

Early computation of part structure: Evidence from visual search

YAODA XU

Harvard University, Cambridge, Massachusetts

and

MANISH SINGH

Rutgers University, New Brunswick, New Jersey

The visual system represents object shapes in terms of intermediate-level parts. The minima rule proposes that the visual system uses negative minima of curvature to define boundaries between parts. We used visual search to test whether part structures consistent with the minima rule are computed *preattentively*—or at least, rapidly and early in visual processing. The results of Experiments 1 and 2 showed that whereas the search for a non-minima-segmented shape is fast and efficient among minima-segmented shapes, the reverse search is slow and inefficient. This asymmetry is expected if parsing at negative minima occurs obligatorily. The results of Experiments 3 and 4 showed that although both minima- and non-minima-segmented shapes pop out among unsegmented shapes, the search for minima-segmented shapes is significantly slower. Together, these results demonstrate that the visual system segments shapes into parts, using negative minima of curvature, and that it does so rapidly in early stages of visual processing.

A basic problem in object perception is segmenting retinal arrays into regions that are likely to have arisen from coherent objects—in other words, the formation of *perceptual units*. Perceptual units exist at many levels, however. The objects and surfaces we see are not unstructured wholes; they themselves have further part structure. When we look at a chair, for example, we see not only a coherent object, but also a spatial arrangement of clearly defined parts: a back, a seat, and four legs. Perceptual units thus exist at the level of whole objects, at the level of parts, and possibly at the level of smaller parts nested hierarchically within larger ones (Marr & Nishihara, 1978; Palmer, 1977).

In principle, although any physical subset of an object may be considered to be its part, only some of these “parts” are perceptually natural. Parts, like objects (Spelke, 1994), must be cohesive, connected units¹: The bottom inch of a chair’s leg along with the top half of its back hardly constitutes a natural part. Connectedness by itself is not sufficient, however. For instance, the upper half of a leg along with a small adjoining portion of the seat is a connected piece. But it seems rather contrived to consider it a *part* of

the chair. What stimulus features, then, serve to characterize perceptually natural parts?

Hoffman and Richards (1984) noted that whenever two independent parts connect or interpenetrate to form a complete object (as occurs with manufactured objects; Figure 1A) or when a part grows out of a body (as occurs with biological objects; Figure 1B; see Leyton, 1992), the boundaries between these parts typically lie in negative minima of curvature.² Intuitively, these are points of a high (locally maximal) magnitude of curvature that lie in concave regions of the shape. Hoffman and Richards thus formulated the *minima rule: Human vision segments shapes by using negative minima of curvature as boundaries between parts*.

A number of experimental studies have demonstrated the perceptual reality of part-based representations (Biederman, 1987; Biederman & Cooper, 1991; Driver & Baylis, 1995) and, in particular, of the minima rule. Indeed, researchers have used the minima rule to explain a variety of visual phenomena, including the perception of shape similarity (Attneave, 1971; Hoffman, 1983), short-term memory for shapes (Braunstein, Hoffman, & Saidpour, 1989), the perception of symmetry and repetition (Baylis & Driver, 1994, 1995; Driver & Baylis, 1995, 1996), the assignment of figure and ground (Hoffman & Singh, 1997), the perception of transparency (Singh & Hoffman, 1998), and visual search (Hulleman, te Winkel, & Boselie, 2000; Wolfe & Bennett, 1997).

In this paper, we will address the question of whether parsing at negative minima of curvature occurs *preattentively*—or at least, rapidly and early in visual processing. Traditionally, the processing of certain visual *features* has been

Y.X. is supported by a McDonnell-Pew Investigator Initiated Grant and M.S. by NSF BCS-0216944. For many useful discussions, we thank Bart Anderson, Roland Fleming, Don Hoffman, and Mary C. Potter. We also thank Bart Anderson for his kind support at M.I.T. Correspondence concerning this article may be addressed either to Y. Xu, Vision Sciences Laboratory, Psychology Department, Harvard University, Cambridge, MA 02138 (e-mail: yaoda@wjh.harvard.edu) or to M. Singh, Department of Psychology, Rutgers University, New Brunswick Campus, 152 Frelinghuysen Road, Piscataway, NJ 08854-8020 (e-mail: manish@ruccs.rutgers.edu).

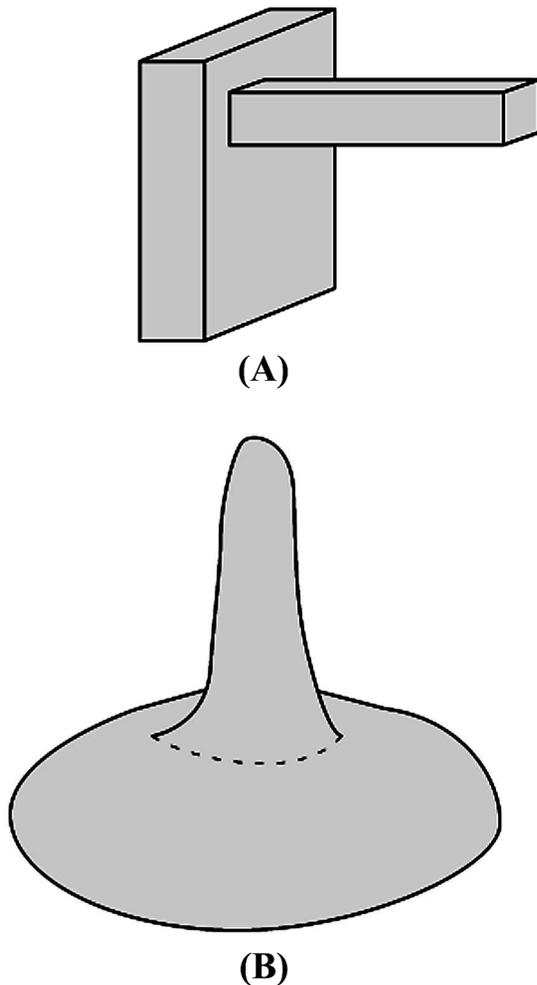


Figure 1. Demonstration of the motivation behind the minima rule: Whenever two independent parts connect or interpenetrate generically (A) or when a part grows out of a body (B), the boundaries between parts lie in negative minima of curvature. Intuitively, these are points of locally highest magnitude of curvature that lie in concave regions of the shape.

considered to be completely free of attentional demands—hence, *preattentive* (Neisser, 1967)—and asymmetry in visual search tasks has been taken to be an important sign of such processing (Treisman & Gormican, 1988; Treisman & Souther, 1985; Wolfe, 1994). In particular, if visual search for a target with Feature X is fast and efficient among distractors that lack Feature X but the reverse search is slow and inefficient, Feature X is considered to be computed preattentively.³ Some recent work has disputed whether *any* visual task can really be considered to be free from attentional demands (see, e.g., Joseph, Chun, & Nakayama, 1997). On the other hand, Wolfe's (1994) guided search model suggests that there are a limited number of preattentive features that can guide attention toward likely targets. Flat slopes (and *pop-out*) occur when the preattentive features can guide attention to the target on the first deployment of attention. If a single fea-

ture defines the target, that feature is a good candidate for a preattentive feature. Despite these differing views on preattentive processing, there is nevertheless a consensus that pop-out and search asymmetry are indicative of features that are computed rapidly and early in visual processing.

Some evidence for the early computation of negative minima has come from such visual search studies. Wolfe and Bennett (1997), for example, found that a target with a sharp negative minimum of curvature (Figure 2A) pops out among a set of distractors with no such negative minimum (Figure 2B), but not vice versa—hence, suggesting that negative minima are computed early and in parallel across the visual field. Using stimuli like those in Figure 3, however, Elder and Zucker (1993) found that the search for their target shape with a negative minimum (Figure 3A) did not quite reach the criterion for parallel search. There are two possible interpretations of this apparent inconsistency (Hulleman et al., 2000). First, it may be that negative minima have no special status in visual search and that the results of Wolfe and Bennett (1997) simply reflected the fact that the shape in Figure 2A has a point of high curvature, whereas the shape in Figure 2B does not. (Curvature has indeed been shown previously to be a basic feature; see Wolfe, Yee, & Friedman-Hill, 1992.) Second, it may be that negative minima do have a special status but that the results of Elder and Zucker did not reveal this because their curvature extrema had a low turning angle, and hence low visual salience⁴ (see, e.g., Hoffman & Singh, 1997).

To distinguish between these two possibilities, Hulleman et al. (2000) conducted visual search experiments using shapes such as those shown in Figure 4. The shape in Figure 4B has a sharp negative minimum of curvature,

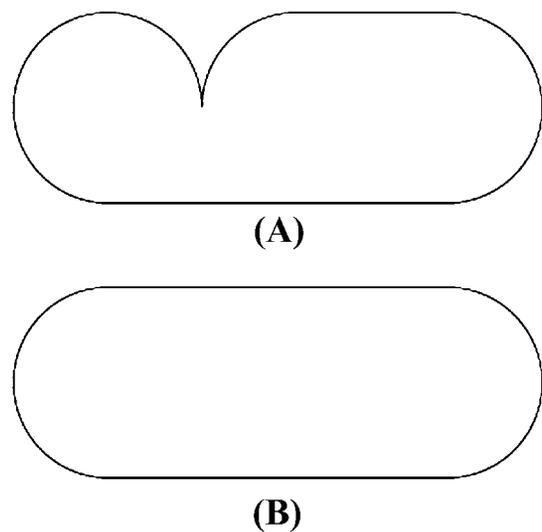


Figure 2. Stimuli adapted from the study by Wolfe and Bennett (1997). Their experiment showed that whereas the shape with the sharp negative minimum (A) pops out among distractors that lack a negative minimum (B), the reverse search is slow and inefficient.

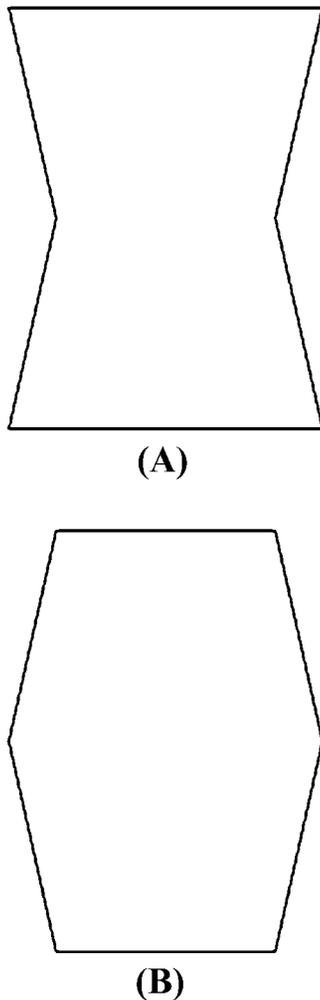


Figure 3. Stimuli adapted from the study by Elder and Zucker (1993). In their study, the search for the shape with the negative minimum (A) did not quite reach the criterion for pop-out. This indicates either that negative minima have no privileged status of visual search or that the curvature extrema on the shapes used in this study were perceptually too weak.

whereas the one in Figure 4A does not. However, the shape in Figure 4A does have an equally sharp point that lies in a convex region of the shape (i.e., a positive maximum of curvature). In a series of experiments, Hulleman et al. found that a shape with a negative minimum (such as Figure 4B) pops out among distractors, each with a positive maximum (such as Figure 4A), but not vice versa—even though the magnitude of curvature is the same in both extrema. They thus argued that negative minima (and, more generally, salient concavities) are computed early and in parallel across the visual field. More evidence in support of this interpretation has come from experiments by Humphreys and Müller (2000). They not only found a search asymmetry in favor of concave targets over convex targets, but also showed that this asymmetry can be reversed by switching the assignment of figure and ground.

Granting that *negative minima* are computed early in visual processing does not in itself imply, however, that part structures themselves are computed early. On a two-dimensional shape, negative minima are simply points that lie on the contour of the shape, whereas the computation of part structure requires computing *part cuts* that segment the shape into parts. In this paper, we will study the simple case of two-part shapes that can be segmented by joining two negative minima of curvature on the contour of the shape. In addition to negative minima, recent work has also pointed to the role of more global, region-based, geometric factors in determining perceived shape in general (Burbeck & Pizer, 1995; Kimia, Tannenbaum, & Zucker, 1995) and part structure in particular (Siddiqi & Kimia, 1995; Siddiqi, Tresness, & Kimia, 1996; Singh & Hoffman, 2001; Singh, Seyranian, & Hoffman, 1999). The method developed in this paper will allow us to study more complex cases as well. We will return to this issue in the General Discussion section.

Some evidence for the early computation of parts comes from judgments of symmetry and repetition. It has often been observed (at least as far back as Mach; see Baylis & Driver, 1994, 1995) that humans are more sensitive to symmetry within a pattern than to repetition. This observation has been confirmed in many experiments (e.g., Bertamini, Friedenber, & Kubovy, 1997; Bruce & Morgan, 1975; Corballis & Roldan, 1974; Friedenber & Bertamini, 2000). More recently, Baylis and Driver (1994) have argued that whereas symmetry within a shape can be detected in parallel, repetition seems to involve a serial process. They had participants judge whether or not a given shape was perfectly symmetric (see Figure 5A) or whether or not it was perfectly repeated (see Figure 5B). They found that, for symmetry judgments, response times (RTs) were hardly affected by the number of *steps* added to the sides of the shapes (RT slopes < 6 msec/step). How-

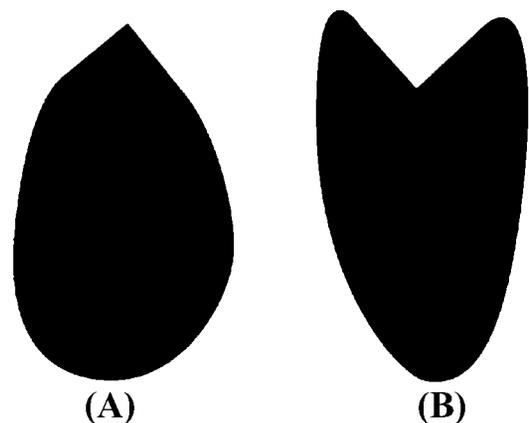


Figure 4. Stimuli adapted from the study by Hulleman, te Winkel, and Boselie (2000). Although each of the two shapes contains a sharp corner with the same turning angle, the shape with the negative minimum (B) popped out among distractors with the positive maximum (A), but not vice versa. This suggests that negative minima—and salient concavities, more generally—do have a privileged status in visual search.

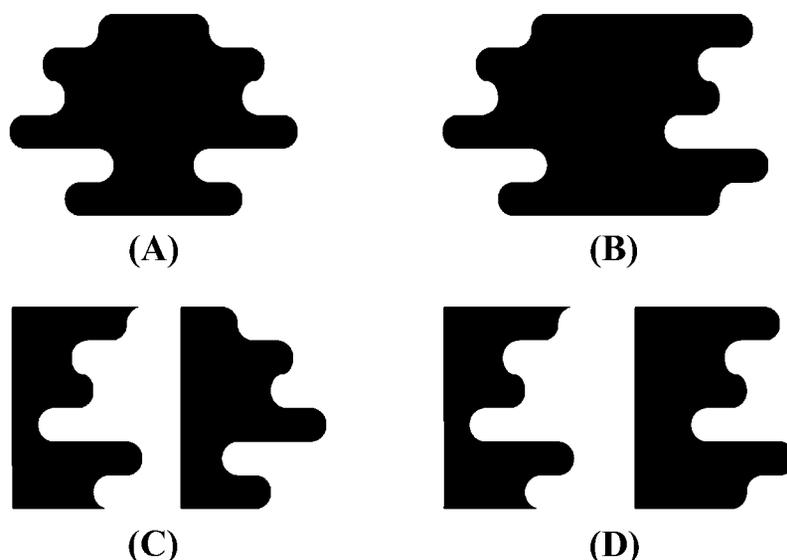


Figure 5. Stimuli adapted from the study by Baylis and Driver (1994, 1995). For the stimuli in the top row, judgments of symmetry (A) were faster and more accurate than judgments of repetition (B). Baylis and Driver argued that this is because the shape in A has matching parts, whereas the one in B has mismatching parts on the two sides. Consistent with this account, they found that the advantage for symmetry could be reversed if the symmetric contours resulted in mismatching parts (C) and the repeated contours resulted in matching parts (D).

ever, for repetition judgments, response times increased steadily with the number of steps (RT slopes > 26 msec/step). Baylis and Driver (1994) argued that this occurs because a symmetric shape has matching negative minima of curvature on the two sides—and hence, matching parts—whereas a repeated shape has mismatching parts (because concavities on one side correspond to convexities on the other). If the visual system represents shapes in terms of parts and compares them at the level of parts (rather than point by point), it would indeed make sense for symmetry to be an easier judgment than repetition. Moreover, if parts mediate the parallel detection of symmetry, they must themselves be computed in parallel.

As further evidence for their claim, Baylis and Driver (1995) demonstrated that the advantage for symmetry can be reversed by reversing the figure and ground relationships on one of the two curves (see Figures 5C and 5D). For these modified stimuli, they found that repetition became an easier judgment than symmetry, consistent with the fact that the repeated pattern (Figure 5D) now has matching parts, whereas the symmetric pattern (Figure 5C) has mismatching parts.

The results of Baylis and Driver (1994, 1995) thus provide some evidence for the early computation of parts. These results do not demonstrate, however, that perceived parts are necessarily delineated at negative minima of curvature. The symmetric shape in Figure 5A, for example, also has matching positive maxima of curvature (points of locally highest magnitude of curvature that lie in *convex* regions) and matching inflection points (points of

transition from convex to concave, and vice versa), and in principle, these could also have been used to make the symmetry judgments. Similarly, the repeated shape in Figure 5D has matching negative minima, as well as matching positive maxima. Thus, this methodology does not single out negative minima as the critical determiners of perceived parts. In addition, the results were obtained in the context of judgments of symmetry and repetition, so the experimental method does not allow one to test shapes that are *not* near-symmetric or near-repeated. It therefore remains unclear to what extent these results generalize to arbitrary shapes. Finally, this experimental method does not allow one to compare the perceptual naturalness of two different ways of segmenting a shape (e.g., by comparing part structures that are consistent with the minima rule with those that are not).

In the present study, our goal was to test directly whether part structures consistent with the minima rule are computed preattentively—or at least, rapidly and early in visual processing. In Experiments 1 and 2, we used visual search asymmetry, resulting from different parsings, to assess the early computation of part structure (Treisman & Souther, 1985). In Experiments 3 and 4, we compared visual search RTs for differently segmented shapes among unsegmented shapes and obtained further evidence for the early computation of perceived part structures. All the experiments used the standard visual search paradigm; however, unlike previous visual search studies, we used the same shape for targets and distractors, but segmented in two different ways.

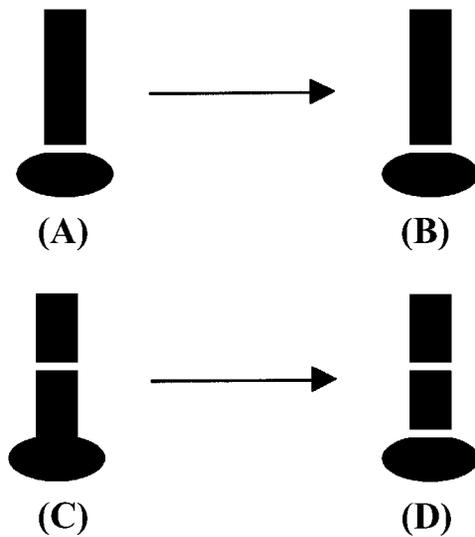


Figure 6. Stimuli used in Experiment 1: (A) minima-parsed shape; (B) no further parsing occurs for Shape A; (C) non-minima-parsed shape; and (D) the resulting shape if parsing at negative minima occurs automatically for the shape presented in panel C.

EXPERIMENT 1

In Experiment 1, participants searched for either a shape segmented at negative minima of curvature (Figure 6A) among distractor shapes parsed elsewhere (Figure 6C) or the reverse. Our logic was as follows: If the visual system segments shapes at negative minima of curvature, then the bottom “part” of the shape in Figure 6C would be segmented further into two parts, yielding the effective parsing shown in Figure 6D. As a result, the shape in Figure 6C would end up having an extra “feature” that the shape in Figure 6A does not have—namely, the cut in the rectangular part of the shape (or, equivalently, an additional part—three rather than two). Previous visual search studies have shown that when a target is defined by the presence of a unique feature, search is usually fast and efficient, whereas when the target is distinguished by the absence of a feature, search becomes slow and self-terminating (e.g., Treisman & Gormican, 1988; Treisman & Souther, 1985). If the visual system in fact segments shapes at negative minima, we should expect the search for a non-minima-segmented shape among minima-segmented shapes to be fast and efficient and the search for a minima-segmented shape among non-minima-segmented shapes to be slow and inefficient.

In designing the stimuli, we minimized differences along other dimensions that might distinguish between the natural and the unnatural parsings. For example, the lengths of the part cuts in the two parsings were equated, because cut length has previously been shown to be an important factor in parsing shapes (Singh et al., 1999). Moreover, in the context of visual search, the difference in the lengths of these line segments might itself provide a distinguishing

feature, quite independently of the fact that these line segments are part cuts. Similarly, the relative areas of the two parts created by these cuts were also equated: The area of the bottom part in Figure 6A is equal to the area of the top part in Figure 6C. Finally, the part cuts in both parsings had the same orientation (i.e., horizontal), so that any observed search asymmetry could not be attributed to a difference in orientation.

Method

Participants. Twelve participants, 7 males and 5 females, from the Massachusetts Institute of Technology were recruited. They were between 17 and 40 years of age, all had normal or corrected-to-normal vision, and they were paid for their participation.

Materials and Design. The shapes used in Experiment 1 are shown in Figures 6A and 6C. They are naturally seen as a bar attached to an oval, with either a cut in the middle of the bar (non-minima parsing) or a cut at the junction of the bar and the oval (minima parsing). The size of the (uncut) bar was $0.57^\circ \times 1.83^\circ$, and that of the oval was $0.57^\circ \times 1.15^\circ$. The thickness of the cut was 0.11° . The area of the oval in the minima-parsed shape was equal to the area of the detached bar piece in the non-minima-parsed shape. Hence, the two “parts” in both the minima-parsed and the non-minima-parsed shapes had equal relative areas.

Each display consisted of either 6 or 12 shapes distributed in a 5×3 grid ($9.38^\circ \times 9.72^\circ$). Each position was offset slightly from the grid so that the shapes were not perfectly aligned with each other. The experiment was run in two sessions. In one session, the participants searched for a non-minima-parsed shape among minima-parsed shapes (Figure 7A), and in the other session, they searched for a minima-parsed shape among non-minima-parsed shapes (Figure 7B). The order of the two sessions was counterbalanced across participants. Within each session, displays with set size 6 and set size 12 were intermixed randomly. These resulted in a total of 96 trials for each set size: 48 target-present trials and 48 target-absent trials. Trials from different conditions were divided evenly into three 64-trial blocks. Thirty-two practice trials preceded the experimental trials in each session.

Apparatus. The displays were presented on the 15-in. monitor of a 350-MHz G3 iMac, using the MacProbe experimental software. The participants were seated 50 cm away from the monitor.

Procedure. Each trial began with the presentation of a fixation dot for 500 msec, which was then followed by the search display. The search display remained on the screen until the participants had pressed one of the two prespecified keys to indicate the presence or absence of the target. The participants were instructed to use either their thumbs or their index fingers to press the response keys. The response keys were the Control key on the extreme left of the keyboard and the Enter key on the extreme right of the keyboard. Key assignments were counterbalanced across participants, and the assignment was indicated by appropriately labeled stickers on the keys. Within a block, the next trial began automatically 500 msec after the participant had responded to the previous trial. An incorrect response was indicated by a single short beep. The participants were allowed to take a break for as long as they wished between blocks.

Results

Reaction time for correct trials. Of the total data collected, 7.46% were removed owing to response errors. Two further data points were removed because the RTs were greater than 3,000 msec. The remaining data were analyzed with a within-subjects analysis of variance (ANOVA). The mean RTs averaged across participants are plotted in Figure 8.

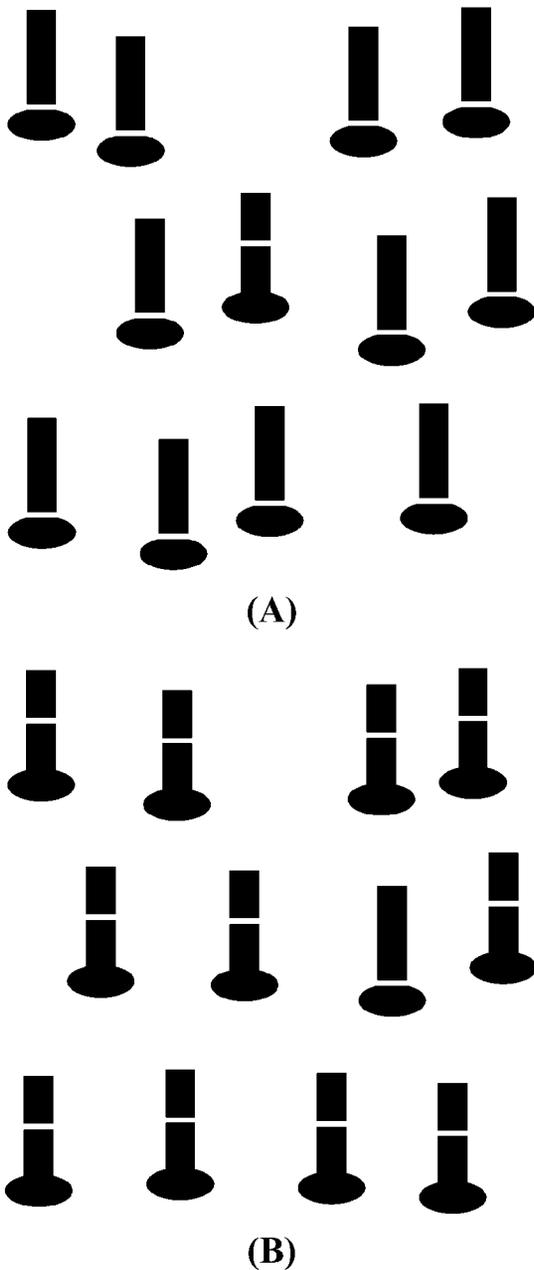


Figure 7. Examples of search displays used in Experiment 1. The participants searched either for a non-minima-parsed target among minima-parsed distractors (A) or for a minima-parsed target among non-minima-parsed distractors (B).

Overall, the two search types differed from each other significantly [$F(1,11) = 82.24, p < .001$], such that participants were much faster at detecting the presence or absence of a non-minima-parsed shape among minima-parsed shapes than they were at detecting a minima-parsed shape among non-minima-parsed shapes. The effect of set size was also significant [$F(1,11) = 51.87, p < .001$], and it interacted significantly with search type [$F(1,11) = 33.53, p < .001$]. The participants were much faster at detecting the presence, rather than the absence, of a target

[$F(1,11) = 51.96, p < .001$], and this effect interacted significantly with search type [$F(1,11) = 31.13, p < .001$]. In addition, the interaction between set size and target presence/absence, as well as the three-way interaction (among search type, set size, and target presence/absence) were both significant [$F(1,11) = 5.11, p = .045$, and $F(1,11) = 5.22, p = 0.043$, respectively].

A separate ANOVA was also carried out for each search type. When the participants searched for a non-minima-parsed shape among minima-parsed shapes, the main effects of set size and target presence/absence were both significant [$F(1,11) = 11.21, p = .006$, and $F(1,11) = 17.90, p = .001$, respectively]. However, the interaction of the two effects did not reach significance ($F < 1$). When the participants searched for a minima-parsed shape among non-minima-parsed shapes, the effect of set size, the effect of target presence/absence, and the interaction of the two were all significant [$F(1,11) = 54.31, p < .001$, $F(1,11) = 51.44, p < .001$, and $F(1,11) = 7.43, p = .020$, respectively]. These results show that the two types of search differed from each other qualitatively, with the search for a non-minima-parsed shape among minima-parsed shapes being consistent with a feature search and the reverse search being consistent with a serial self-terminating search (Wolfe, 1998).

Error rates. The mean error rates are presented in Table 1. The participants made more errors searching for a minima-parsed shape among non-minima-parsed shapes than they did in the reverse search [$F(1,11) = 6.73, p = .025$]. They also made more errors with larger set sizes [$F(1,11) = 16.99, p = .002$] and more errors when the target shape was present than when it was absent [$F(1,11) = 28.48, p < .001$].

Discussion

When the participants searched for a non-minima-parsed shape among minima-parsed shapes, the search slopes were very shallow for both target-present and target-absent trials (both were less than 10 msec/item: 5.5 msec/item for target-present trials, 6.7 msec/item for target-absent trials). Moreover, the ratio between target-present and target-absent slopes was significantly less than 2.0 ($F < 1$). These results thus fulfilled the criteria proposed by Wolfe (1998) for feature search and indicated that when a part cut occurred at a non-negative-minima location on the target shape, it was considered as a unique feature among distractors with part cuts at negative minima. The parsing from Figure 6C to Figure 6D must, therefore, occur rapidly and early in visual processing.

When the roles of the target and the distractors were reversed—such that the participants searched for a minima-parsed shape among non-minima-parsed shapes—search became slow and inefficient (19.3 and 32.2 msec/item for target present and absent, respectively). In this case, the obligatory parsing at negative minima in the distractor shapes made the cut in the target shape much less prominent. As a result, the target was now defined by the *absence* of a cut in the non-minima location. As is well

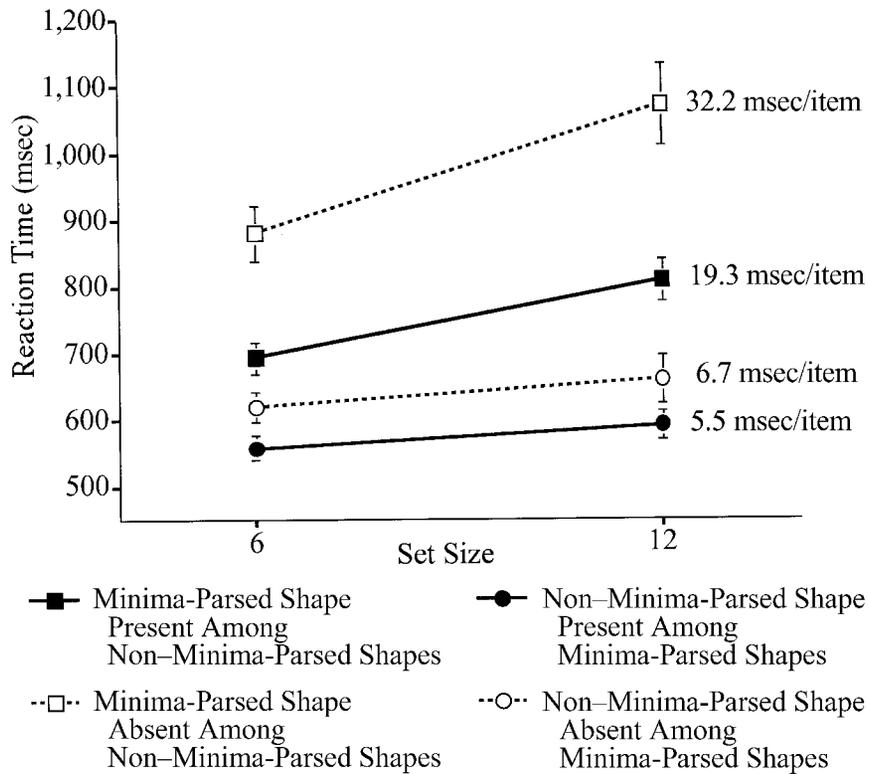


Figure 8. The mean reaction times for correct trials in Experiment 1. When the participants searched for a non-minima-parsed shape among minima-parsed shapes, the search was fast and efficient for both target-present and target-absent trials; when the roles of the target and the distractors were reversed, search became slow and inefficient. These results indicated that parsing at negative minima of curvature occurs rapidly and early in visual processing.

known, search becomes slow and inefficient when the target is defined by the absence of a feature (e.g., Treisman & Gormican, 1988; Treisman & Souther, 1985).

It is worth noting that this search asymmetry cannot be attributed to the difference in the lengths of the top vertical bars, because a search based on bar length would, in fact,

predict the *opposite* pattern of results. Search experiments involving differences in bar length have shown that longer bars pop out among shorter ones, but not vice versa (Treisman & Gormican, 1988). Since the minima-segmented shapes have longer bars on top than the non-minima-segmented shapes, a search based on the lengths of these

Table 1
Error Rates for Experiments 1-4

Experiment	Cut	Target							
		Present				Absent			
		Size 6		Size 12		Size 10		Size 20	
		M	SE	M	SE	M	SE	M	SE
1	minima parsed	.09	.02	.04	.01	.15	.02	.06	.01
	non-minima parsed	.08	.01	.04	.01	.10	.02	.05	.01
2	minima parsed	.10	.01	.02	.01	.14	.03	.01	.003
	non-minima parsed	.08	.01	.03	.01	.13	.02	.02	.01
3	minima parsed	.05	.02	.02	.004	.06	.02	.01	.004
	non-minima parsed	.04	.01	.02	.01	.04	.01	.02	.01
4	minima parsed	.04	.01	.01	.005	.06	.01	.02	.01
	non-minima parsed	.04	.01	.03	.01	.03	.01	.02	.01

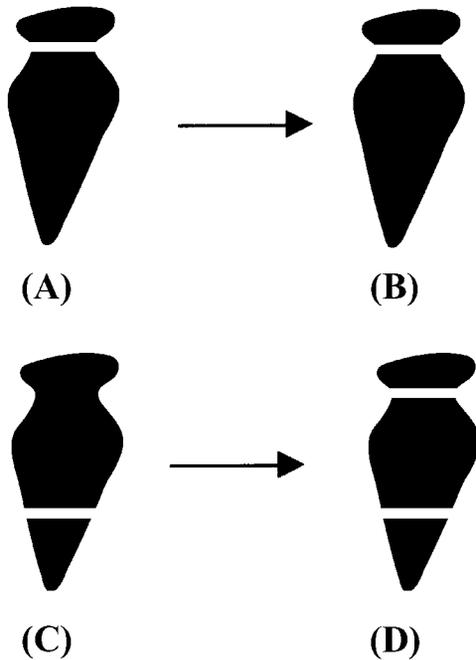


Figure 9. Stimuli used in Experiment 2: (A) minima-parsed shape; (B) no further parsing occurs for Shape A; (C) non-minima-parsed shape; and (D) the resulting shape if parsing at negative minima occurs automatically for the shape presented in panel C.

bars would make the search for minima-segmented shapes easier—the opposite of what we obtained.

EXPERIMENT 2

In Experiment 2, we sought to test the generality of the results obtained in Experiment 1. Note that the shape we used in Experiment 1 had the following two characteristics. First, each of its two parts had a symmetric and familiar form (i.e., a rectangular bar and an oval). Hence, in principle, the parsing of this shape into parts could have occurred simply owing to the familiarity and symmetry of its constituent parts. Second, the part boundaries on the shape used in Experiment 1 were both points of tangent discontinuity—that is, sharp corners—whereas, in general, part boundaries can be smooth as well. In order to test the generality of our results, in Experiment 2, we used a shape whose parts were both unfamiliar and nonsymmetric and whose part boundaries were smoothed rather than sharp corners (see Figure 9). As is well known, smoothed part boundaries are visually less salient than are corresponding concave-corner boundaries and lead to a weaker part structure (see, e.g., Hoffman & Singh, 1997). Thus, one might expect a somewhat weaker effect with such part boundaries. However, it is still reasonable to ask whether an asymmetry will be obtained in visual search. As in Experiment 1, the part cuts in the two parsings were equated in their lengths, orientations, and the areas of the parts that they created.

Method

Participants. The same 12 participants who took part in Experiment 1 also participated in this experiment. The order of the two experiments was counterbalanced across participants.

Materials and Design. The shapes used in Experiment 2 are shown in Figures 9A and 9C. The size of the shape was $1.26^\circ \times 2.75^\circ$, and that of the cut was $0.69^\circ \times 0.11^\circ$. For the parsing shown in Figure 9A, the vertical length of the top part was 0.40° , and that of the bottom part was 0.88° . As in Experiment 1, the relative areas of the parts created by the minima parsing and the non-minima parsing were equated.

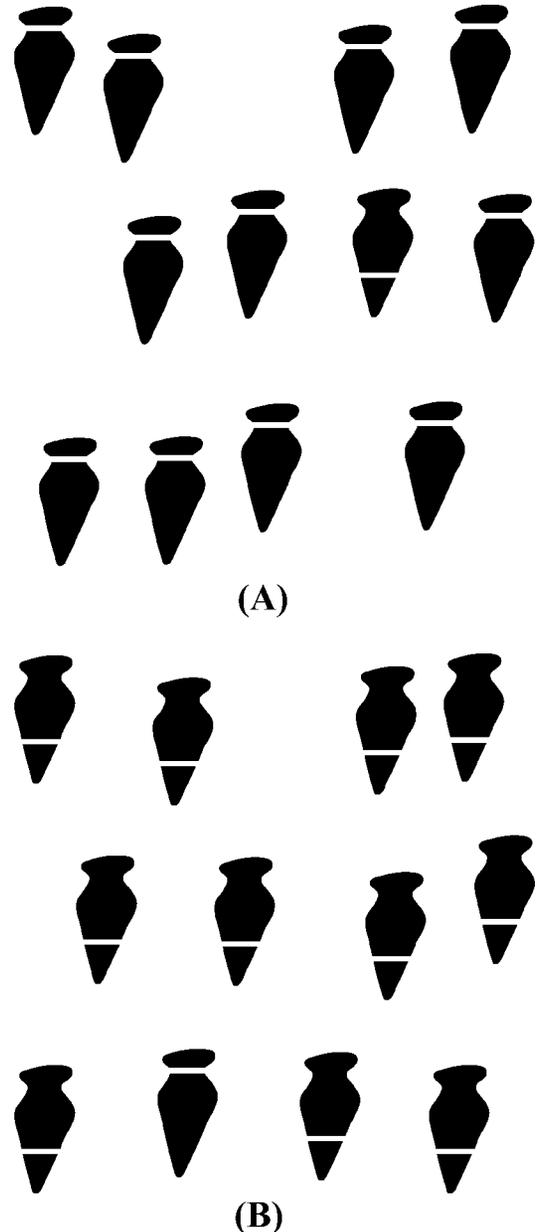


Figure 10. Examples of search displays used in Experiment 2. The participants searched either for a non-minima-parsed target among minima-parsed distractors (A) or for a minima-parsed target among non-minima-parsed distractors (B).

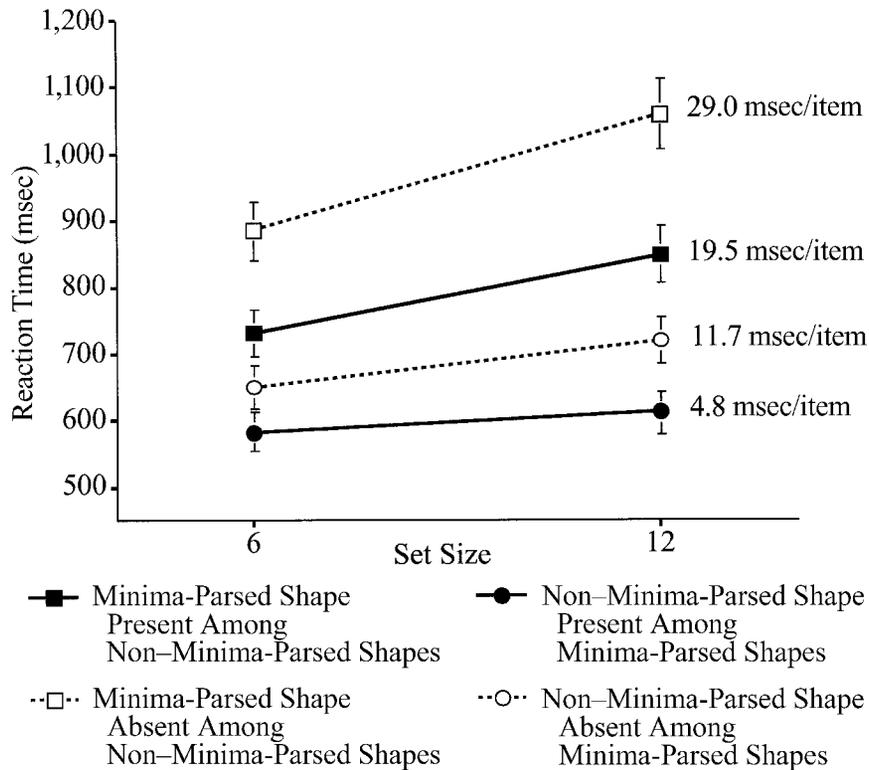


Figure 11. The mean reaction times for correct trials in Experiment 2. As in Experiment 1, when the participants searched for a non-minima-parsed shape among minima-parsed shapes, the search was fast and efficient; when the roles of the target and the distractors were reversed, search became slow and inefficient. These results indicated that parsing at negative minima of curvature occurs rapidly and early in visual processing even for shapes whose parts are neither symmetric nor simple familiar shapes and whose part boundaries are smoothed rather than concave corners.

The experiment was run in two sessions. In one session, the participants searched for a non-minima-parsed shape among minima-parsed shapes (Figure 10A), and in the other session, they searched for a minima-parsed shape among non-minima-parsed shapes (Figure 10B). Other aspects of the design and procedure were identical to those of Experiment 1.

Results

Reaction time for correct trials. Of the total data collected, 6.51% were removed owing to response errors. Two further data points were removed because the RTs were longer than 3,000 msec. The remaining data were analyzed with a within-subjects ANOVA. The mean RTs averaged across participants are plotted in Figure 11.

Overall, as in Experiment 1, the two search types differed from each other significantly [$F(1,11) = 38.27, p < .001$], such that the participants were much faster at detecting the presence/absence of a non-minima-parsed target among minima-parsed distractors than they were at detecting the presence/absence of a minima-parsed target among non-minima-parsed distractors. The effect of set size was significant [$F(1,11) = 38.44, p < .001$], and it interacted significantly with search type [$F(1,11) = 13.42, p = .004$]. The participants were faster at detecting the

presence, rather than the absence, of the target [$F(1,11) = 52.46, p < .001$], and this effect interacted significantly with search type [$F(1,11) = 5.99, p = .032$]. The interaction between set size and target presence/absence was also significant [$F(1,11) = 7.37, p = .020$]. However, the three way interaction among search type, set size, and target presence/absence was not significant ($F < 1$).

A separate ANOVA was also carried out for each search type. When the participants searched for a non-minima-parsed shape among minima-parsed shapes, the effect of set size and target presence/absence and the interaction between the two were all significant [$F(1,11) = 26.89, p < .001, F(1,11) = 15.44, p = .002$, and $F(1,11) = 6.89, p = .024$, respectively]. When the participants searched for a minima-parsed shape among non-minima-parsed shapes, the effect of set size and the effect of target presence/absence were both significant [$F(1,11) = 28.40, p < .001$, and $F(1,11) = 35.81, p < .001$, respectively]. The interaction between set size and target presence/absence was marginally significant [$F(1,11) = 3.76, p = .079$].

Error rates. The mean error rates are presented in Table 1. Only the effect of target presence/absence [$F(1,11) = 52.00, p < .001$] and the interaction between target pres-

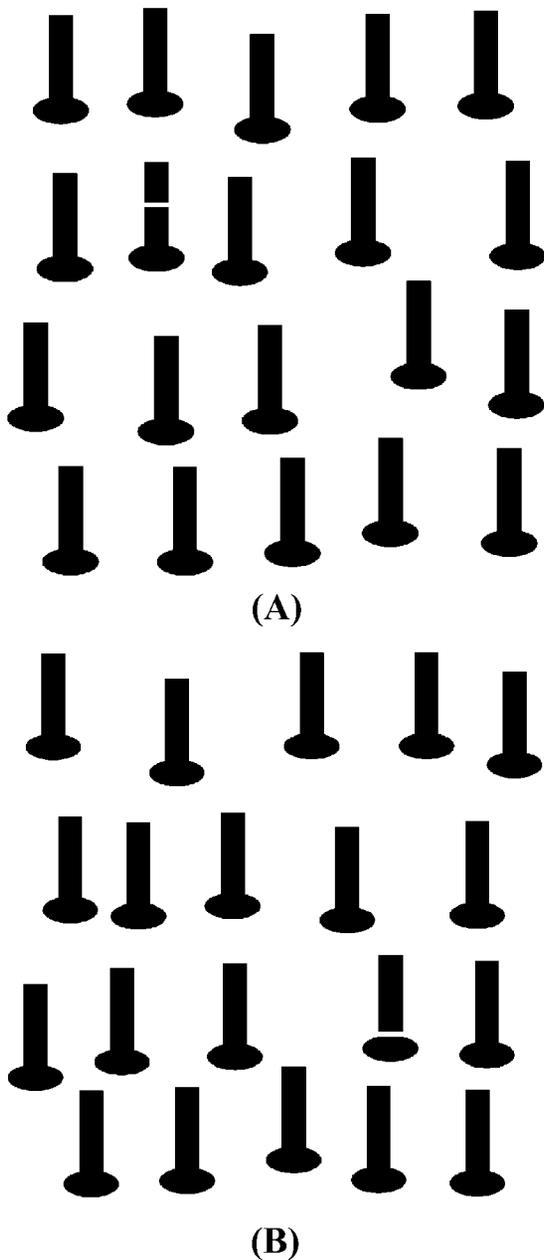


Figure 12. Examples of search displays used in Experiment 3. The participants searched either for a non-minima-parsed target among uncut distractors (A) or for a minima-parsed target among uncut distractors (B).

ence/absence and set size [$F(1,11) = 7.30, p = .021$] were significant.

Discussion

As in Experiment 1, there was a clear search asymmetry between the search for a non-minima-parsed target among minima-parsed distractors and the search for a minima-parsed target among non-minima-parsed distractors. When the part cut was located at a non-minima lo-

cation on the target shape, search was very fast and efficient (4.8 and 11.7 msec/item for target-present and target-absent trials, respectively). However, when the cut occurred at the negative minima on the target shape, search became slow and inefficient (19.5 and 29.0 msec/item for target present and target absent, respectively).

When the participants searched for the non-minima-parsed target among minima-parsed distractors, although search was still very fast, the ratio between target-present and target-absent slopes was higher than 2.0 (recall Wolfe's, 1998, criteria for a feature search). Note, however, that there is an inherent difficulty in interpreting the slope ratios for these shallow slopes, because the confidence intervals around the average slope ratios are quite wide (± 1.7), thus adding uncertainty to the estimate of the slope ratio.

In sum, although the magnitude of the effect was slightly weaker than it was in Experiment 1 (reflecting the fact that smoothed part boundaries typically lead to perceptually weaker part structures), the asymmetry in visual search remained highly significant. This indicates that parsing at negative minima occurs rapidly and early in visual processing even for shapes whose parts are neither symmetric nor simple familiar shapes and whose part boundaries are smoothed rather than concave corners.

EXPERIMENTS 3–4

In the next two experiments, our goal was to get at the issue of early parsing of shapes at negative minima by using a slightly different method. Using the same basic shapes as in Experiments 1 and 2, we asked the participants to search for a parsed shape among unparsed shapes. A pilot study showed that a parsed shape pops out among unparsed shapes, largely irrespective of the location of the part cut (since the target is defined by the presence of a unique feature, the part cut, that the distractors lack). We predicted, however, that if parsing at negative minima occurs obligatorily in early stages of visual processing, a cut located at negative minima would be somewhat less efficacious as a feature that distinguishes the target from the distractors. As a result, a cut located at negative minima should be harder to detect—and therefore, slower to search for (Figure 12B)—than a cut located at a non-negative-minima location (Figure 12A).

Obvious shapes were used as those in Experiment 1 (Figures 6A and 6C). The participants searched either for a non-minima-parsed shape among uncut shapes (Figure 12A) or for a minima-parsed shape among uncut shapes (Figure 12B).

Method

Participants. Twelve new participants, 6 males and 6 females, from the same participant pool were recruited. They were paid for their participation.

Materials and Design. The stimuli used were the shapes in Figures 6A and 6C and the uncut shape. There were either 10 or 20 search items in a given display. In one session of the experiment, the participants searched for a non-minima-parsed shape among uncut shapes (Figure 12A); in the other, they searched for a minima-parsed

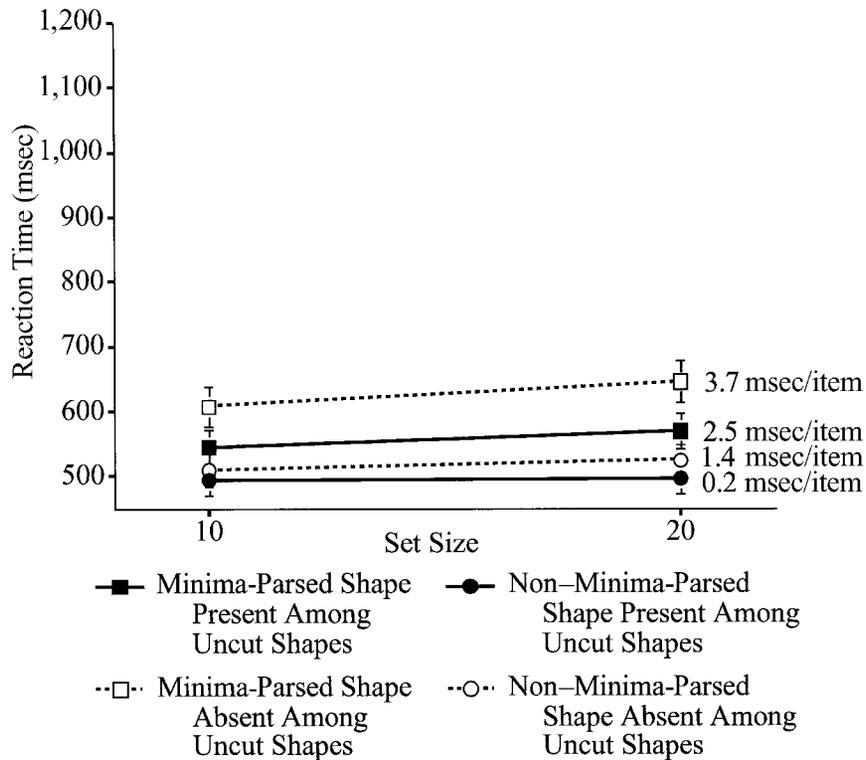


Figure 13. The mean reaction times for correct trials in Experiment 3. The parsed target popped out among uncut distractor shapes, both when the target was parsed using a minima cut and when it was parsed using a non-minima cut. However, the participants were significantly faster (on average, by 86 msec) at detecting the presence or absence of the parsed target when it was parsed with a non-minima cut than when it was parsed with a minima cut.

shape among uncut shapes (Figure 12B). The order for these two sessions was counterbalanced across participants. Other aspects of the design and procedure were identical to those in Experiment 1.

Results

Reaction time for correct trials. Of the total data collected, 3.10% were removed owing to response errors. Two further data points were removed because the RTs were longer than 2,000 msec.⁵ The remaining data were analyzed with a within-subjects ANOVA. The mean RTs averaged across participants are plotted in Figure 13.

Overall, the two search types differed significantly from each other, such that it took participants 86 msec longer, on average, to detect the presence or absence of the target shape when it was parsed with a minima cut than when it was parsed with a non-minima cut [$F(1,11) = 20.41, p = .001$]. The effect of set size was significant [$F(1,11) = 8.55, p = .014$], and it interacted significantly with search type [$F(1,11) = 5.27, p = .042$]. The participants were faster at detecting the presence, rather than the absence, of the target [$F(1,11) = 15.89, p = .002$], and this effect interacted significantly with search type [$F(1,11) = 8.78, p = .013$]. The interaction between set size and target presence/absence was marginally significant [$F(1,11) = 3.97, p = .072$]. However, the three-way interaction among

search type, set size, and target presence/absence was not significant ($F < 1$).

A separate ANOVA was also carried out for each search type. When the participants searched for a non-minima-parsed shape among uncut shapes, only the effect of target presence/absence was significant [$F(1,11) = 7.12, p = .022$]. When the participants searched for a minima-parsed shape among uncut shapes, both the effect of set size and the effect of target presence/absence were significant [$F(1,11) = 12.06, p = .005$, and $F(1,11) = 15.04, p = .003$, respectively].

Error rates. The mean error rates are presented in Table 1. The only significant effect was target presence/absence [$F(1,11) = 6.89, p = .024$]. The interaction between target presence/absence and search type was marginally significant [$F(1,11) = 3.76, p = .078$].

Discussion

As was expected, the parsed target popped out among uncut distractor shapes, both when the target was parsed using a minima cut and when it was parsed using a non-minima cut. (The search slopes for the minima-parsed target were 2.5 msec/item for target present and 3.7 msec/item for target absent. The search slopes for the non-

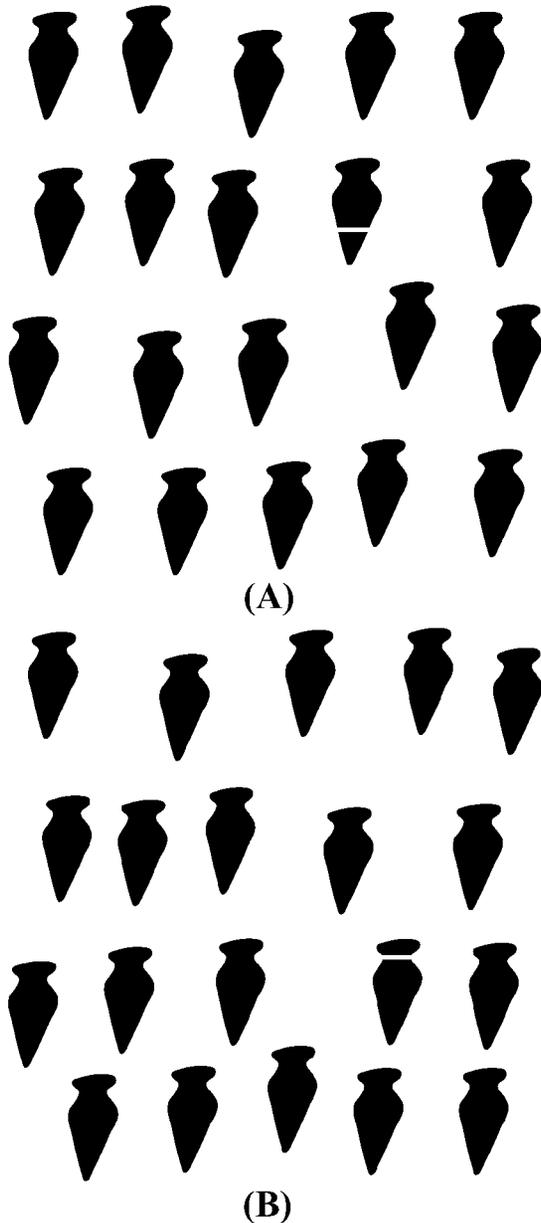


Figure 14. Examples of search displays used in Experiment 3. The participants searched either for a non-minima-parsed target among uncut distractors (A) or for a minima-parsed target among uncut distractors (B).

minima-parsed target were 0.2 msec/item for target present and 1.4 msec/item for target absent.)

However, participants were significantly faster (on average, by 86 msec) at detecting the presence or absence of the parsed target when it was segmented with a non-minima cut than when it was segmented with a minima cut. In other words, a cut located at negative minima was more difficult and slower to detect than a cut located elsewhere. Thus, the presence of a *perceptual* part cut at that location appears to have a masking influence on the detection of a

physically present cut. This effect was most striking for target-absent displays: Even though the search displays were identical in both cases (consisting only of uncut shapes), it took the participants 100 msec longer to determine the absence of a part cut when they were looking for it at negative minima. The fact that parsing at negative minima influenced visual search even when search was already fast and efficient provides further evidence for the claim that such parsing occurs very early in visual processing.

Experiment 4

In this experiment, our goal was to extend the results of Experiment 3 to the shapes used in Experiment 2 (Figure 9). Participants searched either for a non-minima-parsed shape among uncut shapes (Figure 14A) or for a minima-parsed target shape among uncut shapes (Figure 14B). In all other respects, the experimental design and procedure were identical to those of Experiment 3.

Method

Participants. The same 12 participants who took part in Experiment 3 also participated in this experiment. The order of the two experiments was counterbalanced across participants.

Materials and Design. The stimuli shown in Figures 9A and 9C and the uncut shape were used in the experiment. In one session, the participants searched for a non-minima-parsed shape among uncut shapes (Figure 14A), and in the other, the participants searched for a minima-parsed shape among uncut shapes (Figure 14B). The order for the two sessions was counterbalanced across participants.

Results

Reaction time for correct trials. Of the total data collected, 3.10% were removed owing to response errors. One further data point was removed because the RT was longer than 2,000 msec (see note 5). The remaining data were analyzed with a within-subjects ANOVA. The mean RTs averaged across participants are plotted in Figure 15.

Overall, the two search types differed from each other significantly, such that it took participants significantly longer (49 msec, on average) to detect the presence or absence of the target shape when it was parsed with a minima cut than when it was parsed with a non-minima cut [$F(1,11) = 11.69, p = .006$]. The effect of set size was marginally significant [$F(1,11) = 3.31, p = .096$]. The participants were faster at detecting the presence, rather than the absence, of the target [$F(1,11) = 13.83, p = .003$], and this effect interacted significantly with search type [$F(1,11) = 22.72, p = .001$]. As in Experiment 3, the three-way interaction among search type, set size, and target presence/absence was not significant ($F < 1$).

A separate ANOVA was also carried out for each search type. When participants searched for a non-minima-parsed target among uncut distractors, none of the effects reached significance. When participants searched for a minima-parsed target among uncut distractors, only the effect of target presence/absence was significant [$F(1,11) = 27.03, p < .001$].

Error rates. The mean error rates are presented in Table 1. The effect of target presence/absence was signifi-

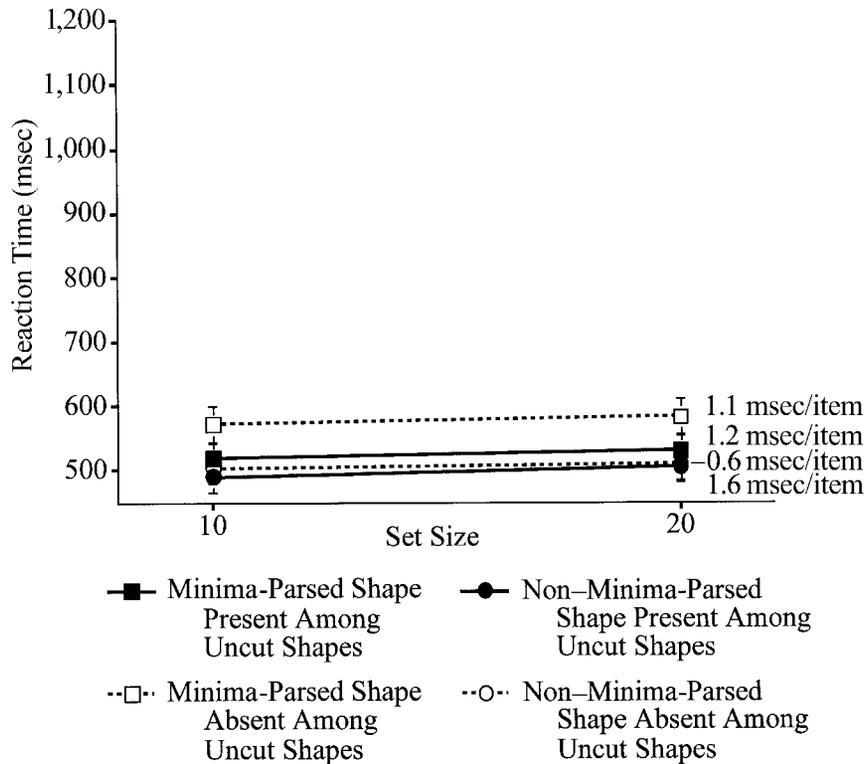


Figure 15. The mean reaction times for correct trials in Experiment 4. As in Experiment 3, the parsed target popped out among uncut distractor shapes, both when the target was parsed using a minima cut and when it was parsed using a non-minima cut. However, even though shapes are neither symmetric nor simple familiar shapes and their part boundaries are smoothed rather than concave corners, the participants were significantly faster (on average, by 49 msec) at detecting the presence or absence of the parsed target when it was parsed with a non-minima cut than when it was parsed with a minima cut.

cant [$F(1,11) = 22.56, p = .001$], and this effect interacted significantly with search type [$F(1,11) = 13.55, p = .004$].

Discussion

As in Experiment 3, the search for a parsed target among unparsed distractors was quite fast and efficient, both for the minima-parsed and the non-minima-parsed targets. (The search slopes for the minima-parsed target were 1.2 msec/item for target present and 1.1 msec/item for target absent. The search slopes for the non-minima-parsed target were 1.6 msec/item for target present and -0.6 msec/item for target absent.)

However, the participants were significantly faster (on average, by 49 msec) at detecting the presence or absence of the target shape when it was parsed with a non-minima cut than when it was parsed with a minima cut. This replicates the results of Experiment 3 with shapes whose parts were neither symmetric nor simple familiar shapes and whose part boundaries were smooth rather than tangent discontinuities. As in Experiment 3, these results were most striking for the target-absent trials: It took the participants 70 msec longer, on average, to determine the ab-

sence of a part cut when they were looking for it at negative minima than when they were looking for it elsewhere, even though the search displays were identical in both cases (consisting simply of identical uncut shapes).

The magnitude of the effect was somewhat smaller than it was in Experiment 3. As in Experiment 2, this difference reflects the fact that the part boundaries in Figure 9 are smooth with low curvature, as compared with the sharp corners in Figure 6, a change that has been demonstrated to reduce the salience of perceived part structure (Hoffman & Singh, 1997).

GENERAL DISCUSSION

We set out to address the question of whether parsing at negative minima occurs *preattentively*—or at least, rapidly and early in visual processing. In Experiment 1, we found that whereas the search for a non-minima-parsed target among minima-parsed distractors was fast and efficient, the search for a minima-parsed target among non-minima-parsed distractors was slow and inefficient. This asymmetry demonstrates that parsing at negative minima occurs

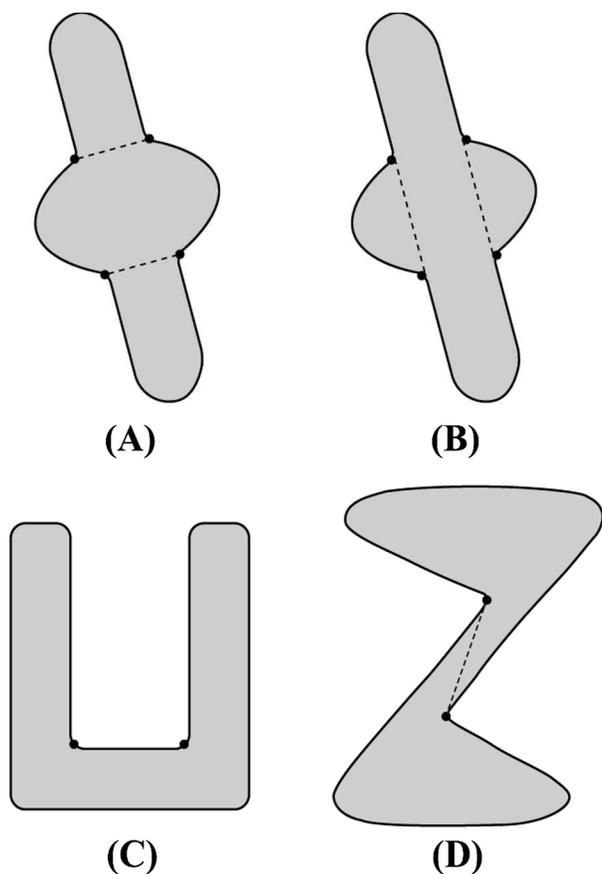


Figure 16. Examples demonstrating that points of negative minima are not by themselves sufficient to define part structures. If a shape contains more than two negative minima, these may be paired in different ways to give part cuts (A and B). Even if a shape has only two negative minima, joining them may not necessarily give perceptually natural part structures (C and D).

rapidly and early in visual processing. Experiment 2 replicated the search asymmetry of Experiment 1 with shapes whose parts are neither symmetric nor simple familiar geometric shapes and whose part boundaries are smooth rather than sharp corners.

In Experiments 3 and 4, we used the same shapes as in Experiments 1 and 2, respectively, but asked the participants to search either for a minima-parsed target among uncut distractors or for a non-minima-parsed target among uncut distractors. We found that although the parsed shape popped out in both cases, participants were significantly slower at detecting the presence or absence of the minima-parsed target than that of the non-minima-parsed target. It is interesting to note that an even larger difference was obtained for the target-absent trials, even though the search displays were identical in the two cases, consisting only of uncut distractor shapes. The fact that a cut located at negative minima was systematically harder to detect and slower to search for provides further evidence for the claim that parsing at negative minima occurs obligatorily. Moreover, parsing at negative minima influenced visual

search even when the search was already fast and efficient, indicating that this parsing must occur rapidly and early in visual processing.

Our findings are also consistent with recent work on the part-based nature of visual attention, which has shown that attention shifts more readily within a single part than across two parts of a single object (Barenholtz & Feldman, 2002; Singh & Scholl, 2000; Vecera, Behrmann, & Filapek, 2001; Vecera, Behrmann, & McGoldrick, 2000). The fact that parts of objects can act as units of attentional selection also suggests that they are computed obligatorily, requiring few attentional resources for their computation.

One recent study (Donnelly, Found, & Müller, 2000) tentatively reached the opposite conclusion on the basis of visual search. In particular, Donnelly et al. postulated that it is *easier* for the visual system to search for convex shapes than to search for shapes with negative minima (or concavities, more generally). Although it is not entirely clear why Donnelly et al.'s conclusion differs from ours as well as those of other researchers (recall the results of Elder & Zucker, 1993; Hulleman et al., 2000; Humphreys & Müller, 2000; Wolfe & Bennett, 1997), it should be noted, as Donnelly et al. acknowledged, that their analysis of concavities was post hoc. The focus of their experiments was on the role of *contour segments* in distinguishing shapes *preattentively*, and these experiments did not systematically manipulate the presence of concavities or negative minima.

In addition to the specific findings of the present study, our experiments also provide a general method for studying how the visual system parses shapes into parts. Rather than using two distinct shapes for targets and distractors, one uses the same shape, but shown with two different parsings. On the assumption that other local factors (such as cut orientation, cut length, etc.) have been controlled for, any asymmetry obtained in visual search is then indicative of a difference in the perceived naturalness of the two parsings.

For the shapes we studied in this article, the perceptually natural part cuts were always obtained by joining the two negative minima on the shapes. In general, however, the minima rule is not by itself sufficient to define the part structure of a shape. The minima rule gives negative minima as boundaries between parts; for two-dimensional shapes, these are points on the bounding contour of the shape. It does not give part cuts that segment the shape, however, which are needed to determine part structure. For example, if a shape has more than two negative minima, they may be joined in more than one way to give part cuts (see Figures 16A and 16B; Beusmans, Hoffman, & Bennett, 1987; Siddiqi & Kimia, 1995). In addition, even if a shape has precisely two negative minima, joining them may not *always* give a natural parsing of the shape (see Figures 16C and 16D; Singh et al., 1999). Geometric factors other than negative minima have been shown to play an important role in determining part structures (e.g., cut length, boundary strength, good continuation, local sym-

metry, etc.; see Singh & Hoffman, 2001; Singh et al., 1999). The methodology described in this article will allow us to investigate in future studies whether parsings determined by other geometric factors also occur rapidly, in early stages of visual processing.

Because the shapes we used in this article were quite simple and their part structures were unambiguous, other schemes of shape representation and parsing would also predict the same perceived part structures for our stimuli. For example, Burbeck and Pizer's (1995) model of object representation by *cores* and Siddiqi and Kimia's (1995) parsing scheme involving *limbs* and *necks* would also predict the perceptually natural parsings depicted in Figures 6A and 9A. Our goal in this article was simply to test whether such perceptually natural parsings are computed rapidly and early in visual processing, for which we found ample evidence. In future work, the visual search methodology can be used to compare directly the relative merits of these alternative parsing schemes.

Finally, our data speak to a broader issue in the visual computation of parts. Historically, two main approaches have been taken toward the problem of the computing of parts. According to the *primitive-shapes approach* to parts, the visual system parses shapes by looking for a pre-defined set of basic shapes in images (e.g., Biederman, 1987; Marr & Nishihara, 1978; Winston, 1975). Essentially, it finds parts by fitting these basic shapes to images. On the other hand, according to the *geometric-constraints approach*, the visual system does not store any such set of primitive shapes but, rather, uses general geometric constraints—based on intrinsic shape geometry alone—to define boundaries and cuts between parts (Hoffman & Richards, 1984; Hoffman & Singh, 1997; Singh & Hoffman, 2001; Singh et al., 1999). One problem faced by the basic-shapes approach is that it can find only those parts that belong to the prespecified set of basic shapes. For the geometric-constraints approach, however, parts do not need to belong to a prespecified set of shape primitives.⁶ In the present set of experiments, visual search asymmetry and feature search differences resulting from parsing at negative minima were observed for shapes regardless of whether the resulting parts were simple familiar shapes or not—hence, providing empirical evidence in favor of the geometric-constraints approach to shape parsing.

REFERENCES

- ATTNEAVE, F. (1971, December). Multistability in perception. *Scientific American*, **225**, 63-71.
- BARENHOLTZ, E., & FELDMAN, J. (2002). *Visual comparisons within and between object-parts: Evidence for a single-part superiority effect*. Manuscript submitted for publication.
- BAYLIS, G. C., & DRIVER, J. (1994). Parallel computation of symmetry but not repetition in single visual objects. *Visual Cognition*, **1**, 377-400.
- BAYLIS, G. C., & DRIVER, J. (1995). Obligatory edge assignment in vision: The role of figure and part segmentation in symmetry detection. *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 1323-1342.
- BERTAMINI, M., FRIEDENBERG, J. D., & KUBOVY, M. (1997). Detection of symmetry and perceptual organization: The way a lock-and-key process works. *Acta Psychologica*, **95**, 119-140.
- BEUSMANS, J., HOFFMAN, D. D., & BENNETT, B. M. (1987). Description of solid shape and its inference from occluding contours. *Journal of the Optical Society of America A*, **4**, 1155-1167.
- BIEDERMAN, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, **94**, 115-147.
- BIEDERMAN, I., & COOPER, E. E. (1991). Priming contour-deleted images: Evidence for intermediate representations in visual object recognition. *Cognitive Psychology*, **23**, 393-419.
- BRAUNSTEIN, M. L., HOFFMANN, D. D., & SAIDPOUR, A. (1989). Parts of visual objects: An experimental test of the minima rule. *Perception*, **18**, 817-826.
- BRUCE, V., & MORGAN, M. J. (1975). Violations of symmetry and repetitions in visual patterns. *Perception*, **4**, 239-249.
- BURBECK, C. A., & PIZER, S. M. (1995). Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research*, **35**, 1917-1930.
- CORBALLIS, M. C., & ROLDAN, C. E. (1974). On the perception of symmetrical and repeated patterns. *Perception & Psychophysics*, **16**, 136-142.
- DONNELLY, N., FOUND, A., & MÜLLER, H. J. (2000). Are shape differences detected in early vision? *Visual Cognition*, **7**, 719-741.
- 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**, 201-206.
- DRIVER, J., & BAYLIS, G. C. (1996). Edge-assignment and figure-ground segmentation in short-term visual matching. *Cognitive Psychology*, **31**, 248-306.
- ELDER, J., & ZUCKER, S. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, **33**, 981-991.
- FRIEDENBERG, J., & BERTAMINI, M. (2000). Contour symmetry detection: The influence of axis orientation and number of objects. *Acta Psychologica*, **105**, 107-118.
- HOFFMAN, D. D. (1983, December). The interpretation of visual illusions. *Scientific American*, **249**, 154-162.
- HOFFMAN, D. D., & RICHARDS, W. A. (1984). Parts of recognition. *Cognition*, **18**, 65-96.
- HOFFMAN, D. D., & SINGH, M. (1997). Saliency of visual parts. *Cognition*, **63**, 29-78.
- HULLEMAN, J., TE WINKEL, W., & BOSELIE, F. (2000). Concavities as basic features in visual search: Evidence from search asymmetries. *Perception & Psychophysics*, **62**, 162-174.
- HUMPHREYS, G., & MÜLLER, H. (2000). A search asymmetry reversed by figure-ground assignment. *Psychological Science*, **11**, 196-201.
- JOSEPH, J. S., CHUN, M. M., & NAKAYAMA, K. (1997). Attentional requirements in a "preattentive" feature search task. *Nature*, **387**, 805-807.
- KIMIA, B. B., TANNENBAUM, A., & ZUCKER, S. W. (1995). Shape, shocks, and deformations: I. The components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision*, **15**, 189-224.
- LEYTON, M. (1992). *Symmetry, causality, mind*. Cambridge, MA: MIT Press.
- MARR, D., & NISHIHARA, H. K. (1978). Representation and recognition of three-dimensional shapes. *Proceedings of the Royal Society of London: Series B*, **200**, 269-294.
- NEISSER, U. (1967). *Cognitive psychology*. New York: Appleton-Century-Crofts.
- PALMER, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, **9**, 441-474.
- RENSINK, R., & ENNS, J. (1998). Early completion of occluded objects. *Vision Research*, **38**, 2489-2505.
- SIDDIQI, K., & KIMIA, B. B. (1995). Parts of visual form: Computational aspects. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **17**, 239-251.
- SIDDIQI, K., TRESNESS, K., & KIMIA, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception*, **25**, 399-424.
- SINGH, M., & HOFFMAN, D. D. (1998). Part boundaries alter the perception of transparency. *Psychological Science*, **9**, 370-378.
- SINGH, M., & HOFFMAN, D. D. (2001). Part-based representations of visual shape and implications for visual cognition. In T. Shipley & P. Kellman (Eds.), *From fragments to objects: Segmentation and*

- grouping in vision* (Advances in Psychology, Vol. 130, pp. 401-459). New York: Elsevier, North-Holland.
- SINGH, M., & SCHOLL, B. (2000, November). *Using attentional cueing to explore part structure*. Poster presented at the Annual Symposium of Object Perception and Memory, New Orleans.
- SINGH, M., SEYRANIAN, G. D., & HOFFMAN, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception & Psychophysics*, **61**, 636-660.
- SPELKE, E. S. (1994). Initial knowledge: Six suggestions. *Cognition*, **50**, 431-445.
- TREISMAN, A., & GORMICAN, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, **95**, 15-48.
- TREISMAN, A., & SOUTHER, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, **114**, 285-310.
- VECERA, S. P., BEHRMANN, M., & FILAPEK, J. C. (2001). Attending to the parts of a single object: Part-based attentional limitations. *Perception & Psychophysics*, **63**, 308-321.
- VECERA, S. P., BEHRMANN, M., & MCGOLDRICK, J. (2000). Selective attention to the parts of an object. *Psychonomic Bulletin & Review*, **7**, 301-308.
- WINSTON, P. A. (1975). Learning structural descriptions from examples. In P. H. Winston (Ed.), *The psychology of computer vision* (pp. 157-209). New York: McGraw-Hill.
- WOLFE, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, **1**, 202-238.
- WOLFE, J. M. (1998). What can 1 million trials tell us about visual search? *Psychological Science*, **9**, 33-39.
- WOLFE, J. M., & BENNETT, S. C. (1997). Preattentive object files: Shapeless bundles of basic features. *Vision Research*, **37**, 25-43.
- WOLFE, J. M., YEE, A., & FRIEDMAN-HILL, S. R. (1992). Curvature is a basic feature for visual search tasks. *Perception*, **21**, 465-480.
- dots or line segments), whose "parts" can, of course, consist of disconnected elements.
2. The differential-geometric principle of *transversality* ensures that whenever 2 three-dimensional shapes intersect generically, the locus of their intersection lies, with a probability of 1, in a concave crease—that is, a concave tangent discontinuity and smoothing such a concave crease yields a locus of negative minima of curvature. (see Hoffman & Richards, 1984, and Hoffman & Singh, 1997, for details).
3. It should be noted that rapid computation can also be indicated by a *slowdown* in search performance, owing to the *preemption* of features that were previously easy to access (e.g., Rensink & Enns, 1998).
4. For any attribute that has *preattentive* status (say, color or orientation), it is clear that a pop-out will occur only if the target and the distractors differ sufficiently along that dimension. If they are too close along that dimension (e.g., very similar hues or orientations), search slopes will get large, despite the featural status of the attribute. In the case of convex and concave corners, lower turning angles simply translate into smaller differences between the two kinds of corners.
5. A pilot study showed that both types of search (i.e., search for minima-parsed targets and for non-minima-parsed targets among uncut distractors) were fast and efficient. As a result, we decided to use 2,000 msec as the cutoff threshold for RT truncation, instead of the 3,000 msec used previously. This change in cutoff threshold led to the elimination of only two data points in Experiment 3 and of a single data point in Experiment 4 and did not significantly alter the results.
6. It should be noted that the theories of Marr and Nishihara (1978) and Biederman (1987) have much broader scope than the minima rule, since they are general theories of shape representation and object recognition, not just of shape segmentation. Comparison of them with the minima rule is valid, however, insofar as these theories also make specific proposals regarding how the visual system segments objects into parts.

NOTES

1. Throughout, our concern is with parts of *objects*, bounded by continuous surfaces, rather than, say, with Gestalt groups (such as arrays of

(Manuscript received July 31, 2001;
revision accepted for publication February 26, 2002.)