The Artful Eye

EDITED BY

RICHARD GREGORY

Department of Psychology, University of Bristol

JOHN HARRIS

Department of Psychology, University of Reading

PRISCILLA HEARD

Department of Psychology, University of Bristol

DAVID ROSE

Department of Psychology, University of Surrey

Oxford New York Tokyo

OXFORD UNIVERSITY PRESS

1995

10 At the edge of movement

STUART ANSTIS AND V. S. RAMACHANDRAN

The edges of objects in the real world are revealed by abrupt changes in retinal images of several kinds: brightness, texture, motion, colour, depth, and so on. The two-dimensional artist cannot depict edges in motion or stereopsis, but plenty of cues are still available to his or her palette. Here we describe what happens when one throws away brightness, colour, and everything else, until only motion is left. Some things get worse — edges cannot be localized accurately, but new and interesting perceptual effects emerge. Now that the computer is becoming an artist's tool, there may be new areas to be opened up using kinetically revealed objects, or objects visible only by colour, or only by depth and so on.

For objects to be visible they must differ from their surround in some way. There is usually a brightness difference, as in the white square in Fig. 10.1(a). But recently there has been much interest in equiluminous objects (Livingstone and Hubel 1987) which differ from their surround along some non-luminous visual dimension, such as stereo depth (Fig. 10.1(d)), or texture (Fig. 10.1(b) and (c)), where the square and the surround have different textures but share the same space-averaged brightness. In particular, motion can make an object stand out from its surround. Figure 10.1(e) shows a motion-defined square on a surround of stationary random dots, like a sheet of sandpaper with a small foreground square of sandpaper moving in front of it. As a result the square is seen as moving. Thus motion itself can reveal the objects, even when there is no difference in brightness. Of course the edges of the square move with the texture, because the whole thing is moving. But one could equally well have a large sheet of stationary sandpaper with a small square hole cut in it, through which is seen part of a large background sheet of moving sandpaper (Fig. 10.1(f)). This gives a stationary window of moving texture.

We have studied these stationary motion-defined windows that contain drifting dots (Ramachandran and Anstis 1987, 1990; Anstis 1989). All stimuli were displayed on a computer-controlled TV

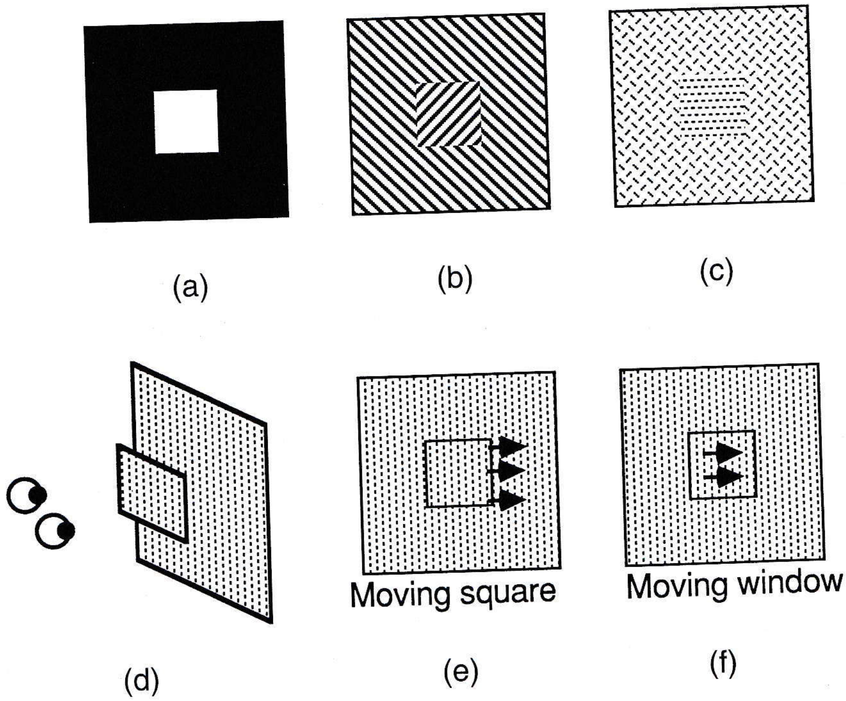

Fig. 10.1. Ways to see a square. (a) brightness. (b,c) texture. (d) stereoscopic depth. (e,f) motion. In (e) the whole square drifted to the right (arrows), including the edges. In (f) the dot texture drifted to the right behind a window with stationary edges. All the stimuli we used were windows like (f).

screen (Anstis 1986; Anstis and Paradiso 1989). We used three types of random-dot textures. The random dots could be stationary, or drifting together (coherently) in one direction or else jumping around incoherently in all directions, rather like the 'snow storm' on a detuned TV receiver (random dynamic noise). If any two of these textures are butted together they produce a clearly visible edge that is revealed only by motion. Notice that the motion edge does not exist in any single snapshot, but only as a correlation between successive frames of a moving picture over time. There is no edge revealed by brightness, because the average brightness of the dots is the same on both sides of the contour. If the dots are made a little darker on one side of the border, we have an edge that is revealed both by motion and by brightness.

We have discovered several illusory effects produced by such windows of drifting texture (Fig. 10.2). First, the position of each window appears *displaced* in the direction of the drifting dots that it contains. This is true only when the textures inside and outside the windows have the same mean brightness, so that the edges are revealed only by motion. When we add the brightness edges the apparent displacement is no longer seen. Second, the window also appears to *move along* with the dots even though it is stationary — an example of motion capture (Ramachandran 1987). Third, if the stationary window is embedded in dynamic visual noise and viewed

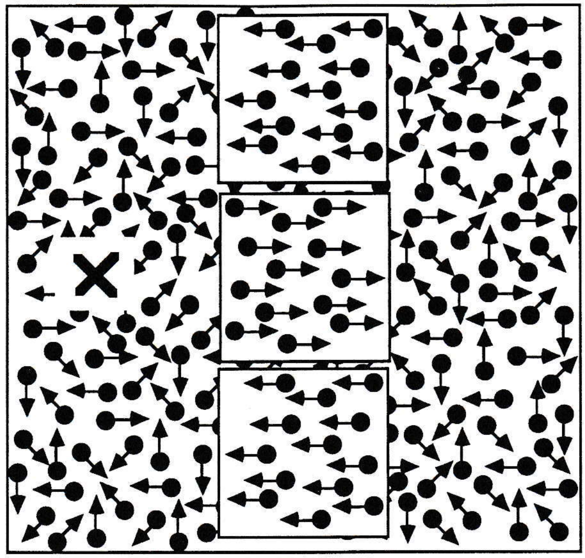

Fig. 10.2. Three stationary windows were filled with random dots that drifted to the right in the central window and to the left in the top and bottom windows. The surround was filled with randomly twinkling dots. Result: the central window appeared subjectively displaced to the right (illusory displacement). Also, if the subject fixated the cross then the windows could disappear within 4–8 s.

peripherally during steady fixation, it rapidly vanishes from sight, becoming invisible within about 5 s and reappearing only when the eyes are moved (Anstis 1989; Ramachandran and Anstis 1990). This rather alarming artificial blindspot or 'galloping glaucoma' is probably caused by the twinkling background rather than by motion within the window, since uniform grey patches embedded in visual noise also fade rapidly when viewed peripherally (Spillman and Kurtenbach 1990; Ramachandran and Gregory 1991). We have noticed that small coloured spots also fade quickly in peripheral vision, even when they have an average brightness different from that of the dynamic noise. Whereas 'Troxler' fading, as described by Pirenne (1967), generally works best for blurred or very stable edges, the twinkling noise in our displays constantly refreshed the edge, yet the perceptual vanishing was very rapid.

In this chapter we extend these studies to stationary windows that contain not drifting dots, but drifting stripes. We shall describe three perceptual phenomena:

(1) illusory offsets;

(2) motion segregation — seeing that objects are different by their different movements;

(3) the aperture problem — the fact that the shape of the frame surrounding movement affects the perceived direction of the movement.

It has long been known that steadily fixated edges perceptually fade (the Troxler effect — see Pirenne 1967). However, this effect has been thought to work best for blurred stationary edges given by luminance change. In our displays, the twinkling noise constantly redefined the edge, yet the perceptual fading was very rapid.

The illusory offsets obtained with the windows filled with drifting stripes are similar to those obtained with random-dot textures. However, rapid disappearance in peripheral vision was not observed with stripe-filled windows and the motion segregation and aperture effects are peculiar to windows filled with stripes.

Illusory offsets

Figure 10.3(a) shows, on a uniform grey surround, a stationary window within which stripes drift to the right. The change of brightness between adjacent black and white stripes was not abrupt but gradual — a graph of brightness against position would produce a sine wave. All windows in this paper were a whole even number of stripes wide, so that at all times a stripe appearing at one edge had the same brightness as the stripe that was disappearing at the other edge. There was almost always a brightness difference between the edge of the window and the grey surround, because at an instant when drifting black (or white) stripes reached the edges of the stationary window these stripes were darker (or lighter) than the grey surround. Only on the way up and on the way down did the grey of the stripes match the grey of the surround, and only at these brief instants was there no brightness-defined edge.

There are three ways to get rid of these brightness edges.

1. The brightness of the uniform surround can be swept up and down smoothly (sinusoidally) over time, in such a way that the surround brightness always matches the stripes at the edges of the window (Fig. 10.3(a)).

2. The mean brightness of the striped pattern can be varied, rather than that of the surround (Fig. 10.3(b)). When the stripes at the edge of the window were light, we lowered their average brightness until these stripes were mid-grey to match the unchanging surround. When the stripes at the edge of the window were dark, we raised their average brightness until these stripes were mid-grey to match the surround. This method is used when many windows, each with stripes in different positions, were viewed against a common surround. Thus, at one instant one window might have

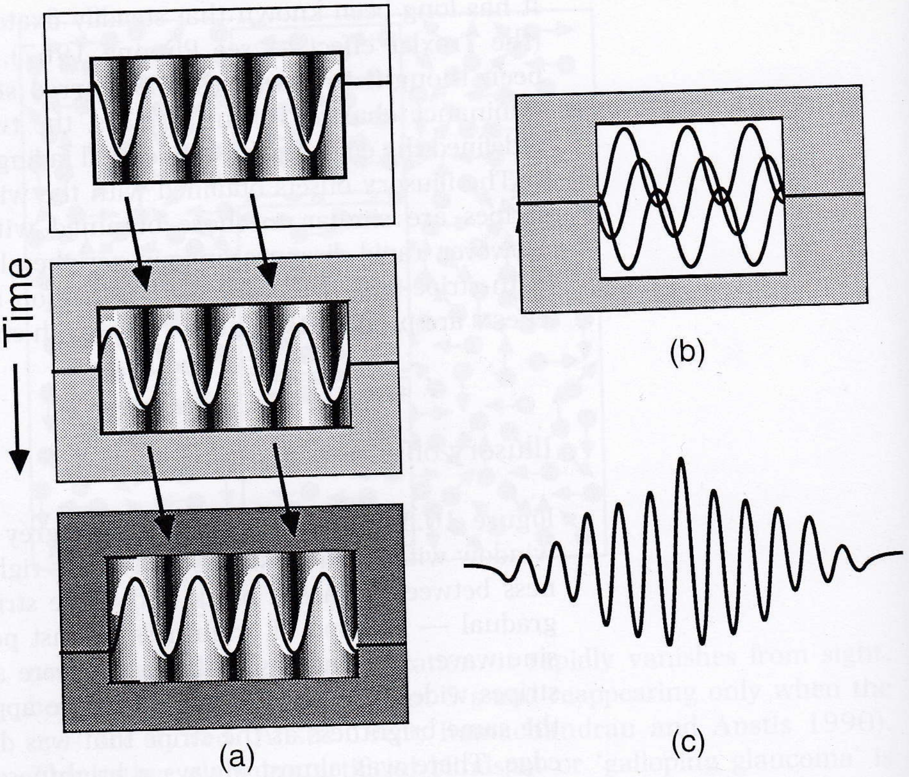

Fig. 10.3. Three methods for removing brightness edges from stationary windows filled with drifting gratings (repetitive patterns of black and white stripes). (a) The brightness of the grey surround was varied smoothly up and down over time in step with the stripes as these reached the sides of the windows. Surround was made white when white parts of the bars were at the edges of the windows (top picture), mid-grey when the grey in the edges of the bars was at the window edges (middle picture), and black when the black troughs of the bars were at the window edges (bottom picture). Brightness profiles of the grating and surround are shown on these examples of the stimuli. (b) The average brightness of the stripes, not of the surround, was varied. Three brightness profiles of the stripes at successive times are shown within the window. In the leftmost, upper profile, dark bars were at the window edges and the mean brightness of the grating was raised to make these dark bars match the mid-grey surround. The middle trace shows the mid-grey of the bar edges at the window edges. In the rightmost, bottom trace, light bars were at the window edges and the average brightness of the bars was lowered to make these light bars match the mid-grey surround. For clarity, the bars themselves are omitted. (c) The contrast of the stripes differs with their position within the window. As stripes enter the window (say from the left) their contrast is zero. Contrast gradually increases until the stripes lie in the middle of the window and then gradually decreases to zero as they reach the right-hand edge of the window (after De Valois and De Valois 1990).

dark stripes at its edges while another had light stripes, but both were matched to the common mid-grey surround.

3. De Valois and De Valois (1990, 1991) have independently discovered the illusory offset phenomenon. The blurry edges of the

windows within which the stripes were displayed effectively removed the brightness edges. The contrast of the stripes was not the same throughout the window, but was high in the middle and tapered off to zero at the edges of the window (Fig. 10.3(c)).

De Valois and De Valois found, as we did, that the apparent position of the window was shifted in the direction of motion of the stripes. This effect became greater as the stripes were moved out of central vision, reaching half a degree at 8° eccentricity.

All three methods produced the same results: with the brightness edges removed, the windows containing stripes showed the same illusory displacement as did the windows containing drifting random-dot textures. In Fig. 10.4(a) three stationary windows containing drifting stripes are lined up one above the other. The stripes in the middle window moved to the right and the stripes in the top and bottom windows moved to the left. If the brightness edges were removed by any of the three methods just described, a strong illusory offset was seen, in which the middle window appeared to be displaced to the right by up to half a degree of visual angle.

In two other conditions in which the brightness edges were not removed (not illustrated) the three windows were seen as being lined up vertically, as indeed they were. In one control condition the surround was simply a steady grey, and in another its brightness changed up and down smoothly, but out of step, not in step, with the varying brightness of the inner edges of the window. This condition kept the brightness edges intact and abolished the illusory offset.

The illusory offsets could produce an impression of stereoscopic depth, just as physical offsets can. If the left eye saw the stimulus while the right eye saw its mirror image (Fig. 10.4(b)), then the binocularly fused central window appeared to lie nearer in depth than the top and bottom windows (Fig. 10.4(c)). This is just what would happen if physical offsets were introduced (as in the right-hand panels of Fig. 10.4(b)).

What is true for position is true for size. In Fig. 10.4(d) the windows of stripes were wrapped around into two disks or bull's-eyes, with expanding concentric rings in one of them and contracting rings in the other. When the brightness edges were removed by our flicker technique, these edges were revealed purely by motion. Result: the expanding rings within one round window made the window itself appear to expand, so that it looked about 10 per cent larger than the window containing the contracting rings.

Why should a stationary window appear to be displaced just because it contains a drifting pattern? A motion-signalled edge must be reported by motion sensors, since no other sensors (for brightness, colour, etc.) are stimulated by it. We conjecture that motion

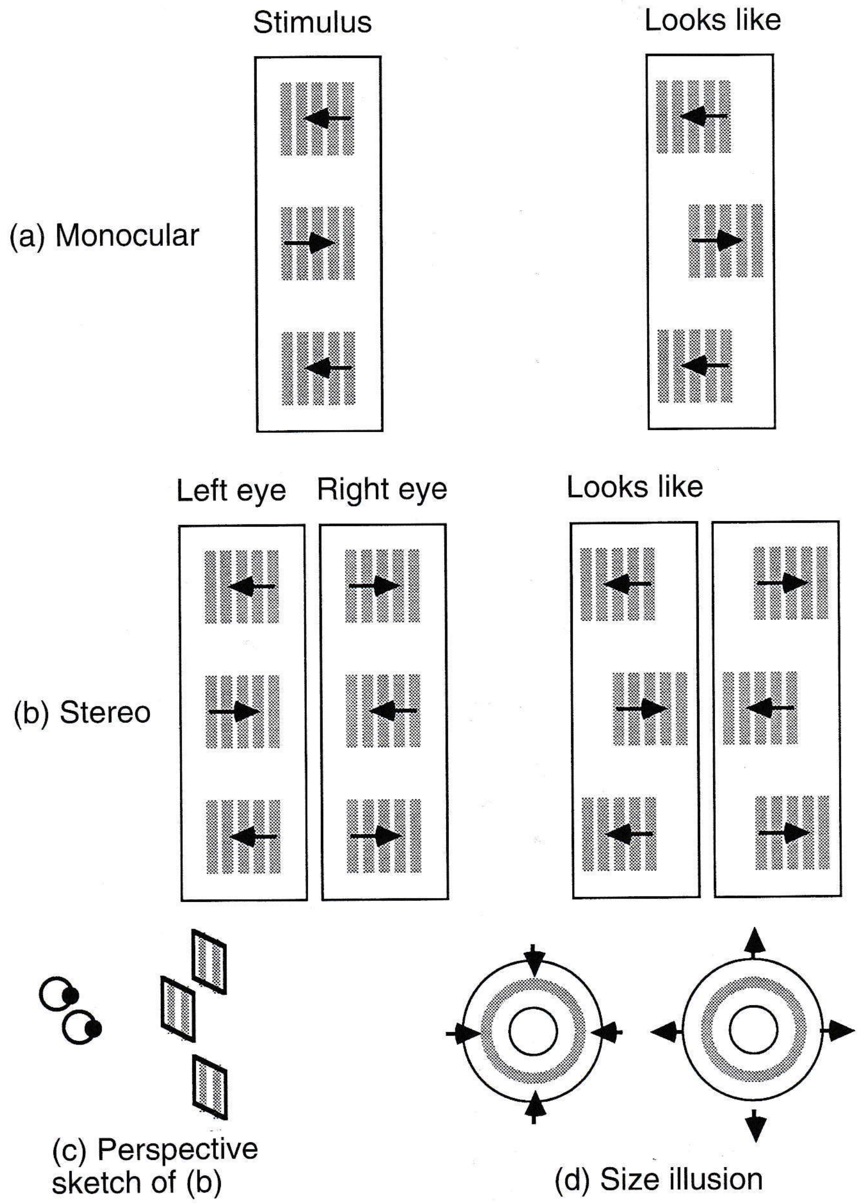

Fig. 10.4. Same as Fig. 10.2., but with gratings instead of random dots. (a) Three stationary windows, aligned vertically, contained drifting stripes. Top and bottom stripes drifted to the left, middle stripes drifted to the right. There were *no* brightness edges, because the surround brightness varied with the window edges (not shown). Result: middle window appeared to be displaced to the right. (b) Left eye saw same as (a), right eye saw its mirror image. Result: in the binocularly-fused image the central window appeared nearer than the others. (c) Perspective sketch of the apparent depth seen in (b). (d) Size illusion in annular windows that contained expanding or contracting circular stripes.

sensors systematically misreport the position of such an edge, assigning to it an offset in the sensor's preferred direction. This offset would help to anticipate the future location of the edge and would compensate for the actual displacement of the object that would occur during the inevitable neural delays of visual processing. Thus, the perceived position is not the position of the stimulus when it actually triggers the sensor, but its *anticipated* position by the time the signal must interact, at some more central neural site, with other sensory or motor processes. This offset played a role in our experiments because signals from directionally selective units are used to localize a kinetic edge and the edge would therefore appear substantially displaced. Patrick Cavanagh (personal communication) and De Valois and De Valois (1990, 1991) have made essentially the same suggestion. When there is also a brightness contrast across the kinetic edge, however, sensors which are stimulated by a stationary luminance change dominate the location of an edge and give accurate information about its position.

In summary, we conjecture that brightness discontinuities give accurate local position information, but motion discontinuities do not. Motion sensors 'aim off' in the direction of the motion to allow for neural processing times, thus leading to illusory displacements. Motion sensors may also respond to rather large areas of the stimulus and so give only crude location information.

Thus, there seem to be two separate but related perceptual effects with kinetic (motion-defined) edges — an apparent offset underlying the perceived offset and also an apparent drift underlying motion segregation of kinetic windows. Both of these phenomena disappear when there is also a brightness contrast across the kinetic edge, because brightness pathways supply signals about position which countermand those of motion to the effect that the edge is stationary and not displaced. Clearly, not all edge cues are created equal. Brightness-edge signals are stronger than motion-edge signals, so they win out when put into competition and they countermand the motion-based illusions. However, colour is probably no stronger than motion, since the illusory offset and drift still occur if the spots or stripes on each side of the kinetic edge are of different colours but of the same brightness. Indeed, a colour border can actually appear to be dragged along by brightness-defined spots (Ramachandran 1987).

Motion segregation

The way in which motion sensors can divide the visual array into regions is readily demonstrated with a field of sparse, coarse spots which move to the right, while a subset of sparse spots arranged in a T move in unison to the left (Fig. 10.5(a)). Subjects can immediately read the T which is revealed by common motion (Fig. 10.5(b)). The segregation fades away as soon as the motion is stopped. We now report that motion segregation can be supported by stationary windows containing drifting gratings — but only if static brightness edges are removed.

Figure 10.5(c) shows an array of striped blobs, in other words stationary windows filled with drifting stripes. Eight of these windows, arranged in a T, contain stripes that drifted to the left, while all the

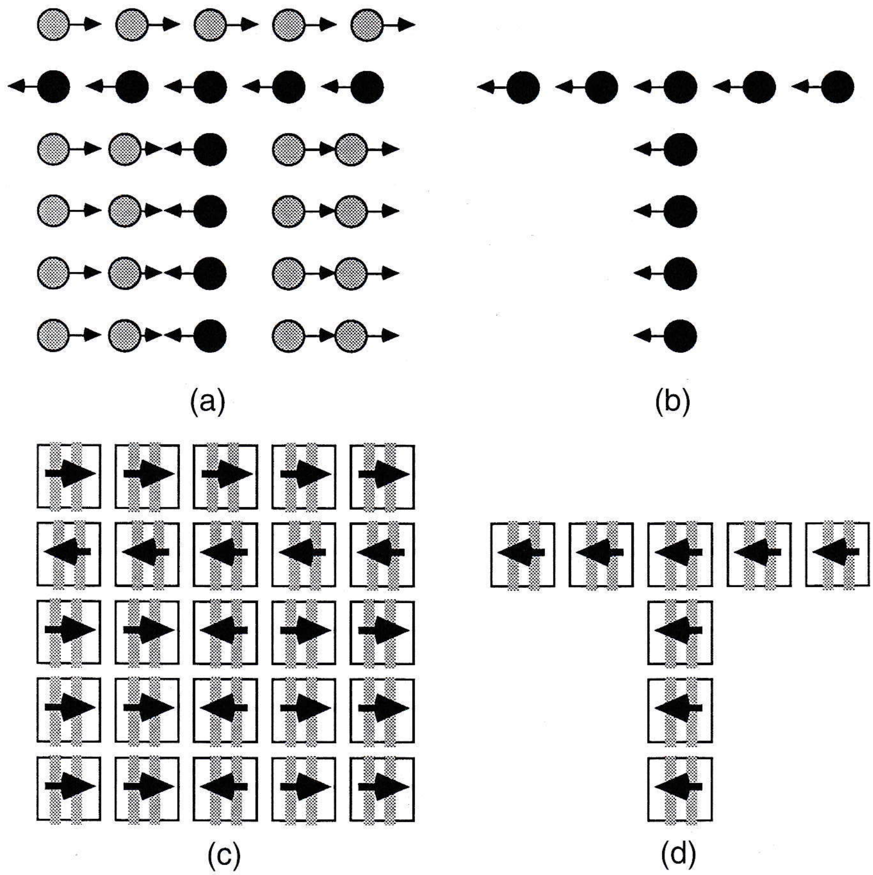

Fig. 10.5. Motion segregation. (a) Grey dots moved to the right, black dots to the left. Actually all dots were black; some are made grey here for exposition only. Result: the T-shaped region of leftward dots detached itself from the surround (b) in immediate, pre-attentive motion segregation. (c) Stationary windows were filled with stripes that drifted to the left in a T-shaped subset of windows and to the right in the remainder. Result depended on the surround, which is not shown. If surround brightness varied simultaneously with that of the stripes at the window edges, giving no brightness edges, there was strong pre-attentive motion segregation, as in (d). If brightness edges were intact, there was little or no motion segregation.

remaining stripes drifted to the right. Incidentally, it was noticed that if the bars in all the windows had the same spatial relationship, say white bars appearing in all the windows at the same moment, all the windows appeared to pulsate together in brightness. This was felt undesirable, since it might affect the motion segregation for better or for worse, so it was removed by randomizing the spatial positions of the bars across windows (not shown). To accommodate this, the brightness of the windows was flickered by method 2 in order to get rid of brightness edges. When this was done an observer immediately saw the eight windows, revealed by their kinetic edges, 'pop out' in pre-attentive perceptual segregation and move to the left as a unit, spelling out a letter T (Fig. 10.5(d)). However, if the brightness edges were not removed the motion segregation collapsed. An observer could still find the leftward-moving blobs given time, but had to search for them with attentive scrutiny.

Why do the windows segregate when their edges are revealed only by motion, but not when we add brightness edges? Fig. 10.6 shows our conjecture that the stimulus of Fig. 10.5(c) is analysed in paral-

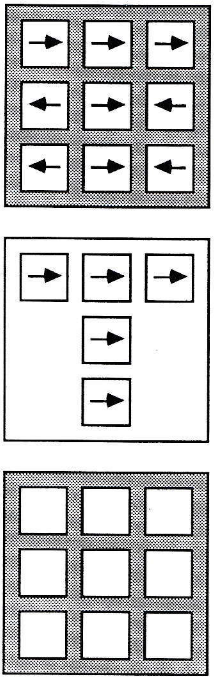

Fig. 10.6. Results shown in Fig. 10.5. can be explained by competition between brightness and motion signals. See text.

lel (that is, separately) for brightness and motion cues. There are nine little windows, with a T-shaped subset of five windows containing motion to the right. The motion cues signal that all the windows in the T-shaped set are moving to the right as a unit. However, if there are brightness edges around each window, brightness cues signal that all nine windows are 'the same'. When only motion edges are present we see the T-shaped subset. The motion signals identify the T as different, but the brightness signals identify all the windows as the same. This cue conflict is resolved in favour of brightness, which is a stronger cue than motion and motion segregation collapses. It is not clear why the observer seems to throw away the motion information. One might think that once the segregating information was extracted it would be conserved, but in fact it was not.

The aperture problem

Imagine a long vertical line moving sideways behind a circular window, its ends being invisible. If the line were moving down at the same time as it moved sideways, it would still look exactly the same. In fact there is a whole family of possible directions and speeds that would give the same retinal stimulus. The movement of drifting

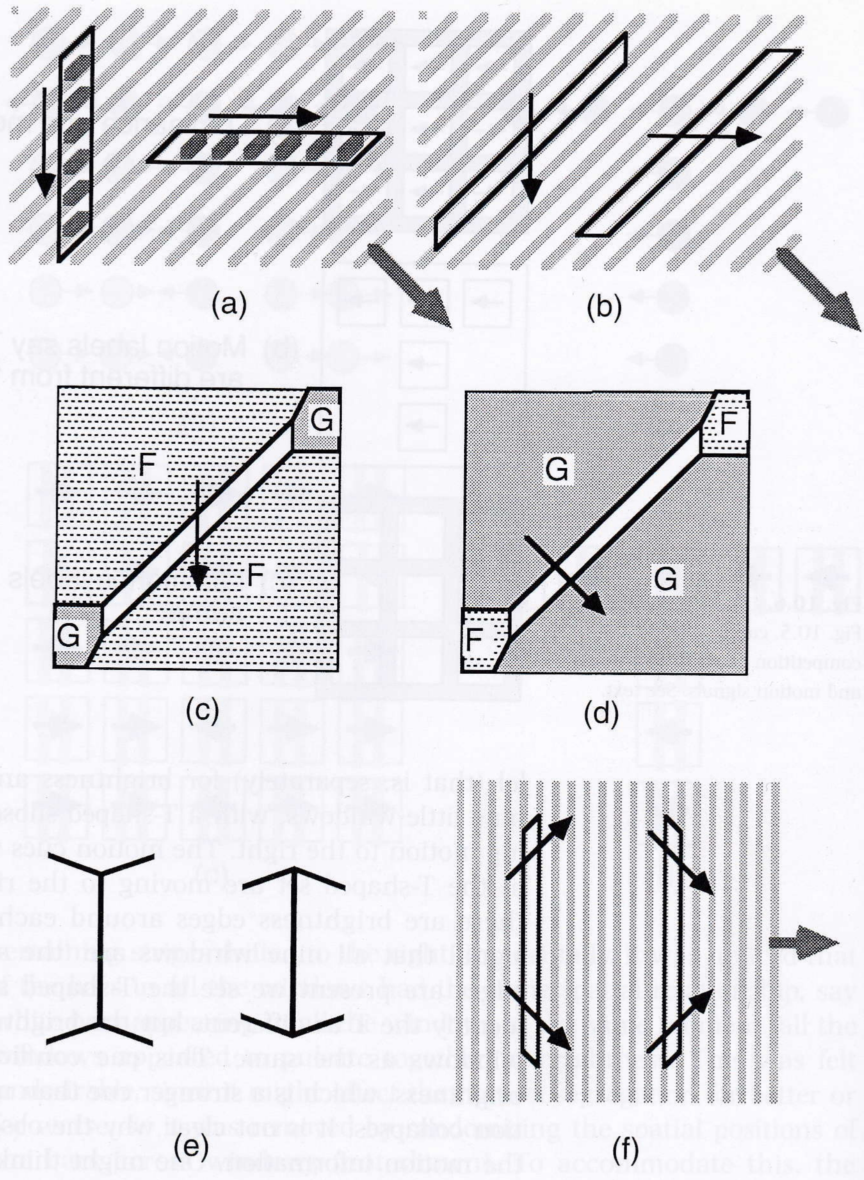

Fig. 10.7. Aperture problem with oblique gratings. (a) Oblique stripes (pale grey) drifted down to the right (thick grey arrow). The stripes were visible only through a vertical and a horizontal slit, through which they appeared to drift vertically and horizontally respectively (thin arrows). The oblique stripes shown in pale grey were actually hidden. (b) Aspect ratios of slits in (a) are radically changed, so that now both slits are oblique, but ends are cut vertically or horizontally. Result: the slits, and the gratings seen through them, appear to drift vertically or horizontally respectively (arrows) — but only if brightness edges are removed. (c) Brightness edges along sides are important. Flicker (F) near long sides keeps the vertical motion shown in (b), but (d) flicker near short ends (with grey (G) along sides) is ineffective. See text. (e) Conventional Müller–Lyer illusion. The line with the 'outgoing' fins appears much longer than the line with the 'ingoing' fins. (f) Dynamic Müller–Lyer illusion. Ends appear to drift outwards on left-hand slit, inwards on right-hand slit. Result: left-hand slit looks longer.

stripes is ambiguous (Adelson and Movshon 1982), because only the motion component at right angles to the stripes produces any visible change. In Fig. 10.7 oblique stripes drifting down to the right (thick grey arrows) were actually visible only through a vertical slit and a horizontal slit, through which they appeared to be drifting vertically or horizontally respectively (thin arrows in Fig. 10.7(a)). These perceived motions are consistent with the intersections that move along the edges of the slits (see Nakayama and Silverman 1988).

If you view drifting vertical stripes through a long thin vertical window, then not surprisingly you see the stripes moving horizon-

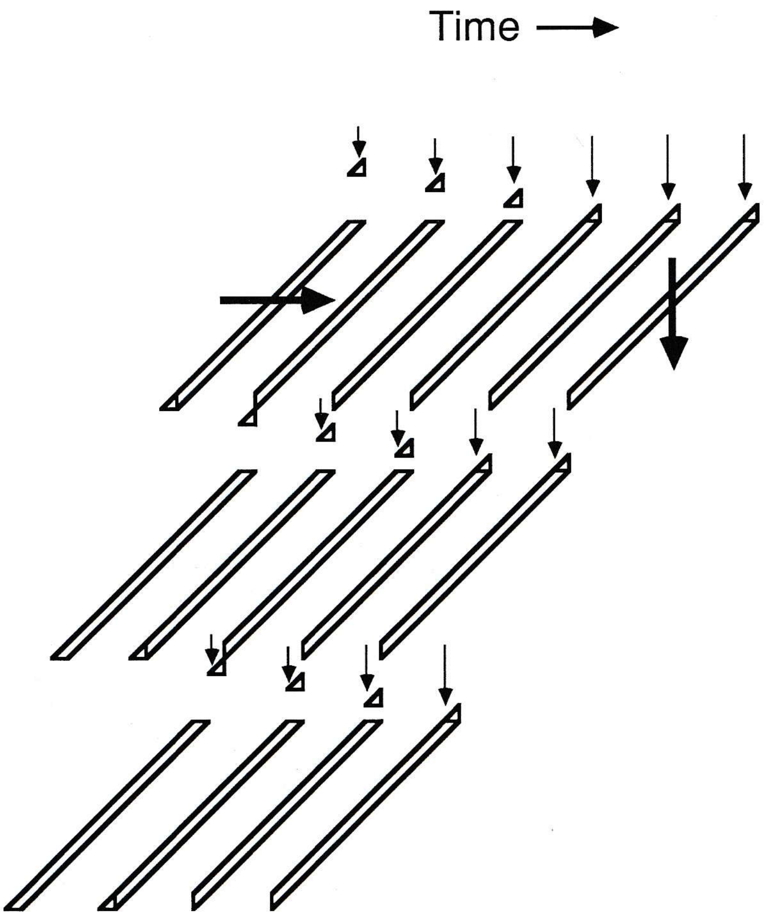

Fig. 10.8. A display of many parallel, oblique stationary slits, behind which oblique stripes (not shown) drift down to the right. Initially the short ends of the slits are cut horizontally, as shown on the left, so that the slits appear to drift horizontally (long horizontal arrow). Then a tiny triangular piece migrates slowly from the bottom of each slit to the top of the slit below. When they arrive they make the slit ends vertical, so that now the slits appear to drift vertically (long vertical arrow).

tally. But if you now make the top and bottom edges of the window slope obliquely, the stripes instantly appear to move obliquely up or down according to the slope of the tiny ends of the window. They appear to be dragged up or down by the tips. This effect is seen only if the brightness differences at the long edges of the window are eliminated. Furthermore, as in the case of the dot pattern, if the average brightness of the stripes is the same as the surround, the entire window appears to drift with the stripes. On the other hand if the brightness difference is introduced the window appears stationary. This phenomenon was demonstrated dramatically by filling the entire screen with oblique slit-shaped windows, all with the same orientation (Fig. 10.8).

The flicker in the surround was manipulated so that it bordered only the sides or only the ends of the slits. Result: it was the sides that were important. When the flicker was confined to the long sides of the slits, removing the luminous edges (Fig. 10.7(c)), the vertically cut slit still appeared to drift vertically and the horizontally cut slit