VISUAL COGNITION, 2001, *8* (2), 163–196

Perception of symmetry and repetition within and across visual shapes: Part-descriptions and object-based attention

Gordon C. Baylis
University of South Carolina, Columbia, USA
Jon Driver
University College London, UK

In five experiments, we investigated the detection of symmetry (i.e., translation plus reflection) or repetition (i.e., translation alone) between two vertical jagged contours. The complexity of the two contours was manipulated, as was their figure-ground assignment; the two contours either belonged to a common object "inside" them, to two separate objects "outside" them, or to two separate objects each to the right of one contour. Replicating Baylis and Driver (1994), symmetry judgements were unaffected by contour complexity when made within a single shape, implying a parallel process operating efficiently across contour discontinuities. However, all the other conditions yielded substantially slower judgements as complexity increased, suggesting either effortful point-by-point comparisons, or a highly inefficient parallel process. In agreement with Baylis and Driver (1995a), symmetry perception was harder when figure-ground assignment turned convexities along one contour into concavities along the other contour; and likewise for repetition detection. However, even when convex parts matched between the two contours, judgements were still affected by complexity unless they belonged to a common object. This supports Baylis and Driver's (1993) proposal that effortless comparisons for the layout of multiple convex parts can only be made within single perceptual objects.

From a very young age, people are extremely good at detecting visual symmetry, especially about a vertical axis (e.g., Baylis, 1998; Bornstein & Krinsky, 1985). Symmetry judgements often occur swiftly and without apparent effort.

Please address all correspondence to G. Baylis, Department of Psychology, University of South Carolina, Columbia, SC 29208, USA.

Gordon C. Baylis was supported by a grant from the National Science Foundation (SBR 96-16555). Jon Driver was supported by project grants from the BBSRC and ESRC (United Kingdom).

http://www.tandf.co.uk/journals/pp/13506285.html DOI:10.1080/13506280042000126

As Pascal (1888/1950) noted, “symmetry is what we see at a glance”, and many subsequent authors have similarly commented on the high efficiency of symmetry perception (e.g., Barlow & Reeves, 1979; Baylis & Driver, 1994; Bruce & Morgan, 1975; Corballis & Roldan, 1974; Julesz, 1971; Koffka, 1935; Mach, 1885/1959). This has led to substantial interest in the possible underlying perceptual mechanisms (see Tyler, 1996, and Wagemans, 1995, for reviews).

A primary task for the visual system is to recognize visual objects across considerable variation in the retinal input, and despite distortions introduced by factors such as non-linearities in the visual system itself. Many workers (e.g., Biederman, 1987; Lowe, 1987; Marr, 1982) have pointed out that the visual system might usefully exploit so-called “non-accidental” image properties in accomplishing this difficult task. Non-accidental properties are aspects of the image which are unlikely to have resulted from coincidences in the current projection of the visual scene onto the retina, or from distortions introduced by the visual system itself. In other words, they are image properties that are particularly likely to reflect aspects of the distal image source, and hence are of prime importance in visual perception.

Bilateral mirror symmetry seems to be a paradigm case of a non-accidental property, since it is extremely unlikely that such regularity could occur in the retinal image in the absence of any corresponding regularity in the image source. However, mirror symmetry is only one of a number of possible regularities that are equivalent in terms of the degree of redundancy they introduce to the image, or their formal information content. Formally, amounts of information are mathematically defined in terms of the uncertainty reduced by an event, with less probable events thus yielding more information (Shannon & Weaver, 1949). Because non-accidental properties are precisely those that are unlikely to arise by chance in an image, they tend to be properties that are high in terms of the formal information they can convey. However, regularities that are equivalent in terms of their mathematical information content may not be equivalent psychologically, and this seems particularly so for visual symmetry. For instance, compare the experience of looking at Figure 1(a) versus Figure 1(b).

As Mach (1885/1959) originally pointed out, the symmetry in Figure 1(a) is striking, yet the regularity in Figure 1(b) is far less so. This difference between the ease of detecting mirror symmetry (i.e., translation plus reflection of a contour, as in Figure 1a), versus the difficulty of detecting mere translation without reflection (as in Figure 1b; for convenience we shall refer to this form of regularity as “repetition”) has now been experimentally confirmed by many researchers (e.g., Baylis & Driver, 1995a, b; Bruce & Morgan, 1975; Corballis & Roldan, 1974). We shall term this finding the “Mach effect”, to acknowledge its original discoverer. The issue raised by this Mach effect is as follows. When considering a human visual system that is thought to have evolved particular sensitivity to non-accidental properties, one might wonder why regularity such

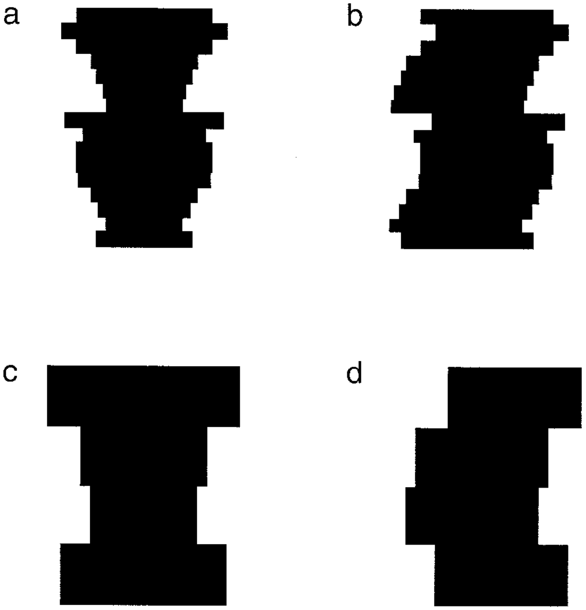

Figure 1. Example shapes from Driver and Baylis (1994). (a) and (c) show symmetrical stimuli that have 16 and 4 "steps" (i.e., discontinuities along their border) respectively; whereas (b) and (d) show repeated stimuli (i.e., with translated edges) having a comparable number of steps.

as literal repetition (Figure 1b) is not as salient as mirror symmetry (Figure 1a). Surely repetition of a contour is just as unlikely to have occurred by chance, and thus should be equally "non-accidental", and equally informative?

In fact, a substantial proportion of real-world objects are themselves mirror symmetrical, and hence are able to project an image that is mirror symmetrical (or nearly so). As a result, detection of mirror symmetry may provide one useful way of segmenting objects from their background (see Bahnsen, 1928; Driver, Baylis, & Rafal, 1992; Rock, 1983). By contrast, we are unable to think of many real-world examples of objects with contours that are repeated on either side, in the present sense of the contours undergoing mere translation. A degree of such repetition may sometimes arise between the occluding contours of objects and their corresponding shadow, but note that the regularity in the image here would be an "accident" of the current lighting conditions that cast the shadows, unlike the regularity between the two sides of a symmetrical object.

Thus, one can generate possible suggestions for why the perception of mirror symmetry might be of greater functional use than the perception of repetition between contours. However, this alone does not reveal the computational basis of the difference between the perceptual ease of detecting symmetry versus repetition. Baylis and Driver (1994) examined whether the mechanisms underlying symmetry perception may be qualitatively different from those underlying repetition detection, in addition to the obvious difference in subjective difficulty. In particular, they examined whether symmetry perception may reflect a parallel or co-operative process, which can apply efficiently to all the various parts of an object at the same time, whereas repetition detection may force observers to resort to an inefficient or effortful (and perhaps serial, point-by-point) comparison.

In Baylis and Driver's (1994) first experiment, observers judged whether the two sides of a vertically elongated 2D object were symmetrical or asymmetrical, while the "complexity" of the object's shape was manipulated, in terms of the number of discontinuities along each side (there were either four, eight, or sixteen discrete "steps" along each side; e.g., compare the different number of discontinuities in Figure 1a versus 1c). The two sides were either perfectly mirror-symmetrical, or else 25 per cent of the "steps" along one side were randomly unrelated to the other side. Baylis and Driver (1994) proposed that any serial comparison between discontinuities along the two sides should lead to slower judgements for shapes with additional discontinuities. In contrast, an efficiently parallel process of symmetry detection (or one which can apply co-operatively across many features at the same time), should not show this slowing. In fact, judgements were only very slightly affected by shape complexity (being slowed by only 3 ms per additional discontinuity that had to be considered).

The outcome was very different when perfect repetition rather than symmetry had to be detected between the two sides (Figure 1b or 1d). Now every additional discontinuity in the shape substantially delayed detection of the regularity, suggesting (by analogy with the visual search literature; e.g., Treisman & Gelade, 1980) that an effortful point-by-point comparison of each discontinuity along the two sides (or a very inefficient "noisy" parallel process) may be required for repetition detection. This result provides some initial evidence for qualitative differences in the mechanisms that underlie symmetry perception versus repetition detection. However, it does not offer any direct explanation for why effortful comparisons seems to be required across multiple discontinuities in a shape for repetition detection, yet not for symmetry perception; the result merely documents this fact.

Baylis and Driver (1995a, b) suggested that the difference may originate in the representational nature of the shape-descriptions that are routinely derived for visual objects. A number of workers (e.g., Biederman, 1987; Hoffman & Richards, 1984) have suggested on computational grounds that the shape of an object's occluding contour may be represented in terms of the layout of convex parts, with adjacent parts being separated by points of concavity, or more formally by negative minima of curvature along the occluding contour. Driver and Baylis (1995) review evidence for the psychological reality of this computational scheme in the human perception of visual shapes.

Once shape is considered in terms of such convex parts, an important difference between mirror symmetry and repetition becomes apparent. Consider the visual objects in Figure 2. In the case of the symmetrical object in Figure 2(a), the parts (i.e., convexities) on the two sides of the object are equivalent. In contrast, the shape descriptions in terms of convex parts for the two repeated sides of the object in Figure 2(b) do not match; convex parts on one side now correspond to concavities (that is, to boundaries between parts) on the other side. Baylis and Driver (1995a) suggested that this may explain the usual Mach effect, whereby mirror symmetry is more salient than repetition between the two sides of a shape. In the case of a symmetric object, the regularity between the two sides may be apparent simply on the basis of the proposed routine shape-description, which represents the layout of convex parts. By contrast, such a shape-description would only serve to obscure the correspondence between the two sides for the case of an object with repeated contours, as convex parts do not correspond for the two repeated sides. On this account, symmetry perception might seem so effortless because it can be subserved by routine

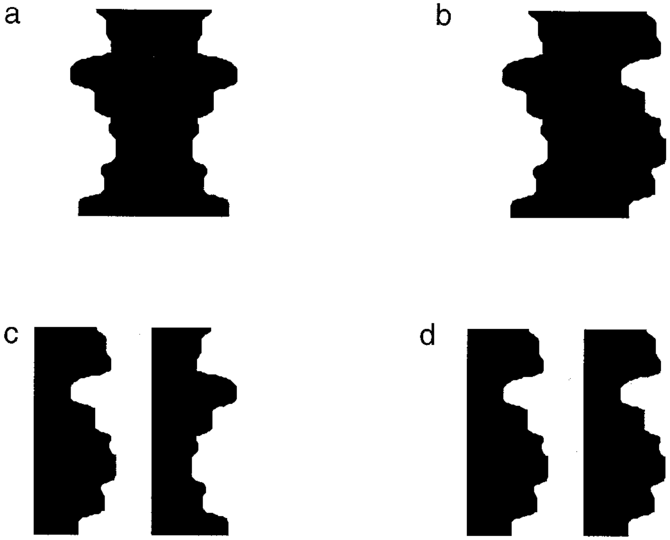

Figure 2. Four examples of the displays from Baylis and Driver (1995a, Exps. 1 and 4: (a) symmetrical shape (note that each convexity on one side is matched by a corresponding convexity on the other side); (b) shape with repeated contours (note that a convex region along one side becomes concave along the other side, and vice-versa, yielding non-corresponding convex parts for the two sides); (c) The curved contours of these two shapes have a symmetrical relation, but convex regions along one contour are now concave along the other, yielding non-corresponding convex parts due to the figure-ground reversal for the left-most of the two contours, as compared with (a); (d) The curved contours in these two shapes are repeated, and now have matching convex parts. Baylis and Driver (1995) found that symmetry judgements were easier than repetition judgements with the format exemplified by (a) and (c). The reverse was found for the format shown in (b) and (d).

shape-descriptions in terms of convex parts, such as those proposed by Biederman (1987), Driver and Baylis (1995), and Hoffman and Richards (1984). By contrast, only repetition detection would require deliberate and effortful comparisons, since the regularity in this case would not be apparent from the standard shape-description.

Baylis and Driver (1995a) tested this account by changing the format of the displays so as to reverse figure-ground assignment (and hence reverse which regions should be coded as concave versus convex) for just one of the two contours that had to be judged as symmetrical versus unrelated, or as repeated versus unrelated. In addition to the formats shown in Figures 2(a) and 2(b), their participants also judged displays like those in Figure 2(c), where the shape-descriptions in terms of convex parts no longer match for two symmetric contours, because of the figure-ground reversal for one contour. A similar figure-ground manipulation produced displays like Figure 2(d), where repetition detection can now be based on the match in convex part-descriptions for the two repeated contours. With the formats shown in Figure 2(c) and 2(d), symmetry judgements actually became harder than repetition judgements, reversing the usual Mach effect. This implies that the Mach effect does not depend solely on the presence versus absence of mirror symmetry between contours, as previously thought, but rather on whether the convex parts yielded after figure-ground assignment match or mismatch for the two contours that must be compared.

Thus, the emerging account is as follows: the perception of symmetry within a shape seems so effortless because the correspondence between the two sides is apparent in the standard shape-description in terms of convex parts, which yields matching convexities along each side. From this perspective, Baylis and Driver's (1994) finding that symmetry perception within a shape is scarcely affected by the number of convex parts implies that the relative position of all the component parts of a single shape are accessible simultaneously, in an efficiently parallel manner. Baylis and Driver (1993; see also Baylis, 1994) had previously argued that this is effectively a logical requirement for useful shape perception, since shapes are defined precisely by the relative layout of all their parts, and thus a true appreciation of shape implies a simultaneous appreciation of the spatial disposition of all of the component parts which make up that shape.

Repetition detection within a shape is by contrast harder and more effortful than symmetry perception, because the regular relationship between the two sides is not apparent in the standard shape-description in terms of convex parts, as this yields a mismatching decomposition into convex parts for two repeated sides. Hence, an effortful point-by-point comparison (or some highly inefficient parallel process) is required to reveal repetition within a shape. However, repetition detection can be made easier by reversing figure-ground assignment for one of the two contours, so that convex parts now match between repeated contours (Figure 2d). Conversely, symmetry perception can be rendered harder by the same figure-ground manipulation (Figure 2c), which produces convex parts that no longer match for symmetric cases. In the latter case, even symmetry perception may come to depend on the effortful point-by-point comparison (or highly inefficient parallel process) that Baylis and Driver (1994) observed for repetition detection, when they varied the number of discontinuities along the contours of a single shape (see Figure 1b versus 1d).

This leads to the hitherto untested prediction that symmetry perception should likewise become dependent on the "complexity" of the contours to be compared for formats such as Figure 2c, where convex parts no longer match. Recall that Baylis and Driver (1994) found that additional discontinuities along each contour scarcely affect symmetry perception within one object, a case where two related contours will have matching convex parts. This was taken to imply that symmetry judgements can be based on efficient parallel processing of all the parts of an object simultaneously. Our first experiment tested the new prediction that such processing should become ineffective when convex parts no longer matched between the symmetrical contours, forcing observers to rely on an effortful point-by-point comparison instead (or some highly inefficient parallel process) in order to detect the regularity.

EXPERIMENT 1

As in Baylis and Driver (1994), this study examined how the number of discontinuities along two contours affects the ease of judging whether or not they are symmetrical. Our earlier study on this issue had found that symmetry judgements can apparently be based on a highly efficient parallel match of all the parts within a shape, with little effect of the number of discontinuities along each side. Following Baylis and Driver (1993; 1995a), we now hypothesized that this may only be possible when the symmetry judgement can rely on the standard shape-description of an object in terms of its convex parts. When convex parts are made to mismatch for symmetrical contours, as with the format in Figure 2c, a more effortful comparison of the two contours should now be required instead, leading to slower judgements with additional discontinuities.

Method

Participants. The 16 participants, 10 female and 6 male, were psychology undergraduates with normal or corrected acuity by self-report. Participants received course-credit for their participation.

Apparatus and materials. The experiment was conducted on an Optimal 286 microcomputer with colour VGA graphics. Display onsets and offsets were provided by altering the colour lookup table to ensure that they occurred within a single frame.

Four example displays are shown in Figure 3. They were presented in bright red on a black background. Viewing distance was 70 cm. The shapes were vertically elongated and approximately rectangular, with completely straight horizontal edges along their short ends, but with some markedly jagged "stepped" edges along their vertical axis, just as in Baylis and Driver (1994). These jagged contours were 5.5° in height, and were on average equally spaced around the left and right of the vertical midline of the screen, being separated by a mean 3.1° horizontally (each step along the jagged edges varied randomly between 0.5° and 2.6° from the vertical midline of the screen, in 0.03° units). The number of jagged steps along each of these contours was varied as in Baylis and Driver (1994), to provide the manipulation of visual complexity. There were either 4, 8, or 16 discontinuous steps along each side, always at equal intervals (e.g., there are 16 steps in Figures 3a and 3c, but only 4 steps in Figures 3b and 3d).

Once generated, these jagged contours were blocked out with colour to produce two different types of display which were analogous to the single-object format (Figures 3a and 3b) and the non-corresponding two-object format

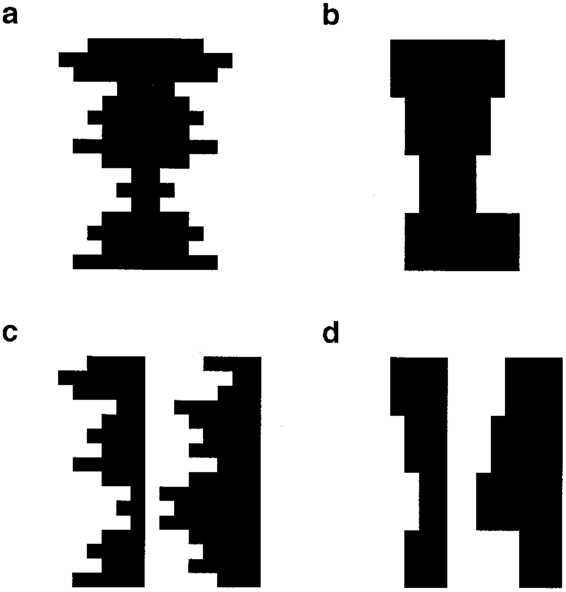

Figure 3. Four examples of stimuli from Experiment 1. These were shown in red on a black background. Top: Examples of displays for the single-object condition: a symmetric display (a) and an asymmetric display (b) with 16 steps or 4 steps respectively along each elongated side. Bottom: Examples of displays for the non-corresponding condition: a symmetric display (c) and an asymmetric display (d), with 16 or 4 steps on each elongated side.

(Figures 3c and 3d) previously used by Baylis and Driver (1995a; see Figures 2a and 2c). In the non-corresponding format, two separate figures were formed, one comprising just the left half from the single-object condition, and the other figure, which was on average equally wide, appearing to the right of the rightmost jagged contour (see Figures 3c and 3d). The asymmetrical contours (e.g., Figures 3b and 3d) were constructed exactly like the symmetrical contours, except that 25 per cent of the jagged steps on one edge were altered to a new random value that was at least 0.5° different (in horizontal distance from the screen centre) from the value on the other side.

The participants' task was to respond by pressing one of two buttons (the "Z" key or the "/" key on the standard extended keyboard) with one or other index finger as rapidly as possible, depending on whether the jagged contours in the display were perfectly symmetrical or 25 per cent asymmetrical. The mapping of particular hands to symmetrical and asymmetrical responses was counterbalanced across participants.

Design. The design was within-subject with three independent variables. The first was whether the jagged edges were blocked out with colour to yield the single-object format or the non-corresponding format (henceforth, the display format factor). The second factor was whether displays were perfectly symmetrical or partly asymmetrical (e.g., Figures 3a and 3c, versus 3b and 3d). The third factor was the number of pseudorandom steps along the elongated sides of each shapes (either 4, 8, or 16 on each side, exactly as in the previous study by Baylis & Driver, 1994; there are 4 steps in Figures 3b and 3d, but 16 in Figures 3a and 3c). These three factors were crossed to yield 12 equiprobable and randomly intermingled conditions.

Procedure. Participants were shown a diagram of typical displays, similar to Figure 3, and told that they must judge whether each presented shape was perfectly symmetrical or not, responding as quickly and accurately as possible by pressing the appropriate key. The sequence of events on any one trial was then explained to them. A fixation cross was presented for 500 ms, followed by a display centred at fixation until the participant responded, whereupon the screen became blank. An incorrect response produced a loud beep, whereas no feedback was given on correct trials. Following an intertrial period of 800 ms, this sequence was repeated to produce the next trial. Reaction times (RTs) were recorded in ms.

Participants were given 8 blocks of 150 trials. Within each block the conditions were interleaved in a different pseudorandom sequence for each participant. At the end of each block, the participant's mean RT for correct responses was displayed on the monitor, together with their mean error rate, and a message that requested that they be more accurate in the next block if their error rate had exceeded 15 per cent, or that they respond more quickly if their error rate had been below 5 per cent. Participants rested for as long as they wished between blocks, pressing the space-bar to continue.

Treatment of results. The first block of 150 trials was discarded as practice, as were the first 2 trials of each block. Thus 1036 trials were available for each participant. All these trials contributed to the accuracy analyses. However, the data were trimmed for RT analysis by removing incorrect responses, plus trials immediately following an error because of the variability they typically introduce (Rabbitt, 1966). Upper and lower RT cut-offs for acceptable trials were also calculated for each participant using the method of Driver and Baylis (1991), which excludes responses with very long or short latencies beyond which performance is at chance. This procedure derives the lowest upper bound beyond which responses are no more likely to be correct than incorrect (as tested by chi-square at the .05 level), and derives the highest lower bound in a similar manner. The combination of these upper and lower RT criteria excluded 2.0 per cent of the recorded data (2.1 per cent for Experiment 2, 1.9 per cent for Experiment 3, 2.3 per cent for Experiment 4, and 4.2 per cent for Experiment 5). For all experiments in this paper, all further analyses were carried out using SYSTAT (Wilkinson, 1990).

Results

The mean results are shown in Figure 4, with the means of participants' median RTs appearing in the upper graph, and the error rates in the lower graph, plotted against the number of discontinuous steps along the jagged vertical contours (i.e., complexity). Separate functions are shown for the single-object format, and for the non-corresponding format, and likewise the data are separated for symmetrical versus asymmetrical responses. It is immediately apparent that the number of steps had a substantially larger effect for the non-corresponding format than the single-object condition.

A three-way within-subjects analysis of variance (ANOVA) on the RT data showed significant effects of: display format, $F(1, 15) = 74.24$, $p < .001$, with slower responses for non-corresponding displays; symmetry, ($F(1, 15) = 4.8$, $p < .05$, with faster responses overall for symmetrical contours; and the number of steps, $F(2, 30) = 41.2$, $p < .001$, with slower responses when there were more steps. The interaction between display type and the number of steps was also highly significant, $F(2, 30) = 9.6$, $p < .001$, with a much greater effect of steps for the non-corresponding displays. Finally, the three-way interaction also reached significance $F(2, 30) = 5.4$, $p < .01$, because the steepest function against number of steps was for symmetrical responses only in the non-corresponding condition. Neither of the other two-way interactions approached significance.

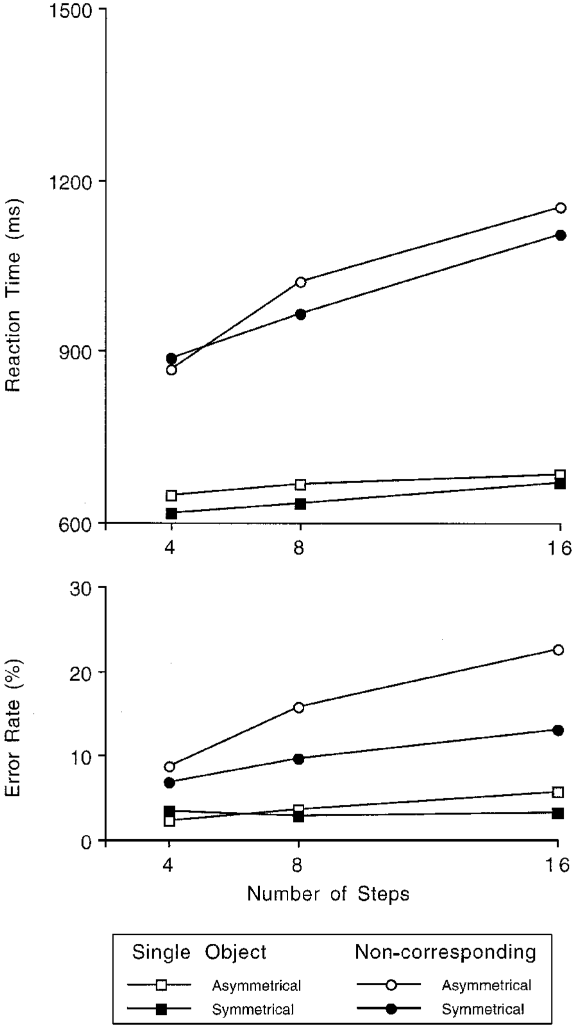

Figure 4. Means of participants' median RTs (top) and error rates (bottom) for the different display types in Experiment 1. The extensive scale is to allow a direct comparison with the results of Experiments 2–5, which are plotted against similar scales in Figures 6, 7, 9, and 10.

An analogous three-way ANOVA on the error data showed a largely similar pattern of results. There were significant effects of display format, $F(1,15) = 78.0$, $p < .001$, with less accurate responses for non-corresponding displays, and of symmetry, $F(1, 15) = 15.8$, $p < .001$, with more accurate responses overall for symmetrical contours; and a robust effect of the number of steps, $F(2, 30) = 38.7$, $p < .001$, with less accurate responses when there were more steps. The interaction between display type and symmetry was significant, $F(1, 15) = 13.4$, $p < .001$, indicating a larger difference between symmetrical and asymmetrical responses for non-corresponding displays, as was the interaction of display type with the number of steps, $F(2, 30) = 26.3$, $p < .001$, with a greater effect of steps for the non-corresponding displays, as in the RT data, plus the interaction of symmetry with number of steps, $F(2, 30) = 5.6$, $p < .01$, with more effect of the number of steps in the asymmetrical displays. The three-way interaction was again significant, $F(2, 30) = 5.4$, $p < .01$.

The most important result from this study was the dramatic increase in the impact of the number of steps for the non-corresponding format as compared with the single-object format. This outcome was robust in both the RT data and in the error data, with both measures increasing substantially as the number of steps increased only in the non-corresponding format.

Discussion

The results from the single-object format support previous evidence (Baylis & Driver, 1994) for efficiently "parallel" processing of the discontinuities along the contours of a single visual shape which determine whether or not that shape is symmetrical. There was admittedly a slight increase in RT with additional steps even in the single-object format, but this influence was very small. Indeed, linear regression of RT against the number of steps for the single-object condition gave slopes of only 4.6 ms/step for symmetrical responses and 3.7 ms/step for asymmetrical responses (with R^2 of 1.0 and 0.98 respectively). This is in close agreement with the previously reported values of 3.2 and 2.5 ms/step respectively for this condition in Baylis and Driver (1994). Such slope values are within the range normally taken to indicate efficient parallel processing of items in the visual search literature (e.g., Treisman & Gelade, 1980). In the present case, the suggested parallel processing would apply to all the discontinuities along a single shape's borders.

When the jagged contours were blocked out in such a way as to reverse figure-ground assignment for one of them, so that symmetrical contours should no longer receive matching descriptions in terms of convex parts (see Figures 3c and 3d), performance became much poorer overall. This confirms Baylis and Driver's (1995a) initial report of the difficulty of this format for symmetry judgements, presumably due to the mismatch in convex parts. The novel finding was that with these non-corresponding displays, increasing the complexity of the visual shapes now led to a marked increase in reaction time, and also in the number of errors. These increases are suggestive of an effortful point-by-point comparison rather than an efficient parallel process, or at least of a very inefficient parallel process that becomes increasingly degraded as extra discontinuities must be considered.

At the suggestion of a referee, we also analysed variation in reaction times for the different conditions in this experiment. There was a general trend for standard errors of the means (SEMs) to increase along with overall RT. To test for an deviation from such a pattern, we ran an ANOVA (with the same design as those earlier) on proportional variation in RT (i.e., each SEM divided by the corresponding mean RT). This showed that no factors or interactions approached significance. Those conditions with longer mean RTs (see Figure 4) were associated with larger SEMs; for instance, those conditions with a larger number of steps for the non-corresponding displays. This could be reconciled either with a serial process underlying symmetry judgements for such displays, or with an inefficient parallel process that operates less well when there are more steps.

Regardless of whether or not a strictly serial process was used to judge the non-corresponding displays, it remains clear that a much less efficient process was used than for the single-object displays. This outcome fits our previous proposals that judgements of symmetry between contours are only efficient when they can be based on the suggested "standard" shape description of a single object in terms of its convex parts, which should make the regularity in symmetrical objects explicit (in the sense of Marr, 1982) due to the matching convexities. When the format is changed such that the convexities do not match (as in the non-corresponding displays), the usual shape description is no longer sufficient to specify the regularity between symmetric contours, and so performance becomes less efficient even for the very same two contours, and performance is now increasingly degraded as the number of discontinuities to be considered rises.

Recall that Baylis and Driver (1994) had previously found that repetition detection between the two contours on either side of a shape also declined substantially as the number of discontinuities along each contour was increased (see Figure 5b versus 5a), just as found here for the non-corresponding symmetrical format. Presumably, this previous result for repetition detection may also have been due to non-correspondence between the convex parts of the contours that had to be compared, by the same argument given earlier for non-corresponding symmetry displays. If so, then reversing figure-ground assignment such that the part-descriptions of repeated contours now correspond (as in Figure 5c) might reveal efficiently parallel judgements of repetition for the very first time.

On the other hand, as mentioned earlier, Baylis and Driver (1993) have proposed that the relative layout of all component parts may only be available in parallel within individual objects, by means of a shape-description that simultaneously represents all the convex parts whose relative position constitutes the shape (see also Donnelly, Humphreys, & Riddoch, 1991). If only these within-object shape-descriptions allow for parallel comparisons among multiple parts, then judging the layout of component parts from separate shapes (as in Figure

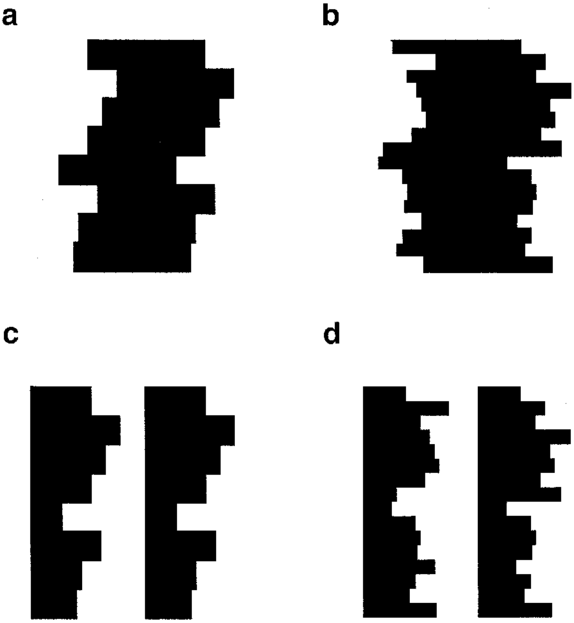

Figure 5. Four examples of stimuli from Experiments 2 and 3. These were shown in red on a black background. Top: Examples of displays for the single-object condition: (a) a repeated shape (i.e., with translated jagged contours), and (b) an unrelated shape, with 8 or 16 steps respectively along each elongated side. Bottom: Examples of displays for the corresponding condition: a repeated shape (c) and an unrelated shape (d) with 8 or 16 steps respectively along each elongated side. All these types of display were used in Experiment 2, but only the corresponding displays (c and d) were employed in Experiment 3.

5c) might always lead to inefficient comparisons, regardless of whether convexities correspond across the shapes. Effortful comparison would be required because the routine shape descriptions that we have proposed, in terms of the layout of convex parts within an object, would not directly specify the relative layout of parts from distinct objects, as such arrangements do not constitute the shape of any individual object. If this account is correct, then repetition judgements might still show substantial effects of part complexity (i.e., of the number of steps to be compared) even in a two-object format where convex parts now correspond (e.g., Figure 5c), because of the requirement to compare parts across separate objects. The next study tested these two opposing predictions.

EXPERIMENT 2

This experiment examined whether judgements of contour repetition always produce an inefficient comparison of the contours (which degrades when an increasing number of discontinuities must be considered), or whether judgements can proceed efficiently and in parallel, with little effect from the number of discontinuities along each contour, provided convex parts are made to correspond for the repeated contours (as in Figure 5c).

If the only factor determining whether discontinuities along two contours can be compared in an efficiently parallel manner is whether or not their convex parts correspond, then such efficient parallel performance should be found even for repetition detection, when the format yields matching convexities (Figure 5c). However, if performance of this judgement across two objects is inefficient and affected by the number of discontinuities that need to be considered whenever two separate objects must be compared, then a substantial effect from the number of steps should still be apparent even in this format.

Method

Unless otherwise stated, the method followed that of Experiment 1. Participants now judged whether the two elongated, jagged contours were exact repetitions of each other or not.

Participants. The 16 new participants, 9 female and 7 male, were again undergraduates with normal or corrected acuity by self-report. They were rewarded for participation as before.

Apparatus and materials. Example displays are shown in Figure 5. The task was to respond by pressing one of two buttons (the "Z" key or the "/" key on the standard extended keyboard, as before) with one or other index finger as rapidly as possible, depending on whether the jagged contours in the display were perfect repetitions of each other (e.g., Figures 5a and 5c) or not (e.g., Figures 5b and 5d). Repeated contours were generated just as for symmetrical contours in the previous experiment, except that the two jagged contours were now merely translated, rather than being translated and reflected as before. Unrepeated contours (as in Figures 5b and 5d) were generated analogously to the asymmetrical contours from Experiment 1 (i.e., 25 per cent of the steps mismatched randomly while the remainder were repetitions of each other, as defined earlier).

Half the displays had the jagged contours blocked out with colour in the intervening region to yield the single-object format (Figures 5a and 5b), with non-corresponding convex parts on the two repeated sides. The other half were blocked out to yield corresponding two-object displays (see Figures 5c and 5d), for which the convex parts of the two contours now match exactly when they are repeated, yielding two separate but identical objects. Within these formats, half the displays had perfectly repeated jagged edges, half unrepeated. The mapping of particular hands to repeated and unrepeated responses was counterbalanced across participants. The procedure was as before except for the change in task from symmetry detection to repetition detection, and appropriate changes in the instructions to accommodate this.

Results

The means of participants' median RTs, together with their associated mean error rates, are shown in Figure 6, with RT data shown in the upper graph, and error data in the lower graph. The striking result in these graphs is that *both*

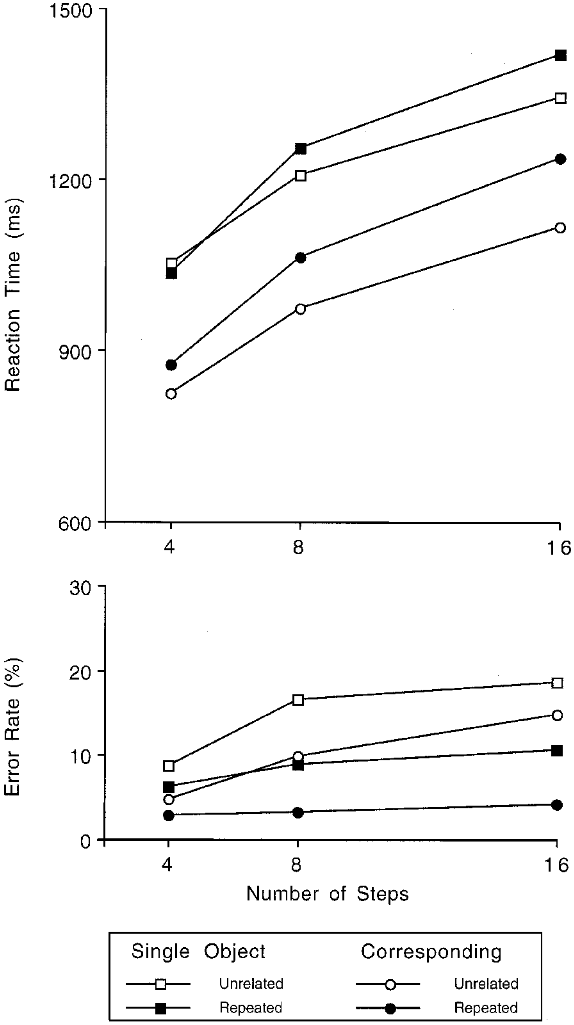

Figure 6. Means of participants' median RTs (top) and error rates (bottom) for the different display types in Experiment 2.

formats now showed substantial increases in RT as the number of steps was increased. The effect of increasing the complexity (i.e., the number of steps in the figures) was similar for the different formats. For the single-object displays, the slope was 23.4 ms/step for unrelated, and 30.6 ms per step for repeated displays. For corresponding displays the slopes were 23.5 and 29.1 ms per step respectively.

A three-way within-subjects ANOVA on the RT data showed significant effects of display type, $F(1, 15) = 46.9$, $p < .001$, with faster responses for the corresponding two-object displays, of repetition, $F(1, 15) = 6.0$, $p < .05$, with faster overall responses for unrepeated shapes, and of the number of steps, $F(2, 30) = 119.6$, $p < .001$, with much slower responses when there were more steps. None of the interactions approached significance.

An analogous three-way ANOVA on the error data showed a similar pattern of results. There was a significant effect of display format, $F(1, 15) = 38.2$, $p < .001$, with more accurate responses for corresponding displays overall, of repetition, $F(1, 15) = 6.0$, $p < .05$, with more accurate responses for repeated shapes, suggesting a possible speed/error trade-off along this factor since unrepeated shapes yielded faster RTs, and a main effect of the number of steps, $F(2, 30) = 25.4$, $p < .001$, with more errors as the number of steps increased, in accordance with the RT increase against steps. The interaction of repetition and number of steps was significant, $F(2, 30) = 7.3$, $p < .01$, with a greater influence of the number of steps for unrepeated displays. None of the other interactions approached significance.

As in Experiment 1, we also tested whether variation in RTs differed between conditions. As in that experiment, the SEMs increased only in proportion to the overall RTs; an ANOVA on the ratios (SEM divided by the corresponding mean RT) showed no factors or interactions that approached significance.

Discussion

In this experiment, judgements of contour repetition showed a marked increase in RT and error rates as the number of discontinuities to be compared was increased. The results for the single-object format are in agreement with the previous findings of Baylis and Driver (1994), who similarly observed that repetition detection for the contours on either side of a single object requires either a serial point-by-point comparison, or at best a highly inefficient and noisy parallel process, which degrades substantially as additional discontinuities along each contour must be considered. They proposed that this arises because, in the single-object format, the regularity between contours in repeated displays is not apparent from the routine shape-description of the object in terms of convex parts, since convexities do not match between the two sides.

In agreement with the previous findings of Baylis and Driver (1995), the present corresponding format led to better performance than the single-object condition overall. This fits with the suggested importance of convex parts, as these match for repeated contours in the corresponding condition but not in the single-object condition.

However, the novel and in some respects surprising finding from this experiment was that judgements of contour repetition still became substantially impaired as the number of discontinuities increased, even in the corresponding condition, where figure-ground factors were manipulated to create matching convex parts for two repeated contours. Indeed, the detrimental effect of the number of steps was just as great for the corresponding condition, with matching convex parts, as for the single-object condition, with mismatching parts (see Figure 5). This suggests that correspondence between matching convexities along two contours can only be efficiently accessed in an efficiently parallel manner when that correspondence arises within a single shape (as for the single-object symmetrical displays of Experiment 1, but not for the two-object corresponding displays in Experiment 2). This interpretation would fit with Baylis and Driver's (1993) previous proposals that routine shape description in terms of the relative layout of convex parts only operates within objects, and does not directly specify the relative layout of parts from separate objects.

However, it might be that, despite the requirement for a speeded response, some strategy of making effortful (possibly point-by-point) comparisons for the single-object repeated displays could somehow have carried over to the corresponding two-object displays, since these two formats were randomly intermingled within each block of trials. This possibility seems somewhat unlikely, especially given that Experiment 1 had found entirely different effects from the number of steps upon judgements for two formats that were randomly intermingled (i.e., the number of steps affected the two-object format much more than the single-object format in that symmetry experiment). Nevertheless, we considered such strategic possibilities for the present repetition experiment to merit further investigation, especially since the substantial effect from the number of steps upon even the corresponding repetition condition seems of considerable theoretical importance. If this result could be confirmed, it would imply that observers cannot efficiently compare all the convex parts of two separate objects in parallel, even when the parts correspond, despite the fact that they can make such comparisons in parallel for corresponding parts within a single object.

In the next experiment, participants again judged whether contours were repeated or not. In contrast to Experiment 2, only the corresponding two-object format was now used throughout. This was done to maximize the salience for observers of the fact that the convex parts of the two contours matched exactly whenever they were repeated; and also to remove any necessity of switching between the different perceptual strategies that may have been appropriate for the single-object displays versus the corresponding displays in Experiment 2. Our question was whether observers would now be able to judge repetition efficiently between two objects, regardless of the number of discontinuities to be considered along each contour (thus producing an efficiently "parallel" result like that found for symmetry within one object in Experiment 1). This was encouraged by making it clear to observers that the convex parts of the contours would always match exactly whenever they were repeated. In this way, we hoped to optimize the observers' comparison of the two separate objects presented in the corresponding condition.

If Baylis and Driver (1993) were correct in proposing that the relative layout of convex parts is only routinely coded in parallel within individual shapes, then performance should still be substantially affected by the number of discontinuities to be compared, since a between-object comparison is required in the present task. On the other hand, if all that matters is whether convex parts consistently match when the two contours are repeated, then we might now find little if any effect of contour complexity, as for the within-object symmetry judgements in Experiment 1.

Method

Unless otherwise stated, the method followed that of Experiment 2.

Participants. The 10 new participants, 6 female and 4 male, were recruited as before.

Apparatus and materials. The only difference in the task materials from Experiment 2 was that only the corresponding format (e.g., Figures 5c and 5d) was now used throughout; the previous single-object format was dropped. There were now 5 blocks of 130 trials, with the first block discarded as practice. Thus, 512 trials were available for analysis.

Results and discussion

The averaged results are shown in Figure 7, with the means of participants' median RTs in the upper graph, and the error rates below. It is immediately apparent that, as in Experiment 2, repetition judgements became dramatically slower as the number of discontinuities along the contours to be compared was increased, even though the convex parts of the jagged contours now always matched whenever they were repeated.

A two-way within-subjects ANOVA on the RT data showed no overall effect of repetition, $F(1, 9) = 3.2$, n.s.), but a highly significant effect of the number of steps, $F(2, 18) = 28.0$, $p < .001$, with much slower responses when there were more steps; plus a significant interaction of the two factors, $F(2, 18)$

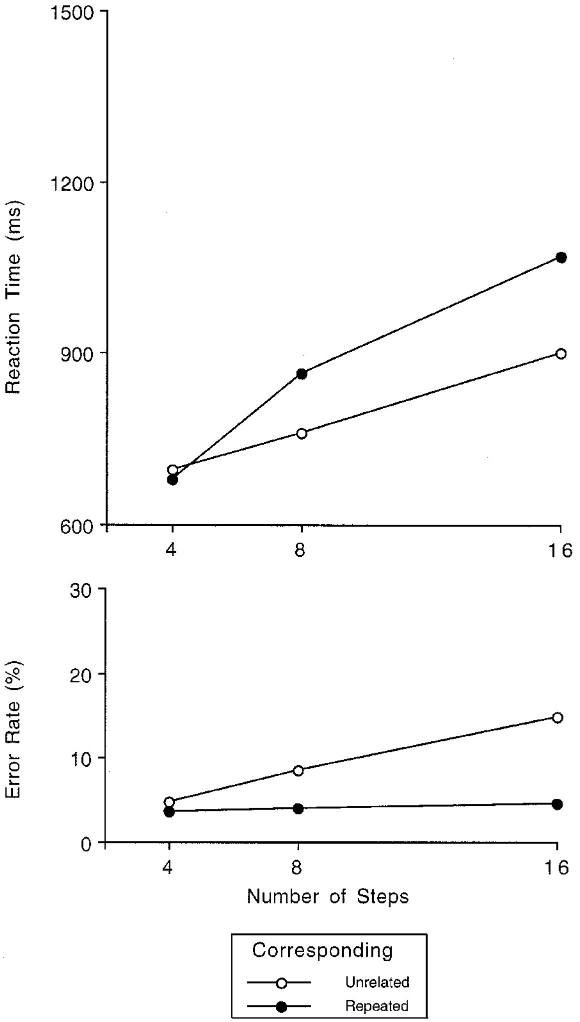

Figure 7. Means of participants' median RTs (top) and error rates (bottom) for the different display-types in Experiment 3.

= 5.7, $p < .05$, because the effect of steps was somewhat more pronounced for repeated responses. The effect of increasing the number of steps was similar to that seen in Experiment 2. The slope was 17.2 ms/step (compared to 23.5 in Experiment 2) for unrelated contours, and 31.5 ms/step (compared to 29.1 in Experiment 2) for repeated contours.

An analogous two-way ANOVA on the error data showed the following pattern of results. There was a significant effect of relatedness, $F(1, 9) = 31.8$, $p < .001$, with more accurate responses for repeated shapes overall, and of the number of steps, $F(2, 18) = 20.9$, $p < .001$, with more errors when there were more steps. The interaction was also significant, $F(2, 18) = 9.0$, $p < .01$, with the effect of steps now being more pronounced for unrepeated responses, whereas the opposite had applied in the RT data. This minor discrepancy between the RT and error data suggests caution in interpreting the two-way interactions, which may be partly due to speed–accuracy trade-offs. However, such possible trade-offs do not undermine the main finding of this study, namely the robust detrimental effect of increasing the number of steps. This effect of steps appears to be as large as that found for the same corresponding condition in Experiment 2. However, comparison of Figures 6 and 7 suggests that overall RTs may be somewhat lower in the present study, perhaps because only one format was now presented, whereas the corresponding format and single-object format had been randomly intermingled in Experiment 2. In order to test this, a further between-experiment analysis were carried out on the RT and accuracy data.

Three-way mixed design ANOVAs (Experiment × Repetition × Number of steps) were used to compare the data from Experiment 3 with the data from the corresponding-format condition in Experiment 2. The ANOVA on the RT data showed a significant main effect of experiment, $F(1, 24) = 8.5$, $p < .01$, with faster responses overall in Experiment 3, but no interaction involving this factor approached significance. A similar ANOVA for the error data showed no effect of experiment, and no interactions involving this factor ($F < 1$ in all cases). Thus, we may conclude that the critical influence of the number of steps was equivalent in Experiment 2 and 3.

The results of this experiment thus confirm that judgements of repetition between two contours cannot be made in the highly efficient parallel fashion found for symmetry within one shape (cf., Experiment 1), even when the convex parts of two repeated contours (on two separate shapes) correspond. Moreover, the substantial decline in performance as the number of discontinuities along the contours increased was still found even though the corresponding format, with its matching convex parts for repeated contours, could now be anticipated on every single trial.

Note that the present decline in performance when an increased number of discontinuities has to be considered cannot be attributed to mere acuity limits becoming more critical as the number of steps increases, nor to any changes in the spatial frequency composition of the steps themselves. All such factors were shared with the equivalent manipulation of steps in Experiment 1, where symmetry judgements within objects were scarcely affected by varying the number of steps along the contours. Thus, the difficulty with increased contour complexity in the present experiment is psychologically meaningful, rather than being trivially due to low-level visual factors, as shown by the contrasting result found for within-object symmetry judgements in Experiment 1 (compare the pair of RT lines in Figure 7, to the lower pair of RT lines in Figure 4).

However, this difference in outcome might be attributed to either of two factors which distinguished the efficient conditions of Experiment 1 from the present highly inefficient conditions. First, Experiment 1 involved symmetry judgements, whereas the present task involved repetition detection. Perhaps performance can be efficient only when convex parts correspond *and* symmetry rather than repetition is judged? A second potentially important difference is that the present corresponding repetition condition always involved a comparison of contours from two separate objects, whereas the efficient parallel conditions in Experiment 1 involved symmetry judgements within just a single object.

On our own account, it should be the latter factor that is critical in determining whether a substantial effect from the number of steps is found. According to Baylis and Driver (1995), the fundamental difference between symmetry versus repetition judgements for contours lies in whether or not convex parts match; and convex parts matched for both the conditions we are currently considering (i.e., the single-object symmetry condition from Experiment 1, see Figure 3a; versus the corresponding repetition condition from Experiments 2 and 3, see Figure 5c). Thus, the dramatic difference between these two conditions, in terms of the effect produced by the number of steps, cannot be attributed to whether or not convex parts match. Instead, we believe it is due to whether or not convex parts can be related within a single shape-description for a common object (as in Figure 3a), or must be compared across distinct objects which each have their own separate shape description in terms of convex parts (as in Figure 5c).

If our account is correct, then even symmetry judgements for contours in two separate objects should depend substantially on the number of steps, even when the convex parts on the two objects match. This prediction was tested in our next experiment.

EXPERIMENT 4

This experiment examined whether judgements of symmetry always proceed in an efficiently parallel manner across the display as a whole, provided convex parts correspond for the contours that must be compared; or whether such a match based on the corresponding convex parts of mirror-symmetric contours can only occur efficiently within single objects.

We returned to the 100 per cent symmetrical versus 25 per cent asymmetrical contours used in Experiment 1. As in that experiment, one of the possible display types was the non-corresponding two-object format (see Figures 8a and 8b), which obscures symmetry by using figure-ground segmentation to turn convexities along one contour into concavities along the other, and vice-versa. Recall that in Experiment 1, we had found that symmetry judgements for these displays are inefficient and substantially affected by the number of steps. The

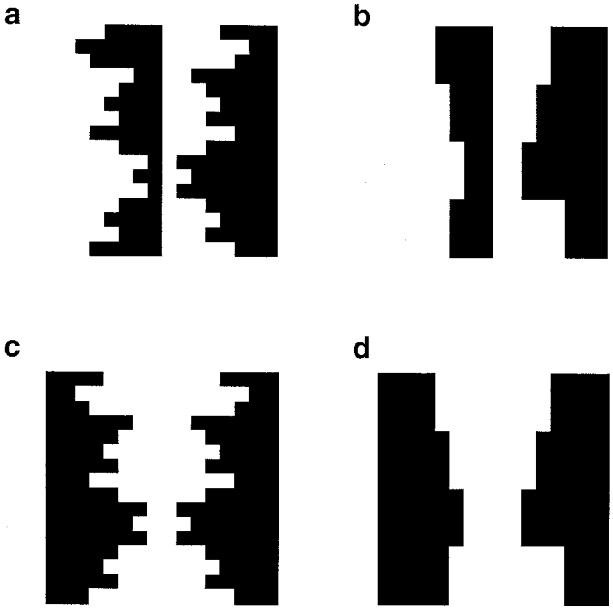

Figure 8. Four examples of displays from Experiment 4. Top: Examples of displays for the non-corresponding condition: a symmetric display (a) and an asymmetric display (b), with 16 or 4 steps respectively along each elongated side. Bottom: Examples of displays for the outer condition: a symmetric display (c) and an asymmetric display (d) with 16 or 4 steps respectively along each elongated side.

other display type was a new format; outer two-object displays were formed by blocking with colour the regions that lay immediately outside each jagged contour (see Figures 8c and 8d). The important point about such displays is that the convex parts for symmetric contours still correspond, as in the original single-object symmetric condition from Experiment 1 (cf. Figure 3a). However, these convex parts now get assigned by bottom-up figure-ground segmentation to two separate objects in the outer format.

If convex parts can be efficiently compared in parallel whenever symmetry must be judged for matching convexities (as in the efficient single-object condition in Experiment 1), then performance in the new outer format should be scarcely affected by the number of steps. However, if Baylis and Driver (1993) were correct in proposing that the relative layout of convex parts can only be efficiently judged within a single object, then the new outer format should lead to inefficient performance, declining as the number of discontinuities to be considered increases, since parts from two separate objects have to be compared for the outer format.

Note that the latter prediction of inefficient performance for the outer format is, in many respects, counterintuitive from any perspective other than the notion of within-object shape descriptions that we have been advocating. For instance, there is, if anything, even more reflective symmetry and greater redundancy within the image for the present outer format (Figure 8c) than for the single-object symmetric format from Experiment 1 (Figure 3c), due to the added symmetry between the additional outer contours (i.e., the two straight vertical lines) in the outer format. Thus, if the ease of the single-object condition from Experiment 1 were merely due to some combination of matching convex parts plus the special perceptual redundancy that may be inherent in mirror symmetry, then the present outer format should be at least as easy, and perhaps even more efficient due to the added redundancy. However, we predicted that the outer format should be inefficient, with performance declining sharply against the number of steps. This counterintuitive prediction follows directly from Baylis and Driver's (1993) proposal that the relative layout of convex parts is only available in an efficiently parallel manner within individual objects, by virtue of the standard shape description that we propose to be derived for each object.

Note also that since the present experiment randomly intermingled the new outer format with the non-corresponding format, the conditions of presentation for the outer format were identical to those previously used for the single-object format in Experiment 1. This allowed us to make a meaningful between-experiment comparison of performance for the single-object symmetric format of Experiment 1 against the new outer format.

Method

Unless otherwise stated, the method followed that of Experiment 1.

Participants. The 15 new participants, 9 female and 6 male, were again undergraduates with normal or corrected acuity by self-report.

Apparatus and materials. Example displays are shown in Figure 8. The task was as in Experiment 1, namely to respond by pressing one of two buttons with one or other index finger as rapidly as possible, depending on whether the jagged contours in the display were perfectly symmetrical (e.g., Figures 8a and 8c) or not (e.g., Figures 8b and 8d). Half the displays were blocked out with colour like those in Figures 8a and 8b, to yield the non-corresponding format, in which convex parts for symmetric contours do not correspond. The other half were blocked out to yield the new outer format (see Figures 8c and 8d) for which convex parts correspond between symmetric contours on separate objects. The lateral position of the two jagged contours was the same (on average) across these two formats, as was the width of each individual red shape.

Within the two formats, a random half of the displays had perfectly symmetrical jagged edges, while half were 25 per cent asymmetric.

Results and discussion

The average results are shown in Figure 9, with the means of participants' median RTs in the upper graph, and the error rates plotted below. It can be seen that performance declined substantially as the number of steps was increased for the non-corresponding format (as found in Experiment 1); and that this also applied for the new outer format, in accordance with our prediction. The effect of increasing the complexity in the non-corresponding format was similar to

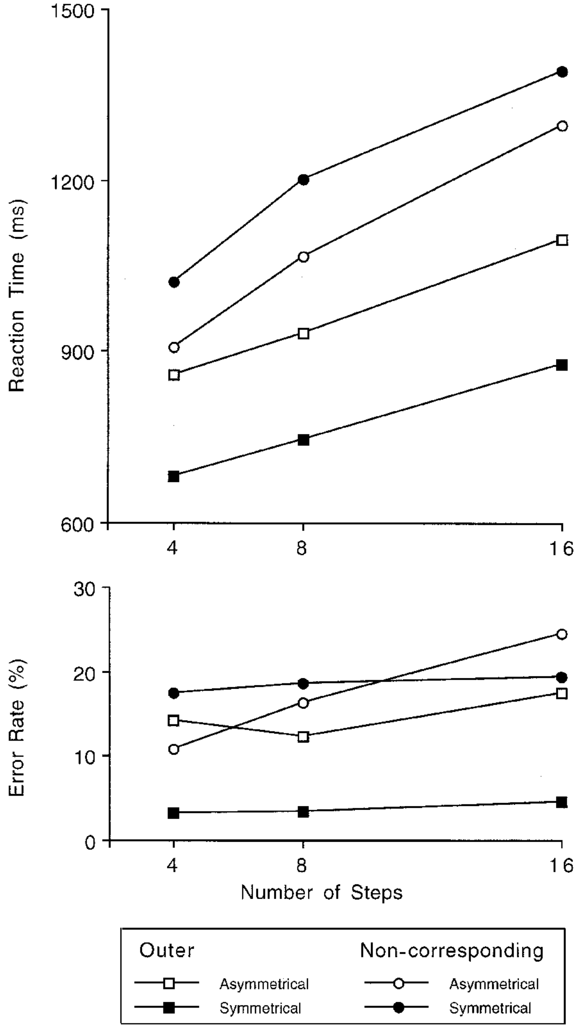

Figure 9. Means of participants' median RTs (top) and error rates (bottom) for the different display types in Experiment 4.

that seen for judgements of repetition in Experiment 2. The slope was 30.0 ms/step for asymmetric displays, and 32.1 ms/step for symmetric displays. The slopes for the outer conditions, while somewhat lower, still did not approach the shallow slopes seen for single object symmetry judgements in Experiment 1. For asymmetric displays the slope was 20.0 ms/step, whereas for symmetric displays the slope was 16.1 ms/step.

A three-way within-subjects ANOVA on the RT data showed a significant effect of display type, $F(1, 14) = 15.9$, $p < .001$, with faster responses for outer displays overall, no effect of symmetry, $F(1, 15) = 2.1$, n.s., but a highly significant effect of the number of steps, $F(2, 28) = 31.1$, $p < .001$, with much slower responses when there were more steps. There was a significant interaction of display type with symmetry, $F(1, 14) = 18.1$, $p < .001$, with faster response to symmetric outer, but to asymmetric non-corresponding displays. No other interactions approached significance.[1]

An analogous three-way ANOVA on the error data showed a similar pattern of results. There was significant effects of display type, $F(1, 14) = 34.0$, $p < .001$, with more accurate responses for outer displays, of relatedness, $F(1, 14) = 9.6$, $p < .01$, with more accurate responses for symmetrical shapes, and of the number of steps, $F(2, 30) = 14.7$, $p < .001$, with more errors when there were more steps). The interaction of symmetry with number of steps was also significant, $F(2, 28) = 7.0$, $p < .01$, with a sharper increase in errors for asymmetric responses. None of the other interactions were significant.

Thus, increasing the number of discontinuities along the contours that had to be compared always impaired performance substantially, even for the outer format, where convex parts matched for symmetric contours. Clearly, the presence of mirror symmetry between matching convex parts is insufficient, on its own, to allow efficient parallel consideration of all the discontinuities in the contours. It is evidently also vital that the symmetrical convex parts should belong to a common object (as for the single-object condition in Experiment 1; see Figure 3c) so that their relative layout may be assessed by means of the routine shape description we propose.

To confirm the difference in outcome for the present outer format, versus the single-object condition from Experiment 1, we ran a between-experiment analysis to compare the effect from the number of steps for these two formats directly. As noted earlier, this is a legitimate comparison, since the outer format and the single-object format had both been presented in an identical context (i.e., with randomly intermingled non-corresponding-format trials). A mixed ANOVA on the RT data had the between-subject factor of experiment (i.e., single-object data from Experiment 1, versus outer condition data from Experiment 4), and the within subject factors of the number of steps, and of symmetrical versus asymmetrical displays. There was a main effect of experiment, $F(1,$

---

[1] Once again, no terms were reliable in an ANOVA on SEMs as a ratio with overall RT.

26) = 5.8, $p < .04$, with slower RTs for the outer condition of Experiment 4, a main effect of the number of steps, $F(2, 52) = 25.1$, $p < .001$, with slower RTs when there were additional steps, and an interaction of experiment and symmetry, $F(1, 26) = 25.7$, $p < .001$, with symmetrical responses being very much faster than asymmetric in the outer condition of Experiment 4. Critically, there was also a significant interaction between experiment and the number of steps, $F(2, 52) = 9.4$, $p < .001$), confirming the important finding that RT increased substantially more with additional steps in the outer condition of Experiment 4 than in the single-object condition of Experiment 1 (compare the slope of the RT function depicted by square symbols in Figure 9 with that for the square symbols in Figure 4).

This analysis confirms that the outer format of Experiment 4 led to significantly less efficient judgements of symmetry than the single-object format from Experiment 1, with a significantly larger effect from the number of discontinuities that must be considered along each contour. In both formats, convex parts matched for symmetric contours, so this alone is clearly insufficient to ensure efficiently parallel performance. The observed difficulty with the outer format offers specific support for our hypothesis that the layout of multiple convex parts can only be efficiently accessed in parallel by means of a shape description that represents all parts within a single object.

A further mixed ANOVA compared the data from the non-corresponding conditions of the two experiments. No term approached significance, confirming that the results for the non-corresponding condition were unaffected by whether the outer format or the single object format had been randomly intermingled with it.

The difficulty in judging symmetry in the outer condition may seem surprising. After all, if a display like Figure 8c is inspected, it seems possible to effortfully reverse figure and ground, such that a single central white object may be seen. If such a single "virtual" central object could ever be perceived, then one might expect that symmetry detection would then proceed in a more efficiently parallel manner, as for the single-object condition in Experiment 1. The possibility of such occasional figure-ground reversals may have contributed to the faster judgements for outer displays than non-corresponding displays. This possibility might particularly apply for the symmetrical outer displays, as a central symmetrical region should be easier to assign as figural than an asymmetrical region, even when other factors (here, the contrast polarity of the shapes and their surround) are biased against this (Bahnsen, 1928; Driver et al., 1992). Indeed the symmetrical outer displays were judged very much faster than the asymmetric outer displays (see Figure 10), suggesting a possible contribution from such reversals of the experimentally intended figure-ground assignment, on at least some trials.

In order to further explore this possibility, in the next experiment we introduce a preview display that should specifically encourage such figure-ground reversal for the outer displays, on a larger proportion of trials. By preceding outer displays with a homogeneous, filled-in red rectangle, the curved contours were then generated by making the central region offset (i.e., reverting this to the black of the computer screen) to produce the outer display. This central transient should somewhat increase the ease with which the central region may be seen as a single (virtual) object, on at least some trials. In such cases, we would predict that judgements of symmetry should become faster and less affected by the complexity of the contours to be judged (i.e., the number of steps), since on our account a more efficient symmetry detection can proceed within single objects.

EXPERIMENT 5

Method

Unless otherwise stated, the method followed that of Experiment 4.

Participants. The 9 new participants, 6 female and 3 male, were again undergraduates with normal or corrected acuity by self-report.

Apparatus and materials. The only difference in terms of materials from experiment 4 was that each experimental display was now immediately preceded by a homogenous red rectangle, presented for 500 s, and covering the full extent of the pair of shapes which were subsequently shown (e.g., see Figures 8c and 8d), including the central region which was subsequently offset to define and separate those shapes on the black computer screen. This had the effect of making the central black region appear more like a single illusory object.

Design. The only other difference from Experiment 4 was that only displays in the outer format were shown. There were now 5 blocks of 130 trials with the first block discarded as practice. Thus, 512 trials were available for analysis.

Results and discussion

The average results are shown in Figure 10, with the means of participants' median RTs in the upper graph, and the error rates plotted below. A two-way within-subjects ANOVA on the RT data showed a significant effect of symmetry, $F(1, 15) = 6.0$, $p < .05$, with faster responses for symmetric contours, plus an effect from the number of steps, $F(2, 30) = 119.6$, $p < .001$, with slower responses when there were more steps. The interaction did not approach significance. An analogous two-way ANOVA on the error data showed a similar pattern of results. There was a significant effect of symmetry, $F(1, 15) = 6.0$, $p <$ .05, with faster responses for symmetric contours, plus a significant effect from

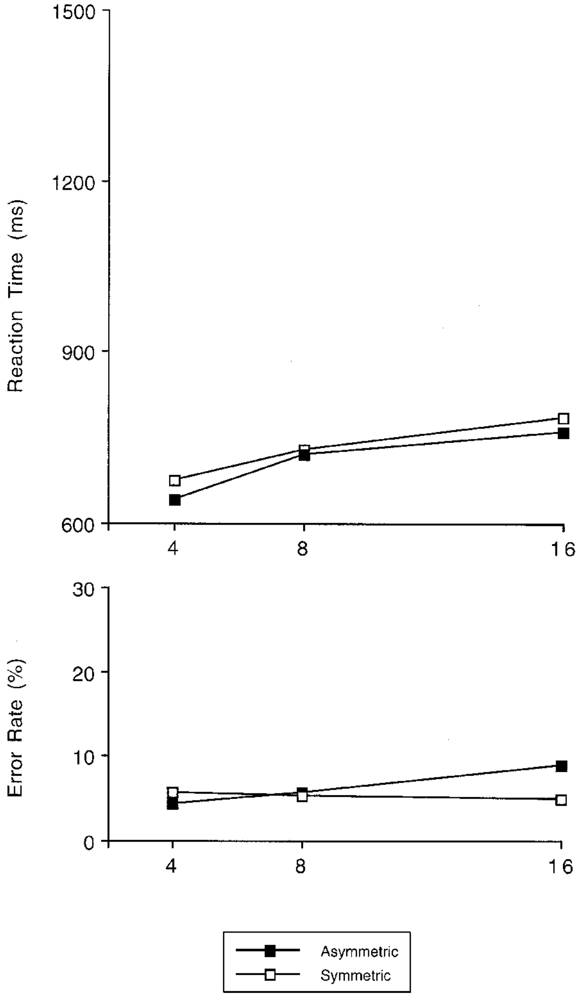

Figure 10. Means of participants' median RTs (top) and error rates (bottom) for the different conditions in Experiment 5.

the number of steps, $F(2, 30) = 119.6$, $p < .001$, with more errors when there were more steps. The interaction did not approach significance.

We compared the effect from the number of steps in this study, with that in Experiment 4, by a three-way mixed design ANOVA, considering the data from this experiment, and those for the analogous outer conditions of Experiment 4. There was no effect of experiment, $F(1, 22) = 1.4$, $p < .1$, but a significant interaction of experiment with symmetry, $F(1, 22) = 16.0$, $p < .001$, as judgements of asymmetric displays improved more in the present experiment (relative to Experiment 4) than judgements of symmetrical displays (perhaps because only the symmetrical displays could ever be figure-ground reversed in Experiment 4, as discussed earlier). Critically, there was a significant interaction of number of steps with experiment, $F(2, 44) = 3.1$, $p < .05$, because judgements of symmetry were less affected by increasing steps in Experiment 5 (slopes of 8.9 ms/step for asymmetric displays, 9.2 ms/step for symmetric displays in Experiment 5, compared with 20.0 and 16.1 respectively). The three-way interaction was non-significant, $F(2, 44) < 1$. A similar three-way ANOVA on the error data showed the same pattern, with no effect of experiment, $F(1, 22) = 2.1$, $p < .1$, an interaction of experiment with symmetry, $F(1, 22) = 10.1$, $p < .001$, and of steps with experiment, $F(2, 44) = 2.9$, $p < .06$, but no three-way interaction, $F(2, 44) < 1$.

These results suggest that making the central region offset (so that this changed region then appeared somewhat more figural) reduced the effect of increasing number of steps for outer displays, so that it now became more like Experiment 1, in the condition where judgements of single central objects had been made. Indeed the slopes approached those seen for such single-object judgements in Experiment 1. For asymmetric displays the slope was 8.9 ms/step (compared to 3.7 in Experiment 1), whereas for symmetric displays the slope was 9.2 ms/step (compared to 4.6).

Encouraging the percept of a single central object, by means of a preview rectangle, thus led to more efficient detection of symmetry than when judgements had to be made for obviously separate objects, as presented in Experiment 4. However, the efficiency of symmetry judgement was still not quite as high as in Experiment 1, where explicit single objects were judged, presumably as the percept of a single central object was less compelling in the "offset" conditions of Experiment 5. These results support the notion that as the tendency to assign the critical curved edges to a single central figure increases, then judgements can be made in an increasingly efficient parallel manner, which encompasses that figure as a whole.

GENERAL DISCUSSION

This series of experiments replicates and extends our previous work (Baylis & Driver, 1994, 1995a) on symmetry and repetition perception for visual shapes, and provides several new insights into the underlying mechanisms. These involve the nature of the decomposition into component parts which takes place for visual shapes, and the conditions under which the described component parts can all be available to the observer simultaneously, in an efficiently parallel manner.

Experiment 1 examined symmetry perception as a function of the number of discontinuities along the border of outline 2D shapes, and of whether figure-ground segmentation yielded matching convex parts for the symmetric contours (whereby a convexity on one contour corresponded to a convexity on the other contour) or instead yielded mismatching convex parts (whereby, for the same two symmetric contours, a convexity on one contour now became a concave region along the other contour). In agreement with Baylis and Driver (1995a), symmetry detection was much more efficient when convex parts matched between related contours, providing further evidence for the psychological reality of the convex-part scheme for shape perception that was originally proposed by Hoffman and Richards (1984; see also Biederman, 1987; Driver & Baylis, 1995). These results also provide further support for Baylis and Driver's proposals (1995a, b; Driver & Baylis, 1995, 1996) that edge comparisons tend to be based upon perceived shapes following figure-ground assignment, rather than upon any more literal representation of edges prior to their figural assignment.

In agreement with Baylis and Driver (1994), Experiment 1 also found that symmetry judgements within a single object were scarcely affected by the number of discontinuities along the contours that had to be compared, suggesting availability of all the convex parts of a single shape simultaneously. By contrast, when figure-ground segmentation assigned the symmetric contours to separate objects with mismatching parts the very same manipulation in the number of discontinuities led to a substantial decline in performance as more convex parts were added. This breakdown in efficient parallel processing might be attributed to the mismatch in convex parts, to the requirement of comparing two separate objects, or to a combination of these factors. Our subsequent experiments sought to disentangle these.

Experiment 2 used analogous displays to Experiment 1, except that any related contours were now repeated rather than symmetric. In replication of Baylis and Driver (1994), it was found that judgements of repetition between the two sides of a single object declined substantially as the number of discontinuities along the sides was increased (even though the acuity requirements and so on were exactly the same as for the case that had produced efficient parallel detection of symmetry in Experiment 1). In agreement with Baylis and Driver (1995a), repetition judgements became easier overall when figure-ground segmentation assigned the repeated contours to separate objects with matching convex parts, rather than to a single object with mismatching parts. However, the novel finding was that even in the two-object case with matching parts, performance still declined as the number of parts to be considered was increased. This result for repetition detection was replicated in Experiment 3, when the format of two separate objects with matching parts could be anticipated on every trial. These findings suggest that the critical factor determining whether an increased numbers of parts will impair performance is not simply whether convex parts match or not, but either: (1) whether convex parts match to produce mirror symmetry in particular, or (2) whether convex parts match within a single perceptual object. The final experiments sought to distinguish these two possibilities.

In Experiment 4, observers judged symmetry in displays where the convex parts always belonged to two separate objects, but either matched or mismatched for symmetrical contours in terms of convex parts. Overall performance was better with matching parts, as in the previous experiments, but performance still declined substantially when an increased number of component parts had to be considered. This implies that even a combination of matching convex parts, plus mirror symmetry, is insufficient to yield efficiently parallel performance across the number of parts. It is evidently critical that those parts must belong to a single common object, as for the symmetrical shapes in Experiment 1 which had yielded the most efficient parallel performance.

Experiment 5 provided further evidence that the difficulty faced by observers in Experiment 4 was due to the fact that the parts were segmented as belonging to separate objects in the displays. Here the displays to be judged were the same as in Experiment 4, but were now preceded with a preview display that should somewhat increase the tendency to perceive a single central figure. This led to a reduction in reaction time, and also a significant reduction in the dependence of reaction time on the number of steps. This suggests that the very same visual information can be processed in a more efficiently parallel manner when it is easier to parse as a single central figure than when parsed as two flanking figures instead (see also Baylis & Driver, 1993).

These results support the importance of figure-ground assignment and convex-part decomposition in determining the ease of edge comparisons, as we have argued for elsewhere (see Baylis & Driver, 1994, 1995a, b; Driver & Baylis, 1995, 1996). However, the present findings also illustrate the importance of whether or not the convex parts in question all belong to a common object, a factor which we had previously emphasized on the basis of an entirely different set of phenomena, involving relative-location judgements for the parts of non-symmetric and non-repeated shapes (see Baylis, 1994; Baylis & Driver, 1993). The present results suggest that a comparison of the layout of component parts can proceed in an efficiently parallel manner only within single objects. As argued by Baylis and Driver (1993), this may arise because the relative layout of component parts within an object is made explicit in the routine shape description that is normally derived for each object; but these routine shape descriptions do not specify the relative layout of parts from separate objects, and so additional (evidently time-consuming) operations are necessary to recover the latter between-object relations.

Individual shapes are essentially constituted by the relative layout of their component parts, and so one can advance the argument that any full appreciation of a shape must require parallel access to the relative location of all its component parts simultaneously. Our results suggest that the human visual system is indeed capable of this for individual objects; but not for separate objects. Relating the parts of separate objects to each other evidently requires effortful or inefficient comparisons, even for the very same set of convex parts that can be effortlessly related within one object.

REFERENCES

Bahnsen, P. (1928). Eine untersuchung uber Symmetrie une Asymmetrie bei visuellen Wahrnehmungen. *Zeitschrift fur Psychologie, 108,* 355–361.

Barlow, H.B., & Reeves, B.C. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research, 19,* 783–793.

Baylis, G.C. (1994). Visual attention and objects: Two-object cost with equal convexity. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 208–212.

Baylis, G.C. (1998). Visual parsing and object-based attention: A developmental perspective. In J. Richards (Ed.), *The cognitive neuroscience of attention: A developmental perspective.* Mahwah, NJ: Lawrence Erlbaum Associates Inc.

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

Baylis, G.C., & Driver, J. (1994). Parallel computation of symmetry but not repetition in single visual shapes. *Visual Cognition, 1,* 377–400.

Baylis, G.C., & Driver, J. (1995a). Obligatory edge-assignment in vision: The role of figure- and part-segmentation in symmetry detection. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 1323–1342.

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

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94,* 115–147.

Bornstein, M.H., & Krinsky, S.J. (1985). Perception of symmetry in infancy: The salience of vertical symmetry and the perception of pattern wholes. *Journal of Experimental Child Psychology, 39,* 1–19.

Bruce, V., & Morgan, M.J. (1975). Violations of symmetry and repetitions in visual patterns. *Perception, 4,* 239–249.

Corballis, M.C., & Roldan, C.E. (1974). On the perception of symmetrical and repeated patterns. *Perception and Psychophysics, 16,* 136–142.

Donnelly, N., Humphreys, G.W., & Riddoch, M.J. (1991). Parallel computation of primitive shape descriptions. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 561–570.

Driver, J., & Baylis, G.C. (1991). Target-distractor separation and feature integration in visual attention to letters. *Acta Psychologica, 76,* 101–119.

Driver, J., & Baylis, G.C. (1995). One-sided edge-assignment in vision. 2: Part decomposition, shape description, and attention to objects. *Current Directions in Psychological Science, 4,* 140–145.

Driver, J., & Baylis, G.C. (1996). Edge-assignment and figure-ground segmentation in visual matching. *Cognitive Psychology, 31,* 248–306.

Driver, J., Baylis, G.C., & Rafal, R. (1992). Preserved figure-ground segregation and symmetry detection in visual neglect. *Nature, 360, 73–75.*

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

Julesz, B. (1971). *Foundations of Cyclopean perception.* Chicago: University of Chicago Press.

Koffka, K. (1935). *Principles of Gestalt psychology.* London: Routledge & Kegan Paul.

Lowe, D.G. (1987). Three-dimensional object recognition from two-dimensional images. *Artificial Intelligence, 31,* 355–395.

Mach, E. (1959). *The analysis of sensation*. New York: Dover. (Original work published 1885).

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.

Pascal, B. (1950). *Pascal's pensées* (H.F. Stewart, Trans.). London: Routledge & Kegan Paul. (Original work published 1888).

Rabbitt, P.M.A. (1966). Errors and error correction in choice response tasks. *Journal of Experimental Psychology, 71,* 264–272.

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

Shannon, C., & Weaver, W. (1949). *The mathematical theory of communication*. Urbana: University of Illinois Press.

Treisman, A., & Gelade, G. (1980) A feature-integration theory of attention. *Cognitive Psychology, 12,* 97–136.

Tyler, C.W. (1996). *Human symmetry perception and its computational analysis*. Zeist, Holland: VSP.

Wagemans, J. (1995). Detection of visual symmetries. *Spatial Vision, 9*(1), 9–32.

Wilkinson, L. (1990). *SYSTAT: The system for statistics*. Evanston, IL: Systat.

Manuscript received October 1997
Revised manuscript received March 2000