

# Medial-Node Models to Identify and Measure Objects in Real-Time 3-D Echocardiography

George D. Stetten,\* *Associate, IEEE*, and Stephen M. Pizer, *Senior Member, IEEE*

**Abstract**— A method is proposed for the automatic, rapid, and stable identification and measurement of objects in three-dimensional (3-D) images. It is based on local shape properties derived statistically from populations of medial primitives sought throughout the image space. These shape properties are measured at medial locations within the object and include scale, orientation, endness, and medial dimensionality. Medial dimensionality is a local shape property differentiating sphere-like, cylinder-like, and slab-like structures, with intermediate dimensionality also possible. Endness is a property found at the cap of a cylinder or the edge of a slab. In terms of an application, the cardiac left ventricle (LV) during systole is modeled as a large dark cylinder with an apical cap, terminated at the other end by a thin bright slab-like mitral valve (MV). Such a model, containing medial shape properties at just a few locations, along with the relative distances and orientations between these locations, is intuitive and robust and permits automated detection of the LV axis *in vivo*, using real-time 3-D (RT3D) echocardiography. The statistical nature of these shape properties allows their extraction, even in the presence of noise, and permits statistical geometric measurements without exact delineation of boundaries, as demonstrated in determining the volume of balloons in RT3D scans. The inherent high speed of the method is appropriate for real-time clinical use.

**Index Terms**—Core atom, left ventricle, medial dimensionality, 3-D ultrasound.

## I. INTRODUCTION

SHAPE identification and measurement in medical images are difficult but worthwhile goals. Nowhere is this more evident than in the automated analysis of cardiac left ventricular (LV) volume using real time three-dimensional (RT3D) ultrasound [1]–[6]. RT3D ultrasound is a new imaging modality that electronically scans a volume in 3-D, using a matrix array [Fig. 12(b)] instead of a conventional linear array [Fig. 12(a)]. Because of the large number of 3-D images produced every second by RT3D ultrasound, manual analysis of the data is extremely labor intensive. This makes RT3D ultrasound a prime candidate for automated analysis.

Manuscript received March 15, 1999; revised July 13, 1999. This work was supported in part by a grant from Whitaker Biomedical Engineering, and in part by the NSF under Grant CDR8622201 and the NIH under Grants 1K08HL03220, P01CA47982, and HL46242. The Associate Editors responsible for coordinating the review of this paper and recommending its publication were M. Viergever and W. Niessen. *Asterisk indicates corresponding author.*

\*G. D. Stetten was with the Medical Image Display and Analysis Group, University of North Carolina, Chapel Hill 27599-3165 USA, and with the Department of Biomedical Engineering, Duke University, Durham, NC 27708 USA. He is now with the Department of Bioengineering, University of Pittsburgh, Pittsburgh, PA 15261 USA.

S. M. Pizer is with Medical Image Display and Analysis Group, University of North Carolina, Chapel Hill 27599-3165 USA

Publisher Item Identifier S 0278-0062(99)09104-1.

Ultrasound, in general, presents a significant challenge to automated analysis because of high noise, low resolution, path dependence, and nonrectilinear coordinates. RT3D ultrasound, in particular, suffers from especially high noise and low resolution.

While an exhaustive comparison of the numerous methods developed for object recognition and measurement is beyond the scope of this paper, the majority of researchers in recent years have chosen deformable models for locating the cardiac ventricle, based either on contours [two-dimensional (2-D)] or surfaces (3-D). Deformable models constitute a top-down approach, in which an iterative search for likely boundaries is constrained by prior knowledge about the expected shape. For a review of deformable models, see McInerney and Terzopoulos [7]. Methods of finding and measuring the LV using 2-D data from conventional ultrasound scanners have concentrated on deformable contours [8]–[12]. Deformable surfaces have been applied to mechanically scanned 3-D echocardiographic data by at least one researcher [13], but our own experience with applying deformable methods to RT3D ultrasound, with its even lower data quality, has not produced encouraging results [14]. The difficulties we encountered along this avenue led us to review bottom-up approaches based on the measurement of geometric image properties.

One such approach has yielded particularly successful fully automatic detection of the LV in 2-D ultrasound data. In this method circular arc matched filters are used to find cross sections of the ventricle [15]. Other successes have been reported in 2-D ultrasound using a version of template matching to find the center of the ventricle and then fuzzy reasoning to find the boundary [16]. In recent research along the same lines, we have developed a Hough transform approach, using circular edge filters, that yields fully automated measurement of balloons in RT3D data [17], [18]. All of these approaches possess fundamentally medial aspects, that is, they relate multiple boundary points to common central points deep within the object. Our preliminary success with the RT3D ultrasound data led us to explore a generalized medial approach that extends to more complicated shapes and that differentiates medial relationships between boundary points into several fundamental types.

The lineage of the medial approach may be traced to the medial axis (otherwise known as the symmetric axis or skeleton) introduced on binary images by Blum and developed by Nagel, Nackman, and others [19]–[21]. Pizer has extended the medial axis to gray-scale images, producing a graded measure called medialness, which links the aperture of the

boundary measurement to the radius of the medial axis to produce what has been labeled a core. A core is a locus in a space whose coordinates are position, radius, and associated orientations [22], [23]. Methods involving these continuous loci of medial primitives have proven particularly robust against noise and variation in target shapes [24]. Determining locations with high medialness and relating them to a core has been accomplished by analyzing the geometry of loci resulting from ridge extraction [25]. Models including discrete loci of medial primitives have also provided the framework for a class of active shape models known as deformable m-reps (sampled medial representations) [26], [27].

The objective of the work reported here is to build on these ideas to produce a method capable of analyzing 3-D shapes to identify and measure structures in the heart, using RT3D ultrasound. The approach aims to extract the scale, orientation, and dimensionality (shape type) of sections of anatomical structures by statistical analysis of populations of medial primitives. In particular, the primitives are identified by first searching for individual boundary points throughout the image in an initial sweep and then matching pairs of boundary points to form what are called core atoms. A core atom is the smallest indivisible unit for detecting medialness, namely, a single pair of boundary points. Core atoms tend to cluster along a medial ridge with greater stability than individual boundary detectors, allowing for robust statistical analysis of the core and the underlying figure. Core atoms have already been developed for analysis of uncluttered 2-D shapes [28]. The techniques are generalized here to 3-D and augmented to provide for spatially sampled populations of core atoms to differentiate local portions of the core.

## II. METHODS

The methods developed in this paper will be discussed in terms of a particular goal, the automated determination of the apex-to-mitral-valve (AMV) axis of the systolic LV. During systole, the LV is basically a large cylinder with an apical cap at one end and a slab-like mitral valve (MV) at the other (we limit ourselves here to apical scans and to times when the MV is closed). This model of the LV is shown in Fig. 13(a). The overall method, which we will develop in detail below, is outlined in Fig. 1. Although applied here to detecting the AMV axis, the underlying method is generally adaptable and expandable to other more complicated anatomical structures.

### A. Finding Core Atoms

A core atom is defined as two boundary points  $\mathbf{b}_1$  and  $\mathbf{b}_2$  that satisfy particular requirements (described in detail below), guaranteeing that the boundaries face each other. A core atom can be represented by a single vector  $\vec{\mathbf{c}}_{1,2}$  from the first boundary point to the second. A core atom is said to be located at a center point midway between the boundary points [see Fig. 2(a)]. The medialness at the center point is high because the boundariness at both boundary points is high and because the boundary normals face each other. Core atoms carry information about orientation, width and position,

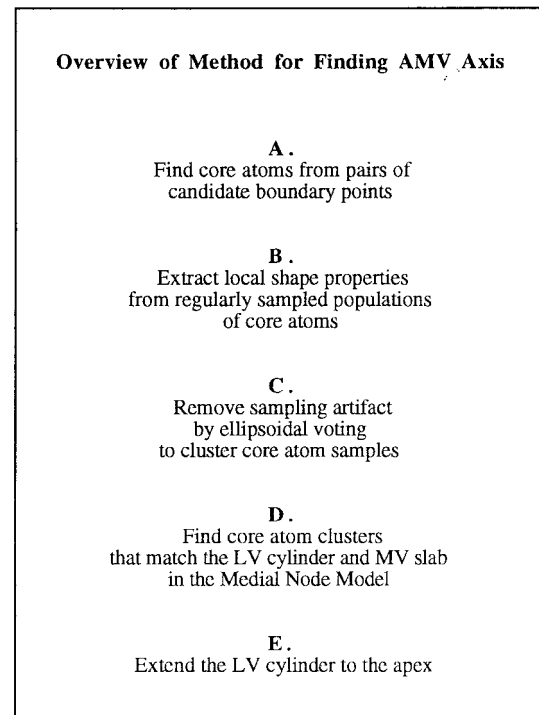


Fig. 1. Overview of method to find the apex-to-mitral valve axis. Letters correspond to subsections in Section II.




Fig. 2. (a) A core atom consists of two boundary points that face each other across an acceptable distance and a center point at which the core atom is said to be located. (b) The search area (gray) for boundary point  $\mathbf{b}_2$  depends on boundary normal  $\hat{\mathbf{n}}_1$  and the expected distance between the boundaries.

permitting populations of core atoms to be analyzed in these terms.

Unlike some medial models, the angle between the lines from the center point to each respective boundary point for a core atom is fixed at  $180^\circ$ , as with Brady [29]. A certain flexibility in relative orientations of the associated boundaries away from parallel is allowed. Boundariness is sampled on a rectilinear grid to select a population of boundary points  $\mathbf{b}_i$  at locations  $\mathbf{x}_i$  with orientations  $\hat{\mathbf{n}}_i$ . Any kind of boundariness, including those based on gradient, variance, or texture analysis, can be used to form core atoms, provided an orientation is established for each boundary point. In general, the aperture of the boundariness detector is held proportional to the distance between the boundaries of the desired core atoms. The strength inherent to the statistics of populations is meant to counteract the weakness of the preselection of boundary points.

Core atoms are created from a population of candidate boundary points by finding pairs that satisfy the following three criteria.

- 1) The magnitude of the core atom vector  $\vec{c}_{1,2}$ , i.e., the distance from one boundary point to the other, must be between  $c_{\min}$  and  $c_{\max}$

$$\vec{c}_{1,2} = \mathbf{x}_2 - \mathbf{x}_1 \quad c_{\min} \leq \|\vec{c}_{1,2}\| < c_{\max}. \quad (1)$$

The core atom vector can be oriented either way since the order of the boundary points is arbitrary.

- 2) The boundary points must have sufficient face-to-faceness  $F$  defined as

$$F(\mathbf{b}_1, \mathbf{b}_2) = f_1 \cdot f_2 \quad f_1 = \hat{c}_{1,2} \cdot \hat{\mathbf{n}}_1 \quad f_2 = \hat{c}_{2,1} \cdot \hat{\mathbf{n}}_2 \quad (2)$$

( $\hat{\cdot}$  denotes normalization,  $\hat{\mathbf{v}} \equiv \vec{\mathbf{v}}/\|\vec{\mathbf{v}}\|$ .) Since  $f_1$  and  $f_2$  are normalized to lie between +1 and -1, their product  $F$  must also lie between +1 and -1. Values for  $F$  near +1 occur when the boundaries face toward (or away from) each other across the distance between them. A threshold for acceptable face-to-faceness is set within some error  $\varepsilon_f$  such that  $F(\mathbf{b}_1, \mathbf{b}_2) > 1 - \varepsilon_f$ .

- 3)  $F(\mathbf{b}_1, \mathbf{b}_2) > 0$  implies that  $f_1$  and  $f_2$  are both positive, or both negative. The sign of  $f_1$  (or  $f_2$ ) is called the polarity. The appropriate polarity is either + or -, depending on whether the expected target is lighter or darker than the background.

A single boundary point can be involved in a number of core atoms, each linking to a different partner on the other side of the object. Although at first glance the search for pairs of boundary points appears to be  $O(n^2)$ , hashing individual boundary points beforehand by location yields a large reduction in computation time [see Fig. 2(b)]. The search area for boundary point  $\mathbf{b}_2$  is limited to a solid sector surrounding the orientation  $\hat{\mathbf{n}}_1$  of boundary point  $\mathbf{b}_1$  and to a range between  $c_{\min}$  and  $c_{\max}$ . The width of the sector depends on  $\varepsilon_f$ .

### B. Extracting Local Shape Properties

Observe that collections of core atoms can group in three basic ways, corresponding to the fundamental geometric shapes shown Fig. 3. The object boundaries are shown in dark gray with the corresponding cores shown in light gray. Beneath each object is the population of core atoms that would be expected to form with such objects, the core atoms now being depicted as simple line segments.

The sphere generates a Koosh-ball-like cloud of core atoms with spherical symmetry and with the core atom centers clustered at the center of the sphere. The cylinder generates a spokes-of-a-wheel arrangement with radial symmetry along the axis of the cylinder and the core atom centers clustered along the axis of the cylinder. The slab results in a bed-of-nails configuration across the slab, with the core atom centers clustered in the midplane of the slab. The cores of these basic objects form, respectively, a point, a line, and a plane. As shown in Fig. 3, a system of local coordinate axes, namely,  $\hat{\mathbf{a}}_1, \hat{\mathbf{a}}_2$ , and  $\hat{\mathbf{a}}_3$ , can be assigned in each case, although not all the axes are unique given the symmetries involved. For example, in the slab,  $\hat{\mathbf{a}}_1$  and  $\hat{\mathbf{a}}_2$  can rotate freely about  $\hat{\mathbf{a}}_3$ . Such a set of coordinate axes can be found for any population




Fig. 3. Fundamental shapes (dark gray), corresponding cores (light gray), core atom populations (line segments) and eigenvectors  $\hat{\mathbf{a}}_1, \hat{\mathbf{a}}_2$ , and  $\hat{\mathbf{a}}_3$ .




Fig. 4. The lambda triangle defines the domain of possible eigenvalues that determine medial dimensionality.

of core atoms using eigenanalysis, as will be shown below. Furthermore, the extent to which a core atom population resembles one of the three basic configurations depends on the corresponding eigenvalues.

Given a population of  $m$  core atoms  $\vec{c}_i, i = 1, 2, 3, \dots, m$ , the analysis of a given core atom population begins by separating each core atom vector  $\vec{c}_i$  into its magnitude  $c_i$  and its orientation  $\hat{c}_i$ . We ignore, for the moment, the location of the core atom. The analysis of the magnitude  $c_i$  over a population of core atoms yields a mean and standard deviation for the measurement of width in the underlying figure. The orientation  $\hat{c}_i$  of core atoms in a population lends itself to eigenanalysis, yielding measures of dimensionality and overall orientation for the population. We develop the eigenanalysis here in  $n$  dimensions, although for the remainder of the paper  $n$  will be 3.

Given a population of  $m$  core atom orientations  $\hat{c}_i$  in  $n$  dimensions, it is possible to find an  $n$ -dimensional vector  $\hat{\mathbf{a}}_1$  that is most orthogonal to that population as a whole by minimizing the sum of squares of the dot product between  $\hat{\mathbf{a}}$  and each individual unit vector  $\hat{c}_i$

$$\hat{\mathbf{a}}_1 = \arg \min_{\hat{\mathbf{a}}} \frac{1}{m} \sum_{i=1}^m (\hat{\mathbf{a}} \cdot \hat{c}_i)^2 = \arg \min_{\hat{\mathbf{a}}} (\hat{\mathbf{a}}^T \mathbf{C} \hat{\mathbf{a}}), \quad \text{where} \quad (3)$$

$$\mathbf{C} = \frac{1}{m} \sum_{i=1}^m \hat{c}_i \hat{c}_i^T.$$

The  $\mathbf{C}$  matrix is positive definite symmetric and has a unit trace. Therefore, its eigenvalues are positive and sum to one, and its eigenvectors are orthogonal. If the eigenvalues




Fig. 5. (a) Sphere. (b) Core atom cloud. (c) Sample displaced by  $\vec{p}$ .




Fig. 6. (a) Cylinder. (b) Core atom cloud. (c) Sample displaced by  $\vec{p}$ .

of  $\mathbf{C}$  are sorted  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ , the corresponding eigenvectors  $\hat{\mathbf{a}}_1 \dots \hat{\mathbf{a}}_n$  are the axes of a coordinate system in which  $\hat{\mathbf{a}}_1$  is the most orthogonal to the population  $\hat{\mathbf{c}}_i$  as a whole. For example, it would be the axis of the cylinder in Fig. 3. Furthermore, the eigenanalysis guarantees that  $\hat{\mathbf{a}}_2$  is the most orthogonal to the population  $\hat{\mathbf{c}}_i$  among those directions that are already orthogonal to  $\hat{\mathbf{a}}_1$ . This process can be repeated until  $\hat{\mathbf{a}}_n$  remains the least orthogonal to the population  $\hat{\mathbf{c}}_i$ , representing a form of average orientation for  $\hat{\mathbf{c}}_i$ . The axes  $\hat{\mathbf{a}}_1 \dots \hat{\mathbf{a}}_n$  are thus ordered from codimensional (orthogonal to the vector set) to dimensional (collinear with the vector set). In 3-D, the codimensional space is that of the core itself. That is, the space most orthogonal to the core atoms is the point, line, or plane of the core, as shown in Fig. 3.

Returning now specifically to 3-D, the previous analysis yields three eigenvalues which describe the dimensionality of the core

$$\lambda_i \geq 0 \quad \lambda_1 + \lambda_2 + \lambda_3 = 1 \quad (4)$$

An eigenvalue of zero means that the corresponding eigenvector is perfectly orthogonal to every core atom orientation  $\hat{\mathbf{c}}_i$ . Such is the case for  $\hat{\mathbf{a}}_1$  in the cylinder, and for both  $\hat{\mathbf{a}}_1$  and  $\hat{\mathbf{a}}_2$  in the slab. In the sphere none of the eigenvectors is completely orthogonal to the core atom population. Given the symmetries of the three basic shapes, the eigenvalues shown in Fig. 4 result. Notice that they sum to 1 for each shape. Since  $\lambda_3$  is dependent on the other two, the system may be viewed as having only two independent variables,  $\lambda_1$  and  $\lambda_2$ . Possible values for  $\lambda_1$  and  $\lambda_2$  are limited by  $\lambda_1 \leq \lambda_2$  and  $\lambda_2 \leq (1 - \lambda_1)/2$  which define a triangular domain we call the *lambda triangle* (Fig. 4).

The vertices of the lambda triangle correspond to the three basic shapes in Fig. 3, and all possible eigenvalues for any population of core atoms fall within the triangle. A rather crude simplification of dimensionality is possible by dividing

the triangle into three compartments to provide an integer description of dimensionality. Arbitrary thresholds of  $\lambda_1 = 1/5$  and  $\lambda_2 = 1/3$  will be used to divide the triangle into such areas of integer dimensionality to clarify our experimental results. However, it should be remembered that the underlying parameterization of medial dimensionality is not an integer or even a single scalar, but rather two independent scalars  $\lambda_1$  and  $\lambda_2$  whose values are constrained by the lambda triangle.

### C. Removing the Sampling Artifact

To incorporate location, we sort core atoms by their center points into bins on a regular 3-D lattice. Each bin thus represents a spatial sampling of medialness. The number of core atoms in a sample volume can be thought of as the medial density at that location.

How do we choose an appropriate size for the sample volume? As we shall see, the local distribution of core atoms can have a significant cross section, and the density within that distribution may not be uniform. To preserve resolution, the sample volume needs to be smaller than the typical cross-section of a core atom cloud. Generally, the core atom population within a sample volume will not contain the entire thickness of a cloud. Furthermore, it will demonstrate a distortion in dimensionality as one moves out from the center of the cloud. This is shown in Figs. 5 and 6. The vector from the theoretical core (center of the sphere, axis of the cylinder) to the sample volume is called the displacement vector  $\vec{p}$ . As shown in Fig. 5, a zero-dimensional core at the center of a sphere will appear to be one-dimensional (1D) (cylindrical) when sampled off-center. Likewise, in Fig. 6, the 1-D core of a cylinder will appear to be 2-D (slab-like). We call a sample that may be off-center a coronal sample.

We cannot know the actual location of the center of the cloud from a single coronal sample, but we can combine multiple samples in the cloud to determine its center. Again,




Fig. 7. Ellipsoids of three coronal core atom samples coalescing at the true center.

consider the shapes in Figs. 5 and 6. The population of core atoms in each coronal sample will be flattened in a plane orthogonal to  $\vec{p}$  and thus develop orthogonality to that direction. We already have seen a method using eigenanalysis to determine the direction of maximum orthogonality to a core atom population, namely the first eigenvector  $\hat{a}_1$ . The results of the eigenanalysis for each sample may thus be used in a Hough-like fashion simultaneously to vote for its own center of mass, as well as for possible samples whose corona it may inhabit. The voting takes place within ellipsoids around each sample volume. The axes of each ellipsoid are long in direction(s) orthogonal to the sample of core atoms. Thus, each ellipsoid can be expected to extend in the  $\vec{p}$  direction for that sample.

Fig. 7 demonstrates this concept. A circular cross section through an object is shown with three coronal sample (each containing three core atoms) displaced from the center. An ellipsoid is associated with each sample, with the major axis of each ellipsoid along the eigenvector most orthogonal to the corresponding core atoms. The three ellipsoids intersect at the center the circle. Fig. 7 can be interpreted as the cross section of a sphere with the populations of core atoms being cylindrical (seen in cross section) and the ellipsoids intersecting at the center of the sphere (as in Fig. 5). Alternatively, it can be interpreted as the cross section of a cylinder with the populations of core atoms being slab like and the ellipsoids intersecting along the axis of the cylinder (as in Fig. 6).

Various ways of constructing such ellipsoids are possible. We have chosen the following heuristic for its simplicity. The axes of our ellipsoid are the eigenvectors of the sample's  $\mathbf{C}$  matrix. The lengths  $a_i$  of the ellipsoid's three axes ( $i = 1, 2, 3$ ) are related to the eigenvalues  $\lambda_i$  as follows:

$$\begin{aligned} a_1 &= \gamma \bar{c} & a_2 &= \frac{\alpha_2}{\alpha_1} a_1 & a_3 &= \frac{\alpha_3}{\alpha_1} a_1, & \text{where} \\ \alpha_i &= 1 - \lambda_i, & \gamma &= 1/2. \end{aligned} \quad (5)$$

The scalar distance  $\bar{c}$  is the mean diameter of the core atoms in the sample and the dimensionless number  $\gamma$  relates  $\bar{c}$  to the size of the ellipsoid, determining how many neighbors will be reached. The ellipsoids make it possible to cluster the core atoms for a given cloud, in effect, to coalesce the corona. Each sample (the votee) receives votes from all the neighboring




Fig. 8. Distribution of samples in the lambda triangle for parametric test objects.

samples whose ellipsoids overlap it. The votes from those ellipsoids are assigned a strength  $v$  where  $v = m \cdot \exp(-d_e^2)$   $m$ , being the number of core atoms in the voting sample and  $d_e$  being the ellipsoidal distance

$$d_e = \sqrt{\sum_{i=1}^3 \left( \frac{\hat{a}_i \cdot \vec{d}}{a_i} \right)^2} \quad (6)$$

from the center of the voter ellipsoid to the votee,  $\vec{d}$  being the vector from voter to the votee. Votes are constructed to contain information about the constituent core atom population of the voter, including its  $\mathbf{C}$  matrix which may simply be summed (scaled by  $v$ ) over all votes for eigenanalysis of the entire constituent core atom population of a particular votee. Thus are formed clusters of core atoms that no longer suffer from coronal distortion. The center of mass for the constituent core atoms in a cluster, as well as their scale and dimensionality, will tend to reflect the true core rather than any particular coronal sample.

One expects the displacement vector  $\vec{p}$  to be one of the eigenvectors at the closest point on the theoretical core because 1) the displacement vector will be orthogonal to the core at that point and 2) the normal to the core is always one of its eigenvectors. In 3-D, the medial manifold can have at most two dimensions and, thus, will always have such a normal. At the theoretical core, the eigenvalue in the direction of the displacement should be 1/3 for a sphere, 1/2 for a cylinder, and 1 for a slab. As one moves out along the displacement vector, the corresponding eigenvalue should drop toward zero as the sample develops orthogonality to  $\vec{p}$ , except in the case of the slab, in which the core atom population can be expected to fall off rather abruptly without significant flattening.

We demonstrate these concepts using three parametric test objects with simple geometries: a sphere, a torus, and a spherical shell. The torus is basically a cylinder of varied and known orientation and the spherical shell is likewise a slab of varied and known orientation. Fig. 8 shows the eigenvalues of all coronal samples containing greater than 1% of the entire core atom population, plotted on the lambda triangle. The sphere shows two groups of samples, one near the top (sphere) vertex of the triangle and another near the right (cylinder) vertex, consistent with the dimensional effects of the corona predicted in Fig. 5. The torus shows clustering near the right (cylinder) vertex, with some spreading toward the left (slab), consistent with the dimensional effects of the corona predicted in Fig. 6. The spherical shell shows tight clustering at the left (slab) vertex consistent with the observation that core atoms

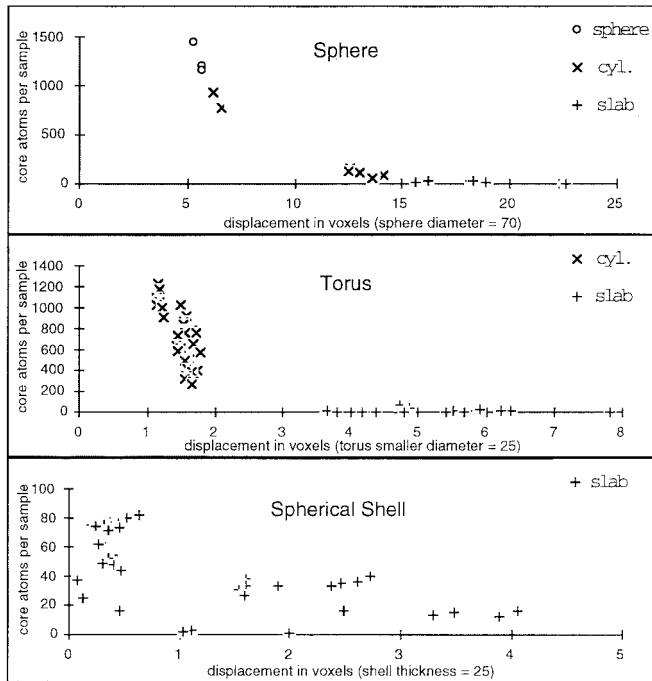


Fig. 9. Number of core atoms per sample versus displacement from the theoretical core, showing dimensional distortion in the corona.




Fig. 10. Core atom samples (small symbols) and clustered samples (large symbols) for parametric objects (line = slab, cross = cylinder, three-axis symbol = sphere).

in a slab are collinear with  $\vec{p}$  and therefore cannot develop significant orthogonality.

Spatial information about the samples is displayed in Fig. 9, which plots the number of core atoms in each bin as a function of distance from the theoretical core. In all cases, the core atoms are concentrated near the theoretical core, i.e., the center of the sphere, axis of the cylinder, etc. Integer dimensionality (determined by the simple partition of the lambda triangle shown in Fig. 4) is labeled as follows:  $o$  = sphere,  $x$  = cylinder,  $+$  = slab. Dimensionality behaves as expected, clearly showing the predicted distortion with displacement from the theoretical core of the sphere and the cylinder. As expected, the slab shows no such distortion.

The spatial distribution of core atom samples is visualized in a different manner in Fig. 10. Each sample volume that contains more than 1% of the total number of core atoms is shown as a thin-lined symbol. Here, the partition of the lambda triangle in Fig. 4 is used to decide between three possible symbols: a slab is represented as a single line, a cylinder as




Fig. 11. Endness. (a) Manifested as a cap on a cylinder. (b) Manifested as the edge of a slab. Boundary points contributing to endness are labeled  $b$  and eigenvectors of core atom cluster labeled  $a_1$ ,  $a_2$ , and  $a_3$  as in Fig. 2.

a cross, and a sphere as three intersecting axes. The length of the thin lines is constant in each test object, chosen for clarity. The orientation of the thin lines indicates the predominant direction(s) of core atoms in each sample, i.e., across the slab or orthogonal to the axis of the cylinder, keeping in mind that perfect spheres have no predominant orientation and perfect cylinders allow for arbitrary rotation around the axis.

As expected, the sphere shows cylindrical samples in its corona oriented toward its center. Further out from the center a few slab-like samples reflect simply the paucity of core atoms in those sample volumes. Near the center of the sphere, one true spherical sample (a small three-axis symbol) may be discerned.

The thick-lined symbols show the results of ellipsoidal voting, i.e., they represent clusters of samples. To prevent a cluttered illustration, votes have been tallied for all samples, a single winner declared, and then all votes by members of that constituency for other candidates removed. Using this Hough credit attribution strategy, the election is then repeated until all votes are gone. The winning clusters thus represent nonoverlapping constituencies, and are easier to see in the illustration. The winning clusters are represented by thick lines in a manner similar to the samples, except the length of the axes now corresponds to the actual mean scale of the constituent core atoms. Thus, the thick-lined three-axis cross in the sphere indicates its diameter. For the sphere there is only one predominant winning cluster, with virtually every core atom in its constituency. The torus shows cylindrical initial samples properly oriented, but dispersed throughout the corona. At the outer regions of the corona a few slab-like samples are visible. The clusters, by contrast, are centered on the circular midline of the torus. The spherical shell shows only slab-like samples, which coalesce with ellipsoidal voting into slab-like clusters. The orientation of the initial samples and clustered samples are both across the local slab. Ellipsoidal voting is seen here to perform another function, that of connecting samples that share a core along the midplane of a slab or along the axis of a cylinder.

#### D. Matching Local Shape Properties to a Medial-Node Model

The local shape properties derived from core atom clusters are now matched against the expected properties for the object




Fig. 12. (a) Conventional 2-D ultrasound uses a linear array to steer within a slice. (b) RT3D ultrasound uses a matrix array to scan a volume without moving the transducer. (c) Two orthogonal B-mode slices, and one C-mode slice (parallel to the transducer) within the 3-D ultrasound pyramid.

sought, here the cardiac LV. These can be viewed as being nodes in a medial-node model, in particular, the large dark cylinder of the LV and the thin bright slab of the MV [see again Fig. 13(a)]. Matching occurs both between individual nodes and clusters (clusters that appear to be candidates individually for the LV or MV) and between pairs of nodes and pairs of clusters (pairs of clusters must have appropriate relative distance and orientation to be an LV, MV pair). From this population of cluster pairs, an optimum location for the LV cylinder and MV slab can be determined, as described in more detail in the experimental section below.

#### E. Extending the Axis to the LV Apex

Some attention must be paid to cases where a cylinder ends at a cap, or a slab ends at an edge. The property of endness has been described by Clary *et al.* [30]. Endness, as viewed from the core atom perspective, is illustrated in Fig. 11. To detect endness, clusters of core atoms can be used as vantage points. Once a local cylinder has been established, unpaired boundary points can be sought along the axis of the cylinder, in either direction, as evidence of a cap on the cylinder. Similarly, once a local slab has been found, unpaired boundary points indicating an edge to the slab can be sought.

In the particular application of detecting the AMV axis, the direction of the cylinder is determined by the line between the optimally placed LV and MV nodes. Extending this line toward the apex of the LV permits detection of boundary points contributing to the cap of the cylinder, as described in more detail in the experimental section below.

### III. EXPERIMENTAL VALIDATION

#### A. Identifying the AMV Axis

In this section we experimentally validate the above method and demonstrate its ability to automatically identify the AMV axis of the LV. No preprocessing of the data was performed. Boundariness was found using a difference of Gaussian measurement of intensity gradient, with Gaussian application accomplished by repeated convolution with a  $2 \times 2 \times 2$  binomial kernel and with further constraints applied as to the absolute intensity of the boundary point.

To identify the cylinder in the image data, boundary points were determined with four applications of the binomial kernel. Core atoms with diameters 0.8–4.6 cm and face to faceness greater than 0.88 [see (2)] were collected in bins on a regular lattice and ellipsoidal voting applied. An example of the resulting clusters is displayed in Fig. 13(b). Crosses are shown in the cylindrical chamber of the ventricle. Due to the preselection of core atoms by scale, no other significant densities of core atoms were found. A single intensity constraint could not be found to reliably identify the endocardial boundary, because the intensity varied between images extending into the range of the outer boundary between the myocardium and surrounding connective tissue (epicardium or septum). However, since this outer boundary also formed a cylinder roughly concentric to that of the endocardial boundary, core atoms forming from the outer boundary established approximately the same axis for the LV.

Next, the MV was sought. Boundary points were determined as above, but with only two applications of the binomial kernel to accommodate the finer structure of the MV. Core atoms with diameters 0.0–0.8 cm and face-to-faceness greater than 0.55 were collected. The lower threshold for face to faceness was necessary because of the smaller size of the core atoms. As shown in Fig. 13(c), the densest clusters formed at the center of the MV, although weaker false targets were also detected in the myocardium off to the side of the ventricle. To eliminate these false targets, a criterion was established for the formation of appropriate pairs of clusters. Only slab-like clusters appropriately located and oriented with respect to cylindrical clusters were accepted. These pairs were allowed to vote for their constituent clusters and the mean location of the winning clusters used to establish a single MV location and a single LV cylinder location. The vector between these two locations was used to establish a cone for expected unpaired boundary points at the apex of the LV and the mean distance to such apical boundary points used to determine the location of the apical cap along that vector. Thus, an AMV axis was established, as shown in Fig. 13(d).

The entire procedure that produced Fig. 13(b)–(d) was automatic and required approximately 10 s per 3-D scan on a 400-MHz Pentium computer (each scan holds approximately 2 million 8-bit voxels).

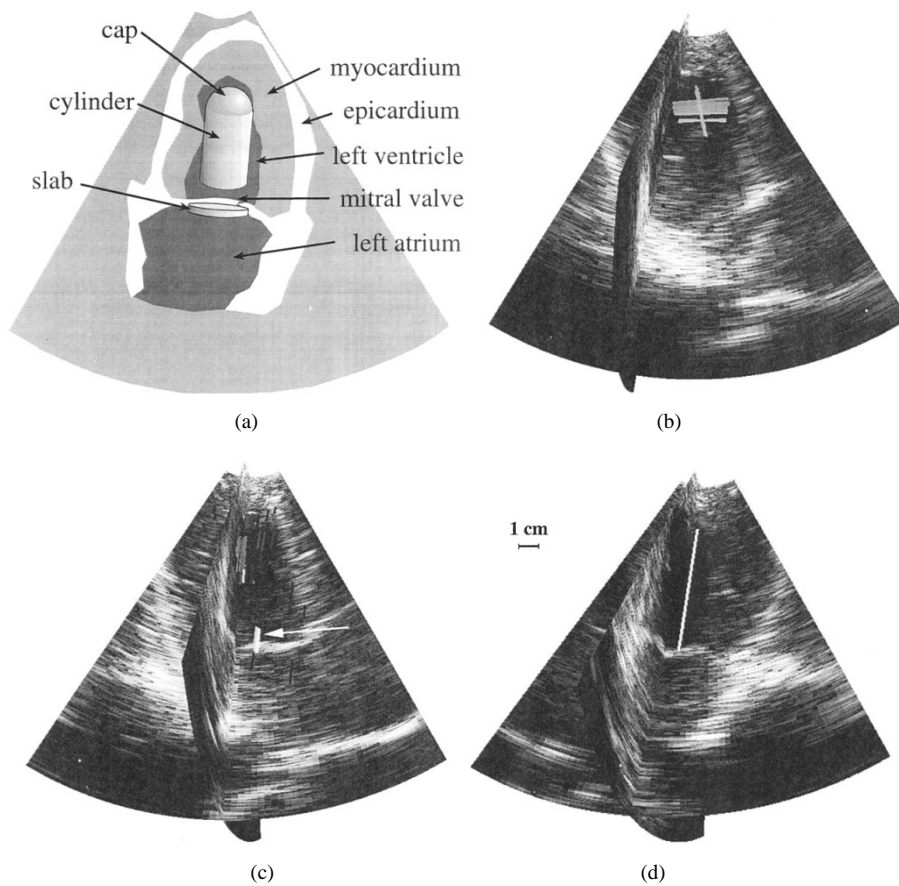


Fig. 13. Using a statistical model of medial primitives to automatically identify the axis of the cardiac left ventricle in RT3D ultrasound data. (a) Model of left ventricle and MV. (b) Cylinder of ventricle. (c) Slab of MV. (d) Automated AMV axis (scale of 1 cm is shown).

The anatomical end-points of the AMV axis (the ventricular apex and the center of the MV) were also determined manually. A human operator was instructed to follow the general cylindrical shape of the left ventricle, marking the ventricular apex and the center of the MV on B-mode and C-mode slices. Manual identification of these landmarks was conducted first, and then the automated analysis performed just once without any adjustment of parameters.

The data included 18 RT3D scan sequences of *in vivo* human hearts, using a Volumetrics Model-1 scanner operating at 2.5 or 3.5 MHz. Of the 18 sequences, 12 were described as normal, four as dilated cardiomyopathy, one as akinetic, and one as pericardial effusion. In each sequence, only scans in which the MV was closed were used, for a total of 155 scans. All parameters for the method were established during its developing using several normal sequences, none of which were included in these 18 sequences.

The locations of the manual end points were compared to those determined automatically for all 155 scans, as shown in Fig. 14(a) and (b). The error reported is simply the total physical distance (cm) in three dimensions between the manual and automated end points. For reference, a scale of 1 cm is marked on the cardiac scan in Fig. 13(d), with the pyramid of a typical scan having a height of approximately 15 cm. As can be seen in Fig. 14(a), for most scans the center of the MV was correctly located within 2 cm (rms error 1.2 cm). The greatest error was slightly more than 3 cm. For the apex of




Fig. 14. (a) Error (cm) between manual and automated placement of MV end-point of AMV axis for all 155 3-D ultrasound scans. (b) Error between manual and automated placement of apical end-point of AMV axis, for all 155 scans. (c) Error for the apical end-point for a random subset of 65 scans corrected for bias measured in the remaining 90 scans, reducing the rms error from 1.4 to 0.7 cm. (See Fig. 15(d) for scale of 1 cm.)

the left ventricle, the greatest error was approximately 2.5 cm (rms error 1.4 cm).

These errors were improved by eliminating a consistent bias between manual and automated measurements. The scans were divided blindly into two groups (training and experimental) without regard to image quality, distributing normal and abnormal hearts evenly, and placing the akinetic and pericardial effusion scans in the experimental group. Error for the LV apex is displayed in Fig. 14(c), with only the experimental group (65 scans) shown, corrected for bias observed in the training group (90 scans). Correcting for bias yielded an rms error of 0.7 cm ( $x = 0.37$  cm,  $y = 0.47$  cm,  $z = 0.41$  cm) for placement of the LV apex, and 1.1 cm ( $x = 0.48$  cm,  $y = 0.62$  cm,  $z = 0.72$  cm) for the MV. The automated



placement of the LV apex tended to be further into the ventricular chamber (average of 1.2 cm further away from the transducer) than the manual measurement. This bias may be due to greater local curvature at the LV apex than would have been expected from the simple model of an untapered cylinder with a hemispherical cap, since a typical LV actually narrows considerably as it approaches the apex. Another possible explanation for the bias is that the axis, as determined by the LV and MV clusters, is not centered perfectly as it extends into the apex, leading to contact along the wall rather than at the apex. Other sources of error include sampling in the boundary detection and ambiguities in the correct location of the landmarks by manual placement.

*B. Measuring Balloon Volume in 3-D Ultrasound Data*

The previous section demonstrates the core atoms' ability to identify anatomical structures in noisy cluttered images. It is expected that core atoms will also be useful in making measurements of diameter and related geometric parameters. In this section, we demonstrate the ability to accurately measure volume of real objects in 3-D ultrasound data.

We used a series of seven balloons filled with an ethanol-water mixture of known density. The volume of the balloons was determined by weight, ranging from 58.0 ml to 83.1 ml. The balloons were scanned in a bath of the same mixture using the prototype RT3D ultrasound scanner developed at Duke University known as T4. Boundariness was found, as in the previous section, using a difference of Gaussian, with six applications of the binomial kernel. Core atoms with lengths ranging from 2.3 to 7.8 cm and face to faceness of greater than 0.88 were collected from each scan. A center of mass was computed for all core atoms in each scan. The majority of core atoms formed a roughly spherical Koosh-ball configuration as predicted in Fig. 3. An actual example of clusters from a balloon is shown in Fig. 16. Volume measurement could be performed simply by selecting core atoms within 1 cm of the center of mass of all core atoms in the image, since the image was uncluttered by other targets. The mean length of these core atoms (divided by two) yielded an effective radius  $r$  from which a volume was calculated using  $v = (4/3)\pi r^3$ . This automatically determined volume was compared to balloon volume determined by weight.

A complication arose because the skin of a balloon in ultrasound presents a dark-light-dark ridge in intensity rather than a single dark-light transition. This causes boundary points to change orientation as one crosses the skin of the balloon. Using the inward facing boundary points to form core atoms yields a smaller radius than expected, while using the outward facing boundary points (by switching the polarity, defined in Section II-A) yields a larger radius than expected. These results for the two cases are marked + and x, respectively, in Fig. 15.

An optimum weighted average radius  $r = kr_i + (1 - k)r_o$  was computed from the inner radius  $r_i$  and the outer radius  $r_o$ , by minimizing the rms percent error between the resulting calculated volume and that determined by weight, for the set of seven balloons. The optimum value for  $k$  was 0.38, favoring the outward facing boundary points, yielding an rms percent




Fig. 15. Radius of fluid-filled balloons determined automatically using core atoms to analyze 3-D ultrasound images, compared to radius of the same balloons determined by weight assuming a spherical shape. (x) outer facing boundary, (+) inner facing boundary, (o) weighted average of inner and outer boundaries, with weight determined to minimize rms percent error by volume to 6.5%.




Fig. 16. Core atom clusters in balloon (from the inside boundary of the intensity ridge) identified as cylindrical along the axis of the balloon, although a significant number of core atoms actually formed vertically through the balloon as well.

error by volume of 6.5%, which agrees favorably with previous measurements on the same balloons using a Hough transform [18]. The weighted average radii (marked o in Fig. 15) yielded a linear regression (dotted line) close to unity (slope = 1.03, intercept = 0.081 cm).

This measurement of balloon volume was completely automated and took advantage of the fact that many individual measurements were combined statistically. In a sense, each core atom serves as a yardstick crossing the sphere near the center. Since, generally, the selected core atoms are chords of the sphere and not true diameters [as shown in Fig. 5(c)], they can be expected to underestimate the actual diameter. This may explain the bias toward the outer surface in optimum

value for  $k$ . Another reason for this bias may be that the balloons were generally ellipsoidal with a single major axis along the ultrasound beam, producing a propensity for shorter core atoms across the two minor axes.

#### IV. CONCLUSIONS

We have described a new method for identifying anatomical structures using fundamental properties of local shape, extracted statistically from populations of medial primitives, and have demonstrated the feasibility of using this method to identify and measure shapes under challenging conditions. The balloon experiment demonstrated the ability of core atoms to accurately measure a roughly spherical target with an uncluttered background. Identification of the left ventricular axis in 3-D ultrasound data using a model with three nodes has also been demonstrated.

The bottom-up approach of our method based on image properties may offer advantages over top-down approaches such as deformable surfaces. We still employ a model, but simply apply it at a higher more abstract level. This may provide a more robust and efficient fully automated system especially in noisy data.

Present directions in our research include determining ventricular volume using the medial-node model of the LV axis, and constructing more complicated medial-node models for the heart and other anatomical structures, as well as using other imaging modalities besides RT3D ultrasound. We are introducing variability into the model to reflect normal and pathologic variation in anatomy, extending the method to the spatio-temporal domain, and applying the approach to new methods of visualization. Finally, work is underway to increase the flexibility of initial boundary detection by providing multiple adaptable boundary profiles and to automatically develop optimal medial-node models directly from training sets.

#### ACKNOWLEDGMENT

The authors thank T. Shiota, Department of Cardiology, the Cleveland Clinic Foundation, and Volumetrics Medical Imaging, Inc., Durham, NC, for data generously supplied.

#### REFERENCES

- [1] G. Stetten, T. Ota, C. Ohazama, C. Fleishman, J. Castellucci, J. Oxaal, T. Ryan, J. Kisslo, and O. T. v. Ramm, "Real-time 3D ultrasound: A new look at the heart," *J. Cardiovascular Diagnosis Procedures*, vol. 15, no. 2, pp. 73–84, 1998.
- [2] D. P. Shattuck, M. D. Weinschenker, S. W. Smith, and O. T. v. Ramm, "Explososcan: A parallel processing technique for high speed ultrasound imaging with linear phased arrays," *J. Acoust. Soc. Amer.*, vol. 75, no. 4, pp. 1273–1282, 1984.
- [3] S. W. Smith, J. H. G. Pavy, and O. T. von Ramm, "High-speed ultrasound volumetric imaging system—Part I: Transducer design and beam steering," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 38, pp. 100–108, Mar. 1991.
- [4] O. T. von Ramm, S. W. Smith, and H. G. Pavy Jr., "High-speed ultrasound volumetric imaging system—Part II: Parallel processing and image display," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 38, pp. 109–115, Mar. 1991.
- [5] O. T. von Ramm, H. A. Pavy, S. W. Smith, and J. Kisslo, "Real-time, three-dimensional echocardiography: the first human images," *Circulation*, vol. 84, suppl. II, no. 4, p. 685, 1991.
- [6] S. W. Smith, H. G. Pavy, J. A. Miglin, and O. T. v. Ramm, "Improved real-time volumetric ultrasonic imaging," *Ultrasonic Imag.*, vol. 14, pp. 186–211, 1992.
- [7] T. McInerney and D. Terzopoulos, "Deformable models in medical image analysis: a survey," *Med. Image Anal.*, vol. 1, no. 2, pp. 91–108, 1996.
- [8] V. Chalana, D. T. Linker, D. R. Haynor, and Y. Kim, "A multiple active contour model for cardiac boundary detection on echocardiographic sequence," *IEEE Trans. Med. Imag.*, vol. 15, pp. 290–298, June 1996.
- [9] T. Hozumi, K. Yoshida, H. Yoshioka, T. Yagi, T. Akasaka, T. Tagaki, M. Nishiura, M. Watanabe, and J. Yoshikawa, "Echocardiographic estimation of left ventricular cavity area with a newly developed automated contour tracking method," *J. American Society Echocardiography*, vol. 10, no. 8, pp. 822–829, 1997.
- [10] I. Mikic, S. Krucinski, and J. D. Thomas, "Segmentation and tracking in echocardiographic sequences: Active contours guided by optical flow estimates," *IEEE Trans. Med. Imag.*, vol. 17, pp. 274–284, Apr. 1998.
- [11] S. Malassiotis and M. G. Strintzis, "Tracking the left ventricle in echocardiographic images by learning heart dynamics," *IEEE Trans. Med. Imag.*, vol. 18, pp. 282–290, Mar. 1999.
- [12] I. A. Hunter, J. J. Soraghan, J. Christie, and T. S. Durrani, "Detection of echocardiographic left ventricle boundaries using neural networks," in *Computers in Cardiology*. Washington, DC: IEEE Computer Society Press, 1993.
- [13] G. Coppini, R. Poli, and G. Valli, "Recovery of the 3-D shape of the left ventricle from echocardiographic images," *IEEE Trans. Med. Imag.*, vol. 14, pp. 301–317, June 1995.
- [14] G. Stetten, T. Irvine, D. Ritscher, O. T. v. Ramm, J. Panza, V. Sachdev, J. Castellucci, M. Jones, and D. Sahn, "Improved accuracy for a semi-automated method for computing right ventricle (RV) cavity volumes from Real Time 3D Echo: Comparison studies to ultrasonic crystals in an open-chest animal model," in *Proc. American College Cardiology 48th Scientific Sessions*, New Orleans, LA: 1999, to be published.
- [15] D. C. Wilson, E. A. Geiser, and J. Li, "Feature extraction in two-dimensional short-axis echocardiographic images," *J. Math. Imag. Vision*, vol. 3, pp. 285–298, 1993.
- [16] J. Feng, W. C. Lin, and C. T. Chen, "Epicardial boundary detection using fuzzy reasoning," *IEEE Trans. Med. Imag.*, vol. 10, pp. 187–199, June 1991.
- [17] G. Stetten and R. Morris, "Shape detection with the flow integration transform," *Inform. Sci.*, vol. 85, pp. 203–221, 1995.
- [18] G. Stetten, M. Caines, C. Ohazama, and O. T. v. Ramm, "The volumetric cardiogram (VCG): Volume determination of cardiac chambers using 3D matrix-array ultrasound," in *Proc. SPIE Symp. Medical Imaging*, 1995, vol. 2432, pp. 185–196.
- [19] H. Blum and R. N. Nagel, "Shape description using weighted symmetric axis features," *Pattern Recognit.*, vol. 10, pp. 167–180, 1978.
- [20] L. R. Nackman, "Curvature relations in 3D symmetric axes," *CVGIP*, vol. 20, pp. 43–57, 1982.
- [21] L. R. Nackman and S. M. Pizer, "Three-dimensional shape description using the symmetric axis transform I: Theory," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 2, pp. 187–202, Mar. 1985.
- [22] C. A. Burbeck and S. M. Pizer, "Object representation by cores: Identifying and representing primitive spatial regions," *Vision Research*, vol. 35, no. 13, pp. 1917–1930, 1995.
- [23] S. M. Pizer, D. H. Eberly, B. S. Morse, and D. S. Fritsch, "Zoom invariant vision of figural shape: The mathematics of cores," *Comp. Vision Image Understanding*, vol. 69, no. 1, pp. 55–71, 1998.
- [24] B. S. Morse, S. M. Pizer, D. T. Puff, and C. Gu, "Zoom-Invariant vision of figural shape: Effect on cores of image disturbances," *Comp. Vision Image Understanding*, vol. 69, pp. 72–86, 1998.
- [25] J. D. Furst and S. M. Pizer, "Marching cores: A method for extracting cores from 3D medical images," in *Proc. Workshop Mathematical Methods Biomedical Image Analysis*, San Francisco, CA, June 1996.
- [26] D. Fritsch, S. Pizer, L. Yu, V. Johnson, and E. Chaney, "Segmentation of medical image objects using deformable shape loci," *Information Processing in Medical Imaging (Lecture Notes in Computer Science)*. Berlin, Germany: Springer-Verlag, 1997.
- [27] S. M. Pizer, D. S. Fritsch, P. A. Yushkevich, V. E. Johnson, and E. L. Chaney, "Segmentation, registration, and measurement of shape variation via image object shape," *IEEE Trans. Med. Imag.*, vol. 18, pp. 851–865, Oct. 1999.
- [28] G. Stetten, R. Landesman, and S. Pizer, "Core-Atoms and the spectra of scale," in *Proc. SPIE Medical Imaging Conf.*, 1997, vol. 3034, pt. 2, pp. 642–652.
- [29] M. Brady, *Criteria for Representations of Shape*. New York: Academic, 1983.
- [30] G. J. Clary, S. M. Pizer, D. S. Fritsch, and J. R. Perry, "Left ventricular wall motion tracking via deformable shape loci," in *Proc. CAR Int. Symp. Exhibition*, Berlin, Germany, 1997.