Chapter 9 Conclusions

This dissertation has described a new method for using local medial properties of shape to identify and measure anatomical structures. A *bottom up* approach based on image properties has been presented, combining statistics with geometry at an early stage in the analysis. The approach differs from *top down* approaches, such as deformable surfaces, in that initial placement of a model is not required. A model is still applied, but at a higher, more abstract, level, where relatively few and robust features permit a Hough-like voting strategy to replace the iterative search in deformable models. Given the automated placement of a local medial framework, statistical determination of boundary location provides geometric measurements without requiring explicit boundaries, in keeping with the statistical nature of core atoms.

The methods put forth here show encouraging results under challenging conditions and are efficient compared to other techniques for 3D analysis. Methods involving *core atoms*, *core ions*, and *medial node models* have been implemented and tested in terms of their effectiveness at identifying and measuring graphical test objects, balloons, and left ventricles.

This chapter re-examines the results in light of the original thesis and the claims made in Chapter 1, discussing the strengths and weaknesses of the new methods. It concludes with an exploration of possible future directions for the research.

9A. Claims Revisited

This section reviews the claims made in Chapter 1, verifying that each has been addressed.

Claim #1. Automated measurement of ventricular volume is possible in RT3D echocardiographic data using a statistical analysis of voxel intensity and location with respect to a medial node model (described in Claim #2).

The original purpose of core atoms included direct measurement of geometrical parameters. Since each core atom represents a distance measurement, populations of core atoms can, under certain circumstances, be analyzed statistically to produce measurements such as slab thickness, cylinder diameter, or sphere diameter. This ability was demonstrated on the graphical test objects (Section 4C) and balloons (Section 8B). When applied to the LV, core atoms were capable of finding the LV axis, but tended to overestimate LV volume because they formed at both the epicardial and endocardial boundaries (Section 8D).

An alternate approach used the medial analysis to provide position, orientation, and size parameters for a fuzzy segmentation of voxels by location and intensity (Sections 7B and 8E). The fuzzy segmentation assigned a probability to each voxel of its inclusion in the LV, yielding stable results without requiring an explicit boundary location. The need for the fuzzy segmentation approach reflects more the difficulty in distinguishing the endocardial border in the RT3D ultrasound images (even by a human observer) than a fundamental problem with core atoms. Core atoms appear to remain a viable approach for direct measurement in other applications.

The fuzzy segmentation method distinguishes itself from standard methods that produce a discrete list of positions in the interior or at the boundary of an object. The fuzzy method generates geometric information, instead, from statistical analysis of populations of local probabilities. Any use of the fuzzy segmentation must therefore require neither an explicit surface nor a distinct topology. The results from the LV suggest that fuzzy segmentation may produce a more robust measure of volume and boundary location than other methods, especially when stabilized by an underlying statistical analysis of medial properties using core atoms, and especially in applications such as RT3D echocardiography where boundaries do not always exist. These results, however, leave much to be desired. Further refinements such as optimizing the segmentation over multiple iterations or determining parameters by training on populations of hearts would be helpful to see if improved results could be achieved.

In both the direct use of core atoms and the fuzzy segmentation of voxels, a general approach to geometric measurement has been explored in which measurements are organized by the medial manifold, converting the cardinal coordinate system into a radial/tangential one. This coordinate system reflects the local shape and permits sorting boundary points, or the voxels themselves, along normals to the postulated boundary. The resulting measurements from the balloons show that high accuracy is possible when conditions are favorable, and those from the LV demonstrate robust behavior under challenging conditions.

Claim #2. Automated identification of anatomical structures is possible in RT3D echocardiographic data using correspondences between a medial node model of the heart and a set of medial measurements extracted empirically from the image data (described in Claim #3).

A theory was developed in Chapter 5 for matching measurements of medial properties to nodes in a model. Measurements at individual locations were matched to individual nodes, as well as pairs of locations to pairs of nodes. The resulting medial node model thus included expected geometric relationships between local medial properties as well as the expected measurement of the local properties themselves.

Methods based on this approach were verified empirically on RT3D ultrasound data in two separate experiments. Fluid-filled balloons were identified using a medial node model with a single node, demonstrating the ability of core atoms to identify a roughly spherical target in an uncluttered background (Section 8B). A medial node model containing 3 nodes was demonstrated to be capable of automatically identifying the apex-to-mitral-valve axis of the human LV *in vivo* (Section 8A).

Claim #3. The medial properties of scale, orientation, dimensionality, and endness can be extracted from cardiac image data by clustering regularly sampled populations of medial primitives (described in Claim #4).

A theory for clustering samples of core atoms to reflect the underlying core was developed in Chapter 4. It derived from understanding the dimensional distortion that occurs in the corona where a sample is displaced from the center of the core. The method of clustering core atoms was demonstrated to reduce sampling artifacts in graphical test images and to yield anatomically correct scale, orientation and dimensionality in the cylinder of the LV and the slab of the mitral valve. The clustering technique provides the added benefit of reinforcing samples that share a local core, thus providing mutual support between multiple samples along the axis of the LV and at the center of the mitral valve.

Pairs of clusters were shown capable of providing a vantage point for measuring endness at the apex of the LV, each pair containing one cylindrical cluster from the LV and one slab-like cluster from the mitral valve. The stability of pairs of clusters provided sufficient constraint to orientation such for core ions to find the apex of the LV. The orientation of individual clusters was found to be less stable for measuring endness, at least in the context of the LV. Other applications may be able to make use of endness measurements, however, from individual core atom clusters.

Claim #4. A set of candidate boundary points can be collected in an image and medial primitives formed from pairs of these boundary points with appropriate relative distance and orientation. The medial primitives can be sorted by location to provide for statistical analysis of local medial properties.

The formation of medial primitives (core atoms) was demonstrated on graphical test objects both in 2D and 3D. The center points of core atoms were shown to provide a statistical measure of the presence of a medial ridge. In 2D, sorting sub-populations of core atoms by their scale was shown capable of specifying orientation and dimensionality as a function of scale in uncluttered targets. In 3D, localizing core atom populations demonstrated the capability of providing local measurements of medialness and its properties of scale, orientation, and dimensionality in the presence of clutter.

Claim #5. Evaluation of these methods on parametric test objects, fluid-filled balloons and *in vivo* human hearts establishes the ability of the method to match a model to the image data and determine the volume of the underlying anatomical structure.

Evaluation with parametric test objects was based on foreknowledge of the geometric parameters. For the fluid filled balloons, evaluation was based on measurement of balloon weight from which volume was calculated. Evaluation for the human hearts employed three sets of manual traces on stacks of slices orthogonal to a manually identified AMV axis. From these traces, localized boundary error as well as volume error were computed.

9B. Strengths of this Approach

The nature of RT3D ultrasound data makes it particularly applicable to critical situations where changes in heart function are monitored from one minute to the next. For the analysis of

RT3D echocardiographic data to be clinically useful under these circumstances, a reasonably fast performance is required. One set of manual tracings on the 155 frames of *in vivo* data used in this dissertation required approximately 4 hours of intensive labor. At 22 frames per second, that much data is collected by the RT3D scanner every 7 seconds. By comparison, the automated analysis on a 400 MHz Pentium computer required approximately 10 seconds per frame, or about 3 minutes for a single cardiac cycle. This is still about 200 times slower than real-time, but within the realm of clinical utility. The inevitable increase in computer speed, further optimization of the algorithms, and implementation in hardware could be expected to make real-time analysis possible in the not-too-distant future.

The core atom technique differs from other methods of finding the medial manifold in that isolated boundary points are found first and then paired with appropriate mates. This permits both the gathering of core atoms from the entire data space in an initial sweep, and the subsequent statistical analysis to be completed in a short amount of time. Other methods, such as ridge tracking or deformable models that look for both boundaries simultaneously, must cover multiple orientations and scales. They therefore typically require initialization to localize the search if reasonably short computation times are to be achieved.

Core atoms owe their robustness to statistical analysis and to the fundamental nature of the shape measurements being made. In particular, dimensionality, scale, relative orientation, and endness are not easily mutated by variations in anatomy. In the example here, almost any ventricle is roughly a cylinder within a particular range of scales and orientations. There simply are no other large dark cylinders in the neighborhood.

The robust nature of the method has been demonstrated by its application to a particularly difficult data set. The subjective nature of the manual tracing underscores this appraisal of the poor quality of the RT3D data. The correlation between manual traces speaks primarily to the operator's knowledge of expected cardiac anatomy, since the endocardial border was often impossible to distinguish in the ultrasound data.

9C. Weaknesses of This Approach

A major weakness in the method would appear to be the pre-thresholding of boundary points. Once chosen, stronger boundary points are no longer differentiated from weaker ones. Once passed over by the initial thresholding, a boundary point is lost forever even though its boundariness may have fallen just below the threshold.

No provision presently exists for adapting the boundary thresholding criteria to an individual image or to a particular object in the image. Obviously, real boundaries come in many different types. A single object may even have different types of boundaries on opposite sides. For example, the myocardium is basically a slab with the epicardium on one side and the ventricle on the other. Ideally, core atoms designed to find the myocardium should have two different types of boundary points at either end. Such developments are possible, but beyond the scope of this dissertation.

A weakness in the medial-based approach in general is that both boundaries must be present. If one of the boundaries is occluded or out of range, medialness will not be high. Core atoms did surprisingly well with some of the LV images in which much of the ventricular border was out of range, apparently because the statistical nature of the analysis allowed the cylinder of the ventricle be extrapolated from those boundary pairs which did exist. It is less clear how to improve the model so that it would be able to extrapolate entire boundaries out of range the way a human observer can.

Performing the manual traces for this dissertation made it clear, in general, just how simplistic the particular medial node model in this dissertation is, compared to human capabilities for image analysis. Multiple clues were employed during the manual tracing beyond what the human operator understood explicitly, for example, seeing the mitral valve in arbitrary cross-section, understanding the anatomy of the aorta relative to the mitral valve, and knowing where to put the boundary when there simply was none. The human prior is far more sophisticated, flexible, and reliable than the prior used by the methods in this dissertation.

9D. Future Directions

Perhaps the most crucial improvement needed for core atoms in their present state is to increase the flexibility of initial boundariness detection. Ideally, a system of multiple adaptable boundary profiles would extend from the core atom framework back into the image, permitting identification of local boundary type in conjunction with medial properties.

Other future directions in this particular research involve constructing more complicated medial-node models for the heart. One improvement would identify the "crux" of the heart, i.e., the intersection of the annulus and the septum. This would free the model from depending on the mitral valve with the dual benefit of permitting analysis throughout the entire cardiac cycle, thus freeing the mitral valve to be analyzed independently. The crux is one of the strongest visual landmarks for human operators, being the central structure of the familiar "four chamber" view. It represents the mechanical center of the heart from which forces are exerted and as such may be expected to be stable.

An extension of the method into the spatio-temporal domain would permit a model to match the entire cardiac cycle instead of a single frame. Core atoms themselves could be made spatio-temporal, permitting them to correlate across time or track moving objects.

Many steps in the methods presented here contain parameters whose values were set ad hoc. These could undoubtedly be optimized in terms of the overall ability of the method to identify and measure specific anatomical structures. Optimizing these parameters will require larger data sets that reflect the general population of normal and pathologic hearts. These larger populations will, in turn, require variability in the model to encompass and distinguish between the many different types of hearts in the clinical population.

A fair test of the capabilities of the core atom approach will include its application to other imaging modalities besides RT3D ultrasound. MRI, CT, and PET offer greater constancy in voxel intensity and less noise than ultrasound, while avoiding the issues of non-rectilinear coordinates. A wide range of applications is possible for core atoms outside of medical imaging, both in 2D and 3D, such as computer vision and laser-based range imaging.

Core atoms may prove useful in the initialization of models or other geometrical computations such as the LV surface model in this dissertation, permitting presently semi-automated algorithms to become fully automated. In some cases, global geometric properties may be measured by core atoms without segmentation, such as the size and anisotropy of cell populations imaged by 3D modalities operating at microscopic scales.

There is great promise in applying the core atom approach to new methods of visualization. Core atoms provide a statistical method of establishing boundary normals for lighting purposes and for determining which of two surfaces in an object is closer to the viewer for removal of one or the other surface.

Finally, future research in core atoms should include developing methods to automatically create medial node models directly from training sets, perhaps without any, or only minimal, human assistance. It is possible that the properties of shape that best identify a particular anatomical structure may indeed be different from those a human would consider obvious. Automating the construction of medial node models would facilitate their application to the expanding repository of medical images, thereby expediting the discovery of correlations between images and diseases and leading to further understanding of the diseases themselves.