

Parsing silhouettes: The short-cut rule

Manish Singh Gregory D. Seyranian Donald D. Hoffman
Department of Cognitive Sciences
University of California, Irvine

Many researchers have proposed that, for the purpose of recognition, human vision parses shapes into component parts. Precisely how is not yet known. The minima rule for silhouettes (Hoffman & Richards, 1984) defines boundary points at which to parse, but does not tell how to use these points to cut silhouettes, and therefore does not tell what the parts are. In this paper, we propose the short-cut rule, which states that, other things being equal, human vision prefers to use the shortest possible cuts to parse silhouettes. We motivate this rule, and the well-known Petters rule for modal completion, by the principle of transversality. We present five psychophysical experiments that test the short-cut rule, show that it successfully predicts part cuts which connect boundary points given by the minima rule, and show that it can also create new boundary points.

Introduction

The ease with which we recognize visual objects is deceptive: chess programs can now compete with chess masters, but no computer-vision system can compete with the vision of a toddler. Object recognition is complex and computationally demanding, and typically uses cues such as shape, color, texture, motion, and context. However, the ease with which we can, in many cases, recognize an object without any cues but shape suggests that shape is a key aspect of recognition. This raises the question: How does human vision represent shape for the purpose of recognition?

One proposal is that human vision uses parts combined into structural representations (Biederman, 1987; 1990; Biederman & Cooper, 1991; Hoffman & Richards, 1984; Marr & Nishihara, 1978; Palmer, 1977). It represents a shape in terms of (i) the shapes of its component parts and (ii) the spatial relationships between these parts.

On this approach, representing a shape involves parsing it into sub-units—not unlike parsing a sentence in natural language—and describing the relationships between these

Manish Singh is now at the Department of Brain and Cognitive Sciences, Perceptual Science Group, E10-120, Massachusetts Institute of Technology, Cambridge, MA 02139.

For discussions and suggestions we thank Marc Albert, Bruce Bennett, Myron Braunstein, Mike D’Zmura, Ki-Ho Jeon, Jin Kim, Jeff Liter, Scott Richman, Asad Saidpour, and Jessica Turner. For assistance in running the experiments we thank My Bui, Rizwana Mela, Elizabeth Sullivan, and Michelle Young. And we thank Stephen E. Palmer, Steven Yantis, and two anonymous reviewers for helpful comments on previous versions of the paper. This research was supported by National Science Foundation Grant DIR-9014278.

Correspondence concerning this article can be addressed to Manish Singh at manish@psyche.mit.edu, Greg Seyranian at gregs@aris.ss.uci.edu, or Donald D. Hoffman at ddhoff@uci.edu.




Figure 1. Part cut on a silhouette. The silhouette in (a) is unambiguously parsed in (b) with a part cut (depicted by a dashed line).




Figure 2. A silhouette with more than one parse. The silhouette in (a) is naturally parsed as in (b) or in (c).

subunits—again, not unlike the syntactical structure of a sentence. And, just as some sentences admit multiple parses, so too do some shapes. For example, compare the shape in Figure 1a (which is nicely parsed by the part cut, indicated by a dashed line, in Figure 1b) with the one in Figure 2a (which can be parsed either by the part cuts shown in Figure 2b, or by the part cuts shown in Figure 2c). The parsing of

shapes, when it occurs, is quick and effortless—perhaps even preattentive (Baylis & Driver, 1994; 1995a; 1995a; Driver & Baylis, 1995; Hoffman & Singh, 1997).

Part-based representations have been studied extensively by psychologists¹ and computer scientists.² Parts, unlike template and Fourier approaches, can deal effectively with occlusion, self-occlusion, and certain types of non-rigidity in which rigid parts move relative to each other (Hoffman & Richards, 1984; Pinker, 1985). They can explain the classic visual phenomenon, first noted by Mach, that symmetry is easier to detect than repetition (Baylis & Driver, 1995b; Driver & Baylis, 1995; Mach, 1885/1959), and the phenomenon discovered by Attneave that a piece of curve looks different depending on which side is taken to be “figure” (Attneave, 1971; Hoffman & Richards, 1984; see Figure 5). They can also alter the perception of transparency (Singh & Hoffman, in press). However, parts may be less important in the visual recognition of faces (Farah, 1996; Tanaka & Farah, 1993; Turk & Pentland, 1991; Yuille, 1991).

A natural question arises: How does human vision parse shapes into parts? Some theorists postulate that there is a set of basic shapes, or primitives, which human vision searches for in images—primitives such as generalized cones (Binford, 1971; Marr, 1977) or geons (Biederman, 1987). According to these theories, human vision parses a shape by finding these primitives in the shape. Hence the primitives are responsible for (i) finding parts, and (ii) describing them.

Other theorists postulate that there are rules, based on geometric properties alone, by which human vision computes the boundaries between parts for any given shape. The minima rule (Hoffman & Richards, 1984) is a step in that direction. For a 2D silhouette, the minima rule provides boundary points on the silhouette outline, through which part cuts must pass (see Figure 3a). And for 3D shapes, the minima rule provides boundary curves on the surface of the shape through which part cuts must pass (see Figure 3b). However, the minima rule does not define the part cuts themselves—it only constrains them by requiring them to pass through the boundary points it provides.

We deal here only with the parsing of silhouettes, not of 3D shapes. The relationship between 3D shapes and silhouettes is complex, and beyond the scope of this paper (but see, e.g., Richards, Koenderink, & Hoffman, 1987). However, human subjects do see parts in silhouettes—so the parsing of silhouettes is of psychological interest.

In this paper, we propose that human vision parses silhouettes according to the short-cut rule.

Short-cut rule: *Divide silhouettes into parts using the shortest possible cuts.*

In other words, if boundary points can be joined in more than one way to parse a silhouette, human vision prefers that parsing which uses the shortest cuts. For the shape in Figure 2a, for instance, the short cut rule gives the cuts shown in Figure 2b rather than the cuts shown in Figure 2c. In this paper a cut is (1) a straight line which (2) crosses an axis of local symmetry (see section entitled Short Cuts), (3) joins two points on the outline of a silhouette, such that (4) at least one




Figure 3. Minima rule gives part boundaries which are (a) points on 2D shapes, and (b) curves on 3D shapes.

of the two points has negative curvature. For some shapes such as elbows, the short-cut rule can create boundary points that are not negative minima of curvature.

We begin by reviewing the minima rule and related research on shape partitioning. We then motivate the short-cut rule by the geometry of transversal intersections, and relate it to Petter’s rule for modal completion (Petter, 1956; Kanizsa, 1979, p. 40). Finally, we present five experiments that test, and support, the short-cut rule.

A few caveats and disclaimers. There is much to the parsing of visual shapes that we cannot explore here. First, top-down factors can influence visual parsing (e.g., Schyns, Goldstone, & Thibaut, 1998; Schyns & Murphy, 1994; Schyns & Rodet, 1997; Singh & Landau, 1998). These are no less important than the short-cut rule that we examine. But to keep this work to manageable size, we focus here on the short-cut rule. Once geometric factors like the short-cut rule are understood we have a better chance to integrate them with the top-down. Second, as mentioned earlier, many factors besides parts influence visual recognition. Motion, texture, color, and surface characteristics can also be important (Biederman, Glass, & Stacy, 1973; Bruce & Humphreys, 1994). These cannot be ignored. Ultimately we want to know how these interact with geometric factors in visual parsing and object recognition. Third, parts appear at many spatial scales. Smaller parts nest within larger ones to form hierarchies. Information at multiple scales affects the cutting of parts. The issue of scale and how it affects part cuts is complex and important. But here we assume that some scale of resolution, and some piecewise continuous approximation of the shapes, has been fixed.

¹ A partial list includes Baylis & Driver (1994; 1995a; 1995b), Bennett and Hoffman (1987), Beusmans, Hoffman, and Bennett (1987), Biederman (1987), Biederman and Cooper (1991), Braustein, Hoffman, and Saidpour (1989), Driver & Baylis (1995), Hoffman ((1983a); (1983b)), Hoffman and Richards (1984), Hoffman and Singh (1997), Marr ((1977); (1982)), Marr & Nishihara, 1978), Palmer, 1977), Pentland, 1986), Stevens & Brookes, 1988), Terzopoulos, Witkin, & Kass (1987), Todd, Koenderink, van Doorn, and Kappers (1995), Tversky and Hemenway (1984).

² A partial list includes Binford (1971, December), Brooks (1981), Dickinson, Pentland, & Rosenfeld (1992), Guzman (1971), Siddiqi and Kimia (1995), Winston (1975).

The Minima Rule

Our starting point is the minima rule (Hoffman & Richards, 1984). In this section we briefly review the minima rule, its philosophy, strengths, and motivations. We then discuss its limitations. These limitations are the point of departure for our work here.

The minima rule embodies a fundamental shift in philosophy toward the problem of object parts. As noted earlier, the dominant philosophy of most researchers has been to specify the shapes that parts may take, and then to look for these shapes in images. We call this the shape primitives approach. Among the part shapes that have been proposed are polyhedra (Roberts, 1965; Waltz, 1975; Winston, 1975), generalized cylinders and cones (Binford, 1971; Marr & Nishihara, 1978), geons (Biederman, 1987), and superquadrics (Pentland, 1986). Each such proposal works well on a special class of objects. None comes close to capturing the variety and complexity of object parts in general. And each, if viewed as a general theory of the human perception of parts, is clearly ad hoc.

The minima rule separates the issue of finding parts from the issue of describing them. In the shape primitives approach, finding and describing parts is done in a single process: trying to fit primitives to a given image. This entails, however, that parts whose shapes are not in the predefined set of primitives cannot be found. Hoffman & Richards (1984) have argued, instead, that the mechanisms which find parts are more basic, and operate regardless of the shapes of the parts. The minima rule thus has different strengths and limitations. Rather than defining part shapes it defines part boundaries. And rather than looking for part shapes in images, it looks for part boundaries.³ One advantage is that its definition of part boundaries is expressed solely in the language of differential geometry and therefore applies to any shape defined by a piecewise differentiable function. In other words, it applies quite generally. We will discuss its limitations shortly.

But first we examine the rule, its motivations, and its strengths. The minima rule comes in two versions, one for partitioning 3D shapes and one for partitioning 2D silhouettes. Here we need only the version for silhouettes, since our task is to define part cuts on silhouettes. We need not discuss the 3D version, which can be found elsewhere (Hoffman, 1983; Hoffman & Richards, 1984; Bennett & Hoffman, 1987; Beusmans et al., 1987; Hoffman & Singh, 1997) and which requires prior discussion of the differential geometry of surfaces. The version for silhouettes is as follows.

Minima Rule for Silhouettes: *Cut each silhouette into parts using concave cusps and negative minima of curvature of its bounding contour.*

Figure 4 illustrates this rule. Figure 4a shows a silhouette, with regions of positive curvature and negative curvature marked. (The shaded side is “figure,” and the white is “ground.”) Note that positive regions are convex and negative regions concave. This figure also has two cusps on it, one concave and labeled A, one convex and labeled B. Figure




Figure 4. Geometry related to the minima rule. In (a) is shown a silhouette with “figure” indicated by stippling. Regions of positive and negative curvature are indicated by + and – respectively. The cusp labeled A is concave, i.e., pointing into the figure. The cusp labeled B is convex, i.e., pointing into the ground. In (b) the negative minima of curvature are indicated by dots. These are the part boundary points according to the minima rule. In (c) is shown the same silhouette as in (a), but with figure and ground reversed. Note the change in signs of curvature. Now the cusp labeled A is convex whereas B is concave. In (d) the new negative minima of curvature are indicated by dots. A reversal of figure and ground changes the part boundaries.




Figure 5. A demonstration of the minima rule. The two wiggles on the right look different even though they are identical. Why? According to the minima rule you break them into parts differently.

4b shows the boundary points, as defined by the minima rule, between parts. If figure and ground reverse, as in Figure 4c, then regions of negative curvature become positive and vice versa. Moreover concave cusps become convex, and vice versa. This entails, according to the minima rule, that there are now different boundary points, as shown in Figure 4d.

This shift of boundary points, when figure and ground re-

³This is what makes the minima rule a “boundary-based” approach. The “boundaries” in question are boundaries between parts—regions where one part ends and another begins. They are not the bounding contours that define 2D silhouettes—an easy misinterpretation (see Siddiqi & Kimia, 1995, p. 239–240).




Figure 6. The minima rule does not say to form part cuts by connecting consecutive negative minima of curvature. It is clear that doing so can produce very strange parts, as in (a), (b). Sometimes, this does not even give legitimate part cuts, as in (c).

verse, can explain some interesting perceptual effects. One, discussed by Attneave (1971), is shown in Figure 5. On the left is a disk with a wiggle through it. On the right the two halves are pulled apart at the wiggle. By construction each half has an identical wiggle. But notice that the two wiggles look quite different—a fact easily confirmed by experiment (Hoffman, 1983a). The reason, according to the minima rule, is this: The two halves induce opposite assignments of figure and ground on the two wiggles. Therefore regions of positive curvature for one wiggle have negative curvature for the other, and vice versa. And therefore the parts, whose boundaries are (by the minima rule) at negative minima of curvature, must be different for the two wiggles. Thus the reason the wiggles look different is that you divide them differently into parts. A similar account can be given for the well-known face-goblet illusion, and for some 3D illusions as well (Hoffman & Richards, 1984; Hoffman & Singh, 1997). The minima rule has also fared well in several psychophysical tests of its implications for the perception of shape similarity (Hoffman, 1983a), short-term memory for shapes (Braunstein et al., 1989), naming pictures with deleted contours (Biederman & Cooper, 1991), the detection of symmetry and repetition (Baylis & Driver, 1994, 1995a; 1995b; Driver & Baylis, 1995), structure from motion (Saidpour, 1983a), figure-ground perception in 2D and 3D (Hoffman & Singh, 1997), the perception of transparency (Singh & Hoffman, in press), and preattentive popout (Wolfe & Bennett, 1997).

It is clear from the above treatment that we are using curvature as a signed quantity. This differs, for example, from Attneave's (1954) treatment in which curvature is taken to mean magnitude of curvature—and is therefore a positive quantity. Attneave's observation that information along a contour is concentrated at points of 'maxima of curvature' is a statement, in our framework, about the visual importance of both negative minima and positive maxima of curvature. Within the context of parsing, however, the minima rule makes a distinction between these two kinds of curvature extrema, assigning a special status to the negative minima.

The minima rule is distinct from the theory of codons (Richards & Hoffman, 1985; Richards, Dawson, & Whittington, 1986). The domain of codons is plane curves, whereas




Figure 7. Transversality and concave creases. On the left are two generic surfaces. On the right they interpenetrate generically to form a composite object. By transversality, their surfaces form concave creases (depicted by dashed curves) at almost every point where they intersect.

the domain of the minima rule is, as we have seen, silhouettes and three-dimensional shapes. The theory of codons uses minima of curvature, both positive and negative, to parse plane curves. The minima rule, by contrast, uses only negative minima of curvature as boundary points on silhouettes. The theory of codons parses any given plane curve into 'codons', and then represents that curve using a classification scheme for codons. The minima rule for silhouettes, by contrast, only provides boundary points through which part-cuts must pass. It does not define the cuts—and hence the parts—themselves. In particular, it does not join consecutive minima to form parts. Indeed, doing so can give strange parts (Figure 6a, b) or no parts at all (Figure 6c).

The minima rule is based on a principle from the field of differential topology called transversality (Guillemin & Pollack, 1974; Hoffman & Richards, 1984). The relevant case of transversality is shown in Figure 7. On the left are two arbitrary shapes in 3D, labeled S_1 and S_2 . On the right S_1 and S_2 interpenetrate to form a single composite object. S_1 and S_2 are distinct elements of the (sparse) visual scene on the left, and would be natural candidates for parts of the composite object on the right, if only we could distinguish them. Transversality says that we can—if S_1 and S_2 are generic shapes and if they interpenetrate at random. In this generic case, at almost every point where the surface of S_1 intersects that of S_2 the tangent planes to the two surfaces have different orientations in space. Therefore the two surfaces meet, at each such point, in a concave crease. These points are indicated by the dashed curve on the composite object. This crease is concave because it points into the object (i.e., into the "figure"). By contrast, the edges of a cube are convex creases, because they point out of the cube (i.e., out of the "figure"). Another way to understand it is this: If S_1 and S_2 each have smooth surfaces, then the dashed contour indicates the only points on the surface of the composite object that are not smooth. Thus transversality directly motivates a strategy for dividing 3D shapes into parts along concave creases. Combining this with processes of smoothing gives the minima rule for 3D shapes, and combining it further with projection onto an image plane leads to the minima rule for




Figure 8. A limitation of the minima rule. The two crosses in (a) and (b) have the same number of negative minima, and at roughly the same locations. However, their natural partitionings, as given in (c) and (d), are very different. Clearly, some property other than the presence of negative minima of curvature is required to explain the difference in partitionings of these crosses.

silhouettes (Hoffman & Richards, 1984).

Transversality, however, is not an account of part genesis. As we have seen, it applies when two separate objects are joined to form a new object. But it applies equally well, in smoothed form, when an object protrudes from another, as when a branch grows out of a stem. Transversality is an abstract principle of mathematics and so applies regardless of the genesis of the part.

The minima rule explains some aspects of our perceptions of parts. But it has notable limitations. As we mentioned earlier, for a silhouette the minima rule gives precise boundary points at which to cut. But, as noted by Beusmans et al. (1987), the minima rule does not tell how to pair these boundary points to define part cuts. Consider, for instance, the two shapes in Figures 8a and 8b—versions of which appear in Rom and Medioni (1993), Siddiqi & Kimia (1995), and Kimia, Tannenbaum, and Zucker (1991; 1995). The shape in Figure 8a can be seen as the result of pulling the shape in Figure 8b outwards and to each side. Clearly, such a transformation has little effect on the negative minima of the shape: There are still four negative minima and at roughly the same locations. And in both cases, the minima rule simply provides these four negative minima as boundary points, and is silent about how to connect them to form part cuts.⁴ However, the natural perceptual organizations of these two shapes are different—see Figures 8c and 8d. The shape in 8a is most naturally perceived as a small central core surrounded by four small parts, whereas the shape in 8b is most naturally perceived as a large vertical body with two small parts protruding on the sides. Hence the minima rule, in itself, is unable to account for the difference in perceptual organizations of the two shapes.⁵ There are two reasons. First, the minima




Figure 9. Another limitation of the minima rule: Some part boundaries are not negative minima of curvature.

rule uses only properties of the contour that outlines the silhouette. And second, the minima rule uses only differential properties of the contour (namely, the presence of negative minima of curvature). To define part cuts, we must add to the minima rule (i) properties of the region enclosed by the contour, and (ii) properties that are more global to the shape.

Another limitation of the minima rule is that it does not indicate which points in addition to negative minima of curvature are good part boundaries, even though there are surely such points (Hoffman & Richards, 1984, p. 72). Figure 9, for instance, shows an elbow which can be naturally cut as indicated by the dotted line. This line terminates at one end in a negative minimum of curvature, and at the other in a point with zero curvature.

The short-cut rule, as we shall see, augments the minima rule in a way that repairs some of these limitations. But there is relevant prior work on parts and part cuts by Marr (1977), Biederman (1987, 1990), and by Siddiqi & Kimia (1995), to which we now turn.

Generalized Cones

A silhouette can convey a rich sense of shape in three dimensions. To explain this, Marr (1977) suggested that human vision interprets silhouettes as being the images of generalized cones (first defined by Binford, 1971). A generalized cone is the three-dimensional surface “swept out by moving a simple smooth cross-section along some axis, at the same time magnifying or contracting it in a smoothly varying way” (Marr, 1977, p. 447)—see Figure 10. Marr showed that generalized cones are a powerful tool for interpreting simple silhouettes. Marr further suggested that human vision interprets complex silhouettes as the composition of two or more generalized cones. To do so it must cut the silhouette into parts, each of which corresponds to a single generalized cone, and then analyze each cone.

These cuts, Marr proposed, are straight lines that lie entirely within the silhouette and divide it into regions. Each such region has a qualitative symmetry whose axis is the image of the axis of the corresponding generalized cone. The endpoints of each cut are at curvature inflections, and sometimes in concave or convex regions. The precise position of

⁴ Considerations of genericity, however, allow us to rule out part cuts that cross each other—see Beusmans et al. (1987).

⁵ It is important to note, however, that since the minima rule does not make cuts, it cannot, a fortiori, make unnatural ones (cf. Kimia et al., 1995, p. 212).




Figure 10. The definition of a generalized cone.




Figure 11. (a) Some geons, and (b) some objects made by combining geons. Adapted from Biederman (1987).

the endpoints is chosen to minimize the portion of the silhouette's contour left unmatched by the qualitative symmetries.

Marr's analysis is elegant. But as he was well aware the assumptions it requires are restrictive: The silhouette must be the image of generalized cones whose axes are coplanar, and the viewing direction must not significantly foreshorten the axes of the generalized cones. A more general account is needed.

Geons

To explain the speed and accuracy with which human vision recognizes objects at the entry-level—for example, as a table, a horse, or a car—Biederman (1987, 1990) proposed a theory called recognition-by-components or RBC. RBC postulates 24 primitive volumetric shapes, called geons⁶ (see Figure 11a for some examples), and claims that any visual object can be represented as an arrangement of these geons in specific spatial relationships (see Figure 11b). Furthermore, the geon representations are, by construction, stable over changes in viewpoint, so that the same representation is activated from almost any viewing direction.

RBC faces the same difficulties that primitive-based schemes typically face, viz., limited generality. However, it has the advantage that geons derive from the principle of “nonaccidental properties” (Lowe, 1985; Witkin & Tenenbaum, 1983), whereas polyhedra, superquadrics, and generalized cones make no appeal to a first principle. Nonaccidental properties are properties of 3D shape which, generically, survive projection onto an image plane. An example is the distinction between straight and curved: a curved edge in 3D will, generically, project to a curved edge in the image plane. It takes a special viewpoint to make a curved edge in 3D project to a straight edge in an image—and if such a special viewpoint happens to occur, human vision often misinterprets the resulting image because it interprets the image as arising from a generic, rather than special, viewpoint




Figure 12. The four nonaccidental properties used in deriving the set of geons: (a) curved versus straight cross section; (b) constant versus expanding only versus expanding and then contracting cross-section; (c) symmetrical versus asymmetrical cross-section; (d) curved versus straight axis.




Figure 13. Geons have (a) pointed tips, or (b) truncated tips, but not (c) rounded tips, even though the difference pointed, truncated, and rounded tips is one which survives projection, and is therefore a nonaccidental property.

(Freeman, 1994). Figure 12 shows the four nonaccidental properties used to generate the set of geons.

Although geons appeal to the principle of nonaccidental properties, they are not the complete set of shape primitives that follow from this principle. As an example, geons end either in pointed tips (as in Figure 13a) or in truncations (as in Figure 13b); there are no geons with rounded tips (Figure 13c). But clearly the distinction between rounded tips, truncated tips, and pointed tips is one that generically survives projection—and is therefore a nonaccidental property. Indeed it is required for the proper recognition of toes, fingers, peeled bananas, and aircraft fuselages, which cannot be approximated by geons (i.e., with truncated or pointed tips). One might say this problem is easy to fix: Just add to the list of geons a couple more shapes with rounded tips. And no doubt this is easy to do. But there are many other nonaccidental properties that have been omitted from the list as well. Why are some nonaccidental properties used to define geons and others not? No principle has been given for choosing some nonaccidental properties and not others.

⁶ The 1987 version of RBC had 36 geons, but the 1990 version has 24.




Figure 14. Siddiqi & Kimia's (1995) definition of limb. In (a) the limb is depicted by the dashed curve. The cocircular tangents are depicted by arrows. In (b) and (c) are examples where the definition of limb fails, since the tangents are not cocircular. This failure of cocircularity is generic.

Limbs and Necks

Siddiqi and Kimia (1995), building on the work of Kimia et al. (1991; 1995), recognized the limits of the minima rule, and proposed that silhouettes are parsed in two ways: limbs and necks. Siddiqi & Kimia (1995) define a limb as follows: "A limb is a part-line going through a pair of negative curvature minima with co-circular boundary tangents on (at least) one side of the part-line" (p. 243). A "part-line" is a part cut. Two tangent vectors are "co-circular" if both are tangent to one circle (Parent & Zucker, 1989, p. 829). Figure 14a illustrates their definition of limb, with the part cut depicted by a dashed line, and co-circular tangents depicted by arrows. This definition almost never applies to real parts since its requirement of co-circular tangents almost surely⁷ never holds. Figures 14b and c, for instance, show examples in which tangents at negative minima of curvature are not co-circular. It is easy to concoct such examples since the failure of co-circularity is generic.⁸ This means that allowing tolerance in the computation of co-circularity, which is useful in other contexts (e.g., Parent & Zucker, 1989), cannot fix the problem here: limbs would just go from measure zero to highly unlikely.

The definition of limb, furthermore, has a counterintuitive implication: many part cuts that one normally calls limbs, e.g., cuts for the arms and legs of the human body (Figure 15), fail both the co-circularity condition and the condition of passing through two negative minima of curvature, and thus fail to be limbs according to the above definition. The definition of limb, therefore, is much too restrictive. The cuts it defines are almost surely never found on any real shape.

Siddiqi & Kimia (1995) define a neck as follows: "A neck is a part-line which is also a local minimum of the diameter of an inscribed circle" (p. 243). Figure 16a illustrates this definition, as does the central cut in Figure 16b and the two cuts in Figure 16c. Unfortunately the definition of neck given by Siddiqi & Kimia fails for a large class of shapes that should be classified as necks. In Figure 16d, for example, the dashed line indicates a natural cut, and should be made. But it is not captured by the definition of a neck. The problem is that the circle cannot be inscribed: the cut is the diameter of the circle, but the circle is too big to be inscribed. Thus the definition of neck is too restrictive.




Figure 15. Siddiqi & Kimia's (1995) definition of limb excludes human (and animal) limbs; such limbs are not cocircular, and many do not contain two negative minima.

Short Cuts

The situation seems to be this. Negative minima of curvature are powerful geometric determinants of perceived part boundaries on 2D silhouettes. But they are not the only such determinants. And a comprehensive (bottom-up) account of part perception must describe all geometric factors and their interactions in determining not just boundary points, but entire part cuts. Here, we propose an important such factor: cut length.

Consider the elbow depicted in Figure 17a. Which cut seems most natural—cut xy or cut xz ? Casual inspection suggests that xy is by far more natural. (This is also, as we will see shortly in our experiments, the verdict of subjects in experiments with similar figures.)

Why is xy the preferred cut? Even in a shape as simple as this elbow, several factors are at play. We have designed this elbow, however, to minimize the effects of most factors other than cut length. For example, the two lengths labeled L in Figure 17a are identical, so that this length is not a factor. The points y and z both have identical curvature (viz., zero), so that this also is not a factor. One difference is the area of the two parts defined by the two cuts, but this difference

⁷ Something is said to be true "almost surely" if it is true "everywhere except possibly on sets of measure zero." See, for example, Guillemin & Pollack (1974) and Halmos (1950).

⁸ A simple proof of this is the following: Pick at random two points in the plane. Pick at random a line passing through the first point. Draw the circle passing through both points and tangent to this line (this circle is defined uniquely). Then, of the countless lines through the second point, just one will be tangent to this circle. Thus cocircularity obtains with measure zero, and limbs almost surely don't occur.




Figure 16. Siddiqi & Kimia's (1995) definition of neck. In (a) the neck is depicted by the dashed curve. In (b) the middle dashed line is a neck. In (c) both dashed lines are necks (Siddiqi et al., 1996). In (d) is an example where the definition of neck fails, since the dashed line is not the diameter of an inscribed circle.




Figure 17. The distance factor. In both (a) and (b) the cut xy is preferred to the cut xz because xy is shorter.

is eliminated in Figure 17b, and still the cut xy is preferred. This suggests that a key factor here is that the Euclidean distance from x to y is shorter than that from x to z .

Therefore it appears that the goodness of a cut joining two points x and y involves the Euclidean distance $d_{x,y} = |x - y|$. We say the goodness “involves the distance” rather than “is the distance.” There is more to this factor. In some cases closer is not better, as shown in Figure 18. Points x and w in this figure are certainly close to each other, but we are not tempted to cut the figure from x to w .

Why? Figure 18 suggests that the answer lies in the geometry of the silhouette between x and w . In this figure it seems that points x and z could have a cut between them, but points x and w could not, even though the distance between x and w is smaller than that between x and z . So Euclidean distance alone is not the key to determine how close two points must be before a cut between them should be impossible. One consideration, of course, is that the cut xz involves two negative minima, whereas cut xw involves only one. But there is another factor at play here. To state it, we need the notion of local symmetry. Local symmetry is a weak form of symmetry that




Figure 18. Closer is not always better: A part cut must also cross an axis of local symmetry.




Figure 19. The axes of local symmetry for two shapes.

allows for the axes of symmetry to be curved, and also for axes that span only local subshapes of an entire shape. What the axes of local symmetry provide, in effect, is the skeletal axial structure of any given 2D shape. Figure 19 displays the axes of symmetry for two silhouettes. Various schemes have been proposed to compute the axes of local symmetry (Blum & Nagel, 1978; Brady & Asada, 1984; Leyton, 1992). We will use Brady & Asada's (1984) definition.

What seems to be key in Figure 18, then, is that the straight-line cut between x and z passes through an axis of local symmetry of the silhouette, whereas the straight cut between x and w does not. A cut that fails to cross a local-symmetry axis simply does not chop off a region natural enough to be considered for parthood. In other words, as long as a cut crosses a local symmetry axis, shorter cuts are (other things being equal) better. If a cut does not cross an axis of local symmetry, it is simply no good. This is not a differential-geometric property of the contour of the silhouette, but rather a more global geometric property of the region enclosed by it.

Why should cut length be such an important factor in determining part cuts on a silhouette? The answer lies in the geometry of transversal intersections. Consider a generic intersection, in 3D space, of two cylinders such that the radius of cross section of one cylinder is larger than that of the other (see Figure 20). Only two kinds of transversal intersections are possible that will produce ambiguities in parsing the projected silhouette: a complete intersection, as shown in Figure 20a, and a partial intersection, as shown in Figure 20b. A complete intersection can be characterized by the property that it leads to two contours of intersection, whereas a partial intersection leads to a single contour of intersection.

Let us consider the relative probability of obtaining these




Figure 20. The two kinds of transversal intersections for two cylinders with *unequal* radii of cross section: (a) complete intersection, and (b) partial intersection.




Figure 21. The two kinds of transversal intersections for two cylinders with *equal* radii of cross section: (a) complete intersection, and (b) partial intersection.

two kinds of intersections, as a function of the ratio of the radii of the two cylinders (larger to smaller). As this ratio gets larger (say, as we keep the radius of the thicker cylinder fixed, and gradually decrease the radius of the thinner one), the probability of obtaining a partial intersection gradually decreases to 0, and the probability of obtaining a complete intersection increases to 1. (The limiting case, of course, is achieved when the thin cylinder is just a line piercing through the thicker cylinder.) On the other hand, as this ratio gets smaller and approaches 1 (say, as we keep the radius of the thicker cylinder fixed, and gradually increase the radius of the thinner one), the probability of obtaining a partial intersection gradually increases to 1, and the probability of obtaining a complete intersection decreases to 0. In the limiting case, where both cylinders have the same radius (see Figure 21), the probability of obtaining a complete intersection is, in fact, 0. In other words, the set of relative orientations and translations that give rise to a complete intersection (shown in Figure 21a) has measure zero in the set of all possible orientations and translations that yield an intersection.⁹

Consider now the shape of the concave creases produced in each of the two intersection types. In the case of a complete intersection, so long as the two cylinders have unequal radii, the concave creases always encircle the thinner cylinder, and never the thicker one (see Figure 20a). In the case of a partial intersection, the concave crease encircles neither of the two cylinders (see Figure 20b)—so in this case the

parsing of the projected image is left ambiguous.

The outcome of this analysis, then, is as follows: As the ratio of the radii of the two cylinders (larger to smaller) increases, the probability that the two cylinders will meet in a complete intersection gets closer to 1. Therefore, the probability that the concave crease of the intersection goes around the thinner cylinder gets closer to 1. Hence, a projected silhouette of this intersection should be naturally parsed using the shorter cuts (which correspond to the projections of the concave creases produced by the intersection). On the other hand, as the ratio of the radii of the cylinders approaches 1, there is a high probability that the cylinders will meet in a partial intersection—therefore, the parsing of the projected silhouette is ambiguous.

Given any silhouette, whether cylindrical or not, whose corresponding 3D geometry is unknown, the principle of genericity (in the form used by Freeman, 1994) dictates that the silhouette should be interpreted as deriving from a 3D shape that is roughly as deep as it is wide in the image. Therefore, as we saw with the cylinders above, the concave crease will generically go around the part with the thinner silhouette. For this reason, we hypothesize that:

1. given a silhouette in which the part boundaries can be paired in more than one way to yield cuts, human vision will prefer to make the shorter cuts, and
2. the probability of making the shorter cuts will increase as the ratio of the longer to the shorter cut gets more extreme.

The short-cut rule differs from necks (Siddiqi & Kimia, 1995) in that the short-cut rule, unlike necks, does not require the use of an inscribed circle, a requirement that we noted earlier is quite restrictive. Therefore the short cut rule is not restricted to measuring distances only along diameters of inscribed circles, but instead takes into account distances between all pairs of points on the contour of the silhouette which are separated by an axis of local symmetry. For example, in Figure 17a and b the short-cut rule can explain why subjects prefer the cut *xy* over the cut *xz*; however neither cut is a limb or a neck (Siddiqi & Kimia, 1995), so limbs and necks cannot explain this preference.

This ecological motivation for the short-cut rule also motivates the well-known Petter's rule (Petter, 1956; Kanizsa, 1979, p. 40) for modal completion of contours, which states that human vision prefers to make modal completions as short as possible. For instance, for the cross in Figure 23b (which is rendered in a homogenous color), human vision prefers to make a modal completion along contour *y* rather than contour *x*, so that the vertical bar is seen as occluding the horizontal one. This rule has been motivated by the heuristic that, due to perspective projection, closer objects tend, *ceteris paribus*, to have larger retinal images (Shipley & Kellman, 1992b; Stoner & Albright, 1993). But this motivation is admittedly very rough (Tommasi, Bressan, & Vallortigara,

⁹ Recall that our current discussion includes only those intersections that produce ambiguities in parsing the projected silhouette.




Figure 22. Some simple figures used in our experiments.

1995). However, the ecological motivation we have already given for the short-cut rule can be applied to Petter’s rule as follows. When human vision is presented with a stimulus, like Figure 23b, of homogeneous color, it finds the negative minima of curvature and pairs them using the short-cut rule. Because the color is homogeneous, it is initially ambiguous whether these pairings should be taken to be part cuts or modal contours (since it is not known whether the silhouette arises from a single object with parts, or from two different objects separated in depth). Further processing, possibly going on in parallel with the pairing process, is required to make this decision. Once this decision is made, the pairing decided on by the short-cut rule is then taken to be either a part cut or a modal contour. In this way, Petter’s rule for modal contours inherits the ecological motivation for the short-cut rule.

The Experiments

We chose two classes of shapes for the experiments: crosses (see Figure 22a) and elbows (see Figures 22b and 22c). These are natural candidates because, as we have seen, they are simple, they display key limitations of the minima rule, and they allow a clean test of the short-cut rule.

The crosses have four negative minima of curvature. While the minima rule states that these are part boundaries, it does not specify how they should be joined. The short-cut rule predicts that subjects will join these part boundaries so as to make the shortest cuts possible. Because these figures have straight lines and 90° angles, they are easily parametrized and thus allow us to test a truly representative sample of crosses, not just those that might favor the short-cut rule. The elbows have but one negative minimum. The minima rule states that this point must be a part boundary but does not indicate what the other part boundary should be. The short-cut rule predicts that subjects will choose that part boundary which will create the shortest possible cut. These figures, like the crosses, are easily parametrized and thus allow us to test a representative sample of elbows.

Experiments 1 and 3 required subjects to hand draw cuts on shapes, and contained fewer trials, whereas Experiments 2, 4, and 5 involved a 3AFC task on a computer, and contained many more trials. These experiments do not test whether the subjects would spontaneously parse such shapes if not asked to do so: Our task required subjects either to




Figure 23. (a) The parameters of the cross-shaped figures used in Experiments 1 and 2. The stimuli were drawn with black outlines and gray interiors in order to suppress the perception of illusory contours, which are quite striking in (b), for example. We wanted subjects to perceive a single object with parts, and not two objects partially occluding the other.

draw a fixed number of cuts, or to choose among a few given cuts. However, as we note in the Introduction, there is prior empirical work which suggests that human vision does parse shapes spontaneously (Biederman, 1987; Biederman & Cooper, 1991; Braunstein et al., 1989; Hoffman, 1983a; 1983b; Hoffman & Singh, 1997), and perhaps even preattentively (Baylis & Driver, 1995a; 1995b). What our current experiments investigate is how the preferred cuts change with various parameters of the shapes. The ‘Crosses’ Experiments The crosses were symmetric about the vertical and horizontal axes. This allows them to be parametrized by four (orthogonal) parameters x , y , z , and w (see Figure 23a). The space of such crosses thus has four parameters. However, since we assume, for now, that the parsing of shapes is scale invariant, we do not wish to distinguish between a cross and scaled versions of it. Therefore, we can factor out scaling, so that the space of such (scale-invariant) crosses has three parameters—for example, the three (orthogonal) parameters $\frac{x}{y}$, $\frac{z}{y}$, and $\frac{w}{y}$.

The independent variables we chose to parametrize the space of crosses were:

1. *distance ratio*,

$$d = \frac{\text{length of horizontal cut}}{\text{length of vertical cut}} = \frac{x}{y};$$

2. *area ratio*,

$$A = \frac{\text{area of part produced by horizontal cut}}{\text{area of part produced by vertical cut}} = \frac{xz}{wy};$$

3. *horizontal protrusion*,¹⁰

$$h = \frac{\text{length of horizontal part}}{\text{width of horizontal part}} = \frac{w}{y}.$$

Note that the area ratio, $\frac{xz}{wy}$, is orthogonal to the two other variables and is scale invariant. We chose this variable, instead of $\frac{z}{y}$, for its more intuitive and geometric interpretation. For example, it allows us to test whether the area of a part created by a given cut influences parsing.

¹⁰ We are using the term “protrusion” somewhat loosely. See Hoffman & Singh (1997) for a precise definition that applies generally.




Figure 24. Three types of cuts on a cross: (a) horizontal, (b) vertical, and (c) multiple.

Based on the minima rule, the short-cut rule, the principle of genericity (e.g., Beusmans et al., 1987), and considerations of symmetry, we expected three ways of parsing a cross to be most natural. We called these horizontal cuts (Figure 24a: two horizontal cuts, both passing through negative minima), vertical cuts (Figure 24b: two vertical cuts, both passing through negative minima), and multiple cuts (Figure 24c: both horizontal and vertical cuts are present).

We predicted that subjects would make the shortest cuts between parse points. As an example, for the cross in Figure 23a, the vertical cuts are shortest and seem most natural.

The stimuli were drawn with a black outline and gray interior, as illustrated in Figure 23a. This was done to suppress the perception of illusory contours—which are quite striking, for example, when the figures are filled uniformly with black, as in Figure 23b (see, for example, Kanizsa, 1979, and Shipley & Kellman, 1992, for work on self-splitting figures). We wanted subjects to perceive a single object, and not two rectangular objects—one partially occluding the other.

Experiment 1

As discussed above, we expected the crosses to be parsed with vertical cuts, horizontal cuts, or a combination of these. To check this, the first experiment tested whether subjects would use any other cuts to parse the crosses, in a free-drawing task.

Method

Subjects. Forty-five undergraduate students at the University of California, Irvine, volunteered to participate for course credit. Data from two subjects were excluded from analysis because they failed to follow instructions.

Materials. The stimulus set consisted of a packet of 113 sheets of paper (27 different stimuli, each repeated 4 times, plus 5 practice trials.) Each sheet had a cross-shaped figure printed on it, that was symmetric about both of its axes. To avoid any vertical or horizontal biasing effects, half of the figures were presented with a tilt of 15 degrees to the left of vertical, and the other half with a tilt of 15 degrees to the right of vertical. The first five sheets were practice.

Design. This experiment was a $3 \times 3 \times 3$ within-subjects factorial design. The three independent variables were: (i) The distance ratio, d , (ii) the area ratio, A , and (iii) the horizontal protrusion, h . Each of the three independent variables had three levels: $1/2$, 1 , and 2 . Figure 25 illustrates all the stimuli used for the experiment. The dependent measure was the percentage of vertical cuts. (“Multiple cuts” was not an




Figure 25. All figures used in Experiments 1 and 2 organized by distance ratio, area ratio, and horizontal protrusion.

option in this experiment.) The short cut rule predicts that as d increases the percentage of vertical cuts increases, because as d increases the vertical cuts become shorter than the horizontal.

Procedure. The experiment was run in six separate sessions. Each subject was seated at a desk, in a separate cubicle room. On the desk was the stimulus set packet (placed face down), a pencil, and a ruler. The subjects were instructed to (a) pick up the top sheet from the packet, (b) turn it face up, (c) decide how they would cut the figure most naturally into three parts, and (d) draw two straight-line cuts, with the ruler and pencil provided, to achieve this partitioning. (We limited subjects to two cuts because this is the simplest task that adequately tests our theoretical expectations.) After they were done with a figure, they were to place that sheet face-down, in a separate stack. It was stressed that they should




Figure 26. The proportion of vertical cuts made by subjects in Experiment 1, as a function of distance ratio and area ratio.

not look forward or backward in the two stacks. The first five sheets were practice trials. The ordering of the experimental trials was randomized for each subject. The subjects were monitored continuously to make sure they were following instructions. Subjects were debriefed and thanked for their participation.

Results and Discussion

Each response was classified as a vertical cut, horizontal cut, or other. A response was classified as vertical if both cuts made were vertical, and passed within 5 mm of the negative minima. It was classified as horizontal if both cuts made were horizontal, and passed within 5 mm of the negative minima (see Figure 24). It was classified as other if it was neither horizontal, nor vertical. Out of all responses from the 43 subjects whose data were included in the analysis, less than 0.1% were other. These were excluded from the analysis.

An alpha level of .05 was used for all statistical tests. A three-way analysis of variance (ANOVA) showed a main effect of the distance ratio, $F(2, 84) = 44.743, p < .0001$, but no main effect of the area ratio, $F(2, 84) = 1.839, ns$, and no main effect of the horizontal protrusion, $F(2, 84) = 0.933, ns$. There was a significant interaction between the distance ratio and the area ratio, $F(4, 168) = .381, p < .05$. Post-hoc comparisons revealed that the area ratio had an effect only when the distance ratio was 1—and for these stimuli, subjects preferred to cut parts with larger areas. (See Figure 26 for a graphed summary of the results.) These results suggest that in these stimuli the distance ratio is the dominant geometric factor used by human vision for parsing. As predicted, shorter cuts are much preferred to longer.

These results also indicate that in forced-choice experiments on crosses it is legitimate to restrict the alternatives to horizontal cuts, vertical cuts, or a combination of these.

Experiment 2

Because of the tedious task in Experiment 1, each stimulus was presented only 4 times to each subject. This precluded reliable modeling of individual subjects' data, or comparing of trends across subjects. However the second experiment

used a forced-choice task on a computer, allowing many more trials, and allowing us to model each subject's data individually.

Method

Subjects. Ten graduate students at the University of California, Irvine, volunteered to participate. Subjects were not paid and received no course credit.

Materials. The experiment was run on a Macintosh Quadra 840AV, using the program SuperLab. The stimuli used were the same 27 cross-shaped figures as in Experiment 1 (see Figure 25), but each was presented 24 times.

Design. This experiment had the same $3 \times 3 \times 3$ within-subjects factorial design as Experiment 1: the independent variables were d , A , and h , and each had the levels 1/2, 1, and 2. Each stimulus figure was repeated in 24 trials, resulting in a total of 648 (+ 5 practice) trials. There were two dependent measures: the percentage of vertical cuts responses, and the percentage of multiple cuts responses.

Procedure. The subjects were seated at a desk, 0.6 meters from a computer screen. The computer displayed the instructions for the experiment. The subjects pushed a key to indicate that they were done with the instructions. Each trial that followed was structured as follows: First there appeared on the screen, for 2 seconds, a cross-shaped figure, presented either with a tilt of 15 degrees to the left of vertical, or 15 degrees to the right of vertical. During this time, the subjects were to decide how they would partition the shape most naturally into parts. There followed a blank screen for 500 ms, and then a screen with the initial cross-shaped figure (in its original location) along with three possible partitionings of that figure presented, in small, at the bottom of the screen, and numbered 1, 2, and 3. The subjects indicated, by pressing the corresponding number key, which of the three choices corresponded to their partitioning. The three choices were vertical cuts, horizontal cuts, and multiple cuts (see Figure 24). The numbering of the three different partitioning choices was counterbalanced across trials. The subject's response terminated the trial.

The first 5 trials were practice. The experimental trials were divided into 4 blocks, each consisting of $27 \times 6 = 162$ trials, for a total of 648 trials.

Results and Discussion

Subjects' data were initially analyzed for internal consistency. Data for the first 12 instances of each figure were correlated to data for the last 12. Prior to the experiment we chose to reject all data from any subject whose correlation was less than 0.5. No data were eliminated. Subjects' data were analyzed individually. First the multiple cuts responses were analyzed. In general, subjects made few multiple cuts (mean = 4.23%), and did so primarily when the distance ratio and area ratio were both 1. Five of the ten subjects made virtually no multiple cuts (mean = 1.48%). The remaining five subjects made multiple cuts only when both d and A were 1 (mean = 6.98%). Because of the low overall occurrence of the multiple cuts responses, the percentage of vertical cuts responses was taken as the primary dependent measure.

Linear Regression. For vertical cuts responses, each subject's

data was fitted to a number of linear regression models. The percentage of vertical cuts, c_v , was transformed to a new variable, c'_v , using an arc sine function (Kendall & Stuart, 1963) to improve the normality of its distribution. The precise transform was

$$c'_v = 2 \text{Sin}^{-1} \sqrt{c_v} \quad (1)$$

The terms corresponding to the distance ratio d , and the area ratio A , were taken to be, respectively, $\log(d)$, and $\log(A)$, (instead of simply d and A).¹¹

From some preliminary modeling, it was clear that the distance ratio was doing most of the work. The models we tested, therefore, all included the distance ratio—along with all possible combinations of the other two independent variables. In these models, the horizontal protrusion, h , explained almost none of the variability in the data. (This remained true, also, when we replaced h with $\log(h)$.)

Therefore, we needed to consider only those models that involved the variables $\log(d)$, and $\log(A)$. Also, in almost every case, adding $\log(A)$ to the model containing only the distance ratio did significantly improve the model's fit.

For these reasons, our conclusion was that the linear regression equation,

$$c'_v = \beta_0 + \beta_1 \log(d) + \beta_2 \log(A),$$

provided the best model for our data. Figure 27 shows the summary of this model, with the β_1 and β_2 parameters plotted along the x -axis and y -axis, respectively, with 95% confidence intervals for the estimates of these parameters, and with r^2 -values of the model's fit to each individual's data.

In sum, distance was again the strongest factor, with subjects preferring shorter, rather than longer, cuts as predicted. Although area was also used by individual subjects (out of 10 subjects, 8 had a coefficient for A that was significantly different from zero), it was used differently by different subjects: 7 out of the 10 subjects cut off parts with the smallest area, while 3 cut off parts with the largest area. Overall, however, there was no main effect of area.

The 'Elbows' Experiments

Experiment 3

Unlike the crosses from Experiments 1 and 2, the elbow in Figure 22b has but one negative minimum of curvature, labeled g . The minima rule states that this point is one end of a part cut, but does not state which point is the other end. Several cuts seem plausible. Joining at point i seems reasonable as this is the only other extremum of curvature in the figure, albeit a positive maximum (we call this a diagonal cut). However, joining at either j or h also seems reasonable as these are the locally shortest cuts possible that pass through the figure's axis of symmetry (we call these horizontal cut and vertical cut, respectively). The short-cut rule makes a clear prediction: subjects will choose the shortest part cut, i.e., the cut gh . Experiment 3, like Experiment 1, studies the cuts subjects make in a free-hand task. Experiment 4, like




Figure 27. A plot of the coefficients β_1 and β_2 for the final linear model $c'_v = \beta_0 + \beta_1 \log(d) + \beta_2 \log(A)$ in Experiment 2. Each point represents a different subject. Also shown are 95% confidence intervals, and r^2 -values for the model.

Experiment 2, studies cuts in more detail. As in Experiments 1 and 2, three factors are systematically varied: the distance ratio, d , the area ratio, A , and the horizontal protrusion, h (see Figure 28). Experiment 5 investigates whether smoothing the

¹¹ The reason for this is as follows: When a cross-shaped figure with distance ratio d_0 , and area ratio A_0 , is rotated through an angle of 90° , the distance and area ratios of the resulting cross become, respectively, $\frac{1}{d_0}$, and $\frac{1}{A_0}$. At the same time, by our very convention, vertical cuts on the original cross become horizontal cuts on the rotated cross, and vice versa. Also (ignoring multiple cuts), we have,

$$c_v = 1 - c_h,$$

where c_v is percent vertical cuts, and c_h is percent horizontal cuts. Now, if our model were, for example, $c'_v = \beta_0 + \beta_1 \log(d) + \beta_2 \log(A)$, (i.e., without the \log 's), the equation above would give us, $\beta_0 + \beta_1 d_0 + \beta_2 A_0 = 1 - (\beta_0 + \beta_1 \frac{1}{d_0} + \beta_2 \frac{1}{A_0})$, which cannot hold, except in the degenerate case where $d_0 = A_0 = 1$. On the other hand, if d and A , are replaced by $\log(d)$ and $\log(A)$ respectively, this problem is solved:

$$\beta_0 + \beta_1 \log(d_0) + \beta_2 \log(A_0) = 1 - (\beta_0 + \beta_1 \frac{1}{\log(d_0)} + \beta_2 \frac{1}{\log(A_0)}),$$

In fact, this equation now tells us that for a perfect equality, we must have $\beta_0 = 0.5$. (This argument, however, is clearly not applicable to the horizontal protrusion, h .)

The assumption behind this argument ($c_v = 1 - c_h$) holds for about half of the subjects. For the others, the transform can be thought of as a simple fitting technique.




Figure 28. The parameters of the L-shaped figures used in Experiments 3–5.

cusps affects the short-cut rule. In this experiment the horizontal protrusion is held constant, but a new variable, level of curvature, κ , is added. For all experiments we expect the distance ratio to be a strong factor, with subjects preferring to make shorter as opposed to longer cuts.

Based on the theoretical considerations discussed above, we expected the elbows to be parsed with three kinds of cuts: vertical, horizontal, and diagonal. To check this, Experiment 3, like Experiment 1, allowed subjects to make freehand cuts.

Method

Participants. Twenty five undergraduates at the University of California, Irvine, volunteered to participate for course credit. The data of one subject, who failed to follow instructions, was eliminated.

Materials. Testing and design were as in Experiment 1, with one difference: the stimuli were elbows rather than crosses. The values for distance ratio and area ratio were as in Experiment 1. Figures were randomly rotated between positive 15–35 degrees and negative 15–35 degrees to minimize orientation effects. Subjects saw each figure 4 times.

Design. This experiment was a $3 \times 3 \times 3$ within-subjects factorial design. The three independent variables of the experiment were: (i) The distance ratio, d , (ii) the area ratio, A , and (iii) the horizontal protrusion, h . Each of the three independent variables had three levels: 1/2, 1, and 2. The dependent measure was the number of vertical cuts. Figure 29 illustrates all the stimuli used for Experiment 3.

Procedure. Each subject was seated at a desk, in a separate cubicle. On the desk was a ruler, a pencil, and a stack of 113 sheets of $8\frac{1}{2}$ by 11 inch paper. Each sheet displayed an elbow from the set shown in Figure 29. The stack was face down. Subjects were instructed to (a) pick up the top sheet from the stack, (b) turn it face up, (c) decide how they would cut the figure into two parts most naturally, and (d) draw that cut with the ruler and pencil provided. After they were done with a figure, they were instructed to place that sheet face-down, in a separate stack. The first five sheets were practice trials. The ordering of experimental trials was randomized for each subject. The subjects were monitored continuously to make sure they were following instructions. Subjects were debriefed and thanked for their participation.




Figure 29. All elbow stimuli used in Experiments 3 and 4.




Figure 30. Examples of a vertical, horizontal, and diagonal cut in Experiments 3 and 4.

Results and Discussion

Each response was classified as a vertical cut, horizontal cut, or diagonal cut. Any cut with an endpoint more than 5 mm away from a “perfect” example of each cut (see Figure 30) was classified as other. Out of 2700 cuts made by all subjects, only 88 could be classified as other. Of these 88, seventy-four were made by one subject, and 12 by another. Out of all 24 subjects whose data were included in the analysis, only 2.4% were diagonal. Furthermore, diagonal cuts were made primarily when the distance ratio and area ratio were both equal to 1, as illustrated in Figure 31. Therefore, we analyzed the percentage of vertical cuts.

A three-way ANOVA showed a main effect of the distance ratio, $F(2,46) = 74.46, p < .0001$, and a main effect




Figure 31. The proportion of diagonal cuts made by subjects in Experiment 3, as a function of distance ratio and area ratio.




Figure 32. The proportion of vertical cuts made by subjects in Experiment 3, as a function of distance ratio and area ratio.

of the area ratio, $F(2,46) = 5.65, p < .01$, but no main effect of the horizontal protrusion, $F(2,46) = 0.06, ns$. Post-hoc comparisons with a Tukey’s HSD revealed that, as predicted by the short-cut rule, subjects made significantly more vertical cuts when the distance ratio was 2 than when 1, or 1/2. The number of vertical cuts made at a distance ratio of 1 was significantly greater than at a distance ratio of 1/2. Furthermore, subjects made significantly more vertical cuts when the area ratio was 1/2 than when 1 or 2. There was an interaction between the distance ratio and the area ratio, $F(4,92) = 4.50, p < .01$. (See Figure 32 for a graphed summary of the results.) As in Experiment 1, post-hoc analysis revealed that the area ratio had an effect only when the distance ratio was 1; for these stimuli, subjects preferred to cut parts with larger areas. As predicted, subjects preferred to make shorter cuts for these figures.

Experiment 4

Experiment 3, as in Experiment 1, had only 4 repetitions per stimulus per subject. This meant that we could not reliably model individual subjects’ data, or compare trends

across subjects. However, in Experiment 4, as in Experiment 2, we used a 3AFC design presented by computer. We therefore had many more trials.

Method

Participants. Twelve graduate students volunteered to participate. Subjects were not paid and received no course credit.

Design and Procedure. Testing and design were as in Experiment 2. Stimuli were the same as those used in Experiment 3. Values for the distance ratio, area ratio, and horizontal protrusion were the same as those used in Experiment 3. Figures were randomly rotated between positive 15 – 35 degrees and negative 15 – 35 degrees to minimize orientation effects. Figures were presented mirrored about the vertical axis to control for left/right bias. The display of horizontal, vertical and diagonal choices in the key were counterbalanced for each figure. We presented subjects with each of the 27 figures 24 times, 12 mirrored and 12 non-mirrored, for a total of 648 trials. The procedure was identical to that in Experiment 2. On each trial subjects saw one large figure. Subjects were instructed to cut the figure into two parts most naturally, as quickly and carefully as possible. A key then appeared below the figure. Subjects were to select the option that best represented their choice. We limited the choices to these three cuts based on our findings in Experiment 3. The instructions were designed to encourage subjects to decide how they would parse the figure before the key appeared.

Results and Discussion

Subjects’ individual data were first analyzed for internal consistency. Responses to the first 12 instances of each figure were correlated to those for the last 12. Prior to the experiment we chose to reject data from any subject whose correlation was less than 0.5. One subject’s data was eliminated. The diagonal cut responses were analyzed first. As in Experiment 3, subjects made few diagonal cuts and did so primarily when the distance ratio and area ratio were both 1. Again, we chose to analyze only the percentage of vertical cuts.

Linear Regression Analysis. Each subjects’ data was fitted to several linear models as in Experiment 2. We performed a *log* transform of each factor (distance ratio, area ratio and horizontal protrusion) for the same reason as in Experiment 2 (see footnote 11). The percentage of vertical cuts, c_v , was transformed to a new variable, c'_v , using the arc sine function in Equation 1 to improve the normality of its distribution. Model comparisons revealed that the horizontal protrusion was a significant factor for only 2 of the 12 subjects. Therefore, the final model used was:

$$c'_v = \beta_0 + \beta_1 \log(d) + \beta_2 \log(A).$$

Figure 33 summarizes the individual fits to the first model by plotting the coefficients for the distance and area ratios. Coefficients for the distance ratios are plotted on the abscissa, and coefficients for the area ratio are plotted on the ordinate. Error bars are given for 95% confidence intervals. Factors whose confidence intervals include 0 are not significant. The r^2 for the model is also given near each plot point. As can be seen, the models are quite predictive for almost all subjects.




Figure 33. Plot of the coefficients β_1 and β_2 for the linear model, $c'_v = \beta_0 + \beta_1 \log(d) + \beta_2 \log(A)$ in Experiment 4. Each point represents a different subject. Also shown are 95% confidence intervals, and r^2 -values for the model.




Figure 34. An example of a smoothed elbow used in Experiment 5. The elbows were created by two hyperbolic functions $y_1 = \frac{\epsilon}{x}$ and $y_2 = \frac{\epsilon}{x+\beta} + \lambda$.

Average r^2 for the distance and area ratio model was 0.772, $SD = 0.199$. This includes one outlying r^2 of 0.197. Without this score, the average rises to 0.829.

The distance ratio was by far the most important factor, with significant coefficients ranging from 0.4 to 2.2. The coefficients are all positive, indicating that subjects prefer shorter cuts as predicted.

Experiment 5

In Experiment 5 we sought to determine if the short-cut rule still holds when part-boundaries are smooth (not cusps as in all our previous experiments). We smoothed the elbows of Experiments 3 and 4 by using two hyperbolic functions with identical curvature, as illustrated in Figure 34.¹² The distance and area ratios can therefore be precisely computed.¹³ Smoothing the elbows has an interesting effect: the positive maximum becomes the unique point locally symmetric to the negative minimum (see Figure 22c). By contrast, in the cusp case (see Figure 22b), the negative minimum, g , is locally symmetric to all points between h and j . We chose to study smoothing on elbows rather than on crosses because smoothing crosses does not change the local symmetry relationships between the part boundaries. Horizontal protrusion remained the same, $\frac{w}{y}$. However, because this variable proved insignificant in Experiments 1, 2, 3 and 4, we chose to hold this variable constant at a value of 2. This allowed for a larger factorial design with the remaining variables. In addition to the distance ratio and area ratio, we included a variable for level of curvature, κ , with the values of 'high', 'low', and 'cusp' (or infinite).

Method

Participants. Ten graduate students volunteered to participate. Subjects were not paid and received no course credit. Four of these subjects had previously participated in either Experiment 4 or Experiment 2, but had not been debriefed.

Design and Procedure. Testing and design were as in Experiment 4. This experiment was a $3 \times 3 \times 3$ within-subjects factorial design. The three independent variables of the experiment were as follows: (i) The distance ratio, d , (ii) the area ratio, A , and (iii) the level of curvature, κ . Distance and area ratios were the same as in Experiments 3 and 4, viz., levels of $\frac{1}{2}$, 1, and 2. The horizontal protrusion was held constant at 2, a value previously used in Experiments 3 and 4. Figure 35 illustrates all the stimuli that were used for Experiment 5. The forced choices for cut types remained the

¹² The functions were as follows:

$$y_1 = \frac{\epsilon}{x} \quad \text{and} \quad y_2 = \frac{\epsilon}{x+\beta} + \lambda.$$

In these figures the negative minimum occurs at point:

$$a = (\sqrt{\epsilon} + \beta, \sqrt{\epsilon} + \lambda),$$

and the positive maximum at:

$$b = (\sqrt{\epsilon}, \sqrt{\epsilon}).$$

¹³ The functions were:

$$\text{Distance Ratio} = \frac{\sqrt{\epsilon} + \beta}{\sqrt{\epsilon} + \lambda}, \quad \text{and}$$

$$\text{Area Ratio} = \frac{\int_{\sqrt{\epsilon} + \lambda}^z \frac{\epsilon}{\sqrt{\epsilon} + y} + \lambda + \frac{\epsilon}{\sqrt{\epsilon} + y + \beta} dy}{\int_{\sqrt{\epsilon} + \beta}^w \frac{\epsilon}{\sqrt{\epsilon} + x} + \beta + \frac{\epsilon}{\sqrt{\epsilon} + x + \lambda} dx}.$$




Figure 35. All elbow stimuli used in Experiment 5.

same as in Experiments 3 and 4: horizontal, vertical, and diagonal. Each cut originated at the negative minimum of curvature. Diagonal cuts were joined to the positive maximum, horizontal cuts were drawn parallel to the horizontal axis (of the picture plane), and vertical cuts were drawn parallel to the vertical axis. Figures were then randomly rotated between positive 15 – 35 degrees and negative 15 – 35 degrees to minimize orientation effects. Figures were also presented mirrored about the vertical axis to control for left/right bias. The display of horizontal, vertical and diagonal choices in the key were counterbalanced across the presentation of each figure. We presented subjects with each figure 24 times, 12 mirrored and 12 non-mirrored for a total of 648 trials. Figures were displayed on a Macintosh Quadra 840 AV running SuperLab for Macintosh.

Results and Discussion

Subjects' data were initially analyzed for internal consistency. Data for the first 12 instances of each figure were correlated to data for the last 12. Prior to the experiment we chose to reject all data from any subject whose correlation




Figure 36. Plot of the coefficients β_1 and β_2 for the linear model, $c'_v = \beta_0 + \beta_1 \log(d) + \beta_2 \log(A)$ in Experiment 5. Each point represents a different subject. Also shown are 95% confidence intervals, and r^2 -values for the model.

was less than 0.5. No data were eliminated.

The diagonal cut responses were analyzed first. Only 2 of 10 subjects made more than 3% diagonal cuts. As in Experiments 3 and 4, diagonal cuts were made primarily when the distance ratio and area ratio were both 1.

Linear Regression Analysis. Subjects' data were fitted to the following model:

$$c'_v = \beta_0 + \beta_1 \log(d) + \beta_2 \log(A).$$

The term for curvature was not found to be significant for any subject and was therefore not included in the final. As in previous modeling, an arc sine transform was performed on the percentage of vertical cuts. Figure 36 summarizes the individual fits to the vertical cuts model by plotting the coefficients for the distance-ratio and area-ratio terms. Coefficients for the distance ratio are plotted on the abscissa, and coefficients for the area ratio are plotted on the ordinate. Error bars are plotted for 95% confidence intervals. Factors whose confidence intervals include 0 are not significant. The r^2 for the model is also given near each plot point. This graph demonstrates that, even for the smoothed elbows, subjects preferred to make shorter cuts.

Concluding Remarks

Human vision constructs visual objects, including their shapes and surface properties (Hoffman, in press; Singh & Hoffman, 1997). Decomposing these shapes into parts facilitates the recognition and manipulation of objects. To date

there is no comprehensive theory of shape parsing. The minima rule, for instance, gives precise boundary points at which to parse, but does not state how to join these points to form part cuts. In this paper, we examine a geometric rule for determining part cuts: the short-cut rule. This rule follows from the geometry of transversal intersections in three dimensions, and is supported by our experiments. In experiments with crosses, subjects prefer to join negative minima of curvature so as to produce the shortest cuts possible. In experiments with elbows, subjects prefer the shortest cut passing through the single negative minimum of curvature; thus the short-cut rule not only joins existing boundary points given by the minima rule, but it also creates new boundary points that are not negative minima of curvature. Indeed, as shown in Figure 37a, the short-cut rule can dictate a natural cut which terminates at neither end in a negative minimum of curvature (Siddiqi & Kimia, 1995); the short-cut rule pulls the boundary points of the cut slightly away from negative minima of curvature that have low salience. Sharp negative minima of curvature are stronger attractors of part cuts than weak negative minima (Hoffman & Singh, 1997). So while the short-cut rule might pull the boundary points of a part cut away from weak negative minima, as in Figure 37a, it might not when the negative minima are sharp, as in Figure 37b. The interaction between the short-cut rule and the salience of negative minima is an interesting topic for further research. Another example of this interaction can be seen in the dumbbell of Figure 16b, which has a neck in the middle of the shape. The center region of the dumbbell is composed of concave arcs of circles. Since each concave arc of a circle is a region of negative minima of curvature,¹⁴ the minima rule does not, by itself, pick out a unique point of the arc as a part boundary. Moreover, the magnitude of (normalized) curvature in these arcs is low, so that they have low salience as part boundaries. At the ends of each arc are sharp negative minima of curvature, which have high salience. Thus salience would dictate that part cuts should join these sharp points (cutting off the ends of the dumbbell), whereas the short-cut rule would dictate a part cut in the middle of the dumbbell. In this competition the short-cut rule loses (see, e.g., Braunstein et al., 1989, figure 6); subjects prefer to cut at the salient sharp points. In other competitions, salience might lose to the short-cut rule. Salience and the short-cut rule can also agree, rather than compete, as in Figure 16c. Here the part cut indicated at the top of the shape is longer and connects negative minima with weaker salience, whereas the part cut indicated at the bottom is shorter and connects negative minima of higher (normalized) curvature and therefore higher salience. As a result, subjects prefer the bottom cut (Siddiqi, Tresness, & Kimia, 1996).

Our motivation for the short-cut rule can also motivate Petter's (1956) rule. Recall that the short-cut rule can be related to Petter's rule because, given a chromatically homogeneous silhouette, it is initially ambiguous whether the silhouette arises from a single 3D object with parts, or from two different 3D objects separated in space. Hence, given such a silhouette, human vision finds the negative minima of curvature and pairs them using the short-cut rule. De-




Figure 37. (a) An example of a part cut in which neither end is a negative minimum of curvature (Siddiqi & Kimia, 1995). This example shows that the short-cut rule can sometimes pull the endpoints of a part cut slightly away from negative minima of curvature that have low salience (Hoffman & Singh, 1997). However, as shown in (b), the short-cut rule might fail to pull the endpoints of part cuts away from negative minima of curvature that have high salience.

pending on further processing, these pairings may become either part cuts or modal completions. Under this assumption, the results of our experiments suggest that the “larger means closer” heuristic (Shipley & Kellman, 1992; Stoner & Albright, 1993) fails to provide an adequate motivation for Petter's rule: Recall that with the crosses, the distance ratio consistently predicted subjects' preferred choices of cuts, whereas the area ratio failed to do so. However, further work is required to confirm this result for modal completions.

The short-cut rule and the salience of negative minima of curvature together determine many part cuts on silhouettes; but there are undoubtedly other factors as well. Part boundaries that are locally symmetric to each other might be better candidates for part cuts. Cuts that have “good continuation” with tangents at part boundaries might be preferred to those that do not (Kellman & Shipley, 1991; Siddiqi et al., 1996; Singh & Hoffman, in press). The orientation of a silhouette might affect part cuts. One set of cuts might be preferred because it yields fewer parts, or parts with no negative minima on their contours, or just simpler descriptions. And for some shapes, subjects might see no part cuts at all. Future work on part cuts will need to address these issues.

References

- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*, 183–193.
- Attneave, F. (1971). Multistability in perception. *Scientific American*, *225*, 63–71.
- 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). Obligatory edge assignment in vision – the role of figure and part segmentation in symmetry

¹⁴ Circular concavities are nongeneric: Almost any infinitesimal perturbation of such a concavity induces in it a locally unique negative minimum of curvature. So examples like Figure 16b almost surely never appear in nature. Hence, it is not surprising that human treats circular concavities no differently than negative minima of curvature.

- detection. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 1323–1342.
- Bennett, B. M., & Hoffman, D. D. (1987). Shape decompositions for visual shape recognition: The role of transversality. In W. A. Richards & S. Ullman (Eds.), *Image understanding* (pp. 215–256). Norwood, New Jersey: Ablex.
- 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. (1990). Higher-level vision. In D. N. Osherson & J. M. Kosslyn, S. M. and Hollerbach (Eds.), *Visual cognition and action: An invitation to cognitive science, vol.2* (pp. 41–72). Cambridge, MA: MIT Press.
- Biederman, I., & Cooper, E. E. (1991). Priming contour-deleted images: evidence for intermediate representations in visual object recognition. *Cognitive Psychology*, 23, 393–419.
- Biederman, I., Glass, A. L., & Stacy, J., E. W. (1973). Scanning for objects in real-world scenes. *Journal of Experimental Psychology*, 97, 22–27.
- Binford, T. O. (1971, December). Visual perception by computer. *IEEE Systems Science and Cybernetics Conference, Miami, FL*.
- 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., & Humphreys, G. W. (1994). Recognizing objects and faces. In V. Bruce & G. W. Humphreys (Eds.), *Object and face recognition* (pp. 141–180). East Sussex: Lawrence Erlbaum Associates Ltd.
- 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.
- Farah, M. (1996). Is face recognition ‘special’? evidence from neuropsychology. *Behavioral Brain Research*, 368, 542–545.
- Freeman, W. T. (1994). The generic viewpoint assumption in a framework for visual perception. *Nature*, 76, 181–189.
- Guillemin, V., & Pollack, A. (1974). *Differential topology*. Englewood Cliffs, NJ: Prentice-Hall.
- Guzman, A. (1971). Analysis of curved line drawings using context and global information. In *Machine intelligence*, 6. Edinburgh: Edinburgh University Press.
- Halmos, P. R. (1950). *Measure theory*. New York: Van Nostrand.
- Hoffman, D. D. (1983a). *Representing shapes for visual recognition*. Doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA.
- Hoffman, D. D. (1983b). The interpretation of visual illusions. *Scientific American*, 249, 154–162.
- Hoffman, D. D. (in press). *Visual intelligence*. 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.
- Kanizsa, G. (1979). *Organization in vision: essays on gestalt perception*. New York: Praeger.
- Kendall, M. G., & Stuart, A. (1963). *The advanced theory of statistics*. New York: Hafner Publishing Company.
- Kimia, B. B., Tannenbaum, A. R., & Zucker, S. W. (1991). Entropy scale-space. In (pp. 333–344). New York: Plenum Press.
- Leyton, M. (1992). *Symmetry, causality, mind*. Cambridge, MA: MIT Press.
- Lowe, D. (1985). *Perceptual organization and visual recognition*. Amsterdam: Kluwer.
- Mach, E. (1885/1959). *The analysis of sensations, and the relation of the physical to the psychical*. New York: Dover.
- Marr, D. (1977). Analysis of occluding contour. *Proceedings of the Royal Society of London, B*, 197, 441–475.
- Marr, D. (1982). *Vision*. San Francisco: Freeman Press.
- Marr, D., & Nishihara, H. K. (1978). Representation and recognition of three-dimensional shapes. *Proceedings of the Royal Society of London, B*, 200, 269–294.
- Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, 9, 441–474.
- Parent, P., & Zucker, S. W. (1989). Trace inference, curvature consistency and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11, 823–839.
- 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, 213–217.
- Pinker, S. (1985). Visual cognition: An introduction. In S. Pinker (Ed.), *Visual cognition* (pp. 1–63). Cambridge, MA: MIT Press.
- Richards, W. A., Dawson, B., & Whittington, D. (1986). Encoding contour shape by curvature extrema. *Journal of the Optical Society of America, A*, 3, 1483–1489.
- Richards, W. A., & Hoffman, D. D. (1985). Codon constraints on closed 2d shapes. *Computer Vision, Graphics, and Image Processing*, 31, 156–177.
- Richards, W. A., Koenderink, J., & Hoffman, D. D. (1987). Inferring three-dimensional shapes from two-dimensional silhouettes. *Journal of the Optical Society of America, A*, 4, 1168–1175.
- Roberts, L. (1965). Machine perception of three-dimensional solids. In J. T. Tippett et al. (Eds.), *Optical and electrooptical information processing* (pp. 125–143). Cambridge, MA: MIT Press.
- Rom, H., & Medioni, G. (1993). Hierarchical decomposition and axial shape description. *IEEE Transactions on PAMI*, 11, 823–839.
- Saidpour, A. (1983a). *The saliency of part boundaries in object decomposition*. Doctoral dissertation, University of California, Irvine, Irvine, CA, USA.
- Schyns, P. G., Goldstone, R. L., & Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, 21, 1–54.
- Schyns, P. G., & Murphy, G. L. (1994). The ontogeny of part representation in object concepts. In D. L. Medin (Ed.), *The psychology of learning and motivation: Advances in research and theory*, 31 (pp. 305–349). San Diego, CA: Academic Press, Inc.
- Schyns, P. G., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 681–696.

- Shipley, T. F., & Kellman, P. J. (1992b). Perception of partly occluded objects and illusory figures: Evidence for an identity hypothesis. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 106–120.
- 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., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: psychophysical aspects. *Perception*, *25*, 399–424.
- Singh, M., & Hoffman, D. D. (1997). Constructing and representing visual objects. *Trends in Cognitive Sciences*, *1*, 98–102.
- Singh, M., & Hoffman, D. D. (in press). Part boundaries alter the perception of transparency. *Psychological Science*.
- Singh, M., & Landau, B. (1998). Parts of visual shape as primitives for categorization. *Behavioral and Brain Sciences*, *21*, 36–37.
- 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 in vision. *Journal of Cognitive Neuroscience*, *5*, 129–149.
- Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology*, *46*, 225–245.
- Tommasi, L., Bressan, P., & Vallortigara, G. (1995). Solving occlusion indeterminacy in chromatically homogeneous patterns. *Perception*, *24*, 391–403.
- Turk, M., & Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, *3*, 71–86.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, *113*, 169–193.
- 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. 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., & Bennett, S. C. (1997). Preattentive object files: shapeless bundles of basic features. *Vision Research*, *37*, 25–43.
- Yuille, A. L. (1991). Deformable templates for face recognition. *Journal of Cognitive Neuroscience*, *3*, 59–70.