Description of solid shape and its inference from occluding contours

Jack M. H. Beusmans

Department of Information and Computer Science, University of California, Irvine, California 92717

Donald D. Hoffman

School of Social Sciences, University of California, Irvine, California 92717

Bruce M. Bennett

Department of Mathematics, University of California, Irvine, California 92717

Received May 28, 1986; accepted January 7, 1987

We explore a method of representing solid shape that is useful for visual recognition. We assume that complex shapes are constructed from convex, compact shapes and that construction involves three operations: solid union (to form humps), solid subtraction (to leave dents), and smoothing (to remove discontinuities). The boundaries between shapes joined through these operations are contours of extrema of a principal curvature. Complex objects can be decomposed along these boundaries into convex shapes, the so-called parts. We suggest that this decomposition into parts forms the basis for a shape memory. We show that the part boundaries of an object can be inferred from its occluding contours, at least up to a number of ambiguities.

1. INTRODUCTION

On the basis of studies showing selective impairment in object recognition in persons with unilateral cerebral lesions, Warrington and Taylor suggested that recognition is a two-stage categorization process.[1,2] Persons with right-hemisphere lesions have difficulty in deciding whether pictures representing different views of an object are in fact showing the same physical object. Persons with left-hemisphere lesions are able to classify different aspects of an object as belonging to the same physical stimulus, but their ability to attach meaning to their percept is impaired. This suggests that in the first stage of visual recognition the image is categorized on perceptual grounds only, whereas a perceptual category is given semantic content in the second stage. In this paper we will focus on the perceptual categorization stage and explore a method for representing solid shape, a method useful for general-purpose visual recognition.

We will focus on symbolic representations of shape, which aim at a structured description of complex objects.[3,4] We consider a representation to be structured if it treats a complex object as a configuration of irreducible elements. Thus, one could choose an *a priori* set of shape primitives and use them to approximate a complex object. Blum, for example, introduced the symmetric axis transform, in which two-dimensional (2-D) contours are approximated by a collection of maximal disks, whose centers compose the so-called symmetric axis of the shape.[5] A shape can then be segmented on the basis of the behavior of the symmetric axis and the way in which the contour changes its shape relative to the symmetric axis. This approach has been extended to three-dimensional (3-D) objects by approximating their shape by spheres.[6–8] A more general and higher-level shape primitive is the generalized cone or cylinder.[9] In the vision system ACRONYM, for instance, airplanes are modeled in terms of a set of generalized cones.[10] In this primitive-based approach to shape representation, the description of an object in terms of shape primitives determines its decomposition into parts, either directly as with generalized cones or indirectly as with the disks and spheres in the symmetric axis transform. If objects do indeed consist of generalized cones or some other primitives, as might be the case with industrial products, then this approach results in a satisfactory representation. In general, however, there will always be objects whose shape defies effective description in terms of some set of primitives. And extending the set of primitives does not provide for a more principled theory of shape representation.

Instead of defining parts by their shapes, one could define parts by their boundaries.[11,12] But which class of curves on the surface of an object should one use as boundaries between parts? One choice can be motivated by considering two arbitrary objects that are made to interpenetrate (Fig. 1). In general, the new object will have a closed contour of concave discontinuity, the part boundary, separating the two objects that were joined [Fig. 1(b)]. Thus we can decompose composite objects along contours of concave discontinuity or their smoothed derivatives, contours of negative minima of a principal curvature. In the present paper we will explore this boundary-based approach to shape representation. We will define some additional part boundaries, enabling us to decompose objects into convex, compact shapes—the so-called parts. We refer to such a decomposition as the deep structure of an object, and we suggest a shape memory based on deep structures.

0740-3232/87/071155-13$02.00

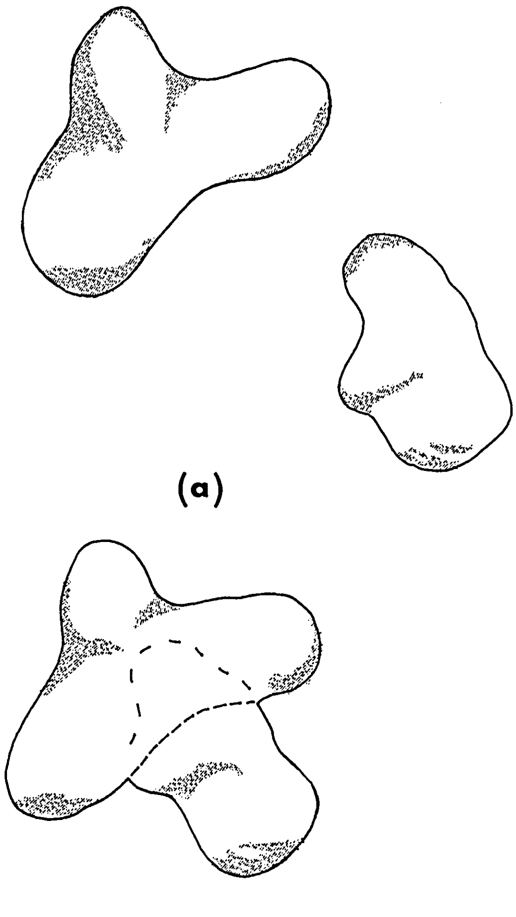

Fig. 1. Whenever an object penetrates another object, the contour of intersection (– –) will be a contour of concave discontinuity of the tangent plane with a probability of 1. Two objects are shown before (a) and after (b) penetration.

Human observers can often recognize familiar objects from their contours or silhouettes, indicating that these 2-D outlines convey sufficient information to distinguish between different 3-D objects. What does this mean in the context of a shape memory built on deep structures? It means that the human visual system can infer deep structure from contours. Koenderink and van Doorn showed that the shape of the occluding contour conveys information about surface curvature at the corresponding points on the surface.[13,14] In this paper we will study how this knowledge of surface curvature can be used to derive deep structures from occluding contours.

The remainder of this paper is organized as follows. In Section 2 we present elements of a theory of 3-D shape representation that is useful for visual recognition. In Section 3 we show how to infer part boundaries and thereby deep structures from the occluding contours of an object. Finally, in Section 4 we discuss the proposed theory.

2. A REPRESENTATION OF THREE-DIMENSIONAL SHAPE

A. Regularities and Vision

It is the task of a general-purpose vision system to assign the most plausible interpretation to images of the external 3-D world. However, any given image can correspond to infinitely many states of the environment. For example, the image of a cylinder under orthographic projection in the direction of its length axis will always be a circle, regardless of the length of the cylinder. In other words, we have a so-called ill-posed problem; i.e., the available input does not warrant a unique interpretation. Formulating what extra information can be used to obtain a well-posed problem has been the focus of much of the recent work on vision.[9]

The extra information needed to interpret an image is usually expressed in terms of constraints on, for instance, the shape or possible deformations of objects in the 3-D world. Marr and Poggio based their theory of stereopsis on the assumptions that a physical marker has a unique position in space at any one time and that matter is cohesive.[15] To recover the shape of an object from a sequence of images, a process variously known as structure-from-motion or kinetic-depth-effect, one can assume that objects are globally rigid[16] or locally rigid.[17]

The problem of representing 3-D objects in terms of parts is similarly ill posed, since there are infinitely many ways in which a complex object could be decomposed.

B. Transversality and Parts

The question we now pose is whether there exist regularities in the structure of 3-D objects that would allow a meaningful definition of the term part of an object. Consider a possible genesis of a composite object, an example of which is given in Fig. 1. A composite object could be obtained by penetrating one object with another, possibly followed by a smoothing of the contour of intersection. This is not to say that the composite actually arose in this manner, an obvious counterexample being limbs of animals. Our only concern is to find a principled representation of the shape of an object. Thus the composite object can be thought to consist of two objects: the penetrant and the penetrated object, separated by a contour of intersection. Since the original objects were chosen arbitrarily, we cannot use their shape to identify the parts of the composite object, and the question becomes whether the contour of intersection has some distinguishing property. The answer is yes and is stated in the following transversality principle:

> *Whenever two arbitrarily shaped, smooth surfaces are made to interpenetrate, they intersect transversally with a probability of 1.*

This means that, generically, at a given point along the contour of intersection, the tangent planes to the penetrating surfaces differ. In other words, this contour will be a contour of discontinuity of the tangent planes (Fig. 1). Assuming the object to be the figure and positive surface normals to be oriented such that they point into the figure, the contour will be a contour of concave discontinuity. It has been proved that smoothing of such discontinuities gives rise to contours of negative minima of a principal curvature.[18] (Readers not familiar with differential geometry can find brief descriptions of relevant terms in Appendix A). One can thus propose that smooth surfaces be partitioned along contours of negative extrema of a principal curvature. This partitioning is unique, and the resulting representation is unambiguous. The sequence of penetrations of objects that produced the composite object is, of course, not uniquely defined.

Generic penetrations, however, do not generate all possible shapes; counterexamples are elbows and dents. In case of an elbow, the intuitively appealing part boundary is an open contour of concave discontinuity. A dent, on the other hand, is separated from the surrounding surface by a closed contour of convex discontinuity. We include these cases by introducing the operations of nongeneric penetration and subtraction.

C. Aspects of a Theory of Shape Representation

In the previous subsection we discussed the term "part of an object." We now formalize that discussion and present the framework for a theory of 3-D shape representation based on the notion of parts.

Any well-defined scheme to decompose complex objects into simpler ones has to terminate; i.e., it has to recognize certain objects as irreducible. We assume that these irreducible objects are convex and compact, in other words, objects without any dents or depressions. By assuming certain shapes to be irreducible or primitive, our approach does not suddenly become primitive based; we are not going to look for convex shapes in a complex object. Instead we will decompose objects along certain part boundaries, a decomposition that terminates with convex, compact objects. We further assume that objects can be combined only through solid union and solid subtraction and that the resulting discontinuities can be smoothed. Thus we postulate the following:

Axiom 2.1
The 3-D shape of a smooth object can be described in terms of primitive shapes or parts, viz., convex and compact objects.

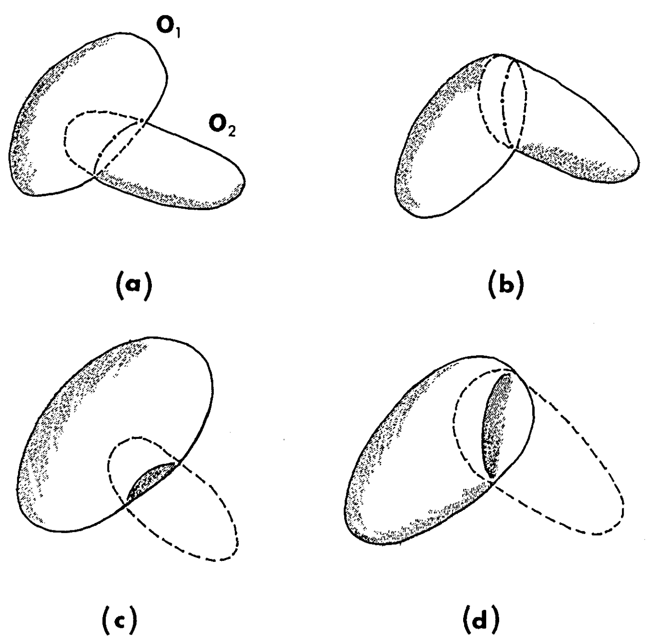

Fig. 2. Primitives can be combined through (a) generic penetration ($O_1 \cup O_2$), resulting in a closed contour of concave discontinuity (– • –); (b) nongeneric penetration ($O_1 \cup O_2$), resulting in an open contour of concave discontinuity; (c) generic subtraction ($O_1 - O_2$), giving rise to a closed contour of convex discontinuity; and nongeneric subtraction (d), which will yield the same result as (c). Negative parts are indicated by dashed outlines.

Axiom 2.2
Parts can be combined only through generic or nongeneric solid union or solid subtraction. Only these operations give rise to contours of convex and concave discontinuities or to their smoothed derivatives, contours of positive maxima and negative minima of a principal curvature, respectively.

These two axioms form the foundation for our theory of shape representation. We next define solid union and solid subtraction precisely and study the part boundaries that they cause.

Let O_i be solid objects, and let the surface of each O_i, denoted surf(O_i), be the zero-level set of an appropriate function $F_i(\mathbf{x})$, where $\mathbf{x} = (x, y, z) \in \mathcal{R}^3$. The interior of O_i, indicated by int(O_i), is defined as $\{\mathbf{x} \in \mathcal{R}^3 | f_i(\mathbf{x}) < 0\}$. For example, $f_i(\mathbf{x}) = x^2 + y^2 + z^2 - 1$ defines a sphere with a radius of 1, whose surface consists of those points for which f_i is zero.

Definition 2.3
The interpenetration or solid union of O_1 and O_2 is defined as follows:

$$O_3 = O_1 \cup O_2 = \{\mathbf{x} \in \mathcal{R}^3 | \mathbf{x} \in O_1 \text{ or } \mathbf{x} \in O_2\} = \{\mathbf{x} \in \mathcal{R}^3 | f_1(\mathbf{x}) \le 0 \text{ or } f_2(\mathbf{x}) \le 0\}. \quad (1)$$

The surface of the new, composite object O_3 is then

$$\text{surf}(O_3) = \text{surf}(O_1 \cup O_2) = \{\mathbf{x} \in \mathcal{R}^3 | \mathbf{x} \in \text{surf}(O_1) \text{ and } \mathbf{x} \notin \text{int}(O_2) \text{ or } \mathbf{x} \in \text{surf}(O_2) \text{ and } \mathbf{x} \notin \text{int}(O_1)\}, \quad (2)$$

or, in terms of the functions f_i,

$$\text{surf}(O_3) = \text{surf}(O_1 \cup O_2) = \{\mathbf{x} \in \mathcal{R}^3 | f_1(\mathbf{x}) = 0 \text{ and } f_2(\mathbf{x}) \ge 0 \text{ or } f_2(\mathbf{x}) = 0 \text{ and } f_1(\mathbf{x}) \ge 0\}. \quad (3)$$

Generically, i.e., with a probability of 1, the contour of intersection of the surfaces of O_1 and O_2 will be a closed contour of concave discontinuity. Bennett and Hoffman showed that smoothing the contour of concave discontinuity results in a contour along which the surface has a locally largest negative principal curvature [Fig. 2(a)].[18] Note that this contour is located in a hyperbolic (saddle-shaped) region of the surface. For nongeneric interpenetrations, the contour of intersection has an interval where the tangent planes of O_1 and O_2 are parallel, resulting in an open contour of concave discontinuity. The shape of the composite object then resembles an elbow [Fig. 2(b)].

Removing O_1 from the composite object O_3 leaves a dent bounded by a contour of convex discontinuity. We can express this operation in terms of closed-set solid subtraction of O_1 from O_2 as follows.

Definition 2.4
The closed-set solid subtraction of O_1 from O_2 is defined as follows:

$$O_4 = O_2 - O_1 = \{\mathbf{x} \in \mathcal{R}^3 | \mathbf{x} \in O_2 \text{ and } \mathbf{x} \notin \text{int}(O_1)\}$$
$$= \{\mathbf{x} \in \mathcal{R}^3 | f_2(\mathbf{x}) \leq 0 \text{ and } f_1(\mathbf{x}) \geq 0\}; \quad (4)$$

and the surface of the new object O_4 is defined as

$$\text{surf}(O_4) = \text{surf}(O_2 - O_1)$$
$$= \{\mathbf{x} \in \mathcal{R}^3 | \mathbf{x} \in \text{surf}(O_2) \text{ and } \mathbf{x} \notin \text{int}(O_1)$$
$$\text{or}$$
$$\mathbf{x} \in \text{surf}(O_1) \text{ and } \mathbf{x} \in O_2\}. \quad (5)$$

We refer to the subtracted part as the negative part to distinguish it from positive parts, which correspond to physical parts of the object. The surface of the composite object has a closed contour of convex discontinuity composed of the points satisfying $f_1(\mathbf{x}) = f_2(\mathbf{x}) = 0$. Smoothing the convex discontinuity gives rise to a contour along which the surface has a locally largest positive principal curvature [Fig. 2(c)]. Gaussian curvature along this contour can be positive (convex but not concave) or negative (saddle shaped), or it can alternate between the two. Note that even if the original penetration of the two convex objects had been nongeneric, subtraction results in a contour of convex discontinuity that is closed [Fig. 2(d)].

We defined two operations on solid objects, solid union and solid subtraction, and mentioned smoothing as an operation to remove discontinuities. Applying these operations results in contours of extrema of a principal curvature, as a consequence of the principle of transversality. We now use contours of extrema to partition an arbitrary composite object. For example, from the presence of a closed contour of negative minima of a principal curvature on the surface of an object, we infer that that contour resulted from solid union of two positive parts. Before we can give the complete rules for finding part boundaries, we have to introduce some terminology regarding the deformation of contours.[19]

Definition 2.5
A contour p is homotopic to a contour q if p can be continuously deformed into q. The deformation is called a homotopy from p to q.

Definition 2.6
A homotopy between contours p and q is permissible if none of the intermediate contours intersects a contour of extrema of a principal curvature other than p or q.

Definition 2.7
The following partitioning rules define part boundaries on a smooth surface of genus zero (that is, without any holes):

(1) A closed contour of negative minima of a principal curvature is a boundary between two positive parts (indicated B_c^-).

(2) An open contour of negative minima of a principal curvature is part of a boundary between two positive parts (B_0^-), unless the contour is located in a concave region.

(3) A closed contour of positive maxima of a principal curvature is a boundary between a positive and a negative part if the contour is permissibly homotopic to a contour lying within the hyperbolic region separating the two prospective parts (B_c^+).

(4) An open contour of positive maxima of a principal curvature is not a part boundary.

Rules (1) and (4) are self-explanatory. The qualification in rule (2) serves to exclude the case of a B_0^- at the bottom of a concavity or dent. The condition in rule (3) is based on the fact that smoothing the contour of convex discontinuity caused by solid subtraction always results in a hyperbolic region. Thus the edges of a cube that have been smoothed to form closed contours of positive maxima of a principal curvature are not B_c^+. This corresponds to the way in which a human observer appreciates the shape of a cube: The cube's edges are not considered to be part boundaries but are considered to be merely features of its shape.

Applying these rules to a complex object results in its decomposition into irreducible objects. What will the shape of these irreducible objects or parts be like? Assume for a moment that a part is convex except for a hyperbolic region. Consider the family of lines of curvature along which principal curvature changes from positive to negative and back to positive as we go from the convex to the hyperbolic region and back to the convex region. Obviously, principal curvature has at least one minimum value along these lines of curvature, resulting in contours of negative minima of a principal curvature, i.e., part boundaries, in the generic case. Consequently a part cannot contain any hyperbolic regions. By reasoning along similar lines, one can also show that a convex object with a concave dent can still be reduced, except possibly in the presence of umbilical points (at umbilics, the principal curvature is the same in all directions; for example, all points on a sphere are umbilical). If the convex and concave regions both have an umbilical point where principal curvatures attain their extreme values, there will be no contours of extrema of principal curvature along which to partition the surface.

Let us examine some applications of the partitioning rules and determine whether the human visual system uses them.

Example 2.8
Figure 3(a) shows a bonelike object having two B_c^-'s. We can therefore partition it into three positive parts, which are

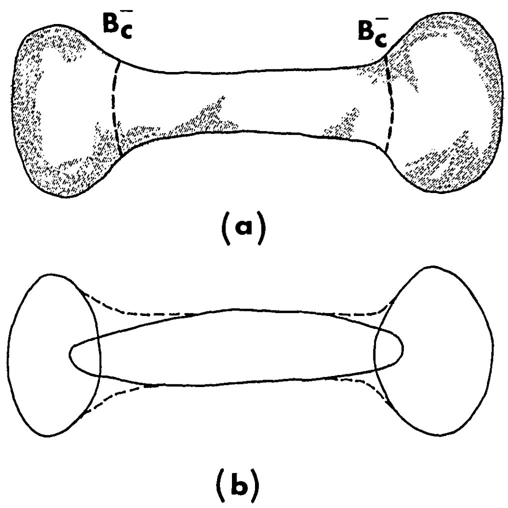

Fig. 3. (a) A bonelike object with two closed contours of negative minima of a principal curvature (B_c^-); (b) its decomposition into parts.

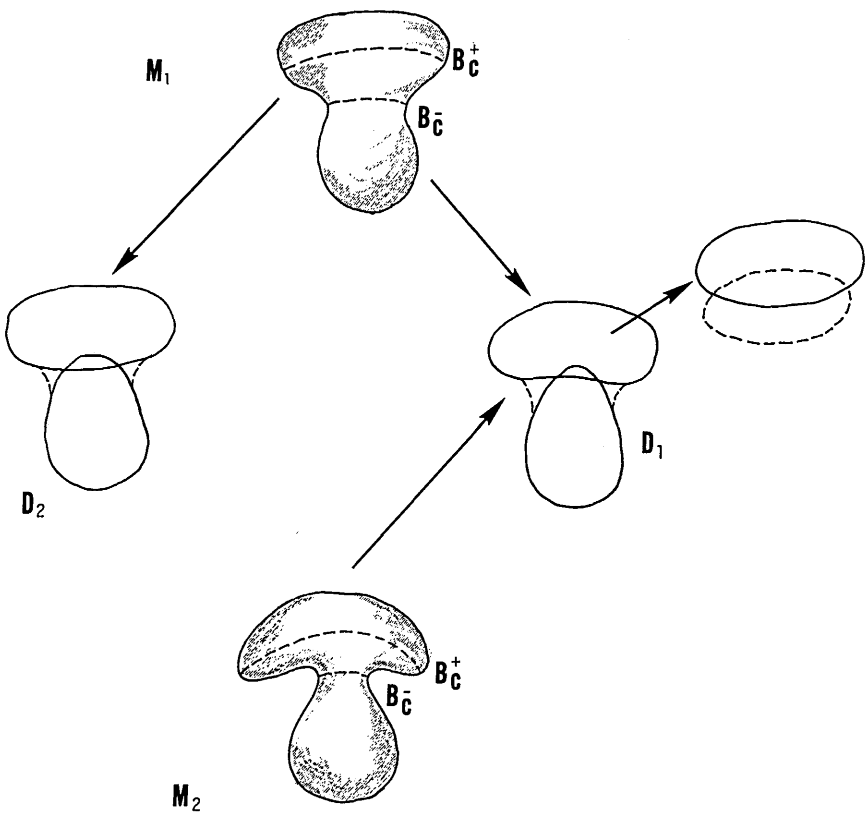

Fig. 4. Two mushroomlike shapes (M_1 and M_2). Depending on the sequence in which partitioning rules are applied, M_1 has two decompositions (D_1 and D_2), whereas M_2 has only one (D_1).

combined through solid union. The concave discontinuities caused by the solid unions have been smoothed to such an extent that one of the parts became hidden [Fig. 3(b)]. This example illustrates that in partitioning a complex object we do not simply look for convex objects.

Example 2.9
The partitioning rules do not always result in a decomposition that corresponds to the human appreciation of an object's shape. Consider, for example, the two mushroomlike objects M_1 and M_2 in Fig. 4. Both have a B_c^- and a B_c^+; hence both can be decomposed into a positive part (the stalk of the mushroom) and a positive part with a negative part (the cap), resulting in decomposition D_1. But human observers would consider M_1 to have the structure D_2 rather than D_1. We can solve this problem by applying the partitioning rules in a certain sequence: First apply the rules for B_c^-'s and B_0^-'s until there are no more contours of negative minima of a principal curvature (resulting in a partial decomposition), and then apply the rules for B_c^+'s to the components of the partial decomposition. In case of M_1 the partial decomposition consists of two positive parts, one having a B_c^+. Rule (3) no longer applies to the latter component, as there is no permissible homotopy to a contour in a hyperbolic region. The partial decomposition of M_2, on the other hand, consists of a positive part and a positive part with a dent. Thus rule (3) still applies.

Example 2.10
Consider the cosine surface shown in Fig. 5. If it is assumed that we are looking down on the surface, partitioning rule (1) predicts negative minima part boundaries (B_c^-) whose locations are indicated by concentric lines on the surface. This decomposition into concentric ridges corresponds to the way in which a human observer would partition the surface. Rule (3), however, predicts another set of part boundaries, namely, the three contours of positive maxima (B_c^+), lying between the B_c^-'s. It seems then that the human visual system prefers B_c^- over B_c^+, at least under these circumstances. To incorporate this preference, the rule for B_c^+'s can be modified to include the necessary condition that there be no permissible homotopies to B_c^-'s on either side of

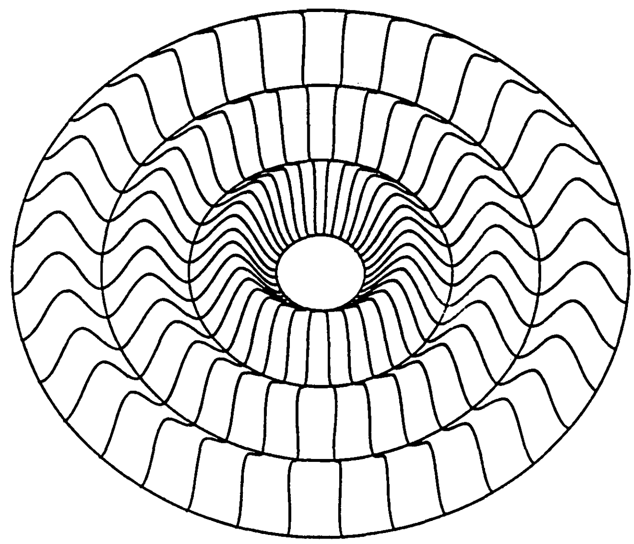

Fig. 5. Cosine surface obtained by rotating a cosine curve about the vertical axis. Assuming that we are looking down on the surface, it is decomposed into three concentric ridges, whose boundaries are indicated by concentric circles.

a tentative B_c^+. Note, however, that the resulting parts are not convex. A similar situation occurs with a torus, which does not have any part boundaries at all. We suspect that we can account for these cases in an extension of the existing theory that includes objects with holes (genus greater than zero). It is, for instance, possible that we would have to consider some nonconvex objects as irreducible.

In summary, we have defined rules to decompose an object into parts, thereby revealing its so-called deep structure. The deep structure of an object specifies its constituent parts and their interrelationships. It indicates, for example, that a dented object consists of a positive and a negative part and that these parts are related through solid subtraction. Note that the exact configuration of the two parts is not represented at the level of deep structure and that negative parts do not correspond to physical parts. By specifying the exact configuration of parts and the smoothing of the contours of discontinuity, the surface structure is obtained. Since these specifications transform deep structure into surface structure, they will be referred to as transformation rules. (We borrowed these terms from Chomsky's extended standard theory of syntax.[20,21] Shepard also used the terms deep structure and surface structure in the context of vision.[3] In his terminology, surface structure refers to the proximal stimulus, i.e., the image of an object, and deep structure refers to the internal representation of that object.) We can use the deep structure of objects to define equivalence classes. Objects in the same deep structure class can be differentiated through different transformation rules. In other words, we propose that objects be categorized first on the basis of their deep structure and then on the basis of their surface structure.

3. SURFACE SHAPE AND PART BOUNDARIES FROM OCCLUDING CONTOURS

In the previous section we developed a theory of shape representation based on the notion of parts that can be combined through penetration and subtraction. We found that the resulting boundaries between parts can be characterized in terms of surface curvature: Smoothed part boundaries are contours of extrema of a principal curvature. Thus, in order to decompose an object into its parts, we have to locate part boundaries, which, in turn, requires some knowledge of surface curvature. Surface curvature can be computed from a number of sources: e.g., depth information obtained through stereo[9,15]; structure-from-motion[16,17]; shading[22]; motion parallax[23]; and occluding contours.[14,24] Some of these sources are richer than others: motion parallax, shading, and occluding contours yield only the sign of Gaussian curvature, whereas the more detailed depth information obtained through stereo or structure-from-motion provides Gaussian curvature.

We will focus our attention on occluding contours because they constitute images that are minimal in the sense that only discontinuities of the mapping from surface to image are indicated and because the human visual system can recognize objects from silhouettes. By studying the minimal conditions in which the human visual system can still operate, we hope to gain the most pertinent insights into its functioning. We shall investigate what occluding contours tell us about surface curvature, and we shall formulate rules for locating part boundaries on the basis of that knowledge.

A. Shape along Folds

As before, we consider only smooth surfaces of genus zero (that is, without any holes). Since we assume objects to be opaque, certain parts of their surface are not visible from a particular vantage point. The curves on the surface separating visible from nonvisible surface patches are called folds [Fig. 6(a)]. Note that surface normals in visible regions point away from the observer (by convention, surface normals are oriented such that they point into compact objects). However, not all regions whose surface normals point away from the observer are in fact visible, because of the interposition of other objects or parts of the same object. A region will, however, always be invisible if its surface normals are pointing towards the observer. We can therefore define the fold locus as follows:

Definition 3.1

A fold is a locus of points on the surface where the unit surface normal $\mathbf{N}$ is orthogonal to a unit vector $\mathbf{V}$ in the direction of the line of sight; $\mathbf{N} \cdot \mathbf{V} = 0$.

Since we are considering only compact and smooth objects, folds form closed, smooth curves. By looking roughly in the direction of the length axis of, say, a pencil, it is immediately clear that folds are generally not planar or orthogonal to the line of sight [see also Fig. 6(a)]. This is not to say that the line of sight and the folds are unrelated. We show next that their directions are, in fact, conjugate directions, something that had already been observed by Koenderink.[13,14]

Proposition 3.2

The line of sight and the direction of folds are conjugate directions.

Proof

Let (x, y, z) be the Cartesian coordinate system centered at a point P on the fold [Fig. 6(a)]; x is along the line of sight, and z is in the direction opposite the surface normal $\mathbf{N}$ at P. The x–y plane is the tangent plane at P. Directions in the x–y plane are given by the direction numbers $x : y$. For example, $1 : 0$ specifies the direction along the x axis. We can approximate the surface $\mathbf{x}(x, y)$ by a Monge patch of the form $[x, y, f(x, y)]$, such that $\mathbf{x}_x = \mathbf{e}_1$, $\mathbf{x}_y = \mathbf{e}_2$, and $f(x, y) = \frac{1}{2}(ax^2 + 2bxy + cy^2)$, where $\mathbf{e}_1$ and $\mathbf{e}_2$ are unit vectors in the direction of the x and y-axes, respectively.

Let the direction of the fold at P be indicated by the direction numbers $\partial x : \partial y$, and let $\partial \mathbf{N} = \mathbf{N}_x \partial x + \mathbf{N}_y \partial y$. By definition $\mathbf{N} \cdot \mathbf{V} = \mathbf{N} \cdot \mathbf{e}_1 = 0$ along the fold, and thus $\partial(\mathbf{N} \cdot \mathbf{e}_1) = \partial \mathbf{N} \cdot \mathbf{e}_1 + \mathbf{N} \cdot \partial \mathbf{e}_1 = \partial \mathbf{N} \cdot \mathbf{e}_1 = 0$. Together with $\mathrm{d}\mathbf{x} = \mathbf{x}_x dx + \mathbf{x}_y dy = \mathbf{e}_1$ for $dx : dy = 1 : 0$ (line of sight), we then have $\partial \mathbf{N} \cdot \mathrm{d}\mathbf{x} = 0$; that is, the line of sight and the direction of folds are conjugate. Intuitively, the difference vector $\partial \mathbf{N}$ between normals at neighboring points on the fold has to be perpendicular to the line of sight, since the normals themselves are by definition perpendicular to the line of sight.

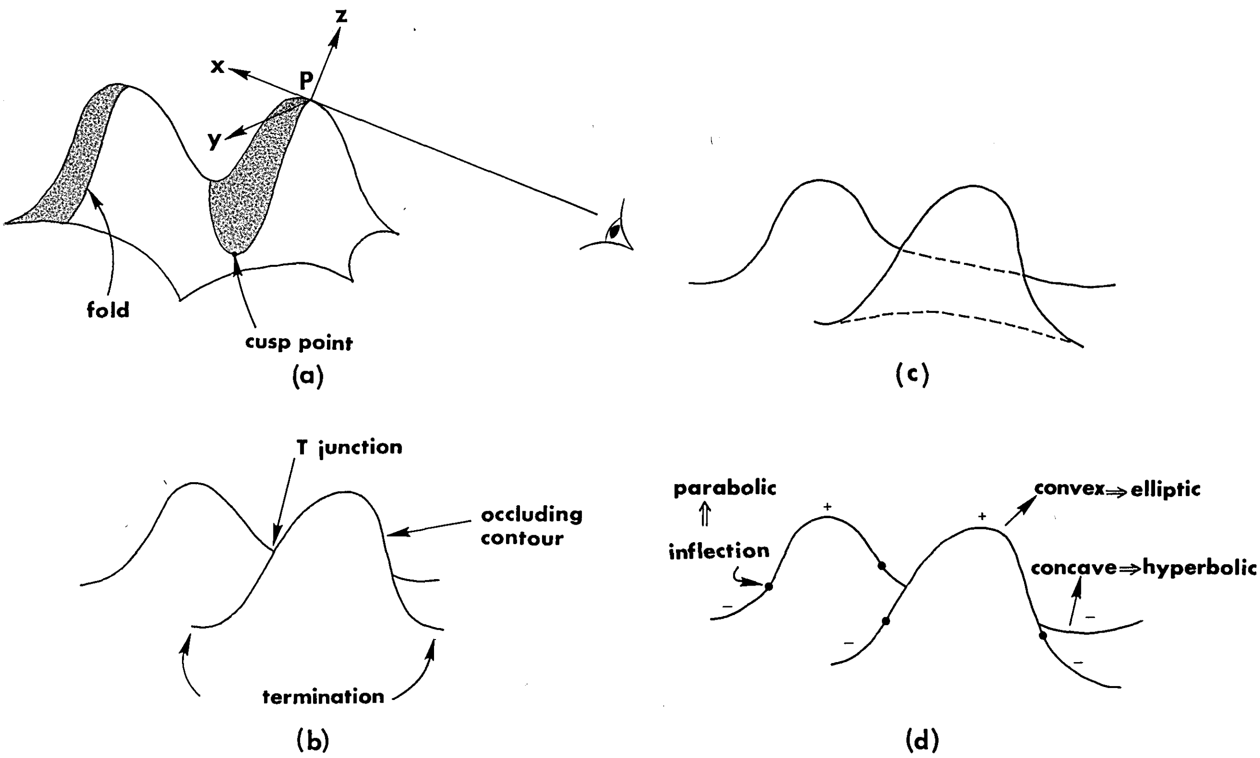

Fig. 6. (a) A 3-D surface, seen from a vantage point at infinity, whose visible and invisible (shaded areas) parts are separated by folds; at cusp points the line of sight is in the direction of the fold. (b) The image of (a) consists of the projections of folds and cusps, that is, occluding contours and terminations, respectively. Folds can be partially occluded by other parts of the surface, giving rise to T junctions in the image. (c) A possible completion of the occluding contours of (b). (d) The sign of curvature of the occluding contour equals the sign of Gaussian curvature at corresponding points on the fold.

Proposition 3.3

For points on a fold, the direction of that fold is the only direction conjugate to the line of sight, and therefore folds do not cross.

Proof

The direction $\partial x : \partial y$, conjugate to some arbitrary direction $dx : dy$, is given by the linear equation $(Ldx + Mdy)\partial x + (Mdx + Ndy)\partial y = 0$, where L, M, and N are the second fundamental coefficients. If we represent the surface as the Monge patch mentioned in Proposition 3.2, $L = a$, $M = b$, and $N = c$. By choosing $dx : dy = 1 : 0$ (the line of sight), we obtain the equation $a\partial x + b\partial y = 0$; that is, the equation for the direction $\partial x : \partial y$ conjugate to the line of sight. But because a cannot be zero along the fold (there would not be any fold otherwise), we conclude that, for points on a fold, there is only one direction conjugate to the line of sight. Since the direction of the fold itself is conjugate to the line of sight, the fold is the only such direction. Now suppose that two folds cross at a point P. Since the direction of folds and the line of sight are conjugate, we have a situation at P in which two directions are conjugate to the line of sight. Since we just showed that there can only be one such direction, we conclude that folds do not intersect.

Summarizing, folds are smooth, closed curves that do not intersect; that is, they are topologically equivalent to nested circles. Folds are projected onto occluding contours in the image [Fig. 6(b)]. As is clear from Figs. 6(a) and 6(b), occluding contours are not always closed; that is, the corresponding folds are not always completely visible. A contour can stop either at a T junction or at a termination [Fig. 6(b)]. In the case of T junctions, the corresponding fold is occluded by another object or a nonneighboring part of the same object. In the case of terminations, on the other hand, the fold does occlude neighboring parts of the same fold. To study the behavior of the contour in a neighborhood of a termination, imagine for a moment that the objects that we are looking at are transparent instead of opaque. Figure 6(c) gives one example of what we might see in that case (contours are no longer occluding). Since the mapping between surface and image has a cusp at terminations, the point on the surface that projects onto the termination is termed the cusp point. At cusp points the fold reverses direction; that is, the observer looks in the direction of the fold: the line of sight and the direction of the fold coincide [Fig. 6(a)]. Combining this with the above observation that line of sight and direction of the fold are conjugate, we conclude that at cusp points we are looking in self-conjugate or asymptotic directions. As asymptotic directions exist only on hyperbolic surface patches, so do cusp points.[24] Whitney showed that folds and cusps are the only stable singularities of mappings between smooth manifolds of dimension 2.[25] Since the image and the surfaces of 3-D objects are both 2-D manifolds, and projection is a mapping between manifolds, we need not be concerned with other singularities.

We now turn to the problem of inferring 3-D shape from occluding contours. Koenderink[14] showed that one can read off a valuable intrinsic property of the surface directly from its occluding contours:

Proposition 3.4
The sign of curvature of the occluding contour equals the sign of Gaussian curvature at the corresponding points on the fold.

That is, an elliptic arc of the contour (one that is convex toward the background) indicates an elliptic (convex) surface, a hyperbolic arc (one that is concave toward the background) indicates a hyperbolic surface, and an inflection point of the contour corresponds to a parabolic point [Fig. 6(d)]. Note that it is assumed that the occluding contour has been assigned an orientation and that it is known which side of the contour is the occluder and which side is being occluded. Also note that extrema of curvature of occluding contours do not necessarily indicate the presence of extrema of a principal curvature on the surface, i.e., possible part boundaries.

B. Shape between Folds

In the previous section we saw that the sign of Gaussian curvature at folds can be extracted from their projections, the occluding contours. We will now investigate how this curvature information can be propagated to regions of the surface between folds. But let us first study how surface regions with different signs of Gaussian curvature are laid out on the surface. Elliptic regions (positive Gaussian curvature) and hyperbolic regions (negative Gaussian curvature) are separated by parabolic lines (zero Gaussian curvature), which are therefore closed curves. Isolated parabolic points are not stable, since the slightest perturbation of the surface results in a closed contour separating elliptic and hyperbolic surface patches. Generically parabolic curves do not intersect themselves or each other.[26] Of course, folds can intersect parabolic lines; their visible intersections show up as inflections of the occluding contour.

We have now characterized two classes of curves on the surface: folds and parabolic lines. Members of each class form smooth, closed curves that do not intersect curves of the same class. But what can we say about the intersections between folds and parabolic lines? To begin with, the only stable intersections of curves on smooth surfaces are transversal intersections. Since folds and parabolic lines are topologically equivalent to circles, every intersection of a fold and a parabolic line has to be matched by another one. Folds will therefore intersect parabolic lines an even number of times; in other words, folds will have an even number of parabolic points. Not all these parabolic points might actually be visible from a particular vantage point; that is, not all these parabolic points show up as inflections of the occluding contours. It is, of course, also quite possible that none of the parabolic points is visible or that a parabolic line does not cross a fold to begin with. In that case we will have to infer the presence of these parabolic lines indirectly from the presence of other parabolic curves.

We can propagate curvature information from folds to regions between folds as follows. Consider a fold with n parabolic points. Each of these parabolic points is paired with another one. By pairing two parabolic points we assert that both lie on the same parabolic curve, and by pairing parabolic points we also pair their projections, the inflection points. Figure 7 shows the three possible pairings of a contour with four inflection points. (In the sequel we consider only pairings on one side of a fold; pairings on the other side

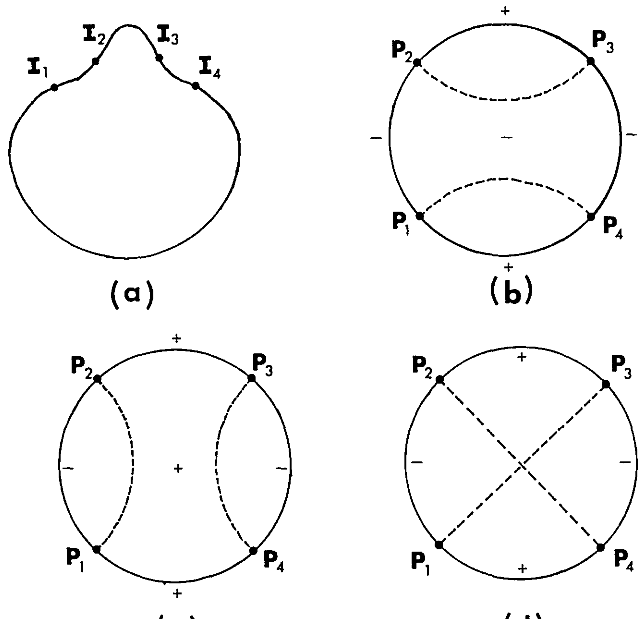

Fig. 7. (a) A contour having four inflection points (I_1, I_2, I_3, I_4). The corresponding four parabolic points (P_1, P_2, P_3, P_4) on the fold can be paired generically (b), (c) or nongenerically (d).

are similar.) These pairings are generic except for the one in Fig. 7(d). This pairing is nongeneric because the parabolic lines connecting the parabolic points on the fold intersect. As each pairing of parabolic points corresponds to a different completion of the partially known surface curvature, it is important to know how many different completions are in fact possible for a given number of parabolic points on a fold. We next derive expressions for the total number of pairings (generic plus nongeneric) and the number of generic pairings only.

Remark 3.5
The number of generic plus nongeneric pairings, t_n, for a fold with n parabolic points is given by

$$t_n = \frac{n!}{2^{n/2}(n/2)!}. \tag{6}$$

Proof
The number of generic plus nongeneric pairings, t_n, for a fold with n parabolic points equals $(n-1)(n-3)\ldots 3\cdot 1$, since the first parabolic point can choose among $n-1$ potential partners, the second can choose among $n-3$, etc. We can simplify this expression as follows:

$$\begin{aligned} t_n &= (n-1)(n-3)\ldots 5\cdot 3\cdot 1, \\ &= \frac{n!}{n(n-2)(n-4)\ldots 4\cdot 2}, \\ &= \frac{n!}{2^{n/2}(n/2)(n/2-1)\ldots 2\cdot 1} \\ &= \frac{n!}{2^{n/2}(n/2)!}. \end{aligned} \tag{7}$$

Remark 3.6
The number of generic pairings, g_n, for a fold with n parabolic points is

$$g_n = \begin{bmatrix} n \\ n/2 \end{bmatrix} \Big/ (n/2 + 1). \quad (8)$$

Proof
We indicate the parabolic points on an oriented fold by P_i, where $0 \le i \le n - 1$, and the parabolic points lying between and including P_i and P_j by the interval $[P_i, P_j]$, where P_i and P_j are beginning and end points, respectively, if the fold is traversed in the positive direction. (P_i, P_j) denotes a pairing between P_i and P_j, and g_n indicates the number of generic pairings for a fold with n parabolic points.

A pairing of the parabolic points on a fold is generic if the following constraints are satisfied: (1) all n parabolic points are paired, and (2) the parabolic curves connecting the parabolic points do not cross. From constraint (1) it follows that a pairing (P_i, P_j) partitions the parabolic points into the two disjoint subsets $[P_{i+1}, P_{j-1}]$ and $[P_{j+1}, P_{j-1}]$, both with an even number of parabolic points, namely, $j - i - 1$ and $n + i - 1$, respectively; and from constraint (2) it follows that no pairings between these subsets are possible. We can now derive a recursive formula for the number of generic pairings. Select some parabolic point P_i, and enumerate all possible pairings with it; for each of these pairings, enumerate all possible pairings within the two disjoint intervals. Thus the number of generic pairings with the pair (P_i, P_j) will be $g_{j-i-1}g_{n-j+i-1}$, assuming that $i < j$. Since i can be any value, we choose $i = 0$ for convenience. The allowed values for j are then $1, 3, \ldots, n - 1$, and the number of pairings is

$$g_n = \sum_{j=1,3,5}^{n-1} g_{j-1}g_{n-j-1} = \sum_{j=0,2,4}^{n-2} g_j g_{n-j-2} = \sum_{j=0}^{n/2-1} g_{2j}g_{n-2j-2}; \quad (9)$$

that is,

$$g_n = g_0 g_{n-2} + g_2 g_{n-4} + \ldots + g_{n-2}g_0, \quad (10)$$

where $g_0 = 1$.

To derive a nonrecursive expression for g_n, divide all subscripts in Eq. (10) by 2, add 1 to each subscript, and let $m = n/2 + 1$; then

$$g_m = g_1 g_{m-1} + g_2 g_{m-2} + \ldots + g_{m-1}g_1, \quad (11)$$

and $g_1 = 1$. Now, consider the so-called generating function[27] $g(x) = \Sigma_{m=1}^{\infty} g_m x^m$ for the sequence $\{g_1, g_2, \ldots\}$. By Eq. (11), $g(x) = x - [g(x)]^2$. Solving this quadratic equation yields $g(x) = 1/2 - 1/2(1 - 4x)^{1/2}$. By expanding $(1 - 4x)^{1/2}$ into a binomial series we obtain

$$g(x) = 1/2 - 1/2 \sum_{m=0}^{\infty} (-1)^m \begin{bmatrix} 1/2 \\ m \end{bmatrix} 4^m x^m = -1/2 \sum_{m=1}^{\infty} (-1)^m \begin{bmatrix} 1/2 \\ m \end{bmatrix} 4^m x^m. \quad (12)$$

Since a function has at most one Taylor series expansion about $x = 0$, it follows that

$$g_m = -1/2(-1)^m \begin{bmatrix} 1/2 \\ m \end{bmatrix} 4^m = \begin{bmatrix} 2m - 2 \\ m - 1 \end{bmatrix} \Big/ m. \quad (13)$$

Substituting $m = n/2 + 1$, we obtain the final result

$$g_n = \begin{bmatrix} n \\ n/2 \end{bmatrix} \Big/ (n/2 + 1). \quad (14)$$

Requiring pairings between parabolic points to be generic obviously decreases their number. Table 1 illustrates this effect but also shows that the number of generic pairings is still quite large, even for moderately many parabolic points. If we were also to consider pairings between parabolic points on different folds, the number of pairings would become even larger, since each parabolic point would have many more potential partners. It is clear, then, that any vision system depending on the pairing of parabolic points has to use constraints other than genericity to arrive at a reasonably small number of pairings. For example, pairing might proceed from coarser to finer scales, or only pairings between neighboring parabolic points might be allowed, as it is likely that they were generated by the same event, i.e., penetration or subtraction. Alternatively, sources of curvature information other than occluding contours might be used: certain extrema in illumination, for instance, occur at parabolic lines.[26]

By pairing the parabolic points of folds, we divide the surface between folds into regions, regions for which we can specify the sign of Gaussian curvature in a globally consistent manner. This leads us to the following definition.

Definition 3.7
A curvature interpretation of the surface of an object is

(1) A generic pairing of the parabolic points on the folds of the surface and also

(2) An assignment, for each region bounded by parabolic curves, of the sign of principal curvatures in a way that is

Table 1. Number of Pairings of *n* Parabolic Points on a Single Fold

Parabolic Points	Total Pairings[a]	Generic Pairings[b]
0	1	1
2	1	1
4	3	2
6	15	5
8	105	14
10	945	42

[a] Total pairings (generic plus nongeneric), $t_n = n!/(n/2)!2^{n/2}$.
[b] Generic pairings, $g_n = [{n \atop n/2}]/(n/2 + 1)$.

globally consistent and consistent with the sign of Gaussian curvature on each fold.

We can derive the sign of Gaussian curvature directly from the sign of the principal curvatures: if the principal curvatures have equal signs, Gaussian curvature is positive; otherwise it is negative. By specifying the two signs of principal curvature we can distinguish among convex (both are positive), concave (both are negative), and hyperbolic (one is positive and one is negative) surface regions. Assignment of the sign of principal curvatures in a globally consistent manner means that parabolic lines must separate regions with different signs of Gaussian curvature and that, conversely, regions with different signs of Gaussian curvature must be separated by parabolic lines. Before we continue we should mention that the above curvature interpretations are minimal in the sense that only those parabolic lines are included for which there is positive visual evidence (in the form of inflection points on the occluding contours) or that are necessary for a consistent assignment of Gaussian curvature.

To illustrate our approach we will take a closer look at a surface with two folds and six visible parabolic points.

Example 3.8

Consider the occluding contours shown in Fig. 8(a). The contour C_2 is the silhouette of the object O with surface S, and the contour C_1 is contained within C_2; that is, the fold F_1 (of C_1) is nested within fold F_2 (of C_2). Inflection points are indicated by I_i, and their corresponding parabolic points are indicated by P_i. We are going to address the following problem: Does the path X between points A and B lie within a convex region of the surface? In other words, does the contour C_1 indicate the presence of a hump separated from the background surface by a hyperbolic region? Note that the first part of path X is actually invisible from our vantage point, whereas the path Y is entirely visible. We will assume that the surface S has finitely many parabolic lines, it being intuitively clear that the contrary is nongeneric. We will examine the possible completions of contour C_1: (1) no parabolic points on the invisible part of fold F_1 and (2) two or more (but finitely many) parabolic points on the invisible portion of F_1.

In case (1), no parabolic points are on the invisible part of F_1 [Fig. 8(b)]. The parabolic line crossing F_1 at parabolic point P_5 has to go through P_6 as well, because crossings of folds and parabolic lines come in pairs and P_6 is the only remaining point on F_1. And since the path X has to cross this parabolic line, X cannot be completely contained within a convex region.

In case (2), two or more parabolic points are on the invisible part of F_1 [Fig. 8(c)]. Consider the Gaussian curvature along F_1: The segment between P_5 and P_6 is convex, the two segments [P_6, P_7] and [P_8, P_5] are hyperbolic, and the remaining segment [P_7, P_8] is concave. In general, any segment of the invisible part of a fold is either hyperbolic or concave, just as any segment of the visible part of a fold is either hyperbolic or convex. The reason for this difference between the visible and invisible parts of folds derives from the fact that radial curvature,[14] that is, curvature of the intersection of the surface and the normal section containing the line of sight, along the visible parts is positive (curvature vector points into the figure, parallel to the surface normal),

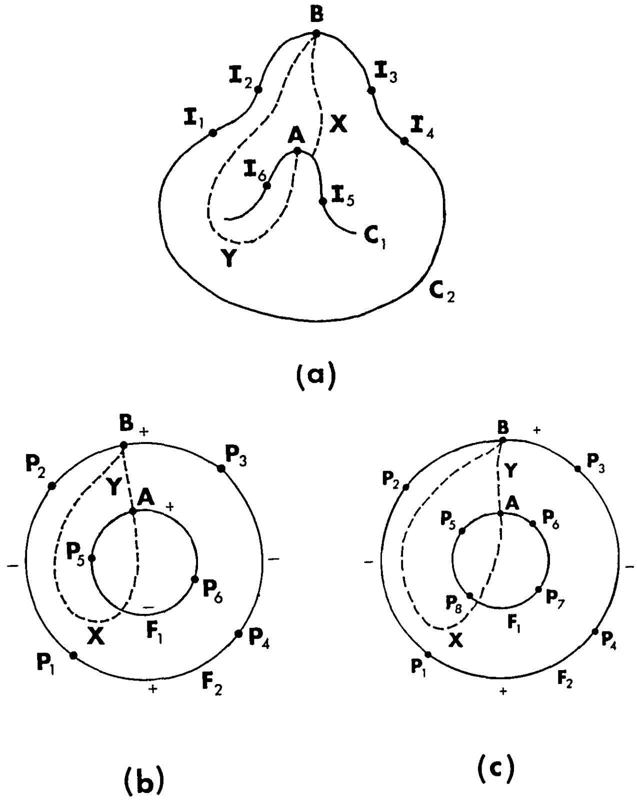

Fig. 8. (a) Is the path X between A and B completely contained within a convex region of the surface? (b) Folds F_1 and F_2, with P_5 and P_6 being the only parabolic points on fold F_1. (c) Folds F_1 and F_2, with the invisible parabolic points P_7 and P_8 on fold F_1.

whereas it is negative along invisible parts. Since transversal curvature,[14] that is, curvature of the occluding contours, can be positive or negative, it follows that there cannot be any concave segments along a visible fold, just as there cannot be any convex segments along the invisible part of a fold. From this it follows immediately that the parabolic line crossing F_1 at P_5 has to continue through P_6, and, again, we conclude that X has to cross a parabolic line and therefore cannot be completely contained within a convex region.

In summary, path X has to pass through a hyperbolic region, whereas it is possible that path Y is completely contained within a convex region.

C. Part Boundaries from Occluding Contours

We showed how to derive a curvature interpretation from the occluding contours of an object. We now turn to the problem of locating part boundaries on a surface when its curvature interpretation is given. According to our theory (Subsection 2.C) composite objects arise through combinations of primitive objects called parts, and boundaries between parts are contours of extrema of a principal curvature. A curvature interpretation, on the other hand, specifies only the signs of the principal curvatures, which is clearly insufficient to locate part boundaries.

Even though we are unable to locate part boundaries precisely on the basis of a curvature interpretation, we are able to decide whether a certain region of the surface contains a part boundary. This is possible because hyperbolic regions can originate only from combinations of parts, as the parts themselves are by definition convex. If we further assume that partitionings having fewest part boundaries are preferred, we can formulate rules for partitioning a surface on the basis of its curvature interpretation. (Note, however, that these rules apply only to cases in which existing part boundaries have not been disrupted by subsequent penetrations or subtractions.)

Definition 3.9
A surface can be partitioning on the basis of its curvature interpretation by the following rules:

(1) A hyperbolic ring separating two convex regions contains a B_c^-.
(2) A hyperbolic region inside a convex region indicates the presence of a B_0^- or a B_c^+.

Example 3.10
Consider the occluding contour shown in Fig. 9(a). It has four inflection points and therefore four curvature interpretations (or, more precisely, it has four interpretations that are minimal in the sense that no extraneous parabolic lines are included). These four curvature interpretations are shown in Figs. 9(b)–9(e): Fig. 9(b) shows a hyperbolic ring separating two convex regions, which implies the existence of a B_c^-; Fig. 9(c) shows a hyperbolic region inside a convex region, which implies the existence of a B_0^- or a B_c^+; Fig. 9(d) is the same as Fig. 9(c); and Fig. 9(e) shows two hyperbolic regions inside a convex region, implying the existence of two B_0^-'s, two B_c^+'s, or one B_0^- and one B_c^+.

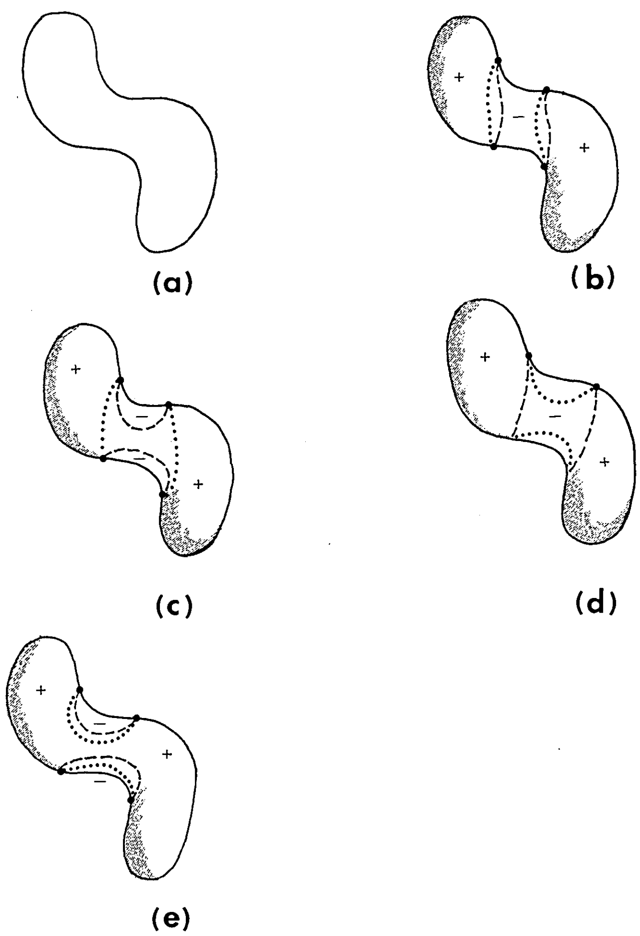

Fig. 9. (a) Occluding contour with four inflection points and four curvature interpretations, which are illustrated in (b)–(e). Dashed and dotted lines indicate parabolic lines on the visible and invisible parts of the surface, respectively.

Summarizing, occluding contours generally allow multiple curvature interpretations and therefore multiple partitionings; that is, observers can infer multiple deep structures from the occluding contours of an object. Given the fact that we can decompose the contours of objects in a number of ways, we can ask whether the human visual system prefers some decompositions over others. Experiments addressing this question could help us to constrain further the definitions of part boundaries.

4. DISCUSSION

If an organism is to interact successfully with its environment, it has to maintain an internal model of certain important properties of the environment. It is, for instance, important for human beings to be able to recognize objects by their shape. This ability implies that human beings somehow store shape descriptions of known objects in a shape memory. We elaborated a representation of 3-D shape in which smooth objects are partitioned along contours of extrema of a principal curvature. We arrived at these particular part boundaries by assuming that an object's shape can be described in terms of convex parts and that parts can be combined only through solid union and subtraction. The decomposition of an object into parts is called the deep structure of the object. Koenderink and van Doorn's observation[24] that "sculptors in many European and Asiatic stylistic periods build their shapes from 'ovoids' (in academic practice), that is, elliptically bounded pot-like volumes" lends support for the hypothesis that the human visual system represents shape in terms of combinations of parts, i.e., as deep structures. In some sense parts are analogous to morphemes, the smallest linguistic units that have meaning and cannot be further divided into even smaller units. We proposed one such unit for shape description: convex objects. It remains to be seen whether other aspects of shape, e.g., spatial relationships, can also be described in terms of irreducible units. We also have to extend the theory to include objects with holes (objects of genus greater than zero).

The description of solid shape in terms of parts has a number of important properties. First, the deep structure of an object is an object-centered description, i.e., is independent of a particular vantage point (but the observer might have to look at the object from different points of view to learn its complete structure). Second, the deep structure of a nonrigid object, such as a hand, is independent of the particular configuration of its rigid components, in this case the fingers of the hand. In other words, the deep structure of a hand is invariant under all possible positions of the fingers. The nonrigid aspect of shape (different positions of the fingers) would be captured by the transformation rules, which derive surface structure from deep structure. A third important property of the proposed representation is that it is incremental: Adding object X to object Y to form a composite object entails adding the description of X to the description of Y without changing the description of either

X or *Y* itself. The same can be said about removing parts of an object.

Describing the shape of an object is not an end in itself but a means toward recognizing that object. We propose a shape memory based on the deep structure of objects. Objects sharing the same deep structure can be distinguished by certain aspects of their surface structure. For instance, all humans fall into the category of human being on the basis of their deep structure. We can distinguish individuals on the basis of their surface structure, such as the exact shape of their eyes. A shape memory built on deep structures that are object-centered descriptions of shape does not adequately model human shape memory. Experiments on the use of coordinate systems in long-term memory showed that humans use both object-centered and viewer-centered coordinate systems.[28] A related aspect of recognition is the fact that humans find it difficult to recognize inverted or otherwise disoriented figures.[29] A face or an outline of a country shown upside-down is suddenly hard to recognize. We suggest that rotating an object does not alter the way in which humans decompose it into parts but does affect the predicates of spatial relationships between the parts. Perhaps (some) spatial relationships are specified with respect to a viewer- or environment-centered coordinate system.

We can now try to interpret the results of the studies on object recognition in persons with cerebral lesions. One interpretation is that persons with right-hemisphere lesions, who have difficulty in deciding whether two views are of the same object, cannot construct the deep structure of the objects that they are looking at. Another interpretation is that these patients can construct deep structures from images but cannot compare them. Persons with left-hemisphere lesions can compare objects on the basis of their shape (i.e., they can construct deep structures), but they cannot access shape memory (for instance, because the matching between a deep structure and the deep structures stored in memory does not function properly). As a consequence, they cannot attach meaning to their percepts.

APPENDIX A: RELEVANT CONCEPTS FROM DIFFERENTIAL GEOMETRY

Consider a plane curve C that is smooth everywhere, that is, has a well-defined tangent everywhere, and let $\mathbf{x} = \mathbf{x}(s)$ be a natural representation of C (where the parameter s indicates length along C). The tangent vector $\mathbf{t}(s)$ indicates the direction of C at $\mathbf{x}(s)$. The rate at which the direction of $\mathbf{t}(s)$ changes as one moves along the curve is a measure of its curvature: the faster it changes, the larger the curve's curvature is. The curvature vector is defined as $\mathbf{k}(s) = \mathrm{d}\mathbf{t}(s)/\mathrm{d}s$, and the principal normal unit vector, $\mathbf{N}(s)$, is a vector parallel to the curvature vector. By convention, the principal normal is positive if it points toward the figure side of the curve and negative if it points toward the ground side. We can define a continuous valued function $\kappa(s)$ along C such that $\mathbf{k}(s) = \kappa(s)\mathbf{N}(s)$. If $\mathbf{N}$ and $\mathbf{k}$ point in the same direction, $\kappa(s)$ will be positive; otherwise it will be negative. $\kappa(s)$ is called the curvature of C at $\mathbf{x}(s)$. Depending on the sign of curvature an arc will be called hyperbolic (negative curvature), elliptic (positive curvature), or parabolic (zero curvature).

The curvature at a point P on a surface in $\mathscr{R}^3$ is, of course, more complex than the curvature of a plane curve, as it depends on the direction in which the surface is traversed. Let C be a curve through P cut out by a plane containing the normal, $\mathbf{N}$, at P; that is, C is a normal section of the surface patch containing P, and its curvature is called normal curvature, denoted κ_n. The normal curvature depends on the direction of C, and the two perpendicular directions for which the value of κ_n is maximal or minimal are called the principal directions. The corresponding normal curvatures, κ_1 and κ_2, are called principal curvatures. A line on the surface whose tangent is everywhere in a principal direction is called a line of curvature.

The Gaussian curvature K of a surface is defined as $K = \kappa_1\kappa_2$. The sign of the Gaussian curvature is a qualitative measure of surface shape: if $K > 0$ the surface is on one side of the tangent plane (synclastic or elliptic) and is either convex ($\kappa_1, \kappa_2 > 0$) or concave ($\kappa_1, \kappa_2 < 0$); if $K < 0$ the surface is on either side of the tangent plane (anticlastic or hyperbolic) and is saddle shaped (one of the principal curvatures is negative and the other is positive); and if $K = 0$, at least one of the principal curvatures is zero (monoclastic or parabolic), and the surface is cylindrical or planar. Since a hyperbolic surface has negative and positive normal curvatures, there must be two directions for which $\kappa_n = 0$, the so-called asymptotic directions. A curve everywhere tangent to an asymptotic direction is an asymptotic line. Two directions, with direction numbers $\mathrm{d}u : \mathrm{d}v$ and $\delta u : \delta v$, are conjugate if $\mathbf{dx} \cdot \delta\mathbf{N} = 0$.

ACKNOWLEDGMENTS

We thank Ron Ball, Steve Reuman, and Whitman Richards for useful discussions; Cheryl Larsson for drawing the figures; and the reviewers for their comments. This material is based on work supported by the National Science Foundation under grant no. IST-8413560 and by the Office of Naval Research, Psychological Sciences Division, Engineering Psychology Group.

REFERENCES

1. E. K. Warrington and A. M. Taylor, "The contribution of the right parietal lobe to object recognition," Cortex **9**, 152–164 (1973).
2. E. K. Warrington and A. M. Taylor, "Two categorical stages of object recognition," Perception **7**, 695–705 (1978).
3. R. N. Shepard, "Psychophysical complementarity," in *Perceptual Organization*, M. Kubovy and J. R. Pomerantz, eds. (Erlbaum, Hillsdale, N.J., 1981), pp. 279–341.
4. M. Minsky, "A framework for representing knowledge," in *The Psychology of Computer Vision*, P. Winston, ed. (McGraw-Hill, New York, 1975), pp. 211–277.
5. H. Blum and R. N. Nagel, "Shape description using weighted symmetric axis features," Pattern Recognition **10**, 167–180 (1978).
6. J. O'Rourke and N. Badler, "Decomposition of three-dimensional objects into spheres," IEEE Trans. Pattern Anal. Machine Intell. **PAMI-1**, 295–305 (1979).
7. L. R. Nackman, "Curvature relations in three-dimensional symmetric axes," Comput. Graph. Image Process. **20**, 43–57 (1982).
8. R. Mohr and R. Bajcsy, "Packing volumes by spheres," IEEE Trans. Pattern Anal. Machine Intell. **PAMI-5**, 111–116 (1983).
9. D. Marr, *Vision* (Freeman, San Francisco, Calif., 1982).

10. R. A. Brooks, "Model-based three-dimensional interpretations of two-dimensional images," IEEE Trans. Pattern Anal. Machine Intell. **PAMI-5**, 140–150 (1983).
11. D. D. Hoffman, "The interpretation of visual illusions," Sci. Am. **249**(6), 154–162 (1983).
12. D. D. Hoffman and W. A. Richards, "Parts of recognition," Cognition **18**, 65–96 (1984).
13. J. J. Koenderink and A. J. van Doorn, "The singularities of the visual mapping," Biol. Cyber. **24**, 51–59 (1976).
14. J. J. Koenderink, "What does the occluding contour tell us about solid shape?" Perception **13**, 321–330 (1984).
15. D. Marr and T. Poggio, "Cooperative computation of stereo disparity," Science **194**, 283–287 (1976).
16. S. Ullman, *The Interprettion of Visual Motion* (MIT Press, Cambridge, Mass., 1979).
17. J. J. Koenderink and A. J. van Doorn, "Depth and shape from differential perspective in the presence of bending deformations," J. Opt. Soc. Am. A **3**, 242–249 (1986).
18. B. M. Bennett and D. D. Hoffman, "Shape decompositions for visual shape recognition: the role of transversality," in *Image Understanding II*, W. A. Richards, ed. (Ablex, Norwood, N.J., 1985).
19. S. Lipschutz, *General Topology* (McGraw-Hill, New York, 1965).
20. N. Chomsky, *Aspects of the Theory of Syntax* (MIT Press, Cambridge, Mass., 1965).
21. N. Chomsky, *Essays on Form and Interpretation* (Elsevier/North-Holland, Amsterdam, 1977).
22. A. P. Pentland, "Local shading analysis," IEEE Trans. Pattern Anal. Machine Intell. **PAMI-6**, 170–187 (1984).
23. J. J. Koenderink and A. J. van Doorn, "Invariant properties of the motion parallax field due to the movement of rigid bodies relative to an observer," Opt. Acta **22**, 773–791 (1975).
24. J. J. Koenderink and A. J. van Doorn, "The shape of smooth objects and the way contours end," Perception **11**, 129–137 (1982).
25. H. Whitney, "On singularities of mappings of Euclidean spaces. I. Mappings of the plane into the plane," Ann. Math. **62**, 374–410 (1955).
26. J. J. Koenderink and A. J. van Doorn, "Photometric invariants related to solid shape," Opt. Acta **27**, 981–996 (1980).
27. G. E. Andrews, *Number Theory* (Saunders, Philadelphia, Pa., 1971).
28. P. Jolicoeur and S. M. Kosslyn, "Coordinate systems in the long-term memory representation of three-dimensional shapes," Cognitive Psychol. **15**, 301–345 (1983).
29. I. Rock, "The perception of disoriented figures," Sci. Am. **230** (1), 78–85 (1974).