Pergamon

0042-6989(95)00196-4

Vision Res., Vol. 36, No. 9, pp. 1265–1279, 1996

Printed in Great Britain
0042-6989/96 $15.00 + .00

Contour Integration with Colour and Luminance Contrast

WILLIAM H. McILHAGGA,* KATHY T. MULLEN*†

Received 25 April 1994; in revised form 30 December 1994; in final form 5 July 1995

In this study, we consider how colour contrast can be used to integrate form and how it interacts with luminance contrast in the task. The performance of form integration was assessed by measuring the detection of a winding "contour" of aligned gabor elements embedded in a background of randomly oriented gabors, using both luminance and isoluminant (red/green) contrast. Performance on the task improves with gabor element contrast, and identical performance for colour and luminance contour detection is achieved at high screen contrasts, showing that colour is able to support a complex form integration task. In a second experiment, we investigate whether colour and luminance contrast can be combined in contour integration by measuring the detection of a path with alternating isoluminant colour and luminance elements. We find that contour detection uses both colour and luminance information cooperatively, but performance is much poorer than would be expected from a single common contour integration process which fails to distinguish the two types of contrast. This suggests that there are specific contour integration processes for colour and luminance. In a third experiment, we measure the effects of variations in colour and luminance contrast on contour detection using elements that combine colour and luminance contrast. We find that varying the colour contrast of elements tends to worsen the detection of a luminance contour, as do luminance contrast variations for colour contour detection. These results suggest no special role for colour in integrating contours, and are discussed with regard to their ecological significance‡.

Colour Isoluminance Form Contour Spatial

INTRODUCTION

The primary goal of vision is to generate descriptions of the world from the retinal image. Our descriptions are largely structured in terms of objects and surfaces, and thus one task of vision is to segment the retinal image into regions, each of which contains points imaged from one object. Image segmentation appears to be based on low-level cues, such as continuity, since there are many demonstrations showing how continuity can override other, presumably higher level, pictorial organizations such as symmetry, repetition and pattern, recognizable shapes, and Prägnanz (Koffka, 1935; Kanisza & Gerbino, 1982; Rock, 1983). Image segmentation may be based on image discontinuities, which often mark points in an image where the projection of one object ends and another begins. In early vision, discontinuities are sensed by local oriented detectors (for convenience we will call these "edge detectors"). Segmentation may occur by integrating the outputs of these local edge detectors into longer contours, which delineate the boundaries of object images. A number of computational approaches to image segmentation are based on this two stage process of local edge detection followed by boundary integration (Marr, 1982; Grossberg & Mingolla, 1985; Parent & Zucker, 1989; Heitger & von der Heydt, 1993). These computational studies have emphasized the importance of common orientation of the edges comprising the contour: two nearby local edges are more likely to come from the same contour if they are aligned.

Image discontinuities are caused by, among other things, the different surface properties of the objects, particularly their spectral reflectance. Measurements of surface reflectance show that the greatest variability occurs in the overall albedo—simply the total amount of light reflected (Parkkinen, *et al.*, 1989). However, in normal scenes albedo is confounded with the illuminant intensity, so it is a poor cue without further complicated processing [some consequences of which can be seen in, e.g. Adelson (1994)]. On the other hand, the ratio of reflected light in two different spectral bands (chromaticity) is invariant with illuminant intensity, and so provides a very simple cue to surface material (Rubin & Richards, 1982; Brill, 1990). Chromaticity or colour may

*Department of Ophthalmology H4-14, McGill Vision Research, 687 Pine Avenue West, Montreal, Canada H3A 1A1.

†To whom all correspondence should be addressed [*Email* kmullen@-violet.vision.mcgill.ca].

‡These results were presented at an annual meeting of the Association for Research in Vision and Ophthalmology (McIlhagga *et al.* 1993).

(a) (b) (c) (d)

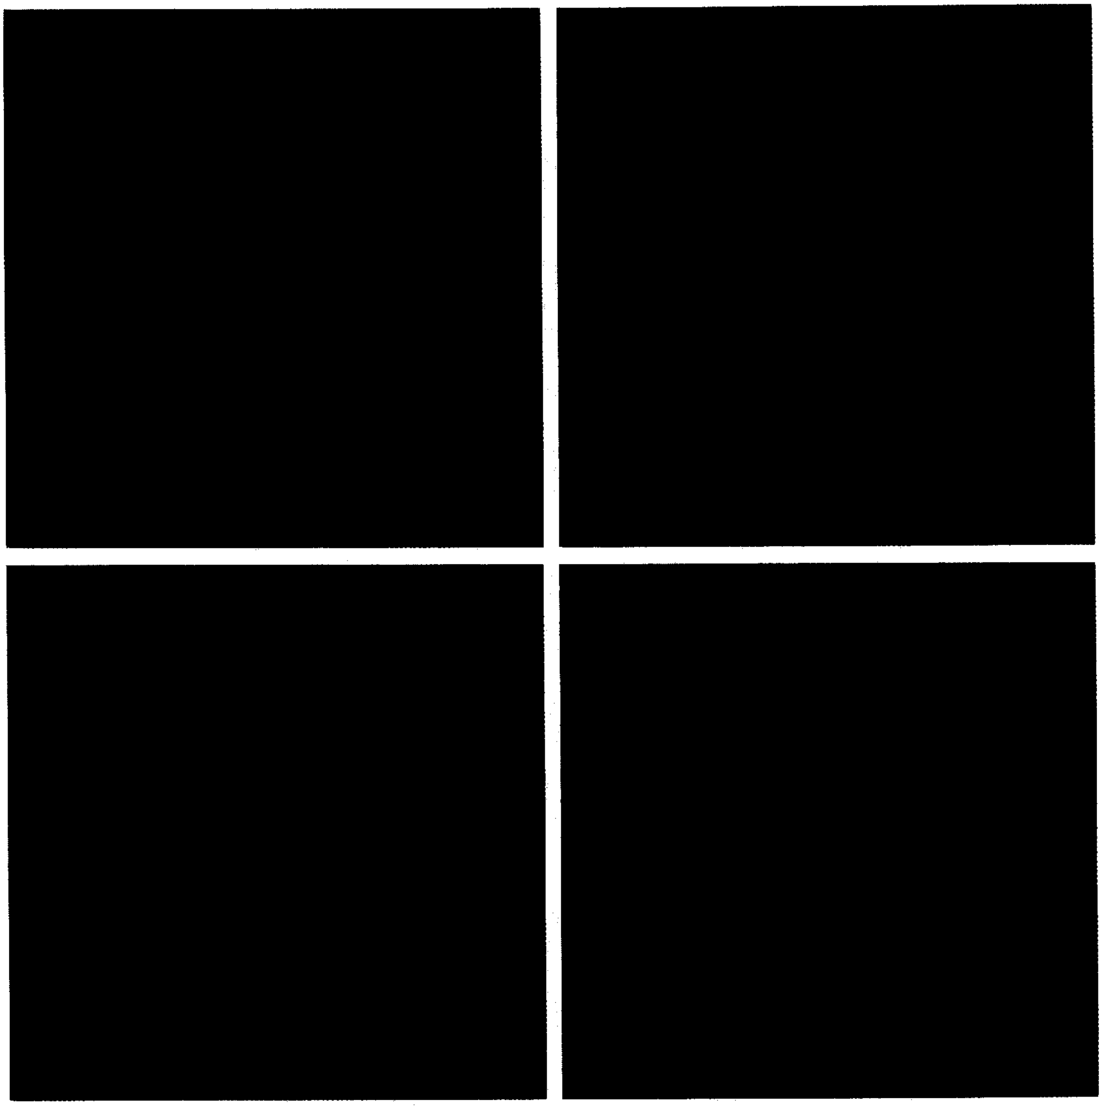

FIGURE 1. Representations of the stimuli used. (a) Luminance gabor elements only, and the "contour" (a string of 10 gabor elements with a contour angle of 0 deg) lies to the left of the figure. (b) The red–green isoluminant condition. The contour is not the same as the one appearing in (a), but also lies to the left of the figure. (c) Both luminance and isoluminant elements alternated in the contour and randomly distributed in the background. The contour is more difficult to see, but lies horizontally across the middle of the figure. (d) Elements that combine colour and luminance contrast. The contour has elements with the same colour contrast, but with the luminance contrast alternating in sign. It curves across the middle of the figure.

thus serve as a "linking feature" (Barlow, 1981) for image segmentation. The potential uses for colour in form perception and image segmentation have been discussed by Mollon (1989) and Mullen and Kingdom (1991). There are at least two ways colour may be involved. First, regions delineated by luminance boundaries may be later grouped on the basis of a common interior colour. This is well established experimentally. In the Ishihara test, spots are grouped together on the basis of a similar colour to form recognizable shapes. Furthermore, the colour of texture elements has been shown to mask other texture segregations, e.g. that based on orientation (Morgan *et al.*, 1992). Some models of perception also use colour this way, relegating it to merely "filling-in" regions delineated by luminance defined boundaries (Grossberg & Mingolla, 1985; Livingstone & Hubel, 1987). The second way colour could be used is at the region boundaries themselves. If each local edge is also sensed by a colour-contrast system, the colours on each side of the edge could assist in the grouping of edges into boundaries.

In the experiments described in this paper, we will be examining whether the local colour contrast at edges influences the integration of those edges into a continuous boundary. The colour system has the characteristics necessary for sensing edges. As with luminance vision, there appear to be bandpass channels in the colour system (Switkes *et al.* 1983; Losada & Mullen, 1994, 1995), and orientation can be successfully discriminated at isoluminance (Webster *et al.*, 1990), both of which properties are sufficient for edge detection. The question is whether this colour edge information can be used in the boundary integration processes of segmentation. To examine boundary integration, we use a stimulus and task introduced by Field *et al.* (1993), based on earlier work by Uttal (1975), Smits *et al.* (1985), and Beck *et al.* (1989). In the stimulus, a set of oriented gabor elements are placed along a winding contour, and embedded in a field of randomly scattered, randomly oriented gabor elements. The task is to detect the presence of the elements of the contour [see e.g. Fig. 1(a)]. The detectability of the contour serves as a measure of how successful the visual system is at integrating the contour elements into a single form. These studies have shown the importance of common orientation in contour integration, and lend support to the computational models mentioned above. All these previous experiments have used luminance stimuli, however, and the contrast of the stimuli was not explored as a potential boundary integration cue.

We perform three experiments. The purpose of the first experiment is to quantify the effects of contrast magnitude on contour detection in both the luminance and colour domain. We find that contour detection is contrast dependent, and at high screen contrasts detection appears to reach asymptotic performance levels. Similar contour detection performance is reached with either luminance or isoluminant colour elements. However, when luminance and isoluminant contrasts are equated for orientation discrimination, there is a relative deficit in colour contour detection. The second set of experiments test whether colour and luminance contours are detected independently, or whether there is some cooperation between colour and luminance pathways, by measuring detection of a contour composed of alternated luminance and isoluminant gabor elements. We find that the detection of such an alternated colour–luminance contour is better than that expected if contour integration is entirely independent for colour and luminance contrast, but worse than expected if it is assumed that the contour integration mechanism does not distinguish between luminance and colour. The third and last experiment looks at the effects of combined colour and luminance contrast in contour detection. From ecological considerations, the integration of a colour contour should not be affected by variations in luminance contrast along it, since this corresponds to the case of an object seen under spatially varying illumination such as dappled light or shadows. On the other hand, one might expect that varying colour contrast should have a major effect on the integration of a luminance defined contour, since the colour of an object usually varies little. We did not however see such a pattern in the detection results to support such a selective role for colour, and in fact they are subject to considerable individual variation.

METHODS

In all experiments, the observer's task was to identify what we call "contour" stimuli and "no-contour" stimuli. A contour stimulus consisted of a set of oriented gabor elements aligned along a common contour, embedded in a background of similar, but randomly oriented gabor elements. A no-contour stimulus consisted of only randomly placed gabor elements. Gabor elements were used to limit the spatial bandwidth of the stimuli. Though real contours have a wide range of spatial frequencies, it is not clear how the information from different spatial frequencies interacts in the visual system. Restricting the stimulus to a limited bandwidth should simplify these interactions. Additionally, the peak spatial frequency of the elements was low (1.5 c/deg), which reduces luminance artifacts due to chromatic aberration. In this section, we first describe the construction of these stimuli, then the experimental protocol, and finally our definition of colour and luminance contrast.

Stimuli

The gabor elements used to construct the stimuli were defined by the equation

$$g(x,y,\theta) = c\sin(2\pi f(x\sin\theta + y\cos\theta))\exp\left(-\frac{x^2+y^2}{2\sigma^2}\right) \quad (1)$$

where θ is the element orientation from 0 to 360 deg, (x, y) is the distance in deg from the element centre, and c is the contrast. The sinusoidal frequency f is 1.5 c/deg, and the space constant σ is 0.17 deg.

A "no-contour" stimulus was constructed with the following algorithm. A 14.1 deg wide square was divided into a 14 × 14 grid of equally sized cells. A gabor element of random orientation was placed in each cell of the display, with the restriction that each grid cell contained the centre of only one gabor element. This prevents the clumping of elements that would occur if

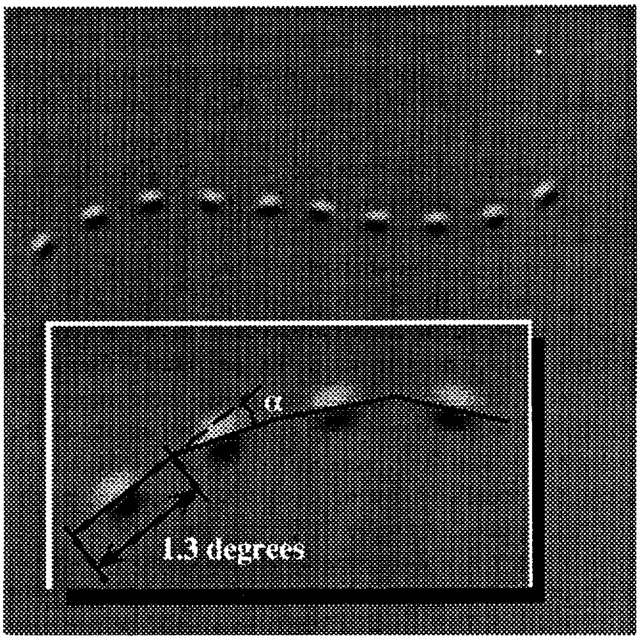

FIGURE 2. Construction of the contour. The main picture shows the contour component of a stimulus. The inset shows the relationship between the invisible line segments of the contour backbone, and the gabor elements. Each line segment averaged 1.3 deg long. The contour angle α is the angle between successive line segments.

they were placed entirely at random. The elements were also placed so that their centres were further than 0.47 deg apart, to avoid overlap. It was sometimes impossible to place a gabor element in its cell because it would be too close to elements previously placed in neighbouring cells; this produced an "empty" cell. No more than eight empty cells were permitted in a display, and the average was four. The distance between neighbouring gabor elements averaged 1.3 deg.

A "contour" stimulus consisted of two parts: the contour itself (shown separately in Fig. 2), and the background. The contour had a backbone of 10 invisible line segments, and the shape of the backbone was controlled by a single parameter α, which we call the contour angle. Each line segment was randomly selected to be between 1.2 and 1.4 deg long, and joined the next at an angle uniformly distributed from $\pm\alpha - 10$ to $\pm\alpha + 10$ deg. Gabor elements were then placed at the middle of each line segment, and their orientation θ was the same as that of the line segment (the orientation of each line segment is ambiguous, within the range 0–360 deg, but traversing the contour from one end to the other imposes a direction, and hence an unambiguous orientation, on each of the component line segments). Finally, to avoid changes in contour detection due to random closure, which can have a dramatic effect on detection (Kovacs & Julesz, 1993; Elder & Zucker, 1993), the contour was checked to ensure that it neither intersected itself, nor looped back on itself. If so, it was discarded and a new contour generated.

The entire contour was pasted into the display at a random location, making sure that the centres of the gabor elements occupied different cells. Finally, the remaining empty cells were filled with randomly oriented gabor elements, as in the no-contour stimulus (Fig. 1). The average length of each backbone line segment (1.3 deg) is the same as the average distance between neighbouring gabor elements in the background. Pilot studies showed that contour detection varied inversely with the length of the backbone line segments, but in a smooth manner, so the choice of segment length was not critical.

We measured a number of statistics of both contour and no-contour stimuli to ensure that there were no irrelevant cues to aid contour detection. Both the average distance from an element to its neighbours, and the number of empty neighbouring cells, were the same whether or not elements were part of the contour or the background. Thus the presence of the contour does not affect the local density of elements. Furthermore, the average number of empty cells was the same for both contour and no-contour stimuli, so the contour does not cause any global changes in density. If neither density nor proximity are cues, contour visibility should be due only to the alignment of the elements in the contour, since nothing else distinguishes contour elements from background. This was confirmed in a control experiment in which the orientation of contour elements was randomized; the contour could not be detected even under extended viewing, and regardless of the contour angle α. The importance of orientation in our experiments is in distinction to the kinds of grouping processes that operate in say the Ishihara plates, or explored by Kingdom *et al.* (1992) in alignment detection. In the Ishihara test, the grouping occurs on the basis of common colour, and the dots making up the test have no intrinsic orientation. In our stimuli, the colour of the elements is no help in delineating the contour, but rather it is their common orientation.

Apparatus and experimental protocol

All stimuli were displayed on a Sony Trinitron monitor attached to a Sun Sparcstation 2 computer, which constructed stimuli on-line and controlled stimulus display and response collection. The monitor was driven by 8-bit D/A converters on a 24-bit frame-buffer. The monitor was gamma-corrected in software with look-up tables. The limited range of the D/A converters means that low contrasts cannot be displayed accurately; accordingly all stimuli had more than 6% contrast. The average luminance of the gamma-corrected monitor measured with a UDT 265 photometric sensor changed when high spatial frequency waveforms (12 c/deg square wave) were displayed at over 50% contrast, indicating a spatial nonlinearity in the monitor. Most stimuli were displayed at 50% contrast or less to avoid these nonlinearities. The monitor was viewed at a distance of 60 cm in a blacked-out room. In all experiments, only the red and green guns of the monitor were modulated and the blue gun was zero. The phosphor chromaticities were $x = 0.6228$, $y = 0.3419$ for the red gun, and $x = 0.2828$, $y = 0.6045$ for the green gun. The average luminance was 15 ft L.

Each experimental run consisted of a block of 25 "contour" stimuli and 25 "no-contour" stimuli randomly interleaved. In each run, the contour angle α was kept fixed at either 0, 15, 30 or 45 deg, and the observer knew which angle was being used. Normally, 4–6 runs were performed consecutively. All stimuli were displayed for 1 sec, cued by a beep. Stimulus onset and offset were abrupt. The observer's task was to decide if the display contained a contour (contour stimulus), or consisted of just randomly oriented elements (no-contour stimulus). This task will usually be referred to as "contour detection". Their choice was communicated by pressing a button, and feedback was given. Two observers, KTM and WHM (the authors), collected a full set of data. In the third experiment we used an additional naïve observer, AW. All had normal colour vision. AW also collected a smaller set of data in the first two experiments, corroborating the results of the first two observers, but these are not shown.

Definition of contrast

Displays were composed of luminance and chromatic gabor elements, generated by modulating the intensities of the red and green guns of the display. A luminance element had in-phase red and green modulations r, g of:

$$r = r_{mean} \cdot (1+g(x,y,\theta)), \text{ and } g = g_{mean} \cdot (1+g(x,y,\theta))$$

where r_{mean}, g_{mean} are the mean gun luminances measured with the UDT 265 photometric sensor, and $g(x, y)$ is defined in Eqn (1). A colour element had counter-phase red and green gun modulations of:

$$r = r_{mean} \cdot (1+g(x,y,\theta)), \text{ and } g = g_{mean} \cdot (1-g(x,y,\theta))$$

The element contrast was defined as the contrast c of the gabor pattern $g(x, y, \theta)$, and corresponds to the Michelson contrast. Luminance and colour contrasts are not however directly comparable. The mean luminances r_{mean} and g_{mean} were selected for each observer so that the chromatic elements were isoluminant; then r_{mean} and g_{mean} are themselves of equal "sensation luminance" (Kaiser, 1988) for that observer. The requisite mean luminances were found using a motion nulling technique (Anstis & Cavanagh, 1983), as follows. When a red sinusoidal grating $r_{mean}(1 + \sin(x + vt))$ and a green sinusoidal grating $g_{mean}(1 + \sin(x - vt))$ are superimposed moving in opposite directions, one of three percepts may be seen: (i) the image appears to drift in the direction of the red grating; (ii) the image appears to drift in the direction of the green grating; (iii) the image appears stationary but flickering. When the third percept is seen, the gratings are isoluminant. To set isoluminance for these experiments, red and green gratings with the same spatial frequency as the gabor elements were superimposed, moving in apparent motion with $\frac{1}{4}$ cycle jumps. A staircase procedure was used to find the levels of r_{mean} and g_{mean} that produced no sensation of drift one way or the other, while keeping the sum $r_{mean} + g_{mean}$ constant at 15 ft L, the same as the average luminance during the contour experiments. At the isoluminance setting, the mean background level $r_{mean} + g_{mean}$ looked yellow.

RESULTS

Contour integration with luminance and chromatic elements

The purpose of the first set of experiments was to determine if contour integration is possible at isoluminance, and to quantify the influence of contrast on contour integration. An additional aim was to compare contour integration for luminance and chromatic stimuli. Detectability of a contour (number of correct responses over total number of responses) was measured as a function of the contour angle (α) and element contrast for both isoluminant and luminance stimuli. A luminance stimulus was composed entirely of luminance (yellow/

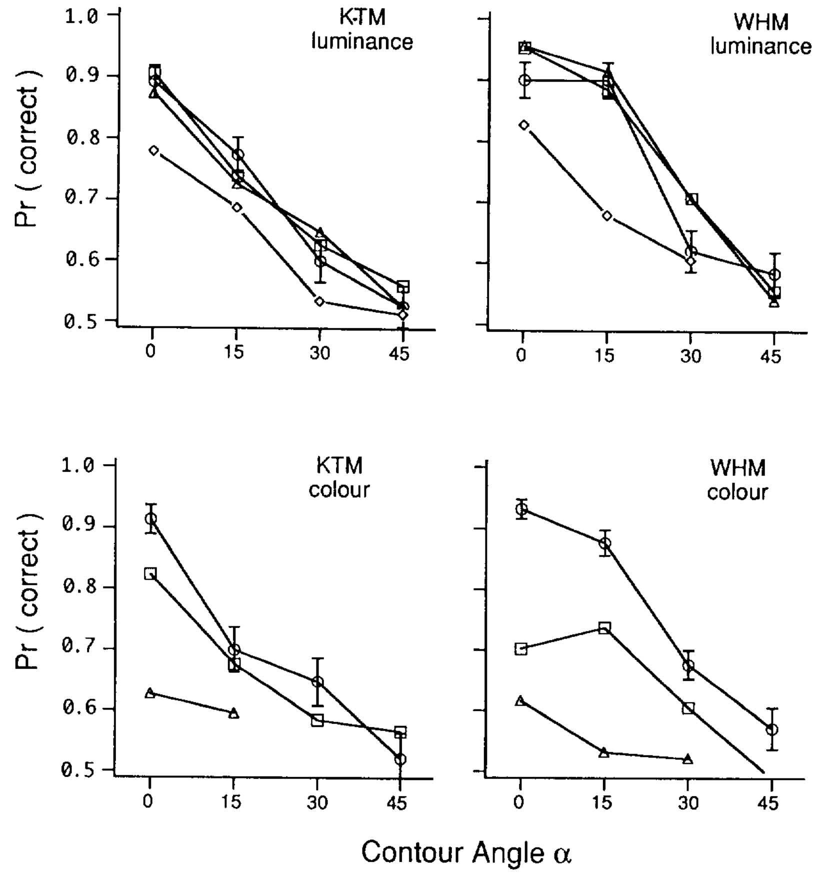

FIGURE 3. Probability of detecting luminance and isoluminant colour contours. The y-axis gives the probability of a correct yes/no response. The x-axis gives the contour angle in degs. Contrasts are indicated by the following symbols: ○ 50%; □ 25%; △ 12%; ◇ 6%. Error bars are attached to the 50% contrast data. Stimuli are shown in Fig. 1(a) and (b).

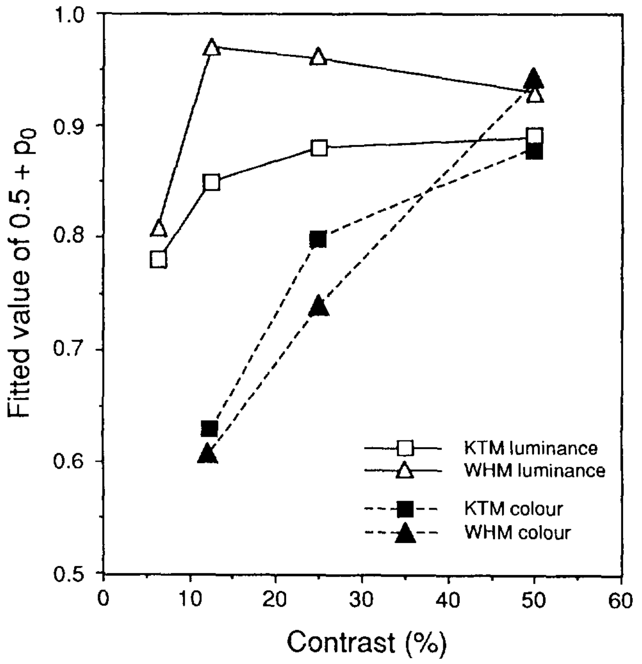

FIGURE 4. The detection data in Fig. 3 were fitted with a smooth curve $pr(correct) = 0.5 + p_0(exp(-\alpha^2/2\sigma^2))$. This graph plots the value of $0.5 + p_0$ obtained at each contrast, for colour and luminance contours.

black) elements [Fig. 1(a)], and an isoluminant stimulus was composed entirely of red/green isoluminant elements [Fig. 1(b)].

We used contour angles of 0, 15, 30 and 45 deg, and Michelson contrasts of 6%, 12%, 25% and 50%, with luminance and isoluminant stimuli. Results from these experiments are plotted in Fig. 3, which shows the proportion of correct responses as a function of contour angle for all contrasts. We have also computed d' values for these data (Green & Swets, 1966), but using d' does not change the interpretation. In all cases, detection declines with increasing contour angle, as reported by Field *et al.* (1993). The detection of luminance and isoluminant contours are remarkably similar, and at 50% contrast performances are virtually identical. To summarize the effect of contrast on detection, we fitted Gaussian curves to the data in Fig. 3, given by:

$$pr(\text{correct with contour angle } \alpha) = (0.5 + p_0 \cdot \exp(-(\alpha^2)/2\sigma^2))$$

with two parameters p_0 and σ. The main effect of contrast is to alter the intercept at $\alpha = 0$, namely $0.5 + p_0$, and σ was nearly constant at around 20 deg. The change in the intercept as a function of contrast is shown in Fig. 4. For luminance elements, performance asymptotes at 12%. For colour the asymptote probably occurs at 50% contrast, although higher contrasts could not be displayed reliably to confirm this.

To compare the detection of colour and luminance contours, colour and luminance contrasts need to be scaled to a common metric. Detection thresholds for individual gabor elements can be used as a measure of equivalent colour and luminance contrast. Unfortunately, the display system did not have enough contrast resolution to measure element detection thresholds, and furthermore it is not clear that detection of the gabor elements is the only contrast-limited task that must be performed before the contour can be detected. For these reasons, we used orientation discrimination to equate colour and luminance contrasts, since orientation is the only cue available for integration of the contours. Variability in the rotation of the gabor elements about the line segments of the contour backbone ("off-path orientation") causes a reduction in detection performance (Field *et al.*, 1993). Thus scaling by orientation sensitivity should eliminate any differences in performance between the colour and luminance system due to accuracy in encoding the orientation of the individual contour elements. Orientation discrimination was measured using a temporal two-alternative forced-choice method. The observer was shown two gabor elements one after the other. The first gabor element was randomly oriented between 85 and 95 deg (vertical = 90 deg). The second element was $\pm x$ deg from the first element. Each interval was displayed for $\frac{1}{3}$ sec, with a 1 sec inter-stimulus interval. The observer was asked to decide if the second element was rotated left or right with respect to the first. The elements were in every respect identical to those used in the contour experiments. Orientation discrimination thresholds (the value of x at which the observer is 80% correct) were measured for luminance and isoluminant elements over a range of contrasts. Orientation sensitivity (reciprocal of threshold) is plotted in Fig. 5. The thresholds are larger than those of Webster *et al.* (1990), probably because of the smaller stimuli used here. Luminance and colour thresholds follow a similar form whereby orientation sensitivity increases with

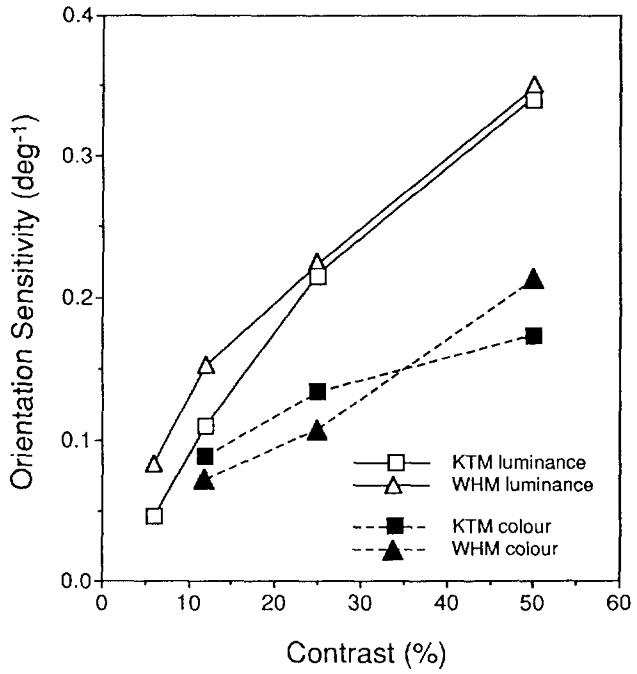

FIGURE 5. Orientation sensitivity (reciprocal of threshold) as a function of contrast for single gabor elements used in the stimuli. Open symbols are for luminance contrast, and solid symbols for colour contrast.

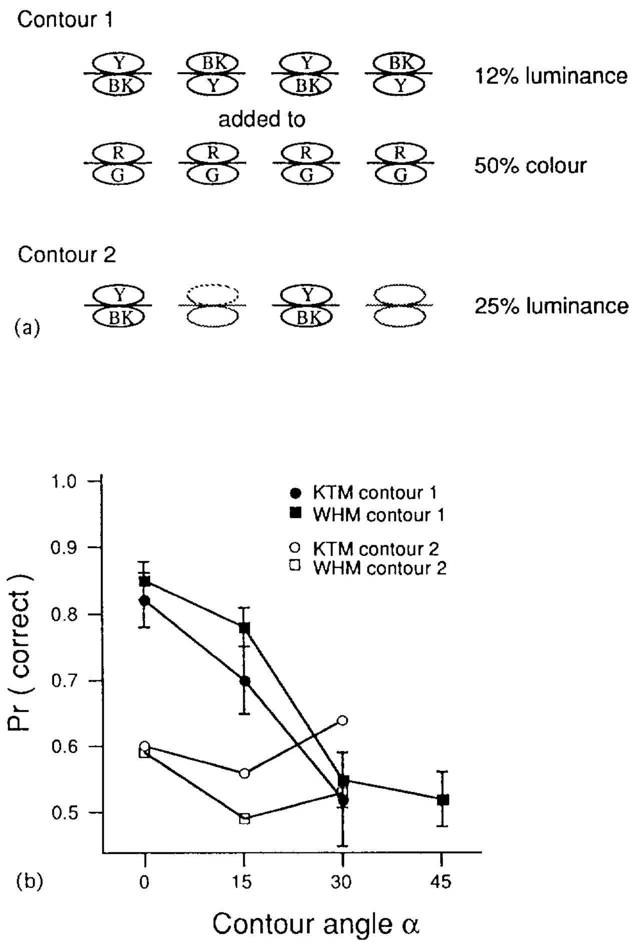

FIGURE 6. (a) Construction of the stimuli used to check isoluminance of the contour displays. Contour 1 is a 50% contrast isoluminant contour, to which has been added a 12% luminance contour, with the sign of luminance contrast alternating. Contour 2 is a 25% luminance contour, with every second element removed. If the colour contour detection is attributable to a 12% luminance artifact, the detection of both contours should be identical. (b) Detection of these contours. Contour 1 is easier to detect than contour 2.

contrast. Measured orientation thresholds range from 22 to 2.8 deg for luminance contrasts of 6–50%, and from 15 to 5 deg for the smaller colour contrast range of 12–50%. Orientation sensitivity with colour elements is equal to or greater than sensitivity with luminance elements when colour contrast is twice luminance contrast. If this comparison provides the appropriate scaling for colour and luminance contrast, a 25% isoluminant contour should be detected as easily as a 12% luminance contour. Instead, we find performance at 25% isoluminance is about the same as 6% luminance. Thus, based on a contrast scaling from orientation discrimination measures, contour detection is worse by approximately a factor of 2 for isoluminant compared to luminance stimuli in the low to middle contrast range.

Apart from the effects of contrast, however, detection performance with luminance and isoluminant stimuli seems to follow the same pattern. One explanation may be that the nominally "isoluminant" stimuli in fact contained sufficient luminance contrast for the luminance system alone to perform the task. Since performance measured with isoluminant stimuli of 50% contrast is about the same as with luminance stimuli of 12% contrast (from Figs 3 and 4), 12% is the smallest luminance contrast artifact in the isoluminant stimuli which could explain the results. To check for a luminance artifact of this size, we measured detection of a 50% contrast isoluminant contour after adding +12% luminance contrast to the first element of the contour, −12% to the second, and +12% to the third, and so on. This was achieved by superimposing 12% contrast luminance elements on top of the 50% contrast isoluminant elements [Fig. 6(a), contour 1]. The same luminance contrast (randomly ±12% contrast) was also added to the isoluminant background elements. If indeed 50% "isoluminant" contrast contains 12% luminance contrast, then adding the alternating luminance contrast should cancel the luminance contrast of every second element in the contour and background, and double the luminance contrast of the remaining elements. That is, detection of the 50% "isoluminant" ±12% luminance contour should be the same as detection of a 25% luminance contour, in which half the elements in the contour and background have been erased [Fig. 6(a), contour 2]. The results for these two cases are shown in Fig. 6(b). Detection of the five-element contour (No. 2) is close to chance levels, whereas detection of the colour contour with the added alternating luminance contrast is good. The difference between these two sets of results indicates that luminance artifacts cannot account for performance with isoluminant stimuli, and colour contrast alone is a sufficient basis for contour integration. As a final note, all colour elements with the added ±12% luminance contrast appeared clearly non-isoluminant. An approximation to the appearance of this stimulus is shown in Fig. 1(d).

Linking luminance and chromatic elements

The results of the previous section show that the integration of gabor elements into a contour is similar with both luminance and colour contrast, and indeed can reach identical performances. There are a number of possible explanations for this similarity. The first possibility is that the colour and luminance elements are encoded by a single common low-level pathway, which differs only in its sensitivity to colour and luminance contrast and in other respects does not distinguish between the two kinds of element contrast. If the low-level path fails to distinguish between colour and luminance, neither can subsequent contour integration processes, and so detection performance would be similar with both colour and luminance stimuli. This is equivalent, in effect, to a contour integration process which ignores whether the elements have colour or luminance contrast, and treats them all identically. As a second possibility, the colour and luminance pathways may be separate, and the similarity of colour and luminance contour integration arises because each pathway has its own private integration mechanism, which operates in a nearly identical manner in both pathways. In

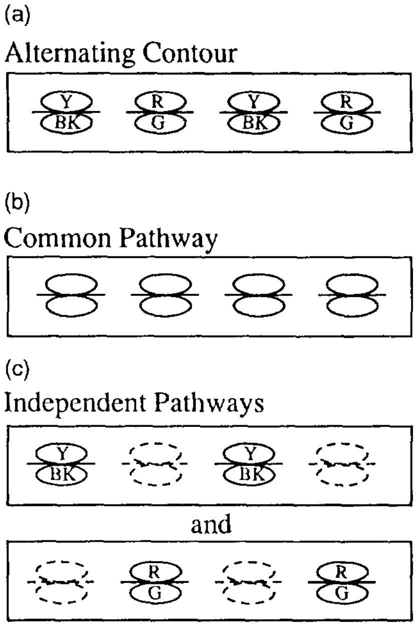

FIGURE 7. (a) A diagram of a colour–luminance alternating contour. See also Fig. 1(c). (b) How this contour would be seen by a pathway which does not distinguish between colour and luminance contrast. (c) How the contour would be seen if there are separate colour and luminance pathways. If the pathways are completely independent, the dashed elements will be invisible leaving two five-element contours, whereas if the pathways display a cross-sensitivity to the other contrast the dashed elements will be visible but with a reduced effective contrast.

this section we describe experiments to test these possibilities.

Consider a contour composed of alternating colour and luminance elements, as shown in Fig. 7(a), and illustrated in Fig. 1(c). This contour is created by giving even numbered elements in the contour a luminance contrast, and odd numbered elements an isoluminant colour contrast (numbering the elements from one end of the contour). The background elements are also randomly luminant or isoluminant. If the first hypothesis is correct, namely that colour and luminance contrasts are either indistinguishable or are ignored, the alternating contour should be as detectable as a pure luminance or pure isoluminant contour with performances similar to those in Figs 3 and 4 [illustrated in Fig. 7(b)]. We measured detection of an alternating-element contour with colour and luminance contrasts of 50%. Results (Fig. 8, ○) show that the detection of the alternating contour was worse than detection of either a pure luminance or a pure colour contour of the same contrast (dashed lines). We repeated the experiments using lower luminance contrasts of 25 and 12% (Fig. 8, □ and △ respectively). Altering the contrast of the luminance elements, however, produced virtually no change in the results. In all cases, the detection of the alternating colour–luminance contour was worse than the detection of a pure luminance-only or colour-only contour. If the first hypothesis were correct, the alternating colour–luminance contour would be indistinguishable from a pure contour, and detection should be the same. Thus the first hypothesis is not supported, and colour and luminance contrast are distinguished by the contour integration process.

We now address the second hypothesis that colour and luminance contrasts are analysed in separate parallel pathways. If the analysis is profoundly separate, the alternating 10-element contour should be seen as two separate contours of five luminance elements alone, and five colour elements alone [illustrated in Fig. 7(c)]. In this case, the detection of an alternating colour–luminance contour should be simply the probability summation of the detection of these two sub-contours. To evaluate the probability summation prediction, we used a two-alternative forced-choice method, since this makes summation easy to calculate. In each trial, the observer was shown two stimuli in temporal succession. One stimulus was a contour, the other a no-contour stimulus. Each of the two stimuli was generated independently. The observer's task was to identify which of the two intervals contained the contour stimulus. Display time in the two intervals was reduced to 0.5 sec, so that performance would not saturate. Despite this, performances are generally higher than obtained from the yes/no procedure.

The detection probability for the alternating colour–luminance contour (described above) was compared with the detection probabilities of its colour and luminance components. The luminance component was generated by making the contrast of all the colour elements in both the contour and the background equal to zero, and the colour component was generated by making all luminance contrasts zero. If the second hypothesis is correct, detection of the alternating contour should be the probability summation of the detection of the luminance and colour components. That is:

pr(correct on alternating contour)=

1-2(1-*pr*(correct on luminance component))·(1-*pr*(correct on isoluminant component)).

This equation should hold regardless of the contrasts of the luminance and chromatic elements. We used colour

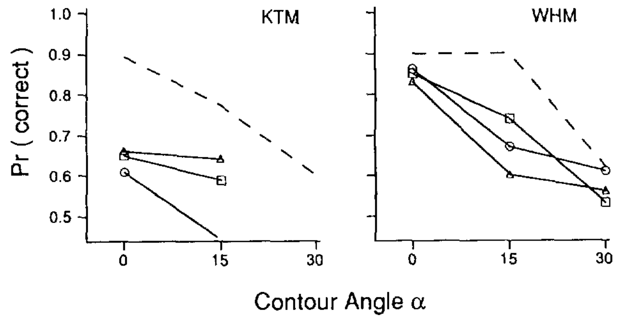

FIGURE 8. Detection probabilities for contours of elements alternated in colour and luminance contrast, as a function of contour angle. The contrast of the colour elements was fixed at 50% and the contrast of the luminance elements was 50% (○), 25% (□) or 12% (△). In terms of perceived contrast, the luminance elements of 25% contrast best matched the 50% colour elements. The dashed line shows the detection of a pure (50% contrast) luminance contour (from Fig. 3).

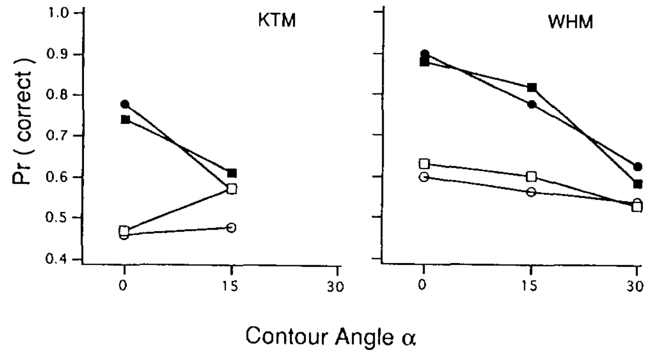

FIGURE 9. Detection probabilities measured using a two-alternative forced-choice procedure for contours of elements alternated in colour and luminance contrast, and their luminance and colour components (five-element) contours presented alone. The solid symbols give the measured two-alternative forced-choice detection probabilities for the alternated contour [Fig. 7(a)]; circles are 50% colour contrast with 50% luminance contrast, squares are 50% colour contrast with 25% luminance contrast. The open symbols give the predicted two-alternative forced-choice detection probability based on the probability summation of detection of the luminance component alone, and the colour component alone. Performances are different from those in Fig. 8 since a two-alternative forced-choice procedure was used in this case.

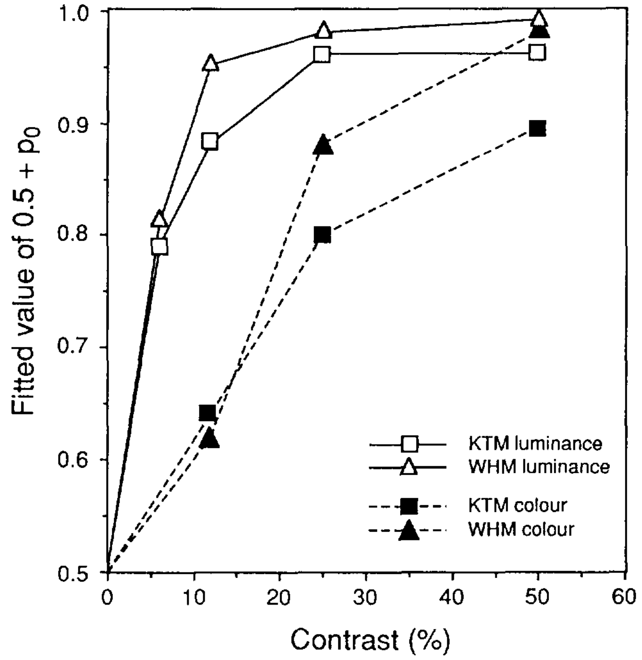

FIGURE 10. Results are for the detection of contrast alternating luminance-only contour (open symbols) or colour-only contour (solid symbols). Contours consisted of a gabor element of 50% contrast alternated with a gabor element of a different contrast magnitude. For the luminance contour, 50% contrast was alternated with elements of 50%, 25%, 12% or 6%, and for the colour contour 50% contrast was alternated with elements of 50%, 25% or 12%. The value of the alternating contrast is plotted on the x-axis. Full psychometric functions were collected for each contrast alternating contour (detection probability vs contour angle) as in Fig. 3. The functions were fitted using the equation given in Fig. 4. The graph plots the value of $0.5 + p_0$ (performance for a contour angle of 0) obtained for each contrast alternating condition. Results for KTM and WHM.

contrast of 50% alternated with luminance contrasts of 50 or 25% in separate experiments. The results are shown in Fig. 9 which plots the two-alternative forced-choice detection probabilities of the alternated contour (solid symbols), together with the probability summation prediction (open symbols), against the contour.angle. Clearly, the detection of the alternating path is better than the probability summation of its colour and luminance components. The probability summation predictions are low, reflecting the fact that the detection of a five-element path is close to chance levels. [For results obtained for a five-element contour with a yes/no procedure, see Fig. 6 (open symbols).] Thus idea that colour and luminance contours are analysed entirely independently cannot be supported.

Another version of the second hypothesis is that the colour and luminance contour integration processes remain largely separate, but exhibit some degree of cross-sensitivity between the contrast types, perhaps due to cross-sensitivity of low-level detectors [as might be found with some detection/discrimination experiments (DeValois & Switkes, 1983; Cole *et al.*, 1990). Thus when presented with the colour–luminance alternating contours, each selective process would "see" elements that alternate between higher and lower contrasts [illustrated in Fig. 7(c)]. The effect of this may be to reduce contour detection significantly. To test whether this is a likely explanation we measured detection performance for luminance-only contours and colour-only contours composed of elements which alternated in their contrast magnitude. Luminance contrast elements of 50% were alternated with elements of 50%, 25%, 12% and 6% in separate experiments. Colour contrast elements of 50% were alternated with elements of 50%, 25% and 12%. All methods and the data fitting procedure were the same as those used in the first experiments (Figs 3 and 4). Based on the fitted results, performances were obtained at a contour angle of 0 deg ($0.5 + p0$) and are plotted as a function of the contrast of the alternating elements in Fig. 10. Performance with this type of contour is contrast dependent, resembling the contrast

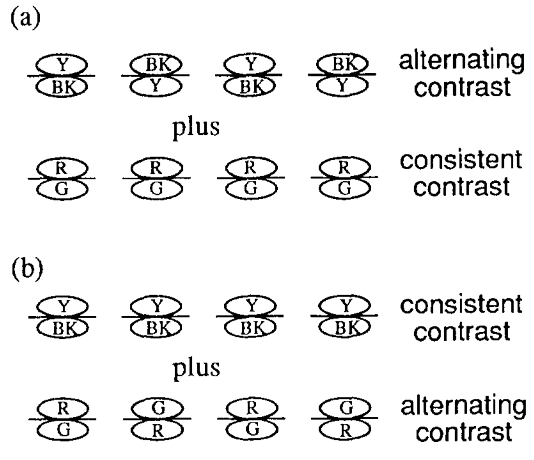

FIGURE 11. Contours composed of combinations of colour and luminance contrast. (a) A contour with consistent colour contrast but alternating luminance contrast, constructed by superimposing a luminance and a colour contour. See also Fig. 1(d). (b) A contour with consistent luminance contrast and alternating colour contrast.

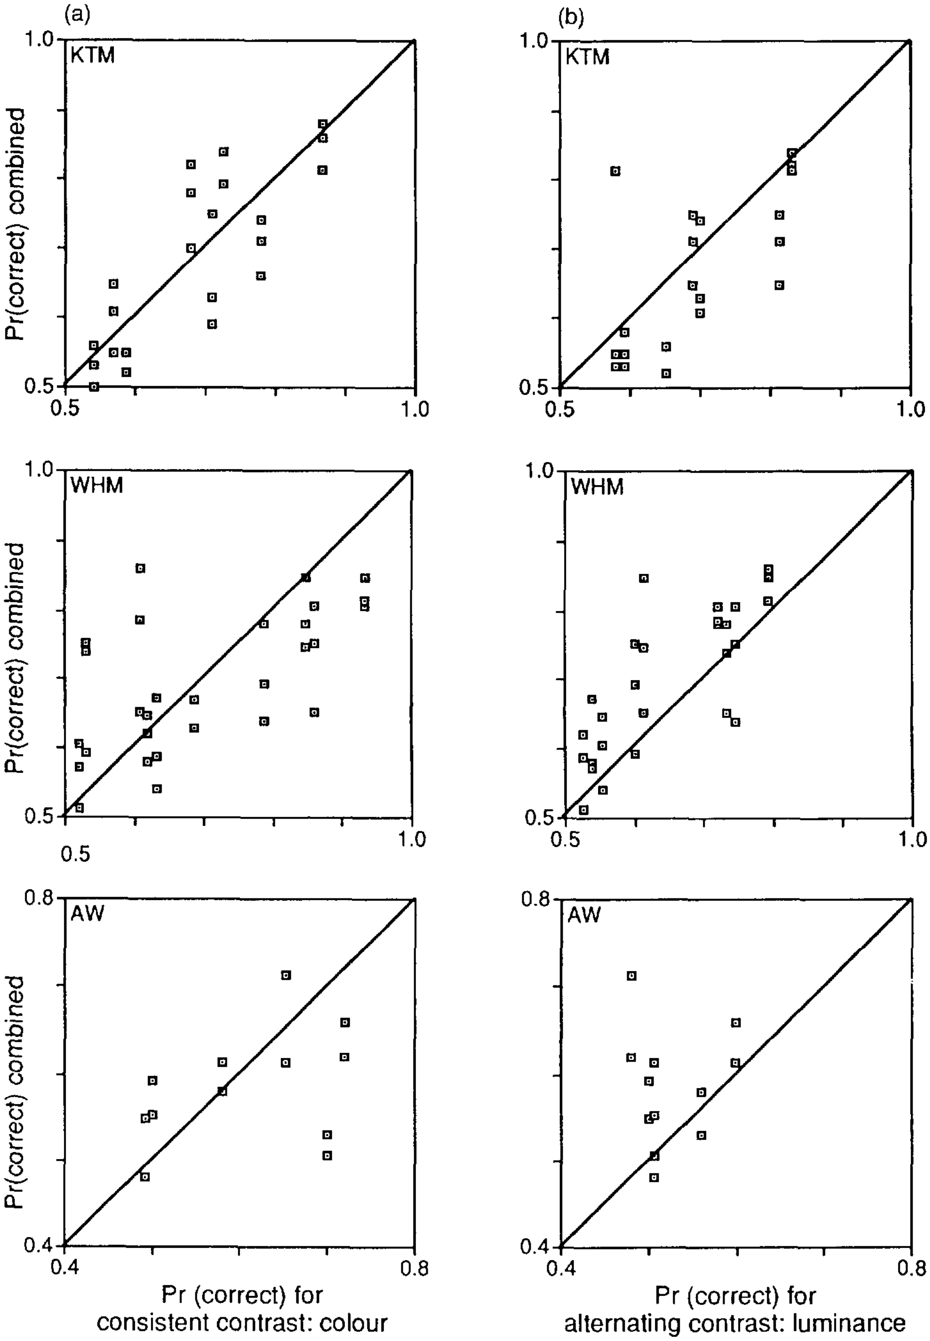

FIGURE 12. Results of detection experiments using contours like those shown in Fig. 11(a) and Fig. 1(d). All contours had fixed colour contrast but varying luminance contrast. (a) Plot of the detection probability of the combined contour (y-axis) against the detection probability of a contour of the same angle α and with the same colour contrast, but no luminance contrast. (b) Plot of the detection probability of the combined contour against the detection probability for a contour of the same angle and alternating luminance contrast, with no colour contrast. Note that the scale on the lower two graphs is different.

dependence for contours composed of one invariant contrast (Fig. 4). Detection probabilities on the alternating colour–luminance contour were between 0.61 and 0.66 for KTM and 0.83 and 0.87 for WHM (see Fig. 8). It can be seen that these would lie on part of the colour and luminance functions in Fig. 10 that are strongly contrast dependent. Thus the cross-sensitivity model predicts that performance should be affected by the contrast magnitudes of the alternating elements, as well as their contrast type. Yet, as Fig. 8 shows, the relative contrasts of the colour and luminance elements in the colour–luminance alternating contours are unimportant for performance. This difference in the contrast dependence of the two types of task (colour–luminance alternations vs high–low contrast alternations) is evidence against a cross-sensitivity model of colour and luminance contour integration. Possible explanations for these results are considered in the Discussion.

Combined luminance and chromatic contrast

As raised in the Introduction, colour can be a useful cue for object segregation, since it is largely invariant with changes in lighting intensity. Along the boundary of an object, the luminance contrast can vary considerably

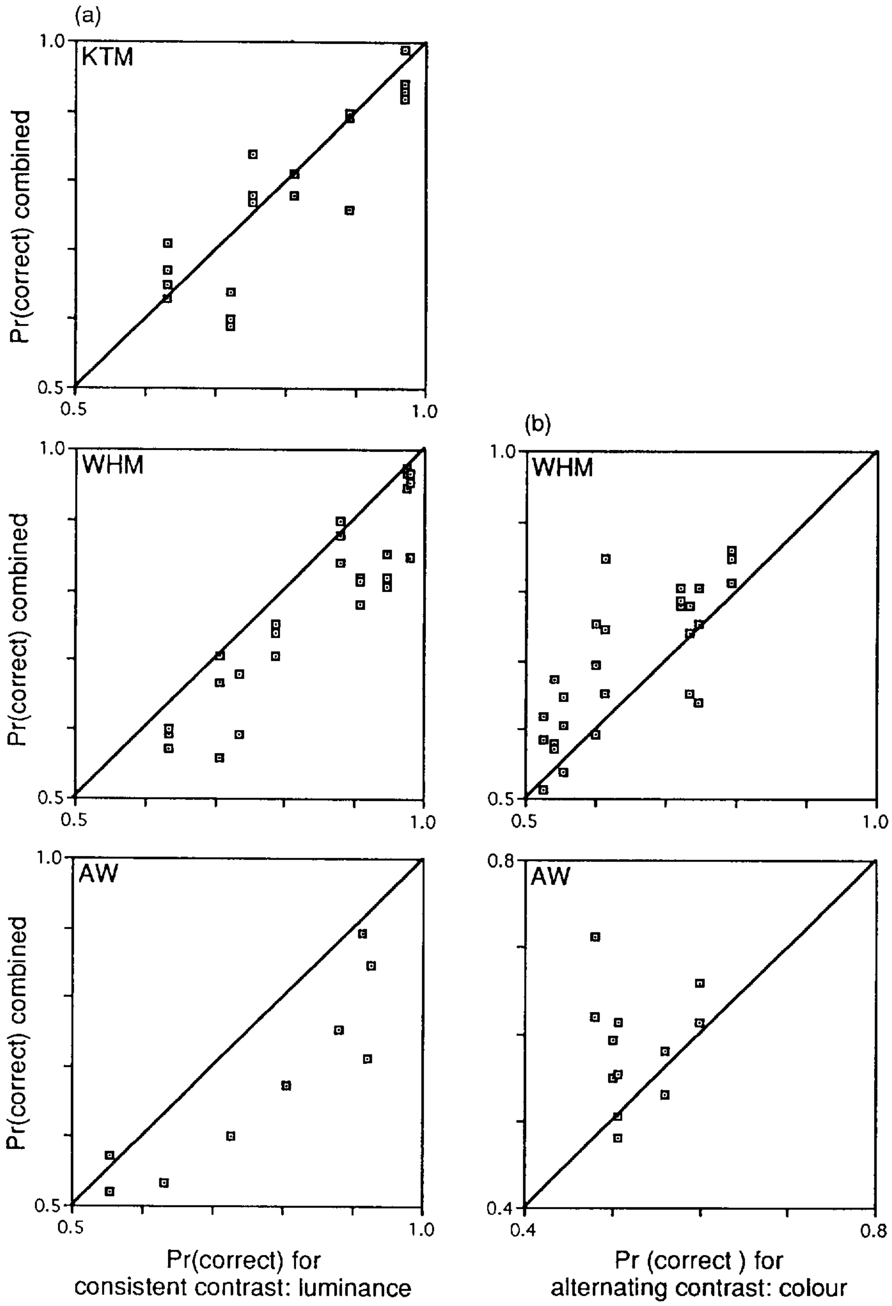

FIGURE 13. Results of detection experiments using contours like those shown in Fig. 11(b). All contours had fixed luminance contrast but varying colour contrast. (a) Plot of the detection probability of the combined contour (y-axis) against the detection probability of a contour of the same angle α and luminance contrast, but with no colour contrast. (b) Plot of the detection probability of the combined contour against the detection probability for a contour of the same angle and alternating colour contrast, but no luminance contrast. Note that the scale on the lower graph of (b) is different.

depending on the illumination incident on the object itself, and on the luminance of the image adjacent to the object boundary. The colour of the object, however, is more likely to be consistent along the boundary. These ideas imply that (i) the integration of colour contours should be relatively unaffected by variations in luminance contrast along the contour, and (ii) adding consistent colour contrast should improve the integration of a contour which has varying luminance contrast, by increasing the likelihood that the contour is seen as a cohesive border. Conversely, one would expect that (iii) integration of a luminance contour would be adversely affected by variations in colour contrast, since this would tend to indicate a change of material (Rubin & Richards, 1982), and (iv) adding consistent luminance contrast should not improve the integration of a contour in which the colour contrast alternates.

The final series of experiments was designed to test these predictions. We compare the detection of contours composed of elements of both luminance and colour contrast, to the detection of contours having the same angle α but with colour-only or luminance-only contrast. The elements in the combined colour–luminance stimulus were each the sum of a colour element and a luminance element. In the first of these experiments, we looked at combined colour–luminance contours which had a consistent colour contrast, but varying luminance contrast. Within the contour, all elements had the same colour contrast and contrast polarity, so that one "side" of the contour was uniformly red, and the other "side" uniformly green. The luminance contrast, however, was the same sign for even-numbered elements, and opposite sign for odd-numbered elements [Fig. 11(a) and Fig. 1(d)], so that it alternated along the contour. This produced a contour alternating in two kinds of elements: one which has a light red to dark green contrast (same sign), and the other which has a dark red to light green contrast (opposite sign). For background elements, the luminance contrast was randomly selected to be the same sign or the opposite sign as the colour contrast. We performed yes/no detection experiments using colour contrasts of 50%, 25% and 12%, with luminance contrasts of 25%, 12% and 6% in all combinations. We used contour angles $\alpha = 0$, 15 and 30 deg. Each combination of contrasts and angle was evaluated in a separate experiment. We also performed experiments with a luminance contrast of 0% (an isoluminant colour contour), and a colour contrast of 0% (a pure luminance contour with its contrast alternating in sign along the contour).

The results of these experiments are shown in Fig. 12. In Fig. 12(a) we have plotted the detection probability for a combined colour–luminance contour against the detection probability for the colour-only contour of the same colour contrast and contour angle. If prediction (i) is correct, the results should cluster along the diagonal line of unity in the graphs, indicating that the varying luminance has little effect on contour detection. For AW and KTM, the hypothesis is broadly supported, but there is considerable scatter. It is rejected at the 10% level for WHM (the method for this and all succeeding statistical tests is given in the Appendix). Instead, varying luminance contrast generally worsens detection of the colour contour. In Fig. 12(b) we have plotted the detection probability for a combined colour–luminance contour against the detection probability for an alternating sign luminance-only contour of the same luminance contrast and contour angle. If prediction (ii) is correct, the points should fall above the diagonal line, indicating that the consistent colour improves the detection of a contour of varying luminance. This is so for WHM and AW; there are significantly more data above the diagonal at the 5% level. The hypothesis is rejected for KTM. The ecological hypotheses are not consistently supported, and there is considerable individual variation.

In the second of these experiments, we reversed the roles of colour and luminance. This time, the contours all had a consistent luminance contrast, but the colour contrast alternated in sign along the contour. Within the contour all elements had the same luminance contrast, so that one side of the contour was uniformly light, and the other side uniformly dark. This time, the colour contrast was the same sign for even-numbered elements, and opposite sign for odd-numbered elements [Fig. 11(b)], so that it alternated along the contour. We again used luminance contrasts of 25%, 12% and 6%, and colour contrasts of 50%, 25% and 12%. Background elements were the same as the previous case. The results of these experiments are shown in Fig. 13. In Fig. 13(a) we have plotted the detection probability for these combined colour–luminance contours against the detection probability for a luminance-only contour of the same contrast and contour angle. If prediction (iii) is correct, the results should cluster below the diagonal line in the graphs, indicating that the addition of varying colour contrast worsens contour detection. For WHM and AW, there are significantly more data below the diagonal (at the 5% level), but not for KTM. In the right column we have plotted the detection probability for the combined colour–luminance contour against the detection probability for an alternating sign colour-only contour of the same colour contrast and contour angle. If prediction (iv) is correct, the points should fall around the diagonal line, indicating that the consistent luminance has little effect on the detection of a contour of varying colour. This is not so for WHM and AW. Instead, significantly many data show improvement (at the 1% level). Data were not collected for KTM, as detection of the alternated colour contour alone was poor. Again, these experiments have failed to consistently support the ecological hypotheses.

Although not explicitly represented, the results for luminance-only and colour-only contours composed of elements alternating in their contrast sign are included in Fig. 12(b) and Fig. 13(b) respectively. Alternating the sign of the contrast of the odd symmetric gabors is detrimental to contour detection for both colour and luminance contrast. This is represented by the clustering of the data points in these figures to the lower probability regions of the abscissa. Good performances (detection probabilities above 0.80) cannot be achieved with this type of contour, demonstrating that contour integration is sensitive to the sign of the element contrast.

DISCUSSION

We have examined the respective roles of colour and luminance contrast in the integration of contours, and their interactions in this task. Colour contrast alone is effective in delineating a contour, provided the contrast is sufficiently high. We chose to scale colour and luminance contrasts according to orientation discrimination thresholds, since orientation is the only cue to the detection of the contour. Scaling by orientation sensitivity should account for any differences between colour and luminance contour processes due to inaccuracy in encoding the orientation of the individual contour elements. With this scaling, colour is not as effective as luminance contrast at contour integration, except at high contrasts where similar performances are reached. In terms of Fig. 4, there is about a four-fold difference between colour and luminance contour detection in units of screen contrast. About a two-fold difference would remain after scaling contrast for the differences in orientation sensitivity between colour and luminance vision (Fig. 5). We conclude that colour is able to support a complex form detection task, although it requires a somewhat (two-fold) greater contrast gain to match the performances based on luminance contrast.

Contours can be detected even when the elements switch from luminance to isoluminant, but detection is degraded relative to luminance-only or colour-only contour detection. We have shown that detection is not preserved to the extent that would be expected if it was based on a single common contour integrating process which failed to distinguish between colour and luminance contrast. Neither, however, is detection reduced as far as would be expected if there was an entirely independent encoding of the colour and luminance contours, with probability summation determining performance. We have also considered a third possibility in which the colour and luminance contour integration processes are largely separate, but each has a low gain for the contrast of the other type, exhibiting a cross-sensitivity between the contrast types. However, performance on the alternating colour–luminance contours does not display the contrast dependence expected from this type of model. Overall, these results suggest that models of contour integration must include specific processes that can distinguish colour and luminance contrast.

The reduction in detectability with contours made from alternated colour and luminance elements could conceivably be due to a competing organization in the stimulus. Colour appears to play a role in camouflaging texture boundaries (Morgan *et al.*, 1992) by segmenting the image into larger more global regions. In out task there is a tendency to group all isoluminant elements together, and to group all luminance elements together. The contour may be masked by this grouping, thus degrading performance. However, similar reductions in detectability are obtained by alternating the phase of the elements in a purely luminance or purely chromatic contour, but in this case there is no competing grouping based on element phase. The apparent strength of the colour–luminance grouping also varies with the relative salience of the elements, but the contour detectability does not. Finally, when we varied the contrast of a luminance or of a colour contour (Fig. 10), there was no reduction in performance, although there was a clear segregation of the stimulus into high-contrast and low-contrast elements, similar to that found by Morgan *et al.* (1992). It is thus unlikely that competing groupings can offer an explanation for the reduction in detectability of a contour made from alternated colour and luminance elements.

The results of these experiments are compatible with a simple two-stage model of contour detection. In the first stage, the colour and luminance elements are encoded by independent low-level processes which detect the position and orientation of the elements. In the second stage, these elements whether colour, luminance or both, are integrated to form a contour. The ability of the second stage to integrate elements depends on a number of factors, including proximity and alignment, but also the comparative contrasts of the elements. Integration is most successful when the elements have the same contrast type, and is much less successful when the contrasts differ. Differences can either be in the sign of the contrast, as in the contour with elements alternating between positive and negative luminance contrasts of one type (Expt 3), or in the type of contrast, as in the contour with elements alternating between colour and luminance contrast (Expt 2). In the terminology of Field *et al.* (1992), the "association field" around each element is sensitive to the contrasts of the elements. This association field could be implemented by neurons similar to those found by Peterhans and von der Heydt (1991). It would be interesting to know whether these neurons are sensitive to colour, or whether there are colour-sensitive variants of this class of neurons.

Grossberg and Mingolla (1985) have devised a neural network which seems capable of detecting contours like the ones used in this study. In their network the responses of local orientation-sensitive neurons inputs to the receptive fields of second-stage "bipole cells". The contour emerges as a result of competition at the bipole level, together with a feedback loop between bipole cells and the lower level orientation detectors. Grossberg and Mingolla's boundary-completion network discards information about the contrast (via a full-wave rectifying "complex cell") before oriented cell responses are fed to the bipole cells. Another model has been advanced by Heitger and von der Heydt (1993), but this differs most from Grossberg and Mingolla's model in the algorithm, and the output of both models is very similar; so too is the use of "complex cells" which discard the contrast polarity across an edge. Clearly, from the experiments described here, contrast polarity is not discarded, and colour and luminance contrasts also remain distinguished in contour integration. Grossberg and Mingolla's network could be modified, however, to include two types of bipole cells: contrast selective as well as contrast unselective. (A similar modification can be proposed for Heitger and von der Heydt's network.) Contours which are built from elements of the same type of contrast, whether colour or luminance, will activate both selective and unselective bipoles, whereas contours built from inhomogeneous contrast elements will only activate the unselective bipoles. This may lead to a reduction in the detectability of the colour–luminance alternating contours, as observed here (Expts 2 and 3). This proposal is reminiscent of the specific and unspecific pathways proposed by Gorea *et al.* (1993) to explain the contributions of colour and luminance to motion mechanisms; neurons sensitive to both colour and luminance (Lennie *et al.*, 1990) could form the substrate for the unspecific mechanism.

The results of the last set of experiments, testing the effects of combinations of colour and luminance contrast, are surprising despite their considerable individual variability. For observers WHM and AW, adding a sign-alternating contrast of one type to a contour of a consistent contrast of the other type generally reduced performance, regardless of whether a varying luminance contrast was added to a colour contour of consistent contrast [Fig. 12(a)], or the other way around [Fig. 13(a)]. In keeping with this, these two subjects also showed an improvement in contour detection when a consistent contrast of one type was added to a sign-alternating contrast of the other type, again regardless of the contrast combination used [Fig. 12(b) and Fig. 13(b)]. Observer KTM, on the other hand, displays results which suggest that the detection of colour and luminance contrasts are virtually independent; for her, contrast variations failed to "camouflage" contours of consistent contrast regardless of the colour/luminance combination used, and surprisingly, consistent colour contrast failed to improve the detection of a sign alternating luminance contour [Fig. 12(a)]. Regardless of the individual variability in the detection of combined colour plus luminance contours, the pattern of results suggests that there is no selective power of colour contrast either to disrupt or improve contour detection in these experiments. Thus the results do not square with the ecological hypotheses that varying luminance should have little effect on a colour contour, but not *vice versa*, or with the related hypotheses that adding consistent colour contrast should improve integration of varying luminance contours, but not *vice versa*.

The ecological hypothesis however cannot yet be discarded. The luminance of an object is a product of the luminance of the incident light and the albedo of the object, and separating them perceptually is the result of a number of complex and little understood processes (Adelson, 1993). We designed the experiment under the assumption that the luminance changes would be interpreted as changes in incident light from shadows or other variations in the illuminant, but some observers may have attributed them to albedo changes. If so, a luminance change would indicate a change in surface reflectance as strongly as a colour change, so their detection of varying colour contours and varying luminance contours should be similar (as was the case for WHM and AW). A clear test of the ecological hypothesis would need a stimulus where the observer could only attribute luminance changes to the illuminant; such a stimulus would have to include junctions (T-shape or X-shape) associated with illuminant changes in the real world, and which have been left out of our stimulus.

REFERENCES

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

Anstis, S. & Cavanagh, P. (1983). A minimum motion technique for judging equiluminance. In Mollon, J. D. & Sharpe, L. T. (Eds), *Color vision: Physiology and psychophysics* (pp. 156–166). London: Academic Press.

Barlow, H. B. (1981). Critical limiting factors in the design of the eye and visual cortex. *Proceedings of the Royal Society of London B*, *212*, 1–35.

Beck, J., Rosenfeld, A. & Ivry, R. (1989). Line segregation. *Spatial Vision*, *4*, 75–101.

Brill, M. H. (1990). Image segmentation by object color: A unifying framework and connection to color constancy. *Journal of the Optical Society of America A*, *7*, 2041–2047.

Cole, G. R., Stromeyer, C. F. & Kronauer, R. E. (1990). Visual interactions with luminance and chromatic stimuli. *Journal of the Optical Society of America A*, *7*, 128–140.

DeValois, K. K. & Switkes, E. (1983). Simultaneous masking interactions between chromatic and luminance gratings. *Journal of the Optical Society of America*, *73*, 11–18.

Elder, J. & Zucker, S. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, *33*, 981–991.

Field, D., Hayes, A. & Hess, R. (1993). Contour integration by the human visual system: Evidence for a local "Association field". *Vision Research*, *32*, 173–193.

Gorea, A., Papathomas, T. V. & Kovacs, I. (1993). Two motion systems with common and separate pathways for color and luminance. *Proceedings of the National Academy of Science U.S.A.*, *90*, 11,197–11, 201.

Green, D. M. & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Grossberg, S. & Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception & Psychophysics*, *38*, 141–171.

Heitger, F. & von der Heydt, R. (1993). A computational model of neural contour processing: Figure-ground segregation and illusory contours. *Proceedings of the 4th International Conference on Computer Vision*, Berlin 1993. IEEE Computer Society Press, Los Almitos, Calif.

Kaiser, P. K. (1988). Sensation luminance: A new name to distinguish CIE luminance from luminance dependent on an individual's spectral sensitivity. *Vision Research*, *29*, 455–456.

Kanisza, G. & Gerbino, W. (1982). Amodal completion: Seeing or thinking? In Beck, J. (Ed.), *Organisation and representation in perception* (pp. 167–190). London: Lawrence Erlbaum.

Kingdom, F. A. A., Moulden, B. & Collyer, S. (1992). A comparison between colour and luminance contrast in a spatial linking task. *Vision Research*, *32*, 709–717.

Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt Brace & World.

Kovacs, I. & Julesz, B. (1993). A closed curve is much more than an incomplete one: Effect of closure in completion of segmented contours. *Investigative Ophthalmology & Visual Science (Suppl.)*, *33*, 1084.

Lennie, P., Krauskopf, J. & Sclar G. (1990). Chromatic mechanisms in striate cortex of macaque. *Journal of Neuroscience*, *10*, 649–669.

Livingstone, M. & Hubel, D. (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience*, *7*, 3416–3468.

Losada, M. A. & Mullen, K. T. (1994). The spatial tuning of chromatic mechanisms identified by simultaneous masking. *Vision Research*, *34*, 331–341.

Losada, M. A. & Mullen K. T. (1995). Colour and luminance spatial tuning estimated in the absence of off-frequency looking. *Journal of the Optical Society of America A*, *12*, 250–260.

Marr, D. (1982). *Vision*. San Francisco, Calif.: W.H. Freeman & Co.

McIlhagga, W. H., Mullen, K. T. & Field, D. J. (1993). Roles of color and luminance contour integration. *Investigative Ophthalmology & Visual Science (Suppl.)*, *34*, 766.

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.

Morgan, M. J., Adam, A. & Mollon, J. D. (1992). Dichromats detect colour-camouflaged objects that are not detected by trichromats. *Proceedings of the Royal Society of London B*, *248*, 291–295.

Mullen, K. T. & Kingdom, F. A. A. (1991). Colour contrast in form perception. In Gouras, P. (Ed.), *Vision and visual dysfunction 6: The perception of colour* (pp. 198–217). London: Macmillan.

Parent, P. & Zucker, S. W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *11*, 823–839.

Parkkinen, J. P. S., Hallikainen, J. & Jaaskelainen, T. (1989). Characteristic spectra of Munsell colors. *Journal of the Optical Society of America A*, *6*, 318–322.

Peterhans, E. & von der Heydt, R. (1989). Mechanisms of contour perception in monkey visual cortex II: Contours bridging gaps. *Journal of Neuroscience*, *9*, 1749–1763.

Rock, I. (1983). *The logic of perception*. Cambridge, Mass.: MIT Press.

Rubin, J. M. & Richards, W. A. (1982). Color vision and image intensities: When are changes material? *Biological Cybernetics*, *45*, 215–226.

Smits, J. T. S., Vos, P. G. & van Oeffelen, M. P. (1985). The perception of a dotted line in noise. *Spatial Vision*, *1*, 163–177.

Switkes, E., Bradley, A. & DeValois, K. K. (1983). Contrast dependency and mechanisms of masking interactions among chromatic and luminance mechanisms. *Journal of the Optical Society of America A*, *5*, 1149–1162.

Uttal, W. R. (1975). *An autocorrelation theory of form detection*. Hillsdale, N.J.: Lawrence Earlbaum.

Webster, M. A., DeValois, K. K. & Switkes, E. (1990). Orientation and spatial frequency discrimination for luminance and chromatic gratings. *Journal of the Optical Society of America A*, *7*, 1034–1049.

APPENDIX

This Appendix describes the nonparametric statistical test used in analysing the results of Expt 3. Suppose we have a set of data (x_i, y_i), $i = 1...n$. Under the null hypothesis H_0: $E(x_i) = E(y_i)$ [where $E(\ )$ is the expected value] we would expect $\frac{pr(x_i<y_i)=pr(x_i>y_i)=1}{2}$. Let r be the number of data points where $x_i < y_i$. Then r follows a binomial distribution with parameters n and $p = 0.5$. We will accept the alternative hypothesis H_1: $E(x_i) < E(y_i)$ if the value of r is significantly different from $\frac{n}{2}$. From the binomial distribution, the probability of observing r or more data with $x_i < y_i$ is

$$\sum_{i=r}^{n} \binom{n}{i} 0.5^n$$

If this probability is less than the chosen significance level, we reject H_0 in favour of H_1. For example, with 10 data points ($n = 10$) we reject H_0 at the 10% significance level if we observe more than seven data points with $x_i < y_i$. Note that this test, while simple, lacks power. It can accept H_0 even when there is considerable scatter of data.

Acknowledgements—We are grateful to D. J. Field for helpful discussion. This work was supported by a grant from the Medical Research Council of Canada (No. MT-10819), and Human Frontier Science Program Fellowship LT-285/93.