Inferring three-dimensional shapes from two-dimensional silhouettes

W. A. Richards

Natural Computation Group, E10-120, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Jan J. Koenderink

Fysisch Laboratorium, Rijksuniversiteit Utrecht, Utrecht, 3584 CC, The Netherlands

D. D. Hoffman

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

Received June 2, 1986; accepted March 4, 1987

Although an infinity of three-dimensional (3-D) objects could generate any given silhouette, we usually infer only one 3-D object from its two-dimensional (2-D) projection. What are the constraints that restrict this infinity of choices? We identify three mathematical properties of smooth surfaces plus one simple viewing constraint that seem to drive our preferred interpretation of 3-D shape from 2-D contour. The constraint is an extension of the notion of general position. Taken together, our interpretation rules predict that "dents" in a 3-D surface should never be inferred from a smooth 2-D silhouette.

INTRODUCTION

Our aim is to understand how unique three-dimensional (3-D) interpretations can be made from two-dimensional (2-D) silhouettes. For example, outline P3 of Fig. 2 below looks like a dumbbell, whereas T3 of Fig. 3 looks like a croissant, and T4 looks like a pear. Yet each of these silhouettes has an infinity of 3-D interpretations, considering that we are given no information about the bumps and dents on either the back or the front side of the surface. Why, then, do we tend to pick only one or two 3-D shapes? Clearly, some powerful constraints must be imposed on our interpretations. One of these, to be elaborated below, is that we do not propose protrusions or indentations of a surface without evidence for such. However, this rule by itself is not sufficient to drive unique 3-D interpretations of these silhouettes. To this end, we identify some intrinsic properties of smooth surfaces that are implicitly understood when 3-D shape interpretations are made.

Before embarking on our analysis of inferring 3-D shape from 2-D silhouettes, we first introduce a method for enumerating all possible silhouettes. Without such an enumeration scheme, our selection of outlines might be rather *ad hoc* and arbitrary. We choose a representation of plane curves that is based on curvature. In this way we can capture the general form of all types of invaginations and protrusions but not at the expense of carrying scale and metrical information. For our purposes, our choice of using extrema or singularities of curvature has perceptual relevance, because these extrema make explicit the parts of an outline. Arguments for this choice are presented elsewhere, including in the companion paper.[1-3] The most primitive parts of a 2-D shape are called codons and are illustrated in Fig. 1. This set provides a complete basis for describing any wiggly curve, such as a silhouette, and hence can be used to enumerate a class of silhouettes. The representation has the further advantage of making explicit certain features of an outline that permit 3-D interpretations, such as the Gaussian curvature.

THE CODON REPRESENTATION

The codon shape primitives are defined in terms of the relations between the maxima, minima, and zeros of curvature encountered as one traverses a plane curve.[1,2] To specify a codon type, one first must assign a direction to the curve, in essence defining which side of the curve corresponds to a figure. Our convention is to keep the figure to the left of the direction of traversal. Clockwise rotation of the curve now corresponds to negative curvature, and counterclockwise rotation is positive curvature. Minima of curvature are then used to break the curve into segments, whereas maxima and zeros are used to describe the shape of each segment. With this scheme, there are only five basic types of segments, or codons, which are shown in Fig. 1. Referring to the figure, a codon type is specified by the presence or absence of an inflection (zero of curvature) and by whether the positive extremum of curvature(i.e., the maximum of curvature) occurs before or after the inflection, when present. For example, a type 2 codon has two inflections, and a 1^+ has only one inflection that follows the positive extremum of curvature, whereas for a 1^- codon the inflection precedes the positive extremum. There are also two codon types without inflections. The 0^+ and 0^- differ in that in the first case (0^+) the codon boundary is a positive minimum of curvature, whereas in the second (0^-) the boundary is negative extrema. The set of primitives shown in Fig. 1 provides a complete descrip-

0740-3232/87/071168-08$02.00

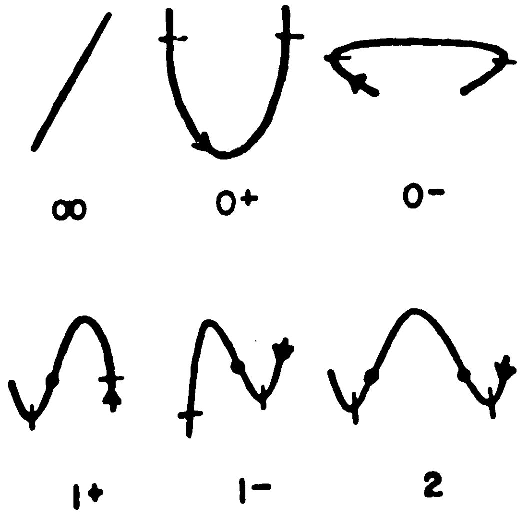

Fig. 1. The primitive codon types. Zeros of curvature are indicated by dots; minima are indicated by slashes. The straight line (∞) is a degenerate case included for completeness, although it is not treated in the text.

tion of a curve in terms of the singularities of curvature (maxima, minima, and zeros).

Our procedure for enumerating all possible silhouettes is simply to build closed codon strings of increasing length. Without constraint, we can expect 5^N possible sequences of length N. However, silhouettes are closed outlines that are not self-intersecting. If we also impose the constraint that the silhouette be smooth (i.e., without cusps), then the number of possible sequences is reduced to roughly $3 \times 2^{N-2}$. For example, of the 3125 combinations of codon strings of length 5, only 25 will satisfy the smooth-silhouette constraint.[4–6] Allowing a single cusp raises the number to 457. In order to keep our silhouettes manageable, therefore, we consider only smooth silhouettes of codon length 4 or less. All these possibilities are shown in Figs. 2–4, and they include hints of animallike shapes such as Q12. Their construction is discussed elsewhere.[4–6] These figures thus represent the silhouettes that we wish to analyze. Although this set may appear limited, it will be shown that our techniques can be readily generalized.

THE CANONICAL VIEW

We begin by examining the simple outlines of Fig. 2, the ellipse, the peanut, and the dumbbell, for the analysis of these simple silhouettes provides us with the tools needed to interpret the more complex shapes of Figs. 3 and 4. The simplest of these three outlines is the ellipse, which we naturally interpret as the silhouette of an ellipsoid, or egg. But why? If the outline is a special view of an object, such as the end-on view of the dumbbell or peanut, we could be fooled. Our interpretation thus assumes that our view is such that none of the bumps or dents of the object is occluded or invisible. To capture this notion, we propose our first interpretation rule:

(R1) Do not propose undulations of the 3-D surface without evidence for such.

The above rule is an extension of the general position restriction, which requires that the view of an object is not a special one and is stable under perturbation. For our purposes the restriction states that a slight shift in viewpoint should not change the topology of the viewed structure, such as by suddenly revealing a bump or dent in the surface that was previously hidden by occlusion. We define such views of a surface as being generic. Our interpretation rule thus implicitly assumes that the observer's view is generic. It also implies that all the undulations of the surface that are needed to infer a plausible 3-D shape are visible. The view of the silhouette is thus assumed to be a special generic view,

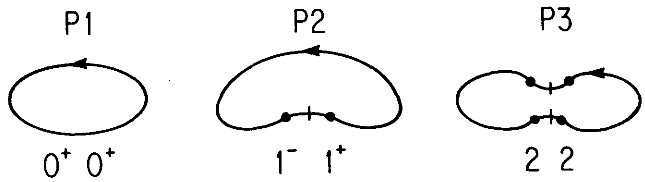

Fig. 2. Legal, smooth, closed codon pairs. Part boundaries are indicated by the slashes; inflections are indicated by dots.

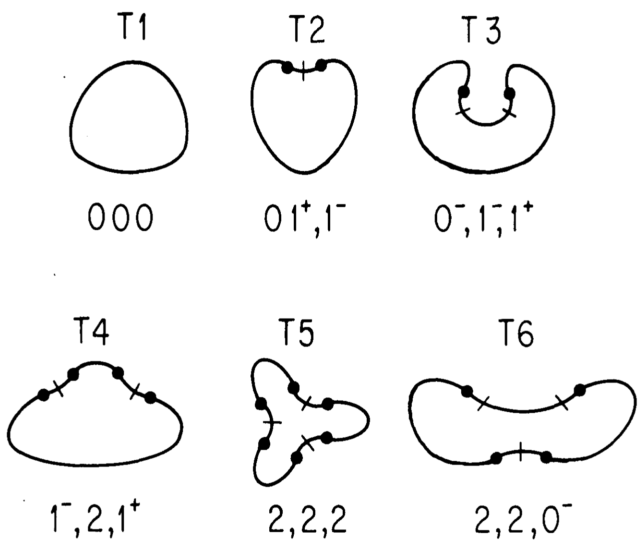

Fig. 3. Legal, smooth, closed codon triples. The tick marks indicate the extrema of negative curvature, which are generally the part boundaries, whereas the dots show the inflections.

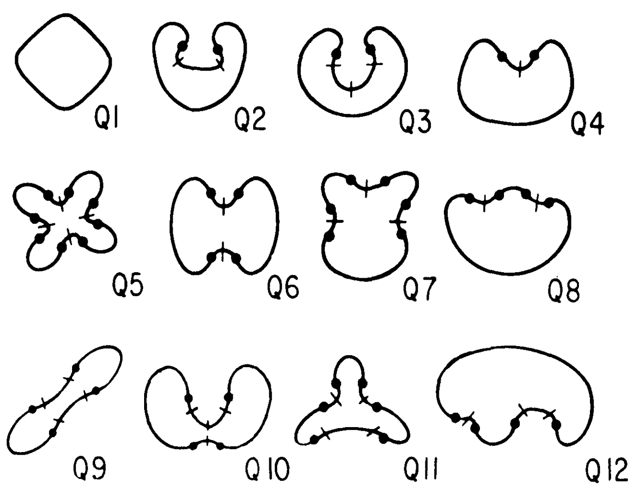

Fig. 4. Legal, smooth, closed codon quadruples. The tick marks indicate the extrema of negative curvature. The dots show inflections through which the flexional loci must pass.

namely, one that might be called prototypic. As more information is added about the 3-D surface, it might be expected that the inferred shape will evolve in a graceful manner. This captures the notion expressed by Marr in his principle of least commitment.[7] Thus, in the case of the ellipse outline of Fig. 2, the most plausible 3-D interpretation according to the rule is an ellipsoid.

GAUSSIAN CURVATURE

The ellipse is a special outline because it has no undulations and hence its sign of curvature everywhere is the same, namely, positive. The peanut and the dumbbell are more complex, however, with bumps and dents. Clearly, we need a means of describing undulations on 3-D surfaces so that we can enumerate all possible 3-D interpretations of the 2-D outlines. We choose for this purpose an intrinsic property of 3-D surfaces, namely, the Gaussian curvature.

At any point on a smooth (nonplanar) surface, there is a direction where the surface curves the most and another direction where the surface curves the least. These two directions are the directions of principal curvature, and they are always perpendicular.[8] The Gaussian curvature is simply the product of these two curvatures. Of interest to us is the sign of the Gaussian curvature, which permits a qualitative description of the topology of a surface. When the directions of both principal curvatures are identical, such as on an ellipse, the Gaussian curvature is positive; when the principal curvatures are in opposite directions, such as on a saddle, the Gaussian curvature is negative. If one of the principal curvatures is zero, such as on a cylinder, then the Gaussian curvature will be zero also. Any point on a surface will thus have positive, negative, or zero Gaussian curvature, depending on whether it is locally elliptical, hyperbolic (saddle), or cylindrical. A 3-D dumbbell may now be defined as a single hyperbolic region of negative Gaussian curvature (the neck) joining two protrusions of positive Gaussian curvature (the two ovoids). A 3-D peanut is a hyperbolic region (saddle) lying within an ellipsoid of positive Gaussian curvature.

In addition to providing the basis for a taxonomy of 3-D shapes, Gaussian curvature has another distinct advantage for our purposes. Consider a point on the 3-D surface that projects into the 2-D silhouette. The following is then true:

(C1) The sign of the Gaussian curvature of points on the 3-D surface that project into the silhouette is the same as the sign of curvature of those projections.

This theorem by Koenderink and van Doorn[9] thus assures us that the Gaussian curvature of the 3-D shape is positive at points on the surface that project into regions of positive curvature on the silhouette. Thus both the peanut and the dumbbell outlines of Fig. 2 require that the corresponding 3-D shapes have hyperbolic (saddle) regions of negative Gaussian curvature within a region (or two) of positive Gaussian curvature. Note that this theorem also implies that dents, which have positive Gaussian curvature but are concavities in the surface, will never appear in the generic silhouette.

At this point, one might be misled to the false conclusion that our problem is essentially solved. However, as is shown in Table 1 of the companion paper by Beusmans *et al.*,[3] for even the simple rabbit-head silhouette Q7 there are 15 possible 3-D interpretations if the embedding of regions of positive and negative Gaussian curvature is unconstrained. For the jack shape Q5 there are 105 possibilities. Yet we see both of these 2-D shapes in only one or two ways as three-dimensional objects. Clearly, the constraints and rules that we invoke must be quite powerful. In the next section we introduce one important mathematical constraint, also elaborated in the companion paper,[3] and subsequently add another interpretation rule to constrain the 3-D possibilities further.

TYPES OF SURFACE UNDULATIONS

Consider a surface of hills and valleys with perhaps also a depression that might collect water after a rainfall. The hill, being a bump on an elliptic surface, will have positive Gaussian curvature (such as the most probable interpretation of T4), but so will its inverted shape that creates the dent or depression, for the product of two negative principal curvatures will be positive. The positive sign of Gaussian curvature thus does not tell us whether the surface is convex or concave.[10] Similarly, there will be two kinds of regions associated with negative (hyperbolic) Gaussian curvature, which we will call furrows and ridges (these regions may also be called saddles and humps). A furrow (or saddle) is a region of negative Gaussian curvature within an elliptical region, whereas a ridge (or hump) is an elliptic region within a hyperbolic one. Shapes T2 and Q8 tend to be given these two interpretations. We thus have four types of surface undulations: bumps, dents, furrows (saddles), and ridges (humps). Together, they form a complete qualitative description of any smooth surface.[11–14]

We next need a simple scheme for representing the relations between the four types of surface undulations. For this we choose the Gaussian sphere. Out intent is to project the silhouette onto the Gaussian sphere and to use this projection, together with the topology of the silhouette, to constrain the possibilities of 3-D shapes.

The Gaussian sphere is simply a parallel mapping of all the surface normals into the unit sphere, with the tail of each normal placed at the center.[8] Each point on the surface of this sphere thus corresponds to a particular orientation (see Ref. 15 for an extended discussion). In the case of a convex object with positive Gaussian curvature everywhere, no two points on the surface will have the same projection onto the sphere. However, for objects with concavities, the same point on the Gaussian sphere may represent two or more points on the object's surface.

Consider now the mapping of a silhouette, such as that of the peanut, onto the Gaussian sphere. Assume for the moment parallel projection; then each visual ray that gives rise to the silhouette must strike the surface in such a way to be perpendicular to the surface normal at the point of contact. All surface normals on the surface contours that give rise to the silhouette must therefore lie in the frontal plane, parallel to the image plane. This plane, when mapped into the Gaussian sphere, will pass through the origin of the sphere. Hence the locus of any silhouette seen under parallel projection will be a great circle on the Gaussian sphere (Fig. 5).

Let us now trace the surface normals A–E of the silhouette

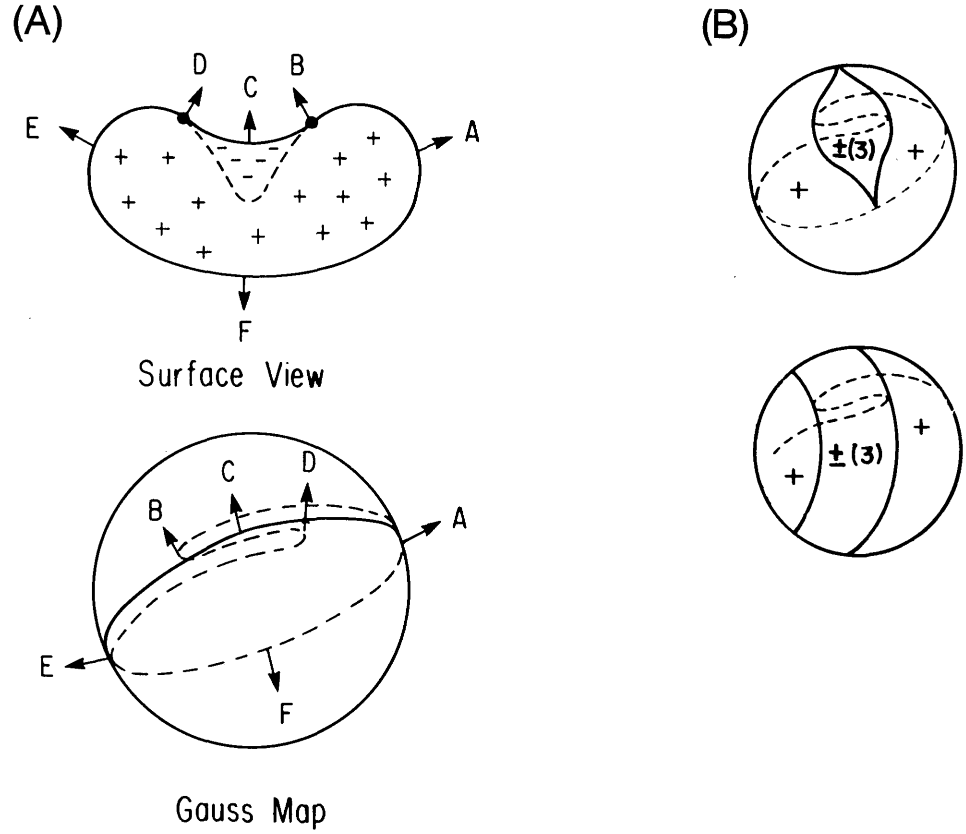

Fig. 5. The outline of a 3-D shape seen under parallel projection maps onto a great circle on the Gaussian sphere. (A) A peanut-shaped outline with two inflections will have two folds on the Gaussian map, as shown by the dashed line. The extremities of these folds must lie on flexional lines of zero Gaussian curvature, as suggested in the 3-D rendition of the silhouette (top left). (B) There are only two possible ways of closing the two flexional loci that intersect the folds on the great circle.[16] The lower example with two loops is not appropriate for the silhouette of the peanut, which has only two inflections. Thus the upper example correctly describes the Gaussian map of the peanut. We indicate the triple covering of the Gaussian map within the flexional lines by the ±(3) notation.

of the peanut onto the appropriate great circle of the Gaussian sphere. Starting at A, we move through the vertical to position B, which is one of the two inflections. Passing through this point to the second inflection at D requires that we traverse a position on the silhouette with a vertical normal C. The surface normal at the bottom of the well at C thus has the same direction as the normal to the two bumps at the top of the peanut. Thus this point C on the Gaussian sphere actually corresponds to three points on the object's surface. In Fig. 5, the dashed contours on the Gaussian spheres represent the locus of the silhouette of the peanut. The ends of the two folds on these lines are the inflection points (dots) on the silhouette (B and D) where the direction of rotation of the surface normals changes as one moves along the contour. Of necessity, the end of any such fold is a point of zero Gaussian curvature, where the outline (or surface) goes from a region of positive to negative Gaussian curvature. Between the folds the Gaussian map is said to be triply covered in this case because each point on the sphere corresponds to three points on the surface.

For the class of generic surfaces that we are considering, it is not possible to have an isolated point of zero Gaussian curvature. Rather, all points of zero Gaussian curvature must lie on closed lines.[14] We call these lines the flexional lines of a surface, for they are the boundaries between regions of positive and negative Gaussian curvature and produce an inflection on the surface (or silhouette). (Flexional lines are also called parabolic lines in some texts.) Koenderink and van Doorn[13,14] have proven the following important property of flexional lines:

(C2) For generic surfaces, the flexional (parabolic) lines are closed and nonintersecting.

Thus we now know that on the Gaussian sphere the flexional lines must also be closed.

Returning to our peanut example in Fig. 5, this means that the ends of the two folds on the Gaussian sphere must lie on closed curves. As shown by Whitney[16] and others,[14,17] there are only two ways that we can close flexional curves on the Gaussian map: either we can join the ends in a smooth loop or we can create a cusp. If the flexional curve is closed with a cusp on the Gaussian map, as illustrated in the upper example of Fig. 5(B), then this specifies a wrinkle on the surface (i.e., a furrow or ridge), and there must be a cusp on the opposite (invisible) side of the great circle locus of the silhouette (dashed lines). We call this a pleat. If the flexional curves are closed by simple loops, as illustrated in the lower example of Fig. 5(B), then we now are left with only the possibility of joining the folds on the Gaussian map with two loops.[16] However, two loops on the Gaussian map indicate four folds because each loop will cross the great circle of the silhouette twice. Hence the remaining possibility of two loops is also excluded because four folds on the Gaussian map require four inflections on the silhouette and the peanut has only two. The peanut silhouette thus must correspond to an ovoid of positive Gaussian curvature with a single furrow (saddle) of negative Gaussian curvature, which is the 3-D peanut.[18]

Notice that a dent in an ellipsoid is not included in our interpretations of the peanut-shaped outline. If a dent is to appear on the smooth silhouette, then its hyperbolic lip must be visible. However, this would require two visible flexional lines on the silhouette in the region of the dent—an impossibility. Hence the following constraint appears for interpreting smooth silhouettes:

(C3) A region of negative curvature on a silhouette is always interpreted in three dimensions as a furrow (or neck), never as a dent.

Inspection of the shapes illustrated in Figs. 2–4 shows that this is the case. Note that this constraint follows directly from rule (R1) and constraint (C1).

THE DUMBBELL (OR PEAR)

The analysis of the dumbbell silhouette now follows quite simply. Its Gaussian map is shown in Fig. 6. There are two pairs of folds, each pair corresponding to the upper and lower views of the bar of the dumbbell. The flexional lines through the extremities of these folds may be closed in two ways, as illustrated: either we can fuse the pairs as we did previously with the peanut to create an egg with either one or two furrows or we can form two rings or loops to create a 3-D dumbbell. These are the only two possibilities, given our interpretation rule (R1) that no undulations of the surface should be proposed without evidence for such. But of these two possibilities, which do we pick? Again, we invoke a natural extension to our interpretation rule. To show two furrows, the 3-D object must be oriented more carefully with respect to the viewer than in the case of the dumbbell. For example, if the furrows were on the front and back faces of the shape, then the silhouette could be an ellipse. The dumbbell shape is thus a more general position interpretation and should be preferred. A corollary to our interpretation rule is thus

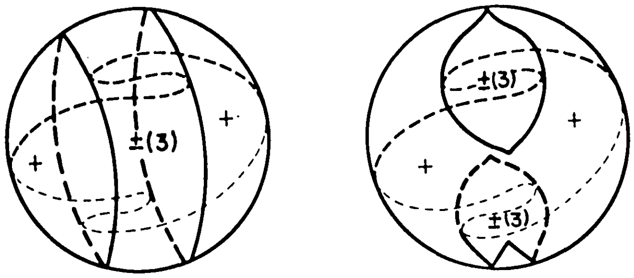

Fig. 6. The dumbbell outline has two pairs of folds when mapped onto the Gaussian sphere (dashed lines). Closed nonintersecting flexional lines can thus be created in three ways: by fusing the folds on each side of the sphere, creating two pleats with two cusps each (right); by creating two loops and no cusps (left); or by fusing the cusps on one side of the great circle, leaving a gap in the fused contours to create one big pleat with two cusps. The fusing versions are like an ellipsoid with one or two separate furrows, whereas the loop version is the true dumbbell. Again, as in Fig. 5, the ±(3) regions are triply covered.

(R2) Pick the most general position 3-D interpretation, namely, that 3-D shape that preserves the signs of the curvature of the silhouette over the widest range of viewpoints.

This corollary to our interpretation rule now excludes the egg with one or two furrows, for such an interpretation requires that the egg be viewed in a somewhat restricted manner. Thus the preferred 3-D interpretation for the dumbbell outline should be the dumbbell.

LOOPS AND PLEATS

A strategy for enumerating the legal 3-D shapes now emerges. When the silhouette is mapped onto the Gaussian sphere, the inflections of the silhouette map into folds on the Gaussian map. These points on the folds occur in pairs and delimit a region of either positive or negative Gaussian curvature. (The silhouette gives us the sign of curvature of the region.) We have two ways of joining the flexional loci that intersect the extremity of these folds: by forming either loops or pleats that create cusps on the Gaussian map. The loop choice will correspond either to a bump or dent on the surface or to a neck or knuckle in the 3-D shape; the latter cases require an inflection to appear at (roughly) opposite sides of the parts of the silhouette. The second (pleat) scheme corresponds to either a ridge or a furrow on one side of the 3-D shape. Whether the region between the flexional loci on the Gaussian map is a ridge, furrow, etc. is given by the sign of curvature of the silhouette [constraint (C2)]. In total, our parts of a surface will be a bump, a ridge, or a knuckle or one of their complements, a dent, a furrow, or a neck.[19]

If we have only two inflections, as in the peanut, then there is only one possibility. This is the pair of cusps or pleat on the Gaussian map (i.e., furrow or saddle in the 3-D shape) because the loops would create two more inflections on the silhouette that are not seen and hence are inferred not to be present. Referring to Figs. 3 and 4, the same argument applies to the codon triples T2 and T3 (croissant) or the three quadruples Q2 (bib), Q3, and Q4 (apron).

If we have four inflections positioned in the silhouette as in shapes T4, T6, Q6, Q9, and Q10, we still have only the loop and pleat possibilities to consider, as we did before with the dumbbell. The preferred choice is loops on the Gaussian map rather than two pleats, because the first solution yields a surface of revolution whose canonical silhouette will always be the same. The argument is identical to that given previously for the dumbbell, although in this case shape T4 would be a pear with the indentations corresponding to a neck.

Figure 7 summarizes the interpretations given to the silhouettes. Rather than showing all silhouettes, we have illustrated only the major classes. For example, the dumbbell shapes T6, Q7, Q9, and Q10 are topologically identical to P3; therefore only P3 is shown. Likewise T2, T3, Q2, Q3, and Q4 are similar to the simple peanut P2. The remaining outlines to be interpreted thus fall into the class containing T4, T5, Q7, Q8, and Q12.

Shape T4 is now trivial. The Gaussian mapping of its flexional loci is the same as for the dumbbell. Hence the preferred interpretation will be a pear as indicated. For

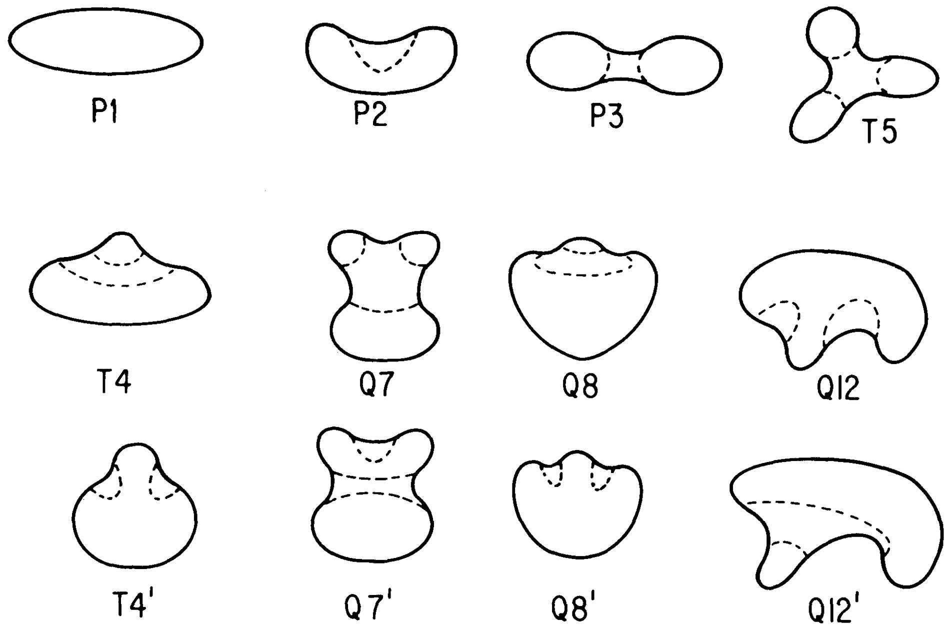

Fig. 7. Different classes of the outlines to be interpreted. Dashed lines are the preferred flexional (parabolic) loci. Primes indicate alternative interpretations.

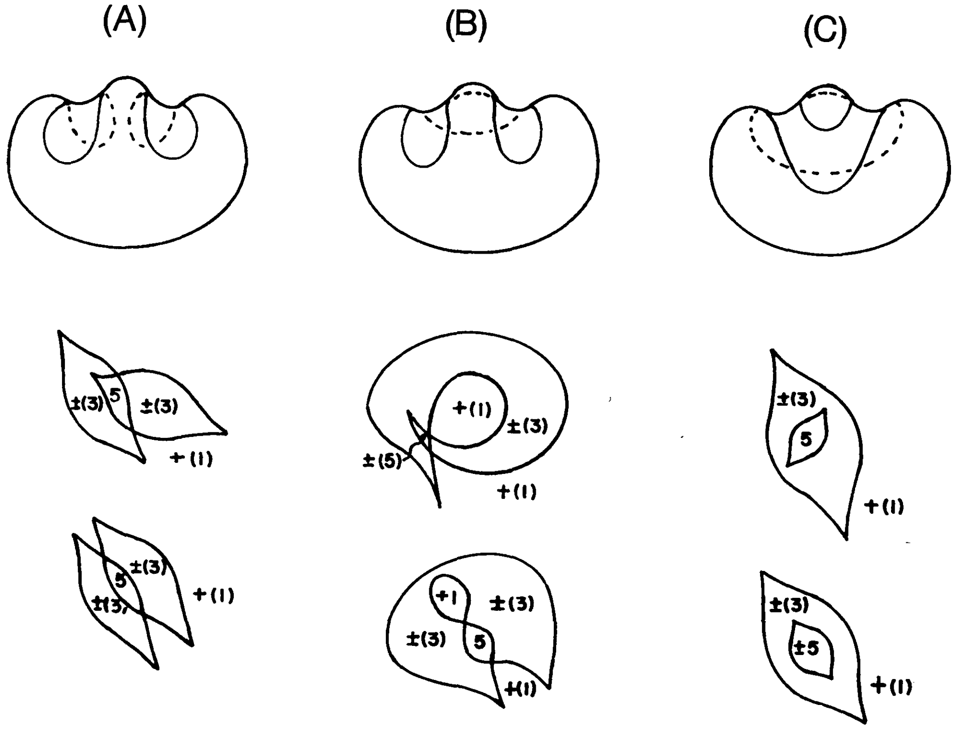

Fig. 8. Three possible interpretations of the silhouette Q8. The flexional (parabolic) lines on the surface are illustrated in the top row. Beneath each of these possibilities is sketched the shape of the flexional loci on the Gaussian map. The numbers in parentheses indicate the coverings. The first case is two furrows (A), the second case is a single furrow (B), and the third case is a ridge in a furrow (C).

similar reasons, the jack silhouette T5 should also be interpreted as having three fingers, as illustrated. However, Q7 can be similarly interpreted as a body with two limbs, for its Gaussian mapping is similar to that of T5. Hence we are left with only two remaining silhouettes, Q8 and Q12.

Our present constraints are not powerful enough to find a preferred interpretation for shape Q8. Unlike with shapes Q6 and T4, which also have four inflections, there is no symmetry axis about which the silhouette Q8 can be rotated to produce the same 3-D outline. Thus the two indentations of Q8 can not correspond to a neck in a surface of revolution. Instead, because all four inflections are on one side of the outline, we will have a region on the Gaussian map that will be five times covered (rather than the triple-covering characteristic of the neck in Fig. 6). After reviewing the possible pairings once more and excluding dents by constraint (C3), we have on the Gaussian map the following three possibilities, as illustrated in Fig. 8: two adjacent pleated contours (furrows), as in Fig. 8A; a single pleat (a furrow), as in Fig. 8B; and a pleat enclosed by a pleat (a ridge in a furrow) as in Fig. 8C. Unfortunately, our present interpretation rules do not force a unique choice among the three options illustrated in Fig. 8. Although there may be a preference for a ridge in a furrow [Fig. 8(C)], our general position rule (R2) is not sufficiently formulated to exclude those shown in Figs. 8(A) and 8(B). Underlying our choice of the ridge in a furrow seems to be the notion that the probability of two furrows' being aligned with the viewer as in Figs. 8(A) and 8(B) is less than if one single furrow is so aligned as in Fig. 8(C). A similar argument applies to shape Q12.

DISCUSSION

Effect of Constraints

Our four interpretation rules and constraints have allowed us to assign one preferred 3-D interpretation to each of the silhouettes shown in Figs. 2–4. These silhouettes are all the possibilities for smooth codon strings of length 4 or less. Without constraints, these are 625 different sequences of four codons, or 625 possible 2-D outlines that are generically different (i.e., have different sequences of the extrema of curvature). Closure and smoothness reduce this number to the 12 outlines shown in Fig. 4. In a similar manner we can ask, given a single outline of Fig. 4: How many different 3-D shapes are possible? Clearly, there is a very large number, which increases as smaller and smaller undulations are tolerated.

If we now introduce the simple mathematical constraint that the flexional loci can not intersect on the surface [constraint (C2)], the possibilities are markedly reduced but still substantial, for example, for eight inflection points. Koenderink's rule reduces the number of pairings that must be considered from 105 to 14, as discussed by Beusmans *et al.*[3] However, when our viewing constraint is added, together with the stipulation that loops be given preference over pleats, provided that constraint (C3) is not violated, our preferred interpretations become almost unique, at least for the 2-D outlines examined. Therefore two simple interpretation rules plus two mathematical theorems provide powerful constraints on the interpretation of silhouettes.

Instant Psychophysics

Of course, the 2-D outlines of Fig. 7 may not always give a unique interpretation. Our constraining rules are not rigid. For example, the spade outline Q8, by our rules, should be seen as a ridge (hump) in a shallow furrow (saddle). However, it is easy to imagine Q8 as a convex blob with two furrows, an image that corresponds to two pleats on the Gaussian map, although this is not our constrained solution. The outlines given in the lower panel of Fig. 7 illustrate such less-preferred interpretations. Similarly, outlines T5 and Q7 have the alternative 3-D interpretation of a furrow in the base of a necked shape. Instead of all loops, the alternate choice would then be two loops and a pleat on the Gaussian map. Because these are viable alternatives, we conclude that our preference rules are tentative and seek verification from shading, stereopsis, etc. Also it is clear that, at some level, the metrics of the outline, its orientation, and the familiarity of the shape come into play in our judgment. At present, we are only proposing interpretation rules, given no information other than the general topology of the Gaussian map.[20]

CONCLUSIONS

We invoke four simple constraints for interpreting smooth 2-D outlines as 3-D shapes. Three are simply mathematical properties of smooth surfaces, namely, (1) that the sign of curvature of the silhouette reflects the sign of Gaussian curvature, (2) that the flexional loci are closed and nonintersecting, and (3) that these loci form either loops or pleats on the Gaussian map. The fourth constraint is an interpretation rule, stating (4) that undulations not seen in the 2-D outline are not present in three dimensions. This rule is an extension of the notion of general position. A corollary to this interpretation rule is that necks on the surface, which correspond to loops on the Gaussian map, are preferred over furrows, which in turn are preferred over dents.

ACKNOWLEDGMENTS

This work was supported by a North Atlantic Treaty Organization grant for international collaboration and by the U.S. Air Force Office of Scientific Research under an Image Understanding contract F49620-83-C-0135. Comments by Michael Leyton and Jim Little were greatly appreciated. William Gilson provided technical help in the preparation of the manuscript.

REFERENCES AND NOTES

1. D. Hoffman and W. Richards, "Representing smooth plane curves for recognition: implications for figure-ground reversal," *Proceedings of the National Conference on Artificial Intelligence* (American Association for Artificial Intelligence, Menlo Park, Calif., 1982), pp. 5–8.
2. D. Hoffman and W. Richards, "Parts of recognition," A.I. Memo No. 732 (Massachusetts Institute of Technology, Cambridge, Mass., 1984); Cognition 18, 65–96 (1984).
3. J. M. H. Beusmans, D. D. Hoffman, and B. M. Bennett, "Description of solid shape and its inference from occluding contours," J. Opt. Soc. Am. A **4**, 1155–1167 (1987).
4. W. Richards and D. Hoffman, "Codon constraints on closed 2D shapes," A.I. Memo No. 769 (Massachusetts Institute of Technology, Cambridge, Mass., 1984); Comput. Vis. Graphics Image Process. **31**, 265–281 (1985).
5. W. Hamscher, MIT Artifical Intelligence Laboratory, Cambridge, Massachusetts, 02139, "Codon constraints on 2D cusps," (personal communication, 1986).
6. M. Leyton, "A process-grammar for shape," Artificial Intell. **23** (to be published).
7. D. Marr, "Early processing of visual information," Philos. Trans. R. Soc. London Ser. B **275**, 483–524 (1976).
8. D. Hilbert and S. Cohn-Vossen, *Geometry and the Imagination* (Chelsea, New York, 1952) (translated by P. Nemeny).
9. J. J. Koenderink and A. J. Van Doorn, "The singularities of the visual mapping," Biol. Cybernet. **24**, 51–59 (1976).
10. Little has proposed a notation that removes this ambiguity. We will simply use the word bump or dent. See J. J. Little, "Recovering shape and determining attitude from extended Gaussian images," Tech. Rep. TN85-2 (Department of Computer Science, University of British Columbia, Vancouver, British Columbia, Canada, 1985).
11. A Cayley, "On contour and slope lines," London Edinburgh Dublin Philos. Mag. J. Sci. 18(120), 264–268 (1859).

12. J. C. Maxwell, "On hills and dales," London Edinburgh Dublin Philos. Mag. J. Sci. **40**(269), 421–427 (1870).
13. J. J. Koenderink and A. J. van Doorn, "Photometric invariants related to solid shape," Opt. Acta **27**, 981–996 (1980).
14. J. J. Koenderink and A. J. van Doorn, "A description of the structure of visual images in terms of an ordered hierarchy of light and dark blobs," in *Proceedings of the Second International Visual Psychophysics and Medical Imaging Conference* (Institute of Electrical and Electronics Engineers, New York, 1981).
15. B. K. P. Horn, "Extended Gaussian images," Proc. IEEE **72**, 1671–1686 (1984).
16. H. Whitney, "On singularities of mappings of Euclidian spaces. I. Mappings of the plane into the plane," Ann. Math. **62**, 374–410 (1955).
17. T. Banchoff, T. Gaffney, and C. McCrory, *Cusps of Gauss Mappings* (Pitman, Boston, 1982).
18. Of course, other interpretations of the peanut outline are still possible, given our interpretation rule and constraints, namely, a planar-shaped pond or a potato chip. Why do we tend to infer the 3-D object over these two other possibilities? Perhaps one explanation is that neither of the two alternatives is truly a general position view of a 3-D surface. For example, the outline of the potato chip or that of a planar region has little to do with the actual shape of the surface. This is not the case for opaque 3-D objects, however. Hence, if one is given a 2-D outline from which one is to infer 3-D shape, only opaque, volumetric 3-D shapes have powerful enough constraints on their projections to permit a unique inference.
19. We can continue this taxonomy of 3-D parts further to define knobs, caps, cavities, tongues, grooves, etc. by placing restrictions on the relations between the normals at the folds, as was done for the neck and knuckle. Such a scheme is an extension of that proposed by Maxwell[12] but is beyond the scope of this paper.
20. J. J. Koenderink, "An internal representation for solid shape based on the topological properties of the apparent contour," in *Image Understanding 1985–86*, W. Richards and S. Ullman, eds. (Ablex, Norwood, N.J., 1987).