

HARVARD UNIVERSITY
Graduate School of Arts and Sciences



DISSERTATION ACCEPTANCE CERTIFICATE


The undersigned, appointed by the
School of Engineering and Applied Sciences

have examined a dissertation entitled:


“Enhanced Ultrasound Visualization for Procedure Guidance”

presented by : Laura Brattain

candidate for the degree of Doctor of Philosophy and here by
certify that it is worthy of acceptance.

Signature  _____

Typed name: Professor R. Howe

Signature  _____

Typed name: Professor H. Pfister

Signature  _____

Typed name: Professor T. Zickler

Date: May 6, 2014

Enhanced Ultrasound Visualization for Procedure Guidance

A dissertation presented

by

Laura Brattain

to

The School of Engineering and Applied Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Engineering Sciences

Harvard University

Cambridge, Massachusetts

May 2014

© 2014 Laura Brattain

All rights reserved.

Dissertation Advisor:
Professor Robert D. Howe

Author:
Laura Brattain

Enhanced Ultrasound Visualization for Procedure Guidance

Abstract

Intra-cardiac procedures often involve fast-moving anatomic structures with large spatial extent and high geometrical complexity. Real-time visualization of the moving structures and instrument-tissue contact is crucial to the success of these procedures. Real-time 3D ultrasound is a promising modality for procedure guidance as it offers improved spatial orientation information relative to 2D ultrasound. Imaging rates at 30 fps enable good visualization of instrument-tissue interactions, far faster than the volumetric imaging alternatives (MR/CT). Unlike fluoroscopy, 3D ultrasound also allows better contrast of soft tissues, and avoids the use of ionizing radiation.

The current key challenges in increasing the clinical adoption of 3D ultrasound include limited field of view, signal dropout, low resolution and poor rendering. A cardiac gated 3D mosaicing and visualization system is developed with an aim to address the limitations above. The system integrates electromagnetic position tracking with cardiac gating and leverages parallel computing for real-time instrument tip tracking in mosaiced 3D ultrasound volumes. User studies in clinical tasks show a 46% reduction in task completion time.

To further improve the mosaicing results, a novel 3D ultrasound compounding method based on structure tensors and weighted by ultrasound incident angles is developed. This method addresses the challenge in integrating multiview US images where the intensity values at overlapping regions could be significantly different due to differences in local structure orientations with regards to ultrasound beam propagation directions. Water tank and *in vivo* studies show 20% improvement in contrast to noise ratio when compared to the commonly used intensity averaging.

In clinical settings, a system that offers ease of use is essential. A robotic control and visualization system is built to facilitate the use of automated ultrasound catheter imaging for procedure guidance. The system creates 3D panorama by controlled sweep while keeping the catheter tip at a fixed and safe location. Instrument tracking and tip identification techniques are also developed, both require complex joint coordination thus are difficult for manual control. We are able to achieve an overall accuracy of $2.2mm$ RMS in position and 0.8° in tracking. Such system has a great potential in improving the overall clinical work flow.

This work was sponsored by the United States Department of the Air Force under United States Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

Contents

Abstract	iii
Acknowledgments	ix
1 Introduction	1
1.1 Motivation	1
1.2 3D Ultrasound	2
1.3 Intra-cardiac Ultrasound	3
1.4 Procedure Guidance	3
1.4.1 Interventional cardiology	4
1.4.2 Cardiac electrophysiology	4
1.5 Thesis Outline	6
2 Real-time 3D Ultrasound Enhanced Visualization	8
2.1 Introduction	8
2.2 Materials and Methods	12
2.2.1 System design	12
2.2.2 Mosaicing process	12
2.2.3 3D volume mosaicing	13
2.2.4 ECG gated mosaicing	16
2.2.5 Instrument tip tracking	18
2.2.6 Slice views	19
2.3 Experiments and Validations	19
2.3.1 Water tank string phantom study	19
2.3.2 Mosaicing <i>in vivo</i> experiments	24
2.3.3 Instrument tip tracking	28
2.4 User Study	29
2.4.1 Study design	29
2.4.2 Results	33
2.5 Discussion	33
2.5.1 Temporal resolution and imaging time	35
2.5.2 Real-time overlay	35

2.5.3	System extensions	36
2.5.4	Application to beating heart procedures	36
2.6	Conclusion	37
3	Structure Tensor for 3DUS Multiview Compounding	38
3.1	Introduction	38
3.2	Methods	45
3.2.1	3DUS structure tensor	45
3.2.2	Design of an US incident angle weighted coherence measure	47
3.2.3	Algorithm for the compounding	49
3.3	Design Validation	50
3.4	Results	57
3.4.1	Water tank results	57
3.4.2	<i>in vivo</i> results	60
3.4.3	Contrast to noise ratio	60
3.5	Discussion and Conclusion	63
3.5.1	Multiscale compounding	63
3.5.2	GPU implementation	65
4	Automated Sweeping and Instrument Tracking with Ultrasound Catheters	66
4.1	Introduction	66
4.2	System Design	69
4.2.1	Overview	69
4.2.2	Sweeping	72
4.2.3	Instrument tracking	73
4.2.4	3D panorama of 2D ICE images acquired from sweeping	73
4.2.5	Instrument tip localization and visualization	78
4.3	Experimental Results	78
4.3.1	Sweeping	78
4.3.2	Instrument tracking results	80
4.3.3	Instrument tip visualization	81
4.3.4	ICE imaging plane thickness vs. system accuracy	81
4.4	Discussion	84
5	Conclusion and Future Work	85
5.1	Conclusion	85
5.2	Future Work	86
5.2.1	3DUS surface boundary detection for volume rendering	86
5.2.2	US catheters and robotics	88

5.3 Future of Procedure Guidance	89
Bibliography	90

Acknowledgments

It has been an incredible journey. I am grateful and honored to have had the opportunity to work with many dedicated, talented and supportive individuals. This work would not have been possible without their contributions.

First, I would like to express my deep gratitude to Professor Robert Howe, my research supervisor, for his outstanding guidance and mentorship. Through exploring project ideas, discussing research results, and writing technical papers together, he encouraged me to step out of my comfort zone and think with a critical mind. His enthusiasm for research and dedication to his students will serve as a long lasting inspiration for me in my future career.

I want to thank my advisory committee members Professor Hanspeter Pfister and Professor Todd Zickler for their time and expertise. I also want to thank Professor Robert Wood and Professor Conor Walsh for the valuable interactions.

I would like to thank MIT Lincoln Scholars Program and my group leaders, mentors and colleagues at MIT Lincoln Laboratory for their support on this research.

I could not have made much progress in my research without the help from my clinical collaborators. I want to extend my special thanks to Dr. Nikolay Vasilyev and Dr. Pedro del Nido from Boston Children's Hospital, and Cory Tschabrunn and Dr. Elad Anter at Beth Israel Deaconess Medical Center. They provided me with the valuable opportunity to valid engineering ideas in clinical settings. Dr. Vasilyev was instrumental in the *in vivo* studies I presented in this thesis. I also like to express my thank to Scott Settlemier at Philips Medical, for his extraordinary technical support and valuable knowledge on Philips ultrasound machines.

I wish to thank various former and current colleagues at Harvard Biorobotics Laboratory for their friendship and support. They include Dr. Mohsen Moradi Dalvand, Alperen Degirmenci, Molly Finn, Dr. Peter Hammer, Dr. Frank Hammond III, Leif Jentoft, Dr. Sam Kesner, Paul Loschak, Dr. Douglas Perrin, Dr. Robert Schneider, Neil Tenenholtz, Dr. Yaroslav Tenzer, Qian Wan, Dr. Raye Yeow, and Dr. Shelten Yuen. I was fortunate to have had the opportunity to work closely with one member of the laboratory, Paul Loschak. Paul

and I collaborated on the project presented in Chapter 4 of this thesis. I will not forget the days when we braved the winter storms to collect data at the hospital.

Finally, I am thankful to my family for their unwavering support, patience and understanding as I pursue this goal. Their encouragement and unconditional love helped me through the many challenges I faced during this journey. For that I am forever grateful.

To my family

Chapter 1

Introduction

1.1 Motivation

This dissertation aims to facilitate ultrasound guidance for procedures within the beating heart. Beating heart procedures are a minimally invasive alternative to the conventional cardiac surgeries, which require the use of the cardiopulmonary bypass (CPB). CPB is associated with elevated risks of stroke [80], inflammation [20], and long-term neurocognitive damage [19, 61]. Beating heart procedures have shown a significant reduction of these adverse affects while offering the advantages of reduced post-operative discomfort, faster healing times and lowered risk of infections or complications [3, 12, 53, 76, 78]. More importantly, they allow cardiologists to assess the quality of a procedure across cardiac cycles [34].

One key limiting factor in enabling the beating heart approach to a broader range of conditions is image guidance capabilities. Cardiac CT and MRI have good resolution and can produce 2D or 3D information, but are limited to pre-operative imaging due to their slow frame rate, high cost or large footprint. X-ray based fluoroscopy is the current standard for procedure guidance, but it is unable to distinguish soft tissue. Significant efforts have been devoted to the registration of pre-operative CT or MRI images with real-time data, nonetheless, this approach often leads to mismatch of the two as a result of the patient

motion or change of patient conditions. Ultrasound (US) is an imaging modality that has been widely used for the diagnosis of cardiac diseases. US is non-invasive, inexpensive, high contrast between cardiac tissue and blood, and can provide real-time diagnostic information of the anatomical region of interest. Unlike CT or MRI, in which 2D images are usually acquired at a slow rate as a stack of parallel slices and in a fixed orientation, US provides images at higher rates, and in arbitrary orientations. The high acquisition rate and arbitrary orientation of the images provide unique problems to overcome and opportunities to exploit in volume reconstruction and dynamic 3D visualization [31].

While conventional US produces 2D images, in the last decade, 3D ultrasound (3DUS) has emerged as a new imaging approach that is rapidly achieving widespread use with numerous applications. Real-time 3DUS imaging offers improved spatial orientation information relative to 2DUS, therefore is capable of visualizing complex 3D structures [21, 74]. It is a promising modality for enabling intra-cardiac beating heart procedures.

1.2 3D Ultrasound

Until recently 3D ultrasound imaging systems were based on acquiring and then stitching a series of 2D ultrasound images to form a 3D volume. The three common approaches include mechanical scanners, free-hand techniques with position sensing, and free-hand techniques without position sensing. In the past decade, phased array based probes start to gain popularity. It consists of multiple small ultrasound crystal elements, each of which can be pulsed individually. In contrast to a conventional probe which must be physically scanned to sweep the beam through an area of interest, the beam from a phased array probe can be moved electronically. A diverging pyramidal beam and returned echoes can be displayed in real time as multiple planes [31], and that is the basic principle of real-time 3DUS.

Real-time 3D echocardiography can be performed using a transthoracic (TTE) probe from outside the chest or a transesophageal (TEE) probe from inside the esophagus. TTE is the most common type of cardiac echocardiography, during which the transducer head must be

aimed through 'acoustic windows' in the thorax in order to get around the interference from the rib cage or soft tissue. The four typical acoustic windows for TTE include parasternal, apical, subcostal, and suprasternal notch. TEE, on the other hand, has the advantage of closer proximity to the heart, thus eliminates the need to use acoustic windows, minimizes interference from the lungs, and increases the resolution of the images obtained. TEE is useful for intra-operative monitoring of anesthetized patients and for patients in which TTE did not produce quality images, but it is more invasive, typically requiring anesthesia.

1.3 Intra-cardiac Ultrasound

Ultrasound catheter, or intra-cardiac echocardiography (ICE) is used to visualize the soft tissue from inside the heart. It has an ultrasound crystal array at its distal tip with side firing imaging plane that can be manipulated in two orthogonal directions through controlling the two knobs on the handle. An ICE catheter is inserted into the patient's vasculature (e.g. femoral vein) and navigated to the heart, where it acquire 2D B-mode images of cardiac structures. Compared to external probes, ICE can achieve higher quality views of targets in the near-field with higher acoustic frequencies, reducing aberration and attenuation due to layers of intervening muscle, fat, and other tissues. When compared to TEE, ICE is also associated with decreased procedure length while eliminating the risks of endotracheal or esophageal intubation and general anesthesia [9]. Recently 3D ICE catheters were reported by Siemens as being capable of acquiring ultrasound volumes from within the heart [94].

1.4 Procedure Guidance

The overall objective of this thesis work is to develop solutions that could be useful for a broad range of cardiac procedures. This section provides a brief overview of the common procedures in cardiology, which can be largely divided into two categories: interventional cardiology and cardiac eletrophysiology.

1.4.1 Interventional cardiology

Interventional cardiology is a branch of cardiology that specializes in the catheter based treatment of structural heart diseases [49, 82]. The procedures are done by initially placing a long, thin, catheter into one of the femoral veins. These vessels are accessed through a small puncture site in either the left or right groin. Local anesthesia is used in the groin to permit relatively painless introduction of a small sheath into the blood vessel. A catheter is then advanced through the sheath, up to the heart. This process is called cardiac catheterization.

A large number of procedures, including coronary angioplasty (stent placement), patent foramen ovale defect (PFO) closures and valve replacement can be performed on the heart by catheterization. For example, the interventional cardiology procedure of primary angioplasty is now the gold standard of care for an acute myocardial infarction, which occurs when an atherosclerotic plaque slowly builds up in the inner lining of a coronary artery and then suddenly ruptures, causing catastrophic thrombus formation, occluding the artery and preventing blood flow downstream [91]. Coronary angioplasty restores blood flow to the heart muscle by extracting clots from occluded coronary arteries and deploying stents and balloons through a small hole made in a major artery.

Majority of these procedures rely on X-ray based fluoroscopy visualization in combination with contrast agents. Although it is able to show instruments and electrodes, it is unable to distinguish soft tissue. In procedures such as valve replacements and PFO closures, when instrument-tissue contact is crucial, better visualization alternatives are in great need.

1.4.2 Cardiac electrophysiology

Cardiac electrophysiology (EP) deals with the diagnosis and treatment of rhythm disorders of the heart such as atrial fibrillation (AF) or ventricular tachycardia. To assess the arrhythmia, electrophysiologists study the electrocardiograms of intra-cardiac and external spontaneous electrical activity, as well as the cardiac responses to programmed electrical stimulations using catheters with electrodes at tip. Catheter based ablation procedures are

often performed as part of the treatment for arrhythmias.

AF is one of the most common arrhythmias, affecting nearly 1,000,000 adults in the US and contributing substantially to the incidence of stroke and cardiovascular mortality [94]. It involves the two upper chambers of the heart, when the heart muscles quiver instead of moving with coordinated contractions. From electrocardiography point of view, it is characterized by the presence of rapid and irregular waves that vary in size, shape, and timing [24, 30]. This results in ineffective pumping of blood to the ventricle. In a large proportion of patients with AF, the primary arrhythmia substrate is thought to be located at the openings of the pulmonary veins into the left atrium. Catheter based procedures aim to create linear destructive lesions in the tissue around the pulmonary veins to prevent propagation of the abnormal rhythm to the rest of the atrial tissue. The procedure involves inserting catheters through the femoral artery in the groin, threading them up into the heart. Through a hole pierced on the septum, the catheters get access to the four blood vessels that are conduits of the uncontrolled impulses that cause fibrillation. Ablation catheters emit radio frequency energy at the tip to cauterize the tissue and thus stop the random electrical impulses.

Similar to interventional cardiology, in current EP suite, guidance is largely provided by fluoroscopy, which is unable to image soft tissues. To compensate for this, the widely adopted approach is to generate a three-dimensional electrophysiological model resembling the shape of the atria. This is done by using an electromagnetic (EM) position sensing system that records the locations of the ablation catheter's tip in space. The point clouds are then registered to a CT or MRI based pre-operative anatomic model. The EP software suite CARTO Electroanatomical Mapping System (Biosense Webster, Diamond Bar, CA USA) and EnSite NavX (St. Jude Medical, St. Paul, MN, USA) are two leading commercial systems on the market.

The main deficiencies of such systems currently include: (1) the mapping and model creation can be a tedious process. (2) The complexity of the system requires extensive experience from the user. (3) The pre-operative CT or MRI model may not accurately

capture the intra-operative condition. To our knowledge, there is no commercial system currently capable of producing real-time patient-specific anatomical visual guidance for ablation procedures.

The quality of the ablation procedures such as AF ablation largely depends on the ability of the surgeon to create continuous linear lesions without gaps using ablation catheters [29]. However, the anatomy of the pulmonary veins is quite variable and areas such as the junction of the left pulmonary veins and left atrial appendage (LAA) present a difficult and anatomically complex structure for the surgeon to maneuver [37]. The complex 3D curvature of the pulmonary vein area, rapid cardiac motion, and limitations in real-time imaging and catheter navigation account for a significant proportion of the failures in AF ablation [4].

1.5 Thesis Outline

Unlike fluoroscopy, real-time 3DUS allows visualization of the soft tissues at a high frame rate, and avoids the use of ionizing radiation. It can also be easily integrated into procedures as the small probe can be readily placed at the point of interest (Section 1.2).

Despite these evident advantages, however, a decade after its commercial introduction, 3DUS is rarely used clinically for procedure guidance. The adoption of 3DUS has been slow due to its limited field of view (FOV), low signal to noise ratio, missing anatomic information and tedious analysis process [77]. 3D imaging in ultrasound at a given frequency is limited by the acoustic time of flight. These constraints can be represented by a triangle which encloses the vertices of frame rate, resolution and field of view (frustum size). For instance, broadening the frustum size for the same frame rate blurs the point spread function for a constant number of transmitted scan lines and this reduces spatial resolution. We hypothesize that developing a real-time 3DUS mosaic to visualize the complex intra-cardiac anatomy, together with the abilities to determine precise ablation catheter position and tip-tissue contact would greatly improve the speed and accuracy of the procedures. In addition, the development of 3DUS specific image processing techniques would also assist

the real-time interpretation and diagnosis, lowering the threshold in the required training and skills for 3DUS analysis. With these goals in mind, this thesis is organized as the following:

Chapter 2 discusses the development of an ECG gated 3DUS mosaicing and enhanced visualization system as an effective way to provide a real-time wide-view of the target anatomy in both space and time. This aims to address 3DUS inherent limitation of limited field of view.

When building a large FOV 3D volume, complementary information from different US probe positions are often needed for the creation of complete anatomic structures. Chapter 3 discusses a novel technique that uses structure tensor for multiview 3DUS compounding.

Chapter 4 focuses on ICE catheters. We argue that automated pointing of ICE catheters will make ICE far more useful, effective, and broadly applicable. Based on this assumption, we developed a robotic ICE catheter control and visualization system. In this chapter, I will present a number of visualization techniques which include automated sweeping for 3D panorama creation and instrument tracking and tool tip identification.

The research effort above has broad applications beyond any specific beating heart procedure. For example, Real-time 3D mosaicing addresses the inherent narrow field of view presented in ultrasound images. The algorithms developed in Chapter 3 can be extended to general multiview image composition to improve the preservation of salient features from multiple contributing images while avoiding the introduction of noise at the same time. The methods developed in Chapter 4 can be used for a variety of non cardiac procedures such as those involving fluid-filled or solid organs.

Chapter 2

Real-time 3D Ultrasound Enhanced Visualization

2.1 Introduction

Real-time 3D ultrasound (3DUS) offers important advantages for guiding diverse medical procedures. Foremost is the ability to visualize complex 3D structures [74]. Studies have shown that real-time 3DUS is more efficient and accurate than 2DUS for basic surgical tasks and can enable more complex procedures [21]. Imaging rates up to 30 volumes per second also enable good visualization of instrument-tissue interactions, far faster than the volumetric imaging alternatives (MR and CT scans). Fluoroscopy provides fast frame rates, but only has a limited number of 2D views, requiring the clinician to mentally combine them to derive 3D structure. Unlike fluoroscopy, 3DUS also allows visualization of soft tissues, and avoids the use of ionizing radiation. 3DUS is easily integrated into procedures as the small probe can be readily placed at the point of interest. Finally, costs are also far lower, with top-of-the-line 3D ultrasound machines costing far less than comparable fluoroscopy, CT, or MR systems.

Despite these evident advantages, a decade after its commercial introduction, 3DUS is rarely used clinically for procedure guidance. There has been a broad spectrum of

research in 3DUS guidance, in diverse areas including liver surgery [50, 62], liver ablation [5], kidney imaging [10] and cardiac imaging [42, 57, 96]. Nonetheless, 2D ultrasound is still the prevailing choice in hospitals [74]. The reasons for this surprising lack of acceptance of 3DUS are diverse. One clear drawback of 3DUS is limited resolution. While voxel sizes are less than one millimeter, noise and distortion typically make it hard to discern features smaller than a few millimeters. In addition, 3DUS images are typically displayed as volume-rendered images. While this is effective for visualizing tissue surfaces surrounded by fluids as in obstetrics and cardiology, volume rendering can accentuate the distortion and noise inherent in 3DUS imaging [41], resulting in irregular surfaces and difficulty in distinguishing instrument artifacts [60, 65]. Volume rendering is also problematic for visualizing the internal features of solid organs like liver and kidney, where the entire organ produces textured reflections that fill the imaging volume. Another limitation is the small field of view. Because of the tradeoff in ultrasound imaging between volume size, resolution and frame rate, the volume size is inherently limited.

We hypothesize that enhanced displays can overcome key limitations in current 3DUS guidance, and bring the benefits of 3DUS to a broad range of procedures. In this chapter, we aim to address (1) the limited field of view; (2) the difficulty in distinguishing instrument from tissue in volume rendered images.

To overcome 3DUS inherent characteristics of small field of view, we mosaic multiple 3DUS volumes together to create a panoramic view. Volume registration is a key step for the success of the mosaicing. The approaches for volume registration can be largely divided into image-based, tracking-based or hybrid approaches. Image-based 3DUS registration methods can be voxel-based [77, 81], feature-based [58, 64, 73, 87, 92] or phased-based [35, 75, 97]. Voxel-based methods compute a metric between two volumes in an iterative fashion until a match meeting a certain threshold is found. These methods can handle relatively large displacement, but can also be sensitive to noise. Feature-based methods first find salient features and then compute a registration matrix using the matching features from the two volumes. They work well when there are sufficient overlapping regions between the

two volumes, therefore can only tolerate small translations and rotations and can be slow. Schneider *et al.* reported a real-time image based rigid registration method using 3D SIFT features [83]. Their work takes advantage of the fact that the scale of objects does not change in real-time 3DUS volumes with small displacement. Phased-based methods use a combination of local phase information and manually selected landmarks to find the registration matrix. They are computationally intensive. Tracking-based registration method involves attaching one or multiple position sensors to the ultrasound probe. Poon and Rohling used an optical position tracker and developed two different registration methods to correct the errors in the overlapping regions [72]. Yao *et al.* and Zhuang *et al.* used optical tracking for multiview 3DUS dataset compounding [97, 100]. Magnetic tracking is often used as well, especially when optical tracking is not feasible such as in the case when a clear line of sight from the optical sensor to the ultrasound probe cannot be guaranteed [11, 13, 98]. Recently, Housden *et al.* developed a hybrid method that leverages the availability of x-ray in EP suite and use x-ray images to track a TEE probe [40]. By registering the TEE probe positions to x-ray coordinate system, they were able to create an extended field of view of the 3DUS. Kutter *et al.* presented a multi-modal 3DUS registration and mosaicing approach by incorporating information from co-registered CT [48].

In the context of beating heart procedures, electrocardiograph (ECG) gating is commonly used in the 3D reconstruction of a time-series of 3DUS volumes (3D+time). Efforts have been reported in the reconstruction of gated 3DUS by mosaicing tracked 2DUS imaging [67, 71, 89], where 2DUS frames are acquired at the same phase of the cardiac cycle and composited into 3D volumes, which in turn are assembled into a time series for temporal display. With the arrival of real-time 3DUS machines, 3DUS volumes can be obtained directly and they offer improved spatial information compared to 2DUS, especially for geometrically complex targets with rapid motion. Brekke *et al.* developed an algorithm that assembles ECG gated 3D cardiac ultrasound sub volumes into a symmetric pyramidal without using image registration [15]. They reported decreased geometric distortions compared to the non-stitched 3D volume. One major limitation of such a system is that sub volumes must

be acquired at specific azimuth and elevation angles, thus users have limited control over probe placement and orientation. For real-time procedure guidance, ECG gated 3D freehand mosaicing is desired.

To address the lack of surface definition and the difficulty in distinguishing instrument from tissue in volume rendered images, we propose to display cut plane images or ‘slice views’ that contain the instrument tip. Because these cross-sectional views show the point of contact of the instrument with the tissue, as well as adjacent tissue regions, the clinician can readily determine the specifics of the tool-tissue interaction. Manually selecting these cut planes from within the 3DUS volume is a highly challenging task, especially when the instrument moves around within the volume. The ability to automatically visualize these slice views would greatly enhance the usability of 3DUS.

In this chapter, I describe the design of a system for enhanced real-time 3DUS visualization in the context of beating heart procedural guidance. The system performs freehand real-time 3DUS reconstruction and visualization with ECG gating. The freehand capability is important as it allows clinicians to optimize the ultrasound acquisition geometry. Volume registration is performed through tracking the 3DUS probe using an electromagnetic (EM) tracking system. This method is similar to the reconstruction of 2D images as previously described, with a key difference being that the data throughput is one or two orders of magnitudes higher in 3D imaging. I address this issue by leveraging the parallel computing power of graphics cards. In addition to the 4DUS of the heart, which can be used as the navigation map, the system is also capable of tracking the working catheter tip through the continuous display of the instrument tool tip and thick slices that go through instrument shaft. I report the results of a user study that demonstrates the potential of such enhanced displays in improving the efficacy of real-time 3DUS guided procedures.

2.2 Materials and Methods

2.2.1 System design

To investigate the potential benefits of mosaicing and slice views for procedure guidance, we implemented a prototype visualization system. We used Philips 3DUS scanner iE33 with an X7-2 2D/3D probe, imaging at 8.1cm and 35Hz with a volume size of $112 \times 48 \times 112$ voxels (Philips Healthcare, Andover, MA, USA). The resolution of the voxel is primarily determined by the imaging depth. An electromagnetic (EM) tracking system (3D Guidance trakStar System, Ascension Technology Corporation, Shelburne, VT, USA, 1.4mm position and 0.5° RMS accuracy) [88] tracked the trajectories of the 3DUS probe and the instrument tip. Image processing and rendering was done on a graphics processing unit (GPU) enabled computer (Dell Alienware Aurora, Intel Core i7 processor at 2.67GHz , 6GB RAM, NVIDIA GTX260 graphics card).

2.2.2 Mosaicing process

The software application is implemented in C/C++, using multi-threading for data acquisition, NVidia CUDA C (Driver v5.0.0) for parallel computing, and OpenGL (v3.3) for visualization. CUDA is a parallel computing platform and programming model invented by NVIDIA. It enables significant increase in computing performance by harnessing the power of GPU. Although originally designed to render computer graphics, GPUs are increasingly being used for programs in science, engineering, medicine and finance, among other domains. These non-graphics domain applications are referred to as general-purpose GPU programming. NVIDIA GPUs are built on the CUDA Architecture, in which CPUs are often referred to as the host, and GPUs as the device.

In our system, three threads run in parallel on the CPU: 3DUS volumes from the 3DUS machine, EM sensor readings from the EM tracker, and ECG waveform streaming. The streamed data then goes into their respective circular buffers which maintain a user specified duration of data. Here we choose to save the latest 2 seconds of data from each streaming




Figure 2.1: Data flow between host and device

source. All the data is time-stamped. ECG peak detection is performed online to identify the latest cardiac cycle in the ECG circular buffer. Using the matching timestamps, we then put the acquired 3DUS volumes and EM sensor readings into the right phases of the cardiac cycle, and compute the transformation of the incoming frame in terms of the reference frame. The transformation matrix and the gated 3DUS are then transferred to GPU for mosaicing and visualization. Figure 2.1 illustrates the data flow between CPU and GPU.

For clarity, in the next two sub sections, we divide our mosaicing approach into two steps. First, we present our algorithms for 3D volume registration and compounding of static object. Next we expand our algorithms to include ECG gated 3D volumes (3D + time).

2.2.3 3D volume mosaicing

Each 3DUS volume is represented as a 3D matrix V with intensity value $v_{(i,j,k)}$. Once a reference ultrasound frame is selected (by default, the first frame is used as the reference

frame, V^{ref}), it is uploaded to GPU and zero-padded to the size of the final mosaic volume (V^{mosaic}). Each new input volume V^{input} is subsequently padded, registered, interpolated, and compounded into V^{mosaic} as it comes in. The transformation matrix T_{input}^{mosaic} between the input volume and the reference volume is determined by the outputs of the EM tracker and the calibration matrix between the EM tracker and the 3DUS coordinate frame

$$T_{EM_{ref}}^{Transmitter} \cdot T_{US}^{EM} \cdot T_s \cdot P^{ref} = T_{EM_{input}}^{Transmitter} \cdot T_{US}^{EM} \cdot T_s \cdot P^{input}$$

Therefore

$$T_{input}^{mosaic} = T_s^{-1} \cdot T_{US}^{EM^{-1}} \cdot T_{EM_{ref}}^{Transmitter^{-1}} \cdot T_{EM_{input}}^{Transmitter} \cdot T_{US}^{EM} \cdot T_s$$

T_{US}^{EM} is the transformation matrix between the 3DUS probe and EM sensor attached to the probe. It is derived through a calibration procedure where we scan a triangle wire frame phantom with known geometric dimensions from different probe positions and orientations and perform intensity based registration. P^{input} represents a voxel location in the input volume to be registered to the reference volume, which is V^{mosaic} in this case. P^{ref} is the corresponding registered voxel location in V^{mosaic} . T_s is the scaling matrix that converts the ultrasound volume from voxel unit to a physical unit. $T_{EM_{ref}}^{Transmitter}$ and $T_{EM_{input}}^{Transmitter}$ are the respective EM sensor positions for V^{mosaic} and V^{input} .

The computation of volume registration, compounding and ray-casting are all performed on the GPU using a combination of CUDA linear arrays, texture memory and constant memory. CUDA texture memory is a specifically organized and cached part of the CUDA global memory that provides fast linear, bi-linear or tri-linear lookups. It can be in the form of 1D, 2D or 3D arrays. In our program, we use 3D textures to store the mosaiced volume and the current input volume, therefore taking advantage of their built-in tri-linear spatial locality for volume registration. 3D texture also facilitate its direct utility in volume rendering via OpenGL and CUDA pixel buffer object (PBO). Since data transfer between host (CPU) and device (GPU) is an expensive operation, and it is especially costly to update a texture from the host, we keep data transfer between host and device at minimum. Once the original input volume is transferred to the device, the subsequent operations are

performed without transferring the data back to the host. Our algorithms also interpolate and composite the two volumes in one function, which further speeds up the mosaicing process. Algorithm 1 and 2 describe the registration and compounding steps on GPU.

We also utilized CUDA constant memory to further reduce the memory bandwidth. Constant Memory is a piece of read only memory (does not change over the course of a CUDA kernel execution) that offers the benefit of short latency and high bandwidth when all threads from half of the warp (16 threads) access the same location. This provides a large memory access reduction when compared to using global memory. However, the total size of constant variables in an application has a current limit of 64 KB, making constant memory not suitable for large image data. In our program, we load the volume transformation matrix T_{input}^{mosaic} to a constant memory for an added performance gain.

Once the input volume is transformed to the reference volume, the interpolated volume needs to be compounded into the mosaiced volume. Intensity averaging of the overlapping voxels is the most commonly used method. In our current implementation, we use the weighted averaging of the existing voxel intensity and the contributing non-zero voxel intensity [81].

Algorithm 1 3D Volume Reconstruction and Visualization on GPU

Transfer the reference volume from CPU host to GPU device
Zero pad to a linear array V^{mosaic}
Initialize 3D texture $3DTex^{display}$ for volume rendering
Store the initial non-zero voxel locations in N , where $n_{(i,j,k)} = 1$ if $V_{(i,j,k)}^{mosaic} > 0$
for each input volume to be registered to V^{mosaic} **do**
 Copy the transformation matrix T_{input}^{mosaic} to a constant memory
 Update 3D texture $3DTex^{interp}$ for interpolation
 Call Registration and Compositing Kernel
 Update $3DTex^{display}$
 Call Ray-casting and update OpenGL pixel buffer object for rendering
end for

Algorithm 2 CUDA 3D Registration and Compounding Kernel

```
for each thread/voxel  $(m, n, p)$  in  $3D\text{Tex}^{interp}$  do
   $v = 3D\text{Tex}_{(m,n,p)}^{interp}$ 
  if  $v > 0$  then
    Compute its voxel location  $(i, j, k)$  in  $V^{mosaic}$  with 3D texture tri-linear interpolation
    Update the voxel value  $V_{(i,j,k)}^{mosaic}$  such that
    if  $V_{(i,j,k)}^{mosaic} == 0$  then
       $V_{(i,j,k)}^{mosaic} \leftarrow v$ 
    else
       $V_{(i,j,k)}^{mosaic} \leftarrow (n_{(i,j,k)} \times V_{(i,j,k)}^{mosaic} + v) / (n_{(i,j,k)} + 1)$ 
    end if
     $n_{(i,j,k)} \leftarrow n_{(i,j,k)} + 1$ 
  end if
end for
```

2.2.4 ECG gated mosaicing

Cardiac gating divides the motion of the heart into multiple phases across a complete cardiac cycle. Typical ECG tracing of a cardiac cycle consists of P wave, QRS complex, and T wave. The R wave peak in the QRS complex is typically used for cardiac gating as it can be readily identified in the ECG signal. The interval between R wave peaks, the RR interval, defines one cardiac cycle as shown in the top row of Figure 2.2. In our study, we perform R wave peak detection from the latest ECG circular buffer and divide the latest RR interval into a user-specified N number of equally spaced cardiac bins. The time stamps of these cardiac bins are then used to place the matching 3DUS and EM tracker data into the right bins. Figure 2.2 illustrates the data collection process. For simplicity, only four ECG bins and three spatial locations are shown here. Color corresponds to spatial location, and number indicates the bin. Transformation matrix T is a 4×4 homogeneous matrix containing rotation and translation. $T_{i,j}$ represents T acquired at spatial location i at bin j . First, the user chooses a location to image and sets it as the reference location. N bins of gated 3DUS and EM tracker data are selected from the circular buffers as describe above and saved as the reference frames. The four blue volumes in Figure 2.2 represent the reference volumes of the four bins. The user then moves the probe to another spatial




Figure 2.2: ECG gated 3DUS mosaicing

location and triggers the data acquisition at that location. The subsequent volumes acquired (purple and green volumes in Figure 2.2) are added to the mosaiced volumes by registering with the corresponding reference frames of the matching bin numbers. The mosaicing is performed immediately at each data acquisition and can be visualized at real time. The final time-series consist of N mosaiced volumes as illustrated in the last row of Figure 2.2. To implement the gated mosaicing on GPU, we modify Algorithm 1 and 2 such that V_{4D} is a cascade of N number of 3D volume V along the depth direction. For example, if one 3D volume V is of the dimension $width \times height \times depth$, then V_{4D} will have a dimension of $width \times height \times depth \times N$. The block and grid dimension of the GPU kernels are also adjusted accordingly. V_{4D} is uploaded to GPU in one data transfer. The transformation matrices are also cascaded in a similar fashion and uploaded to a constant memory. In this way, we are able to perform 4D volume registration and compounding without adding significant computation time. The only limiting factor here is the global memory bandwidth, one of the limitations in current CUDA architecture.

2.2.5 Instrument tip tracking

To track the tip of an instrument inside a 3DUS volume, an EM sensor is attached rigidly at the tip of the instrument, with the x-axis of the EM sensor aligned with the shaft of the instrument. The instrument tip location can then be registered to the mosaiced volume or to a live input volume through coordinate transforms. There are three coordinate frames involved in the system: US, EM sensor, and EM transmitter (Figure 2.3). Similar to what is described in Section 2.2, assuming p^{US} is the voxel location in the ultrasound volume that corresponds to the instrument tip, $p^{Transmitter}$ is the EM sensor output from the tip of instrument in terms of the EM transmitter coordinate frame, and T_s is the scaling matrix that converts the ultrasound volume from voxel unit to a physical unit, the overall transformation can then be established as the multiplication of a series of homogenous transformation matrices

$$p^{Transmitter} = T_{EM}^{Transmitter} \cdot T_{US}^{EM} \cdot T_s \cdot p^{US}$$

The tip location in the ultrasound volume can then be derived as

$$p^{US} = T_s^{-1} \cdot T_{US}^{EM^{-1}} \cdot T_{EM}^{Transmitter^{-1}} \cdot p^{Transmitter}$$




Figure 2.3: Coordinate frames in instrument tip tracking

2.2.6 Slice views

EM tracker provides six degree of freedom (6 DOF) sensor position and orientation outputs. The orientations of the three orthogonal axes ($X_{tip}, Y_{tip}, Z_{tip}$) at the instrument tip are used to generate the initial orthogonal slice views that contain the tool tip. Users then have the option to further adjust the orientation and thickness of each of the slice.

To render the slice views, each of EM sensor's orthogonal axes is used as a normal of a cut plane. Given a normal \hat{n} and a point $Tip = (Tip_x, Tip_y, Tip_z)$ on the plane, the plane is described as

$$\hat{n} \cdot (P - Tip) = 0$$

where P is any point on the plane. Given that the instrument has a thickness, to be able to visualize the instrument shaft and the tip-tissue contact, we define a slice with a thickness ϵ as

$$\hat{n} \cdot (P - Tip) < \epsilon$$

Any points inside the ultrasound volume that meet this equation will be displayed. The system tracks the instrument tip position and three orthogonal orientations at real-time.

2.3 Experiments and Validations

2.3.1 Water tank string phantom study

Method

We use phantom data in a water tank to develop and test the system. The phantom consists of vertical and diagonal strings suspended across a rigid frame with a dimension of $60mm \times 60mm$ (Figure 2.4(a)). A string phantom provides an easy means to validate the mosaicing result by confirming the mosaiced colinearity of the segments from each input volume. The phantom is suspended in the water tank with one side connected rigidly to a cam device. The cam device simulates the 1D periodic motion of the mitral annulus and gives a motion size of $\sim 30mm$ (Figure 2.4(b)). A signal from the cam device simulates the

output of an ECG monitor and is used for gating.

Results

Figure 2.5(a) shows the mosaiced result of the string phantom at one bin, and Figure 2.5(b)-(f) shows five 3DUS input volumes used for the creation of the mosaiced volume. It is evident that the mosaiced volume provides a wider field of view. The input volume has a dimension of $128 \times 48 \times 208$ voxels, and the dimension of the mosaiced volume is $256 \times 256 \times 256$ voxels. On average, the time used to perform every two volume registration and compounding is $25ms$. This mosaicing process is applied to each bin of the motion cycle. Figure 2.6 shows the mosaiced string phantom at different phases of the RR interval with $N = 6$. The red dot is the fiducial point on the string phantom. The yellow dotted line indicates the mid-line of the cross section. The varying distance between the fiducial point and the mid-line is the result of the motion of the phantom.




Figure 2.4: Experiment setup. (a) Water tank setup. The string phantom is connected to a cam device that provides a periodic motion simulating the 1D motion of the mitral annulus. (b) Typical periodic motion trajectory from the cam device.




Figure 2.5: *Mosaiced string phantom. (a) Mosaiced volume at one of the ECG bins. (b)-(f) Five input volumes. Note that the mosaiced volume provides a much wider field of view.*

Validation

We use the string phantom to validate the colinearity of the mosaiced lines. For a given mosaiced line, we first identify the reference line segment that the other line segments are registered to (Figure 2.7(a)) and use it as the centroid of each of the mosaiced lines in (Figure 2.7(b)). We then calculate the number of non-trivial voxels around the centroid at different radius sizes. We expect to see most of the contributing voxels to be around the centroid, and the number of contributing voxels falls off quickly when the radius increases. Figure 2.8 is a plot of the percentage of the contributing voxels vs. radius from the centroid calculated using the four mosaiced lines and their respective reference segments. It can be seen that most of the voxels are located within 1 to 2 voxel radius. The reference segments taper off at 3 voxel radius, while the mosaiced lines fall off at 4 voxel radius, thus the mosaiced lines follow the reference segments closely. The mosaiced lines have slight mis-alignments at a few junctions of the segments, which resulted in a small percentage of




Figure 2.6: Mosaiced string phantom at different phases of the RR interval with total of $N = 6$ bins. The red dot is the fiducial point on the string phantom. The yellow dotted line indicates the mid-line of the cross section. The varying distance between the fiducial point and the mid line is the result of the motion of the phantom.

contributing voxels at a bigger radius.




Figure 2.7: Reference line segments and the matching mosaiced lines. The number and color correspond to a matching pair.




Figure 2.8: Colinearity of mosaiced line segments

2.3.2 Mosaicing *in vivo* experiments





Method

We conducted *in vivo* experiments on a porcine heart using Philips 3DUS scanner iE33 with an X7-2 2D/3D probe, imaging at 11cm with a volume size of $112 \times 48 \times 128$ voxels (Philips Healthcare, Andover, MA, USA). Similar to the water tank study, the probe was tracked by an EM tracker. The imaging subject was a 40kg anesthetized Yorkshire pig with the chest opened to permit access to the epicardial surface. A coupling gel standoff approximately 1cm thick was attached to the front of the probe to enlarge the field of view. The probe was pressed by hand against the the heart in the vicinity of the right atrium. Images were acquired for 3 seconds at four locations.

Results

Six ECG bins ($N = 6$) are generated with each mosaiced volume of $128 \times 96 \times 128$ voxels. Figure 2.9 compares the mosaiced volume at one of the bins to the input volumes. Figure 2.9(a)-(b) show two contributing volumes. Figure 2.9(c)-(d) are the mosaiced volume from two different views with key anatomical structures identified. We can see an extended view of the right atrium (RA) and more complete structure of the mitral valve (MV). Greater spatial information is provided by the extended field of view.

Figure 2.10 and Figure 2.11 display the mosaiced volumes in Figure 2.9(c)-(d) at six different bins. MV opening and closing in conjunction with RA and Left Atrium (LA) contraction and expansion are visible across bins and from different view points. The gated mosaiced volumes provide comprehensive spatial views and more complete motions at the same time, information unavailable in one single volume.



RA - Right Atrium LA - Left Atrium MV - Mitral Valve

Figure 2.9: Mosaiced volume from two different views - *in vivo* study. (a)-(b) Two contributing volumes. (c)-(d) Mosaiced volume from two different views with key anatomical structures identified. Greater spatial information is provided by the extended field of view.




Figure 2.10: ECG gated 3DUS mosaicing in vivo study - view from Figure 2.9(c). Mosaiced volumes at different phases of a cardiac cycle.




Figure 2.11: ECG gated 3DUS mosaicing *in vivo* study - view from Figure 2.9(d). Mosaiced volumes at different phases of a cardiac cycle.

2.3.3 Instrument tip tracking

Method

The tip and slice views can be displayed in reference to a single input volume or within a mosaiced volume. In water tank studies, we imaged phantom objects made of gelatin with fiber supplement, which closely mimicks animal tissue features [59]. Figure 2.12(a) is the phantom shaped into a rectangular form with openings at the corners. The openings in the phantom can be used to simulate tool insertions and tool tip tracking. In some minimally invasive heart procedures, the instrument is introduced into the left atrium through a small hole made in the left atrial appendage (LAA). The instrument we used is a flexible catheter of 3mm in diameter with an EM sensor built-in at its tip, as shown in Figure 2.12(b). It closely resembles the dimension of typical catheters used in minimally invasive procedures. Figure 2.12(c) illustrate the orientations of the cut planes, which are rendered as thick slices in Figure 2.12(d)-(f). Once the cut planes are identified, the user can further adjust the thickness and orientation of each of them.

Results

Figure 2.13 are the instrument tracking and mosaicing results. The first row of the figure shows the instrument in single volumes of the atrium phantom. Only part of the instrument is captured in these volumes due to 3DUS limited field of view. Note the difficulty in identifying the tip location. Some volumes may not contain the tip. The second row of the figure shows the mosaiced volume containing the entire instrument tip from three different point of views, from which we can see the extent of the tool shaft. In the third row of the figure, orthogonal slice views generated from the mosaiced volume show the exact instrument tip and surrounding tissue.

In addition, a mosaiced volume of a ROI can be created first, after which the instrument tip is registered to the volume and tracked without having to have to be imaged. This is illustrated by the fourth row of Figure 2.13, where the instrument tip is being tracked as a




Figure 2.12: *Instrument tip and orthogonal cut planes through the instrument. (a) Atrium phantom with openings at the corners. (b) Instrument tip at one of the phantom openings. (c) Cut planes that go through the instrument tip. (d)-(f) Orthogonal thick slices that contain the instrument tip. (e) and (f) also go through the instrument shaft.*

green dot in mosaiced volume of simulated pulmonary vein ostia (PVO).

2.4 User Study

We conducted a user study, following a protocol approved by our institutional review board, comparing performance in clinical tasks with and without the enhanced displays. To provide a specific clinical focus, the tasks are taken from intra-cardiac procedures as this specialty shows strong potential for benefiting from 3DUS guidance, however, as discussed below, the proposed display system can apply to a range of other procedures.

2.4.1 Study design

Five interventional cardiologists with experience in minimally invasive cardiac procedures were recruited for the user study. Their experience in conducting procedures ranged from 3.5 years to 11 years after the beginning of postgraduate training. Each subject performed




Figure 2.13: *Instrument tracking and mosaicing results. First row: Instrument in single volumes of the atrium phantom. Only part of the instrument is captured in these volumes. Note the difficulty in identifying the tip location. Second row: Mosaiced volume contains the entire instrument tip from three different point of views. Third row: Orthogonal slice views generated from the mosaiced volume showing the exact instrument tip and surrounding tissue. Fourth row: Instrument tip in mosaiced volume of simulated PVO.*

two instrument navigation tasks in a water tank. Subjects could not directly see the task by eye. There were three 3DUS display conditions: the volume rendered display on the 3DUS machine, Slice views in single volume, and slice views in mosaiced volume. The order of the tasks and displays for each subject was randomized. In the first two display conditions, subjects moved the probe with one hand to follow the moving instrument and keep it in the 3DUS field of view. With the mosaiced display, there were two options. If the imaging object was static, the user could simply use the EM tracker generated slice views in the mosaic. With moving imaging object, such as a beating heart, the real-time 3DUS volume would be superimposed on the mosaic. In this study, since the image objects were static, subjects did not hold the US probe during the third testing conditions, instead the instrument tip location is tracked by EM tracking and registered and displayed on the static mosaic volume.

Catheter based ablation procedures such as atrial fibrillation ablations aim to create linear destructive lesions in the tissue around the pulmonary veins to prevent propagation of abnormal electrical signals to the rest of the atrial tissue [15]. The first task was to trace the perimeter of the rectangle created by four rubber bands (Figure 2.14(a)) with the tip of a catheter in a water tank. The dimensions of the rectangle were roughly $40\text{mm} \times 30\text{mm}$. The goal of this task was to evaluate the effectiveness of the slice views in improving user's ability to maneuver the catheter tip along a predefined path. Rubber bands were chosen because their 3DUS images (Figure 2.14(b)) show noticeable noise and blurring along the edges, representative of *in vivo* conditions where noise and artifacts are common. The subjects were instructed to hold the catheter roughly 5cm above its tip, and move the tip along the loop ABCD as outlined in the blue dotted lines in Figure 2.14(a).

The second task was to trace the mitral annulus of a porcine heart in a water tank using a rigid instrument. An instrument with straight and rigid tip was inserted through a puncture created on the left atrial appendage to access the mitral valve annulus (Figure 2.15). In minimally invasive beating heart mitral valve repairs such as mitral annuloplasty, an anchor driver can be inserted through the left atrial appendage to reach the mitral annulus under

3DUS guidance [16]. Instrument tip and tissue often blend together along the atrial wall, making it difficult to discern the instrument tip. This task aims to assess the value of slice views and mosaicing for such procedures. Task completion time was recorded for each subject trial. The instrument tip trajectories were also recorded with EM tracker.




Figure 2.14: Navigation task 1. (a) Photograph of the rectangle (40mm \times 30mm) created by four rubber bands. (b) Mosaiced 3DUS image of the four rubber bands containing the rectangle. (c) Example input volumes as a comparison.




Figure 2.15: Navigation task 2. (a) Top down view of the mitral valve and mitral annulus. (b) Instrument inserted through the left atrial appendage (LAA). (c) 3DUS volume rendered showing the left atrium (LA) and part of the instrument.

2.4.2 Results

Typical trajectories for both tasks are shown in Figure 2.16. The completion times from all subject trials were analyzed using Wilcoxon rank sum test (Figure 2.17). The criteria for statistical significance was $p \leq 0.05$. The average deviation of the trajectory $\mu_{distance}$ for task 1 was also calculated by first measuring each line segment AB, BC, CD and DA in Figure 2.14(a), and then calculating the mean of the distance from each point on the trajectory to the corresponding line segment. A smaller $\mu_{distance}$ indicates a closer match to the original path.

In task 1, compared to the completion time under 3DUS volume without slice views, the completion time decreased by 46% with slice views ($p \leq 0.0159$), and $\mu_{distance}$ decreased by 20%. The completion time decreased by 69% using mosaiced volume with slice views ($p \leq 0.0079$), and $\mu_{distance}$ decreased by 25%.

In task 2, compared to the completion time under 3DUS volume without slice views, the completion time decreased by 36% with slice views in single volume ($p \leq 0.0317$). The completion time decreased by 46% using slice views in mosaiced volumes ($p \leq 0.0159$). In both tasks, the variability of the results (standard error) also decreased with enhanced displays. This suggests that with slice views and mosaic displays, the surgical task is less dependent on each user's individual echocardiography skills and experience. Subjects reported that the three orthogonal views were intuitive. All subjects stated that mosaiced volumes combined with slice views have great potential for 3DUS guided interventions.

2.5 Discussion

The work in this chapter aims to address the following three key issues in current 3DUS guidance: (1) Volume rendering alone is not adequate to visualize tool-tissue interaction; (2) Slice view avoids problem with volume rendering, but is difficult to align manually; (3) 3DUS has a small field of view and is difficult to navigate. Our prototype system and user study demonstrate that improved visualization techniques could mitigate key limitations in




Figure 2.16: Typical instrument tip trajectories. Left column: Trajectory of tracing the rectangle of rubber bands. Right column: Trajectory of tracing the mitral annulus of a porcine heart. Trajectories are smoothest with slice views in mosaiced volumes, as shown in the last row.




Figure 2.17: Task completion time comparison. (a) Task 1 – Tracing a rectangle of rubber bands. (b) Task 2 – Tracing the mitral annulus of a porcine heart.

3DUS. A mosaiced volume overcomes 3DUS limited field of view and can be used for broad navigation. Computer assisted instrument tip tracking in slice views facilitates real-time instrument navigation and visual feedback.

The tasks used here are simpler than clinical catheter based procedures and 3DUS imaging in a water tank has better quality than *in vivo* situations. Thus, we expect even better user improvement with slice views and mosaicing in *in vivo* and clinical procedures.

A number of image guidance systems share features with the approach presented here. For example, 3D Slicer [85], was originally applied to neurosurgery, where tracked instrument positions are superimposed on static preoperative brain images. These systems necessarily used slice views. Similarly, ultrasound visualization tools such as Stradx [75] work with prerecorded data. The system proposed here, however, uses real-time volumetric data, which presents significant challenges for registration, mosaicing, and tracking.

2.5.1 Temporal resolution and imaging time

There is a tradeoff between imaging time to build or update the mosaic and time resolution (i.e. the number of time bins N). Temporal resolution is particularly important for accurately capturing the motion of fast-moving structures such as cardiac valves. For example, the peak velocity of the mitral valve annulus has been estimated as $v = 210\text{mm/s}$ [46], so assuming a typical RR interval of $T_{RR} = 1\text{s}$ and with number of bins $N = 10$, the spatial blur within one time bin can be as high as $dx = v \cdot (T_{RR}/N) = 21\text{mm}$. This shows that high temporal resolution reduces the effect of spatial blurring. On the other hand, higher temporal resolution will result in longer data acquisition and processing time and higher demand on computer memory size. In addition, the current 3DUS streaming is at approximately 30Hz , which limits the data acquisition to 30 volumes per second.

2.5.2 Real-time overlay

One of the important features that this system can offer is a ‘big picture’ view of a large section of the heart, combined with a ‘fovea’ real-time view of the smaller region where

the procedure is performed. The mosaiced time series (4D mosaic) can be viewed as a map in the background, which can be extremely useful to the user for navigating to a specific location and for planning subsequent surgical steps. The specific location where the working catheter is can then be automatically overlaid on the 4D mosaic in real time. Color coding of the real-time and mosaiced data can inform the clinician of which regions represent the current versus historical information. This allows for a simple and reliable control of the image acquisition and display process.

2.5.3 System extensions

Novotny *et al.* developed a GPU based real-time instrument detection algorithm that uses a generalized Radon transform [66]. Any such image-based tracking approach can be advantageously integrated with the current EM tracker based slice views. The EM sensor provides the initial estimate of the tip location, which reduces the image search space hence further speed up the algorithm.

Accuracy of the volume mosaicing could also be further improved using image based methods. Schneider *et al.* developed a real-time feature-based 3DUS registration framework on GPU [83]. Grau *et al.* reported a structure orientation and phase based algorithm to register apical and parasternal 3DUS datasets of the heart [35]. EM based volume registration can be used as an initial estimate for either of these two algorithms.

2.5.4 Application to beating heart procedures

The slice views presented here can be integrated with an ECG gated mosaicing system for beating heart intra-cardiac procedures.

Currently in a typical catheter based intra-cardiac ablation procedure, a 3D electrophysiological model resembling the shape of the atria is generated by recording the ablation catheter tip locations in space. The point clouds are then registered to a CT or MRI based pre-operative anatomic model (e.g. CARTO, Biosense Webster, Diamond Bar, CA). Mapping and model creation can be a tedious process and fluoroscopy is routinely used. An ECG

gated 3DUS mosaicing system combined with enhanced user display such as slice views could shorten the procedure time, reduce human exposure to fluoroscopy, and provide improved visualization of tool-tissue interaction.

2.6 Conclusion

In this chapter, we presented a set of enhanced display techniques for real-time 3DUS visualization. The system integrates EM tracking systems and GPU implementation for real-time 3D mosaicing and instrument tip cut plane tracking. This mitigates the 3DUS distortions, an adverse effect inherent in conventional volume rendering. The experimental studies and user study, combined with feedback from participants demonstrate the potential of such enhanced visualization in instrument navigation and procedure execution with 3DUS guidance.

Chapter 3

Structure Tensor for 3DUS Multiview Compounding

3.1 Introduction

Real-time 3D ultrasounds (3DUS) holds great promise in enabling a broad range of minimally invasive cardiac procedures. It is non-invasive, inexpensive, able to image through the blood, and can provide real-time diagnostic information of the anatomical region of interest. 3DUS also mitigates the difficulties in spatial perception associated with traditional 2DUS [21]. But as discussed in previous chapters, there exist key limiting factors in the clinical adoption of 3DUS. In addition to the limited field of views (FOV) and low signal to noise ratio, signal attenuation and dropout can also affect the quality of the data. Ultrasound images are constructed by sound reflection and scattering at tissue interfaces. The strength of the reflections is highly depending on the propagation of the ultrasound wave in regards to the orientation of the tissue boundary.

Furthermore, signal dropout can occur due to occlusion of underlying tissue by structures that reflect or absorb the acoustic signal. For instance, when the muscle fibers lies perpendicular to the US beam, they reflect most of the wave back, little information from below the muscle can be gathered, resulting in much attenuated signals, shadows, or missing




Figure 3.1: Example 3DUS images showing missing anatomical structures. (a) A TEE image of the aortic valve area. (b) A TTE image of the mitral valve region. Arrows indicate missing anatomical structures.

anatomical structure. This phenomenon is illustrated by Figure 3.1. Two examples of 3D echocardiography images of part of the heart are shown here, where important anatomical structures are missing as a result of signal dropout or limited FOV. Figure 3.1(a) was acquired using a transesophageal (TEE) probe from inside the esophagus and the view is limited by the esophagus. Figure 3.1(b) was acquired using a transthoracic (TTE) probe from outside the chest through one of the acoustic windows in the thorax. The arrows indicate missing anatomical structures.

Two local structures of importance are lines and surfaces. The strength of the returned US signals vary depending on the incident angle in regards to the orientations of the structures. For example, in Figure 3.2, a wooden toothpick is imaged from two different angles. From angle (a), the US probe is almost parallel to the orientation of toothpick, thus little signal is reflected. Note the artifact and distortion at the tip of the target. From angle (b), a much higher percentage of the US incident wave is reflected by the toothpick, note the improvement in the quality of the image.

Similarly, when imaging a surface, if the incident angle is oblique to the surface normal, as illustrated by Figure 3.3(a), the resulting image contains more artifacts and noise than if




Figure 3.2: A wooden toothpick imaged from two different angles indicated by arrows (a) and (b). (a) The toothpick is imaged straight down. Note the artifacts and distortion at the tip. (b) The toothpick is imaged from a large angle. Note the improvement in the quality of the image.

the surface is imaged from straight down, where the surface boundary reflects stronger to the US incident wave (Figure 3.3(b)).

One of the ways to mitigate this adverse affect of signal dropout is to image the region of interest (ROI) from different view points and reconstruct a more complete image. This is called multiview fusion. It also serves to increase the FOV. In the TTE case, images from different acoustic windows (e.g. apical and parasternal views) can be fused together, whereas in the TEE case, images can be acquired by carefully rotating the probe inside the esophagus and then fused together for complementary information.

3DUS multiview fusion is a subset of the 3DUS mosaicing, where it aims to increase the FOV through interpolating and compounding multiple volumes into a single large composite volume. Compounding happens after one or all incoming volumes are registered to the reference volume, and it is a key step in determining the final intensity values of the composite volume.




Figure 3.3: A tissue mimicking phantom surface imaged from two different angles indicated by arrows (a) and (b). (a) The surface is imaged from an oblique angle. Note the noise on the surface. (b) The surface is imaged straight down. Note the improvement in the quality of the image.

In contrast to spatial mosaicing for increasing FOV, where the probe moves at small displacement, during the acquisition of multiview images, probe orientations can vary significantly from view to view resulting in large intensity variations of the same target. This presents both challenges and opportunities in multiview compounding.

A simple compounding approach is to choose the maximum intensity pixel/voxel from the contributing images. This method is easy to implement, but has a tendency to increase noise levels, highlight speckle effects and often should not be used alone. Intensity averaging for spatial compounding [74, 81] is another commonly used technique. Averaging the intensity of the voxels in overlapping regions could improve the signal and noise ratio (SNR) in some local regions, and decrease the speckle effect. This method works well in areas of homogeneous tissue, but can create blurs or decrease the final intensity at a region where features with a prominent orientation, such as the endocardial surface or septum wall surface have significantly different contributing intensity values from different viewing angles.

As an alternative to intensity-based methods, an adaptive multiscale compounding algorithm for using the phase information was recently reported by Grau *et al.* [36]. This method is extremely computationally intensive and was only applied to a small number of apical and parasternal images. Yao and Penney compared the performance of different compounding methods with a larger number of images [97]. Four compounding methods were investigated, mean, maximum, and two methods derived from phase-based compounding. They concluded that the phase-based method performed better and the signal to noise ratio and contrast could be improved by using increased number of images (10+).

Rajpoot *et al.* proposed a wavelet based fusion algorithm to exploit the low- and high-frequency separation capability of wavelet analysis [77]. Their approach is based on the assumption that in a 3DUS image, high frequency components represent speckle regions while low frequency components contains features such as the myocardium. To suppress the noise while maintaining the speckle appearance, information in the high frequency sub bands are averaged. To preserve features, a voxel-wise maximum approach is chosen in the low frequency sub bands. Although this is a conceptually simple approach, the registration and fusion process is reported as being computationally intensive and time consuming, which is not suitable for any real-time purpose.

We propose the use of structure tensor for improved compounding. The structure tensor, also known as second moment matrix, is an important tool in computer vision for the purpose of orientation estimation and local structure analysis. It is based on the integration of data from a local neighborhood. Typically this neighborhood is defined by a Gaussian window function and the structure tensor is computed by the weighted sum within this window [8, 16, 32].

The initial matrix of a 2D structure tensor is expressed as the outer product of the image gradient

$$S_0 = \begin{pmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{pmatrix}$$

where I_x and I_y are the partial derivatives of image intensity I with respect to x and y . Its

discrete form can be expressed as

$$S_0[p] = \begin{pmatrix} (I_x[p])^2 & I_x[p]I_y[p] \\ I_x[p]I_y[p] & (I_y[p])^2 \end{pmatrix}$$

with $p = (p_x, p_y)$ representing a pair of integer indices (p_x, p_y) , and $I[p]$ a series of discrete samples of I .

The structure tensor for a certain neighborhood of scale τ is then computed by convolution of the components of S_0 with a Gaussian kernel G_τ

$$S_\tau = G_\tau * S_0 \quad (3.1)$$

where

$$G_\tau(x, y) = \frac{1}{2\pi\tau^2} e^{-(x^2+y^2)/2\tau^2}$$

Its discrete form can be expressed as

$$S_\tau[p] = \sum_r G[r] S_0[p - r] \quad (3.2)$$

where r ranges over a finite set of index pairs of the window (the neighborhood) and $G[r]$ is the weight of the window depending on r , such that the sum of all weights is 1. The structure tensor presents integrated information of a local neighborhood without cancellation effects, whereas summing gradient alone in a window, gradient vectors in the opposing directions would cancel each other out. This makes structure tensor robust against noise and a more reliable estimator of orientations in noise-prone data, such as ultrasound [17].

Smoothing also distributes the information about the orientation into the areas between edges. Thus we can estimate the dominant orientation at locations where the gradient is close to zero [16]. The dominant orientation can be obtained from the structure tensor as the eigenvector of the largest eigenvalue.

By studying the relationship of the eigenvalues, a local region coherence map can be inferred, by which we can distinguish areas where structures are oriented uniformly (i.e. edges) from areas where structures have different orientations (i.e corners).

Because a structure tensor is the outer multiplication of the image gradient, it is symmetric and semi positive definite, thus, there exists non-negative eigenvalues such that $\lambda_1 \geq \lambda_2 \geq 0$. Their respective eigenvectors are \hat{e}_1 and \hat{e}_2 . We can distinguish different local structures in 2D by analyzing eigenvalues:

- (1) $\lambda_1 \gg \lambda_2 \approx 0$: there exists a preferred local orientation represented by λ_1 and \hat{e}_1 .
- (2) $\lambda_1 \approx \lambda_2 \gg 0$: the gradient in the window has no predominant direction, indication of rotational symmetry within the window.
- (3) $\lambda_1 \approx \lambda_2 \approx 0$: homogeneous regions.

At scale τ , the anisotropic measure, also known as the coherence measure, can be expressed by the eigenvalues of the local structure tensor [44, 56]

$$C_\tau = \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} \right)^2 \quad (3.3)$$

when $\lambda_2 > 0$.

When $C_\tau \approx 1$, the gradient is aligned and there is a predominant direction, and when $C_\tau \approx 0$, there is no preferred orientation.

As discussed in the introduction section, the incident angle of ultrasound wave and the muscle fiber orientations has non-negligible affects on the outcome of an ultrasound image. We aim to design a measure to describe the orientation of a local structure in regards to the ultrasound beam direction \hat{d}_{beam} . From eigenvalues, we can estimate local coherence. From eigenvectors, we can gather the orientations of the local structures. Figure 3.4 shows that in the 2D case, θ , the angle between the gradient vector (the larger of the two eigenvectors, \hat{e}_1) and the US beam direction \hat{d}_{beam} tells us how a local structure is aligned with the beam propagation direction. When the gradient vector is well aligned with the incident direction, $\cos(\theta) \approx 1$. When the gradient vector is perpendicular to the US propagation direction, $\cos(\theta) \approx 0$. The concept can be expanded to 3D cases which we will discuss in the following sections.




Figure 3.4: Angle between the US beam direction and the orientation of a 2D structure. $\cos(\theta)$ can be used as an indicator of the alignment of the gradient vector and the US incident beam.

This chapter is organized as follows. In the next sections, I will start with brief background information on 3D structure tensor, including the calculation of the coherence measures. I will then discuss how to construct a weighting function using the structure tensor and US beam incident angle, for the purpose of volume compounding. Results from both water tank and *in vivo* studies will be presented. I will conclude this chapter with a summary and discussions on multiscale compounding and GPU implementation for real-time performance.

3.2 Methods

3.2.1 3DUS structure tensor

The initial matrix of a 3D structure tensor is expressed as the outer product of the image gradient

$$S_0 = \begin{pmatrix} I_x^2 & I_x I_y & I_x I_z \\ I_y I_x & I_y^2 & I_y I_z \\ I_z I_x & I_z I_y & I_z^2 \end{pmatrix}$$

where I_x , I_y and I_z are the partial derivatives of image intensity I with respect to x , y and z .

Its discrete form can be expressed as

$$S_0[p] = \begin{pmatrix} (I_x[p])^2 & I_x[p]I_y[p] & I_x[p]I_z[p] \\ I_y[p]I_x[p] & (I_y[p])^2 & I_y[p]I_z[p] \\ I_z[p]I_x[p] & I_z[p]I_y[p] & (I_z[p])^2 \end{pmatrix}$$

where $p = (p_x, p_y, p_z)$ now represents integer indices of a 3D grid.

Similar to Equation 3.1 and Equation 3.2, given a 3D Gaussian kernel G_τ , structure tensor S_τ for a certain neighborhood of scale τ is the convolution of S_0 with G_τ .

Typically in computing the partial derivatives, a pre-smoothing is applied. This is especially important in dealing with US data due to its noisy nature. For computational efficiency, we combine the the smoothing and partial derivatives in one step by using the Difference of Gaussian (DoG), which is a separable kernel and can be expressed as the convolutions of multiple 1D kernels.

In this paper, we use σ to denote the scale of the pre-smoothing Gaussian kernel G_σ , and τ the scale of the structure tensor windowing function G_τ .

Similar to the 2D case, a structure tensor for a 3D function is also a symmetric and semi positive definite matrix, therefore, there exists non-negative eigenvalues $(\lambda_1, \lambda_2, \lambda_3)$ and an orthogonal system of eigenvectors $(\hat{e}_1, \hat{e}_2, \hat{e}_3)$. The eigenvalues and eigenvectors summarize the distribution of gradient directions within the neighborhood of p defined by the window τ . For ease of explanation, we order them such that $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$. Different local structures can be gleaned from the eigenvalues:

- (1) $\lambda_1 \gg \lambda_2 \approx \lambda_3$ infers a surface-like (surfel) neighborhood. If we use an ellipsoid to represent the 3D structure tensor, the ellipsoid in this case is stretched along one axis (\hat{e}_1) only, like a tube.
- (2) $\lambda_1 \approx \lambda_2 \gg \lambda_3$ represents a line-like (curvel) neighborhood. The ellipsoid of the 3D structure is a flat disk.
- (3) $\lambda_1 \approx \lambda_2 \approx \lambda_3$ indicates an isotropic neighborhood. The ellipsoid of the 3D structure is a sphere.

3.2.2 Design of an US incident angle weighted coherence measure

The eigenvalues of S_τ provide useful information on the coherence of local structures, i.e. the amount of anisotropy. Several different coherence measures are considered here for the purpose of improving the compounding of the registered 3DUS volumes. The goal here is to:

- (a) Identify isotropic regions where all three eigenvalues are roughly the same sizes. In these regions, the non-zero intensity values are averaged.
- (b) Identify regions of high anisotropy, in which we further distinguish line-like and surface-like structures and their orientation in regards to US beam incident angle.
- (c) Use the measure above to construct a weighted compounding.

Weickert *et al.* proposed a measure of local coherence [93]

$$\kappa := \sum_{i=1}^3 \sum_{j=i+1}^3 (\lambda_i - \lambda_j)^2 \quad (3.4)$$

κ becomes large for strongly differing eigenvalues, and it goes to zero for isotropic structures. To keep the coherence value to $[0, 1]$, κ was further transformed to

$$C_{anisotropy} = \begin{cases} \alpha, & \text{if } \kappa = 0, \\ \alpha + (1 - \alpha) \exp(-\frac{K}{\kappa}), & \text{otherwise.} \end{cases} \quad (3.5)$$

where K is a scaling factor of positive value. When $\kappa \gg K$, $C_{anisotropy} \approx 1$, and when $\kappa \ll K$, $C_{anisotropy} \approx \alpha$.

A different coherence measure proposed by Brox *et al.* evaluates when the two larger eigenvalues of the structure tensor are large enough compared to the smallest one, i.e. that the ellipsoid associated to the tensor is flat [18]

$$C_1 = \left(\frac{\lambda_1 - \lambda_3}{\lambda_1 + \lambda_3 + \epsilon} \right)^2 - \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \epsilon} \right)^2 \quad (3.6)$$

ϵ is a small positive constant as a regulation term. When $\lambda_2 \approx \lambda_3$, this measure yields a value close to 0, and when $\lambda_1 \approx \lambda_2 \gg \lambda_3$, it becomes close to 1. This measure detects

line-like ('curvel') structure in a neighborhood.

For completeness, we propose the following to measure when the largest eigenvalue is much larger than the other two

$$C_2 = \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \epsilon} \right)^2 \quad (3.7)$$

When $\lambda_1 \approx \lambda_2$, $C_2 \approx 0$. When $\lambda_1 \gg \lambda_2$, $C_2 \approx 1$. This identifies surface-like ('surfel') neighborhood.

In practice, both C_1 and C_2 are sensitive to noise and can be large even when the eigenvalues are small. We apply $C_{anisotropy}$ to each of them such that they only measure line-like and surface-like structures at areas of high anisotropy

$$\begin{aligned} C_{curvel} &= C_{anisotropy} \cdot C_1 \\ &= C_{anisotropy} \cdot \left(\left(\frac{\lambda_1 - \lambda_3}{\lambda_1 + \lambda_3 + \epsilon} \right)^2 - \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \epsilon} \right)^2 \right) \end{aligned} \quad (3.8)$$

$$\begin{aligned} C_{surfel} &= C_{anisotropy} \cdot C_2 \\ &= C_{anisotropy} \cdot \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \epsilon} \right)^2 \end{aligned} \quad (3.9)$$

All three measure $C_{anisotropy}$, C_{curvel} , and C_{surfel} from above aim to infer the coherence of local structures, but with slightly different goals. $C_{anisotropy}$ measures the anisotropy and does not try to infer the shapes of the local structures. C_{curvel} measures specifically the 'flatness' of the tensor for the purpose of identifying line-like structures. C_{surfel} measures the strength of the largest eigenvalue against the smaller two for surface-like structures.

In regards to ultrasound images, we would also like to consider the US beam incident angle for the reasons discussed in Section 3.1. We choose to design a measure W for this purpose. We break W into two terms:

- (1) In a surfel neighborhood, the largest eigenvector \hat{e}_1 represents the surface normal. θ is the angle of the incident wave unit direction vector \hat{d}_{beam} and \hat{e}_1 (Figure 3.5(a)). $\cos(\theta)$ equals to the dot product of \hat{e}_1 and \hat{d}_{beam} , and thus is chosen to evaluate the orientation of a surfel in regards to the incident angle. Taking the dot product of $\cos(\theta)$ and

Equation 3.9, we get

$$W_{surfel} = C_{surfel} \cdot (|\hat{e}_1 \cdot \hat{d}_{beam}|) \quad (3.10)$$

When \hat{e}_1 is well aligned with \hat{d}_{beam} in a surfel neighborhood, (i.e. the surface is perpendicular to \hat{d}_{beam} and $\cos(\theta) \approx 1$), $W_{surfel} \approx 1$. When it is not a surfel, or when \hat{e}_1 is perpendicular to \hat{d}_{beam} , $W_{surfel} \approx 0$.

- (2) In a curvel neighborhood, the third eigenvector \hat{e}_3 indicates the orientation of the line-like structure. $\theta + \phi = \pi/2$ as shown in Figure 3.5(b). We construct the following measure

$$W_{curvel} = C_{curvel} \cdot (1 - |\hat{e}_3 \cdot \hat{d}_{beam}|) \quad (3.11)$$

When $|\hat{e}_3 \cdot \hat{d}_{beam}|$ is small in a curvel neighborhood, i.e. the line-like structure is perpendicular to \hat{d}_{beam} , $W_{curvel} \approx 1$. When it is not a curvel, or when $|\hat{e}_3 \cdot \hat{d}_{beam}|$ is large, $W_{curvel} \approx 0$.

Combining the two terms above, we have a measure that includes both the orientations of the local structures and its relationship with the US incident wave

$$\begin{aligned} W &= W_{surfel} + W_{curvel} \\ &= C_{surfel} \cdot (|\hat{e}_1 \cdot \hat{d}_{beam}|) + C_{curvel} \cdot (1 - |\hat{e}_3 \cdot \hat{d}_{beam}|) \end{aligned} \quad (3.12)$$

Here we use the US axial direction as \hat{d}_{beam} , as shown in Figure 3.5.

3.2.3 Algorithm for the compounding

Measure W combines the local structure information with the incident wave direction and is a measure specifically designed for signals such as ultrasound. The steps to use W for the compounding of multiple registered volumes are summarized in Figure 3.6 and in Algorithm 3.

Let $W_i(p)$ represents the measure W at voxel location $p = (x, y, z)$ in a registered volume V_i , $i = 1, 2, \dots, N$. I_i is the intensity value at voxel p in volume V_i . N is the total number of registered volumes to be compounded. Figure 3.6 outlines the three key conditions for the




Figure 3.5: Measure of surfel and curvel with US beam incident angle. (a) Surfel. (b) Curvel.

weighted compounding.

- (1) If $W_i(p)$ values are small (below a pre-defined threshold) in all the volumes, the non-zero intensity values of I_i are averaged.
- (2) If there is only one large $W_i(p)$, the intensity value from that volume is chosen to be the final intensity value at voxel p .
- (3) If $W_i(p)$ values are large in multiple volumes, the intensity values of those volumes are weighted by their respective $W_i(p)$ for a composite value.

3.3 Design Validation

To validate our design, we applied our weighting function to 3DUS images with line-like structures in Figure 3.2 and surface-like structures in Figure 3.3. Figure 3.7(a)-(b) and Figure 3.8(a)-(b) show the respective source 3DUS images smoothed with $\sigma = 2$.




Figure 3.6: Steps for US incident angle weighted compounding.

Algorithm 3 US incident angle weighted volume compounding

Define a threshold, below which W scores are considered small. Define $I_f(p)$ as the intensity value of the final compounded volume at voxel p .

```

for each voxel  $p = p(x, y, z)$  do
  for each  $i = 1 : N$  do
     $Ws(i) \leftarrow W_i(p)$ 
     $Is(i) \leftarrow I_i(p)$ 
  end for
   $index \leftarrow Ws > threshold$ 
  if length of  $index == 0$  then
     $I_f(p) \leftarrow Average(Is)$ 
  else if length of  $index == 1$  then
     $I_f(p) \leftarrow Is(index)$ 
  else if length of  $index > 1$  then
     $I_f(p) \leftarrow Average(Ws(index) * Is(index))$ 
  end if
end for

```

In Figure 3.2(a), because the US beam was almost in parallel to the orientation of the wooden toothpick, the resulting US image was low in SNR. The tip of the toothpick has noticeable artifact. If this image is registered and compounded using the intensity averaging method, the artifact would be introduced into the final composite images. When we measure this image using the coherence measures without accounting for the US beam direction, the tip of the toothpick is given high score as can be seen in the top row of Figure 3.7(c), which shows an XZ -plane slice of the measures of C_{curvel} , C_{surfel} and $C_{anisotropy}$ in the form 2D gray scale images. With our proposed measure, the tip of the toothpick is barely seen, as it should be for the purpose of improving compounding (Bottom row in Figure 3.7(c)). Also compare this to Figure 3.8(c), in which the orientation of the toothpick has a large angle with the incident beam and our measure assigns higher scores to them, but they are lower than the scores without considering the US beam incident angle.

Similar observations can be drawn from the comparison of the coherence measures of a surface. Figure 3.9(c) shows the measures of 3.9(b), where the US beam is oblique to the surface. Figure 3.10(c) shows the measures of 3.10(b), where the US beam is almost perpendicular to the surface. Across the board, the scores for the surfel components from Figure 3.10(c) are higher than those in Figure 3.9(c). In addition, when compare the bottom row of the two figures, we see the differences in the surfel measures with respect to the incident wave. The W value of the surface almost perpendicular to the US beam angle is much higher.

Based on the comprehensive comparisons above, we demonstrated that our measure is able to capture the combination of the local structure orientation and its relationship with the US incident wave. This prevents artifacts from being introduced from views of small angle and help enhance features from views of large angle during the compounding.






Figure 3.7: Comparison of coherence measures using 3DUS source image in Figure 3.2(a). (a) Source 3DUS smoothed with $\sigma = 2$. (b) 2D XZ-plane slice of (a). (c) Coherence measure of (b). Top row: C_{curvel} , C_{surfel} , and $C_{anisotropy}$. Bottom row: W_{curvel} , W_{surfel} and W measure weighted by US incident angle. The tip of the toothpick region is assigned a much lower score in bottom rows.



(a) Source 3DUS ($\sigma = 2$)



(b) 2D XZ-plane slice from (a)






Figure 3.8: Comparison of coherence measures using 3DUS source image in Figure 3.2(b). (a) Source 3DUS smoothed with $\sigma = 2$. (b) 2D XZ-plane slice of (a). (c) Coherence measure of (b). Top row: C_{curvel} , C_{surfel} , and $C_{anisotropy}$. Bottom row: W_{curvel} , W_{surfel} and W measure weighted by US incident angle.



(a) Source 3DUS ($\sigma = 2$)



(b) 2D XZ-plane slice from (a)





Figure 3.9: Comparison of coherence measures using 3DUS source image in Figure 3.3(a). (a) Source 3DUS smoothed with $\sigma = 2$. (b) 2D XZ-plane slice of (a). (c) Coherence measure of (b). Top row: C_{curvel} , C_{curvel} , and $C_{anisotropy}$. Bottom row: W_{curvel} , W_{curvel} and W measure weighted by US incident angle.



(a) Source 3DUS ($\sigma = 2$)



(b) 2D XZ-plane slice from (a)




Figure 3.10: Comparison of coherence measures using 3DUS source image in Figure 3.3(b). (a) Source 3DUS smoothed with $\sigma = 2$. (b) 2D XZ-plane slice of (a). (c) Coherence measure of (b). Top row: C_{curvel} , C_{curvel} , and $C_{anisotropy}$. Bottom row: W_{curvel} , W_{curvel} and W measure weighted by US incident angle.

3.4 Results

3.4.1 Water tank results

We conducted water tank experiment to evaluate our method. We used a Philips Sonos 7500 3DUS machine (Philips Healthcare, Andover, MA, USA) with an Xt-7 3D matrix probe. The probe was tracked by an electromagnetic (EM) tracking system (3D Guidance trakStar System, Ascension Technology Corporation, Shelburne, VT, USA, $1.4mm$ and 0.5° RMS accuracy) [88] for registration purpose. The imaging phantom was a wooden ball embedded in a gelatin-based mixture that mimics the echogenic properties of animal tissue [59] (Figure 3.11(a)). We scanned the phantom at depth $80mm$ and frequency $5MHz$ from a number of different directions - two orthogonal views from the top, and two opposing directions from the sides, as illustrated by the outlines in Figure 3.11(b).

Figure 3.12(a)-(d) show typical volumes acquired from different view points. Note the signal dropout due to reflection from the surface of the wooden ball and missing surface definition where the surface is parallel to the US beam. Each single volume only provides very limited information of the ball. The mosaiced volume is two times the size of the single volume and shows the entire top half of the inner surface and cavity of the ball. Figure 3.13 compares a complete cross section from the mosaiced volume to the contributing single slices. Figure 3.13(d) shows the complete contour of a cross section of the wooden ball. When compared to intensity averaging compounding result (Figure 3.13(e)), our method preserves the details of the inner surface from each contributing single volumes much better.

The compounding of two volumes of size $128 \times 128 \times 256$, including the computation of the coherence measures takes about 60s on a CPU (Intel Core i7 CPU 920 @ 2.67 GHz, 12 GB Memory, 64-bit operating system). Volume registration is based on our previously reported methods [13, 83] and can already achieve real-time performance of $\sim 30ms$ per every two volumes of $256 \times 256 \times 256$.




Figure 3.11: Water tank imaging phantom. (a) Wooden ball embedded inside a gelatin-based phantom. (b) Close up view of the wooden ball. The colored slices indicate representative imaging angles.






Figure 3.12: Typical single volumes of the wooden ball phantom from different views. Note the signal dropout and missing surface definition where the surface is parallel to the US beam.



(a) Segments of the cross section



Complete cross section

Figure 3.13: An XY -plane 2D slice of the mosaiced volume showing a complete cross section of the wooden ball. (a) Example XY -plane slices from contributing single volumes. (b) An XY -plane 2D slice from our method. (c) The same slice from the intensity averaging compounding method. Note that the features are better preserved in (b).




Figure 3.14: Example of contributing 3D volumes

3.4.2 *in vivo* results

The *in vivo* data was collected on a human subject using an X7-2t 3D matrix TEE probe connected to a Philips iE33 3DUS machine with ECG gating. Total of five different views at the same cardiac phase were registered using a real-time image-based rigid registration method developed by Schneider *et al.* [83]. Each single contributing volume is of the size $160 \times 48 \times 208$ but only offers a partial and incomplete view of the mitral valve area. (Figure 3.14).

The FOV of the mosaiced volume is three times the original volume. It can be seen from Figure 3.15(a) that important anatomical structures such as the mitral leaflets are much more complete in the mosaic volume, and preserved better than that in Figure 3.15(b). It can also be noted that our method performs better at the boundary of the input volumes. Less artifact is introduced from the input volume boundaries.

3.4.3 Contrast to noise ratio

To quantify the results, we computed the contrast to noise ratio (CNR) defined as the difference between the mean values of the feature and background, divided by the standard




Figure 3.15: Mosaiced volume slice views of the YZ-plane. (a) US beam direction weighted compounding. (b) Intensity averaging. Result in (a) shows better definition of the mitral valve leaflets and better performance at the boundary of input volumes.

deviation of the background and the feature

$$CNR = \frac{\mu^{feature} - \mu^{background}}{\sqrt{(\sigma^{feature})^2 + (\sigma^{background})^2}} \quad (3.13)$$

where μ represents the mean, and σ the standard deviation. CNR indicates how well a feature is shown against its surrounding background.

Figure 3.16 shows the areas selected for computing the CNR scores. Blue arrows indicate the features and white arrows indicate the background. The features in Figure 3.16(a) are the walls of the wooden ball, and the background includes the wooden ball cavity and the tissue around the surface of the ball. CNR scores from the intensity averaging compounded volume and from the weighted compounding volume are compared. The results are tabulated in Table 3.1. Two areas are selected in Figure 3.16(b). The dotted rectangle on top is mitral leaflets against the blood pool, while the dotted rectangle at bottom is the left ventricular wall against the blood pool. The CNR scores for these two areas are tabulated in Table 3.2 and Table 3.2. The range of the intensity values for all




Figure 3.16: Features and backgrounds selected for computing CNR scores. Blue arrows indicate features and white arrows indicate background. (a) A cross-section of the wooden ball phantom. (b) Left atrium and left ventricle of the *in vivo* TEE data.

the data sets is from 0 – 255. When compared to intensity averaging, we achieve $\sim 20\%$ increase in CNR. Our weighted compounding method performs better in preserving the feature-to-background ratio.

Table 3.1: Wooden Ball Contrast to Noise Ratio

Method	$\mu^{feature}$	$\mu^{background}$	$\sigma^{feature}$	$\sigma^{background}$	CNR	$\Delta CNR\%$
Intensity Averaging	101.84	68.45	14.87	8.28	1.96	—
Weighted Compounding	136.49	70.75	25.72	10.68	2.36	20.41

Table 3.2: *in vivo* Mitral Leaflets Contrast to Noise Ratio

Method	$\mu^{feature}$	$\mu^{background}$	$\sigma^{feature}$	$\sigma^{background}$	CNR	$\Delta CNR\%$
Intensity Averaging	150.13	41.59	53.88	12.99	1.95	—
Weighted Compounding	160.36	36.23	49.70	10.68	2.44	20.10

Table 3.3: *in vivo* Left Ventricular Wall Contrast to Noise Ratio

Method	$\mu^{feature}$	$\mu^{background}$	$\sigma^{feature}$	$\sigma^{background}$	CNR	$\Delta CNR\%$
Intensity Averaging	161.15	27.80	46.40	21.75	2.60	—
Weighted Compounding	177.69	22.56	46.22	17.10	3.15	21.15

3.5 Discussion and Conclusion

In this chapter, I presented a novel 3DUS compounding method based on a coherence measure derived from local structure tensors and weighted by the US incident angle. This method addresses the issue of integrating US images from multiple viewpoints where the intensity values of the overlapping region could be significantly different due to difference in the orientations of the local tissue structures with regards to US beam incident angles. The commonly used approach of intensity averaging tends to decrease the contrast at the regions of prominent features, an outcome we would like to avoid. Our weighted coherence measure takes into account of the alignment of the feature orientations with respect to the US beam, and assigns a higher score to those that have a larger angle and a lower score to those with a smaller angle. A compounding algorithm based on this measure was also presented. Both water tank and *in vivo* studies demonstrated a $\sim 20\%$ increase in CNR when compared to intensity averaging.

3.5.1 Multiscale compounding

The scale of the window function for generating the structure tensor largely determines the size of the features to be measured, and has an impact on the assessment of local structures. For example, Figure 3.17 shows the effects of different scales on the structure detected. At $\tau = 4$, the line-like toothpick also has relatively high score as surfels and is detected as regions of two surfaces with opposing normal. At a higher scale $\tau = 10$, the curvel score becomes noticeably higher than that of the surfel. Choosing the right scale for a feature of interest is an important part of the method.

In cases where features of different sizes are of interest, the compounding method




Figure 3.17: *Different scales for the window function. Top row: Source 3DUS image. Second row: W_{curvel} , W_{surfel} , and W at scale $\tau = 4$. Third row: W_{curvel} , W_{surfel} , and W at scale $\tau = 10$. Note that at $\tau = 4$, the line-like toothpick is also detected as regions of two surfaces with opposing normal. At scale $\tau = 10$, the surfel score becomes small when compared to that of the curvel.*

presented in Section 3.2 can be extended to include multiple scales. For example, first W is computed at different scale levels, and then the values of W at different scales are compared to determine the final compounded image using the compounding steps outlined in Section 3.2.3.

3.5.2 GPU implementation

Convolution with Gaussian filters and the computations of the eigenvalues and eigenvectors are the two time consuming steps in the implementation of structure tensor based compounding. Fortunately, both of them are highly parallelizable.

We implemented a 3D convolution kernel accelerated on a computer graphics card by means of the CUDA parallel computing model described in Chapter 2. To compute a structure tensor of a given 3D image, we upload the 3DUS volume to GPU and store it as a 3D texture or linear array. The Gaussian kernels are stored in constant memory. DoG convolution can be done in 1D for X, then Y and Z direction. The resulting structure tensor is stored in a linear array. Next, the computation of eigenvalues and eigenvectors can be implemented such that GPU threads are mapped to individual voxel, providing coalesced memory accesses and massive parallelism.

This approach is simple and scalable, only limited by the global memory bandwidth. We expect GPU acceleration will reduce the computation time significantly and bring it to real-time performance.

Chapter 4

Automated Sweeping and Instrument Tracking with Ultrasound Catheters

4.1 Introduction

Catheters enable many diagnostic and repair procedures to be accomplished with minimal collateral damage to the patient's healthy tissues. In complex catheter procedures, workflow is often limited by visualization capabilities, which contributes to operator's inability to prevent and assess complications, as well as facilitation of key procedural components. In electrophysiological (EP) cardiac procedures, guidance is largely provided by fluoroscopy, which is unable to image soft tissues. To compensate for this, a widely-adopted approach is to generate a 3D electrophysiological model resembling the shape of the cardiac chamber by sampling surface points with a magnetic position sensing system that records the locations of the ablation catheter tip in space. The point clouds may then be registered to a CT or MRI based pre-operative anatomic model and displayed to the clinician with the real-time positions of the catheters superimposed [22, 23]. Although catheter-based geometry is acquired real-time, the pre-operative rendered anatomy is not, and thus may result in an anatomic mismatch at the time of the procedure.

Ultrasound (US) imaging catheters (intra-cardiac echocardiography, or ICE) have been

routinely used in EP procedures for over a decade [54]. These catheters are inserted into the patient's vasculature (e.g. femoral vein) and navigated to the heart, where they acquire 2D B-mode images of cardiac structures.

Compared to transthoracic (TTE) and transesophageal echocardiography (TEE) probes, US catheters produce higher resolution images with their proximity to cardiac structures in the near-field, before acoustic lines diverge and image plane thickness increases, and where attenuation of higher frequencies (up to 10MHz) is not a limiting factor. In addition, US catheters provide access to a far greater range of viewing angles than from the limited windows available for TTE and TEE imaging. For example, attempts to use TEE for guiding atrial ablation procedures have found that visualization of the inferior pulmonary veins with TEE is not possible in a significant fraction of patients [68, 68]. The versatility of ICE imaging is particularly important during EP procedures, as it provides excellent visualization of all cardiac chambers when the probe is placed in the appropriate anatomic position. Recent studies suggest that ICE monitoring of lesion formation may increase the effectiveness of ablation procedures [54, 95].

Unfortunately, controlling ICE catheters requires the clinician to aim the imaging plane by manually turning control knobs and rotating and advancing the catheter handle. This makes it highly challenging to align the image plane with the target, so moving between targets requires extensive time and skill to obtain an adequate view. This has largely limited the use of ICE to a few critical tasks such as transeptal puncture. Recently-developed 3D ICE catheters can acquire only a limited number of slices in the third dimension due to the small catheter diameter. This 'thick slice' remains challenging to manually aim due to the non-intuitive controls and narrow field of view.

Automated transducer pointing will make ICE far more useful, effective, and broadly applicable. Cardiologists presently use a combination of pre-operative images, fluoroscopic imaging, electroanatomic mapping, and haptic feedback through the catheter handle to navigate ablation catheters. However, the actual catheter tip-to-tissue interaction can only be visualized in real time with the use of US imaging. This presents a challenge for the

operator because significant training and time are required to manually maneuver US imaging catheters.

We hypothesize that automatic tracking of the ablation tool tip with direct visual feedback will better facilitate ablation procedures, including confirmation of adequate catheter tip-to-tissue contact. Real-time monitoring also enables rapid lesion assessment and may aid in the detection of impending complications. Automatic panoramic US imaging and enhanced displays also promise to decrease the need for fluoroscopy, reducing ionizing radiation exposure to patients and medical personnel. To our knowledge, no similar capability has been reported in the literature. Instrument tracking and real-time visual feedback using ICE catheters are unique contributions of the system presented here.

The system provides different functionality than current commercial catheter robots, such as the Amigo from Catheter Robotics [2], CorPath from Corindus [25], Artisan from Hansen Medical [5], and EPOCH from Stereotaxis [28]. They primarily enable teleoperation of catheter controls to increase operator comfort and reduce exposure to radiation from fluoroscopic imaging. These systems simply replicate manual control knob, which does not mitigate the difficulties of aiming imaging catheters in joint space. While some systems offer limited Cartesian control, these systems do not feature orientation control, which is essential for aiming an imager.

The system described in this chapter is a collaboration between myself and my colleague Paul Loschak at Harvard Biorobotics Laboratory. Paul was responsible for the mechanical engineering and kinematic modeling portion of the work while I was responsible for ultrasound image processing and visualization, in addition to the overall software design and development.

In this chapter, I begin with an overview of our robotic US catheter steering system. Next, I describe 3D panorama creation, instrument tracking and tip identification, capabilities enabled through our robotic system. I will present our experimental results and conclude this chapter with discussions on both the contributions and limitations of the current system.

4.2 System Design

4.2.1 Overview

We developed a robotic ICE catheter steering system to automate the pointing of ICE catheters. The ICE catheter we used for testing and validation was a 10Fr (3.30mm diameter) 110cm long catheter with a 64-element 2D ultrasound transducer at its distal tip (AcuNav, Biosense Webster, Diamond Bar, CA, USA). The AcuNav ICE catheter is the most common ICE catheter in clinical use at present.

An AcuNav ICE catheter can be viewed as a four degree of freedom (DOF) system that has two orthogonal bending directions (Pitch and Yaw) and can translate along and rotate about its body axis (Figure 4.1). It is unique in that it features an extra bending direction for achieving desired tip orientations for imaging purposes. The fourth DOF offers a number of functionality. For example, using 3 DOFs to reach a position and the extra 1 DOF to steer the imaging plane to a specific orientation, or spinning the catheter tip about its own axis, thereby sweeping the imaging plane across a desired region, while keeping the catheter tip in the same location.

Figure 4.2 is a block diagram of our system. At the center of the system is a robotics manipulator (Figure 4.3) that has four pairs of motor/EPOS controller (Maxon Precision Motors, inc., Fall River, MA, USA), each controlling one of the 4 DOFs. Closed form solutions for forward and inverse kinematics were developed to enable position control of the catheter tip. Additional kinematic calculations enables 1-DOF angular control of the imaging plane. The robot controls an ICE catheter. An electromagnetic (EM) tracker system (3D Guidance trakStar System, Ascension Technology Corporation, Shelburne, VT, USA, 1.4mm position and 0.5° RMS accuracy) [88] is rigidly attached at the distal end of the ICE catheter, and is used for closed loop control, performance validation, and safety functions. The other end of the ICE catheter connects to an ICE compatible US machine, the output of which is sent to our image processing pipeline, and the results are sent to display. The user interacts with the robot and the display through a user interface module. The combination




Figure 4.1: ICE catheter - a four DOF system. Yaw and pitch of imaging planes are controlled by the two bending knobs. Rotation is achieved by rotating around the Z axis. Translation is done by moving the catheter body along its base Z axis. (Image courtesy of Paul Loschak)

of positioning with imaging plane rotation enables a wide range of visualization capabilities, such as quick sweeps from inside the heart for 3D/4D volumes reconstruction or ablation catheter tracking [52].

Getting a US catheter robotic and visualization system together required both mechanical engineering, kinematics modeling and careful software design and implementation. The software needs to be able to process multiple data feeds, such as motor encoder readings, EM position tracking outputs, and ICE image streaming, simultaneously and responds with proper control commands at a timely fashion. Based on these requirements, I implemented a C/C++ based software program using multi-threaded programming, with an user interface that offers direct interaction with critical modules that include motor controls, EM tracking, kinematics, and imaging tasks. Parallel computing on GPU is also used here for the computationally intensive tasks such as volume registration and rendering. Figure 4.4 provides an overview of the major software modules and the interaction among them.

Safety is always a priority in a clinical environment, thus being able to perform imaging tasks while keeping the ICE catheter tip stationary is a very useful capability. It prevents potential unwanted collisions with important anatomical structures. Based on this important assumption, we developed the following key capabilities: (1) Automated sweep while keeping the ICE catheter tip at a fixed location, and build a 3D ‘panorama’ that shows the




Figure 4.2: *System Overview.* A robotic manipulator controls an ICE catheter, which is connected to an ICE compatible ultrasound machine. The distal tip of the ICE catheter is tracked by an EM tracker. Image processing and user interface are developed to support a number of visualization functionality.




Figure 4.3: *System Hardware* (a) Imaging phantom and ICE catheter tip with an EM sensor attached. (b) Robotics manipulator. (Image courtesy of Paul Loschak)

tissue structure across a treatment area and beneath the surface; (2) Follow the working instrument (e.g. ablation catheter) by keeping its tip constantly in the imaging plane; and (3) Sweep locally (dither) near the ablation tool tip for a detailed view of the work region and provide real-time visual cues on tip-tissue interaction.




Figure 4.4: *Software Modules and Data Flow Diagram. The CPU handles data inputs, kinematics and control algorithms while the GPU handles computational intensive tasks such as volume registration and rendering.*

4.2.2 Sweeping

During manual manipulation, a clinician may wish to position the ICE catheter in a desired region of the heart and sweep the imaging plane to get a comprehensive view of the region. While bending, it is extremely difficult to intuitively and manually spin the catheter about its own axis while keeping the tip in place. Our system uses a combination of yaw, pitch, and handle roll to rotate the catheter tip about its distal tip (Figure 4.5(a)). This is different from simply rotating the catheter handle or body around its rotational axis.

The automated sweeping mode allows the user to input a desired range of angles to sweep with a specified angular resolution at the setup. During the sweeping, the imager is rotated through a sequence of angles in the range defined by the user, while having its tip continuously position-controlled to remain at the fixed point. Several sweeps can be done at a few different user specified locations to generate a large patient specific anatomical map for real-time navigation guidance.




Figure 4.5: Schematic diagram of (a) Sweeping while keeping the tip stationary. (b) Instrument tracking. (Image courtesy of Paul Loschak)

4.2.3 Instrument tracking

In instrument tracking mode the system aims the imaging plane at a moving target by computing the angle between the target and the ICE imaging plane and commanding a specific sweep angle. Simultaneously, the position controller makes small adjustments of the tip position to maintain the ICE catheter tip position (Figure 4.5(b)).

4.2.4 3D panorama of 2D ICE images acquired from sweeping

The sweeping functionality enables acquisition of closely-spaced 2D images across a user-specified region of interest (ROI). The 2D slices are non-parallel sections which need to be spatially registered into a common Cartesian coordinate frame using catheter tip EM tracker locations and then interpolated and compounded into a gridded 3D volume. There are several leading methods for 3D reconstruction of ultrasound images [7, 33, 75, 81, 86]. Our method is similar to the voxel based approach. In order to achieve real-time performance, we implemented the 3D reconstruction and visualization on a GPU, based on our previous real-time mosaicing technique for 3D/4D ultrasound [13, 14].

A 2D US image has an infinite thickness due to US out-of-plane beam width. The

thickness of 2D ICE images ranges from 3 – 6mm primarily depending on the imaging depth [99]. To interpolate an input 2D ICE image to the reference volume, we first add a small thickness of 1mm to represent the mid-plane. This reduces the edge artifacts during the interpolation. We also add some empty space to the slice (zero-padding) to avoid losing some part of the image near the border during registration. The padded image is then registered to the reference volume. For high resolution reconstruction, we choose a voxel resolution same as the pixel spacing of the ICE image (0.16mm at imaging depth 90mm). The transformation matrix is determined by the EM sensor outputs and a series of coordinate transforms from the 2D ICE image frame to the EM transmitter coordinate frame, which is the reference frame. The registered volume is then compounded to the reference volume using weighted averaging. The registration and compounding are done on GPU using the algorithms similar to those outlined in Section 2.2.3 of Chapter 2. Figure 4.6 summarizes the steps described in this section.

Figure 4.7 shows the coordinate frames and transformation matrices that determine the registration matrix. T_{EMS}^{CT} is the transformation matrix from EM sensor to ICE catheter tip, and T_{CT}^{IMG} is the transformation matrix from ICE catheter tip to the image frame. Both are known *a priori* by attaching the EM sensor rigidly at a known position and orientation on the catheter tip, and by knowing the image resolution and dimension. We further define T_{EMS} as the EM sensor outputs, $P = [i, j, k, 1]^T$ as a voxel at location (i, j, k) in its homogeneous vector form, and T_s as the scaling matrix that converts voxel unit to a physical unit. We can establish the following between an input image P_{input} and its location P_{ref} in the reference coordinate

$$T_{EMS_{ref}} \cdot T_{EMS}^{CT} \cdot T_s \cdot T_{CT}^{IMG} \cdot P_{ref} = T_{EMS_{input}} \cdot T_{EMS}^{CT} \cdot T_s \cdot T_{CT}^{IMG} \cdot P_{input}$$

The registration matrix T_{input}^{ref} thus is

$$T_{input}^{ref} = T_{CT}^{IMG^{-1}} \cdot T_s^{-1} \cdot T_{EMS}^{CT^{-1}} \cdot T_{EMS_{ref}}^{-1} \cdot T_{EMS_{input}} \cdot T_{EMS}^{CT} \cdot T_s \cdot T_{CT}^{IMG}$$

The registered 2D slices fan out in the lateral direction, thus are not well aligned with




Figure 4.6: 3D Panorama Creation Pipeline




Figure 4.7: Coordinate frames and transformation matrices. T_{EMS}^{CT} is the transformation matrix from EM sensor to ICE catheter tip, and T_{CT}^{IMG} is the transformation matrix from ICE catheter tip to the image frame. Both are known a priori by attaching the EM sensor rigidly at a known position and orientation on the catheter tip and by knowing the image resolution and dimension. EM transmitter coordinate frame is chosen as the reference frame.

the gridded space, resulting in gaps in the reconstructed 3D voxel array. In our system, the 2D slices can be acquired at a small spacing in order to keep the gaps small, at the expense of longer sweeping times. To fill the gaps, a number of gap filling methods have been reported, including voxel averaging [55, 63], adjacent voxel interpolation [39], Voxel Nearest Neighbor (VNN) [75] and Distance Weighted interpolation (DW) [90]. Coupe [26] reported a 3D freehand US reconstruction method using probe trajectory, which handles the gaps at the same time. They concluded that since the virtual slice was generated by using the information from the closest two frames, this method outperformed traditional reconstruction approaches such as VNN and DW. The main limitation of their method was the assumption of constant probe speed between two slices and the added time on probe trajectory estimation.

This limitation does not exist in our system. Both the actual catheter tip locations at two adjacent frames and the commanded tip trajectory are known through our control module. This allows us to generate an image through a ‘virtual’ tip position on the trajectory by projecting images from the two closest 2D frames to the position. A virtual image is then generated as the composite of the two projected images and interpolated onto the 3D volume. For instance, in Figure 4.8, I_i^{ref} and I_{i+1}^{ref} represent two adjacent image frames i and $i + 1$ acquired at an angular spacing of θ , and registered to the reference frame, which is the EM transmitter coordinate frame as shown in Figure 4.7. T_i^{ref} and T_{i+1}^{ref} are their respective transformation matrices in terms of the reference coordinate. The red arrow shows the catheter tip Z axis Z_{CT} , which coincide with EM sensor X axis X_{EM} . The commanded catheter tip trajectory from frame i to $i + 1$ is to roll about Z_{CT} by θ .

In Figure 4.8, $I_{v_i}^{ref}$ is a virtual slice computed by rotating the catheter tip at frame i counter clockwise by an angle of ϕ around Z_{CT} . Here we choose $\phi = \frac{\theta}{2}$. This is the same as rotating around X_{EM} by ϕ in the EM transmitter reference coordinate frame. We use $T_{v_i}^i$ to




Figure 4.8: Virtual slice $I_{v_i}^{ref}$ represents the image formed by rotating the catheter tip at frame i by an angle of ϕ around Z_{CT} from. I_i^{ref} and I_{i+1}^{ref} are two adjacent image frames i and $i + 1$ acquired at an angular spacing of θ , and registered to the reference frame. T_i^{ref} and T_{i+1}^{ref} are their respective transformation matrices in terms of the reference coordinate.

represent this transformation, and it is expressed as

$$T_{v_i}^i = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\phi) & -\sin(\phi) & 0 \\ 0 & \sin(\phi) & \cos(\phi) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Therefore

$$I_{v_i}^{ref} = T_{v_i}^i \cdot I_i^{ref}$$

and its location in terms of the reference is

$$T_{v_i}^{ref} = T_{v_i}^i \cdot T_i^{ref}$$

Similarly, a virtual slice $I_{v_{i+1}}^{ref}$ can be created by rotating frame $i + 1$ around Z_{CT} clockwise by ϕ . A final virtual slice I_v^{ref} can be computed by averaging $I_{v_i}^{ref}$ and $I_{v_{i+1}}^{ref}$.

4.2.5 Instrument tip localization and visualization

Localizing the instrument tip in ultrasound is a challenging task due to noise, artifacts and US out-of-plane beam width. Here, we present a novel but simple technique to identify and visualize the instrument tip. First, we find the rough location of the tip through EM based tracking, next we dither near the tip region across a user defined angular range (e.g. $\pm 5^\circ$) so that we sweep past the tip. This ensures that the actual tip location will be part of the reconstructed volume with an accuracy only limited by the spacing of the adjacent frames. The 3D volume shows the extend of the tip and the tip location is identified by registering the EM sensor on the tip to the 3D volume. Slice views that cut through the instrument shaft and tip can also be extracted automatically.

4.3 Experimental Results

4.3.1 Sweeping

We conducted water tank experiments using gelatin-based phantoms that closely mimic animal tissue features. The ICE catheter was connected to a Siemens Acuson X300 US imaging system (Siemens Medical Solutions USA, Inc. Malvern, PA, USA) and introduced through the side of the water tank where the imaging phantom was located.

The first experiment was to sweep across a specified ROI and build a 3D panorama. The imaging phantom is shaped to resemble the left atrium with four openings that simulate the pulmonary vein ostia (PVO), which are the critical areas for imaging during an atrial fibrillation ablation procedure. The atrium area (the opening in the phantom) is roughly $40mm \times 40mm$. The ICE tip was directly in front of the phantom (Figure 4.9(a)). The images were acquired at $90mm$ depth and $6.7MHz$ frequency. The system swept across the phantom in 1° increments over 40° while the ICE tip position remained stationary. Volume rendering is done as described in 4.2.4. Figure 4.9(b)-(d) shows the reconstructed 3D volume from three different view points. The PVOs are easily seen from Figure 4.9(c).




Figure 4.9: Sweeping result shows the reconstructed volume of the atrium phantom from three different views. The simulated PVO can be clearly seen from (c).

4.3.2 Instrument tracking results

In the instrument tracking experiment, the instrument was a 3mm diameter catheter with an EM sensor attached at the tip and calibrated to the catheter tip position. It closely resembled the dimension and echogenic properties of an ablation catheter tool tip. The ICE catheter tracked the tool tip as it moved around the simulated PVO in the atrium phantom. Figure 4.10(a) shows the experiment setup and 4.10(b) plots the imaging plane as it followed the tool tip. The lines are imaging planes, and the circles represent tool tip positions. Colors indicate corresponding imaging planes and tool tip positions. We calculated the distance of each of the targets to the center line of their respective image plane. On average, the targets are within $\pm 1\text{mm}$ of the imaging plane center line thus appear in the images. The angular tracking accuracy reported by the system is 0.3° , adding EM sensor's angular error of 0.5° [88], the overall angular tracking accuracy of the system is 0.8° .




Figure 4.10: Instrument tracking (a) Instrument tip at PVO. (b) Imaging plane following the tool tip. The lines show imaging planes, the circles represent tool tip positions. Colors indicate corresponding imaging planes and tool tip positions. Targets are within $\pm 1\text{mm}$ of the imaging plane center line, thus appear in the images.

4.3.3 Instrument tip visualization

To compensate for misalignment and artifacts, we implemented a correction method, based on the sweeping capability developed. We dither locally near the instrument tip region, first sweeping the tip and then continuing to sweep until the tip is no longer in view (Figure 4.11(a)). This way, the actual tip is guaranteed to be included in the swept volume. Figure 4.11(b) is an example of a 2D ICE image showing the ablation tool tip, while Figure 4.11(d) contains no tip. Note the artifacts around the ablation tool tip which accentuate both the importance and challenge in tool tip identification. In Figure 4.11(c), the entire distal tip section is visualized in 3D volume rendering and the tip is identified.

The tip is imaged at $60mm$ depth with an angular spacing of 0.5° between the acquired frames (Figure 4.12). The accuracy of the tip identification is derived by adding EM tracking position accuracy $1.4mm$ [88] and our system reported angular accuracy of 0.3° to d_2 . Thus the overall tip identification accuracy of our system is $\sim 2.2mm$.

4.3.4 ICE imaging plane thickness vs. system accuracy

Inaccuracy of the EM trackers in ICE and the working instrument may result in misalignment of the US imaging plane with the instrument. In this case the image may not show the tool tip. To analyze accuracy limits, the thickness of the ICE imaging plane can be compared to the possible positioning errors of the tool tip. Figure 4.13 illustrates a simplified case considering only misalignment error in the sweep angle, using a typical ICE image plane depth $90mm$ and thickness $6mm$ [99]. RMS accuracy specifications for the EM tracker are $1.4mm$ in position and 0.5° in angle [88] and our system tracking accuracy is 0.3° . A simple geometric analysis of the worst case accuracy scenario shows that an ablation catheter of diameter $3mm$ would be visible in the US image, although it may not be in the mid-plane. Metal interference in clinical procedure rooms will likely decrease the accuracy of the EM trackers, nonetheless the close proximity of the imaging catheter to the ablation catheter will minimize the relative tracking errors.




Figure 4.11: Dither for tool tip identification (a) Region to be swept includes the tip and beyond. (b) 2D ICE image containing part of the tip. Note the artifacts around the tip. (d) 2D ICE image containing no tip. (c) Reconstructed 3D volume showing the instrument shaft, tip and surrounding tissue.




Figure 4.12: Tool tip identification accuracy. The tip is imaged at $d_1 = 60\text{mm}$ depth with an angular spacing of $\theta = 0.5^\circ$ between the acquired frames. The accuracy of the tip identification is derived by adding EM tracking position accuracy 1.4mm [88] and our system reported angular accuracy of 0.3° to d_2 . The overall tip identification accuracy of our system is $\sim 2.2\text{mm}$.




Figure 4.13: (left) Typical ICE image plane depth; (right) cross section view of imaging plane for calculation of imaging accuracy based on EM measurement accuracy of sweep angle: solid line is EM-estimated boundary of image plane, dotted line is true image plane; circle is target ablation catheter cross section.

4.4 Discussion

Situational awareness plays a critical role in intra-operative procedure guidance. The current environment lacks real-time direct visual feedback of instrument-tissue interactions, which in part contributes to the low success rate of ablation procedures. The system described in this chapter addresses this issue with the following capabilities: (1) Build real-time panorama of a ROI by spinning the ICE catheter at desired angles while keeping its tip at a fixed and safe location; (2) Instrument tracking and tool tip location identification; (3) Volume rendering of the instrument tip and the surrounding tissue. All these tasks would be difficult to achieve manually. The current system can be easily extended to 4D (3D + time) based on the work discussed in Chapter 2 on 3D US mosaicing and visualization with ECG gating. In addition to better spatial localization than 2D views, the reconstructed 3D volume can also be used for surface and instrument segmentation, facilitating the integration with the current clinical EM mapping system or with an augmented virtual reality environment.

The robotic system in combination with US image processing and visualization capabilities has the potential to further improve the intra-operative procedure guidance.

Chapter 5

Conclusion and Future Work

5.1 Conclusion

3DUS is a promising modality for procedure guidance. Its ability to visualize soft tissue at real-time rates is an important aspect in enabling real-time visual feedback, which is lacking in current clinical environment. However, the adoption of 3DUS has been slow due to its small field of view, limited resolution, signal dropout and high learning curves in interpretation and analysis. This thesis aims to provide new engineering technologies to mitigate these key challenges and ultimately improve the quality of 3DUS guidance.

In Chapter 2, I presented a set of enhanced display techniques for real-time 3DUS visualization. The system integrates EM tracking system with GPU implementation for real-time ECG gated 3D mosaicing and instrument tip cut plane tracking. The user study with cardiologists demonstrated the potential of such enhanced visualization in instrument navigation and procedure execution with 3DUS guidance.

Ultrasound incident angle can affect the outcome of the image, and this could result in missing anatomical features from certain data acquisition angles. In many situations, images from complimentary views are acquired in order to reconstruct a complete view of the region of interest. To further improve the mosaicing results, in Chapter 3, I described a novel 3DUS volume multiview compounding technique using structure tensors. Structure

tensors are commonly used in computer vision for local coherence estimation. Here I designed a structure tensor based measure that incorporates information from both the local structures and their orientations in regards to US incident waves. The measure was validated by water tank and *in vivo* experiments and has shown a 20% increase in contrast to noise ratio when compared to intensity averaging. When compared to other recent work on 3DUS multiview compounding, such as phase based and wavelet based methods, which are computational intensive, this approach is highly parallelizable, thus can be accelerated with parallel programming on computer graphics cards with minimal effort.

ICE catheters have been used for visualization in beating heart procedures such as atrial fibrillation ablation. Although it provides better resolution from imaging inside the heart and reduces the exposure to radiation, its application has been limited to few due to the high learning curve in manual control of the imaging planes. We argue that a robotic steering system could broaden its role in procedure guidance. In Chapter 4, I presented visualization techniques developed alongside the robotic system we built. By controlled imaging near the instrument tip and its surrounding tissue, the system provides volumetric rendering of the local work region and helps identify the exact tip location. We are able to achieve an accuracy of $2.2mm$ RMS in position and 0.8° in angle during instrument tracking.

5.2 Future Work

5.2.1 3DUS surface boundary detection for volume rendering

Facilitated by improvements in computer graphics cards, direct volume rendering has become a popular visualization method, by which, volumetric data is visualized without an explicit surface, instead, the delineation of the surface is replaced by applying an opacity transfer function. This approach has worked well for medical CT and MRI data, but not for ultrasound data due to its speckle artifacts, low dynamic range, low SNR, and often incomplete surface boundary [43].

Assigning the voxels with non-binary, continuous opacity values is essential for dis-

playing smooth and anti-aliased surfaces in volume rendering. Different opacity transfer function (OTF) approaches have been suggested. The histogram volumes method [47] generates semi-automated transfer function by using the gradient vector to find the direction that passes perpendicularly through the object boundary and then use the statistical properties of the histogram to provide the overall picture of the boundary characteristics. An adaptive design of OTF is also proposed based on the evaluation of collections of voxels gathered by transversing the volume in the rendering direction (i.e. tube cores) [38]. The VolumePro real-time volume rendering system offers an interactive transfer function design and uses gradient magnitude modulation of opacity and illumination [70].

All the approaches above make the same underlying assumption that the features of interest in the volume are the boundary regions between areas of relatively homogeneous material [6]. Gradient and directional derivatives along the gradient are mostly commonly used to extract or infer boundaries.

I propose to design an OTF based on coherence measures from structure tensors. Structure tensor presents integrated information of a local neighborhood without cancellation effects, thus making it robust against noise and a more reliable estimator of orientations in noise-prone data, such as ultrasound [17]. Smoothing also distributes the information about the orientation into the areas between edges. Thus we can estimate the dominant orientation at locations where the gradient is close to zero [16]. US incident angle can also be taken into account to remove the artifacts resulted from less optimal viewing angles.

The coherence measure modulated OTF can be implemented using a parallel programming approach by utilizing 3D texture mapping on GPU at voxel level. Real-time user interaction to adjust the weighting function is also feasible. I hypothesize that this method is more robust in dealing with noisy US data than existing OTF approaches and has the added benefit of reducing incident angle dependent artifacts.

5.2.2 US catheters and robotics

In clinical settings, a system that offers ease for use is essential. An ICE catheter, such as the AcuNav (Biosense Webster, Diamond Bar, CA, USA) [1] can be viewed as a four DOF robotic system. By building a robotic manipulator, we provide means for a simplified work flow in clinical settings when time is critical. Imaging tasks such as controlled quick sweep and automatic instrument tracking require complex coordination among the joints, which are difficult or impossible for manual control but can be achieved with high precision by our robotic system.

The system presented in the thesis provides a good foundation for future robotic controlled catheter systems. The next step is to have the system tested under *in vivo* conditions for comprehensive evaluations of the engineering capabilities developed.

In reconstructing 3D volumes from the 2D images acquired through automated sweeping, ‘virtual slice’ technique is used to fill the gaps during interpolation. Other interpolation method such as natural neighbors could also be explored. Natural neighbors interpolation operates locally and does not infer trends or produce extrema that are not present in the data [84].

We have developed the control module for ICE catheter sweeping from multiple views. As discussed in Chapter 3, multiview imaging and compounding are effective in mitigating signal dropout resulted from occlusion of underlying tissue by structures that reflect or absorb the acoustic signal. The multiview compounding method presented in Chapter 3 can be readily integrated with the robotic system and applied to multiview sweeping for improved target visualization.

The current user interface, although functional and informative for system implementation and debugging, as the system involves, a user interface that has a strong clinical focus should be developed and tested with user studies. For example, the system could keep track of no-fly zones for safety and alert the clinicians when the instrument is approaching a dangerous zone. Automated surface and instrument segmentation methods could also be developed for integration with the current EP suites or augmented virtual reality

environment.

Besides cardiac procedures, the visualization technique developed here could also be applied to procedures in a variety of fluid-filled and solid organs. For example, robotic US catheters could be used during the incremental resection process of transurethral resection of bladder tumors (TURBT), for augmenting endoscopic visualization with high-resolution US panoramic views of subsurface tumor margins. Similarly, in laparoscopic and robotic partial nephrectomy, it can be introduced to the renal pelvis via the ureter for tumor margin detection [27, 45, 51, 69]. It is hypothesized that continuous visualization of instrument-tissue contact and tumor margins can improve work flow and decrease complications across a range of resection procedures.

5.3 Future of Procedure Guidance

Seamless integration of multidisciplinary technologies, such as robotics, telecommunication, virtual reality, play an important role in the future of procedure guidance. Robotic technologies have shown to improve the control in the operating room environment and facilitate the performance of complex procedures [79]. Intra-operative imaging, such as real-time 3DUS provides patient-specific information and offers surgeons the opportunity to make better treatment decisions when time is of essence. Virtual reality provides a means for spatial, temporal and tactical immersion. The culmination of the interdisciplinary efforts have the potential to greatly improve clinical work flow and bring better patient outcome.

Bibliography

- [1] AcuNav. Biosense webster. <https://www.biosensewebster.com/acunav.php>, March 2014.
- [2] Amigo. Catheter robotics. <http://catheterrobotics.com/rcs-main.htm>, March 2014.
- [3] K. V. Arom, T. F. Flavin, R. W. Emery, V. R. Kshetry, P. A. Janey, and R. J. Petersen. Safety and efficacy of off-pump coronary artery bypass grafting. *The Annals of Thoracic Surgery*, 69(3):704–710, 2000.
- [4] M. Arruda, H. Mlcochova, S. K. Prasad, F. Kilicaslan, W. Saliba, D. Patel, T. Fahmy, L. S. Morales, R. Schweikert, and D. Martin. Electrical isolation of the superior vena cava: an adjunctive strategy to pulmonary vein antrum isolation improving the outcome of af ablation. *Journal of cardiovascular electrophysiology*, 18(12):1261–1266, 2007.
- [5] Artisan. Hansen medical. <http://www.hansenmedical.com>, March 2014.
- [6] C. L. Bajaj, V. Pascucci, and D. R. Schikore. The contour spectrum. In *Proceedings of the 8th conference on Visualization'97*, pages 167–ff. IEEE Computer Society Press, 1997.
- [7] C. Barry, C. Allott, N. John, P. Mellor, P. Arundel, D. Thomson, and J. Waterton. Three-dimensional freehand ultrasound: image reconstruction and volume analysis. *Ultrasound in medicine & biology*, 23(8):1209–1224, 1997.
- [8] J. Bigun, G. H. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 13(8):775–790, 1991.
- [9] F. Boccalandro, E. Baptista, A. Muench, C. Carter, and R. W. Smalling. Comparison of intracardiac echocardiography versus transesophageal echocardiography guidance for percutaneous transcatheter closure of atrial septal defect. *The American Journal of Cardiology*, 93(4):437–440, 2004.
- [10] E. M. Boctor, G. Fichtinger, R. H. Taylor, and M. A. Choti. Tracked 3d ultrasound in radio-frequency liver ablation. In *Medical Imaging 2003*, pages 174–182. International Society for Optics and Photonics, 2003.
- [11] E. M. Boctor, A. Jain, M. A. Choti, R. H. Taylor, and G. Fichtinger. Rapid calibration method for registration and 3d tracking of ultrasound images using spatial localizer. In *Medical Imaging 2003*, pages 521–532. International Society for Optics and Photonics, 2003.

- [12] W. D. Boyd, N. D. Desai, D. F. D. Rizzo, R. J. Novick, F. N. McKenzie, and A. H. Menkis. Off-pump surgery decreases postoperative complications and resource utilization in the elderly. *The Annals of Thoracic Surgery*, 68(4):1490–1493, 1999.
- [13] L. J. Brattain and R. D. Howe. *Real-time 4d ultrasound mosaicing and visualization*, pages 105–112. Medical Image Computing and Computer-Assisted Intervention MICCAI 2011. Springer, 2011.
- [14] L. J. Brattain, N. V. Vasilyev, and R. D. Howe. *Enabling 3d ultrasound procedure guidance through enhanced visualization*, pages 115–124. Information Processing in Computer-Assisted Interventions. Springer, 2012.
- [15] S. Brekke, S. I. Rabben, A. Staylen, A. Haugen, G. U. Haugen, E. N. Steen, and H. Torp. Volume stitching in three-dimensional echocardiography: distortion analysis and extension to real time. *Ultrasound in medicine & biology*, 33(5):782–796, 2007.
- [16] T. Brox, R. V. D. Boomgaard, F. Lauze, J. V. D. Weijer, J. Weickert, P. Mrafazek, and P. Kornprobst. *Adaptive structure tensors and their applications*, pages 17–47. Visualization and processing of tensor fields. Springer, 2006.
- [17] T. Brox and J. Weickert. *Nonlinear matrix diffusion for optic flow estimation*, pages 446–453. Pattern Recognition. Springer, 2002.
- [18] T. Brox, J. Weickert, B. Burgeth, and P. Mrafazek. Nonlinear structure tensors. *Image and Vision Computing*, 24(1):41–55, 2006.
- [19] J. Bucerius, J. F. Gummert, M. A. Borger, T. Walther, N. Doll, J. F. Onnasch, S. Metz, V. Falk, and F. W. Mohr. Stroke after cardiac surgery: a risk factor analysis of 16,184 consecutive adult patients. *The Annals of Thoracic Surgery*, 75(2):472–478, 2003.
- [20] J. Butler, D. Parker, R. Pillai, S. Westaby, D. J. Shale, and G. M. Rocker. Effect of cardiopulmonary bypass on systemic release of neutrophil elastase and tumor necrosis factor. *The Journal of thoracic and cardiovascular surgery*, 105(1):25–30, Jan 1993. LR: 20071115; JID: 0376343; 0 (Tumor Necrosis Factor-alpha); 9007-41-4 (C-Reactive Protein); EC 3.4.21.36 (Pancreatic Elastase); EC 3.4.21.37 (Leukocyte Elastase); CIN: J Thorac Cardiovasc Surg. 1994 Apr;107(4):1160-1. PMID: 8159042; EIN: J Thorac Cardiovasc Surg 1993 Jun;105(6):1056; ppublish.
- [21] J. W. Cannon, J. A. Stoll, I. S. Salgo, H. B. Knowles, R. D. Howe, P. E. Dupont, G. R. Marx, and P. J. del Nido. Real-time three-dimensional ultrasound for guiding surgical tasks. *Computer aided surgery*, 8(2):82–90, 2003.
- [22] CardiacMapping. St. jude medical. <http://professional.sjm.com/products/ep/mapping-visualization/cardiac-mapping-system>, March 2014.
- [23] CartoSound. Biosense webster. <http://www.biosensewebster.com/cartosound.php>, March 2014.
- [24] S. S. Chugh, J. L. Blackshear, W.-K. Shen, S. C. Hammill, and B. J. Gersh. Epidemiology and natural history of atrial fibrillation: clinical implications. *Journal of the American College of Cardiology*, 37(2):371–378, 2001.

- [25] CorPath. Corindus. <http://www.corindus.com/physician/corpath-clinical-research>, March 2014.
- [26] P. Coupe, P. Hellier, X. Morandi, and C. Barillot. Probe trajectory interpolation for 3d reconstruction of freehand ultrasound. *Medical image analysis*, 11(6):604–615, 2007.
- [27] A. Doerfler, Y. Cerantola, J.-Y. Meuwly, B. Lhermitte, H. Bensadoun, and P. Jichlinski. Ex vivo ultrasound control of resection margins during partial nephrectomy. *The Journal of urology*, 186(6):2188–2193, 2011.
- [28] EPOCH. Stereotaxis. <http://www.stereotaxis.com/physicians/the-lab/>, March 2014.
- [29] S. Ernst, F. Ouyang, F. Löber, M. Antz, and K.-H. Kuck. Catheter-induced linear lesions in the left atrium in patients with atrial fibrillation: an electroanatomic study. *Journal of the American College of Cardiology*, 42(7):1271–1282, 2003.
- [30] R. H. Falk. Diagnosis and management of the cardiac amyloidoses. *Circulation*, 112(13):2047–2060, Sep 27 2005. JID: 0147763; 0 (Amyloid); RF: 123; ppublish.
- [31] A. Fenster, D. B. Downey, and H. N. Cardinal. Three-dimensional ultrasound imaging. *Physics in Medicine and Biology*, 46(5):R67, 2001.
- [32] W. Forstner and E. Gulch. A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *Proc. ISPRS intercommission conference on fast processing of photogrammetric data*, pages 281–305, 1987.
- [33] A. H. Gee, G. M. Treece, R. W. Prager, C. J. Cash, and L. Berman. Rapid registration for wide field of view freehand three-dimensional ultrasound. *Medical Imaging, IEEE Transactions on*, 22(11):1344–1357, 2003.
- [34] B. Gersak and Z. Sutlic. Aortic and mitral valve surgery on the beating heart is lowering cardiopulmonary bypass and aortic cross clamp time. In *Heart Surg Forum*, volume 5, pages 182–186, 2002.
- [35] V. Grau, H. Becher, and J. A. Noble. Registration of multiview real-time 3-d echocardiographic sequences. *Medical Imaging, IEEE Transactions on*, 26(9):1154–1165, 2007.
- [36] V. Grau and J. A. Noble. *Adaptive multiscale ultrasound compounding using phase information*, pages 589–596. Medical Image Computing and Computer-Assisted Intervention MICCAI 2005. Springer, 2005.
- [37] S. Y. Ho, J. A. Cabrera, V. H. Tran, J. Farre, R. H. Anderson, and D. Sanchez-Quintana. Architecture of the pulmonary veins: relevance to radiofrequency ablation. *Heart (British Cardiac Society)*, 86(3):265–270, Sep 2001. LR: 20130915; JID: 9602087; OID: NLM: PMC1729909; ppublish.
- [38] D. Honigmann, J. Ruisz, and C. Haider. Adaptive design of a global opacity transfer function for direct volume rendering of ultrasound data. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, page 64. IEEE Computer Society, 2003.

- [39] F. Hottier and A. C. Billon. *3D echography: status and perspective*, pages 21–41. 3D imaging in medicine. Springer, 1990.
- [40] R. J. Housden, A. Arujuna, Y. Ma, N. Nijhof, G. Gijsbers, R. Bullens, M. O’Neill, M. Cooklin, C. A. Rinaldi, and J. Gill. *Evaluation of a real-time hybrid three-dimensional echo and x-ray imaging system for guidance of cardiac catheterisation procedures*, pages 25–32. Medical Image Computing and Computer-Assisted Intervention, MICCAI 2012. Springer, 2012.
- [41] J. Huang, J. K. Triedman, N. V. Vasilyev, Y. Suematsu, R. O. Cleveland, and P. E. Dupont. Imaging artifacts of medical instruments in ultrasound-guided interventions. *Journal of Ultrasound in Medicine*, 26(10):1303–1322, 2007.
- [42] X. Huang, N. A. Hill, J. Ren, G. Guiraudon, D. Boughner, and T. M. Peters. *Dynamic 3D ultrasound and MR image registration of the beating heart*, pages 171–178. Medical Image Computing and Computer-Assisted Intervention MICCAI 2005. Springer, 2005.
- [43] B. Jafoanne and H. Haussecker. *Computer vision and applications: a guide for students and practitioners*. Academic Press, 2000.
- [44] B. Jähne. *Spatio-temporal image processing: theory and scientific applications*, volume 751. Springer, 1993.
- [45] B. F. Kaczmarek and C. G. Rogers. *Robotic Partial Nephrectomy: Basic Principles and Techniques*, pages 63–75. Robotic Renal Surgery. Springer, 2013.
- [46] D. T. Kettler, R. D. Plowes, P. M. Novotny, N. V. Vasilyev, P. J. del Nido, and R. D. Howe. An active motion compensation instrument for beating heart mitral valve surgery. In *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, pages 1290–1295. IEEE, 2007.
- [47] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Moller. Curvature-based transfer functions for direct volume rendering: Methods and applications. In *Visualization, 2003. VIS 2003. IEEE*, pages 513–520. IEEE, 2003.
- [48] O. Kutter, W. Wein, and N. Navab. *Multi-modal registration based ultrasound mosaicing*, pages 763–770. Medical Image Computing and Computer-Assisted Intervention, MICCAI 2009. Springer, 2009.
- [49] S. E. Lakhan, A. Kaplan, C. Laird, and Y. Leiter. The interventionalism of medicine: interventional radiology, cardiology, and neuroradiology. *Int Arch Med*, 2(1):27, 2009.
- [50] T. Lange, S. Eulenstein, M. Hafnerbein, H. Lamecker, and P.-M. Schlag. *Augmenting intraoperative 3D ultrasound with preoperative models for navigation in liver surgery*, pages 534–541. Medical Image Computing and Computer-Assisted Intervention, MICCAI 2004. Springer, 2004.
- [51] P. Liang, Y. Wang, D. Zhang, X. Yu, Y. Gao, and X. Ni. Ultrasound guided percutaneous microwave ablation for small renal cancer: initial experience. *The Journal of urology*, 180(3):844–848, 2008.

- [52] P. M. Loschak, L. J. Brattain, and R. D. Howe. Automated pointing of cardiac imaging catheters. In *Proc. IEEE Int. Conf. Robotics and Automation*, 2013. NIHMS485450.
- [53] M. J. Mack. Pro: beating-heart surgery for coronary revascularization: is it the most important development since the introduction of the heart-lung machine? *The Annals of Thoracic Surgery*, 70(5):1774–1778, 2000.
- [54] N. F. Marrouche, D. O. Martin, O. Wazni, A. M. Gillinov, A. Klein, M. Bhargava, E. Saad, D. Bash, H. Yamada, and W. Jaber. Phased-array intracardiac echocardiography monitoring during pulmonary vein isolation in patients with atrial fibrillation impact on outcome and complications. *Circulation*, 107(21):2710–2716, 2003.
- [55] H. A. McCann, J. C. Sharp, T. M. Kinter, C. N. McEwan, C. Barillot, and J. F. Greenleaf. Multidimensional ultrasonic imaging for cardiology. *Proceedings of the IEEE*, 76(9):1063–1073, 1988.
- [56] G. Medioni, M.-S. Lee, and C.-K. Tang. *A computational framework for segmentation and grouping*. Elsevier, 2000.
- [57] V. Mor-Avi, L. Sugeng, and R. M. Lang. Three-dimensional adult echocardiography: where the hidden dimension helps. *Current cardiology reports*, 10(3):218–225, 2008.
- [58] M. Moradi, P. Abolmaesoumi, and P. Mousavi. Deformable registration using scale space keypoints. In *Medical Imaging*, pages 61442G–61442G–8. International Society for Optics and Photonics, 2006.
- [59] H. Morehouse, H. P. Thaker, and C. Persaud. Addition of metamucil to gelatin for a realistic breast biopsy phantom. *Journal of ultrasound in medicine*, 26(8):1123–1126, 2007. pmid:17646379.
- [60] J. Mung, F. Vignon, and A. Jain. *A non-disruptive technology for robust 3D tool tracking for ultrasound-guided interventions*, pages 153–160. Medical Image Computing and Computer-Assisted Intervention, MICCAI 2011. Springer, 2011.
- [61] J. M. Murkin, W. D. Boyd, S. Ganapathy, S. J. Adams, and R. C. Peterson. Beating heart surgery: why expect less central nervous system morbidity? *The Annals of Thoracic Surgery*, 68(4):1498–1501, 1999.
- [62] M. Nakamoto, Y. Sato, M. Miyamoto, Y. Nakamjima, K. Konishi, M. Shimada, M. Hashizume, and S. Tamura. *3D ultrasound system using a magneto-optic hybrid tracker for augmented reality visualization in laparoscopic liver surgery*, pages 148–155. Medical Image Computing and Computer-Assisted Intervention, MICCAI 2002. Springer, 2002.
- [63] T. R. Nelson and D. H. Pretorius. Three-dimensional ultrasound imaging. *Ultrasound in medicine and biology*, 24(9):1243–1270, 1998.
- [64] D. Ni, Y. P. Chui, Y. Qu, X. Yang, J. Qin, T.-T. Wong, S. S. Ho, and P. A. Heng. Reconstruction of volumetric ultrasound panorama based on improved 3d sift. *Computerized Medical Imaging and Graphics*, 33(7):559–566, 2009.

- [65] P. M. Novotny, S. K. Jacobsen, N. V. Vasilyev, D. T. Kettler, I. S. Salgo, P. E. Dupont, P. J. D. Nido, and R. D. Howe. 3d ultrasound in robotic surgery: performance evaluation with stereo displays. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 2(3):279–285, 2006.
- [66] P. M. Novotny, J. A. Stoll, N. V. Vasilyev, P. J. D. Nido, P. E. Dupont, T. E. Zickler, and R. D. Howe. Gpu based real-time instrument tracking with three-dimensional ultrasound. *Medical image analysis*, 11(5):458–464, 2007.
- [67] D. F. Pace, D. G. Gobbi, C. Wedlake, J. Gumprecht, J. Boisvert, J. Tokuda, N. Hata, and T. M. Peters. An open-source real-time ultrasound reconstruction system for four-dimensional imaging of moving organs. *Insight Journal*, pages 1–8, 2009.
- [68] D. L. Packer, P. Keelan, T. M. Munger, J. F. Breen, S. Asirvatham, L. A. Peterson, K. H. Monahan, M. F. Hauser, K. Chandrasekaran, L. J. Sinak, and D. R. H. Jr. Clinical presentation, investigation, and management of pulmonary vein stenosis complicating ablation for atrial fibrillation. *Circulation*, 111(5):546–554, Feb 8 2005. JID: 0147763; ppublish.
- [69] C. P. Pavlovich, M. M. Walther, P. L. Choyke, S. E. Pautler, R. Chang, W. Linehan, and B. J. Wood. Percutaneous radio frequency ablation of small renal tumors: initial results. *The Journal of urology*, 167(1):10–15, 2002.
- [70] H. Pfister, J. Hardenbergh, J. Knittel, H. Lauer, and L. Seiler. The volumepro real-time ray-casting system. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 251–260. ACM Press/Addison-Wesley Publishing Co., 1999.
- [71] S. Pieper, M. Halle, and R. Kikinis. 3d slicer. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 632–635. IEEE, 2004.
- [72] T. C. Poon and R. N. Rohling. Three-dimensional extended field-of-view ultrasound. *Ultrasound in medicine & biology*, 32(3):357–369, 2006.
- [73] B. C. Porter, D. J. Rubens, J. G. Strang, J. Smith, S. Totterman, and K. J. Parker. Three-dimensional registration and fusion of ultrasound and mri using major vessels as fiducial markers. *Medical Imaging, IEEE Transactions on*, 20(4):354–359, 2001.
- [74] R. Prager, U. Ijaz, A. Gee, and G. Treece. Three-dimensional ultrasound imaging. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 224(2):193–223, 2010.
- [75] R. W. Prager, A. Gee, and L. Berman. Stradx: real-time acquisition and visualization of freehand three-dimensional ultrasound. *Medical image analysis*, 3(2):129–140, 1999.
- [76] J. Puskas, W. Williams, P. Duke, J. Staples, K. Glas, J. Marshall, M. Leimbach, P. Huber, S. Garas, and B. Sammons. Off-pump coronary artery bypass grafting provides complete revascularization with reduced myocardial injury, transfusion requirements, and length of stay: a prospective randomized comparison of two hundred unselected

patients undergoing off-pump versus conventional coronary artery bypass grafting. *The Journal of thoracic and cardiovascular surgery*, 125(4):797–808, 2003.

- [77] K. Rajpoot, V. Grau, J. A. Noble, C. Szmigielski, and H. Becher. Multiview fusion 3-d echocardiography: Improving the information and quality of real-time 3-d echocardiography. *Ultrasound in medicine & biology*, 37(7):1056–1072, 7 2011.
- [78] A. J. Rastan, H. B. Bittner, J. F. Gummert, T. Walther, C. V. Schewick, E. Girdauskas, and F. W. Mohr. On-pump beating heart versus off-pump coronary artery bypass surgery-evidence of pump-induced myocardial injury. *European journal of cardio-thoracic surgery : official journal of the European Association for Cardio-thoracic Surgery*, 27(6):1057–1064, Jun 2005.
- [79] M. Reijnen, C. J. Zeebregts, and W. Meijerink. Future of operating rooms. *Surgical technology international*, 14:21–27, 2004.
- [80] G. W. Roach, M. Kanchuger, C. M. Mangano, M. Newman, N. Nussmeier, R. Wolman, A. Aggarwal, K. Marschall, S. H. Graham, and C. Ley. Adverse cerebral outcomes after coronary bypass surgery. *New England Journal of Medicine*, 335(25):1857–1864, 1996.
- [81] R. Rohling, A. Gee, and L. Berman. A comparison of freehand three-dimensional ultrasound reconstruction techniques. *Medical image analysis*, 3(4):339–359, 1999.
- [82] O. Rourke. *Hurst's the heart*. Tata McGraw-Hill Education, 2006.
- [83] R. J. Schneider, D. P. Perrin, N. V. Vasilyev, G. R. Marx, P. J. del Nido, and R. D. Howe. Real-time image-based rigid registration of three-dimensional ultrasound. *Medical image analysis*, 16(2):402–414, 2012.
- [84] R. Sibson. A brief description of natural neighbour interpolation. *Interpreting multivariate data*, 21, 1981.
- [85] Slicer. 3d slicer. <http://www.slicer.org/>, March 2014.
- [86] O. V. Solberg, F. Lindseth, H. Torp, R. E. Blake, and T. A. N. Hernes. Freehand 3d ultrasound reconstruction algorithms, a review. *Ultrasound in medicine & biology*, 33(7):991–1009, 2007.
- [87] P. Soler, O. Gerard, P. Allain, E. Saloux, E. Angelini, and I. Bloch. Comparison of fusion techniques for 3d+t echocardiography acquisitions from different acoustic windows. In *Computers in Cardiology, 2005*, pages 141–144. IEEE, 2005.
- [88] trakSTAR. Ascension. <http://www.ascension-tech.com/medical/trakSTAR.php>, March 2014.
- [89] G. M. Treece, R. W. Prager, A. H. Gee, C. J. Cash, and L. Berman. Grey-scale gating for freehand 3d ultrasound. In *Biomedical Imaging, 2002. Proceedings. 2002 IEEE International Symposium on*, pages 993–996. IEEE, 2002.

- [90] J. W. Trobaugh, D. J. Trobaugh, and W. D. Richard. Three-dimensional imaging with stereotactic ultrasonography. *Computerized Medical Imaging and Graphics*, 18(5):315–323, 1994.
- [91] P. Valensi, L. Lorgis, and Y. Cottin. Prevalence, incidence, predictive factors and prognosis of silent myocardial infarction: a review of the literature. *Archives of cardiovascular diseases*, 104(3):178–188, 2011.
- [92] Z. Wang, G. Slabaugh, G. Unal, and T. Fang. Registration of ultrasound images using an information-theoretic feature detector. In *Biomedical Imaging: From Nano to Macro, 2007. ISBI 2007. 4th IEEE International Symposium on*, pages 736–739. IEEE, 2007.
- [93] J. Weickert and H. Scharr. A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance. *Journal of Visual Communication and Image Representation*, 13(1):103–118, 2002.
- [94] W. Wein, E. Camus, M. John, M. Diallo, C. Duong, A. Al-Ahmad, R. Fahrig, A. Khamene, and C. Xu. *Towards guidance of electrophysiological procedures with real-time 3D intracardiac echocardiography fusion to C-arm CT*, pages 9–16. Medical Image Computing and Computer-assisted Intervention, MICCAI 2009. Springer, 2009.
- [95] M. Wright, E. Harks, S. Deladi, F. Suijver, M. Barley, A. van Dusschoten, S. Fokkenrood, F. Zuo, F. Sacher, and M. Hocini. Real-time lesion assessment using a novel combined ultrasound and radiofrequency ablation catheter. *Heart Rhythm*, 8(2):304–312, 2011.
- [96] S. Yagel, S. Cohen, I. Shapiro, and D. Valsky. 3d and 4d ultrasound in fetal cardiac scanning: a new look at the fetal heart. *Ultrasound in obstetrics & gynecology*, 29(1):81–95, 2007.
- [97] C. Yao, J. M. Simpson, C. H. Jansen, A. P. King, and G. P. Penney. Spatial compounding of large sets of 3d echocardiography images. In *SPIE Medical Imaging*, pages 726515–726515–8. International Society for Optics and Photonics, 2009.
- [98] H. Zhang, F. Banovac, A. White, and K. Cleary. Freehand 3d ultrasound calibration using an electromagnetically tracked needle. In *Medical Imaging*, pages 61412M–61412M–9. International Society for Optics and Photonics, 2006.
- [99] H. Zhong, T. Kanade, and D. Schwartzman. *Image Thickness Correction for Navigation with 3D Intra-cardiac Ultrasound Catheter*, pages 485–492. Medical Image Computing and Computer-Assisted Intervention, MICCAI 2008. Springer, 2008.
- [100] X. Zhuang, C. Yao, Y. Ma, D. Hawkes, G. Penney, and S. Ourselin. Registration-based propagation for whole heart segmentation from compounded 3d echocardiography. In *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*, pages 1093–1096. IEEE, 2010.