

Chapter 9 in “*From Fragments to Objects: Grouping and Segmentation in Vision*,” T. F. Shipley & P. J. Kellman (Eds.) Elsevier Science, North Holland Press (2001; pp. 401–459).

9

PART-BASED REPRESENTATIONS OF VISUAL SHAPE AND IMPLICATIONS FOR VISUAL COGNITION

*Manish Singh, Department of Psychology and Center for Cognitive Science, Rutgers University,
New Brunswick*

Donald D. Hoffman, Department of Cognitive Sciences, University of California, Irvine

ABSTRACT

Human vision organizes object shapes in terms of parts and their spatial relationships. Converging experimental evidence suggests that parts are computed rapidly and early in visual processing. We review theories of how human vision parses shapes. In particular, we discuss the minima rule for finding part boundaries on shapes, geometric factors for creating part cuts, and a theory of part salience. We review empirical evidence that human vision parses shapes into parts, and show that parts-based representations explain various aspects of our visual cognition, including figure-ground assignment, judgments of shape similarity, memory for shapes, visual search for shapes, the perception of transparency, and the allocation of visual attention to objects.

INTRODUCTION

Our visual world is replete with structure, in the form of objects and surfaces. The inputs to our visual systems, however, contain no objects or surfaces—only discrete arrays of intensity values that can vary dramatically with small changes in viewing conditions. The objects and surfaces we see thus need to be constructed by our visual systems from these inputs. A critical aspect of this construction is parsing the input arrays into regions that are likely to have arisen from distinct objects; in other words, the formation of ‘perceptual units.’

Entire objects, however, are not the only perceptual units. Indeed, the objects we see are not unstructured wholes—they themselves have further part structure. We perceive a table, for example, as a coherent perceptual object; but also as a spatial arrangement of clearly defined parts: four legs and a top. Hence, ‘perceptual units’ exist at many levels: at the level of whole objects, at the level of parts, and possibly smaller parts nested within larger ones (Palmer, 1977; Marr & Nishihara, 1978). Moreover, as we will see, mechanisms of part segmentation interact with mechanisms of object and surface segmentation in a number of ways.

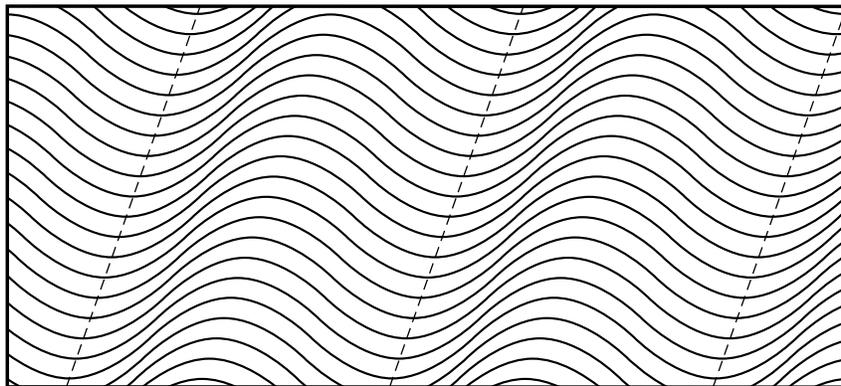


Figure 1. *This display is seen as hills separated by dashed lines. If you turn the page upside down, convex and concave reverse—leading to a switch in perceived part boundaries. As a result, the dashed lines now appear to lie on top of hills, rather than along the boundaries between hills.*

Many theorists believe that, because human object recognition is often sensitive to viewpoint, shape descriptions employed by human vision are not part-based. However, this belief is based on the assumption that a part-based approach is synonymous with (or, at least, commits one to) a volumetric, object-centered, approach to shape representation in which a set of basic 3D parts is specified in advance (e.g., Marr & Nishihara, 1978; Biederman, 1987). This assumption is false. Parts are equally consistent with, and just as useful for, 2D (image-based)

and 2- $\frac{1}{2}$ D (surface-based) representations as they are for 3D (volumetric) representations. And using parts does not entail that one must commit to a fixed repertoire of primitive shapes.

In this chapter we argue, both on computational and empirical grounds, that the human representation of visual shape is in fact part-based. Human vision parses shapes into component parts, and it organizes them using these parts and their spatial relationships. This places strong constraints on theories of human visual shape description.

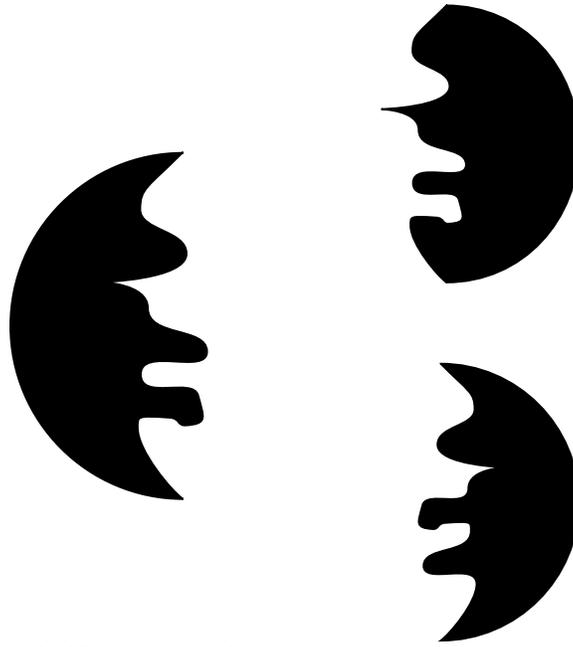


Figure 2. Which of the two half moons on the right is more similar to the one on the left? Most observers pick the bottom one, even though it is the one on top, not bottom, whose wiggly contour is identical to the one on the left.

From a computational perspective, parts are useful for many reasons. First, many objects are articulated: they consist of parts that move non-rigidly with respect to each other. Fingers, for instance, can move independently. A template for an outstretched hand correlates poorly with a template for a pointing hand. A part-based description allows one to decouple the shapes of the parts from the spatial relationships between the parts—hence providing a natural way to represent and recognize articulated objects. Second, one never sees all of an object in one view: the back of an object is not visible due to self-occlusion, and its front may be occluded by other objects. Representing shapes by parts allows the recognition process to proceed with those parts that are visible.

From an empirical perspective, recent evidence suggests that parts are computed quickly, automatically, and in parallel over the visual field (Baylis & Driver, 1994; 95a; 95b; Hulleman, te Winkel, & Boselie, 2000; Wolfe & Bennett, 1997, Xu & Singh, 2001). With ambiguous shapes flashed for just 50 ms and followed by a mask, parts affect the perception of figure and ground

4 From Fragments to Objects: Segmentation and Grouping in Vision

(Hoffman & Singh, 1997). Hence, whatever further transformations of shape into other representational formats may take place later in visual processing, these are likely to be

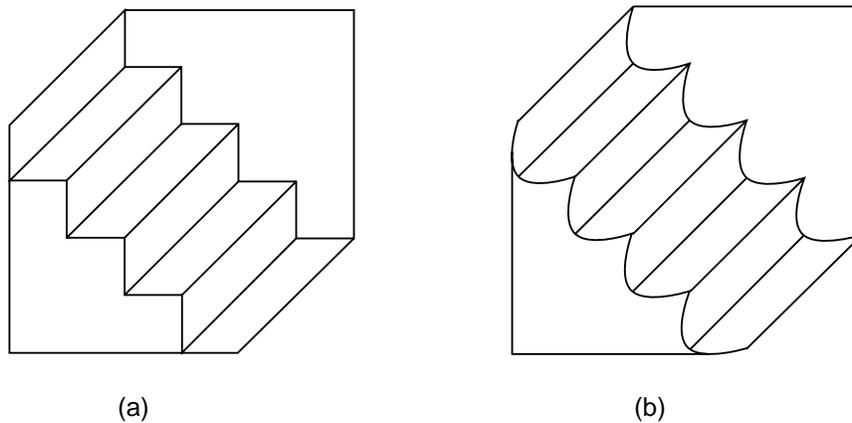


Figure 3. The staircase on the left can be seen as either right-side up, or as upside down. The staircase on the right, however, is more likely to be seen in the upside-down interpretation.

influenced by the early formatting of shapes into parts. Indeed, we show that parts explain a remarkable variety of visual phenomena, including the following.

1. Phenomenology. In Figure 1, we see hill-shaped parts with dashed lines in the valleys between them. Turn the figure upside down, and new hills appear; the dashed lines lie on top of the new hills.

2. Similarity. Of the two half moons on the right of Figure 2, the bottom looks more similar to the half moon on the left (Attneave, 1971; Hoffman, 1983)—even though the top, not the bottom, has the same wiggle as the one on the left.

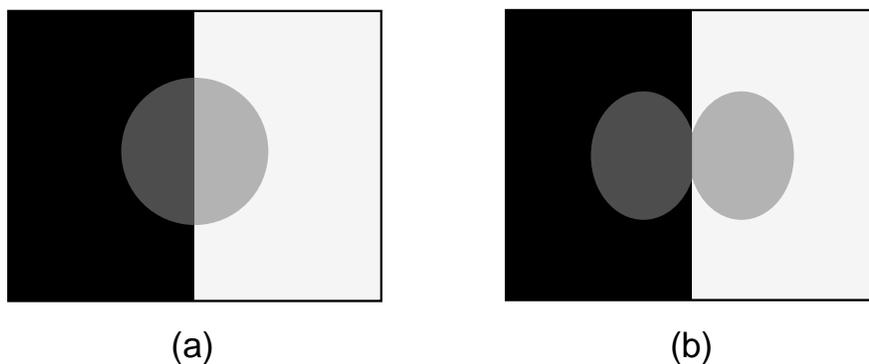


Figure 4. The display on the left is readily seen as a partially transmissive, i.e., transparent, disk over a bipartite background. The display on the right is more difficult to see as containing a transparent overlay.

3. Figure-ground. In Figure 3, the staircase on the right looks upside down, whereas the one on the left can be seen either as right side up or as upside down (Hoffman & Singh, 1997).

4. Transparency. In Figure 4, the display on the left looks transparent, but the one on the right does not (Singh & Hoffman, 1998).

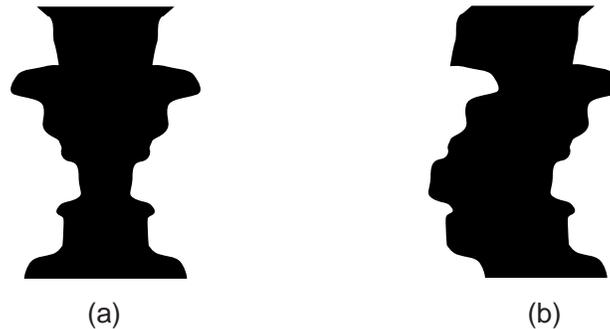


Figure 5. *Observers are faster and more accurate at detecting symmetry within a figure—as in (a)—rather than repetition within a figure—as in (b).*

5. Symmetry. Symmetry (Figure 5a) within an object is easier to detect than repetition (Figure 5b; Baylis & Driver, 1994).

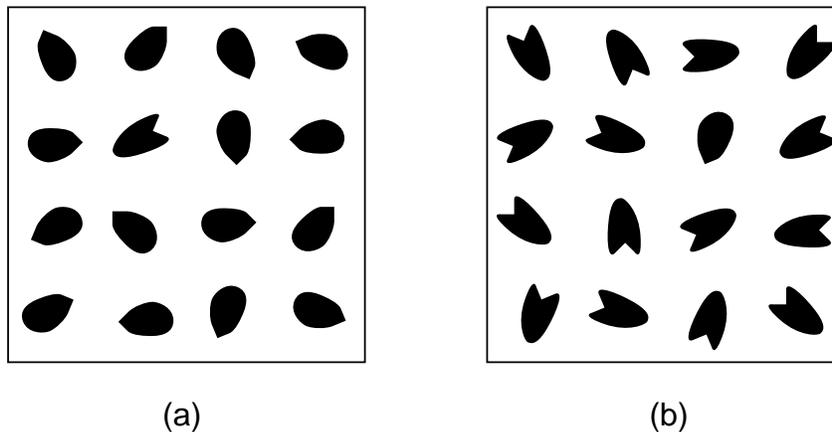


Figure 6. (a) *A shape with a concave corner pops out among distractors with convex corners.*
(b) *However, the reverse visual search is slow and inefficient.*

6. Pop out. A shape with a concave corner pops out among a set of distractors with convex corners (see Figure 6a), but not vice versa (Figure 6b) (Hulleman, te Winkel, & Boselie, 2000).

7. Recognition memory. Some “parts” of shapes are better remembered than others (Braunstein, Hoffman & Saidpour, 1989).

8. Attention. The allocation of visual attention to an object can be part-based, so that attention shifts more readily within a part than across parts of an object (Singh & Scholl, 2000; Vecera, Behrmann, & McGoldrick, 2000; Vecera, Behrmann, & Filapek, 2001; Barenholtz & Feldman, 2001).

We begin by discussing the *minima rule* for finding part boundaries on shapes (Hoffman & Richards, 1984), and review experimental work that demonstrates its psychological reality. However, part boundaries alone do not define parts. So next we consider how human vision uses part boundaries defined by the minima rule, together with other geometric factors, to create cuts on shapes that specify parts. In particular, we discuss the role of cut length (Singh, Seyranian, & Hoffman, 1999), part-boundary strength (Hoffman & Singh, 1997), good continuation, and local symmetry (Singh, Seyranian, & Hoffman, 1996) in determining part cuts. We then review work on part salience, i.e., geometric factors that make some parts more visually salient than others. In discussing the implications of perceptual part structure for visual cognition, we demonstrate that part salience can influence the assignment of perceived figure and ground (Hoffman & Singh, 1997). We show how the minima rule and part salience allow us to refine and extend figural conditions for the perception of transparency proposed by the Gestalt psychologists Metelli and Kanizsa (Singh & Hoffman, 1998). We review results of visual-search experiments that demonstrate that part structures consistent with the minima rule are computed rapidly and early in visual processing (Xu & Singh, 2001). And, we review results of attentional-cueing experiments that reveal that part structure and part salience modulate the allocation of visual attention to multi-part objects (Singh & Scholl, 2000). We end with a discussion of an alternative approach to parts and shape description involving curve evolution and shocks, and relate it to our geometric-constraints approach.

DEFINING BOUNDARIES BETWEEN PARTS

Geometric Considerations

Physically speaking, any subset of an object may be considered to be a part of that object. But human observers clearly find some parts more natural, and others rather contrived. For instance, the legs and the top of a table are clearly seen as distinct natural parts; whereas the upper half of

a leg along with a small adjoining piece of the top seems rather contrived. Similarly, the partitioning shown in Figure 7b looks quite natural, whereas the one in Figure 7c looks contrived. How does the visual system compute such psychologically natural parts?

From a computational perspective, computed parts must satisfy certain principles in order to be useful (Sutherland, 1968; Marr & Nishihara, 1978; Hoffman & Richards, 1984). They must be stable over small generic perturbations of viewpoint, and over small generic changes in object shape. They must be computable from retinal images. And they must be defined for all shapes, and not just a small class of shapes. These principles suggest that to define parts we must use the intrinsic geometry of shapes.

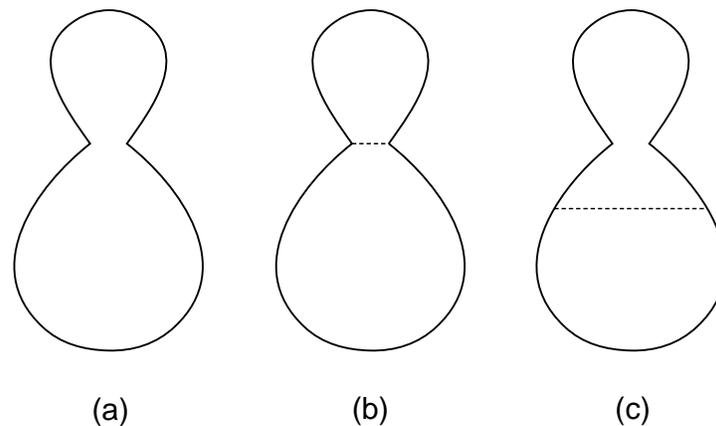


Figure 7. Although any subset of an object is physically a “part” of it, human observers clearly find some parts perceptually natural (b), whereas others seem rather contrived (c).

There are two primary approaches that theorists have taken in defining parts. One theoretical approach defines, in advance, a set of shape primitives that are the possible parts. Proponents of this approach postulate that human vision parses shapes by finding these primitives in images; hence it uses the shape primitives both to find parts and to describe them. Examples of shape primitives that have been proposed include *polyhedra* (Roberts, 1965; Waltz, 1975; Winston, 1975), *superquadrics* (Pentland, 1986), *generalized cones and cylinders* (Binford, 1971; Marr, 1977; Marr & Nishihara, 1978), and *geons* (Biederman, 1987). Although each set of primitives works well on a special class of objects, none can capture the variety of object parts we see—since part shapes that are not in the predefined set of primitives cannot be represented adequately. Nevertheless, if well motivated, shape primitives may be useful as qualitative descriptors of parts for recognizing objects at the basic level, such as *car*, rather than at a superordinate level, such as *vehicle*, or a subordinate level, such as *BMW 325i*. Biederman (1987), for example, uses nonaccidental properties (Binford, 1981; Lowe, 1987; Witkin &

Tenenbaum, 1983), i.e., 3D properties that generically survive projection onto an image plane¹ (like straight versus curved), to motivate his set of geons—because human vision is more sensitive to differences in nonaccidental properties than to metric differences. However, the list of geons only includes a small subset of those shapes that follow from the principle of nonaccidental properties. No geon, for instance, has a rounded tip—even though the difference between a rounded tip and a truncated or pointed tip is nonaccidental. Similarly, no geon can have a cross section that changes its shape or rotates as it sweeps along the axis, or even be at any angle other than 90 degrees to the axis. Indeed, there is as yet no complete, well-motivated, set of shape primitives—and it remains unclear whether there can be.

A different approach to parts, and one that we espouse, is to postulate that human vision uses general computational rules, based on the intrinsic geometry of shapes, to parse visual objects. This approach separates the issue of finding parts from the issue of describing them. For example, Hoffman & Richards (1984) have argued that mechanisms which find parts are more basic, and operate regardless of the shapes of the parts (i.e., independently of whether or not the parts belong to any particular class of shape primitives). They motivate their minima rule, by the principle of transversality, from differential topology.

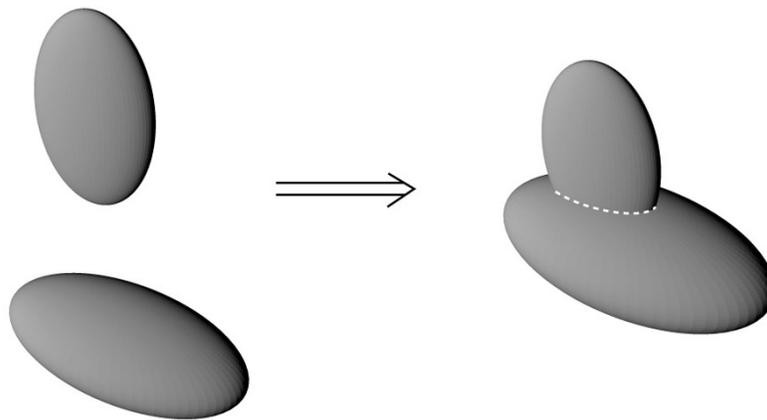


Figure 8. *The principle of transversality: When two 3D shapes intersect, they generically create a concave crease at the locus of intersection.*

¹ Strictly speaking, nonaccidental properties are simply spatiotemporal regularities that are unlikely to arise by chance alone: “In its bare form, the argument says that what *looks* parallel or rigid really *is* parallel or rigid because the spurious appearance of parallelism is extremely unlikely to arise among causally unrelated curves” (Witkin & Tenenbaum, 1983, p. 505). 3D properties that generically survive projection onto an image plane are a special case of this general definition.

Principle of Transversality: Two generic 3D shapes intersect almost surely in a concave crease.²

Figure 8 demonstrates this principle: When the two shapes on the left are joined at random, their intersection creates a concave crease (i.e., a crease that points *into* an object—marked on the right by the dashed contour). If we now view the composite shape on the right as a single object, the two original shapes are natural parts of this object, and the concave crease defines their part boundary. Hence transversality motivates the following rule for parsing 3D shapes.

Rule of Concave Creases: Divide 3D shapes into parts at concave creases.

The Schröder staircase in Figure 9 illustrates the rule of concave creases. When you see the “upright staircase” interpretation, the two dots lie on the same step of the staircase. And this step is bounded on both sides by concave creases. Now if you reverse figure and ground, and see the “inverted staircase” interpretation, the two dots lie on different steps. These steps are nevertheless bounded by the new concave creases.

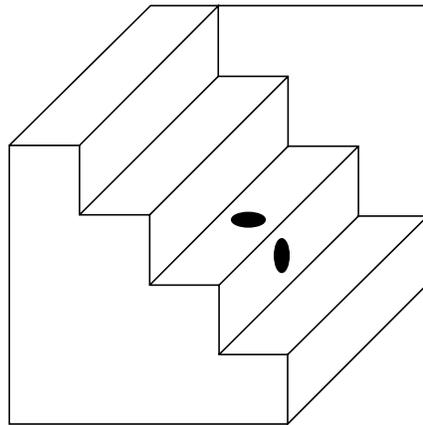


Figure 9. *Are the two dots on the same step, or on different steps? When you see the staircase as upright, the two dots appear to lie on the same step. When you see the staircase as upside down, the dots appear to lie on different steps. Note that, in both interpretations, the boundaries between steps lie along the perceived concave creases.*

As we saw in Figure 1, however, smooth surfaces with no creases can also be seen as having parts. It is easy to generalize the rule of concave creases so that it applies to smooth surfaces as well. Just note that if one smooths a surface in the neighborhood of a concave crease then, intuitively, the concave crease becomes a locus of negative minima of curvature (i.e., a concave region with locally highest magnitude of curvature). This intuition can be made precise

² Something is true “almost surely” if it is true everywhere except (possibly) on a set of measure zero. See,

using the tools of differential geometry (Bennett & Hoffman, 1987). Similarly, if a branch or limb grows out of a main body, its boundary lies along a locus of negative minima (Leyton, 1987). Such considerations lead to the minima rule for 3D shapes.

Minima Rule for 3D shapes: Divide 3D shapes into parts along negative minima of the principal curvatures, along their associated lines of curvature. ³

The minima rule can explain the parts we see in Figure 1. The sinusoidal curves are lines of curvature on the ruled surface. And the dashed lines are the loci of negative minima of the principal curvatures along these lines of curvature. As the minima rule predicts, these dashed lines are the part boundaries we perceive. If you turn the page upside down, then figure and ground reverse (owing to a preference by human vision to see figure below). This turns concave into convex, and hence negative minima into positive maxima, and vice versa. And the new part boundaries we see are the new negative minima, as predicted by the minima rule.

By considering 2D silhouettes as projections of 3D shapes, and using the tools of differential geometry, we obtain the minima rule for 2D silhouettes.

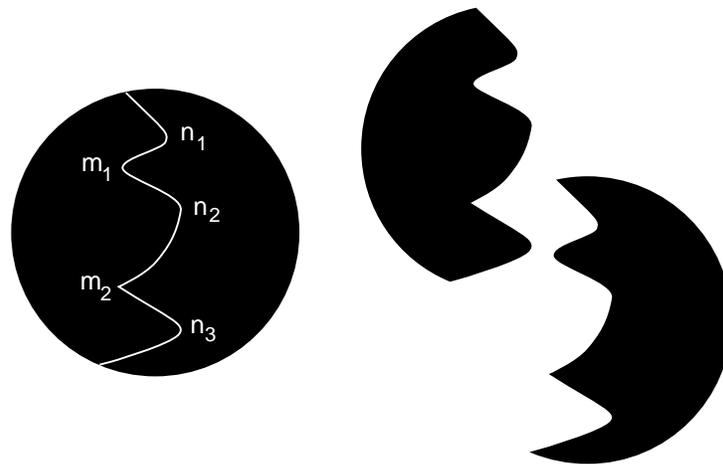


Figure 10. *The wiggle drawn down the middle of the disk appears different, depending on whether it is seen as belonging to the half moon on the left, or the half moon on the right. Note that the part boundaries are different for these two half moons, because switching figure and ground switches convex and concave. Whereas the negative minima of curvature lie at points m_1 and m_2 for the left half moon, they lie at the points n_1 , n_2 , and n_3 for the right half moon.*

e.g., Guillemin & Pollack (1974).

³ Recall that for any point on a surface, there are associated two canonical directions: one in which the surface curves the most, and the other in which the surface curves the least. These are called, respectively, the directions of greatest, and least, curvature. The curvatures in these two directions are called the principal curvatures, and the curves obtained by always moving along these directions are called the lines of curvature (DoCarmo, 1976).

Minima Rule for Silhouettes: Divide silhouettes into parts using points of negative minima of curvature on their bounding contour as boundaries between parts.

This rule explains an interesting effect described by Attneave (1971). Draw a random curve down the middle of a disk (as in Figure 10 on the left) and pull the two halves apart (as in Figure 10 on the right). The wiggles on the two halves look different, even though they are, by construction, identical. The reason, according to the minima rule, is that because border-ownership (or figure-ground) is reversed in the two halves, these induce different parsings on the identical wiggle. For the half disk on the left, the negative minima of curvature are labeled m_1 and m_2 , and these are, by the minima rule, the boundaries between parts. For the half disk on the right, however, the points m_1 and m_2 are not negative minima, but positive maxima of curvature (since they now lie in convex, rather than concave, regions). The negative minima, and hence the part boundaries, for this half are labeled n_1 , n_2 , and n_3 . The wiggles on the two half disks therefore have different parts. Since these parts are the basic units of our perceptual representations of shape, the two half disks receive different representations, and hence they look different.

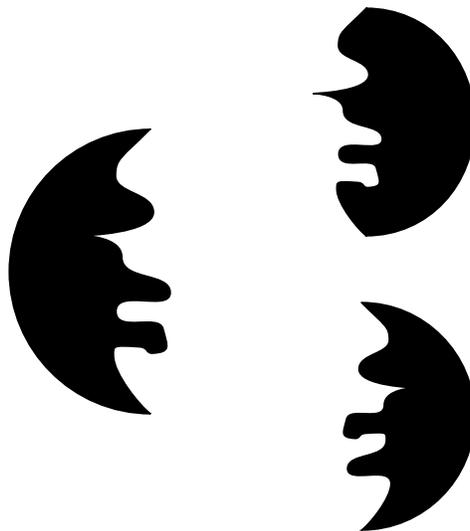


Figure 11. Most observers see the bottom half moon as being more similar to the one on the left, even though it is the one on top whose wiggle is identical to it. Note, however, that the half moon at the bottom has the same parts as the one on the left—even though two of them have been swapped.

It is also instructive to relate the minima rule to Attneave's (1954) claim that information along a contour is concentrated on points of 'maxima of curvature.' Attneave used 'curvature' to mean magnitude of curvature (i.e., as an unsigned quantity). Hence his statement about the importance of 'maxima of curvature' translates, in our framework, to the visual importance of

both negative minima *and* positive maxima of curvature. In the context of parsing, however, the minima rule distinguishes between these two types of extrema, and assigns a special status to the negative minima (i.e., those points of locally highest curvature that lie in concave regions). Indeed, we shall see ample evidence that the visual system treats negative minima quite differently from positive maxima of the same magnitude of curvature.

Psychophysical Evidence

Consider again the large half moon on the left of Figure 11. Which of the two half moons on the right looks most similar to it? In an experiment with such stimuli, Hoffman (1983) found that subjects chose the bottom one, or its analogue, 94% of the time. Note, however, that the half moon at the top has the same wiggle as the one on the left, whereas the half moon at the bottom does not. In fact, the wiggle at the bottom was created by taking the one on the left, mirror reversing it, and then swapping two of its parts.

So why does the half moon at the bottom look more similar to the one on the left? Note that it has the same parts as the one on the left—even though two of these parts have been swapped. The half moon at the top, however, does not. Even though it has the same wiggle as the large shape on the left, it induces an opposite assignment of figure and ground on it; hence the visual system carves it differently into parts. It has different parts, so it looks different.

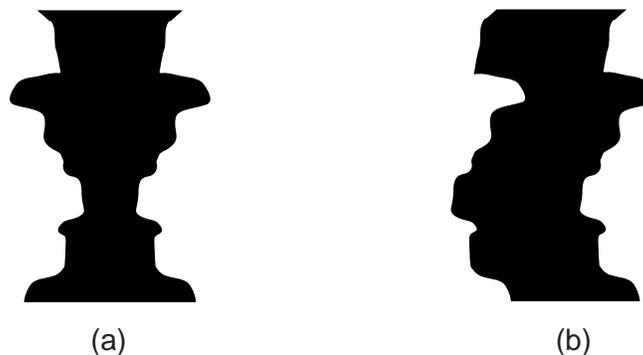


Figure 12. *Human observers can detect symmetry (a) but not repetition (b) in parallel. This is surprising because repetitions involve only a translation, whereas symmetry involves a reflection and a translation. Note, however, that the symmetric figure has matching parts on the two sides, whereas the repeated figure has mismatching parts.*

It has long been observed (at least as far back as Mach; see Baylis & Driver, 1994; 1995), that humans are more sensitive to symmetry in a pattern than to repetition (compare Figure 12a

with 12b). This observation has since been confirmed by researchers (Baylis & Driver, 1994; Bruce & Morgan, 1975; Corballis, 1974). Indeed, Baylis & Driver (1994) found that human vision detects symmetry, but not repetition, in parallel. This advantage for symmetry has been something of a paradox because repetition involves simply a translation of the curves to be compared, whereas symmetry involves a translation *plus* a reflection. Hence any account based on a point-by-point comparison of the two curves should predict an advantage for repetition, not symmetry.

Baylis & Driver (1995a) noted that in a symmetric shape, such as Figure 12a, the two sides have matching negative minima of curvature—and hence matching parts. In a repeated pattern, such as Figure 12b, however, the two sides do not have matching negative minima, because they get opposite assignments of figure and ground. Hence the two sides have mismatching parts. Therefore, if human vision represents the two sides in terms of their parts, as per the minima rule, and then compares them at the level of parts, one would expect it to be more sensitive to symmetry (as in Figure 12a) than to repetition (as in Figure 12b).

In support of their claim, Baylis & Driver (1995a) argued that one should be able to reverse the advantage for symmetry using displays such as Figure 13, in which symmetries have mismatching parts (13a) and repetitions have matching parts (13b). It is easy to create such displays because, as we noted above, the parts on a curve depend on which side ‘owns’ the curve, and is hence assigned the status of ‘figure,’ rather than ‘ground.’ For these new displays, Baylis & Driver (1995a) found that subjects were faster and more accurate at detecting repetition than symmetry. Hence it is matching versus mismatching parts that determines the speed and accuracy of such judgments—not symmetry versus repetition per se.

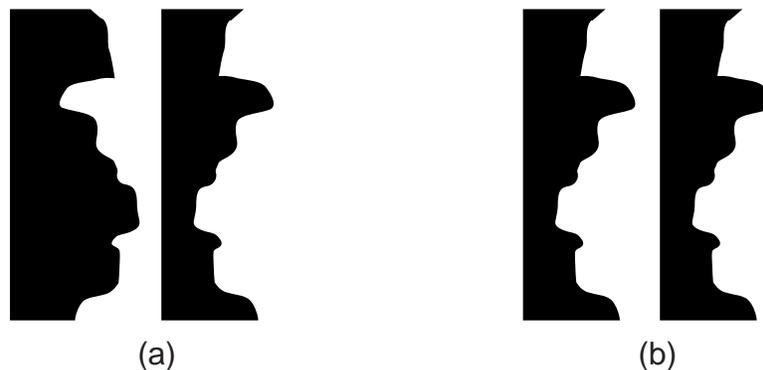


Figure 13. *When the symmetric curves are made to have mismatching parts, and the repeated curves are made to have matching parts, then repetition becomes easier to judge than symmetry. Hence it is matching parts that give symmetry its advantage in displays such as Figure 12. (Adapted from Baylis & Driver, 1995)*

Combined with their earlier study (Baylis & Driver, 1994) which showed that human vision detects symmetry (with matching parts) in parallel—but not repetition (with mismatching parts)—these results suggest that human vision computes parts in parallel. These experiments also demonstrate that human vision compares shapes, not by a process of point-by-point matching of contours, but by a process of matching parts. Hence the basic units of shape description are intermediate-level parts, and not individual pixels and edge segments, or unstructured wholes.

Evidence for the early computation of part boundaries also comes from visual search experiments. In such experiments, subjects search for a target object in a field of distractor objects. If the time to find the target is independent, or nearly independent, of the number of distractors, then the target is said to “pop out” of the field of distractors. In a series of experiments on visual search, Wolfe & Bennett (1997) found that preattentive processing for *overall* shape is limited. They found, however, that a shape having a salient negative minimum of curvature (Figure 14, top) pops out among distractors that do not (Figure 14, bottom), but not vice versa. This kind of asymmetry has traditionally been interpreted as a sign of preattentive processing (Treisman & Souther, 1985), suggesting that part boundaries are computed preattentively—or at least rapidly and early in visual processing.

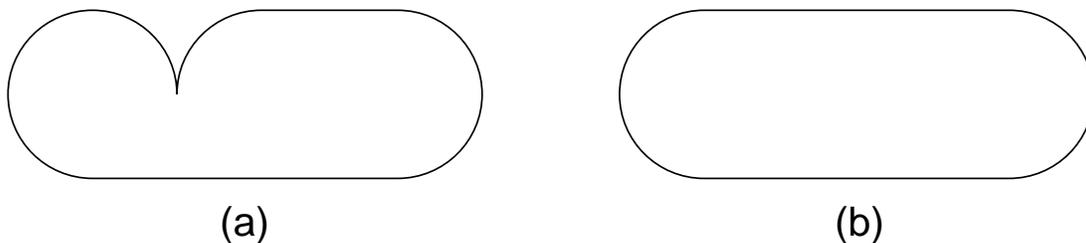


Figure 14. A shape with a sharp negative minimum (a) pops out among distractors that lack a negative minimum (b)—but not vice versa. (Adapted from Wolfe & Bennett, 1997)

One might attribute Wolfe & Bennett's (1997) results to the fact that the shape in Figure 14a has a point of high curvature, whereas the shape in Figure 14b does not. To control for this factor, Hulleman et al. (2000) conducted visual search experiments with shapes as in Figure 15. The shape in Figure 15a has a negative minimum of curvature, whereas the one in Figure 15b has a positive maximum of curvature. These two extrema of curvature are equally sharp; they differ only in that one lies in a concave region, whereas the other lies in a convex region. Hulleman et al. (2000) found that the shape with the negative minimum of curvature (as in Figure 15a) pops out among distractors, each with a positive maximum of curvature (as in Figure 15b)—but not vice

versa (recall Figure 6). Hence not all points of high curvature are processed rapidly, only those that lie in concave regions—in other words, negative minima.⁴

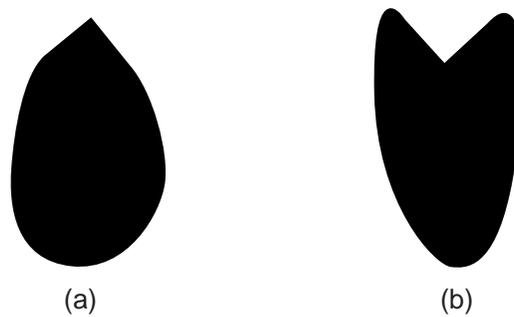


Figure 15. A shape with a sharp negative minimum of curvature (b) pops out among distractors that have a positive maximum of curvature (a)—but not vice versa—even though the two curvature extrema are equally sharp. (Adapted from Hulleman, te Winkel, & Boselie, 2000.)

Further evidence for the minima rule comes from the experiments of Braunstein et al. (1989). Their subjects viewed 3D surfaces of revolution presented in structure from motion, i.e., as computer-generated dot displays set in motion (see Figure 16a). After viewing a surface for 20 seconds, subjects had to choose, from four alternatives, a part that belonged to the object just

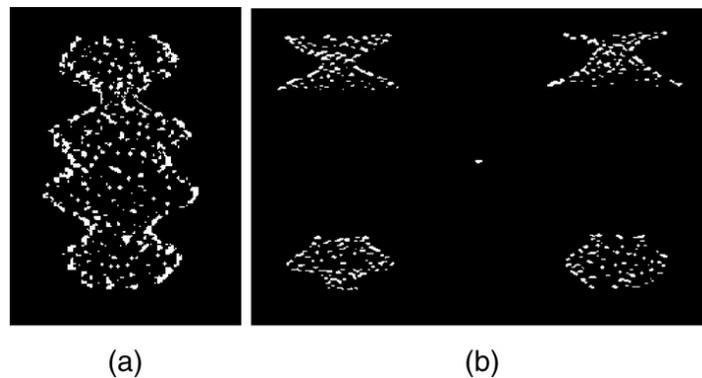


Figure 16. When asked to indicate which of four “parts” (shown in b) was in the initial display (a), subjects are twice as likely to respond with negative-minima parts than positive-maxima parts. (Adapted from Braunstein, Hoffman, & Saidpour, 1989.)

viewed. Of these four alternatives, one was a part defined by negative-minima boundaries, one was a “part” defined by positive-maxima boundaries (i.e., points of locally highest curvature in convex regions), and the other two were distractors (see Figure 16b). 77% of the subjects

⁴ Although Wolfe & Bennett’s and Hulleman et al.’s experiments suggest that part *boundaries* are computed pre-attentively, they do not in themselves show that parts or part structures are also computed pre-attentively. We will return to this issue later (see section on “Implications for Visual Cognition.”)

responded more frequently with negative-minima parts than with positive-maxima parts, and 65% of all correct responses were negative-minima responses. Hence subjects were twice as likely to remember parts defined by negative minima, than those defined by positive maxima. This demonstrates that parts also affect recognition memory for shapes.

In a different experiment, subjects had to mark part boundaries on 3D surfaces, using a joystick to move a red line on the computer monitor. The surfaces had relatively narrow points defined by circular concavities (of low curvature), as well as (sharp) negative minima (see Figure 17). 81% of all responses were made at negative minima of curvature, 10% at positive maxima, and 9% at other places—again supporting the minima rule. (We will later discuss the roles of cut length and boundary strength in parsing, as well as the status of circular concavities.)

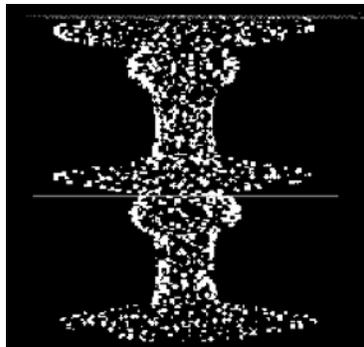


Figure 17. *Subjects are more likely in these display to place part boundaries at the sharp negative minima of curvature than at relatively narrow points with low-curvature. (Adapted from Braunstein, Hoffman, & Saidpour, 1989.)*

Biederman and Blicke (cited in Biederman, 1987) asked subjects to identify objects in line drawings that had some of their contours deleted. They deleted contours in one of two ways: one

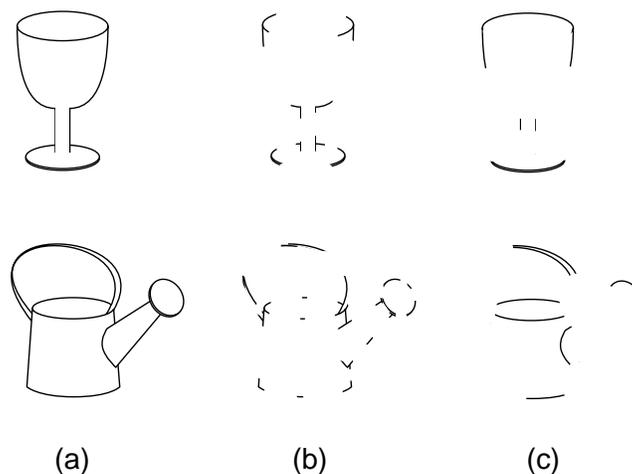


Figure 18. *Subjects can easily recognize objects when contours are deleted in a way that allows parts to be recovered, as in (b). However, recognition becomes extremely difficult when contours are deleted in regions around negative minima—which makes it difficult to recover the parts. (Adapted from Biederman, 1987.)*

which allowed for the recovery of the object's parts (see Figure 18b), and the other which did not (see Figure 18c). They found that subjects had no difficulty identifying objects in the recoverable-parts condition, with error rates approaching zero for long (5s) exposures. But subjects were largely unable to identify objects in the nonrecoverable-parts condition, even with long exposures. They found this difference despite the fact that the proportion of contour deleted in the two conditions was comparable. Although these results do not support any particular set of shape primitives, they do indicate that recovering parts is important for recognition. They also support the minima rule because many of the nonrecoverable-parts stimuli were in fact created by deleting regions around negative minima of curvature (see Figure 18c).

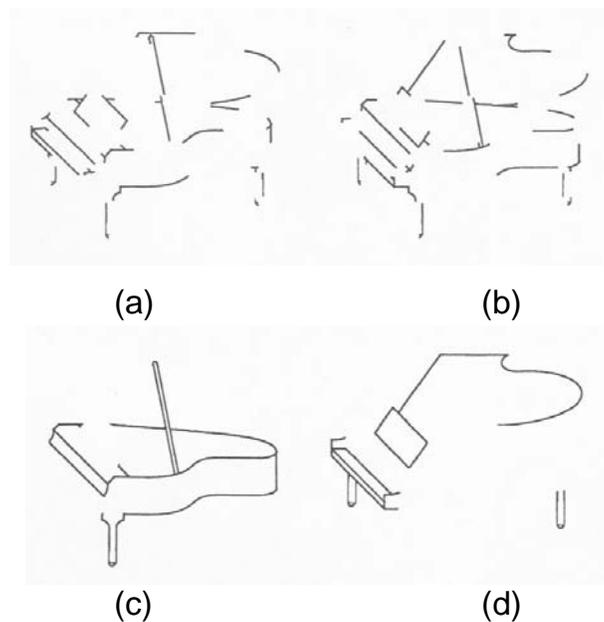


Figure 19. *Priming with complementary contours is as effective as priming with identical contours, when roughly half of the contours in each part is deleted—so that all parts can be recovered in each case (a and b). However, priming with complementary contours is much less effective when the contours of half of the parts are deleted (c and d). (Adapted from Biederman & Cooper, 1991.)*

Biederman & Cooper (1991) first presented their subjects with contour-deleted line drawings. They deleted contours in one of two possible ways: either by deleting every other feature (i.e., edge or vertex), or by deleting every other part. The first manipulation resulted, roughly, in the deletion of half of each part (see Figure 19a), whereas the second resulted in the deletion of half of the parts in the drawing (see Figure 19c). In a subsequent block, they showed subjects contour-deleted images that were either identical to the ones in the first block, or complementary to them (Figures 19b and d). Subjects had to name the line drawings with basic-level names as quickly as possible. Having seen a line drawing earlier produces an advantage in

speed and accuracy the second time through, i.e., *priming*. In the feature-deletion case, the level of priming was as high for complementary images as for identical images—since all parts could be recovered in either case. But in the part-deletion case, the level of priming was much lower for complementary images, in which altogether different parts were seen. Biederman & Cooper (1991) concluded that all of the visual priming in object naming is due to the activation of part-based representations, and none due to the activation of specific edges and vertices. This again points to the importance of parts in object identification and naming, though not to any specific set of pre-defined parts.

Somewhat indirect evidence for part-based descriptions of shape comes from the experiment of Todd, Koenderink, van Doorn, & Kappers (1996). These researchers studied the effects of changing viewing conditions on perceived 3D structure. They presented observers with images depicting smoothly curved surfaces (male and female torsos) that were lit either frontally or obliquely. Observers made judgments of local surface orientation at a large number of probe points over these surfaces. From these local measurements, these researchers could estimate the 3D structure perceived by the observers (Koenderink, van Doorn, & Kappers, 1992). Their results showed that surfaces perceived in monocular and stereoscopic conditions were related by an affine stretching, and similarly, surfaces perceived under frontal and oblique illumination contained a strong affine component (with stereoscopic viewing and oblique illumination eliciting greater perceived depths). In addition, they found that the change in illumination condition did not lead to a *uniform* stretching of the entire surface, but rather, to a stretching that was piecewise uniform. In other words, different “pieces” of the surface appeared perceptually stretched by different amounts—and these “pieces” corresponded to the natural and namable parts of the human torso. We will see further evidence for part-based representations, and their role in visual cognition, in the following sections.

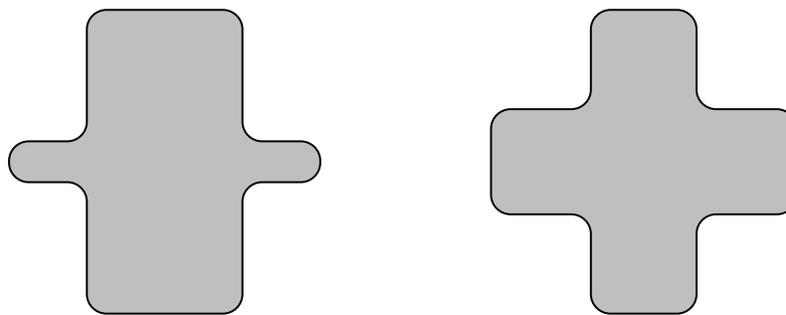


Figure 20. Although negative minima of curvature define part boundaries, they do not determine part cuts. Both shapes in this figure have four negative minima. However, the shape on the left is naturally segmented using two part cuts (into a large vertical body and two parts on the sides), whereas the one on the right is naturally segmented using four part cuts (into a central core and four parts).

FROM PART BOUNDARIES TO PART CUTS

The minima rule defines precise part boundaries on shapes. As we have seen, for a 2D shape these part boundaries are points of negative minima of curvature on the outline of the shape. The minima rule does not, however, define part cuts: It does not tell how to pair negative minima of curvature to create cuts that parse the shape into parts (Beusmans et al., 1987; Siddiqi & Kimia, 1995). For example, for the cross on the left in Figure 20, the minima rule gives the four negative minima as part boundaries, but is silent on how to join these to give cuts. Note that this cross is naturally perceived as a large vertical body with two small parts on the sides. The cross on the right also has four negative minima. However, it is naturally perceived as a small core surrounded by four small parts.

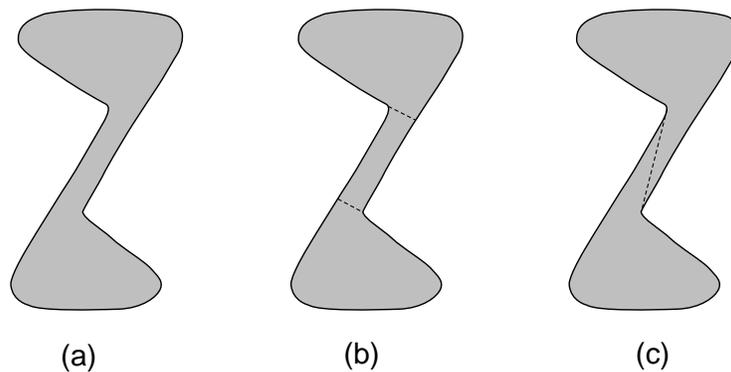


Figure 21. *The natural part cuts for the shape in (a) are shown in (b). Note that each of these cuts joins a negative minimum of curvature to a point of zero curvature. Simply joining the two negative minima, on the other hand—as in (c)—leads to a perceptually unnatural parsing. Thus geometric constraints in addition to the minima rule are needed to define cuts, and hence the parts themselves. (Adapted from Singh, Seyranian, & Hoffman, 1999.)*

In addition, the minima rule does not tell when to use part boundaries other than negative minima of curvature. The shape in Figure 21a, for example, is naturally parsed using the two cuts shown in 21b. Each of these two cuts terminates at one end in a negative minimum of curvature, but at the other end, in a point of zero curvature. Simply joining the two negative minima, on the other hand, produces perceptually unnatural parts (Figure 21c).⁵ Clearly, some rules are needed, in addition to the minima rule, to define part cuts—and hence the parts themselves.

For our current purposes, we take a part cut to be a straight-line segment which joins two points on the outline of a silhouette such that (1) at least one of the two points has negative

⁵ This example also shows that any scheme that creates part cuts by joining *consecutive* negative minima along a contour will give unnatural cuts.

curvature, (2) the entire segment lies in the interior of the shape, and (3) the segment crosses an axis of local symmetry (Brady & Asada, 1984).⁶ Condition (1) builds on the intuition behind the minima rule that concave regions with high magnitude of curvature are important for parsing. However, it allows for the possibility that other geometric factors may also play a role in parsing, thereby pulling part cuts slightly away from negative minima of curvature. Condition (3) ensures the region being separated by the cut is natural enough to be considered for parthood. In Figure 22, for example, the segment ab crosses an axis of local symmetry, whereas the segment bc does not. As a result, ab is a legitimate potential cut, but bc is not.

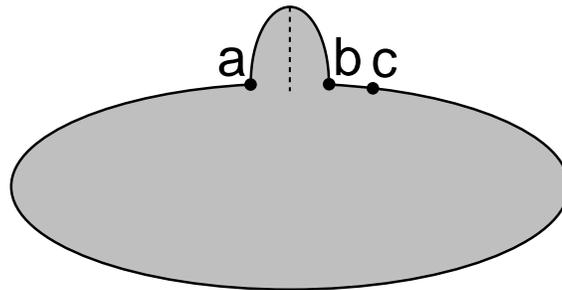


Figure 22. *Only points along the contour of a shape that are separated by an axis of local symmetry are valid candidates for being the end points of a part cut. Thus the join of **a** and **b** is a legitimate part cut, whereas the join of **b** and **c** is not.*

In this section, we consider four geometric factors that determine how human vision creates part cuts. These factors are: the length of a cut, the strength of its boundary points, good continuation, and local symmetry. We consider also the role of ‘simple descriptions.’

Cut Length

Consider the elbow in Figure 23a. Cut pq on this elbow looks far more natural than cut pr. In Figure 23b, we have made the areas of the two segments equal, and pq is still the preferred cut, suggesting that the area of the parts is not determining the cuts in these figures. Instead, examples like these suggest that human vision prefers to divide shapes into parts using the shortest cuts possible.⁷

⁶ Local symmetry is a weak form of symmetry that allows for the axes of symmetry to be curved, and to span only local subregions of the shape rather than the entire shape. The axes of local symmetry provide, in effect, the skeletal axial structure of any given silhouette. See the section on ‘Local Symmetry’ for details.

⁷ These cuts must, of course, satisfy conditions (1)–(3) above. In Figure 22, for example, cut bc is shorter than cut ab—but it is nevertheless unnatural because it does not cross an axis of local symmetry.

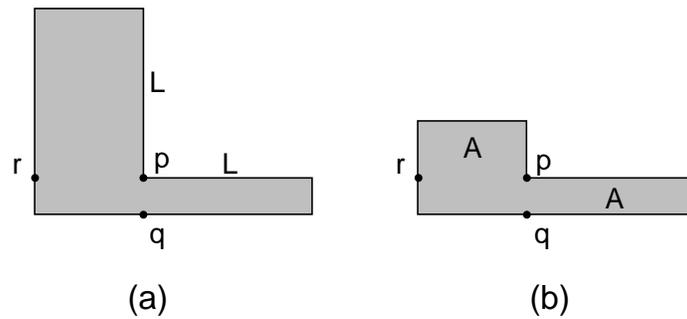


Figure 23. The role of cut length in determining part cuts. *The cut pq in (a) appears far more natural than the cut pr . This is also true in (b) where the areas of the two candidate parts have been equated.*

This “short-cut rule” (Singh, Seyranian, & Hoffman, 1999) is motivated using the geometry of transversal intersections in three dimensions. Consider first the simple case of two cylinders with unequal radii. If these cylinders intersect *completely* (i.e., in a way that produces

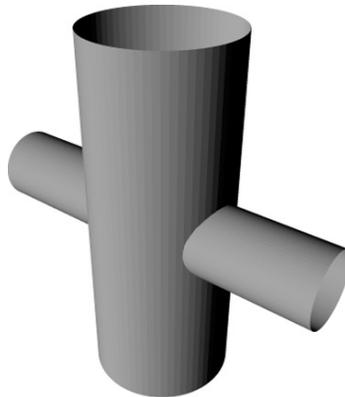


Figure 24. *Motivating the visual system’s preference for shorter part cuts using the transversal intersection of volumes in 3D. When two cylinders of unequal radii intersect completely, the concave creases always encircle the thinner cylinder, and never the thicker one. (Adapted from Singh, Seyranian, & Hoffman, 1999.)*

two, rather than a single, concave creases), these creases always encircle the thinner cylinder and never the thicker cylinder (see Figure 24). If this intersection is now projected onto an image plane, these concave creases project onto the shorter cuts, and not the longer ones. Hence a silhouette of this intersection is naturally parsed using these shorter cuts. Furthermore, as the ratio of the radii (larger to smaller) of the two cylinders increases—say, by keeping the radius of the thicker cylinder fixed, and decreasing the radius of the thinner one—complete intersections become more and more likely (indeed, the probability of obtaining a complete intersection

approaches 1). Hence, the probability assigned to the shorter cuts should increase as the ratio of their radii gets more extreme.

Now, given a silhouette produced by a 3D shape of unknown geometry, the principle of genericity (e.g., Freeman, 1994) assigns high probability to those 3D interpretations in which the shape is about as deep as it is wide in the image. Therefore, as in the case of the cylinders above, the concave creases will encircle the thinner shape, and hence project onto the shorter cuts. Thus the silhouette is naturally parsed using these shorter cuts.

In a series of experiments using crosses and elbows Singh, Seyranian & Hoffman (1999) studied subjects' preferences for making part cuts, as a function of relative cut lengths and relative part sizes induced by the cuts. They found that subjects strongly and consistently preferred to parse shapes using shorter cuts, rather than longer ones. However, subjects did not show a consistent preference for either smaller or larger parts. In addition, their results demonstrated that the short-cut rule can create part boundaries that are not negative minima of curvature (see, for example, the two elbows in Figure 23, and the shape in Figure 21).

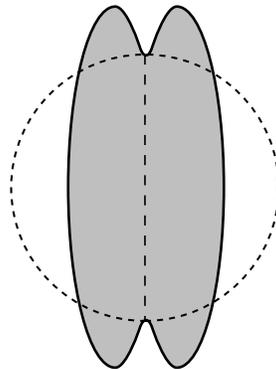


Figure 25. *Siddiqi & Kimia's definition of "neck" fails to capture cases in which a circle of locally minimal diameter cannot be inscribed within the shape.*

Siddiqi & Kimia (1995; Siddiqi et al. 1996) proposed a method for parsing shapes, called 'necks.' A neck is a part cut "which is also a local minimum of the diameter of an inscribed circle" (p. 243). Although this method prefers locally shorter cuts, it measures distances only along diameters of circles that can be inscribed within the shape. This requirement turns out to be too restrictive. Figure 25, for example, shows a shape with a natural cut that should be made; but this cut is not captured by the definition of a neck. The problem is that the circle whose diameter is the cut cannot be inscribed in the shape. The short-cut rule, on the other hand, considers distances between *all* pairs of points on the silhouette outline, as long as these are separated by an axis of local symmetry (recall the three conditions that part cuts must satisfy). For example, in

Figure 23, the short-cut rule explains the preference for the cut pq over the cut pr , but necks fail to explain this preference.

Boundary Strength

Negative minima of curvature are attractors of part cuts: A cut which passes through a negative minimum is typically preferred to one which passes near, but not through, it. All negative minima are not created equal, however: Some are more salient than others (Hoffman & Singh, 1997).

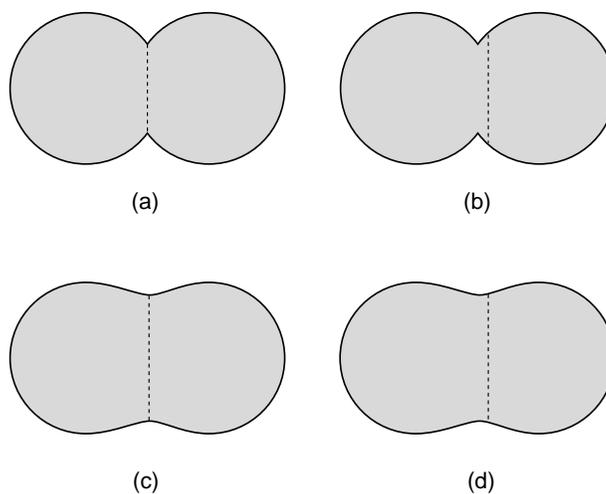


Figure 26. *Sharper negative minima are stronger attractors of parts cuts than weaker negative minima. In (b), a slight deviation of the part cut from negative minima looks clearly wrong. However, in (d) a deviation of identical magnitude appears less contrived.*

Figure 26a, for example, shows a 2D shape with two part boundaries of extremely high strength, together with a natural part cut that joins them. If this part cut is displaced just slightly, as in Figure 26b, it looks clearly wrong. Figure 26c shows a similar silhouette, but with weak part boundaries. Note that the displaced part cut in Figure 26d does not look nearly as wrong as in Figure 26b, even though the displacements are equal in magnitude. Examples such as these suggest that sharper extrema of curvature are more powerful attractors of part cuts.⁸

Boundary strength can interact with the short-cut rule in the following way: For weak part boundaries, the preference for short cuts can sometimes pull part cuts away from negative minima of curvature, as in Figure 27a (figure modified from Siddiqi & Kimia, 1995). Note,

however, that if the part boundaries are sharp, they force the part cuts to pass through them, even if this means making slightly longer cuts, or making two cuts instead of one (see Figure 27b).

Another interaction between boundary strength and the short-cut rule can be seen in Figure 28a. This shape has a narrow region in the middle defined by concave arcs of circles. Each of these arcs is a *region* of negative minima of curvature—so the minima rule by itself does not specify any unique boundary point on them. Furthermore, these concave arcs have low curvature, and hence low boundary strength. At the endpoints of these arcs are negative minima

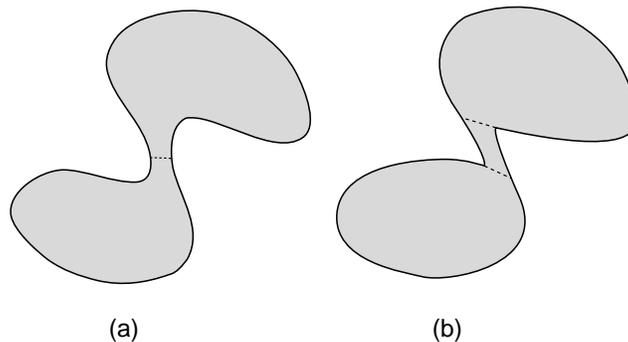


Figure 27. (a) When negative minima are weak, other factors such as cut length can sometimes pull part cuts away from negative minima. (Adapted from Siddiqi & Kimia, 1995). (b) However, when negative minima are sharp, they force the cuts to pass through them—even if this means making two cuts instead of one.

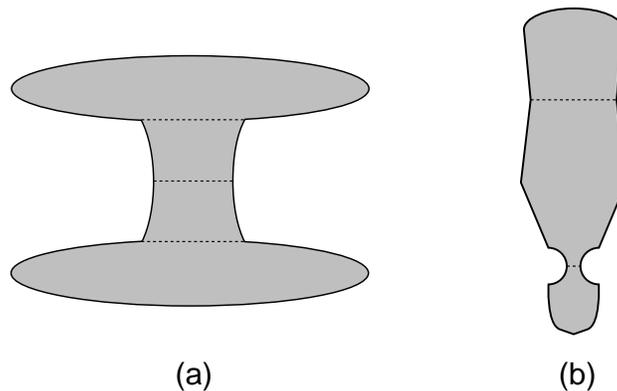


Figure 28. Demonstrating the interaction between cut length and the strength of part boundaries. In (a), the cuts at the sharp negative minima are preferred to the shorter cut at the low-curvature arcs of circles. (Adapted from Braunstein et al., 1989.) In (b), the cut at the arcs of circles is preferred because it is both shorter and involves boundaries with higher salience.

of curvature with high boundary strength. The cuts joining these sharp negative minima are slightly longer than the ‘neck’ cut in the middle; but these cuts are nevertheless preferred by

⁸ We will discuss precise geometric factors that determine the strength of part boundaries in the section on “Part Saliency.”

subjects (Braunstein et al., 1989). So, in this example, the short-cut rule loses to boundary strength. In other cases, the short-cut rule can also win over boundary strength (for example, if the shorter cut joining the two weak boundaries in Figure 28a were to be made extremely short).

Boundary strength and the short-cut rule can also cooperate, rather than compete. In Figure 28b, the ‘neck’ cut at the bottom is short and joins two circular part boundaries with high curvature—and hence high boundary strength. The cut at the top, on the other hand, is longer and joins negative minima with low boundary strength. Hence, subjects in this case prefer the bottom cut. Siddiqi et al. (1996) claim that this provides an example of ‘necks’ winning over the minima rule. But, in fact, neither cut violates the minima rule, and the bottom cut wins simply because *both* the short-cut rule and boundary strength are in its favor.⁹

Good Continuation

Consider the shape in Figure 29. Here the parsing induced by the shorter cuts (shown in Figure 29b) appears less natural than the one induced by the longer cuts (shown in Figure 29c).

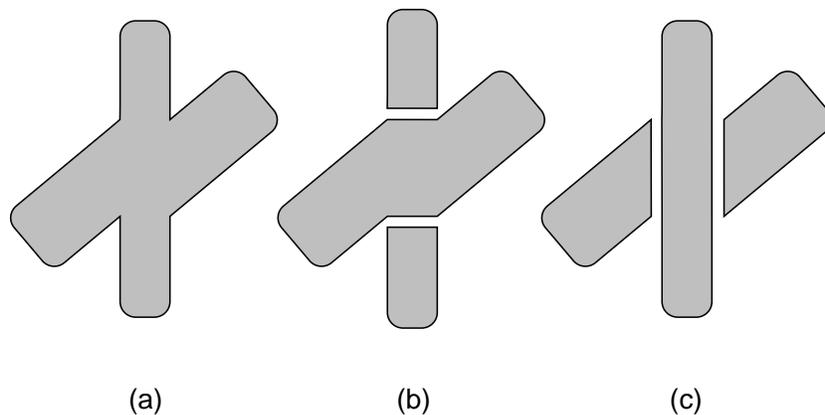


Figure 29. An example of the role of good continuation in parsing. The horizontal cuts in (b) appear less natural than the vertical cuts in (c), even though the vertical cuts are longer.

However, there is another factor at play here, in addition to minimizing cut length: In Figure 29c each cut continues the directions of two tangents at the negative minima of curvature—but not in

⁹ As we mentioned earlier, circular concavities are appropriately thought of as regions, rather than points, of negative minima of curvature. Moreover, circular concavities are nongeneric in that almost any perturbation will induce an isolated point of negative minimum of curvature. Hence, it is not surprising that human vision treats circular concavities no differently than negative minima of curvature.

Figure 29b. Hence good continuation between a pair of tangents (one at each of the two part boundaries) is an important geometric factor for determining part cuts.

Recognizing the role of good continuation in parsing shapes, Siddiqi & Kimia (1995) proposed a scheme for parsing known as “*limbs*.” A limb is a part-line going through a pair of negative minima with co-circular boundary tangents on (at least) one side of the part-line” (p. 243). In other words, Siddiqi & Kimia define good continuation between two vectors by requiring that both be tangent to the same circle. This requirement is much too restrictive, however, since the slightest perturbation of either vector will result in the violation of co-circularity. Indeed, given two random vectors in a plane, the probability is zero that they will be co-circular (Singh, Seyranian & Hoffman, 1999). Hence, this definition almost never applies to real parts.

Although the requirement of co-circularity fails, the intuition of good continuation holds, and other formalizations may be used to capture this intuition. For example, in the context of illusory contour formation, the requirement of contour *relatability* (Kellman & Shipley, 1991) is often used to determine good continuation between two tangents: The extensions of these tangents must intersect, and their exterior angle of intersection must be acute. As it stands, relatability is an all-or-none property, but it is easily extended to a graded measure of good continuation (see Singh & Hoffman, 1999). Intuitively, this graded measure has two components:

- (1) The smoothness of the smoothest possible curve that can be interpolated between the two tangents (captured by *variation in curvature*), and
- (2) The absolute value of the angle through which this smoothest curve turns between the first tangent and the second tangent (captured by *total curvature*).

The smoother the curve, and the smaller the angle it needs to turn through, the better is the continuation between the two tangents. Recently, psychophysical evidence in support of these two measures of good continuation has been provided by Kubovy & Gephstein (2000).

Two additional considerations, however, constrain the role of good continuation in determining part cuts. First, good continuation can play a role in defining part cuts in which only one of the two part boundaries is a negative minimum of curvature. In this case, only the tangent at the negative minimum constrains good continuation—not the tangent at the other part boundary. For example, the perceptually natural cuts in Figure 30a simply continue the directions of (one of) the segments that goes into each negative minimum of curvature. Since the definition of “limb” requires two negative minima, this has the consequence that many parts that are defined by a single negative minimum of curvature—including arms, elbows, and knees—fail to be “limbs” according to Siddiqi & Kimia’s definition.

Second, parts are often defined by smooth negative minima of curvature. At each such negative minimum, there is, by definition, only one tangent. And requiring good continuation between these two tangents can give counterintuitive results. For example, in Figure 30b, a curve that smoothly continues the tangent at one negative minimum into the tangent at the other negative minimum is a poor candidate for the “good continuation” of either part—even though these two tangents are co-circular. This suggests that, for smooth part boundaries, the tangents that constrain good continuation should not be the ones at the negative minima, but at the inflection points closest in arc length to the smooth negative minima. These inflections are likely to provide a stable estimate of the tangent direction of the contour on each side of the negative minimum (see Figure 30c).¹⁰

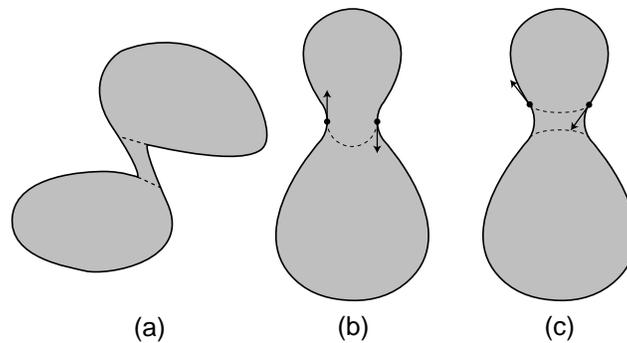


Figure 30. The figure in (a) demonstrates that good continuation plays a role in also defining those part cuts, only one of whose boundary points is a negative minimum of curvature. (b) In the case of smooth negative minima, computing good continuation from tangents at the negative minima themselves yields unintuitive results. (c) In this case, the tangents that constrain good continuation should be the ones at the inflection points closest to the negative minima.

Local Symmetry

The relevance of local symmetry for segmenting shapes was first noted by Blum (1973). The Symmetry Axis Transform (SAT; see below for a precise definition) was put forward as a formalism for local symmetry, and it was suggested that a shape should be divided into parts at branch points of axes of local symmetry. For example, the two branch points in the axes of local symmetry of the "dog bone" in Figure 31a partition it into five parts. Unfortunately, the SAT

¹⁰ It is also possible, however, that when part boundaries are highly smoothed, good continuation simply becomes less relevant, because the estimation of the tangent direction on each side of the negative minimum becomes significantly more difficult.

scheme gives the same axes to a rectangle (Figure 31b), and therefore provides an unintuitive part structure for it (as acknowledged by Blum & Nagel, 1978).

To motivate our own intuitions on the relevance of local symmetry to part cuts, consider the two points A and B in Figure 32a: It is, in part, because these two negative minima of curvature are locally symmetric to each other that they form such a natural candidate pair for a

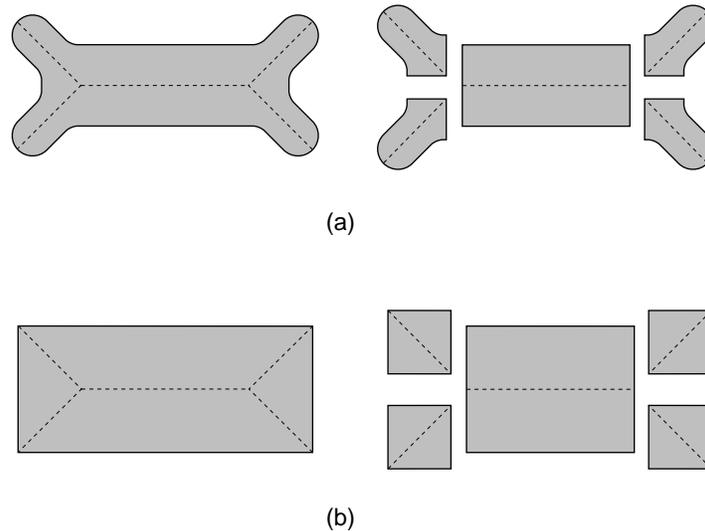


Figure 31. (a) The parsing of the “dog bone” into parts at the branch points of the Symmetric Axis Transform. (Adapted from Blum, 1973.) (b) Unfortunately, this parsing rule gives the same part structure to a rectangle.

cut. As another example, consider the smooth elbow in Figure 32b. There is only one negative minimum of curvature on the shape (located at the point A), and so the minima rule tells us that one endpoint of a part cut must be point A . However, the minima rule does not dictate where the other endpoint of this cut will lie. One could imagine, for example, the cut passing through any of the points between B_1 and B_3 . The most natural cut appears to be the one through B_2 . Interestingly, B_2 is the point (of all the points between B_1 and B_3) that is *furthest* from A . Hence, in this case, some factor other than minimizing cut length is at work. We suggest that this factor is the local symmetry between B_2 and A . Note, similarly, that the natural cuts in Figure 32c are determined by points that are locally symmetric to each of the two negative minima of curvature.

What is local symmetry? Local symmetry may be thought of as a weak form of symmetry. For example, take an object that is perfectly symmetric (Figure 33a), and imagine a process that curves or bends its axis of symmetry (Figure 33b). Observers still see some “symmetry” in the resulting shape. This is one aspect of what local symmetry is intended to capture. Note, though, that this shape has an axis of symmetry that spans the entire shape. The

notion of local symmetry allows, in addition, for axes that span only local subshapes of the entire shape, as illustrated in Figure 33d. What the axes of local symmetry provide, in effect, is the skeletal axial structure of any given 2D shape.

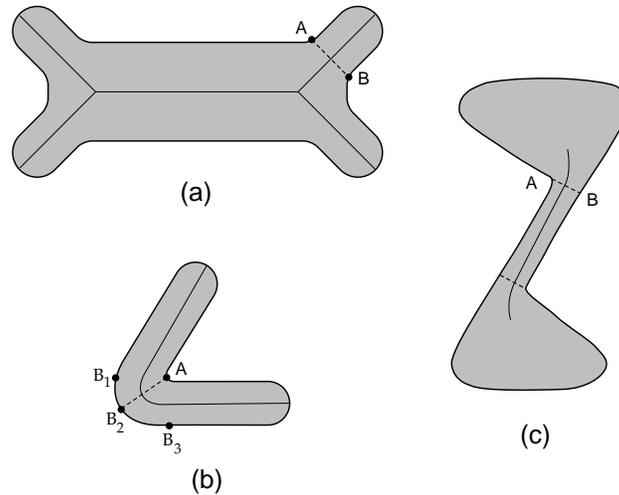


Figure 32. Demonstrating the role of local symmetry in determining part cuts. There is a strong perceptual preference for cuts that join locally-symmetric pairs of boundary points.

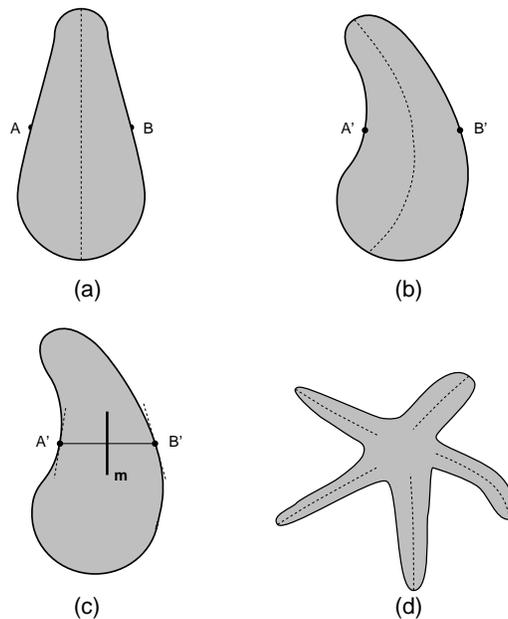


Figure 33. Clarifying the notion of local symmetry. Curving the axis of a symmetric shape (a) leads to a weaker form of symmetry (b). The axis of local symmetry locally reflects the tangent at A' into the tangent at B' (c). The axes of local symmetry span local subshapes, providing a skeletal structure of the shape (d).

There are at least two respects, therefore, in which the notion of local symmetry can be said to be "local." First consider the two points A and B in Figure 31a that are symmetric with

respect to each other. In what sense can their corresponding points, A' and B' , in the transformed shape (Figure 33b), be said to be locally symmetric? Imagine a mirror being placed at the point where the curved symmetry axis intersects the join of A' and B' (see Figure 33c). And suppose that the mirror is oriented perpendicular to this join. Then not only will this mirror reflect the point A' into the point B' , but it will also reflect the tangent to the curve at A' to the tangent to

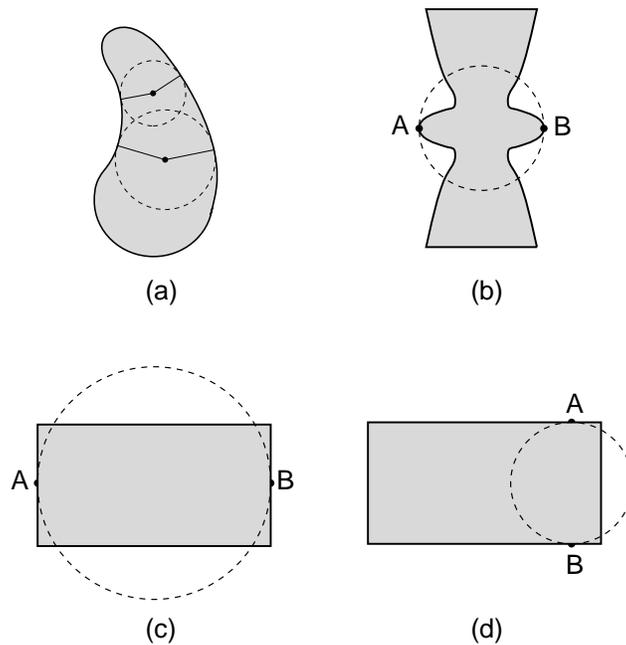


Figure 34. (a) Blum's SAT is defined by taking the locus of the centers of circles that can be inscribed within the shape. (b)–(d) demonstrate some shortcomings of Blum's definition. The requirement that the circle must lie within the shape means that the pairs of points A and B shown in these displays are not locally symmetric according to this definition.

the curve at B' (Leyton, 1992). In other words, the curve in an (infinitesimal) neighborhood of the point A' is reflected into the curve in an infinitesimal neighborhood of the point B' : hence the "local" symmetry. The other sense in which local symmetry is "local" is simply that no global symmetry axis is required: one may have local axes spanning local subshapes (as in Figure 33d).

Many different schemes for local symmetry have been put forward to capture these intuitions. The most prominent among them are the SAT scheme mentioned above (Blum, 1973; Blum & Nagel, 1978) and Smoothed Local Symmetry (SLS; Brady, 1983; Brady and Asada, 1984).

Blum's SAT is obtained by considering maximal disks that lie entirely in the interior of a given shape, and taking the locus of their centers as an axis (as in Figure 34a). This scheme has a number of shortcomings (Brady, 1983; Leyton, 1992). For example, in Figure 32b, the points A

and B are clearly symmetric. However, it is impossible to draw a circle that lies entirely in the interior of the shape and is tangent to it precisely at A and B . For the same reason, the SAT scheme does not give the minor axis of a rectangle (Figure 34c), nor does it give all of the major axis (Figure 34d). In fact, as we noted above (in Figure 31), the SAT of the rectangle is the same as that of the “dog bone.”

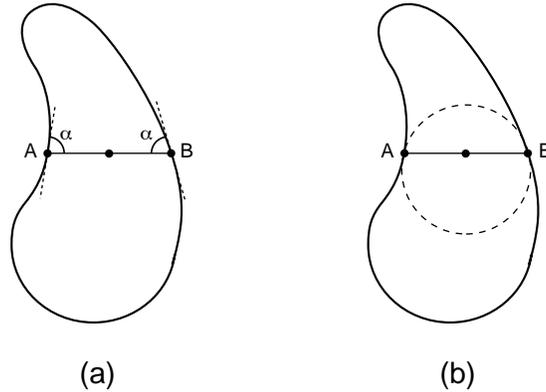


Figure 35. Two points are locally symmetric according to Brady’s SLS scheme if the tangents at these points make equal angles with their join. The axis of local symmetry is then defined by the locus of midpoints of the joins of such pairs of points. This scheme preserves all the advantages of Blum’s SAT, but does not require a circle to be inscribed within the shape.

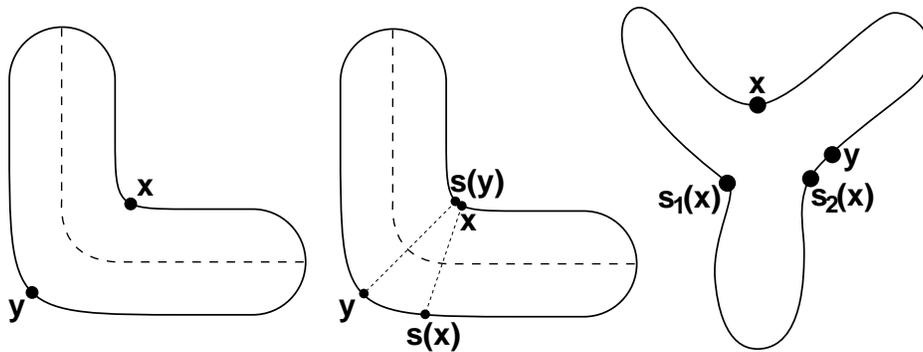


Figure 36. Computing the extent to which two points on a shape deviate from being locally symmetric.

Brady’s SLS scheme is formulated as follows: two points A and B are said to be locally symmetric if their tangents make equal angles with the line segment that joins them (Figure 35a). The axis of symmetry is then defined by taking the locus of the *midpoint* of the join \underline{AB} (Figure 35b). The SLS scheme maintains all the advantages of the SAT, but eliminates its shortcomings. In particular, it does not require the circle that is simultaneously tangent to the contour at points

A and B to lie in the interior of the shape (the existence of such a circle is guaranteed by the SLS construction; see Leyton, 1992). We adopt Brady's SLS formalism for local symmetry.

Given that pairs of points that enjoy local symmetry yield more natural cuts than those that do not, how should we motivate a formal definition of this term. Consider two points x and y on the shape in Figure 36a. Let $s(x)$ denote the point on the shape that is locally symmetric to x . Then the arc-length distance (i.e., distance measured along the contour) between the points y and $s(x)$, denoted $\|y - s(x)\|_a$, measures how much the point y deviates from local symmetry to x .¹¹ Similarly, the arc-length distance between the points $s(y)$ and x , denoted $\|x - s(y)\|_a$, measures how far the point x deviates from local symmetry to y (see Figure 36b). In general these two distances are not equal. Their mean, then, provides a measure of how much the pair $\{x,y\}$ deviates from local symmetry (Singh, Seyranian, & Hoffman, 1996).

Simple Descriptions: Number of parts and convexity

The idea behind parsing shapes is to allow for efficient descriptions of these shapes in terms of simpler components. Hence the most natural parsings are those that lead to the simplest parts, and the simplest descriptions in terms of parts. This suggests two constraints: (1) Other things being equal, parsings with fewer parts should be preferred to those with more parts. If a simple

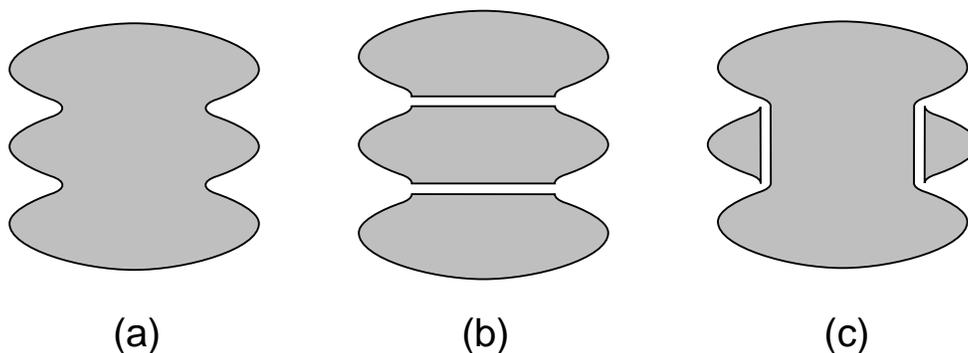


Figure 37. *The parsing depicted in (b) is more natural than the one in (c) even though it requires longer part cuts. Note that the parts in (b) have no further negative minima, whereas the central part in (c) has four negative minima and can thus be parsed further.*

¹¹ In general, if the point y has n points locally symmetric to it, namely, $s_1(x), \dots, s_n(x)$, the point that is relevant in evaluating the cut xy is clearly the one that is closest, in arc-length distance, to y (see Figure 36c). In this case, the term $\|y - s(x)\|_a$ needs to be replaced by $\min_i \{\|y - s_i(x)\|_a\}$.

description of a shape has already been achieved with a small number of parts, there is no need for further parsing. (2) Parsings that create parts with *no* negative minima on their outlines should be preferred to parsings that create parts *with* negative minima—since the presence of negative minima is indicative of further part structure. The simplest shapes are convex shapes—whose outlines have positive curvature throughout (see, e.g., Rosin, 2000, for the role of convexity in parsing). If the outline of a shape has regions of negative curvature, especially if these regions contain salient negative minima of curvature, this usually indicates that the shape can be further parsed to give simpler subshapes. As an example, consider the shape in Fig 37a. The parsing in 37b looks more natural than the one in 37c, even though it involves making slightly longer cuts. The reason is clear: The resulting central “part” in 37c has four negative minima of curvature, and can be parsed further into three different parts; whereas the parsing in 37b leads to parts with no further part structure.

PART SALIENCE

As the shape in Figure 38a indicates, perceptual parthood is a graded—rather than an all-or-none—notation. In other words, parts are not all created equal. For example, in Figure 38a, the part labeled **A** is more visually salient than the part labeled **B**. Three main geometrical factors determine the perceptual salience of a part (Hoffman & Singh, 1997): (1) its protrusion, (2) its relative area, and (3) the strength of its boundaries. Salience of a part increases as its protrusion, relative area, or boundary strength increases. In this section we briefly consider protrusion and

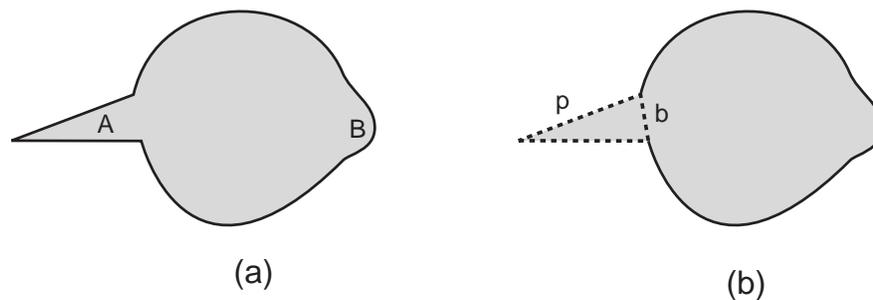


Figure 38. *Perceptual parthood is a graded, rather than an all-or-none, notion. Some parts (such as **A**) are clearly more perceptually salient than others (such as **B**).*

relative area, and then discuss in more detail the strength of part boundaries. We restrict attention to 2D shapes; the theory for 3D shapes is more complex and discussed elsewhere (Hoffman & Singh, 1997).

The protrusion of a part is, roughly, how much the part sticks out. This can be made more precise using terminology illustrated in Figure 38b. Here the outline of part **A** is divided into a base b and a perimeter p . The protrusion of **A** is simply the ratio p/b (Hoffman & Singh, 1997). This definition of protrusion is scale invariant (i.e., it gives the same value if the entire shape is uniformly shrunk or expanded), and it gives larger values to parts long and narrow than to parts short and wide. The greater the protrusion of a part, the greater is its visual salience. In Figure 38a, part **A** has greater protrusion than part **B**.

The effect of relative area on visual salience has long been recognized by Gestalt psychologists (e.g., Rubin, 1958). The relative area of a part is the area of that part divided by the total area of the whole object, a definition that is scale invariant. Again, the greater the relative area of a part, the greater is its salience. In Figure 38a, part **A** has greater relative area than part **B**.

The strength of a part boundary depends first on how “sharp” the boundary is; sharper boundaries have greater strength, and lead to more salient parts. Concave corners, such as those that define part *A* in Figure 38a, are sharper than smooth boundaries, such as those that define part *B*. Intuitively, the difference is in degree of curvature. Concave corners have infinite curvature whereas smooth boundaries have finite curvature. And various smooth boundaries can differ in their magnitudes of curvature, leading to differences in their strengths. The only caveat here is that curvature is not a scale invariant quantity, so that curvature simpliciter is not an appropriate measure of boundary sharpness. But this is easily fixed by using relative curvatures, or by introducing a normalization factor (Hoffman & Singh, 1997). We can state this as the following hypothesis:

Hypothesis of normalized curvature: The salience of a part boundary increases as the magnitude of normalized curvature at the boundary increases.

One rationale for this hypothesis follows from the principle of transversality. Recall that two shapes will generically intersect transversally, and that a transversal intersection creates concave corners at each point of intersection. These concave corners are the boundaries where one shape stops and the next shape begins, and are therefore part boundaries on the composite body formed by the two intersecting shapes. Boundaries that are not concave corners can be obtained from these corner boundaries by various amounts of smoothing. The greater the smoothing, the lower is the (magnitude of) curvature of the resulting boundary, and the further from the primary transversal case. This is illustrated in Figure 39a, where two rectangular shapes, **A** and **B**, intersect transversally at points p and q . If we smooth the transversal intersection at p

just a little, to obtain curve 1, we induce an extremum of curvature with a higher magnitude of curvature; if we smooth a lot, to obtain curve 2, we induce an extremum of curvature with a lower magnitude of curvature (Bennett & Hoffman, 1987). In short, lower-curvature boundaries are further from (i.e., more smoothed from) the original transversal case than higher-curvature boundaries, and therefore have lower salience as boundaries.

We can be more precise about the smoothing in Figure 39 and its effect on the magnitude of curvature at the induced smooth part boundary. We represent the transversal intersection at point p in Figure 39 by the equation $xy = 0$. Various hyperbolic smoothings of this transversal intersection can be written $xy - \varepsilon = 0$, where $\varepsilon > 0$ parametrizes the hyperbolic smoothings such that larger values of ε correspond to greater smoothings. The curvature on the smoothing $xy - \varepsilon = 0$ is given by

$$\kappa(x, \varepsilon) = \frac{-2\varepsilon x^{-3}}{\sqrt{1 + \varepsilon^2 x^{-4}}}$$

and the extremum of curvature is given by

$$\kappa_{ext}(\varepsilon) = \frac{1}{\sqrt{2\varepsilon}},$$

so that as $\varepsilon \rightarrow 0$, i.e., as the smoothing decreases, the magnitude of curvature at the extremum increases (Bennett & Hoffman, 1987).

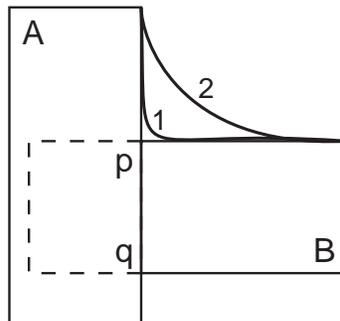


Figure 39. *Demonstrating the role of curvature in determining the perceived salience of a part boundary. Greater degrees of smoothing of a transversal intersection lead to negative minima of curvature with lower magnitudes of curvature.*

Because curvature is a purely local notion, however, it fails to capture another determinant of perceived boundary strength, namely, the angle through which the contour turns in a neighborhood of the negative minimum (Hoffman & Singh, 1997). All concave-corner boundaries, for example, have infinite curvature, but their perceptual salience varies greatly as a function of the angle through which the contour turns around the boundary: Sharp corners require the contour to turn through a large angle (Fig 40a), whereas shallow corners require small turning

angles (Fig 40c). This notion can also be generalized to smooth negative minima by considering them to be the result of smoothing concave corners of different turning angles. Given a smooth part boundary, a stable estimate of the turning angle is obtained by measuring the angle through which the contour turns between the nearest inflection point on each side of the negative minimum—a region we term the *locale* of the boundary (Hoffman & Singh, 1997). This yields the following hypothesis:

Hypothesis of Turning Angle: The salience of a negative-minimum boundary increases as the magnitude of the turning angle around the boundary increases.

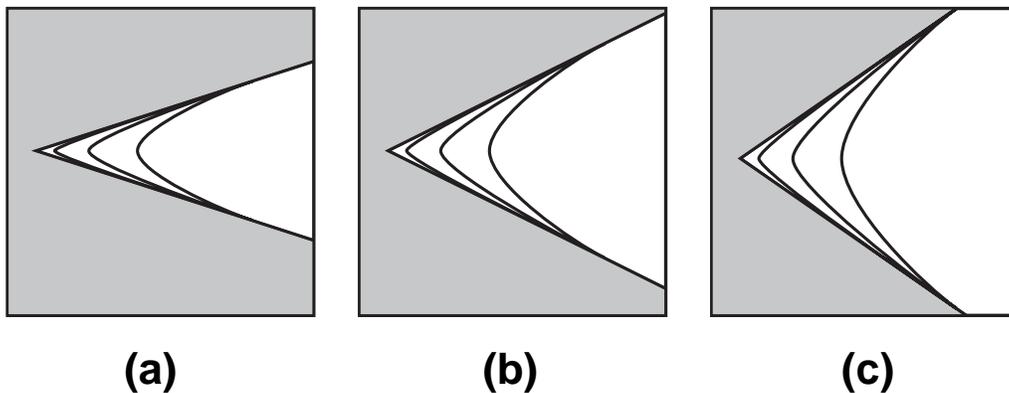


Figure 40. *The turning angle captures how much the tangent must turn in a neighborhood of the part boundary. This parameter can vary independently of the magnitude of curvature at the point of negative minimum.*

As is clear from Figure 40, the curvature at a negative minimum and its turning angle can vary independently of each other. Smooth negative minima with the same turning angle can have different magnitudes of local curvature, and vice versa. Both variables must thus be taken into account to determine the perceived boundary salience (see the section on “Parts and Transparency” for an experimental design involving both variables).

IMPLICATIONS FOR VISUAL COGNITION

In earlier sections, we have seen how parts can explain many aspects of our visual phenomenology. In this section, we consider in more detail how mechanisms that compute part

structure and part salience interact with mechanisms of object and surface segmentation, and with the allocation of visual attention.

Parts and Figure-Ground Assignment

An important early function of human vision is to parse the retinal image into regions that are likely to correspond to distinct objects. To do so, the visual system must decide which of the two sides of a given image contour corresponds to an object, and hence ‘owns’ that contour, and which corresponds to the underlying background. We hypothesize that this choice of figure and ground is intimately related to the salience of parts created on both sides of the contour (Hoffman & Singh, 1997).

Hypothesis of salient figures: Other things being equal, that choice of figure and ground is preferred which leads to the most salient parts for the figure object.

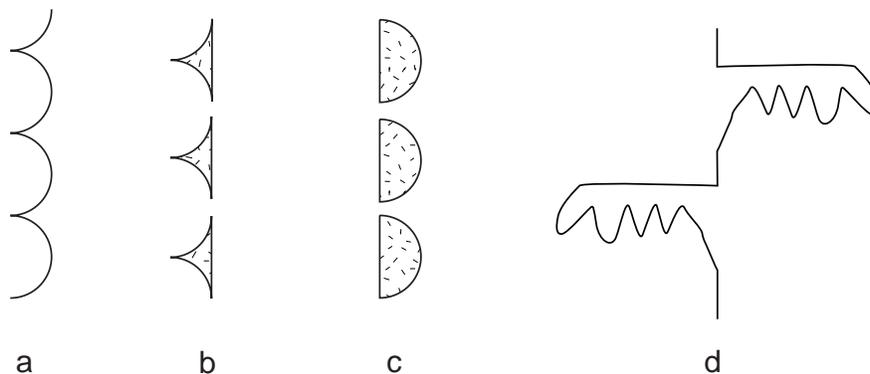


Figure 41. *Illustrating the Hypothesis of Salient Figures. The left of the display in (a) is more likely to be seen as figure because it leads to more salient parts (c), than the parts obtained if the right side were taken to be figure (b). This rule can lead to globally inconsistent assignments of figure and ground (d).*

Many factors, of course, affect the perception of figure and ground, such as symmetry (Bahnsen, 1928; Hochberg, 1964), size and contrast (Rubin, 1958; Koffka, 1935), and convexity (Kanizsa & Gerbino, 1976; Stevens & Brookes, 1988). Part salience is another factor. To see how it works, consider the curve in Figure 41a. It has two possible choices of figure and ground, one in which the right side of the curve is seen as figure, which would lead to the parts shown in Figure 41b, and one in which the left side is seen as figure, which would lead to the parts shown in Figure

41c. Although both sets of parts have the same value of protrusion, the parts in Figure 41c are more salient than the parts in Figure 41b because the parts in 41c have sharper boundaries and greater area. Thus the hypothesis of salient figures predicts that the parts in Figure 41c will be seen more frequently—and hence the left side of the contour in Fig 41a is more likely to be seen as figure.

It has often been argued that no shape description occurs (or even, can occur) prior to the assignment of figure and ground—and that only the figure side of a contour is ever given a shape description (e.g., Koffka, 1935; Baylis & Driver, 1995; Palmer 1999). Contrary to this, the hypothesis of salient figures suggests that human vision analyzes *both* sides of a visual contour, and it does so prior to assigning one of them as “figure.” This analysis takes place at least to the extent that both sides of a contour are segmented into parts, and the perceptual salience of these parts is compared (indeed, such analysis is a critical component of the process of assigning border ownership). Moreover, according to the hypothesis, the assignment of figure and ground need not be globally consistent—because, e.g., the left side of a contour may have more salient parts along some portion of a contour, but the right side may have more salient parts along other portions (see Figure 41d).

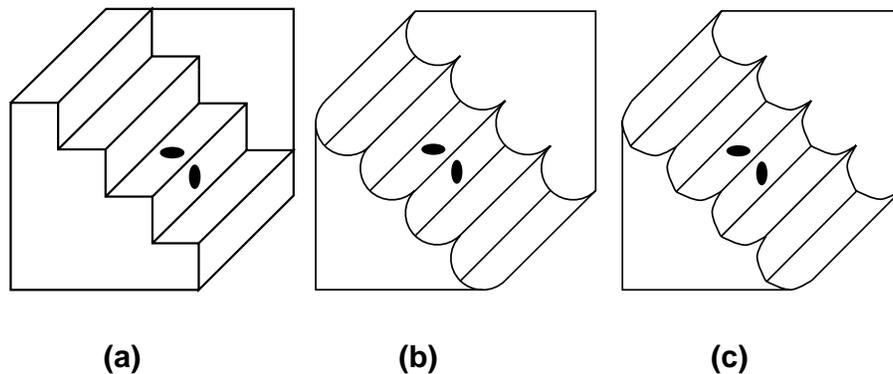


Figure 42. (a) depicts the standard Schröder staircase which can be seen either in the upright or the inverted interpretation. There is usually a bias to see the upright interpretation. This bias is reversed in (b) and (c), which are more likely to be seen in the inverted interpretation, since this choice yields more salient parts for the figural side.

One experimental test of the theory of part salience used three stimuli shown in Figure 42 (Hoffman & Singh, 1997). In Figure 42a is the standard Schröder staircase. In Figures 42b and 42c the staircase is modified so that the steps for the inverted figure-ground interpretation are more salient. As you can see, the inverted interpretation is most easily seen in Figures 42b and 42c, as predicted by the theory of part salience. Subjects were shown these figures using brief 50 millisecond flashes followed by a mask, and were asked to report whether they saw the two dots

as on the same step or on different steps. From these reports it was possible to determine which choice of figure and ground the subjects saw. The results showed a significant effect of part salience in the predicted direction.

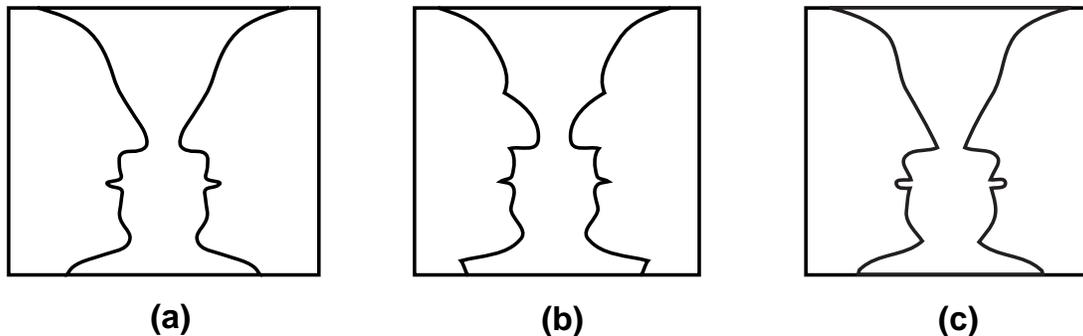


Figure 43. (a) Rubin's standard "face-vase" illusion, which can be seen either as a vase or as two faces. One can bias the display either toward the face interpretation (b), or toward the vase interpretation (c), by modulating the relative salience of the parts on the two sides of the curves.

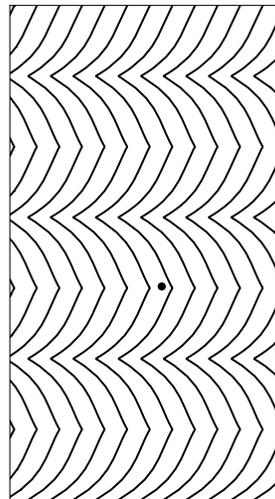


Figure 44. The dot is seen as lying on top of a hill, rather than in a valley. This is consistent with the fact that assigning figure to the left side of the repeated curves leads to more salient parts.

Another experimental test used three stimuli shown in Figure 43 (Hoffman & Singh, 1997). Figure 43a is the standard face-goblet figure, which can be seen either as a goblet in the center or as two faces looking at each other. In Figure 43b the face-goblet figure is modified so that the parts of the faces have the most salient part boundaries (sharp corners) and the parts of the goblet have less salient part boundaries (smoothed boundaries). Figure 43c shows the opposite modification. The theory of part salience predicts a roughly equal balance of face interpretations and goblet interpretations in Figure 43a, a preference for seeing the faces in Figure

43b, and a preference for seeing the goblet in Figure 43c. Subjects were shown these stimuli using 250 millisecond displays followed by a mask, and were asked to report which interpretation, faces or goblet, they saw. The results again showed a clear effect of part salience in determining figure and ground.

A simple demonstration of the power of part salience to influence choice of figure and ground is shown in Figure 44 (Hoffman & Singh, 1997). The dot in this figure can in principle be seen as either lying in a valley or as lying on top of a hill. The hypothesis of part salience predicts that we should see the dot lying on top of a hill. You can check the prediction for yourself; and you can see if it still holds as you turn the page upside down.

Parts and Transparency

The visual parsing that we have considered so far involved a “sideways” parsing of the image into regions that correspond to distinct parts of objects. In other words, one *side* of a perceived part boundary is assigned to one part, and the other *side* to another part. The perception of transparency, on the other hand, provides a striking example of parsing in depth. Here a single intensity value in the image must be parsed into two distinct surfaces along the same line of sight—one of which is seen *through* the other. For example, the dark gray patch in Figure 45 is parsed perceptually into two surfaces: an underlying black surface, and an overlying mid-grey transparent filter. Similarly, the light grey patch in Fig 45 is parsed into an underlying white surface, and a mid-grey filter. Thus, in order to construct the percept of transparency, the visual system must *scission* (Koffka, 1935) or decompose the luminance in a given image region and assign the different components to distinct surfaces. In this section, we show that this parsing in depth interacts in an interesting way with the "sideways" parsing into parts.

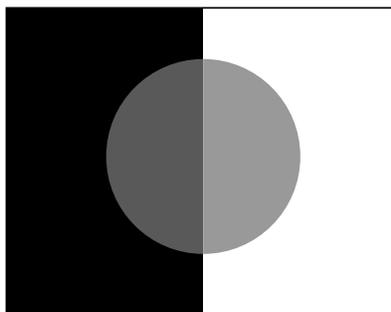


Figure 45. *Computing perceived transparency requires the depthwise segmentation of an image region (e.g., the mid-grey half disks) into two surfaces—one seen through the other.*

When does the visual system initiate the scission of an image region to create the percept of transparency? Two main kinds of image conditions have been studied that lead to perceptual transparency: (a) photometric conditions, involving relative luminance (or color) values, and (b) figural conditions, involving geometric or configural factors.

In the context of simple displays like Figure 45, two luminance conditions are critical (Metelli, 1974; Gerbino, Stultiens, Troost, & de Weert, 1990; Beck, Prazdny, & Ivry, 1984; Anderson, 1997). First, the two grey patches in the center must have the same contrast polarity as the two halves of the surround. For example, in Fig 46a, the contrast polarity of the two grey patches has been reversed and, as a result, the central region no longer scissions into multiple layers. Second, the central region (consisting of the two grey patches) must have a lower “contrast” than the surround. This lowering in “contrast” has typically been defined in terms of lowered reflectance differences (Metelli, 1974), lightness differences (Beck et al., 1984), and luminance differences (Gerbino et al., 1990) across a contour. However, Singh & Anderson (2001) have recently shown that the critical variable is Michelson contrast.¹² The central region in Figure 46b, for example, has been given higher contrast and, as a result, the central region no longer appears transparent. (Note, however, that now the surrounding region may appear transparent.) A limiting case of transparency occurs when the contrast in the central region is zero: now the central regions appear simply as an opaque occluding surface (see Figure 46c).

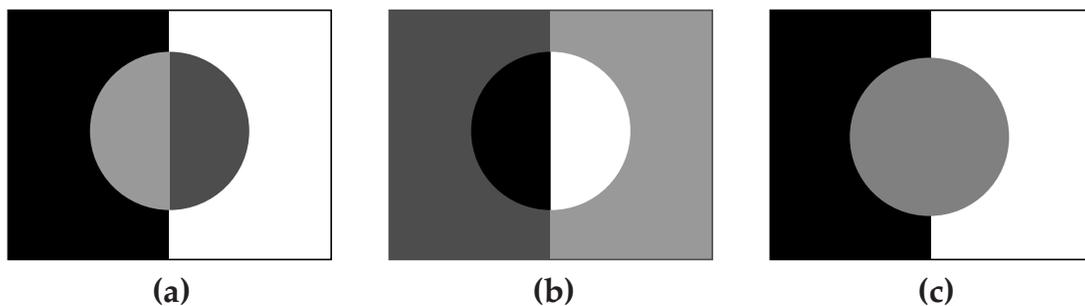


Figure 46. *Illustrating the luminance conditions for perceptual transparency. (a) The central region must preserve contrast polarity with respect to the surround. (b) the central region must have lower contrast. (c) Occlusion is a limiting case of transparency.*

In addition to these luminance conditions, figural conditions for transparency have also been studied—mostly by the Gestalt psychologists Kanizsa (1979) and Metelli (1974). These figural conditions may be divided naturally into two kinds. The first requires continuity of

¹² The fact that the visual system uses Michelson contrast has the perceptual consequence that observers systematically underestimate the opacity of darkening transparent filters and overestimate the opacity of lightening transparent filters (see Singh and Anderson, 2001, for details).

contours (and, more generally, textures) on the underlying surface. For example, if the grey patches are shifted together so that their mutual border no longer aligns with the black-white border (see Figure 47a), the percept of transparency is lost. This is consistent with the fact that although transparent filters reduce the contrast across underlying contours, they cannot produce large shifts in the positions of these contours.

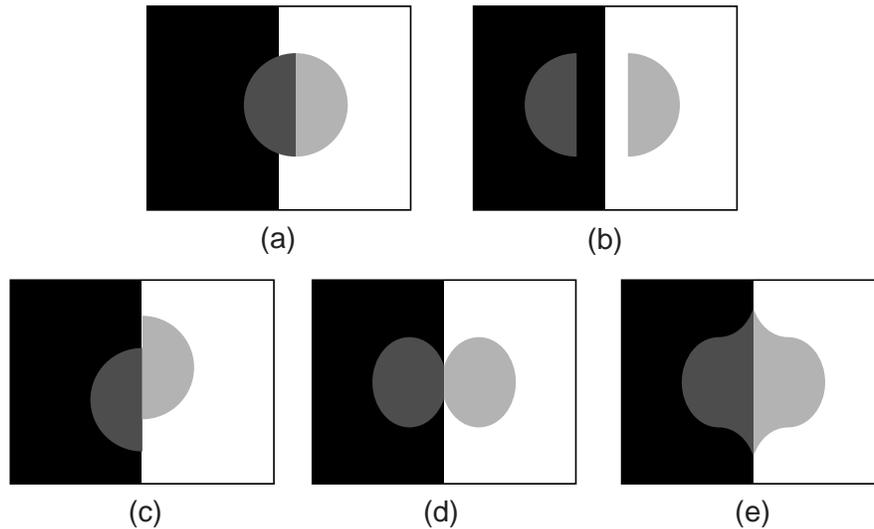


Figure 47. *Demonstrating the role of figural conditions in initiating percepts of transparency. (a) The border between the two grey patches must be continuous with the boundary separating the bipartite background. (b) The two grey patches must unite into a single surface. There must be no discontinuous jumps (c) or tangent discontinuities (d,e) along the contour of the putative transparent surface. Note that, although (d) and (e) both involve tangent discontinuities, the presence of negative minima has a more deleterious effect on perceived transparency.*

The second kind of figural condition requires continuity and grouping of the image regions that define the putative transparent surface. In the extreme, if the two grey patches are separated—so that they no longer unite into a single surface—they no longer appear transparent (see Figure 47b). The percept of transparency is also lost in Figures 47c and 47d, even though there could be a transparent filter with the odd shape shown in Figure 47c or in Figure 47d. One reason why transparency is not seen is provided by the principle of genericity. According to this principle, the visual system rejects unstable scene interpretations of image data. In Figures 47c and 47d, the interpretation of transparency is unstable because it assumes a special (or “accidental”) placement of the transparent filter in front of the bipartite background, and a special vantage point of the observer—thus resulting in the precise alignment of the discontinuities on the filter with the contour separating the bipartite background. As a consequence, a slight displacement of either the transparent filter or the observer’s vantage can lead to a large change in the image.

In addition, the evidence presented above, that part boundaries are computed early in visual processing, suggests another explanation (Singh & Hoffman, 1998). In both Figures 47c and 47d the luminance boundary between the two grey regions precisely aligns with the part boundaries between the regions. If the visual system permits different parts of an object to have different reflectances, then the luminance change from light grey to dark grey would be interpreted as a reflectance change on an opaque surface—thus leading to a suppression of the percept of transparency. Hence, the visual system might operate according to the following “part-coloring rule”: Interpret changes in image color that align with part boundaries as changes in surface reflectance.

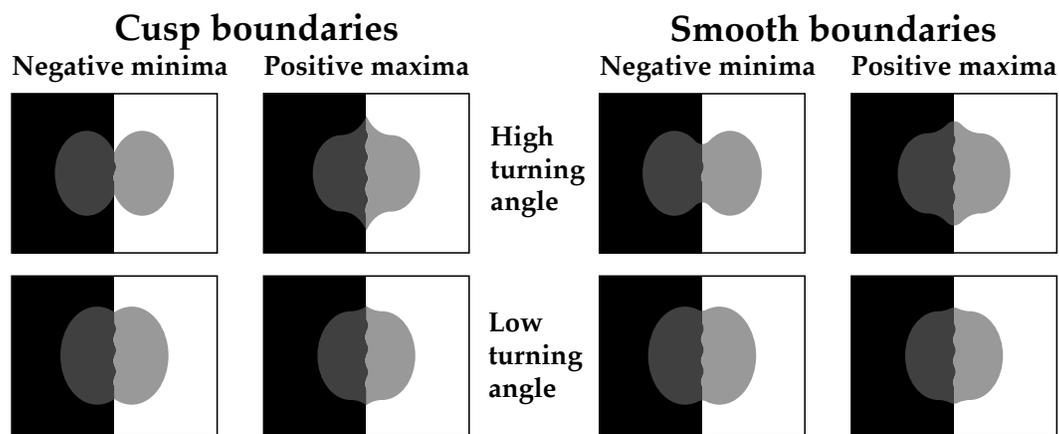


Figure 48. The stimuli used in Singh & Hoffman’s (1998) Experiment 1 to test the role of perceived part structure on perceptual transparency.

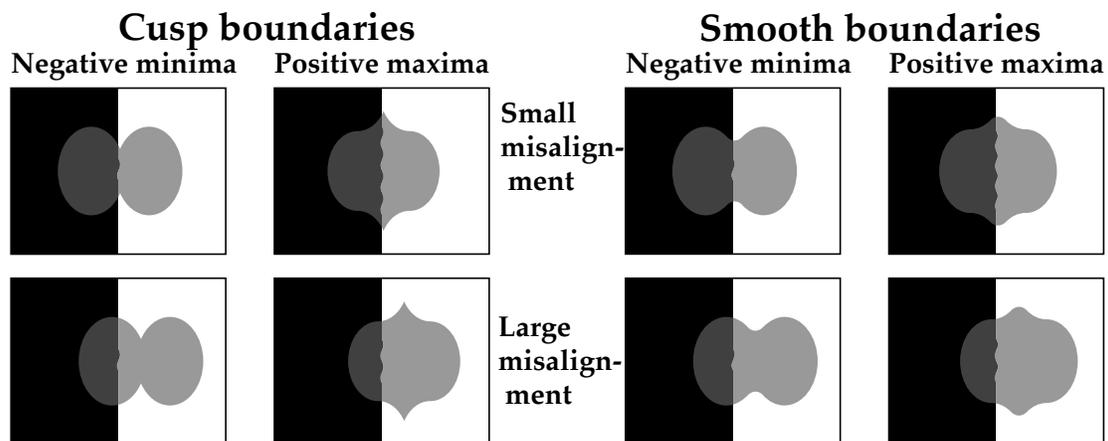


Figure 49. The stimuli used in Singh & Hoffman’s (1998) Experiment 2 to study the role of the precise alignment of the extrema of curvature with the bipartite background.

We thus have two different accounts for the loss of transparency in Figures 47c and 47d—one based on the principle of genericity, and the other on the part-coloring rule. The

relative contributions of these two explanations can be separated using a display like Figure 47e. In this display, we have tangent discontinuities of the same magnitude as in Figure 47d, and these are also precisely aligned with the luminance boundary—hence leading to a non-generic interpretation of transparency. However, these tangent discontinuities are not part boundaries because they lie in convex, rather than concave, regions. Therefore, whereas the part-coloring rule distinguishes between Figure 47d and Figure 47e—predicting a greater loss of transparency in Figure 47d—the genericity principle does not.

Singh and Hoffman (1998) tested the predictions of the part-coloring rule and genericity rule in an experiment in which observers rated the perceived transparency of displays like the ones shown in Figures 48. Using the notion of part salience reviewed in the previous section, Singh and Hoffman designed the figures to differ in three independent dimensions: the sign of curvature at the extrema (convex or concave), the turning angle at the extrema (42° , 72° , 102° , 132°), and the degree of smoothing of the extrema (corner, small smoothing, large smoothing). They compared the ratings of transparency for these figures to the ratings of transparency for a figure whose central gray region had no curvature extrema at all (i.e., a disk). They found that ratings for the positive-maxima cases (convex corners and smoothed convex corners) were

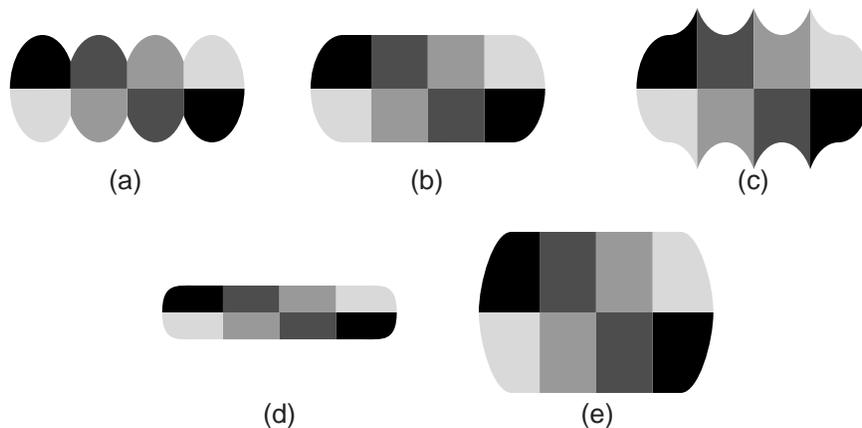


Figure 50. *Another demonstration of the interaction between the sideways segmentation of a shape into parts, and the depth-wise segmentation into multiple layers involved in transparency. Relative to (b) and (c), the display in (a) is least likely to be seen as containing transparency. The displays in (d) and (e) demonstrate that this effect cannot be attributed simply to differences in the contour length of the putative transparent filter.*

significantly lower than ratings for the disk, supporting the prediction of the genericity rule. Moreover they found that ratings for the negative-minima cases (concave corners and smoothed concave corners) were significantly lower than those for the positive maxima—indicating that the part-coloring rule operates in addition to the genericity rule. They also found that ratings

decreased with increasing salience of the extrema, supporting the predictions of the salience theory of the previous section.

The genericity and part-coloring rules also predict that if we misalign the extrema from the color changes in the image, then ratings of transparency should go up. This was tested in a second experiment by Singh and Hoffman (1998), using stimuli such as those shown in Figure 49. They found that for both convex and concave extrema, ratings of transparency increased as the misalignment between the extrema and color changes increased, supporting the predictions of both the genericity and part-coloring rules.

Hence the “sideways” parsing of an image region into parts interacts with the parsing in depth that leads to the perception of transparency, via the part-coloring rule: If the change in luminance in an image region coincides with the boundary between two distinct parts, the visual system tends to assign this to a reflectance change, thereby allowing the two parts to have distinct colors. This, in turn, leads to a suppression in the percept of transparency. The displays in Figure 50 again demonstrate this rule at work: Whereas Figures 50b–c are easily seen as containing transparency, Fig 50a is not. (Figures 50d–e demonstrate that this difference is not simply due to differences in the contour length of the putative transparent filter.)

Parts and Pre-attentive Processing

We reviewed experiments earlier, demonstrating that a target shape having a negative minimum of curvature pops out among distractors that do not have a negative minimum—but not vice versa (recall Figs 14 and 15; Wolfe & Bennett, 1997; Hulleman et al., 2000). Such an asymmetry has traditionally been interpreted to be a sign of pre-attentive processing, or of a basic feature (Treisman & Souther, 1985; Treisman & Gormican, 1988; Duncan & Humphreys, 1992; Wolfe & Bennett, 1997; Wolfe, 1998). Although more recent work has disputed whether *any* visual search can really be considered to be free from attentional demands (see, e.g., Joseph, Chun, & Nakayama, 1996), it is nevertheless true that pop-out and search asymmetry are indicative of features that are computed rapidly and early in visual processing.

In particular, the above experiments suggest that negative minima (and, more generally, perceptually salient concavities) are computed pre-attentively, or at least rapidly and early in visual processing. They leave open the question, however, of whether part structures themselves

are computed early. As we have seen, on a 2D shape, negative minima are points on the bounding contour of the shape—whereas determining part structures requires computing *part cuts* that divide the shape into parts. Some evidence for the early computation of parts comes from Baylis & Driver's (1994; 1995) experiments on judging symmetry and repetition (recall Figs 12 and 13). Note, however, that because this experimental method can only use shapes that are near-

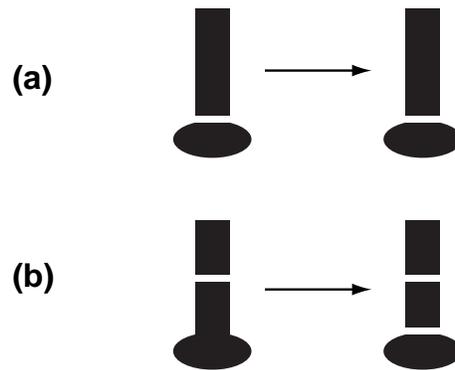


Figure 51. *One of the shapes used by Xu & Singh (2001) to test whether part structures consistent with the minima rule are computed pre-attentively. The targets and distractors for visual search consisted of two different parsings of the same shape—one had cuts at negative minima (a), the other had cuts elsewhere (b). If parsing at negative minima occurs automatically, the stimulus in (b) should be parsed further, hence yielding an additional feature that can guide visual search.*

symmetric or near-repeated, it is unclear how the results generalize to arbitrary shapes. Furthermore, this method does not allow for a comparison of the perceptual naturalness of alternative part structures. It does not, in fact, uniquely point to negative minima of curvature as the critical loci of part boundaries. For example, the shapes in Fig 12a and Fig 13b have not only matching negative minima on the two sides, they also have matching positive maxima and matching inflection points. Hence, the advantage enjoyed by these shapes could, in principle, be due to the matching positive maxima or inflections as well.

To get at the issue of early processing of part structure, Xu & Singh (2001) used the visual search paradigm. However, unlike previous visual search studies—which have used different shapes for targets and distractors—they used the same shape, but parsed in two different ways. One of these parsings was always consistent with the minima rule, the other not (an example of a shape used is shown in Fig 51). Xu & Singh argued that, if parsing at negative minima occurs automatically, then the bottom “part” in Fig 51b would be parsed further, yielding the effective parsing shown on the right. As a result, the non-minima-parsed shape (Fig 51b) would end up having an additional feature that the minima-parsed shape (Fig 51a) does not have. Hence, the search for the non-minima-parsed shape should be fast and efficient, whereas the

reverse search for the minima-parsed shape should be slow and inefficient (defined, as it is, by the *absence* of a feature). Note that the lengths of the cuts, the orientations of the cuts, and the areas of the parts created are all equated in the two parsings, so any search asymmetry cannot be attributed to these factors.

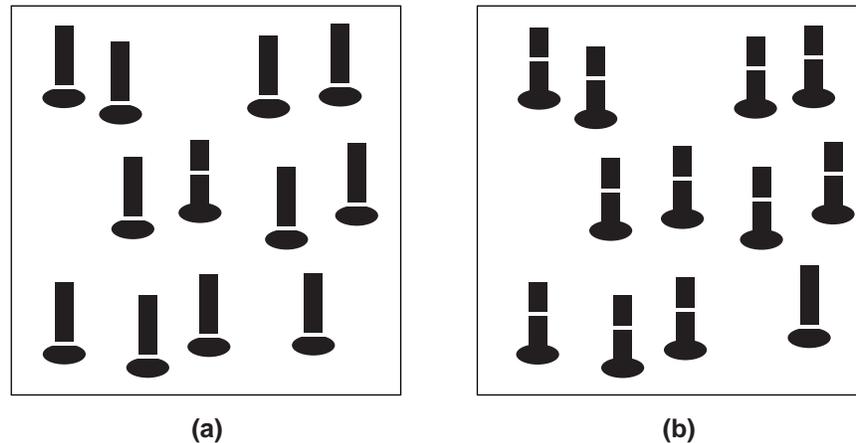


Figure 52. Stimuli from Xu & Singh's (2001) experiment. Subjects search for either (a) a target parsed with non-minima cuts among distractors parsed with minima cuts, or (b) vice versa. As predicted, the first search was a pop out, whereas the second was slow and inefficient.

This is exactly what their data revealed. When subjects searched for a non-minima-parsed target among minima-parsed distractors (see Fig 52a), search slopes satisfied all criteria for a pop-out (or feature) search (search slopes were 5.5 ms/item for target-present trials, and 6.7 ms/item for target-absent trials). This indicates that when a part cut appeared at a non-negative-minima location on the target shape, it was considered as a unique feature among the cuts that occurred at negative minima on the distractor shapes. However, when subjects searched for a minima-parsed target among non-minima-parsed distractors (see Fig 52b), the search became slow and inefficient (slopes were 19.3 ms/item for target present, and 32.2 ms/item for target absent). Thus, the automatic parsing at negative minima in the distractor shapes made the negative-minima cut in the target shape much less prominent.

To provide further evidence that parsing at negative minima occurs early in visual processing, Xu & Singh (2001) conducted another experiment in which subjects searched for parsed shapes among unparsed shapes (Fig 53). In different blocks, subjects searched either for a minima-parsed shape or a non-minima-parsed shape. Since the target can be distinguished from the distractors by the presence of a cut, quite independent of its location, the parsed shape was expected to pop out in both cases. Xu & Singh argued, nevertheless, that if parsing at negative minima occurs automatically in early stages of visual processing, a cut located at negative minima

would still be less efficacious as a feature that distinguishes the target from the unparsed distractors. Indeed, their results showed that although the search for a parsed shape among unparsed shapes was fast and efficient in both cases, the search for the minima-parsed shape was nevertheless significantly slower. This result was most striking for target-absent displays. In these cases, both search displays were identical (since the parsed target shape was absent in both cases). Nevertheless, it took subjects 100 ms longer to determine the absence of a cut when they were looking for it at negative minima, than when they were looking for it elsewhere. The fact that parsing at negative minima influenced visual search even when the search was already fast and efficient thus provides further evidence that such parsing must occur very early in visual processing.

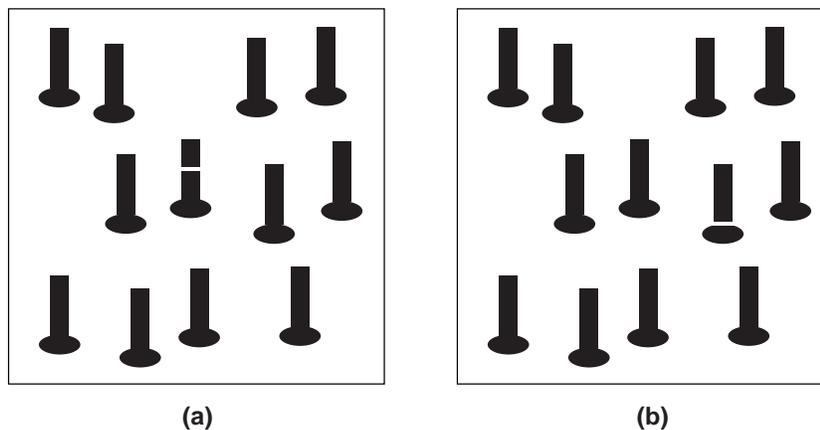


Figure 53. *Stimuli from a subsequent experiment by Xu & Singh (2001). Subjects searched either for (a) a non-minima parsing, or (b) a minima parsing among unparsed shapes. As expected, both searches were pop out. (In both cases, the cut on the target shape acts as unique feature.) Nevertheless, the search for the minima-parsed target was slower by about 100ms.*

Beyond these specific findings, Xu and Singh's experiments also provide a general methodology for studying how the visual system parses shapes into parts. In particular, to compare the perceptual naturalness of two different ways of parsing a given shape, one conducts visual search experiments using one parsing as the target and the other as distractors; and then reverses their roles. Any asymmetry between these two searches can be interpreted as a difference in the perceptual naturalness of the two parsings—assuming that other local factors have been controlled for.

Parts and Attentive Processing

A great deal of recent work has indicated that it is insufficient to characterize visual attention in purely spatial terms, for example, as a ‘spotlight’ that can move around the visual field, enhancing visual processing within the spatial region ‘lit’ by the spotlight (e.g., Posner, Nissen, & Ogden, 1978). A number of experiments have demonstrated that visual attention can be object-based as well. In other words, attention shifts more readily within a single object than across separate objects—so that attending to one location on an object tends to enhance processing

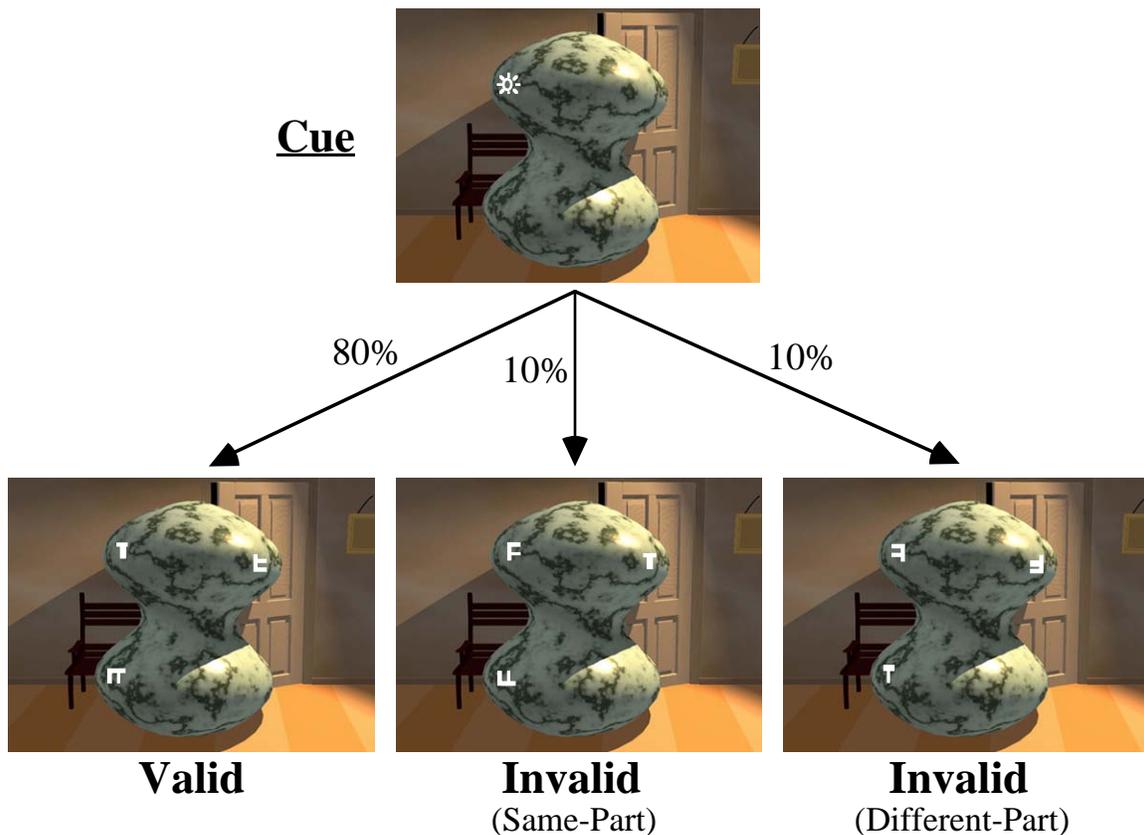


Figure 54. A schematic of the trial sequence used in Singh & Scholl’s (2000) attentional-cueing experiment. (The cue and probes are not to scale.) The experiment demonstrated that the allocation of visual attention to an object can be part-based.

resources for the entire object. One source of evidence for object-based attention comes from the attentional-cueing paradigm. Egly, Driver, & Rafal (1994), for example, presented participants with a display containing two vertically elongated rectangles. They cued one end of a rectangle, and asked participants to make speeded responses to the (possible) appearance of a probe. On most of the trials, the probe appeared at the cued location. However, on a critical subset of trials,

the probe appeared elsewhere. Amongst these “invalid-cue trials,” participants were faster at detecting probes that appeared on the opposite end of the same rectangle, than those that appeared on the same end of the other rectangle. Since both of these locations were equidistant from the (invalid) cue, this response-time advantage for the same-object trials indicated that visual attention shifts more readily within a single object, than across two different objects.



Figure 55. *The three shapes used in Singh & Scholl’s attentional-cueing experiment. These differ in the perceptual saliency of the part structure. Part saliency was found to modulate the magnitude of the part-based cueing effect.*

Moore, Yantis, & Vaughan (1998) extended this finding to amodally completed objects, using a discrimination version of the attentional-cueing paradigm. They placed a third rectangle that partially occluded the middle portions of the two elongated rectangles (so that the two rectangular objects could only be defined after amodal completion had taken place) and asked participants to make speeded responses to whether the probe was a **T** or an **L**. They again found a response-time advantage for same-object trials relative to different-object trials.

We have argued that ‘perceptual units’ exist not only at the level of objects, but also at the level of parts within an object. If object parts do indeed constitute perceptually natural units in visual representation, might we not expect a part-based advantage, analogous to the object-based advantage revealed by attentional cueing? To address this question, Singh & Scholl (2000) used (the discrimination version of) the attentional-cueing paradigm, on 3D rendered objects that had two parts (see Fig 54). On each trial, one of four locations was cued, followed by three characters: a probe (either **T** or **L**) and two distracters (rotated **F**'s). On 80% of the trials, the cue was valid (i.e., it correctly predicted the probe location), and on 20% of the trials, the cue was invalid (i.e., the probe did not appear at the cued location). On half of the invalid-cue trials, the probe appeared on the same part as the cue; on the other half, the probe appeared on the other part. Since the same-part and different-part probe locations were equidistant from the (invalid) cue, any systematic difference in response-time between these two types of trials would provide evidence for a part-based effect.

Singh & Scholl (2000) performed this attentional-cueing experiment with three different objects. All were two-part objects, but they differed in the perceptual salience of their parts. In one case (the “high-salience” object), the boundary between the two parts was a crease (i.e., surface tangent discontinuity), in the other two cases (the “mid-salience” and “low-salience” objects), the part boundary was smooth, and the turning angle at the boundary was shallower (see Fig 55). They found that, for all three salience levels, the mean RT's were significantly faster for same-part trials than for different-part trials. Moreover, the magnitude of this part-based effect was modulated by the salience of the parts, so that the effect size for “low-salience” object (30.6 ms) was less than half that for the “high-salience” object (65.8 ms).¹³

These results provide evidence for part-based attention: Visual attention shifts more readily within a part, than across parts of a single object. Moreover, geometric factors that determine perceived part salience modulate the strength of this part-based cueing effect. Finally, these results indicate that attentional cueing can be used as a tool—providing an objective performance measure—to study perceptual part structure.

Similar results have also been obtained on 2D shapes using a different experimental paradigm in which participants are asked to report, or to compare, two attributes of either the same part or of two different parts of an object. Participants are faster (Barenholtz & Feldman, 2001) or more accurate (Vecera, Behrmann, & McGoldrick, 2000; Vecera, Behrmann, & Filapek, 2001) when the two attributes are on the same part, than when they are on two different parts. In addition, the experiments of Barenholtz & Feldman (2001) demonstrate that this part-based advantage cannot be attributed simply to the magnitude of the curvature of the object contour that intervenes between the two features to be compared. In particular, the effect is greatly reduced if the two features are separated by a convex contour (i.e., no part boundary) of the same curvature magnitude as a concave contour (i.e., with a part boundary).

Parts and Illusory Contours

Consider the displays in Figure 56a–c. Each of these displays could be seen, in principle, either as the projection of a single object with multiple parts, or as the projection of two distinct objects—one partly occluding the other. Surprisingly, even though these displays are

¹³ Note that the magnitude of the part-based effect for the high-salience object is larger than the object-based effect obtained by Egly et al. (1994) and Moore et al., (1998). It has been shown recently that that 3D rendered shapes yield

chromatically homogeneous, and therefore lack any T-junctions to signal occlusion, each of 56a–c is seen as *two* overlapping objects. Because of the tendency of such displays to *split* perceptually into two distinct objects, they have been termed “spontaneously-splitting figures” or “self-splitting objects” (Kanizsa, 1979; Kellman & Loukides, 1987). Constructing this percept of two overlapping objects requires that one of the two objects (the occluding object) be modally completed in front, and the other (occluded object) be amodally completed behind it.

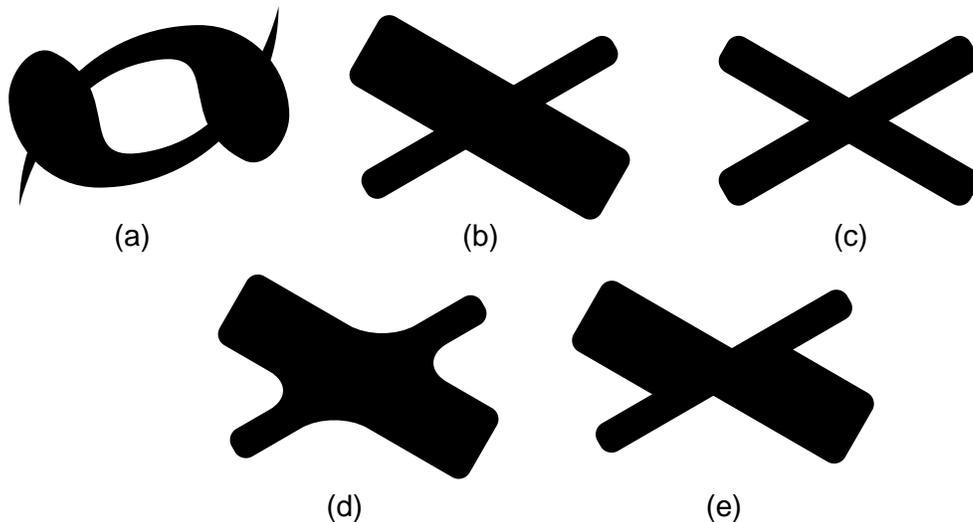


Figure 56. (a)—(c) tend to be seen as two overlapping objects— even though these figures lack any T-junctions to signal occlusion. Smoothing the negative minima (d), or disrupting the good-continuation between one set of contours (e), switches the percept to that of a single object with multiple parts.

When does the visual system split a single figure of homogenous color into two distinct objects? And, how does it assign relative depth ordering to these objects? It is clear that certain image properties *must* be satisfied in order for the perceptual split to occur. For example, if the negative minima are smooth (Fig 56d), rather than tangent discontinuities, the display is seen as a single object with parts, rather than two objects separated in depth. This makes sense because the display is no longer consistent with a generic view of an object interposed between the viewer and a second object (Shipley & Kellman, 1992). Similarly, if the good continuation (or relatability) between one of the two sets of contours is disrupted, the display is again seen as a single object (Fig 56e).

Concerning the second question, note that the figure (or portion of the figure) that requires the shorter contour completions is the one that is perceived to be modally completed in

larger effects (Atchley & Kramer, 2000). Indeed, this was a primary motivation for using 3D stimuli—since starting with an object-based effect of ~30-40ms would not leave much room to study the effect of parts, and part salience.

front (see Figs 56a–b). For the cross in Fig 56b, for example, observers are more likely to see the thicker bar as being in front of, and occluding, the thinner bar. For the cross in Fig 56c, on the other hand, both bars require equally long completions and this leads to a perceptual bistability: One switches between seeing one bar in front, and seeing the other one in front.

The tendency of human vision to see modal completions as short as possible is known as Petter's rule (Petter, 1956). This rule has often been motivated by the heuristic that closer objects tend to project larger retinal images (e.g., Stoner & Albright, 1993). However, Petter's rule cannot be explained by a “larger is closer” heuristic because relative lengths of interpolated contours can vary independently of relative areas of the corresponding figures (Tommasi, Bressan, & Vallortigara, 1995). Moreover, Petter's rule works even when it contradicts the “larger is closer” hypothesis (Singh, Hoffman, & Albert, 1999).

However, given that it is initially ambiguous whether a display such as Fig 56b is the silhouette of a single object with multiple parts or of two objects separated in depth, it is likely that the visual system applies the same segmentation rules in both cases (Singh, Seyranian, & Hoffman, 1999). In particular, when presented with such a silhouette, the visual system finds the negative minima of curvature and pairs them according to the short-cut rule. Depending on whether the silhouette is interpreted to be a single object (Fig 56d–e) or multiple objects (Fig 56a–c), these pairings are then taken to be either part cuts or modal completions. In this way, Petter's rule can be understood in terms of the same visual mechanisms that create parts cuts—and it thus inherits the ecological motivation for the short-cut rule.

AN ALTERNATIVE APPROACH TO PARTS AND SHAPE DESCRIPTION: CURVE EVOLUTION AND SHOCKS

An alternative approach to parts and shape description has been proposed by Kimia and colleagues (e.g., Kimia, Tannenbaum, & Zucker, 1995). Although we have previously mentioned some shortcomings in the definitions of limbs and necks as proposed by Siddiqi & Kimia (1995) and Siddiqi, Tressness, & Kimia (1996), there is actually a more general framework on which these intuitions are based. In this section, we briefly review this framework and contrast it with our geometric-constraints approach to parts.

This framework involves deriving a structural description of 2D shapes from the singularities (or “shocks”) of a curve evolution process in which every point on the shape’s contour moves inward at constant speed. This process is identical to Blum’s grassfire (SAT) transformation, and the final skeletal structure obtained is the same as Blum’s medial axis.¹⁴ However, in addition to storing the final skeletal structure, Kimia et al. also classify the singularities that are formed, and store the order of shock, as well as the location and time of its formation. These together are taken as a representation of the 2D shape.

Figure 57 shows examples of different orders of shocks. In particular, at loci of 1-shocks, the radius around the axis varies monotonically (a “protrusion”), whereas 3-shocks occur in

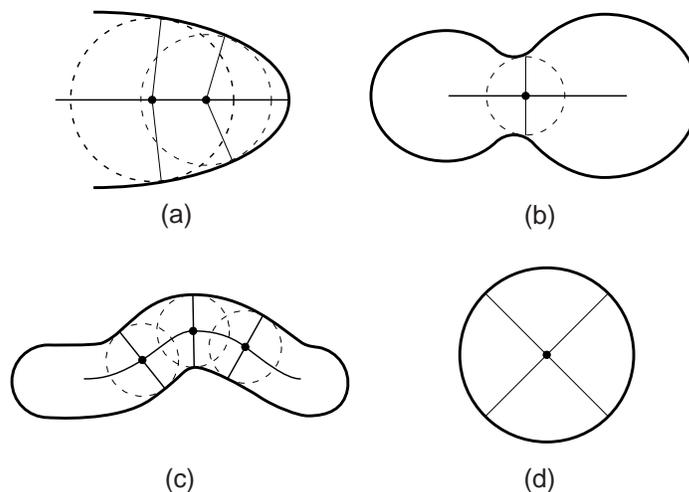


Figure 57. *Different orders of “shocks” obtained from the singularities of a curve evolution process. (Adapted from Kimia, Tannenbaum, & Zucker, 1995).*

regions where the radius function is constant, i.e., the contours on the two sides are parallel (“bends”). 2-shocks correspond to local minima of the radius function (“necks”), and 4-shocks corresponds to local centers of mass (“seeds”).

Critical to the context of parsing, however, note that the notion of “protrusion” as a connected locus of 1-shocks does not distinguish between a part such as the one shown in Figure 58a, and a convex region that ends in a positive maximum of curvature, as shown in Figure 58b. (This is reminiscent of the shortcoming of Blum’s proposal that shapes should be parsed at the branch points of the axes of local symmetry. Recall that the medial axis gives the same skeletal

¹⁴ Kimia et al. study a more general class of shape deformations, which are describable in terms of two basic kinds: (1) deformations in which a curve evolves at a constant speed in a direction normal to the curve (“reaction”), and (2) deformations in which the speed of evolution is proportional to the curvature at a given point (“diffusion”). However, for the purpose of computing shocks and shock-graphs, they use only the first kind.

structure to a rectangle and a dog-bone shape, and is therefore unable to find parts solely on this basis; see Figure 31). Making the distinction between parts and convex tips requires the notion of a negative minimum of curvature—so that the shock-based approach also needs to incorporate negative minima to fully capture part structure. Indeed, it is likely that this consideration is what motivated Siddiqi & Kimia (1995) to define a notion of “limb” using further geometric criteria—which, as we noted earlier, turn out to be too restrictive.

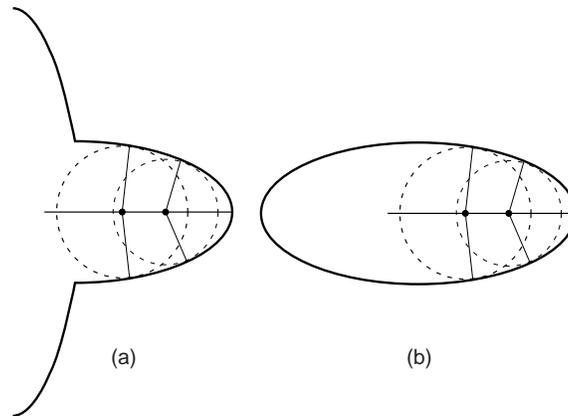


Figure 58. *The notion of “protrusion” as a connected locus of first-order shocks does not distinguish between a part (a) and a convex region ending in a positive maximum (b). Making this distinction requires the notion of a negative minimum of curvature.*

Based on the geometry of shocks, Kimia et al. proposed a “shape triangle” which contains three different continua along which shapes can be defined: a parts-protrusions continuum, a bends-parts continuum, and a bends-protrusions continuum (see Figure 59). For example, along the parts-protrusions continuum, the ratio of minimum (“neck”) width to maximal width is gradually increased so that the shape looks more like a vertical body with small “protrusions” on the sides, rather than parts stacked vertically, one on top of the other. Within our framework, protrusions and bends are different varieties of parts with varying perceptual salience and/or involving different ways of pairing part boundaries (i.e., neg min to neg min, or neg min to pos max). For example, in moving from “parts” to “protrusions,” the length of the horizontal cuts increases relative to the vertical ones, and as a result, the vertical cuts become more likely to be perceived. Furthermore, the salience (e.g., the strength of the part boundaries and protrusion) of the parts gradually decreases. Hence, one moves from a part structure consisting of salient parts stacked on top of each other to one consisting of a main vertical body with low-salience parts on the sides. Similarly, along the parts-bends continuum, the local symmetry between pairs of negative minima is gradually offset, so that the negative minima on one side of the curve eventually become locally symmetric to the positive maxima on the other side, rather than the

negative minima. Hence, using the two variables of cut length and local symmetry, one can span the space defined by the three continua that define the shape triangle (see Figure 60a).

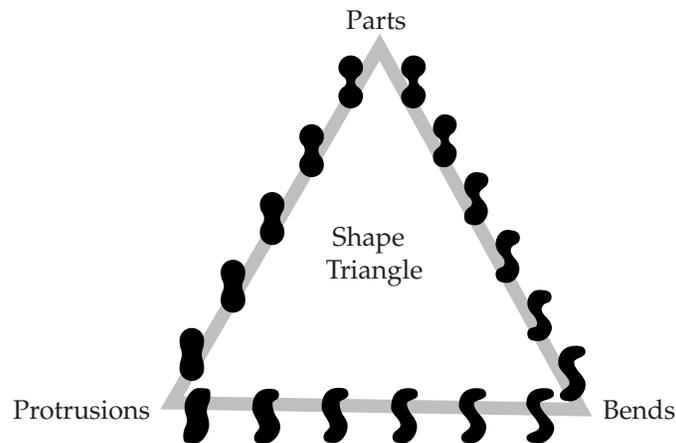


Figure 59. *Kimia et al.’s shape triangle embodies three different continua along which a shape can vary. These continua are also captured by the geometric constraints of cut length, local symmetry, and part-boundary strength.*

Note, moreover, that moving along the part-protrusion continuum, and the bend-protrusion continuum involves a change in another variable that is not made explicit within the shape triangle, namely, the curvature at the extrema of curvature (see Siddiqi, Kimia, Tannenbaum, & Zucker; Figure 6). Within our geometric-constraints framework, the strength of curvature extrema is a further variable that is orthogonal to cut length and local symmetry (recall the section on part salience). Figure 60b, for example, shows sharp-extrema versions of the shapes shown in Figure 60a (in the two-parameter space spanned by local symmetry and cut length). Note that although low-curvature “bends” have the appearance of a single-part object, “bends” with sharp corners are more likely to appear as distinct parts.

Thus, within the geometric-constraints approach, cut length, local symmetry and the strength of part boundaries interact to determine perceived parts and their visual salience: Human vision prefers shorter cuts, cuts that join locally-symmetric negative minima, and cuts that join points of high magnitude of curvature. These geometric constraints capture the important aspects of perceived part structure. We argue that the visual system employs simple geometric constraints such as these, rather than compute detailed shock graphs. In particular, visual search tasks that require distinguishing “parts,” “protrusions,” and “bends” (e.g., Siddiqi et al., 2000) can be performed on the basis of the visual system’s sensitivity to simple geometric properties such as cut length, strength of curvature and local symmetry—rather than on specific differences in computed shock graphs.

From the point of view of shape description, the shock-based approach does represent a great deal more information than the geometric-constraints approach—which provides simply the locations of the part boundaries and part cuts on a shape, and the relative visual salience of the resulting parts. From the point of view of parsing, however, we have noted that loci of first-order shocks do not distinguish between convex tips and parts (recall Figures 58 and 31)—and thus need to incorporate the notion of negative minima in order to make this distinction. Moreover,

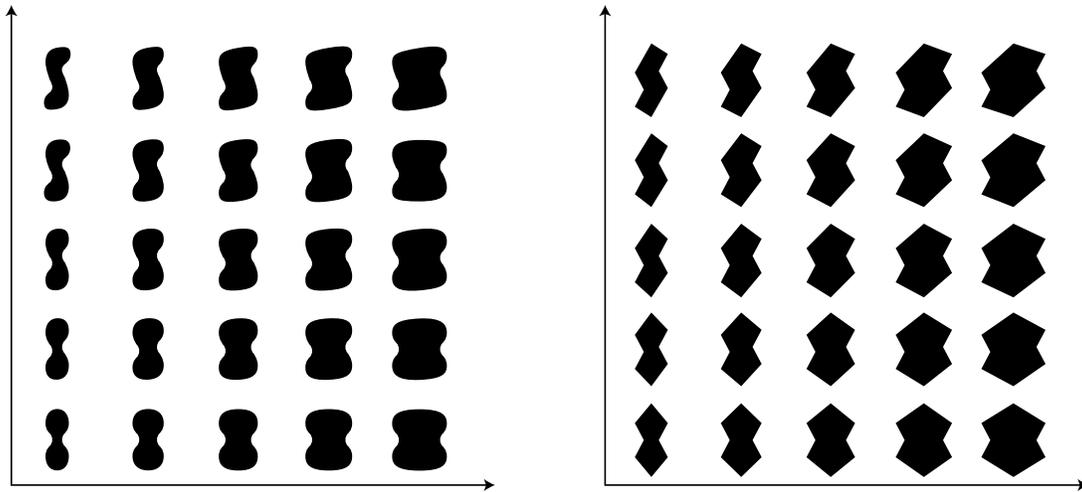


Figure 60. Each diagram represents a two-parameter space involving changes in cut length and local symmetry. The two diagrams portray different levels of strength of curvature extrema.

the number of possible shock types undergoes a combinatorial explosion with the addition of the third dimension—so that basing 3D-shape description on a classification of shock types becomes very difficult. The geometric-constraints approach, on the other hand, generalizes easily and naturally to three-dimensional surfaces and volumes, because the notions of negative minima, strength of curvature extrema, cut length, and local symmetry all have natural counterparts in 3D.

CONCLUSIONS

We have reviewed evidence that the human visual representation of shape is part-based, and we have discussed several geometric factors that determine how human vision parses shapes into parts, and determines their perceived salience. Moreover, we have shown that mechanisms of part segmentation interact with mechanisms of object and surface segmentation in a number of ways: in assigning border ownership, in computing surface structure in transparency, and in

dividing spontaneously-splitting figures. Finally, we have reviewed evidence suggesting that part segmentation is carried out rapidly and early in visual processing, and that the allocation of visual attention to objects can be part-based.

REFERENCES

- Atchley, P. & Kramer, A. (2000). Object and space-based attentional selection in three-dimensional space. *Visual Cognition*, 8, 1–32.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61, 183-193.
- Attneave, F. (1971). Multistability in perception. *Scientific American*, 225, 6, 63–71.
- Anderson, B.L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: the role of junctions. *Perception*, 26, 49-453.
- Bahnsen, P. (1928). Eine untersuchung ueber symmetrie und asymmetrie bei visuellen wahrnehmungen. *Zeitschrift fuer Psychologie*, 108, 355-361.
- Barenholtz, E., & Feldman, J. (2001). Interpretation of part boundaries and the movement of attention. Poster presented at the 1st Annual Meeting of the Vision Sciences Society, Sarasota, Florida (May 2001).
- Baylis, G.C., & Driver, J. (1994). Parallel computation of symmetry but not repetition in single visual objects. *Visual Cognition*, 1, 377-400.
- Baylis, G.C., & Driver, J. (1995a). One-sided edge assignment in vision. 1. Figure-ground segmentation and attention to objects. *Current Directions in Psychological Science*, 4, 140-146.
- Baylis, G.C., & Driver, J. (1995b). Obligatory edge assignment in vision - The role of figure and part segmentation in symmetry detection. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 1323-1342.
- Beck, J., Prazdny, K., & Ivry, R. (1984). The perception of transparency with achromatic colors. *Perception & Psychophysics*, 35, 407-422.
- Bennett, B.M. & Hoffman, D.D. 1987. Shape decompositions for visual recognition: the role of transversality. IN: W. Richards and S. Ullman (Eds.), *Image Understanding 1985-86*, Norwood, New Jersey: Ablex Publishing, 215-256.
- Beusmans, J., Hoffman, D.D., & Bennett, B.M. (1987). Description of solid shape and its inference from occluding contours. *Journal of the Optical Society of America, A*, 4, 1155-1167.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Biederman, I., & Cooper, E.E. (1991). Priming contour-deleted images: evidence for intermediate representations in visual object recognition. *Cognitive Psychology*, 23, 393-419.

- Binford, T.O. (1971, December). Visual perception by computer. *IEEE Systems Science and Cybernetics Conference*, Miami, FL.
- Binford, T. O. (1981). Inferring surfaces from images. *Artificial Intelligence*, *17*, 205-244.
- Blum, H. & Nagel, R.N. (1978). Shape description using weighted symmetric axis features. *Pattern Recognition*, *10*, 167-180.
- Brady, M., & Asada, H. (1984). Smoothed local symmetries and their implementation. *The International Journal of Robotics Research*, *3*, 36-61.
- Braunstein, M.L., Hoffman, D.D., & Saidpour, A. (1989). Parts of visual objects: an experimental test of the minima rule. *Perception*, *18*, 817-826.
- Bruce, V. and Morgan, M. J. (1975). Violations of symmetry and repetitions in visual patterns. *Perception*, *4*, 239-249.
- Casati, R. & Varzi, A. (1999). *Parts and places*. Cambridge, MA, MIT Press.
- Corballis, M. C. & Roldan, C. E. (1974). On the perception of symmetrical and repeated patterns. *Perception and Psychophysics*, *16*, 136-142.
- Do Carmo, M. (1976). *Differential geometry of curves and surfaces*. Englewood Cliffs, NJ, Prentice-Hall.
- Driver, J., & Baylis, G.C. (1995). One-sided edge assignment in vision. 2. Part decomposition, shape description, and attention to objects. *Current Directions in Psychological Science*, *4*, 201-206.
- Egly, R., Driver, J., & Rafal, R. (1994). Shifting visual attention between objects and locations: Evidence for normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, *123*, 161-177.
- Freeman, W. T. (1994). The generic viewpoint assumption in a framework for visual perception. *Nature*, *368*, 542-545.
- Gerbino, W., Stultiens, C. I., Troost, J. M., & de Weert, C. M. (1990). Transparent layer constancy. *Journal of Experimental Psychology: Human Perception & Performance*, *16*, 3-20.
- Hochberg, J. (1964). *Perception*. Englewood Cliffs, NJ: Prentice-Hall.
- Hoffman, D.D. (1983a). *Representing shapes for visual recognition*. PhD Thesis, Massachusetts Institute of Technology.
- Hoffman, D.D. (1983b). The interpretation of visual illusions. *Scientific American*, *249*, 6, 154-162.
- Hoffman, D. D. (1998). *Visual Intelligence: How we create what we see*. New York: Norton.
- Hoffman, D.D., & Richards, W.A. (1984). Parts of recognition. *Cognition*, *18*, 65-96.

- Hoffman, D. D. & Singh, M. (1997). Saliency of visual parts. *Cognition*, 63, 29-78.
- Hulleman, J., te Winkel, W., & Boselie, F. (2000). Concavities as basic features in visual search: Evidence from search asymmetries. *Perception & Psychophysics*, 62, 162-174.
- Joseph, J. S., Chun, M. M., & Nakayama, K. (1996). Attentional requirements in a “preattentive” feature search task. *Nature*, 379, 805–807.
- Kanizsa, G. (1979). Organization in vision: essays on Gestalt perception. New York: Praeger.
- Kanisza, G., & Gerbino, W. (1976). Convexity and symmetry in figure-ground organization. In M. Henle (Ed.), *Art and artefacts* (pp. 25-32). New York: Springer.
- Kellman, P. J. & Loukides, (1987). An object perception approach to static and kinetic subjective contours. IN: S. Petry & G. E. Meyer (Eds.), *The perception of illusory contours* (pp. 151–164). New York: Springer-Verlag.
- Kellman, P. J. & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, 23, 141-221.
- Kimia, B. B., Tannenbaum, A. R. & Zucker, S. W. (1995). Shapes, shocks, and deformations. 1. The components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision*, 15, 189-224.
- Koffka, K. (1935). *Principles of gestalt psychology*. New York: Harcourt, Brace and World.
- Koenderink, J. J., Van Doorn, A. J., & Kappers, A. M. (1992). Surface perception in pictures. *Perception & Psychophysics*, 52, 487–496.
- Kubovy, M., & Gepstein, S. (2000). From good continuation to best continuation. Talk given at the 41st Annual Meeting of the Psychonomic Society, New Orleans, LA.
- Leyton, M. (1987). Symmetry-curvature duality. *Computer vision, graphics, and image processing*, 38, 327–341.
- Leyton, M. (1992). *Symmetry, causality, mind*. Cambridge, MA: MIT Press.
- Lowe, D. (1985). Perceptual organization and visual recognition. Amsterdam: Kluwer.
- Marr, D. (1977). Analysis of occluding contour. *Proceedings of the Royal Society of London, B*, 197, 441–475.
- Marr, D. & Nishihara, H.K. (1978). Representation and recognition of three-dimensional shapes. *Proceedings of the Royal Society of London, B*, 200, 269-294.
- Metelli, F. (1974). The perception of transparency. *Scientific American*, 230, 90-98.
- Metelli, F., Da Pos, O., & Cavedon, A. (1985). Balanced and unbalanced, complete and partial transparency. *Perception & Psychophysics*, 38, 354-366.
- Moore, C., Yantis, S., & Vaughan, B. (1998). Object-based visual selection: Evidence from perceptual completion. *Psychological Science*, 9, 104–110.

- Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, 9, 441-474.
- Palmer, S. E. (1999). *Vision Science: From photons to phenomenology*. Cambridge, MA: MIT Press.
- Pentland, A.P. (1986). Perceptual organization and the representation of natural form. *Artificial Intelligence*, 28, 293-331.
- Petter, G. (1956). Nuove ricerche sperimentali sulla totalizzazione percettiva. *Rivista di Psicologia*, 50, 213-227.
- Posner, M. I., Nissen, M. J., & Ogden, W. C. (1978). Attended and unattended processing modes: The role of set for spatial location. IN: H. J. Pick and I. J. Saltzman (Eds.), *Modes of perceiving and processing information* (pp. 137-157). Hillsdale, NJ: Erlbaum.
- Rosin, P. L. (2000). Shape partitioning by convexity. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 30, 202-210.
- Rubin, E. (1958). Figure and ground. In D.C. Beardslee (Ed), *Readings in perception*, Princeton, New Jersey: Van Nostrand, 194-203. (Reprinted from *Visuell wahrgenommene figuren*, 1915, Copenhagen: Gyldenalske Boghandel.)
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, 80, 1-46.
- Shipley, T. F., & Kellman, P. J. (1990). The role of discontinuities in the perception of subjective figures. *Perception & Psychophysics*, 48, 259-270.
- Siddiqi, K. & Kimia, B.B. (1995). Parts of visual form: computational aspects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17, 239-251.
- Siddiqi, K., Kimia, B. B., Tannenbaum, A. R., & Zucker, S. W. (2000). On the Psychophysics of the Shape Triangle. *Vision Research*, In Press.
- Siddiqi, K., Tresness, K.J., & Kimia, B.B. (1996). Parts of visual form — Psychophysical aspects. *Perception*, 25, 399-424.
- Singh, M. & Anderson, B. L. (2001). Toward a perceptual theory of transparency. *Psychological Review*, in press.
- Singh, M. & Hoffman, D. D. (1998). Part boundaries alter the perception of transparency. *Psychological Science*, 9, 370-378.
- Singh, M. & Hoffman, D. D. (1999). Completing visual contours: The relationship between relatability and minimizing inflections. *Perception and Psychophysics*, 61, 943-951.
- Singh, M., Hoffman, D. D., & Albert, M. K. (1999). Contour completion and relative depth: Petter's rule and support ratio. *Psychological Science*, 10, 423-428.
- Singh, M. & Scholl, B. (2000). Using attentional cueing to explore part structure. Poster

- presented at the *Annual Symposium of Object Perception and Memory*, New Orleans, Louisiana (November 2000).
- Singh, M., Seyraninan, G. D., & Hoffman (1996). Cuts for parsing visual shapes. *Institute for Mathematical Behavioral Sciences, Technical Report MBS96-33*.
- Singh, M., Seyraninan, G. D., & Hoffman (1999). Parsing silhouettes: The short-cut rule. *Perception and Psychophysics*, *61*, 636-660.
- Stevens, K.A. & Brookes, A. (1988). The concave cusp as a determiner of figure-ground. *Perception*, *17*, 35-42.
- Stoner, G. R. & Albright, T. D. (1993). Image segmentation cues in motion processing: Implications for modularity vision. *Journal of Cognitive Neuroscience*, *5*, 129-149.
- Sutherland, N.S. (1968). Outlines of a theory of visual pattern recognition in animals and man. *Proceedings of the Royal Society of London, B*, *171*, 297--317.
- Todd, J. T., Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. L. (1996). Effects of changing viewing conditions on the perceived structure of smoothly curved surfaces. *Journal of Experimental Psychology: Human Perception & Performance*, *22*, 695-706.
- Tommasi, L., Bressan, P., & Vallortigara, G. (1995). Solving occlusion indeterminacy in chromatically homogeneous patterns. *Perception*, *24*, 391-403.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, *95*, 15-48.
- Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, *114*, 285-310.
- Vecera, S.P, Behrmann, M., & McGoldrick, J. (2000). Selective attention to the parts of an object. *Psychonomic Bulletin & Review*, *7*, 301-308.
- Vecera, S.P, Behrmann, M., & Filapek, J.C. (2001). Attending to the parts of a single object: Part-based selection limitations. *Perception & Psychophysics*, *63*, 308-321.
- Waltz, D. (1975). Generating semantic descriptions from drawings of scenes with shadows. In P. Winston (Ed.), *The psychology of computer vision* (pp. 19-91). New York: McGraw-Hill.
- Winston, P.A. (1975). Learning structural descriptions from examples. In P.H. Winston (Ed.), *The psychology of computer vision* (pp. 157-209). New York: McGraw-Hill.
- Witkin, A.P., & Tenenbaum, J.M. (1983). On the role of structure in vision. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and machine vision* (pp. 481-543). New York: Academic Press.
- Wolfe, J. M. (1998). What can 1 million trials tell us about visual search? *Psychological*

Science, 9, 33-39.

Wolfe, J. M., Bennett, S. C. (1997). Preattentive object files: Shapeless bundles of basic features. *Vision Research*, 37, 25-43.

Xu, Y., & Singh, M. (2001; Manuscript submitted for publication). Pre-attentive computation of part structure: Evidence from visual search.