

Perception of 3-D Surfaces from 2-D Contours

Fatih Ulupinar and Ramakant Nevatia, *Fellow, IEEE*

Abstract— Inference of 3-D shape from 2-D contours in a single image is an important problem in machine vision. We survey classes of techniques proposed in the past and provide a critical analysis. We propose that two kinds of symmetries in figures, which are known as parallel and skew symmetries, give significant information about surface shape for a variety of objects. We derive the constraints imposed by these symmetries and show how to use them to infer 3-D shape. We discuss the zero Gaussian curvature (ZGC) surfaces in depth and show results on the recovery of surface orientation for various ZGC surfaces.

Index Terms— Shape from contour, symmetry analysis, zero Gaussian curvature surfaces.

I. INTRODUCTION

ONE OF THE BASIC goals of midlevel vision is to recover the local orientations of the surfaces of the objects in a scene. The basic difficulty, of course, is that an image is a 2-D projection of the 3-D scene, and hence, the 3-D structure cannot be recovered without some assumptions. In spite of the inherent ambiguities in a single view, humans are able to perceive 3-D surfaces in single images. Much effort has been devoted in the past few years to understanding this ability and has led to development of techniques such as shape from shading, shape from texture, and shape from contour (sometimes also known as shape from shape).

We believe that of all the monocular cues, the shape of the 2-D contours is the most important one for the shape of the 3-D surfaces. Strictly speaking, not only is such interpretation infinitely ambiguous, but the contours can only give shape information near the contours; shape of the surface in between can vary smoothly without producing other contours. Nonetheless, humans, when presented with contours of various and not necessarily familiar objects, perceive complete surfaces (and even solids); some examples are given in Fig. 1.

Barrow and Tenenbaum [2] have shown, by some examples, that in case of conflict between contour and shading information, humans use the contour information. Beiderman [1] claims that in the experiments with humans, the recognition of a full-colored image of an object is not faster than the

recognition of the line drawing of the object. He also shows that we do not necessarily need to have any familiarity with the object in order to perceive a shape from its boundary. We conjecture that the reason for preferring shape from contour over other cues, such as shading, may be that even though shape from contour methods require some assumptions, other methods require even stronger assumptions. Shape from shading methods, for example, need to assume that the reflectance properties of the surface are known, that the albedo is constant, and that the light sources are known precisely.

These observations are not to argue that only shape from contour is important but that it is an essential element in monocular perception that cannot be ignored. We believe that such an ability will also be essential for machine vision systems if they are to work with monocular images in the absence of highly specific models.

The early work on inferring 3-D structure from a 2-D shape was focused on analysis of line drawings of polyhedra [6], [4], [12], [7], [9]. In polyhedral scenes, the problem is that of segmentation and estimating orientations of faces. Recently, some papers (e.g., [17], [15]) have been published that study the projection and geometry of certain classes of curved surfaces. However, the analysis is not sufficient to recover the 3-D surfaces. Techniques for nonpolyhedral scenes have been proposed in [2], [26], [3], [20], [27], [5]. We discuss these in more detail in Section II and contrast them with our work later. Most of these techniques examine a single surface in the scene, whereas human perception of one surface can be strongly influenced by the perception of the other surfaces comprising the entire object.

Our technique is based on the analysis of symmetries in a scene. We show how certain symmetries can be used to infer surface orientations. We also conjecture that humans perceive only slight surface orientation information in figures lacking symmetries. Our work may be viewed as being based on generalizations of concepts that have been used previously such as by Kanade [7] for polyhedral scenes and by Stevens [20] and Xu and Tsuji [27] for curved surfaces. We believe that our method is of wide applicability. In particular, we provide a rather complete analysis for the case of “zero-Gaussian curvature” (ZGC) surfaces, but many parts of the technique also extend to nonZGC surfaces [24]. We believe that the universe of ZGC surfaces is significant (all the examples in Fig. 1 are included) and that they define a reasonable step up in complexity from planar surfaces on which much of previous work in the field has focused.

Our method assumes that clean, closed boundaries are given (or can be extracted from the real image). We do not address the issue of separating object boundaries from

Manuscript received August 20, 1990; revised November 12, 1991. This work was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Air Force Office of Scientific Research under Contract F49 620-90-C-0078. The United States Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

F. Ulupinar was with the Institute for Robotics and Intelligent Systems, School of Engineering, University of Southern California, Los Angeles, CA 90089-0273. He is now with the Department of Computer Engineering, Bilkent University, Ankara, Turkey.

R. Nevatia is with the Institute for Robotics and Intelligent Systems, School of Engineering, University of Southern California, Los Angeles, CA 90089-0273.

IEEE Log Number 9205812.




Fig. 1. Some figures for which we readily perceive the shape from contour alone.

surface markings or other *perceptual grouping* operations here, although we believe that some of the previous work in our group on this topic is relevant [13]. In addition, we believe that the specific conditions needed for an object surface to be reconstructed by our method will provide further constraints for the perceptual grouping operations when surface markings and other noisy boundaries are present.

In Section II, we discuss previous related methods. In Section III, we define two kinds of symmetries and discuss the qualitative shape inferences that can be drawn from them. In Section IV, we describe our technique for quantitative shape recovery of Zero-Gaussian Curvature surfaces and state our conclusions in Section V.

We assume orthographic projection throughout the paper unless specifically mentioned otherwise (in a separate paper, we have shown how many of the constraints for orthographic projection can be transformed to the case of perspective projection [25]).

In this paper, we will use *gradient space* to represent the orientation of surfaces (given by their normals). To review, the normal N of a plane, $ax + by + cz + d = 0$ is given by the vector $N = (a, b, c)$. This can be rewritten as $(p, q, 1)$, where $p = a/c$ and $q = b/c$ (note that this excludes cases where $c = 0$; however, such planes are parallel to the line of sight and are not imaged as planes under orthographic projection anyway). (p, q) can be thought of as defining a 2-D space (the gradient space) such that every point in this space corresponds to the normal of a plane in 3-D.

II. PREVIOUS WORK

Here, we present an overview of the important classes of previous methods. We also give our view of their strengths and weaknesses. As our method builds on some of the previous work, this section will also help provide some of the relevant background for describing our work.

These previous methods fall into two broad classes. In the first, some property of the 3-D surface is made *extremal*. In the second class, which is known as *constraint based*, some geometric constraints are satisfied.

A. Extremal Methods

In this class of methods, a measure of some desirable property, such as smoothness or compactness, is associated with each interpretation (of a curve or a figure), and the one with the best value is chosen.




Fig. 2. Shape perception: (a) smooth curve; (b) two symmetric curves; (c) two nonsymmetric curves

Barrow and Tenenbaum [2] proposed to use the smoothness of the curve as the preference criterion. Their method is not restricted to planar curves. The smoothness function they propose is an integral function of curvature and torsion of the 3-D curve. Over all possible 3-D curves that can generate a given 2-D image curve, they choose the one that minimizes the aggregate curvature and the torsion of the 3-D curve. Besides being highly sensitive to noise, this method fails to utilize structural information given in the image. For example, in Fig. 2(a), the single curve does not give any strong 3-D shape impression, but in Fig. 2(b), where another curve that is perfectly symmetric to the first one is added, a definite shape is perceived. In Fig. 2(c), two nonsymmetric curves are displayed, and again, there is no definite shape perception.

Weiss [26] has proposed a modified measure that uses curvature rather than its derivative and handles polygons in a cleaner way. For planar curves, he proposes to minimize the integral of the square of the curvature, and for polygons, he proposes using the square of the angles of corners for polygonal scenes as a minimization criteria. This method also does not utilize the structure information available in the image.

Brady and Yuille [3] used the “compactness” of a figure as their preference criterion. The measure of compactness is chosen to be $(area)/(perimeter)^2$. This measure implicitly assumes that the curve is planar. This method is insensitive to small amounts of noise and processes smooth curves and polygons in a unified way. Although this method has many nice features from a computational point of view, it does not always give answers consistent with human perception. The method correctly interprets an ellipse as being a projection of a circle, but it interprets a rectangle in the picture as a slanted square, which is not the typical human perception. In addition, when the boundary is not complete, this method is undefined.

B. Constraint-Based Methods

In this class of methods, constraints on 3-D surface orientations are obtained by a variety of observations, with the expectation of getting unique (or a few) solutions when the various constraints are combined. In general, the constraints are based on some *assumption* that an observed regularity in the image corresponds to a regularity in the 3-D scene. We briefly survey the important techniques below.

A symmetric figure in 3-D, with a straight axis and lines of symmetry that are orthogonal to it, projects in a figure that is *skew symmetric* (under orthographic projection), i.e., the lines of symmetry are no longer orthogonal to the axis but are at a constant angle to it. Kanade [7] showed that if we assume that the inverse also holds, i.e., the observed skew symmetry in the image is due to orthogonal symmetry in 3-D, then the orientation of the plane is constrained to be along a curve (a

hyperbola) in the gradient space. Kanade combined the skew symmetry constraints with a constraint that is derived from the projection of a line formed by the intersection of two planes (we call this constraint *shared boundary constraint* in Section IV-A-1). The two constraints together are often sufficient to fix the orientation of the surface in the scene, and the answers given by this method appear to be consistent with human interpretation. Of course, this method applies only when skew symmetric objects are present.

Stevens [20] studied cylindrical surfaces using the orthogonality property. A cylindrical surface is one where one of the principal curvatures is zero, and the lines of zero curvature (the *rulings* of the surface) are parallel to each other. For such a surface, the lines of maximum curvature are planar and parallel to each other. Stevens assumes that the lines of maximum curvature are given. The rulings can be obtained from these by connecting points with the same tangent. The surface is thus covered by a grid of curves, with the property that on the actual surface, the curves are orthogonal at the points of intersection. Thus, constraints similar to those of the skew symmetry analysis can be applied. Stevens chooses to use the *slant*, σ and *tilt*, τ representation¹ instead of (p, q) representation.

As before, the constraint is not enough to give unique orientations. However, Stevens observes that slant and tilt parameters can be bounded and that the bound depends on the angle β between the two intersecting curves, with error in tilt approaching zero as β approaches π . This happens near the occlusion boundaries of a cylindrical surface. Starting from these points, where tilt can be fixed accurately, Stevens gives a method of propagating the estimates along the lines of maximum curvature by the following formula (given without proof here):

$$\tan \tau_1 \tan \beta_1 = \tan \tau_2 \tan \beta_2 \quad (1)$$

where τ_i are the tilt angles, and β_i are the angles between the lines of maximum and minimum curvature at two points along a line of maximum curvature.

This method, however, does not always give correct results, even when a circular cylinder is given to it, as can be shown by a simple example. Consider the two cylinders in Fig. 3. The points on the limb edges for both cases have β_1 equal to π , which give unambiguous values of zero tilt. Using Stevens method to extrapolate along the cross section, we will get the same orientations for the midpoints of the two cylinders where β_2 is $\pi/2$, which is in clear contradiction with our perception (which indicates that the top surface of (b) appears to be much more slanted to us than that of (a)).

Xu and Tsuji [27] have described an extension of this method for more general curved surfaces. In addition, their method does not require that all lines of curvature be given but rather that the surface be cut along these lines of curvature. Given a figure with four sides, such that two of the opposite sides are lines of maximum curvature and the other two are

¹Orientation of a surface, having gradient (p, q) , can be alternately described in terms of its *tilt*, τ and *slant*, σ , which can be viewed as the polar coordinates of a point in gradient space, specifically $\tau = \arctan \frac{q}{p}$, $\sigma = \arctan \sqrt{p^2 + q^2}$




Fig. 3. Two cylinders resting at different slant angles.

lines of minimum curvature, they show how a *net* (or a grid) over the figure can be constructed with the expected property that the corresponding net on the 3-D surface is orthogonal (this construction is strictly valid only for a restricted class of surfaces). Surface orientations are first computed at special points on the net where error in tilt is small (as in Stevens's method) and then propagated to other points on the net.

This method seems to work well in some cases but has some drawbacks. The propagation scheme for noncylindrical surfaces is only approximate, and errors can add up. In addition, it applies no test for whether the four sides of a figure *could* be lines of curvature, and the method may give answers completely inconsistent with human interpretation.

In a recent paper, Horaud and Brady [55] present a method for interpreting linear straight homogeneous generalized cylinders (LSHGC's). Their method makes the following assumptions: a) the axis of the LSHGC projects as the axis of the ribbon formed by the two limb contours in the image plane, b) the cross section of the LSHGC is planar, and c) the cross section is orthogonal to the axis in 3-D. Satisfying assumption b) above gives a constraint that the orientation of the cross section must be along a certain curve in the orientation space (the curve is shaped like the character "s" and, hence, is called an "s curve"). They also require that the back-projected cross section satisfy the Brady-Yuille compactness measure. If an orientation satisfies both constraints, then that orientation is chosen. They do not specify what should be done if this is not the case, although a natural extension would be to take the most compact shape constrained by the s curve. Finally, they suggest sweeping the reconstructed cross section along the 3-D axis to reconstruct the surface of the LSHGC.

This method has the attractive property that it attempts to combine the constraints from two surfaces. However, it has several deficiencies. The compactness measure can only be applied to complete cross sections. More seriously, assumption a) above is incorrect. Given the image of an LSHGC, we can choose any axis that passes through the apex of the LSHGC in the image plane and reconstruct an *orthogonal* LSHGC in 3-D, including the ends, that would project like the LSHGC in the image plane, as shown in the following. Take any back projection of the chosen axis, and back project the two cross sections (the top and the bottom) on any two planes orthogonal to the back-projected axis; the orthogonal LSHGC can be completed by joining the points on the cross sections such that




Fig. 4. Examples (a) and (c) show parallel symmetry with curved contours, and (b) shows parallel symmetry with straight contours. The dotted curves are axes of symmetry, and the dashed lines are lines of symmetry.

the lines joining these points pass through the back-projected apex. Finally, even if an axis is chosen in the image plane, the point through which the 3-D axis pierces the reconstructed 3-D cross section must be chosen in order to reconstruct the LSHGC. This point is not addressed in the Horaud and Brady paper.

In other work, Nalwa [15] has derived a symmetry condition that must be satisfied by the limb boundaries of a solid of revolution (sufficiency of these conditions is also claimed under certain general-viewing conditions). However, this paper does not show how to actually reconstruct the surface. Ponce *et al.* [17] have derived properties that must be satisfied by a broader class of surfaces known as straight homogeneous generalized cylinders (SHGC's). Again, this property by itself is not sufficient to reconstruct the 3-D surface of SHGC's.

III. SYMMETRIES AND QUALITATIVE INFERENCES

Our proposed technique is based on observations of symmetries in figures. We believe that symmetries play an important role in shape perception; this also has been noted and used by many researchers [16, [15], [8], [7], [20], [13]. We define two types of symmetries that we call *parallel symmetry* and *skew symmetry* and discuss how they can be used to infer surface shape.

A. Parallel and Skew Symmetries

For curves to be symmetric (parallel or skew), certain point-wise correspondences between two curves must exist. We will call the lines joining the corresponding points on the curves the *lines of symmetry*, the locus of the midpoints of these lines the *axis of symmetry*, and the curves forming the symmetry as the *curves of symmetry*.

1) *Parallel Symmetry*: Consider two curves $X_i(s) = (x_i(s), y_i(s))$, for $i = 1, 2$, parameterized by arc length s . Let $\theta_i(s) = \arctan((dy_i(s)/ds)/(dx_i(s)/ds))$. Then, $X_1(s)$ and $X_2(s)$ are said to be parallel symmetric if there exists a correspondence function $f(s)$ between them such that

$$\theta_1(s) = \theta_2(f(s)) \quad (2)$$

for all values of s for which X_1 and X_2 are defined, and $f(s)$ is a continuous monotonic function. Note that computing symmetry between two curves using this definition requires estimating the function $f(s)$ as well. A useful special case is when $f(s)$ is restricted to be a linear function. In that case, the symmetry condition becomes

$$\theta_1(s) = \theta_2(as + b) \quad (3)$$




Fig. 5. Example (a) shows skew symmetry with curved contours, and (b) shows skew symmetry with straight contours. The dotted curves are the axes of symmetry, and the dashed lines are lines of symmetry.

where a and b are constant (a may be thought of as a scale parameter). Some examples of parallel symmetry are given in Fig 4; the correspondence function $f(s)$ is linear for (a) and (b) but not for (c). Note that the above definition of parallel symmetry also holds for curves consisting of straight lines and corners as in Fig. 4(b).

2) *Skew Symmetry*: In this symmetry, the point-wise correspondence should be such that the axis of the symmetry is straight, and the lines of symmetry are at a constant angle (not necessarily orthogonal) to the axis of symmetry. Skew symmetry was first proposed by Kanade [7] and used in the analysis of scenes of polyhedral objects.

Ponce [18] has given point-wise conditions for two curves to be skew symmetric. We state these here without proof. For the case when the lines of symmetry are orthogonal to the axis of symmetry, the criterion for the two curves $X_1(s)$ and $X_2(s)$ to be skew symmetric is

$$\kappa_1(s) = -\kappa_2(s + b) \quad (4)$$

where $\kappa(s)$ is the curvature, and b is the offset. In general, let $\alpha_i(s)$ be the angle between the line of symmetry and the tangent to the curve i at the corresponding points. Then, the necessary condition for the two curves to be skew symmetric is

$$\kappa_1(s) \sin(\alpha_2(s))^3 = -\kappa_2(s + b) \sin(\alpha_2(s + b))^3 \quad (5)$$

An example is given in Fig. 5(a). The above conditions are only valid for curves and not for lines, that is, curvature should be nonzero. For lines, the first definition of skew symmetry can simply be applied as two lines are skew symmetric if another set of two lines that joins the end points of the given lines are parallel to each other. In this case, the two new lines are the lines of symmetry, and the line joining the end points of these lines is the axis of symmetry. An example of skew symmetry for straight lines is given in Fig. 5(b).

We believe that these two symmetries are major sources of information for extracting shape from contour. We discuss this process next.

B. Qualitative Shape Inferences from Symmetries

We now describe some *qualitative* inferences about the shape of surfaces from their symmetries. We also prove some

of the inferences that we make. We use the assumption of general viewpoint defined as the following:

Definition 1—General Viewpoint: A scene is said to be imaged from a general viewpoint if perceptual properties of the image are preserved under slight variations of the viewing direction.

Specifically, the perceptual properties we are interested in are straight lines, parallelism of lines, and the symmetry of curves.

Here, we discuss the interpretation of *individual* surfaces independently. In an object, of course, several surfaces may be visible, and their interpretations must be mutually consistent. This can provide a mechanism for either reinforcing individual surface interpretations or choosing among possible multiple interpretations for an individual surface.

It will be useful to consider figures as belonging to one of the following three classes:

1) *Case I:* Here, *one* skew symmetry covers the entire boundary of the surface. We allow more than one *alternative* description for a figure (Fig. 6 shows two examples). For example, the ellipse in Fig. 6 can be described as being skew symmetric about any axis that passes through its center, and all symmetries include all the points on the ellipse boundary.

Surfaces belonging to case I are generally perceived to be planar. We prove that if a contour belongs to case I bounded by nonlimb edges, then the contour has to be planar under the assumption of general viewpoint and if the correspondence is static with respect to changing viewpoint. Limb edges (or limbs) of a surface are generated by points on the surface whose normal is orthogonal to the viewing direction. Such an edge changes its position on the surface as the viewpoint changes. Nonlimb edges, on the other hand, do not change their position on the surface as the viewpoint moves; they include creases and wireframes.

Lemma 1: A 3-D skew symmetric figure projects as a skew symmetric figure under orthographic projection.

Proof: It is a direct result of the property of the orthographic projection that parallel lines project as parallel lines and that midpoints of lines project as midpoints of the projected lines. Therefore, the 3-D lines of symmetry project as the lines of symmetry on the image plane, and the projection of the 3-D axis is the line joining the midpoints of the lines of symmetry on the image plane.

Theorem 1: If a 3-D contour formed by nonlimb edges produces a skew symmetric line drawing in the image plane such that the 3-D correspondence is invariant² under small perturbations of the viewpoint, then the 3-D contour must be planar (under the assumption of general viewpoint).

Proof: Since the 3-D correspondence is invariant with respect to variations of the viewpoint (that is, the projection of the same set of 3-D points correspond from different viewpoints) then, the assumption of general viewpoint implies that parallel lines in the image plane must be the projection of parallel 3-D lines; otherwise, they would not project parallel from nearby viewpoints. Therefore, we conclude that the 3-D

²We thank Dr. V. Nalva for pointing out this assumption, which was omitted in an earlier version of the paper.




Fig. 6. Some examples of case I.

lines, say l_i , that project as the lines of skew symmetry on the image plane must be parallel to each other in 3-D because lines of skew symmetry are parallel to each other in the image plane. The axis of symmetry in 3-D, which can be obtained by joining the midpoints of the 3-D lines l_i , must be straight because its projection on the image plane, which is the axis of skew symmetry, is straight. Therefore, the lines l_i have to lie on a plane because they are parallel to each other, and a single line, which is the 3-D axis of symmetry, intersects them. Hence, the 3-D contour, which encloses the lines l_i , is planar. \square

Lemma 1 shows that planar skew symmetric figures project as skew symmetric figures on the image plane. Theorem 1 shows the reverse is also true under the stated conditions. We conjecture that invariance of the correspondence is not necessary but have not proved it.

Note that if the 3-D contour forming the skew symmetry on the image plane is a limb edge, then the 3-D contour could be nonplanar. For example, limb edges on surfaces of revolution produce an orthogonal skew symmetry [15]. Generally, such surfaces also produce a parallel symmetry together with skew symmetry, and they belong to case II, which is defined below.

Case II: Here, the boundary of the figure is covered by exactly *two* symmetries, and furthermore, at least one of them must be a parallel symmetry. We will argue that the case II figures are the ones that give us the most information about the surface shape and that such cases are common in everyday scenes. Fig. 7 shows some examples of this case. We will show that if one of the two symmetries is a skew symmetry with straight curves of symmetry (as in Fig. 7(a) and (b)), then the surface must be a zero Gaussian curvature surface. Otherwise, if the curves of skew symmetry are not straight or have two curved parallel symmetries, then we perceive a doubly curved surface (i.e., both of the principal curvatures are nonzero). Surfaces of revolution are such cases where the limb boundaries project as orthogonal skew symmetry as shown in Fig. 7(c). A double bent paper-like surface, which is shown in Fig. 7(d), illustrates this case with two parallel symmetries.

1) *ZGC Surfaces:* A zero Gaussian curvature (ZGC) surface is one where the the Gaussian curvature (the product of the maximum and minimum curvatures) of the surface is zero everywhere. Cylinders and cones are examples. These surfaces are also called developable surfaces since they can be generated from a piece of paper by rolling and/or bending

without cutting. Lines of minimum curvature for a ZGC surface, which are also called rulings, are straight, i.e., it is possible to embed straight lines on a ZGC surface along these rulings.

The following theorem asserts that case II figures satisfying specific properties must have ZGC along its skew symmetry contours.

Theorem 2: If a surface generates one parallel symmetry and one skew symmetry, with straight curves of skew symmetry on the image plane, and the straight curves of skew symmetry are also the lines of symmetry for the parallel symmetry, then the Gaussian curvature of the surface must be zero along the curves of skew symmetry.

Proof: There are two subcases, depending on whether the curves of skew symmetry are limb edges or not.

- a) The straight curves of skew symmetry are produced by limb edges: In this case, just the straightness of the limb is sufficient for the surface to have ZGC along the limb. This can be inferred as a special case of theorem given by Koenderink [10]. We give an alternative proof here that does not need the additional assumptions used in Koenderink's proof. Let the surface $X(u, v)$ be parameterized such that a u parameter curve is along the limb boundary for $v = v_o$. Since the curve $X(u, v_o)$ is along the limb boundary and the projection of this curve is straight, the surface normal \mathcal{N} along this curve is constant, that is, $\mathcal{N}_u(u, v_o) = 0$. This condition is a sufficient condition for the Gaussian curvature of the surface along the $X(u, v_o)$ to be zero. The Gaussian curvature κ of a surface is given by

$$\kappa = \frac{LN - M^2}{EG - F^2} \quad (6)$$

where L, M, N are the coefficients of the second fundamental form of the surface, and E, F, G are the coefficients of the first fundamental form. The equations of these coefficients are given in (30). Particularly, the coefficients L and M can be written as

$$L = -X_u \cdot \mathcal{N}_u \quad M = -X_v \cdot \mathcal{N}_u. \quad (7)$$

Since $\mathcal{N}_u(u, v_o) = 0$, the Gaussian curvature κ must be zero along the limb.

The above proof does not require the assumption of general viewpoint; hence, it only shows that along the curve of the limb boundary, the surface has ZGC. With the assumption of general viewpoint, we conclude that an open region surrounding the limb boundary also has ZGC.

- b) The straight curves of skew symmetry are cut edges. Consider Fig. 8, where a pair of parallel symmetric curves on a ZGC surface cut along a *ruling* is shown. Since, in the image plane, the tangents t_1 and t_2 of the top and bottom curves are parallel; by the assumption of general viewpoint, they must be parallel in 3-D. In addition, since the skew symmetry curves (one of which is the ruling in Fig. 8) are straight on the image plane, the 3-D corresponding curves must also be straight, that is,




Fig. 7. Examples of case II surfaces: (a) and (b) are ZGC surfaces and (c) and (d) are doubly curved surfaces.

the surface embeds straight lines. Therefore, the surface can locally be represented as a ruled surface having

$$X(u, v) = f(v) + ug(v) \quad (8)$$

where $f(v)$ and $g(v)$ are arbitrary vector functions of the parameter v only. The vector function $g(v)$ indicates the direction of the ruling that are also the u parameter curves. The normal of this surface is

$$\mathcal{N}(u, v) = \frac{(f'(v) + ug'(v)) \times g(v)}{|(f'(v) + ug'(v)) \times g(v)|}. \quad (9)$$

For Fig. 8, let the dotted line (ruling) be the cut boundary for $v = v_o$, and let $\mathcal{N}_1(u_1, v_o)$ and $\mathcal{N}_2(u_2, v_o)$ be the normals of the surface at points where the ruling intersects the parallel symmetry curves. Since the tangents t_1 and t_2 are the same and, of course, the tangent of the ruling is constant along it, the surface normals \mathcal{N}_1 and \mathcal{N}_2 , which are the cross products of t_1 and t_2 with the tangent of the ruling, must be the same, that is

$$\mathcal{N}(u_1, v_o) = \mathcal{N}(u_2, v_o) \\ (f'(v_o) + u_1g'(v_o)) \times g(v_o) \equiv (f'(v_o) + u_2g'(v_o)) \times g(v_o)$$

where \equiv indicates parallelism of the vectors. Then, we have that either $u_1 = u_2$ or the three vectors f', g', g are dependent. Clearly, $u_1 \neq u_2$; therefore, f', g', g are dependent, and hence, the surface normal \mathcal{N} is independent of the u parameter curve, that is, $\mathcal{N}_u(u, v_o) = 0$. As in the case (a) above, this condition is sufficient so that the Gaussian curvature of the surface along the $X(u, v_o)$ curve is zero. \square

- 2) *Generalization:* If we assume that the type of the surface does not change as we go from its boundary to the inside, then we conclude that the whole surface must be a ZGC surface if it satisfies the property given in theorem 6.

This generalization may appear to be a rather sweeping one. However, it is no more so than the common assumption that a polygonal line drawing corresponds to polyhedral objects.

It follows that if the parallel symmetry has a linear correspondence function, then the surface is conic, and if the correspondence function is an identity, then the surface is cylindrical. We now show how we can infer the *rulings* and the *cross sections* of the ZGC surface. Rulings are the lines along which the curvature of the surface is zero. Cross sections are the transverse (not necessarily orthogonal) curves, specifically, the curves that project into parallel symmetric curves. We first give a theorem that is key to inferring properties of cross sections and rulings.




Fig. 8. ZGC surface cut along the “ruling.”




Fig. 9. Objects with cross sections having (a) only one skew symmetry and (b) two skew symmetries

Theorem 3: Curves obtained by intersecting a ZGC surface with two parallel planes are parallel symmetric such that the lines of symmetry are the rulings of the surface.

The proof of this theorem is given in Appendix A-1. Note that the reverse of this theorem, that parallel symmetry curves *must* come from parallel planar cuts, is not valid. In Appendix A-2, we show that lines of maximum curvature that are not necessarily planar can also be projected as parallel symmetric curves. However, we believe that it is reasonable to infer that parallel symmetry curves are planar, unless we have evidence to the contrary. In general, lines of curvature of a ZGC can be very complex, and it is unlikely that an observed surface would be cut in this way. If the curves are neither planar nor along the curves of maximum curvature, it is quite difficult to obtain parallel symmetry. For example, in order to obtain parallel symmetry for a conic surface (a subcase of ZGC surfaces) by cutting with nonplanar cross sections, the cuts must be translated along the axis of the cone and scaled exactly with the scaling function of the cone.

Our interpretation does allow for piecewise planar cross sections, as indicated by multiple skew symmetries. Fig. 9 shows an example. The cross section of the object in Fig. 9(a) has a single skew symmetry and is perceived as planar, whereas the cross section of the object in Fig. 9(b) has two skew symmetries, and the perception is that the cross section has two planar parts, that is, if the cross section has multiple skew symmetries, then it will be piecewise planar such that each planar section has one skew symmetry.

3) *Recovering Rulings:* We can infer the rulings of the surface by joining the corresponding points on the two curves forming the parallel symmetry by straight lines, as shown in Fig. 4(c) (the corresponding points on the two curves have the same tangent). Note that the orientation of a ZGC surface does




Fig. 10. (a) Figure with two skew symmetries; (b) addition of an extra curve clarifies the perceived shape

not change along a ruling (this is also proved as a byproduct of the above proofs in the Appendix). Therefore, if we find the orientation of the surface at a single point on a ruling, we can extend it along the ruling. A quantitative analysis for ZGC surfaces is presented in Section IV.

Case III: This class includes all remaining cases. Three interesting subclasses occur here.

- The contours satisfy specific properties of some special objects (e.g., straight homogeneous generalized cylinders [17], [23]), in which case, the analysis can use these special properties.
 - We hypothesize the presence of some boundaries not present in the image to convert the figure into case I, II, or III(a) above. For example, consider Fig. 10(a) with two skew symmetries. Note that we have no strong feel for the 3-D shape of this surface. However, if we assume that there is one missing boundary that would introduce a parallel symmetry (and an additional skew symmetry) as shown in Fig. 10(b), the surface shape becomes very distinct. Of course, more than one such construction may be possible, where each gives an alternative interpretation; some constructions may be preferable according to some *heuristic* criteria.
- Another interesting case is where two symmetries cover most of the boundary but not all of it; an example is given in Fig. 11(a). Here, two choices are available. Either we can inscribe a smaller figure inside the larger one, or we can extrapolate some of the boundaries to meet the requirements of case II. The two choices are shown in Fig. 11(b) and (c). Note that the extrapolation is preferred if the “top” surface is also shown, as in Fig. 11(d).
- All other cases. Such figures are out of the scope of this paper. We also believe that such shapes are difficult for humans to perceive.

IV. QUANTITATIVE SHAPE RECOVERY

We now describe our technique of quantitative shape recovery. We will focus on ZGC surfaces, although some of our analysis applies to more general cases. Remember that presence of ZGC surfaces is indicated by observing the properties given in Theorem 2.

Our method is dependent on the use of three constraints:

- Curved shared boundary constraint* (CSBC): This constraint relates the orientation of the points on two inter-




Fig. 11. (a) Face of a cylinder with a clipped corner; (b) parallel symmetry cover only part of the surface; (c) top curve is extended for the parallel symmetry to cover the whole face; (d) top of the surface is also included

secting surfaces. A simpler constraint is obtained if we can assume the planarity of the curve produced by the intersection of two surfaces.

- **Inner surface constraint (ISC):** This constraint restricts the orientation of the neighboring points on a surface.
- **Orthogonality constraint (OC):** This constraint uses the assumption that the lines of parallel symmetry and the axis of parallel symmetry (i.e., the rulings and the cross sections) are orthogonal to each other in 3-D.

We first describe these constraints in detail and then discuss the combination of these constraints for different classes of ZGC surfaces. These constraints, in general, are not sufficient (even though a sufficient number of equations are obtained) to give unique surface orientations; they typically leave one degree of freedom unconstrained. To fix this degree of freedom, we use the shape of the parallel symmetry curves (or the cross sections).

A. Constraints

We now give some constraints that derive from observations of the symmetries and other boundaries in the image. We formulate three constraints discussed in subsections below and then discuss how to combine them.

1) **Curved Shared Boundary Constraint (CSBC):** This constraint relates the orientations of the two surfaces on opposite sides of an edge. The planar version has been used since the early days of polyhedral scene analysis [12]. Shafer *et al.* [19] extended it to the case of intersection of curved surfaces.

Consider two surfaces $X_1(u, v)$ and $X_2(u, v)$ meeting at a curve $\Gamma(s) = (x(s), y(s), z(s))$ as in Fig. 12. Let $N_1(u, v)$ and $N_2(u, v)$ be the normals of X_1 and X_2 , respectively. Along the curve $\Gamma(s)$, we can represent the normals N_1 and N_2 as $N_i(s) = N_i(u_i(s), v_i(s))$. Let the normals $N_i(s)$ be represented in p - q space as $N_i(s) = (p_i(s), q_i(s), 1)$. Then, the curved shared boundary constraint (CSBC) states that along the curve $\Gamma(s)$, the orientation of the surfaces X_1 and X_2 are constrained by the tangent $(x'(s), y'(s))$ of the image of the curve $\Gamma(s)$ by the following equation:

$$(x'(s), y'(s), z'(s))$$




Fig. 12. Two curved surfaces meeting at a curve Γ .

$$\begin{aligned} & \cdot ((p_2(s), q_2(s), 1) - (p_1(s), q_1(s), 1)) \\ & = 0x'(s)(p_2(s) - p_1(s)) \\ & + y'(s)(q_2(s) - q_1(s)) = 0. \end{aligned} \quad (11)$$

Proof of this constraint is omitted but follows from results given in [19]. A stronger constraint can be obtained if we can assume that the intersection curve Γ is planar. Say Γ lies in a plane with orientation (p_c, q_c) . With the assumption of planarity, the constraint equation becomes

$$x'(s)(p_c - p(s)) + y'(s)(q_c - q(s)) = 0. \quad (12)$$

For ZGC surfaces, we will assume that the parallel symmetric curves (or cross sections) are planar (based on an argument given in Section III-B-2).

2) **Inner Surface Constraint (ISC):** The inner surface constraint restricts the relative orientations of the neighboring points within a surface. For ZGC surfaces, the image of the rulings of the surface is used to constrain the surface orientation of neighboring points.

Let $X(u, v) = (x(u, v), y(u, v), z(u, v))$ be a (u, v) parametric representation of the surface X , and let v be along the direction of minimum curvature (rulings for ZGC surfaces). We can form an orientation function in terms of the parameters u and v : $O(u, v) = (p(u, v), q(u, v))$. The inner surface constraint (ISC) states that for a constant value of the parameter v , say v_0 , as the parameter u changes the direction of the function O in the p - q plane, $O_u = (p_u, q_u)$ should be orthogonal to the direction of the tangent of the rulings, that is, the lines of symmetry (x_v, y_v) under orthographic projection, that is

$$\begin{aligned} (p_u, q_u) \cdot (x_v, y_v) &= 0 \\ p_u x_v + q_u y_v &= 0. \end{aligned} \quad (13)$$

$$\boxed{\frac{q_u}{p_u} = -\frac{x_v}{y_v}}$$

The proof of this property is given in Appendix A-3. Geometrically, the ISC can be described as follows: As we move along the axis of parallel symmetry (the u parameter curve), the surface orientation should move in the p - q plane in a direction orthogonal to the image of the rulings (the lines of parallel symmetry). For example, for cylindrical surfaces, this ISC curve is a straight line since all rulings are parallel




Fig. 13. Inner surface constraint.

to each other. Note that this constraint does not require any regularity assumptions about the contour.

The above equation expresses the inner surface constraint in a continuous domain. In the discrete domain, suppose the constraint equation relates the change in orientation along the surface orientation is to be computed at n points for a ZGC surface (these n points are along the axis of the parallel symmetry since the surface orientation for a ZGC does not change along the rulings). We have $2n$ unknowns (p_i, q_i) for n points. This constraint provides us with $n - 1$ constraint equations as shown below.

Let the image of the ruling r_i between the i th and $(i + 1)$ st points make an angle γ_i with the horizontal, as in Fig. 13. The constraint equation relates the change in orientation along the axis of symmetry (p_u, q_u) to the tangent of the ruling (x_v, y_v) . Here, the tangent of the ruling is $(x_v, y_v) = (\cos(\gamma_i), \sin(\gamma_i))$, and the derivatives (p_u, q_u) can be approximated by the first-order difference as $(p_u, q_u) = (p_{i+1} - p_i, q_{i+1} - q_i)$. Substituting these in (13) gives

$$(p_{i+1} - p_i) \cos(\gamma_i) + (q_{i+1} - q_i) \sin(\gamma_i) = 0. \quad (14)$$

3) *Combination of ISC and CSBC*: In the discrete domain, we need to quantize $(p(s), q(s))$ as (p_i, q_i) and estimate $(x'(s), y'(s))$ from the image of $\Gamma(s)$, which is $(x(s), y(s))$ under orthographic projection. If the ZGC surface is to be described at n points, then there are $2n + 2$ unknowns, $2n$ for the surface orientations (p_i, q_i) and 2 for the cross section plane (p_c, q_c) . This constraint provides us with n constraint equations. By using the CSBC in conjunction with the ISC, we get $2n - 1$ equations. This leaves us with three degrees of freedom for describing a ZGC surface totally.

The two constraints are shown graphically in Fig. 14. A ZGC surface (a frustum) is shown in (a) with rulings and the axis of the symmetry marked on the surface. The ISC curve is shown on the p - q plane. Here, the section of the ISC curve from the point (p_i, q_i) to (p_{i+1}, q_{i+1}) is orthogonal to the ruling r_i . The straight lines on the p - q plane are the CSBC's such that at each point i , the tangent of the axis of symmetry (the dotted curve on the surface) is orthogonal to the corresponding CSBC line on the p - q plane. Three parameters required to fix all the orientations (p_i, q_i) are the orientation of the plane containing the intersection curve (p_c, q_c) and the quantity shown as d in Fig. 14, which we call the *angle parameter*. The angle parameter can be described as distance of the ISC curve from the point (p_c, q_c) , which corresponds to an angle in 3-D. Specifying the length of one of the CSBC lines is enough to fix the angle parameter d .




 Fig. 14. Three degrees of freedom present p_c, q_c, d in a ZGC surface after applying the constraints ISC and CSBC.


Fig. 15. Two cylinders (a) cut along the curves of maximal curvature; (b) cut in an arbitrary direction while preserving parallel symmetry. Now, we have the perception of an elliptical cylinder.

4) *Orthogonality Constraint (OC)*: We will assume orthogonality between the axis of parallel symmetry and the lines of parallel symmetry. This is equivalent to slicing the surface along rulings to obtain thin skew symmetric planar strips and assuming that these strips are orthogonally symmetric in 3-D, as in Kanade's analysis for polyhedra [7]. This preference is illustrated in Fig. 15, where in (a), we see a circular cylinder, but in (b), we see an orthogonal elliptical cylinder rather than a slanted cylinder. Note that for ZGC surfaces, where the lines of symmetry are the rulings of the surface (or the lines of least curvature), the orthogonality constraint implies that the cross sections must be along lines of maximum curvature. This is, in general, in conflict with our drive to perceive the cross section as being planar (except for cylinders and cones). We discuss how to resolve these conflicting constraints in Section IV-B.

For a ZGC surface, say, the tangent of the axis of symmetry, makes an angle α with the horizontal, and the ruling makes an angle β at some point on the surface, as in Fig. 16. Let the normal of the surface be $N = (p, q, 1)$ at that point. Since the 3-D tangent vectors A and B are on the tangent plane of the surface, they can be represented as

$$\begin{aligned} A &= (\cos(\alpha), \sin(\alpha), p \cos(\alpha) + q \sin(\alpha)) \\ B &= (\cos(\beta), \sin(\beta), p \cos(\beta) + q \sin(\beta)). \end{aligned} \quad (15)$$

and from the orthogonality of the 3-D vectors A and B , we get $A \cdot B = 0$ or

$$\cos(\alpha - \beta) + (p \cos \alpha + q \sin \alpha)(p \cos \beta + q \sin \beta) = 0. \quad (16)$$

This is the equation of a hyperbola in the p - q space, constraining possible orientations for the surface normal N . In discrete domain, we need to digitize α , β , p , and q above as $\alpha_i, \beta_i, p_i, q_i$ for each point on the axis symmetry. This constraint provides us with n equations if the surface orientation is to be computed at n points.




Fig. 16. Orthogonality constraint.

B. Combining the Constraints

The three different constraints of the previous sections provide $3n - 1$ constraint equations for n points producing $2n + 2$ unknowns (including (p_c, q_c)). This suggests that the system of equations is overconstrained (for $n > 3$). Thus, in general, it may not be possible to find an interpretation for the contours such that the surface obeys all the given constraints exactly. Furthermore, the orthogonality constraints and planarity of cross sections assumed by CSBC are usually in conflict, as discussed in Section IV-A-4. However, for special but important cases, these sets of constraints are dependent and may give a unique answer or even leave one degree of freedom unconstrained.

1) *Cylindrical Surfaces*: A cylindrical surface is a ZGC surface for which rulings are parallel to each other in 3-D. An example is given in Fig. 17(a). Let this surface be parameterized by $X(u, v) = (x(u, v), y(u, v), z(u, v))$ such that u is along the axis of symmetry, and v is along the rulings. As we move along the axis of symmetry, let the angle between the tangent of the axis of symmetry and the horizontal be $\alpha(u)$. Note that α is a function of u only, and let the angle between the ruling and the horizontal be β , as in Fig. 17. Note that, since all rulings are parallel, β is constant. We can always rotate the coordinate system to make β equal to $\Pi/2$, as in the figure. With these angles, we have

$$\begin{aligned} X_u &= (x_u, y_u, z_u) = (\cos \alpha, \sin \alpha, z_u) \\ X_v &= (x_v, y_v, z_v) = (\cos \beta, \sin \beta, z_v) = (0, 1, z_v). \end{aligned} \quad (17)$$

Our purpose is to compute the surface orientation $(p(u), q(u))$ along the axis of symmetry. Applying the inner surface constraint in (13) gives

$$p_u x_v + q_u y_v = 0 \Rightarrow q(u) = q_o(\text{constant}) \quad (18)$$

that is, the ISC curve is a horizontal line on the p - q plane as shown by dotted line in Fig. 17(b).

Say that the orientation of the cross-section plane is (p_c, q_c) . Then, the curved-shared boundary constraint gives

$$\begin{aligned} x_u(p_c - p(u)) + y_u(q_c - q_o) &= 0 \\ \cos(\alpha(u))(p_c - p(u)) + \sin(\alpha(u))(q_c - q_o) &= 0 \\ p(u) &= \frac{\sin(\alpha(u))(q_c - q_o)}{\cos(\alpha(u))} + p_c. \end{aligned} \quad (19)$$

that is, if we fix p_c, q_c and q_o , then the surface's orientation $(p(u), q_o)$ for all values of u is fixed. The last constraint is




Fig. 17. (a) Cylindrical surface with axis of symmetry and the rulings marked; (b) constraints ISC, CSBC, and the orthogonality for the cylindrical surface.

given by orthogonality, as given in (16). Since $\beta = \pi/2$, we have

$$\sin(\alpha(u)) + q_o p(u) \cos(\alpha(u)) + q_o^2 \sin(\alpha(u)) = 0. \quad (20)$$

Substituting $p(u)$, by (19), in the above equation gives

$$\sin(\alpha(u))(1 + q_o q_c) + p_c \cos(\alpha(u)) = 0. \quad (21)$$

Since the above equation is equal to zero for all values of u , then we have both $p_c = 0$ and

$$1 + q_o q_c = 0 \Rightarrow q_o = -1/q_c. \quad (22)$$

With the orthogonality constraint, we have $p_c = 0$ and $q_o = -1/q_c$, leaving q_c as a variable, that is, the three constraints CSBC, ISC, and OC are satisfied for a cylindrical surface, and still one degree of freedom, namely q_c , remains. In Section IV-C, we describe a method to estimate q_c . The method uses the shape of the parallel symmetry curves.

2) *Circular Cones*: A circular cone is a linear straight homogeneous generalized cylinder (LSHGC), whose cross section is a circle. The importance of circular cones is that these are the only ZGC surface that have a unique (two including the Necker's reversal) solution to the three constraints (ISC, CSBC, and OC) given before. In [22], we have analyzed the image of a cone under these constraints, and a unique solution, which is also in agreement with the assumption that the ellipse of the cross section in the image plane is the projection of a circle in 3-D, is found.

3) *General ZGC Surfaces*: For surfaces other than cylindrical surfaces and the circular cone, the three constraints cannot be satisfied exactly. We believe that in most cases, the planarity assumption is stronger than the orthogonality assumption. Therefore, the following process tries to maximize the orthogonality while keeping the constraints ISC and CSBC satisfied exactly.

As discussed in Section IV-A-3, there are three degrees of freedom left for reconstructing a ZGC surface. The free variables are (p_c, q_c) and d . We choose the values for these free variables that minimize the orthogonality error

$$\Xi = \sum_{i=1}^n \cos \theta_i \quad (23)$$

where θ_i is the angle between the two 3-D vectors (A and B) in Fig. 16, whose projection on the image plane make angles α_i and β_i with the horizontal. $\cos \theta_i$ is given by

$$\frac{(\cos(\alpha_i - \beta_i) + (p_i \cos \alpha_i + q_i \sin \alpha_i)(p_i \cos \beta_i + q_i \sin \beta_i))^2}{(1 + (p_i \cos \alpha_i + q_i \sin \alpha_i)^2)(1 + (p_i \cos \beta_i + q_i \sin \beta_i)^2)}.$$

Here, (p_i, q_i) are dependent on (p_c, q_c) and d as given by constraints ISC and CSBC. We want to maximize the orthogonality by minimizing the above function Ξ for (p_c, q_c) and d . We can convert this problem into a 2-D minimization problem by associating a d value to each choice of (p_c, q_c) that minimizes Ξ .

Unfortunately, for a general conic surface, the global minimum for Ξ occurs when $(p_c, q_c) = (0, 0)$ and $d = \infty$. This is an infeasible interpretation. However, function Ξ , in terms of (p_c, q_c) , has a ‘‘valley’’ of local minima (passing through the origin of the p - q space), and the valley is typically a straight line. Any choice of (p_c, q_c) along this valley is essentially equally acceptable, i.e., we have one degree of freedom to fix. In Section IV-C, we discuss how to choose a specific value of (p_c, q_c) on this line using the shape of the cross section.

C. Estimating (p_c, q_c)

As discussed in Section IV-B, the previous three constraints (ISC, CSBC, OC) leave one degree of freedom such that the orientation of the cross section plane (p_c, q_c) is constrained to be along the minimum line of the function Ξ . It is expensive to compute this minimum line. Instead, we use the following gradient descent algorithm to compute (p_c, q_c) :

- 1) Choose a starting line l_o passing through the origin in the p - q space in the direction of the skew symmetry axis. Set the current line to $l = l_o$.
- 2) Compute the (p_c, q_c) for the line l using the method described below.
- 3) Compute the value of Ξ for (p_c, q_c) , and check if (p_c, q_c) is along the minimum line of Ξ by repeating the above process for lines $\pm \delta \theta$ degrees off the line l , and by comparing the Ξ values for these lines.
- 4) If (p_c, q_c) is along the minimum line of Ξ , then stop. Otherwise, choose another line by rotating the line l by $\delta \theta$ degrees in the direction of descending Ξ , and go to step 2.

1) *Computing (p_c, q_c) Given a Line l :* We rotate the coordinate system such that the line l is aligned with the q axis of the p - q plane. Then, we have $p_c = 0$, and q_c is the unknown quantity.

To fix q_c , we use the shape of the cross section. We propose a method based on perceptual properties rather than on mathematical constraints.

Our method is based on the following observations of human perception:

- We prefer compact shapes (as also observed in [3]).
- We prefer medium slant to very high or very low slant.
- We have a large range of uncertainty for the perceived slants.

Based on these observations, we propose a two-stage method for determining q_c . First, we estimate a value for q_c , and then,




Fig. 18. (a) Cylindrical object and the ellipse fitted to the cross section; (b) orientation (p_e, q_e) that would make the ellipse a circle and its projection on the q axis gives q_e , which is first approximation to q_c .

we update it with a bias toward 45° . For the first estimation, an ellipse is fit to the cross section and then backprojected to an orientation that makes it a circle (apart from being much faster, this has an advantage over the $(area)/(perimeter)^2$ measure used by Brady and Yuille [3] in that it does not require that closed contours be given). Since this is a rather heuristic approach, we have performed a psychological study validating this approach. The details of this study are provided in [21], and a summary of the results is given at the end of this subsection. The two steps are described in detail below.

2) *First Estimation of q_c :* An ellipse-fitting process is utilized as a first approximation for q_c . An ellipse is fit to the cross section contour, and then, the orientation of the circle (p_e, q_e) that would be projected as the fitted ellipse is projected on the q axis on the p - q plane to obtain the first approximation of q_c ; call it q_e . Fig. 18 shows an example. Note that there are two values of (p_e, q_e) that make a circle project as the ellipse in the image plane. These correspond to Necker’s reversal. We choose the one that gives a solid shape interpretation in preference to the one that gives the interpretation of a hole.

The behavior of the method is dependent on the choice of the ellipse fitting algorithm used. We have experimented with two different ellipse fitting algorithms. The first one is based on the scattering of the boundary points. A covariance matrix of the equally spaced contour points is computed by

$$C = \begin{bmatrix} \frac{1}{n} \sum (x_i - \bar{x})^2 & \frac{1}{n} \sum (x_i - \bar{x})(y_i - \bar{y}) \\ \frac{1}{n} \sum (x_i - \bar{x})(y_i - \bar{y}) & \frac{1}{n} \sum (y_i - \bar{y})^2 \end{bmatrix} \quad (24)$$

where (x_i, y_i) are the equally spaced boundary points, and (\bar{x}, \bar{y}) is the mean. The scattering of these contour points is given by the eigenvalues e_1 and e_2 of the covariance matrix C . Say that the unit vectors v_1 and v_2 are eigenvectors of C corresponding to e_1 and e_2 , respectively. Then, we can approximate the cluster of points (x_i, y_i) with an ellipse whose major and minor axes are in the directions v_1 and v_2 with magnitudes $\sqrt{e_1/2}$ and $\sqrt{e_2/2}$. This method is quite robust when the contour is closed. However, for open contours, the method consistently underestimates the eccentricity of the ellipse.

The second method is a regular least squares fit of the parameters of the quadratic representation of the ellipse to the boundary points. This method is robust when the contour is similar to an ellipse whether it is closed or not but may




Fig. 19. Objects and ellipses fit for their cross sections. The cross sections of the objects are segmented based on their concavities (or inflection points), and the whole cross section has the same slant.

give a bad fit if the contour is not similar to an ellipse. We apply both methods to a contour and choose the one having the smaller fit error (the ε in (26)).

If the cross section has repetitive parts, as in Fig. 19, then the slant perception is governed by the shape of individual parts rather than the overall figure. For closed cross sections, we segment them by finding the concavities on the two sides and then matching them. For open cross sections (as in Fig. 19(c)), we simply segment it at inflection points. An ellipse is fit to each part, and corresponding q_e 's are computed for each part. The q_e for the whole cross section is given by a weighted average of the q_e values, where the weight is given by the length of the curve to which the ellipse was fit. Note that this is different from the segmentation of the cross section described in Section III-B-2, which is based on the broken skew symmetry axis and results in a different slant computation for each segmented part as in Fig. 9.

Fig. 19 shows the ellipses fit to the cross sections of various objects. The objects in Fig. 19 are segmented by the above method, and an ellipse is fit to each part.

3) *Updating q_c* : The purpose of this updating process is to simulate the behavior that humans have in preferring medium slant to very high and very low slant. We update q_e to obtain the final q_c as follows (after converting q_e into degrees):

$$q_c = 45^\circ + \lambda(q_e - 45^\circ) \quad (25)$$

where λ is a confidence factor in the range $[0, 1]$ and is a function of how well the ellipse approximates the cross-section curve. Intuition suggests that the better the approximation of the ellipse, the higher the value of λ should be, and the closer the q_c is to the 45° , the less the correction should be. The λ we are using is

$$\lambda(\varepsilon) = (1 - \varepsilon^2) \quad (26)$$

where ε is the ellipse fit error given by \bar{d}/\sqrt{ab} , \bar{d} is the average distance of the contour points from the fitted ellipse, and a and b are the half lengths of the major and minor axes of the ellipse. We use an approximation (not described here) to compute \bar{d} . Note that ε is in the range $[0, 1]$.

We believe that the exact form of the function is not critical. Small changes in q_c do not radically affect the perceived surface shape, and humans estimate q_c rather imprecisely.

4) *Validation*: We have conducted a psychological experiment with human subjects on the perception of q_c for cylindrical and conic objects [21]. Results of the experiment show that the standard deviation of the perceived angle for the top plane is quite high, with an average standard deviation σ of 8° . The

interval of uncertainty for the slant of each object, which is the angle interval that contains the 90% of the responses given for that object, is 24° . The comparison shows that the algorithm performs quite well for a variety of shapes. The average of the differences between the mean of the human response, and the computed slants are only 6° (which is smaller than the average standard deviation of human responses).

D. Computational Results

The inputs to our program are the segmented curves that define the contour of each object. These segmented curves are grouped into closed regions using continuity. Each closed region is taken to correspond to an object surface. Next, we find symmetries among segments of a surface. Every segment in a surface is checked for parallel symmetry against every other segment in the surface. Two segments are considered to be parallel symmetric if they return a *low* parallel symmetry error, which is defined as

$$\frac{1}{l_2} \int_0^{l_2} \sqrt{(x_2'(s) - x_1'(\rho s))^2 + (y_2'(s) - y_1'(\rho s))^2} ds \quad (27)$$

where segment $C_1(s) = (x_1(s), y_1(s))$ and segment $C_2(s) = (x_2(s), y_2(s))$ are parametrized in terms of their arclength s , and $\rho = l_1/l_2$ is a scaling parameter where l_1 and l_2 are the lengths of the segments C_1 and C_2 .

The above error measure is effective only if the entire lengths of two segments are parallel symmetric to each other. In addition, this measure is limited to *linear* parallel symmetry (which is found in cylindrical and conic surfaces).

Segments are also checked for having the same curvature sign at the corresponding points. This measure is especially useful when the segments are almost straight, in which case, the error measure given in (27) may be low even if the segments are not parallel symmetric.

The surfaces containing a parallel symmetric segment pair are treated as curved, and others are treated as planar. For curved surfaces, the curves joining parallel symmetric curves are checked if they are straight, which confirms that the surface is a ZGC. The curved surfaces are associated with their planar neighbors, which will be treated as the cross section having the normal $(p_c, q_c, 1)$.

For each object, the orientation of the planar cross section (p_c, q_c) is computed using the method described in Section IV-C. Then, the angle parameter d is computed by minimizing the orthogonality error Ξ given in (23). The surface orientation (p_i, q_i) at each point is then computed by using constraints ISC and CSBC as illustrated in Fig. 14.

In the following, we show results on some synthetic examples as well as a real image. Evaluation of shape from contour results is difficult as there is no real "ground truth." Even when contours are derived from a real object or from projection of a synthetic object, the same boundaries *could* have been derived from a projection of infinitely many other real or synthetic objects. Thus, in a sense, the only good measure of the performance of our algorithms is a comparison with human performance. We use two graphical methods to display the computed orientations. The first one shows the surface




Fig. 20. Sample contours, the needle images computed, and their images after shading the object with the computed orientation at every point on the surface.

normals as oriented *needles* along one cross section. The orientations along the other cross sections are the same because the orientations are translated along rulings. The orientation and length of a needle is the projection of a unit surface normal at that point (for us to perceive 3-D orientation from this requires solving a shape from contour problem in itself). The second method is to display the surface orientations by constructing a synthetic shaded image from the reconstructed surface by assuming Lambertian reflection and a point source of light (for humans to perceive this requires solving a shape-from-shading problem). We believe the needle diagrams to be more effective than shaded images for this purpose. We present graphical results for the reader to make his/her own judgment. Unfortunately, this can only give qualitative rather than quantitative evaluation. For the real image, we also give a comparison with the real object.

1) *Synthetic Images:* Fig. 20 shows computed surface orientations from contours of objects in Fig. 1 using our method. The inputs to the algorithm are the curves (given as a list of points) defining the contour for each object. In our judgment, the reconstruction is consistent with human perception of the given figures. It is worth noting that for all the objects, the computed orientation at the limb boundaries of the objects is orthogonal to the boundary, even though this is not an explicit constraint in our method.

The cross section of the object in the bottom last row is segmented into two planar sections based on the observation of the skew symmetry of the cross section. Each section is processed individually, but the inner surface constraint is required to be applied between the two sections of the object.

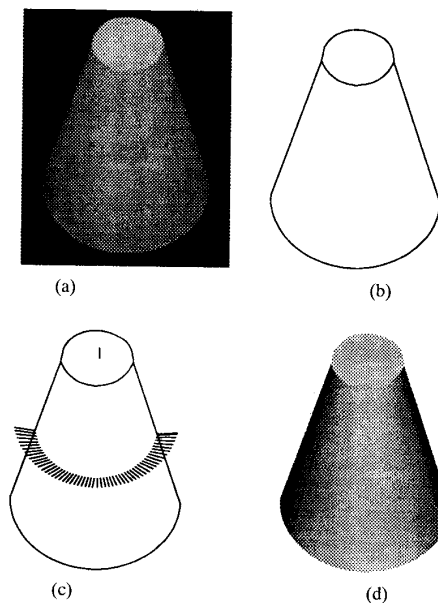


Fig. 21. Processing of a real image; the cone image, edges, computed surface normals, and the shaded image with the computed surface normals.

2) *Real Image Example:* To apply our method to real images, we need to first find the boundaries of the objects and then the symmetries, if any, contained in them. In general, we can expect object boundaries to be fragmented and several intensity boundaries that correspond to surface markings, shadows, and noise to be present. To separate the object

boundaries from these other boundaries, and to fill in the gaps in object boundaries as appropriate, is a difficult problem in monocular image analysis, but this paper is not about such analysis. Perceptual grouping has been suggested as one solution to such problems, and in our group, we have developed such techniques that we believe are part of the solution to these problems [13]. In addition, the fact that we are seeking certain specific relations between curves should help in the process of perceptual organization. We hope to explore these issues in further research. Here, we only show an example where boundaries can be detected cleanly, and no surface markings, shadows, or highlights are present. Nonetheless, we still deal with the noise in the location of the detected edges and the effects of this noise on computing tangents.

In the example to be shown, we first detect edges in the image using a Canny edge detector. Edges are then linked into curves. The curves are segmented into smaller curves by detecting corners using a multiscale version of the curvature-based corner finder described in [14]. Then, the segmented curves are given to the reconstruction system as before. The tangent of the axis of the parallel symmetric curves (which are necessary for the curved shared boundary constraint) is computed by convolving the larger of the symmetric curves with a first derivative of Gaussian kernel having a large standard deviation ($\sigma = 10.0$) to smooth out the noise.

In Fig. 21, we show the results on a real image. The image ($245 \times 300 \times 8$) is that of a circular cone. Fig. 21(b) shows edges, Fig. 21(c) shows the recovered surface normals as needles, and Fig. 21(d) shows a reconstructed image assuming Lambertian reflection and point source of light. We believe that the results agree well with human perception of the original image. The average error of the surface normal from the actual cone is about 5° (the average error in tangent estimates of the image curves is also about 5°). Most of this error is concentrated near the limb boundaries where the image tangent estimates also have higher error. Note that the image was obtained by perspective projection, but we have processed it as if it were obtained by an orthographic projection. For this example, the difference in the two projections apparently does not create a large error.

This example is intended to demonstrate that our reconstruction algorithm is robust enough to work with real image computations, at least in a controlled laboratory setting. We do not, however, claim to have solved other problems of monocular image analysis.

V. CONCLUSIONS

We have presented a theory of how to infer 3-D shape from contour of curved surfaces if certain symmetry properties are present. We have given a detailed analysis for ZGC surfaces, but the essentials of the theory also extend to more complex shapes [24]. We believe that ZGC surfaces form a large and useful class of objects and that studying them is an appropriate step up in complexity from planar surfaces, where much of the previous work in the field has been focused.

Our theory does make certain assumptions, as must all methods, to extract shape from contour, but we believe that our assumptions are minimal and that the results agree with human perception. We have argued that, in a certain sense, this is the only evaluation that can be made for shape-from-contour methods, as many shapes can produce the same contour.

The method we have presented only exploits interaction between a curved surface and a planar surface. For more complex objects, it may be beneficial (even necessary) to exploit the interaction between more than two surfaces.

We have presented results on several synthetic examples and one real image. To apply our method to complex real images, where surface markings, shadows, and highlights may be present, will require solution of other monocular image analysis problems. However, we believe that our approach can help provide a model for such analysis. We intend to pursue this in our future research.

VI. APPENDIX A PROOFS OF THEOREMS

In this section, we give three proofs; two are related to the existence of parallel symmetries on ZGC surfaces (Theorems 4 and 3) and the other proves the ISC. All of the proofs uses the following surface representation.

Let $X(u, v) = (x(u, v), y(u, v), z(u, v))$ be a (u, v) parametric representation of the class C^2 ZGC surface X . Let us assume that the v parameter curves are along the lines of minimum curvature (rulings) of the surface.

The normal \mathcal{N} of this surface at any point is given by

$$\mathcal{N} = \frac{X_u \times X_v}{|X_u \times X_v|} \quad (28)$$

where \times is the vector product operator, and $|V|$ is the length of the vector V . Note here that $|\mathcal{N}| = 1$. First I and second II fundamental forms of such a surface are given by

$$\begin{aligned} I(X_u du + X_v dv) &= E du^2 + 2F dudv + G dv^2 \\ II(X_u du + X_v dv) &= L du^2 + 2M dudv + N dv^2 \end{aligned} \quad (29)$$

where

$$\begin{aligned} E &= X_u \cdot X_u & F &= X_u \cdot X_v & G &= X_v \cdot X_v \\ L &= X_{uu} \cdot \mathcal{N} & M &= X_{uv} \cdot \mathcal{N} & N &= X_{vv} \cdot \mathcal{N} \end{aligned} \quad (30)$$

Since the parameter v is along the ruling (a line), the normal curvature of the surface in the direction X_v , given by $II(X_v)$, should be zero; then, we have

$$II(X_v) = N = 0. \quad (31)$$

Gaussian curvature κ of such a surface is given by [11]

$$\kappa = \frac{LN - M^2}{EG - F^2}. \quad (32)$$

Since the Gaussian curvature of the surface is zero, setting $\kappa = 0$, with substituting 0 for N by (31), gives

$$M = 0. \quad (33)$$

1. Proof of Theorem 3

Consider the surface X , as given above. In addition, assume that the u parameter curves on the surface X are planar and parallel to each other. We have to show that the tangent of the u parameter curves $\frac{X_u}{|X_u|}$ is constant with respect to v (i.e., $\frac{X_u}{|X_u|}$ is a function of u only).

Let the planes on which the u parameter curves lie have the normal \mathcal{P} (\mathcal{P} is constant). Then, we have

$$X_u \cdot \mathcal{P} = 0 \Rightarrow 0 = \frac{\partial(X_u \cdot \mathcal{P})}{\partial v} = X_{uv} \cdot \mathcal{P} + X_u \cdot \mathcal{P}_v = X_{uv} \cdot \mathcal{P} \quad (34)$$

that is, $X_u \perp \mathcal{P}$ and $X_{uv} \perp \mathcal{P}$ for all u and v . In addition, $X_u \perp \mathcal{N}$ by (28) and $X_{uv} \perp \mathcal{N}$ since $M = 0$; therefore, unless $\mathcal{N} // \mathcal{P}$, we have

$$X_{uv} = c_1 \mathcal{N} \times \mathcal{P} \quad \text{and} \quad X_u = c_2 \mathcal{N} \times \mathcal{P} \quad (35)$$

for some constants c_1 and c_2 , that is, $X_u // X_{uv}$, and the derivative of $\frac{X_u}{|X_u|}$ with respect to v is

$$\frac{\partial}{\partial v} \left(\frac{X_u}{|X_u|} \right) = \frac{X_{uv}|X_u| - X_u \frac{X_u \cdot X_{uv}}{|X_u|}}{|X_u|^2}. \quad (36)$$

Since $X_u // X_{uv}$, we can substitute X_{uv} by $\frac{|X_{uv}|}{|X_u|} X_u$ in the above equation:

$$\frac{\partial}{\partial v} \left(\frac{X_u}{|X_u|} \right) = X_u \left(\frac{|X_{uv}|}{|X_u|^2} - \frac{(X_u \cdot X_u)|X_{uv}|}{|X_u|^4} \right) = 0. \quad (37)$$

Therefore, the tangents of the u parameter curves are parallel to each other at the points at which they meet a particular ruling, resulting in u parameter curves projecting as parallel symmetric with the lines of symmetry corresponding to the rulings. \square

2. Curves of Maximum Curvature for ZGC Surfaces

In this section, we show that curves of maximum curvature also project as parallel symmetric curves such that the rulings corresponds to the lines of symmetry.

Theorem 4: Curves of maximum curvature on a ZGC surface project into parallel symmetric curves in the image plane. Furthermore, the lines of parallel symmetry are the projections of the rulings.

For this section, the term cross section is used to refer to the curves of maximum curvature on the surface. We have to show that the 3-D tangent of every cross section at the point that they intersect with a particular ruling is equal to each other.

Consider the surface X given previously. In addition, assume that the parameter u is along the cross sections of the surface. Then, we have

$$X_u \cdot X_v = 0 \quad (38)$$

for all u and v . By (28), we have $\mathcal{N} \perp X_v$; then

$$\begin{aligned} 0 &= \frac{\partial(X_v \cdot \mathcal{N})}{\partial v} \\ &= X_{vv} \cdot \mathcal{N} + X_v \cdot \mathcal{N}_v \\ &= N + X_v \cdot \mathcal{N}_v = X_v \cdot \mathcal{N}_v \end{aligned} \quad (39)$$

that is $X_v \perp \mathcal{N}_v$, and $X_u \perp \mathcal{N}_v$ since

$$\begin{aligned} 0 &= \frac{\partial(X_u \cdot \mathcal{N})}{\partial v} \\ &= X_{uv} \cdot \mathcal{N} + X_u \cdot \mathcal{N}_v = M + X_u \cdot \mathcal{N}_v \\ &= X_u \cdot \mathcal{N}_v. \end{aligned} \quad (40)$$

Now, we have $X_u \perp \mathcal{N}$, $X_v \perp \mathcal{N}$, $X_v \perp \mathcal{N}_v$ and $X_u \perp \mathcal{N}_v$ and $X_u \perp X_v$. Therefore, \mathcal{N}_v should be parallel to \mathcal{N} , but by definition, $d\mathcal{N} \perp \mathcal{N}$. Since any first-order derivative $d\mathcal{N} = \mathcal{N}_u du + \mathcal{N}_v dv$ of \mathcal{N} has to be orthogonal to \mathcal{N} , observe the following:

$$0 = d1 = d|\mathcal{N}| = d(\mathcal{N} \cdot \mathcal{N}) = 2d\mathcal{N} \cdot \mathcal{N}. \quad (41)$$

Hence, \mathcal{N}_v has to be equal to zero. That is that the surface normal \mathcal{N} is only a function of the parameter u . Since $\mathcal{N} = \frac{X_u \times X_v}{|X_u \times X_v|}$, $X_u \cdot X_v = 0$, and the tangent vector of v parameter curves $\frac{X_v}{|X_v|}$ is a function of u only (because v parameter curves are straight lines), then $\frac{X_u}{|X_u|}$ is a function of the parameter u only. $\frac{X_u}{|X_u|}$ is the tangent of the cross section, and it is constant for all values of v . Consequently, at the points of intersections of cross sections with a particular ruling, cross sections have the same 3-D tangents. Therefore, cross sections form parallel symmetry in the image plane under orthographic projection (i.e., a point (x, y, z) projects to the point (x, y)). \square

3. Proof of the Inner Surface Constraint

Here, we will prove the inner surface constraint asserted by (13).

Consider a ZGC surface with v parameter curves along the rulings, and u parameter curves are arbitrary. Here, we have $X_v \cdot \mathcal{N}_u = 0$ since

$$\begin{aligned} 0 &= \frac{\partial(X_v \cdot \mathcal{N})}{\partial u} \\ &= X_{uv} \cdot \mathcal{N} + X_v \cdot \mathcal{N}_u \\ &= M + X_v \cdot \mathcal{N}_u = X_v \cdot \mathcal{N}_u. \end{aligned} \quad (42)$$

We can write \mathcal{N} , in terms of the gradient (p, q) , as

$$\mathcal{N} = c(p, q, 1) \quad (43)$$

where c is the scale coefficient and is equal to $(p^2 + q^2 + 1)^{-1/2}$. Differentiation of \mathcal{N} with respect to the parameter u gives

$$\begin{aligned} \mathcal{N}_u &= c_u(p, q, 1) + c(p_u, q_u, 0) \\ &= \frac{c_u}{c} \mathcal{N} + c(p_u, q_u, 0). \end{aligned} \quad (44)$$

If we set $X_v \cdot \mathcal{N}_u = 0$ from (42), where $X_v = (x_v, y_v, z_v)$ and \mathcal{N}_u is given in (44), we get

$$X_v \cdot \mathcal{N}_u = \frac{c_u}{c} X_v \cdot \mathcal{N} + c(x_v, y_v, z_v) \cdot (p_u, q_u, 0) = 0. \quad (45)$$

We also have $\mathcal{N} \cdot X_v = 0$ from (28). Therefore,

$$\begin{aligned} x_v p_u + y_v q_u &= 0 \\ \frac{q_u}{p_u} &= \frac{x_v}{y_v}. \end{aligned} \quad (46)$$


Q.E.D.

ACKNOWLEDGMENT

We thank Dr. V. Nalwa for many insightful and valuable comments. We also thank the anonymous reviewers for many suggestions that have improved the quality of this paper.

REFERENCES

- [1] I. Beiderman, "Recognition by components: A theory of human image understanding," *Psycholog. Rev.*, vol. 94, no. 2, pp. 115-147, 1987.
- [2] H. G. Barrow and J. M. Tenenbaum, "Interpreting line drawings as three dimensional surfaces," *Artif. Intell.*, vol. 17, pp. 75-116, 1981.
- [3] M. Brady and A. Yuille, "An extremum principle for shape from contour," *IEEE Trans. Patt. Anal. Machine Intell.*, vol. PAMI-6, pp. 288-301, 1984.
- [4] M. B. Clowes, "On seeing things," *Artif. Intell.*, vol. 2, no. 1, pp. 79-116, 1971.
- [5] R. Horaud and M. Brady, "On the geometric interpretation of image contours," *Artif. Intell.*, vol. 37, pp. 333-353, 1988.
- [6] D. A. Huffman, "Impossible objects as nonsense sentences," *Machine Intell.*, vol. 6, pp. 295-323, 1971.
- [7] T. Kanade, "Recovery of the three-dimensional shape of an object from a single view," *Artif. Intell.*, vol. 17, pp. 409-460, 1981.
- [8] R. Kashipati, "Shape description from sparse and imperfect data," Ph.D. thesis, Univ. of Southern California, 1988.
- [9] T. Kanade and J. R. Kender, "Mapping image properties into shape constraints: Skew symmetry, affine transformable patterns, and the shape from texture paradigm," in *Human and Machine Vision* (J. Beck, B. Hope, and A. Rosenfeld, Eds.). New York: Academic, 1983, pp. 237-257.
- [10] J. J. Koenderink, "What does the occluding contour tell us about solid shape," *Perception*, vol. 13, pp. 321-330, 1984.
- [11] M. Lipschutz, *Differential Geometry*. New York: McGraw-Hill, 1969.
- [12] A. K. Mackworth, "Interpreting pictures of polyhedral scenes," *Artif. Intell.*, vol. 4, pp. 121-137, 1973.
- [13] R. Mohan and R. Nevatia, "Segmentation and description based on perceptual organization," in *Proc. Comput. Vision Patt. Recogn. Conf.* (San Diego), 1989, pp. 333-341, 1989.
- [14] G. Medioni and Y. Ysumoto, "Corner detection and curve representation using cubic b-splines," *Comput. Vision Graphics Image Processing*, vol. 39, pp. 267-278, 1987.
- [15] V. Nalwa, "Line-drawing interpretation: Bilateral symmetry," *IEEE Trans. Patt. Anal. Machine Intell.*, vol. 11, pp. 1117-1120, 1989.
- [16] R. Nevatia and T. O. Binford, "Description and recognition of complex-curved objects," *Artif. Intell.*, vol. 8, pp. 77-98, 1977.
- [17] J. Ponce, D. Chelberg, and W. B. Mann, "Invariant properties of straight homogeneous generalized cylinders and their contours," *IEEE Trans. Patt. Anal. Machine Intell.*, vol. 11, no. 9, pp. 951-966, 1989.
- [18] J. Ponce, "Ribbons, symmetries and skew symmetries, in *Proc. Image Understanding Workshop*, 1988, pp. 1074-1079.
- [19] S. A. Shafer, T. Kanade, and J. R. Kender, "Gradient space under orthography and perspective," *Comput. Vision Graphics Image Processing*, vol. 24, pp. 182-199, 1983.
- [20] K. A. Stevens, "The visual interpretations of surface contours," *Artif. Intell.*, vol. 17, pp. 47-73, 1981.
- [21] F. Ulupinar, "Perception of 3-D shape from 2-D images of contours," Rep. IRIS 278, Univ. of Southern California, 1991.
- [22] F. Ulupinar and R. Nevatia, "Using symmetries for analysis of shape from contour," in *Proc. 2nd ICCV*, 1988, pp. 414-426, 1988.
- [23] ———, "Perception of 3-D surfaces from 2-D contours," in *Proc. 10th Int Conf. Patt. Recogn.* (Atlantic City, NJ), 1990, pp. 147-154.
- [24] ———, "Shape from contour: Shgcs," in *Proc. 3rd ICCV*, 1990, pp. 582-586.
- [25] ———, "Constraint for interpretation of line drawings under perspective projection," *Comput. Vision Graphics Image Processing*, vol. 52, pp. 674-676, 1991.
- [26] I. Weiss, "3-D shape representation by contours," *Computer Vision Graphics Image Processing*, vol. 41, pp. 80-100, 1988.
- [27] G. Xu and S. Tsuji, "Inferring surfaces from boundaries," in *Proc. 1st ICCV* (London), 1987, pp. 716-720.



Fatih Ulupinar received the B.Sc. degree in computer engineering in 1985 from the Middle East Technical University, Ankara, Turkey, and the M.Sc. and Ph.D. degrees in computer engineering from the University of Southern California (USC), Los Angeles, in December 1986 and August 1991, respectively.

He has conducted research on various areas of computer vision at the Institute of Robotics and Intelligent Systems at USC since 1986. He has authored various papers in computer vision during his stay at this Institute. His research interests include computer vision, robotics, neural networks, artificial intelligence, and geometry as applied to computer vision. He is currently with the Faculty of Computer Engineering, Bilkent University, Ankara, Turkey.



Ramakant Nevatia (F'92) received the B.S. degree from the University of Bombay, India, and the M.S. and Ph.D. degrees from Stanford University, Stanford, CA, all in electrical engineering.

He has been with the University of Southern California (USC), Los Angeles, since 1975, where he is currently Professor of Computer Science and Electrical Engineering and Director of the Institute for Robotics and Intelligent Systems. He spent the academic year 1981-1982 at Stanford University as a Visiting Professor. He has authored two books: *Machine Perception* and *Computer Analysis of 3-D Curved Objects*. His research interests include computer vision, artificial intelligence, and robotics.

Dr. Nevatia is a member of the Association for Computing Machinery and the American Association for Artificial Intelligence. He is an Associate Editor of IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, *Pattern Recognition*, AND *Computer Vision, Graphics and Image Processing: Image Understanding* and was a Technical Editor in the areas of robot vision and inspection systems for the IEEE JOURNAL OF ROBOTICS AND AUTOMATION.