CIS-TR 85-08

The Concave Cusp as a Determiner of Figure-Ground

Kent A. Stevens
Allen Brookes

Department of Computer and Information Science
University of Oregon

THE CONCAVE CUSP AS A DETERMINER OF FIGURE-GROUND[1]

Kent A. Stevens and Allen Brookes

Department of Computer and Information Science
University of Oregon
Eugene, OR 97403

Abstract. There are known tendencies to interpret as figure against background those regions that are lighter, or smaller or, especially, regions that are more convex. Wherever convex, opaque objects abut or partially occlude one another in an image, the points of contact between the silhouettes form concave cusps, each indicating the local assignment of figure versus ground across the contour segments. We propose that this local geometric feature is a preattentive determiner of figure-ground, and that it contributes to the previously-observed tendency to prefer convexity in general. Evidence is given that (1) figure-ground can be determined solely on the basis of the concave cusp feature, and (2) the salience of the cusp derives from the local geometry and not from the adjacent contour convexity.

1. Introduction

Figure-ground, the process of distinguishing a figure relative to its surround, is often exemplified by Rubin's [1958] vase-face illusion. The contours which, in one figure-ground interpretation, comprise the silhouette of a vase, can also be seen as the edges of two faces in silhouette. The figure-ground interpretation in such illustrations is readily affected by focussed attention and verbal suggestion. Figure-ground presumably has a substantial preattentive component as well, with factors such as size, brightness, symmetry, closure, regularity, and convexity contributing to an "immediate" impression of figure against ground. Random texture tends to segregate perceptually on the basis of contrast, such that the white patches might appear as figure against the grey and black patches, or the black patches against the white and grey [Julesz 1965; Richards & Purks 1978]. For such random texture there is a tendency to see white-as-figure and to see smaller-as-figure [Rubin 1958]. These tendencies are balanced when roughly 0.4 of the area is black [Frisch & Julesz 1966]. Symmetry, if present in the texture, dominates size (and contrast), causing one to prefer the symmetric forms even if larger (or darker) [Hochberg 1964]. Still more important is convexity, which is favored over both regularity and symmetry [Kaniza & Gerbino 1976].

The importance of convexity in determining figure-ground is demonstrated in figure 1, generated by placing convex blobs of various shapes at random locations. The immediate interpretation is usually of a texture comprised of small convex objects, regardless of the contrast sign (compare figures 1a and 1b). Note that virtually all of the texture area (81%) is seen as figure; only the small irregular fragments are seen as ground, in strong distinction to the trend reported by Frisch and Julesz [1966]. Of course, while the placement of the blobs is random in figure 1, the resulting texture is not random, owing to the regular, convex, geometry of the individual blobs. For comparison, the texture in figure 2 is composed of irregular, concave, blobs. Here the figure-ground sense is more ambiguous and exhibits the tendencies Frisch and Julesz [1966] discussed. It is easily shown that convexity becomes important only when the two figure-ground alternatives are spatially adjacent, so that the assignment of contour to figure is locally ambiguous. When the same elements as comprise figure 2 are sufficiently isolated they can be seen as figure even though they are concave.

[1] The research reported in this article was supported by Air Force Office of Scientific Research contract number F49620-83-C-0093.

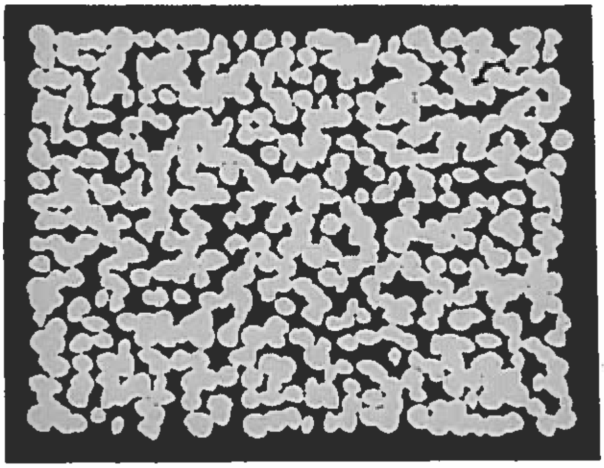

A

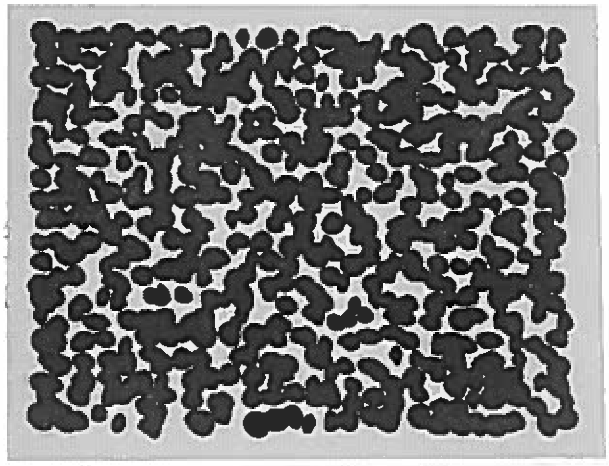

B

Figure 1. Texture perceived as small convex blobs in close packing, independent of contrast sign.

Convexity, as a determiner of figure-ground, derives from the fact that most physical textures are composed of compact surfaces that are convex almost everywhere. They give rise to image texture that, in general, have convex outlines or silhouette contours. But how is convexity measured and processed by the visual system?

2. Computational Issues

2.1 The Texture Parsing Problem

Most textures are the image projections of discrete and compact physical objects. An image of leaves against the bright sky, for instance, results in a mottled texture of leaves silhouetted against a bright background. In the early visual processing of such an image, there is presumably a representation of the intensity changes corresponding to the leaf silhouettes. Locally, a binary decision is necessary as to which side of the edge corresponds to figure. This decision effectively parses, two-dimensionally, the set of edges into one of two possible figure-ground interpretations. For the silhouetted leaves, the correct parsing treats the dark regions as figure; the alternative (and incorrect) parsing treats as figure the fragments of sky visible between the leaves. The texture elements that result from the latter parsing, of course, do not correspond to individual physical objects; the shape of each is generated by the placement, shape, and orientation of the surrounding leaves.

The shape properties of the visible portion of the silhouette depend on which side is seen as figure [Hoffman & Richards 1982, 1984], and thus the visual description of texture depends on the figure-ground interpetation. Moreover, geometric shape properties are meaningfully attributed only to the images of physical objects (e.g. the leaves and not the random shapes of the

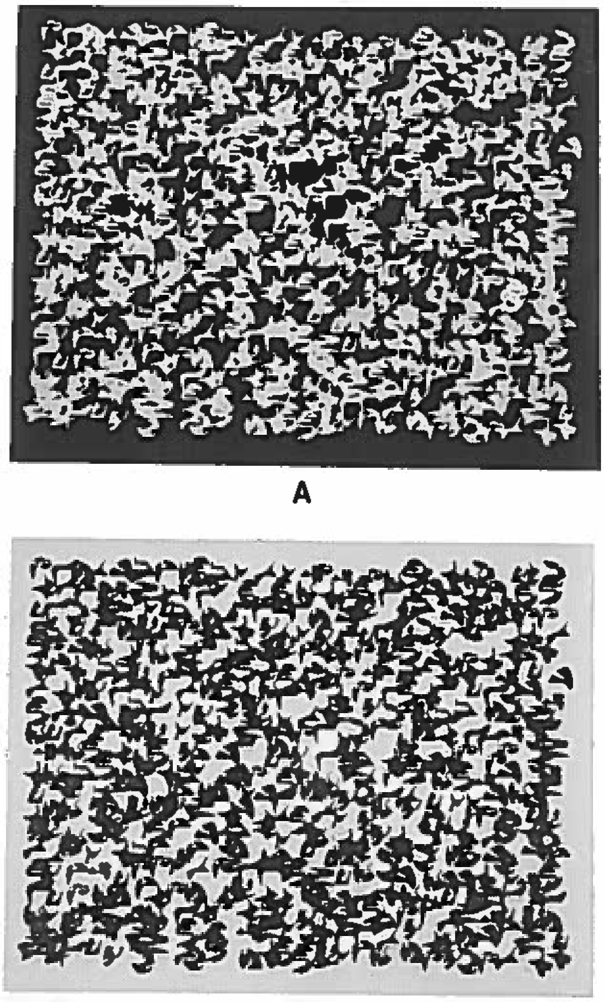

A

B

Figure 2. Texture composed of concave, randomly shaped blobs in close packing. Observe the figure-ground dependence on contrast sign (see text).

interstices). Therefore, visual processes that take as input a geometrical description of image texture (such as processes that lead to texture segmentation and recognition) depend on first achieving the correct figure-ground parsing. The visual system has apparently developed robust strategies for deriving texture descriptions which reflect the order, regularity and structure of the objects comprising the physical texture, and which ignore the randomness of the interstices. Note that two explanations can be forwarded: either the visual system derives two texture descriptions in each locality (corresponding to the two texture parsings) and selects that which reflects the structure and recognizable shapes of the physical texture, or it determines the correct parsing in a primarily "bottom-up" manner.

The consequences of figure-ground ambiguity in texture is not widely recognized. Two probable reasons are that we are very adept at deciding figure-ground and seldom see reversals of figure-ground in natural scenes, and for synthetic textures in particular, the discrete constituent elements (the bars, dots, etc.) are usually sufficiently separated that the figure-ground distinction is again unambiguous and stable. It is in consideration of natural images that the significance of the computational problem becomes apparent.

2.2 Computing Convexity in Texture

The computational problems associated with figure-ground raises substantial questions regarding how convexity might actually be measured. Convexity is a global property of a closed curve, which has several mathematical definitions, e.g. *i)* a straight line connecting any two points on the curve lies entirely within the curve, *ii)* the curve has no inflection points, and equivalently, *iii)* the curvature has everywhere the same sign. Note that if (i) is rephrased that a straight line connecting any two points on the curve does not intersect the curve at any intermediate point, then all three definitions can be applied to open as well as closed curves. Each of the above definitions might suggest a variety of algorithms for determining convexity.

Consider a smooth arc of curve without inflection points. Assuming it corresponds to the physical edge of a convex object, it immediately follows that the convex side of the curve is figure. The evidence provided by curvature sign alone is very local, as if the curve were examined through a small aperature. It is akin to deciding figure-ground on the basis of contrast sign, where presumably the lighter side of an edge is more likely to correspond to a physical object (and the darker to be shadow or background). Just as one can demonstrate the tendency to interpret lighter as figure one can demonstrate the tendency to interpret as figure the convex side of a curve (see below). Since this tendency persists in stimuli that are devoid of other relevant information, one may conclude that some measure of curvature sign is a salient property as regards figure-ground. But we will also show that curvature sign is not the only factor underlying the convexity preference.

The preference for convexity likely stems from several causes, all of which are consequences of convex objects imaging as convex silhouettes. In addition to the straightforward notion of curvature sign just discussed, there are localized events that arise where convex silhouettes overlap which also indicate the appropriate figure-ground assignment. Since the physical objects that comprise a texture are usually distributed three-dimensionally in space, their silhouettes often overlap in the image. At each point of overlap the two silhouettes conjoin to produce a sharply discontinuous concave cusp. The figure-ground interpretation of a concave cusp is straightforward: the region within the cusp is ground, and the two component arcs correspond to two physically distinct objects. The cusp point itself has no physical significance, however, it is merely the point where they overlap from the given perspective.

The figure-ground stimuli that Kaniza and Gerbino [1976] used to show the preference for convexity have, in addition to smoothly convex curve arcs that define and enclose convex shapes cusp-like discontinuities in the curves. We suggest that these sharp cusps contribute strongly to figure-ground. The observed convexity preference is due not only to curvature sign along smooth arcs of figures, but to the geometry of the sharp discontinuities where the smooth arcs conjoin.

2.3 Figure-ground Determination and Part Boundaries

The geometry of the concave cusp suggests the figure-ground relationship across each of the two arcs. It also demarks a point where two distinct silhouettes intersect, the physical interpretation of which is that distinct physical parts, either separated in space or abutting, project so that their silhouette contours intersect. Not only are the two figures distinguished from their common background, but from each other. This latter role of the concave cusp, as *a priori* evidence for a part boundary, has been recognized by other researchers, but for a distinctly different physical interpretation. Specifically, we are concerned with the cusp as evidence that two convex objects partly overlap or abut; the other work has concerned parts that interpenetrate or join to form a common object and moreover the parts need not be convex.

The silhouette of an object composed of distinct parts generally carries information about where the parts conjoin. Deep concavities in the silhouette outline are good candidate part boundaries, specifically at the point where the curvature is most negative at the base of the concavity [Marr & Nishihara 1978]. Hoffman and Richards [1984] similarly observe that points of minimum curvature along a silhouette curve often correspond to part boundaries, citing the fact that for two arbitrarily shaped surfaces that interpenetrate, the locus of intersection is a contour of concave discontinuity of their tangent planes. Consequently they propose parsing surfaces in 3-D along loci of negative minima of each principal curvature, and in 2-D, at points of negative minima of curvature [Hoffman & Richards 1982, 1984]. Their treatment of 2-D silhouette contours is derived from the 3-D case. Their proposal results in apparently psychologically valid predictions both for the apparent parts of surfaces in 3-D and of curves in 2-D. They define curvature sign relative to the side of the curve that is regarded as figure (a protrusion in the silhouette is associated with positive curvature by their convention, and an indentation with negative curvature). If the figure-ground sense across the contour reverses, the curvature sign also reverses (so that minima of curvature become maxima of curvature and vice versa). Since what is regarded as a minimum of curvature depends on the figure-ground interpretation, they predict that the subjective parts of a curve will appear delimited by minima of curvature as dependent on the figure-ground sense, a prediction that is upheld in e.g. the vase-face illusion.

The interpretation of a silhouette curve thus depends on the figure-ground and assumes that figure has been distinguished from the background. We propose the following extension: that certain types of cusp are early *a priori* evidence for which side is figure, as a precursor to subsequent analysis of the silhouette's shape.

2.4 Figure-ground Interpretation of Cusps of Different Type

Not all cusp discontinuities in tangent along a curve can be interpreted equally as evidence for figure-ground. Even the concave cusp as we define it has an alternative interpretation, that being of a sharp concave physical object such as a thorn (see figure 3a) rather than a gap between two convex objects. That either interpretation might be valid cannot be ignored, but in the absence of other figure-ground evidence, a concave cusp more likely corresponds to two overlapping or abutting convex objects than to a single sharp concave spike or thorn. By this argument, the observed bias or preference for convexity reflects a statistical fact about our visual world.

The concave cusp, formed by two convex silhouettes, is but one of six types of cusp, the geometry of each depending on the combination of curvatures of the two intersecting silhouette curves at their point of intersection. Since each arc may have either positive, negative, or negligible (zero) curvature, the concave cusp is termed negative/negative and the other cusps are: negative/zero, negative/positive, positive/zero, positive/positive, and zero/zero (see figure 3). These five cusp types are weaker figure-ground evidence than the negative/negative, and can be rank-ordered roughly. The negative/zero cusp is somewhat weaker, as only one convex object would be involved, and the straight edge providing no additional information. Still weaker evidence would be the negative/positive cusp which, by the above interpretation would correspond to the intersection of a concave and a convex silhouette. The more probable interpretation of the negative/positive cusp seemingly would be a sharply pointed figure such as a leaf tip. Similarly,

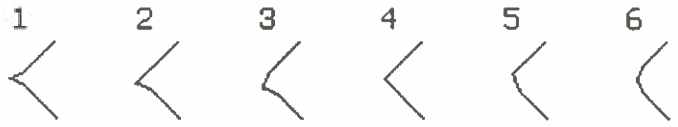

Figure 3. Six types of cusps: *1*) negative/negative curvature, *2*) negative/zero curvature, *3*) negative/positive curvature, *4*) zero/zero curvature, *5*) zero/positive curvature, *6*) positive/positive curvature.

---

the positive/positive and positive/zero cusps are less unlikely to be part boundaries than sharp convexities in the silhouette of a single figure. Finally, the zero/zero discontinuity, a simple corner defined by two straight edges, is clearly the weakest evidence.

In the following we show that the known figure-ground preference for convexity (e.g.

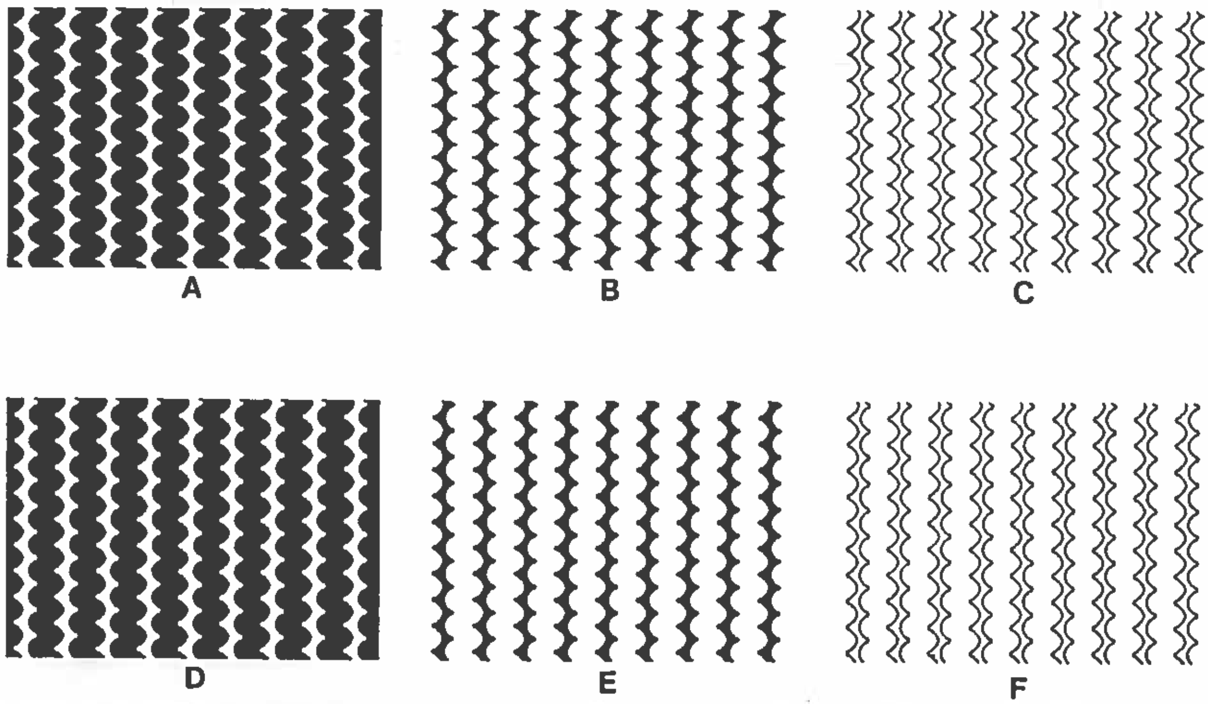

Figure 4. Asymmetric convex (sinuous) shapes vs. asymmetric concave shapes. In *a,b,c* there is preference for the sinusous shapes due to the sharp cusps. In *d,e,f* there is still a preference for sinuous shapes but the alternative parsing can also be seen, especially in *e*.

reported by Kaniza and Gerbino [1976]) derives in part from the presence of concave (negative/negative) cusps. Since the concave cusp is comprised of two convex arcs, and we recognize that the convexity of an individual arc induces a figure-ground preference in and of itself, we must, in the process, show that the particular geometric arrangement of cusp, and not merely the curvature of the arcs, is effective.

3. Demonstrations

Figure 4a presents a texture with two distinct figure-ground parsings. When white is regarded as figure one sees sinuous shapes resembling coiled telephone cords; when black is figure one sees concave shapes such as thorns. There is a preference to see the sinuous shapes, wherein the concave cusps demark convex segments along the cords. The contrast is reversed in figure 4b, and as expected, the concave thorns, now white, may be seen as figure more readily. The competing contribution of contrast sign is removed in the line-drawn version in figure 4c and the sinuous shapes are again dominant. Neither figure-ground interpretation is absolute in figure 4a-c, however. Spontaneous reversals are frequent, nonetheless the initial impression is generally to see the sinuous, convex shapes, and this interpretation is held the greater fraction of the time. In figure 4d-f the pattern has been subtly modified to remove the sharp concave discontinuities; the figure-ground interpretation is now more ambiguous. Compare the line drawings in figures 3c and 4f to observe the effect of the sharp cusps. We suggest that the relatively greater stability of the sinuous (coiled telephone cord) interpretation in figure 4c over figure 4f to the presence of sharp the concave cusps in figure 3c and their absence in figure 4f — the change in terms of total curvature being negligible, and other factors remaining constant.

Convexity and symmetry can be placed in opposition to further increase the figure-ground ambiguity. In figure 5a, for example, one may see white sinuous cord-like figures, as before, where the convexity dominates. But observe that the black background shapes are symmetric, and if seen as figure resemble strings of concave beads. The concave but symmetric figures are more readily seen when the contrast is reversed as in figure 5b; the line-drawn version is shown in figure 5c. In figure 6 both figure-ground interpretations are symmetric, and the tendency is to see convex beads rather than the alternative concave beads. The convex figure interpretation is strong even in figure 6d where the spacing favors the alternative concave-figure interpretation. There is still a substantial tendency to see strings of convex beads despite the wide separation of the convex arcs that define their silhouettes. In figure 6e, where the sharp concave cusps are slightly rounded, the alternative interpretation of widely separated strings of concave beads is more readily achieved.

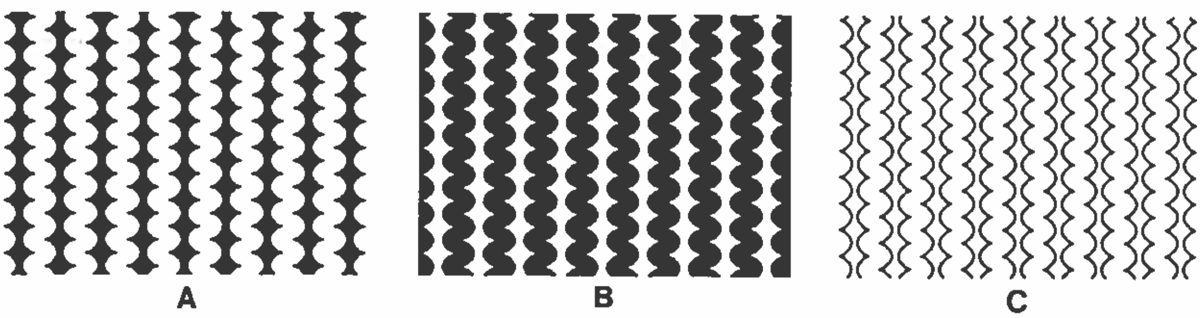

Figure 5. See text.

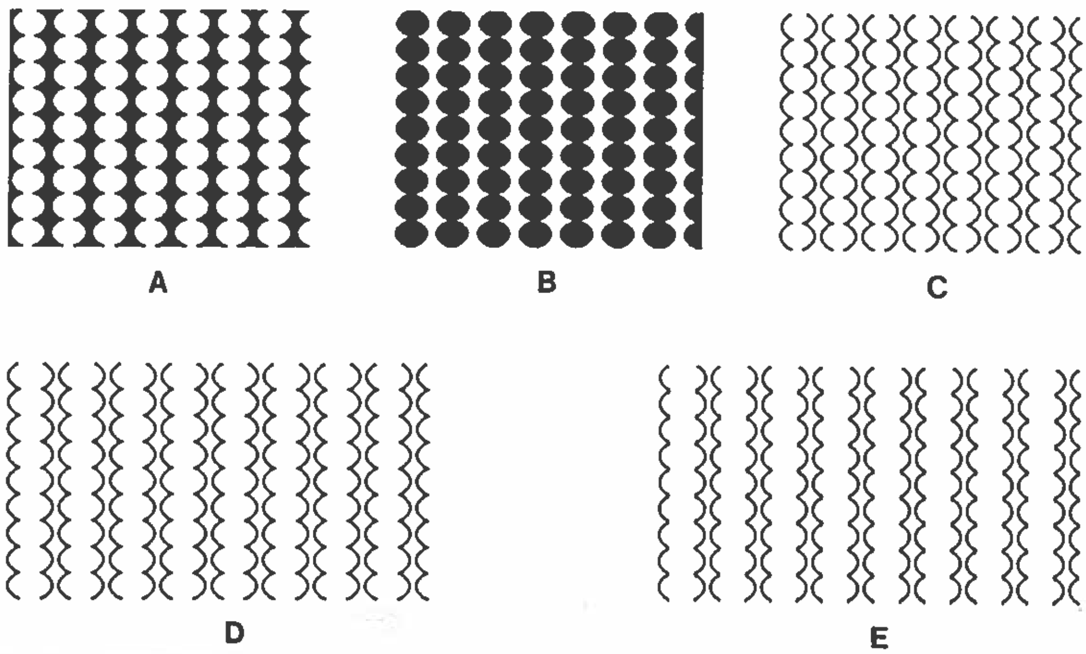

Figure 6. See text.

While an effect due to the sharp concave cusp is evident in these demonstrations, the figure-ground interpretations are influenced by many uncontrolled factors, particularly since the resulting shapes are two-dimensional. To further simplify the stimulus, therefore, we next turn to one-dimensional patterns (see figures 7 and 8). These patterns, being line-drawings, remove the influence of edge contrast on figure-ground, but are interpreted in terms of occluding edges nonetheless. Moreover, rather than defining competing two-dimensional shapes, the patterns of curves merely suggest a one-dimensional arrangement of overlapping edges. The local figure-ground problem is thus reduced to determining whether the curve corresponds to a left or right edge.

Consider first the symmetrical pattern in figure 7a, which can be seen as serrated edges (sawteeth) overlapping either to the left or right with equal ease. In figure 7b a very slight concave cusp is introduced by appending a minute line segment on the left-hand vertices, with the result that one sees an arrangement of occluding edges, each partly overlapping the surface below it to the right and in turn occluded by the nearer surface to the left. (This will be termed "overlapping to the right". The left side of each curve is seen as figure against the background

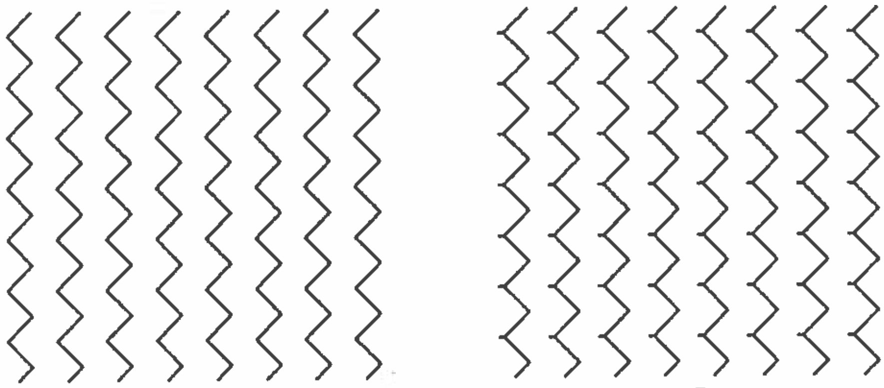

Figure 7. The symmetrical pattern in *a* can be seen as a series of serrated edges that successively overlap. There is roughly equal preference for overlap to the left as to the right. The slight cusp introduced in *b* causes strong preference for edges to overlap to the right.

immediately to its right.[2]

The geometric arrangement in figure 7b, which influences apparent figure-ground so strongly, is composed of straight line segments instead of continuously curved arcs (as in figure 4a). This arrangement serves us in two ways. First, it permits patterns to present distinct and effective cusp-like features while controlling for contour curvature. Second, they provide insight into the defining geometry of the cusp feature, as will be discussed later.

It was mentioned that since a concave cusp is formed by the intersection of two convex (positive curvature) arcs, we need to show that contour curvature alone is not determining figure-ground. This was demonstrated, in part, by figures 4c versus 4f where blunting the concave cusp adds little to the convexity but extinguishes the discontinuity, and hence, we argue, the concave cusp feature. Also consider figure 8. In figure 8b one prefers overlap to the left on the basis of convexity; note that only a slight amount of curvature is needed to produce this bias (compare with the symmetrical sawtooth pattern in figure 8a). In figure 8c the overlap is at least as strongly biased towards the left by the straight-line approximations to concave cusps. These cusp shapes induce a strong impression of figure-ground in the absence of smooth convexity. The following reports on experiments that further explores these effects.

4. Experiments

Two experiments were performed. The first involved subjects judging the direction of overlap, as in figures 7 and 8 above. From the direction of overlap one can infer which side of the curve was regarded as figure. It can also be determined by presenting a probe dot on one side or the other of a given curve within the stimulus figure and, since the curve is interpreted as an occluding edge, have the subject respond whether or not the dot was on the edge (figure) or on the

[2] Note that if the patterns were oriented horizontally the judgments of overlap direction would be confounded by the independent tendency to interpret depth as increasing as one scans from bottom to top. The apparent overlap is biased towards interpreting the each edge as occluding that which lies above it (rotate figure 7 so that the sawtooth curves are horizontal).

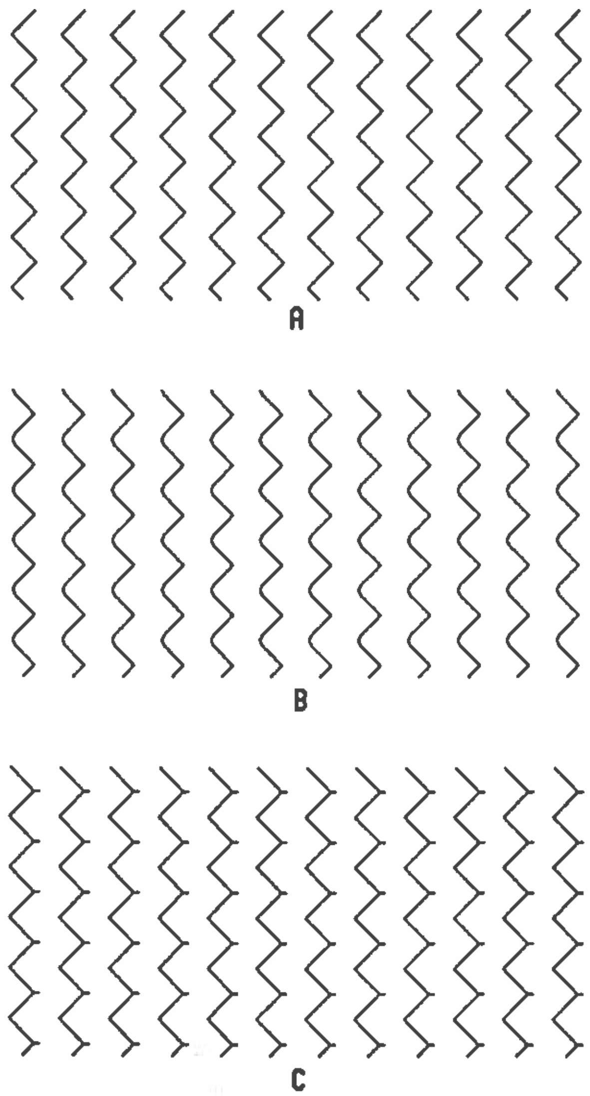

Figure 8. The symmetrical pattern in *a* is ambiguous. The shallow convexities in *b* induce a preference for overlap to the left. The primitive concave cusps in *c* have similar effect, and are more effective in short presentations.

background. The second experiment employed this task.

2.1 Experiment 1: Direction of Overlap

Method

Stimuli: The stimuli consisted of regularly-spaced sawtooth curves as shown in figure 8, which suggested a series of serrated edges overlapping either to the left or the right. The stimulus configurations (figure 9) consisted of ten configurations: three straight-line approximations to concave cusps (shapes 1, 2, and 3), three concave cusps defined by continous curves (shapes 4, 5, and 6), three smoothly convex corners (shapes 7, 8, and 9) and a symmetrical sawtooth (shape 10). Note that the basic sawtooth curve consists of either (i) a cusp feature on the left and a sharp corner on the right, or (ii) a sharp corner on the left and a smoothly rounded corner on the other. The stimuli were displayed as white lines of intensity 5.5 ft.-L. on a grey background of 1.2 ft.-L., in a darkened room. The stimuli were generated by a Symbolics 3670 Lisp Machine and displayed on a Tektronix 690SR color monitor.

The various sawtooth curves were composed of bitmaps. When viewed from 76 inches, one pixel subtended 1′. The .4 mm dot pitch of the CRT permitted individual pixels to be resolved. The smallest continuous cusp (figure 9, shape 4) blended into the straight segment of the sawtooth over an extent of 4 pixels (4′), and the overall amplitude of the sawtooth was roughly 20′. Note that the smallest cusp differed from a sharp corner by the addition of a single pixel (subtending 1′). The difference was just visible from the subject's viewing distance. All configurations (except shape 10 in figure 9) suggest overlap to the right. When a sawtooth stimulus was presented on the display, either the bitmap or its mirror reflection (about the vertical) was projected in order to balance for direction of overlap.

Procedure: The subject's task was to report, by pressing buttons on a mouse (an interactive x-y pointing device), which direction the surfaces overlapped. The task started with a one-second presentation of a fixation point, a cross subtending 8x8′. The fixation point appeared against a blank background at the location where a sawtooth curve would momentarily appear. The fixation point was placed along a given sawtooth curve at a point midway between two extrema. The sawtooth pattern was then presented for 500 msec during a training run and subsequently at 250 msec during runs for which data were collected. After the stimulus interval or the subject's response, which ever occurred first, the stimulus was replaced by a random dot masking pattern of moderate density chosen to roughly match the mean luminance of the stimulus display. The

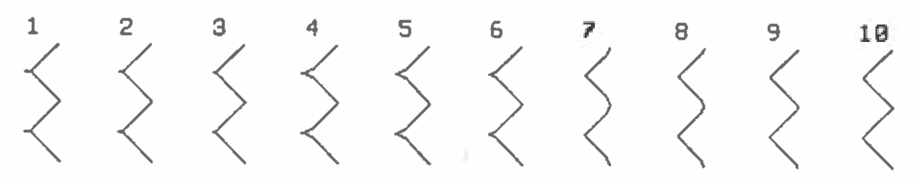

Figure 9. Cusp stimuli for the experiments. Shapes 1-3 are straight-line approximations to cusps. Shapes 4-6 are the continuously curved cusps. Shapes 7-9 are the smooth corners and shape 10 the symmetrical sawtooth.

subject was instructed to make rapid responses of overlap direction while maintaining accuracy. Overlap direction was explained with reference to the edges of the pages of an open book, e.g. the pages on the left "overlap to the left" and if the subject sees this direction of overlap in the stimulus the left mouse button should be depressed. The response time was measured relative to stimulus onset. The experiment consisted a sequence of 100 trials for which overlap direction and reaction times were recorded. The sequence consisted of five repetitions of randomized presentations of the ten stimuli and their mirror reflections. Prior to collecting the data, subjects were given a practice run of 20 trials. The subjects consisted of two females and five males with normal or corrected vision; all were unpaid volunteer subjects.

Results

Table 1 shows the mean and individual error counts for each stimulus shape across subjects. Data was collapsed across mirror reflections, and the error counts for each of the 10 stimulus shapes were computed. A judgment was regarded as an "error" if the subjective direction of overlap did not match that predicted by the shape. Note that for shape 10, the symmetrical sawtooth, there is no correct or incorrect judgment, but that data was included to reveal any subject bias towards interpreting overlap to either the left or right. The mean error would have been 5.0 if no bias existed; the observed mean of 6.2 across subjects shows a bias towards judging overlap towards the right.

Table 1. Error counts by subject.

Shape	A	B	C	D	E	Mean
1	0	1	1	0	4	1.2
2	0	2	2	0	4	1.6
3	1	3	4	0	4	2.4
4	0	1	0	0	3	0.8
5	0	1	3	0	2	1.2
6	0	0	0	0	5	1.0
7	0	1	2	0	1	0.8
8	0	2	4	0	0	1.2
9	4	2	3	3	4	3.2
10	7	6	6	7	5	6.2

Comparing the results across shape type, the least mean error rate occurred with the continuously-curved cusps (shapes 4-6) in general. Comparing mean error rates within shape or across shapes few differences reached significance. All were significantly different from shape 10 (at the .05 level). Within shape type there was a weakly significant trend for the error rates to increase with decreasing size of the feature. Regarding the potential for left-right bias, note that subjects A and D had the strongest bias in shape 10 but the fewest errors, and subject E, with the most errors, had no bias in shape 10. The mean reaction times revealed no significant differences.

Most subjects were able make the judgments both quickly and accurately with little practice. With practice, we found that very accurate judgements were possible with very short presentation times. For example, one of the authors (AB) produced the results shown in table 2 with 80 msec presentation time. Note that shape 10 was not incorporated in this experiment.

Table 2: Subject AB (80 msec)

Shape	1	2	3	4	5	6	7	8	9
Error counts	0	1	0	0	0	0	0	0	1
Reaction times	428	558	589	438	457	475	526	481	722

Discussion

The first point to draw from this experiment is that all shapes other than the simple corner (shape 10) were roughly equivalent in defining figure-ground. That is, figure-ground can be determined on the basis of convexity in the absence of a cusp (shapes 7-9), or independently, on the basis of a cusp in the absence of convexity (shapes 1-3). Moreover, the straight-line approximations to concave cusps (shapes 1-3) were similar to the continuously-curved concave cusps (shapes 4-6) despite their visually obvious differences. With regard to the size of the cusp feature for a given shape, the error means varied little between the smallest to largest (e.g. between shapes 1 and 3 or between 4 and 6). The smallest feature was shape 3, which differed from a sharp corner by the addition of a single pixel, was quite effective in defining figure-ground despite it being barely discriminable from a sharp corner (shape 10) at the subject's viewing distance.

The results for subject E are significantly different from those of the other subjects. The task of seeing the contours as edges of overlapping surfaces was difficult for some subjects and required a greater number of learning trials. A few candidate subjects reported that the patterns merely looked like flat, two-dimensional arrangements of lines whose corners and angles, like arrowheads, pointed to either the left or right. For these subjects the impression of the sawtooth curve as a serrated edge was not natural. Consequently we decided to perform an experiment in which apparent figure-ground was probed not by direction of overlap, but more directly by asking whether a dot presented on one or the other side of an indicated curve was on the edge (in the figure suggested by the curve) or to the side, that is, on the background adjacent to the edge. This second task still required subjects to hold an interpretation of the sawtooth curve as the edge of a serrated surface. The subjects' instruction, however, encouraged a more local judgment of the figure-ground relationship of the probe dot relative to the immediately adjacent curve. In reflection on the first experiment, we suspect that although the judgement could be made locally on the basis of whether the indicated curve was a left or right edge, the instructions encouraged a more diffuse, global judgement of whether the pattern represented a cascading series of overlapping edges, and then to report the direction of overlap. a task that some found difficult. The following experiment, it turns out, produced much more uniform results and required fewer introductory trials.

4.2 Experiment 2: Judging Figure-ground

Method

Stimuli: The stimuli consisted of patterns of sawtooth curves as in Experiment 1. In addition, each stimulus was presented with an additional dot placed on one or the other side of the sawtooth curve indicated by the fixation point. The dot would be judged as either on or off the given edge, as an independent means of probing apparent figure-ground. The symmetric sawtooth (shape 10) was eliminated in this experiment.

Procedure: Subjects were instructed, as before, to direct attention to the fixation point, then to the relatioship between the indicated sawtooth curve and the dot (which appeared randomly to the left or right of the curve). As rapidly as possible while preserving accuracy the subject was then to indicate whether the dot appeared to be on the given edge or not, by depressing a mouse button. The sequence consisted of three repetitions of randomized presentations of the nine stimuli with each of the permutations of overlap and dot position. Prior to data collection, subjects were given a practice run of one sequence of 36 trials; data was then collected for a sequence of 108 trials (3 sequences of 36). Again the viewing distance in all cases was 76 inches, so that one pixel subtended 1′. Four female and five male graduate students, all unpaid volunteers, participated as subjects.

Results

Table 3 shows the mean and individual error counts for each stimulus shape across subjects. Again data was collapsed across mirror reflections, and the error counts for each stimulus shape was computed.

Table 3. Error counts by subject

Shape	A	B	C	D	E	F	G	H	I	Mean
1	0	0	1	1	2	1	1	1	2	1.0
2	1	0	1	1	0	2	2	0	2	1.00
3	0	0	0	4	1	2	1	0	0	0.89
4	0	0	0	0	1	1	1	2	0	0.56
5	0	0	0	1	0	1	1	1	0	0.44
6	1	0	2	2	1	1	0	0	1	0.89
7	0	1	0	1	0	4	0	0	0	0.67
8	0	0	0	0	0	2	1	0	0	0.33
9	0	2	3	6	0	2	1	2	2	2.00

Pooling across type (i.e. shapes 1-3, 4-6, and 7-9) revealed no significant differences. Within type, the straight-line approximated cusps (shapes 1-3) were insignificantly different, as were the continuously-curved cusps (shapes 4-6). We found the convex shape 9 (having the smallest radius of curvature) to be different from the other convex shapes (shape 8, at the .05 level and shape 7 between the .1 and .05 level).

Discussion

Generally one is struck more by the similarity of the results across type and within type than with the differences. Also, in light of the very low error rates for all conditions tested (where the worst mean of 2 out of 12 occurred for shape 9) shows that the visual system very efficiently uses minute evidence along a curve in determining figure-ground. The only differences concerned shape 9, the smallest of the rounded corners. Seemingly, the introduction of even a slight amount of curvature to the sawtooth pattern biases apparent figure-ground. With reference to figure 9, note that the curvature is induced by only 2 to 3 pixels of "rounding" which subtends only 2-3′. Likewise, the introduction of a minute cusp-like feature is effective at nearly the limit of visual resolution, such as the single pixel in shape 3, which amounts to a 1′ tip added to convert the sharp corner to a minute crack-like cusp subtends a similarly small locality.

5 General Discussion

Starting from earlier observations that convexity, as a general property, influences apparent figure-ground, we have examined how convexity might be measured. Earlier we pointed out (section 2.3) that convexity has many geometric definitions in principle, each suggesting visual algorithms of varying tractablility and utility. We wish to emphasize here that global measures of convexity, particularly those that require an isolated and closed contour are probably inappropriate, and that instead local evidence for figure-ground (still based on convexity) would be preferable.

We proposed that the concave cusp discontinuity would be a useful feature of a curve on which to base an early decision of figure-ground. The experiments we performed confirmed the proposal, to the extent that our stimuli successfully isolate the concave cusp from the already-expected contribution of curvature convexity. We designed shapes 4-6 as representative examples of concave cusps, shapes 7-9 as representative examples of convexity without cusps, and shapes 1-3 as examples of cusps without convexity. We designed a straight-line approximation to a cusp, to see if the particular local geometry, and not the curvature of individual curved arcs, may induce the figure-ground decision. We found that all three shape types were roughly equally effective, from which we conclude that while convexity alone is sufficient (shapes 7-9), as expected, so is the distinctive geometry of the concave cusp, even when one contrives to minimize the contribution due to contour curvature (shapes 1-3 as opposed to shapes 4-6). Recall the arguments to this effect in the demonstrations involving figures 7.

The concave cusp seemingly behaves as a geometrically-defined feature, and while we do not offer a concise definition, the apparent equivalence of the cusps composed of straight-line segments and those composed of smooth arcs suggests the definition is rather primitive — the component arcs need not curve continuously in the vicinity of the cusp for the arrangement to be effective (refer back to figures 7b and 8c as well). Earlier we noted that the concave (negative-negative) cusp is but one of six arrangements of arcs at a discontinuity of tangent along a curve, and predicted that it should be the most effective. Variations such as those suggested in figure 4 should permit a more specific description of what geometric aspects constitute this figure-ground determiner.

REFERENCES

Frisch, H.L, & Julesz, B. 1966 Figure-ground perception and random geometry. *Perception and Psychophysics* **1**, 389-398.

Hochberg, J. 1964 *Perception*. New Jersey: Prentice-Hall.

Hoffman, D.D. 1983 The interpretation of visual illusions. *Scientific American* **249**, December.

Hoffman, D.D. & Richards, W. 1982 Representing smooth plane curves for recognition: Implementations for figure-ground reversal. *Proceedings of the National Conference of the American Association for Artificial Intelligence*, 5-8.

Hoffman, D.D. & Richards, W. 1984 Parts of recognition. *Cognition* **18**, 65-96.

Julesz, B. 1965 Texture and visual perception. *Scientific American* **212**, 38-48.

Julesz, B. 1981 Figure and ground perception in briefly presented isodipole textures. In *Perceptual organization*, Kubovy, M. & Pomerantz, J.R., eds. Hillsdale, New Yersey: Lawrence Erlbaum Associates, 27-54.

Kaniza, G. 1979 *Organization in Vision*. New York: Praeger.

Kaniza, G. & Gerbino, W. 1976 Convexity and symmetry in figure-ground organization. In *Art and Artefacts*, M. Henle, ed. 25-32. New York: Springer.

Marr, D. 1976 Early processing of visual information. *Phil. Trans. R. Soc. Lond. B* **275**, 483-524.

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

Marr, D. & Hildreth, E. 1980 Theory of edge detection. *Proc. Roy. Soc. Lond. B.* **207**, 187-217.

Marr, D., Poggio, T. & Hildreth, E. 1980 Smallest channel in early human vision. *J. Opt. Soc. Am.* **70**, 868-870.

Richards, W. & Purks, S.R. 1978 On random-dot texture discrimination. *J. Opt. Soc. AM.* **68**, 270-271.

Rubin, E. 1958 Figure and Ground. In *Readings in perception*, D.C. Beardslee and M. Wertheimer, eds. Princeton, N.J.: Van Nostrand.

Schiller, P.H., Findlay, B.L. & Volman, S.F. Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *J. Neurophysiol.* **39**, 1288-1319.

Ullman, S. 1984 Visual Routines. *Cognition* **18**, 97-159.

Waltz, D.I. 1975 Generating semantic descriptions from drawings of scenes with shadows. Ph.D. dissertation, AI Lab, MIT, 1972.

Watt, R.J. & Morgan, M.J. 1983 The recognition and representation of edge blur: evidence for spatial primitives in human vision. *Vision Research* **12** 1465-1477.

Wilson, H.R. & Bergen, J.R. 1970 A four mechanism model for spatial vision. *Vision Research* **19**, 19-32.