Asymmetries in Different Feature Spaces

Curvature discontinuities are cues for rapid shape analysis

ÁRNI KRISTJÁNSSON

Harvard University, Cambridge, Massachusetts

and

PETER ULRIC TSE

Max Planck Institute for Biological Cybernetics, Tübingen, Germany

Using the visual search method, we show that stimuli that contain curvature discontinuities (i.e., points where the second derivative along an image contour is not defined) are easily found among stimuli containing only smooth changes in curvature. Curved stimuli that lack curvature discontinuities, however, are difficult to find among distractors that have them. These results suggest that the visual system detects and analyzes abrupt changes in curvature in the image quickly to extract vital information about the 3-D structure of the visual environment.

The visual system is confronted with numerous sources of ambiguity when interpreting the 2-D retinal image. A major ambiguity arises because every 2-D retinal image is consistent with many 3-D world layouts (e.g., Hoffman, 1998; Richards, Koenderink, & Hoffman, 1987). In order to recover the 3-D structure of the world, the visual system must make assumptions about the nature of the image-to-world mapping and assign probabilities to likely correspondences between image information and world structure. Several researchers have noted that the visual system operates as if it assumes that it is not perceiving an object from one of the few “accidental” viewpoints for which an object’s surface layout is not derivable from its projected contours and other image cues (Albert & Hoffman, 1995; Barrow & Tenenbaum, 1981; Binford, 1981; Lowe, 1987; Nakayama & Shimojo, 1992; Richards et al., 1987; Rock, 1984). If the visual system assumes such a nonaccidental viewpoint, a straight contour in the image will correspond to a straight edge or border in the world, because a curved edge will project to a straight line only from a very limited number of views. Slight deviations from such accidental views will cause a straight contour in the image to appear curved if the edge in the world is in fact curved. Similarly, a curved contour in the image

can be assumed to project from a curved edge, border, or surface in the world.

A corollary of the “nonaccidental viewpoint principle” is that certain image cues are more informative about world structure than others. A T-junction, for example, is more likely to indicate occlusion than is a straight contour in the image, because T-junctions generically arise in the image when there is a depth discontinuity between world surfaces lying along the line of sight (Clowes, 1971; Huffman, 1971; Kellman & Shipley, 1991; Lowe, 1987; Malik, 1987; Nakayama, Shimojo, & Silverman, 1989; Tse & Albert, 1998; Waltz, 1975). In contrast, straight contours do not generically arise in the image when there is occlusion in the visible scene. Moreover, straight contours can arise in the image from a host of world conditions that have nothing to do with occlusion. The visual system may have evolved to be especially sensitive to highly predictive cues, because rapid detection and processing of these cues may provide the most efficient route to recovering world structure from the image (Gibson, 1950).

Fast search rates in the visual search paradigm have generally been regarded as evidence that the visual system is predisposed to rapidly detect a given image feature, event, or configuration. It is well known that tangent discontinuities (i.e., points where the first derivative along an image contour is not defined) and curvature “pop out” (i.e., are rapidly detected, as determined by flat search slopes as a function of distractor set size) when the distractors are straight lines (Treisman & Gormican, 1988; Wolfe, Yee, & Friedman-Hill, 1992; see also Zucker, Dobbins, & Iverson, 1989). For example, an L or a C will pop out among | distractors regardless of the number of distractors. This has been taken as evidence that tangent discontinuities and curvature are “primitive” features to

We thank Marc K. Albert, Donald D. Hoffman, Jeremy Wolfe, and an anonymous reviewer for helpful comments. This research was supported by a Fulbright graduate study fellowship to Á.K., a McDonnell-Pew postdoctoral fellowship to P.U.T., and by a U.S. Air Force Office of Scientific Research grant to Ken Nakayama. Correspondence should be addressed to Á. Kristjánsson, Vision Sciences Laboratory, Harvard University, William James Hall, 33 Kirkland St., Cambridge, MA 02138 (e-mail: kristjan@wjh.harvard.edu).

which the visual system is highly sensitive. If the logic of the previous paragraph is correct, tangent discontinuities and curvature are likely to be among those image features that are particularly informative about world structure.

Curvature and Search Asymmetries

Curved contours in the 2-D image convey important information about the shape of objects. Wolfe et al. (1992) reached the conclusion that curvature is a basic feature for visual search tasks. They showed that curved contours pop out among straight contours and that this result was not accounted for by differences in local variation of orientation. Treisman and Gormican (1988) had previously found that curved contours were easily found among straight contours, whereas the converse did not hold: Straight contours did not pop out among curved contours.

One way to think about a "search asymmetry" of this sort is the following: Image cues to which the visual system is highly sensitive will pop out among distractors when those distractors do *not* contain much of the type of information that makes the target pop out. Not only must the target contain a type of information that the distractors lack or nearly lack, but also this type of information must be one that the visual system is predisposed to detect. When the experiment is inverted, and a single distractor is now used as a target among the previous targets, now serving as distractors, the new target may not pop out. The asymmetry occurs because the new target (i.e., what was a distractor in the previous experiment) does not contain information to which the visual system is highly tuned. So, for example, curved contours pop out among straight lines presumably because curvedness is a particularly informative image cue to which the visual system is highly sensitive. However, straight lines do not pop out among curves because the visual system is not particularly sensitive to their existence in the image. If this analysis is correct, this is because straight lines are less informative about world structure than curves are.

The search asymmetry for curved targets among straight line distractors attests to the visual system's bias to preferentially detect information about curvature rapidly. However, there are several different classes of information about world structure conveyed by different contour curvature cues. It may be that the visual system is more predisposed to detect one type of curvature cue over others.

A number of theorists have highlighted the importance of changes in curvature in the 2-D image. Hoffman and Richards (1984), for example, conjectured that 2-D silhouettes are segmented into parts at points of negative minima of contour curvature (i.e., concavities). Similarly, Richards, Dawson, and Whittington (1986) emphasized extrema of curvature as clues to the shape of surfaces. Others (Beusmans, Hoffman, & Bennett, 1987; Richards et al., 1987) have noted that points of contour curvature inflection correspond to points on a surface separating positive from negative surface curvature, for volumes with smooth surfaces. Tse (in press) has suggested that points of curvature discontinuity play a role in estimating the

likely cross section of a volume given only that volume's 2-D silhouette. According to these authors, visual processes that analyze form use minima, maxima, inflections, and discontinuities of curvature to derive 3-D shape from the image. Curved regions of a contour that lie between these particularly informative regions are presumably less important 2-D cues to the 3-D shape of an object. In general, regions in the retinal image where sudden changes or discontinuities occur, contain more information about the spatiotemporal structure of the environment than uniform regions (Attneave, 1954). If this is correct, then contour regions where the curvature of a figure changes contain more information about its shape than regions where the curvature is uniform. In particular, curvature discontinuities may offer a very strong cue to the 3-D shape of an object (Tse, in press).

It should come as no surprise that the visual system has evolved multiple visual "shape-from-x" systems, because the need to rapidly recover the 3-D structure of the world is essential for an animal to accomplish almost any task in its environment. In most cases there are multiple cues to 3-D form in an image. There are shadows, shading, disparity, highlights, occlusion, motion, contours, and so forth. Certain classes of stimuli activate only one of these shape-from-x systems. Because we are interested in the role played by contour curvature discontinuities in shape analysis, we need a type of stimulus that activates only the shape-from-contour system. There are two types of stimuli that satisfy this criterion: silhouettes and outlines. Outlines are never seen in the natural environment. However, an object may look very much like a silhouette under certain natural lighting conditions, such as in twilight or when lit from behind. Many researchers have discussed the importance of 2-D silhouettes as probes to the workings of the visual system (Hoffman & Richards, 1984; Kellman & Shipley, 1991; Richards et al., 1987; Singh, Seyranian, & Hoffman, 1999). In this paper we follow the tradition of using silhouettes and outline figures as the best means of probing the properties of the shape-from-contour system.

Discontinuities in Curvature

First-order or tangent discontinuities are important cues for scene segmentation and border ownership. First-order discontinuities appear in the image as, for example, T-, L-, and X-junctions. They are mathematically defined as points where the first derivative along a contour is not defined. In addition to first-order discontinuities, Tse and Albert (1998) and Tse (in press) proposed that contour curvature (or "second-order") discontinuities in the 2-D image are important clues to the 3-D structure of the visible environment. Consider two circles arranged as in Figure 1a. The smaller circle touches the larger circle at only one point, P. The curvature at every point along the larger circle has a value of 1 over the radius of that circle, $1/R$ (Courant & John, 1989). Similarly, the curvature at every point along the smaller circle has curvature $1/r$. Now consider the curve in the image defined by merging half

Figure 1. P is the point of intersection between two circles, as shown in (a). When a single curve is built from two corresponding semicircles, as in (b), P is a point of abrupt curvature change. Note that both circles have the same tangent (i.e., first derivative) at P, even though they have different second derivatives at P.

of one circle with half of the other circle at P (Figure 1b). Because the two circles have the same tangent at P, the tangent at every point along this curve is well defined (excluding the endpoints). However, the curvature to the right of P is $1/R$ and to the left of P is $1/r$. Therefore P is a point of abrupt curvature change along the image contour. The larger the difference between R and r, the greater the curvature discontinuity. Whereas the first derivative taken along the contour at P is continuous, a discontinuity exists in the second derivative at P. In general, a curvature discontinuity is defined as a point along a contour where the second derivative is not well defined. In contrast, an ellipse has smooth curvature changes everywhere along its contour, because the radius of curvature changes smoothly as one traverses the contour. There are no abrupt or discontinuous changes in curvature, and the second derivative is everywhere defined along the contour of an ellipse. The same principles hold for arbitrary image curves.

Figure 2a shows how such curvature discontinuities can dramatically alter the 3-D interpretation of silhouettes. The “bump” on the left contains a curvature discontinuity, whereas the ellipse on the right does not. If, as Tse and Albert (1998) proposed, curvature discontinuities are important features used to analyze the 2-D image in order to recover 3-D world structure, “abrupt” changes in curvature might pop out among smooth changes in a visual search task. It is thus possible that a search asymmetry, similar to the one Treisman and Gormican (1988) found for curved versus straight lines, might exist between abrupt and smooth changes in curvature. That is, curvature discontinuities might pop out among smooth changes in curvature, whereas the converse might not be true.

Using the visual search method, we contrasted stimuli that contain abrupt curvature changes with ones that contain only smooth changes in curvature. We show here that stimuli that contain curvature discontinuities (i.e., points where the second derivative along an image contour is not defined or where curvature changes abruptly) are easily found among stimuli containing only smooth changes in

curvature. These results suggest that the visual system detects abrupt changes in curvature in the image quickly to extract vital information about the structure of the visual environment.

EXPERIMENT 1 Bumps Versus Ellipses

In Experiment 1 we used two different stimuli adapted from Tse and Albert (1998; see also Albert & Tse, 2000; Tse, in press) that differed in the 2-D image only in that one contained an abrupt change in curvature whereas the other one did not (Figure 2). Given a 3-D interpretation, however, these stimuli tend to be interpreted as volumetric bumps on the one hand and flat ellipses on the other. The bumps were constructed by merging two half-ellipses that had identical major axes and different minor axes along their common major axis. The bumps contain a curvature discontinuity at the two points where the contours of the two half-ellipses meet (Figure 2). Note that the areas, heights, and widths of the bumps and matched ellipses are exactly the same. If curvature discontinuities are important cues to shape, we should observe patterns in search rates that reflect efficient processing of these stimuli—that is, that search rates are relatively unaffected by the number of distractors in the image to be searched, when the task is to find a bump among ellipses serving as distractors. In the converse search task, where an ellipse is to be found among bumps, we might not find a pop-out search pattern if smooth curvature changes are not an image property to which the visual apparatus is highly tuned.

Method

The 6 participants were undergraduates at Harvard University participating in the experiment for course credit. At a viewing distance of 64 cm, observers responded by pressing a key to indicate whether the target was present or absent in a search array of 4, 8, 12, or 16 display items, presented on a 75-Hz screen controlled by a G3 Macintosh computer. Observers were instructed to respond as

Figure 2. (a) The two stimuli used in Experiment 1. Observers either searched for a "bump" among "ellipses" or for an ellipse among bumps. (b) The bumps were constructed by combining the bottom half of an ellipse of major axis a and minor axis b where $(b = 2c)$ and the top half of an ellipse of major axis a and minor axis a where $(a = 2b)$. The bumps contain a curvature discontinuity along the contour where the two component half-ellipses meet. (c) The ellipses have major axis a and minor axis $b + c$ where $(b + c = 2d)$. Note that the widths and heights of the bumps and ellipses are exactly the same since $2d = b + c$. Moreover, the areas of the bumps and ellipses are identical because $\pi a(2d) = \pi a(b + c)$. (d) Two representative frames from Experiment 1 with a set size of 16. The target is present in both cases.

quickly as possible whether the target was present or not, while maintaining a high degree of accuracy. The display remained visible until a response key was pressed. Auditory feedback was given to indicate whether the response was correct or incorrect. Data associated with mistakes were not used in the analysis.

Two different stimuli were used (Figure 2). When a bump was the target, the ellipses were the distractors; the roles reversed when an ellipse was the target. These conditions were run in separate

blocks of 150 trials, two blocks under each condition, 600 trials in all, apart from 20 practice trials before each block.

The height of the bumps and ellipses at a 64-cm viewing distance was 1.68° (corresponding to $2d$ or $b + c$ in Figure 2) and the width (corresponding to a in Figure 2) was 2.23° . Stimuli were presented within an invisible 4×4 matrix on the screen subtending $12.3^\circ \times 12.3^\circ$, and each stimulus item was confined to a $2.9^\circ \times 2.9^\circ$ square. The stimulus was centered on one of nine random positions

Figure 3. Overall response times for the four different conditions of Experiment 1. The error bars show the largest standard error of the mean for each condition.

within the square, which resulted in a slight irregularity in the appearance of the array. The stimuli were black (0.5 cd/m^2) and appeared on a white (50.9 cd/m^2) background.

To prevent any effects of jaggedness of the stimuli due to aliasing on the screen, the stimuli were convolved with a 2-D Gaussian envelope with an SD of 1.5 pixels. The orientation of the stimuli was in every case as shown in Figure 2a. Figure 2d gives two examples of a search display from the experiment.

Results and Discussion

The results are presented in Figure 3 and the error rates are depicted in Table 1. The bumps were found very easily among the ellipses. The slopes of search times versus set size were -1.9 msec/item for the target trials and 8.0 msec/item on blank trials. The search is extremely efficient, appearing to be completely independent of set size, at least when the target is present. On the other hand, observers found it much more difficult to find the ellipses among the bumps. The slopes of search times versus set size were 28.7 msec/item for the target trials and 29.5 msec/item on blank trials. The slope differences were significant, for both the target-present trials (paired $t_5 = 4.8, p < .005$) and the target-absent trials (paired $t_5 = 5.8, p < .005$). The results show that the bumps are found very efficiently among the ellipses, whereas the converse does not hold. There is thus a classic search asymmetry between the two stimuli.

Table 1
Error Rates (in Percentages) in Experiment 1

Set size	Bumps Among Ellipses		Ellipses Among Bumps	
	Present	Absent	Present	Absent
4	3.2	2.6	3.5	4.3
8	3.4	4.2	3.5	4.4
12	4.5	4.2	4.7	5.3
16	4.4	5.3	6.3	4.7

There are at least two possible interpretations of this result: It could be that the curvature discontinuities on the bumps pop out. Another possibility is that the fast search rates for the bumps among the ellipses are due to the different 3-D interpretations of the two stimuli. The latter is a distinct possibility since many studies show that stimuli differing only in a 3-D interpretation pop out, such as stimuli that appear to be lit from below among ones that appear to be lit from above (Enns & Rensink, 1990; Ramachandran, 1988; Sun & Perona, 1996), or stimuli at different depths (He & Nakayama, 1992; Nakayama & Silverman, 1986). In fact, Sun and Perona (1996) showed that fast parallel processing of certain patterns is *dependent* upon a 3-D-shape representation. These studies indicate that some type of 3-D interpretation of the visual environment is arrived at relatively early in visual processing. Thus, observers in our study might have quickly constructed a 3-D representation of the bumps and ellipses and performed the search over these higher level representations, meaning that the curvature discontinuities themselves did not pop out.

In Experiments 2A and 2B, we investigated whether contour curvature discontinuities in the image would pop out among smooth changes in curvature when a 3-D interpretation is not as easily attainable as in Experiment 1.

EXPERIMENT 2A Filled Half-Bumps and Half-Ellipses

Experiment 2A controlled for the possibility that it is not curvature discontinuities per se that pop out, but rather a 3-D interpretation that might have derived from them. The stimuli used in Experiment 2A did not have as strong 3-D interpretations as did those in Experiment 1 (Figure 4a). The "half-bumps" in Experiment 2A were simply one half of the bumps of Experiment 1 cut vertically. If the results of Experiment 1 were due to efficient search for abrupt changes in curvature, we would expect a similar pattern of results for Experiments 1 and 2A. However, if the pop-out of bump stimuli in Experiment 1 was simply due to different 3-D interpretations of the stimuli, a different pattern of search times should be observed in the two experiments. In particular, there should be no pop-out of the half-bump stimuli if pop-out of the bumps in Experiment 1 was due to their 3-D interpretation and not the presence of curvature discontinuities.

Method

The 6 participants were undergraduates at Harvard University participating in the experiment for course credit. The stimuli that were used are presented in Figure 4a. The observers were told to search for either the stimulus with abrupt changes in curvature or the one with smooth changes in curvature, run in separate blocks. The meanings of the terms *curvature* and *abrupt change in curvature* were explained to participants if necessary, and examples of corresponding bump and ellipse stimuli were given. The height of the half-bumps and half-ellipses at a 64-cm viewing distance was 1.68° , and their width was 1.12° . The heights, widths, and areas of the targets and distractors were identical. The orientation of the

Figure 4. The stimuli used in Experiments 2A and 2B. (a) The filled half-bumps and half-ellipses from Experiment 2A and (b) the corresponding outline versions from Experiment 2B.

Figure 5. Overall response times for the four different conditions of Experiment 2A. The error bars show the largest standard error of the mean for each condition.

stimuli was in every case as shown in Figure 4. In all other respects methods were similar to those described for Experiment 1.

Results and Discussion

The results of Experiment 2A are presented in Figure 5 and the error rates are depicted in Table 2. They replicate the search asymmetry of Experiment 1. A similar search asymmetry was observed between the stimuli with and without curvature discontinuities. The stimuli with curvature discontinuities pop out among the ones with no curvature discontinuities, whereas the converse does not hold. The slopes of search times versus set size were 12.1 msec/item for the target trials and 28 msec/item on blank trials when the target contained a curvature discontinuity, but the same slopes were 18.9 msec/item for the target trials and 80.2 msec/item on blank trials when the target contained no curvature discontinuities. The slope differences were significant, both for the target trials (paired $t_5 = 3.9$, $p < .005$) and for the absent trials

(paired $t_5 = 6.2$, $p < .005$). Even though the half-bumps appear to be less easily found than the full bumps in Experiment 1, the same search asymmetry is still present and the general pattern is similar. The differences in absolute search rates and slopes might be due to the fact that different groups of participants were used for Experiment 1 and Experiments 2A and 2B.

The results of Experiment 2A suggest that the presence of curvature discontinuities was directly responsible for the fast search rates for the bumps among the ellipses in Experiment 1. If the results of Experiment 1 were primarily due to the existence of a 3-D interpretation, we would not have found a similar pattern of results for the stimuli that do not have such a salient 3-D interpretation.

EXPERIMENT 2B

Outline Half-Bumps and Half-Ellipses

To reinforce our conclusions from Experiment 2A, we also tested performance on the "outline half-bumps" in Figure 4b, which are simply outlined presentations of the half-bumps in Figure 4a. The reason for using the outline stimuli is that a 3-D interpretation similar to the one in Experiment 1 might still be possible for the half-bumps in Figure 4a, whereas a 3-D interpretation of the outline half-bumps is difficult if not impossible. Moreover, both the filled half-bumps and half-ellipses have tangent, and therefore also curvature, discontinuities on their right sides (see Figure 4a), where the original bumps and ellipses have been cut vertically in half. It could be that these spurious curvature discontinuities affected search times for half-ellipses differently than for half-bumps. Outline half-ellipses and half-bumps are a good control for this possible confound because, ignoring endpoints, only the outline half-bumps contain curvature discontinuities. In contrast to filled half-ellipses, outline half-ellipses do not contain curvature discontinuities at all.

Method

The procedures were in every respect the same as those of Experiment 2A except for the difference in stimuli: In Experiment 2B the target and distractors were the stimuli in Figure 4b. Their height and width were the same as those described for the stimuli in Experiment 2A. The same observers participated in Experiment 2B as in Experiment 2A.

Results and Discussion

The results of Experiment 2B, shown in Figure 6, are almost identical to the results of Experiment 2A (see Table 3 for error rates). The slopes of search times versus

Table 2
Error Rates (in Percentages) in Experiment 2A

Set size	Half-Bumps Among Ellipses		Half-Ellipses Among Bumps	
	Present	Absent	Present	Absent
4	2.9	3.5	3.6	4.6
8	3.4	4.6	3.8	5.1
12	4.6	3.9	5.6	5.4
16	5.3	5.1	5.8	5.5

Figure 6. Overall response times for the four different conditions of Experiment 2B. The error bars show the largest standard error of the mean for each condition.

Table 3
Error Rates (in Percentages) in Experiment 2B

Set size	Outlined Half-Bumps Among Ellipses		Outlined Half-Ellipses Among Bumps	
	Present	Absent	Present	Absent
4	2.3	3.7	3.1	5.1
8	2.9	3.9	3.9	4.1
12	4.9	5.1	4.3	4.9
16	4.8	5.6	4.7	5.8

set size were 10.5 msec/item for the target trials and 24 msec/item on blank trials when the target contained a curvature discontinuity, but the same slopes were 37 msec/item for the target trials and 71.6 msec/item on blank trials when the target contained no curvature discontinuities. The slope differences were significant, both for the target trials (paired $t_5 = 5.7, p < .005$) and for the absent trials (paired $t_5 = 5.9, p < .005$). The asymmetry between the stimuli that contained curvature discontinuities and those that did not is even larger in this experiment than in Experiment 2A, and thus the result supports our conclusion that it is curvature discontinuities that are easily found in Experiment 1, even though the slopes of reaction time versus set size were somewhat smaller in Experiment 1.

EXPERIMENT 3
High Versus Low Rates of Change of Curvature

Even though Experiments 1 and 2 strongly suggested that curvature discontinuities pop out among smooth changes in curvature, another interpretation is still possible. It could be that what popped out was not the abrupt changes in curvature among smooth changes in curvature, but rather a higher *rate of change* in curvature among lower rates of change in curvature. Experiment 3 was designed to decide between these two possibilities.

Imagine driving along a contour at a constant distance per unit of time, as if the contour were a road. The sharpness with which one must turn the driving wheel corresponds to the sharpness of a curvature change along the contour. The bumps have sharper curvature changes than the ellipses around the region of their curvature discontinuities. A curvature change can be sharp without being a curvature discontinuity. For example, an ellipse whose major axis is much larger than its minor axis will have a small radius of curvature at its maxima of curvature without having any curvature discontinuities along its contour. So it may not be that abrupt curvature changes per se pop out. It may merely be that the figure with a smaller radius of curvature at its maxima of curvature pops out among distractors with larger radii of curvature at their maxima of curvature.

Experiment 3 contrasts arcs of ellipses that have either high or low rates of change of curvature. Ellipses can have high or low rates of change in curvature, but no curvature discontinuities, as we proceed along their contour. The first and second derivative are well defined for all points along an ellipse's contour. If the explanation for the results of Experiments 1 and 2 is simply that high rates of curvature change pop out among lower rates of curvature change, we should expect to see a search asymmetry similar to the one observed between the bumps and ellipses in Experiment 1, and the half-bumps and half-ellipses in Experiments 2A and 2B when arcs of ellipses that have different aspect ratios (the ratio of the length of the major over the minor axes) are contrasted.

To control for possible confounding effects of the size of the stimuli in Figure 7 (arcs of ellipses with high aspect ratios are larger than those with low aspect ratios when other parameters are constant), we equated the length of the lines making up the arcs so that the area of the rectangle that each arc could fit into was equal. Furthermore, each arc could be of five different sizes on any trial. Randomizing the size in this way limits the influence of the size of the stimuli on the results, since size becomes irrelevant to the task.

Figure 7. The arcs of ellipses used in Experiment 3. The one on the left has a high aspect ratio (height/width), whereas the one on the right has a lower one. Note that the low aspect ratio ellipse has a smaller radius of curvature at its maximum of curvature. Note also that the areas of the rectangles surrounding the two arcs are equal.

Figure 8. Overall response times for the four different conditions of Experiment 3. The error bars show the largest standard error of the mean for each condition.

Method

The 6 participants were members of the Vision Sciences Laboratory at Harvard University and were unaware of the purpose of the experiment. The stimuli are presented in Figure 7. The heights and widths of the rectangles that the arcs in the experiment fit into varied randomly on any given trial from $1.35^\circ \times 0.9^\circ$ to $1.89^\circ \times 1.26^\circ$. The size of each display item varied independently on each trial. Observers either searched for a high aspect ratio arc (low rate of change of curvature) among low aspect ratio arcs (high rates of change of curvature) or the converse—they searched for a low aspect ratio arc among arcs with high aspect ratios. In the experiment the stimuli were presented without the rectangles shown in Figure 7. Three different set sizes (4, 10, and 16) were tested. Otherwise, methods were similar to those described for the previous experiments.

Results and Discussion

The results from Experiment 3, shown in Figure 8 (cf. Table 4), reveal that the search asymmetries we observed in Experiments 1 and 2 were not due to the fact that the stimuli with curvature discontinuities simply have a higher rate of change of curvature per unit contour distance. It is also not the case that a target with a smaller radius of curvature at its maximum of curvature pops out among distractors with larger radii of curvature at their maxima of curvature. If either high curvature change or small radii of curvature at maxima of curvature pop out, we should have seen faster search speeds and flatter search slopes for the arcs with the low aspect ratio than for the ones with the high aspect ratio, but that pattern is

not present in the data. If anything, the high aspect ratio arcs are easier to find than the low aspect ratio ones. The slopes of response times versus set size for the high aspect ratio arcs were 4.9 msec/item for target trials and 29.9 msec/item for blank trials. The same slopes for the low aspect ratio arcs were 15.6 msec/item for the target trials and 24 msec/item for the blank trials. The slope differences for the high versus low aspect ratio arcs were significant for the present trials (paired $t_5 = 4.6$, $p < .01$) but not for the absent trials (paired $t_5 = 0.75$, $p > .25$).

The results rule out differences in the rate of curvature change as the explanation for the results of the previous experiments. The results reinforce our conclusion that the crucial image information that was easily found in our first two experiments was contour curvature discontinuity.

EXPERIMENT 4 The Role of Symmetry

There is another possible confound that could account for the results of the previous experiments. The stimuli that contain curvature discontinuities are less symmetrical than the ones with smooth curvature changes. For example, the ellipse in Figure 2a is symmetrical about both its major and minor axes, whereas the bump in Figure 2a is symmetrical only about its minor axis. Research has shown that symmetry can play an important role in the analysis of the visual environment (Gurnsey & Browse, 1989; Treisman & Gormican, 1988).

In order to assess the role played by symmetry in our previous results, we tested search performance for the stimuli depicted in Figure 9. The S in panel a contains two curvature discontinuities. It is made by rearranging the parts of a bump obtained by cuts along its major and minor axes. The S in panel b contains no curvature discontinuities. It is constructed from the two parts of an ellipse that result from a cut along its minor axis. Both

Table 4
Error Rates (in Percentages) in Experiment 3

Set size	High Aspect Ratio Among Low		Low Aspect Ratio Among High	
	Present	Absent	Present	Absent
4	3.5	5.3	4.0	5.4
10	3.7	4.3	3.8	6.0
16	4.1	6.5	7.1	7.5

Figure 9. The stimuli used in Experiment 4. Note that the “sharp S” in panel a contains two curvature discontinuities, whereas the “smooth S” in panel b contains no curvature discontinuities. Also note that the overall length of the two Ss is equal.

types of stimuli are invariant under a 180° rotation. Most importantly, neither type of S is bilaterally symmetric. If the results of Experiments 1 and 2 are merely due to the differences in symmetry between targets and distractors, and are not due to the presence of curvature discontinuities, we should not find a search asymmetry for the stimuli shown in Figure 9. If we do find an asymmetry, it must be due to the presence of curvature discontinuities, and not to the presence of symmetry.

Method

The procedures were the same as in previous experiments. The same observers participated in this experiment as in Experiment 3.

Results and Discussion

As can be seen in Figure 10 (cf. Table 5), a strong search asymmetry was found between the stimuli with and without curvature discontinuities in Experiment 4. It is far easier to find a sharp S with curvature discontinuities among smooth Ss without them, than vice versa. The slopes for the sharp Ss among smooth ones were 26.3 msec/item for present trials and 65.6 msec/item for absent trials, but the same slopes for the smooth Ss among sharp ones were 63.1 msec/item and 133.4 msec/item, respectively. These slope differences were significant (paired $t_5 = 2.94$, $p < .05$ for present trials and paired $t_5 = 5.81$, $p < .005$ for absent trials). These results are therefore not due to symmetry, but to the presence or absence of curvature discontinuities in the stimuli. They support our conclusion that the reason for the search asymmetries in Experiments 1 and 2 is the presence or absence of curvature discontinuities, not differences in symmetry between the stimuli.

The slopes of set size versus response time are, however, slightly higher than in the previous experiments, but this may be due to the fact that the two types of S's used in the experiment share many features, making the search harder. The important result is that there is a similar search asymmetry between stimuli with and without curvature discontinuities, as in the previous experiments.

EXPERIMENT 5 Occluded Curvature Discontinuities Are Easily Found

What is search performance like when the curvature discontinuities on the bumps are occluded (Figure 11)? This question is interesting because it may be relevant to whether curvature discontinuities are amodally completed behind occluders, given global configurations that suggest the presence of such a discontinuity in an oc-

Figure 10. Overall response times for the four different conditions of Experiment 4. The error bars show the largest standard error of the mean for each condition.

Table 5
Error Rates (in Percentages) in Experiment 4

Set size	Smooth Ss Among Sharp Ss		Sharp Ss Among Smooth Ss	
	Present	Absent	Present	Absent
4	7.6	7.5	3.8	4.8
10	9.1	8.9	3.4	5.1
16	9.6	10.2	4.4	5.6

cluded part of the visual environment. In Experiment 5 we tested performance with “partially occluded bumps and ellipses” on which the curvature maxima were occluded.

Method

The 6 participants were undergraduates at Harvard University and took part in the experiment for course credit. They searched for either the partially occluded bump in Figure 11 among the partially occluded ellipses, or vice versa. The height and width of the stimuli were the same as in Experiment 1 except for the circular black occluders, whose diameter was 0.75° . Otherwise, the procedures were similar to those described for Experiment 1.

Results and Discussion

Once again the pattern of search rates was similar to that of Experiment 1 (Figure 12; see Table 6 for error rates). The partially occluded bumps were quickly and easily found among the partially occluded ellipses. However, when a partially occluded ellipse was to be found among partially occluded bumps, search was a lot harder. The slopes of search times versus set size were 11.28 msec/item (target trials) and 52 msec/item (blank trials) when observers searched for the bump among ellipses and 41.5 msec/item (target trials) and 88.3 msec/item (blank trials) when observers searched for the ellipse among bumps. These differences were significant (paired

Figure 11. The partially occluded outline bumps and ellipses used in Experiment 5. Note that no curvature discontinuities are visible.

Table 6
Error Rates (in Percentages) in Experiment 5

Set size	Occluded Bumps Among Ellipses		Occluded Ellipses Among Bumps	
	Present	Absent	Present	Absent
4	3.1	3.0	3.4	4.2
8	3.5	4.1	3.6	4.5
12	4.6	4.3	4.6	5.2
16	5.1	5.2	5.5	4.8

$t_5 = 7.1, p < .001$) for present trials and (paired $t_5 = 8.3, p < .001$) for absent trials.

Although these results are by no means conclusive, they do suggest that analysis of contour curvature and curvature (or tangent) continuity and discontinuity happens at least in some cases after amodal completion has taken place. This result is consistent with the finding of He and Nakayama (1992) that an L-shaped stimulus is available for visual search only after a stage of amodal completion (see also Baylis & Driver, 1995). The results of this experiment suggest that contour curvature discontinuities do indeed pop out, but that they need not be visible in the image in order to do so. Their existence need only be implied by other image cues in order to pop out.

EXPERIMENT 6 The Degree of Pop-Out Due to Curvature Discontinuities

Experiments 1–5 established that curvature discontinuities in silhouettes pop out. This is the stage where visual search studies often conclude, as if pop-out were an all-or-nothing affair. Although search for an image feature may lead to flat search slopes when the strength of that feature present in the target is much larger than the strength of that feature present in the distractors, pop-out is not necessarily an all-or-nothing affair. Pop-out may occur above a certain threshold of target–distractor difference along the feature dimension that defines the target as a target (Doshier, 1998; Nakayama & Joseph, 1998; Verghese & Nakayama, 1994). Below this threshold, however, search slopes may not be flat. In the present case, the degree of curvature discontinuity, corresponding to the degree to which r and R differ in Figure 1, may affect search slopes. The two limiting cases of curvature discontinuities are first-order (or tangent) discontinuities on the one hand and smooth changes of curvature on the other. It is known that corners pop out among straight line distractors, and that curves also pop out among straight line distractors (Treisman & Gormican, 1988; Wolfe et al., 1992). However, we do not know how search slopes are affected by changing the degree of a parameter that links corners and smooth curves along the single dimension of curvature discontinuity abruptness. If the abruptness of a curvature discontinuity (defined, say, by R/r ; see Figure 1) increases, we would expect performance on a task where the target contains a curvature discontinuity to improve up to a certain “pop-out threshold” beyond which all search slopes would be flat as a function of set size. In Experiment 6 we tested search performance when observers searched for a bump among ellipses, while varying the abruptness of the curvature discontinuity from one trial to the next.

Method

We used five different targets while the distractors stayed the same (see Figure 13). The areas, heights, and widths of all the tar-

Figure 12. The response times for the four conditions of Experiment 5. The error bars denote the largest standard error of the mean for each condition.

gets and distractors were the same, and the ratio between the minor axes of the two half-ellipses that made up the bump (Figure 2) varied between the five different targets (Figure 13). As this ratio increased (the ratio R/r in Figure 1 or the ratio b/c in Figure 2B), the sharpness of the curvature discontinuity increased. The height of the bumps and ellipses was 1.88° and their width was 2.23° at a viewing distance of 64 cm. Three different set sizes were tested (4, 10, and 16). The target was present on 60% of the trials and could be one of the five targets in Figure 13 picked randomly from trial to trial. Six observers participated in 300 experimental trials preceded by 20 practice trials. The procedures were similar in all other respects to those of previous experiments.

Results and Discussion

As the curvature discontinuity on the bumps got sharper and approached a first-order discontinuity, performance on the search task improved. Both the slopes of response

times versus set size and the y-intercepts decreased as the sharpness of the curvature discontinuity increased. Performance on target trials for the five different targets is plotted in Figure 14. The slopes of set size versus response times are plotted as black squares against the left axis in Figure 15 for the five targets and the blank trials; error rates are plotted as bars (axis on right). Error rates were quite low for all the conditions (0.5%–2%), except when the target was the one with the least abrupt curvature discontinuity (ratio = 1.49), where observers missed the target on 15% of the trials. The search rate for that target was also by far the slowest. Since performance is worst for the least abrupt discontinuity, it is tempting to conclude that this particular discontinuity falls outside the range of discontinuities to which the visual system is highly tuned. This experiment also suggests that pop-out

Figure 13. The stimuli used for Experiment 6. (a) The five different targets in the experiment. The numbers underneath refer to the ratio of the larger over the smaller half-ellipse minor axes that constitute the upper and lower halves of the bump stimulus. (b) The distractors were ellipses equal in area, width, and height to the bumps serving as targets. Since the heights and widths of all five target bumps are the same, the same ellipse distractor was used with all five targets.

Figure 14. The results of the target trials of Experiment 6. The results for the five different targets (differing in ratio, see Figure 13) are shown. The error bars show the largest standard error of the mean for each condition.

is not an all-or-nothing affair. Rather, pop-out occurs in degrees depending on how similar or different targets and distractors are along the relevant search dimension.

GENERAL DISCUSSION

The experiments reported in this paper suggest the following conclusions:

1. Stimuli that contain curvature discontinuities are easily found among stimuli that contain only smooth changes in curvature.
2. The search results are highly asymmetrical; that is, stimuli lacking such discontinuities are hard to find among those that do contain curvature discontinuities.

3. The results cannot be accounted for by the differences in apparent 3-D shape between the bumps and ellipses.

4. The search asymmetry is not due to the fact that the stimuli with the curvature discontinuities have a higher rate of change of curvature in the neighborhood of their maxima of curvature. Furthermore, the results cannot be accounted for by differences in symmetry between the stimuli that do and do not contain curvature discontinuities.

5. Curvature discontinuities may not need to be visible in the image in order to pop out. They may merely need to be implied in the scene. Indeed, curvature discontinuities may be subject to search only after a stage of amodal completion.

6. Effortless search is not an all-or-nothing affair. Rather, it depends on the degree to which the relevant search dimension is present in the targets and in the distractors. There may be a "pop-out threshold" along that featural dimension beyond which visual search is suddenly very hard and prone to errors.

Perhaps the most important result is that we find a search asymmetry similar to the one reported by Treisman and Gormican (1988). Whereas they found that curves are easily found among straight lines, but that the opposite is not true, we find that abrupt curvature changes are easily found among smooth curvature changes, but that the opposite is not true. We have hypothesized that search asymmetries exist because the visual system is highly sensitive to only certain types of information. If it is "looking for" curvature information, but not for information about straight lines, then a curve will be found rapidly among straight line distractors. But the opposite will not be true. Straight lines will not be found rapidly among curves if the visual system is not "on the lookout" for straight lines. Similarly, if the visual system is highly tuned to detect curvature discontinuities, a stimulus with an abrupt curvature change will pop out among distractors that lack

Figure 15. The slopes and error rates for the six different conditions of Experiment 6; the five different target ratios (see Method section and Figure 13) and the target-absent trials. The black squares denote the slopes (axis on left) and the striped bars denote the error rates (axis on right).

such a discontinuity. However, a target with smooth curvature change as one progresses along the contour will not pop out among distractors that do possess curvature discontinuities. This asymmetry can, moreover, be taken as strong evidence that the visual system is highly sensitive to curvature discontinuities, which in turn suggests that curvature discontinuities are image cues that are particularly informative about world structure.

The stimuli in Figure 2a clearly show the dramatic influence that the addition of curvature discontinuities can have on the 3-D interpretation of what might otherwise just look like a flat silhouette. This suggests that the visual system seeks out discontinuities in curvature and uses them, among other cues, to obtain a 3-D representation of the 2-D retinal image. Curvature discontinuities may thus be added to the list of important pictorial information that the visual system uses to recover the third dimension in the visual environment. A demonstration of the role played by subtle curvature discontinuity differences in generating a 3-D shape and scene interpretation can be found in Figure 13. First, the presence of curvature discontinuities makes the silhouettes in Figure 13a appear more volumetric than the ellipse in Figure 13b. Tse (in press) has argued that curvature discontinuities are an important cue to a volume's shape because they can reveal potential information about the shape of cross sections of a volume. Second, assuming a round cross section, the sharpness of the curvature discontinuity can reveal information about the viewing angle or the angle of the plane that supports an object, or in which an object is embedded. For example, the five silhouettes in Figure 13a might all derive from the same 3-D object viewed from different positions. Or they might be the image projections of different volumes viewed from the same position. Lastly, the supporting ground plane that supports these volumes might have different slants (see Tse, 2000). Although the sharpness of curvature discontinuities can raise several ecologically plausible interpretations, in isolation they cannot disambiguate between them. To accomplish this, certain assumptions about ground plane orientations, viewing angles, and the most probable shapes of objects in the world may be required (Albert & Tse, 2000; Tse, 1998, 1999a, 1999b, 2000).

Curvature Discontinuities and Tangent Discontinuities

If the visual system is highly tuned to detect curvature discontinuities, it will always detect a tangent discontinuity because a tangent discontinuity is always a curvature discontinuity. However, if the visual system were only highly tuned to detect tangent discontinuities, it would not detect curvature discontinuities below some threshold of abruptness. As our experiments show, the visual system is very sensitive to curvature discontinuities. Although it may be that the visual system is tuned to both tangent and curvature discontinuities, it may only be sensitive to curvature discontinuities, since it would then detect tangent discontinuities, as it were, "for free." Indeed,

in peripheral vision, tangent and curvature discontinuities can be indistinguishable because of poor spatial resolution. Experiment 6, moreover, suggests that tangent discontinuities may merely be a limiting case of the curvature discontinuity dimension, rather than a feature dimension existing independently of the curvature discontinuity dimension. Future work will have to determine whether tangent discontinuity detection is governed by its own set of filters with their own properties and detection thresholds, or whether tangent discontinuity detection is accomplished using the same filters that govern contour curvature discontinuity detection.

Different Degrees of Pop-Out

We defined "pop-out threshold" as the point along some feature dimension that defines a target as a target beyond which search slopes are flat as a function of set size. This might be an unrealistic criterion, because search slopes as a function of set size are rarely truly flat. A more interesting criterion to specify the pop-out (rapid and accurate search) domain is suggested by our data. Figure 15 shows a nonlinear jump in the error rates for the target with the least abrupt contour curvature discontinuity. This suggests that this degree of curvature discontinuity falls outside the range of curvature discontinuities easily detected by the visual system. Future work will have to determine whether there is a critical threshold beyond which error rates and search slopes are relatively low. If so, this might be a useful criterion for defining a pop-out threshold. We have attempted to demonstrate the existence of such a threshold for the particular feature dimension defined by curvature discontinuities; parametric studies similar to Experiment 6 could be done for other feature dimensions as well, such as color, motion, size, and so forth. Indeed, all the traditional "primitive" features that pop out could be parametrically studied to determine their pop-out domain.

Conclusions

Stimuli that contain curvature discontinuities are easily found among stimuli containing only smooth changes in curvature. On the other hand, stimuli with smooth curvature changes are hard to find among stimuli containing curvature discontinuities. These results suggest that the visual system detects and analyzes abrupt changes in curvature in the image quickly to extract vital information about the 3-D structure of the visual environment.

REFERENCES

- ALBERT, M. K., & HOFFMAN, D. D. (1995). Genericity in spatial vision. In R. D. Luce (Ed.), *Geometric representations of perceptual phenomena: Papers in honor of Tarow Indow on his 70th birthday* (pp. 95-112). Mahwah, NJ: Erlbaum.
- ALBERT, M. K., & TSE, P. U. (2000). The role of surface attachment in perceived volumetric shape. *Perception*, **29**, 409-420.
- ATTNEAVE, F. (1954). Some informational aspects of visual perception. *Psychological Review*, **61**, 183-193.
- BARROW, H. G., & TENENBAUM, J. M. (1981). Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence*, **17**, 75-116.
- BAYLIS, G., & DRIVER, J. (1995). One-sided edge assignment in vision:

1. Figure-ground segmentation and attention to objects. *Current Directions in Psychological Science*, **4**, 140-146.
- BEUSMANS, J. M. H., 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.
- BINFORD, T. O. (1981). Inferring surfaces from images. *Artificial Intelligence*, **17**, 205-244.
- CLOWES, M. B. (1971). On seeing things. *Artificial Intelligence*, **17**, 79-116.
- COURANT, R., & JOHN, F. (1989). *Introduction to calculus and analysis I*. New York: Springer-Verlag.
- DOSHER, B. (1998). Models of visual search: Finding a face in a crowd. In D. Scarborough & S. Sternberg (Eds.), *An invitation to cognitive science: Vol. 4. Methods, models and conceptual issues* (pp. 455-521). Cambridge, MA: MIT Press.
- ENNS, J. T., & RENSINK, R. A. (1990). Sensitivity to three-dimensional orientation in visual search. *Psychological Science*, **1**, 323-326.
- GIBSON, J. J. (1950). *The perception of the visual world*. Boston: Houghton Mifflin.
- GURNSEY, R., & BROWSE, R. A. (1989). Asymmetries in visual texture discrimination. *Spatial Vision*, **4**, 31-44.
- HE, Z. J., & NAKAYAMA, K. (1992). Surfaces versus features in visual search. *Nature*, **359**, 231-233.
- HOFFMAN, D. D. (1998). *Visual intelligence: How we create what we see*. New York: W. W. Norton.
- HOFFMAN, D. D., & RICHARDS, W. A. (1984). Parts of recognition. *Cognition*, **18**, 65-96.
- HUFFMAN, D. A. (1971). Impossible objects as nonsense sentences. *Machine Intelligence*, **5**, 295-323.
- KELLMAN, P. J., & SHIPLEY, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, **23**, 141-221.
- LOWE, D. G. (1987). Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, **31**, 355-395.
- MALIK, J. (1987). Interpreting line drawings of objects. *International Journal of Computer Vision*, **1**, 73-107.
- NAKAYAMA, K., & JOSEPH, J. S. (1998). Attention, pattern recognition, and pop-out in visual search. In R. Parasuraman (Ed.), *The attentive brain* (pp. 279-298). Cambridge, MA: MIT Press.
- NAKAYAMA, K., & SHIMOJO, S. (1992). Experiencing and perceiving visual surfaces. *Science*, **257**, 1357-1363.
- NAKAYAMA, K., SHIMOJO, S., & SILVERMAN, G. (1989). Stereoscopic depth: Its relation to image fragmentation, grouping, and the recognition of occluded objects. *Perception*, **18**, 55-68.
- NAKAYAMA, K., & SILVERMAN, G. (1986). Serial and parallel processing of visual surface conjunctions. *Nature*, **320**, 264-265.
- RAMACHANDRAN, V. S. (1988). Perceiving shape from shading. *Scientific American*, **259**, 76-83.
- RICHARDS, W. A., DAWSON, B., & WHITTINGTON, D. (1986). Encoding contour shape by curvature extrema. *Journal of the Optical Society of America A*, **3**, 1483-1491.
- RICHARDS, W. A., KOENDERINK, J. J., & HOFFMAN, D. D. (1987). Inferring three-dimensional shapes from two-dimensional silhouettes. *Journal of the Optical Society of America A*, **4**, 1168-1175.
- ROCK, I. (1984). *The logic of perception*. Cambridge, MA: MIT Press.
- SINGH, M., SEYRANIAN, G. D., & HOFFMAN, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception & Psychophysics*, **61**, 636-660.
- SUN, J. Y., & PERONA, P. (1996). Preattentive perception of elementary three-dimensional shapes. *Vision Research*, **36**, 2515-2529.
- TREISMAN, A., & GORMICAN, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, **95**, 15-48.
- TSE, P. U. (1998). Illusory volumes from conformation. *Perception*, **27**, 977-994.
- TSE, P. U. (1999a). Complete mergeability and amodal completion. *Acta Psychologica*, **102**, 165-201.
- TSE, P. U. (1999b). Volume completion. *Cognitive Psychology*, **39**, 37-68.
- TSE, P. U. (2000). The sawtooth illusion. *Perception*, **29**, 874-875.
- TSE, P. U. (in press). A contour propagation approach to surface filling-in and volume formation. *Psychological Review*.
- TSE, P. U., & ALBERT, M. K. (1998). Amodal completion in the absence of image tangent discontinuities. *Perception*, **27**, 455-464.
- VERGHESE, P., & NAKAYAMA, K. (1994). Stimulus discriminability in visual search. *Vision Research*, **34**, 2453-2467.
- WALTZ, D. (1975). Understanding line drawings of scenes with shadows. In P. H. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill.
- WOLFE, J. M., YEE, A., & FRIEDMAN-HILL, S. R. (1992). Curvature is a basic feature for visual search tasks. *Perception*, **21**, 465-480.
- ZUCKER, S. W., DOBBINS, A., & IVERSON, L. (1989). Two stages of curve detection suggest two styles of visual computation. *Neural Computation*, **1**, 68-81.

(Manuscript received March 2, 2000;
revision accepted for publication October 10, 2000.)