Automated Identification and Measurement of Objects via Populations of Medial Primitives, with Application to Real Time 3D Echocardiography

George D. Stetten and Stephen M. Pizer
Medical Image Display and Analysis Group, UNC, Chapel Hill.
NSF/ERC for Emerging Cardiovascular Technologies, Duke University.
updates available at **http://www.duke.edu/~stetten/IPMI**

Abstract

We suggest that identification and measurement of objects in 3D images can be automatic, rapid and stable, based on the statistical properties of populations of medial primitives sought throughout the image space. These properties include scale, orientation, endness, and medial dimensionality. The property of medial dimensionality is 0.0 for the sphere, 1.0 for the cylinder, and 2.0 for the slab, with non-integer dimensionality also possible. Endness results at the cap of a cylinder or the edge of a slab. The values of these medial properties at just a few locations provide an intuitive and robust model for complex shape. For example, the left ventricle during systole can be described as a large cylinder with an apical cap at one end, a slab-like mitral valve at the other (closed during systole), and appropriate interrelations among components in terms of their scale, orientation, and location. To look for such geometric components we extract populations of medial primitives we call *core atoms*. A core atom consists of two boundary points that face each other across an appropriate distance. A core atom may be grouped with others into a population by the location of its center, midway between its boundary points. Populations of core atoms thus grouped form a cloud around the medial manifold, and provide the basis for statistical analysis to yield a measure of the properties listed above. We demonstrate our method on simple geometric test objects, and show it capable of automatically identifying the left ventricle and measuring its volume *in vivo* using Real-Time 3D echocardiography.

Introduction

The lineage of the medial approach may be traced to the medial axis (otherwise known as symmetric axis or skeleton) introduced on binary images by Blum and developed by Nagel, Nackman, and others [1-3]. 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 locus in a space of position, radius, and associated orientations [4, 5]**.** Methods involving these continuous loci of medial primitives have proven particularly robust against noise and variation in target shape [6]. Determining locations with high medialness and relating them to a core has been accomplished by analyzing the geometry of loci resulting from ridge extraction [7]. Models including discrete loci of medial primitives have also provided the framework for a class of active shape models known as *Deformable Shape Loci* [8].

The objective of the work reported here is to build on these ideas to produce a method for analyzing the shape of the heart in Real Time 3D ultrasound, a new imaging modality that uses a matrix array of transducers elements to scan the moving heart in 3D at more than 20 frames/second [9]. The approach to analyzing this data aims to extract the scale, orientation and dimensionality (shape type) of sections of cardiac anatomy 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 by matching pairs of boundary points to form what are called *core atoms*. Core atoms tend to cluster along a medial ridge and allow for statistical analysis of the core and its underlying figure. Core atoms have already been developed for analysis of 2D shape [10] and are generalized here to 3D. The analysis is also extended to

spatially sampled populations of core atoms. The research presented here is part of a Ph.D. thesis nearing completion, which covers many of the aspects in greater detail [11].

What is a Core Atom?

A *core atom* is defined as two boundary points $\mathbf{b}_1$ and $\mathbf{b}_2$ that satisfy particular requirements (described in detail below) that guarantee 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. The core atom is said to be "located" at a *center point* midway between the boundary points (see Fig. 1). The *medialness* at the center point is high because the *boundariness* at both boundary points is high and because the boundary points face each other. Core atoms carry orientation, width and position, providing the ability for populations of core atoms to be analyzed in these terms.

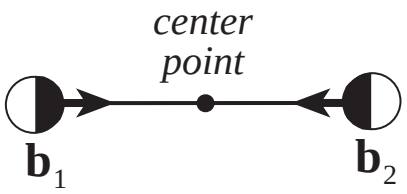

Fig. 1. A core atom consists of two boundary points that face each other across an acceptable distance, and a center point at which the core atom is said to be located.

Unlike medial models where *object angle* (half the angle between the lines from the center point to each respective boundary point) is permitted to vary, the object angle of a core atom is fixed at 90°. Core atoms thus follow in the tradition of Brady [12]. As in Brady, the underlying figure is not required to have parallel boundaries.

In the experiments presented below, boundariness is based on a Difference of Gaussian (DOG) measurement of intensity gradient, accomplished by repeated application of the binomial kernel. The number of applications determines the aperture of the boundariness detector, and is generally proportional to the size of the core atom. Further constraints are placed on the levels of intensity along the gradient direction. Other forms of boundariness, such as those based on texture analysis, would also be appropriate for core atoms, provided an orientation is established for the boundary point.

For efficiency, boundariness vectors are sampled on a rectilinear grid, and their magnitude compared to a threshold to select a population of boundary points $\mathbf{b}_i$ at locations $\underline{\mathbf{x}}_i$ with orientations $\vec{\mathbf{n}}_i$. The strength inherent to the statistics of populations is meant to counteract the weaknesses of the rapid thresholding. Core atoms are created from this population by finding pairs of candidate boundary points that satisfy the following three criteria:

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

$$\vec{\mathbf{c}}_{1,2} = \underline{\mathbf{x}}_2 - \underline{\mathbf{x}}_1 \qquad c_{\mathbf{min}} \le \left\|\vec{\mathbf{c}}_{1,2}\right\| < c_{\mathbf{max}}$$

The core atom vector can be oriented either way since the order of the boundary points is arbitrary. As mentioned already, $\left\|\vec{\mathbf{c}}_{1,2}\right\|$ is roughly proportional to the aperture used to find the boundaries.

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

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

("^" 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 towards (or away from) each other across the distance between them. A threshold for acceptable face-to-faceness is set within some error ε_f

$$\mathrm{F}(\mathbf{b}_1, \mathbf{b}_2) > 1 - \varepsilon_f$$

(3) Assuming $\mathrm{F}(\mathbf{b}_1, \mathbf{b}_2) > 0$, it follows 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.

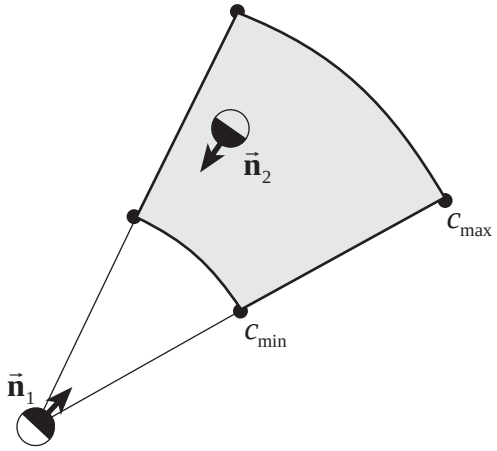

Fig. 2. The search area for the second boundary point (gray) is constrained by $\vec{\mathbf{n}}_1$, $c_{\min}$, and $c_{\max}$. The actual search area is 3D.

Although at first glance the search for pairs of boundary points appears to be $\mathrm{O}(n^2)$, hashing individual boundary points beforehand by location yields a large reduction in computation time. The search area for the *second* boundary point(s) is limited to a solid sector surrounding the orientation $\vec{\mathbf{n}}_1$ of the *first* boundary point, and to a range between $c_{\min}$ and $c_{\max}$. The width of the sector depends on ε_f (see Fig. 2).

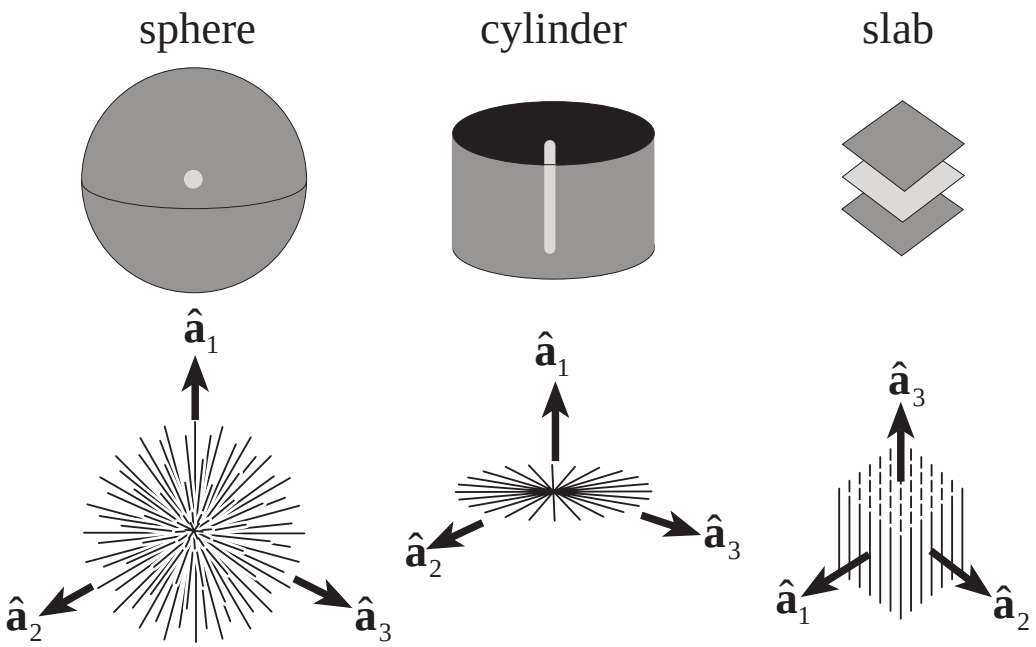

Fig. 3 - Three fundamental shapes (dark gray), corresponding cores (light gray), corresponding core atom populations (line segments) and eigenvectors $\hat{\mathbf{a}}_1$, $\hat{\mathbf{a}}_2$ and $\hat{\mathbf{a}}_3$.

Three Basic Configurations: Sphere, Cylinder, and Slab

Observe that collections of core atoms can group in three basic ways corresponding to the fundamental geometric shapes shown Fig. 3. The surfaces 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 being depicted as simple line segments.

The sphere generates a "Koosh ball" like cloud of core atoms with spherical symmetry, 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 core atom centers clustered along a plane in the middle of the slab. It is reassuring to find that the cores of these objects assume the most basic 3D shapes of all: the point, line, and plane.

As shown in Fig. 3, a system of shape- specific coordinate axes, namely $\hat{\mathbf{a}}_1$, $\hat{\mathbf{a}}_2$, and $\hat{\mathbf{a}}_3$, can be assigned in each case, although the axes are not all 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 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 $\vec{\mathbf{c}}_i$, $i = 1,2,3,....m$, the analysis of a core atom population begins by separating each core atom vector $\vec{\mathbf{c}}_i$ into its magnitude c_i and its orientation $\hat{\mathbf{c}}_i$. We ignore, for the moment, the location of the core atom. The analysis of magnitude c_i over a population of core atoms is straightforward, yielding a mean and standard deviation for the

measurement of width in the underlying figure. The orientation $\hat{\mathbf{c}}_i$ of core atoms in a population lend themselves 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 the population of m vectors 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 $\hat{\mathbf{c}}_i$.

$$\hat{\mathbf{a}}_1 = \arg\min_{\hat{\mathbf{a}}} \frac{1}{m}\sum_{i=1}^{m}\left(\hat{\mathbf{a}}\cdot\hat{\mathbf{c}}_i\right)^2 \qquad \text{which simplifies to} \quad \hat{\mathbf{a}}^{\mathbf{T}}\mathbf{C}\hat{\mathbf{a}} \qquad \text{where} \quad \mathbf{C} = \frac{1}{m}\sum_{i=1}^{m}\hat{\mathbf{c}}_i\hat{\mathbf{c}}_i^{\mathrm{T}}$$

The **C** matrix is positive definite, symmetric, and has a unit trace. Therefore, its eigenvalues are positive and sum to 1, and its eigenvectors are orthogonal. If the eigenvalues of **C** are sorted $\lambda_1 \le \lambda_2 \le \ldots \le \lambda_n$, the corresponding eigenvectors $\hat{\mathbf{a}}_1 \ldots \hat{\mathbf{a}}_n$ 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 eigenanalyis 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 $\hat{\mathbf{a}}_n$ is left as the least orthogonal to the population $\hat{\mathbf{c}}_i$, representing a form of *average* orientation for $\hat{\mathbf{c}}_i$. The axes $\hat{\mathbf{a}}_1 \ldots \hat{\mathbf{a}}_n$ are thus ordered from *codimensional* (orthogonal to the vector set) to *dimensional* (collinear with the vector set). In 3D, 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 itself as shown in Fig. 3.

The Lambda Triangle

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

$$\lambda_i \ge 0 \qquad \lambda_1 + \lambda_2 + \lambda_3 = 1$$

An eigenvalue of zero means that the corresponding eigenvector is perfectly orthogonal to every core atom $\hat{\mathbf{c}}_i$. 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 are completely orthogonal to the core atom population. Given the symmetries of the three basic shapes, the following eigenvalues result.

sphere	cylinder	slab
$\lambda_1 = 1/3$	$\lambda_1 = 0$	$\lambda_1 = 0$
$\lambda_2 = 1/3$	$\lambda_2 = 1/2$	$\lambda_2 = 0$
$\lambda_3 = 1/3$	$\lambda_3 = 1/2$	$\lambda_3 = 1$

Since λ_3 is dependent on the other two, the system may be viewed as having only two independent variables, λ_1 and λ_2. Because of constraints already mentioned, possible values for

λ_1 and λ_2 are limited by $\lambda_1 \le \lambda_2$ and $\lambda_2 \le \frac{1-\lambda_1}{2}$ which define a triangular domain we call the *lambda triangle* (See Fig. 4).

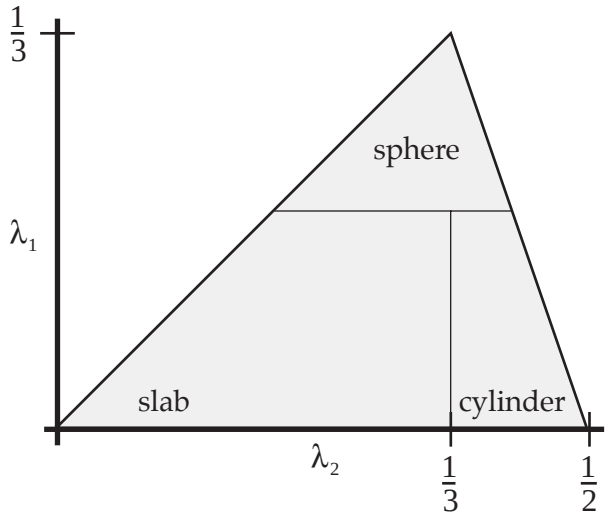

Fig. 4. The lambda triangle defines the domain of possible eigenvalues.

The vertices of the lambda triangle correspond to the three basic shapes in Fig. 3, with all possible eigenvalues falling 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 $\lambda_1 = 0.2$ and $\lambda_2 = 0.33\overline{3}$ 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 dimensionality is not an integer or even a single scalar, but rather two independent scalars, λ_1 and λ_2 whose values are constrained by the lambda triangle.

Spatial Sampling in the Corona

We now return to the question of core atom location, which we have so far ignored. To incorporate location into the analysis of core atoms, we separate them into bins on a regular 3D lattice by the location of their center points. Thus each bin represents a spatial sampling of medialness. The number of core atoms in a sample volume can be thought of as their *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 cloud may not be uniform. To preserve resolution, the sample volume may need to be smaller than the typical cross section of the core atom cloud. When a core is sampled off-center, it will demonstrate a distortion in its dimensionality. For example, the zero-dimensional core at the center of a sphere will appear to be one-dimensional (cylindrical) when sampled off center, as shown in Fig. 5. The vector from the theoretical core (center of the sphere) to the center of the core atom density in the sample volume is called the *displacement vector* $\vec{\mathbf{p}}$ (See Fig. 5C). The core atom population within a sample volume may not contain the entire thickness of the core, but rather a sub-sampling of the core called a *coronal density*. We can expect to be sampling coronal densities in general. It would be helpful know where in a given cloud around the core a sample was collected, but that presupposes knowledge about the overall distribution of core atoms.

We can predict certain relationships to exist between the distribution of core atoms over the entire core and that of a sample volume displaced from the center of the core. The population of core atoms will be flattened in a plane perpendicular to $\vec{\mathbf{p}}$, and thus develop orthogonality to that direction. This can be seen in Fig. 5, where the spherical distribution of core atoms in 5B has been flattened into a cylindrical distribution in 5C. The same effect can be seen in the case of the cylinder in Fig. 6, where, displaced off the central axis of the cylinder by $\vec{\mathbf{p}}$, the population of core atoms becomes slab-like and orthogonal to $\vec{\mathbf{p}}$. One expects the displacement vector to be one of the eigenvectors of any density at the closest point to 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 3D, the medial manifold can have at most 2 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 towards zero as the density develops orthogonality to $\vec{\mathbf{p}}$, except in the case of the slab, in which the core atom population can be expected to fall off rather abruptly without significant flattening. These expectations have been confirmed by experiments on graphical test figures as described in detail elsewhere [13]

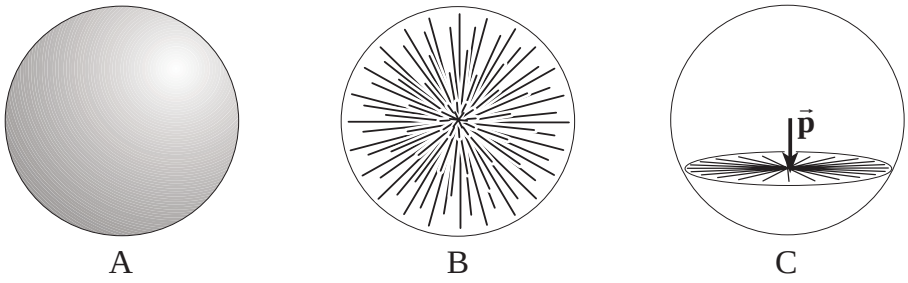

Fig. 5. **A.** spherical object. **B.** complete set of core atoms **C.** cylindrical coronal density displaced by $\vec{\mathbf{p}}$

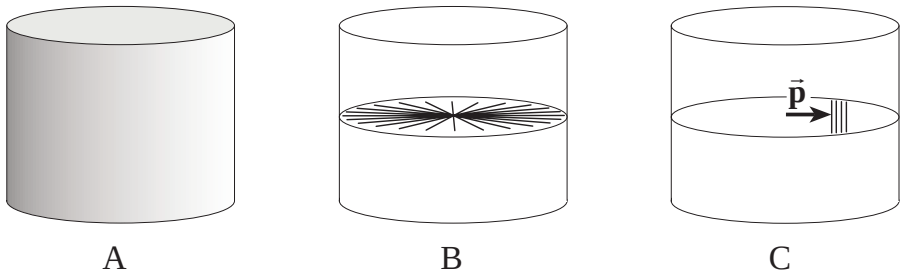

Fig. 6. **A.** cylindrical object. **B.** complete set of core atoms **C.** slab-like coronal density displaced by $\vec{\mathbf{p}}$

Figures 5 and 6 suggest that the displacement vector $\vec{\mathbf{p}}$ could somehow be used to compensate for the dimensional distortion in the corona. However, an isolated density that is, for example, cylindrical cannot know whether it represents the true center of a cylinder or simply the corona of a sphere. The results of the eigenanalysis for the density in a sample volume may be used in a Hough-like fashion simultaneously to vote for its own dimensionality and center of mass, as well as for possible higher-order dimensional densities whose corona it may inhabit. The voting takes place within ellipsoids around each sample volume. The axes of the ellipsoid are long in directions

orthogonal to the core atom population in its density. Thus the ellipsoid can be expected to extend in the $\vec{\mathbf{p}}$ direction, which is always orthogonal to the coronal density, .

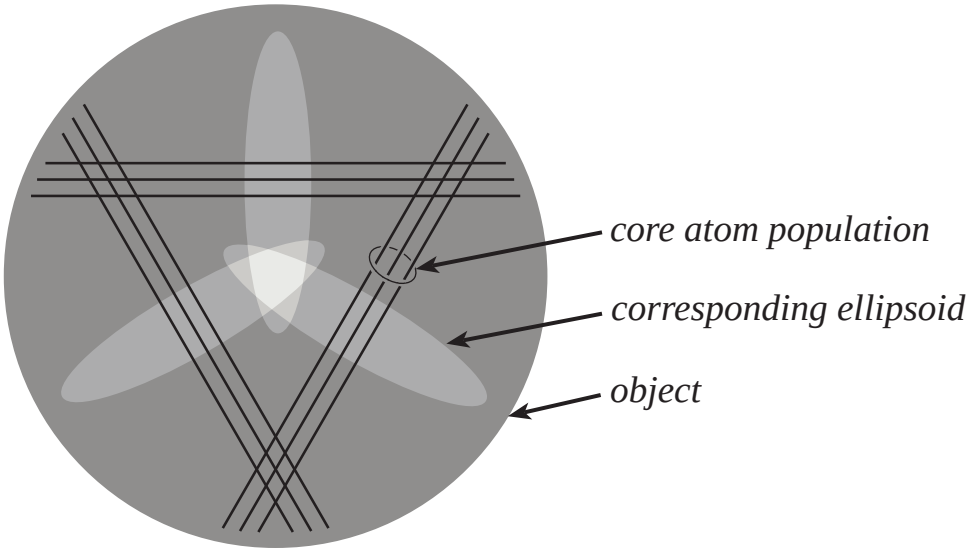

Fig. 7. Ellipsoids of three coronal core atom densities coalescing at the true center.

Fig. 7 demonstrates this concept. A circular cross-section through an object is shown with three coronal densities (each containing 3 core atoms) displaced from the center. An ellipsoid is associated with each density, 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. The figure can be interpreted as the cross-section of a spherical object 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 cylindrical object with the populations of core atoms being slab-like and the ellipsoids intersecting along the axis of the cylinder (as in Fig. 6).

There are various ways to construct such ellipsoids. We have chosen the following heuristic for its simplicity. The axes of our ellipsoid are the eigenvectors of the density's **C** matrix. The lengths of the axes are related to the eigenvalues λ_i as follows:

$$a_1 = \gamma \bar{c}$$

$$a_2 = \frac{\alpha_2}{\alpha_1} a_1 \qquad \text{where} \qquad \alpha_i = 1 - \lambda_i \qquad \text{and} \qquad \gamma = 1/2$$

$$a_3 = \frac{\alpha_3}{\alpha_1} a_1$$

The scalar distance $\bar{c}$ is the mean diameter of the core atoms in the density, and the dimensionless number γ 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, and in effect to coalesce the corona. Each sample volume (the *votee*) receives votes from all the neighboring sample volumes whose ellipsoids overlap it. The votes from those ellipsoids are assigned a strength v, where

$$v = m \cdot \exp\left(-d_e{}^2\right),$$

m is the number of core atoms in the voting density and d_e is an *ellipsoidal distance*

$$d_e = \sum_{i=1}^{3}\left(\frac{\left(\hat{\mathbf{e}}_i \bullet \vec{\mathbf{d}}\right)^2}{a_i}\right)$$

from the center of the voter ellipsoid to the votee, $\vec{\mathbf{d}}$ being the vector from the voter to the votee, and $\hat{\mathbf{e}}_i$ the cardinal coordinate system. Votes are constructed to contain information about the constituent core atom population of the voter, including its **C** matrix which may simply be summed (scaled by v) to produce an eigenanalysis of the entire constituent core atom population of a particular candidate. Thus are formed what we call *superdensities*, clusters of core atoms that no longer suffer from distortion. The center of mass for the constituent core atom population of a superdensity will tend to be at the true core, rather than in the corona.

Tests with Parametric Objects

To validate the methods described above, we apply them to 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. (The sphere is simply itself.)

Eigenanalysis of the coronal densities collected in a rectilinear lattice of sample volumes yielded the following results. Fig. 7 shows all densities containing greater than 1% of the entire core atom population plotted on the lambda triangle. The sphere shows two groups of densities, 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, which is locally a cylinder, shows clustering near the right (cylinder) vertex, with some spreading towards the left (slab) consistent with the dimensional effects of the corona predicted in Fig. 6. The spherical shell, which is locally a slab, shows tight clustering at the left (slab) vertex consistent with the observation that core atoms in a slab are collinear with $\vec{\mathbf{p}}$ and will therefore not develop significant orthogonality.

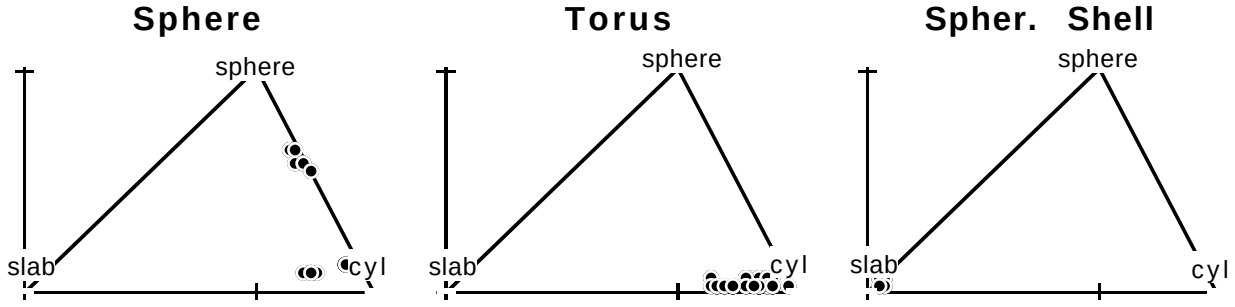

Fig. 8 Distribution of sampled densities on lambda triangles, for parametric test objects.

Unfortunately, Figure 8 does not show any spatial information about the sampled densities. To show this we resort here to color diagrams, with the understanding that IPMI proceedings are

historically printed without color. The reviewers are asked to accept this excursion as an attempt to demonstrate clearly the validity of the method, and with the understanding that black-and-white versions of these diagrams are being explored for the final manuscript. The color diagrams are included here in the text, beginning with Fig. 9, which maps an RGB color scheme onto the lambda triangle. Thus each possible dimensionality is assigned a unique color.

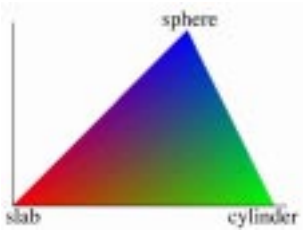

Fig. 9 Lambda Triangle with color coding.

The spatial distribution of densities for the test objects is shown Figure 10. Each sample volume whose density contains more than 1% of the total core atoms is shown as a thin-lined symbol. A simple partition of the lambda triangle (Fig. 4) was used to decide between three possible symbols: a slab is represented as a single red line, a cylinder as a green cross, and a sphere as 3 intersecting blue axes. The colors come from the map in Fig. 9. The length of the thin lines are constant, and chosen for clarity in each test object. The orientation of the thin lines indicates the predominant direction(s) of core atoms in each density, 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 arbitrary rotation around the axis.

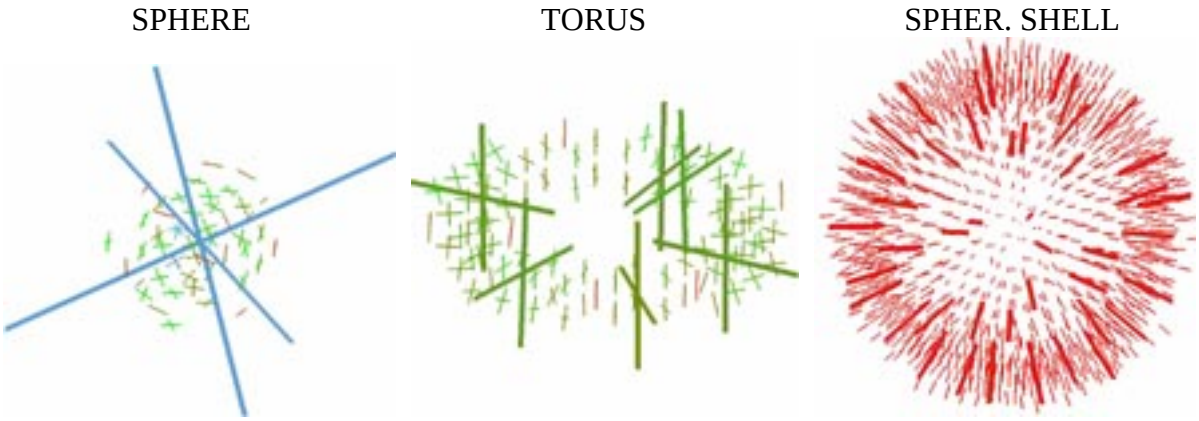

Fig. 10 Densities and superdensities for parametric objects (see Fig. 9 for color scheme).

As expected the sphere shows cylindrical densities in its corona oriented towards the center (see Fig 10). Further out from the center a few slabs-like densities reflect simply the paucity of core atoms in those sample volumes. Near the center one true spherical density (a small blue 3-axis symbol) may be discerned.

The thick-lined symbols show the results of ellipsoidal voting, i.e., they represent superdensities. To prevent a cluttered illustration, votes have been tallied for all densities, a single winner declared, and then all votes for other candidates removed from its constituency. Using this Hough credit attribution strategy, the election is then repeated until all votes are gone. These chosen superdesities thus represent non-overlapping constituencies, and are easier to see in the

illustration. They are represented by thick lines in a manner similar to the densities, except the length of the axes now corresponds to the actual mean scale of the constituent core atoms. Thus the thick-lined blue 3-axis cross are actual diameters of the spherical object. For the sphere there is only one predominant winning superdensity; with virtually every core atom in its constituency.

The torus shows cylindrical densities properly oriented but dispersed throughout the corona. At the outer regions of the corona a few slab-like densities are visible. The superdensities, however, are centered on the circular mid-line of the torus.

The spherical shell shows only slab-like densities, which coalesce with ellipsoidal voting into slab-like superdensities. The orientation of both are across the local slab. Here ellipsoidal voting is seen to perform another function, that of connecting densities that share a core. The same will happen along the axis of a cylinder.

Endness

Some attention must be paid to fundamental geometric forms where a cylinder ends at a hemispherical cap, or a slab ends at a hemicylindrical edge. The property of endness has been described by Clary, et al.[14]. Endness as viewed from the core atom perspective is illustrated in Fig. 11.

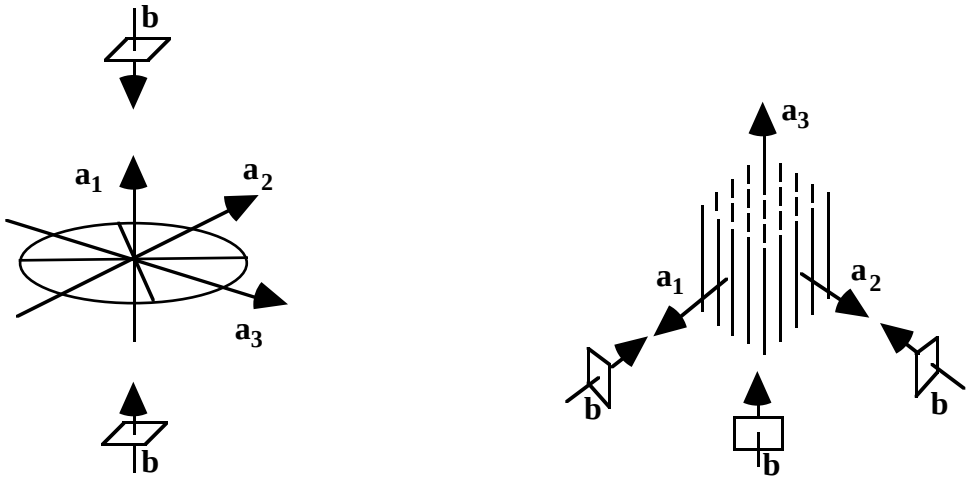

Fig. 11 Endness, manifested as a cap on a cylinder (A) and the edge of a slab (B).

To detect endness, densities of core atoms can be used as a starting points. Once a local cylinder has been established, 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, boundary points indicating an edge to the slab can be sought. A mathematics for this is derived elsewhere [11].

Identifying and measuring the Cardiac Left Ventricle

We now turn to a useful clinical application, the automated determination of left ventricular volume using Real Time 3D (RT3D) echocardiography. RT3D is a new imaging modality that electronically scans a volume in 3D using a matrix array (Fig. 12a) instead of the conventional linear array (Fig. 12b). The resulting data "pyramid" can be sliced into sets of orthogonal "B-mode" images or "C-mode" images that are parallel to the transducer. RT3D is described in detail

elsewhere [9], but its primary novelty is the ability to capture a single cardiac cycle at 22 frames/second, which no other available imaging modality can accomplish.

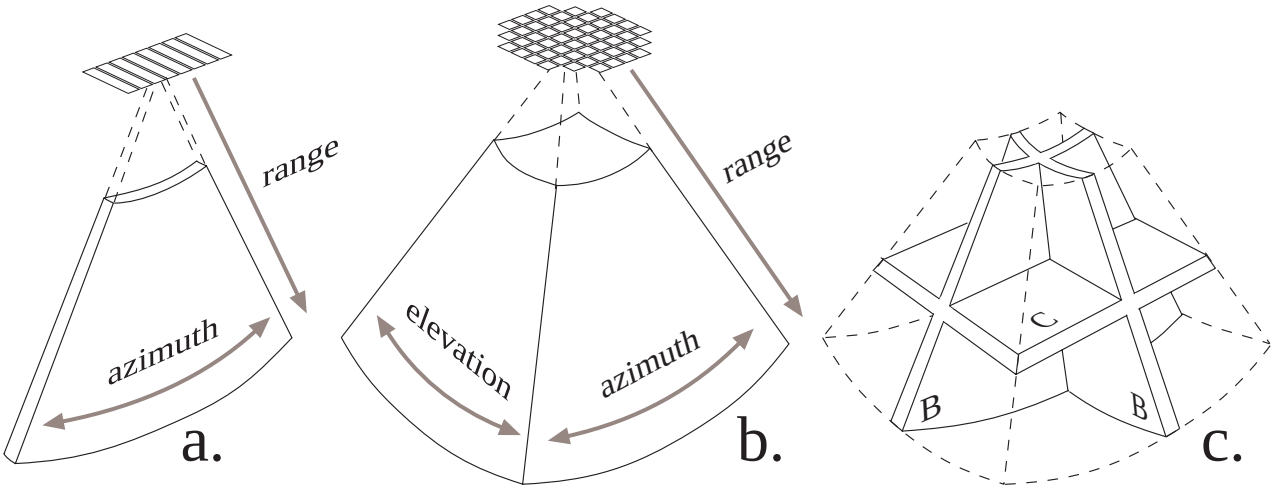

Fig. 12. Real Time 3D echocardiography uses a matrix array to scan a 3D volume.

RT3D images of an *in vivo* human heart present a significant challenge to image analysis techniques, including high noise, low resolution, path dependence, and a non-rectilinear data space. These problems are addressed elsewhere [11], but the suggestion that the statistical nature of our method yields robustness is severely tested in its application to RT3D echocardiography.

We now expand on the example from the abstract: The left ventricle during systole is basically a large cylinder with an apical cap at one end, and a slab-like mitral valve at the other (we limit ourselves here to apical scans, and to times when the mitral valve is closed). The model is shown in Fig. 13.

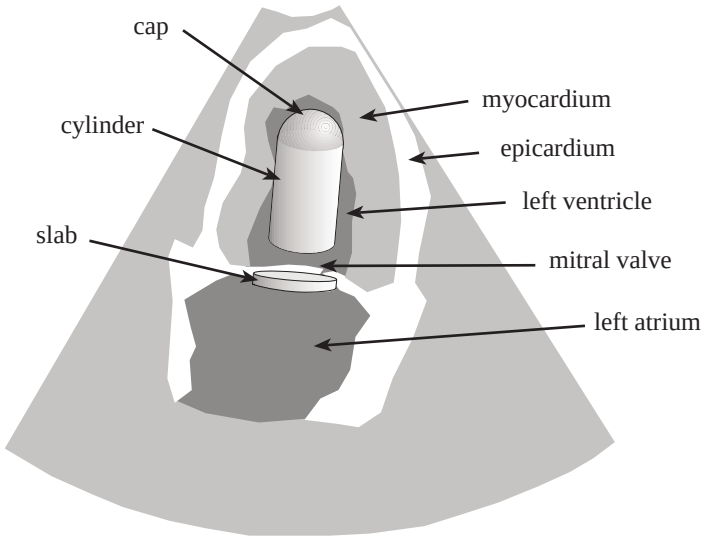

Fig. 13 Model of left ventricle as cylinder with cap and mitral valve as slab.

To identify the cylinder in the image data, core atoms of an appropriate range of diameters were collected in sample volumes on a regular lattice, and ellipsoidal voting was applied. An example of the resulting superdensities is displayed in Fig. 14b. Green crosses are shown in the cylindrical chamber of the ventricle. Due to the pre-selection of core atoms by scale, no other significant densities of core atoms were found.

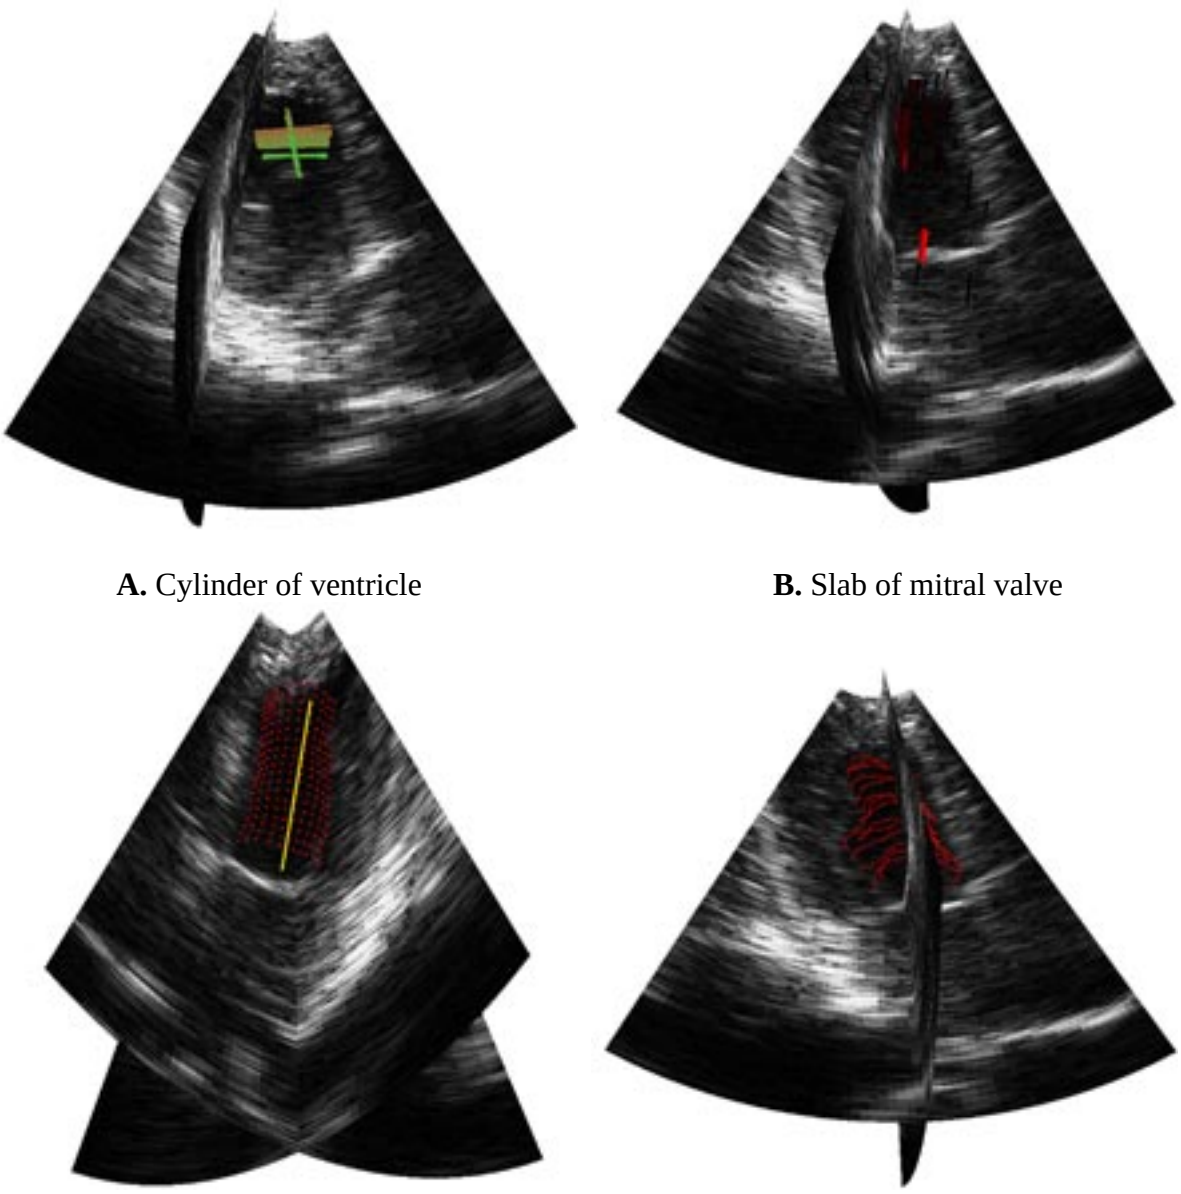

A. Cylinder of ventricle **B.** Slab of mitral valve

C. Automated axis and surface map **D.** Manual tracings of ventricle

Fig. 14 Real Time 3D ultrasound with automated and manual identification of left ventricle.

Next, the mitral valve was sought, by limiting the formation of core atoms to an appropriately smaller scale, and to orientations nearly parallel to the transducer. As shown in Fig. 14a, the

strongest superdensities (short red lines) were clustered around the center of the mitral valve, although weaker false targets were detected in the myocardium. To eliminate these false targets, a criterion was established for the formation of appropriate pairs of superdensities, in the spirit of core atoms. Only slab-like densities appropriately located and oriented with respect to cylindrical densities were accepted. These pairs were allowed to vote for their constituent superdensities, and the mean location of the winning superdensities used to establish a single mitral valve location and a single LV cylinder location. The vector between these two locations was used to establish a cone for expected boundary points at the apex of the LV, and the mean distance to the resulting boundary points used to determine the location of the apical cap along that vector. Thus an axis between the apex and the mitral valve was established

Given the axis between the apex and the mitral valve, LV volume was estimating by collecting boundary points around the axis. Only boundaries that faced the axis were accepted. The boundary points were organized into bins using cylindrical coordinates, in other words, disks along the axis and sectors within each disk. An average radius from the axis was established for the boundary points in each bin, creating a *surface map*. Fig. 14c shows such a surface map (red dots) and the underlying axis (yellow). The problem of empty bins was avoided by convolving the surface map with the binomial kernel in 2D until each bin had some contribution to its average radius. Volumes were then calculated by summing over all sectors. The entire procedure including identification and volume measurement of the LV was automated, and required approximately 15 seconds on a 200 MHz Silicon Graphics O2 computer.

The automated volumes were compared to manual tracings performed on a stack of flat slices orthogonal to a manually-placed axis (see Fig 14d). This axis employed the same anatomical end-points (the ventricular apex and the center of the mitral valve) as the axis determined automatically above. The volumes and locations of the end-points were compared to those determined automatically. Results are shown in Fig. 15. They are very encouraging, particularly the automated placement of the axis end points, which had an RMS error of approximately 1 cm. Volume calculations introduces additional errors of their own, but were still reasonable for ultrasound. Only four cases have been tried, and all are shown. The method worked in all cases. Further experiments are immediately forthcoming and would of course be included in a final IPMI paper.

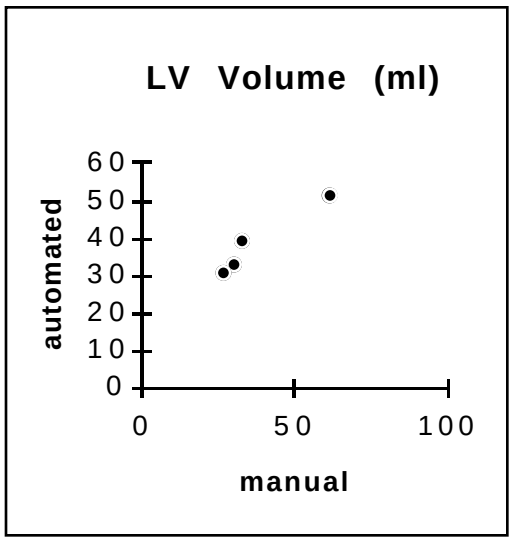

RMS error
Axis End-Points
and LV volume

LV apex	1.3 cm
Mitral Valve	0.98 cm
LV volume	6.6 ml

Fig. 15. Volume measurement and axis end-point location, manual vs. automatic.

Conclusions

We have described a new method for identifying anatomical structures using fundamental properties of shape extracted statistically from populations of medial primitives, and have demonstrated its feasibility by applying it under challenging conditions. Further studies are presently underway to establish its reliability over a range of data. Future directions include introducing greater specificity and adaptability in the boundary thresholding, incorporating more than 2 nodes into the model, introducing variability into the model to reflect normal variation and pathologic anatomy, extending the method to the spatio-temporal domain, and applying it to visualization.

Acknowledgments

Supported through a Whitaker Biomedical Engineering grant, NSF grant CDR8622201 NIH grants 1K08HL03220, P01CA47982, and HL46242, with data supplied by Volumetric, Inc.

References

[1] H. Blum and R. N. Nagel, "Shape description using weighted symmetric axis features," *Pattern Recognition*, vol. 10, no. pp. 167-180, 1978.

[2] L. R. Nackman, "Curvature relations in 3D symmetric axes," *CVGIP*, vol. 20, no. pp. 43-57, 1982.

[3] L. R. Nackman and S. M. Pizer, "Three-Dimmensional shape description using the symmetric axis transform I:Theory," *IEEE Trans. PAMI*, vol. 2, no. 2, pp. 187-202, 1985.

[4] 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.

[5] S. M. Pizer, D. H. Eberly, B. S. Morse and D. S. Fritsch, "Zoom invariant vision of figural shape: the mathematics of cores.," *Computer Vision and Image Understanding*, vol. 69, no. 1, pp. 55-71, 1998.

[6] B. S. Morse, S. M. Pizer, D. T. Puff and C. Gu, "Zoom-Invariant vision of figural shape: Effect on cores of image disturbances.," *draft, Department of Computer Science, UNC Chapel Hill.*, vol. no. pp. 1996.

[7] J. D. Furst and S. M. Pizer, "Marching Cores: A method for extracting cores from 3D medical images.," Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis, San Francisco, CA, June 1996.

[8] 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, Poultney, Vermont, 1997.

[9] G. Stetten, T. Ota, C. Ohazama, C. Fleishman, J. Castelucci, J. Oxaal, J. Kisslo and O. T. v. Ramm, "Real-Time 3D Ultrasound: A New Look at the Heart," *J. Cardiovasc. Diagnosis and Proc.*, vol. 15, no. 2, pp. 73-84, August, 1998.

[10] G. Stetten, R. Landesman and S. Pizer, "Core-Atoms and the spectra of scale, (Technical Report TR97-006, UNC Dept. Computer Science," SPIE Medical Imaging Conference, 1997.

[11] G. Stetten, Automated Identification and Measurement of Cardiac Anatomy Via Analysis of Medial Primitives, Ph.D. Thesis, Stephen Pizer, advisor, UNC Chapel Hill., (exected, March 1999).

[12] M. Brady, Criteria for representations of shape., New York, Academic Press, 1983.

[13] G. Stetten and S. Pizer. Extracting Shape Properties via Populations of Medial Primitives. University of North Carolina, Department of Computer Science, (or www.stetten.com), 1998.

[14] G. J. Clary, S. M. Pizer, D. S. Fritsch and J. R. Perry, "Left ventricular wall motion tracking vi deformable shape loci," CAR International Symposium and Exhibition, Amsterdam, 1997.