were SR = 118°, MT = 125° and FK = 120° in the blue direction. These were the values used in Experiment 3. The average value across subjects is 121°, the value shown in Fig. 3.

3.3. Contrast detection thresholds

In order to equate the contrasts of the *L–M* and *S* gratings (a procedure not necessary for the BY gratings, since we never compared the efficacy of the BY relative to the *L–M* and *S* gratings) we defined contrast in terms of multiples of detection threshold. Contrast detection thresholds were measured using a conventional 2IFC (two-interval-forced-choice) procedure, in which the target appeared in one of the intervals, the other being blank. A standard ‘two-up-one-down’ staircase procedure established the threshold at the 70.7% correct level. Stimulus exposure duration was 500 ms. The staircase procedure increased or decreased contrast by a factor of 1.25 trial-by-trial. After ten reversals the staircase was terminated, and the threshold contrast calculated as the geometric average contrast over the last eight reversals. Three or four thresholds were gathered for each condition and averaged. The values obtained were, for SR: $L–M = 0.0024$, $S = 0.013$; for MT: $L–M = 0.0019$, $S = 0.012$; for HW: $L–M = 0.0029$, $S = 0.027$.

3.4. Matching perceived depth

Subjects used the keys on the response box to adjust the amplitude of the depth corrugations in the stereograting until they matched those in the plaid. There was no time limit. Some subjects experienced fading of one or more of the plaid’s components during prolonged

fixation, so all subjects were encouraged to let their eyes roam freely around the stimuli. During each experimental session all the conditions of an experiment were presented in random order, and for each experiment there were ten repeat sessions and therefore ten measurements per condition.

4. Results

4.1. Experiment 1. Do *L–M* and *S* gratings differ in their capacity to depth-enhance?

We wished to compare the depth-enhancing capabilities of colour gratings defined along the two cardinal directions. Example stimuli are shown in Fig. 1a and b. There were 36 conditions: six colour contrasts for each of the *L–M* and *S* cardinal directions, and three contrasts for the *LUM* grating. From now on we will refer to the *LUM* grating simply as the shading grating. The six colour contrasts were: $L–M = 0.0, 0.005, 0.01, 0.02, 0.04, 0.08$; $S = 0.0, 0.025, 0.05, 0.1, 0.2, 0.4$. The three shading contrasts were 0.05, 0.15 and 0.45. All 36 conditions were presented in random order during a session, and there were ten repeat sessions.

Fig. 5a–c shows the results for the 0.15 shading contrast condition. Each graph plots perceived depth as a function of colour contrast, with colour contrast given in multiples of detection threshold. The horizontal dashed line shows the magnitude of perceived depth in the absence of colour contrast, so all points that lie above this line indicate depth enhancement due to the added colour contrast. As can be seen perceived depth rises systematically with colour contrast. Similar results were found for the two other shading contrasts (graphs not shown).

In order to obtain an overall measure of the amount of depth enhancement across conditions we performed the following analysis. We first fitted a sigmoidal function to the data when plotted against the log of colour contrast. The choice of a sigmoidal function was not based on any theoretical grounds; it appeared to be the best function to capture the shape of the data. The sigmoidal function was:

$$D = a + b / (1.0 + \exp[(c - \log C) / d]) \quad (2)$$

where D is matched depth, C colour contrast (in multiples of detection threshold) and a , b , c and d free parameters. Example fits to HW’s 0.45 shading contrast condition are shown in Fig. 3d. We next estimated the average amount of depth enhancement under each sigmoidal curve. To do this we calculated the area under each curve that was bounded by the zero-colour-contrast horizontal dotted line, and the two vertical dotted lines positioned at the minimum ($\log C_{\min}$) and

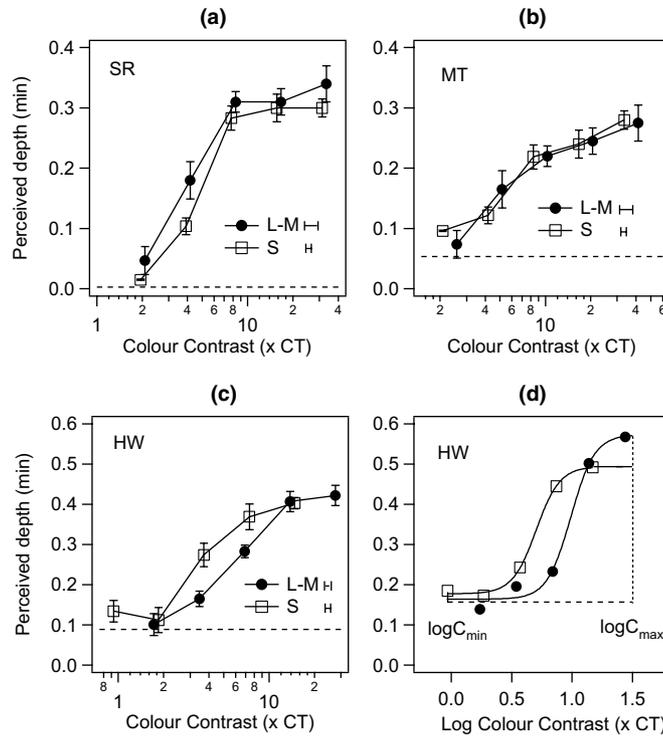


Fig. 5. Sample results from Experiment 1. (a)–(c) show the three subjects’ results for the 0.15 shading contrast condition. Perceived depth is plotted against the colour contrast of the right-oblique chromatic grating, expressed in multiples of detection threshold. Separate plots are shown for the *L–M* and *S* chromatic gratings. The horizontal dashed line shows the amount of perceived depth in the absence of the chromatic grating, so all points that lie above this line show depth-enhancement from the chromatic grating. The two horizontal error bars next to the legend symbols show the standard errors of the contrast-detection thresholds used to normalise the contrast thresholds, and thus give an indication of the likely error in the positioning of each set of data along the abscissa. (d) shows the method for estimating the average amount of depth enhancement in HW’s 0.45 shading contrast data. Sigmoidal functions have been fitted to the *L–M* and *S* data when plotted against log colour contrast, and the area under each curve between the dashed horizontal (representing perceived depth with zero colour contrast) and dotted vertical lines calculated. See text for details.

maximum ($\log C_{\max}$) abscissa values established across both the *L–M* and *S* data. The average depth enhancement was then obtained by dividing the area by the difference between $\log C_{\min}$ and $\log C_{\max}$. The integral of the sigmoidal function, G , is given by:

$$G = (a + b) \log C - bc + bd \ln(1 + (\log C - c)/d) \quad (3)$$

and the average depth enhancement by:

$$\frac{G_{\max} - G_{\min} - D_0(\log C_{\max} - \log C_{\min})}{\log C_{\max} - \log C_{\min}} \quad (4)$$

where G_{\min} and G_{\max} are the areas bounded by $\log C_{\min}$ and $\log C_{\max}$ and D_0 is the depth in the absence of colour contrast.

The resulting estimates of average depth enhancement are shown in Fig. 6, plotted against shading contrast. Although there is some between-subject variation, there appears to be no consistent difference in depth-enhancement between the two cardinal directions. A two-factor within-subjects analysis-of-variance (ANOVA), with Cardinal direction and Shading contrast as factors confirms this impression. The main effect of Cardinal direction was not significant ($F(1,2) = 0.05$; $p = 0.84$). Shading contrast was significant

($F(2,4) = 10.02$; $p < 0.05$). The interaction between Cardinal direction and Shading contrast was not significant ($F(2,4) = 0.1$; $p = 0.91$). Thus we can conclude that no difference has been shown to exist between the two cardinal directions in terms of their potency for depth enhancement.

4.2. Experiment 2. Chromatic properties of depth suppression

The second experiment concentrated on the depth-suppressing capabilities of colour variations defined along the cardinal axes. Example stimuli are shown in Fig. 1c–f. In order to explore the factors affecting depth-suppression it was necessary to use plaids with already strong impressions of depth. We therefore used plaids containing a relatively high contrast depth-enhancing colour grating—a 0.027 contrast *L–M* grating, or a 0.133 contrast *S* grating. The depth-suppressing colour gratings were always added in phase with the shading grating, but the polarity of the phase relationship (0° vs. 180°) was randomized. There were 64 conditions: four combinations of depth-enhancing and depth-suppressing cardinal directions, four contrasts of

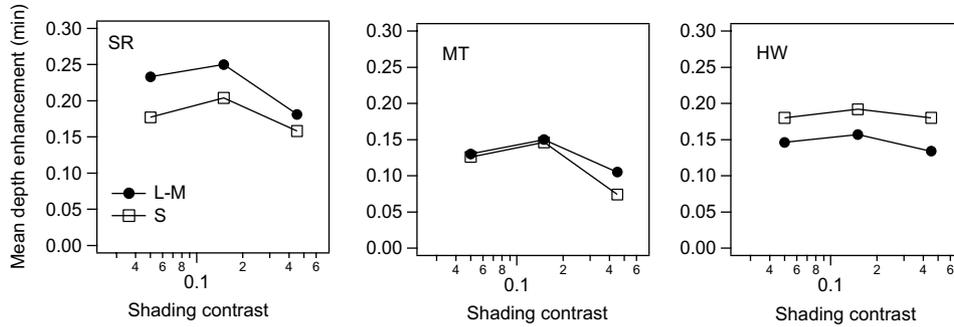


Fig. 6. Results of the curve-fitting procedure as illustrated in Fig. 5d applied to all shading contrast conditions in Experiment 1. The average amount of depth-enhancement is plotted as a function of shading contrast, for both *L–M* and *S* chromatic gratings, and for all three subjects.

the depth-suppressing grating and four shading contrasts. The four combinations of depth-enhancing and depth-suppressing cardinal directions were: *L–M* & *L–M*; *L–M* & *S*; *S* & *L–M*; *S* & *S*. The four contrasts of the depth-suppressing gratings were, for *L–M*: 0.0, 0.013, 0.027 and 0.053; for *S*: 0.0, 0.067, 0.133 and 0.267. The four shading contrasts were 0.037, 0.075, 0.15 and 0.3.

Results for the 0.075 shading contrast condition are shown in Fig. 7. The four plots are for the different combinations of depth-enhancing and depth-suppressing cardinal directions. Colour contrast is again given in multiples of detection threshold. The horizontal lines show the amount of perceived depth in the absence of any depth-suppressing colour contrast. Note that the horizontal lines in this experiment correspond to much greater perceived depths than in the previous experiment, due to the presence of the right-oblique depth-enhancing colour gratings. All plots show that perceived depth is reduced by the presence of the left-oblique, aligned chromatic gratings.

In order to estimate the magnitude of depth suppression for each condition we used the same procedure as

that employed in the first experiment, with the difference that instead of fitting the data with a sigmoidal function, we fitted it with a straight line. The reason for using a straight line is that there was no indication of any systematic acceleration or deceleration along the three data points making up each plot when plotted against log colour contrast.

The resulting estimates of average depth-suppression as a function of shading contrast are shown in Fig. 8. All the data points are greater than zero, showing that depth suppression from aligned colour contrasts was found at all shading contrasts and for all combinations of depth-enhancing and depth-suppressing cardinal directions. As with the data from Experiment 1, there are between-subject differences, but there appears to be no systematic difference in the magnitude of depth-suppression between the different conditions. This conclusion is supported by a 3-factor within-subjects ANOVA, with factors Enhancing cardinal direction, Suppressing cardinal direction and Shading contrast. There was no significant effect of Enhancing cardinal direction ($F(1,2) = 0.403, p = 0.6$), nor Suppressing cardinal direction ($F(3,6) = 1.29; p = 0.36$), nor Shading contrast

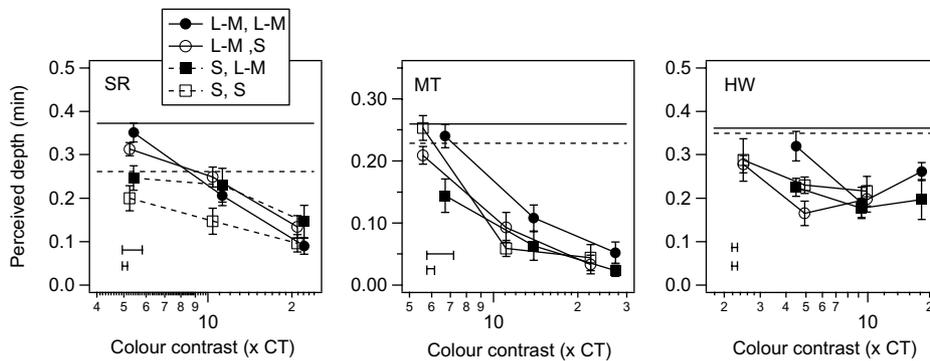


Fig. 7. Results for Experiment 2 for the 0.075 shading contrast conditions. Perceived depth is plotted against colour contrast (in multiples of detection threshold) of the left-oblique chromatic grating added in-phase to the shading grating. In the legend, the first term refers to the cardinal direction of the right-oblique depth-enhancing chromatic grating, while the second term refers to the cardinal direction of the left-oblique depth-suppressing chromatic grating. The horizontal continuous and dashed lines represent perceived depth in the absence of right-oblique colour contrast, for respectively the *L–M* and *S* conditions. All points below these lines show depth-suppression due to the left-oblique chromatic grating. The two horizontal error bars on the bottom left of each figure correspond to those in Fig. 5.

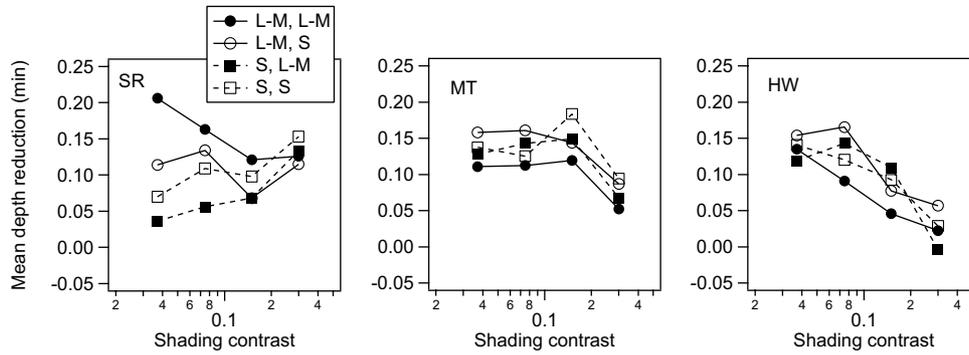


Fig. 8. Results from Experiment 2. The average amount of depth-suppression is plotted against shading contrast. See text for details.

($F(1,2) = 1.904$; $p = 0.30$). The interaction between the depth-enhancing and depth-suppressing cardinal direction was not significant ($F(1,2) = 0.021$; $p = 0.88$), and neither were any other interactions.

4.3. Experiment 3. Depth suppression using BY gratings

In the introduction we suggested that blue-yellow gratings might produce a phase-dependent depth-suppression. The aim of this experiment was to test this idea. Example stimuli are illustrated in Fig. 2.

As in the previous experiment, we wanted plaids with an already strong impression of depth, so all plaids contained a right-oblique depth-enhancing 0.027 contrast *L–M* grating. There were 64 conditions: four phase relationships between the BY and shading gratings (0° , 90° , 180° , 270°), four BY colour contrasts (0.0, 0.012, 0.023, 0.047) and four shading contrasts (0.037, 0.075, 0.15, 0.3).

The results for the 0.037 shading contrast conditions are shown in Fig. 9. The horizontal dashed lines show the amount of perceived depth in the absence of the BY grating, so points below this line indicate depth-suppression, points above it depth-enhancement. For the two in-phase conditions (0° and 180°) the BY grating

produced pronounced depth-suppression in all subjects. For the two out-of-phase conditions (90° and 270°) the addition of the BY grating had little effect on perceived depth except at the highest BY contrast, where in subjects SR and MT depth-suppression can be observed. For subject FK, a small amount of depth-enhancement was observed with the out-of-phase conditions.

Using the same method as described for Experiment 2 we calculated the average amount of depth-suppression for each plot. The resulting estimates of average depth-suppression are plotted against shading contrast in Fig. 10. To test the main hypothesis of this experiment, namely that the amount of depth suppression would be less for the 0° compared to 180° phase conditions, we performed a 2-factor, within-subjects ANOVA with Polarity-of-aligned-phase (0 vs. 180) and Shading contrast as factors. Neither factor, nor factor interaction, was significant (Polarity-of-aligned-phase: $F(1,2) = 3.268$, $p = 0.212$; Shading Contrast: $F(3,6) = 0.54$, $p = 0.672$; Shading Contrast \times Polarity-of-aligned-phase: $F(3,6) = 0.198$, $p = 0.89$).

We expected that the amount of depth-suppression would be less when the BY gratings were out-of-phase than when in-phase. Indeed in one subject (FK), out-of-phase BY gratings produced a small amount of

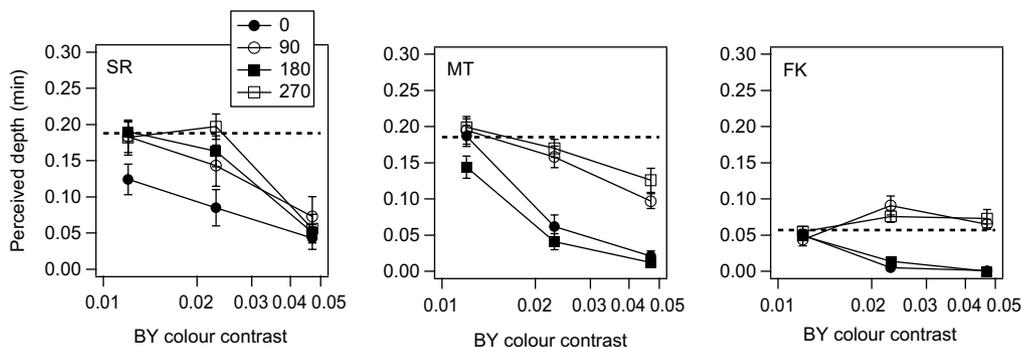


Fig. 9. Results for Experiment 3 for the 0.037 shading contrast conditions. Perceived depth is plotted as a function of the colour contrast of the left-oblique BY (blue-yellow) grating added to the shading grating in various relative phases. The horizontal dashed line represents the amount of perceived depth in the absence of the BY grating. Points above this line show depth-enhancement resulting from the BY grating, points below depth-suppression. The numbers in the legend show the four phase relationships, which are illustrated in Fig. 2.

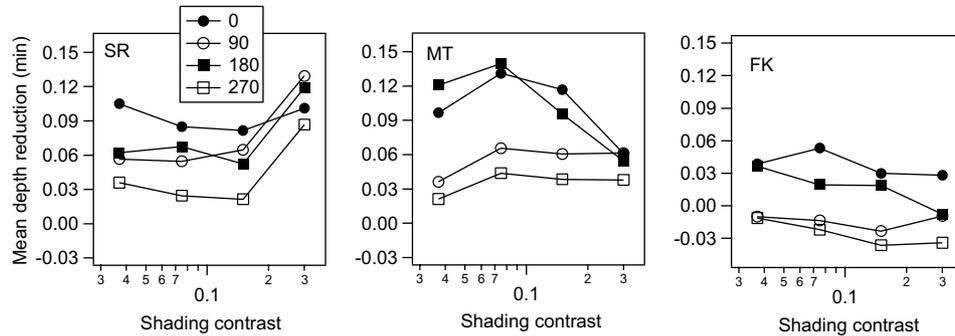


Fig. 10. Results from Experiment 3. Average depth suppression is plotted against shading contrast for the four phase relationship conditions. Note that in FKs 90° and 270° data the points are negative, showing depth enhancement.

depth-enhancement. A 2-factor within-subjects ANOVA with Shading contrast and Phase-alignment (in-phase vs. out-of-phase) showed that Phase-alignment was significant ($F(1,2) = 21.08$; $p < 0.05$).

A prominent feature of the data, and one that was quite unexpected, is the difference in the magnitude of depth suppression between the 90° and 270° phase conditions. The 270° condition produced less depth suppression (and in FK's data more depth enhancement) in every condition tested. However, the effect was not quite significant using a 2-factor within-subjects ANOVA, with Non-aligned-phase (90 vs. 270) and Shading contrast as factors (Non-aligned-phase: $F(1,2) = 11.75$; $p = 0.076$).

5. Discussion

The results of the present study can be summarised as follows.

1. L - M and S gratings are similar in their capacity to both enhance and suppress perceived depth in mixed colour-plus-luminance plaids.
2. The capacity of L - M or S gratings to suppress perceived depth does not depend on whether the depth-enhancing gratings are themselves L - M or S .
3. The capacity of BY (blue-yellow) gratings to suppress perceived depth does not depend on whether the blue phase falls in the dark or the bright part of the shading grating.
4. Out-of-phase BY gratings are less effective depth suppressors than in-phase BY gratings (and in some cases are depth-enhancers).

Using colour gratings defined along, and between, the cardinal directions of colour space, the results of the present study generalise the findings of Kingdom (2003) to new colour directions. That the two cardinal directions are no different in their capacity to enhance or suppress perceived depth leads to the main conclusion of this study: while the colour-shading effect is highly

dependent on colour contrast, it is agnostic to colour direction. The corollary to this conclusion is that if indeed the colour-shading effect reflects the influence of the visual system's assumptions about the colour-luminance nature of surfaces and spatially non-uniform illumination, as Kingdom (2003) has argued, then the assumptions are about colour contrast, not colour direction. We have not of course explored the full gamut of colour directions, and therefore cannot rule out the possibility that there are colour directions that are especially effective, or especially ineffective, at depth-enhancement and/or depth-suppression. However, our present results suggest that this is unlikely.

Could the depth-enhancing capabilities of colour contrast be a luminance artifact? Kingdom (2003) showed that a plaid consisting of two orthogonal-in-orientation, equal-contrast luminance gratings elicited little perceived depth. Perhaps though low contrast luminance gratings act as potent depth-enhancers for higher contrast orthogonal-in-orientation gratings, making it possible that our ostensibly isoluminant depth-enhancing colour gratings were simply providing a low contrast luminance input. Two reasons however rule out this possibility. First, Kingdom (2003) showed that adding various amounts of luminance contrast to the depth-enhancing colour gratings made little difference to perceived depth, while colour contrast was a highly salient factor. This surely implies that it was the chromatic, not luminance content of the grating that was responsible for the depth-enhancement. Second, casual observations of pure-luminance plaids indicate that a low contrast luminance component only ever suppresses perceived depth in the orthogonal-in-orientation luminance component.

In the Introduction we put forward the hypothesis that the magnitude of depth-suppression would be lower when the depth-enhancing and depth-suppressing colour directions were the same compared to when they were different. The suggestion was that the visual system might bind plaid components with similar colour composition into a single surface, and interpret any residual luminance component as shading. However, there was

no support for this hypothesis, and therefore no support for the colour-binding idea.

We also hypothesised that blue-yellow chromatic gratings would be less effective depth-suppressors when the blue phase of the grating fell in the dark rather than bright part of the shading grating, on the grounds that the visual system might have knowledge of the fact that shaded regions in the natural visual world tend to be bluer than their non-shaded surrounds. However, we found no support for this hypothesis either. Perhaps the bluing in shadows is not as robust a physical phenomena as assumed here, or perhaps it is normally of such low contrast relative to the shading that it is simply ignored by shape-from-shading mechanisms. It should be noted of course that our choice of blue-yellow was purely subjective and not based on any physical analysis of the spectral characteristics of shading in natural scenes. Interestingly however, the average direction chosen by our subjects for “sky-blue to sunny-yellow” was 121° in the blue direction (see Fig. 3), which is very close to the value of 125° that best captures the range of colours in scenes with significant amounts of blue sky, measured using a similarly-scaled MacLeod–Boynton colour space (Webster & Mollon, 1997).

We found that out-of-phase BY gratings produced less depth-suppression (and in one subject a small amount of depth-enhancement) than in-phase BY gratings, consistent with previous results using red-green gratings (Kingdom, 2003). Why though should there be any depth-suppression with out-of-phase BY gratings, given our claim that non-aligned-with-colour luminance variations are invariably interpreted as shading? One possibility is that relative phase shifts between sinusoidal modulations of colour and luminance do not produce in every subject a categorical shift from ‘aligned’ to ‘non-aligned’ in the same way that, say, a shift in relative orientation from 0° to 90° does, or that a comparable phase shift in a square-wave pattern might. Instead, the shift might be from ‘aligned’ to ‘less-aligned’. Another possibility is that the depth-enhancing $L-M$ grating used in the BY experiment produced a ceiling effect in perceived depth, so that if the BY grating were to have any influence it would only be in the depth-suppressive direction. Previous unreported data showed that when no other chromatic grating is present, depth enhancement is generally observed with out-of-phase red-green gratings. Future experiments are needed to establish the conditions that give rise to depth-enhancement as opposed to depth-suppression with out-of-phase colour gratings.

In our BY experiment we found that in all conditions depth-suppression was greater (or depth-enhancement less) with the 90° compared to 270° condition (see Fig. 2). The effect was not however quite significant using our within-subjects ANOVA. Although there are non-

parametric statistical tests that might elicit a significant result here, parametric tests are more powerful and so we prefer to err on the side of caution and consider this result as a hint of something that needs to be investigated further with a larger number of subjects and conditions. If one compares the two conditions in Fig. 2, one can see that in the 90° case the yellow phase falls in the trough of the perceived depth corrugations, whereas in the 270° case it falls on the peak. Perhaps the latter situation is consistent with the interpretation that sunlight comes from above, and that therefore the luminance variations in the 270° case are more likely to be shading, thus producing less depth-suppression/more depth-enhancement.

Acknowledgement

This research was supported by Canadian Institute of Health Research grant #MOP-11554 given to F.K., and by bursaries awarded to S.R. and K.M. by the Faculty of Medicine at McGill.

References

- Attick, J. J., Griffin, P. A., & Redlich, A. N. (1996). Statistical approach to shape from shading: reconstruction of three-dimensional face surfaces from single two-dimensional images. *Neural Computation*, 8, 1321–1340.
- Cavanagh, P. (1991). Vision at equiluminance. In J. J. Kulikowski, I. J. Murray, & V. Walsh (Eds.). *Vision and visual dysfunction Vol. V: limits of vision* (pp. 234–250). Boca Raton, FL: CRC Press.
- Churma, M. E. (1994). Blue shadows: Physical, physiological, and psychological causes. *Applied Optics*, 33, 4719–4722.
- Cole, G. R., Hine, T., & McIlhagga, W. (1993). Detection mechanisms in L-, M-, and S-cone contrast space. *Journal of the Optical Society of America A*, 10, 38–51.
- Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology*, 357, 241–265.
- DeValois, R. L. (1965). Analysis and coding of color vision in the primate visual system. *Cold Spring Harbor Symposia on Quantitative Biology*, 30, 567–579.
- DeValois, R. L., & DeValois, K. K. (1975). Neural coding of color. In E. C. Carterette & M. P. Friedman (Eds.). *Handbook of perception* (Vol. 5, pp. 117–166). New York, NY: Academic Press.
- Domini, N. J., & Lucas, P. W. (2001). Ecological importance of trichromatic vision to primates. *Nature*, 410, 363–365.
- Eskew, R. T., Jr., McLellan, J. S., & Giulianini, F. (1999). Chromatic detection and discrimination. In K. R. Gegenfurtner & L. T. Sharpe (Eds.), *Color vision, from genes to perception*. Cambridge, UK: Cambridge University Press.
- Gegenfurtner, K. R., & Rieger, J. (2000). Sensory and cognitive contributions of color to the recognition of natural scenes. *Current Biology*, 10, 805–808.
- Gregory, R. L. (1977). Vision with isoluminant colour contrast: 1. A projection technique and observations. *Perception*, 6, 113–119.
- Kingdom, F. A. A. (2003). Colour brings relief to human vision. *Nature Neuroscience*, 6, 641–644.
- Kingdom, F. A. A., & Simmons, D. R. (2000). The relationship between colour vision and stereoscopic depth perception (Invited

- Paper). *Journal of the Society for 3D Broadcasting and Imaging*, 1, 10–19.
- Kingdom, F. A. A., Beauce, C., & Hunter, L. (2004). Colour vision brings clarity to shadows. *Perception*, 33, 907–914.
- Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of colour space. *Vision Research*, 22, 1123–1131.
- Lehky, S. R., & Sejnowski, T. J. (1988). Network model of shape-from-shading: neural function arises from both receptive and projective fields. *Nature*, 333, 452–454.
- Livingstone, M. S., & Hubel, D. H. (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience*, 7, 3416–3468.
- Lu, C., & Fender, D. H. (1972). The interaction of color and luminance in stereoscopic vision. *Investigative Ophthalmology*, 11, 482–489.
- MacLeod, D. I. A., & Boynton, R. M. (1979). Chromaticity diagram showing cone excitation stimuli by stimuli of equal luminance. *Journal of the Optical Society of America*, 69, 1183–1186.
- Mollon, J. (2000). Cherries among the leaves: The evolutionary origins of color vision. In Steven David (Ed.), *Color perception: philosophical, psychological, artistic and computational perspectives* (pp. 10–30). New York: Oxford University Press.
- 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 P. Gouras (Ed.), *The Perception of colour*. In & J. Cronly-Dillon (Eds.). *Vision and visual dysfunction* (Vol. 6, pp. 198–217). Oxford: Macmillan.
- Mullen, K. T., & Kingdom, F. A. A. (2002). A differential distribution of red-green and blue-yellow colour vision across the visual field: evidence for different rules of connectivity? *Visual Neuroscience*, 19, 1–10.
- Mullen, K. T., & Sankeralli, M. J. (1999). Evidence for the stochastic independence of the blue-yellow, red-green and luminance detection mechanisms revealed by subthreshold summation. *Vision Research*, 39, 733–745.
- Norlander, C., & Koenderink, J. J. (1983). Spatial and temporal discrimination ellipsoids in color space. *Journal of the Optical Society of America*, 73, 1533–1543.
- Parraga, C. A., Troscianko, T., & Tolhurst, D. J. (2002). Spatiochromatic properties of natural images and human vision. *Current Biology*, 483–487.
- Ramachandran, V. S. (1988). Perception of shape form shading. *Nature*, 331, 163–166.
- Regan, D. (2000). *Human perception of objects*. Massachusetts: Sinauer Associates: Sunderland.
- Rubin, J. M., & Richards, W. A. (1982). Color vision and image intensities: When are changes material?. *Biological Cybernetics*, 45, 215–226.
- Sankeralli, M. J., & Mullen, K. T. (1996). Estimation of the L-, M-, and S-cone weights of the postreceptoral detection mechanisms. *Journal of the Optical Society of America A*, 14, 2633–2646.
- Sankeralli, M. J., & Mullen, K. T. (1997). Postreceptoral chromatic detection mechanisms revealed by noise masking in three-dimensional cone contrast space. *Journal of the Optical Society of America A*, 14, 906–915.
- Smith, V. C., & Pokorny, J. (1975). Spectral sensitivity of the foveal cone photopigments between 400 and 700 nm. *Vision Research*, 15, 161–171.
- Stromeyer, C. F., III, Cole, G. R., & Kronauer, R. E. (1985). Second-site adaptation in the red-green chromatic pathways. *Vision Research*, 25, 219–237.
- Sumner, P., & Mollon, J. D. (2000). Catarrhine photopigments are optimised for detecting targets against a foliage background. *Journal of Experimental Biology*, 23, 1963–1986.
- Sun, J., & Perona, P. (1997). Shading and stereo in early perception of shape and reflectance. *Perception*, 26, 519–529.
- Webster, M. A., & Mollon, J. D. (1997). Adaptation and color statistics of natural images. *Vision Research*, 37, 3283–3298.