Motion edges and regions guide image segmentation by colour

PER MØLLER AND ANYA C. HURLBERT

*Department of Physiological Sciences, Medical School, University of Newcastle upon Tyne, Newcastle upon Tyne NE2 4HH, UK
(per.moeller@ncl.ac.uk, anya.hurlbert@ncl.ac.uk)*

SUMMARY

Image segmentation is an important early stage in visual processing, in which the visual system groups together parts of the image that belong together, prior to or in conjunction with object recognition. Two principal processes may be involved in image segmentation: an edge-based process that uses feature contrasts to mark boundaries of coherent regions, and a region-based process that groups similar features over a larger scale. Earlier, we have shown that motion and colour interact strongly in image segmentation by the human visual system. Here we explore the nature of this interaction in terms of edge- and region-based processes. We measure performance on a region-based colour segmentation task in the presence of distinct types of motion information, in the form of edges and regions which in themselves do not reveal the location of the colour target. The results show that both motion edges and regions may guide the integrative process required for this colour segmentation task. Motion edges appear to act by delimiting areas over which to integrate colour information, whereas motion similarities define primitive surfaces within which colour grouping and segmentation processes are deployed.

1. INTRODUCTION

Image segmentation is an important early stage in visual processing, in which the visual system groups together parts of the image that are likely to belong to the same object. In the natural world multiple cues to segmentation are usually present simultaneously—e.g. colour, motion, texture, and depth—any two of which are generally independent properties of objects. It is therefore important to determine how distinct cues interact in segmenting an image. Earlier, we showed in psychophysical studies that motion and colour interact strongly in image segmentation (Møller & Hayhoe 1993; Møller & Hurlbert 1997). In the latter study, we showed that different types of motion noise, spread uniformly across a stimulus in which colour cues only signalled the presence of a target figure, could either facilitate or inhibit detection of the colour target. The form of these interactions, together with the presumed ‘early’ pre-attentive nature of segmentation, support the hypothesis that the interactions occur at an early level before a global segmentation of the entire image is decided.

Here we explore in more detail the nature of the interactions between colour and motion segmentation processes, specifically in terms of edge-based and region-based contributions. Although in one sense the latter two cannot fully be disentangled, since processes that find regions must also mark their boundaries, and closed edges necessarily define regions, there are clear distinctions between them in terms of theory and

implementation. Edge-based segmentation processes rely on detection of local feature contrasts to segregate surfaces, while region-based processes exploit the similarity of features within surfaces to delimit them (e.g. Ballard & Brown 1982). Traditionally, in the study of both animal and machine vision the emphasis has been on edge-based processes, from edge-detection algorithms (e.g. Canny 1986) to models of centre-surround neuronal receptive fields as Laplacian operators (e.g. Marr 1982). Psychophysical studies of scene segmentation and feature pop-out have preferentially highlighted the role played by edge-based processes (e.g. Nothdurft 1994; Sagi 1991), although evidence has recently emerged for region-based segmentation processes within individual cues, e.g. colour, motion (Møller & Hurlbert 1996), luminance (Mumford *et al.* 1987), and texture (Caputo 1996).

Recent successful segmentation algorithms for machine vision and models of biological vision combine edge-based and region-based processes, typically within the same cue (e.g. Chen & Lin 1997; Haddon & Boyce 1990; Zhu & Yuille 1996; Grossberg & Mingolla 1985). Other algorithms use information from multiple attributes to localize edges (e.g. Poggio *et al.* 1988; Baraldi & Parmiggiani 1996), or combine different types of segmentation process across modalities (e.g. Hurlbert & Poggio 1989). There are several advantages of combining information from different attributes and processes: for example, (i) noise and ambiguity in one cue may be alleviated or resolved by information from another; (ii) a sharp discontinuity in one cue may serve to delimit a

region-based process in another cue; and (iii) the entire segmentation process may be speeded up by interactions between cues. The extent to which the human visual system combines information across modalities and processes to enhance image segmentation is still unknown.

Having previously established that there are both region-based and edge-based processes in segmentation by motion (Møller & Hurlbert 1996), and that motion information does influence segmentation by colour (Møller & Hurlbert 1997), we now ask explicitly whether this influence is via an edge- or region-based motion process, or both. In the following experiments, we use a colour segmentation task that is chiefly region-based (Møller & Hurlbert 1996). In this task, detection and localization of the segmentation target relies solely on colour information which must be integrated over space. To the colour stimulus, we add motion information which in itself does not reveal the location of the colour target. We analyse the contribution of distinct motion processes to colour segmentation by comparing conditions in which (i) both motion edges and regions, (ii) edges only, or (iii) neither, are aligned with the colour target. One hypothesis we examine is that motion edges might guide the formation of colour regions by providing boundaries for colour integration.

2. METHODS

The stimuli were precomputed for the start of each session with a Silicon Graphics IRIS computer which also controlled the sequence of events during data collection. Images were displayed on a linearized EIZO flexscan T560i-T RGB monitor at a vertical refresh rate of 60 Hz. The three observers, supported by chin-rests, viewed the stimuli binocularly in a darkened room from a distance of 1.14 m. The stimulus was a 6.5° square grey field of luminance 1.0 cd m^{-2} (CIE chromaticity coordinates $x=y=0.33$) displayed against a black surround. The square grey field was filled with 2048 randomly positioned square dots, each with side length 3.9 min of arc and luminance 6.0 cd m^{-2} (see figure 1a). The target was a 0.78° wide vertical band of dots appearing to the left or right of the midline, positioned with its most central edge at 1.43° from the central fixation target. Dot motion was generated by assigning each dot a particular velocity which determined its position in each precomputed frame. At the speed used of 2.4° s^{-1} , each dot moved 3 pixels (2.4°) between frames.

(a) Colour parameters

Target dots were distinguished from background dots by a different distribution of colours, and target detection thresholds were measured by varying the difference between the two distributions. In all the colour segmentation measurements reported here, background dots were assigned colours from a broad rectangular distribution in colour space with a mean CIE chromaticity x value of 0.38 (range 0.28–0.48) and a mean CIE chromaticity y value of 0.35 (range 0.29–0.41) (see figure 1b). Within the limits of the distribution, all colours had equal probability of being assigned to dots in the specified stimulus area. Target dots were assigned colours from rectangular distributions of which each had the same mean and range of CIE y values as the background distribution, but a different range $[\pm 0.06]$ and different mean CIE x values (see figure 1b). The mean x values were 0.385, 0.395, 0.405, 0.415, 0.425, and 0.435. The target distribution was completely

Figure 1. (a) Snapshot of the colour segmentation stimulus. Dot density, dot size, and placement and size of colour target relative to the grey field are all proportional by the same factor to those in the actual stimuli. The difference between the background and target colour distributions has been markedly enhanced by shifts in both mean and variance in order to ensure that the target is visible under the limitations of colour resolution and dynamic range of the printed version. The difference seen here is therefore supra-threshold. (b) Target and background colour distributions in the CIE chromaticity plane, for a target distribution of $x = 0.405 \pm 0.06$, $y = 0.35 \pm 0.06$, and background distribution of $x = 0.38 \pm 0.1$, $y = 0.35 \pm 0.06$. The dashed triangle gives the gamut of colours displayable on the CRT monitor.

contained within the background distribution for most stimuli and extended only very little outside the background distribution for the rest. The design of the stimuli prohibits target detection based on the colours of individual dots.

(b) Motion parameters

Initially we measured colour segmentation thresholds under three motion conditions. (i) In the broad-broad

Figure 2. Schematic drawing of colour-segmentation stimuli. The broad distributions defining target and background areas are here represented by black dots for the target and white dots for the background. The colour target has the same width of 0.78° in all five motion conditions. In the actual stimulus there are approximately four broad motion stripes on each side of the midline. During the brief presentation of the stimuli observers fixated on the middle of the display. In (a) the stimulus is divided into alternating upward and downward moving stripes. Each motion stripe has the same width as the colour target, and their respective boundaries are aligned. In (e) the stimulus is divided into the same alternating motion stripes as in (a) but here the boundaries of the colour target are located at the middle of two neighbouring motion stripes. In the uniform motion condition, (c), all dots move upward with the same speed. In (b) the stimulus is divided into alternating narrow (0.13°) and broad (0.78°) motion stripes. The boundaries of the colour target are aligned with the boundaries of one of the broad motion stripes. In (d) the stimulus is divided into the same alternating motion stripes as in (b), but here one of the narrow stripes is located at the middle of the colour target.

in-phase condition, the stimulus was divided into alternating upward and downward motion stripes, each uniformly moving with a speed of 2.4° s^{-1} . Each motion stripe had the same width as the colour target (0.78°), with a total of 8.3 motion stripes in the stimulus. The boundary of the colour target was aligned with the motion discontinuities. The task was to detect whether the colour target was to the left or right of the midline. The motion information in itself conveys no information about the location of the colour target (see figure 2a). (ii) In the broad–broad out-of-phase condition, the stimulus was divided into upward and downward moving stripes with the same width as the colour target, as in the in-phase condition, but with the motion discontinuities displaced 180° out of phase with the colour target. That is, the boundaries of the colour target were located at the middle of each neighbouring motion stripe (see figure 2e). (iii) In the control condition, all dots moved upwards with a speed of 2.4° s^{-1} (see figure 2c).

In order to disentangle the effects of motion edges and regions, we added two motion conditions in a second experiment. (iv) In the narrow–broad in-phase condition, half the motion stripes were broad (0.78°) and the other half narrow (0.13°), with the narrow motion stripes aligned with the edges of the colour target (see figure 2b). All motion discontinuities

were fully detectable, as verified by control experiments. (v) In the narrow–broad out-of-phase condition, the narrow motion stripes were displaced such that a narrow motion stripe was located in the middle of the colour target (see figure 2d). The same speed values as in the above conditions were used, so that the local motion discontinuities were the same as in the broad–broad conditions.

(c) Timings and task

Observers performed a spatial two-alternative forced choice (2AFC) task using the method of constant stimuli. Data were collected separately for the five motion conditions. After initiating each trial with a button-press, the observers fixated on a small white cross that appeared in the centre of the grey field and were instructed to stay fixated there until each trial ended. After 500 ms the fixation point was extinguished and the grey field displayed for 300 ms, after which the target/background dot stimulus was displayed for 50 ms (three frames). The stimulus presentation was followed by a 250 ms blank interval (displaying the grey square field devoid of dots), followed by a mask for 67 ms. The mask was designed to interfere with further neural processing of the

Figure 3. Colour segmentation thresholds (mean target CIE x —mean background CIE x , at 75% correct responses) for three observers for five motion conditions. (a): broad–broad in-phase condition in which the colour target is aligned with motion stripes of the same width (0.78°). (b): narrow–broad in-phase condition. The stimulus is divided into alternating narrow (0.13°) and broad (0.78°) motion stripes and the boundaries of the colour target are aligned with one of the broad motion stripes. (c): uniform motion condition. (d): narrow–broad out-of-phase condition. The stimulus is divided into narrow and broad motion stripes and one of the narrow motion stripes is located at the middle of the colour target. (e): broad–broad out-of-phase condition in which the boundaries of the colour target are located at the middles of two broad neighbouring motion stripes. Error bars show the s.e.m.

stimulus, and consisted of 2048 dots of the same size as the stimulus dots, but with randomly assigned colours and motions. The mask duration was selected based on results from earlier experiments (Møller & Hurlbert 1997), and was kept constant throughout all the trials. When the mask was extinguished, observers had unlimited time to indicate, with an appropriate button-press, whether they perceived the target to the left or right of the midline. Each run of a particular experimental condition consisted of 20 trials per target signal value. At least five runs were completed for each condition, with a total of at least 100 measurements for each target signal value for each condition. We determined thresholds by fitting a cumulative Gaussian function using probit analysis to the raw per cent correct values as a function of target signal value. The threshold was chosen as the target signal value corresponding to 75% correct responses.

3. RESULTS

The results for all five motion conditions are shown in figure 3 as 75% correct thresholds; the error bars show the standard errors of the mean. Despite inter-individual differences in absolute thresholds (observer AL performing worst in any one condition), all three observers show the same relative behaviour. Focusing first on the broad–broad motion conditions (bars *a* and *e*), we see that for each observer, the in-phase condition produces a significantly lower threshold than does the out-of-phase condition, with the uniform motion control threshold (bar *c*) in between the two. (Broad–broad in-phase versus out-of-phase significance levels, for the three observers PM, AL, and MB, respectively:

$F_{1,70} = 43.26$, $p < 0.0001$; $F_{1,48} = 29.57$, $p < 0.0001$; $F_{1,40} = 10.70$, $p < 0.0023$. Broad–broad in-phase versus uniform motion significance levels: $F_{1,60} = 16.97$, $p < 0.0002$; $F_{1,40} = 7.127$, $p < 0.011$; $F_{1,40} = 6.672$, $p < 0.014$. Broad–broad out-of-phase versus uniform motion significance levels: $F_{1,70} = 6.531$, $p < 0.013$; $F_{1,48} = 17.72$, $p < 0.0002$; $F_{1,40} = 3.099$, $p < 0.086$.)

These results show that colour segmentation in this task depends on the motion state of the stimulus, and are consistent with the hypothesis that motion discontinuities are shared with the colour segmentation process. It is plausible that motion discontinuities act here by defining boundaries within which the colour signals are integrated, yet the in-phase and out-of-phase stimuli differ not only in the placement of motion edges relative to the colour target, but also in the placement of similar motion signals within stripes. All dots within a motion stripe move in the same way and these dots are flanked by two motion stripes whose dots move in the opposite direction. It is possible that the similarity of motion signals within a stripe defines the area over which to integrate colour information rather than the motion discontinuities alone.

The results from the two narrow–broad configurations of motion stripes are therefore crucial in distinguishing between the influence of motion edges and regions on colour segmentation. In the narrow–broad in-phase condition, the dots in the broad motion stripes flanking the colour target, have the same motion as those within the target, with the opposing motion signals localized to the edges of the

colour target. Although the motion discontinuities in this configuration are essentially identical to those in the broad–broad configuration, the thresholds for the narrow–broad in-phase condition (bar *b*) are higher than those for the broad–broad condition (bar *a*) (significance levels for PM, AL, and MB, respectively: $F_{1,50} = 3.091$, $p < 0.085$; $F_{1,40} = 9.452$, $p < 0.004$; $F_{1,40} = 0.7273$, $p < 0.399$). Thresholds in the broad–broad out-of-phase condition (bar *e*) are larger than the narrow–broad out-of-phase thresholds (bar *d*). For two of the three observers these differences were highly significant (significance levels for PM, AL, and MB, respectively: $F_{1,70} = 4.324$, $p < 0.297$; $F_{1,48} = 23.7$, $p < 0.0001$; $F_{1,40} = 11.56$, $p < 0.002$). For all three observers, the narrow–broad in-phase thresholds are significantly lower than the uniform motion thresholds (PM: $F_{1,50} = 5.152$, $p < 0.028$; AL: $F_{1,40} = 4.306$, $p < 0.045$; MB: $F_{1,40} = 9.611$, $p < 0.004$). Narrow–broad in-phase thresholds (bar *b*) are also lower than narrow–broad out-of-phase thresholds (bar *d*), although the difference is significant for only one observer. (For observers PM, AL, and MB, respectively: $F_{1,50} = 11.03$, $p < 0.002$; $F_{1,40} = 0.9497$, $p < 0.336$; $F_{1,40} = 0.1011$, $p < 0.753$). There is no consistent difference across observers between the uniform motion thresholds and thresholds in the out-of-phase narrow–broad condition.

4. DISCUSSION

The results show that both motion edges and regions may guide the integrative process required for this colour segmentation task. (Motion in itself has no effect on the task, in that colour segmentation thresholds for static stimuli are indistinguishable from those for the uniform motion condition.) The fact that thresholds for both in-phase conditions are significantly smaller than for the uniform motion condition demonstrates that motion edges are shared with the colour segmentation process, particularly in the narrow–broad in-phase condition, where only motion edges exist to facilitate colour segmentation. We surmise that the facilitatory effect arises because the motion edges delimit an area over which to integrate colour information, which corresponds exactly to the colour target. The second experiment shows that motion regions also play a role: the thresholds in the narrow–broad in-phase condition are consistently higher than in the broad–broad in-phase condition. This result implies that the colour segmentation process is guided also by the similarity of motion signals within the broad stripes. In the narrow–broad in-phase condition, the motion similarity extends across the entire stimulus, interrupted only by the narrow motion discontinuities; hence, we hypothesize that colour signals are integrated or ‘smeared’ across a larger surface than the target area. In the broad–broad in-phase condition, the motion similarity stops at the borders of the target area, where the opposing-motion stripes begin.

Conversely, when motion edges or regions are misaligned with the colour target, performance on the colour segmentation task is, in general, degraded. This

effect is weakest when only motion edges are misaligned with the colour target—only one out of three observers shows a significant decline in performance for the narrow–broad out-of-phase condition as compared with the narrow–broad in-phase condition. Thresholds in the broad–broad out-of-phase condition, on the other hand, are larger than thresholds in all other conditions (highly significantly so for at least two out of three observers for each comparison). We attribute this extreme degradation of performance on the colour segmentation task to the fact that in this condition both motion edges and regions are misaligned with the colour target. If motion edges supplied the only influence on colour segmentation, then we would expect comparable performance for the out-of-phase broad–broad and narrow–broad conditions, but in fact the former is far worse. When the broad–broad motion stripes are out of phase with the colour target, the dots in the colour target are divided between two areas moving in opposite directions. Half the dots in each of these two areas come from the target colour distribution, the other half from the background colour distribution. We hypothesize that integrating colour signals over these two areas leads to a weaker colour segmentation signal compared to the signals in both the uniform motion condition and the broad–broad in-phase condition, causing higher colour segmentation thresholds.

These results suggest that motion-defined surfaces may serve as primitive regions over which colour information is integrated, in agreement with earlier work (Møller & Hurlbert 1997). Our implicit assumption is that grouping of similar motion signals occurs more readily for nearby motion signals (stripes separated only by thin discontinuities in the narrow–broad condition) than for distant signals (stripes separated by broad regions of opposing motion in the broad–broad condition). Control experiments support this assumption: an odd-man-out broad stripe, in which dots move upwards faster than dots in all other upwards-moving stripes, is significantly more easily detected in the narrow–broad configuration than in the broad–broad configuration (in both conditions, all dots in the stimulus have colours drawn from the same broad distribution). Easier detection of the odd-man-out suggests that the broad stripes with similar motion group together more strongly in the narrow–broad configuration.

This interpretation of our results also assumes that motion discontinuity and similarity signals influence colour segmentation at an early stage. An alternative explanation is that the presence of the clearly visible motion discontinuities in all conditions allows observers to make eye movements and shift spatial attention to a roughly correct spatial location of the target, thereby bringing the target area into the fovea or spotlight of attention. But since the stimulus duration was only 50 ms and it takes about 200 ms to programme a saccade (Carpenter 1988), eye movements cannot explain any of the effects. Spatial attention, though, might be attracted to a strong motion discontinuity within 50 ms. We may also rule out this possibility, since in all stimuli, motion discontinuities were symmetrical with respect to fixation, excluding any selective

deployment of attention to the right or left of fixation based on their presence.

Furthermore, if the facilitatory effects in the in-phase conditions are caused by enhanced attention, we would expect the same facilitatory effect in the narrow-broad condition as in the broad-broad condition. We would also expect to find similar thresholds in the broad-broad in-phase and out-of-phase conditions, since the motion discontinuities in these two conditions are displaced by only 0.39° relative to each other. This small difference in eccentricity is not sufficient to explain the large difference in thresholds between the conditions.

We therefore conclude that colour and motion signals interact at an early level during image segmentation. Both motion edges and motion similarities within motion regions facilitate the primarily region-based colour segmentation process studied here, even when the motion signals convey no information about the location of the colour target, as the stimulus design insures. These and other results (Møller & Hurlbert 1997) suggest that motion similarities define discrete surfaces within which colour grouping and segmentation processes are deployed. Motion edges may mark boundaries of territories within which to integrate colour information, particularly in this task, where there are no obvious colour edges in the stimulus. This dual interaction between different types of segmentation process in the two cues confers a clear performance advantage, as predicted by computational studies (e.g. Poggio *et al.* 1988) and consistent with other recent psychophysical and neurophysiological findings.

In their study on contour localization, Rivest & Cavanagh (1996) found interactions between all pairings of contours defined by luminance, colour, motion and texture and concluded that neuronal responses to the different attributes must converge in a common neural site before a decision about localization is made. Other evidence supports the conclusion that at some level, contours are represented in an abstract, attribute-independent way. Berkley *et al.* (1994) demonstrated a tilt after-effect for all combinations of adapting and test contours defined by luminance, motion and illusory contours, indicating contour cue-invariance in the orientation domain. As further evidence for the importance of primitive surfaces in guiding early visual processes, He & Nakayama (1994) concluded that for a rapid texture discrimination task, the visual system cannot ignore information regarding surface layout defined by binocular disparity.

Where in the visual system might the interactions responsible for these effects take place? Although anatomical and physiological data indicate that colour and motion information are conveyed in distinct pathways in the primate visual system (Livingstone & Hubel 1988; Merigan & Maunsell 1993; Zeki & Shipp 1988), recent physiological evidence of interactions between colour and motion in the macaque monkey has been found in areas as early as V1 (Sawatari & Callaway 1996), V2 (Gegenfurtner *et al.* 1996) and V4 (Ferrera *et al.* 1994). Of particular importance to our results is the finding that V4 neurons display the same

orientation selectivity for gratings defined either by random dot motion or luminance (Logothetis & Charles 1990), suggesting an integration of visual cues for the extraction of shape information at this early level. At higher levels, it is known that neurons in macaque inferotemporal (IT) cortex are involved in 2D shape perception, and respond invariantly to representations of the same shape in different modalities (Sary *et al.* 1993, 1995). Whether the interaction between motion and colour information in image segmentation demonstrated in this psychophysical study occurs at an early level in the human visual system could be revealed by imaging studies (fMRI/VEP/MEG) using similar stimuli and tasks.

We thank David Bramwell for his extensive input into the experimental design and for all computer programming. Supported by the Wellcome Trust (Project Grant 040075) and the European Community Human Capital and Mobility Programme (Contract CHRX-CT93-0183).

REFERENCES

- Ballard, D. H. & Brown, C. M. 1982 *Computer vision*. Englewood Cliffs, NJ: Prentice-Hall.
- Baraldi, A. & Parmiggiani, F. 1996 Combined detection of intensity and chromatic contours in color images. *Opt. Eng.* **35**(5), 1413–1439.
- Berkley, M. A., Debruyn, B. & Orban, G. 1994 Illusory, motion, and luminance-defined contours interact in the human visual system. *Vision Res.* **34**, 209–216.
- Canny, J. F. 1986 A computational approach to edge detection. *IEEE Pattern Anal. Mach. Intell.*, (PAMI-) **8**, 679–698.
- Caputo, G. 1996 The role of the background: texture segregation and figure-ground segmentation. *Vision Res.* **36**, 2815–2826.
- Carpenter, R. H. S. 1988 *Movements of the eyes*. London: Pion Limited.
- Chen, L. H. & Lin, W. C. 1997 Visual surface segmentation from stereo. *Image Vision Comput.* **15**(2), 95–106.
- Ferrera, V. P., Nealey, T. A. & Maunsell, J. H. R. 1994 Responses in macaque visual area V4 following inactivation of the parvocellular and magnocellular LGN pathways. *J. Neurosci.* **14**, 2080–2088.
- Gegenfurtner, K. R., Kiper, D. C. & Fenstemaker, S. B. 1996 Processing of color, form and motion in macaque area V2. *Vis. Neurosci.* **13**, 161–172.
- Grossberg, S. & Mingolla, E. 1985 Neural dynamics of perceptual grouping: texture, boundaries, and emergent segmentations. *Percept. Psychophys.* **38**, 141–171.
- Haddon, J. F. & Boyce, J. F. 1990 Image segmentation by unifying region and boundary information. *IEEE Trans. Patt. Anal. Machine Intell.* **12**(10), 929–948.
- He, Z. J. & Nakayama, K. 1994 Perceiving textures: beyond filtering. *Vision Res.* **34**, 151–162.
- Hurlbert, A. & Poggio, T. 1989 A network for image segmentation using color. In *Neural information processing systems. I.* (ed. D. S. Touretzky), pp. 297–303. California: Morgan Kaufman Publishers.
- Livingstone, M. & Hubel, D. H. 1988 Segregation of color, form, movement and depth: anatomy, physiology and perception. *Science* **240**, 740–749.
- Logothetis, N. K. & Charles, E. R. 1990 V4 responses to gratings defined by random dot motion. *Invest. Ophthalmol. Vis. Sci.* **31**(4), S90.
- Marr, D. 1982 *Vision: a computational investigation into the human representation and processing of visual information*. San Francisco, CA: Freeman.

- Merigan, W. H. & Maunsell, J. H. R. 1993 How parallel are the primate visual pathways? *A. Rev. Neurosci.* **16**, 369–402.
- Møller, P. & Hayhoe, M. 1993 Influence of colour on motion segmentation. *Invest. Ophthalmol. Vis. Sci.* **34**(4), SI230.
- Møller, P. & Hurlbert, A. C. 1996 Psychophysical evidence for fast region-based segmentation processes in motion and color. *Proc. Natn. Acad. Sci. USA* **93**, 7421–7426.
- Møller, P. & Hurlbert, A. C. 1997 Interactions between colour and motion in image segmentation. *Curr. Biol.* **7**, 105–111.
- Mumford, D., Kosslyn, S. M., Hilger, L. A. & Herrnstein, R. J. 1987 Discriminating figure from ground: the role of edge detection and region growing. *Proc. Natn. Acad. Sci. USA* **84**, 7354–7358.
- Nothdurft, H.-C. 1994 Common properties of visual segmentation. In *Higher-order processing in the visual system, Ciba Foundation Symposium 184* (ed. M. J. Morgan), pp. 245–268. Chichester, UK: Wiley.
- Poggio, T. & Gamble, E. B. & Little, J. J. 1988 Parallel integration of vision modules. *Science*, **242**, 436–440.
- Rivest, J. & Cavanagh, P. 1996 Localizing contours defined by more than one attribute. *Vision Res.* **36**, 53–66.
- Sagi, D. 1991 Spatial filters in texture segmentation tasks. In *Channels in the visual nervous system: neurophysiology, psychophysics and models* (ed. B. Blum). London: Freund.
- Sary, G., Vogels, R. & Orban, G. 1993 Cue-invariant shape selectivity of macaque inferior temporal neurons. *Science* **260**, 995–997.
- Sary, G., Vogels, R., Kovacs, G. & Orban, G. 1995 Responses of monkey inferior temporal neurons to luminance-, motion-, and texture-defined gratings. *J. Neurophys.* **73**, 1341–1354.
- Sawatari, A. & Callaway, E. M. 1996 Convergence of magno- and parvocellular pathways in layer 4B of macaque primary visual cortex. *Nature* **380**, 442–446.
- Zeki, S. & Shipp, S. 1988 The functional logic of cortical connections. *Nature* **335**, 311–317.
- Zhu, S. C. & Yuille, A. L. 1996 Region competition—unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *IEEE Trans. Patt. Anal. Machine Intell.* **18**(9), 884–900.

Received 22 July 1997; accepted 25 July 1997