



The role of chromatic contrast and luminance polarity in stereoscopic segmentation

Frederick A.A. Kingdom ^{a,*}, Hyung-Chul O. Li ^b, Elizabeth J. MacAulay ^a

^a McGill Vision Research Unit, 687 Pine Avenue West, Room H4-14, Montréal, Québec, Canada H3A 1A1

^b Department of Industrial Psychology, Nowon Gu, Wolgae Dong, Kwangoon University, Seoul, South Korea

Received 17 March 1998; received in revised form 22 February 2000

Abstract

We have investigated whether our ability to discriminate the stereoscopic depth of random-dot targets set amongst random-depth distractors is facilitated when target and distractors differ in particular combinations of colour and luminance polarity. For flat-plane targets, stereo-thresholds were found to be lower with a target-distractor colour/luminance difference, but only when that difference enabled the target elements to be identified in the monocular image, either by virtue of being less numerous than the distractors, or because the subject knew beforehand the target's colour/luminance. If neither of these conditions prevailed, stereoscopic thresholds were no different when target and distractors were segregated by colour/luminance than if they were not. For sine-wave disparity grating stimuli, in which subjects were required to discriminate the orientation of the depth corrugations, no advantage was found when target and distractors were segregated by colour/luminance under any condition. These results suggest that segregation by colour/luminance is only beneficial to the stereoscopic processing of random-element stimuli when the task can be performed by attending to a small number of target elements. A corollary to this conclusion is that stereopsis mechanisms do not automatically pre-filter the image into different colour/luminance maps. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Colour; Stereopsis; Image segmentation

1. Introduction

Stereoscopic vision is an important means of determining the three-dimensional (3-D) structure of the world around us. In many natural scenes, such as forests, the visual system is confronted with dense arrays of elements spread across multiple depth planes. Stereopsis is useful for perceptually breaking up such scenes, often revealing otherwise camouflaged objects. Typically, figure and ground differ not only in depth, but also along other dimensions, such as brightness, colour, size and orientation. Are stereopsis judgements facilitated by the presence of these other cues?

The random-dot-stereogram (Julesz, 1971) is the laboratory version of 'figure-ground' segregation in stereopsis. In this study, we consider whether figure-ground differences in the colour and luminance polarity

(i.e. whether the element is darker or brighter than the background) of elements making up random-dot stereograms facilitate stereopsis. This is a different question from the one normally asked with regard to colour and stereopsis, namely whether stereopsis is possible at isoluminance, i.e. when only chromatic information is present (see Simmons & Kingdom, 1997; Kingdom & Simmons, 2000, for recent summaries). In principle, it is possible that colour could help stereoscopic judgements even if those judgements were impossible at isoluminance (that many stereoscopic judgements are possible at isoluminance is for the present discussion besides the point) by 'colour-labelling' the signals emanating from luminance-contrast-tuned disparity detectors. Chromatic information could help constrain the way outputs of luminance-based disparity detectors were integrated.

There are a variety of ways in which chromatic information could do this. One would be by helping solve stereo-correspondence, for example if bright red

* Corresponding author. Fax: +1-514-8431691.

E-mail address: fkingd@po-box.mcgill.ca (F.A.A. Kingdom).

was only allowed to match with bright red, dark green with dark green, etc. There is evidence that with stimuli possessing luminance contrast, both colour (Treisman, 1962; Julesz, 1971; Ramachandran, Rao, Sriram, & Vidyasagar, 1973; Akerstrom & Todd, 1988; Jordan, Geisler, & Bovik, 1990; Kovacs & Julesz, 1992) and luminance polarity (see review in Howard & Rogers, 1995) can serve to reduce the ‘false-target’ problem in stereopsis. A second means by which colour/luminance differences could facilitate luminance-based stereopsis would be to help identify the target elements, even if the observer had no prior knowledge of the actual target colour/luminance. For example, a target consisting of a few red dots on a background of a larger number of green dots (or vice-versa) will tend to ‘pop-out’ from the background; this monocular form cue to the target might enable its elements to be selected for depth analysis. A third means would be if elements with common colours/luminances were automatically grouped, or pre-filtered, into separate colour/luminance maps. This mechanism could operate in the absence of any unambiguous monocular cues to the target, for example if target and background elements were differently coloured but equally numerous, and one did not know which colour was which. A fourth means would be if the observer had prior knowledge of the target colour/luminance, and selectively attended to it. Again, this mechanism could operate for those targets that, as in the previous example, were not identifiable monocularly without such prior knowledge. It is with the last three of these potential means, namely monocular form cues, automatic grouping/pre-filtering, and prior-knowledge-based selective attention, that this study is concerned.

Some recent studies have considered whether colour/luminance differences can facilitate motion processing in noise (Edwards & Badcock, 1994, 1996; Croner & Albright, 1997; Li & Kingdom, 1998, 2000; Snowden & Edmunds, 1999) and the results of these studies partly motivate the present study on stereopsis. Croner and Albright employed the ‘global motion’ paradigm, in which subjects were required to discriminate the direction of motion of a set of coherently moving target dots amongst incoherently moving distractor dots. They found that motion-coherence thresholds (the minimum number of target dots required for correct motion-direction identification) were lower when target and distractors differed in colour, even though the subject did not know which colour was target and which distractor on each trial. Li and Kingdom (1998, 2000) suggested that in this paradigm the differently-coloured target dots could be picked out in the static view of the stimulus by virtue of their relatively small number (typically motion-coherence thresholds are of the order of 10–15%). In other words a static form cue to the target produced the lower thresholds. Li and Kingdom

(2000) went on to show that when target and distractor dots were equal in number (50% each) and the target thus not identifiable in the static view of the stimulus (the target colour was again unknown on each trial), performance was no better in the segregated compared to the non-segregated condition¹. They concluded that motion mechanisms do not automatically group, or pre-filter, elements with similar colours/luminances for motion processing. One of the aims of this study is to consider whether a similar conclusion holds for stereopsis.

In the experiments described below we measured stereoscopic discrimination in random-dot-stereograms whose target and distractor dots were either segregated or non-segregated on the basis of colour/luminance. The results have helped us to clarify the circumstances where colour/luminance differences do, and do not, facilitate stereoscopic discrimination in random-dot displays.

2. Methods

2.1. Subjects

Two of the authors, FK and EM, acted as subjects. Both were emmetropic with normal colour vision. FK was a highly experienced psychophysical observer. EM was naïve as to the purpose of the experiments, when she acted as an observer.

2.2. Stimuli

2.2.1. Generation and calibration

The stimuli were generated by VSG2/3F video-graphics card (Cambridge Research Systems) hosted by a Gateway 2000 P5 computer, and displayed on a BARCO Calibrator monitor. The monitor frame-rate was 107 Hz. The VSG2/3F displays images with 8-bit-per-gun (256 level) resolution, but the pixel intensities can be selected from 12-bit (4096 levels) linearised colour look-up-tables (CLUTs). Each gun on the monitor was calibrated using the Optical system (Cambridge Research Systems), which generates the 12-bit gamma-corrected CLUTs. The 12-bit CLUTs provided a contrast resolution of about 0.05%. Whatever the contrast of the stimulus, it was always displayed with the full

¹ Because motion coherence thresholds are typically around 10–15% target dots, Li and Kingdom (2000) changed the stimulus parameters in the 50% target dot condition to reduce performance below the ceiling of 100% correct. They reduced the between-frame dot displacement size from 0.1 to 0.034° per frame, with the result that performance for the 50% target dot conditions was around 80% correct.

8-bits, the intensities of which were suitably selected from the 12-bit CLUTs. Only the red R, and green G phosphors were employed, their chromaticity coordinates being, respectively, $x = 0.623$, $y = 0.340$; $x = 0.278$, $y = 0.584$.

2.2.2. Gaussian micropatterns — ‘dots’

We employed Gaussian, rather than hard-edged, micropatterns to minimise the effects of chromatic aberration, which with hard-edged stimuli can introduce spurious luminance contrasts. These ‘dots’ were generated using the function:

$$L(x, y) = M + A \exp\left[\frac{-(x^2 + y^2)}{2\sigma^2}\right]$$

with M background luminance (which depended on whether the red or green gun was modulated) A amplitude (depending on the density of dots) and σ , the space constant of 0.15° . The function was clipped at a diameter of 1.0° . Six types of dot were generated, by combining modulations of the red and green guns in various ways, as illustrated in Fig. 1. The six types of dots were ‘red’, ‘green’, ‘bright red’, ‘bright green’, ‘dark red’ and ‘dark green’.

Two types of random-dot-stereogram were employed and are shown in achromatic form in Figs. 2 and 3. In all stimuli the x - y position of each dot was

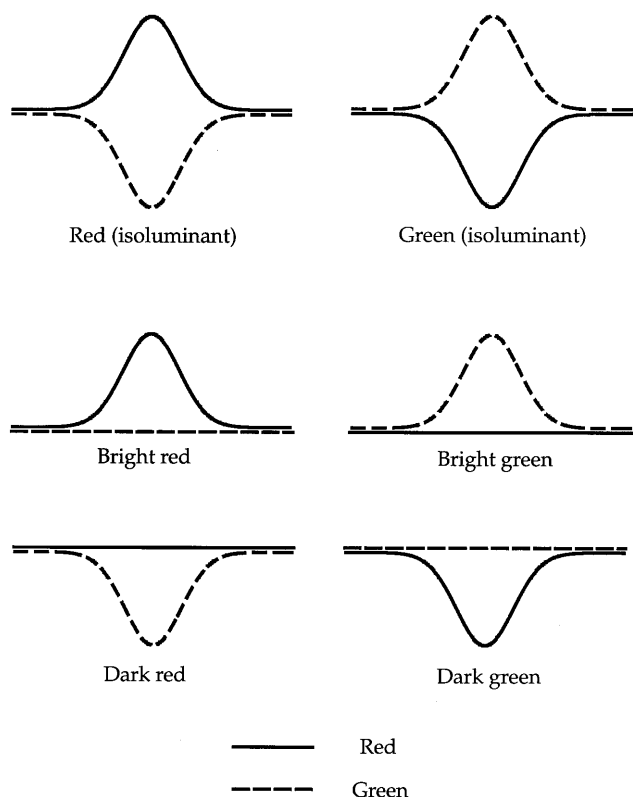


Fig. 1. Red and green gun modulation profiles of the various Gaussian dots employed.




Fig. 2. Flat-target stimuli in black-and-white. The actual stimuli consisted of a mixture of bright red, dark red, bright green and dark green dots. The target is a flat plane of dots lying either towards the front or back of the random-depth distractor dots (depending on how the figure is free-fused). The top figure shows the segregated condition; the target, comprising 15% of the dots, are bright dots, while the distractors are dark dots. Because the target dots are in a minority there is a monocular form cue to the target which does not depend on prior knowledge of the actual target dot type. The bottom figure shows the non-segregated condition; the bright and dark dots are equal in number, and the 15% of target dots are made from equal proportions of both. In this figure, the target is effectively camouflaged, making it much more difficult to determine its depth.

chosen randomly within a $10.8 \times 10.8^\circ$ stimulus window. For the stimuli in the main part of the experiment, four types of dot were used — bright red, dark red, bright green and dark green. When two dots overlapped, their respective red and green modulations (though not dc levels) were added separately (i.e. red with red, green with green), which resulted in their colour and luminance contrasts also adding separately. For example, if a bright red dot fell precisely on a dark green dot, both the colour and the luminance contrasts of the two dots would cancel to produce a uniform mid-yellow the same colour as the background. On the other hand, if a bright red dot fell precisely on a bright green dot the result would be a bright yellow dot. This form of additivity was achieved through the creation of two separate arrays of dots for each stereo-half, each array associated with its own CLUT. One array/CLUT was designated for the dark red and bright green dots, the other array/CLUT was for the dark green and bright red dots. The two arrays/CLUTs were presented on alternate frames of the monitor, resulting in the perception of a single stimulus with all four types of dot.

The contrast of each display was measured in terms of the component Gaussian dots. The colour and luminance contrasts of the dots were based on the Weber contrasts of their R (red) and G (green) modulations, $\Delta R/R_b$ and $\Delta G/G_b$, where ΔR and ΔG were the amplitudes and R_b and G_b background luminances. The colour contrast of a dot was then calculated as $(\Delta R/R_b - \Delta G/G_b)/2$ and its luminance contrast as $(\Delta R/R_b + \Delta G/G_b)/2$. In the main experiments whenever the red or green component of a dot was Gaussian modulated it was always with a Weber contrast of 50%. Thus the four main types of dot (dark red, bright red, dark green, bright green) had each a luminance contrast of 25% and a colour contrast of 25%. R_b and G_b (and hence also ΔR and ΔG as contrast was constant) were set according to the $R/(R+G)$ value determined at isoluminance for each subject (see below for details). This produced a yellow background with mean luminance 8.0 cd m^{-2} .

2.2.3. Gaussian dot disparity

In order to define the disparity of the dots with sub-pixel accuracy, the dots were drawn from a selection of ‘templates’, stored off-screen in the VSG’s video memory. Each template was a square patch of pixels within which a dot was positioned. The Gaussian profile of each dot was offset from the centre of the template with sub-pixel precision by half its disparity, and hence the total number of templates was twice the number of required disparities. When the position of a




Fig. 3. Sinusoidal-target stimuli in black-and-white. The target is an obliquely oriented disparity-grating with about four cycles, set amongst an equal number of random-depth distractors. In the top figure the target and distractors are segregated by dot-type, whereas in the bottom they are non-segregated.

dot had been randomly selected, its disparity was calculated and the appropriate template pair selected and copied into their positions in each stereo-half.

2.2.4. Flat-target stereograms

Fig. 2 shows example stimuli. Each stereo-half consisted of 300 dots (in the cyclopean view), each defined with 25% colour and 25% luminance contrast. The ‘target’ was a flat plane of dots embedded in random-disparity ‘distractor’ dots. The target dots were positioned at a fixed disparity either in front or behind the average disparity of the distractor dots. The disparity of the distractor dots was chosen from a random (flat) distribution in the range -50 to $+50$ arcmin (total range 100 arcmin). Flat-target stereograms with different proportions of target dots were employed. In the achromatic versions of the stimuli shown in Fig. 2a the target and noise dots are of opposite luminance polarity — the ‘segregated’ condition, while in Fig. 2b both target and noise are made from an equal mixture of luminance polarities — the ‘non-segregated’ condition.

2.2.5. Sinusoidal-target stereograms

Fig. 3 shows example stimuli. The target was a random-dot, sine-wave disparity grating, embedded in random-depth distractors. The disparities of the target dots were defined according to a sinusoidal function, whose mean was zero, whose orientation was left or right oblique (-45 or $+45^\circ$), and whose spatial frequency was 0.3 c deg^{-1} . The phase of disparity modulation was always randomised on each presentation. The disparities of the distractor dots were chosen randomly from a specified range, which was either -12.5 to $+12.5$ arcmin (range 25 arcmin) or, -50 to $+50$ arcmin (range 100 arcmin).

2.2.6. Stereo-presentation

The stereo-pairs were presented on either side of the monitor screen, separated by 55 arcmin. They were combined optically by a modified 8-mirror Wheatstone stereoscope. All mirrors were cemented into position except for the two front mirrors whose position along the line of sight could be adjusted until fusion was accomplished. Viewing distance to the stimulus along the mirror path was 55 cm.

2.3. Procedure

We measured stereo-thresholds in all experiments except the experiment measuring the isoluminant point, where contrast thresholds were obtained for a criterion depth judgement (see below for details). In the case of the flat-target stimuli, the stereo-threshold was defined as the minimum detectable difference in disparity between the two targets in a two-interval-forced-choice (2IFC) pair, one presented in front of the mean distrac-

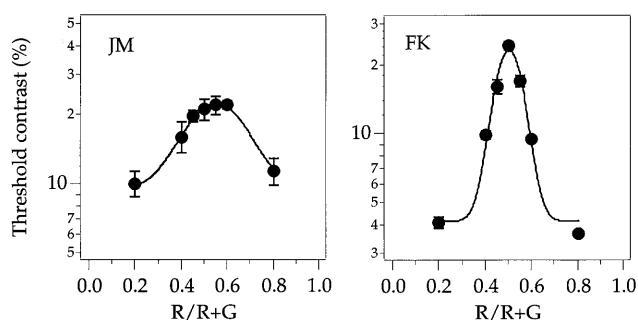


Fig. 4. Measurement of the isoluminant point. Contrast thresholds for identifying the orientation of a single-surface disparity grating are shown as a function of the $R/(R+G)$ ratio. The continuous curves are the best fitting Gaussian functions, used to estimate the $R/(R+G)$ ratio producing highest thresholds.

tor disparity, the other behind. For the sinusoidal-target stimuli, the stereo-threshold was the minimum amplitude of disparity modulation required to identify correctly the orientation (left or right oblique) of the corrugations in a single stimulus presentation. Before each session the subject was required to adapt to a blank yellow screen at the isoluminant $R/(R+G)$ ratio for 1 min. A button press recorded the subjects' responses. A tone accompanied each stimulus presentation, and a different tone indicated an incorrect decision. Stimulus exposure duration was 500 ms. A standard 'two-up, one-down' staircase procedure was employed (Levitt, 1971) to obtain stereo-thresholds at the 70.7% correct level. Target disparities were altered by a factor of 1.25 (increase or decrease) during the staircase. The staircase was terminated after 12 reversals and the threshold calculated as the geometric mean disparity over the last ten reversals. At least three thresholds were measured for each condition, and unless otherwise stated, the data points in the figures give the geometric means and geometric standard errors of these measurements.

3. Experiments

3.1. Measurement of isoluminant point

Although the aim of the study was to examine stereoscopic performance using stimuli with both luminance and chromatic contrast, we wanted to minimise the possibility that differences in the chromatic contrasts of dots introduced artifactual differences in their luminance contrasts. We, therefore, established the relative red R, and green G, mean luminances — the $R/(R+G)$ ratio — which resulted in an 'isoluminant' stimulus according to the criterion of 'worst performance'. To do this we measured contrast thresholds for determining the orientation of a disparity grating similar to the

one shown in Fig. 3 as a function of the $R/(R+G)$ ratio. The disparity grating consisted of 1500 dots, and had a spatial frequency of disparity modulation of 0.075 c deg^{-1} , with a fixed amplitude of 15 arcmin. The procedure was the same as that used for measuring stereo-thresholds (see Section 2). The results are shown in Fig. 4. At low and high $R/(R+G)$ ratios contrast thresholds were low, and rose steeply to a point roughly mid-way along the $R/(R+G)$ axis. To obtain an accurate estimate of the $R/(R+G)$ value producing the highest contrast thresholds, we fitted a Gaussian function to the logarithmically transformed ordinate values, and estimated the $R/(R+G)$ ratio at the peak of the fitted function. The result was 0.49 for FK and 0.55 for EM. These isoluminant $R/(R+G)$ ratios were then used throughout the study.

3.2. Experiment 1. Stereo-segmentation by colour/luminance in the presence of monocular form cues

The aim of this experiment was to determine whether stereo-thresholds for a target consisting of a relatively small number of dots amongst random-depth distractors was superior when target and distractors differed in colour/luminance. With the target dots in a minority there is a monocular form cue to the target in the segregated-by-colour/luminance conditions; targets are potentially identifiable in the monocular image without prior knowledge of the target colour/luminance. Fig. 2 shows achromatic versions of the flat-target, and Fig. 3 the sinusoidal-target stimuli. All stimuli comprised bright red, dark red, bright green and dark green dots. The flat-target stimuli consisted of a flat plane of dots set amongst random-depth distractors, and thresholds for discriminating the disparity of the target (front vs. behind) were obtained. The sinusoidal-target stimuli consisted of sinusoidally-modulated-in-depth dots set amongst random-depth distractors, and thresholds for identifying the orientation of the corrugations (left- or right-oblique) were obtained. For both types of target, performance was measured as a function of the proportion of target dots, which varied from 10 to 50%. There were four test conditions: 'colour-segregated', 'luminance-segregated', 'colour-by-luminance-segregated', and 'non-segregated'. In the colour-segregated condition, the target consisted of one colour of dot (e.g. red, both bright and dark), while the distractors were of the opposite colour (e.g. green, both bright and dark), the choice of target colour being randomised for each forced-choice pair. In the luminance-segregated condition, the target consisted of one luminance polarity of dot (e.g. bright, both red and green), with opposite luminance polarity distractors (e.g. dark, both red and green), again with target luminance polarity randomised. In the colour-by-luminance-segregated condition the target differed from the distractors in a

combination of colour and luminance polarity. For example on some trials the target would consist of bright red and dark green dots, the distractors dark red and bright green dots, while on other trials it would be vice versa. Finally in the non-segregated condition all four types of dot were present in equal proportions in both target and distractors.

The results are shown in Fig. 5a and b for the flat-target and sinusoidal-target stimuli. For the flat-target stimuli, thresholds in the non-segregated condition (filled squares) are highest when the proportion of target dots is 10%, falling precipitously with an increase in target dot proportion to near-asymptotic levels at about 20–30%. For the three classes of segregated condition, performance is more-or-less constant across target dot proportion. The pattern of performance is different for the sinusoidal-target stimuli; for both segregated and non-segregated conditions thresholds fall in parallel as the proportion of target dots increase.

The results for the flat-target stimuli are easy to interpret. In the non-segregated condition there was no monocular form cue to the target, and this would be expected to be most detrimental when the target was very sparse and hence heavily camouflaged (< 50%). For the sinusoidal-target stimuli, however, it appears

that being able to identify the target dots in the segregated condition confers no benefit, as thresholds were no different for segregated and non-segregated conditions. We will return to a consideration of the sinusoidal-target results in Section 4.

Consistent with our previous arguments concerning the analogous situation in the motion domain (Li & Kingdom, 2000), we argue here that the results found with the flat-target stimuli do not imply that stereopsis mechanisms automatically group, or pre-filter elements into separate colour/luminance maps in order to process their depth properties. We argued previously that evidence for automatic pre-filtering can only be obtained under conditions where there are no extraneous cues to help identify the target. In the stereopsis domain, this requires a stimulus with no monocular cues to the target. In the experiment just reported, only one of the conditions satisfies this constraint — the 50% target condition (the rightmost data point on each graph in Fig. 5) — and for neither subject did there appear to be any difference between the segregated and non-segregated thresholds in this condition. This suggests that dots with similar colour/luminance properties are not automatically pre-filtered for stereoscopic processing. However, it is possible that with the 50% target dot condition, performance was at near-asymptotic levels, and the difference between the segregated and non-segregated conditions was obscured by a ceiling effect. The next experiment was designed to test for automatic pre-filtering at target dot proportions where there was no possible ceiling effect.

3.3. Experiment 2. Stereoscopic segmentation by colour/luminance in the absence of monocular form cues

The purpose of this experiment was to test whether stereopsis mechanisms automatically pre-filter dots into similar colour/luminance maps in the absence of monocular cues to the target. To avoid possible ceiling effects associated with target dot proportions of 50% while still eliminating monocular target cues, we used the following arrangement. Targets were formed from 25% of the dots, with the remaining 75% of dots distractors. All four dot types — bright red, bright green, dark red or dark green — were present in equal proportions in all conditions. In the segregated conditions, one of the four dot types was the target, the other three distractors. In the non-segregated condition, both target (again 25%) and distractors (again 75%) were made from equal numbers of all four types of dot. It is important to realise that even though the target dots were in a minority in this experiment, they could not be identified unambiguously in the monocular view of the segregated condition, as all four dot types were present in equal numbers, and the target colour was ran-

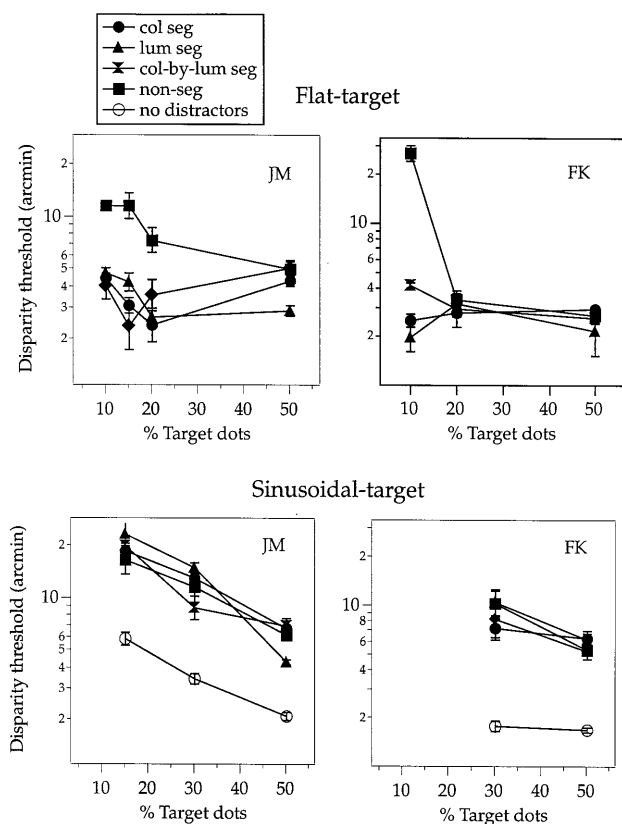


Fig. 5. Results from Section 3.2. Stereo-thresholds are plotted as a function of the proportion of target dots. Seg, segregated; col, colour; lum, luminance; non-seg, non-segregated.

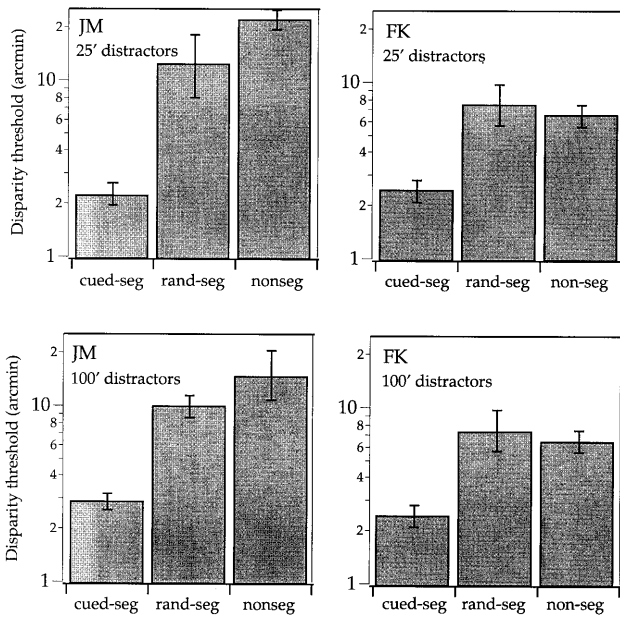


Fig. 6. Results from flat-target stimuli in Section 3.3, in which the target was 25% of the dots. Results are shown for two subjects and two distractor depth ranges (25 and 100 arcmin). In the cued-segregated (cued-seg) condition the target was one of four types of dot and subjects knew which type was the target before each session and was encouraged to attend to it. The distractors were made from equal numbers of the remaining three types of dot. In the randomly-selected (rand-seg) condition the target dots were again of just one type of dot, but the type of dot was randomly selected on each trial. In the non-segregated (non-seg) condition the target comprised all four dot types in equal proportions, as were the distractors.




Fig. 7. Results from sinusoidal-target stimuli in Section 3.3, with distractors set to a range of 25 arcmin. See Fig. 6 legend for further details.

domised on each trial. This arrangement, however, raises an interesting question. What if the subject did know the colour of the target dots in the segregated condition, and was encouraged to attend to that colour? Would performance in the segregated condition be better than in the non-segregated condition? We therefore included conditions where in the segregated condition the target colour/luminance was fixed throughout the session, and subjects were cued as to the particular target colour/luminance.

The three test conditions were as follows. In the

'cued-segregated' condition, only one of the four types of dot was the target, the remaining being distractors, and the subject was informed prior to the session which dot type made up the target. Data were collected for all four target dot types. In the 'random-segregated' condition, again only one dot type was the target, but the target dot type was randomly selected on each trial from the four available. So while the subject was aware that on all trials the target was segregated from the distractors by dot type, he/she did not know which dot type was the target on each trial. In the 'non-segregated' condition, both target and distractors were made from all four dot types in equal proportions; the subject was made aware of the nature of the condition.

The results are shown in Fig. 6 for the flat-target stimuli, and Fig. 7 for the sinusoidal-target stimuli. First consider the flat-target results in Fig. 6. Data are shown for distractors set to two disparity ranges, 25 and 100 arcmin. There is no consistent difference between the random-segregated and non-segregated conditions for either disparity range of distractor. However, thresholds are significantly lower for all cued-segregated conditions.

Turning now to the results for the sinusoidal stimuli in Fig. 7, for which only data for the 25 arcmin distractor range was gathered, the results also showed no difference in thresholds between the random-segregated and non-segregated test conditions. However, unlike with the flat-target stimuli, the cued-segregated sinusoidal-target thresholds are no lower than in the other two conditions. For the sinusoidal-target stimuli we also measured thresholds in the absence of distractors (for the flat-target stimuli the distractors were necessary as the reference frame), to demonstrate that the absence of any cued-segregation superiority in performance was not a result of the distractors being irrelevant to the task. As the white bars show, the no-distractor comparison conditions produced significantly lower thresholds than all three test conditions, showing that the distractors indeed degraded performance when present.

4. Discussion

The results of the study can be summarised as follows.

1. For random-dot, flat targets set amongst random-depth distractors, stereo-depth discrimination is facilitated when target and distractors are segregated by colour/luminance under two circumstances; (a) if segregation provides a monocular form cue to the target dots by virtue of the target dots being in a minority, or (b), if the subject has prior knowledge of the target dot type, and selectively attends to it.

2. For random-dot, sinusoidal targets set amongst random-depth distractors, there are no apparent circumstances in which segregation of target and distractors by colour/luminance facilitates stereopsis.

The results with flat-target stimuli complement recent studies involving the analogous situation in the motion domain. Discriminating the direction of motion of coherently moving dots set amongst a larger number of incoherently moving distractors, is improved when the target and distractors differ in colour/luminance (Croner & Albright, 1997; Snowden & Edmunds, 1999; Li & Kingdom, 2000). We have suggested previously that static form cues to the target underlie this result (Li & Kingdom, 1998, 2000). By the same token, a monocular form cue to the target was present in the < 50% target dot conditions in our first experiment, and for the flat-target stimuli this appeared to benefit performance. However, as Li and Kingdom (2000) showed with the global motion paradigm, segregation of the target by colour/luminance did not benefit motion direction discrimination in the absence of a static target cue, unless the target dots were pre-cued and selectively attended to (see also Edwards & Badcock, 1994, 1996; Snowden & Edmunds, 1999). Similarly here, in the absence of monocular form cues to the target there were no differences in stereoscopic discrimination thresholds between segregated and non-segregated conditions.

Unlike with the flat-target stimuli, however, our experiments with sinusoidal-target stimuli revealed no circumstances in which segregation by colour/luminance facilitated stereoscopic discrimination. Why might this be so? A possible clue emerges from a comparison of the nature of the flat-target and sinusoidal-target stimuli. With the flat-target stimulus, the task could in principle be performed by noting the depth of just one target dot in each forced-choice pair. Indeed, inspection of Fig. 5a shows that in the segregated conditions of Section 3.2, there was little effect of the proportion of target dots on performance; subjects performed the task just as well with 10% as with 50% target dots. In fact, the superior performance of the segregated over the non-segregated conditions in Section 3.2 was primarily a result of the rapid rise in thresholds of the non-segregated conditions as the proportion of dots decreased. With the sinusoidal-target stimulus, on the other hand, Fig. 5b shows that reducing the proportion of target dots was equally detrimental for both segregated and non-segregated conditions. This suggests that with the sinusoidal-target stimuli, subjects invariably integrated as much of the dot-depth information as possible, irrespective of the colours/luminances of the dots, and consequently incurred no benefit from their segregation. Thus a likely reason for the different pattern of results with the flat-target and

sinusoidal-target stimuli lies in the different requirements for successfully performing the task; only when a few dots are needed to perform the task does segregation by colour/luminance incur a benefit (and see Snowden & Edmunds, 1999, for similar arguments applied to the global motion paradigm). Thus segregation by colour/luminance is beneficial to stereoscopic discrimination only when it is possible to perform the task by attending to a small number of target elements. This becomes possible either because there are monocular form cues to the target, or because of prior knowledge about the target colour/luminance; in either case this enables attention to be directed to the target.

4.1. No automatic pre-filtering by colour/luminance for stereoscopic processing

The finding common to both the flat-target and sinusoidal-target stimuli, was that in the absence of either monocular target cues, or the possibility of selective attention to the target, stereopsis was no better in the segregated compared to non-segregated conditions. This appears to demonstrate that the visual system does not automatically pre-filter the stimulus into separate colour/luminance maps for stereoscopic processing. If automatic pre-filtering occurred, performance should have been better in the segregated compared with non-segregated conditions, because in the segregated condition one of the hypothetical pre-filtered maps had a near-perfect signal. Why not automatic pre-filtering? One reason that we can reject is that local disparity detectors are not sensitive to either colour or luminance polarity; the evidence against this was given in Section 1. A more plausible explanation comes from an examination of the perceptual needs of animals living in densely textured environments, such as forests. In these environments it does not necessarily make sense for the visual system to pre-filter the image into separate colour/luminance maps prior to stereopsis and motion processing. Consider a multi-coloured object such as a butterfly or leopard lying behind foliage. The animal's multi-colouring provides camouflage by 'breaking up' its structure; stereopsis and motion perception are important mechanisms for *breaking* this camouflage. If motion and stereopsis mechanisms operated on separate colour/luminance maps, extra processing would be necessary to integrate the various maps to reveal the object's structure; better simply to concentrate on those features common to the whole object, namely its motion and depth properties, and ignore colour/luminance (a caveat to this argument is that in those situations in which a rapid response is required, stereopsis is possibly too slow as a 'camouflage-breaker' to be useful).

The situation, however, would be quite different if the object's shape was clearly visible by virtue of a

difference in colour/luminance with its background, or if there was prior knowledge of the object's colours/luminances. Then it would make sense for the visual system to first segregate the object from its background before processing its motion and depth properties, as we have demonstrated in this study.

Acknowledgements

This research was supported by a grant from the MRC (Canada), # MT11554, given to Frederick A.A. Kingdom, and a McGill University Medical Faculty award given to Elizabeth J. MacAulay.

References

- Akerstrom, R. A., & Todd, J. T. (1988). The perception of stereoscopic transparency. *Perception and Psychophysics*, *44*, 421–432.
- Croner, L. J., & Albright, T. D. (1997). Image segmentation enhances discrimination of motion in visual noise. *Vision Research*, *37*, 1415–1427.
- Edwards, M., & Badcock, D. R. (1994). Global motion perception: interaction of on and off pathways. *Vision Research*, *34*, 2849–2858.
- Edwards, M., & Badcock, D. R. (1996). Global-motion perception: interaction of chromatic and luminance signals. *Vision Research*, *36*, 2423–2431.
- Howard, I. P., & Rogers, B. J. (1995). *Binocular vision and stereopsis*. Oxford: Oxford University Press.
- Jordan, J. R., Geisler, W. S., & Bovik, A. C. (1990). Color as a source of information in the stereo correspondence process. *Vision Research*, *30*, 1955–1970.
- Julesz, B. (1971). *Foundations of cyclopean perception*. Chicago: University of Chicago Press.
- Kingdom, F. A. A., & Simmons, D. R. (2000). The relationship between colour vision and stereoscopic depth perception. *Journal of the Society for 3D Broadcasting and Imaging*, *1*, 10–19.
- Kovacs, I., & Julesz, B. (1992). Depth, motion, and static-flow perception at metaisoluminant color contrast. *Proceedings of the National Academy of Sciences of USA*, *89*, 10390–10394.
- Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America*, *49*, 467–477.
- Li, H.-C. O., & Kingdom, F. A. A. (1998). Does segregation by colour/luminance facilitate the detection of structure-from-motion in noise? *Perception*, *27*, 769–784.
- Li, H.-C. O., & Kingdom, F. A. A. (2000). Segregation by colour/luminance does not necessarily facilitate motion discrimination in the presence of motion distractors. *Perception and Psychophysics*, in press.
- Ramachandran, V. S., Rao, V. M., Sriram, S., & Vidyasagar, T. R. (1973). The role of colour perception and 'pattern' recognition in stereopsis. *Vision Research*, *13*, 505–509.
- Simmons, D. R., & Kingdom, F. A. A. (1997). On the independence of chromatic and luminance stereopsis mechanisms. *Vision Research*, *37*, 1271–1280.
- Snowden, R. J., & Edmunds, R. (1999). Colour and polarity contributions to global motion perception. *Vision Research*, *39*, 1813–1822.
- Treisman, A. (1962). Binocular rivalry and stereoscopic depth perception. *Quarterly Journal of Experimental Psychology*, *14*, 23–37.