The Royal Society

Motion perception: from phi to omega

David Rose[1]* and Randolph Blake[2]

[1]*Department of Psychology, University of Surrey, Guildford, Surrey GU2 5XH, UK*
[2]*Department of Psychology, Vanderbilt University, Nashville, TN 37240, USA*

When human observers view dynamic random noise, such as television 'snow', through a curved or annular aperture, they experience a compelling illusion that the noise is moving smoothly and coherently around the curve (the 'omega effect'). In several series of experiments, we have investigated the conditions under which this effect occurs and the possible mechanisms that might cause it. We contrast the omega effect with 'phi motion', seen when an object suddenly changes position. Our conclusions are that the visual scene is first segmented into objects before a coherent velocity is assigned to the texture on each object's surface. The omega effect arises because there are motion mechanisms that deal specifically with object rotation and these interact with pattern mechanisms sensitive to curvature.

Keywords: motion perception; visual perception; omega effect; phi motion; random dots

1. INTRODUCTION

Human visual perception continuously seeks structure and meaning in the dynamic patterns of light imaged on the retina. Even when the spatio-temporal information in that pattern of light is impoverished or underspecified, constructive, synthetic processes fill in the gaps (Gregory 1970; Rock 1983), whether those gaps exist in space (as exemplified by subjective contours: Kanizsa 1955) or in time ('phi motion': Wertheimer 1912; Kolers 1972). The visual system assumes, in other words, that the world is orderly and structured, and, when confronted with unstructured or partly structured input, the brain literally completes the picture. These constructive propensities of human vision are so powerful that they even operate when the retinal input is completely random. For instance, people report seeing regular and repetitive patterns after a few seconds of viewing a dot pattern that is genuinely random (MacKay 1965).

The phi motion effect is one of the simplest of these phenomena. A single dot is displayed in one location at time t_1 and at a different location at time t_2. Within particular limits of the spatial and temporal changes, observers see a single dot moving from one location to the other, rather than one dot disappearing and another dot appearing. A more complex example of such constructive properties is the peculiar and striking 'omega effect' which was first discovered and briefly described decades ago by MacKay (1961, 1965) but which has since been wholly neglected. This effect occurs when viewing dynamic visual noise through a narrow, curved aperture. Dynamic visual noise consists essentially of randomly arranged dots whose spatial positions change haphazardly from time-to-time. An example is the 'snow' seen on a detuned television set. In fact, the effect can be obtained most simply by placing a cardboard mask with a slit cut in it over the face of a detuned television. When the slit is a narrow, elongated rectangle, 1 cm or 2 cm across, the dots visible through the slit, although moving in random directions, tend to appear to stream along the length of the aperture, as if they were bubbles or sparks flowing in a tube. However, if the slit forms a circle or annulus instead of a rectangle, the illusory streaming now becomes highly coherent, and the dots appear to move around the circle at a steady, uniform velocity. The direction of flow appears to reverse spontaneously between clockwise and anticlockwise at irregular intervals of several seconds duration. This compelling illusion occurs even though all directions and velocities of motion are present in the underlying dynamic noise. It is this coherent motion that constitutes and defines the omega effect.

Here, we first present further observations that confirm and extend MacKay's brief original reports, and then we discuss how these may provide indications to the origins of the effect and the mechanisms of motion perception. Preliminary results and videotaped demonstrations have been presented (Rose *et al.* 1994*a*,*b*; Rose & Blake 1995).

2. METHODS

Many of our basic observations and quantitative experiments were done with masks placed over the face of a detuned television set. The masks consisted of computer-generated figures, laser printed and then photocopied onto overhead projector transparencies. Additional studies were made with masks superimposed digitally over dynamic random noise with a computer videographics board. In other experiments we generated noise dots directly within a defined area of the face of a computer screen by using customized software. Full details are given here with the quantitative results.

*Author for correspondence.

3. RESULTS

(a) *Basic observations of the phenomenon*

For our initial explorations, we generated a 'bullseye' pattern of four concentric rings of dynamic noise, of varying width, separated and surrounded by a black mask (figure 1a). The authors and many other visitors and staff viewed the display and made comments. The omega effect, i.e. circular streaming, was readily apparent in all four rings without prompting or instructing the observers. It was very robust and did not depend crucially on any particular fixation strategy or deliberate pattern of eye movements. The direction of flow reversed at irregular intervals of several seconds duration (comparable to the behaviour of other bistable visual phenomena including the Necker cube and binocular rivalry). To the observers, it seemed these reversals could also be triggered by an act of will; however the direction of rotation could not be held constant indefinitely by volition. For example, our attempts to generate motion after-effects from the omega effect had to be abandoned because observers could not keep the rotation going in the same direction.

Flow in all four rings was often in the same direction at any one time; however, it was possible to perceive different directions of rotation in different rings simultaneously. Also, the omega effect did not depend crucially on viewing distance; as we moved away from the display, the stimulus took on a 'movie wagon wheel' appearance, with large blurred blobs moving around the rings at approximately the same speed. Thus the rotatory motion persisted despite the drop out of high spatial frequencies.

Parallel straight lines (figure 1b) did not generate such coherent motion: although there was a tendency for the noise dots to stream along parallel with the lines, in one direction or the other, this motion was weak, unsystematic and irregular. Similarly, a radial pattern (figure 1c) very rarely generated any sense of coherent expansion or contraction: instead, the dominant percept was of chaotic linear flow, especially near the centre, with no cohesion of direction between the various arms. Occasionally, there was a percept of rapid rotation of the dots around the centre. Both these phenomena will be dealt with further in subsequent sections.

(b) *Quantitative experiments: series 1 and 2*

Masks were placed over the screen of a detuned 20-inch Sony television monitor. The mean luminance of the dynamic dots was $12 \, \mathrm{cd \, m^{-2}}$, of the black masks $0.7 \, \mathrm{cd \, m^{-2}}$, and the experiments were done in a room partly illuminated by fluorescent lighting (background *ca.* $7 \, \mathrm{cd \, m^{-2}}$). The 'standard' mask covered the left side of the screen and a 'test' mask was shown on the right. In series 1 the standard was the bullseye pattern whose outer diameter subtended $4.2°$ at the viewing distance of 1.83 m; in series 2 the standard was a single annulus of inner diameter $1.7°$, outer diameter $2.4°$. A total of eight observers naive as to the purposes of the experiment were used in each series. The observers were asked to rate the strength of the apparent flow in the test stimulus, relative to the standard stimulus, which was given a rating of 10. The coherence of the flow was stressed as
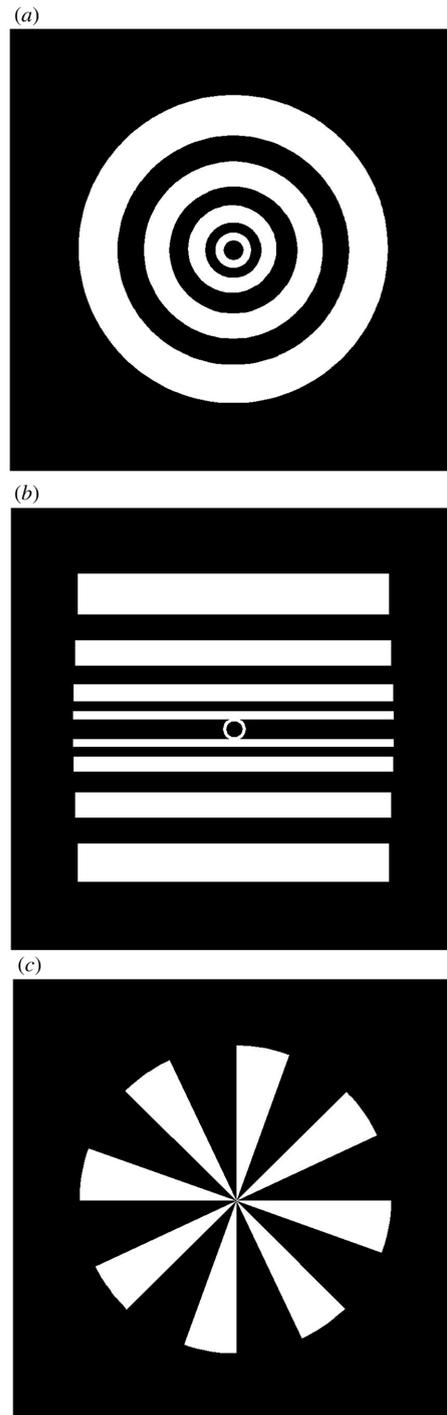

Figure 1. Masks used for preliminary studies. These were printed and photocopied onto transparencies, which were positioned over the screen of a detuned television. The 'snow' on the screen was thus visible within the white areas surrounded by the black mask.

the relevant feature to judge. The test stimuli were exposed for 30 s each in random order within each series. A rating procedure was used to quantify the effect rather than nulling because the phenomenon shows reversals of apparent direction at unpredictable intervals. Also, ratings could be obtained across the entire variety of mask shapes we used, which would not have been possible

with other techniques such as nulling or velocity estimation.

(i) *Series 1*

The first issue we addressed was whether the effect depends purely on the curvature of the edges or whether it is determined by the overall configuration of the display. Accordingly we tested a series of polygonal figures varying between near-circles and near-straight lines. As figure 2a shows, the effect declines only slightly for octagonal and hexagonal stimuli, indicating that local curvature is not essential. However, the square stimulus was not significantly more effective than the straight line stimulus, showing that continuity around a loop is not sufficient by itself to generate the effect.

The rightmost two stimuli in figure 2a demonstrate that the endings of the noise alleyways do not have any significant influence, as open- and closed-ended stripes engendered similar weak effects. (A 'noise alleyway' is any long, thin, unoccluded region which is circumscribed by the mask and within which the random motion stimulus is visible.)

Could the effect be generated by motion vectors at one location momentarily signalling one direction of motion, which then spreads its dominance throughout the remainder of the noise alleyway? (Processes that spread across the visual field between constraining edges, and which thus 'fill in' between borders, have been postulated by, for example, Walls (1954), Gerrits & Vendrik (1970), Grossberg & Mingolla (1985), Paradiso & Nakayama (1991) and Lee (1995).) The existence of a continuous loop in the display could enable this process to reverberate and thus to self-reinforce by temporal summation. As figure 2b shows, however, the effect is not destroyed by breaking the loop, as it should be if this self-reinforcement hypothesis is true. In fact, the strength of the effect is almost monotonically related to the amount of the bullseye visible, regardless of whether the bullseye is divided into one, two or four sectors. Even semicircles generate a very strong effect. (Incidentally, with four sectors, observers reported that although these mostly rotated all in the same direction, sometimes they perceived different arcs of the same annulus to be rotating in opposite directions.)

Another issue we addressed was whether the strength of the perceived coherent motion depends crucially on the total length of local edge between the mask and the dynamic noise (and the ratio between the noise dot size and the circle's width and diameter). We presented circles of the same mean diameter but of different widths, and thus with the same total length of edge contour (figure 2c). The omega effect was very weak when the display area was a wide annulus with only a small central spot, and it increased for narrower noise alleyways. These displays also gave rise to some useful ancillary observations. With the very wide annulus, the location of the isolated central black spot (left part of figure 2c) appeared to drift about randomly. With the medium diameter annulus, the flow effect began most clearly near the edges and then spread slowly out to capture the centre of the annulus, so that within about 20 s the entire enclosed region of noise appeared to be rotating coherently. Evidently, therefore, some kind of cooperative
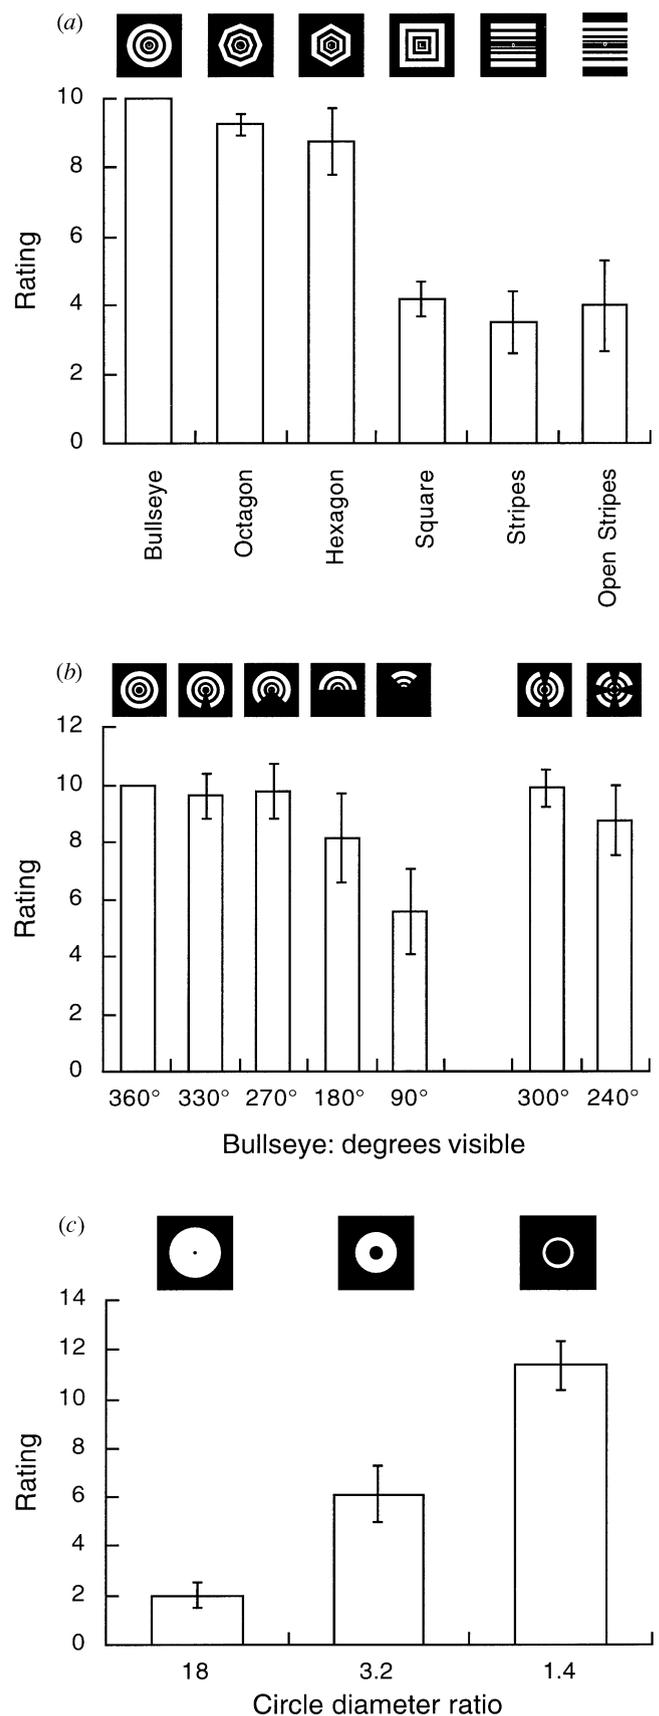
Figure 2. (*a*) Mean ratings assigned to stimuli in series 1 designed to test the importance of curvature. Bars show standard errors. Insets above each column illustrate the displays, which contained dynamic random 'snow' in the white regions within the black masks. (*b*) Further data from series 1, testing the effect of interrupting the bullseye. (*c*) Further data from series 1, investigating the importance of annulus diameter.

spreading does occur; the local mask edges themselves are not the sole determining factor.

(ii) *Series 2*

In this series we used a single annulus as the standard, and all the test stimuli comprised strips of noise of the same width as the standard (0.33°).

A figure 'S' of exactly the same total length, area and curvature as the standard ring evoked a clear, but weaker, impression of coherent motion (leftmost column 1 of figure 3a). The remaining data in figure 3a examine the role of curvature in more detail. All the test displays had identical length and area of noise and also identical edge contour length. As the noise area changed from semicircular to straight there was a progressive weakening of the apparent flow (Page's L-statistic $= 385$, $p < 0.05$; data from columns 2–6 of figure 3a).

The oval 'running-track' figures illustrated in figure 3b (insets 2–4) were presented to see whether the flow induced in a semicircular region would be in any way diluted or slowed by running into the lesser flow in a straight segment. It is clear that the effect was not diminished by inserting straight sectors into an annulus. Indeed even the long 'motor race-track' (rightmost inset 5 in figure 3b) gave a similar strength rating (the perception of this display is discussed in more detail below in § 3c,i). Thus the total length of the exposed strip of noise does not itself determine the apparent flow of the noise.

However, when annuli of different diameters were tested (figure 3c) it was apparent that the ratings of motion strength were affected. Taking these data together with those of figure 3a, it seems that the radius of curvature is an important factor, with some suggestion of 'tuning' for arcs with radii close to 1–2° (figure 4).

(c) *Further observations*

We also examined numerous other variations on the display pattern, and although our observations were more informal than those in the previous section, they were nevertheless robust, as witnessed by numerous visitors to our laboratories and observers of our videotaped demonstrations at conferences (Rose *et al.* 1994a,b; Rose & Blake 1995).

We aimed to understand the mechanisms that generate the omega effect, and our experiments were therefore guided by a number of ideas as to the origin: (i) neural channels for motion analysis (based on magno cells and/or cortical areas MT–V5 or MST); (ii) computational mechanisms of dot matching; (iii) optic flow mechanisms; and (iv) (cognitive) processes of scene segmentation. These are not mutually exclusive categories and we are not committed to any of them, but they guided and inspired our choice of stimuli.

(i) *Scene segmentation*

The 'motor race-track' stimulus (rightmost inset in figure 3b) contained two intersections where the flow crossed itself at an angle. At each intersection, one track was normally perceived as dominant, with flow proceeding uninterrupted across the intersection. This track appeared in front, with dots in the other track flowing behind, to give a 'fly-over' appearance to the intersection. At each intersection, either track could dominate
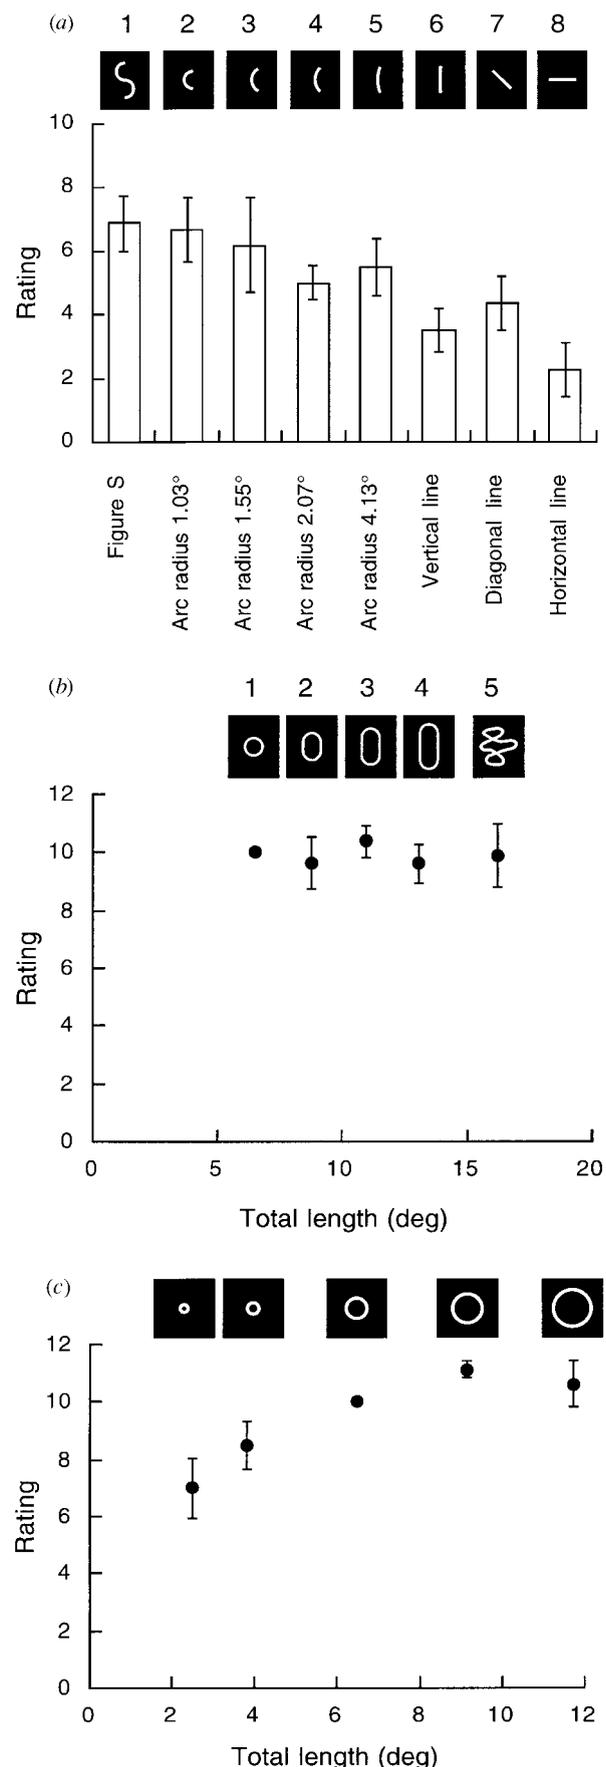

Figure 3. (a) Mean ratings assigned in series 2 to test the importance of annular closure (leftmost column 1), curvature (columns 2–6) and orientation (columns 6–8). (b) Further data from series 2 to test the effects of path length. (c) Further data from series 2 to investigate the effect of circle diameter.

in alternation, with irregular periods of a few seconds between the changes in dominance.

A more dramatic example is given in figure 5*a*. The two rings in this 'Olympiad' figure intersect twice, and as before, at each intersection, flow dominated along one ring and that dominance alternated between the rings. When a given ring dominated both intersections, that ring was seen as lying in depth in front of the other ring. When each ring dominated one intersection, the two rings were seen as interlocking and slanted in depth, much as two real interlocking rings would be.

Our attempts to induce coherent flow purely by 'cognitive contours' were less clear-cut. In figure 5*b*, rotational streaming could be seen around the circular channel in a narrow annulus, much as with the previously described real annuli (see, for example, figure 3*c*). However, the flow tended to escape and run up the radial alleyways as well, and could sometimes even be seen circling around the outside of the whole figure. One observer reported seeing capture of the whole inner part of the figure by the annulus, so the dots appeared to move coherently clockwise or anticlockwise, even behind the inner black spokes.

The 'Celtic cross' figure 5*c* consists of four open triangles forming a fan plus our original bullseye pattern. Viewing random motion behind this mask, one saw an incoherent streaming of isolated dots shooting outwards from the central narrow apex of each triangle, the whole resembling a Roman-candle firework. This was phenomenally distinct from the coherent flow that was seen around the annuli (see §4). Observers reported the triangles to appear mainly in front of the annuli, and only less frequently could a circle be seen running across in front of a triangle. Different arcs in the same annulus could flow in opposite directions simultaneously, just as occurred when the fan was black (as in figure 2*b*, rightmost column inset). The outer blind endings of the triangles (i.e. their bases) were never captured by motion around the outermost annulus, and the display never resembled a sheet of sandpaper rotating behind the mask.

When two annuli intersected to form a 'figure 8' pattern (figure 5*d*), the flow sometimes ran along an S-shaped path from one circle to the other, and sometimes ran round within each circle, with one or the other ring dominating at the intersection and thus appearing in front. The rotation in each ring was often in opposite directions, 'like gear wheels', as one observer described it. When two rings were presented either abutting (just touching one another) or separated by a clear gap (not illustrated) they could rotate in the same or in opposite directions. Quantitative measures showed however that there was no change across these three displays: in each, the two rings were seen rotating in the same direction for 48% of the time on average (s.e. 1.94%; between the three displays: $F_{2,27}=1.33$, n.s.; data from three observers). There was thus no Gestalt grouping across two annuli whether isolated, continuous or contiguous.

We considered earlier whether a filling-in or cooperative process sweeping across the internal spatial map of the image was responsible for the predominance of flow in closed loops, and we concluded that it was not (figure 2*b*). At the same time, we asked whether such a process might 'bounce' off the endings of closed alleyways of dynamic noise. The rightmost two columns in figure 2*a*
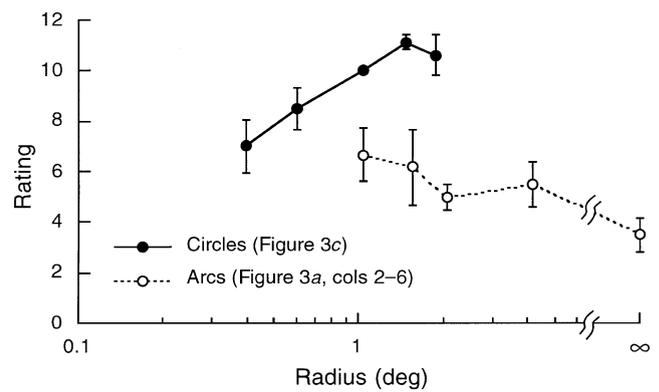
Figure 4. Summary of data from figure 3*a,c* to examine the tuning of the effect for stimulus curvature.
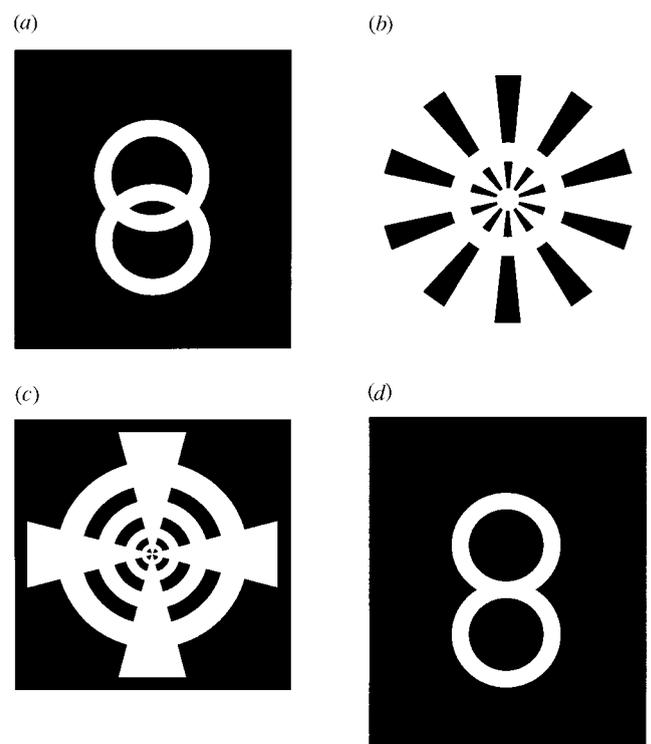

Figure 5. Some further stimuli used to investigate the conditions that give rise to coherent streaming. (*a*) The Olympiad. (*b*) An annulus formed by cognitive contours. (*c*) The Celtic cross. (*d*) The figure eight.

show that closed- and open-ended straight stripes evoke equivalent, low degrees of coherent motion. Figure 6 illustrates that closed endings to curved alleys are also not crucial: flow can readily be seen along the open-ended semicircle, as it can with closed semicircles (figures 2*b* and 3*a*). Figure 6 provides a particularly compelling mask for demonstrating the omega effect, and we urge readers to photocopy it onto a transparency (enlarging it if necessary) and observe its effect on the snow produced on a detuned television. Streaming is readily seen to be much stronger along the semicircle than along the straight line, despite them having identical lengths, widths, areas and open endings. Placing a dark rod across the endings of the semicircles does not affect the apparent flow therein.

Unless one postulates that cognitive contours are imposed across the endings, and that these segment the scene as effectively as real contours, we must conclude that the closed endings are not responsible for the effect. Note also that the flow around the semicircle does not extend out into the open area of the display; there is certainly no phenomenal completion of the rotational motion to form a complete annulus.

Realizing that depth is a potent source of information for scene segmentation, we wondered whether placement of the mask in a distinctly different plane of depth from the dots would uncouple their interactions and thus abolish the omega effect. To our surprise, bringing an annular mask forward away from the TV screen did not destroy the apparent rotational flow within the annulus. Indeed, if the mask was advanced so far that the two eyes were seeing separate areas of the display, and hence totally uncorrelated noise, the flow still persisted.

However, flow could not be induced interocularly. When we presented the dynamic noise to one eye only and a photocopy of the mask (i.e. a white annulus on a solid black background, or the black and white bullseye) to the other eye, via a mirror, the flow of the noise appeared random, irrespective of the state of dominance, transparency or mixed rivalry between the stimuli (see the paper by Yang *et al.* (1992) for a discussion of the many possible states of binocular rivalry and fusion).

(ii) *Local edge and contrast effects*

We were also interested in testing more systematically the role of the edges of the masks in constraining the possible range of matches between pairs of dots. To accomplish this, we generated variations of the bullseye mask by using commercial graphics programs on an Apple Macintosh Quadra 950 computer and displayed them on a JVC 20-inch TV monitor via a Truevision NuVista+ video-graphics board. Truevision's Blender software was used to superimpose the mask image on dynamic snow, which had been videotaped previously by superimposing an outer border on a detuned TV signal in a Panasonic WJ-MX12 mixer.

With this method we were able to show that many characteristics of the mask are not crucial. Black, white, red or grey masks (of the same mean luminance as the dynamic noise) were all effective in generating robust apparent flow. With the equiluminant grey mask, we also reduced the contrast of the noise to levels which, according to some views, should have rendered it invisible to parvocellular channels (Shapley *et al.* 1981; Sclar *et al.* 1990). However, the flow was still readily apparent at this low contrast, supporting the association between motion perception and magnocellular mechanisms (see, for example, Livingstone & Hubel 1987; Tootell *et al.* 1995).

We also tested a mask made of static dots identical in size, luminance and density to those comprising the dynamic noise (figure 7a). We reasoned that false matches made between the dots forming and surrounding an annulus might impair the rotation. (The noise mask was obtained from a freeze-frame of the videotaped snow. We checked its validity by videotaping the blended signal, i.e. the static noise mask with the dynamic noise filling the annuli, and freeze-framing that videotape. The boundary between the mask and annular regions was then not
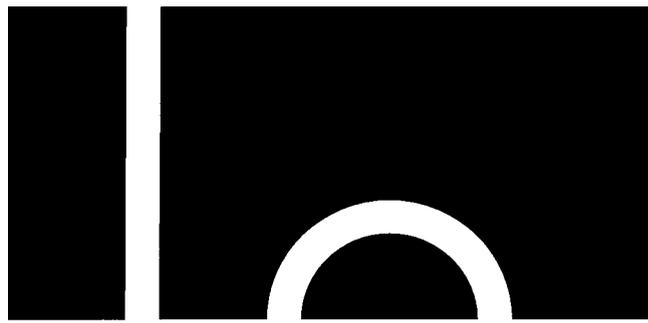

Figure 6. Mask used to compare curved and straight alleyways of noise, with equal lengths, widths, areas and ending types for both alleyways.

visible, the whole screen being filled with uniform static noise.) However, the static noise mask readily induced steady flow around the annuli, with sharp borders seen between the dynamic and static regions.

Next, to test the role of the mask edges, we dithered them to simulate blurring. With a black dithered mask (figure 7b) the rotational flow in the dynamic noise regions of the display was still readily apparent. The annuli took on a tube-like appearance, as the edges resembled those of a shadowed three-dimensional solid. Further, the flow was also clearly visible with a static noise dithered mask (figure 7c). These demonstrations show that sharp edges are not necessary. They also show that static dots in the surround do not interfere significantly with the process that generates perceived coherent motion.

Finally, and most dramatically, we generated a display in which the mask itself consisted of dynamic noise. This was accomplished electronically by making the dots within a virtual annulus (or within the alternate annuli of a bullseye) brighter and of higher contrast relative to those within the rest of a display screen filled with dynamic dots. (A similar display, with rings of brighter but lower contrast dots, was also created by making a slide of a white annulus on a black background, or a black and white bullseye, and projecting its image optically onto the screen of a detuned TV. Adjusting the brightness of the projected image appropriately enabled the same effects to be seen as with the purely electronic display.) When the annuli of the bullseye were visible as bright rings in the dynamic noise, the noise readily appeared to rotate around inside them, and also around the dark rings between the bright ones (and around the outside of the outermost annulus). These tended all to rotate in the same direction at the same time at the same angular velocity. When the display contained only a single bright annulus, the circular centre within the annulus was captured by the motion within the annulus and rotated in synchrony with it; the dots immediately outside the annulus also rotated in synchrony, but this effect faded with distance from the annulus.

When the bullseye was made less salient (i.e. the dynamic noise rings differed little in brightness and contrast from the dynamic noise background), we obtained an unstable percept. At times the bullseye was visible and the dynamic noise rotated around within it.
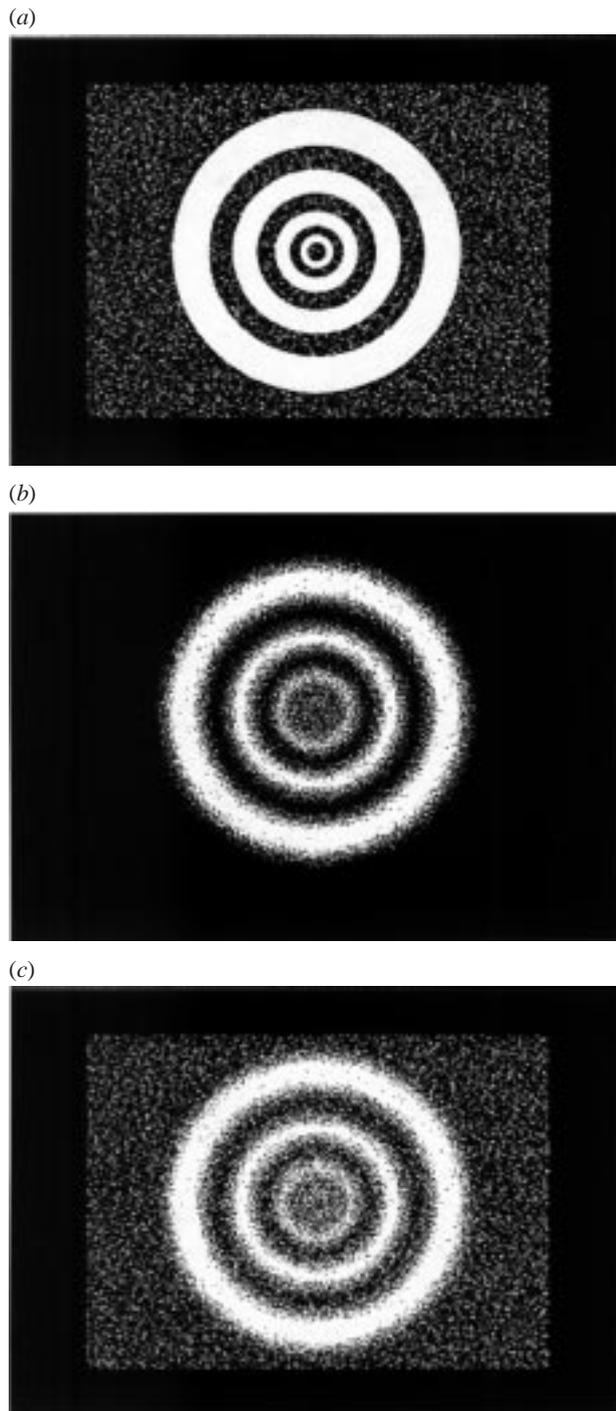
Figure 7. Electronically generated masks to test the role of edge sharpness and of dot matching. Dynamic random noise was visible in the central white annular areas. (*a*) Mask filled with static noise. (*b*) Mask with dithered edges. (*c*) Mask with static noise and dithered edges.

At other times the bullseye or annulus disappeared and the motion appeared random with no clear direction. At other times we could perceive piecemeal dominance, with shifting patches of random noise intervening between arcs and sectors of the bullseye, in which the flow was rotational. The scene resembled very much the dynamic phenomena obtained under binocular rivalry, although with larger patches (Blake *et al*. 1992; Yang *et al*. 1992). There seemed to be an interaction between the pattern
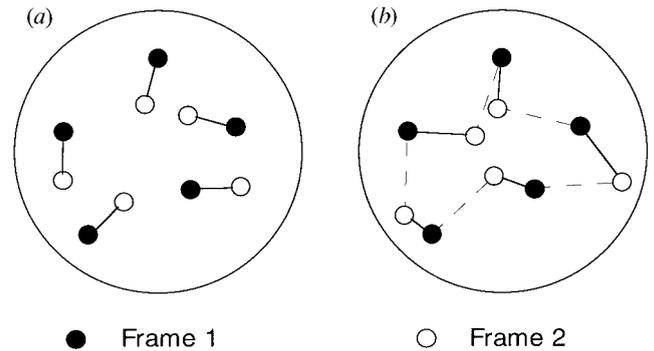
Figure 8. Dot matching from one frame (solid circles) to the next (open circles). (*a*) With short steps, of equal lengths but random directions, the matches are clear. (*b*) With longer steps (dashed lines), the nearest-neighbour matches (solid lines) are no longer the 'correct' ones.

and motion systems in that rotational flow was seen only when the pattern of the bullseye was visible. Whether the pattern's visibility was necessary for (preceded in a causal sense) the rotational flow, or vice versa, remains to be determined. This issue will be addressed in §4.

The bias between pattern and motion salience in this display could also be altered by changing viewing distance: moving away from the TV monitor (or defocusing the eyes) made the high spatial frequencies in the dots drop out, and the bullseye became relatively more visible, with rotational motion of blurred blobs within it. Closer to the screen, the dynamic dots became more dominant and the bullseye disappeared.

All the phenomena mentioned here were seen whether the edges between the dynamic noise background and the brighter dynamic noise annuli were sharp or dithered.

(iii) *Dot matching*

Dynamic television snow contains motion in all directions at all velocities on average, but the phi motion phenomenon is traditionally investigated by using a single pair of dots that appear at different locations at different times. Here, we explore the territory between these two types of stimulus to investigate the question of what conditions determine the switch from seeing individual dots moving in phi fashion to seeing the global coherent motion of the omega effect.

Consider the five dots depicted in the two-frame apparent motion sequence in figure 8*a*. Each dot makes a small step, and it is unproblematic which directions of motion should be perceived; the array amounts virtually to five independent phi movements. However, if step size is increased, as shown in figure 8*b*, then for each of the five dots in frame one, the nearest dot in the second frame is no longer its 'true' partner. Instead, one of the other dots has jumped into closer proximity with it, and is likely to capture the pairing. Williams & Sekuler (1984) have given a formula for calculating the probability of occurrence of such 'false matches'. The probability is $1-\exp(-\pi ds^2)$, where d is the dot density and s is the step size. Williams & Sekuler (1984) used this formula to predict when coherent translational flow would be seen among dots whose local motion vectors varied randomly within a limited range of directions. (The formula is described as 'well-known' by Diggle (1983, p. 16) but we

will continue to refer to it as the 'Williams & Sekuler formula' as they introduced it to the vision research community.)

For our display, the formula is only an approximation because it ignores: (i) dots moving on and off the edges of the display; (ii) dots landing on top of one another, which can occur because they have finite size; and (iii) dots matching over more than one frame of multi-frame sequences. Nevertheless, with our display parameters the approximation is close enough to enable us to test whether the basic principle underlying it has predictive power. We therefore compare the predictions of Williams' & Sekuler's formula with data on the threshold for perceiving the omega effect as a function of dot density and frame-to-frame step size.

Methods

For these measurements television snow was inappropriate because we needed to be able to control the possible pairings of dots that can be matched across space and time to generate motion signals. This required the use of computer-generated dynamic noise stimuli. We therefore generated these on an Apple Macintosh computer by using custom software. Black dots (2.9×2.9 min square, 4 cd m^{-2}) were displayed on a white screen (Apple 13-inch colour monitor, 80 cd m^{-2}). Initially, various numbers of dots were randomly positioned within a square region $4.23° \times 4.23°$. However, this region was masked in software so the only dots displayed on the screen were those within an annular region of $1.21°$ inner diameter and $4.23°$ outer diameter. The dots were animated at 10 frames s^{-1}. Between each frame, every dot jumped the same distance in a direction chosen randomly for each dot. Over a series of frames, each dot therefore performed an independent random walk in space. Dots moving outside the square were wrapped-around.

Observers viewed two such displays side-by-side from a distance of 79 cm. The right-hand screen was used to exhibit a 'standard' stimulus and the left screen exhibited a 'test' stimulus. Each stimulus was displayed for 10 s. The standard stimulus contained on average 16.8 dots deg^{-2}, which, given their size, meant that dot density (the percentage of the annulus area consisting of dots) was 3.9%. Each dot moved 14.5 min arc between frames, equivalent to 2.4 deg s^{-1} velocity. The test stimuli had dot densities that varied parametrically between 2.8 and 16.8 dots deg^{-2}, and inter-frame movement steps of between 1.45 and 43.5 min arc.

The observers were instructed to fixate a small cross that was positioned in the exact centre of each annulus, but were allowed free movement of the eyes from one display screen to the other. They were asked to rate the strength of the coherent movement seen in the test stimulus, given that the standard stimulus had strength '10'. A total of three repetitions of each test stimulus were presented, randomly intermixed with the other test stimuli. Each test stimulus was generated anew before each trial. Author D.R. and four naive observers participated.

Results

Preliminary qualitative observations with this type of stimulus showed us that with five or more dots per degree

squared and inter-frame jump sizes of 10 min or more, rotational flow was readily apparent, despite the absence of any sharp contrast edges demarcating the annular region that contained the dots. The direction of rotation reversed periodically, as it does with television snow (§3a). At lower dot densities the global coherence of the motion tended to be lost, and decreasing the jump size also reduced, and eventually abolished, the sense of systematic directional flow. Instead of flow the display appeared to consist of many squirming or writhing specks with, at most, occasional brief rotational oscillations. The two parameters, density and step size, interacted: higher values of one permitted lower values of the other, while still inducing coherent rotational motion. Conversely, as either dot density or jump size was increased the omega effect became stronger.

The data from the quantitative rating studies essentially verify these qualitative observations. We first plotted the data as a function of inter-frame step size seperately for each dot density (figure 9).

Then we aimed to fit a smoothing function to these data. First we tried empirically to fit logarithmic, polynomial and exponential functions; of these, logarithmic and second-order polynomials gave the best fits overall. However, we found the following theoretical treatment to give fits as good as, or better than, any of these (and on p. 976 we give reasons why this particular theory is actually preferable to those empirical fits).

If Williams & Sekuler's formula is relevant, the transition from seeing local phi motion vectors to seeing global omega motion should occur at some threshold value of $1 - \exp(-\pi d s^2)$, in other words at some constant probability of obtaining false matches. This theory is qualitatively consistent with what we have observed of the dependence on dot density and step size; but is the theory also correct quantitatively? The problem we have is that our data may be contaminated by an unknown nonlinearity, namely the way observers were using the rating scale. In other words, their perceptual experiences may have been determined by mechanisms that reflected the number of false matches, in accordance with the Williams & Sekuler formula, but their response mechanisms might not have attached a number to each perceptual experience in a linear fashion.

To cope with this problem, we assumed any such nonlinearity was a power function. This assumption seemed reasonable because observers generally give ratings that are power functions of stimulus magnitude (Stevens 1951). The prediction therefore becomes that the observers' ratings would equal $a \times (1 - \exp(-\pi d s^2))^b$, where a and b are constants. As we knew the values of d and s for each data point, we could use curve-fitting algorithms to find the two unknowns a and b and to test whether the fits were good.

In figure 9 the results of fitting this theory are shown for each observer and for each dot density; the fits are on the whole statistically satisfactory (table 1).

To examine our power function hypothesis we looked at the values of b, the best-fitting exponent. Figure 10 shows that for four observers b was, on the whole, close to unity, suggesting that if the Williams & Sekuler formula is appropriate it is accompanied by linear response mechanisms. Only for observer D.R. was there any systematic deviation from unity i.e. b rising as dot density increased.
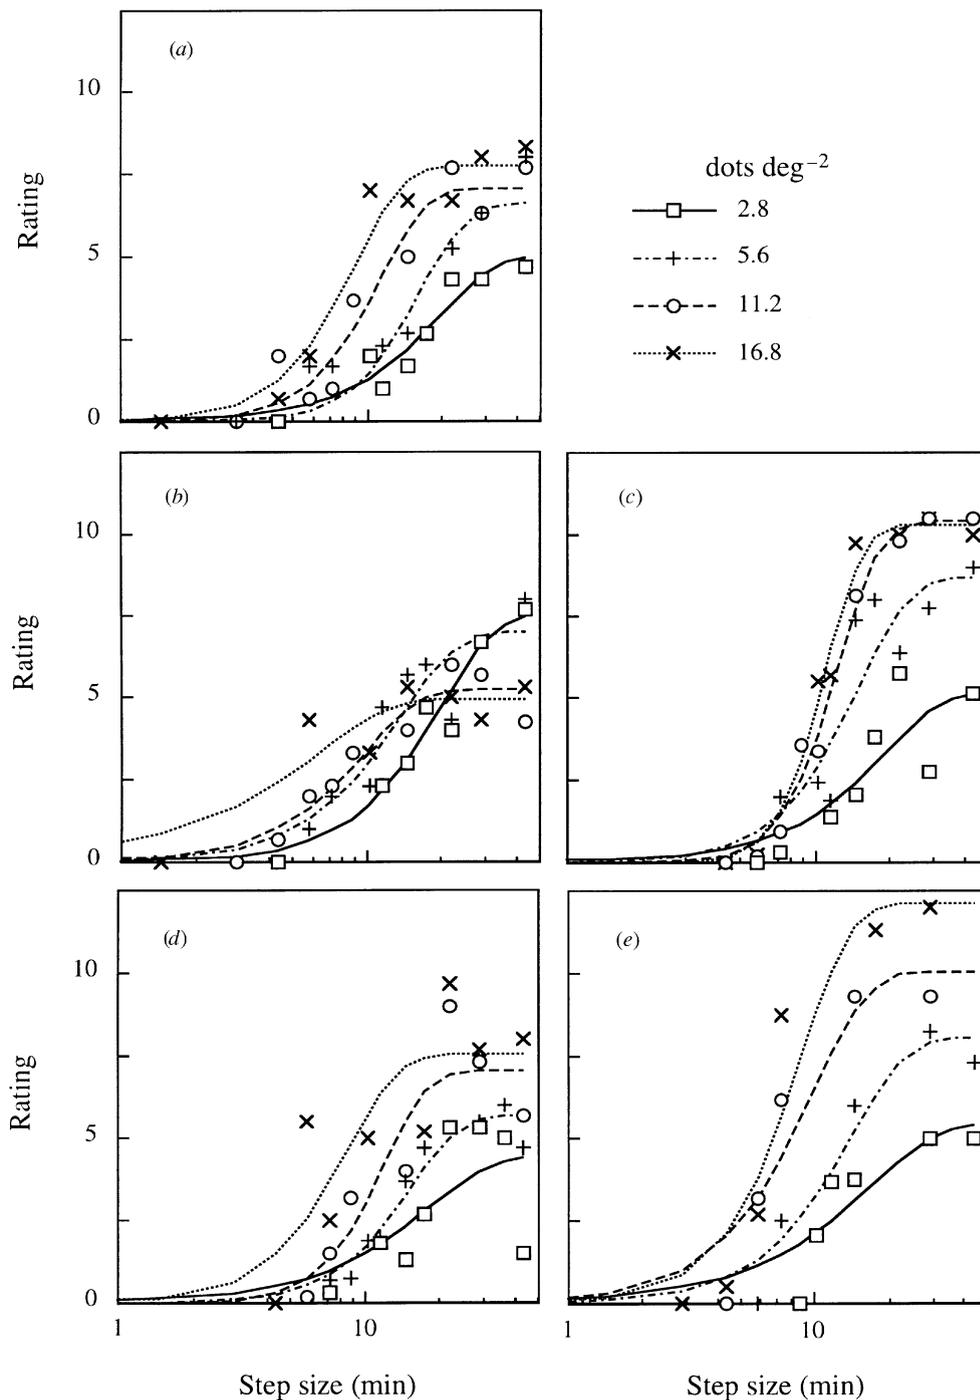
Figure 9. Ratings of the strength of the omega effect as a function of frame-to-frame step size, with dot density as parameter, for five subjects (*a*) A.B.; (*b*) C.F.S.; (*c*) D.R.; (*d*) P.S.; (*e*) S.A. The curves are the best fits of the Williams & Sekuler formula passed through a power function (see text for details and table 1 for goodnesses-of-fit).

As mentioned here, the fits obtained with our equation were overall as good as, or better, than any other simple empirical fit. As a further justification, we can look at the asymptotic behaviour of the theoretical curves (see, for example, figure 9). These imply that as step size increases towards infinity the ratings should level off. The rival logarithmic and second-order polynomial fits however would predict respectively that ratings should continue to rise or should start to fall again. The data in figure 9 clearly show a levelling off, as predicted by our equation.

We also did an even stronger, empirical test. As step size increases, the chances of matching a dot correctly across successive frames decreases, as other dots are increasingly likely to have jumped into the intervening space and captured the pairing (figure 8*b*). If the step size continues to increase towards infinity, these displays therefore asymptotically approach total chaos or randomness. Accordingly, we included an experimental condition in which all the dot locations were chosen entirely randomly on each frame, regardless of the dots' positions in previous frames. A total of four of our observers were asked to rate any coherent motion in such displays, with the same 'standard' stimulus as before as the benchmark for a '10' rating.

Table 1. *Results of goodness-of-fit tests for the modified Williams & Sekuler formula applied to the rating data in figure 9*

| | dot density (dots deg^{-2}) | | | | | | | |
| | 2.8 | | 5.6 | | 11.2 | | 16.8 | |
	χ^2	r	χ^2	r	χ^2	r	χ^2	r
observer								
A.B.	1.96	0.953	8.55	0.919	6.00	0.958	4.28	0.977
C.F.S.	2.10	0.974	8.68	0.928	2.92	0.957	4.14	0.897
D.R.	9.20	0.842	14.37	0.917	2.14	0.993	2.24	0.992
P.S.	15.80[a]	0.667	2.18	0.970	9.84	0.913	22.59[b]	0.816
S.A.	5.30	0.845	3.23	0.966	7.89	0.938	14.68[a]	0.950

[a] $p < 0.05$
[b] $p < 0.01$.

These completely random displays produced clear rotational perceived flow, and furthermore this flow was still apparent even when dot density was made very low (for some observers, with as few as five dots in the whole display). The ratings given are shown in figure 11a. There were some individual differences in the willingness to rate coherence at low densities, but the overall data clearly support the existence of a strong omega effect in these totally random arrays.

We have not attempted to apply any theoretical smoothing function to the data of figure 11a, but have simply applied logarithmic fits. Now, we can read off from this figure the predicted ratings for displays of a given dot density and infinite step size. These predictions are plotted in figure 11b against the actual asymptotic values, a, of the curves fitted previously with our equation (see, for example, figure 9). The predictions match the asymptotes very well, supporting the use of our theoretical fit rather than the logarithmic or polynomial alternatives.

Having justified our theoretical smoothing function, we can now test whether Williams & Sekuler's (1984) false matching formula actually predicts the omega effect. One way of circumventing any uncorrected response non-linearities is to maintain a constant level of response while varying the stimulus input. For example, let us ask what stimulus parameter combinations will all evoke the same rating, say '2', from our observers. Figure 12a shows how this criterion rating is applied to one data set from figure 9. Particular combinations of dot density and step size can thus be found for each data set that elicit a rating of '2'. If false matches determine the incidence of the omega effect, the values of $1-\exp(-\pi ds^2)$ should be the same in all cases. This hypothesis is tested for all data sets (five observers and four dot densities) in figure 12b. The hypothesis is clearly supported: observers rated the strength of the omega effect as '2' whenever the probability of a false match was 0.34 (on average). There was no change in that probability with dot density ($F_{3,16}=0.38$) and no linear trend in the data ($F_{1,16}=0.43$). Figure 12b also shows the outcome for a higher rating of '4'. This rating was elicited by a false match probability of 0.66, regardless of the particular values of s and d ($F_{3,16}=1.16$, n.s.; linear trend: $F_{1,16}=2.29$, n.s.).

4. DISCUSSION

(a) *Comparison with previous studies of illusory motion phenomena*

Our observations have confirmed and extended those of MacKay (1961, 1965). MacKay asked his observers to estimate the apparent angular velocity of the streaming, which was typically about one revolution every 2–4 s, increasing with the ratio of annulus diameter to thickness. The flow's velocity was unaffected by retinal angular image size, brightness, focus, noise-grain size or noise frame-frequency. Interruption of the annulus had little effect, but polygonal alleyways (including triangular) generated reduced rates of flow, especially near the corners. Several of these observations are confirmed by the data in our figures 1 and 2. We have extended MacKay's observations (figure 3 ff of this paper) to test various hypotheses concerning the origins of this illusion.

We concur with MacKay's (1961) conclusion that the omega effect is 'a quite distinct phenomenon' in comparison with the 'complementary images' generated by repetitive patterns of high contrast bars; radial lines fanning out from a point (MacKay's rays) generate 'rosettes' and concentric circles generate 'petalloid' patterns. These repetitive bars, as well as straight parallel lines, also induce motion at right angles to the contours and a prominent motion after-effect (MacKay 1957; Georgeson 1985). A similar mechanism is almost certainly involved in the Leviant illusion studied by Zeki *et al.* (1993). This illusion (see, for example, figure 13a) generates apparent motion within the annulus which is much more rapid than that seen in the omega effect. Furthermore it is shimmery and incoherent, it disappears near the fovea (cf. Georgeson 1980; Rose & Lowe 1982), and it is strongly visible also in linear alleyways (figure 13b; MacKay 1957). To study this further, we presented figure 13a and figure 13b to ten observers and asked them to compare the magnitudes of the effects generated in each display. (The figures were printed on white paper; figure 13a was 135 mm in diameter and figure 13b was 175 mm long.) For three observers, the motion was equally salient in the circle and the straight line; for four observers the motion was stronger in the circle; two observers reported the converse; and another could perceive no motion in either
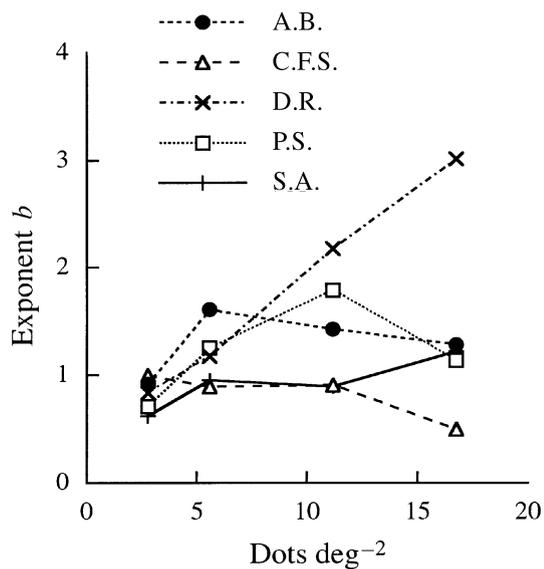
Figure 10. Exponents *b* of the best fitting equations for the data in figure 9.

figure. Quantitative comparisons are however difficult because the retinal eccentricity of the circle and the line cannot be equalized.

Placing a transparency of figure 13*a* over television snow creates a strong impression of rapid rotational motion in the snow behind the radial lines (as first observed by MacKay (1957); a similar but less prominent motion was seen with figure 1*c*). A transparency of figure 13*b* also generates motion of the snow at right angles to the fine lines (MacKay 1957). In an attempt to test whether there are two separate mechanisms at work, we changed the black annulus in figure 13*a* to a transparent one, to give us figure 13*c*. The question was whether the motion of the snow in the annulus would be at the normal speed for the omega effect with the same sized annulus (figure 13*d*) or would it be at the faster speed of the snow behind the radiating lines? The answer proved to be intermediate: the snow visible within the annulus rotated at the same speed as the snow behind the radial lines, and that speed was intermediate between the speed within the annulus alone (figure 13*d*) and that behind the radial lines with a black annulus (figure 13*a*).

The Roman-candle effect we observed at the apices of the triangles in our Celtic cross (figure 5*c*; also figure 1*c*) also consists of an incoherent streaming, and may be a vector sum of the motions generated by the two straight edges forming the apex. Similar streaming is seen in the equivalent parts of figure 13*e* (together with slow but clearly coherent rotation around the concentric arcs, as described by MacKay (1957, 1961)).

Whether these percepts of incoherent motion stem from inhibition between pattern and motion detectors in area V1 (Georgeson 1985), from area V5 (Zeki *et al.* 1993) and/ or from optical factors (Gregory 1993) remains to be determined. Certainly, neither they nor the omega effect are destroyed by viewing through a pin-hole, and are thus unlike the shimmering of MacKay's complementary images (Gregory 1993).

(b) *What is the origin of the omega effect?*

Since MacKay's original observations, motion perception theory has advanced greatly. However, an explanation of the omega effect consistent with all the observations detailed here is not readily apparent. We have considered theories based on: (i) how the masks constrain the possible dot matchings; (ii) comparisons of the optic flow induced by self-motion and by object movement; (iii) knowledge of the receptive field properties in visual areas V1, V5/MT and MST; (iv) interactions among a network of low-level motion detectors (including loop arrangements that would respond selectively to rotational motion); and (v) how multiple velocity vectors can be grouped to give single or multiple planes of coherent apparent motion. We do not have space here to explain all our reasons for rejecting these, but the only promising explanation seems to us to involve interactions between motion detectors and mechanisms sensitive to pattern or form. One clue comes from the experiments described in §3*c*,ii in which the omega effect was found to emerge from a uniform field of dynamic noise whenever an annular region is differentiated from the background by a slight difference in brightness. Most significantly, when the annulus differs only slightly from the background, there is rivalry between the perception of purely random motion and the omega effect. This implies an interaction between rotational motion mechanisms and curved-contour detectors. Such detectors have previously been postulated in early vision (see, for example, Koenderink & Richards 1988; Wilson & Richards 1989; Dobbins *et al.* 1989; Versavel *et al.* 1990; Zetsche & Barth 1990; Dobson & Payne 1992; Wolfe *et al.* 1992; Noss 1994; cf. de Haan 1995; Kramer & Fahle 1996), but the most relevant evidence for our studies comes from the discovery of cells in the higher visual area V4 which respond selectively to bullseye patterns (Gallant *et al.* 1993). Given, then, that mutual antagonism exists between orthogonal motion and contour mechanisms, at least for straight edges (see, for example, Georgeson 1985), and between orthogonal motions (see, for example, Snowden 1989; Grunewald & Lankheet 1996), it thus seems reasonable to propose that curved-contour detectors would be associated positively with rotational motion along those contours.

In other words, the presence of dynamic noise will activate mechanisms that signal the existence of motion, while at the same time the presence of stationary contours in the display will be activating mechanisms that signal lack of motion. When these contours are curved, the conflict can be resolved by assuming the existence of a circular or spherical rotating object in the visual field. Even if only a semicircle or an arc is visible (see, for example, figures 2*b*, 3*a* and 6), the conclusion is still valid because rotating objects can be partly obscured by nearer objects. Segmentation of the visual scene must therefore precede attribution of motion to the texture on the surfaces of objects. This explanation is also consistent with the induction of omega motion by some polygonal stimuli (figure 2*a*), as these contain curvature in their low spatial frequencies, and the omega effect can be seen with only low spatial frequencies (figure 7*b*,*c*).

With straight-line stimuli, however, linear motion of the object along its long axis should be accompanied by visible motion of the ends of the object; yet this information was
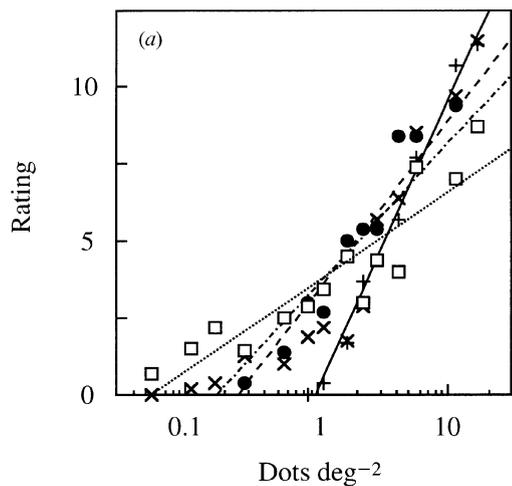
--●-- A.B. $y = 5.566 \log(x) + 3.310$, $r = 0.964$

--✕-- D.R. $y = 4.622 \log(x) + 3.506$, $r = 0.903$

......□...... P.S. $y = 2.905 \log(x) + 3.631$, $r = 0.913$

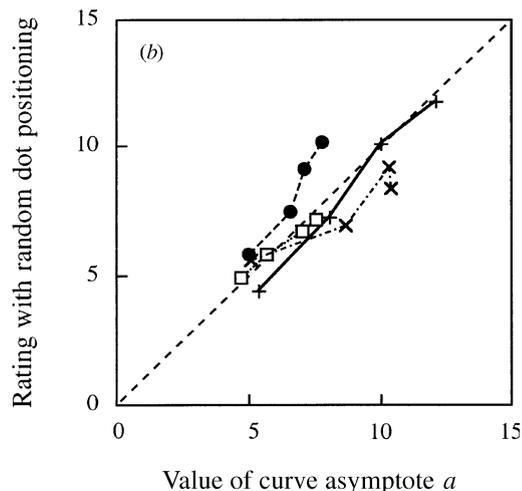—+— S.A. $y = 9.482 \log(x) + 0.158$, $r = 0.994$

- - - - Slope = 1

Figure 11. (*a*) Ratings assigned to displays where dots were randomly repositioned on every frame. Straight lines and equations show logarithmic fits for each of the four subjects who completed this part of the experiment. (*b*) Relation between the asymptotic values *a* (from figure 9) and the empirically measured ratings for dots randomly positioned on every frame (from figure 11*a*). A total of four dot densities were used and four subjects completed all experiments. The diagonal line of slope 1 indicates parity.

not present in our displays, hence weakening any sense of coherent motion with straight alleyways. This theory can easily be tested further: when we view a straight alleyway of TV snow with clearly demarcated ends (see, for example, figure 3*a*, rightmost column), synchronous movement of the ends in the same direction, at the same speed, causes the snow to appear to move coherently with the same velocity as the endings, i.e. in a manner consistent with the presence in the visual field of a moving textured
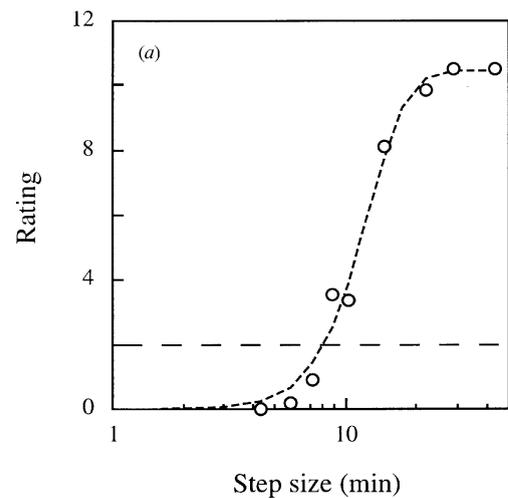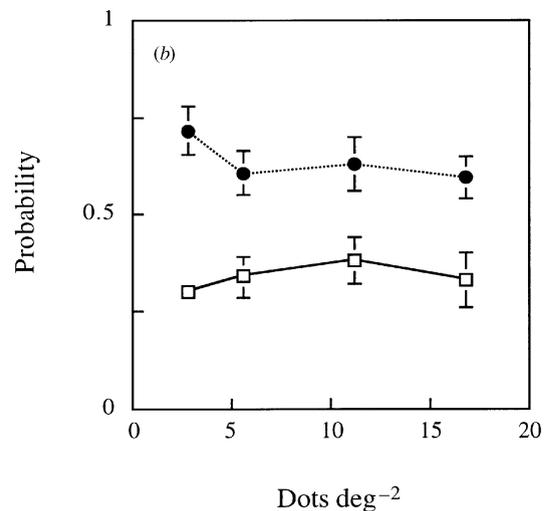

Figure 12. (*a*) Application of a criterion rating of '2' to one of the data sets from figure 9. The point where the horizontal line at rating '2' meets the fitted curve enables us to calculate the value of $s^2 d$ predicted to elicit that rating. (*b*) Probabilities of false matches occurring, for two different criterion ratings, as a function of dot density. Each point is the mean (± 1 s.e.) for five observers.

bar or rod. Similar motion 'capture' was first described by MacKay (1961, 1965) and has interested researchers ever since (see, for example, Ramachandran 1985; Murakami & Shimojo 1993; Zhang *et al.* 1993). This phenomenon may now be explicable under the present theory.

In the real world, a ball or sphere spinning on an axis that does not pass through the eye will possess mainly curved trajectories of motion in its surface texture. A stationary circular outline cannot therefore be associated strongly with any particular direction or type of motion. So if dynamic random noise is seen through a circular aperture, the direction of motion expected is ambiguous in all directions and there is no reason to suppose any one direction should become dominant; indeed no coherent motion percept emerges with such a display (cf. figure 2*c*, left column). The coherent omega effect only
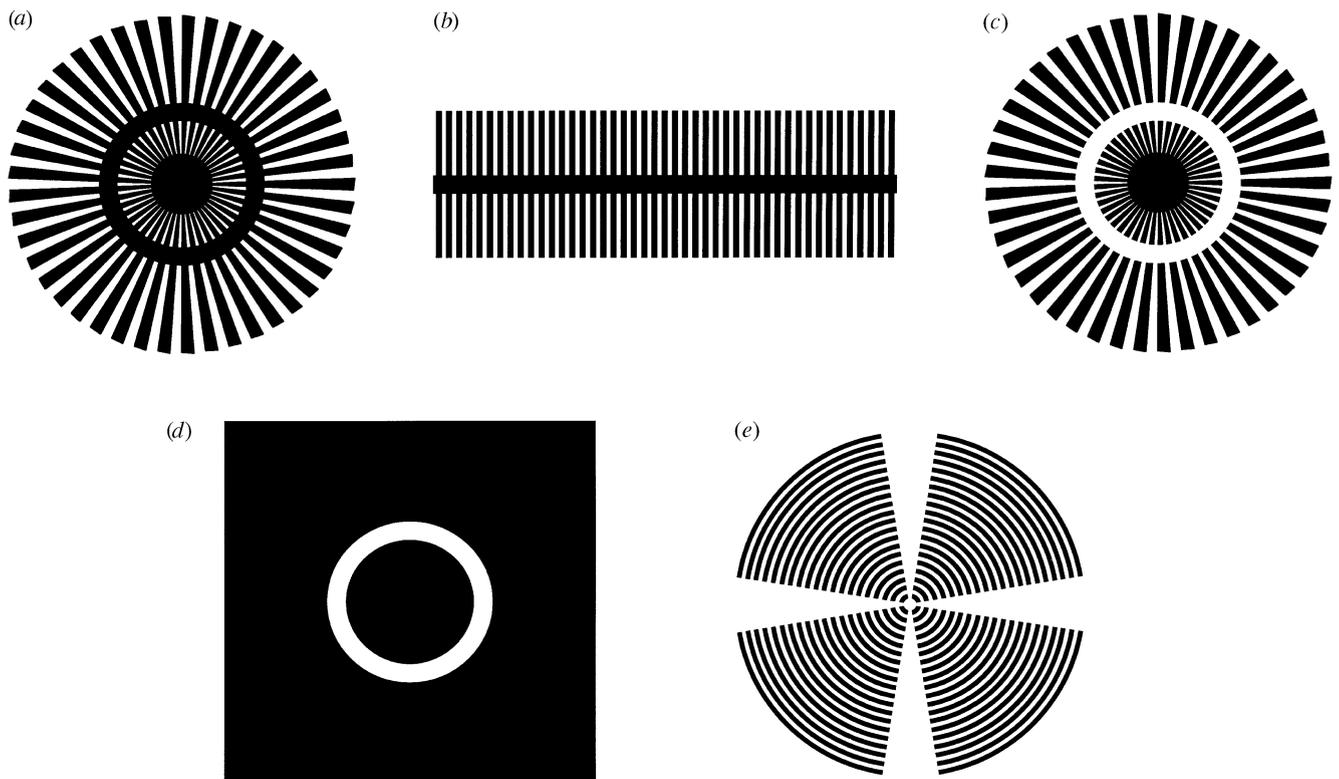
Figure 13. (*a*) A variant of the Leviant illusion (Zeki *et al.* 1993), in which rapid motion around the black annulus is perceived, especially away from the point of direct regard. (*b*) The straight counterpart of the Leviant illusion. The horizontal black line has the same width and total length as the black annulus in part *a*, and the spacing of the narrow black and white repeating bars is similar at the point where it meets them. (*c*) The mask used to compare the Leviant illusion and the omega effect by placing it over television snow. (*d*) An annulus the same size as those in *a* and *c* that was used as a control. (*e*) A pattern mask where the contours are orthogonal to those in *c*. Inspection of part *d*, or of a blank sheet of paper, after a few seconds of fixating parts *a*, *b*, or *c* reveals a motion after-effect that resembles the incoherent streaming seen when those patterns are used to mask television snow (MacKay 1957, 1961); fixation of part *e* however generates a much stronger radial after-effect than is obtained from masked snow, where omega rotation competes with the radial percept.

emerges with an inner, as well as an outer, stationary curved contour; such a configuration would only be generated if a solid three-dimensional annulus were spinning about the axis passing through its centre and the observer's eye.

(c) *Dot matching*

What is the clue that enables the visual system to decide if it is looking at many small moving objects or at the texture on a single surface? In other words, when do we see independent phi movements and when do we see omega motion? The data in § 3c,ii support the idea that the probability of false matches in the display predicts the occurrence and strength of the omega effect. The most obvious difference between the local vectors in figure 8*a*, and those in figure 8*b*, is that when false matches occur (figure 8*b*) the local vectors cover a range of velocities (while maintaining random directions). Indeed, if we try to follow any individual dot, its path will still take random directions but its velocity will also alter irregularly over time. It may well be this latter factor that is crucial in preventing individual dot tracking and in initiating global texture analysis. How a particular velocity is then assigned to a texture surface remains for future analysis (cf. Mingolla *et al.* 1992; Watamaniuk & Duchon 1992; Qian *et al.* 1994; Bravo & Watamaniuk 1995; Todd & Norman 1995).

REFERENCES

Blake, R., O'Shea, R. P. & Mueller, T. J. 1992 Spatial zones of binocular rivalry in central and peripheral vision. *Vis. Neurosci.* **8**, 469–478.

Bravo, M. J. & Watamaniuk, S. N. J. 1995 Evidence for two speed signals: a coarse local signal for segregation and a precise global signal for discrimination. *Vis. Res.* **35**, 1691–1697.

de Haan, E. 1995 Edge-curvature discriminability argues against explicit curvature detectors. *J. Opt. Soc. Am.* A **12**, 202–213.

Diggle, P. J. 1983 *Statistical analysis of spatial point patterns.* London: Academic Press.

Dobbins, A., Zucker, S. W. & Cynader, M. S. 1989 Endstopping and curvature. *Vis. Res.* **29**, 1371–1387.

Dobson, V. G. & Payne, R. D. W. 1992 Representing objects by location-implicit coding of their visual shape constancies. *Ir. J. Psychol.* **13**, 393–411.

Gallant, J. L., Braun, J. & Van Essen, D. C. 1993 Selectivity for polar, hyperbolic, and Cartesian gratings in macaque visual cortex. *Science* **259**, 100–103.

Georgeson, M. A. 1980 The perceived spatial frequency, contrast, and orientation of illusory gratings. *Perception* **9**, 695–712.

Georgeson, M. A. 1985 Inferring cortical organization from subjective visual patterns. In *Models of the visual cortex* (ed. D. Rose & V. G. Dobson), pp. 223–232. Chichester, UK: Wiley.

Gerrits, H. J. M. & Vendrik, A. J. H. 1970 Simultaneous contrast, filling-in process and information processing in man's visual system. *Exp. Brain Res.* **11**, 411–430.

Gregory, R. L. 1970 *The intelligent eye.* London: Weidenfeld & Nicolson.

Gregory, R. L. 1993 A comment: MacKay's rays shimmer due to accommodation changes. *Proc. R. Soc. Lond.* B **253**, 123.

Grossberg, S. & Mingolla, E. 1985 Neural dynamics of perceptual grouping: textures, boundaries and emergent segmentations. *Percep. Psychophysiol.* **38**, 141–171.

Grunewald, A. & Lankheet, M. J. M. 1996 The orthogonal motion aftereffect. *Perception* **25** (suppl.), 33.

Kanizsa, G. 1955 Margini quasi-percettivi in campi con stimulatione omogenea. *Rev. Psicol.* **49**, 7–30.

Koenderink, J. J. & Richards, W. 1988 Two-dimensional curvature operators. *J. Opt. Soc. Am.* A **5**, 1136–1141.

Kolers, P. A. 1972 *Aspects of motion perception.* Oxford: Pergamon Press.

Kramer, D. & Fahle, M. 1996 A simple mechanism for detecting low curvatures. *Vis. Res.* **36**, 1411–1419.

Lee, T. S. 1995 A Bayesian framework for understanding texture segmentation in the primary visual cortex. *Vis. Res.* **35**, 2643–2657.

Livingstone, M. S. & Hubel, D. H. 1987 Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *J. Neurosci.* **7**, 3416–3468.

MacKay, D. M. 1957 Moving visual images produced by regular stationary contours. *Nature* **180**, 849–859.

MacKay, D. M. 1961 Visual effects of non-redundant stimulation. *Nature* **192**, 739–740.

MacKay, D. M. 1965 Visual noise as a tool of research. *J. Gen. Psychol.* **72**, 181–197.

Mingolla, E., Todd, J. T. & Norman, J. F. 1992 The perception of globally coherent motion. *Vis. Res.* **32**, 1015–1031.

Murakami, I. & Shimojo, S. 1993 Motion capture changes to induced motion at higher luminance contrasts, smaller eccentricities, and larger inducer sizes. *Vis. Res.* **33**, 2091–2107.

Noss, R. S. 1994 Line intensity affects perceived shape. *Vis. Res.* **34**, 497–510.

Paradiso, M. A. & Nakayama, K. 1991 Brightness perception and filling-in. *Vis. Res.* **31**, 1221–1236.

Qian, N., Andersen, R. A. & Adelson, E. H. 1994 Transparent motion perception as detection of unbalanced motion signals. III. Modeling. *J. Neurosci.* **14**, 7381–7392.

Ramachandran, V. S. 1985 Apparent motion of subjective surfaces. *Perception* **14**, 127–134.

Rock, I. 1983 *The logic of perception.* Cambridge, MA: MIT Press.

Rose, D. & Blake, R. 1995 Interactions between curved pattern and rotational motion mechanisms. *Invest. Ophthalmol. Vis. Sci.* **36**, S395.

Rose, D. & Lowe, I. 1982 Dynamics of adaptation to contrast. *Perception* **11**, 505–528.

Rose, D., Blake, R. & Steiner, V. 1994a *Perceived flow in dynamic random dot stimuli—the omega effect revisited.* Paper presented to the Experimental Psychology Society, London.

Rose, D., Blake, R. & Steiner, V. 1994b Coherent perceived flow in dynamic random noise. *Invest. Ophthalmol. Vis. Sci.* **35**, S1269.

Sclar, G., Maunsell, J. H. R. & Lennie, P. 1990 Coding of image contrast in central visual pathways of the macaque monkey. *Vis. Res.* **30**, 1–10.

Shapley, R., Kaplan, E. & Soodak, R. 1981 Spatial summation and contrast sensitivity of X and Y cells in the lateral geniculate nucleus of the macaque. *Nature* **292**, 543–545.

Snowden, R. J. 1989 Motions in orthogonal directions are mutually suppressive. *J. Opt. Soc. Am.* A **6**, 1096–1101.

Stevens, S. S. (ed.) 1951 Mathematics, measurement, and psychophysics. In *Handbook of experimental psychology*, pp. 1–49. New York: Wiley.

Todd, J. T. & Norman, J. F. 1995 The effects of spatiotemporal integration on maximum displacement thresholds for the detection of coherent motion. *Vis. Res.* **35**, 2287–2302.

Tootell, R. B. H., Reppas, J. B., Kwong, K. K., Malach, R., Born, R. T., Brady, T. J., Rosen, B. R. & Belliveau, J. W. 1995 Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *J. Neurosci.* **15**, 3215–3230.

Versavel, M., Orban, G. A. & Lagae, L. 1990 Responses of visual cortical neurons to curved stimuli and chevrons. *Vis. Res.* **30**, 235–248.

Walls, G. L. 1954 The filling-in process. *Am. J. Opt. Arch. Am. Acad. Opt.* **31**, 329–341.

Watamaniuk, S. N. J. & Duchon, A. 1992 The human visual system averages speed information. *Vis. Res.* **32**, 931–941.

Wertheimer, M. 1912 Experimentelle Studien über das Sehen von Bewegung. *Z. Psychol.* **61**, 161–265.

Williams, D. W. & Sekuler, R. 1984 Coherent global motion percepts from stochastic local motions. *Vis. Res.* **24**, 55–62.

Wilson, H. R. & Richards, W. A. 1989 Mechanisms of contour curvature discrimination. *J. Opt. Soc. Am.* A **6**, 106–115.

Wolfe, J. M., Yee, A. & Friedman-Hill, S. R. 1992 Curvature is a basic feature for visual search tasks. *Perception* **21**, 465–480.

Yang, Y., Rose, D. & Blake, R. 1992 On the variety of percepts associated with dichoptic viewing of dissimilar monocular stimuli. *Perception* **21**, 47–62.

Zeki, S., Watson, J. D. G. & Frackowiak, R. S. J. 1993 Going beyond the information given: the relation of illusory visual motion to brain activity. *Proc. R. Soc. Lond.* B **252**, 215–222.

Zetsche, C. & Barth, E. 1990 Fundamental limits of linear filters in the visual processing of two dimensional signals. *Vis. Res.* **30**, 1111–1117.

Zhang, J., Yeh, S. & De Valois, K. K. 1993 Motion contrast and motion integration. *Vis. Res.* **33**, 2721–2732.