Available online at www.sciencedirect.com

SCIENCE DIRECT®

Vision Research 46 (2006) 814–822

Vision Research

www.elsevier.com/locate/visres

Colour unmasks dark targets in complex displays

Frederick A.A. Kingdom *, Reza Kasrai

McGill Vision Research Unit, Department of Ophthalmology, McGill University, 687 Pine Avenue W, Room H4-14, Montréal, PQ, Que., Canada H3A 1A1

Received 7 June 2005; received in revised form 1 August 2005

Abstract

Recent studies have suggested that colour (meaning chromatic) variations help the visual system segment luminance-variegated displays into their illumination and reflectance layers. This leads to the prediction that colour variations should unmask partially camouflaged achromatic transparencies on luminance-variegated backgrounds. We used 'Mondrian-like' backgrounds that were either achromatic, i.e., varying only in luminance, or chromatic, which in our stimuli meant varying in both luminance and colour. Both achromatic and chromatic backgrounds had the same luminance distribution. Thresholds for detecting simulated transparency targets were found to be lower when on the chromatic compared to achromatic backgrounds. We hypothesised that the chromatic-background advantage resulted from the extra cue provided by colour as to which borders were background and which transparency, predicting that (a) randomising the colours on either side of the transparency border, (b) rotating the target to destroy its X-junctions, and (c) viewing the target eccentrically, would each destroy the chromatic-background advantage. However, none of these predictions was upheld. We suggest therefore that the chromatic-background advantage is due to a low-level, rather than border-disambiguation mechanism. We suggest that chromatic variations reduce the noise, but not the signal, in the mechanism that detects dark targets in complex displays.

Keywords: Colour vision; Transparency; Masking; Contrast detection

1. Introduction

Transparency, described phenomenally as 'seeing through', has long interested the vision science community. The interest is often motivated by quite general questions about scene analysis because a transparency, allthough a material medium, can be considered to be a member of a broad range of illumination effects that include shadows, shading and highlights, all of which cause abrupt changes in luminance and sometimes chromaticity. Many psychophysical studies have sought to determine the necessary and sufficient conditions for generating an impression of transparency (Beck, Prazdny, & Ivry, 1984; Gerbino, 1994; Kasrai & Kingdom, 2001a, 2002; Masin, 1997; Metelli, 1974; Ripamonti & Westland, 2003; Robilotto, Khang, & Zaidi, 2002; Singh & Anderson, 2002; D'Zmura, Colantoni, Knoblauch, & Laget, 1997), and these conditions can be roughly divided into two classes: surface and figural. Surface conditions refer to the luminance and colour (meaning chromatic) requirements for perceived transparency, while figural conditions refer to its geometric requirements, i.e., the arrangement and orientation of its contours.

In this communication we are concerned with the surface conditions for transparency, but not as they concern its impression but its detectability. Since we wish the present study to be relevant to shadows, we have restricted ourselves to transparencies that, like shadows, are achromatic and have no additive component (transparencies with an additive component, sometimes termed transluscencies, are characterised by their milky appearance). When presented on densely variegated achromatic backgrounds, achromatic transparencies may become partly camouflaged. In Fig. 1A, it is easy to appreciate that the achromatic transparency in the upper left of the display would be difficult to detect if it were either fainter, presented briefly or viewed peripherally. The question we ask is whether

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

0042-6989/$ - see front matter
doi:10.1016/j.visres.2005.08.018

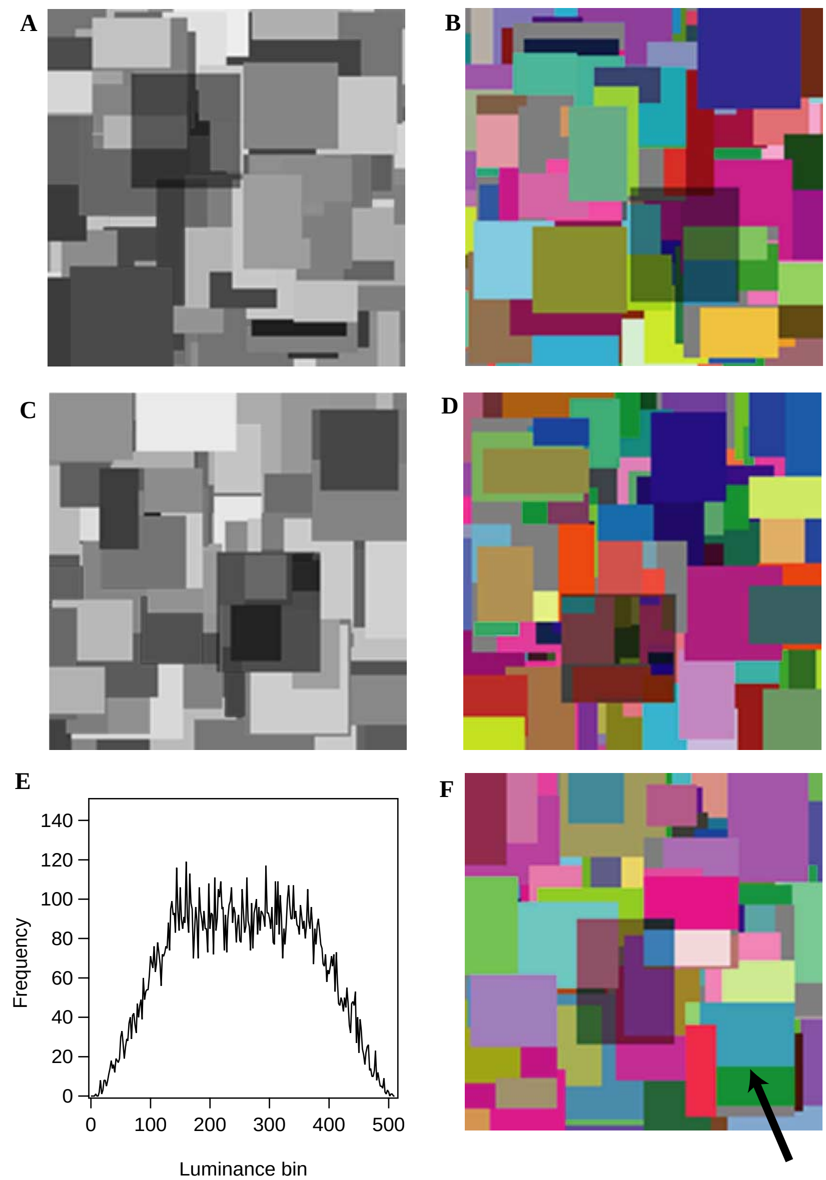

Fig. 1. Example test stimuli. (A) Achromatic background with transparency target; (B) chromatic background with transparency target; (C) achromatic background with rotated target; (D) chromatic background with rotated target; (F) chromatic background with random-colour target. The arrow points to an example distractor—there are six of these in all stimuli (see Section 2); (E) shows a frequency histogram of the luminances of the background rectangles for both chromatic and achromatic backgrounds—see Section 2.

the addition of colour variations to the display—Fig. 1B shows the stimulus—helps to unmask the achromatic transparency under these conditions.

Why might one expect colour variations to unmask achromatic transparencies on luminance-variegated backgrounds? A number of recent studies have demonstrated that colour variations help the visual system to segment displays into what Barrow and Tennebaum (1978) have termed 'intrinsic images',[1] that is their reflectance and illumination layers (Kingdom, 2003; Kingdom, Beauce, & Hunter, 2004; Kingdom, Rangwala, & Hammamji, 2005). The reflectance layer is a representation of the surfaces in a scene, and since most surfaces vary in both spectral as well as intensive reflectance, this layer is in general both colour- and luminance-defined. Shadows and achromatic transparencies on the other hand, which form part of the illumination layer of a scene, impose a divisive reduction in the luminances of surfaces upon which they fall, and are thus luminance-defined.[2] In complex scenes colour is therefore a potential cue for helping discriminate changes in illumination from changes in reflectance via the following rule: luminance variations that are accompanied by col-

[1] Gilchrist (1979, and see Gilchrist et al., 1983) have termed this 'edge classification'.

[2] Shadows can be tinted with colour; for example deep shadows on sunny days tend to be bluish.

our variations are variations in reflectance, while luminance variations that are unaccompanied by colour variations are variations in illumination (Cavanagh, 1991; Mullen & Kingdom, 1991; Rubin & Richards, 1982). In the image processing domain, this rule has recently been applied to help separate the shading and reflectance components of natural images (Marshall, Tappen, William, Freeman, & Adelson, 2003; Olmos & Kingdom, 2004), and psychophysical evidence has revealed that the rule is exploited by the human visual system for 'shape-from-shading' (Kingdom, 2003; Kingdom et al., 2005; Kingdom, Wong, Yoonessi, & Malkoc, 2005), and correctly identifying simulated shadows (Kingdom et al., 2004).

In the last of the aforementioned studies, Kingdom et al. (2004) presented subjects with a stimulus comprising three background sectors that were overlaid with a shadow/transparency that was either correctly or incorrectly simulated. In the achromatic condition, the background sectors differed only in luminance, whereas in the chromatic condition the sectors differed in luminance *and* colour. Subjects were better able to discriminate the correct from incorrect shadows on the chromatic compared to achromatic background. However, when the colours on either side of the border of the shadow/transparency with the background in the chromatic condition were randomised, performance was impaired. Kingdom et al. suggested that the colour variations, provided they were continuous across the border of the shadow/transparency with the background, helped the visual system segregate the shadow/transparency from the background, enabling it to be more efficiently processed. If colour variations improve our ability to *discriminate* correct from incorrect shadows/transparencies, we might expect colour variations to help our ability to *detect* shadows/transparencies that are camouflaged. The first aim of this study is to test this prediction. The targets used here, lacking penumbra, arguably look more like achromatic transparencies than shadows, and so we will refer to them as transparencies from now on. However, we assume our results are equally applicable to shadows.

If our prediction that achromatic transparencies are more easily detected on chromatic compared to achromatic backgrounds is born out, then other predictions follow. Following the idea that continuity of colour across a luminance border is indicative of an illumination change, we should expect that randomising the colours across the transparency border should eliminate the chromatic background advantage. By the same argument, rotating the transparency target, which also destroys the colour continuity, should also eliminate the chromatic background advantage. Finally, viewing the stimuli eccentrically, which renders imperceptible the precise colour composition of the transparency border, should also eliminate the chromatic background advantage.

To determine whether achromatic transparency targets are detected more easily on chromatic (Fig. 1B) compared to achromatic (Fig. 1A) backgrounds, it is essential, we argue, that the two types of background have the same luminance distribution, otherwise any superiority in performance found with the chromatic display could not unequivocally be attributed to its chromatic content. Therefore, we emphasize at the outset that the chromatic backgrounds used in this study are *not* isoluminant. The use of achromatic and chromatic stimuli with the same luminance distribution constitutes a defining characteristic of the stimuli used here as well as in our previous study (Kingdom et al., 2004), and we describe the method employed to achieve this in the following section.

Brief accounts of some of the results reported here have been given elsewhere (Kasrai & Kingdom, 2001b; Kingdom & Kasrai, 2001).

2. Methods

2.1. Stimuli

2.1.1. Generation

All stimuli were generated using the VSG2/3F videographics card (Cambridge Research Systems) hosted by a Gateway 2000 P5 computer and displayed on a BARCO CCID 7551 monitor.

2.1.2. Stimulus backgrounds

Fig. 1 shows example stimuli. Each stimulus was $7.3 \times 7.3°$ at the viewing distance of 100 cm and was positioned in the middle of the monitor screen on a mid-grey background of 13.9 cd/m^2. The stimulus background contained 125 randomly positioned rectangles, clipped at the border of the stimulus where necessary. Both the heights and widths of the rectangles were randomly selected from the range 0.43–1.29°, producing rectangles with a variety of areas and aspect ratios.

The luminances and colours of the background rectangles were assigned as follows. For each stimulus three image planes—red, green, and blue—were created in the VSG's video memory. The three planes were alternated in sequence at 160 Hz; thus each plane was refreshed at approximately 53.3 Hz. While the red plane was displayed, only the red phosphor was activated, while the green plane was displayed only the green phosphor was activated, and while the blue plane was displayed only the blue phosphor was activated. There were 256 linearly spaced intensity levels available for each image plane, resulting in a total of 256^3 possible luminance-colours in the combined three-plane image. Let each luminance-colour be described as a 'tristimulus' value, rgb, where r, g, and b are phosphor intensities expressed as proportions of their maximum luminances, $r_{\max}$, $g_{\max}$, and $b_{\max}$. The maximum phosphor luminances were measured using the same frame-alternating sequence employed for the presentation of the stimuli, but with two of the phosphors set to zero luminance. The values were $r_{\max} = 5.5$ cd/m^2, $g_{\max} = 19.7$ cd/m^2, and $b_{\max} = 2.6$ cd/m^2.

For the chromatic background, each rectangle was randomly allocated a tristimulus value $r_c g_c b_c$ from the full gamut available. The *luminance* L_c of the rectangle is thus

$$L_c = r_c r_{max} + g_c g_{max} + b_c b_{max}. \quad (1)$$

In order that the achromatic background rectangles were selected from the same luminance distribution as the chromatic background rectangles, we first generated random tristimulus values as for the chromatic background (i.e., $r_c g_c b_c$), and then calculated the tristimulus value r_a, g_a, b_a that resulted in a luminance L_a equal to L_c but under the constraint that $r_a = g_a = b_a$. The formula is

$$r_a = g_a = b_a = \frac{r_c r_{max} + g_c g_{max} + b_c b_{max}}{(r_{max} + g_{max} + b_{max})}. \quad (2)$$

This method for equating the luminance distributions from which the chromatic and achromatic background rectangles are selected is robust to any differences in spectral sensitivity between observers. Suppose for example that our estimate of r_{max} underestimated the sensitivity to the red phosphor of a subject by a factor of two. Doubling r_{max} however has on the equality between L_c and L_a, nor therefore on the equality of the luminance distributions for the achromatic and chromatic conditions.

Fig. 1E shows the theoretical distribution of the background rectangle luminances obtained from 16,000 randomly generated tristimulus values, using Eq. (1) to calculate the rectangle luminances. As can be seen, the distribution has a flattish top with sloping tails.

2.1.3. Targets

All targets had an area of 1.3 square deg and a height-to-width aspect ratio that could vary randomly between 0.67 and 1.5. The targets were randomly positioned anywhere on the background with the constraint that the edges of the target were at least 0.43° from the edge of the background.

There were three types of target, termed 'transparency', 'rotated', and 'random colour', and examples of each are shown in Fig. 1. The transparency target was formed by multiplying the pixel intensities in the designated region of all three image planes by t, where $0 < t < 1.0$. A value of t close to unity produced a target with a high physical transparency/low opacity, whereas a value of t close to 0 produced a target with low physical transparency/high opacity. Example transparency targets with $t = 0.5$ are shown in Figs. 1A and B.

The rotated target was produced according to the same principle as the transparency target, except that once created it was reflected across both the vertical and horizontal axis before being superimposed on the background. The result was a target with the same luminance and colour composition as the transparency target, but without any X-junctions (except by chance) along the border with the background. Figs. 1C and D show example rotated targets with $t = 0.3$.

The random-colour target was generated according to the same principle as the transparency target, except that once created, its internal colours were re-allocated new random colours, *but without changing their luminances*. This condition was not applicable to the achromatic background, and an example with $t = 0.5$ is shown in Fig. 1F.

2.1.4. Target distractors

A final feature of our stimuli is the addition of rotated target distractors with $t = 1.0$ to all stimulus backgrounds. To understand the reason for these distractors, consider the potential shape of the psychometric function for each of the three types of target. Remember that the aim is to measure proportion correct detections as a function of t (actually $\log(1 - t)$, see below), in order to calculate a threshold. For the transparency target, as t approaches unity, performance must inevitably fall to chance, as the target literally disappears when $t = 1.0$. However, this is not the case for the rotated and random-colour targets. With $t = 1.0$, these targets do not disappear, because they remain defined by a continuous border around their edge. In pilot experiments we found that performance was around 65% for the $t = 1.0$ rotated and random-colour targets on purely random-rectangle backgrounds. To obtain psychometric functions that spanned 50–100% correct, we added six randomly positioned rotated target distractors with $t = 1.0$ to all stimulus backgrounds prior to adding the target. The distractors were added one at a time and therefore some distractors by chance occluded parts of others. The target was always added last so that it would never be occluded. The presence of the target distractors was found to reduce performance to chance for all types of target with $t = 1.0$. An example distractor is shown by the arrow in Fig. 1F.

2.2. Procedure

We used the method of constant stimuli with a 2IFC (two-interval-forced-choice) procedure. Subjects were instructed to '...detect the dark, rectangular targets in the display, which are of various degrees of transparency'. Subjects were given plenty of practice trials to familiarise themselves with the different types of target and background. On each trial, two stimuli were presented for 250 ms each, one with and one without a target. The task for the subject was to indicate by key press the interval containing the target. A tone for an incorrect decision was provided for feedback. In each block, the background was either chromatic or achromatic, one type of target was presented, and six levels of t were presented (in random order). In each block there were 120 trials, so each t was presented 20 times. The t values for each background–target combination were selected to produce performance in the range 50–100%, as determined by pilot experiments. Examples of the ranges of t employed for some of the conditions are given in Section 3, but t ranged overall from 0.95 to 0.05, the extremes corresponding perceptually to 'very faint and highly transparent' and 'near black and opaque'. The spacing of the six t values was arranged such that when converted to $\log(1 - t)$, the values were equally spaced. For each background–target condition six blocks were run, making a total of 720 trials per psychometric function.

Measurements were made at 0°, 5°, and 10° eccentricity. For the latter two eccentricities, subjects were required to fixate a spot positioned to the right of the stimulus on the horizontal meridian. Since the stimulus was 7.3° wide, and the target randomly positioned within the stimulus, the eccentricity of each target varied somewhat around its specified mean value.

2.3. Subjects

A total of seven subjects were employed, the two authors (FK and RK) and five naïve subjects (LH, CB, HW, AY, and SC) who were undergraduate volunteers. All subjects had normal, or corrected-to-normal acuity, and normal colour vision.

2.4. Data analysis

The logistic function, $0.5 + 0.5 * \exp((x - a)/b)/(1.0 + \exp((x - a)/b))$, with $x = \log(1 - t)$, was fitted to the proportion correct data. a and b were free parameters that determined the threshold at 75% correct and the slope of the psychometric function. The functions were fit using Igor (Wavemetrics), with the reciprocal of the binomial standard deviation of each data point used as a weighting function.

3. Results

Example psychometric functions for one of the naïve observers (CB) are shown in Fig. 2. Each plot shows proportion correct target detections against $\log(1 - t)$. As can be seen, all plots show a more-or-less monotonic increase in performance with $\log(1 - t)$. The measure of $\log(1 - t)$ can be considered a measure of the darkness of the targets relative to their background, and in the case of the transparency targets, a measure of their opacity. In terms of t, the abscissa values range 0.95–0.4 for the left, 0.95–0.3 for the middle and 0.8–0.05 for the right graph. The left graph shows psychometric functions for the 0° eccentricity transparency target on both chromatic and achromatic backgrounds; the middle graph psychometric functions for the 5° eccentricity transparency and random-colour targets on the chromatic background; the right graph psychometric functions for the 10° eccentricity rotated targets on both chromatic and achromatic backgrounds.

The thresholds calculated from the psychometric functions are shown in Fig. 3. The top row of histograms shows thresholds for the transparency target on both chromatic and achromatic backgrounds, and the random-colour targets. The bottom row shows the thresholds for the rotated targets on both chromatic and achromatic backgrounds. Different graphs within each row are for different eccentricities.

To illustrate more directly the relationship between viewing eccentricity and stimulus condition, Fig. 4 plots thresholds as a function of eccentricity for the three subjects who produced data for all eccentricities.

It is clear from Figs. 3 and 4 that thresholds for targets on chromatic backgrounds are lower than for targets on achromatic backgrounds, both for transparency and rotated targets. There appears to be little or no difference between thresholds for transparency and random-colour targets in the chromatic background condition. Although thresholds rise with eccentricity, there appears to be no interaction between eccentricity and the magnitude of the chromatic background advantage. To confirm these observations, we performed two types of statistical analysis: first, t tests on the 0° eccentricity conditions, for which data on five subjects was collected (though not identical subjects for transparency and rotated targets); second, ANOVAS for the three subjects who performed at all three eccentricities. The t tests for the 0° eccentricity data were within-subject and tested for significant differences (two-tailed) between (a) transparency targets on chromatic versus achromatic backgrounds, (b) rotated targets on chromatic versus achromatic backgrounds, and (c) transparency versus random-colour targets on chromatic backgrounds. The ANOVAS were also within-subject and tested for the same significant differences as the t tests, but included the factor eccentricity. The results of both analyses are shown in Table 1.

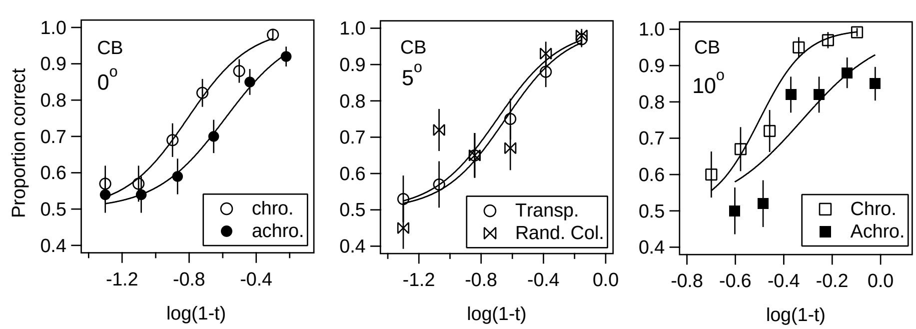

Fig. 2. Example psychometric functions for one of the naïve subjects (CB). Each plot shows the proportion correct detections as a function of $\log(1 - t)$. The continuous lines are best-fitting logistic functions. Error bars are binomial standard deviations. Left 0°, middle 5°, and right, 10° eccentricity. Chro., chromatic; achr., achromatic; transp., transparency; rotated (rot.); rand. col., random colour.

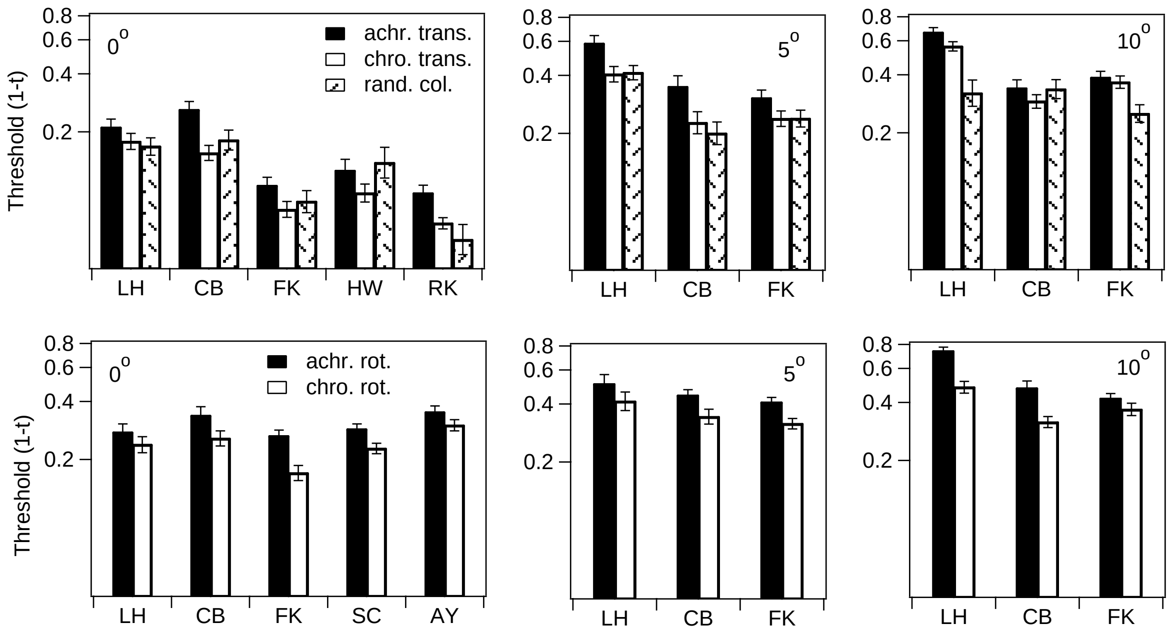

Fig. 3. Thresholds for detecting the dark target. Top: thresholds for chromatic transparency (chro. trans.), achromatic transparency (achr. trans.) and random-colour (rand. col.) conditions. Bottom: thresholds for chromatic rotated and achromatic rotated conditions. Left 0°, middle 5°, and right, 10° eccentricity.

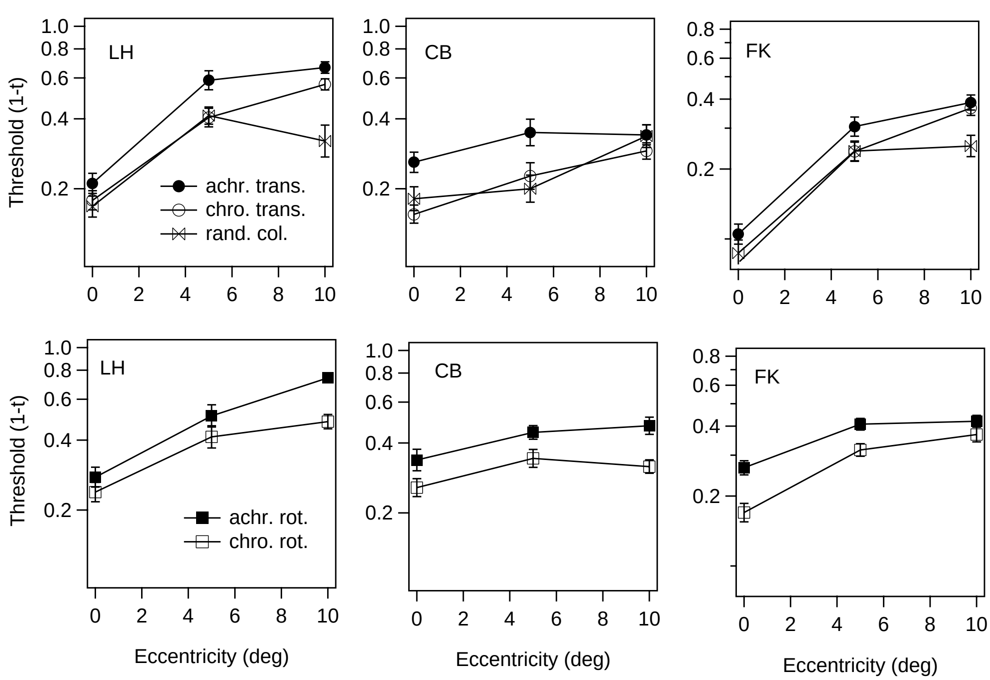

Fig. 4. Thresholds for detecting the dark target as a function of type of background and eccentricity for three subjects. chro., chromatic; achr., achromatic; transp., transparency; rot., rotated; rand. col., random colour.

As can be seen from the right-hand column of Table 1, both analyses reveal significant differences in thresholds between targets presented on chromatic versus achromatic backgrounds, for both transparency and rotated targets. There are no significant differences between the transparency and random-colour targets on the chromatic backgrounds. Eccentricity is a significant main effect, but there are no significant interactions between eccentricity and other factors.

4. Discussion

Achromatic transparency targets on luminance-variegated backgrounds are more easily detected when the back-

Table 1
Statistical analysis of results

Eccentricity	# Ss	Factors and comparisons	Test	Significance $^{*}p < 0.05$
0°	5	Transparency target, chromatic V achromatic background	$t(4) = 5.549$	$p = 0.005^{*}$
0°	5	Rotated target, chromatic V achromatic background	$t(4) = 4.657$	$p = 0.01^{*}$
0°	5	Transparency V random-colour target, chromatic background	$t(4) = 0.75$	$p = 0.495$
All	3	A Transparency target, chromatic V achromatic background	$F(1,2) = 26.1$	$p = 0.036^{*}$
		B Eccentricity	$F(2,4) = 12.1$	$p = 0.02^{*}$
		A × B	$F(2,4) = 4.63$	$p = 0.091$
All	3	A: Rotated target, chromatic V achromatic background	$F(1,2) = 366$	$p = 0.003^{*}$
		B: Eccentricity	$F(2,4) = 13.32$	$p = 0.017^{*}$
		A × B	$F(2,4) = 0.22$	$p = 0.815$
All	3	A Transparency V random-colour target, chromatic background	$F(1,2) = 1.08$	$p = 0.408$
		B: Eccentricity	$F(2,4) = 13.32$	$p = 0.017^{*}$
		A × B	$F(2,4) = 1.653$	$p = 0.3$

ground is also colour-variegated. This confirms the first prediction of the study. For the remaining three predictions, namely that the chromatic background advantage in target sensitivity would disappear when (a) the colours were randomised across the transparency target border, (b) when the transparency targets were rotated, and (c) when the targets were viewed eccentrically, there was no support. We will return to an examination of these unsupported predictions after first considering the significance of the main finding of the study.

How large is the chromatic background advantage, at least as measured in our displays? If we average across subjects and eccentricities, the percentage difference in opacity $(1 - t)$ thresholds between the chromatic and achromatic conditions is 41% for the transparency targets and 31% for the rotated targets. These are both less than a factor of two. Therefore while the chromatic background advantage appears to be robust and significant, it is not large. It remains to be seen whether there are more optimal conditions for eliciting the chromatic background advantage.

Many studies have considered whether luminance- and chromatic-sensitive mechanisms interact in visual tasks. Of these, the studies that have measured contrast detection thresholds for luminance- and colour-defined patterns such as sine-wave gratings or disks are most directly relevant to the present study. The results of these studies have for the most part indicated that chromatic- and luminance-sensitive mechanisms operate independently at detection threshold (e.g., Cole, Hine, & McIlhagga, 1993; Mullen & Losada, 1994; Mullen & Sankeralli, 1999; Wandell, 1985; exceptions are Gur & Akri, 1992; Switkes, Bradley, & DeValois, 1988). For example, Mullen and Losada (1994) found that near-threshold levels of a chromatic grating mask failed to facilitate the detection of an achromatic grating test, and vice versa, provided positional cues to the test were eliminated. This lack of facilitation of an achromatic test by a chromatic mask would seem at odds with the results found here, but it must be borne in mind that the intensive properties of our stimuli are quite different from those employed by Mullen and Losada. In our stimuli, both the chromatic and achromatic backgrounds had the same, suprathreshold level of luminance contrast. For this reason, we consider our findings an instance of chromatic unmasking rather than chromatic facilitation, since the thresholds for our targets are presumably much higher than they would be were the background a uniform field. The analogous experiment to ours using gratings would be to measure contrast thresholds for a luminance grating test in the presence of a suprathreshold luminance grating mask (i.e., a contrast increment threshold), and repeat the measurements with an added chromatic grating mask. However, to our knowledge there is no published data for this stimulus condition, and unpublished experiments in our laboratory using gratings have so far failed to find evidence for chromatic unmasking.

What might be the mechanism of chromatic unmasking? First, consider the information content of our stimuli. In terms of *luminance* content, the chromatic and achromatic conditions have the same information, as we equalised the luminance variance in the two conditions (see Section 2). In terms of *colour* content, this depends on the condition. For the transparency targets, the chromatic stimulus is more information rich than the achromatic stimulus. In the chromatic condition, the background borders change in luminance *and* colour, whereas the transparency borders change in luminance but *not* colour, and this provides a cue to the target not present in the achromatic stimulus. By the same reasoning, the information content in the chromatic transparency condition is greater than that of the chromatic random-colour condition. In the case of the rotated targets, the information content of both chromatic and achromatic stimuli is presumably the same. Therefore, if information content determined performance, we would expect lower thresholds for the transparency targets on chromatic compared to achromatic backgrounds, which we found. However, we would also expect lower thresholds for the transparency compared to random-colour targets on the chromatic background, and similar thresholds for rotated targets on both chromatic and achromatic backgrounds, neither of which was found. Taken together, our results are not consistent with the idea that information content alone determines performance.

The second candidate for chromatic unmasking is reflectance-illumination disambiguation. In Section 1 we discussed how it has been shown that chromatic variations help the visual system to segment complex displays into their illumination (which here includes transparency) and reflectance layers, by reducing the uncertainty as to which borders are illumination and which reflectance. This led us to propose three manipulations that should eliminate the chromatic background advantage. However, none of the manipulations had the desired effect, and we must therefore conclude that it is unlikely that the chromatic background advantage with our stimuli is due to any kind of border disambiguation. In this respect, the data here are quite different from those of Kingdom et al. (2004) where the task was to discriminate correctly from incorrectly simulated shadows; in that study we found that randomising the colours across the shadow border significantly impaired discrimination.

The third possible cause for chromatic unmasking is some kind of low-level contrast gain adjustment. Some years ago, Switkes et al. (1988) 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. At first glance, our results are inconsistent with this idea, as we found that a luminance-defined target was *better* detected when chromatic variations were present than when not. Nevertheless, there is still pertinence to Switkes et al's idea. What might be happening in our stimuli is that the chromatic variations partially suppress the luminance contrasts *within but not between* target and background. In signal detection terms, this translates to the notion that chromatic variations suppress the noise but not the signal. But how might this happen? One can rule out a simple contrast gain control mechanism whereby the chromatic contrasts reduce the variance in luminance-contrast response to the stimuli, as this would impact equally both signal and noise. A possible mechanism is one in which luminance-contrast-sensitive neurones receive inhibitory inputs from similarly tuned neurones outside of their classical receptive fields, as has been demonstrated for a class of neurones in monkey V1 (Kapadia, Westheimer, & Gilbert, 2000; Knierim & van Essen, 1992; Nothdurft, Gallant, & van Essen, 1999). If these neurones were disinhibited by inputs from colour-sensitive neurons in their surrounds, they would become more sensitive to luminance contrast.

What is it about our stimulus arrangement that produces chromatic unmasking? Is the articulated background the critical feature, and if so, is it important that the chromatic and luminance borders are spatially aligned? At this stage we do not have answers to these questions. Such questions will not be easy to answer because of the small size of chromatic unmasking, which leaves little room for exploration. However, we are presently carrying out studies to explore the possible factors.

In conclusion, we have shown a positive role for colour vision in the analysis of image structure that is not explicable in terms of the reflectance–illumination disambiguation role of colour vision that we anticipated might be the cause. Rather, our results point to a role for colour vision in suppressing the luminance noise that in complex displays impairs the detectability of luminance-defined targets. If this is indeed a new role for colour vision then it is one that not only complements colour vision's beneficial role in intrinsic-image segmentation, but also its beneficial role in the analysis of other aspects of image structure, such as the detection of fruit and flowers in dense foliage (Domini & Lucas, 2001; Mollon, 1989; Sumner & Mollon, 2000).

Acknowledgment

This research was supported by Canadian Institute of Health Research Grant #MOP-11554 given to F.K.

References

Barrow, H. G., & Tennebaum, J. M. (1978). Recovering instrinsic scene characteristics from images. In A. R. Hanson & E. M. Riseman (Eds.). *Computer vision systems* (Vol. 99, pp. 3–26). New York: Academic Press.

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

Cavanagh, P. (1991). Vision at equiluminance. In J. J. Kulikowski, I. J. Murray, & V. Walsh (Eds.), *Vision and visual dysfunction volume V: Limits of vision* (pp. 234–250). Boca Raton, FL: CRC Press.

Cole, G. R., Hine, T. J., & McIlhagga, W. H. (1993). Detection mechanisms in L-, M-, and S-cone contrast space. *Journal of the Optical Society of America A, 10*, 38–51.

Domini, N. J., & Lucas, P. W. (2001). Ecological importance of trichromatic vision to primates. *Nature, 410*, 363–365.

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

Gilchrist, A. L. (1979). The perception of surface blacks and whites. *Scientific American, 240*, 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 and Psychophysics, 33*, 425–436.

Gur, M., & Akri, V. (1992). Isoluminant stimuli may not expose the full contribution of color to visual functioning: Spatial contrast sensitivity measurements indicate interaction between color and luminance processing. *Vision Research, 32*, 1253–1262.

Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (2000). Spatial distribution of contextual interactions in primary visual cortex and in visual perception. *Journal of Neurophysiology, 84*, 2048–2062.

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

Kasrai, R., & Kingdom, F. A. A. (2001b). Articulated chromatic backgrounds facilitate detection of illumination boundaries. *Perception, 30*(Suppl.), 3.

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

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

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

Kingdom, F., & Kasrai, R. (2001). Discriminating material from illumination changes in complex scenes. *Proceedings of the Society for Imaging Science and Technology: PICS, 2001*, 5–8.

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

Kingdom, F.A.A., Wong, K., Yoonessi, A. & Malkoc, G. (2005). Colour contrast influences perceived depth in combined shading and texture patterns. *Spatial Vision*, in press.

Knierim, J. J., & van Essen, D. C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *Journal of Neurophysiology, 67*, 961–980.

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

Marshall, F., Tappen William, T., Freeman & Adelson, E. H. (2003). Recovering intrinsic images from a single image in advances. In *Neural Information Processing Systems (NIPS)* (pp. 15). MIT Press.

Metelli, F. (1974). The perception of transparency. *Scientific American, 230*, 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: Gouras, P. (Ed.), The perception of colour. Vol. 6 in Vision and visual dysfunction, by J. Cronly-Dillon (Ed.) Macmillan: Oxford, 198–217.

Mullen, K. T., & Losada, M. A. (1994). Evidence for separate pathways for color and luminance detection mechanisms. *Journal of the Optical Society of America A, 11*, 3136–3151.

Mullen, K. T., & Sankeralli, M. J. (1999). Evidence for the stochastic independence of the blue–yellow, red–green and luminance detection mechanisms revealed by sub-threshold summation. *Vision Research, 39*, 733–745.

Nothdurft, H. C., Gallant, J., & van Essen, D. C. (1999). Response modulation by texture surround in primate area V1: Correlates of "popout" under anesthesia. *Visual Neuroscience, 16*, 15–34.

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

Ripamonti, C., & Westland, S. (2003). Prediction of transparency perception based upon cone excitation ratios. *Journal of the Optical Soceity 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*, 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). Towards a perceptual theory of transparency. *Psychological Review, 109*, 492–519.

Sumner, P., & Mollon, J. D. (2000). Catarrhine photopigments are optimised for detecting targets against a foliage background. *Journal Experiment Biology, 23*, 1963–1986.

Switkes, E., Bradley, A., & DeValois, K. K. (1988). Contrast dependence and mechanisms of masking interactions among chromatic and luminance gratings. *Journal of the Optical Soceity of America A, 5*, 1149–1162.

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

Wandell, B. A. (1985). Color measurement and discrimination. *Journal of the Optical Soceity of America A, 2*, 62–71.