One-Sided Edge Assignment in Vision: 2. Part Decomposition, Shape Description, and Attention to Objects

Jon Driver and Gordon C. Baylis

Rubin originally noted that vision distinguishes figures from background, even within two-dimensional (2-D) displays. He and other Gestaltists[1] suggested on the basis of introspection that figures are seen as having a definite shape, whereas adjoining ground is seen as shapeless despite the dividing contour in common with the figure. In a previous article,[2] we presented recent performance studies showing that the dividing edge between figure and ground is automatically assigned to the figural shape, even when subjects attempt to judge just the dividing edge itself. As a result, recognition of the dividing edge is better if a test stimulus has this edge assigned in the same direction as in the preceding figural shape. We argued that one-sided edge assignment provides an efficient heuristic for deriving the likely three-dimensional (3-D) source of any 2-D image. In the present review, we explore the consequences of such edge assignment for shape representation and suggest an account for why only figures appear shaped.

Jon Driver is University Lecturer in Experimental Psychology at the University of Cambridge, and a Fellow of Gonville and Caius College, Cambridge. **Gordon C. Baylis** is Associate Professor of Cognitive Neuroscience at the Department of Psychology, University of South Carolina. Address correspondence to either author: Jon Driver, Department of Experimental Psychology, University of Cambridge, Downing St., Cambridge, CB2 3EB, U.K.; Gordon Baylis, Department of Psychology, University of South Carolina, Columbia, SC 29208.

SHAPE DESCRIPTION IN TERMS OF CONVEX PARTS

Hoffman and Richards[3] proposed a computational theory of shape description in which edge assignment is critical. Their starting point was from topology rather than psychology: When any two convex bodies intersect, negative minima of curvature (concavities) result at the points of intersection (see Figs. 1a and 1b). Accordingly, the location of such minima (in 2-D, points of sharp concavity) on the occluding contour of an object may be a reliable cue to the boundaries between its component parts. Because the convexity or concavity of each segment along an edge reverses with its figure-ground assignment (see Figs. 1c–1e), the description of convex parts in any given edge depends on the side to which it is assigned. This may explain why adjoining regions that are defined by a common dividing contour can have such different shapes, and thus how recognition of figures can be dissociable from recognition of their adjoining ground.[2]

We have recently confirmed the psychological reality of convex parts, and the influence of figure-ground segmentation on them.[4,5] On each critical trial, subjects were briefly presented with a rectangle that was centrally divided by a ran-

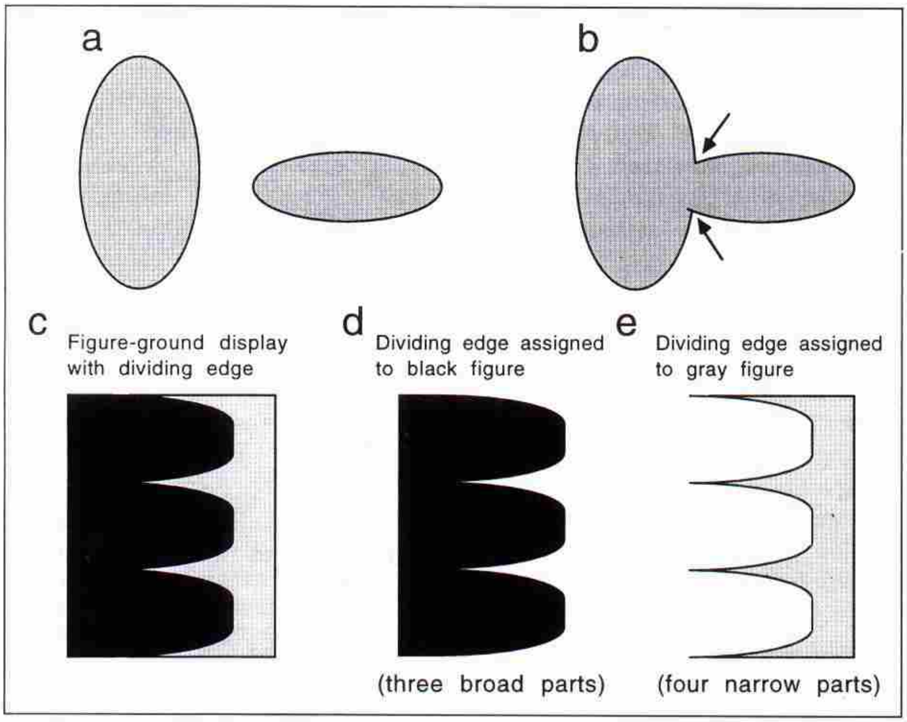

Fig. 1. An explanation of how shape description in terms of convexities and concavities relates to the possible genesis of complex visual objects, and depends on figure-ground assignment. When any two convex bodies (a) intersect (b), the points of intersection (see arrows) are concave. Component parts of a visual object may therefore correspond to convexities (separated by concavities) on the occluding contour of the object. The decomposition of this edge (c) into convex parts depends on which side is segmented as figure (d vs. e).

domly curved contour (Fig. 2a), different on each trial. The contour separated adjoining red and green regions of equivalent size and comparable brightness, to provide an ambiguous figure-ground display. Subjects were instructed to remember either the green shapes or the red shapes throughout the experiment, thus manipulating figure-ground assignment for the curved edges. The side of the relevant color was unpredictable for each brief display, so subjects looked consistently at the center of the display. Compliance with the color instructions was ensured by randomly intermingling the critical trials with biasing trials in which the display comprised separated red and green shapes, each with a unique curved contour. For these biasing trials also, subjects were asked to remember the shape in their instructed color. They could do so correctly only by following their color instructions.

Recognition of the shape in the instructed color was tested 500 ms after each initial exposure, by presenting subjects with a horizontal test slice (see Figs. 2b–2e) that always had adjoining red and green regions divided by a curved contour. For critical trials, with adjoining shapes in the initial display, the test contour was taken either from a section of that preceding display (thus requiring a positive response) or from a comparable section of a curve that had not been shown (requiring a negative response). For biasing trials, with separated, different red and green shapes in the initial display, the test contour was either a section from the initial shape in the instructed color (positive response) or a comparable section of the initial shape in the irrelevant color (negative response).

The factor of interest was whether the test slice showed a convex or concave region from the original contour. For the critical trials, this depended entirely on color instructions because any convexity in one color is necessarily a concavity for the adjoining area in the other color (Figs. 2b–2e). The results for these critical trials showed that recognition was faster and more accurate when the test slice showed a convexity in the instructed color for the original display, rather than a concavity. Those horizontal sections that were easiest for red subjects (i.e., sections with red convexities) were hardest for green subjects, and vice versa. This result demonstrates that the critical factor was the shape to which the dividing edge was assigned. Matching performance was better for convex parts than for the concavities that divided them, and the segments of contour that were recognized best reversed when the figure-ground assignment changed.

PART DECOMPOSITION AND SYMMETRY PERCEPTION

Description of figural shapes in terms of convex parts can explain a long-standing perceptual paradox first noted by Mach.[6] Detecting symmetry between the edges of a figure

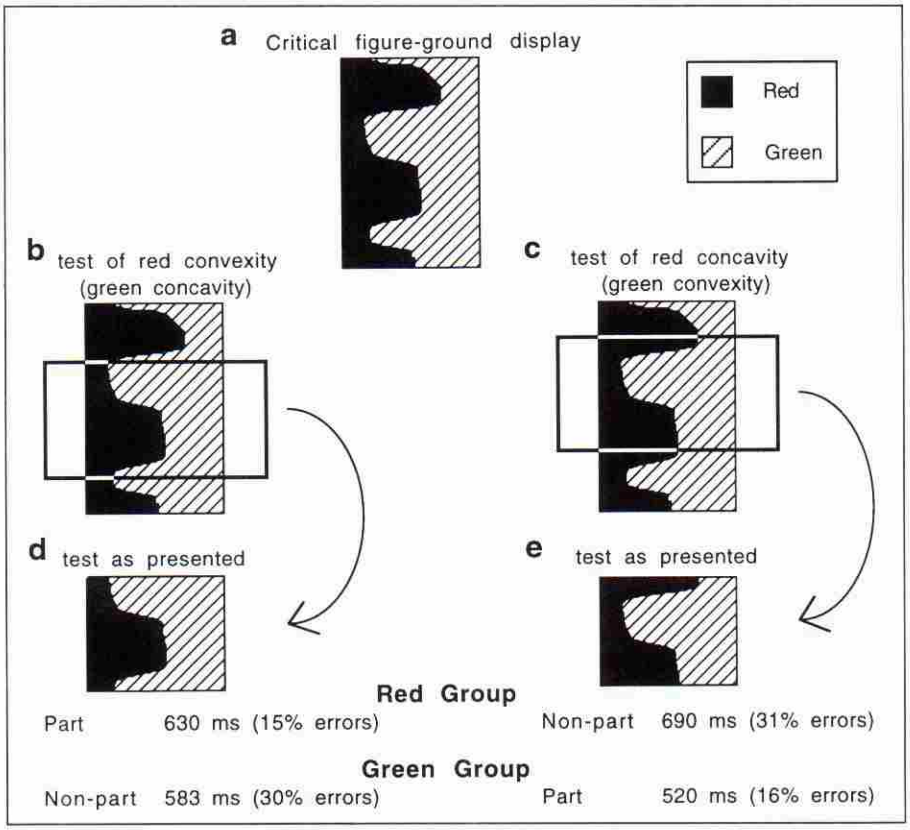

Fig. 2. Schematic illustration of stimuli for a study[5] described in the text. An ambiguous figure-ground display (a) shown at the start of the trial comprised a novel, randomly curved contour dividing equally sized red and green areas of comparable brightness. One group of subjects had to remember the green shape, the other the red shape. They were tested with a horizontal slice that matched (illustrated) or mismatched a section from the preceding display. Each test section was convex in one color (part), and concave in the other (nonpart). Thus, the section shown in (d), derived by sampling as shown in (b), was convex in red and therefore a part for the red group. Conversely, the section shown in (e), derived by sampling as shown in (c), was concave in red and therefore not a part for the red group. For the green group, the section shown in (d) was concave (not a part), whereas that shown in (e) was convex (a part). The mean reaction times and error rates in determining whether these sections matched part of the original display are shown. The crucial factor was whether the test section was convex in the instructed color for each group.

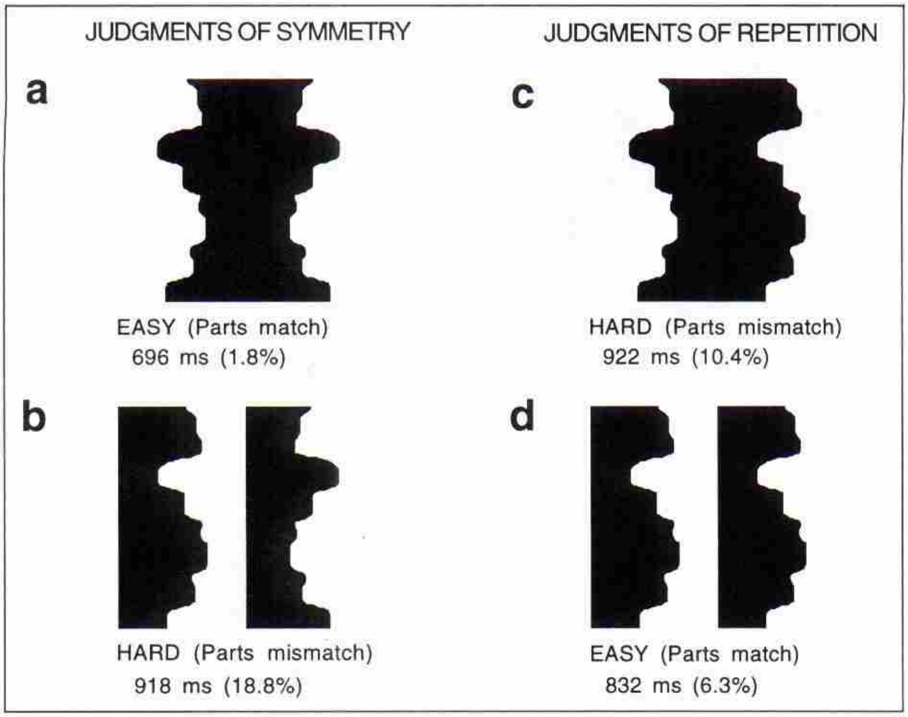

Fig. 3. Example displays from experiments[7] on detection of symmetry (a and b) or repetition (c and d). Judgments of symmetry in (a) and of repetition in (d) can be made by matching the part descriptions from the two sides of the display, whereas corresponding judgments in (b) and (c) cannot. The mean times and error rates for these judgments are shown.

(Fig. 3a) is much easier than detecting repetition of the edges of a figure (Fig. 3c). Perhaps this difference arises because the sides of a symmetrical figure have matching component parts (i.e., each convexity on one side has a corresponding convexity on the other; see Fig. 3a), whereas there is no such match in parts for repeated sides (each convexity on one side has a concavity on the other; see Fig. 3c).

To test this hypothesis,[7] we reversed the bottom-up (image-based) figure-ground assignment for just one of the two edges by altering the way they were colored (Figs. 3b and 3d). Bottom-up factors lead to the assignment of the filled regions (i.e., the parts shown as black in Fig. 3) as figure, with the intervening space seen as ground. In this revised format, convex parts matched for the repeated edges but not for the symmetrical edges. The striking result was that symmetry became harder to detect than repetition (see mean data in Fig. 3). Thus, the critical factor was whether convex parts matched, not simply whether symmetry or repetition was judged. Additional control conditions confirmed that the difficulty of judging symmetry in displays like Figure 3b could not be attributed merely to the presence of two figures rather than a single figure; such a hypothesis also does not explain the greater ease of repetition judgments with the format of Figure 3d versus Figure 3c.

These results support our previous proposal[2,4,8] that one-sided edge assignment is obligatory. If subjects had been able to judge the symmetry or repetition of just the curved edges themselves (as they were instructed), we should have found no difference between the displays in Figure 3a and 3b, or between Figure 3c and 3d, because the curved edges themselves were the same. Over hundreds of trials, subjects were unable to prevent an initial one-sided assignment of the curved edges to the figures in displays like Figure 3b, even when this format could be anticipated as most probable.[9]

These results suggest that symmetry within a figure is coded by the match between convex parts on the two sides. Authors since Pascal[10] have speculated that symmetry is derived effortlessly. Recent theories of visual processing suggest that processes that are effortless operate in parallel across the visual field, whereas effortful visual processes are serial (i.e., require sequential allocation of attention to multiple locations). Therefore, we examined[11] whether the part description on which symmetry is based can be derived in parallel within a shape. Subjects had to judge the presence or absence of symmetry in shapes with 4, 8, or 16 steps on each side (see Fig. 4a for examples). Performance was scarcely affected by the number of steps (and thus component parts) in each shape. This result suggests that the layout of component parts within a figure can be derived in parallel. As a control for the changing size of the steps in the different conditions, a second experiment required subjects to judge whether each shape had repeated contours (see Fig. 4b). As noted before, this task cannot be based on the match between convex parts. In contrast to symmetry judgments, the repetition task showed substantial delays and increased errors as the number of parts in the shapes increased.

PART DESCRIPTIONS ARE OBJECT-BASED

We have also examined[4,12] whether symmetry judgments (and hence the underlying part descriptions) can be derived in parallel across two distinct shapes. As before, subjects had to judge whether the contours on the two sides of shapes were symmetrical or unre-

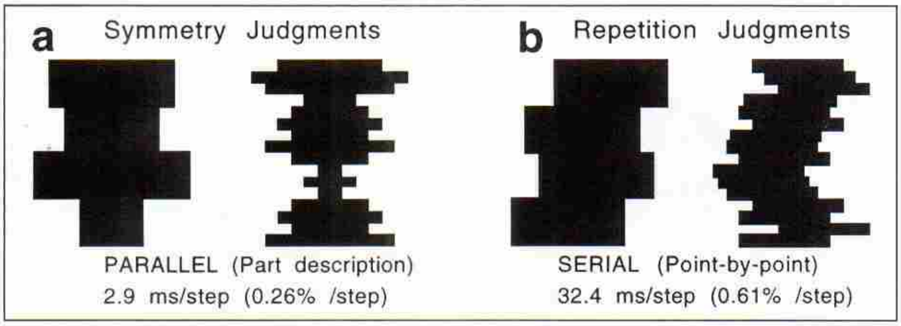

Fig. 4. Example stimuli and results from a study[11] looking at the ease of detecting symmetry (a) and repetition (b). Symmetry, but not repetition, can be detected in parallel within a visual shape. The effect of adding more points to be judged is shown for time and error rates.

lated. However, in addition to displays comprising a single tall shape (Fig. 5a), some displays comprised two short shapes, presented one above the other (Fig. 5b). Each such pair of short shapes was produced by simply deleting an uninformative section (which was always straight, regardless of whether the curved contours were symmetrical) from the middle of a single tall shape. In another condition (Fig. 5c), this middle section was deleted by means of a visible occluder, so that the separated shapes appeared as parts of a single object under partial occlusion. Regardless of display type, the task was to determine whether all the curved contours were symmetrical about the vertical. Even though these contours were equivalent across conditions (except for whether uninformative straight edges were also visible), performance was significantly slower in the two-object condition (Fig. 5b) than in either the unoccluded (Fig. 5a) or the occluded (Fig. 5c) single-object condition.

This finding provides a striking case of a difficulty in attending to two objects simultaneously.[13] The one-sided edge-assignment heuristic that we have advocated here and elsewhere[2] can lead to this difficulty when the edges that must be compared are assigned to distinct figures rather than to a single common shape. This point is further illustrated by recent experiments[14] in which we used ambiguous figure-ground displays analogous to Rubin's famous faces-vase engraving[15] except that they comprised unfamiliar shapes (see Fig. 6a). Our critical manipulation was whether each such display was seen to contain one figure (analogous to Rubin's vase) or two figures (analogous to his faces).

The displays were presented briefly in red and green, with these two colors chosen so that neither was consistently preferred as figure (as confirmed in separate studies). The central shape was in one color, and the two flanking shapes in the other color. In the critical displays, these three shapes were adjoining (as illustrated in Fig. 6a) so that they shared common dividing contours. The task was to compare the height of apices in the pair of dividing contours and indicate the side with the lower apex. One group of subjects was instructed to judge red contours, and another to judge green contours. When the display comprises a central red shape with two flanking green shapes, the critical contours should be seen as belonging to a common figure for red subjects, but to two distinct figures for green subjects. The reverse should apply for a display comprising a central green shape with two flanking red shapes. Adherence to color instructions was encouraged (as for the experiment illustrated in Fig. 2) by randomly intermingling the critical trials, which displayed three adjoining shapes, among biasing trials in which the three shapes in each display were spatially separated. In these biasing trials, the shapes had different contours so that the correct choice of lower apex depended entirely on color instruction.

We found that judgments of apex height on the critical trials were slower when the display contained two outer figures in the instructed color, rather than one central figure. In other words, those displays that were relatively easy for red subjects were relatively hard for green subjects, and vice versa. Thus, performance did not depend on the physical display itself, but on the assignment of the dividing edges that had to be judged, with performance being poorer when they were assigned to two distinct outer figures rather than to a common central shape.

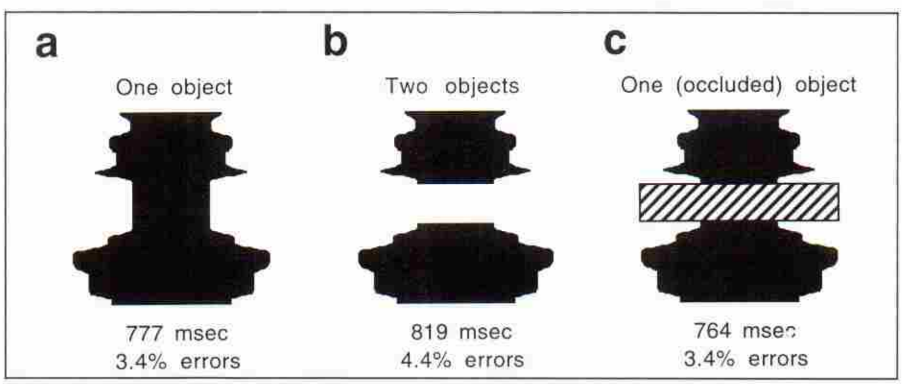

Fig. 5. Example stimuli and results from a study of the ease of deriving shape descriptions for single or multiple objects.[12] The study compared speed and accuracy of symmetry judgments for a single object (a), two separate objects (b), and a single, partially occluded object (c).

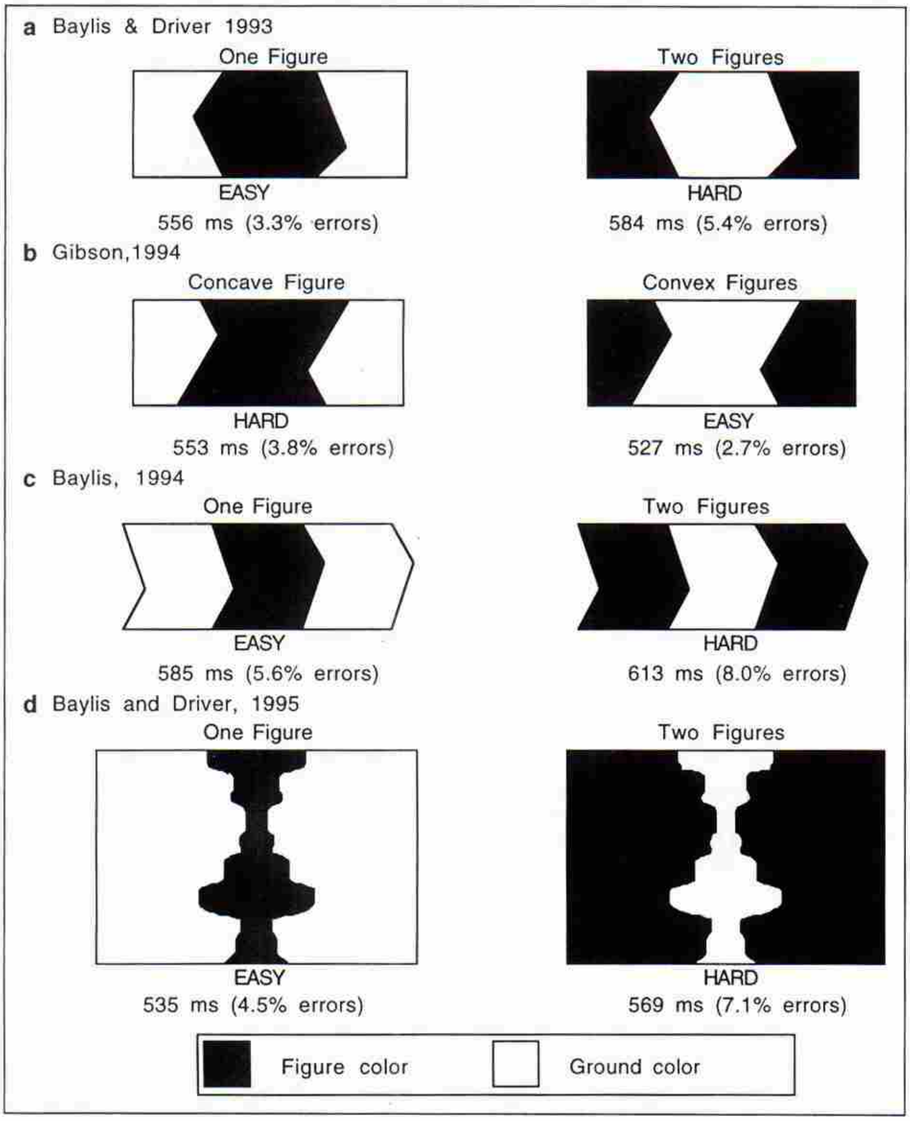

Fig. 6. Studies showing the effect of assigning edges to a single object or to two objects. (a) Two example displays and results from a study[14] in which subjects judged which of two sharp apices was lower. In some displays, apices were segmented as part of two distinct figures (right example), whereas in other displays, they were segmented as parts of a common figure (left example). (b) Example displays and results from a similar study[16] in which the apices were concave for the single-figure interpretation (left) and convex for the two-figure interpretation (right). (c) Example displays and results from a study[17] in which convexity was equated for the two-figure (right) and single-figure (left) interpretations. (d) Example displays and results from a study[7] in which subjects judged whether two curved edges were symmetrical or unrelated. The methods used to generate the curves ensured that the degree of convexity and concavity of the two-figure and single-figure interpretations would be equalized over trials.

Gibson[16] suggested that this finding may reflect convexity differences between the assignments, rather than the number of figures. Indeed, he found an advantage for comparing convexities of two objects versus concavities of a single object (see Fig. 6b), consistent with our claim that convexities provide the constituent parts of shapes for the human visual system. Nonetheless, when all regions have equal convexity, a cost is still found[7,17] for judging edges that are assigned to two distinct figures rather than to a single common figure (see Figs. 6c and 6d).

DOES THE GROUND RECEIVE A PART DESCRIPTION?

The studies we have reviewed show that subjects cannot prevent figure-ground assignment of the edges they judge, and that the figural side of the edge receives a decomposition into convex parts. Rubin[1] originally argued that figures appear shaped, whereas grounds appear shapeless. The computational theory of Hoffman and Richards,[3] together with our performance findings, suggests a straightforward account of this phenomenal difference. Although the figural side of an assigned edge receives a shape description in terms of convex parts, ground may never be decomposed in this way. If only figures receive shape descriptions in terms of convex parts, this would explain their immediate recognition advantage over grounds,[2,4,8] plus the substantial effects of figure-ground segmentation on edge comparisons that we have reviewed here.

However, a recent series of elegant studies[18] by Peterson and colleagues led them to propose that both sides of a dividing edge always receive an initial part decomposition before a one-sided edge assignment is settled upon. Their argument rests on the repeated finding that *denotivity* (roughly speaking, the meaningfulness of a shape) can determine figure-ground segmentation for a dividing edge when other bottom-up factors are fairly ambiguous. For example, when a face profile divides a rectangle into equal areas of black and white against an intermediate gray screen, the face side is usually preferred as the phenomenal figure. Because shape determines which region eventually becomes figural in this case, one interpretation is that shape must be derived for both sides of the edge prior to any figural assignment.

It is possible that shape may influence figure-ground segmentation in this way only under restricted conditions. In studies (such as Peterson's) in which subjects make judgments of which side appears figural in displays with an ambiguous segmentation, subjects may become alerted to the possibility of familiar shapes after a few trials on which the denotative side gets the dividing edge assigned to it by chance.

An influence of denotivity on judgments of figure has also been found[18] when some image-based factors would tend to make the nondenotative side figural. However, it appears that denotivity can oppose only those bottom-up variables that are derived relatively slowly, so that figural assignment may be somewhat ambiguous at least initially. Whether grounds receive any part decomposition at all when bottom-up factors unambiguously assign the dividing edge away from them at the very outset, or when displays comprise only unfamiliar shapes, remains an important question.

The experiments we have reviewed here, together with those considered in our previous article,[2] lead us to the following view. One-sided assignment of dividing edges is an obligatory heuristic in human vision. The direction of assignment, and thus which region becomes figural, is usually determined bottom up, that is, by image-based segmentation factors, but in ambiguous cases may be settled top down, that is, by the strategic allocation of attention. One-sided edge assignment leads to a description of the resulting figural shape in terms of component convex parts. This description may be derived in parallel across parts within a shape, but cannot be readily accessed for two distinct shapes simultaneously.

Edge assignment results in visual attention being object-based[13] in two ways. First, when subjects attempt to judge a dividing edge, they cannot select just the edge. Instead, they end up selecting the entire figural shape to which that edge has been assigned. Second, edge comparisons are less efficient when the edges in question are assigned to distinct objects rather than to a common object. Thus, the allocation of visual attention depends on the parsing of the visual scene. This object-based attention contrasts with purely spatial allocation of attention, which would not be affected by visual parsing.

Acknowledgments—Jon Driver was supported by the Science and Engineering Research Council (United Kingdom). Gordon C. Baylis was supported by U.S. Office of Naval Research Contract N00014-91-J-1735 and Public Health Service Grant R29 NS27296. Our thanks to Randy Gallistel, Steve Palmer, and Mary Peterson for helpful comments.

Notes

1. E. Rubin, Figure and ground, in *Readings in Perception*, D.C. Beardslee and M. Wertheimer, Eds. (Van Nostrand, New York, 1958); K. Koffka, *Principles of Gestalt Psychology* (Routledge & Kegan Paul, London, 1935).

2. G.C. Baylis and J. Driver, One-sided edge assignment in vision: 1. Figure-ground segmentation and attention to objects, *Current Directions in Psychological Science, 4*, 140–146 (1995).

3. D.D. Hoffman and W.A. Richards, Parts of recognition, *Cognition, 18*, 65–96 (1984).

4. J. Driver and G.C. Baylis, Visual attention and object perception, in *The Attentive Brain*, R. Parasuraman, Ed. (MIT Press, Cambridge, MA, in press).

5. J. Driver and G.C. Baylis, Figure-ground segmentation and visual part decomposition, *Perception* (in press).

6. E. Mach, *The Analysis of Sensation* (Dover, New York, 1959) (Original work published 1885).

7. G.C. Baylis and J. Driver, Obligatory edge-assignment in vision: The role of figure- and part-segmentation in symmetry detection, *Journal of Experimental Psychology: Human Perception and Performance* (in press).

8. J. Driver and G.C. Baylis, Edge-assignment and figure-ground segmentation in short-term visual matching, *Cognitive Psychology* (in press).

9. Many observers spontaneously remark that they can perceive symmetry in displays like Figure 3b only by willfully reversing the initial bottom-up organization to yield a single virtual object in which the vertical edge of the rightmost shape provides a central axis between symmetrical halves of different brightness. Such effortful, top-down (i.e., strategic) reversal of the initial figure-ground assignment provides one possible explanation for the substantial time cost induced by this stimulus format. Even the occurrence of such a figure-ground reversal is entirely consistent with our emphasis on the obligatory nature of the initial edge assignment and on the importance of matching between convex parts in symmetry detection.

10. Pascal, *Pascal's Pensées*, H.F. Stewart, Trans. (Routledge & Kegan Paul, London, 1950).

11. G.C. Baylis and J. Driver, Symmetry but not repetition is detected in parallel within single visual shapes, *Visual Cognition, 1*, 377–400 (1994).

12. J. Driver, G.C. Baylis, and M. Price, *Judging visual symmetry in one object versus two objects*, manuscript submitted for publication (1995).

13. N. Kanwisher and J. Driver, Objects, attributes, and visual attention: Which, what, and where, *Current Directions in Psychological Science, 1*, 26–31 (1992); J. Duncan, Selective attention and the organization of visual information, *Journal of Experimental Psychology: General, 113*, 501–517 (1984).

14. G.C. Baylis and J. Driver, Visual attention and objects: Evidence for the hierarchical coding of location, *Journal of Experimental Psychology: Human Perception and Performance, 19*, 451–470 (1993).

15. See Figure 1 in Baylis and Driver, note 2.

16. B. Gibson, Visual attention and objects: Number of objects or convexity? *Journal of Experimental Psychology: Human Perception and Performance, 20*, 203–207 (1994).

17. G.C. Baylis, Visual attention and objects: Two-object cost with equal convexity, *Journal of Experimental Psychology: Human Perception and Performance, 20*, 208–212 (1994). Note that the performance cost when there are two figures in the relevant color may reflect either the time required to shift attention between distinct objects or the time for a willful reversal of the initial figure-ground segmentation to yield a single central figure rather than the initial pair of outer figures. Either possibility implies it is harder to judge edges from two distinct figures than comparable edges within a single figure.

18. M.A. Peterson, Object recognition processes can and do operate before figure-ground organization, *Current Directions in Psychological Science, 3*, 105–111 (1994); M.A. Peterson and B.S. Gibson, Shape recognition inputs to figure-ground organization in three-dimensional displays, *Cognitive Psychology, 25*, 383–429 (1993); M.A. Peterson and B.S. Gibson, Must figure-ground organization precede object recognition? An assumption in peril, *Psychological Science, 5*, 253–259 (1994).