

GRADIENT-ORIENTED BOUNDARY PROFILES FOR SHAPE ANALYSIS USING
MEDIAL FEATURES

by

Robert Joseph Tamburo

BS, Delaware State University, 1999

Submitted to the Graduate Faculty of
University of Pittsburgh in partial fulfillment
of the requirements for the degree of
Master of Science

University of Pittsburgh

2002

UNIVERSITY OF PITTSBURGH

SCHOOL OF ENGINEERING

This dissertation was presented

by

Robert Joseph Tamburo

It was defended on

November 25, 2002

and approved by

Robert J. Sclabassi, MD, PhD,
Professor, Dept. of Neurological
Surgery and Neuroscience

Ching-Chung Li, PhD, Professor,
Dept. of Electrical Engineering

George Stetten, MD, PhD, Assistant
Professor, Dept. of Bioengineering
Thesis Advisor

ABSTRACT

GRADIENT-ORIENTED BOUNDARY PROFILES FOR SHAPE ANALYSIS USING MEDIAL FEATURES

Robert Joseph Tamburo, MS

University of Pittsburgh, 2002

Gradient-oriented boundary profiles have been developed as a novel method to parameterize boundaries. Boundary profiles are created at locations of high gradient magnitude by averaging intensity within a neighborhood of voxels oriented along the image gradient, making them rotationally invariant and relatively insensitive to image noise. A cumulative Gaussian is fit to the collection of averaged voxel intensities yielding estimates of (1) extrapolated intensity values for voxels located far inside and outside of a boundary and (2) anatomical boundary location. Intrinsic measures of confidence have been developed to eliminate low-confidence parameter estimates. Thresholds placed on these measures of confidence allow for high-confidence unsupervised classification of boundaries.

The validity of gradient-oriented profiles is demonstrated on artificially generated three-dimensional test data and shown to accurately parameterize and classify the boundary. Applying the measures of confidence and establishing thresholds, the accuracy of boundary location and intensities estimates improved drastically, making them a high-quality replacement for simpler methods of boundary detection. Towards shape analysis, gradient-oriented boundary profiles are applied to an existing a medial-based approach to shape analysis, known as core atoms. Core atoms in their previous implementation were based on simple gradient direction and unable to form without a priori knowledge of object intensity relative to background. Boundary profiles were applied to core atoms

permitting the formation of so called “core profiles”. Core profiles remove any restriction on the object’s or the background’s intensity, allowing multiple objects of differing intensities to be located with a single application.

Core profiles were applied to 3D computer-generated data, as well as RT3D ultrasound cardiac phantom data. It was shown on computer-generated data that calculating the volume with core profiles is more accurate than calculating the volume with core atoms, because of the improved accuracy of the boundary location. Two new methods of automatically measuring volume on non-parametric data with core profiles are proposed. Future work with includes constructing medial node models improved by gradient-oriented boundary profiles for automated left ventricular identification and measurement.

PREFACE

“Be courteous to all, but intimate with few, and let those few be well tried before you give them your confidence. True friendship is a plant of slow growth, and must undergo and withstand the shocks of adversity before it is entitled to the appellation.” -George Washington

Throughout my Master’s degree graduate studies, several people have played a large role in my completion. I will be forever grateful to my advisor, Dr. George Stetten, for his guidance and encouragement. It was his dedication, patience, and willingness to help that certainly contributed to the completion of my thesis research. I am also thankful to Dr. C. C. Li and Dr. Robert Sclabassi for taking the time to serve on my committee, as well as providing me with insightful comments and suggestions.

I also wish to thank Dr. Harvey Borovetz, for it was his presence and menacing stare in the hallways that sent me fleeing to the lab and Lynette Spataro for her sincerity to help and answering all of my questions. I am also very appreciative to the members of VIA lab for their assistance. Each member of the lab possesses unique subtleties that helped me along the way; Aaron Cois’ programming persistence, Damion Shelton’s quirky sense of humor, and Wilson Chang’s no-nonsense attitude.

Most importantly, I’d like to thank all of my family and friends for love and support. Mom, thanks for being my role model and supporting me throughout my life. Dad, thanks for being so encouraging and caring throughout our years together. Erica, thanks for being such a special and loving person. Last but not least, I’d like to thank all of my friends at #orp who kept me awake and entertained during the early morning hours.

TABLE OF CONTENTS

	Page
ABSTRACT	iii
PREFACE	v
TABLE OF CONTENTS	vi
LIST OF TABLES	viii
LIST OF FIGURES	ix
1.0 INTRODUCTION	1
1.1 Goals and Clinical Motivation	1
1.2 Approach to Automating Image Analysis	3
1.3 Thesis Overview	5
1.4 Guide to Remaining Chapters	5
2.0 BACKGROUND	6
2.1 Boundary Detection	6
2.2 State of Medical Imaging	8
2.3 Methods to Detect and Measure LV Volume in Ultrasound Images	9
2.4 Medial Shape Analysis	10
3.0 GRADIENT-ORIENTED BOUNDARY PROFILES	11
3.1 Finding Boundary Candidates	12
3.2 Generating a Boundary Profile	12
3.3 Fitting the Intensity Profile and Estimating Parameter Values	16
3.4 Eliminating Bad Profiles	18
3.5 Establishing Intrinsic Measures of Confidence	19
3.6 Classifying the Boundary	20

4.0 ASSESSMENT OF BOUNDARY PROFILES ON A SYNTHETIC SPHERE	21
4.1 Testing Sampling Regions and Splat Types on a Sphere	22
4.2 Measuring the Accuracy of Intensity Estimates by Gradient-Oriented Profiles.....	27
4.3 Judging Intrinsic Measures of Confidence.....	29
5.0 BOUNDARY PROFILES APPLIED TO CORE ATOMS	32
5.1 Core Profiles	32
5.2 Measuring Medial Properties With Core Profiles: Medial Densities	35
6.0 CORE PROFILE TEST DATA	39
6.1 Measuring the Accuracy of Gradient-Oriented Profiles.....	43
6.2 Analysis of Core Profile Population.....	46
6.3 Measures of Volume	47
7.0 MEASURING VOLUME WITHOUT GEOMETRICAL ASSUMPTIONS WITH CORE PROFILES ON CONCENTRIC ELLIPSOIDS	49
8.0 APPLICATION OF METHODS TO RT3D ULTRASOUND DATA	54
9.0 INSIGHT TOOL KIT CONTRIBUTION.....	59
10.0 FUTURE WORK	60
10.1 Improving the Performance of Gradient-Oriented Boundary Profiles	60
10.2 Medial-Node Models	61
11.0 CONCLUSION.....	62
BIBLIOGRAPHY	65

LIST OF TABLES

Table 1 Voxel intensity weights for the four bins of a triangle footprint.....	14
Table 2 The voxel intensity weights for the four bins of a Gaussian footprint	15
Table 3 Number of profiles found, eliminated, and total time for discovery and construction.....	23
Table 4 Number of profiles eliminated using a threshold of $z_{\min} > 1.5$	26
Table 5 Effects of neighborhood and splat type on determining the radius of the sphere accurately	27
Table 6 An average core profile scale after sorting them by scale permits the calculation of LV and myocardial volume. The true volumes are found simply by using the equation of a sphere and the known radii of the test data. The percent errors (PE) for volume measure and standard deviation (σ) of the core atom and profile lengths are shown in parentheses.	48

LIST OF FIGURES

Figure 1 An ellipsoidal neighborhood whose center is the boundary point from the DoG gradient detector and major axis is the gradient vector \vec{g} of the boundary point. Two adjoining disks, A and B are shown for which the voxels in the ellipsoid are partitioned and sampled.....	13
Figure 2 Voxel v (white circle) contained in Disc B, splatting its weighted intensity across 4 bins with the most being splatted into Bin B.....	13
Figure 3 (Left) A Gaussian footprint for a voxel. The maximum weight a is splatted to Bin A and Bin A' receives the next largest weight a' . (Right) The equation to calculate σ given a and the bin weight for A'.....	15
Figure 4 Discretized intensity profile from splatting voxels in a sampling neighborhood centered at the boundary point and oriented in the direction of its gradient.....	16
Figure 5 A cumulative Gaussian fit to a profile from a sampled region (shaded area).....	18
Figure 6 Computer generated 3D data of a sphere of radius 30 voxels, inner intensity of 32, and outer intensity of 64.	21
Figure 7 Boundary points found with DoG kernel. The boundary points are shown as yellow points with the data in (a) and without the data on in (b).....	22
Figure 8 Boundary profiles estimating the true location of the sphere's boundary shown with (a) and without (b) the data.....	23
Figure 9 Radius error from the DoG boundary points used to construct boundary profiles.	24
Figure 10 The radius error calculated from μ of profiles generated with differing neighborhoods and splatting footprints after eliminating blatantly "bad" profiles and prior to elimination of low-confidence parameters. The symbol // indicates a break in the axis scale of the graphs.....	25
Figure 11 Measures of confidence for μ plotted against the error in estimating the radius of the sphere. Dashed line shows the threshold that ensures profiles that will estimate the radius within .5 voxels of the true radius.26	26
Figure 12 Radius error from estimated μ after removing low-confidence profiles for profiles generated with an ellipsoidal sampling neighborhood and triangle splatting footprint.....	27
Figure 13 The distribution of error in estimating the intensity values on either side of the boundary as a function of μ (with zero being the center of the ellipsoid).....	29
Figure 14 The error in estimating the intensity for the interior of the sphere I_1 versus its measure of confidence z_1	30

Figure 15 The error in estimating the intensity exterior to the sphere I_2 as a function of its measure of confidence z_2	30
Figure 16 Error in boundary location from gradient-oriented profiles versus measure of confidence.	31
Figure 17 A core profile spanning an object containing voxels of intensity I_1 (shaded area), against a background of intensity I_2 . The core profile consists of two boundary profiles separated by $\vec{s}_{1,2}$ and a center point midway between the boundary locations of the profiles. I_{ij} is the intensity estimate for region i by core profile b_j . The angle between the core vector $\vec{s}_{1,2}$ and the individual gradient directions does not have to be zero (see (ii) of requirements). Also, the exterior voxel intensities I_2 do not have to be the same for each boundary profile (see (iii) of requirements).....	33
Figure 18 A network of four core profiles spanning 3 different regions of homogenous voxel intensity (solid lines) and 1 region of heterogeneous voxel intensity (dashed line). Heterogeneous core profiles span regions of differing intensity.....	35
Figure 19 Top shows basic shapes in dark gray with corresponding core in light gray. Below these are shown the corresponding basic core atom configurations (from left to right): “koosh-ball,” “spokes-of-a-wheel,” and “bed-of-nails.” (courtesy of George Stetten [15])......	36
Figure 20 Lambda triangle.....	37
Figure 21 Test data simulating a 3D ultrasound image of a heart. The interior of the smaller sphere is of intensity 32 and represents the left ventricle, the shell is of intensity 128, representing the myocardium, and the background is of intensity 64.	39
Figure 22 a) The boundary points found with the DOG kernel on a regularly sampled grid overlaid on the data. b) The grid of boundary points seen more clearly without the data. c) The resulting boundary <i>profiles</i> found from the boundary <i>points</i> in (a). d) The same boundary profiles shown in (c) without the data. Notice the profiles in (c) and (d) are more accurately located on the boundaries without the sampling artifact seen in (a) and (b).	40
Figure 23 a) The results of applying core profiles to the test data, displayed with the data. The red dots correspond to boundary profiles of the core profiles. The green dots are the centers of the core profiles. b) Cutaway view exposing the centers of the core profiles. c) The resulting core profiles without the test data.	42
Figure 24 a) Medialness found from core profiles. Red lines indicate slabness, green lines indicate cylindricalness, and blue lines indicate sphereness. b) Cutaway view exposing the lines of medial dimensionality. c) Clustered core atoms indicating sphericalness, displayed as the blue symbol and slabness displayed as the red lines.....	43
Figure 25 Both graphs show the error in estimating intensities on either side of a boundary as a function of $\mu_{relative}$. The left graph shows the error in estimating the intensities on both sides of the inner boundary. The right graph shows the error in estimating the intensities on both sides of the outer boundary.	44
Figure 26 Comparison of errors in estimating boundary locations with the DoG detector versus gradient-oriented profiles. 96% of the total profiles as opposed 29% of the total DoG kernels estimated a boundary location within one voxel.....	44
Figure 27 Intensity measures of confidence for both boundaries.	45

Figure 28 The distribution of center points of the core profiles. The center of the sphere is at 0 and the center of the slab between the spheres is at 22.5.....	46
Figure 29 Distribution of core profile scale consistent with expected results.....	47
Figure 30 Computer generated 3D data of two concentric ellipsoids.....	49
Figure 31 Left shows estimations of boundary locations from gradient-oriented profiles as orange points with the data. Right shows the boundary profiles without the data, conveying the sub-pixel ability of profiles to locate boundaries.....	50
Figure 32 Search and Count Method: Automatically measuring volume with core profiles using an intelligent fill routine.....	50
Figure 33 a) Local medialness found from core profiles. Red lines indicate slabness and green lines indicate cylindricalness. No sphereness has been detected in this case. b) Clustered core atoms indicating cylindricalness for the inner ellipsoid, displayed as the green symbol and slabness displayed as the red lines for the shell around the inner ellipsoid. c) Clustered core profiles without data.....	52
Figure 34 The use of medial nodes to measure volume without explicitly defining a surface with a “medial region fill.” Spherical regions are centered at each medial node with a diameter equal to the average core profile length. The spherical regions expand into ellipsoids to similar voxels until they collide with object boundaries. The final volume measures are found by counting voxels for similar spherical regions.....	53
Figure 35 A 3D scan of the left ventricular balloon model displayed as 2D cross-sections. The transducer is at the top of the image and the necks of the balloons are at the bottom of the image (the circular white blob). The “reflection” at the bottom of the image is an ultrasound artifact caused by reverberations within the balloon. 55	
Figure 36 a) Location of boundaries found by boundary profiles. b) Population of core profiles found (red points are boundary profiles and green points are centers of core profiles).....	56
Figure 37 a) Medial nodes found for the balloon detecting slabness of the simulated myocardium. b) Medial nodes detecting cylindricalness for the LV.....	56
Figure 38 a) Clustered core profiles. b) Applying constraints to clustered profiles.....	57

1.0 INTRODUCTION

Since the birth of medical imaging, doctors and healthcare workers have been able to look beneath the surface of the skin to obtain diagnostic information. Computational techniques to assist in the analysis of medical images have been developed, but have not advanced at the same rate as the imaging technologies themselves. Automating these computational techniques would alleviate the long hours of monotonous work required for manual analysis. This is particularly valuable for large datasets such as those produced by Real-Time Three-Dimensional (RT3D) ultrasound. The work of this thesis takes steps towards automatically making measures of the left ventricle and myocardium beneficial to the diagnosis of cardiovascular diseases. RT3D ultrasound is the imaging modality of interest because of its unique ability to capture three-dimensional images of the heart in real-time. A stable framework is developed in this thesis to automatically identify and measure structures in RT3D ultrasound cardiac phantom data.

1.1 Goals and Clinical Motivation

Dr. Roentgen's inadvertent discovery of X-rays in 1895 was almost immediately applied to medicine when Dr. Edwin Frost performed the first clinical X-ray in early 1896. These projected images dramatically affected the medical community providing a means to look into the body of a patient. The newly acquired visual information assisted doctors and other healthcare professionals in making diagnostic decisions. The clinical value was immediate and spawned the invention of further medical imaging modalities. In 1965, engineers at Siemens created the world's first real-time ultrasound system, using sound to create live two-dimensional (2D) images of tissue displayed on a cathode-ray tube. In 1967, Godfrey Hounsfield built the first Computed Tomography (CT) prototype using X-rays to produce a three-dimensional (3D) tomographic image. Dr. Paul Lauterbur first described Magnetic Resonance Imaging (MRI) in 1972 using radio frequency and magnetization to generate 3D images of

internal structures. These and other imaging modalities continue to rapidly evolve, each offering its unique insight into the human body.

During the development of medical imaging techniques, computational methods of image reconstruction have been essential to image visualization. Doctors need these images to extract information critical to diagnosis and treatment by identifying, measuring, and functionally assessing structures. Unfortunately, image analysis as a means to assist doctors in extracting this information has evolved much slower than the imaging technology itself. Most current forms of clinical image analysis require manual input from the user making it a tedious and time-consuming process. This is of particular concern with large datasets such as those produced by an experimental imaging modality, Real Time Three-Dimensional (RT3D) ultrasound [1, 2]. RT3D ultrasound produces 3D images in real-time at a rate of 22 volumes per second, fast enough to capture a single heart beat, providing a “full picture” of the heart. These unique capabilities of RT3D ultrasound are potentially very beneficial to cardiologists making observations based on cardiac structure and motion for diagnosis and treatment of cardiovascular diseases (CVD) [3-10]. With the limitations in current analysis techniques, RT3D ultrasound data is only useful when processed after the data has been obtained from the patient. In order to take immediate advantage of the data and allow real-time identification and measurement of cardiac structures, a real-time automated method is required.

If automated cardiac image analysis can be improved, its worth to the medical community would immediately impact those suffering from CVD. According to CDC (Centers for Disease Control and Prevention), NCHS (Nation Center for Health Statistics), and the AHA (American Heart Association), CVD contributed to one-third of worldwide deaths in 1999. Although the United States has one of the lowest mortality rates for males and females, CVD ranks as the leading cause of death responsible for 40 percent of the 2.4 million deaths per year, and nearly 61.8 million Americans live with some form of cardiovascular disease. In 2002, it is estimated that Americans will pay about \$330 billion in CVD-related medical costs and disability [11].

Automated techniques exist, but are generally unreliable because of the inherent difficulties with anatomical shapes and image quality, such as anatomy variation, image noise, and discontinuous object boundaries. In particular, ultrasound also suffers from the path dependence of the ultrasound signal and the non-rectilinear coordinate system in which the data are collected. Among ultrasound technologies, RT3D ultrasound has especially high noise and low resolution, a trade-off made to achieve 3D imaging at such a high speed. Despite

these difficulties with RT3D ultrasound, the potential clinical benefit justifies the challenge of automating its analysis.

Continuous measurement of left ventricular volume can be utilized to calculate and monitor cardiac function parameters such as heart rate, stroke volume, and ejection fraction. Myocardial volume, and thickness can be determined and myocardial motion can be tracked. If healthy valves are assumed, stroke volume and cardiac output can be calculated. Each of these measured quantities could be useful in the diagnoses of CVD, including cardiomyopathy, arrhythmia, ischemia, valve disease, myocardial infarction, and congestive heart failure. Abnormal cardiac contractions can be an indication of valve disease and arrhythmia, observed by tracking the cardiac surfaces through the cardiac cycle. Independent of target and imaging modality, an automated technique of boundary classification would serve as a useful tool to recognize any anatomical shape or foreign body.

1.2 Approach to Automating Image Analysis

The work of this thesis focuses on developing a stable framework to automatically (1) locate and identify the left ventricle and myocardium, (2) measure ventricular volume and surface area, (3) measure myocardial volume and thickness, and (4) track cardiac motion on RT3D ultrasound data. Given the difficult goal of automating image analysis on a complex structure in motion in noisy data, a sensible first step is to develop and test the approach on computer-generated data. The next step will be to apply it to RT3D ultrasound images of a left ventricular model consisting of latex balloons. The model presents a simplified structural representation of the heart while introducing the difficulties of RT3D ultrasound data.

Typically, shape detection begins by finding boundary points in an image and then determining their parent shape in some way. Boundary point detection dates back to the origins of image processing and typically involves locating areas with a gradient magnitude above a certain threshold. However, many of these commonly used boundary detectors merely detect the presence of a boundary and discard valuable information that can be useful in shape detection and analysis. Several boundary detectors make use of gradient direction, such as the Canny edge detector [12], but do not further parameterize the boundary or the opposing regions. As part of my Master's research, I have developed a technique called gradient-oriented boundary profiles that can parameterize a boundary by yielding estimates of boundary location, boundary width, and voxel intensities on both sides of the

boundary. Gradient-oriented boundary profiles are rotationally invariant and inherently insensitive to noise. The developmental work of boundary profiles has been presented at a national conference [13] and published in a peer reviewed journal [14].

A subsequent step in my Master's research has been to use these new boundary profiles for shape detection. A shape can be defined by the relationships between the points on its boundaries. Shapes can be found by grouping neighboring boundary points or tracking them along potential boundaries. Generally, such local boundary measurements are highly sensitive to image noise, and a global comparison of boundary points may be more successful. Boundary points can be organized globally using geometric relationships between all boundary points, but this is computationally untenable. A reasonable compromise is to focus on the medial relationship which links boundary points on opposing boundaries through the center of an object. Core atoms [15, 16] are a means to perform this association by locating boundary pairs that face each other across an object. To determine whether boundary points face each other one must compare the orientations of the boundary points. The gradient direction of opposing boundary points may be either towards or away from each other, depending on the intensity of the target and background in the image. Previous implementations of core atoms, based on simple gradient direction were unable to deal with this ambiguity. The application of my new boundary profiles provides an elegant solution. Boundary profiles permit the formation of core atoms without *a priori* knowledge of objects in the image, by comparing the shared intensities of the two boundary points. Therefore, the selection of core atoms is independent of the background intensities. Also, there is no restriction on the objects intensity and multiple objects of differing intensities can be found with the same search. Finally, the profiles yield greater accuracy in determining boundary location than conventional boundary detectors.

Statistical analysis of local core atom populations provides measures of medial properties of local shape. The long-term goal is to combine these simple shapes into complex shapes, to identify and measure the shape parameters of the objects they correspond to. As previously mentioned, abnormal left ventricular volume, myocardial volume, myocardial thickness, and cardiac motion can be indicative of cardiovascular disease, so an accurate means to calculate these measures would be clinically useful. The first step is to locate the cardiac surfaces in 3D real-time data, which is what the majority of this thesis addresses.

1.3 Thesis Overview

- i. Gradient-Oriented Boundary Profiles are developed as a method to parameterize boundaries contributing a robust set of parameters for shape analysis. Profiles are validated on synthetic data and demonstrated on RT3D images of left ventricular phantoms.
- ii. Measures of confidence are developed for the parameters delivered by boundary profiles to eliminate low-accuracy estimates.
- iii. Core atoms are constructed from gradient-oriented profiles and their populations are statistically analyzed to yield local measurements of shape.

1.4 Guide to Remaining Chapters

Chapter 2	Background information on relevant boundary detection techniques, the state of medical imaging, methods to detect and measure left ventricular volume in ultrasound images, and medial shape analysis.
Chapter 3	Gradient-oriented boundary profiles are described in detail, including the algorithm to form them, means to eliminate profiles considered unacceptable, and their ability to classify boundaries.
Chapter 4	Boundary profiles are applied to a computer-generated sphere to investigate different sampling regions and splat type, measure their accuracy in locating the surface of the sphere, and assess their overall performance.
Chapter 5	Core profiles are described as well as methods to measure medial properties.
Chapter 6	Core profiles are applied to data containing two concentric spheres. The accuracy of boundary profiles to locate the boundaries is measured. The population of core profiles is analyzed and used to calculate volumes of the data.
Chapter 7	Measuring volume via core profiles without any geometric assumptions about the data consisting of two concentric ellipsoids.
Chapter 8	Gradient-oriented profiles are applied to real-time three-dimensional ultrasound images.
Chapter 9	Contribution to an open-source image processing toolkit called ITK.
Chapter 10	Work to be done in the future to improve gradient-oriented profiles and more robust methods to automatically identify and measure objects.
Chapter 11	Conclusion

2.0 BACKGROUND

2.1 Boundary Detection

Gradient-oriented profiles require a set of boundary candidates that deliver gradient magnitude and orientation. Boundary detection is a well-established field, described in a number of general references, for example [17-21]. Many methods of image analysis employ a gradient detection scheme that identifies boundary candidates by simple measurement of the gradient magnitude throughout an image and selection of candidates above a pre-determined threshold. Boundaries in images are characterized by abrupt changes in intensity corresponding to a strong gradient. By taking the derivative of the image and locating the maximum derivatives, boundaries can be found. Taking the derivative of the image may proceed by convolving the image with a number of kernels, such as the Roberts Cross and Sobel kernels. The Roberts Cross kernel, when applied to gray-images, yields gradient magnitude and location, but is very sensitive to noise. The Sobel kernel is larger, effectively smoothing the image and reducing high-frequency noise. This makes the Sobel kernel more stable than the Roberts Cross, but it is slower to operate and still somewhat sensitive to noise.

Another class of convolution kernels is the Difference of Gaussian (DoG) gradient detector. The DoG detector is interesting for reasons related to human vision [22]. Human vision relies on first transducing light to a neural signal in the retina, which consists of a layer of photoreceptors connected to retinal ganglion cells through a number of intermediate layers. The axons of the ganglion cells make up the optic nerve, which relays the signals to the lateral geniculate nuclei. From there the signals are transmitted to the visual cortex in the posterior lobe, where the visual information is further processed [23, 24]. The spatial properties of retinal ganglion cells as receptive fields were first described by Hartline in 1938 [25]. He found that these cells would respond even when a light source was moved away from the center of its receptive field. He also noted that the receptive fields were much larger than anticipated of individual photoreceptors, suggesting signal processing and integration through retinal circuitry [26, 27].

Kuffler followed Hartline's work and found that for each ganglion cell there is a small region in which illumination increases electrical activity, and surrounding this excitatory region is a ring-shaped region in which illumination decreases electrical activity. He also found ganglion cells for which the opposite occurs, where illumination in the region inhibits excitatory activity, while illumination in the surrounding ring-shaped region increases electrical activity [28].

Following this pioneering work by Kuffler, neuroscientists mathematically modeled receptive fields. In particular, Rodieck introduced the DoG model for the receptive field of ganglion cells. In this model, each ganglion cell takes a weighted summation of its neighbors, where the weights can be positive or negative [29]. This permits the ganglion cells to act as an edge detector. There are many forms of DoG gradient detectors, and a common version finds the difference between two concentric Gaussian kernels of different scale [6]. The version of the DoG gradient detector used for the research described in this thesis consists of three pairs of same-scale Gaussian filters displaced in location along each of the cardinal axes from the sample point to measure the respective components of the gradient in three dimensions. This particular DoG gradient detector is efficient and delivers gradient direction (unlike the concentric DoG), as well as gradient magnitude, both being required for the construction of boundary profiles. Each boundary candidate that exceeds a pre-determined threshold for gradient magnitude is used to generate an individual boundary profile (details of implementation are in section 3.2).

The noise sensitivity of the Roberts Cross and Sobel kernels arise from taking a derivative, which amplifies the high-frequency image noise. Marr [22, 30] and Canny [12] address this problem by convolving the image with a Gaussian kernel prior to calculating the derivative. Increasing the width of the Gaussian kernel reduces the detector's sensitivity to high-frequency noise, but also eliminates the finer details in the image. Inherent to the DoG detector used here is a Gaussian kernel with a carefully chosen kernel width. After the Gaussian smoothing, Marr's approach proceeds with the application of a Laplacian of Gaussian (LoG) filter locating boundaries wherever the value is zero. The LoG operator can be closely approximated by a concentric DoG operator, which can similarly be used to enhance the visibility of edges while at the same time smoothing small-scale noise [22]. Canny's boundary detector requires a set of boundary detection filter masks that have various orientations. The masks contain a derivative of a Gaussian function to perform a directional derivative across the intended boundary. When applied, a smooth averaging profile appears in the mask along the intended boundary direction in order to reduce noise without compromising the sharpness of the boundary profile [17].

Similar to the Canny edge detector, gradient-oriented profiles operate along the gradient direction averaging voxel intensities to reduce noise sensitivity. The Marr and Canny edge detectors require multiple applications of filters at different scales for a thorough edge detection procedure. Gradient-oriented profiles require a single application, while producing a more thorough and accurate parameterization of the image boundaries than either the Marr or Canny edge detectors.

A useful property of the gradient magnitude is its inherent insensitivity to rotation. More complex parameterizations of the boundary can also be made rotationally invariant by operating along the local gradient vector. Steerable Filters, for example, parameterize the boundary using a basis set of Gaussian derivatives that is easily rotated into the coordinate system of the local boundary [31]. An equivalent rotation is accomplished with gradient-oriented profiles by projecting voxels onto the gradient vector to produce a one-dimensional intensity profile. Fitting a specific function to the profile, the cumulative Gaussian, directly yields physically significant estimates that parameterize the boundary.

2.2 State of Medical Imaging

Each of the primary modalities of medical imaging (Ultrasound, MRI, CT, etc.) offers its own unique insight into the human body. Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) produce high-resolution 3D images, effective for visualizing 3D objects (organs, foreign objects, etc.) in the body, but require significant time for image reconstruction. Conventional ultrasound imaging produces real-time 2D images very useful for observing the behavior dynamic objects, such as cardiac motion. Such ultrasound images can be reconstructed to 3-D after data acquisition, but are no longer real-time and have poor image quality compared to MRI and CT. CT can be used in cardiac studies if scans are gated to the electrocardiogram (ECG), in which case scans are averaged over many cardiac cycles. Cine Computerized Tomography (cine-CT), the fastest form of CT, developed early in the 90's by Imatron [32] is capable of producing a slice every 50 ms. Cine-CT can capture a 3D dataset of the heart in about 400 ms, still not fast enough for analysis during the cardiac cycle without ECG gating.

Real-time three-dimensional (RT3D) ultrasound, developed in the early 90's at Duke University [2, 33] makes use of a matrix array of transducer elements instead of a linear array to capture 3D ultrasound data. The scan rate of the transducer is approximately 22 frames per second (45ms), rapid enough to acquire cardiac data

throughout its cycle. Thus, gating to the ECG is not required to analyze function during a cardiac cycle. Image quality is sacrificed for the increased speed in RT3D data acquisition, resulting in a lower signal-to-noise ratio compared to conventional ultrasound. However, several studies have shown that cardiac assessment with RT3D images is feasible [1, 2, 33-38].

2.3 Methods to Detect and Measure LV Volume in Ultrasound Images

Early methods of measuring LV volume in 2D ultrasound rely on first manually finding the LV boundary, then calculating the volume based on geometric assumptions about the shape and location of the slice. Utilizing Simpson's rule, the LV volume can be approximated by summing the volume of discs in series along the long axis. Using Dodge-Sandler's method, the LV is approximated as an ellipsoid and the LV volume is calculated as the volume of the ellipsoid. The biplane, area/length method requires only 2 cross-sectional areas and the length of the long axis to yield an approximate volume. Since these methods are based on geometric assumptions, they tend not to be very exact. To avoid the need for geometric assumptions, Nadkarni et al. reconstruct 2D images to 3D volumes and then manually segment the LV and calculate its volume [39].

Rather than depending on manual boundary identification in 2D images, deformable contours (snakes) attempt to locate boundaries in a semi-automated fashion. Model based deformable contours attempt to segment objects by iteratively searching for possible boundaries armed with knowledge of the shape [40, 41]. Such deformable contours have been shown capable of finding, as well as measuring, the LV in 2D ultrasound data [42-44]. Hammarneh and Gustavsson have combined active contours with active shape models to segment the left ventricle from 2D ultrasound images after a period of training. The model acquires *a priori* information of the structures by first manually tracing and identifying them, capturing the variability of the LV training set [45]. Setarehdan and Soraghan use a deformable contour to close the boundary after applying fuzzy logic to find a central point in the LV then using a multi-resolution edge detector to locate boundaries points [46].

The success of deformable contours on 2D data has recently been extended to 3D. These deformable models rely on multi-local or global geometric expectations and have met with some success on mechanically scanned 3D ultrasound data [47]. Unfortunately, their application to RT3D ultrasound has not been encouraging due to the inherent low signal-to-noise ratio of the data [48]. Nevertheless, there has been success in semi-

automatic algorithms for assessing wall motion abnormalities [5] and LV ischemic risk volume [49], identifying congenital heart defects in pediatric patients, measuring left ventricular mass in dogs [50, 51], and measuring LV volume [52-58]. Stetten has successfully developed a medial approach to automatically identify the LV and measure its volume by statistically analyzing populations of medial primitives call core atoms [15, 58]. This approach is followed and modified in this thesis to automatically locate the LV and myocardium and measure their volumes in RT3D ultrasound cardiac model data. A brief overview of medial shape analysis is included in the following chapter.

2.4 Medial Shape Analysis

The medial approach to shape analysis originates back to the medial axis first introduced on binary images by Blum [59]. The medial axis is defined as the locus of centers of circles that are at least bi-tangent and fit entirely within the region being considered. Pizer et al. developed a measure called medialness to extend the medial axis to gray level images. This measure links the aperture of the boundary measurement to the radius of the medial axis to produce a “core”, defined as a height ridge of medialness [60, 61]. Methods using medialness as image features have proven robust against image noise and shape variation [62].

The core can be found with ridge tracking, but this requires manual initialization and fails at discontinuities in the ridge. Medialness filters have generally been used to look for pairs of boundaries with a specific orientation from the center. Multiple applications of such filters are necessary to account for different orientations, which is very computationally expensive [63, 64]. As developed by Stetten and Pizer [15], core atoms are more efficient, requiring only a single application to detect a single object. Candidate boundary points are initially found, and then examined pair-wise to see if they have sufficient medialness for core atom formation. Local populations of core atoms are then statistically analyzed to extract medial properties of the core. A brief overview of core atoms is included in chapter 5 where gradient-oriented profiles are used to create a new kind of core atom, a core profile, with the distinct advantages over the original ones developed by Stetten and Pizer.

3.0 GRADIENT-ORIENTED BOUNDARY PROFILES

Gradient-oriented boundary profiles, developed in this thesis, make use of nonlinear regression to parameterize boundaries. Unlike most boundary detectors, which simply detect the presence of a boundary, boundary profiles estimate the location of the boundary, the width of the boundary, and the voxel intensities on each side of the boundary. The methodology of gradient-oriented profiles is presented in this chapter. Also in this chapter, measures of confidence are derived for the boundary parameters, and a scheme to classify boundaries is explored. In chapter 4, gradient-oriented profiles are applied to multiple three-dimensional computer-generated datasets and the performance of boundary profiles is compared to that of the Difference of Gaussian (DoG) boundary detector described in section 2.1.

The process of finding and analyzing gradient-oriented profiles is outlined here and detailed in the following sections (corresponding to section numbers):

- (i) **Finding Boundary Candidates:** Find boundary point candidates with a technique that delivers approximate location, gradient magnitude, and gradient direction. The DoG gradient detector was used for the work of this thesis for reasons explained in section 2.1.
- (ii) **Generating a Boundary Profile:** Define neighborhoods (ellipsoidal or cylindrical regions) around each detected boundary point, and then project voxels within each neighborhood into bins along the major axis (gradient direction) to yield a profile of voxel intensity.
- (iii) **Fitting the Intensity Profile and Estimating Parameter Values:** Fit a cumulative Gaussian to the intensity profile, yielding estimates for intensity on both sides of the boundary as well as the actual location and width of the boundary.
- (iv) **Eliminate Bad Profiles:** Eliminate blatantly “bad” boundary profiles by enforcing restrictions on the parameters based on *a priori* knowledge of the imaging modality.
- (v) **Establishing Intrinsic Measures of Confidence:** For the remaining profiles, calculate measures of confidence for the estimated values and eliminate those falling below a chosen threshold.

- (vi) **Classifying the Boundary:** Classify the boundary using the remaining high-confidence estimated values.

3.1 Finding Boundary Candidates

The initial step in the generation of boundary profiles requires the collection of boundary points. Any method of boundary detection can be used, as long as the gradient direction and magnitude are delivered. The DoG gradient detector is chosen for reasons discussed in chapter 2.1 and implemented in the same manner as Stetten [15] with the use of a binomial kernel. The binomial kernel in 3D is a $2 \times 2 \times 2$ cube filled with ones. When convolved with a 3D image multiple times (and normalized each time by dividing by 8, the Central Limit Theorem dictates that the result will approximate a Gaussian kernel. The number of applications determines the width of the Gaussian. The DoG is achieved by subtracting the resulting Gaussian kernel from itself displaced on either side of the sample location along each of the cardinal axes to yield the corresponding gradient components. A gradient in 3-dimensions is accomplished by applying three separate DoG kernels, one along each cardinal axis.

3.2 Generating a Boundary Profile

Given the detection of boundary points by the DoG kernel, a boundary profile is generated in the direction of the gradient by sampling the voxels in an oriented neighborhood centered at the original boundary point whose major axis is along the gradient vector. Thousands of boundary points are likely to be found, and therefore sampling voxels around each boundary point could be time-consuming. Choosing the shape and size of the neighborhood to minimize computation is important. The shape should be large enough to gather sufficient information by extending just beyond the transition zone of the boundary. Intuitively, a cylindrical or ellipsoidal sampling region would seem reasonable. The effects of choosing sampling regions of different shape on the overall accuracy of boundary profiles are evaluated in section 4.1.

The voxels within the sampling neighborhood are projected onto the major axis where their intensity values are splatted into bins. Splatting is a technique commonly used in computer graphics that projects the footprints of individual voxels onto a plane. Overlapping footprints are collected in bins and normalized to form a

rendered image. Instead of splatting onto a plane, we splat onto a line, that is the major axis of the sampling region. Figure 1 demonstrates the process of sampling voxels within an ellipsoidal region.




Figure 1 An ellipsoidal neighborhood whose center is the boundary point from the DoG gradient detector and major axis is the gradient vector \vec{g} of the boundary point. Two adjoining disks, A and B are shown for which the voxels in the ellipsoid are partitioned and sampled.

The ellipsoid is divided into disks along the major axis, with corresponding bins collecting footprints from voxels in adjoining disks. The footprints are wider than an individual bin, so each voxel contributes to a number of bins. A triangular footprint is has the effect as linearly interpolating the intensity of each voxel between neighboring bins along the axis. Splatting with a triangular footprint proceeds in the following manner (illustrated in figure 2):




Figure 2 Voxel v (white circle) contained in Disc B, splatting its weighted intensity across 4 bins with the most being splatted into Bin B.

Given an ellipsoidal neighborhood such as shown in Figure 1, a voxel v with intensity I_v is contained in disc B. The length of the profile L will be equal to that of the ellipsoid's major axis. Let the width of each bin be

one voxel wide and the length of the profile be an integer value, so the number of bins n is equal to L . The bin assignment for ν is determined by first finding its location on the major axis:

$$d = \vec{\nu} \cdot \vec{g} + \frac{L}{2}, \quad (1)$$

where $\vec{\nu}$ is the vector from the origin of the ellipsoid to ν , and \vec{g} is the unit gradient vector of the boundary point. The bin assignment is then:

$$bin = floor(d), \quad (2)$$

where d is the distance from the edge of the ellipsoid to ν 's projection on the major axis and $floor$ is the greatest integral value less than or equal to d . The displacement b of ν 's projection in its bin, *Bin B* is:

$$b = d - floor(d) \quad (3)$$

The triangle footprint across four bins weights the intensity of ν accordingly:

Table 1
Voxel intensity weights for the four bins of a triangle footprint

bin	Weight
$B - 1$	$1 - b$
B	$2 - b$
$B + 1$	$1 + b$
$B + 2$	b

The weighted intensity of ν for *Bin B* is then:

$$I_{\nu}|_w = (I_{\nu})(weight_{2-b}). \quad (4)$$

All of the voxel weights in each disc are summed in the appropriate four bins. The weight of contribution for each voxel is separately stored so that the total intensity in each bin can be normalized. Thus, the intensity profile represents the average voxel intensity for each disk within the ellipsoidal sampling region. Averaging within each disk reduces the effect of image noise on the boundary profile.

A Gaussian footprint has also been tried shown in Figure 3. The weighting of the Gaussian footprint follows the standard Gaussian

$$g(x) = ae^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (5)$$

The maximum weight is determined by a (amplitude of Gaussian). The parameter σ is calculated by choosing a desired weight for the bin adjacent to the maximum bin (Figure 3).




Figure 3 (Left) A Gaussian footprint for a voxel. The maximum weight a is splatted to Bin A and Bin A' receives the next largest weight a' . (Right) The equation to calculate σ given a and the bin weight for A' .

Equation 5 is modified to be a function of the voxel's position in the bin using the quantity b as described in equation 3. Table 2 shows the weighting for the four bins of the footprint using the notation in Figure 3 with c being equal to σ .

Table 2
The voxel intensity weights for the four bins of a Gaussian footprint

bin	weight
$A - 1$	$ae^{-\frac{(b+1)^2}{2c^2}}$
A	$ae^{-\frac{b^2}{2c^2}}$
$A + 1$	$ae^{-\frac{(b-1)^2}{2c^2}}$
$A + 2$	$ae^{-\frac{(b-2)^2}{2c^2}}$

The performance of the Gaussian footprint is measured against that of the triangle footprint in section 4.1. Regardless of the splatting footprint employed, sampling the voxels in this manner will result in a discretized profile of the intensity in the sampling neighborhood (Figure 4).

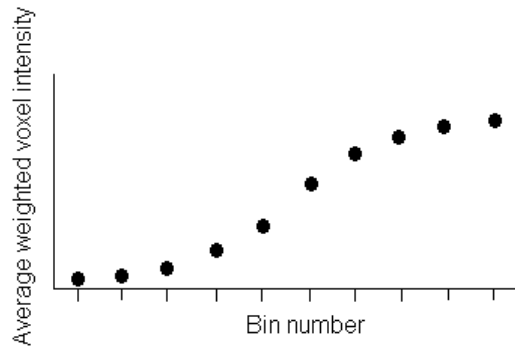


Figure 4 Discretized intensity profile from splatting voxels in a sampling neighborhood centered at the boundary point and oriented in the direction of its gradient.

3.3 Fitting the Intensity Profile and Estimating Parameter Values

The discretized intensity profile captures voxel intensity variation across a boundary. Intuitively, the boundary is located where the largest voxel variation occurs. Voxel intensity begins to stabilize near the tails of the profile, indicating a region in the image away from the boundary, but may not fully stabilize within the extent of the profile. These properties of the profile can be quantified if the profile is fit to a function. A number of different functions could potentially be used to fit to the profile. The cumulative Gaussian was chosen for reasons now described.

Most anatomical boundaries are very abrupt, inherently step functions at the sub-millimeter scale where one tissue ends and another begins. Image acquisition is inevitably limited in resolution, however, with a particular device exhibiting an overall "point spread function" usually at a significantly larger scale than the actual tissue boundary. Additional blurring may be performed intentionally or inevitably, for example, during the image processing that converts the raw data to an image. The result of these sequential convolutions tends to have the effect of convolution with a Gaussian kernel because of the Central Limit Theorem. Thus the step function of the anatomical boundary is expected to reach the image analysis stage as a cumulative Gaussian, i.e., the convolution of a step function with a Gaussian. Therefore, by fitting a cumulative Gaussian, the blurring is, in effect, reversed and the original boundary may be parameterized, as well as the total width of the blurring.

Fitting a cumulative Gaussian to an intensity profile requires optimizing 4 parameters: (1) standard deviation, corresponding to the boundary width, (2) mean, corresponding to boundary location along the major axis, and (3,4) two asymptotic values, corresponding to voxel intensity on either side of the boundary sufficiently far from the boundary to be unaffected by blurring.

The cumulative Gaussian $C(x)$ is derived as follows: The normalized Gaussian,

$$G(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}, \quad (6)$$

is integrated defining the error function (*erf*),

$$\int_{-x}^x G(v)dv = \text{erf}\left(\frac{x-\mu}{\sigma\sqrt{2}}\right), \quad (7)$$

which is scaled and offset to yield the cumulative Gaussian $C(x)$ as follows:

$$C(x) = I_1 + \frac{I_2 - I_1}{2} \left(1 + \text{erf}\left(\frac{x-\mu}{\sigma\sqrt{2}}\right) \right) \quad (8)$$

The four fixed parameters of $C(x)$ are μ (mean), σ (standard deviation), and I_1 and I_2 (asymptotic upper and lower voxel intensities). On one side of the boundary

$$\text{erf}(-\infty) = -1, \quad C(-\infty) = I_1 \quad (9)$$

and on the other side

$$\text{erf}(\infty) = 1, \quad C(\infty) = I_2. \quad (10)$$

The four parameters are labeled below in Figure 5, which shows a particular fit of the cumulative Gaussian along the sampled region of the boundary profile. The optimal boundary location \vec{b} is calculated from μ , the original boundary point location \vec{b}_o , and the gradient \vec{g}_o by

$$\vec{b} = \vec{b}_o + (\mu_{relative})(\vec{g}_o). \quad (11)$$




Figure 5 A cumulative Gaussian fit to a profile from a sampled region (shaded area).

A variety of techniques are available for optimizing the fit of a non-linear function to a set of sample points. A Quasi-Newton non-linear optimization algorithm available as part of the *AD Model Builder* from Otter Research, Inc., [65] was used because it does not require explicit derivation of derivatives and is acceptably rapid and robust.

To make intelligent use of parameters estimated from the curve fitting, measures of confidence for each of the parameters must be established. This takes the form of two steps, as described in the following two sections.

3.4 Eliminating Bad Profiles

The first step involves the elimination of unacceptable parameters by rejecting blatantly “bad” profiles for which a reasonable fit of the cumulative Gaussian cannot be found. The profile is rejected if:

- (i) Either extrapolated voxel intensity value I_1 or I_2 falls outside of the acceptable range of values for the imaging data.
- (ii) The estimated boundary location μ falls outside of the sampling region.

Boundary profiles that are rejected by these criteria are no longer considered. Determining confidence levels for the estimated parameters and eliminating low-confidence estimates on a parameter-by-parameter basis further eliminates bad profiles.

3.5 Establishing Intrinsic Measures of Confidence

If a fit was accepted for a given profile, the next step is to determine a measure of confidence in the individual parameters of the cumulative Gaussian. It is postulated that the confidence in a particular parameter is related to the location of the boundary within the ellipsoid, as well as the boundary width.

In Figure 5, the cumulative Gaussian that best fits an intensity profile is shown. By definition, the underlying function (ignoring the particular region over which it was sampled) is symmetrically distributed on either side of the mean μ (with μ representing the optimum boundary location). However, the function is not generally symmetrically distributed with respect to the sampling region (represented by the shaded area from p_1 to p_2). Given that the best fit has a mean μ and standard deviation σ , two distances d_1 and d_2 can be defined from μ to p_1 and p_2 respectively, from which two normalized distances z_1 and z_2 can be computed as,

$$z_1 = \frac{d_1}{\sigma} \quad \text{and} \quad z_2 = \frac{d_2}{\sigma}. \quad (12)$$

The values of z_1 and z_2 indicate how many standard deviations from μ in each direction the actual samples extend, and serve as measures of confidence for I_1 and I_2 respectively. For each direction, the greater the number of standard deviations the samples extend from μ , the less the effect of the boundary and the greater the confidence the estimated intensity. For example, in Figure 5 there is more confidence in the estimate I_1 than the estimate I_2 , because $d_1 > d_2$, and therefore $z_1 > z_2$. A threshold can be placed on z_1 and z_2 based on the desired accuracy for the intensity estimates.

In addition to a measure of confidence for I_1 and I_2 , a measure of confidence for μ is developed. The measure of confidence for μ indicates whether sufficient samples exist on either side of μ to estimate it accurately. This measure of confidence is defined as

$$z_{\min} = \min(z_1, z_2). \quad (13)$$

If on either side of μ insufficient samples exist to anchor the cumulative Gaussian, z_{\min} will be small, and it is expected there will be difficulty in estimating μ accurately. Like the measures of confidence for profile intensity estimates, a threshold is also placed on z_{\min} to determine whether there is acceptable confidence in μ .

The standard deviation σ determines the width of the transition zone for the cumulative Gaussian. The measure of confidence for σ is a measure of how much the transition zone encompasses the sampling region. If σ is too large, too many samples will anchor the transition zone of the cumulative Gaussian, but not enough samples will anchor the tails resulting in an inaccurate estimate for σ as well as I_1 and I_2 . This measure of confidence is defined as the ratio of σ to the length of the sampling profile L ,

$$z_\sigma = \frac{\sigma}{L}. \quad (14)$$

If too many samples contribute to the transition zone of the cumulative Gaussian, z_σ will be close to 1. A threshold can be placed on z_σ based on the amount of pre-processing blurring. For instance, if a blurring kernel was applied only once then the threshold on z_σ should be low.

3.6 Classifying the Boundary

Armed with high-confidence estimates for I_1 , I_2 , μ , and σ the local boundary can now be classified in terms of these parameters. The parameters I_1 and I_2 estimate voxel intensity beyond the ellipsoid, sufficiently far from the boundary to be stable. The mean μ estimates the anatomical boundary location on the major axis of the sampling neighborhood and σ estimates the total boundary width.

Rather than struggling to extract the last bit of classification from each profile, the measures of confidence permit the exclusion of individual parameters on a case-by-case basis. In the initial gathering of boundary candidates, a sampling interval well below the conventional Nyquist guidelines is used to yield a comprehensive set of profiles from which a thorough examination of a given boundary can be made. By setting thresholds for the intrinsic measures of confidence, individual estimates can be rejected or accepted to make intelligent use of this over-sampled population.

4.0 ASSESSMENT OF BOUNDARY PROFILES ON A SYNTHETIC SPHERE

Computer-generated 3D data containing a sphere was used to assess the performance of gradient-oriented boundary profiles as well as demonstrate the effects of different types of sampling neighborhoods and splatting footprints. The sphere was initially used for testing purposes because of its simple parametric definition. The data was 100 voxels on a side, consisting of 8-bit voxels and the generated sphere had a radius of 30 voxels, an interior value of 32, and an exterior value of 64 (Figure 6). Different combinations of sampling neighborhood shapes and splatting footprints were first explored. Initial thresholds for intrinsic measures of confidence discussed in section 3.5 were established to eliminate low-confidence profiles. Then the radius of the sphere was estimated and compared to its known value. A Root Mean Squared (RMS) error metric was used to determine which combination yielded the more accurate results. After choosing the combination of sampling and splatting methods with least RMS error, the performance of these boundary profiles was assessed in more detail. Optimal thresholds for measures of confidence for each profile parameter were found by using the measured error in the parameters. In section 6.1, generated data is extended to a hollow sphere, representing the LV and myocardium to test the application of profiles to core atoms.



Figure 6 Computer generated 3D data of a sphere of radius 30 voxels, inner intensity of 32, and outer intensity of 64.

4.1 Testing Sampling Regions and Splat Types on a Sphere

For each type of sampling region and splat type, a binomial kernel was applied multiple times on the computer-generated sphere, approximating a Gaussian and simulating the point-spread function of an imaging system. Using the DoG gradient detector kernel as described in section 3.1, 1,264 boundary candidates were found. Figure 7 shows the centers of the DoG kernels that met the gradient threshold requirement. These boundary point locations are located on the sampling grid rather than on the boundaries themselves, often identifying two boundaries on each side of the true boundary.




Figure 7 Boundary points found with DoG kernel. The boundary points are shown as yellow points with the data in (a) and without the data on in (b).

These boundary points were used as the centers of the sampling regions for boundary profile construction. The sampling regions were either an ellipsoid or cylinder oriented along the boundary point's gradient. The sampling region size was chosen *ad hoc* (potential automated methods are discussed in Chapter 11) having a major axis length of 10 voxels and a minor axis length of 6 voxels. Along the major axis, 10 bins were established, each being one voxel wide. The intensity values of the voxels within the neighborhoods were splatted into the bins using either a triangle or Gaussian footprint as described in section 3.2. A cumulative Gaussian was fit to the resulting intensity profile as described in section 3.3. A fit was rejected if either intensity estimates were outside of the range of the 8-bit data or the estimate for boundary location fell outside of the sampling region (section 3.4). Data including the number of profiles found, number of profiles eliminated, and total time for construction is shown below in table 3.

Table 3
Number of profiles found, eliminated, and total time for discovery and construction

Neighborhood Type	Splat Type	# of Blatantly Bad Profiles	Final # of Profiles	Time to Find Profiles (seconds)*
Cylindrical	Gaussian	96	1168	158
Cylindrical	Triangle	64	1200	150
Ellipsoidal	Gaussian	96	1168	158
Ellipsoidal	Triangle	64	1200	152

* Workstation Specifications: 550 MHz Pentium III Xeon Processor and 512 MB Memory.

For these trials, a 550 MHz Pentium III Windows-based machine with 512 MB of memory was used for computation and visualization. The average time for profile construction was approximately 154 seconds, slower than the DoG kernel and too slow for real-time analysis. A more powerful computer (Dual 1.2GHz processors with 2GB memory) was subsequently used and shown to significantly decrease the computation time to approximately 90 seconds for each trial. The boundary location estimates by the profiles are shown in Figure 8 for an ellipsoid sampling neighborhood and triangle footprint. Visually, the points are located on the surface of the sphere and not the sampling grid like they were with the DoG points. The sub-pixel boundary estimation accuracy is demonstrated below. Unfortunately, the computation overhead of constructing profiles makes them insufficient for real-time image analysis, but profiles offer a vastly improved set of boundary points and valuable image information.




Figure 8 Boundary profiles estimating the true location of the sphere's boundary shown with (a) and without (b) the data.

The improvement of estimating boundary location by profiles appeared to be consistent, independent of neighborhood type and splat type, which will be shown quantitatively later in this chapter.

The radius of the sphere was estimated from both DoG boundary points and the boundary profiles by calculating the distance to the center of the sphere. The error in radius estimate was found by comparing the calculated value to the known radius of the sphere. The distance from the center of the sphere to the location on the major axis of the neighborhood corresponding to μ was defined as the estimated radius $R_{estimate}$. This was compared to the known radius R_{true} of the sphere, with R_{error} being defined as

$$R_{error} = |R_{true} - R_{estimate}|. \quad (15)$$

Figure 9 shows the error distribution for the DoG boundary point radius estimation. The radius error ranges from 0 to 8 voxels and the bi-modal distribution is most likely due to sampling error. In contrast, boundary profiles were able to consistently estimate the radius to within 4 voxels, with a majority of the errors within 0.05 voxels of the true radius demonstrating sub-pixel accuracy by boundary profiles. Figure 10 shows the distribution of radius errors estimated from boundary profiles with the four different combinations of sampling regions and splat types.




Figure 9 Radius error from the DoG boundary points used to construct boundary profiles.

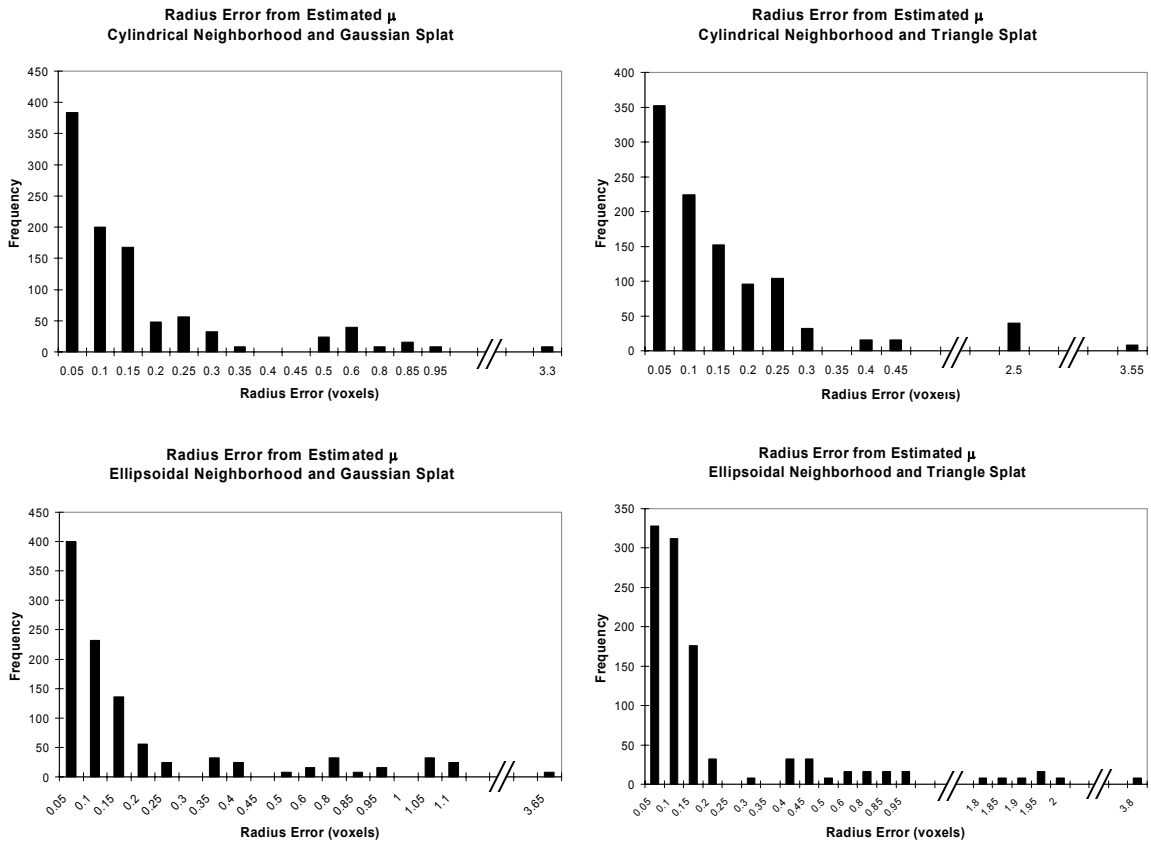


Figure 10 The radius error calculated from μ of profiles generated with differing neighborhoods and splatting footprints after eliminating blatantly “bad” profiles and prior to elimination of low-confidence parameters. The symbol // indicates a break in the axis scale of the graphs.

The measure of confidence for μ , which was defined as $z_{\min}(z_1, z_2)$ in section 3.5, was determined for each trial, and used to eliminate profiles expected to produce a large radius error. By plotting the z_{\min} vs. the radius error a measure of confidence $z_{\min} > 1.5$ standard deviations was chosen to ensure that μ is consistently within .5 voxel of the true boundary and eliminates all outliers across the 4 trials (Figure 11). Table 4 shows that the number of profiles eliminated for each trial was about the same.

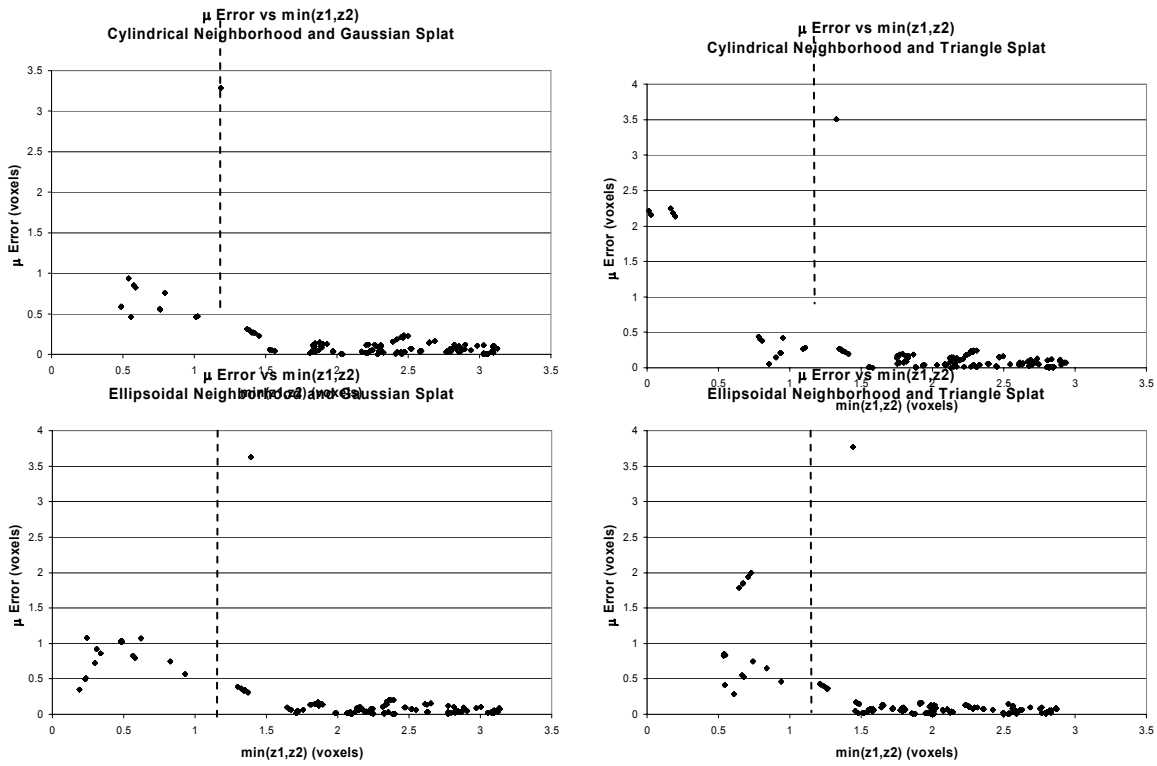


Figure 11 Measures of confidence for μ plotted against the error in estimating the radius of the sphere. Dashed line shows the threshold that ensures profiles that will estimate the radius within .5 voxels of the true radius.

Table 4

Number of profiles eliminated using a threshold of $z_{\min} > 1.5$

Neighborhood Type	Splat Type	# of Eliminated Profiles
Cylindrical	Gaussian	144
Cylindrical	Triangle	192
Ellipsoidal	Gaussian	200
Ellipsoidal	Triangle	240

After removing profiles below the μ measure of confidence threshold, Root Mean Squared (RMS) was used as a metric to assess which sampling neighborhood and footprint combination yields the least error in estimating the radius:

$$RMS = \sqrt{\frac{1}{n} \sum_{i=1}^n R_{error}^2} \quad (16)$$

The resulting RMS values are summarized in Table 5. Identical to standard RMS, a value near 0 indicates least error. Since the ellipsoidal neighborhood and triangle footprint resulted in the smallest RMS value, this combination was used for the remainder of this thesis. The radius errors for the remaining 1,048 profiles after applying the μ measure of confidence threshold are shown in Figure 12. The radius error ranges from 0 to .2 voxels as opposed to 0 to 8 voxels before applying the boundary profiles (compare with Figure 9).

Table 5
Effects of neighborhood and splat type on determining the radius of the sphere accurately

Neighborhood Type	Splat Type	RMS
Cylindrical	Gaussian	0.092
Cylindrical	Triangle	0.104
Ellipsoidal	Gaussian	0.086
Ellipsoidal	Triangle	0.078

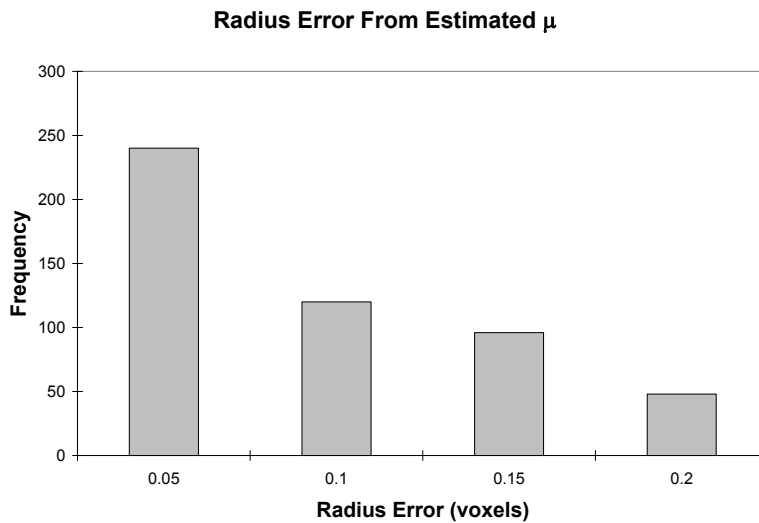


Figure 12 Radius error from estimated μ after removing low-confidence profiles for profiles generated with an ellipsoidal sampling neighborhood and triangle splatting footprint.

4.2 Measuring the Accuracy of Intensity Estimates by Gradient-Oriented Profiles

After eliminating profiles with low-confidence estimates for μ , the accuracy of the intensity parameters of the remaining profiles was measured using *a priori* knowledge of the intensities on both sides of the boundary. I_{true} is the known voxel intensity, denoted I_1 inside the sphere or I_2 outside the sphere. The use of *a priori* knowledge of the intensities can be eliminated and replaced with a statistical analysis of the population of boundary

profiles. The intensity estimates of the boundary profiles can be clustered within a threshold and averaged for a measure of I_{true} . The estimate of intensity derived from the profile for either I_1 or I_2 is denoted $I_{estimate}$, and I_{error} is defined as

$$I_{error} = |I_{true} - I_{estimate}|. \quad (17)$$

Figure 13 is a graph of I_{error} vs. $\mu_{relative}$, where

$$\mu_{relative} = \mu - \frac{L}{2} \quad (18)$$

and L is the length of the sampling region's major axis and

$$-\frac{L}{2} < \mu_{relative} < \frac{L}{2}. \quad (19)$$

Representing μ in this manner as $\mu_{relative}$ defines a coordinate system for the boundary location within the sampling region. If:

$$\mu_{relative} = 0, \text{ the boundary is near the center of the sampling region} \quad (20)$$

$$\mu_{relative} < 0, \text{ the boundary is towards the lower asymptotic end of the sampling region} \quad (21)$$

$$\mu_{relative} > 0, \text{ the boundary is towards the upper asymptotic end of the sampling region.} \quad (22)$$

When $\mu_{relative}$ is close to 0, the profiles are better able to estimate the intensity on either side of the boundary than when μ is located away from the center of the sampling region. When μ is located away from the centers of the sampling region, intensity values are estimated less accurately in that direction. However, such a profile will estimate the intensity in the opposite direction with even better accuracy. For example, in Figure 13, I_2 can be estimated to within one unit of intensity, when $\mu_{relative}$ is much less than zero.

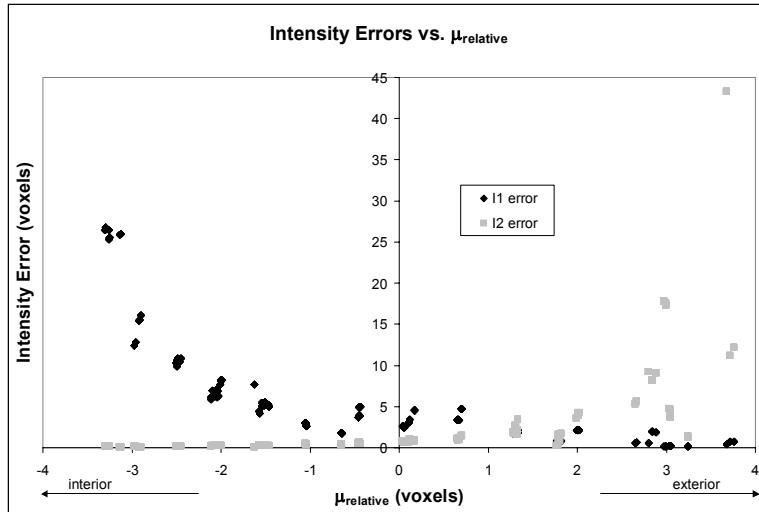


Figure 13 The distribution of error in estimating the intensity values on either side of the boundary as a function of μ (with zero being the center of the ellipsoid).

4.3 Judging Intrinsic Measures of Confidence

Gradient-oriented profiles have been shown to estimate the intensity values and localize the boundary location. The validity of the intrinsic measures of confidence is now demonstrated.

The measures of confidence for the profile parameters were calculated using the methods discussed in section 3.5. Figure 14 shows the error in estimating I_1 (interior of sphere) as a function of its measure of confidence z_1 and for 2 standard deviations ($z_1 \geq 2.0$) the error $I_{error} < 10$. Therefore, a threshold placed on z_1 of 2.0 would guarantee this accuracy.




Figure 14 The error in estimating the intensity for the interior of the sphere I_1 versus its measure of confidence z_1 .

With this threshold, 776 of the total profiles (67%) will remain to estimate the intensity I_1 . Similarly, Figure 15 shows the error in estimating I_2 (exterior of sphere) as a function of its the measure of confidence z_2 , with similar results for $I_{error} < 10$ the threshold for z_2 is 1.5.

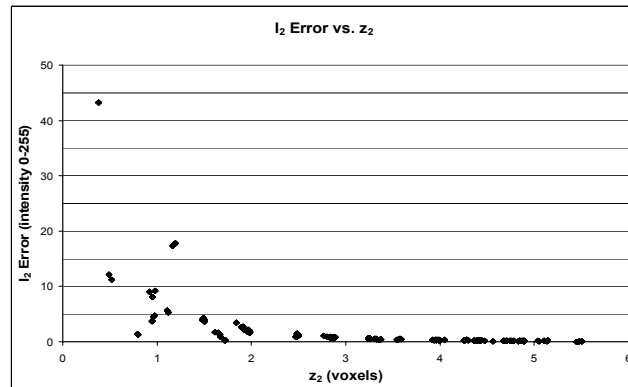


Figure 15 The error in estimating the intensity exterior to the sphere I_2 as a function of its measure of confidence z_2 .

It was predicted earlier that profiles with μ near the center of the ellipsoid should be able to estimate both intensities I_1 and I_2 with high accuracy. There were 632 out of the 1,160 (54%) profiles that could reliably classify both intensities as described above and they with μ near the center of the sampling region.

Profiles with μ near the center of the ellipsoid should also be able to reliably estimate μ . Our measure of confidence for μ was therefore defined as $z_{\min}(z_1, z_2)$. As shown in Figure 16, for $z_{\min} > \frac{1}{2}$ standard deviations, a high confidence seems warranted for μ , since it is consistently within 1 voxel of the true boundary. This threshold places high confidence in the estimates of boundary location for 1,152 out of 1,160 profiles (99%).

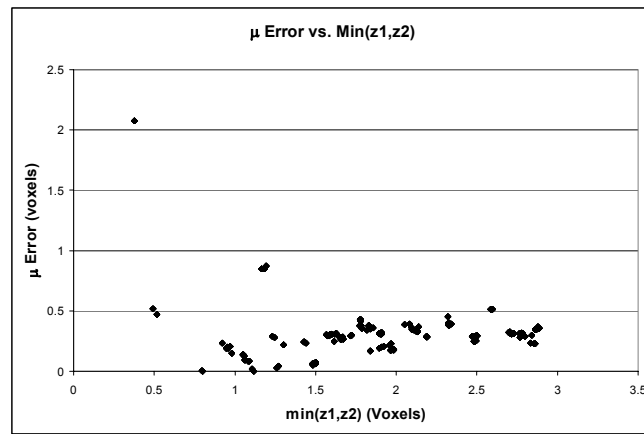


Figure 16 Error in boundary location from gradient-oriented profiles versus measure of confidence.

5.0 BOUNDARY PROFILES APPLIED TO CORE ATOMS

Boundary profiles were applied to a medial approach to shape analysis, called *core atoms*, which relate multiple boundary points to common central points deep within the object [16]. A core atom is formed from pairs of boundary points that face each other across the intervening distance. As already discussed, core atoms as previously implemented employ simple gradient boundary points and therefore have several limitations. The first limitation derives from the ambiguity of gradient direction depending on the background. In addition, core atoms only provide gradient strength rather than absolute intensity, making it impossible to differentiate particular tissue types on either side of the boundary by intensity. These limitations of core atoms can be overcome by providing intensity information. Gradient-oriented profiles do exactly this. Thus it seems natural to apply gradient-oriented profiles to core atoms, enhancing the performance of the core atoms by providing absolute intensity information. Core atoms based on gradient-oriented profiles have the added advantage of greater accuracy in estimating boundary location.

5.1 Core Profiles

For the purposes of this dissertation, the term *core profile* has been given to the next-generation core atom enhanced by gradient-oriented profiles. The reason for this was mostly for the sake of clarity. A core atom is no more than two boundary points connected in image space, if sufficient medialness exists, and is represented by the center point between these two points. Core profiles are still two connected boundary points, represented by a center point, but differ in having new medialness requirements for the two boundary points (boundary profiles). Also, by means of the boundary profiles core profiles contain intensity information. The process of constructing core profiles is presented in section 5.2 and resulting performance of core profiles on three-dimensional cardiac phantom data can be found in section 6.2. These results also include a simple and preliminary method for volume measurement in a cardiac phantom corresponding to the LV and the myocardium.

The theory behind the construction of core profiles is very similar to that of core atoms. A brief review will be included here, but a more detailed explanation can be found elsewhere [15, 16]. Rather than consisting of two boundary points like a core atom, a core profile consists of two boundary profiles. The collection of the boundary profiles follows the methods as discussed in section 3.1. For each boundary profile, search regions for other boundary profiles are established along the image gradient in both directions. Pairs of boundary profiles are then collected that meet the following requirements:

- (i) The distance between the boundary profiles (estimated boundary locations) is within a specified range. This permits savings in computation if there is *a priori* information about the expected width of the object. This distance is termed the *scale* of the resulting core profile and is defined as

$$\|\vec{s}_{1,2}\| = \|b_2 - b_1\|, \quad (23)$$

where $s_{\min} \leq \|\vec{s}_{1,2}\| \leq s_{\max}$ and the expected width is between s_{\min} and s_{\max} . The vector $\vec{s}_{1,2}$ indicates the direction from the first boundary profile location b_1 to the second boundary profile location b_2 and the core profile is defined to be located at the midpoint between the two boundary profiles. Figure 17 illustrates a core profile created across an object of intensity I_1 against a background of intensity I_2 .




Figure 17 A core profile spanning an object containing voxels of intensity I_1 (shaded area), against a background of intensity I_2 . The core profile consists of two boundary profiles separated by $\vec{s}_{1,2}$ and a center point midway between the boundary locations of the profiles. I_{ij} is the intensity estimate for region i by core profile b_j . The angle between the core vector $\vec{s}_{1,2}$ and the individual gradient directions does not have to be zero (see (ii) of requirements). Also, the exterior voxel intensities I_2 do not have to be the same for each boundary profile (see (iii) of requirements).

(ii) The face-to-faceness, defined as

$$F(b_1, b_2) = \left| \left(\frac{\vec{s}_{1,2}}{\|\vec{s}_{1,2}\|} \cdot \vec{n}_1 \right) \left(\frac{\vec{s}_{2,1}}{\|\vec{s}_{2,1}\|} \cdot \vec{n}_2 \right) \right|, \quad (24)$$

where n_i ($i = 1,2$) is the orientation of the i^{th} boundary profile, is sufficiently close to 1, while still allowing some variation in the relative orientations of the boundaries.

(iii) For the candidate boundary profiles, the intensity estimates for the space they traverse must both be high-confidence estimates (established in section 3.5) and have similar values within some pre-determined tolerance. As illustrated in Figure 17, I_{ij} ($i, j = 1,2$) is the intensity estimate for region i by boundary profile b_j . For example, I_{12} is the intensity estimate of region I_1 by boundary profile b_2 and I_{11} is the intensity estimate of region I_1 by boundary profile b_1 . The measure of confidence placed on the intensity estimate for region i by profile b_j is denoted z_{ij} ($i, j = 1,2$). Thus, profile b_1 estimates intensity I_1 with high-confidence if $z_{11} > threshold(z_{11})$. Core profile $\vec{s}_{1,2}$ is acceptable if all of the following are met:

$$(a) \quad z_{11} > threshold(z_{11}) \quad (25)$$

$$(b) \quad z_{12} > threshold(z_{12}) \quad (26)$$

$$(c) \quad |I_{11} - I_{12}| < \tau, \text{ where } \tau \text{ is an intensity tolerance} \quad (27)$$

Equations 25 and 26 guarantee that the intensity estimates for I_1 by both profiles are estimated with high-confidence and equation 27 states that these estimated intensities must be nearly equal.

Any boundary profile may link to several other boundary profiles, thus becoming involved in a multiple number of core profiles. The additional intensity information delivered from gradient-oriented profiles may permit core profiles to be formed in both directions of the gradient, resulting in a network of core profiles as shown in Figure 18. Core profiles $\vec{s}_{1,2}$, $\vec{s}_{2,3}$, and $\vec{s}_{3,4}$ are termed homogeneous core profiles because their entire lengths span local regions homogenous intensity. Core profile $\vec{s}_{1,4}$, is a *heterogeneous* core profile because it spans

multiple regions of varying intensities. Since only homogeneous core profiles are desirable, the additional constraint is required that the intensity must be roughly homogenous throughout the interior region.




Figure 18 A network of four core profiles spanning 3 different regions of homogenous voxel intensity (solid lines) and 1 region of heterogeneous voxel intensity (dashed line). Heterogeneous core profiles span regions of differing intensity.

Core profiles as described above, only pay attention to the *interior* intensity, and are therefore essentially independent of background. However, it is possible to require boundary profiles to only pay attention to *exterior* intensities instead. Instead of *medialness*, such core atoms would detect what could be called *lateralness*. For example, in a cross-section of the human torso, such core profiles would traverse the entire torso, ignoring internal organs, and base their creation on the common intensity of air exterior to the body. This type of core atom will not be explored further here because these core profiles may not be forming on real objects at all.

5.2 Measuring Medial Properties With Core Profiles: Medial Densities

Populations of homogeneous core profiles can be analyzed for a measure of medialness following the same methodology that Stetten has developed for core atoms [15]. A brief overview is included here. For 3D data, there are three basic core atom configurations; “koosh-ball,” “bed-of-nails,” and “spokes-of-a-wheel.” The “koosh-ball” configuration occurs when core atoms are formed within a spherical surface. The core atom vectors form the koosh spikes and the centers of the core atoms form the cluster near the center of the sphere. Data containing a cylinder results in a “spokes-of-a-wheel” arrangement, having core atom centers located along the cylindrical axis. Core atoms that form on the surface of a slab will yield the “bed-of-nails” configuration across the interior of the slab, with core atom centers clustered on a 2D plane between the surfaces of the slab (Figure 19).




Figure 19 Top shows basic shapes in dark gray with corresponding core in light gray. Below these are shown the corresponding basic core atom configurations (from left to right): “koosh-ball,” “spokes-of-a-wheel,” and “bed-of-nails.” (courtesy of George Stetten [15]).

For a given population of core atoms, their orientation vectors can be statistically analyzed with eigenanalysis to yield a measure of dimensionality and an overall estimate of orientation. The resulting eigenvalues $\lambda_1 < \lambda_2 < \lambda_3$ define the dimensionality of the core and eigenvectors $\vec{a}_1, \vec{a}_2, \vec{a}_3$ define a specific coordinate system for orientation. The eigenvectors are ordered such that \vec{a}_1 is most orthogonal to the population of core atom vectors and \vec{a}_3 is least orthogonal to the set of core atom vectors.

The eigenvalues are normalized and greater than or equal to zero. A value of 0 indicates that the corresponding eigenvector is completely orthogonal to every core atom. This is true for \vec{a}_1 of a cylinder and \vec{a}_1 and \vec{a}_2 of a slab. For the sphere, none of the eigenvectors is orthogonal to every core atom since the core is a point. Since $\lambda_1 + \lambda_2 + \lambda_3 = 1$, the eigenvalues can be viewed as a system with only 2 independent variables, λ_1 and λ_2 . Further constraints on the system limit values for λ_1 and λ_2 to $\lambda_1 \leq \lambda_2$ and $\lambda_2 \leq \frac{1-\lambda_1}{2}$, which defines a triangular domain called the lambda triangle (Figure 20).

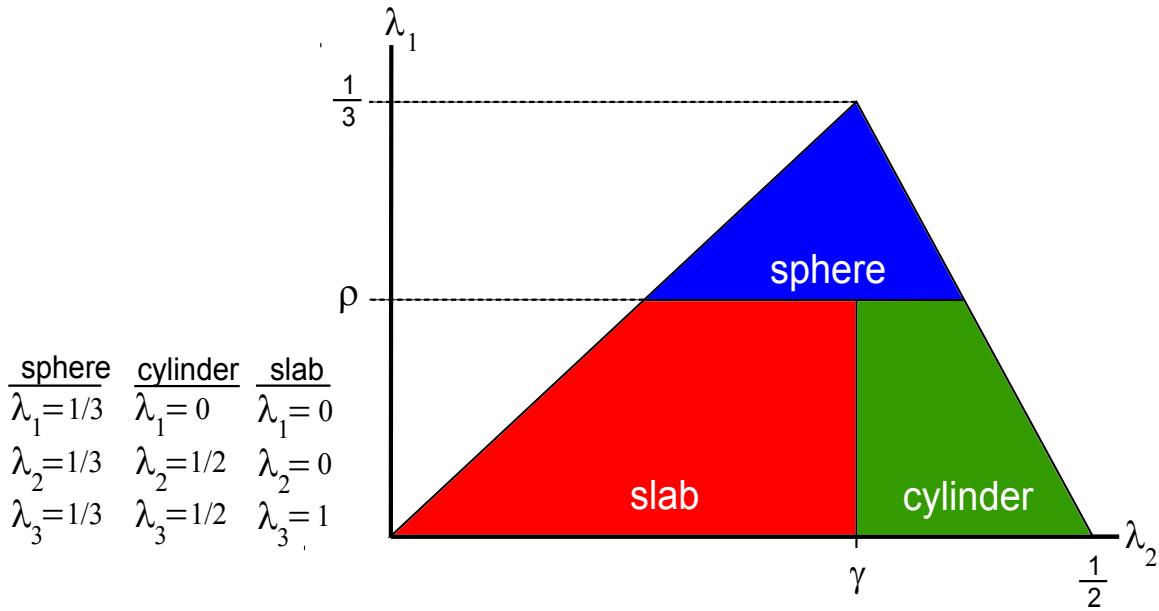


Figure 20 Lambda triangle

The triangle's vertices represent the three basic shapes. All possible sets of eigenvalues are bounded by the triangle. Thresholds placed on λ_1 and λ_2 (ρ and γ), divide the triangle into 3 compartments designating dimensionality for a given population of core atoms as an integer value:

- Sphere = 0
- Cylinder = 1
- Slab = 2

For this work, the thresholds ρ and γ follow the arbitrary values established by Stetten, $\rho = \frac{1}{5}$ and $\gamma = \frac{1}{3}$. This simplified designation of dimensionality is primarily used for visualization, with the medial densities displayed as lines with the following color codes:

- Sphere = blue
- Cylinder = green
- Slab = red

Medial properties measured in this manner are prone to displacement from the true core, a result of the face-to-faceness threshold and sampling artifact. This may cause a misclassification of local dimensionality. For instance, consider the cross section of a cylinder. Core atoms formed should ideally traverse the center axis of the cylinder.

However, if they are displaced from the true core of the cylinder, a slabness measure in the periphery of the true cylindrical core will result. This problem is overcome by taking the location of the core atom populations into consideration and searching for medial neighbors to cluster core atoms, using the method described by Stetten [16].

6.0 CORE PROFILE TEST DATA

Core profiles were tested on a 3D data set, 100 voxels cubed, 8-bit voxels. The test target consisted of two concentric spheres, chosen to crudely represent an ultrasound scan of the heart. The inner sphere represented the endocardium with a radius of 15 voxels. The outer sphere represented the epicardium and had a radius of 30 voxels. The interior intensity of the smaller sphere was 32 (LV cavity), the shell between the spheres had an intensity of 128 (myocardium), while the exterior of the larger sphere had an intensity of 64 (background). This test data is shown below in Figure 21.

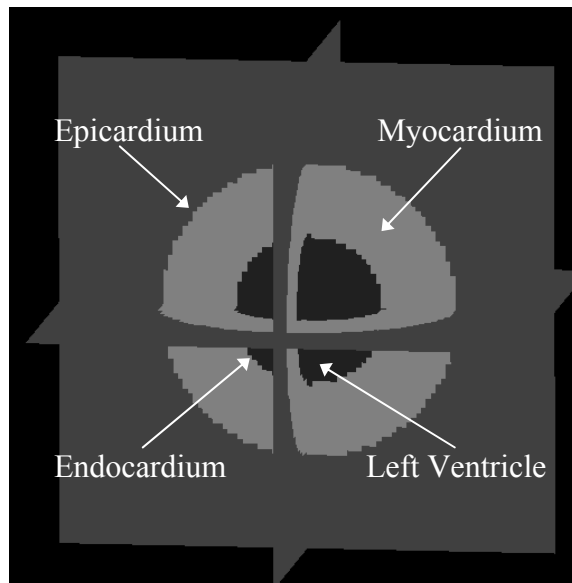


Figure 21 Test data simulating a 3D ultrasound image of a heart. The interior of the smaller sphere is of intensity 32 and represents the left ventricle, the shell is of intensity 128, representing the myocardium, and the background is of intensity 64.

Using the DoG kernel, as described in section 3.1, a total of 1,936 boundary candidates were found (shown in Figs. 22a and 22b). The parameters for the size of the ellipsoid, bin size, and footprint type were as described in

section 4.1. Of the resulting 1,936 profiles, 384 profiles were eliminated because they did not meet the constraints defined in section 3.5, leaving 1,552 profiles for analysis (shown in Figs. 22c and 22d).

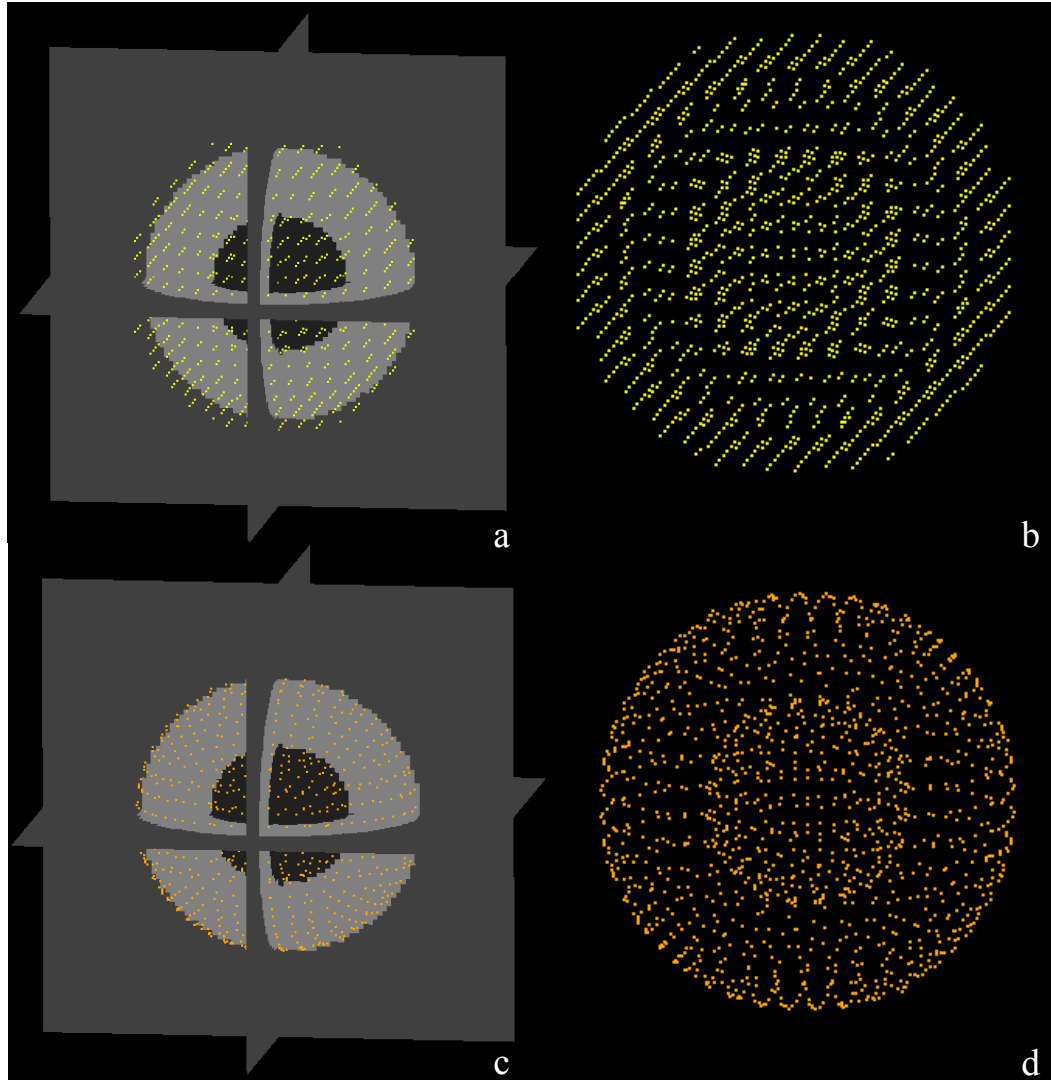
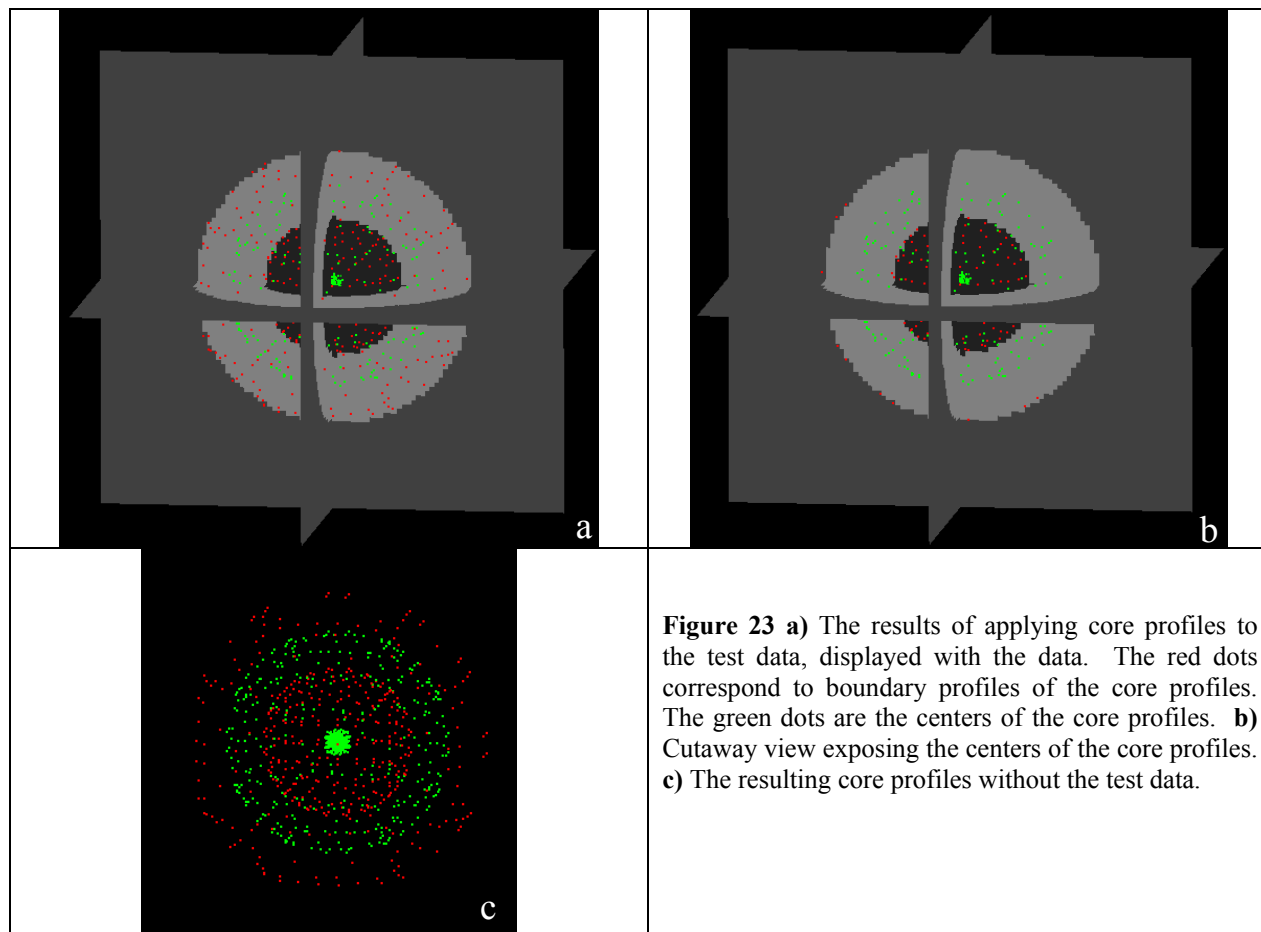
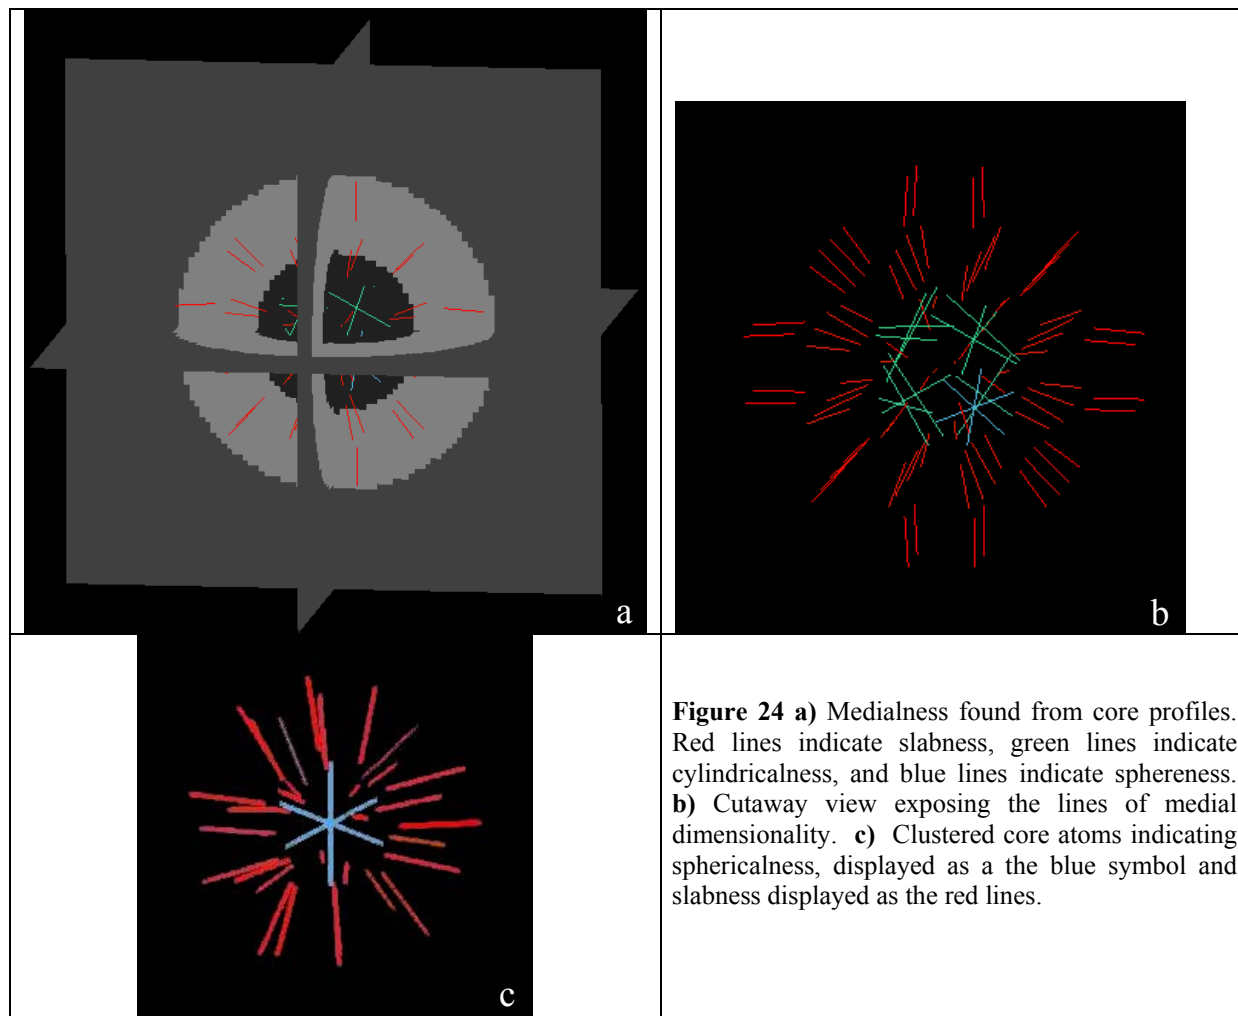


Figure 22 **a)** The boundary points found with the DOG kernel on a regularly sampled grid overlaid on the data. **b)** The grid of boundary points seen more clearly without the data. **c)** The resulting boundary *profiles* found from the boundary *points* in (a). **d)** The same boundary profiles shown in (c) without the data. Notice the profiles in (c) and (d) are more accurately located on the boundaries without the sampling artifact seen in (a) and (b).

For the remaining profiles, the measures of confidence for the classification parameters were calculated as derived in section 3.5. There were a total of 98,771 boundary comparisons of which 10,110 met the requirements as discussed in section 5.1 for core profile formation. In Figs. 23a through 23c, the core profiles are shown by their boundary profiles and center locations. The boundary profiles are shown as green dots and the core profile centers

as red dots. Qualitatively, the centers are located as expected. The core profiles that form across the inner sphere have center points near the center of that sphere. Core profiles that form across the spherical shell have center points midway through the shell. The analysis of these core profiles is performed in section 6.2 along with a measure of the accuracy of the gradient-oriented profiles for this new test model, to ensure that they perform in an effective, consistent manner. Figures 24a and 24b show the medialness of the sphere using the techniques described in section 5.2. As expected, the shell around the inner sphere was found to demonstrate slabness (red lines). The inner sphere demonstrates some expected sphericalness (blue lines), but also some cylindricalness (green lines). The cylindricalness is due to the core profile sample being too far away from the theoretical core of the sphere. By using the clustering methods briefly discussed in section 5.2, the inner sphere is correctly labeled as a sphere. The sphere is displayed as 3 orthogonal blue lines placed in the center of the core profile cluster with lengths equal to the average core profile length of the cluster. Slabs were found between the two spheres as expected, and are shown as thick red lines with length approximately equal to the thickness of the shell.





6.1 Measuring the Accuracy of Gradient-Oriented Profiles

Unlike the single sphere (Chapter 4) the test model in chapter 5 contained two spheres and generated two populations of boundary profiles. The accuracy of these boundary profiles can be examined in the same way as discussed in section 4.1 and 4.2, if the populations of the profiles are first separated by their estimation of the radius of the boundary. There were 280 profiles near the inner boundary and 1,263 profiles near the outer boundary. Figure 25 shows the distribution of error in estimating the intensity values on either side of both boundaries as a function of μ . Results are consistent with previous results, showing that when μ is near the center of the sampling region, profiles are able to estimate the intensity on either side of the boundary better than when μ is located away from the center of the sampling region.




Figure 25 Both graphs show the error in estimating intensities on either side of a boundary as a function of $\mu_{relative}$. The left graph shows the error in estimating the intensities on both sides of the inner boundary. The right graph shows the error in estimating the intensities on both sides of the outer boundary.

Using Eq. 15, the accuracy of the profiles in estimating the true boundary location can be measured. Out of the 1,152 profiles, 1488 (95.87%) were able to estimate the location of a boundary within one voxel. By comparison, the simple gradient (DoG) detector had only 456 out of the 1,152 (29.38%) boundary points within one voxel of a boundary. Figure 26 compares the error in estimating the true boundary locations by profiles and the DoG detector.

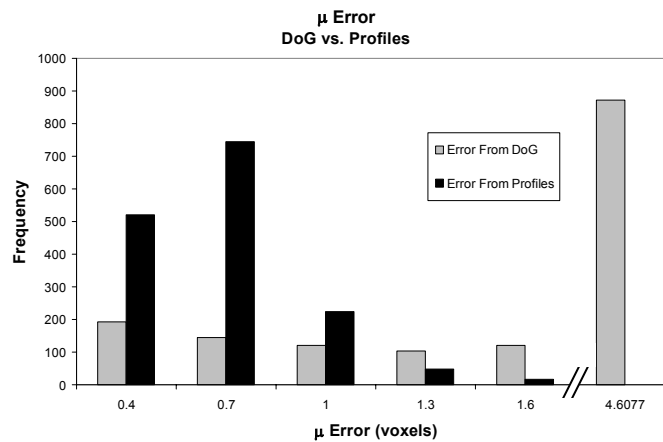


Figure 26 Comparison of errors in estimating boundary locations with the DoG detector versus gradient-oriented profiles. 96% of the total profiles as opposed 29% of the total DoG kernels estimated a boundary location within one voxel.

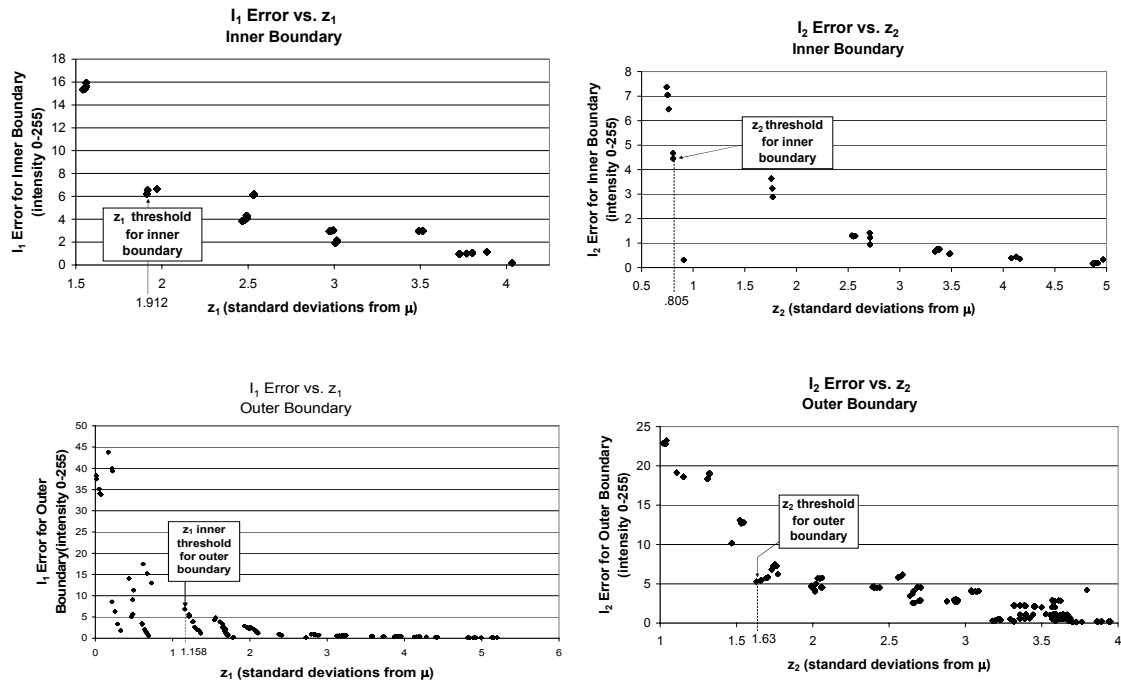


Figure 27 Intensity measures of confidence for both boundaries.

The measures of confidence for the profile parameters were calculated using the methods described in section 3.5. Plotting I_1 versus z_1 (Figure 27) shows the error in estimating the inner intensity of the inner boundary I_1 as a function of its measure of confidence z_1 . When z_1 is greater than 1.9, the error in estimating I_1 is less than 7 units of intensity (8-bit data with a full scale of 0-255). With this threshold, 80% of the inner boundary profiles are considered acceptable for estimating I_1 . The error in estimating the outer intensity of the inner boundary I_2 as a function of its measure of confidence z_2 can be found by plotting I_2 versus z_2 . If an intensity error of 8 is acceptable for estimating I_2 , 100% of the inner boundary profiles can estimate I_2 with high-confidence.

Similar plots for the outer boundary show the error in estimating the inner and outer intensity of the outer boundary as a function of their measures of confidence. A threshold of 1.1 on z_1 accepts profile estimates (84%) of the inner intensity of the outer boundary I_1 to within 7. A threshold of 1.6 on z_2 permits 84% of the outer boundary profiles to estimate, with high-confidence, the outer intensity I_2 .

6.2 Analysis of Core Profile Population

The two concentric spheres generate two populations of core profiles: those that span the inner sphere, and those that span the shell between the spheres. Since only homogenous core profiles are being considered, core profiles that span the outer sphere are not allowed. Recall that each core profile can be described by its two boundary profiles, its center point, and its scale. In the previous section, it was shown that the boundary profiles composing core profiles are able to accurately locate the true boundaries and estimate the intensities of the test model for an ultrasound image of a heart. We now show that the center point and scale of the core profiles created are consistent with the expected results. The spherical nature of the data simplifies this analysis. We denote the center of the two spheres as the “origin”. The center points of the core profiles should either be near the origin spheres or half way between the two spheres. Figure 28 shows the distance between the centers of the core profiles and origin. Half-way between the two spheres is 22.5 units from the origin, and corresponds to the second peak on the graph.

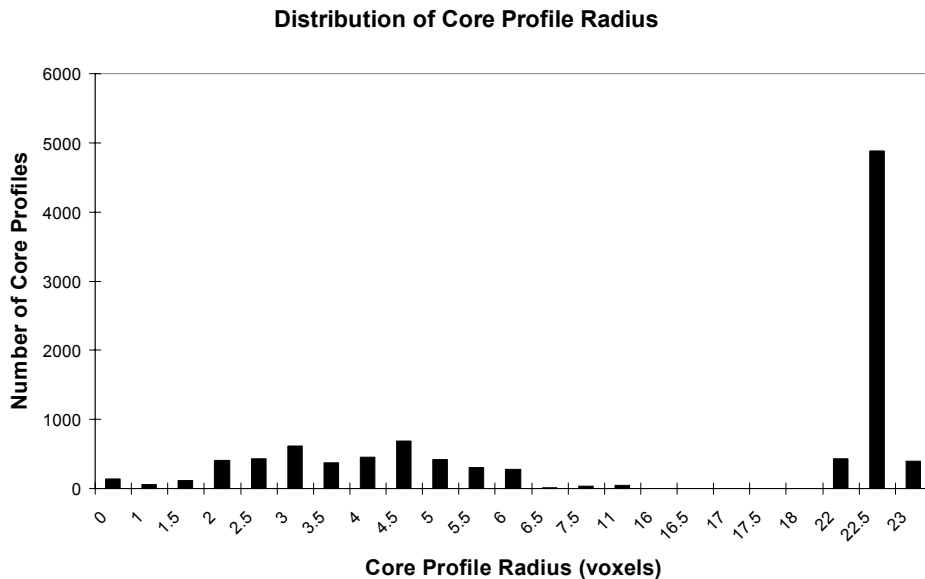


Figure 28 The distribution of center points of the core profiles. The center of the sphere is at 0 and the center of the slab between the spheres is at 22.5.

Many core profiles have their centers located very close to the midpoint of the slab, but there is a greater spread for the centers of core profiles formed across the inner sphere. This spread is expected because the error in gradient orientation has a greater effect at the larger scale of the inner core profiles.

Core profiles (homogeneous) formed across the inner sphere should have a scale of 30, while those formed across the outer sphere (heterogeneous) should have a scale of 60 voxels. Those formed across the slab should have a scale near 15 voxels. Figure 29 shows these three populations of core profiles. Heterogeneous core profiles have been included for the purposes of this experiment, although they are usually excluded.

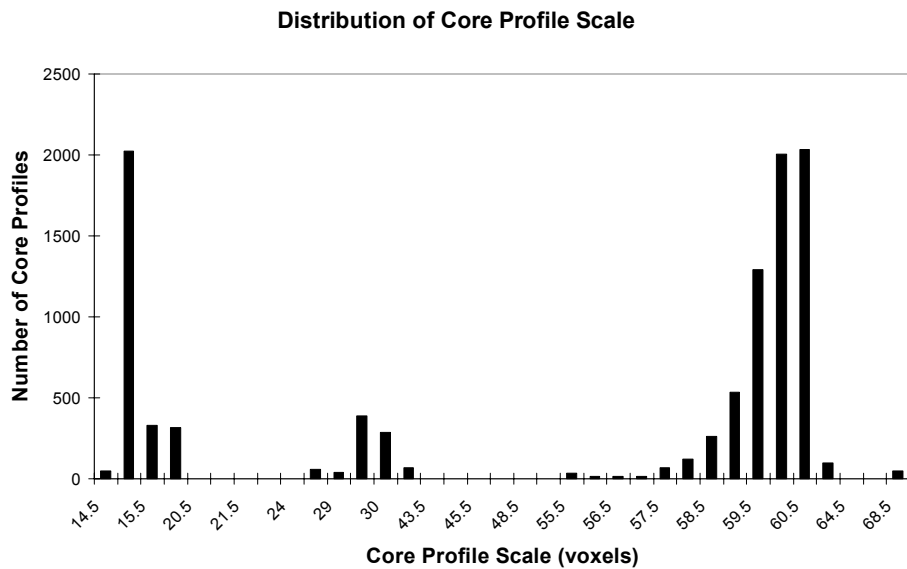


Figure 29 Distribution of core profile scale consistent with expected results.

6.3 Measures of Volume

The expected volume in the test data corresponding to the simulated left ventricle (LV) and myocardium in the concentric spheres can be calculated using the equation for the volume of a sphere,

$$V = \frac{4}{3} \pi r^3. \tag{28}$$

The expected volume of the LV is calculated by using Eq. 28 with a radius of 15 voxels. Likewise, the expected volume of the entire heart (LV and myocardium) is found with Eq. 28 and a radius of 30 voxels. The expected volume of the myocardium is the difference between expected LV volume and total heart volume.

These volumes are measured using the scale of the core profiles. The core profiles with a scale between 20 and 45 voxels are averaged and used to approximate the volume of the LV by using Eq. 28. Similarly, the heart volume is found with core profiles having a scale between 45 and 70 voxels and the volume of the myocardium is the difference between these two measures. Table 6 contains a summary of these volume measures as well as an error measure (percent error) for the volume calculated by the core profiles.

Table 6

An average core profile scale after sorting them by scale permits the calculation of LV and myocardial volume. The true volumes are found simply by using the equation of a sphere and the known radii of the test data. The percent errors (PE) for volume measure and standard deviation (σ) of the core atom and profile lengths are shown in parentheses.

Method of Calculation	LV Volume (voxels)	Heart Volume (voxels)	Myocardium Volume (voxels)
Known Parameters of Data	14,137	113,097	98,960
Average Core Atom Scale	13,158 (PE = 7%, $\sigma = 2.7$)	114,082 (PE = 1%, $\sigma = 5.4$)	100,924 (PE = 2%)
Average Core Profile Scale	13,215 (PE = 6%, $\sigma = 2.1$)	111,002 (PE = 2%, $\sigma = 2.3$)	97,787 (PE = 1%)

The percent error for measuring volume with core atoms and core profiles did not differ by much. However, the standard deviations of the core profile scales used to calculate volume were less than those by core atoms. Thus, it can be concluded that the errors in boundary location seen with the DoG kernel tend to average out, so the volume measure is just as accurate as with boundary profiles, but the individual boundary profiles are more consistent. Calculating the volumes with the core profiles in this manner was possible because only spheres were involved. In real data, volumes cannot be calculated in this manner. Currently, more versatile methods of volume calculation based on core profile features are being developed and these are discussed in the following chapter.

7.0 MEASURING VOLUME WITHOUT GEOMETRICAL ASSUMPTIONS WITH CORE PROFILES ON CONCENTRIC ELLIPSOIDS

In this section computer-generated 3D data consisting of two concentric ellipsoids was used to illustrate two potential methods for automatic volume measure with core profiles without geometric assumptions. The two concentric ellipsoids served as a crude model of the balloon phantom. The data containing the two concentric ellipsoids were 100 voxels on a side, consisting of 8-bit voxels with a background of intensity 64. The larger ellipsoid had a major axis of 70, minor axis of 50, and intensity of 128. A smaller ellipsoid was placed in the data with a major axis of 60, minor axis of 40, and intensity of 32. The result was a small ellipsoid (voxel value of 32) surrounded by a 10 voxel thickness surrounding shell of intensity of 128 (Figure 30).




Figure 30 Computer generated 3D data of two concentric ellipsoids.

Two methods developed for automated volume measures are discussed here that may prove valuable where volume measures are needed, but where explicitly defining a parametric shape is not possible. These methods are theoretically presented, but not empirically evaluated. This represents future work, which will include automated model building from medial nodes.

The first method relies on a dense population of core profiles (Figure 31 shows the population of boundary profiles from which some core profiles are formed). The method samples voxels in regions around individual core

profiles and counts voxels that have intensity similar to inner intensity of the corresponding boundary profiles. For this reason, it is called the “search and count” method. It is shown in Figure 32 and described in more detail below.




Figure 31 Left shows estimations of boundary locations from gradient-oriented profiles as orange points with the data. Right shows the boundary profiles without the data, conveying the sub-pixel ability of profiles to locate boundaries.




Figure 32 Search and Count Method: Automatically measuring volume with core profiles using an intelligent fill routine.

The “search and count” method proceeds with the following steps until all core profiles have been visited:

1. Visit a core profile and define an ellipsoidal voxel-counting region that extends to both boundary profiles (green circles), centered midway between the boundary profiles (red circle). This ellipsoid region is shown as blue in figure 32.
2. Count the number of voxels within the ellipsoid region that have an intensity value similar to the inner intensity of the boundary profiles. For instance, voxel ω is added to the volume count if its intensity I_{ω} satisfies

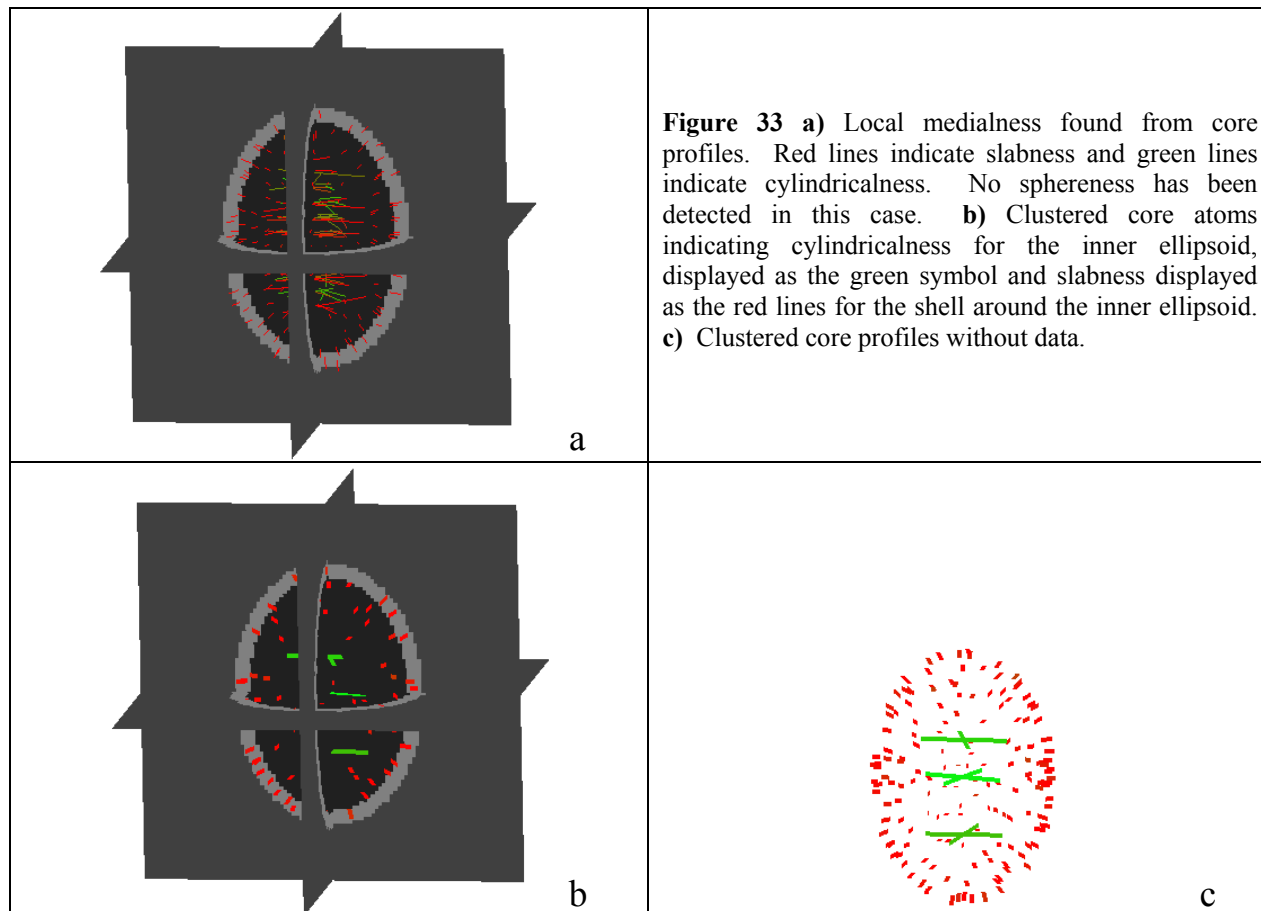
$$\left(\frac{I_{1A} + I_{1B}}{2} - \kappa \right) \leq I_{\omega} \leq \left(\frac{I_{1A} + I_{1B}}{2} + \kappa \right), \quad (29)$$

where I_{1A} is the inner intensity estimated by boundary profile A and I_{1B} is the inner intensity estimated by boundary profile B and κ is an intensity threshold chosen empirically.

3. If a voxel has been added to the volume count, it is no longer considered for the volume count, i.e. all of the voxels in the dark blue region are counted only once.

The “search and count” method will tend to be computationally inefficient as a result of repeatedly visiting voxels in overlapping counting neighborhoods, if only to check whether they have already been counted. If the population of core profiles is sparse, the volume measure will be inaccurate. The effects of boundary discontinuities and image noise are difficult to predict.

The second technique, called “medial region filling” utilizes the medial nodes described in section 5.2. First, medial nodes must be constructed. The nodes are shown below in Figure 33 for the concentric ellipsoids along with their medial densities.



As expected the shell around the inner ellipsoid exhibits slabness and the inner ellipsoid shows cylindricalness. A “medial region fill” is shown in Figure 34.




Figure 34 The use of medial nodes to measure volume without explicitly defining a surface with a “medial region fill.” Spherical regions are centered at each medial node with a diameter equal to the average core profile length. The spherical regions expand into ellipsoids to similar voxels until they collide with object boundaries. The final volume measures are found by counting voxels for similar spherical regions.

Regardless of dimensionality; a spherical fill region (blue circles) is centered at each medial node (gray circle) with a diameter equal to the average core profile length. The spheres are deformed by expanding into ellipsoids along axes orthogonal to the core atom population to include neighboring voxels that have intensity representative of the medial node intensity. The spheres continue to deform until they collide with object boundaries. The volume of the inner volume is measured by counting the number of voxels within the corresponding deformed spheres (ellipsoids). The *medial region fill* method may be more robust than the *search and count* method because it may be less sensitive to individual missing boundaries.

8.0 APPLICATION OF METHODS TO RT3D ULTRASOUND DATA

In this section gradient-oriented boundary profiles and core profiles are formed on real-time 3D ultrasound data, and measures of medialness computed. The data come from David Sahn, M.D., Oregon Health Sciences University, and contain a left ventricular balloon model [66]. The phantom data are used to explore ultrasound-specific difficulties, such as image noise, while minimizing the more complex issues found with cardiac RT3D data.

The balloon model is made up of an inner latex balloon of 10-inch diameter (representing the endocardium) and an outer balloon of 12-inch diameter (representing the epicardium). A mixture of water and ultrasound gel was inserted between the two balloons (representing the myocardium), and pure water was inserted into the inner balloon (representing the LV). Air bubbles were removed with a catheter through the necks of the balloons.

For this dataset, the “myocardial space” was filled with 50 mL of gel solution and the “LV” was filled with 100 mL of water. The balloon phantom was placed in a tank of water and imaged from beneath by a real-time 3D ultrasound machine (Model 1 machine, Volumetric Medical Imaging). The distance from the face of the transducer to the center of the balloon was between 5 and 7 cm. The entire balloon was captured in real-time within a 64x64 pyramidal volume sampled to a range of 18 cm. A representative scan is shown in Figure 35 as 2D cross-sections. The two orthogonal vertical slices are termed B-mode slices and the horizontal slice is called a C-mode slice.

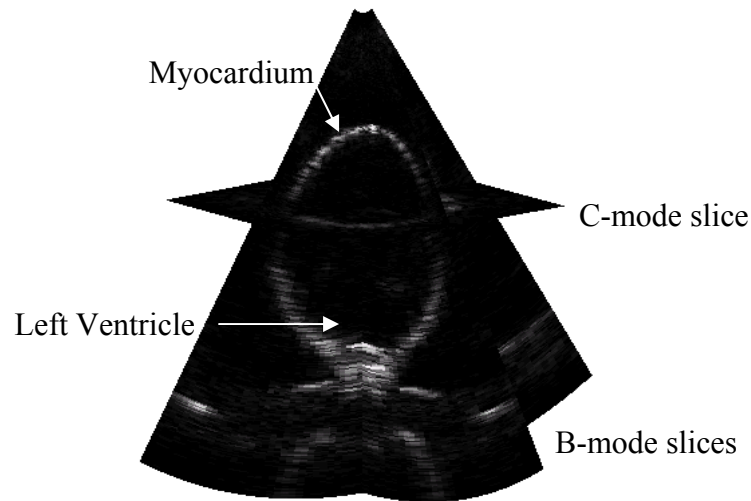


Figure 35 A 3D scan of the left ventricular balloon model displayed as 2D cross-sections. The transducer is at the top of the image and the necks of the balloons are at the bottom of the image (the circular white blob). The “reflection” at the bottom of the image is an ultrasound artifact caused by reverberations within the balloon.

Gradient-oriented boundary profiles were applied to the balloon data. The parameters previously established were used: ellipsoidal sampling neighborhood with a major axis of 10 voxels and a minor axis of 6 voxels and triangle splatting footprint with 4 splatting bins. The resulting boundary profiles are shown in Figure 36a. The profiles form a double layer at the two balloon surfaces. In the upper left corner of Figure 36a one can see the difficulty in analyzing RT3D data. The “hole” in the outer boundary makes it impossible for boundary profiles to form without dramatic over sampling, which would lead to an enormous number of boundary profiles. Boundary profiles that form in the reverberation artifact at the bottom of the scan may also present a problem, but can be neglected during core profile formation by constraining the range to contain only the balloon. Intelligent shape analysis routines, to be devised in ongoing research, will hopefully eliminate such outliers.

There are two differing populations of core profiles once the homogeneous constraint is applied: one population of core profiles across inner balloon and another across the interballoon space. Core atoms in their previous implementation would require a minimum of two passes to detect both populations of core atoms. Each pass would need the scale and gradient direction set to find the particular core atoms desired. An advantage of core profiles is they contain the image intensity information, so a gradient direction does not need to be chosen making them capable of detecting both populations of core profiles in a single pass. Formation of core profiles in 2 passes was compared to formation in 1 pass, as described next.

Core profiles were formed in 2 passes with a face-to-faceness threshold of 12 degrees and a search region either between 1 and 5 voxels or 10 and 15 voxels. The total time to form core profiles for each pass took approximately 120 secs. Both populations of core profiles constructed from the boundary profiles in Figure 36a are shown together in Figure 36b. The red points are the boundary profiles of a core profile and green points are centers of core profiles. The presence of a “hole” in the outer boundary is present and does not permit some core profiles to form across the LV space.

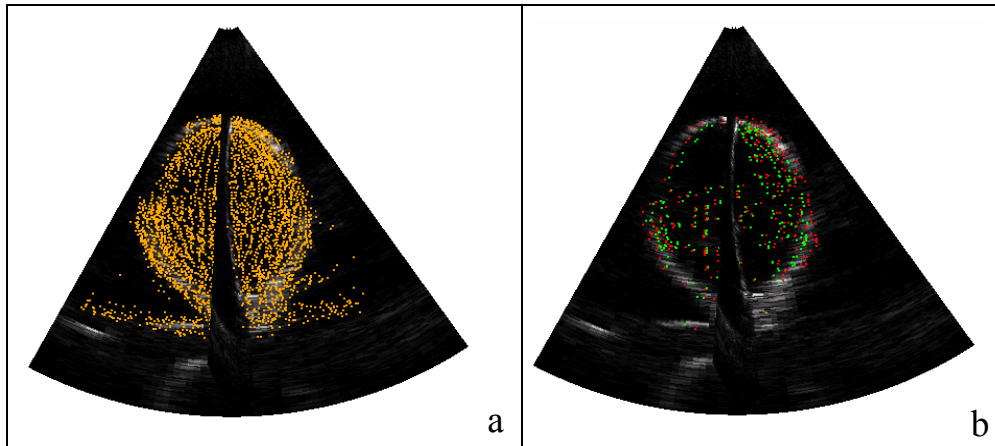


Figure 36 a) Location of boundaries found by boundary profiles. **b)** Population of core profiles found (red points are boundary profiles and green points are centers of core profiles).

After forming the core profiles, each population was clustered according to the method described in section 5.2. Figure 37a shows that the medial nodes found from the short core profiles are slab-like and Figure 37b shows that the long core profile medial nodes cluster as a cylinder.




Figure 37 a) Medial nodes found for the balloon detecting slabness of the simulated myocardium. **b)** Medial nodes detecting cylindricalness for the LV.

Core profiles were formed in a single pass with a face-to-faceness threshold of 12 degrees and a search region between 1 and 15 voxels. The total time to form these core profiles was approximately 140 secs. Core profiles can be clustered at one time by utilizing the classification of boundaries by boundary profiles. Figure 38a shows the resulting clustered profiles. Like the graphical test figure of the two concentric ellipsoids, only slab-like nodes and cylinder-like nodes were found. However, they tended to be dispersed throughout the data rather than localized along the medial cores due to core atoms forming on extraneous noise. This is not a function of the boundary profiles themselves, but an inherent problem to the core profile clustering. Resolving this is beyond the scope of this thesis, but will be addressed in future research.

In the meantime, applying specific constraints on the medial nodes can refine them. Figure 38b shows the medial nodes after applying the following constraints:

1. Cylindrical nodes are discarded if they are not located in some proximity to the long-axis of the data, i.e. through the center of the balloon.
2. Slab nodes are discarded if the distance between them and the nearest cylindrical node is larger than a pre-determine distance.

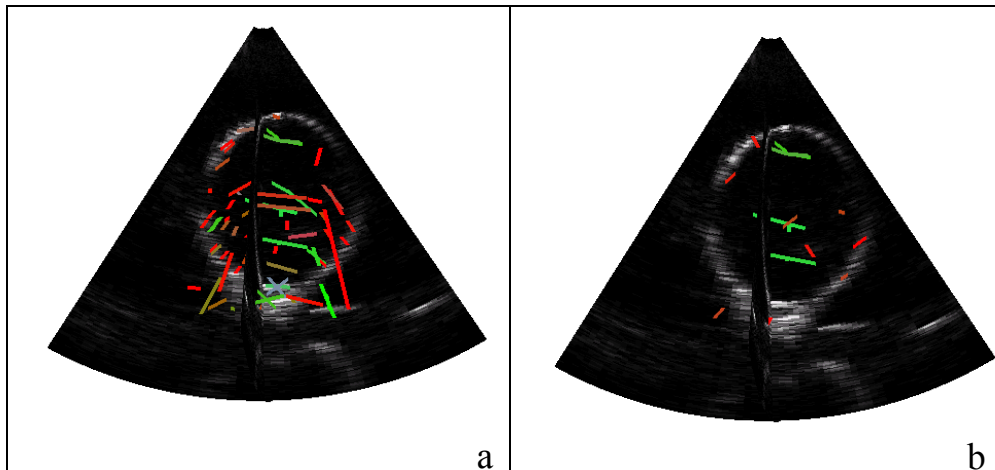


Figure 38 a) Clustered core profiles. **b)** Applying constraints to clustered profiles

Although effective at reducing the number of extraneous medial nodes, extensive *a priori* knowledge of the data is necessary. Future work will involve performing medial analysis on a training dataset then deriving a model based on medial nodes to represent the objects of interest. After the model is trained, a correspondence metric will be

calculated between its nodes and nodes of a given dataset. If sufficient correspondence does not exist, the test node is removed, thus, providing autonomous removal of extraneous medial nodes.

9.0 INSIGHT TOOL KIT CONTRIBUTION

In 1999, the National Library of Medicine of the National Institutes of Health awarded a three-year contract to develop an open-source registration and segmentation toolkit, now known as the Insight Toolkit (ITK). The toolkit, was originally developed and maintained by six primary contractors; three academic organizations, three commercial organizations, and several sub-contractors, including the University of Pittsburgh. A number of new contributors have since joined the consortium.

ITK is an open-source platform-independent software toolkit for performing registration and segmentation, intended to support the Visible Human Project for registration and segmentation. In general, it is hoped that ITK will establish a standard foundation for future image processing research. A repository of fundamental algorithms exists and validation techniques will be available to compare results of different algorithms. The software architecture is multi-dimensional and designed to handle large datasets and serves as a platform for advanced product development. Moreover, it is envisioned that ITK will create a convention for future image processing work and cultivate a self-sustaining community of software users and developers.

I have been involved with the ITK project as a student in George Stetten's lab practically since its conception. During its development, I have learned its architecture, and written base classes and tutorials. After a significant amount of time, I was able to implement the work of this thesis into the toolkit. The source code is currently available for download with the toolkit at the ITK web page <http://www.itk.org>.

10.0 FUTURE WORK

A goal of this thesis was to present methods of shape detection and measurement fast enough to take advantage of the real-time acquisition of RT3D for real-time analysis. Unfortunately, gradient-oriented boundary profiles typically require around 2 minutes to fully execute for a given volume of RT3D data. Even though the information delivered by boundary profiles is valuable for shape analysis, it cannot operate in real-time. Thus, future work will investigate methods to accelerate the formation of boundary profiles. Techniques for volume measures in RT3D were theoretically developed, but rely on either a dense number of boundary profiles or medial nodes. Future work will look into generating models from medial nodes for a more robust detection of shape. These medial node models will also permit automated volume measures based on the registration of the model and the data.

10.1 Improving the Performance of Gradient-Oriented Boundary Profiles

Voxel sampling regions and curve fitting techniques of the intensity profiles will be researched. The actual time to sample voxels in the sampling regions has not been recorded, but it is likely that this contributes to the slow computation speed of profiles in addition to the time to fit the intensity profile. The sampling regions are fixed in size, so for every boundary point the same number of voxels are sampled. If a boundary point is already near the true boundary, fewer voxels actually need to be sampled to generate an intensity profile. If a boundary point is far from the true boundary, it will most likely result in a “bad” profile anyway. Predicting each of these cases beforehand would reduce the number of times voxels in these large sampling regions would occur.

An alternative is to automatically tailor the size of each sampling region specific to the properties of the boundary point. One proposal is to sample voxels linearly along the gradient vector. A cumulative Gaussian would be fit and an “error metric” computed. The length of the sampling line would be increased until the “error metric”

was acceptable. This length would be the length of the sampling region and the width of the sampling region could be some fixed proportion of the length.

The curve fitting takes approximately .01 seconds to completion for each profile. Although this seems rapid, an image may easily contain thousands to hundreds-of-thousands of boundary points, making the total time to curve fit all of the intensity profiles on the order of 10 to 100 seconds. Either the curve fitting process needs to be made faster or the number of intensity profiles needs to be reduced prior to the curve fitting.

10.2 Medial-Node Models

Clustered core profiles, represented by a medial node, demonstrate slabness, cylinderness, or sphereness. These primitive shape features can be used to construct higher-level shapes by calculating a correspondence metric (based on the eigenanalysis results of the core profiles) between nodes in two images or between nodes in an image and a model. Higher-level metrics between pairs and triplets of nodes can also be constructed. This will probably form the basis of my Ph.D. research.

11.0 CONCLUSION

“Reading, after a certain age, diverts the mind too much from its creative pursuits. Any man who reads too much and uses his own brain too little falls into lazy habits of thinking.” -Albert Einstein

In this thesis, gradient-oriented boundary profiles have been developed as a method to parameterize boundaries. Boundary profiles average voxel intensities along the gradient vector, making them rotationally invariant and relatively insensitive to image noise. A cumulative Gaussian was fit to the collection of averaged voxel intensities yielding estimates of (1) extrapolated intensity values for voxels located far inside and outside of the boundary, and (2) anatomical boundary location. Intrinsic measures of confidence were developed to eliminate low-confidence parameter estimates. Thresholds placed on these measures of confidence allowed for high-confidence classification of boundaries.

Gradient-oriented profiles were shown to be capable of accurately parameterizing the boundary on computer-generated data. The radius of a computer-generated sphere was estimated to within one voxel of its known radius. The intensity of the sphere and image background was estimated to within 10 intensity units. Using the measures of confidence and establishing thresholds, the accuracy of boundary location and intensities estimates improved drastically, removing many of the outliers.

A thorough investigation was conducted on the effects of varying the parameters for boundary profile formation, specifically, (1) using ellipsoidal versus cylindrical neighborhoods and (2) using triangle versus Gaussian footprints. The spherical data was used to test the various parameters and it was shown that a combination of ellipsoidal sampling region and triangle footprint yields more accurate results.

Gradient-oriented boundary profiles were shown to improve an existing medial-based approach to shape analysis, called core atoms. Core atoms, a means to measure medial properties, have previously been successful in automated shape detection. However, core atoms in their previous implementation (based on simple gradient direction) were unable to form without *a priori* knowledge of object intensity relative to background. Boundary

profiles were applied to core atoms permitting the formation of core profiles, automatically and independent of background intensity. Core profiles remove any restriction on the object's intensity, allowing multiple objects of differing intensities to be located with a single application.

Core profiles were applied to 3D computer-generated data consisting of two concentric spheres, two concentric ellipsoids, as well as RT3D ultrasound data. The data containing two spheres was chosen so that the volume of the inner sphere and space between the spheres could be easily calculated with core profiles and core atoms. It was shown that calculating the volume with core profiles is more accurate than calculating the volume with core atoms, because of the improved boundary location delivered by boundary profiles.

The data containing two ellipsoids was used as a model for the RT3D ultrasound data to test two new methods of automatically measuring volume on non-parametric data with core profiles. Unfortunately, the population of core profiles and medial nodes were not dense enough for an accurate measure of volume. Methods of volume measures detailed in Chapter 9 will make use of higher-level shape analysis routines developed in my Ph.D. dissertation.

Finally, core profiles were applied to RT3D ultrasound data of a cardiac phantom. Core profiles were able to successfully locate balloon boundaries. Future work with models built from medial nodes will hopefully allow for automated left ventricular identification and measurement.

BIBLIOGRAPHY

BIBLIOGRAPHY

1. Stetten, G., et al., Real-Time 3D Ultrasound: A New Look at the Heart. *Journal of Cardiovascular Diagnosis and Procedures*, 1998. 15(2): p. 73-84.
2. von Ramm, O.T. and S.W. Smith. Real Time Volumetric Ultrasound Imaging System. *Journal of Digital Imaging*. 1990.
3. Ganz, W., et al., Assessment of left ventricular ejection fraction and volumes by real-time, two dimensional echocardiography. *Circulation*, 1971. 27: p. 393-396.
4. Folland, E.D., et al., Assessment of left ventricular ejection fraction and volumes by real-time, two-dimensional echocardiography. *Circulation*, 1979. 60(4): p. 760-766.
5. Collins, M., et al. Assessment of regional wall motion abnormalities using real-time 3D echocardiography. *American Heart Association 69th Scientific Sessions*. 1996.
6. Collins, M., et al., Assessment of regional wall motion abnormalities with Real-Time 3D Echocardiography. *Journal of the American Society of Echocardiography*, 1999. 12(1): p. 7-14.
7. Sahn, D., et al. Calculation of aortic regurgitant stroke volume from Real-Time 3D ultrasound images as the difference between left and right ventricular 3D determined stroke volumes in chronic animal model of aortic regurgitation: A new method with potentially wide application in cardiology. *AHA 71th Scientific Sessions*. 1998.
8. Li, J.S., et al. Evaluation of ventricular outflow tract abnormalities using real-time three dimensional echocardiography. in *American College of Cardiology 46th Annual Scientific Session*. 1998.
9. Sachdev, V., et al. Validation studies for Real-Time 3D Echocardiographic determination of Left Ventricular (LV) diastolic and systolic cavity and stroke volumes: Comparisons of digital sonomicrometric arrays, conductance catheter and EM flow meter measurements. *American College of Cardiology 48th Scientific Sessions*. 1999 (accepted). New Orleans.
10. Xin, J., et al., Validation of Real-Time Three-Dimensional Echocardiography for Quantifying Left Ventricular Volumes in the Presence of a Left Ventricular Aneurysm: In Vitro and In Vivo Studies. *Journal of the American College of Cardiology*, 2000. 36(3): p. 900-907.
11. Centers for Disease Control and Prevention, Nation Center for Health Statistics, and American Heart Association, *Causes of Death for All Americans in the United States, 1999 Final Data*. 2000.
12. Canny, J., A Computational Approach to Edge Detection. *IEEE Pattern Analysis and Machine Intelligence*, 1986. 8(6): p. 679-698.
13. Tamburo, R. and G. Stetten. Gradient oriented profiles for unsupervised boundary classification. *Applied Imagery Pattern Recognition (AIPR) Workshop*. 2000. Washington, DC: IEEE Computer Society.

14. Tamburo, R. and G. Stetten, Gradient-oriented profiles for automated unsupervised boundary classification and their application to core atoms towards shape analysis. *International Journal of Image and Graphics*, 2001. 1(4): p. 659-680.
15. Stetten, G., Automated Identification and Measurement of Cardiac Anatomy Via Analysis of Medial Primitives, in Dept. Biomedical Engineering. 1999, University of North Carolina, Chapel Hill.
16. Stetten, G. and S.M. Pizer, Medial-Node Models to Identify and Measure Objects in Real-Time 3-D Echocardiography. *IEEE Transactions on Medical Imaging*, 1999. 18(10): p. 1025-1034.
17. Mlsna, P. and J. Rodriguez, Gradient and Laplacian-Type Edge Detection. From *Handbook of Video and Image Processing*, 2000. edited by: Al Bovik: p. 415-447.
18. Pratt, W., *Digital Image Processing*. Second ed. 1991, New York: John Wiley and Sons.
19. Petrou, M. and P. Bosdogianni, *Image Processing: The Fundamentals*. 1999, New York: John Wiley and Sons.
20. Umbaugh, S.E., *Computer Vision and image Processing: A Practical Approach Using CVIP Tools*. 1997, New York: Prentice Hall.
21. Jain, A., *Fundamentals of Digital Image Processing*. Prentice Hall Information and System Sciences Series, ed. T. Kailath. 1989, New Jersey: Prentice Hall.
22. Marr, D., *Vision*. 1982: W. H. Freeman & Co.
23. Stetten, G., *Vision System, Guidelines for Biomedical Engineering Handbook*, J.D. Bronzino, Editor. 1995, CRC Press: Boca Raton, Florida. p. 33-42.
24. Silverthorn, D.U., *Human Physiology: An Integrated Approach*. 2nd ed. 2001, Upper Saddle River, NJ: Prentice Hall.
25. Hartline, H.K., The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. *American Journal of Physiology*, 1938. 121: p. 400-415.
26. Hartline, H.K., The receptive fields of optic nerve fibers. *American Journal of Physiology*, 1940. 130: p. 690-699.
27. Hartline, H.K., The effects of spatial summation in the retina on the excitation of the fibers of the optic nerve. *American Journal of Physiology*, 1940. 130: p. 700-711.
28. Kuffler, S.W., Discharge patterns and functional organization of the mammalian retina. *Journal of Neurophysiology*, 1953. 16: p. 37-68.
29. Rodieck, R.W., Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Visual Research*, 1965. 5: p. 583-601.
30. Marr, D. and E. Hildreth, Theory of Edge Detection. *Proceedings of the Royal Society*, 1980. 207(1167): p. 187-217.
31. Freeman, W. and E. Adelson, The Design and Use of Steerable Filters. *IEEE Transactions in Pattern Analysis and Machine Intelligence*, 1991. 13(9): p. 891-906.
32. Imatron, G., *GE Imatron's Electron Beam Tomography (EBT)*. 2002.

33. von Ramm, O.T., et al., Real-time, three-dimensional echocardiography: the first human images. *Circulation*, 1991. 84, suppl. II (4): p. 685.
34. von Ramm, O.T., et al., Real-Time three dimensional echocardiography: The first human images. *Circulation abstract*, 1991.
35. Smith, S.W., et al., Improved real-time volumetric ultrasonic imaging. *Ultrasonic Imaging*, 1992. 14: p. 186-211.
36. Smith, S.W., J. Henry G. Pavy, and O.T. von Ramm, High-speed ultrasound volumetric imaging system - Part I: Transducer design and beam steering. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 1991. 38(2): p. 100-108.
37. von Ramm, O.T., S.W. Smith, and H.G. Pavy, Jr., High-speed ultrasound volumetric imaging system - Part II: Parallel processing and image display. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 1991. 38(2): p. 109-115.
38. Sheikh, K., S. Smith, and O.T. Ramm, Real-Time, Three Dimensional Echocardiography: Feasibility and Initial Use. *Echocardiography*, 1991. 8(1): p. 119-125.
39. Nadkarni, S.K., et al. Theoretical estimation and measurement of accuracy and variability in left ventricular volume and mass quantification using Three-Dimensional Echocardiography. 22nd Annual EMBS International Conference. 2000. Chicago, IL: EMBS.
40. McInerney, T. and D. Terzopoulos, Deformable models in medical image analysis: A survey. *Medical Image Analysis*, 1996. 1(2): p. 91-108.
41. Terzopoulos, D. and D. Metaxas, Dynamic 3D Models with Local and Global Deformations: Deformable Superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1991. 13(7): p. 703-714.
42. Malassiotis, S. and M. Strintzis, Tracking the Left Ventricle in Echocardiographic Images by Learning Heart Dynamics. *IEEE Transactions on Medical Imaging*, 1999. 18(3): p. 282-290.
43. Chalana, V., et al., A Multiple Active Contour Model for Cardiac Boundary Detection on Echocardiographic Sequences. *IEEE Transactions on Medical Imaging*, 1996. 15(3): p. 290-298.
44. Hunter, I.A., et al. Detection of Echocardiographic Left Ventricle Boundaries Using Neural Networks. in *Computers in Cardiology*. 1993. London, UK.
45. Hammarneh, G. and T. Gustavsson, Combining snakes and active shape models for segmenting the left ventricle in echocardiography. *IEEE Computers in Cardiology*, 2000. 27.
46. Setarehdan, S.K. and J.J. Soraghan, Automatic Left Ventricular Feature Extraction and Visualisation from Echocardiographic Images. *Computers in Cardiology*, 1996.
47. Coppini, G., R. Poli, and G. Valli, Recovery of the 3-D shape of the left ventricle from echocardiographic images. *IEEE Transactions on Medical Imaging*, 1995. 14(2): p. 301-317.
48. Stetten, G., et al. 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. *American College of Cardiology 48th Scientific Sessions*. 1999. New Orleans.
49. Fleishman, C.E., et al. Quantitative assessment of LV ischemic risk volume using real-time three-dimensional echocardiography. *American Heart Association 69th Scientific Sessions*. 1996.

50. Fleishman, C.E., et al. Identification of congenital heart defects using real-time three-dimensional echo in pediatric patients. American Heart Association 69th Scientific Sessions. 1996.
51. Fleishman, C.E., et al. Real-time three-dimensional echocardiography: Measurement of left ventricular mass in dogs. American Heart Association 69th Scientific Sessions. 1996.
52. Ota, T., et al. Manual and semiautomated measurement of left ventricular volume using real-time, three-dimensional echocardiography in vivo. American College of Cardiology Annual 46th Scientific Session. 1997.
53. Ota, T., et al. Accuracy of manual and semiautomated methods for volume determination by real-time, three-dimensional echocardiography. American Heart Association 69th Scientific Sessions. 1996.
54. Ota, T., et al. Measurement of left ventricular volume by real-time, three-dimensional echocardiography in dogs. American Heart Association 69th Scientific Sessions. 1996.
55. Angelini, E., et al. Quantification of LV Volumes with 4D Real-Time Echocardiography. EMBS International Conference. 2000. Chicago IL: IEEE.
56. Angelini, E., et al., LV Volume Quantification via Spatiotemporal Analysis of Real-Time 3-D Echocardiography. IEEE Transactions on Medical Imaging, 2001. 20(6): p. 457-469.
57. Corsi, C., et al., Real-Time 3D echocardiographic data analysis for left ventricular volume estimation. Computers in Cardiology, 2000. 27: p. 107-110.
58. Stetten, G. and S. Pizer. Automated Identification and Measurement of Objects via Populations of Medial Primitives, with Application to Real Time 3D Echocardiography. Information Processing in Medical Imaging. 1999. Vesigrad, Hungary: Springer.
59. Blum, H. and R.N. Nagel, Shape description using weighted symmetric axis features. Pattern Recognition, 1978. 10: p. 167-180.
60. Burbeck, C.A. and S.M. Pizer, Object representation by cores: Identifying and representing primitive spatial regions. Vision Research, 1995. 35(13): p. 1917-1930.
61. Pizer, S.M., et al., Zoom invariant vision of figural shape: the mathematics of cores. Computer Vision and Image Understanding, 1998. 69(1): p. 55-71.
62. Morse, B.S., et al., Zoom-Invariant vision of figural shape: Effect on cores of image disturbances. Computer Vision and Image Understanding, 1998. 69: p. 72-86.
63. Furst, J.D. and S.M. Pizer. Marching optimal-parameter ridges: An algorithm to extract shape loci in 3D images. in Medical Image Computing and Computer-Assisted Intervention. 1998. Cambridge, MA: Springer.
64. Furst, J.D. 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. June 1996. San Francisco, CA.
65. Otter Research Ltd, AD Model Builder: <http://otter-rsch.com/admodel.htm>.
66. Rusk, R., et al., Comparison of Ventricular Volume and Mass Measurements from B- and C-Scan Images with the Use of Real-Time 3-Dimensional Echocardiography: Studies in an In Vitro Model. Journal of the American Society of Echocardiography, 2000. 13(10).