MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AI Working Paper 329 **November 1985**

Extending 2-D Smoothed Local Symmetries to 3-D

David J. Braunegg

Abstract: 3-D Smoothed Local Symmetries (3-D SLS's) are presented as a representation for three-dimensional shapes. 3-D SLS's make explicit the perceptually salient features of 3-D objects and are especially suited to representing man-made objects. The definition of the 3-D SLS is given as a natural extension of the 2-D Smoothed Local Symmetry (2-D SLS). Analytic descriptions of the 3-D SLS are derived for objects composed of planar and spherical patches. Results of an implementation of the 3-D SLS are presented, along with suggestions for further research.

1. Introduction

The best motivation for developing a shape representation method for three-dimensional objects is that objects in the real world are three-dimensional. Although our eyes receive two-dimensional images of the surrounding world, we interpret these images in terms of our three-dimensional environment. The knowledge that the world is inherently three-dimensional has motivated various researchers to attempt to recover 3-D information from an image in a variety of ways. Thus we have shape from stereo [Grimson], shape from shading [Horn 1975], shape from motion [Ullman], and various other methods of shape recovery. However clear it may be that we need three-dimensional information of the world around us, it is not clear what the best representation of that three-dimensional world might be. Presented herein is a method for describing three-dimensional shapes which has some perceptual saliency, supports reasoning about three-dimensional objects, and overcomes some of the deficiencies of previous attempts at shape representation.

Motivation

Our goal is to develop a representation scheme which describes manufactured objects in a natural and straightforward manner. The representation should make the salient features of an object explicit while suppressing fine detail. However, the representation should be hierarchical in nature, with the lower levels supplying more detail at the price of verbosity of description. This representation scheme is intended to support reasoning about real-world objects. Specifically, we wish to support the reasoning about mechanical assembly and disassembly operations for the Mechanic's Mate project outlined in [Brady, Agre, Braunegg, and Connell] and about individual objects, such as determining an object's purpose based on its shape [Braunegg]. The coarser descriptions will allow top-level reasoning about actions and interactions of objects, while the finely detailed descriptions will allow the reasoning system to fill in the details of a developing plan or purpose.

The work of [Connell] has already shown that such descriptions can be derived for two-dimensional objects. Connell's program generates symbolic descriptions of the same type employed by ANALOGY [Winston, Binford, Lowry, Katz] that were formerly supplied by the Katz parser [Katz and Winston]. An example symbolic 2-D shape description generated by Connell's program is shown in Figure 1. This figure shows the description generated for a hammer. The semantic network describes the hammer as an object which consists of a *long straight elongated* handle *attached-to* and *centered-in* a *medium-long curved* head. As can be seen in the figure, the full semantic

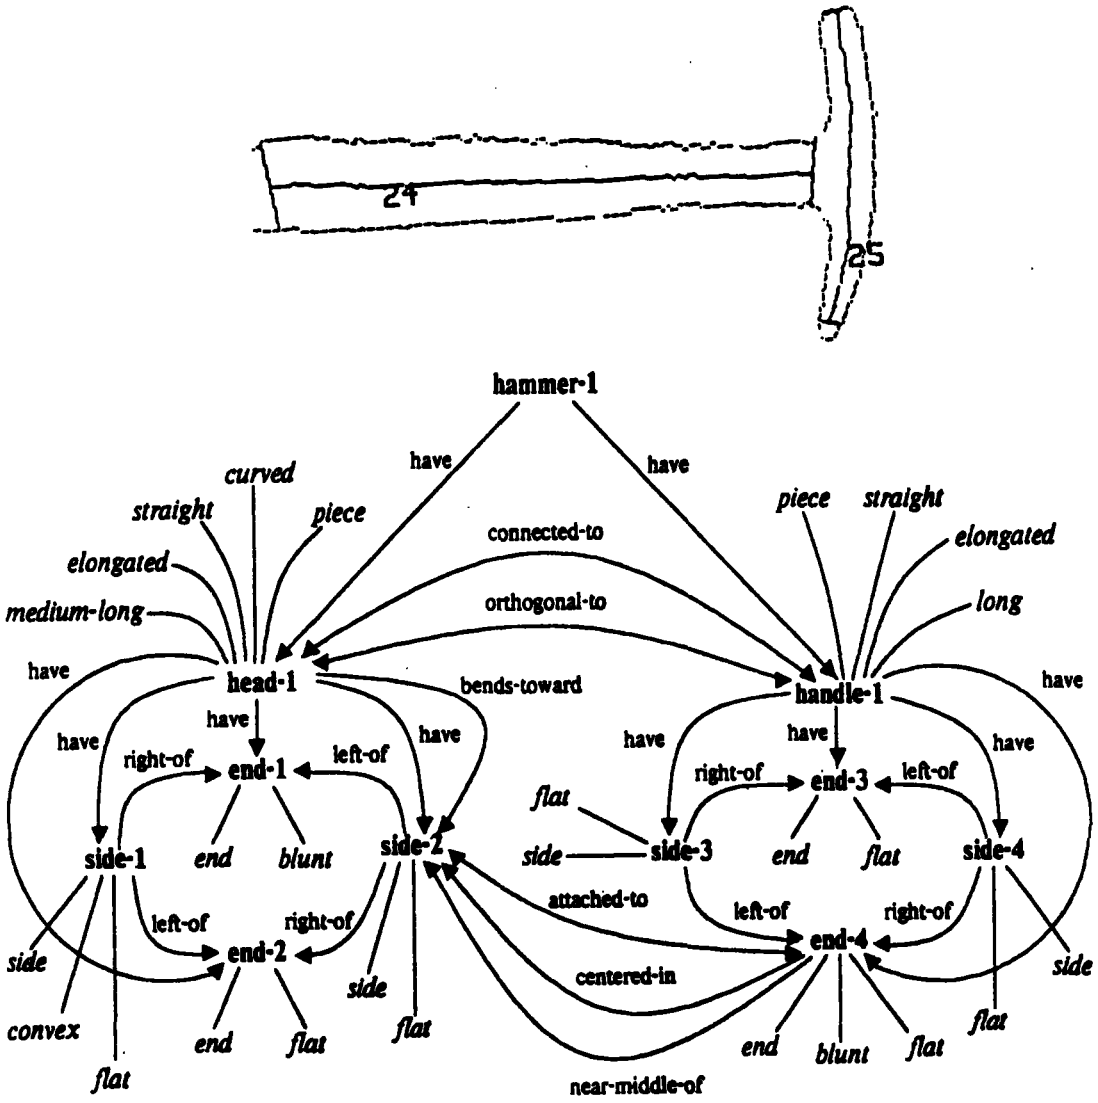

Figure 1. 2-D SLS of a Tack Hammer

a. The main smoothed local symmetries extracted from the edge data of the tack hammer image.

b. The semantic network that is computed from the information in a. (From [Connell])

network contains a rich description of the hammer's shape.

We expect to develop symbolic three-dimensional descriptions from the 3-D Smoothed Local Symmetry representations of objects similar to the two-dimensional descriptions of Connell. For the hammer shown in Figure 1, a similar semantic net would be generated, having additional 3-D descriptors which would, for example, identify the handle as being cylindrical and the head as being box-like. The local symmetries used by the 3-D Smoothed

Local Symmetry make it amenable to describing subparts of an object for a hierarchical symbolic description. Having the subparts described thusly also aids in describing the relations between those subparts.

Extension of 2-D Smoothed Local Symmetries

Brady's two-dimensional Smoothed Local Symmetries (2-D SLS's) [Brady and Asada] generate skeletons of 2-D objects as a means of representing these objects. These skeletons appear to be natural, perceptually salient skeletons for objects composed mainly of elongated regions. (The problem of representing "round" regions in a similar fashion has been addressed by [Fleck] via Local Rotational Symmetries.) With the 3-D Smoothed Local Symmetries (3-D SLS's) described herein, we extend the existing 2-D Smoothed Local Symmetry representation to three-dimensional objects, the intent being to provide a natural skeletonization useful for describing 3-D objects.

As stated above, the 2-D SLS is a skeleton of a 2-D object. It makes explicit the important structural features of a 2-D object while suppressing less-important details (Figure 2). Noting that the skeleton (composed of 1-D line segments) is a representation of co-dimension 1 of the (2-D) object, we are led to the proposition that the skeleton of a 3-D object should consist of two-dimensional entities since such a skeleton would be of co-dimension 1 with the generating (3-D) object. Our intuition tells us that the skeleton of rectangular cutting board should be a plane, and that of a cube should consist of three intersecting planar patches (Figure 3). We will see that our intuition agrees fairly well with the 3-D SLS representation, except that the 3-D SLS provides more information for (i.e., a fuller representation of) the generating 3-D object than does the intuitive skeleton.

We note at this point that although this 3-D SLS research comes chronologically *after* the 2-D SLS work, a hierarchy of shape description methods should have the inverse structure. That is, a good method for describing three-dimensional shapes should yield in a natural way a representation for two-dimensional objects. As will be seen from the definition of the 3-D SLS, the 2-D SLS is a special case of the more general three-dimensional method. Indeed, it has been shown that if part of a 3-D object can be formed by sweeping a 2-D closed curve along a space curve, then a subset of the 3-D SLS surfaces of that 3-D object is simply the 2-D SLS swept along the same space curve [Brady 1985].

Previous Work

[Marr and Nishihara] proposed the use of generalized cylinders to represent three-dimensional objects. However, arguments have been given to show that generalized cylinders are not the skeletons one would want for shape de-

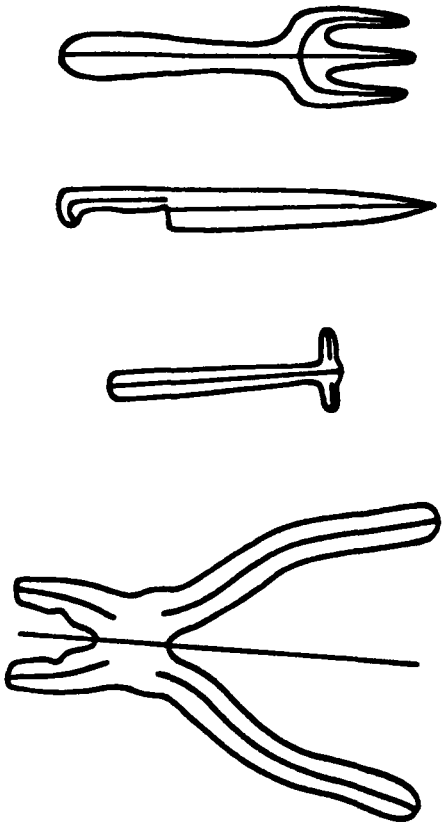

Figure 2. Examples of 2-D Smoothed Local Symmetries
(From [Brady and Asada])

scription [Brady 1983]. Indeed, our intuitive description of a cutting board is not a line along which a rectangle has been swept (the generalized cylinder representation), but rather a plane (the 3-D SLS representation). However, the generalized cylinder representation has been used with some success in a few domains [Brooks].

An objection can be made to the use of generalized cylinders on the grounds that the representation is not unique. Indeed, an object as simple as a wedge has at least two different generalized cylinder representations [Brady]. The representation is best suited to objects which have a natural axis, yet for many everyday objects that we are interested in describing, such as a chair or a table, a salient natural axis is difficult to perceive. As will be seen, the concept of *symmetry* on which the 3-D SLS is based subsumes the concept of axis and thus the representation is applicable to a wider domain than are generalized cylinders.

A different approach to the skeletonization of three-dimensional objects has been presented by [Nackman 1982], this approach based on the Symmetric

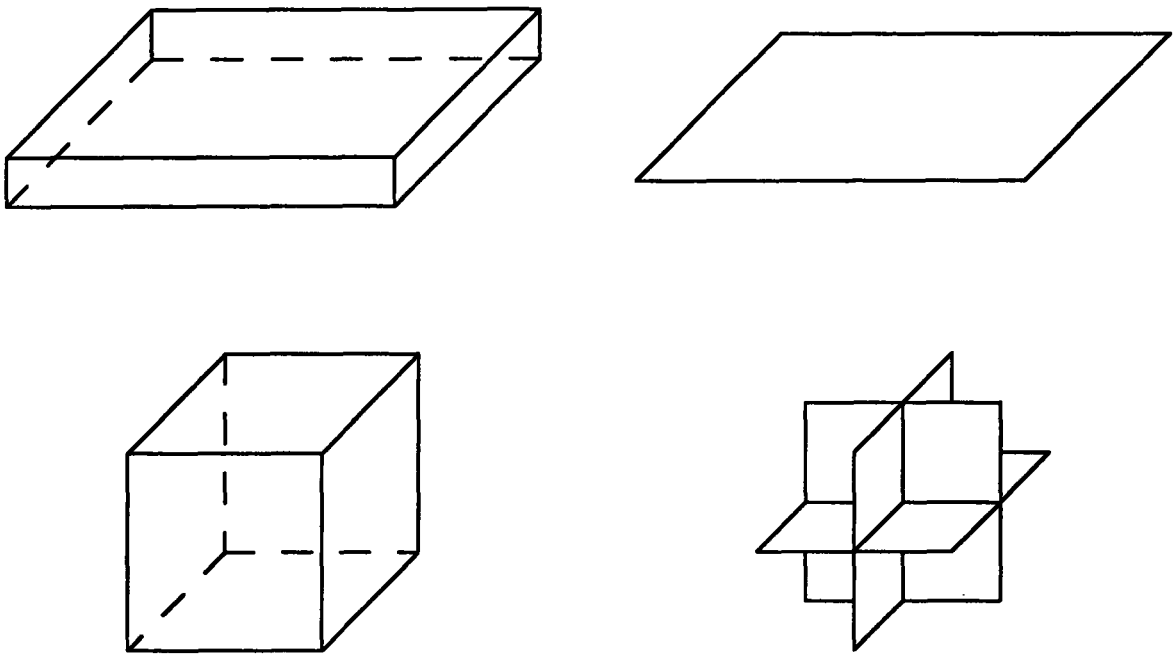

a. b.

Figure 3. Some intuitive skeletons
a. Object. **b.** Skeleton.

Axis Transform (SAT). The same arguments given by [Brady 1983] against the use of the SAT for describing two-dimensional objects carry over to 3-D. The most salient deficiency of the 3-D SAT is its poor performance in describing an object's volume near its boundaries. Because the 3-D SAT of an object is defined by the locus of maximal spheres packed into that object, the 3-D SAT skeleton cannot extend to the boundaries of that object. Thus, Nackman limits the outlines of objects considered for the 3-D SAT to exclude "surfaces with corners, edges, or cusps, such as polyhedra." We also note that, although this method can be used to describe objects containing holes, it cannot describe the holes themselves. It will be seen that the 3-D SLS can be used to describe polyhedra and that the method can also describe objects with holes as well as the holes themselves.

The Extended Gaussian Image (EGI) has also been proposed as a shape representation and has been used with photometric stereo work [Horn 1983]. However, this method goes against our bias toward representations which are in some way perceptually salient. In addition, we feel that the EGI does not make object features explicit in a way that can easily support reasoning about 3-D objects. Finally, the domain of the EGI must exclude objects with holes in order to yield unique object descriptions.

Domain of Description

The concept of two surface elements of an object being related by a symmetry about a plane is essential to the definition of the 3-D Smoothed Local Symmetry (see Chapter 2). The planar symmetries, in turn, influence what types of shapes are best described by the 3-D SLS.

The class of three-dimensional shapes most naturally described by 3-D SLS's are those which for which the whole or subparts of the whole have planar symmetries. However, as the name implies, 3-D Smoothed Local Symmetries describe *local* symmetries. Thus, objects with areas of local planar symmetry are amenable to description via the 3-D SLS. Polyhedra are described naturally since their planar faces give rise to (planar) symmetries. Holes are handled naturally and are not restricted. These characteristics fit in well with our goal of describing manufactured objects. (The *smooth* part of Smoothed Local Symmetries arises from the operation of smoothly joining the local symmetries found in an object.)

The representation works least well on objects which tend to be *round*, i.e., approximate a sphere. The point symmetry of a sphere is a degenerate case of the planar symmetries on which the 3-D SLS is based. The other degenerate case of these symmetries is symmetry with respect to an axis of rotation, exemplified by a cylinder with circular cross section. It should be noted that these two degenerate cases most often occur with natural as opposed to man-made objects. Objects which approximate spheres typically are formed by accretion or growth and cylindrical objects by growth. We have already stated that we are mainly interested in representing objects which are manufactured, so this limitation is not too severe. However, [Fleck] has already provided a solution for the 2-D case with Local Rotational Symmetries. It should be possible to extend the idea of rotational symmetries to 3-D to describe the degenerate cases. Together, Smoothed Local Symmetries and Local Rotational Symmetries should be able to represent most objects in the real world.

The above discussion of symmetry bring up an important distinction between Smoothed Local Symmetries and generalized cylinders. Whereas a generalized cylinder seems to be a natural description for an inflated fire hose, a deflated fire hose is better described as a ribbon, which is just the 3-D SLS representation it would have. The concept of a single *axis* is not natural in the latter case. We would also argue that the inflated hose is also best described by the symmetry about its central axis and that an extension of Smoothed Local Symmetries along the lines of Fleck's Local Rotational Symmetries would naturally fit into the general scheme of representing objects via their symmetries.

Efficiency

An important consideration in shape representation is efficiency of representation. Since the 3-D SLS represents three-dimensional *volumes* with two-dimensional *surfaces*, the representation has the inherent storage efficiency of a co-dimensional 1 representation. In other words, if the complexity of an object in m-dimensional space is n^m, then we would expect the complexity of the representation of co-dimension 1 of that object to be n^{m-1}. Also, since the algorithm for computing the 3-D SLS is essentially a parallel computation, the 3-D SLS should lend itself to a natural implementation on a parallel computer such as the Connection Machine [Hillis].

Support for Useful Generalization

By describing objects via their symmetries, our intent is to make the salient characteristics of the objects explicit while suppressing detail at the top level of representation. In this way, we feel that we can support useful generalization, which is an important part of physical reasoning. For example, the important characteristics of a chair are its back and seat. These would be represented by (possible curved) planes in the Smoothed Local Symmetry scheme. In fact, most chairs have a horizontal plane-like seat with a vertical plane-like back, with an edge from each in close proximity to the other. These characteristic features of a chair, along with the support leg or legs, would be at the top of a symbolic description hierarchy of a chair that could be derived from a 3-D Smoothed Local Symmetry representation of a chair in a straightforward manner.

2. 3-D SLS Definition

To define a Smoothed Local Symmetry, we first make explicit our notion of symmetry. We say that two points on the surface of an object have a symmetry with each other if the points are equidistant from some arbitrary plane and if the surface normal of the object surface at one of the points can be reflected through this plane to form the surface normal at the other point. In other words, we require both the points and their normals to be reflections of each other across the symmetry plane. These symmetries are obviously local since they depend only on the two surface points and their normals (which are based on local surface information).

Generalizing from the 2-D SLS, we offer a formal definition of a 3-D SLS point which can be used to construct SLS points from object surface points.

The 3-D SLS, then, is defined by describing the symmetries between pairs of points on an object's surface. A symmetry exists between two surface points if their normals intersect (at a point known as the SAT[1] point [Blum]) and the distances from each surface point to the SAT point are equal. If this condition holds, then a symmetry exists and the 3-D SLS point is defined as the midpoint between the two surface points. A plane placed at this 3-D SLS point and oriented to be perpendicular to the line connecting the two generating surface points would then represent the symmetry plane described above for these two points. See Figure 4 for the 3-D SLS geometry.

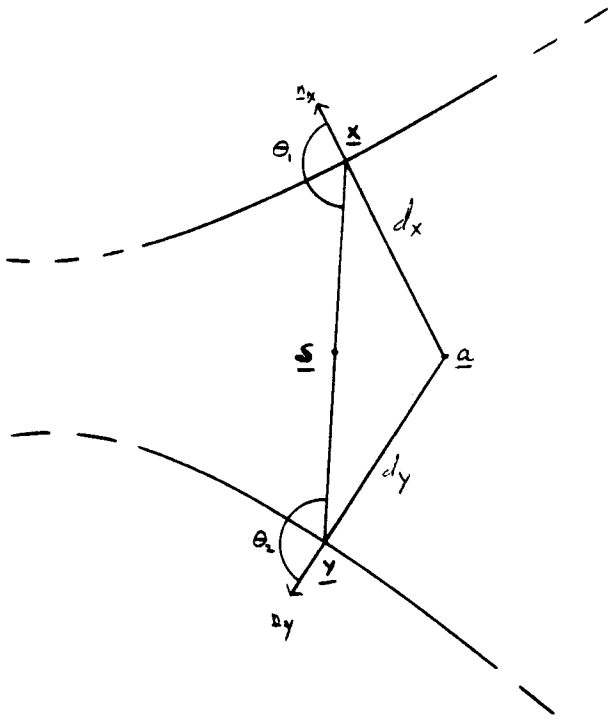

Figure 4. 3-D SLS Geometry
A symmetry exists between surface points **x** and **y** if
a) Their normals $\mathbf{n_x}$ and $\mathbf{n_y}$ intersect (at the SAT point **a**), and
b) The distances d_x and d_y are equal.
If the above conditions hold, the 3-D SLS point is the midpoint **s** between surface points **x** and **y**.

It should be noted that the Symmetric Axis Transform (SAT) has been sug-

[1]SAT, which stands for Symmetric Axis Transform, is actually a misnomer according to our definition of symmetry. Although the SAT point lies on the symmetry plane described above, it only coincides with the 3-D SLS point when the normals of the two surface points involved are colinear. The two surface points do not exhibit a symmetry across the SAT point.

gested for use in describing three-dimensional objects. [Nackman 1982] has offered a method of describing the 3-D SAT surfaces of an object. However, unlike the work presented here, no method has yet been implemented to derive the 3-D SAT from a three-dimensional object ([Nackman 1985], p. 201).

An alternate, but equivalent, method of determining 3-D SLS points exists. If we connect two surface points by a line segment, then a symmetry exists between those two surface points if the angles between their normals and the connecting line segment are equal and if their normals (and thus also the connecting line segment) are coplanar. The 3-D SLS point is again defined as the midpoint of this line segment.

Note that these definitions do not require the 3-D SLS point to be contained inside the object which generated them. (By Blum's definition, SAT points must lie inside the generating object.) Finally, the 3-D SLS points are smoothly joined to form symmetry surfaces which describe the generating object.

As in the 2-D SLS, data can be attached to each symmetry point containing information about the surface points which generated it. This information could include:

- the distance from the symmetry point to the generating points on the object
- the principal curvatures of the object surface at the generating points
- the normal vectors of the generating points
- identification of the patches containing the generating points on the object

Given the symmetry points and the information described above, it would be possible to reconstruct the original object from its 3-D SLS description.

2.1 The 3-D SLS of Some Simple 3-D Objects

Using the definition of the 3-D Smoothed Local Symmetry described above, we have generated the Local Symmetry surfaces of some simple solids. The solids were represented by ($2\frac{1}{2}$-D) depth maps of the 3-D solid objects. The symmetries are shown in Figure 5.

3. Analytic 3-D SLS

We have derived the analytic expressions for the 3-D Local Symmetry between

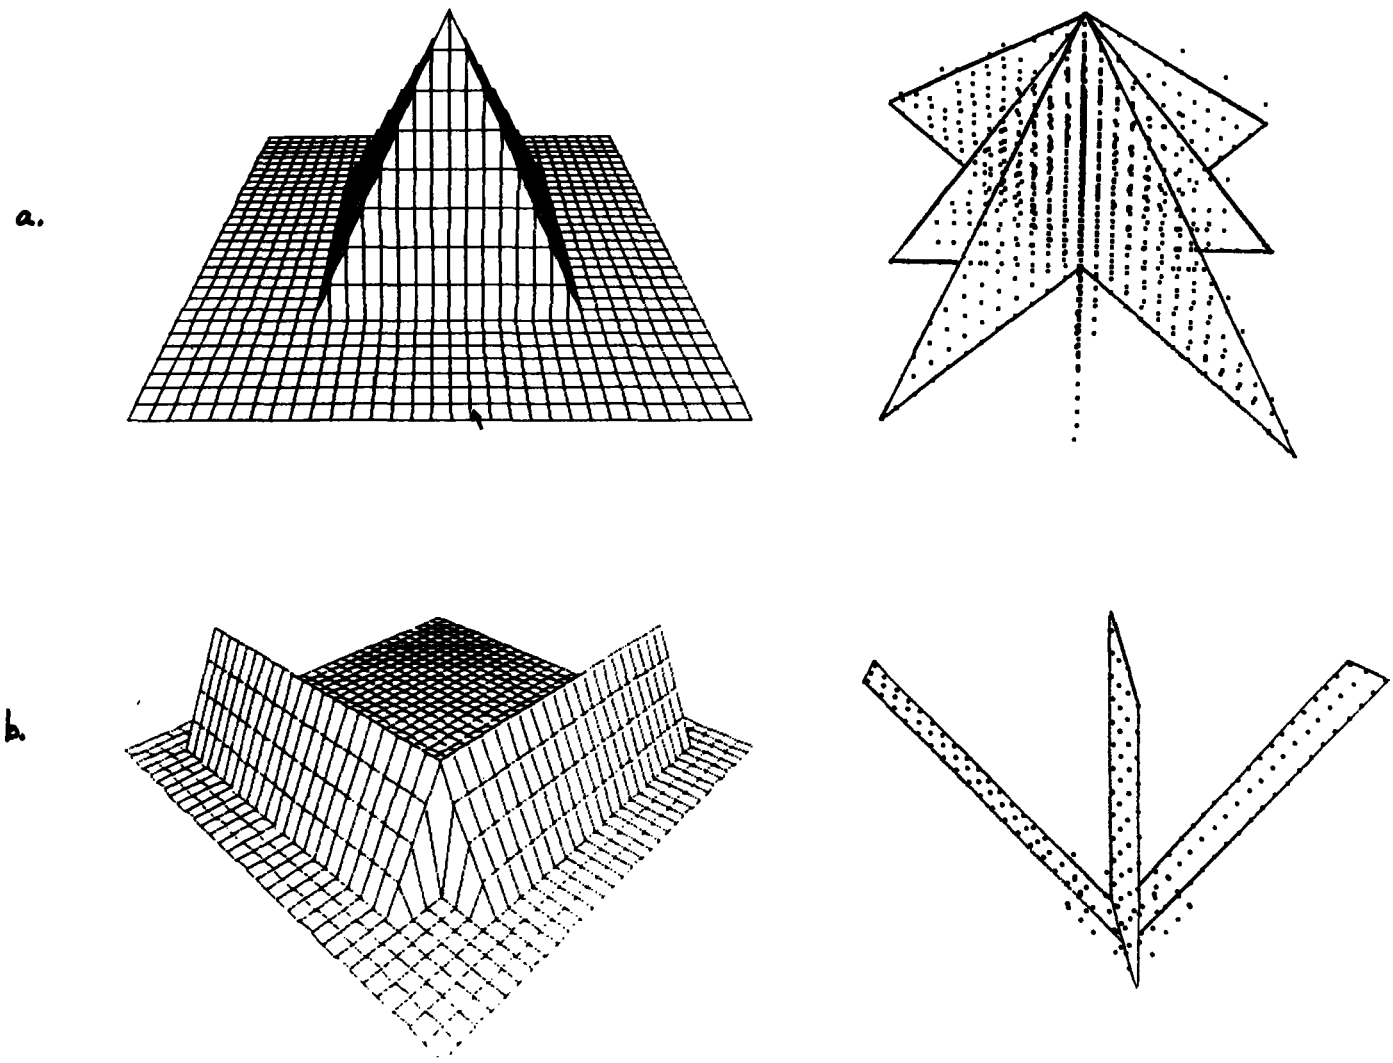

Figure 5. Computed Local Symmetries of Some Simple Objects
The left column shows the object and the right column shows the corresponding Local Symmetries for: **a.** A four-sided pyramid. **b.** An L.

two surfaces for three special cases. These special cases are where the two surfaces are

a) both planes
b) a plane and a spherical surface
c) both spherical surfaces.

The method used in computing the 3-D Local Symmetry is to travel a distance λ along the normal[2] from a surface, then travel the same distance λ

[2]All normals mentioned in this paper are unit normals.

back along the normal of a second surface. Using the equation of the second surface as a constraint (since the final point reached from this "traveling" must lie on the second surface for a symmetry to exist between the surfaces), a solution for λ can be obtained. The equation for points on surface #2 which correspond via symmetry to points on surface #1 is then derived as a function of the coordinates of points on surface #1. Finally, the equation for the symmetry points as a function of the coordinates of points on the first surface is obtained by noting that these symmetry points are simply the midpoints between surface #1 points and their corresponding (via symmetry) points on surface #2. Since we have analytic descriptions of surface #1 and surface #2, the above method yields an analytic description of the symmetry surface.

Local Symmetry points generated from a pair of surfaces will themselves be organized into a surface, though not necessarily one which has a simple equation. However, the Local Symmetry surfaces resulting from two adjacent pairs of object surfaces will not necessarily abut, let alone join smoothly. For this reason, a smoothing method is needed to join the individual Local Symmetry surfaces into one Smoothed Local Symmetry surface for the whole object. Such a smoothing algorithm must be developed for 3-D Smoothed Local Symmetries to allow the representation of complex objects.

3.1 Surface Representations

The surfaces we will be dealing with are planes and spheres. Below are presented some notations and conventions for representing these surfaces [Faugeras and Hebert].

Planes

A plane is described by the equation

$$P: \quad \mathbf{x} \cdot \mathbf{v} - d = 0$$

where

$\mathbf{v}$ is the unit normal vector to the plane

and

d is the distance from the origin to the plane.

Also, we choose the orientation of the normal vector to be toward the outside of the object of which it is a part.

Quadrics

A quadric is described by a symmetric 3×3 matrix $\mathbf{A}$, a vector $\mathbf{v}$, and a scalar d:

$$Q: \quad \mathbf{x}^{\mathrm{T}}\mathbf{A}\mathbf{x} + \mathbf{x} \cdot \mathbf{v} + d = 0$$

The center of the quadric, when it exists, is given by

$$\mathbf{c} = -\frac{1}{2}\mathbf{A}^{-1}\mathbf{v}.$$

Spheres

A sphere is a special case of a quadric where the **A** matrix is a multiple of the identity matrix, i.e., $\mathbf{A} = n\mathbf{I}$. In this case, the vector from the origin to the center of the sphere is given by $\mathbf{c} = -\frac{1}{2n}\mathbf{v}$ and the radius of the sphere by $\sqrt{\mathbf{c}\cdot\mathbf{c} - d/n}$.

3.2 Plane-Plane 3-D SLS

If the two surfaces generating the 3-D SLS are both planes, then the computation of the 3-D SLS surface equation is straightforward. First, we assume that the two planes are described by the equations:

$$P1: \quad \mathbf{x}\cdot\mathbf{v_1} - d_1 = 0$$

$$P2: \quad \mathbf{x}\cdot\mathbf{v_2} - d_2 = 0$$

The corresponding symmetry point $\mathbf{x_2}$ on P2 to any point $\mathbf{x_1}$ on P1 is defined by

$$\mathbf{x_2} = \mathbf{x_1} + \lambda\mathbf{v_1} \pm \lambda\mathbf{v_2} \tag{1}$$

where λ can be positive or negative. Using the fact that $\mathbf{x_2}$ must also lie on P2, i.e.,

$$\mathbf{x_2}\cdot\mathbf{v_2} - d_2 = 0 \tag{2}$$

we can solve for λ by substituting (1) into (2):

$$\lambda = \frac{d_2 - \mathbf{x_1}\cdot\mathbf{v_2}}{\mathbf{v_2}\cdot(\mathbf{v_1} \pm \mathbf{v_2})}. \tag{3}$$

If $\mathbf{v_1} = \pm\mathbf{v_2}$, we can have zero solutions (P1 and P2 are coplanar) or one solution (P1 and P2 are parallel). Otherwise, the planes intersect and we get two solutions (since the SAT point can fall on either side of P1).

When dealing with objects, P1 and P2 are planar patches, not full planes. We can use equation (1) with the λ(s) of (3) to find the (possibly empty) subset of points (from the set of $\mathbf{x_2}$ points found above) which lie on P2 and which correspond via symmetry with the points of P1. With the corresponding subsets of points $\mathbf{x_1}$ from P1 and $\mathbf{x_2}$ from P2, we can find the 3-D SLS patch from

$$\mathbf{x_{SLS}} = \frac{1}{2}(\mathbf{x_1} + \mathbf{x_2}).$$

3.3 Plane-Sphere 3-D SLS

The equations describing the 3-D SLS between a plane and a sphere are a bit more complicated than the plane-plane case, but still tractable. We assume that the plane and the sphere are described, respectively, by

$$P: \quad \mathbf{x} \cdot \mathbf{v_p} - d_p = 0$$
$$S: \quad \mathbf{x}^{\mathbf{T}}\mathbf{A}\mathbf{x} + \mathbf{x} \cdot \mathbf{v_s} + d_s = 0.$$

where $\mathbf{A}$ is a multiple of the 3×3 identity matrix.

The normal vector for any point $\mathbf{x_s}$ on a sphere is simply calculated by subtracting the sphere center coordinates from the point coordinates, $\mathbf{n_s} = \mathbf{x_s} - \mathbf{c}$. Then, for any point $\mathbf{x_s}$ on the sphere, the corresponding symmetry point on the plane is

$$\mathbf{x_p} = \mathbf{x_s} + \lambda \mathbf{n_s} \pm \lambda \mathbf{v_p} \tag{4}$$

where λ can be positive or negative. Using the fact that $\mathbf{x_p}$ must also lie on the plane, i.e.,

$$\mathbf{x_p} \cdot \mathbf{v_p} - d_p = 0, \tag{5}$$

we can solve for λ by substituting (4) into (5):

$$\lambda = \frac{d_p - \mathbf{x_s} \cdot \mathbf{v_p}}{\mathbf{v_p} \cdot (\mathbf{n_s} \pm \mathbf{v_p})}. \tag{6}$$

Note that two solutions are possible since the SAT point can lie on either side of the plane. Finally, we can substitute (6) back into (4) to obtain the coordinates of the symmetry point on the plane, $\mathbf{x_p}$, corresponding to $\mathbf{x_s}$.

Again, when dealing with objects, P and S are planar and spherical patches, respectively. We check the set of $\mathbf{x_p}$ points found above to determine the (possibly empty) subset of points lying on P which correspond via symmetry with the points of S. With the corresponding subsets of points $\mathbf{x_p}$ from P and $\mathbf{x_s}$ from S, we can find the 3-D SLS patch from

$$\mathbf{x_{SLS}} = \frac{1}{2}(\mathbf{x_p} + \mathbf{x_s}).$$

3.4 Sphere-Sphere 3-D SLS

For a 3-D SLS generated by two spheres, we assume that the spheres are described by

$$S1: \quad \mathbf{x}^{\mathbf{T}}\mathbf{A_1}\mathbf{x} + \mathbf{x} \cdot \mathbf{v_1} + d_1 = 0$$
$$S2: \quad \mathbf{x}^{\mathbf{T}}\mathbf{A_2}\mathbf{x} + \mathbf{x} \cdot \mathbf{v_2} + d_2 = 0$$

For point $\mathbf{x_1}$ on S1 with normal $\mathbf{n_1} = (\mathbf{x_1} - \mathbf{c_1})/\|\mathbf{x_1} - \mathbf{c_1}\|$ and point $\mathbf{x_2}$ on S2, the SAT point between $\mathbf{x_1}$ and $\mathbf{x_2}$ is

$$\mathbf{x_{SAT}} = \mathbf{x_1} + \lambda \mathbf{n_1}. \tag{7}$$

For symmetry, we require that the distance from $\mathbf{x_{SAT}}$ to the center $\mathbf{c_2}$ of sphere S2 be either $\lambda + r_2$, $\lambda - r_2$, $-\lambda + r_2$, or $-\lambda - r_2$, where r_2 is the radius of S2. λ is positive when $\mathbf{x_1}$ lies between $\mathbf{c_1}$ and $\mathbf{x_{SAT}}$ while r_2 is positive when $\mathbf{x_2}$ lies between $\mathbf{c_2}$ and $\mathbf{x_{SAT}}$. Using this constraint we find that we must have

$$\|\mathbf{x_{SAT}} - \mathbf{c_2}\| = \lambda \pm r_2, -\lambda \pm r_2,$$

i.e.,

$$(\mathbf{x_{SAT}} - \mathbf{c_2}) \cdot (\mathbf{x_{SAT}} - \mathbf{c_2}) = \lambda^2 \pm 2\lambda r_2 + r_2^2. \tag{8}$$

Substituting (7) into (8), we obtain

$$A\lambda^2 + B\lambda + C = 0$$

where

$$A = \mathbf{n_1} \cdot \mathbf{n_1} - 1$$
$$B = 2[\mathbf{n_1} \cdot (\mathbf{x_1} - \mathbf{c_2}) \pm r_2]$$

and

$$C = (\mathbf{x_1} - \mathbf{c_2}) \cdot (\mathbf{x_1} - \mathbf{c_2}) - r_2^2.$$

Thus,

$$\lambda = \frac{-B \pm \sqrt{B^2 - 4AC}}{2A}. \tag{9}$$

Note that due to the $\pm r_2$ in B, there can be up to four λ's. Because there are different possibilities for the relative positions of S1 and S2 (as shown by the $\pm r_2$), we must check to see which λ's are appropriate for the given situation by verifying that the points that they generate actually lie on S2, i.e., $\|\mathbf{x_2} - \mathbf{c_2}\| = r_2$.

Using equation (7) with the λ(s) from (9), the $\mathbf{x_{SAT}}$ points for the two surfaces can be found. The points of S2 which correspond via symmetry with the points of S1 are then $\mathbf{x_2} = \mathbf{x_1} + \lambda\mathbf{n_1} + \lambda\mathbf{n_2}$ where $\mathbf{n_2} = (\mathbf{x_{SAT}} - \mathbf{c_2})/\|\mathbf{x_{SAT}} - \mathbf{c_2}\|$. After identifying the subset of $(\mathbf{x_1}, \mathbf{x_2})$ symmetry pairs which lie on the actual surface patches of the object partially described by S1 and S2 (respectively), we can find the 3-D SLS patch from

$$\mathbf{x_{SLS}} = \frac{1}{2}(\mathbf{x_1} + \mathbf{x_2}).$$

3.5 Singularities

In the above development of the analytic 3-D SLS, we have ignored the fact that a singularity can occur when dealing with spherical patches. The singularity occurs when the $\mathbf{x_{SAT}}$ point for a spherical patch coincides with the center of that spherical patch (in this situation it is obvious that $\lambda = r_s$). In

this case, every point on the spherical patch is a symmetry point corresponding to $\mathbf{x}_{\mathbf{SAT}}$.

We note that a spherical patch has a self-symmetry at its center point. This is a degenerate case of the 3-D SLS where the symmetry of a sphere is best described by its center point instead of by some 2-D surface. We intuitively expect the analogous situation to occur with a cylinder of circular cross section and its axis. However, the cylindrical surface did not yield to the analysis presented above for planes and spheres. If we assume that such an SLS singularity occurs for cylinders, then we see that the generalized cylinder representation for a cylinder of circular cross section (which would simply be its axis) is a special case of its 3-D SLS representation.

3.6 Implementation

From the segmentation algorithm of Chapter 4, we obtain the analytic descriptions of the surface patches as described above. However, these descriptions are not parameterized. Therefore, rather than generating the symmetry surface as a whole, we generate it in a point-by-point fashion from the above equations. A natural parameterization of the analytic surface descriptions is needed to take full advantage of the equations for the SLS described above.

4. Surface Fitting

We chose to test the above derivation of the 3-D SLS on $2\frac{1}{2}$-D depth maps of various objects instead of on their full 3-D representations. Although methods have been presented for representing the surfaces of an object in its full three dimensions [Baker], the availability of $2\frac{1}{2}$-D depth maps for a wide variety of objects and the ease with which they can be generated (compared with the effort to generate the full 3-D representation) led us to implement the 3-D SLS for depth maps. A depth map provides a sufficiently detailed and rich description of an object to allow testing of the 3-D SLS concepts discussed above. Thus, we proceed with the task of generating the 3-D SLS for objects represented by $2\frac{1}{2}$-D depth maps.

Having derived the 3-D SLS between surfaces, the task of segmenting the object's surface into analytic surfaces remains. Since no closed form solution to the 3-D SLS could be found between general quadric surfaces, the segmentation algorithm works solely with planar and spherical patches. (The

difficulty with quadrics is that, in general, the normals from any number of surface points can pass through any given point in space.)

Following the lead of [Faugeras and Hebert], we implemented a surface segmentation algorithm which will segment the depth map representation of an object into planar and spherical patches. The algorithm uses a region-growing technique to merge neighboring regions until no merge can be performed without exceeding some predetermined global error criterion. These planar and spherical patches are then used in computing the 3-D SLS for an object.

A second segmentation algorithm was also implemented and tested, but did not lead to as good results in the 3-D SLS computation as the above-mentioned technique. In this second method, the object was first partitioned into quad-tree patches. Planes and spheres were then fit to these patches, subdividing the patches where necessary to obtain a fit within the acceptable error. The thought was that this segmentation would be faster than the region-growing one while yielding rectangular bounds for the projections of the patches to make bounds checking easier for the 3-D SLS algorithm. Unfortunately, the artificial patch shape constraints imposed by this method did not yield good results in the calculation of the 3-D SLS from surface patches.

5. Experimental Results

Testing of the segmentation and 3-D SLS algorithms on synthetically generated surfaces yielded good results. The solids were represented by ($2\frac{1}{2}$-D) depth maps of the 3-D solid objects. The symmetries are shown in Figure 5. However, performance on unsmoothed depth maps of real objects was not as pleasing. With real objects, the surface patches tended to be small, yielding many symmetries for the objects. Even the 75° heuristic[3] from the 2-D SLS work did not improve results much. Part of the problem is obviously the profusion of symmetries in the 3-D case compared to the 2-D case.

[3]The 75° heuristic simply states that if the connecting line between the two symmetry points and either of their surface normals form an angle of greater than 75°, then those two points do not provide a valid symmetry. This heuristic prevents, for example, the algorithm from declaring that a point on a planar or near-planar surface forms a symmetry with all the other points on that surface.

6. Problems, Limitations, Future Work

One great limitation in dealing with the 3-D SLS is the lack of a good method for displaying the 3-D SLS points which are generated. We implemented an algorithm to display a perspective view of the 3-D SLS points from any specified viewing direction and distance. We also tried generating stereo pairs for a 3-D view of the 3-D SLS points, but the viewing the 3-D SLS points in this way is not easy.

The algorithms for generating the 3-D SLS will be slow until a way can be found to parameterize the surfaces derived from the surface-fitting (this has been done for the 2-D SLS). The quad-tree-patches segmentation scheme was one attempt at generating patches with easily derivable parameters.

With the 3-D SLS generated from patches, the only way of grouping the 3-D SLS points is by identifying their generating patches. What is really needed is a way to group these points into symmetry surfaces. Parameterized descriptions of the surface patches should help since they would yield analytic descriptions of the symmetry surfaces instead of an aggregation of points as is now obtained.

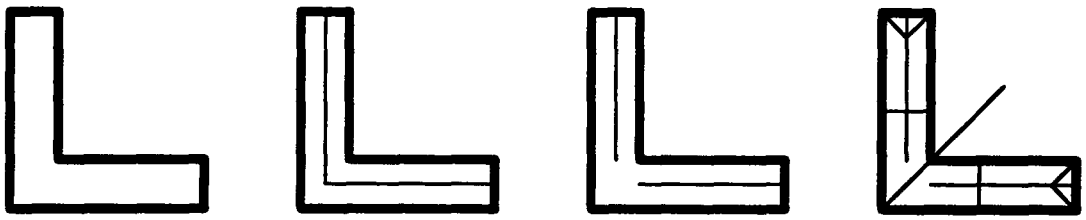

a. b. c. d.

Figure 6. Symmetries for an L shape
a. L. **b.** Expected symmetries. **c.** Gap in major symmetries at join. **d.** Non-trivial symmetries.

When symmetry patches are generated from surface patches, adjacent surface-patch pairs will not necessarily yield continuous symmetry surface patches. This is a direct result of the surface patches being only approximations to the true surface shape that they represent. Smoothing methods will be needed to deal with the symmetry surface patch intersections. Smoothing methods will also have to deal with extending the SLS surfaces across gaps. A simple example of this can be seen in 2-D with the letter "L" (see Figure 6.a). One would expect the symmetries to consist of two lines which form a thinner "L" and which lie in the interior of the original "L" (Figure 6.b). However, due to the definition of the SLS, the symmetry lines of the legs of the "L" do not extend into the area where the legs join (Figure 6.c). (The non-trivial

symmetries are shown in Figure 6.d.) [Bagley] has considered this problem for the 2-D SLS and has demonstrated methods for smoothly extending the SLS lines.

Sequential computation of the 3-D SLS will always cause the generation of a 3-D SLS for a given object to take longer than the generation of the same 3-D SLS using parallel processing. This is due to the fact that the 3-D SLS is generated by an essentially parallel process: the comparison of every surface patch on an object with all other surface patches. This seems to be a natural task for a SIMD machine such as the Connection Machine [Hillis], a possibility which should be explored further.

After the issues mentioned above have been addressed, the next step for this research will be to develop symbolic object descriptions from the 3-D Smoothed Local Symmetry representation. By examining the major and minor symmetries of an object, its salient structure can be derived. These object descriptions would be appropriate for a reasoning system [Brady, Agre, Braunegg, and Connell] [Braunegg], as described previously.

References

Blum, H., "Biological Shape and Visual Science (Part 1)," *J. Theor. Biol.*, **38**, 1973, 205–287.

Bagley, S. C., "Using Models and Axes of Symmetry to Describe Two-Dimensional Polygonal Shapes," Master of Science thesis, Department of Electrical Engineering and Computer Science, MIT, February 1985.

Baker, Henry H., "Building Models of Three-Dimensional Objects," Master of Philosophy thesis, Department of Machine Intelligence, University of Edinburgh, 1976.

Brady, J. Michael, "Criteria for Representations of Shape," *Human and Machine Vision*, Jacob Beck, Barbara Hope, and Azriel Rosenfeld (eds.), Academic Press, New York, 1983.

Brady, J. Michael, personal communication, September 1985.

Brady, J. Michael and Harou Asada, "Smoothed Local Symmetries and Their Implementation," MIT AI Memo 757, February 1984.

Brady, J. Michael, Philip Agre, David J. Braunegg, and Jonathan H. Connell, "The Mechanic's Mate," *European Conference on Artificial Intelligence*, Pisa, Italy, September 1984.

Braunegg, David J., "Function from Form," *First International Symposium on Knowledge Engineering*, Madrid, Spain, November 1985.

Brooks, Rodney A., "Symbolic Reasoning Among 3-D Models and 2-D Images," *Artificial Intelligence*, **17**, 1981, 285–348.

Connell, J. H., "Learning Shape Descriptions: Generating and Generalizing Models of Visual Objects," Master of Science thesis, Department of Electrical Engineering and Computer Science, MIT, September 1985.

Connell, J. and M. Brady, "Learning Shape Descriptions," *IJCAI*, Los Angeles, August 1985.

Faugeras, O. D. and M. Hebert, "The Representation, Recognition, and Positioning of 3-D Shapes from Range Data," *International Journal of Robotics Research*, **5** (3), Fall 1986.

Fleck, M. M., "Local Rotational Symmetries," Master of Science thesis, Department of Electrical Engineering and Computer Science, MIT, September 1985.

Grimson, W. Eric L., *From Images to Surfaces: A Computational Study of the Human Early Visual System*, MIT Press, Cambridge, MA, 1981.

Hillis, W. Daniel, *The Connection Machine*, MIT Press, Cambridge, MA, 1985.

Horn, Berthold K. P., "Obtaining Shape from Shading Information," *The Psychology of Computer Vision*, P. H. Winston (ed.), McGraw-Hill, New York, 1975.

Horn, Berthold K. P., "Extended Gaussian Images," MIT AI Memo 740, July 1983.

Katz, B. and P. H. Winston, "Parsing and Generating English Using Commutative Transformations," MIT AI Memo 677, May 1982.

Marr, D. and H. K. Nishihara, "Representation and Recognition of the Spatial Organization of Three Dimensional Shapes," *Proc. R. Soc. Lond.*, **B 200**, 1978, 269–294.

Nackman, L. R., "Three-Dimensional Shape Description Using the Symmetric Axis Transform," Doctor of Philosophy thesis, Department of Computer Science, University of North Carolina, 1982.

Nackman, L. R., "Three-Dimensional Shape Description Using the Symmetric Axis Transform I: Theory," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAMI–7** (2), March 1985, 187–202.

Ullman, S., *The Interpretation of Visual Motion*, MIT Press, Cambridge, MA, 1979.

Winston, P. H., T. O. Binford, B. Katz, and M. Lowry, "Learning Physical Descriptions from Functional Definitions, Examples, and Precedents," MIT AI Memo 679, January 1983.