Deformable Fourier models for surface finding in 3D images

Lawrence H. Staib[†] and James S. Duncan[†,*]

Departments of [†]Diagnostic Radiology and *Electrical Engineering
Yale University, New Haven, CT 06510, U.S.A.

ABSTRACT

This paper describes a new global shape parametrization for smoothly deformable three-dimensional objects, such as those found in biomedical images, whose diversity and irregularity make them difficult to represent in terms of fixed features or parts. This representation is used for geometric surface matching to three-dimensional image data. The parametrization decomposes the surface into sinusoidal basis functions. Four types of surfaces are modeled: tori, open surfaces, closed surfaces and tubes. This parametrization allows a wide variety of smooth surfaces to be described with a small number of parameters. Surface finding is formulated as an optimization problem. Results of the method applied to synthetic and medical three-dimensional images are presented.

1. INTRODUCTION

This work is aimed at extracting the surfaces of structures found in three-dimensional images. Three-dimensional images are available from the medical imaging modalities: magnetic resonance imaging (MRI), computed tomography (CT), single photon emission computed tomography (SPECT) and positron emission tomography (PET). Confocal microscopy is also a growing source of three-dimensional images. Segmentation in three-dimensional images is a crucial step for visualization either to strip away obscuring structures or to render the surface of the structure delineated. The problems of manual segmentation in two-dimensional images are exacerbated in three-dimensional images when it is done slice by slice. In addition, surfaces formed slice by slice either manually or automatically are likely to contain inconsistencies that will tend to corrupt the rendering.

Smoothly deformable objects do not necessarily have an obvious decomposition that can be exploited. A uniform shape representation that describes the entire shape is therefore needed and it should describe a relatively broad class of shapes. Representations for objects are needed in order to characterize and understand shape. A great deal of research has been done in the area of shape representation. For a representation to be useful for modeling it should be concise. Methods based on explicitly listing points or patches on the surface are verbose because of the implicit redundancy. Parametric representations capture the overall shape in a small number of parameters. This means that the optimization of a match measure between data and a model can occur in a lower dimensional space.

Boundary finding using only local information has often been frustrated by poor-contrast boundary regions due to occluding and occluded objects, adverse viewing conditions and noise. A model-free interpretation is doomed by the underconstrained nature of the problem. Imperfect image data can be augmented with the extrinsic information that a geometric shape model provides. In order to exploit model-based information to the fullest extent, it should be incorporated explicitly, specifically, and early in the analysis. In addition, the boundary can be profitably considered as a whole because it tends to result in a more consistent solution overall.

 0-8194-1008-X/92/$4.00

A powerful property for distinguishing an object from its surroundings in an image is overall or global shape. In order to take full advantage of shape, the problem of object identification is approached as a process of boundary finding or delineation using a boundary measure and incorporating global shape information.

2. RELATED WORK IN BOUNDARY FINDING

Local edge detectors applied to real images produce spurious edges and gaps. These problems can only be overcome by the incorporation of information from higher scale organization of the image and models of the objects sought. Contextual information has been used for boundary determination via grouping [1], relaxation labeling [2] and scale-space methods [3]. These methods, by themselves, will not necessarily find complete boundaries. Pixel search methods associate edge elements by finding an optimal path through a two-dimensional image, based on criteria designed to find boundaries. The typical objective function combines boundary strength and low overall curvature [4]. Pixel search does not generalize obviously to three dimensions because there is no natural ordering of voxels in a surface.

Other investigators have considered whole-surface methods that adjust a tentative surface mesh in order to match to the image. By considering the surface as a whole, a structure is imposed on the problem that bridges gaps and results in overall consistency. Terzopoulos *et al.* [5] used energy-minimizing meshes that are attracted to image features such as lines and edges while internal spline forces impose a smoothness constraint. The goal was to find surfaces implied by silhouettes in two-dimensional images. This idea has also been used for finding symmetry surfaces from scale space stacks of two-dimensional images [6], surfaces in range images [7, 8] and surfaces in three-dimensional images [9, 10].

Other whole-boundary methods optimize in a parameter space. Parametric representations are useful for modeling because they capture the overall shape concisely. This means that the optimization of a match measure between data and a model can occur in a lower dimensional space. Widrow [11] used parametrized templates called rubber masks to model objects. The parameters are sizes and relationships between subparts. Yuille *et al.* [12] used a similar method for finding features in images of faces. Both of these methods describe the overall shape of the structure using very few parameters. However, the object must have sufficient structure to be represented in terms of parts and a new model must be developed for each new object. This approach was also taken by Lipson *et al.* [13] applied to the spine. Pentland and his group have developed a physically-based method for analyzing shape [14]. Shapes are represented by the low-order frequency displacement eigenvectors corresponding to the free vibration modes of the object. Thus, it is similar to a Fourier representation. The shape is recovered using the finite element method.

A number of methods have been developed specifically for identifying structures from MR images. For example, Bomans *et al.* [15] use a boundary-finding method based on a 3D version of the Marr-Hildreth edge operator to find surfaces in MR brain images. Morphological operators are used to remove small holes and thin connections. A number of techniques for segmentation rely on clustering techniques for voxel classification. Cline *et al.* [16] use multispectral voxel classification, in conjunction with connectivity, to segment 3D MR brain images. This method is limited by the assumption of normality in the probability distributions of the tissues. Interactive and semi-automated methods are a compromise between hand tracing and fully automated methods. Kennedy *et al.* [17] describe a number of semi-automated methods for segmenting MR images of the brain. Hohne and Hanson [18] use mathematical morphology, connected components and thresholding to interactively segment 3D images with feedback from rendered displays.

3. SURFACE REPRESENTATIONS

Implicit equations are a traditional and natural representation which define a relationship between coordinates such that all points that satisfy this relationship belong to the structure. Such representations are ideal for determining whether specific points belong to the object but there is no general way for generating such points. Because such operations will be crucial for this work, only explicit parametric representations will be considered further. An arbitrary surface can be represented explicitly by three function of two parameters: $x(u,v)$, $y(u,v)$ and $z(u,v)$. A surface is indexed or parametrized by the two parameters (u,v). While a curve's points are naturally ordered (by arclength), there is no natural ordering of points on an arbitrary surface. Certain classes of surfaces can be represented as a single function. For example, surfaces expressible as a single function of two coordinates, $z(x,y)$, are perpendicular deformations of a plane and thus the points in the plane, (x,y), provide the parametrization. Surfaces expressible as a function of two angles, $r(\theta,\phi)$, are radial deformations of a sphere (also called stellar) and are parametrized by (θ,ϕ).

The main approaches to three-dimensional parametric modeling in computer vision have been polynomials [19], superquadrics [14, 20], spherical harmonics [21, 22] and generalized cylinders [23, 24]. All of these parametrizations are restricted to a limited class of objects.

3.1. Polynomials

Second degree algebraic surfaces have been used extensively because of their simplicity and conciseness. Quadrics are second degree surfaces which include spheres, ellipsoids, cones, cylinders, planes, paraboloids and hyperboloids. Their conciseness, however, greatly limits their expressiveness. More generally, polynomials up to degree m have been used in different forms such as Bernstein or Hermite polynomials for spline surfaces.

3.2. Superquadrics

Superquadrics are an extension of quadrics using an exponent that allows the shape to vary from an ellipsoid to a rectangular parallelepiped. Superquadrics can be expressed parametrically by:

$$\begin{aligned} x(u,v) &= x_0 + a_1 \text{sign}(\cos v \cos u) |\cos u|^{\epsilon_1} |\cos v|^{\epsilon_2} \\ y(u,v) &= y_0 + a_2 \text{sign}(\sin v \cos u) |\cos u|^{\epsilon_1} |\sin v|^{\epsilon_2} \\ z(u,v) &= z_0 + a_3 \text{sign}(\text{siin } u) |\sin u|^{\epsilon_1} \end{aligned} \quad (1)$$

The surface parameters u and v represent latitude and longitude. The exponent ϵ_1 controls the squareness in the u plane and ϵ_2 controls the squareness in the v plane. The parameters a_1, a_2 and a_3 control the size in the x, y and z directions. The basic shape can be altered by such operations as twisting, bending and tapering [25], as can any explicit representation. The main disadvantage of superquadrics is that even with these altering operations, superquadrics are limited by their doubly symmetric cross-section and thus still only represent a very limited family of shapes (without resorting to composition). Superquadrics have been augmented by deformations according to spline models [5] and strain modes [26] in order to increase their expressiveness. Hyperquadrics [27] are a generalization of superquadrics that allow smooth deformations from shapes with convex polyhedral bounds, although no explicit parametrized form is possible.

3.3. Generalized Cylinders

Generalized cylinders (or cones) are a way of representing elongated objects. They are defined by a one-dimensional curve representing the spine of the object and a two-dimensional cross-section that is swept

along the spine to define the surface. This cross-section may vary along the spine. The actual properties of this representation depend on the choices of spine (sweeping rule) and cross-section.

Practical choices usually limit the class of object that is representable. The most common restriction is to straight, homogeneous generalized cylinders (SHGCs) where the spine is straight and the cross-section shape is constant (allowing scaling). These can be defined by [23]:

$$\begin{aligned} x(u,v) &= r(u)x(v) + pz(u) \\ y(u,v) &= r(u)y(v) + qz(u) \\ z(u,v) &= z(u) \end{aligned} \tag{2}$$

where u varies along the spine, v varies along the cross-section, $r(u)$ defines the scaling, $x(t)$ and $y(t)$ define the cross-section shape and $z(u)$, p and q define the spine. If the spine is allowed to bend, the cross-section is usually taken to be perpendicular to the axis. The cylinder radius must therefore be greater than the radius of curvature or else the boundary will cross itself. If the spine and cross-section are represented parametrically, as opposed to directly as an explicit list of coordinates or segments, generalized cylinders can be completely parametric.

An object can be represented by a generalized cylinder only if there exists an axis that a cross-section can sweep along in order to define the surface. The choices for the form of the spine and the cross-section further limit the expressibility of the representation.

3.4. Spherical Harmonics

Spherical harmonics have been used as a type of surface representation for radial or stellar surfaces ($r(\theta,\phi)$). The surface is represented as a weighted sum of spherical harmonics which are orthogonal over the sphere. A surface is represented in polar coordinates by:

$$r(\theta,\phi) = \sum_{m=0}^{M}\sum_{n=0}^{N}(A_{mn}\cos n\theta + B_{mn}\sin n\theta)\sin^n\phi\, P(m,n,\cos\phi) \tag{3}$$

where P(m,n,x) is the nth derivative of the mth Legendre polynomial as a function of x. The parameters of the representation are the weights A_{mn} and B_{mn}.

This is a type of Fourier representation, as defined below, but restricted to stellar surfaces. Stellar surfaces are obtained by deforming a sphere by moving points only in the radial direction. This means that all surface points must be seen from one point in the interior. Thus, spherical harmonics model a somewhat limited class of objects.

4. FOURIER SURFACES

Smoothly deformable objects do not necessarily have an obvious decomposition that can be exploited. A uniform shape representation that describes the entire shape is therefore needed and it should describe a relatively broad class of shapes.

Fourier representations are those that express the function in terms of an orthonormal basis. The motivation for a basis representation is that it allows us to express any object as a weighted sum of a set of known functions. An orthonormal set is desirable because it makes the parameters (weights) distinct.

For example, to express the one-dimensional function $f(t)$ on the interval (a, b) in terms of the basis $\phi_k(t)$, we write:

$$f(t) = \sum_{k=1}^{\infty} p_k \phi_k(t) \qquad \text{where} \quad p_k = \int_a^b f(t)\phi_k(t)\, dt \tag{4}$$

The coefficients p, the projections of the function onto the k basis functions, are the parameters of the representation. In order to use this representation the sum is truncated. In most such representations, the higher indexed basis functions represent higher spatial variation. Therefore, if the function to be represented is expected to have limited spatial variation, as is the case for most real object boundaries, the series can be truncated and still accurately represent the function. The usual basis functions are the sinusoids [28], although others, such as orthogonal polynomials or spherical harmonics in two dimensions, are possible. The sinusoids have the advantage of representing the familiar notion of frequency.

This one-dimensional decomposition can be used as a representation for curves in two (or more) dimensions. A closed curve can be represented by two periodic functions of t, where t varies along the curve from 0 to 2π, $x(t)$ and $y(t)$. A Fourier representation for closed curves can be based on the Fourier decomposition of these two functions using the sinusoidal basis ($\phi = \frac{1}{2\pi}, \frac{\cos x}{\pi}, \frac{\sin x}{\pi}, \frac{\cos 2x}{\pi}, \frac{\sin 2x}{\pi}, \ldots$). Open curves can be represented by having the parameter t start at one end of the line, trace along the contour to the other end, and then retrace the curve in the opposite direction to create a closed path. That is, $x(t) = x(2\pi - t)$ and $y(t) = y(2\pi - t)$ [28]. The resulting functions are even and thus they can be represented by the even sinusoidal basis functions ($\phi = \frac{1}{2\pi}, \frac{\cos x}{\pi}, \frac{\cos 2x}{\pi}, \ldots$). A version of this curve representation has been applied model-based to boundary finding in two-dimensional images [29, 30].

A surface in 3D can be described explicitly by three functions of two surface parameters: $\mathbf{x}(u, v) = (x(u, v), y(u, v), z(u, v))$, where u and v vary over the surface and x, y, and z are the associated Cartesian coordinates. This surface representation imposes no restriction on the class of surfaces representable. Thus, in order to represent surfaces, a basis for functions of two variables is needed; the following can be used [31]:

$$\begin{aligned} \phi \;=\; & \{1, \cos 2\pi mu, \sin 2\pi mu, \cos 2\pi lv, \sin 2\pi lv, \cos 2\pi mu \cos 2\pi lv, \sin 2\pi mu \cos 2\pi lv, \\ & \cos 2\pi mu \sin 2\pi lv, \sin 2\pi mu \sin 2\pi lv, \ldots \qquad (m = 1, 2, \ldots; l = 1, 2, \ldots) \quad \} \end{aligned} \tag{5}$$

The function is then represented by:

$$\begin{aligned} f(u, v) = \sum_{m=0}^{2K} \sum_{l=0}^{2K} \lambda_{m,l} \; [\; & a_{m,l} \cos 2\pi mu \cos 2\pi lv + b_{m,l} \sin 2\pi mu \cos 2\pi lv + \\ & c_{m,l} \cos 2\pi mu \sin 2\pi lv + d_{m,l} \sin 2\pi mu \sin 2\pi lv] \end{aligned} \tag{6}$$

where:

$$\lambda_{m,l} = \begin{cases} 1 & \text{for} \quad m = 0, \; l = 0 \\ 2 & \text{for} \quad m > 0, \; l = 0 \quad \text{or} \quad m = 0, \; l > 0 \\ 4 & \text{for} \quad m > 0, \; l > 0 \end{cases}$$

truncating the series at K. This basis allows the specification of even functions using the cosine terms and odd functions using the sine terms. There are three sets of parameters corresponding to the three coordinate functions: $a_x, b_x, c_x, d_x, a_y, b_y, c_y, d_y, a_z, b_z, c_z, d_z$. While the choice of u and v is obvious for simple surfaces such as spheres (use latitude and longitude) or cylinders (use longitude and height), very complicated surfaces will require some further analysis to determine the appropriate surface parametrization. Axis transforms [6] may provide a way of determining the overall structure on which to base the surface parametrization.

Four classes of simple surfaces in 3D will be described: tori (closed tubes), open surfaces (with one edge), tubes (open surfaces with two edges) and closed surfaces (no edges). The torus is formed using the entire basis

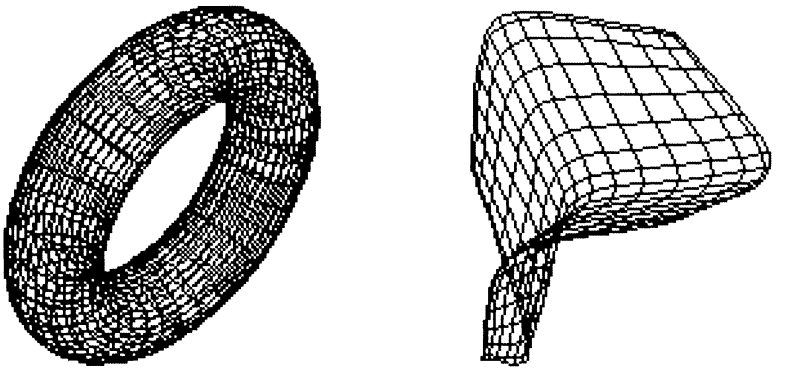

Figure 1: An example torus surface (left) using up to second order harmonics and an example open surface (right) using up to fourth.

shown in Equation 5. The result is a torus because both surface parameters are forced to be periodic. An example torus surface using this parametrization, with terms up to second order, is shown in Figure 1. The other three types of surfaces can be described using subsets of the above basis, which flatten out or constrain the torus in different ways.

4.1 Open Surfaces

Representing open surfaces with the basis in Equation 5 is complicated by the periodicity property. Since the surface is open, a straightforward representation of the surface would result in discontinuities at the boundary. Thus, these discontinuities can be avoided by having the two surface parameters start at one side of the surface, trace along the surface to the other end, and then retrace the surface in the opposite direction to create a closed path.

This results in a function $x(u, v)$ that is even and thus only the purely even terms, $a_{x,0,0}$, $a_{x,m,0}$, $a_{x,0,l}$ and $a_{x,m,l}$ are nonzero. This also holds for $y(u, v)$ and $z(u, v)$. The converse is also true; that is, any expansion with only those terms nonzero for all l and m results in an even function and thus describes an open surface. We are therefore effectively restricting the basis to include only even functions of both u and v.

$$\begin{aligned} \phi_{\text{open}} = \{&1, \cos mu, \cos lv, \ldots, \\ &\cos mu \cos lv, \ldots \qquad (m = 1, 2, \ldots; l = 1, 2, \ldots)\} \end{aligned} \tag{7}$$

Open surfaces are useful for a wide variety of structures including objects with one prominent opening, the bounding surface between two touching objects and flat objects. An open surface using this parametrization, with terms up to fourth order, is shown in Figure 1.

4.2 Tube Surfaces

Tubes require the open representation along one of the surface parameters and the closed representation along the other. This results in the following basis which is even in v and unrestricted in u:

$$\begin{aligned} \phi_{\text{tube}} = \{&1, \cos lv, \sin mu, \cos mu, \ldots, \\ &\cos mu \cos lv, \sin mu \cos lv, \ldots \qquad (m = 1, 2, \ldots; l = 1, 2, \ldots)\} \end{aligned} \tag{8}$$

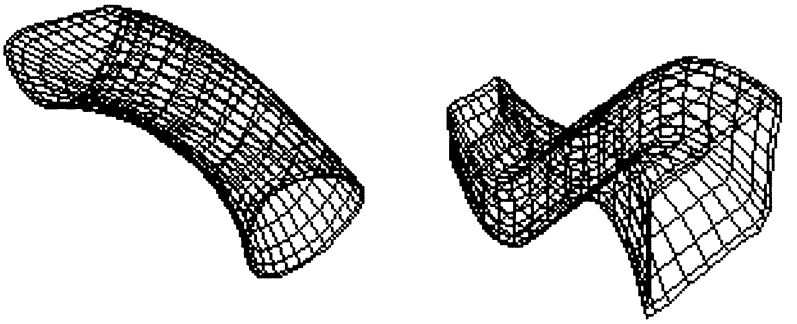

Figure 2: Two tube surface examples using up to fourth order harmonics.

Thus the only nonzero terms are $a_{x,0,0}$, $a_{x,0,l}$, $a_{x,m,0}$, $b_{x,m,0}$, $a_{x,m,l}$ and $b_{x,m,l}$ and the corresponding y and z terms. Tubes are an extension of generalized cylinders where the cross-section is no longer constrained to be planar. This allows for a wider range of shapes to be represented. All of the standard types of generalized cylinders can be represented in a Fourier representation as well. For example, the SHGC defined in Equation 2 can be represented by decomposing the cross-section function ($x(v)$ and $y(v)$) using the closed curve representation, and decomposing the scaling function ($r(u)$) and the spine ($z(u)$) using the open curve representation described above.

Tubes are useful for elongated hollow objects and elongated objects with flat ends. They are also useful for temporal sequences of planar images, where the third dimension is time, and multimodal images, where the third dimension is modality. In this case a simplified tube model would be used where the parameters governing the third dimension are fixed. Two example tube surfaces using this parametrization, with terms up to fourth order, is shown in Figure 2.

4.3 Closed Surfaces

Closed surfaces are the most difficult to represent because they are most dissimilar to tori. One way to represent closed surfaces is by considering tubes whose ends close up to a point at both ends instead of being open. This is done by expressing x and y using the following basis:

$$\begin{aligned} \phi_{\text{closed-xy}} &= \{1, \sin lv, \ldots, \\ &\quad \cos mu \sin lv, \sin mu \sin lv, \ldots \qquad (m = 1, 2, \ldots; l = 1, 2, \ldots)\} \end{aligned} \tag{9}$$

thus forcing both functions to constants at $v = 0, \pi, 2\pi$. This means that z must be expressed using only the cosines:

$$\phi_{\text{closed-z}} = \{1, \cos lv, \ldots, \qquad (l = 1, 2, \ldots)\} \tag{10}$$

This requires that the values for v be repeated as for a open curve but negated for x and y because they are both odd functions. The values for v are just repeated for z because it is an even function. That is:

$$\begin{aligned} x(u, v) &= -x(u, 2\pi - v) \\ y(u, v) &= -y(u, 2\pi - v) \\ z(u, v) &= z(u, 2\pi - v) \end{aligned} \tag{11}$$

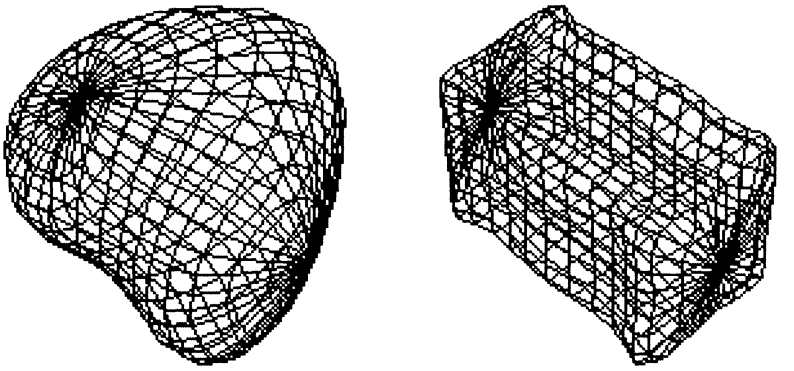

Figure 3: Two closed surface examples using up to fourth order harmonics on the left and eighth order on the right.

This representation is limited in that the axis along z is straight because $z(u,v) = z(v)$. Because the axis is aligned along z, an additional general rotation is necessary to allow for all orientations. Two example closed surfaces using this parametrization are shown in Figure 3, with terms up to fourth order on the left and eighth order on the right. Closed surfaces are useful for objects with no prominent openings. By using higher harmonics, more complex shapes can be represented.

5. BOUNDARY FINDING OBJECTIVE FUNCTION

In order to fit one of these models to the three-dimensional image data, a measure of fit is optimized by varying the model parameters. The surface is expected to be distinguishable by some measure of boundary strength (and direction when useful) computed from the image. The sum or integral of the boundary strength image over a given surface indicates the degree of correspondence between them, and this will be used as the measure of fit.

Any measure that indicates a change in some property that distinguishes the object from the background could be used as a boundary measure. A natural candidate for many images is the gray-level gradient. The magnitude is the strength of the boundary and the direction is the normal to the boundary. The gray level gradient can be calculated by first smoothing with a Gaussian to reduce the effect of noise. This is followed by a finite difference approximation to the partial derivatives in order to control smoothing independently. The smoothed boundary response will also help in the optimization by attracting the surface from further away. In this work, a $3 \times 3 \times 3$ Zucker-Hummel operator [32] is used. It is a centered finite difference, weighted by the distance from the point of computation.

5.1. Function

The sum or integral of the image gradient strength over a given model surface is a measure of the degree of correspondence between them:

$$M(b, \mathbf{p}) = \int\int_{\mathcal{A}} |b(x(\mathbf{p}, u, v), y(\mathbf{p}, u, v), z(\mathbf{p}, u, v))| dA \tag{12}$$

where b is the calculated gradient strength image and $\mathcal{A}$ is the surface. Also, $\mathbf{p}$ is the vector consisting of the parameters, $a_{x,m,l}$, $b_{x,m,l}$, $c_{x,m,l}$, $d_{x,m,l}$, $a_{y,m,l}$, $b_{y,m,l}$, $c_{y,m,l}$, $d_{y,m,l}$, $a_{z,m,l}$, $b_{z,m,l}$, $c_{z,m,l}$, $d_{z,m,l}$, for $m, l \leq K$. Although here we fix the highest order harmonic used, an iterative method for determining the best K using a trade-off between conciseness and fit could be devised.

To include direction, the integral of the inner product of the gradient vector with the normal to the model surface may be used:

$$M(\mathbf{b}, \mathbf{p}) = \iint_{\mathcal{A}} \mathbf{b}(x(\mathbf{p}, u, v), y(\mathbf{p}, u, v), z(\mathbf{p}, u, v)) \cdot \mathbf{n} dA \tag{13}$$

Based on experience with two-dimensional images, the direction information helps discount strong nearby edges in incompatible directions.

Equations 12 and 13 can be evaluated by numerical integration. The boundary strength array, $|b(x, y, z)|$ can be evaluated at each point on the surface using trilinear interpolation. The area element on the surface $\mathcal{A}$ is dA. This is given by:

$$dA = \left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right| du dv \tag{14}$$

5.2. Gradient

The gradient of the objective is necessary for optimization. The derivative of the objective in Equation 12 with respect to the parameters governing x is:

$$\frac{\partial M}{\partial \mathbf{p}_x} = \iint_{\mathcal{A}} \left[|b(x, y, z)| \frac{\partial}{\partial \mathbf{p}_x} \left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right| + \frac{\partial |b(x, y, z)|}{\partial x} \frac{\partial x(\mathbf{p}, u, v)}{\partial \mathbf{p}_x} \left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right| \right] du dv \tag{15}$$

and similarly for y and z. This expression can also be evaluated by numerical integration. The expressions $\frac{\partial |b(x,y,z)|}{\partial x}$, $\frac{\partial |b(x,y,z)|}{\partial y}$, $\frac{\partial |b(x,y,z)|}{\partial z}$ can be determined by discrete derivative calculations at each (x, y, z) point on the surface, again using trilinear interpolation.

The expressions $\frac{\partial x(\mathbf{p}_x, u, v)}{\partial \mathbf{p}_x}$, $\frac{\partial y(\mathbf{p}_y, u, v)}{\partial \mathbf{p}_y}$ and $\frac{\partial z(\mathbf{p}_z, u, v)}{\partial \mathbf{p}_z}$ can be calculated from the expressions for $x(\mathbf{p}, u, v)$, $y(\mathbf{p}, u, v)$, and $z(\mathbf{p}, u, v)$ (shown in Equation). Thus,

$$\begin{aligned}
\frac{\partial x(\mathbf{p}, u, v)}{\partial a_{x,m,l}} &= \lambda_{m,l} \cos 2\pi mu \cos 2\pi lv \qquad \frac{\partial x(\mathbf{p}, u, v)}{\partial b_{x,m,l}} = \lambda_{m,l} \sin 2\pi mu \cos 2\pi lv \\
\frac{\partial x(\mathbf{p}, u, v)}{\partial c_{x,m,l}} &= \lambda_{m,l} \cos 2\pi mu \sin 2\pi lv \qquad \frac{\partial x(\mathbf{p}, u, v)}{\partial d_{x,m,l}} = \lambda_{m,l} \sin 2\pi mu \sin 2\pi lv
\end{aligned} \tag{16}$$

and similarly for y and z.

The partials $\frac{\partial \mathbf{x}}{\partial u}$ and $\frac{\partial \mathbf{x}}{\partial v}$ can be evaluated analytically. For example:

$$\begin{aligned}
\frac{\partial x}{\partial u} = \sum_{m=0}^{2K} \sum_{l=0}^{2K} \lambda_{m,l} \; [\; & -a_{m,l} 2\pi m \sin 2\pi mu \cos 2\pi lv + b_{m,l} 2\pi m \cos 2\pi mu \cos 2\pi lv \\
& -c_{m,l} 2\pi m \sin 2\pi mu \sin 2\pi lv + d_{m,l} 2\pi m \cos 2\pi mu \sin 2\pi lv]
\end{aligned} \tag{17}$$

and similarly for y and z. The expression $\frac{\partial}{\partial \mathbf{p}_x} \left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right|$ can be calculated from Equation 14:

$$\begin{aligned}
\left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right| &= \sqrt{\left(\frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right)^2} \\
&= \sqrt{\left(\frac{dx}{dv}\frac{dy}{du} - \frac{dy}{dv}\frac{dx}{du} \right)^2 + \left(\frac{dz}{dv}\frac{dx}{du} - \frac{dx}{dv}\frac{dz}{du} \right)^2 + \left(\frac{dy}{dv}\frac{dz}{du} - \frac{dz}{dv}\frac{dy}{du} \right)^2} \\
&= \sqrt{(w_z^2 + w_y^2 + w_x^2)}
\end{aligned} \tag{18}$$

Then:

$$\begin{aligned}
\frac{\partial}{\partial \mathbf{p}_x}\left(\left|\frac{\partial \mathbf{x}}{\partial u}\times\frac{\partial \mathbf{x}}{\partial v}\right|\right) &= \frac{1}{2\left|\frac{\partial \mathbf{x}}{\partial u}\times\frac{\partial \mathbf{x}}{\partial v}\right|}\frac{\partial}{\partial \mathbf{p}_x}(w_z^2+w_y^2+w_x^2) \\
&= \frac{w_z}{\left|\frac{\partial \mathbf{x}}{\partial u}\times\frac{\partial \mathbf{x}}{\partial v}\right|}\left(\frac{dy}{du}\frac{\partial}{\partial \mathbf{p}_x}\left(\frac{dx}{dv}\right)-\frac{dy}{dv}\frac{\partial}{\partial \mathbf{p}_x}\left(\frac{dx}{du}\right)\right)+ \\
&\quad \frac{w_y}{\left|\frac{\partial \mathbf{x}}{\partial u}\times\frac{\partial \mathbf{x}}{\partial v}\right|}\left(\frac{dz}{dv}\frac{\partial}{\partial \mathbf{p}_x}\left(\frac{dx}{du}\right)-\frac{dz}{du}\frac{\partial}{\partial \mathbf{p}_x}\left(\frac{dx}{dv}\right)\right) \qquad (19)
\end{aligned}$$

Expressions such as $\frac{\partial}{\partial \mathbf{p}_x}(\frac{\partial x}{\partial u})$ which can be evaluated using:

$$\begin{aligned}
\frac{\partial}{\partial a_{x,m,l}}\left(\frac{\partial x}{\partial u}\right) &= -m\sin 2\pi mu\cos 2\pi lv \qquad \frac{\partial}{\partial b_{x,m,l}}\left(\frac{\partial x}{\partial u}\right) = m\cos 2\pi mu\cos 2\pi lv \\
\frac{\partial}{\partial c_{x,m,l}}\left(\frac{\partial x}{\partial u}\right) &= -m\sin 2\pi mu\sin 2\pi lv \qquad \frac{\partial}{\partial d_{x,m,l}}\left(\frac{\partial x}{\partial u}\right) = m\cos 2\pi mu\sin 2\pi lv \qquad (20)
\end{aligned}$$

All of the above follows similarly for $\frac{\partial}{\partial \mathbf{p}_y}$ and $\frac{\partial}{\partial \mathbf{p}_z}$.

6. OPTIMIZATION

The problem to be solved is that of maximizing the objective function $M(\mathbf{p})$. The objective function we are solving is not in general convex, but depends ultimately on the gray-level surface shape of the image. If the starting point of the optimization is good enough, the global optimum can be found by a local optimization. Thus, an initial position for the surface must be supplied by the user or some initial processing step. Continuous gradient ascent [33] was used to optimize the objective function. This method takes small steps in the direction of the gradient (the direction of greatest increase) until an optimal point is found.

7. EXPERIMENTS

The surface finding system was evaluated using synthetic and medical three-dimensional images. Initial results of the surface finding method using gray-level gradient magnitude (Equation 12) applied to 3D synthetic and medical images are encouraging. The example shown in Figure 4 is a simple synthetic image of an open surface with Gaussian noise added (SNR = 2.5). An initial surface (a disk) was positioned roughly at the target object. The final optimized surface (using up to fourth order terms) matches well with the target. The results of the method applied to the problem of delineating the upper portion of the cerebrum of the human brain from magnetic resonance images are shown in Figure 5. The surface was matched to the gradient magnitude calculated from the image. The final boundary, using up to eighth order parameters, succeeds in delineating the structure properly. In Figure 6, a three-dimensional cardiac image of a dog's heart from the Dynamic Spatial Reconstructor (DSR) is analyzed. The DSR is a dynamic, three-dimensional imaging device based on high-speed x-ray computed tomography capable of imaging the moving heart [34]. As before, the surface, using up to eighth order parameters, was matched to the gradient magnitude calculated from the image. The endocardial (inner) wall of the left ventricle is successfully delineated.

8. SUMMARY

This work presents a general surface finding system for three-dimensional images of smoothly deformable objects. The goal of this work was to incorporate an expressive parametrized model of global shape into surface finding. In addition, a new global shape parametrization for surfaces useful as a representation for computer vision and modeling has been described. This parametrization extends the expressibility of previous parametrizations. Future work includes the application to image sequences and the incorporation of other image features and constraints.

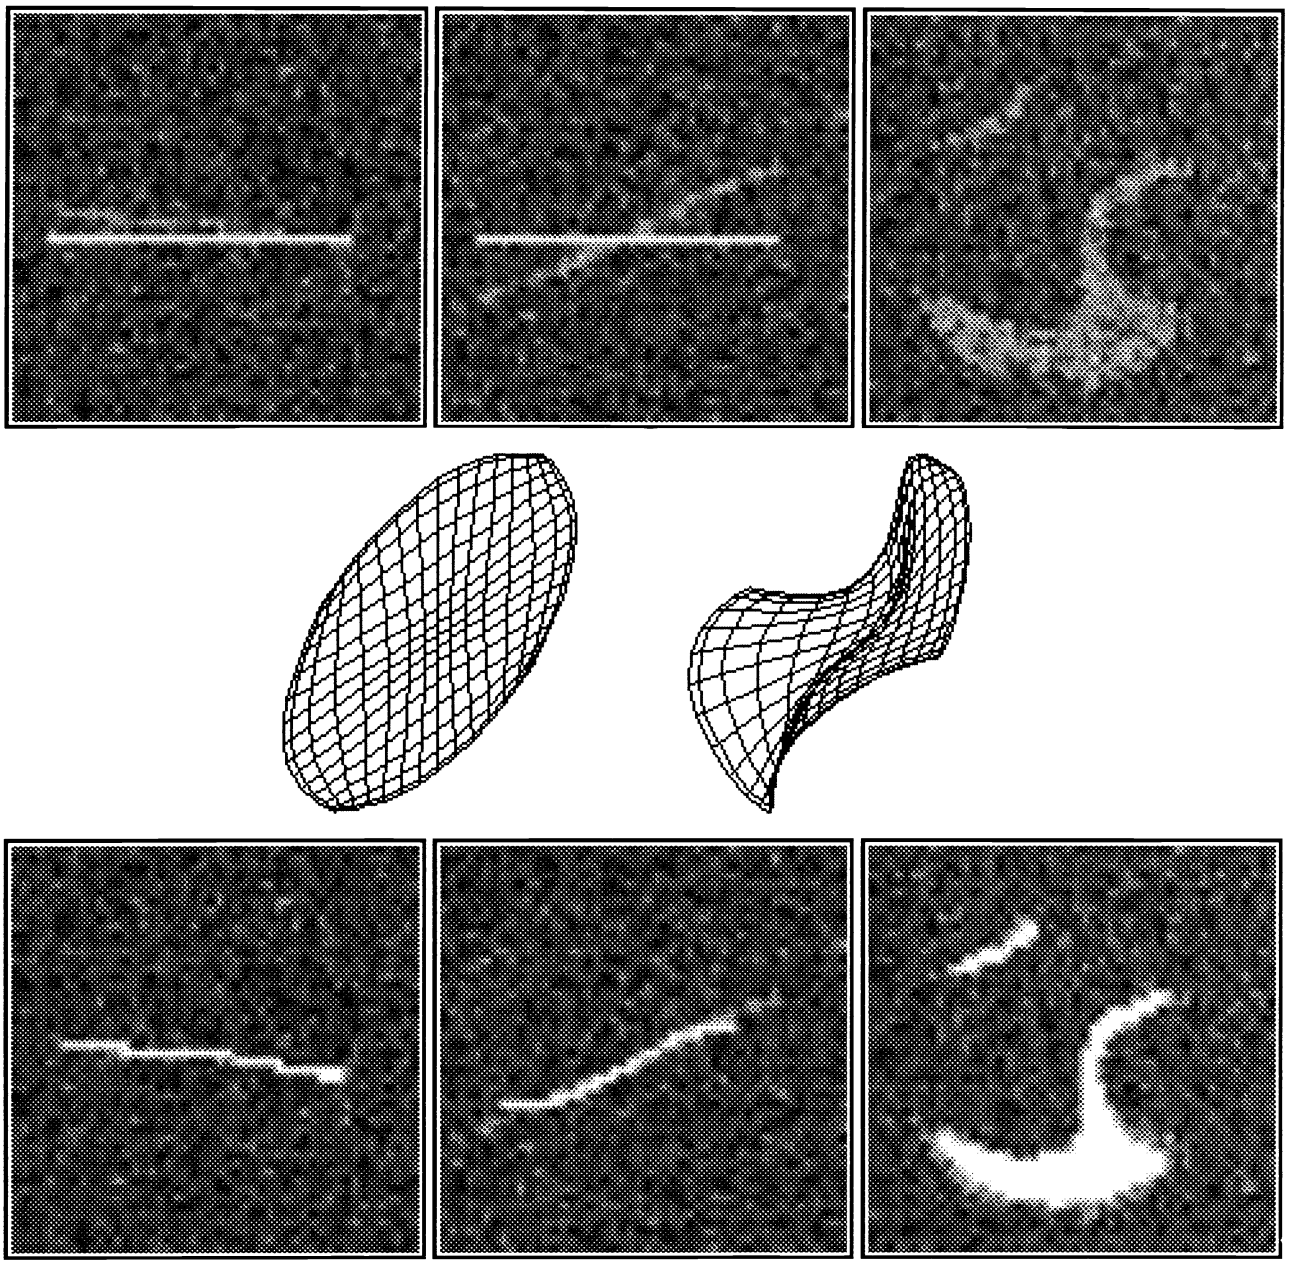

Figure 4: Synthetic image example. Top: Three perpendicular slices through the 3D image ($48 \times 48 \times 48$) are shown with the initial surface. Middle: Wire frame of initial (left) and final (right) surface. Bottom: The same slices shown with the final surface.

References

[1] R. Mohan and R. Nevatia. Segmentation and description based on perceptual organization. In *Proc. IEEE Conf. Computer Vision Pattern Recognition*, pages 333–341, June 1989.

[2] P. T. Sander and S. W. Zucker. Inferring surface trace and differential structure from 3-D images. *IEEE Trans. Pattern Anal. Machine Intell.*, 12(9):833–854, September 1990.

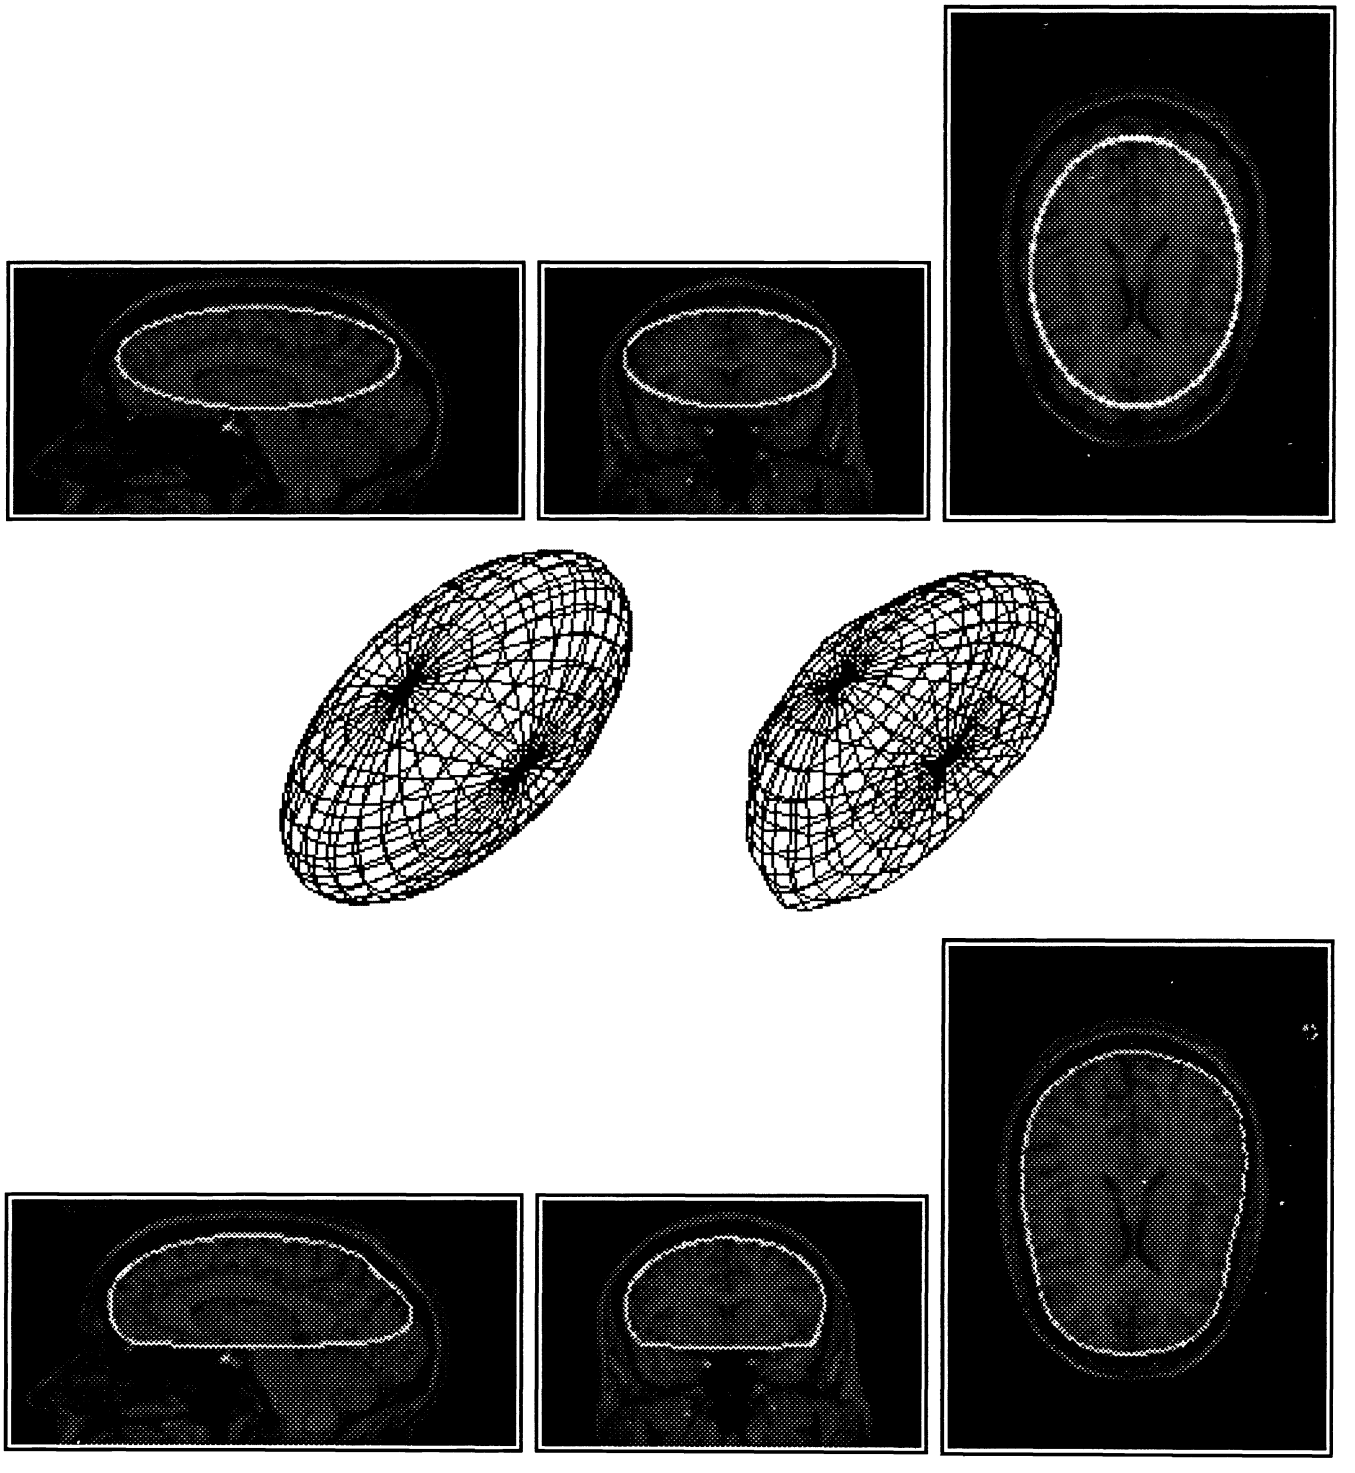

Figure 5: Magnetic resonance brain image example. Top: Three perpendicular slices through the three-dimensional image ($120 \times 160 \times 78$) are shown with the initial surface. Middle: Wire frame of initial (left) and final (right) surface. Bottom: The same slices shown with final surface indicating the upper portion of the cerebrum.

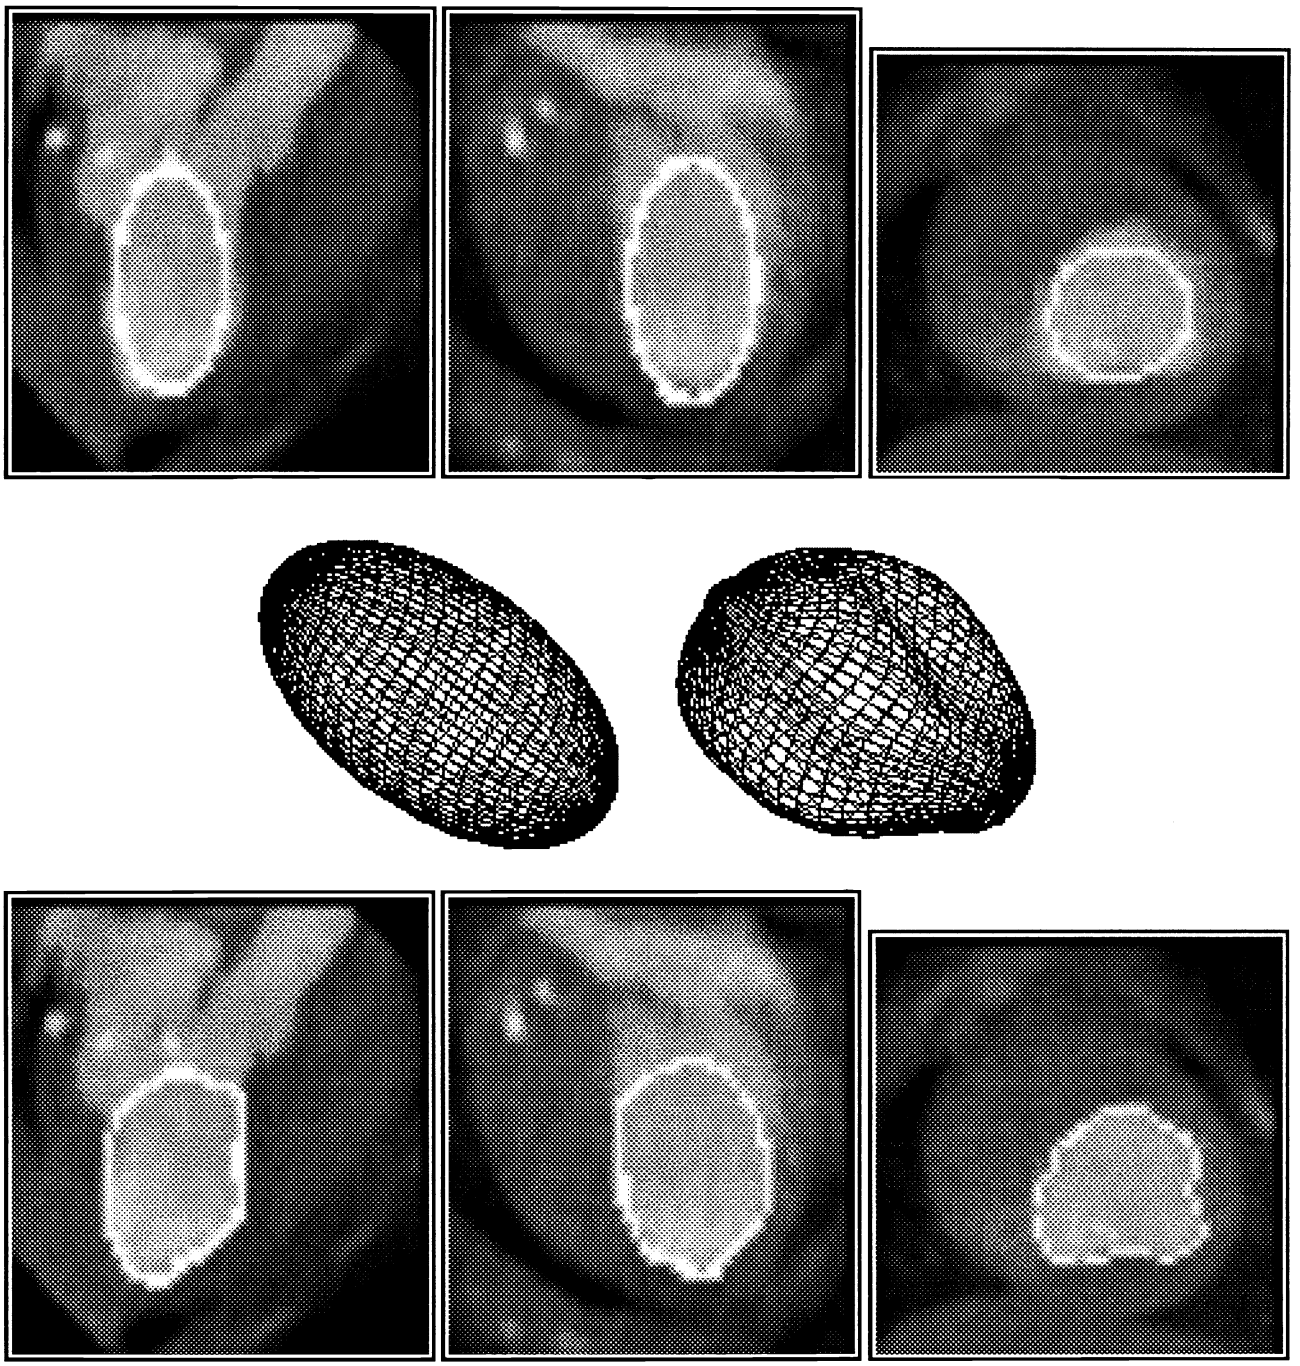

Figure 6: Dynamic Spatial Reconstructor (DSR) cardiac image example. Top: Three perpendicular slices through the three-dimensional image (subsampled to $49 \times 50 \times 55$) are shown with the initial surface. Middle: Wire frame of initial (left) and final (right) surface. Bottom: The same slices shown with final surface at the endocardium.

[3] F. Bergholm. Edge focusing. *IEEE Trans. Pattern Anal. Machine Intell.*, 9(6):726–741, 1987.

[4] A. Martelli. An application of heuristic search methods to edge and contour detection. *Communications of the ACM*, 19(2):73–83, February 1976.

[5] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry-seeking models and 3D object reconstruction. *Int. J. Computer Vision*, 1:211–221, 1987.

[6] J. M. Gauch and S. M. Pizer. Image description via the multiresolution intensity axis of symmetry. In *Proc. Second Int. Conf. Computer Vision*, pages 269–274, 1988.

[7] H. Delingette, M. Hebert, and K. Ikeuchi. Shape representation and image segmentation using deformable surfaces. In *Proc. IEEE Conf. Computer Vision Pattern Recognition*, pages 467–472, June 1991.

[8] Y. F. Wang and J. F. Wang. Surface reconstruction using deformable models with interior and boundary constraints. In *Proc. Third Conf. Computer Vision*, pages 300–303, December 1990.

[9] I. Cohen, L. D. Cohen, and N. Ayache. Introducing deformable surfaces to segment 3D images and infer differential structures. Technical Report 1403, INRIA-Rocquencourt, March 1991.

[10] I. Cohen, L. D. Cohen, and N. Ayache. Introducing new deformable surfaces to segment 3D images. In *Proc. IEEE Conf. Computer Vision Pattern Recognition*, pages 738–739, 1991.

[11] B. Widrow. The "rubber mask" technique. *Pattern Recognition*, 5:175–211, 1973.

[12] A. L. Yuille, D. S. Cohen, and P. W. Hallinan. Feature extraction from faces using deformable templates. In *Proc. IEEE Conf. Computer Vision Pattern Recognition*, pages 104–109, 1989.

[13] P. Lipson, A. L. Yuille, D. O'Keeffe, J. Cavanaugh, J. Taaffe, and D. Rosenthal. Automated bone density calculation using feature extraction by deformable templates. In *Proc. First Conf. Visualization in Biomedical Computing*, pages 477–484, Los Alamitos, California, 1990. IEEE Computer Society.

[14] A. P. Pentland. Perceptual organization and the representation of natural form. *Artificial Intelligence*, 28(3):293–331, 1986.

[15] M. Bomans, K. Hohne, U. Tiede, and M. Riemer. 3-D segmentation of MR images of the head for 3-D display. *IEEE Trans. Medical Imaging*, 9(2):177–183, June 1990.

[16] H. E. Cline, W. E. Lorensen, R. Kikinis, and F. Jolesz. Three-dimensional segmentation of MR images of the head using probability and connectivity. *J. Comput. Assist. Tomogr.*, 14(6):1037–1045, Nov/Dec 1990.

[17] D. N. Kennedy, P. A. Filipek, and V. S. Caviness. Anatomic segmentation and volumetric calculations in nuclear magnetic resonance imaging. *IEEE Trans. Medical Imaging*, 8(1):1–7, March 1989.

[18] K. H. Hohne and W. A. Hanson. Interactive 3D segmentation of MRI and CT volumes using morphological operations. *J. Comput. Assist. Tomogr.*, 16(2):285–294, Mar/Apr 1992.

[19] P. J. Besl. Geometric modeling and computer vision. *Proc. of the IEEE*, 76(8):936–958, 1988.

[20] F. Solina and R. Bajcsy. Recovery of parametric models from range images: The case for superquadrics with global deformations. *IEEE Trans. Pattern Anal. Machine Intell.*, 12(2):131–147, February 1990.

[21] D. H. Ballard and C. M. Brown. *Computer Vision.* Prentice-Hall, Englewood Cliffs, 1982.

[22] R. B. Schudy. Harmonic surfaces and parametric image operators: Their use in locating the moving endocardial surface from three-dimensional cardiac ultrasound data. Computer Science Technical Report 112, University of Rochester, Rochester, New York, March 1981.

[23] J. Ponce. Straight homogeneous generalized cylinders: Differential geometry and uniqueness results. *Int. J. Computer Vision*, 4(1):79–100, 1990.

[24] K. Rao and R. Nevatia. Computing volume descriptions from sparse 3-D data. *Int. J. Computer Vision*, 2(1):33–50, 1988.

[25] A. H. Barr. Global and local deformations of solid primitives. *Computer Graphics*, 18(3):21–30, July 1984.

[26] A. P. Pentland. Automatic extraction of deformable part models. *Int. J. Computer Vision*, 4:107–126, 1990.

[27] A. J. Hanson. Hyperquadrics: Smoothly deformable shapes with convex polyhedral bounds. *Comput. Vision Graphics Image Process.*, 44:191–210, 1988.

[28] E. Persoon and K. S. Fu. Shape discrimination using Fourier descriptors. *IEEE Trans. Pattern Anal. Machine Intell.*, 8(3):388–397, May 1986.

[29] L. H. Staib and J. S. Duncan. Parametrically deformable contour models. In *Proc. IEEE Conf. Computer Vision Pattern Recognition*, pages 98–103, June 1989.

[30] L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformable models. *IEEE Trans. Pattern Anal. Machine Intell.*, 1992. (to appear).

[31] G. P. Tolstov. *Fourier Series.* Prentice-Hall, Englewood Cliffs, 1962.

[32] S. W. Zucker and R. A. Hummel. A three-dimensional edge operator. *IEEE Trans. Pattern Anal. Machine Intell.*, 3(3):324–331, May 1981.

[33] D. A. Pierre. *Optimization Theory with Applications.* Dover, New York, 1986.

[34] R. A. Robb. High speed three-dimensional x-ray computed tomography: The dynamic spatial reconstructor. *Proc. of the IEEE*, 71:308–319, 1983.