

Chapter 7

Boundary Finding with Parametrically Deformable Models

Lawrence H. Staib and James S. Duncan

7.1 Introduction

This work describes an approach to finding objects in images based on deformable shape models. Boundary finding in two and three dimensional images is enhanced both by considering the bounding contour or surface as a whole and by using model-based shape information.

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. 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 bounding curve or surface can be profitably considered as a whole, rather than as curve or surface segments, because it tends to result in a more consistent solution overall.

These models are best suited for objects whose diversity and irregularity of shape make them poorly represented in terms of fixed features or parts. 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.

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 is formulated as a optimization problem using parametric Fourier models which are developed for both curves and surfaces. The model is matched to the image by optimizing in the parameter space the match between the model and a boundary measure applied to the image. Probability distributions on the parameters of the representation can be incorporated to bias the model to a particular overall shape while allowing for deformations. This leads to a maximum *a posteriori* objective function.

7.2 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.

An alternative method for boundary analysis is the Hough transform [5]. The Hough approach is similar to the current method in that it finds shapes by looking for maxima in a parameter space. However, the storage and computational complexity of the Hough method are a great disadvantage, especially if deformations are envisaged.

Other investigators have considered whole-boundary methods that adjust a tentative curve or surface mesh in order to match to the image. By considering the boundary as a whole, a structure is imposed on the problem that bridges gaps and results in overall consistency.

For curve finding, Gritton and Parrish [6] used a flexible bead chain, where the beads are putative boundary points. The beads are attracted towards pixels that have a higher gradient magnitude. Cooper [7] formulated boundary estimation using maximum likelihood. A boundary adjustment scheme similar to the bead chain algorithm [6] is presented to perform the optimization. Kass *et al.* [8] used energy-minimizing snakes that are attracted to image features such as lines and edges while internal spline forces impose a smoothness constraint. The weights of the smoothness and image force terms in the energy functional can be adjusted for different behavior. The solution is found using variational methods.

For surface finding, Terzopoulos *et al.* [9] 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 [10], surfaces in range images [11, 12] and surfaces in three-dimensional images [13].

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 [14] used parametrized templates called rubber masks to model objects. The parameters are sizes and relationships between subparts. Yuille *et al.* [15] 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. Work has also been done developing deformable templates based on Markov models of two-dimensional boundaries incorporating knowledge of shape from statistical features [16]. In the next section we will discuss parametrizations for surfaces in more detail.

Pentland and his group have developed a physically-based method for analyzing shape [17, 18]. 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.

7.3 Curve and 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 curve can be represented explicitly by two functions of one parameter: $x(s)$ and $y(s)$. A 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 curves and surfaces can be represented as a single function. For example, curves expressible as a single function of one parameter, $r(\theta)$, are radial deformations of a circle. Similarly, surfaces expressible as a function of two angles, $r(\theta, \phi)$, are radial deformations of a sphere and are parametrized by (θ, ϕ) . 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.

The main approaches to parametric modeling in computer vision have been polynomials [19], superquadrics [17, 20], spherical harmonics [5, 21] and generalized cylinders [22]. All of these parametrizations are restricted to a limited class of objects.

7.3.1 Polynomials

Second degree algebraic surfaces have been used extensively because of their simplicity and conciseness. Conics are second degree curves including ellipses, parabolas and hyperbolas. Quadrics are second degree surfaces which include spheres, ellipsoids, cones, cylinders, planes, paraboloids and hyperboloids. Their conciseness, however, greatly limits their expressiveness. Higher order polynomial surfaces are expressed using implicit representations.

7.3.2 Superquadrics

Superquadrics are an extension of quadrics using an exponent that allows the shape to vary from an ellipsoid to a rectangular parallelepiped. The two-dimensional analog is the superellipse. Superquadrics can be expressed parametrically by:

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

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 [23], 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 [9] and strain modes [17] in order to increase their expressiveness. Hyperquadrics [24] are a generalization of superquadrics that allow smooth deformations from shapes with convex polyhedral bounds, although no explicit parametrized form is possible.

7.3.3 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 [25]:

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

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.

7.3.4 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:

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

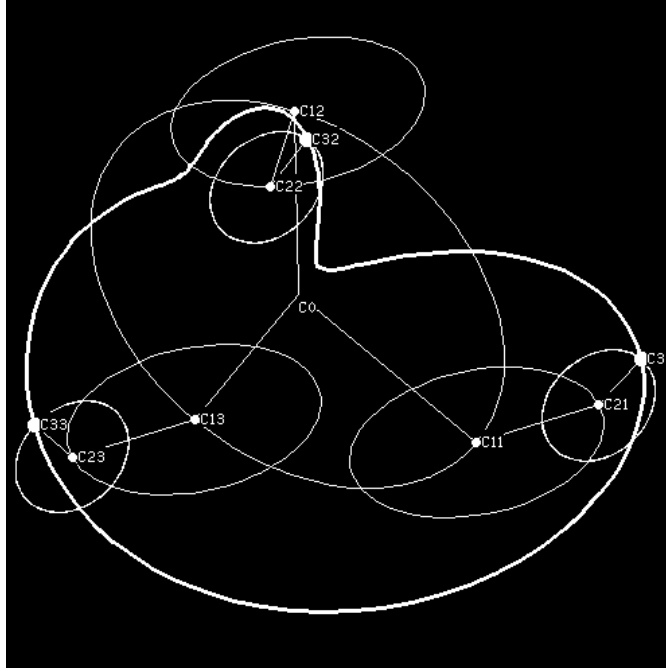


Figure 1: The contour (dark line) at the left is constructed from three component ellipses shown at three different times.

where $P(m,n,x)$ is the n th derivative of the m th 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.

7.4 Fourier Models

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:

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

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 [26], 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.

7.5 Fourier Curves

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

$$\phi = \frac{1}{2\pi}, \frac{\cos x}{\pi}, \frac{\sin x}{\pi}, \frac{\cos 2x}{\pi}, \frac{\sin 2x}{\pi}, \dots \quad (5)$$

If we write the resulting equations in matrix form, we get the elliptic Fourier representation [27], [28], [29]:

$$\begin{bmatrix} x(t) \\ y(t) \end{bmatrix} = \begin{bmatrix} a_0 \\ c_0 \end{bmatrix} + \sum_{k=1}^{\infty} \begin{bmatrix} a_k & b_k \\ c_k & d_k \end{bmatrix} \begin{bmatrix} \cos kt \\ \sin kt \end{bmatrix} \quad (6)$$

where:

$$\begin{aligned} a_0 &= \frac{1}{2\pi} \int_0^{2\pi} x(t) dt & c_0 &= \frac{1}{2\pi} \int_0^{2\pi} y(t) dt \\ a_k &= \frac{1}{\pi} \int_0^{2\pi} x(t) \cos kt dt & b_k &= \frac{1}{\pi} \int_0^{2\pi} x(t) \sin kt dt \\ c_k &= \frac{1}{\pi} \int_0^{2\pi} y(t) \cos kt dt & d_k &= \frac{1}{\pi} \int_0^{2\pi} y(t) \sin kt dt \end{aligned}$$

The closed curve is thus represented by $\mathbf{p}_{\text{raw}} = (a_0, c_0, a_1, b_1, c_1, d_1, \dots)$ which will be referred to as the raw parameter vector. This particular version of Fourier boundary representation has a number of advantages. A geometric interpretation, in terms of ellipses, can be developed from this decomposition. The geometric interpretation will allow for better visualization of the effect of the parameters and invariance to starting point, scale and two-dimensional rotation and translation. Invariance to rotation, scale and translation is important because these parameters are determined not by the object but by the view of the object, which often cannot be held constant.

In Equation 6, the first two coefficients, a_0 and c_0 , determine the overall translation of the shape. Each term in the summation is the parametric form for an ellipse. In the degenerate case $a_k d_k - b_k c_k = 0$ and the parametric form defines

a straight line (a degenerate ellipse). In each term, the matrix determines the characteristics of the ellipse. The contour can be viewed as being decomposed into a sum of rotating phasors, each individually defining an ellipse, and rotating with a speed proportional to their harmonic number, k . This can be seen in Figure 1 where a contour is shown constructed from three component ellipses forming a sort of planetary system. The straight lines represent the phasors for each ellipse shown at three different times. Thus, the point C_{ij} traces out the i th ellipse at time j . Each point is the center of the next higher ellipse. C_0 is the center of the first ellipse. Points C_{31} , C_{32} and C_{33} are three different points on the final curve.

It is important that the curve representation that is decomposed into Fourier components be both continuous and periodic. Discontinuities slow the convergence because of the high frequencies inherent in a step jump. In this representation, both $x(t)$ and $y(t)$ are periodic because the contour is closed, and both $x(t)$ and $y(t)$ are continuous because the contour is continuous.

The geometric properties of each of the component ellipses can be derived from the raw elements of each ellipse matrix. Each ellipse can be described by four geometric properties: semi-major axis length, semi-minor axis length, rotation and phase shift. The rotation is the angle from the x -axis to the major axis of the ellipse, defined from $-\pi/2$ to $\pi/2$. The phase shift is the difference in phase from the major axis to the position of $t = 0$ (the ellipse starting position), defined from $-\pi$ to π .

These ellipse properties can be derived as follows. First consider the general form for an ellipse, which is the product of the raw ellipse matrix and the trigonometric basis function vector:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \cos kt \\ \sin kt \end{bmatrix} \quad (7)$$

In order to determine the ellipse parameters, consider the matrix for an ellipse with its major axis aligned with the x -axis and with no phase shift where A and B are the major and minor semi-axis lengths, respectively. The phasor moves counterclockwise for B positive, clockwise for B negative. The ellipse can be rotated simply by pre-multiplying the ellipse matrix by a rotation matrix. A phase shift of the ellipse by ϕ_0 means replacing t by $t + \phi_0$. This is the same as a pre-multiplication of the basis function vector by a rotation matrix, or equivalently, a post-multiplication of the ellipse matrix. Thus, a rotation of this ellipse by θ and shift by ϕ can be written as a pre-multiplication and a post-multiplication by rotation matrices:

$$\begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix} \begin{bmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{bmatrix} \quad (8)$$

This represents a general ellipse and is thus equivalent to the raw ellipse matrix in Equation 7. Therefore, to find the ellipse parameters given the values of these matrix elements, solve the following four equations that come from identifying corresponding matrix elements for A , B , θ and ϕ .

$$\begin{aligned} a &= +A \cos \theta \cos \phi - B \sin \theta \sin \phi & b &= -A \cos \theta \sin \phi - B \sin \theta \cos \phi \\ c &= +A \sin \theta \cos \phi + B \cos \theta \sin \phi & d &= -A \sin \theta \sin \phi + B \cos \theta \cos \phi \end{aligned} \quad (9)$$

This results in:

$$\begin{aligned}
A^2 &= \frac{\alpha + \sqrt{\alpha^2 - 4\beta^2}}{2} & B^2 &= \frac{2\beta^2}{\alpha + \sqrt{\alpha^2 - 4\beta^2}} \\
\theta &= \tan^{-1} \frac{Ac + Bb}{Aa - Bd} & \phi &= \tan^{-1} \frac{Ba - Ad}{Ac + Bb}
\end{aligned} \tag{10}$$

where:

$$\alpha = a^2 + b^2 + c^2 + d^2, \quad \beta = ad - bc$$

By taking A to be positive and B to agree in sign with j , we get a consistent sign convention. These parameters, $\mathbf{p}_{\text{ref}} = (a_0, c_0, A_1, B_1, \theta_1, \phi_1, \dots)$, represent the shape in terms of the ellipse properties and will be referred to as the refined parameters.

A further conversion can improve this set by converting the rotation and shift parameters from absolute quantities to values relative to the preceding harmonic and by normalizing the axes' lengths [30]. This conversion to relative quantities will allow the isolation of an overall rotation parameter and the removal of the overall phase shift, ϕ_1 , which is arbitrary. Normalizing the axes' lengths creates an overall scale parameter.

Open Curves

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 $y(t) = y(2\pi - t)$ [26]. The resulting functions are even and thus they can be represented by the even sinusoidal basis functions

$$\phi_{\text{even}} = \left\{ \frac{1}{2\pi}, \frac{\cos x}{\pi}, \frac{\cos 2x}{\pi}, \frac{\cos 3x}{\pi}, \dots \right\} \tag{11}$$

This representation can be thought of as decomposing the curve into degenerate ellipses (flattened down to two coincident lines). The equations for the corresponding ellipse parameters are simplified because the sine terms, b_k and d_k , are zero:

$$A^2 = a^2 + c^2 \quad B^2 = 0 \quad \theta = \tan^{-1} \frac{c}{a} \quad \phi = 0 \tag{12}$$

The ellipses are all degenerate with a fixed starting point at one end, thus forcing both the minor semi-axis length, B , and the starting point, ϕ , to be zero.

7.5 Fourier Surfaces

In order to represent surfaces, a function of two variables is needed. Because the parity of the functions will be important, a useful two-dimensional basis is [31]:

$$\begin{aligned}
\phi &= \{1, \cos 2\pi mu, \sin 2\pi mu, \cos 2\pi lv, \sin 2\pi lv, \dots, \\
&\quad \cos 2\pi mu \cos 2\pi lv, \sin 2\pi mu \cos 2\pi lv, \\
&\quad \cos 2\pi mu \sin 2\pi lv, \sin 2\pi mu \sin 2\pi lv, \dots \quad (m = 1, 2, \dots; l = 1, 2, \dots) \}
\end{aligned} \tag{13}$$




Figure 2: An example torus surface (left) using up to second order harmonics and an example open surface (right) using up to fourth.

The function is then represented by:

$$f(u, v) = \sum_{m=0}^{2K} \sum_{l=0}^{2K} \lambda_{m,l} [a_{m,l} \cos 2\pi mu \cos 2\pi lv + b_{m,l} \sin 2\pi mu \cos 2\pi lv + c_{m,l} \cos 2\pi mu \sin 2\pi lv + d_{m,l} \sin 2\pi mu \sin 2\pi lv] \quad (14)$$

where:

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

This allows the specification of even functions (using the cosine terms) and odd functions (using the sine terms). The complex basis is useful for computational purposes because the parameters can be computed in a single transform:

$$\phi = \{1, e^{2\pi i(mu+lv)}, \dots \quad (m = \pm 1, \pm 2, \dots; l = \pm 1, \pm 2, \dots) \} \quad (15)$$

Using Euler's formula, $e^{ix} = \cos x + i \sin x$, we can derive the conversion between the sine-cosine basis parameters and the complex basis parameters.

The bases presented can be used for parametrizing surfaces in three dimensions. Such surfaces can be described explicitly by three functions of two surface parameters:

$$\mathbf{x}(u, v) = (x(u, v), y(u, v), z(u, v)) \quad (16)$$

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. There are three corresponding sets of parameters: $a_x, b_x, c_x, d_x, a_y, b_y, c_y, d_y, a_z, b_z, c_z, d_z$. 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 [10] may provide a way of determining the overall structure on which to base the surface parametrization.

There are four classes of simple surfaces in three dimensions that 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 shown in Equation 14. The result is a torus because both surface parameters are forced to be periodic. An example torus surfaces using this parametrization is shown in Figure 2. 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.

7.5.1 Open Surfaces

Representing open surfaces with the basis in Equation 14 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 .

$$\begin{aligned} \phi_{\text{open}} = \{ & 1, \cos mu, \cos lv, \dots, \\ & \cos mu \cos lv, \dots \quad (m = 1, 2, \dots; l = 1, 2, \dots) \} \end{aligned} \quad (17)$$

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 example open surfaces using this parametrization is shown in Figure 2.

7.5.2 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 :

$$\begin{aligned} \phi_{\text{tube}} = \{ & 1, \cos lv, \sin mu, \cos mu, \dots, \\ & \cos mu \cos lv, \sin mu \cos lv, \dots \quad (m = 1, 2, \dots; l = 1, 2, \dots) \} \end{aligned} \quad (18)$$

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




Figure 3: Two tube surface examples using up to fourth order harmonics.

dimension is time, and multimodal images, where the third dimension is modality. In this case a simplified tube model would be used where the third dimension was independent of the two surface parameters, for example, $z(u, v) = t$. Two example tube surfaces using this parametrization are shown in Figure 3.

7.5.3 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:

$$\begin{aligned} \phi_{\text{closed-xy}} = \{ & 1, \sin lv, \dots, & (19) \\ & \cos mu \sin lv, \sin mu \sin lv, \dots & (m = 1, 2, \dots; l = 1, 2, \dots) \} \end{aligned}$$

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

$$\phi_{\text{closed-z}} = \{1, \cos lv, \dots, \quad (l = 1, 2, \dots)\} \quad (20)$$

This requires that the values for v be repeated as for an 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:

$$\begin{aligned} x(u, v) &= -x(u, 2\pi - v) \\ y(u, v) &= -y(u, 2\pi - v) \\ z(u, v) &= z(u, 2\pi - v) \end{aligned} \quad (21)$$

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 4, with terms up to fourth order on the left and eighth order on the right.




Figure 4: Two closed surface examples using up to fourth order harmonics on the left and eighth order on the right.

7.5.4 Surface geometry

The Fourier surface description makes the calculation of geometric surface properties straightforward because a continuous description of the surface is known. Without an analytic description of the surface, curvature can be calculated based on the computation of derivatives from a local surface patch fit, or from a discrete approximation of the derivatives at each point. These methods are dependent on the proper choice of the size of the patch or neighborhood. For Fourier surfaces, partial derivatives of the surface functions can be calculated from the functional description. Curvature is then calculated directly from these partial derivatives [19]. Surface curvature properties have been used to classify and characterize shape. For example, surface regions can be classified by the sign of the surface curvatures as peaks, ridges, saddles, valleys, pits and flats.

7.6 Boundary Finding Objective Function

In order to fit one of these models to the 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 (direction can also be used) 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 can 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. For two-dimensional images, 2×2 or 3×3 finite differences are used. For three-dimensional images, $2 \times 2 \times 2$ or $3 \times 3 \times 3$ finite differences are used.

The measure of fit for curves can be written as follows, here using only boundary strength:

$$M(b, \mathbf{p}) = \int_0^S |b(x(\mathbf{p}, s), y(\mathbf{p}, s), z(\mathbf{p}, s))| ds \quad (22)$$

where \mathbf{p} is the vector consisting of the basis function parameters. Although this implies fixing the highest order harmonic used, an iterative method for determining the best K using a trade-off between conciseness and fit could be devised. The equivalent measure for surfaces is:

$$M(b, \mathbf{p}) = \iint_{\mathcal{A}} |b(x(\mathbf{p}, u, v), y(\mathbf{p}, u, v), z(\mathbf{p}, u, v))| dA \quad (23)$$

Equation 23 can be evaluated by numerical integration. The boundary strength array, $|b|$ can be evaluated at each point on the surface using linear interpolation.

The length element on the curve is given by:

$$ds = \left| \frac{d\mathbf{x}}{dt} \right| dt = \sqrt{\left(\frac{dx}{dt} \right)^2 + \left(\frac{dy}{dt} \right)^2} dt \quad (24)$$

The area element on the surface \mathcal{A} is given by:

$$dA = \left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right| dudv \quad (25)$$

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

$$\frac{\partial M}{\partial \mathbf{p}_x} = \int_0^S \left[|b(x, y)| \frac{\partial}{\partial \mathbf{p}_x} \left| \frac{d\mathbf{x}}{dt} \right| + \frac{\partial |b(x, y)|}{\partial x} \frac{\partial x(\mathbf{p}, s)}{\partial \mathbf{p}_x} \left| \frac{d\mathbf{x}}{dt} \right| \right] ds \quad (26)$$

The corresponding derivative for the surface objective is:

$$\frac{\partial M}{\partial \mathbf{p}_x} = \iint_{\mathcal{A}} \left[|b(x, y, z)| \frac{\partial}{\partial \mathbf{p}_x} \left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right| + \frac{\partial |b(x, y, z)|}{\partial x} \frac{\partial x(\mathbf{p}, u, v)}{\partial \mathbf{p}_x} \left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right| \right] dudv \quad (27)$$

and similarly for y and z . This expression can also be evaluated by numerical integration. Expressions such as $\frac{\partial |b|}{\partial x}$ can be determined by discrete derivative calculations at each point on the curve or surface, again using linear interpolation.

The expressions such as $\frac{\partial x}{\partial \mathbf{p}_x}$ can be calculated from the expressions for x , y , and z (shown in Equation 6 or Equation 14).

The partials $\frac{\partial \mathbf{x}}{\partial u}$ and $\frac{\partial \mathbf{x}}{\partial v}$ can be evaluated either analytically or from discrete approximation. The expressions $\frac{\partial}{\partial \mathbf{p}_x} \left| \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} \right|$ and $\left| \frac{d\mathbf{x}}{dt} \right|$ can be calculated by expanding and evaluating expressions such as $\frac{\partial}{\partial \mathbf{p}_x} \left(\frac{dx}{dt} \right)$ by discrete approximation. The above follows similarly for $\frac{\partial}{\partial \mathbf{p}_y}$ and $\frac{\partial}{\partial \mathbf{p}_z}$.

7.6.1 Probabilistic Formulation

In order to incorporate probabilistic information into the measure of fit, consider the problem of boundary determination as one in which the data is a two or three dimensional image, $b(\mathbf{x})$, which could be depicting the boundary of any object in the parametric representation and $t_{\mathbf{p}}(\mathbf{x})$ is an image template corresponding to a particular value of the parameter vector \mathbf{p} . In terms of probabilities, if we want to decide which template, $t_{\mathbf{p}}$, an image, b , corresponds to, we need to evaluate the probability of the template given the image, $\Pr(t_{\mathbf{p}}|b)$, and find the maximum over \mathbf{p} . This can be expressed using Bayes rule, where:

$$\Pr(t_{\text{map}}|b) = \max_{\mathbf{p}} \Pr(t_{\mathbf{p}}|b) = \max_{\mathbf{p}} \frac{\Pr(b|t_{\mathbf{p}}) \Pr(t_{\mathbf{p}})}{\Pr(b)} \quad (28)$$

Here, t_{map} is the maximum *a posteriori* solution, $\Pr(t_{\mathbf{p}})$ is the prior probability of template $t_{\mathbf{p}}$ and $\Pr(b|t_{\mathbf{p}})$ is the conditional probability, or likelihood, of the image given the template. This expression can be simplified by taking the logarithm and eliminating $\Pr(b)$, the prior probability of the image data, which is equal for all \mathbf{p} :

$$M(b, t_{\text{map}}) = \max_{\mathbf{p}} M(b, t_{\mathbf{p}}) = \max_{\mathbf{p}} [\ln \Pr(t_{\mathbf{p}}) + \ln \Pr(b|t_{\mathbf{p}})] \quad (29)$$

This maximum *a posteriori* objective function shows the tradeoff or compromise that is made between prior information, $\Pr(t_{\mathbf{p}})$, and image-derived information, $\Pr(b|t_{\mathbf{p}})$. For a uniform prior, this formulation reduces to the maximum likelihood solution.

In order to derive the expression for the likelihood, consider the image b to be a noise-corrupted version of one of these templates with noise that is independent and additive: $b = t_{\mathbf{p}} + n$. This assumption avoids an excessive increase in complexity. Furthermore, Cooper [7] showed, for a related problem, that this assumption did not alter the performance significantly. Then, the likelihood, $\Pr(b|t_{\mathbf{p}})$, is equivalent to $\Pr(n = b - t_{\mathbf{p}})$. The noise at each image point, $n(\mathbf{x})$, equals $b(\mathbf{x}) - t_{\mathbf{p}}(\mathbf{x})$ and is governed by the probability density $\Pr(n)$. These events are independent for each point, so the probability for the noise over the entire region \mathcal{A} is just the product of the individual probabilities. The noise is the combined effect of many factors such as signal degradation, occlusion and boundary measurement which are difficult to model explicitly. We make the assumption that the noise is Gaussian with zero mean and standard deviation σ_n .

The object template, $t_{\mathbf{p}}(\mathbf{x})$, represents the boundary of the object. The boundary can be embedded into the image template by making $t_{\mathbf{p}}(\mathbf{x})$ constant along the boundary of the object it represents and zero everywhere else. In order to match this template with the image, consider $b(\mathbf{x})$ to be a boundary measure applied to the raw image data, $b(\mathbf{x}) = b(i(\mathbf{x}))$. Both $t_{\mathbf{p}}$ and b are image functions that represent boundaries that are summed (or integrated), only along the boundary.

Because the template has support only along the boundary, it is not necessary to sum over the entire image for terms involving the template, but only where the template has support. In addition, the magnitude of $t_{\mathbf{p}}(\mathbf{x})$ is taken to be constant (k), over the boundary that it defines. The function M can be simplified further by removing the terms that do not depend on \mathbf{p} .

The continuous version of this for a curve is:

$$M(b, \mathbf{p}) = \ln \Pr(\mathbf{p}) + \frac{k}{2\sigma_n^2} \int_0^S [b(x(\mathbf{p}, s), y(\mathbf{p}, s)) - k] ds \quad (30)$$




Figure 5: Example mean curve, shown with curves corresponding to parameters plus and minus one standard deviation.

where s is arclength. For a surface, it is:

$$M(b, \mathbf{p}) = \ln \Pr(\mathbf{p}) + \frac{k}{2\sigma_n^2} \iint_{\mathcal{A}} [b(x(\mathbf{p}, u, v), y(\mathbf{p}, u, v), z(\mathbf{p}, u, v)) - k] dA \quad (31)$$

where dA is an area element on the surface \mathcal{A} .

Equations 30 and 31 are the maximum *a posteriori* objective functions for curve and surface finding. In both, the first term is the contribution of the prior probability of the parameter vector. The greater the variance of the prior, the smaller the influence of this term. The second term is equivalent to the objectives in Equations 22 and 23.

The probability distributions associated with the parameters are intended to bias the model towards a particular range of shapes. This prior knowledge comes from experience with a sample of images of the object being delineated, when such a sample is available. When prior information is not available, uniform distributions are used for the prior probabilities of the parameters and an initial estimate of the boundary must be supplied. The images in a sample will differ due to variability in the object shape and the view of the object. The prior probability distributions can then be estimated from the shapes determined from the sample by decomposing the boundaries into their model parameters and collecting statistics. The boundaries of the sample objects are determined either by manual segmentation or, alternatively, this method can go through a training phase on a set of images with manual initialization and uniform distributions. This has been done only for the curve models so far because invariance to the surface parametrization has not been established for the surface models.

An independent, multivariate Gaussian can be used for the parameters. An example distribution is shown in Figure 5. The middle curve corresponds to the mean parameter values. Above and below it are the curves corresponding to the mean parameter values plus and minus one standard deviation.

7.7 Boundary Parameter 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 [32] 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.

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. However, the prior probability term in the objective function is quadratic and it dominates on the tails of the distributions, making distant points in the space non-optimal. The starting point for the optimization will be taken to be the maximum of the prior distributions. The global optimum probably will be near the starting point and thus a local optimum is likely to be a global optimum. The degree to which this is true depends on the width of the distributions. Since a local optimization method is likely to be sufficient, although there is still the possibility of converging to a poor local maximum, the excessive computation involved in finding a global optimum is deemed not necessary. Poor convergence can be identified by a corresponding low objective function value and verified visually. Smoothing can also be used to avoid getting trapped in a local maximum.

7.8 Experiments

From experiments varying the amount of noise added to a synthetic image, this method has been shown to be relatively insensitive to noise [33]. The effect of the initial values of the parameters on the performance was investigated by examining the results of running the same problem from different starting points. Each parameter was found to have a range within which the solution was found reliably [33]. Once the parameters are varied beyond that range, the result will converge to false local minima corresponding to nearby features. This region of success or capture about the true boundary depends on the quality of the image, the degree of smoothing and the particular problem. False minima can be distinguished, however, both visually and by the relative value of the objective function.

The deformable object boundary finding method has been applied to a variety of objects from real images, with an emphasis on heart and brain images using primarily magnetic resonance images. The results of the method applied to the problem of delineating the corpus callosum in the human brain from magnetic resonance images are shown in Figure 6. In these images, the corpus callosum is separated from the rest of the brain by a dark line. In this case, we used the positive magnitude of the Laplacian of the Gaussian as a line detector. The final contour succeeds in delineating the structure properly.

Magnetic resonance is becoming more and more important for cardiac imaging as acquisition rates increase into the range required for imaging the moving heart. In Figure 7, a transaxial cardiac image shows a section through the left ventricular wall. Here, the endocardial (inner) and epicardial (outer) walls of the left ventricle are objects to be delineated. The results of the two separate optimizations are shown.

In Figure 8, a transaxial slice of one frame of a cardiac image of a dog 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




Figure 6: Magnetic resonance mid-brain sagittal image example. Top: Magnetic resonance image (146×106). Middle: Initial contour (6 harmonics). Bottom: Final contour on the corpus callosum of the brain.




Figure 7: Magnetic resonance transaxial cardiac image example. Top: Magnetic resonance image (256×156). Middle: Initial contour on the endocardium and epicardium (4 harmonics). Bottom: Final contour on the endocardium and epicardium of the left ventricle.




Figure 8: Dynamic Spatial Reconstructor (DSR) transaxial cardiac image example. Top: Original transaxial view of left ventricle. Middle: initial contour for epicardium and endocardium. Bottom: Converged result of boundary detection algorithm on epicardium and endocardium.

capable of imaging the moving heart [34]. Both the endocardial (inner) and epicardial (outer) walls of the left ventricle are delineated as the result of two separate optimizations.

Surface finding in three-dimensional images is becoming more important due to the availability of range images and true three-dimensional images from magnetic resonance imaging (MRI), computed tomography (CT), single photon emission computed tomography (SPECT), positron emission tomography (PET) and confocal microscopy. Results of the surface finding method applied to the problem of delineating the upper portion of the cerebrum of the human brain from a three-dimensional magnetic resonance image are shown in Figure 9. The surface was matched to the gradient magnitude calculated from the image. The final boundary succeeds in delineating the structure properly.

In Figure 10, 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 was matched to the gradient magnitude calculated from the image. The endocardial (inner) wall of the left ventricle is successfully delineated.

7.9 Summary

This work presents a general boundary finding system for both two-dimensional and three-dimensional images of simple natural objects. The goal of this work was to incorporate model-based information about global shape into boundary finding for continuously deformable objects. In addition, the shape parametrization can be augmented with probabilistic information. From testing on real and synthetic, the system was found to perform well at delineating structures and to be relatively insensitive to the problems of broken boundaries and spurious edges from nearby objects. The flexibility of the model make this an attractive method for boundary 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. Although the current formulation is for three-dimensional images, these surface models also could be used to model $2\frac{1}{2}$ -D range data where the model would include the hidden surface.

There are, of course, areas of potential improvement for this work. The surface shape parametrization needs invariance to view and choice of surface parametrization u, v . Because the initial estimates of the view parameters may not be very good, an additional process to determine them could be added. This could involve an initial exhaustive coarse search over just those parameters. If this were done at a low resolution, the computation might not be excessive. Additional information, such as other low-level features or constraints between objects, might also help to guide the initial placement.

The framework presented here could perhaps also be used with other shape parametrizations better suited to man-made objects with straight sides and corners. The method could also be extended to object recognition where an image is fit to each of the models for different objects in a database. The correct model will result in the best fit because it will be the closest in the parameter space. The boundary




Figure 9: Magnetic resonance brain image example. Left: Three perpendicular slices through the three-dimensional image ($120 \times 160 \times 78$) are shown with the initial surface. Right: The same slices shown with final surface indicating the upper portion of the cerebrum. Bottom: Wire frame of initial (left) and final (right) surface.




Figure 10: Dynamic Spatial Reconstructor (DSR) cardiac image example. Left: Three perpendicular slices through the three-dimensional image ($49 \times 50 \times 55$) are shown with the initial surface. Right: The same slices shown with final surface at the endocardium. Bottom: Wire frame of initial (left) and final (right) surface.

finder and some of the ideas from this work have also been applied to the problem of contour-based deformable object motion [35].

Spatiotemporal models can be developed and used to measure motion. For two-dimensional objects, the motion can be characterized by the spatiotemporal surface corresponding to the object's moving boundary. The motion of surfaces could be modeled by a manifold in four dimensions. These spatiotemporal surfaces would be parametrized using basis functions. Note that if the correspondence between points on successive boundaries can be determined, this represents an approach to general, non-rigid object motion.

7.10 Acknowledgements

The dynamic spatial reconstructor data was obtained from Richard Robb of the Mayo Clinic. The code for the 3-D image gradient calculation was written by John Mceachen. This work was supported in part by the National Institutes of Health under Grants T15LM07056 and F36LM00007 from the National Library of Medicine and Grant R29HL38333 from the National Heart Lung and Blood Institute.

References

- [1] R. Mohan and R. Nevatia. Segmentation and description based on perceptual organization. In *Proc. Comp. Vision Pattern Recog.*, 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.
- [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. *Comm. ACM*, 19(2):73–83, February 1976.
- [5] D. H. Ballard and C. M. Brown. *Computer Vision*. Prentice-Hall, Englewood Cliffs, 1982.
- [6] C. W. K. Gritton and E. A. Parrish, Jr. Boundary location from an initial plan: The bead chain algorithm. *IEEE Trans. Pattern Anal. Machine Intell.*, 5(1):8–13, January 1983.
- [7] D. B. Cooper. Maximum likelihood estimation of Markov-process blob boundaries in noisy images. *IEEE Trans. Pattern Anal. Machine Intell.*, 1(4):372–384, October 1979.
- [8] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. Computer Vision*, 1(4):321–331, 1988.
- [9] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry-seeking models and 3D object reconstruction. *Int. J. Computer Vision*, 1:211–221, 1987.

- [10] J. M. Gauch and S. M. Pizer. Image description via the multiresolution intensity axis of symmetry. In *Proc. Second Int. Conf. Comp. Vision*, pages 269–274, 1988.
- [11] H. Delingette, M. Hebert, and K. Ikeuchi. Shape representation and image segmentation using deformable surfaces. In *Proc. Comp. Vision Pattern Recog.*, pages 467–472, June 1991.
- [12] Y. F. Wang and J. F. Wang. Surface reconstruction using deformable models with interior and boundary constraints. In *Proc. Third Int. Conf. Comp. Vision*, pages 300–303, December 1990.
- [13] I. Cohen, L. D. Cohen, and N. Ayache. Introducing new deformable surfaces to segment 3D images. In *Proc. Comp. Vision Pattern Recog.*, pages 738–739, 1991.
- [14] B. Widrow. The “rubber mask” technique - I and II. *Pattern Recognition*, 5:175–211, 1973.
- [15] A. L. Yuille, D. S. Cohen, and P. W. Hallinan. Feature extraction from faces using deformable templates. In *Proc. Comp. Vision Pattern Recog.*, pages 104–109, 1989.
- [16] Y. Chow, U. Grenander, and D. M. Keenan. Hands: A pattern theoretic study of biological shapes. Technical report, Div. Appl. Math., Brown University, Providence, Rhode Island, 1989.
- [17] A. P. Pentland. Automatic extraction of deformable part models. *Int. J. Computer Vision*, 4:107–126, 1990.
- [18] A. Pentland and S. Sclaroff. Closed-form solutions for physically based shape modeling and recognition. *IEEE Trans. Pattern Anal. Machine Intell.*, 13(7):715–729, July 1991.
- [19] P. J. Besl. Geometric modeling and computer vision. *Proc. 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] 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 Tech. Rpt. 112, University of Rochester, Rochester, New York, March 1981.
- [22] K. Rao and R. Nevatia. Computing volume descriptions from sparse 3-D data. *Int. J. Computer Vision*, 2(1):33–50, 1988.
- [23] A. H. Barr. Global and local deformations of solid primitives. *Computer Graphics*, 18(3):21–30, July 1984.
- [24] A. J. Hanson. Hyperquadrics: Smoothly deformable shapes with convex polyhedral bounds. *Comp. Vision Graphics Image Proc.*, 44:191–210, 1988.

- [25] J. Ponce. Straight homogeneous generalized cylinders: Differential geometry and uniqueness results. *Int. J. Computer Vision*, 4(1):79–100, 1990.
- [26] E. Persoon and K. S. Fu. Shape discrimination using Fourier descriptors. *IEEE Trans. Pattern Anal. Machine Intell.*, 8(3):388–397, May 1986.
- [27] C. R. Giardina and F. P. Kuhl. Accuracy of curve approximation by harmonically related vectors with elliptical loci. *Comp. Graphics Image Proc.*, 6:277–285, 1977.
- [28] F. P. Kuhl and C. R. Giardina. Elliptic Fourier features of a closed contour. *Comp. Graphics Image Proc.*, 18:236–258, 1982.
- [29] C. S. Lin and C. L. Hwang. New forms of shape invariants from elliptic Fourier descriptors. *Pattern Recognition*, 20(5):535–545, 1987.
- [30] L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformable models. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(11):1061–1075, 1992.
- [31] G. P. Tolstov. *Fourier Series*. Prentice-Hall, Englewood Cliffs, 1962.
- [32] D. A. Pierre. *Optimization Theory with Applications*. Dover, New York, 1986.
- [33] L. H. Staib. *Parametrically Deformable Contour Models for Image Analysis*. PhD thesis, Yale University, New Haven, CT, 1990.
- [34] R. A. Robb. High speed three-dimensional x-ray computed tomography: The dynamic spatial reconstructor. *Proc. IEEE*, 71:308–319, 1983.
- [35] J. S. Duncan, R. L. Owen, L. H. Staib, and P. Anandan. Measurement of non-rigid motion in images using contour shape descriptors. In *Proc. Comp. Vision Pattern Recog.*, pages 318–324, June 1991.