

MASSACHUSETTS INSTITUTE OF TECHNOLOGY  
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 840


April 1985

## INFERRING 3D SHAPES FROM 2D CODONS

Whitman Richards  
Jan J. Koenderink  
D.D. Hoffman

**Abstract:** All plane curves can be described at an abstract level by a sequence of five primitive elemental shapes, called "codons", which capture the sequential relations between the singular points of curvature. The codon description provides a basis for enumerating all smooth 2D curves. Let each of these smooth plane curves be considered as the silhouette of an opaque 3D object. Clearly an infinity of 3D objects can generate any one of our "codon" silhouettes. How then can we predict which 3D object corresponds to a given 2D silhouette? To restrict the infinity of choices, we impose three mathematical properties of smooth surfaces plus one simple viewing constraint. The constraint is an extension of the notion of general position, and seems to drive our preferred inferences of 3D shapes, given only the 2D contour.

**Acknowledgments:** This report describes research done at the Department of Psychology and the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for this work is provided by a NATO grant for international collaboration and by an AFOSR Image Understanding contract F49620-83-C-013J. William Gilson provided technical help with the preparation of the manuscript. J.J. Koenderink's address is the Dept. of Medical Physics, Rijksuniversiteit, Utrecht, the Netherlands; D.D. Hoffman is at the Cognitive Science Group, University of California, Irvine. Helpful comments were received from W.E.L. Grimson and J.J. Little.




**Figure 1** The primitive codon types. Zeroes of curvature are indicated by dots, minima by slashes. The straight line ( $\infty$ ) is a degenerate case included for completeness, although it is not treated in the text.

## 1.0 Introduction: The Codon Representation

Codons are a very abstract and primitive way of describing a wiggly plane curve, such as the silhouette of an opaque object. Given these primitive image descriptors of such a silhouette, our goal is to infer the 3D shape that created it. The codon representation is useful for this task because it makes explicit certain features of an outline that permit qualitative 3D interpretations, such as its Gaussian curvature.

Codons are defined in terms of the relations between the maxima, minima and zeroes of curvature encountered as one traverses a plane curve (Hoffman & Richards, 1982, 1984). The boundaries between the codons are taken at the positive minima and negative maxima of curvature. The sign of curvature is determined by a clockwise ( $-$ ) or counterclockwise ( $+$ ) rotation of the tangent as one traverses the curve, keeping the "figure" to the left. The codon type is then specified by the presence or absence of an inflection (zero of curvature) and whether the positive maxima of curvature occurs before or after the inflection, when present. There are only five possible codon types, which are illustrated in Figure 1. Because any plane curve can be described at an abstract level by the sequence of extrema and inflections of curvature, so it can also be represented by a codon string.

A class of plane curves of special interest are those which are smooth, not self-intersecting and closed, for these can represent the silhouettes of certain simple smooth opaque objects seen in a perspective view. Clearly there are an infinity of such closed curves, even if the scale differences of the bumps and dents are ignored. This infinity can be enumerated simply by constructing closed shapes from a longer and longer sequence of codons. Some examples are shown in Figure 3. Our goal is to determine which smooth, opaque, 3D shapes are represented by each of these 2D figures. Or more generally, what constraints would allow any 2D codon string to have a unique interpretation as a 3D shape?



**Figure 2** Legal smooth, closed codon pairs. Figure is indicated by cross hatching. Part boundaries are noted by the slashes.

## 2.0 Constraints on Codons

We are considering objects whose silhouettes are smooth, and hence which can be represented by a string of codons with smooth joins. Our first step is to enumerate all of these silhouettes, for these will be the data from which the 3D inferences will be made.

Let us define the complexity or order of a closed plane curve or "outline" as the number of codon elements in the description. Then the simplest closed curve will be an ellipse, constructed from two codon elements of type  $0^+$  or  $0^-$ . The  $0^+ 0^+$  ellipse is our typical ellipse, as shown in the upper left panel of Figure 2. The direction of traversal of the curve places "figure" inside, and the two codon boundaries are the positive minima of curvature, one of which is roughly located at the arrowhead. The second elliptical construction,  $0^- 0^-$ , is an elliptical hole with figure outside the outline. Now the codon boundaries are negative minima of curvature, as indicated by the tick marks. (Such negative minima are taken as the part boundaries when "figure" is contained by the outline.)

Continuing our enumeration, only three other codon pairs have smooth joins, namely  $(1^+ 1^-)$ ,  $(2 2)$  or  $(2 0^-)$ . But these pairs yield only two new outlines, namely a "peanut" or a "dumbbell" as shown in Figure 2. These are the only possible smooth outlines of order 2. Figure 3 shows all legal smooth outlines for codon triples and quadruples. Quite surprisingly, we find that when plane curves are constrained to be smooth, closed and not self-intersecting, then there are only a relatively few outlines possible as compared with the total number of combinations of unconstrained codon sequences (Richards & Hoffman, 1984). For example, if we joined codons arbitrarily, then for strings of length (order) 4, we expect  $5^4$  or 625 possibilities of which only nine satisfy our simple constraint that the outline




Figure 3 Legal smooth, closed codon triples and quadruples. The tick marks indicate the minima of negative curvature, which are the part boundaries, whereas the dots show the inflections.

be smooth and closed. These few 2D outlines, then, are the silhouettes from which we wish to infer 3D shape.

### 3.0 The "Canonical" View

We begin by examining the simple outlines of Figure 2, the "ellipse", "peanut" and "dumbbell", for the analysis of these simple silhouettes provides us with the tools needed to interpret the more complex shapes of Figure 3. 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 are occluded or invisible. To capture this notion we propose our first interpretation rule:

R1: Do not propose undulations of the 3D surface without evidence for such. (1)

The above rule is an extension of the “general position” restriction, which requires that the viewpoints under consideration be “generic”. In particular, a slight shift in the viewpoint should not change the topology of the viewed structure, such as suddenly revealing a bump or dent in the surface that was previously hidden by occlusion. Our viewing rule also implies something else about the object and its disposition, namely that *all* of the undulations of the surface are visible. We call this special generic view of an object the “canonical” view.

DEFN: A “canonical” view of a 3D object is a (generic) view with a smooth silhouette (without cusps) that reveals all the undulations of the surface. (2)

By this definition, not all objects will have a canonical view.<sup>1</sup> However, if an object does have a canonical view, then it will be unique, at least in the topological sense. To clarify further what we mean by a canonical view of an object, let us form a dumbbell from the ellipsoid by squeezing it about the middle to create a neck joining two ovoid shapes. Now when this dumbbell is seen end-on from a close vantage point, the silhouette will be a simple circle. This is a generic view under perspective projection, but not a canonical view by our definition because neither the neck of the dumbbell nor the second end are visible. There exist (three) other generic views which reveal both protrusions of the dumbbell. Two of these views produce cusps in the silhouette, similar to that shown in the peanut in Figure 4. The third view yields a smooth silhouette. This view, which is perpendicular to and bisects the symmetry axis of the dumbbell, is non-generic under parallel (or orthographic) projection but becomes generic under perspective viewing. This latter, perspective view is the canonical view of the dumbbell. (Except for the ellipsoid, all canonical views become non-generic under orthographic projection, because a slight change in viewing angle will introduce cusps in the silhouette.)

#### 4.0 Gaussian Curvature

The ellipse is a very special outline, for its sign of curvature everywhere is the same, namely positive, because it has no undulations. By our interpretation rule, the only viable 3D shape will be the ellipsoid, which likewise has no undulations of its surface. The peanut and dumbbell are more complex, however, with bumps and dents. Clearly we need a better

<sup>1</sup>The concept of “canonical” view can be relaxed if necessary for other purposes. In this case, one might speak of a “most general position” view.

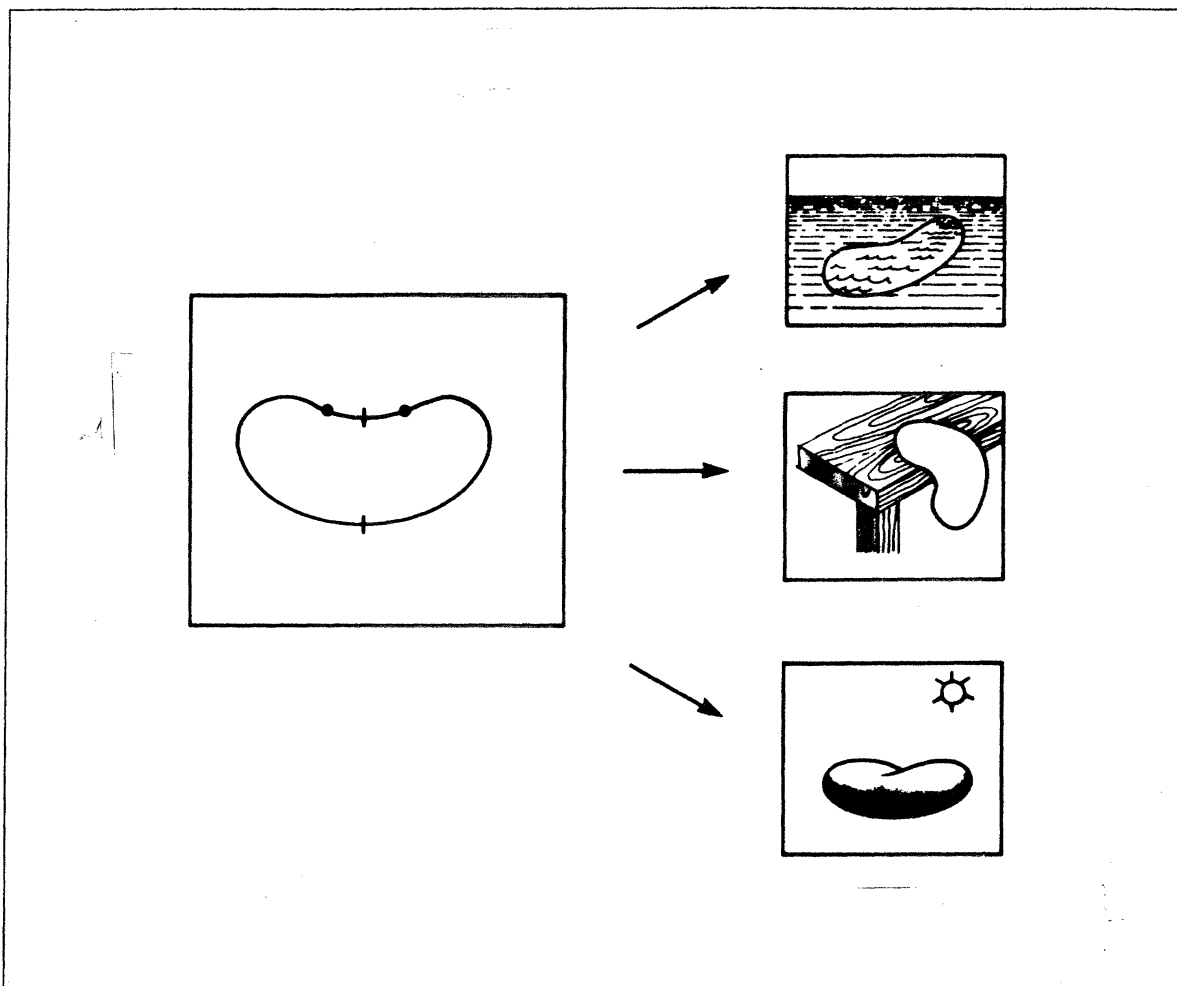


Figure 4 The simple codon shape shown at the left can be the outline of a planar "POND" (upper right), a "POTATO CHIP" (middle right) or a "PEANUT" (lower right). Only the peanut interpretation is plausible, given the constraint of "most general position" viewing.

means of describing these 3D shapes with undulations in order that we can enumerate all possible 3D interpretations of the outlines. We choose for this purpose an intrinsic property of all (non-fractal) 3D surfaces, namely the Gaussian curvature.

At any point on a smooth (generic) surface, there is a direction where the surface curves the most and another direction where the surface curves the least. (A plane is nongeneric.) These two directions are the directions of principal curvature and they are always perpendicular (Hilbert & Cohn-Vossen, 1952). The Gaussian curvature is simply the product of these two curvatures. Of interest to us is the sign of Gaussian curvature, which permits a qualitative description of the topology of a surface. When the direction of both principal curvatures is 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 region of a surface will thus have either positive, negative, or zero Gaussian curvature, depending upon whether it is locally elliptical, hyperbolic (saddle), or cylindrical. A 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 peanut is a hyperbolic region (saddle) lying within an ellipsoid of positive Gaussian curvature.

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

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

This theorem by Koenderink and van Doorn (1976) thus assures us that the Gaussian curvature of the 3D shape is positive at points on the surface that project into regions of positive curvature on the silhouette. Thus, both the peanut and dumbbell outlines of Figure 2 require that the corresponding 3D shapes have hyperbolic (saddle) regions of negative Gaussian curvature within a region (or two) of positive Gaussian curvature. We now need to enumerate the possible ways in which regions of positive and negative Gaussian curvature can be embedded within one another.

## 5.0 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 the surface, will have positive Gaussian curvature. 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. Similarly, there will be two kinds of regions of negative (hyperbolic) Gaussian curvature, which we will call "furrows" and "ridges". These latter undulations may be visualized as the folds and creases of the wrinkles on one's shirt sleeves. A "furrow" is formed when a bent finger is pressed across, but not into, a soft clay ball. A "ridge" is its complement, such as the fold of a wrinkle in a hollow. We thus have four types of surface undulations: bumps and dents, and furrows and ridges, where the latter two should be interpreted as part of "wrinkles". Together, they form a complete qualitative description of any smooth surface (Cayley, 1859; Maxwell, 1870; Koenderink & Van Doorn, 1980, 1981).

We next need a simple scheme for representing the relations between the four types of surface undulations. For this we choose the Gaussian sphere. Our 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 3D 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 (Hilbert & Cohn-Vossen, 1952). Each point on the surface of this sphere thus corresponds to a particular orientation (see Horn, 1984, 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 which gives rise to the silhouette must strike the surface in such a way as to be perpendicular to the surface normal at the point of contact. All surface normals on the surface contours which 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 (see Figure 5).

Let us now trace the surface normals of the silhouette of the peanut onto the appropriate great circle of the Gaussian sphere. Referring again to Figure 4, the outline of the peanut has two inflections (dots) bounding a region of negative curvature. The surface normal at the bottom of the well is vertical, as in the same direction as the normal to the two "bumps" at the top of the peanut. Thus there is one point on the Gaussian sphere that corresponds to three points on the object's surface. In Figure 5 right, the dashed contour is the locus of the silhouette of the peanut. The ends of the two "folds" are the inflection points (dots) on the silhouette where the direction of rotation of the surface normal 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.

For the class of generic surfaces 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 (Koenderink & Van Doorn, 1981). 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 (1980, 1981) have proven the following important property of flexional lines:






**Figure 5** The outline of a 3D shape seen under parallel projection maps onto a great circle on the Gaussian sphere. A "peanut"-shaped outline with two inflections will have two folds on the Gauss map, as shown by the dashed line. The extremities of these folds must lie on flexional lines of zero Gaussian curvature. The three insets at right show possible ways of closing the two flexional loci that intersect the folds on the great circle. The (lower) examples are excluded by the silhouette of the peanut.

**C2:** For generic surfaces, the flexional lines are closed and non-intersecting. (4)

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

Returning to our peanut example in Figure 5, this means that the ends of the two folds must lie on closed curves. There are only two ways we can close flexional curves on the Gaussian sphere: either we can join the ends in a smooth loop, or we can create a cusp (Whitney, 1955; Koenderink & Van Doorn, 1981; Banchoff *et al.*, 1982). If the flexional curve is closed with a cusp on the Gauss map, then this specifies a wrinkle on the surface (*i.e.*, a furrow or ridge), and there must be a cusp on the opposite side of the great circle locus of the silhouette (dashed lines). This case is illustrated in the top right of Figure 5. (We call this a "pleat".) If the flexional curves are closed by simple loops, then we have two other possibilities as shown above, namely either one or two loops. The single loop would have to correspond to the flexional curve of a dent. But a dent has positive curvature and is excluded by Koenderink's theorem (3) because the silhouette has negative curvature in this



**Figure 6** The "dumbbell" outline has two pairs of folds when mapped onto the Gaussian sphere (left). Closed non-intersecting flexional lines can thus be created in three ways: either by "fusing" the folds on each side of the sphere, creating four cusps (upper), by creating two "rings" and no cusps (lower), or by fusing the rings on one side of the great circle, leaving a gap in the fused contours to create two cusps. The "fusing" versions are like an ellipsoid with one or two separate dents, whereas the "ring" version is the true dumbbell.

region. We now are left with only the possibility of joining the folds on the Gauss map with two loops (lower right). But two loops on the Gauss 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 Gauss 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" of negative Gaussian curvature, which is the 3D peanut.<sup>2</sup>

<sup>2</sup>Of course, other interpretations of the "peanut" outline are still possible, given our interpretation rule and constraints. These were shown in Figure 4, namely a planar shaped "pond" or a "potato chip". Why do we tend to infer the 3D object over these two other possibilities? Perhaps one explanation is that neither of the two alternatives are truly general position views of three-dimensional surfaces. A plane is clearly degenerate. But also so is the developable surface of the potato chip, for in each case the "outline" has little to do with the actual shape of the surface. This is not the case for opaque 3D objects, however. Hence given a 2D outline from which one is to infer 3D shape, only opaque, volumetric 3D shapes have powerful enough constraints on their projections to allow a unique inference.

## 6.0 The “Dumbbell” (or “Pear”)

The analysis of the dumbbell silhouette now follows quite simply. Its Gauss map is shown in Figure 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” to create a 3D dumbbell. However, our interpretation rule (1) that the 3D object be seen under a canonical view excludes the egg with one or two dents, for such an interpretation requires that the egg be viewed in a special way. Thus the preferred 3D interpretation for the dumbbell outline is the “dumbbell”.

## 7.0 Loops and Pleats

A strategy for enumerating the 3D shapes now emerges. When the silhouette is mapped onto the Gaussian sphere, the inflections of the silhouette map into folds on the Gauss 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: either by forming “loops” or “pleats” that create cusps on the Gauss map. The “loop” choice will correspond either to a “bump” or “dent” on the surface or to a “neck” or “knuckle” in the 3D shape, the latter cases requiring an inflection to appear at (roughly) opposite sides of “part” of the silhouette. The second, “pleat” scheme corresponds either to a “ridge” or “furrow” on one side of the 3D shape. Whether the region between the flexional loci on the Gauss map is bump or a ridge, dent or furrow, etc., is given by the sign of curvature of the silhouette (3). In total, our “parts” of a surface will be either a “bump (+)”, “ridge (-)”, “knuckle (+)”, or a “neck” (-), or “dent” (+) or “furrow” (-).<sup>3</sup>

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 Gauss map (*i.e.*, furrow in the 3D shape) because the loops would create two more inflections on the silhouette which are not seen and hence are inferred not to be present. Referring to Figure 3, 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 or Q6, we still only have the “loop” and “pleat” possibilities to consider as we did before with the dumbbell. The preferred choice is two “loops” on the Gauss map, rather than two pleats,

<sup>3</sup>We can continue this taxonomy of 3D parts further to define “knobs”, “caps”, “cavities”, “tongues”, “grooves”, etc., by placing restrictions on the relations between the normals at the folds, such as was done for the “neck” and “knuckle”. Such a scheme is an extension of that proposed by Maxwell (1870), but is beyond the scope of this paper.

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.


Our present constraints are not powerful enough to find a preferred interpretation for shape Q8, however. Unlike 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 3D outline. Thus, the two indentations of Q8 can not correspond to a “neck” in a surface of revolution. Reviewing possible pairings once more, we have on the Gauss map the following seven possibilities [case (v) can occur in two different ways]:

- (i) two adjacent pleated contours (two furrows)
- (ii) two loops (a bump in a dent)
- (iii) a pleat enclosed by a loop (a ridge in a dent)
- (iv) a pleat enclosed by a pleat (a ridge in a furrow)
- (v) two adjacent loops (two dents)
- (vi) adjacent loop and pleat (a dent and a furrow)
- (vii) a loop enclosed by a pleat (a bump in a furrow).



Clearly, as before with the peanut silhouette, any interpretation requiring a dent can be excluded, for this interpretation violates theorem (3). Thus, (ii), (iii), (v) and (vi) may be excluded. A pleat in a pleat (iv) is also impossible because both a ridge and a furrow have negative Gaussian curvature. Hence we are left with either (i) or (vii). But (i) seems like an unlikely preference, for the presence of two separate dents or furrows of different orientations would generally require a more special view than a single furrow. Hence the preferred silhouette must be (vii), which is a “bump” in a furrow. Figure 7 (right) depicts the form of this shape, which also applies to quadruple Q9.

We can now complete the preferred 3D shapes for the remaining closed codons. Consider next the “jack” T5 which has six inflections. If inflections of two adjacent parts are paired then only pleats (cusps) on the Gauss map are legal, indicating a sphere with three furrows. If inflections are paired inside the shape, then only loops are legal, indicating a shape with three (symmetrical) limbs. However, loops give surfaces of revolution that are topologically more invariant under changing viewpoints than the furrows associated with pleats on the Gauss map. Hence we choose the three “interior” loops to describe the 3D shape of T5, and also that of quadruple Q7. Each shape corresponds to three



**Figure 7** A top view looking down upon the 3D surface of the "spade" silhouette Q8. The dotted lines are the possible flexional loci, the signs indicate Gaussian curvature. The constraints exclude a dent with a bump (left) as the preferred interpretation. Instead, the preferred interpretation is a furrow with a bump, as shown at the right.

"protrusions", with the third (lower) protrusion of Q7 being its "body". (No strong assertion can be made about the relative disposition of these protrusions to one another.) Similarly, the quadruple Q5 has an analogous 3D shape, but with four "protrusions". The regions of negative curvature on the silhouette then indicate by default that saddle regions join the protrusions.

## 8.0 Discussion

### 8.1 Effect of Constraints

Our rules (1) through (4) have allowed us to assign one preferred 3D interpretation to each of the silhouettes shown in Figures 2 and 3. These silhouettes are all the possibilities for smooth codon strings of length four or less. Without constraints, these are 625 different sequences of four codons, or 625 possible 2D outlines that are generically different (*i.e.*, have different sequences of the extrema of curvature). Closure and smoothness reduce this number to the nine outlines shown in Figure 3. In a similar manner, we can ask, given a single outline in Figure 3, how many different 3D shapes are possible? Clearly there are an innumerable number. However, if we define the "generic" class of 3D shapes as those with no hidden undulations (2), then we can count the generic 3D possibilities.

The Gauss map is the simplest way to describe the counting procedure. An inflection in the silhouette corresponds to a fold in the locus on the Gaussian sphere (Figure 5). These folds must occur in pairs. The pairs are joined in only three ways: by loops or

<u>NUMBER OF INFLECTIONS</u>	<u>NUMBER OF UNCONSTRAINED PAIRINGS</u>	<u>MINIMUM NUMBER OF FLEXIONALLY CONSISTENT 3D SHAPES</u>	<u>NUMBER OF PREFERRED 3D SHAPES</u>
2	3	3	1
4	27	18	1-2
6	405	81	1-2
8	8,505	486	*
10	229,635	2,916	*
12	7,577,955	17,496	*

\*NOT ENUMERATED

Figure 8 Table 1

pleats, or a hemi-loop and a single cusp. Thus for an outline with  $N$  inflections, we have  $(N-1)(N-3)(N-5)\dots$  possible "groups" of pairings, and each group has  $N/2$  pairs. For each group we thus have  $3^{N/2}$  possible *unconstrained* combinations of pleats and loops. Our total possible unconstrained 3D generic shapes is therefore

$$\text{Number of generic 3D shapes} = 3^{N/2} \cdot (N-1)(N-3)(N-5)\dots \quad (6)$$

where  $N$  is the number of inflections on the silhouette. Table 1 (Figure 8) shows how fast these possibilities increase as the outline becomes more complex.

If we now introduce the simple mathematical constraint that the flexional loci can not intersect on the surface (4), the possibilities are markedly reduced, but still substantial, as shown in column 3 (see Appendix 1 for derivation). However, when our canonical viewing constraint is added, together with the stipulation that "loops" are given preference over "pleats", then our preferred interpretations become unique, at least for the 2D outlines examined. Four simple rules, therefore, provide very powerful constraints on the interpretation of silhouettes.

## 8.2 *Instant Psychophysics*

Of course the 2D outlines of Figure 3 do 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 bump in a shallow furrow (Figure 7). This choice follows because "loops" are given preference over "pleats". However, it is easy to imagine outline Q8 as a convex blob with two furrows, which corresponds to two pleats on the Gauss map. Hence we conclude that our rule for preferring loops over pleats is a weak one, and seeks verification from shading information or stereopsis, etc. Similarly outlines T5 or Q7 have alternate 3D interpretations 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 Gauss map. Clearly 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 Gauss map.


## 9.0 Conclusion

We invoke five simple constraints for interpreting smooth 2D outlines as 3D shapes. Three are simply mathematical properties of smooth surfaces, namely (i) that the sign of curvature of the silhouette reflects the sign of Gaussian curvature, (ii) that the flexional loci are closed and non-intersecting, and (iii) that these loci form either loops or pleats on the Gauss map. The fourth constraint is an interpretation rule that the object is seen under a canonical view, which is an extension of the notion of general position. This fourth constraint has the corollary that undulations not seen in the 2D outline are not present in 3D. Finally, our fifth constraint is another interpretation rule that loops on the Gauss map are preferred over pleats.

## REFERENCES

- Banchoff, T., Gaffney, T. & McCrory, C. (1982) *Cusps of Gauss Mappings*. Pitman, Boston.
- Cayley, A. (1859) On contour and slope lines. *Lond. Edinburgh, and Dublin Phil. Mag. and J. Sci.*, 18(120), Oct., Ser. 4, 264-268.
- Hilbert, D. & S. Cohn-Vossen (1952) *Geometry and the Imagination*. Trans. P. Nemeny. Chelsea Publishing Co., New York.
- Hoffman, D. & Richards, W. (1982) Representing smooth plane curves for recognition: Implications for figure-ground reversal. *Proc. National Conference on Artificial Intelligence, August 18-20, 1982, Carnegie-Mellon University-University of Pittsburgh*.
- Hoffman, D. & Richards, W. (1984) Parts of recognition. *MIT AI Memo 732, and Cognition*, 18:65-96 (1984).
- Horn, B.K.P. (1984) Extended Gaussian images. *Proc. IEEE*, 72(12), 1671-1686.
- Koenderink, J.J. & Van Doorn, A.J. (1976) The singularities of the visual mapping. *Biol. Cybernetics*, 24, 51-59.
- Koenderink, J.J. & Van Doorn, A.J. (1980) Photometric invariants related to solid shape. *Optica Acta*, 27, 981-996.
- Koenderink, J.J. & Van Doorn, A.J. (1981) A description of the structure of visual images in terms of an ordered hierarchy of light and dark blobs. *Second Int. Visual Psychophysics and Medical Imaging Conference. IEEE Cat. No. 81CH 1676-6*.
- Koenderink, J.J. (1985) An internal representation for solid shape based on the topological properties of the apparent contour. In *Image Understanding 1985*, W. Richards & S. Ullman (eds.). Ablex, Norwood, NJ.
- Maxwell, J.C. (1870) On hills and dales. *Lond. Edinburgh, and Dublin Phil. Mag. and J. Sci.*, 40 (269), Dec., Ser. 4, 421-427.
- Richards, W. & Hoffman, D. (1984) Codon constraints on closed 2D shapes. *MIT AI Memo 769*. To appear in *Computer Vision, Graphics, and Image Processing* (A. Rosenfeld, ed.).
- Whitney, H. (1955) On singularities of mappings of Euclidian spaces. I. Mappings of the plane into the plane. *Ann. Math.*, 62, 374-410.





**Figure 9** The dots represent the inflections, which are to be paired without the joins intersecting. We always start with link (a). Then consider the remaining  $N - 2$  points. Pairs (b) and (c) are topologically different with respect to pair (a). But (c) is simply a rotation of (b) if (a) is ignored.

## Appendix 1

Given  $N$  inflections on a silhouette, we wish to count the minimum number of topologically different 3D shapes, given that the flexional loci can not cross on the surface. We begin by counting the number of topologically different flexional pairings, for  $N$  inflections. The total number of such constrained 3D shapes will be three times this count, because each pairing can be either a single loop, a pleat, or a hemi-loop and a single cusp on the Gauss map (ignoring front-back similarities).

### Case 1: Regular Tessellation

For  $N$  inflections, there are  $(N - 1)(N - 3)(N - 5) \dots$  possible groups of pairings. Some of these groups may have links that will cross, however. For example, if the inflections lie at the vertices of a square, then one of the three possible groupings will have intersecting joins, namely the one where the links are the diagonals of the square. To enumerate the number of non-intersecting pairings, we will count by induction.

Consider two points. Call this pair (a). For  $N = 2$ , the number of legal pairings  $P_2 = 1$ .

Consider four points, constructed by adding two points to the first pair (a) to form a regular lattice, as shown in Figure 9. Again form pair (a), which will always be our starting pair. Then the only choice is (b). [The top and bottom pairings of the square are excluded because they are simply a rotation of pairings (a)-(b) and do not represent different shapes.] However the flexional link (a) can be either across a concavity or a convexity of the surface, as can be readily seen by completing a smooth "dumbbell" contour through the four points. For  $N = 4$  we thus have as the number of legal pairings  $P_4 = 2$ .




Figure 10 A regular tessellation of points will have the largest number of crossed joins.

Now consider six points, created by adding two more points to the above four to form a hexagon. Again start by joining pair (a). This leaves  $N - 2$  or four points to be paired. But four points can be paired legally in only two ways, when they are in isolation. However pair (a) introduces an axial asymmetry to the remaining four points. With (a) as a reference, link (b) is not equivalent to link (c), which leaves (b) open. Although one 4-point configuration is simply a rotation of the other, they are not the same with respect to link (a). If link (a) is chosen to lie across a convexity, then there are two legal pairings. Similarly if link (a) lies across a concavity.<sup>4</sup> However one of the latter configurations has already been counted. (This can be seen by closing the points with a contour and assigning (+) and (-) labels to the regions to designate the surface curvature). For  $N = 6$ , we have for the number of legal pairings  $P_6 = P_4 + \frac{1}{2}P_4 = 2 + 1 = 3$ .

For eight points, we now have six unpaired points after reference link (a) is established. With link (a) across a convexity, we have only four legal pairings for these six points. As can be seen by referring to Figure 9, two of the legal pairings arise when link (b) is chosen. The other two occur when we rotate the six point configuration (b) by one pair to form configuration (c) that is different with respect to link (a). Thus for link (a) chosen across a convexity, there will be  $2 \cdot P_4$  groups. However, link (a) can also be across a concavity. This will add half the number again, because half the four-point configurations associated with (b) or (c) will have convex links that are equivalent to link (a). For  $N = 8$  we thus have for the number of legal pairings  $P_8 = 2 \cdot P_4 + \frac{1}{2} \cdot 2P_4 = 6$ .

In general, for each new pair, we will double the number of legal pairings for the case where link (a) is across a concavity. To this we add half again the number if link (a) is considered across a convexity.

<sup>4</sup>Without the convex-concave assignments, the number of non-intersecting pairs  $P_N$  for  $N$  inflection points will be  $(3 \cdot P_{N-2} - 1)$ .

Beginning with six points, we have for link (a) chosen convex  $2^{N/2-3} \cdot P_4$  or  $2 \cdot 2^{N/2-3}$  legal pairings. Adding half again, we have for the total number of legal flexional groups

$$P_N = 2 \cdot 2^{N/2-3} + 2^{N/2-3} = 3 \cdot 2^{N/2-3} . \quad (1)$$

with  $P_2 = 1$  and  $P_4 = 2$ . To count the number of possible 3D topologies, the above is multiplied by  $3^{N/2}$  as before. The values are given in the third column of Table 1.

### Case 2: Irregular Tesselations

Any arrangement of points that form a convex hull will be topologically equivalent for our purposes to the regular tessellation considered above. If one of the points lies within the hull of the others, however, the number of possible non-intersecting pairs will increase. The four point case is illustrated in Figure 10. If the fourth point is placed within the triangle formed by the other three, then all linking arrangements become legal. (For six points, formed by one triangle within another, only three out of the possible 15 groups have intersecting links.) The regular tessellation will give the lowest count of groups with non-intersecting links because each link will span the convex hull. Thus the counts given in Table 1 are the minimum number given the constraint that flexional lines don't intersect.