



Robotic Acquisition and Enhancement of Ultrasound Images

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:40050070>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Robotic Acquisition and Enhancement of Ultrasound Images

A dissertation presented

by

Alperen Degirmenci

to

the John A. Paulson 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

April 2018

© 2018 Alperen Degirmenci

All rights reserved.

Dissertation Advisor:
Professor Robert D. Howe

Author:
Alperen Degirmenci

Robotic Acquisition and Enhancement of Ultrasound Images

Abstract

Ultrasound is one of the cheapest, safest, and most ubiquitously found medical imaging technology. However ultrasound imaging is currently underutilized in image-guided interventions due to limited image quality and the lack of automated methods to provide image-based guidance. This work describes methods to improve ultrasound image quality and enables new procedure guidance methods through robotic steering of ultrasound imaging catheters.

The first two chapters present a robotic system that was designed and built to autonomously steer ultrasound-imaging catheters. This system enables autonomous tracking of instruments inside the heart, and can generate 4D (3D + time) ultrasound volumes from the 2D images acquired from the catheter. This work improves visualization of cardiac structures during interventions, potentially leading to a decrease in the use of X-ray fluoroscopy.

Next two chapters present methods for improving ultrasound image quality. One method increases interpolation accuracy during ultrasound scanline conversion using Gaussian process regression. The second method extends the dynamic range of ultrasound images through high dynamic range imaging methods, similar to how smartphones and digital cameras can combine pictures taken at different exposures to improve the visibility of different structures.

The final chapter describes a platform that was developed to characterize the dynamic response of human muscles and internal organs, such as the stomach. The data collected in these studies can help improve the fidelity of finite element models of the human body.

Contents

Abstract	iii
Acknowledgments	xii
1 Introduction	1
1.1 Motivation	1
1.2 Background	3
1.2.1 A Brief History of Medical Imaging	3
1.2.2 Principles of Ultrasound Imaging	5
1.2.3 Use of Ultrasound in Cardiac Interventions	6
1.2.4 Use of Ultrasound in Robotic Interventions	7
1.2.5 Use of Ultrasound in Biomechanics Research	7
1.3 Thesis Contributions	7
1.3.1 Autonomous Steering of Flexible Instruments	8
1.3.2 Real-Time 3D Ultrasound Image Generation and Visualization for Procedure Guidance	8
1.3.3 Improving Ultrasound Image Quality Through Bayesian Methods and Enhanced Dynamic Range	8
1.3.4 Ultrasound Image Acquisition and Processing for Biomechanics Studies	9
1.4 Thesis Overview	9
2 Autonomous Steering of Ultrasound-Imaging Catheters	10
2.1 Background	11
2.1.1 Clinical Problem: Arrhythmia	11
2.1.2 Current Practice	11
2.1.3 Ultrasound Imaging Catheters	12
2.1.4 Robotic Catheters and Image Guidance	12
2.1.5 Goals and Challenges	13
2.2 Catheter Kinematics and Sensing	15
2.2.1 Constrained Catheter Motion	15
2.2.2 Unconstrained Catheter Motion	15
2.2.3 Kinematics	18

2.2.4	Forward Kinematics	20
2.2.5	Inverse Kinematics	20
2.2.6	Joint Mapping	21
2.3	Control Strategy	21
2.3.1	Disturbance Rejection	24
2.3.2	Respiration Compensation	25
2.4	Robot and System Design	28
2.5	Experiments	31
2.5.1	Bench Top Motion	31
2.5.2	Initial in vivo studies	33
2.5.3	Instrument Tracking with Respiration Compensation	36
2.6	Discussion	38
3	Enabling Real-Time Ultrasound Reconstruction for Advanced Procedure Guidance	40
3.1	Introduction	41
3.2	Software Framework	42
3.3	Ultrasound-based Procedure Guidance	46
3.3.1	Ultrasound Volume Reconstruction	46
3.3.2	Rendering	49
3.4	Results	50
3.5	Discussion	53
4	Gaussian Process Regression for Ultrasound Scanline Conversion	55
4.1	Background	56
4.1.1	Previous Work on Improving Ultrasound Image Quality	56
4.1.2	Statistical Models for Ultrasound Signals	57
4.2	Methods	59
4.2.1	Gaussian Process Regression	59
4.2.2	Scanline Conversion - Parallel and Divergent Scanlines	60
4.2.3	Validation - Measuring Interpolation Accuracy	61
4.2.4	Covariance Functions	61
4.2.5	Finding the Optimal Length Scale	64
4.2.6	Reducing Computational Cost - Patched \mathcal{GP} Regression	65
4.3	Experiments and Results	66
4.3.1	Parallel Scanlines	68
4.3.2	Diverging Scanlines	70
4.4	Discussion	72

5	High Dynamic Range Ultrasound Imaging	74
5.1	Motivation	75
5.2	Methods	76
5.2.1	Ultrasound Image Formation	76
5.2.2	High Dynamic Range Formulation	77
5.2.3	Tone Mapping – Display of HDR-US Images	79
5.2.4	Quality Metrics	81
5.3	Experiments and Results	82
5.3.1	<i>Ex Vivo</i>	82
5.3.2	<i>In Vivo</i>	85
5.3.3	Number of Images Required for HDR-US Computations	87
5.4	Discussion	88
5.5	Conclusions	90
6	Quantifying Tissue Dynamics In Biomechanics Studies Through Ultrasound Imaging	91
6.1	Motivation	92
6.2	Experimental Setup	93
6.3	Results	96
6.3.1	Frequency Response	96
6.3.2	3D Reconstruction of the Stomach	97
6.4	Discussion	97
6.5	Conclusion	99
7	Conclusions and Future Work	100
7.1	Contributions	100
7.1.1	Autonomous Steering of Flexible Instruments	100
7.1.2	Real-Time 3D Ultrasound Image Generation and Visualization for Procedure Guidance	101
7.1.3	Improving Ultrasound Image Quality Through Bayesian Methods and Enhanced Dynamic Range	101
7.1.4	Ultrasound Image Acquisition and Processing for Biomechanics Studies	102
7.2	Future Directions	102
7.2.1	Applications to Other Organs	103
7.2.2	Catheter Steering	104
7.2.3	4D Reconstruction, Strain Imaging, Image-Guided Steering	104
7.2.4	GANs for Scanline Conversion	105
7.2.5	Compensation for Brain Shift	105
7.2.6	3D Printed Ultrasound Phantoms	105

7.2.7	Future of Ultrasound	107
	References	108

List of Tables

3.1	Execution time for various stages of the reconstruction pipeline with different input and output sizes.	53
5.1	Acoustic properties of relevant media	77

List of Figures

1.1	The anatomy of an ultrasound probe.	5
1.2	Ultrasound wave propagation phenomena.	6
2.1	An overview of procedure guidance in cardiac interventions	12
2.2	Details of the ICE catheter	13
2.3	Uncontrollable catheter shaft motion	16
2.4	Induced base motion of the ICE catheter	16
2.5	ICE catheter deadzone	17
2.6	Parameter spaces	18
2.7	ICE catheter D-H parameters	19
2.8	Controller diagram	22
2.9	The required configuration space roll adjustments	23
2.10	Disturbance analysis	25
2.11	The prototype system for autonomously steering ICE catheters.	26
2.12	Data flow in respiration motion compensation.	27
2.13	Cross-sectional view showing knob interactions with the catheter	29
2.14	Device diagram showing the communication layers between the physical devices and the software framework components.	30
2.15	Benchtop navigation study	31
2.16	Bench top tests demonstrating the ability of the robotic system to maintain the catheter tip at the target pose despite disturbance step and target step inputs.	32
2.17	Introducer seal is bypassed using a brass tube filled with silicone.	34
2.18	<i>In vivo</i> tests demonstrating the ability of the robotic system to maintain the catheter tip at the target pose despite disturbance step, disturbance ramp, and target step inputs.	35
2.19	In vivo respiration compensation trajectory.	36
2.20	In vivo respiration compensation error histogram.	37
2.21	Automatically pointing the ultrasound imager at a manually-steered target catheter in vivo	37
2.22	In vivo ablation of the left ventricular wall.	38

3.1	Volume rendering graphical user interface	42
3.2	Simplified class diagram	44
3.3	3D reconstruction pipeline	47
3.4	Catheter tip position and imaging plane angle during image acquisition for 4D reconstruction	49
3.5	Fast intraoperative 3D reconstruction result	50
3.6	Intraoperative 4D reconstruction results	52
4.1	<i>In vivo</i> parallel scanline data.	57
4.2	The PDF, CDF, and 1-CDF of the K-distribution.	58
4.3	Polar to Cartesian mapping in scan conversion	60
4.4	Leave-N-out studies	62
4.5	Incorrect choice of the kernel length can lead to reconstruction errors	64
4.6	Scan converted B-mode images of the parallel scanline dataset. Leftmost column shows the ground truth data with 256 scanlines. Rightmost column shows the normalized variance matrix generated during \mathcal{GP} regression, where darker regions indicate higher confidence in the estimates.	67
4.7	Interpolator performance for scan conversion of parallel scanlines	68
4.8	Scan converted B-mode images of the divergent scanline dataset	69
4.9	Interpolator performance for scan conversion of divergent scanlines	70
4.10	Scan converted B-mode images of the divergent scanline dataset	71
5.1	Experimental setup	76
5.2	Response function	78
5.3	<i>Ex vivo</i> ultrasound images	82
5.4	Original, HDR, and tone-mapped <i>ex vivo</i> ultrasound images	83
5.5	Output of five different tone mapping algorithms	84
5.6	Similarity contours for five TMOs	84
5.7	HDR-US image of a human lower leg acquired in the transverse plane	86
5.8	Leave- p -out cross-validation results for <i>ex vivo</i> and <i>in vivo</i> studies	88
6.1	System setup.	93
6.2	Coordinate frames of targets of interest	94
6.3	Velocity fields extracted from ultrasound images	95
6.4	Ultrasound images of the femur and the stomach were acquired using a TEE probe. The tip of the TEE probe is shown at the top.	96
6.5	Frequency response of internal and external anatomy to sinusoidal inputs	97
6.6	3D reconstruction of the stomach at 4 Hz sinusoidal input	98

7.1	Ex vivo ultrasound study of a kidney.	103
7.2	3D printed ultrasound phantoms	106

Acknowledgments

First of all, I would like to thank my advisor Professor Robert Howe for his continuous support and guidance. His focus on engineering meaningful solutions to real-world problems has been inspirational and contagious. He provided me with a lot of intellectual freedom, which gave me the chance to work on (too) many different areas and helped me become an independent researcher. I would also like to thank my committee members Professors Robert Wood and Conor Walsh. Their work and leadership has always been inspiring, and it has been invaluable to get the chance to collaborate with and learn from them.

Next I would like to thank Dr. Douglas Perrin and Dr. Paul Loschak for being my partners in crime through most of this work. Additionally, our clinical collaborators Cory Tschabrunn and Dr. Elad Anter have been instrumental in translating our autonomous catheter steering system from the benchtop to *in vivo* trials.

I would like to thank my undergraduate advisors Professors Allison Okamura, Lester Su, Louis Whitcomb, and Russell Taylor for inspiring, encouraging, and preparing me to pursue my graduate studies.

The path to earning a PhD is quite strenuous and perhaps impossible without fellow graduate students to share the journey with. I feel lucky to have had the chance to work with not just coworkers but also friends at 60 Oxford and at Harvard in general. From the BioRobotics Lab I would like to thank Yashraj Narang, Qian Wan, Buse Aktaş, Dr. Neil Tenenholtz, Dr. Daisuke Yamada, Dr. Richard Nuckols, Dr. Yaroslav Tenzer, Dr. Leif Jentoft, Dr. Lael Odhner, Dr. Laura Brattain, Dr. Mohsen Moradi Dalvand, Prof. Frank L. Hammond III, Dr. Peter Hammer, Dr. Samuel B. Kesner, Dr. Robert Schneider, Alexander Hassan, Annabel Imbrie-Moore, and Melissa Majkut. I would also like to thank many members of the MicroRobotics Lab, the BioDesign Lab, the Weaver Lab, the Wyss Institute, the SEAS Teaching Labs, and the SEAS Graduate Council.

I would like to thank countless friends from across the world who always stayed in contact and provided support. Finally, I would like to thank my parents and grandmother for their love and unending support; without their sacrifices, I would not be where I am today.

To my parents, Zehra Tülin and Mehmet Ali Değirmenci

Chapter 1

Introduction

Medical imaging has revolutionized medicine in the past few decades, eliminating guesswork (i.e. “exploratory surgery”), and enabling quick and accurate diagnosis. Image-guidance in diagnosis and interventions greatly augments the senses of vision and touch that physicians have relied on for centuries [1].

The field of medicine is on the brink of a similar breakthrough with the recent advancements in surgical robotics and the integration of medical imaging with automated procedure guidance. The importance of innovation and research in this area is recognized by many companies and agencies. For instance, image-guided interventions is one of the eight program areas (along with ultrasound, X-ray, and MRI) in the Division of Applied Science & Technology (DAST) at the National Institute of Biomedical Imaging and Bioengineering (NIBIB). So far in the clinic, most robotic systems have lacked the ability to make real-time decisions based on the rich information provided by the imaging systems. Combining real-time image guidance with surgical robots will define the future of robotic surgery and medicine in general.

1.1 Motivation

Image guidance is critical, both in conventional and robotic surgery. From tissue biopsy to cardiac ablation, clinicians rely on medical imaging for (1) preoperative planning, (2) intraoperative guidance during procedures, and (3) postoperative evaluation and validation.

Patient safety and procedure outcome therefore heavily rely on the quality and accuracy of medical images. Thus, improving image-guidance has a direct impact on improving health care.

Among all imaging modalities, ultrasound is an ideal choice for image guidance in a number of procedures. Images are acquired in real time at high frame rates, with no radiation exposure, and at much lower expense than CT and MRI. Its portability and low cost make ultrasound imaging especially valuable for use in developing countries, emergency interventions in the field, and for reducing the cost of robotic systems. Therefore enabling new procedures that utilize ultrasound is important.

Cardiac interventions is one area where medical imaging has greatly improved the standard of care, enabling minimally-invasive procedures to be performed. Minimally invasive procedures are safer and cheaper than conventional open surgery, however issues such as access to the surgical site, limited dexterity, lack of accurate and real time instrument position feedback, and disturbances due to physiological motions limit which organ systems are viable for minimally invasive approaches. Flexible instruments, such as catheters, with small diameters can be used to gain access to the surgical site through the vasculature; however controlling these manipulators manually is highly challenging. This issue is addressed in Chapter 2.

Intraoperative imaging is crucial for procedure guidance during many interventions, however the high computational resources required for real-time image processing limit the availability of robust, automated procedure guidance. For example, during cardiac ablation, clinicians rely on X-ray fluoroscopy to determine the locations of their catheters. Cardiac catheters with integrated ultrasound imaging can provide real time image acquisition from within the heart, however the difficulty in controlling these catheters, and the computational limitations of real-time image processing and visualization has so far prevented ultrasound-based image guidance for robotic and manual interventions. This issue is addressed in Chapter 3.

Compared to MRI and CT, ultrasound has much lower image quality, and the images are usually 2D instead of 3D. Interpreting ultrasound images requires years of training. However it is possible to improve the image quality using computational photography and machine

learning, improving the accuracy of diagnoses and the efficacy of image-guided interventions. Chapters 4 and 5 report two novel methods which help to achieve this goal.

Finally, the use of ultrasound can enable new experimental methods in basic research. For example, in motion studies, biomechanics researchers often rely on noisy EMG signals to determine muscle activity, and otherwise do not have tools to characterize the motion of internal organs. Ultrasound imaging can be used in these studies to acquire real-time tissue motion in a non-invasive fashion. A new tissue characterization system is reported in Chapter 6 to demonstrate this capability.

1.2 Background

1.2.1 A Brief History of Medical Imaging

The history of medical imaging is considered to begin with the discovery of the X-ray in 1895 [2]. Thus we gained the ability to peek inside the human body in a non-invasive fashion, which Hippocrates and other physicians since his times must have wished for dearly.

Even though ultrasound imaging was already being used to detect flaws in metals in the 1940's, its wide adoption in medicine was not until late 1960's. I urge the inquisitive reader to peruse through [3] for a fascinating story about how ultrasound imaging came to be adopted in cardiology, and medicine in general.

The first *in vivo* magnetic resonance imaging (MRI) images were acquired from a mouse in 1976 [4], and from a human in 1977 [5]. Back then, MRI was called nuclear magnetic resonance (NMR) imaging. However in 1984 it was recommended that the word 'nuclear' be removed from the term 'NMR imaging' on the basis that it might (falsely) "suggest that radioactivity is associated with the procedure in the minds of the patients undergoing examination" [6].

Since their conception, all three imaging methods have been rapidly improved (e.g. better imaging quality, faster acquisition speed, lower radiation) and today they are widely used in hospitals. However these devices are still not considered cheap: an MRI machine can cost up to \$3 million, a computed tomography (CT) scanner up to \$2 million, and an ultrasound

machine up to \$250,000.

Once a diagnosis is made noninvasively using medical imaging, there still remains the problem of gaining physical access to the surgical site for intervention. The surgeon's goal is to inflict the least amount of damage to healthy tissue when gaining access. Laparoscopic techniques have been developed where small 'ports' are opened in the patient's abdominal wall and long, rigid metallic instruments are inserted through to reach the surgical site. The abdomen is insufflated with carbon dioxide to distend the abdominal wall from the internal organs and provide a workspace where the surgeon can insert to see the working tools. The first laparoscopic cholecystectomy in humans was demonstrated in 1985 [7]. Although diagnostic laparoscopy was well established by 1937, it took another 50 years for operative laparoscopy to be developed and accepted as a minimally-invasive alternative to open surgery. Even though laparoscopic procedures were quite common in gynecology in the 1970s, the rest of the medical field was quite resistant to the adoption of this technique.

Laparoscopic surgery requires extensive training for surgeons to develop the necessary hand-eye coordination. The long, rigid tools used in laparoscopy reduce dexterity, which can make it inconvenient or impossible to perform certain procedures. Robotic laparoscopy systems were developed to alleviate this problem by adding 'wrists' to the rigid instruments, and having surgeons operate the system through a console. The da Vinci system is the most successful of these systems, with 4,271 systems installed worldwide [8].

Another method of achieving minimally-invasive access is through the vasculature. In 1929 Werner Forssmann performed the first human right-cardiac catheterization on himself by repurposing a urinary catheter [9]. He was fired from the hospital where he worked for his experiments. Later in 1956, he shared the Nobel Prize in Medicine with Cournand and Richards "for their discoveries concerning heart catheterization and pathological changes in the circulatory system" [10].

Catheters, which are long, thin, flexible instruments, are used throughout the body to gain minimally-invasive or non-invasive access to organs. Cardiac catheters are used to perform procedures such as angio, stent placement, valve replacement, and ablation, which previously

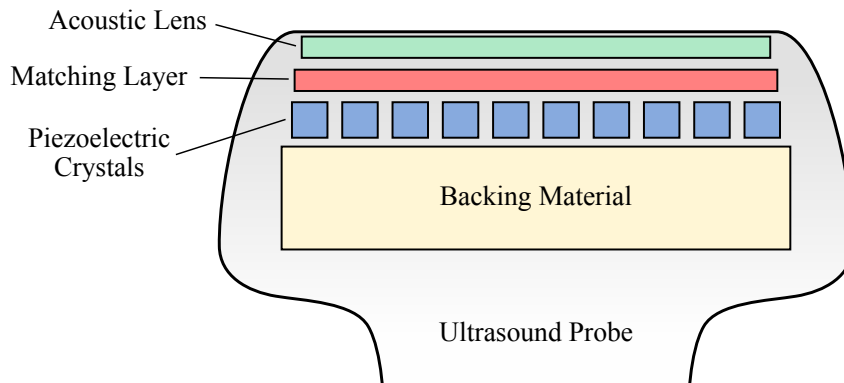


Figure 1.1: *The anatomy of an ultrasound probe.*

required open heart surgery [11]. Two main advantages of these procedures are that they reduce the risk of stroke due to placing the patient on a cardiopulmonary bypass machine, and do not require a thoracotomy – a surgical incision into the chest wall. Image-guidance (usually fluoroscopy) is heavily utilized in these procedures to see the location of the catheter, since there is no direct line of sight.

One problem with image-based procedures is that the biological tissue can shift due to various factors such as respiration, edema, and tissue resection. While a surgeon can compensate for some of these motions, a robotic system can provide more accurate compensation using position information provided by intraoperative imaging. One such system where continuous motion compensation and highly accurate robotic manipulation has enabled complex treatment is in robotic radiosurgery (CyberKnife, AccuRay, Sunnyvale, California, USA) [12]. This system can compensate for respiratory motion of the lungs in order to accurately deliver radiation treatment to tumors.

1.2.2 Principles of Ultrasound Imaging

In ultrasound imaging, a sound wave in the megahertz range is emitted using piezoelectric crystals (Fig. 1.1). The sound wave transmitted through the tissue undergoes a nonlinear process where

1. the wave is attenuated / absorbed by the tissue,




Figure 1.2: *Ultrasound wave propagation phenomena.*

2. the wave is totally reflected or partially reflected and refracted at tissue boundaries, and
3. the wave is scattered (Fig. 1.2).

The reflected signal is recorded by the piezoelectric elements. The data is then processed and represented as a series of intensities along the scanlines.

1.2.3 Use of Ultrasound in Cardiac Interventions

Intracardiac echocardiography (ICE) catheters can acquire 2D ultrasound images from within a patient's heart chambers. Use of ICE catheters is currently limited to critical procedures such as transseptal puncture and Patent Foramen Ovale (PFO) / Atrial Septal Defect (ASD) closure. However many other interventions can benefit from the use of ICE catheters [13]. For example, in ablation treatments for arrhythmia, use of ICE catheters can improve ablation quality [14–18], however this is not performed commonly due to the difficulty in steering the ICE catheter to align the imaging plane with the ablation catheter.

Researchers have shown that 3D ultrasound imaging can guide basic surgical tasks more efficiently and accurately than 2D ultrasound imaging [19]. 3D transesophageal echo (TEE) has been a substantial advance in guiding and monitoring transcatheter heart valve procedures for the treatment of structural heart diseases [20].

1.2.4 Use of Ultrasound in Robotic Interventions

The use of ultrasound-based guidance in robotic interventions is currently fairly limited, however research in this area is fairly active. A number of small ultrasound probes that are compatible with the da Vinci surgical system (Intuitive Surgical, Sunnyvale, CA, USA) are available for use in detecting tumor margins in procedures such as partial nephrectomy [21]. A research system that aimed to bring robotic assistance to treating prostate cancer through brachytherapy is undergoing Phase I clinical trials [22]. A number of robotic systems have been developed to provide passive, semi autonomous, and fully autonomous needle insertion using ultrasound-based image guidance [23–25]. With the development of 3D ultrasound imaging, it became possible to track cardiac tissue in ultrasound volumes. Researchers built handheld and catheter-based robotic devices to compensate for the high speed motion of internal cardiac structures, where real-time tissue tracking is achieved through 3D ultrasound imaging [26, 27].

1.2.5 Use of Ultrasound in Biomechanics Research

Use of ultrasound imaging is not limited to the operating room, but it is also used to advance our understanding of human biomechanics. Ultrasound-based muscle fascicle tracking for biomechanics studies has been demonstrated, including using ultrasound to estimate muscle and tendon dynamics during walking [28–30]. Such work has shown how changes in the aponeuroses and connecting fascicles impact the biomechanics of gait. There are many other organ systems and biomechanics-related research that could benefit from using ultrasound imaging for non-invasive quantification of tissue dynamics, such as for improving the fidelity of finite element simulations of the human body.

1.3 Thesis Contributions

The contributions of this work can be grouped into the four areas below:

1.3.1 Autonomous Steering of Flexible Instruments

In collaboration with another PhD student, Paul Loschak, PhD, I developed a robotic system for autonomous steering of ICE catheters to enable real-time tracking of cardiac structures and other instruments (e.g. ablation catheters) during cardiac interventions. I developed a computationally efficient framework to enable real-time control of the robot. We conducted six *in vivo* live animal trials using our system, successfully steering ICE catheters autonomously inside a beating heart. This enables real-time monitoring of the interaction between instruments and the cardiac tissue, potentially leading to a decrease in the use of X-ray fluoroscopy, and improving the quality of ablations.

1.3.2 Real-Time 3D Ultrasound Image Generation and Visualization for Procedure Guidance

Using the aforementioned robotic catheter system, I developed a framework to enable real-time reconstruction of 3D ultrasound volumes from the 2D images acquired by the ICE catheter. I implemented the computationally-demanding volume reconstruction and rendering algorithms in CUDA. We conducted two *in vivo* trials, in which we generated 3D views of the heart in real time. This work improves visualization of cardiac structures during interventions, potentially leading to a decrease in the recurrence of arrhythmia after ablation and the use of X-ray fluoroscopy.

1.3.3 Improving Ultrasound Image Quality Through Bayesian Methods and Enhanced Dynamic Range

We applied computational photography and machine learning techniques to improve ultrasound image quality. I used Gaussian process regression to increase interpolation accuracy in converting ultrasound scanlines to B-mode images. I also developed methods to acquire and generate high dynamic range ultrasound images, similar to how smartphones and digital cameras can combine pictures taken at different exposures to improve the visibility of different structures. In clinical practice, better quantification leads to more informative clinical

decision-making, and quantification depends heavily on the image quality.

1.3.4 Ultrasound Image Acquisition and Processing for Biomechanics Studies

Thinking about uses for ultrasound imaging outside the confines of surgical robotics, we proposed to use ultrasound imaging to acquire real-time, high frame rate information about the state of muscles and other tissue inside the human body in a non-invasive fashion. Together with collaborators, I developed a platform for characterizing the dynamic response of human muscles and internal organs, such as the stomach, and conducted IRB-approved studies with human subjects. Ultrasound images of the stomach are compounded in 3D space to reconstruct the shape of the organ. The data collected in these studies will help improve the fidelity of finite element models of the human body.

1.4 Thesis Overview

Chapter 2 presents a robotic system that was designed and built to autonomously steer ICE catheters. Chapter 3 describes work on building 3D ultrasound volumes in real-time from 2D images acquired by an ICE catheter. Chapter 4 presents work on using Gaussian process regression to improve ultrasound scanline conversion accuracy. Chapter 5 describes work on extending the dynamic range of ultrasound images to improve the level of visible details. Chapter 6 presents work on applying ultrasound imaging to biomechanics studies, enabling the in situ study of phenomena that could not be observed before. Chapter 7 gives a summary of our contributions and present views for the future of ultrasound imaging, image-guided interventions, and robotic surgery.

Chapter 2

Autonomous Steering of Ultrasound-Imaging Catheters

Cardiac catheterization is useful for performing a variety of diagnostic and interventional tasks in minimally invasive procedures [11]. However procedure guidance and situational awareness in the electrophysiology suite is generally limited to fluoroscopy. Intracardiac echocardiography (ICE) catheters provide real-time ultrasound imaging from within the heart; however their use in cardiac interventions is limited due to the difficulty in manually steering the catheter, and the lack of advanced visualization tools for procedure guidance. Automated robotic steering of ICE catheters can potentially enhance imaging of cardiac structures and working instruments, leading to better visualization and procedure guidance.

In the following two chapters, I present a robotic system that can autonomously steer ICE catheters to track other instruments inside the heart and generate 3D ultrasound volumes of cardiac structures in real time. I describe the software architecture for data acquisition, robot control, and volume reconstruction and rendering. I present results from benchtop experiments and *in vivo* trials. The advanced procedure guidance tools developed in this work can improve intra-procedural treatment by enabling new image guidance capabilities. Even though this work focuses on ICE catheters and cardiac applications, the control strategies developed here can be applied to other flexible instruments, and the imaging methods can be

beneficial in other organ systems. I collaborated closely with Paul Loschak, PhD on developing the robotic system, the control algorithms, and conducting the *in vivo* trials.

2.1 Background

2.1.1 Clinical Problem: Arrhythmia

Cardiac arrhythmia is a heart condition that affects millions of people in the world. The contraction pattern of the heart is regulated by the conduction of electrical signals through pathways found in the heart. In arrhythmia, a disruption in the regular electrical conduction through the heart causes abnormal heart rhythms. Misfiring of cardiac muscles can decrease efficiency in pumping blood to the rest of the body, leading to cardiac failure. There are multiple types of arrhythmia, however with an estimated 33.5 million people affected in the world [31], atrial fibrillation is the most common.

2.1.2 Current Practice

Depending on the type and severity of arrhythmia, patients are given medication to regulate cardiac muscle activity. In patients who do not respond to medication or require immediate intervention, the electrical pathways that lead to the misfiring of the cardiac muscles are ablated by delivering radio frequency energy (*i.e.* heat) to the tissue. This is a minimally invasive procedure that is performed by electrophysiologists using catheters. A successful ablation requires precise localization of the ablation zone, and continuous energy delivery to thoroughly ablate the tissue. Up to 40% of patients have a recurring arrhythmia [32], most of which are caused by incomplete ablation.

During the procedure, an electroanatomical map of the heart is generated using a mapping catheter and analyzed to determine the ablation regions. Currently image guidance is provided through X-ray fluoroscopy during mapping and ablation. Fig. 2.1 shows a fluoroscopy image captured during one of our *in vivo* trials. While catheters and ribs are clearly seen in the image, soft tissue details (e.g. cardiac tissue) are not discernible. Therefore using fluoroscopy

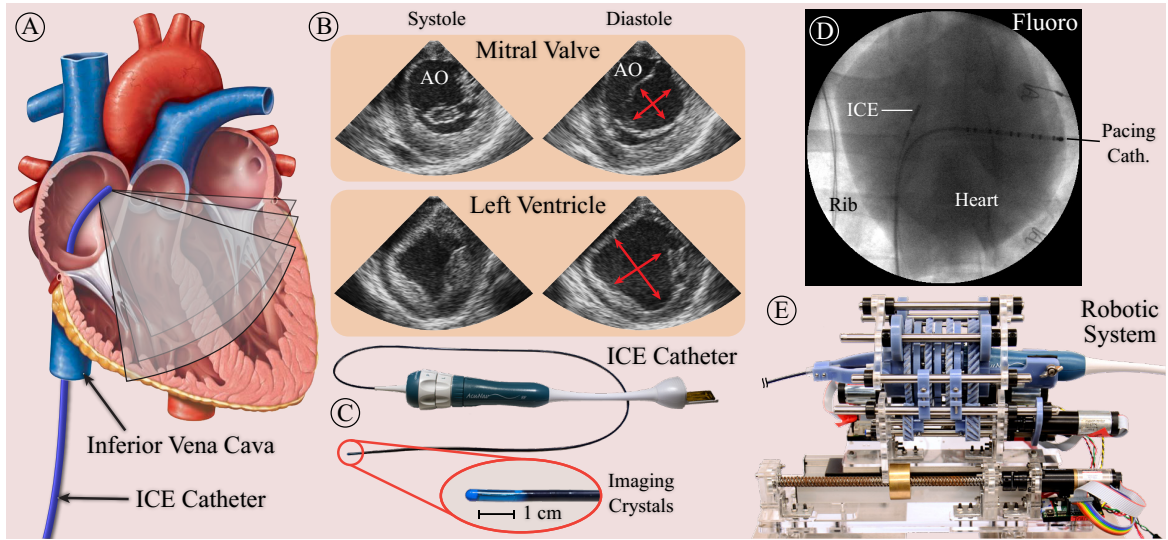


Figure 2.1: (a) A series of ultrasound images are acquired from within the heart at different locations. (The heart diagram is taken from [33]) (b) Ultrasound images of the mitral valve and aortic opening (AO), and the left ventricle at systole and diastole. Images are acquired using an ICE catheter positioned in the right atrium. (c) An ICE catheter. (d) An X-ray fluoroscopy image of the chest, showing the heart, ICE catheter, and a pacing catheter. (e) The robotic catheter steering system.

to confirm ablation quality is inadequate, as the clinician is not able to see if the catheter is in contact with the ablation region, or whether there is sliding between the catheter and the tissue due to cardiac motion.

2.1.3 Ultrasound Imaging Catheters

Intracardiac echocardiography (ICE) catheters provide real-time ultrasound imaging from within the heart. The 2.67 mm (8 Fr) diameter catheter houses a 64-element piezoelectric ultrasound transducer at the distal tip. The utility of ICE catheters in improving ablation quality has been shown in [14–18]. However their use in cardiac interventions has so far been limited due to the significant training and skill required for manually steering the imaging plane, and the lack of advanced visualization tools for procedure guidance.

2.1.4 Robotic Catheters and Image Guidance

There are a number of commercially available robotic catheter steering systems [34–38]. These systems enable the teleoperation of catheters, thereby reducing the X-ray exposure of clinicians.




Figure 2.2: (a) The four controllable degrees of freedom of the ICE catheter. (b) A picture of the ICE catheter handle. (c) Structures and coordinate frames of interest.

However these systems do not offer autonomous steering of cardiac catheters.

A number of research prototypes have demonstrated autonomous steering of cardiac catheters and endoscope-size manipulators in constrained bench top environments [39–45]. These prototypes have not been shown to steer accurately in the unconstrained environments typical of catheter interventions, and do not control the orientation of the manipulator tip, which is crucial for ultrasound imaging. Additionally, most are configured to steer custom built manipulators that are designed for robotic control; however this approach does not translate to clinical use as readily as a system that can robustly manipulate commercially-available catheters.

2.1.5 Goals and Challenges

We recognized the value of an autonomous catheter steering system in delivering image-based guidance during minimally-invasive cardiac procedures. Our major goal was to develop a clinically-viable system that could enable real-time ultrasound-based procedure guidance, and minimize or completely eliminate the use of X-ray fluoroscopy. We focused on ablation monitoring with the end goal of improving ablation quality and accuracy, thereby reducing risk of recurrence.

Our system enables two major capabilities for use in cardiac interventions. First is to

enable real-time tracking of other instruments inside the heart to align the ultrasound imaging plane with the tip of the instrument for monitoring tissue-instrument interactions. The second is to acquire 2D ultrasound images of a section of the heart to reconstruct them into 4D (3D + time) ultrasound volumes, which can provide improved procedure guidance (presented in Chapter 3).

There are a multitude of challenges in translating a system from the benchtop to animal studies. In our case, we faced the following major challenges:

1. Steering a flexible manipulator accurately,
2. Compensating for physiological motion (respiratory and cardiac),
3. Tracking the motion of other working instruments inside a beating heart,
4. Keeping the footprint of the catheter steering system small and compatible with existing devices in the electrophysiology suite,
5. Implementing robust and efficient algorithms to enable reliable and real-time execution (addressed in Chapter 3),
6. Handling data input/output to/from multiple devices asynchronously at different sampling rates (addressed in Chapter 3),
7. Creating an intuitive and informative user-interface that enables quick debugging of the system (addressed in Chapter 3), and
8. Reconstructing and rendering 4D (3D + time) ultrasound volumes (addressed in Chapter 3).

In our previous work we developed a system for automatically pointing cardiac imaging catheters in a constrained bench top environment [46]. While this system has proven to be effective for bench-level steering, it was necessary to implement a more robust controller that can cope with the friction, backlash, and unconstrained catheter shaft motion that occurs during *in vivo* procedures. The following sections describe the motivation for redesigning the controller for improved robustness, the new control strategy, system characterization experiments, and results from *in vivo* catheter steering. This work provides the first method known to us for accurately controlling unconstrained flexible manipulators in the presence of unmodeled external disturbances.

2.2 Catheter Kinematics and Sensing

The ICE catheter is 90 cm long and 2.67 mm (8 Fr) in diameter. A picture of the catheter is shown in Fig. 2.1 and the catheter degrees of freedom (DoF) are indicated in Fig. 2.2. The distal 5 cm of the catheter, which we denote as the ‘bending section’, is built to be more flexible compared to the rest of the catheter body. Four pull wires extend from the distal tip of the bending section to the catheter handle, connecting to two control knobs on the handle in a pull-pull configuration. The pull wires allow the clinician to pitch and yaw the imaging plane. The clinician can also translate and rotate the handle axially to translate and roll the imaging plane (see Fig. 2.2a-b). The clinician manipulates the ultrasound catheter by holding the plastic handle and actuating the four degrees of freedom (4-DOF) to generate the resulting tip motions shown in Fig. 2.2a.

The ICE catheter is an underactuated system with infinite degrees of freedom. Pull wire actuation deflects not only the bending section, but also the entire body of the catheter. This parasitic motion of the catheter body is hard to model or predict as it is highly dependent on the boundary conditions (*e.g.* vasculature geometry and contact, tissue properties) and other unmodeled parameters.

2.2.1 Constrained Catheter Motion

In our previous work, the robotic system for steering cardiac imaging catheters relied on a set of physical constraints that isolated the bending to only the distal 5 cm bending section of the tip [46]. The base of the bending section was fixed with respect to the catheter handle and robot. This catheter tip fixation succeeded in isolating the bending section of the catheter and simplified the control required for steering, which was necessary at the beginning stage of our investigation of catheter position and orientation steering.

2.2.2 Unconstrained Catheter Motion

Accurately maneuvering the 4-DOF catheter *in vivo* requires a controller that is robust to a number of inaccuracies and disturbances, which have been minimally addressed in the robotic




Figure 2.3: Uncontrollable catheter shaft motion. Dashed lines show alternative locations of the catheter shaft due to uncontrolled motion within the vasculature. EM trackers measure the position and orientation of the catheter at two locations: at the base of the bending section, and at the tip.

flexible manipulator literature. Inside the plastic off-the-shelf catheter handle, slack in the pull wires connected to each bending knob produce large backlash and deadzone regions that are non-linearly coupled to the rotations of both bending knobs (Fig. 2.5). Additionally, the long catheter shaft experiences compression, bending, and torsion in the vasculature. This reduces the effectiveness of the bending knobs, increases friction between the catheter and the vascular walls, and causes the location of the base of the bending section to vary, as portrayed in Fig. 2.3. Here the catheter is shown in three potential configurations. The

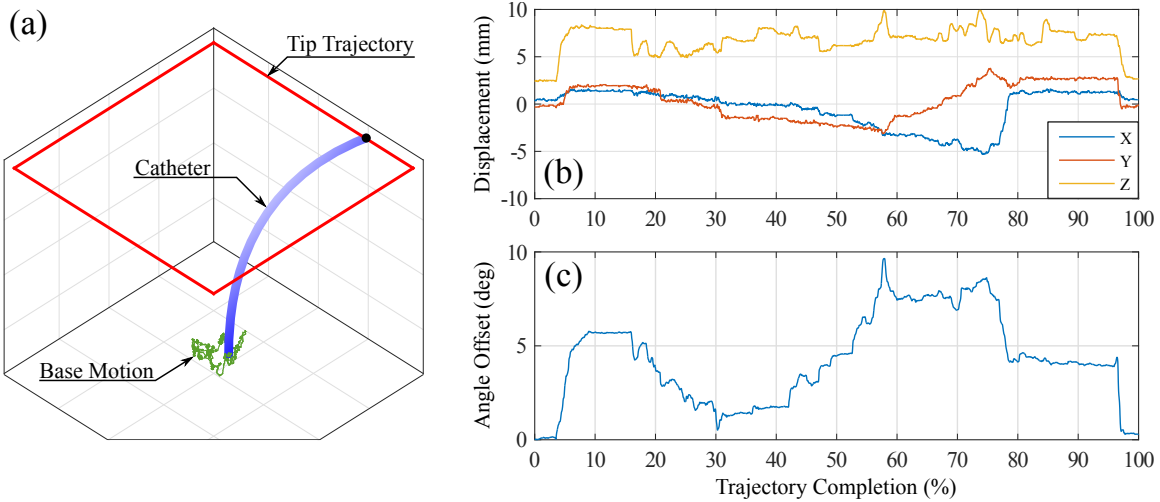


Figure 2.4: Catheter navigation causes uncontrollable motion of the bending section base. (a) The base motion is measured (green) while the catheter tip is navigated along a square trajectory (red). (b) The displacement of base point in xyz coordinates and (c) the change in base point rotation about the catheter axis are measured.




Figure 2.5: ICE catheter deadzone is variable and depends on the current configuration of the catheter.

proposed controller keeps the tip at the target location regardless of where the base is located. The uncontrollable motion of the base point during unconstrained bending significantly affects the ability of the catheter to converge to the desired target pose. To illustrate typical base motion, Fig. 2.4 shows the uncontrollable motion of the bending base during catheter tip navigation of a square trajectory with the shaft unconstrained. The range of displacements of the bending base in the xyz -axes were measured as [8.4 mm, 6.4 mm, 5.9 mm] and the maximum angular deviation from the centerline was 9.65°.

The catheter passes into the patient vasculature through an introducer, which is a plastic tube containing a rubber seal to prevent blood outflow. Friction between the catheter shaft and the introducer seal increases the amount of windup due to catheter rotation and causes the catheter body to buckle during insertion. Respiratory and cardiac motions of the patient also contribute disturbance to the system, resulting in further inaccuracies. Patient vasculature is highly variable in diameter, stiffness, curvature, and length. Catheter translation is subject to

backlash as the catheter body conforms to the near wall or the far wall of the vasculature (examined in [47]).

We decided to interface with commercial ICE catheters instead of building our own, because ICE catheters are commercially available and approved for clinical use, which increases our robotic system’s likelihood of adoption by the medical community. However this comes at the cost of increased controller complexity due to the mechanical limitations of the catheter, as it is not built for robotic control. While some of these limitations may be avoided by altering the mechanical design of a catheter, it is less expensive and more clinically feasible for a robust controller to manipulate an off-the-shelf catheter.




Figure 2.6: Three parameter spaces are used to describe the motion of the catheter. Joint mapping and kinematics map these spaces to each other.

2.2.3 Kinematics

Previous research has mainly focused on 3-DOF continuum robot kinematics and did not consider the extra DOF required to control the heading of the ultrasound imaging plane [40, 41, 43, 48–50]. Our kinematics derivation follows the methodology in [40]. With the use of two sensors (one at the tip and one at the base of the bending section, as shown in Fig. 2.2c) we extend this to enable four DOF control of ultrasound imaging catheters.

Three different parameter spaces are used to control the flexible manipulator: joint, configuration, and task space (Fig. 2.6). The joint space is the four actuator variables corresponding to manual catheter controls, *i.e.* pitch (ϕ_1), yaw (ϕ_2), roll (ϕ_3), and translation (ϕ_4) shown in Fig. 2.2 and 2.2. The configuration space variables ($d, \gamma, \theta, \alpha$) describe the shape of the catheter as shown in Fig. 2.7. Lastly, the task space describes the position of the

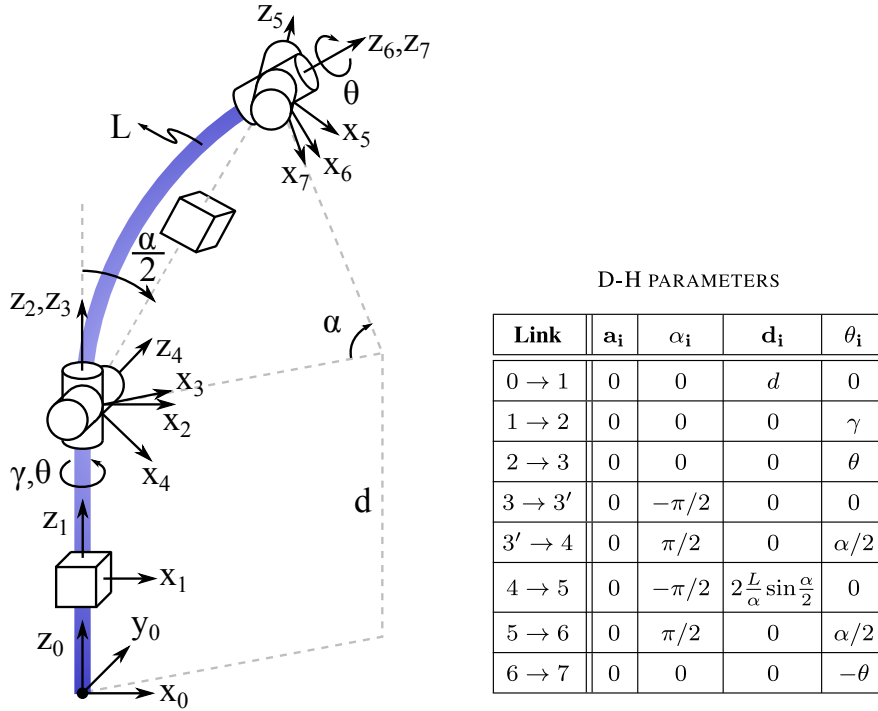


Figure 2.7: Symbolic representation showing D-H parameters: translation, d , catheter handle roll, γ , bending plane angle, θ , bending angle around the global z -axis, α , length of the bending section, L .

catheter tip (x, y, z) and orientation of the imaging plane (ψ). The mapping from joint space to configuration space is based on the mechanical design of the robot and the manipulator, and is therefore system-specific; whereas the mapping from configuration space to task space is a kinematic mapping that is broadly applicable to flexible manipulators.

The accuracy of the mapping between configuration space and task space is based on the following idealized assumptions: the bending section of the catheter has a constant curvature along its length (circular arc); the bending catheter lies on a single plane; and the effects of torsion about the catheter tip are negligible. In reality, these assumptions lead to less accurate results. However, the idealized relationships between task space and configuration space can still provide valuable information regarding how the robot joints should be adjusted for the catheter tip to reach the desired pose, and, coupled with feedback control, can compensate for modeling inaccuracies adequately.

2.2.4 Forward Kinematics

Forward kinematics map the configuration space to task space. The table in Fig. 2.7 shows the D-H parameters we derived for four DOF catheters, following the methodology used in [40]. Frame 0 is the base frame, and Frame 7 is the catheter tip. Fig. 2.7 shows the configuration parameters, where d is the translation, γ is the catheter handle roll, θ is the angle between the x -axis of the catheter handle and the bending plane, and α is the bending from the global z -axis (i.e. the catheter body). L is the length of the bending section. The transformation matrix, $T(d, \gamma, \theta, \alpha, L)$, describes the pose of the catheter tip in base coordinates (Eq. 2.1).

$$T(d, \gamma, \theta, \alpha, L) = \begin{bmatrix} s_{\gamma+\theta}s_\theta + c_{\gamma+\theta}c_\theta c_\alpha & -s_{\gamma+\theta}c_\theta + c_{\gamma+\theta}s_\theta c_\alpha & c_{\gamma+\theta}s_\alpha & Lc_{\gamma+\theta}(1 - c_\alpha)/\alpha \\ -c_{\gamma+\theta}s_\theta + s_{\gamma+\theta}c_\theta c_\alpha & c_{\gamma+\theta}c_\theta + s_{\gamma+\theta}s_\theta c_\alpha & s_{\gamma+\theta}s_\alpha & Ls_{\gamma+\theta}(1 - c_\alpha)/\alpha \\ -c_\theta s_\alpha & -s_\theta s_\alpha & c_\alpha & d + Ls_\alpha/\alpha \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2.1)$$

2.2.5 Inverse Kinematics

The idealized mapping from task space to configuration space begins with T and calculates α , d , θ , and γ . From the translational component of T we can calculate the distance from the catheter tip to the global z -axis as

$$\|x_{tip}, y_{tip}\|_2 = \|T_{14}, T_{24}\|_2 = \frac{(1 - c_\alpha)}{\alpha}, \quad (2.2)$$

where $\alpha \in \mathbb{R}^+$, c_α denotes $\cos \alpha$, and $\|\cdot\|_2$ is the l^2 -norm. Then α can be calculated by computing the root of the nonlinear equation in Eq. 2.2. The idealized d is then calculated from T_{34} . The idealized angle θ is

$$\theta = \text{atan2} \left(-\frac{T_{32}}{s_\alpha}, -\frac{T_{31}}{s_\alpha} \right). \quad (2.3)$$

Lastly, the angle γ can be calculated from

$$\gamma + \theta = \text{atan2}(y_{tip}, x_{tip}) = \text{atan2}(T_{24}, T_{14}). \quad (2.4)$$

2.2.6 Joint Mapping

The desired configuration space motion can be mapped to the corresponding joint space inputs as

$$\phi_1 = 2R_c\alpha \cos \theta / D_{knob} \quad (2.5)$$

$$\phi_2 = -2R_c\alpha \sin \theta / D_{knob} \quad (2.6)$$

$$\phi_3 = \gamma \quad (2.7)$$

$$\phi_4 = d \quad (2.8)$$

where R_c is the radius of the catheter body and D_{knob} is the diameter of the pulley inside the catheter handle.

The following section describes a control strategy that will robustly converge the catheter tip to the desired position and point the ultrasound imager in the desired direction despite these effects and disturbances.

2.3 Control Strategy

Controlling the imaging plane position and angle of an ICE catheter is challenging due to the flexibility (parasitic motion and modeling challenges) of the body, deadzone and friction in the pull wires (Fig. 2.5), and catheter-specific variability in mechanical behavior. We determined that it is necessary to sense not only the pose of the bending tip, but also of the bending base in order to determine the bending curvature and to account for the parasitic motion of the bending base due to catheter actuation. Due to mechanical transmission losses along the flexible body of the catheter, it is not possible to determine the tip configuration using encoders attached to the catheter handle. These transmission losses are catheter-specific, time-varying (materials degrade over time), and dependent on boundary conditions. Therefore




Figure 2.8: The controller (gray box) receives the desired catheter tip pose and iteratively calculates joint angle adjustments to manipulate the catheter.

we attach two electromagnetic (EM) sensors near the bending base and the bending tip (see Fig. 2.2c). EM sensors are commonly used in electrophysiology (EP) labs during cardiac interventions, therefore the use of EM sensors is not a concern for the clinical viability of the system. The EM sensor readings are not affected by the energy delivered during ablation.

A feedforward model is used to calculate the required joint level adjustments to reach the target configuration. These commands are then scaled based on the magnitude of the error, and the catheter position is iteratively adjusted to reach the target.

The measured sensor information and idealized kinematic relationships are used together to calculate the joint space adjustments for tip convergence. The robot navigates the tip of the ultrasound catheter to reach the desired position while aiming the ultrasound imager in the desired direction by the controller shown in Fig. 2.8. It is not possible to control the full 6-DOF pose of the ultrasound imager, therefore the controller aims to adjust the 3-DOF position of the tip and the 1-DOF rotation of the ultrasound imager about the axis of the catheter tip.

Two EM sensors are used to resolve the configuration of the catheter tip. One sensor is placed at the tip of the bending section, proximal to the ultrasound transducer such that the ultrasound beam is not distorted. The second sensor is placed at the base of the bending section to resolve physiological disturbances from the environment and the unmodeled behavior of the catheter body.

Let T_{sensor} be the transformation corresponding to the sensor reading at the current location of the catheter tip. The ultrasound imaging plane is the plane formed by the x - and z -axes of the tip pose. ultrasound imager rotation is measured in rotation of the x -axis about

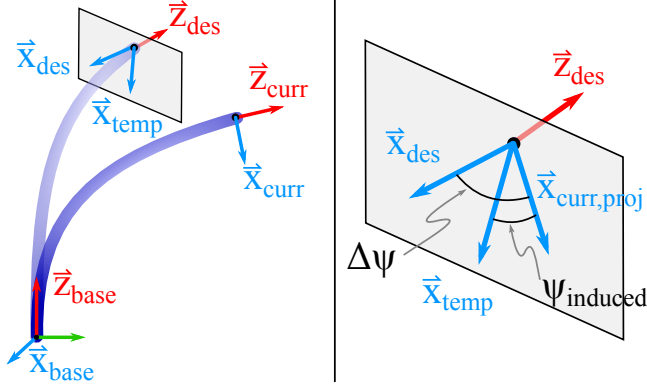


Figure 2.9: The required configuration space roll adjustments is calculated by projecting the \vec{x}_{curr} axis on to the $x - y$ plane of the desired coordinate frame.

the z -axis of the tip pose. T_{sensor} cannot be directly used in the inverse kinematics calculations since the rotational component of the sensor reading will have rotations that are induced by the unmodeled effects in the system; instead a kinematically feasible configuration $C_{curr} = [d_{curr}, \gamma_{curr}, \theta_{curr}, \alpha_{curr}]$ needs to be calculated based on the position $[x_{sensor}, y_{sensor}, z_{sensor}]$, and handle roll γ_{sensor} reported by the base sensor. Let $C_{curr} = [d_{curr}, \gamma_{curr}, \theta_{curr}, \alpha_{curr}]$ represent this kinematically feasible configuration, with pose T_{curr} .

The user indicates a desired position and change in heading, $X_{des} = [x, y, z, \Delta\psi]$. Calculating the desired catheter configuration, $C_{des} = [d_{des}, \gamma_{des}, \theta_{des}, \alpha_{des}]$, requires additional steps beyond using the inverse kinematics equations. Any change in catheter tip position induces a change in the imager heading that is dependent on the current and target positions. In order to arrive at the desired heading, this induced angular change $\psi_{induced}$ has to be calculated and adjusted for in the following manner before outputting commands to the motors (Fig. 2.9).

(1) Calculate inverse kinematics given $X_{des,xyz}$ and the current handle roll, γ_{curr} . The angle γ_{curr} is measured from the base sensor (rather than the current roll joint) because a significant amount of rotation is lost through torsion along the catheter shaft. This calculates the configuration at the target, given γ_{curr} . This pose is defined as T_{temp} , and the configuration parameters are labeled as $C_{temp} = [d_{temp}, \gamma_{temp}, \theta_{temp}, \alpha_{temp}]$.

(2) The position and curvature of the catheter at the temporary and desired configurations are equal. Therefore, the configuration parameters $d_{des} = d_{temp}$ and $\alpha_{des} = \alpha_{temp}$. However,

the temporary pose only represents the target position and requires additional roll calculations to point the ultrasound imager (Fig. 2.9). We begin by calculating the angle between the x -axis at the temporary pose and the x -axis at current pose. This is done by first projecting \vec{x}_{curr} to the plane defined by the vectors \vec{x}_{temp} and \vec{y}_{temp} ,

$$\vec{x}_{curr_{proj}} = \vec{x}_{curr} - (\vec{x}_{curr} \cdot \vec{z}_{temp}) \vec{z}_{temp}, \quad (2.9)$$

and calculating the angle $\psi_{induced}$ between the x -axes,

$$\psi_{induced} = \cos^{-1} \left(\frac{\vec{x}_{curr_{proj}} \cdot \vec{x}_{temp}}{\|\vec{x}_{curr_{proj}}\|} \right). \quad (2.10)$$

To maintain directionality, we calculate the direction of the normal to the plane defined by $\vec{x}_{curr_{proj}}$ and \vec{x}_{temp} ,

$$\xi = \vec{z}_{temp} \cdot \left(\vec{x}_{temp} \times \frac{\vec{x}_{curr_{proj}}}{\|\vec{x}_{curr_{proj}}\|} \right). \quad (2.11)$$

If $\xi < 0$ then the angle $\psi_{induced}$ is negated.

(3) At this point, we can point the ultrasound imager in the target direction by calculating the desired configuration parameters γ_{des} and θ_{des} as

$$\gamma_{des} = \gamma_{temp} + \psi_{induced} + \Delta\psi, \quad (2.12)$$

$$\theta_{des} = \theta_{temp} - \psi_{induced}. \quad (2.13)$$

The desired configuration is then related to the control inputs, ϕ_{1-4} . Convergence to the desired pose is an iterative process, in which the tip of the catheter is sensed, and the remaining relative change required to reach the desired pose is recalculated. Image collection is triggered when the user-defined allowable translational and rotational error thresholds of the tip are satisfied.

2.3.1 Disturbance Rejection

During *in vivo* procedures the catheter additionally experiences external disturbances due to respiratory motion, adjacent catheter motions in the same vessel, tissue motion due to displacement of adjacent organs, etc. We represent external disturbances as U_{ext} in Fig. 2.8.

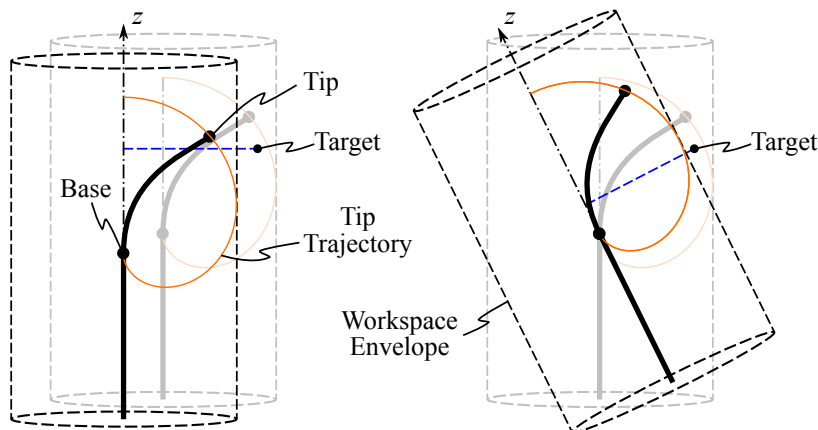


Figure 2.10: Disturbances in the base location of the bending section can cause (left) translations and/or (right) rotations, which can lead to the target exiting the workspace. The workspace of the ultrasound catheter is a cylinder, shown with dotted lines. The gray and black workspace cylinders correspond to the workspace before and after uncontrolled base motion. The tip of the catheter bends along the trajectory shown in orange. The blue dotted line represents the shortest distance between the target and the z -axis of the catheter base.

U_{ext} can cause the catheter shaft to translate (Fig. 2.10 left) and/or rotate (Fig. 2.10 right). Our controller needs to be robust to this external input such that stability is guaranteed and the catheter tip does not diverge from the target location.

The workspace envelope of the catheter is a cylinder (Fig. 2.10). An external disturbance can move the base of the catheter so that the target point is no longer contained in the workspace. The condition for the target to remain within the workspace is to maintain $\|x_{des}, y_{des}\|_2 \leq R_w$, where R_w is the radius of the workspace cylinder. Calculating the workspace boundary from the base tracker pose enables real-time detection of out-of-workspace errors. In the clinical setting, such errors would stop catheter servoing and be reported to the clinician for repositioning the catheter.

2.3.2 Respiration Compensation

During cardiac interventions, the position of the ICE catheter is cyclically altered by physiological disturbances (respiratory and cardiac). These disturbances affect each region of the heart differently, and the ICE imaging plane can become misaligned with the target it is imaging. Due to the mechanical limitations of the ICE catheter (material degradation due to




Figure 2.11: *The prototype system for autonomously steering ICE catheters. The sequence on the bottom shows the order of operations to insert and secure an ICE catheter in to the system. Catheter attachment and detachment takes less than 10 seconds, where the catheter clamp is lifted up to couple with the catheter handle, and a wingnut is tightened to secure the catheter in place.*

cyclic loading, low tensile strength of the pull wires, hysteresis and plastic deformation), we focus on tracking and compensating for the respiratory disturbances.

Before actively steering the catheter, the controller observes the motion induced by respiration on the catheter body using the EM trackers on the ICE catheter, however these also measure the disturbance due to cardiac motion. We first apply a low-pass filter to the EM measurements to filter out the cardiac motion from each of the three translation axes and the four components of the axis-angle representation for the rotation. (Fig. 2.12).

The filtered tracker data now only contains the respiratory disturbance, which is cyclical

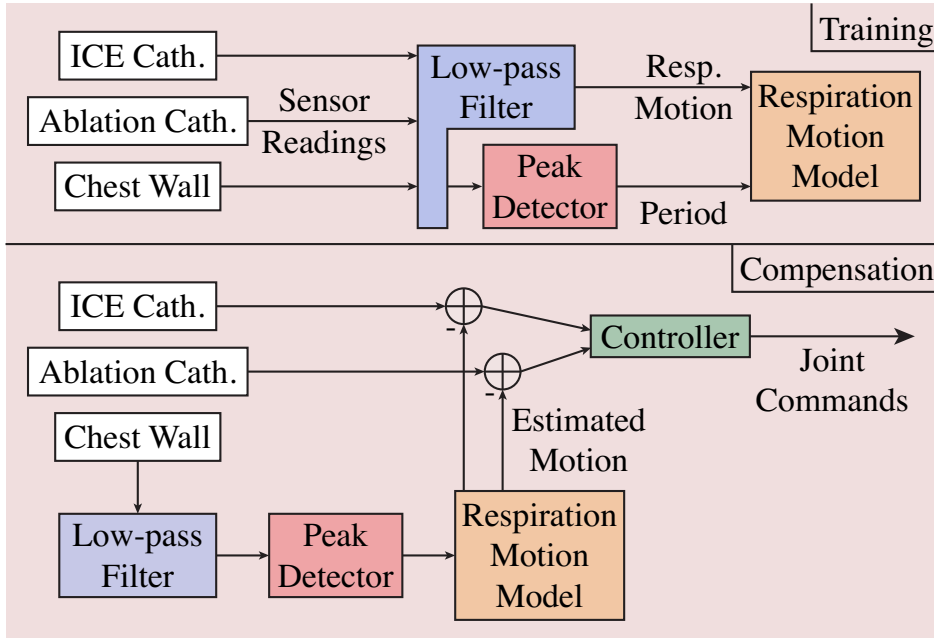


Figure 2.12: Data flow in respiration motion compensation.

and sufficiently regular to be modeled using a fourth order Fourier decomposition. Once the model is trained, the ICE catheter can remain in a fixed position with respect to the heart rather than with respect to the EM emitter simply by adding the modeled disturbance to the catheter’s target pose. Note that while the catheter is actively steered, the respiration model cannot be updated, since motion due to active steering and motion due to respiration can not be accurately separated.

Here we briefly describe the Fourier series modeling approach (a detailed treatment can be found in [51]). A continuous signal can be described as the sum of harmonically related sine waves, expressed as

$$z(t_n) \approx c_0 + \sum_{i=1}^m a_i \sin(\theta_i t_n) + b_i \cos(\theta_i t_n) , \quad (2.14)$$

where t_n is time at step n of N , c_0 is the DC offset, m is the number of harmonics, ϕ_i is the phase offset at each harmonic, a_i and b_i are amplitudes, $\theta_i = 2\pi i \omega_0$, and ω_0 is the fundamental frequency. Eq. 2.14 is given in rectangular form instead of the polar form, because it is linearly

separable, making it possible to write, $\mathbf{z} = \mathbf{A}\mathbf{x}$ where

$$\mathbf{A} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ \sin(\theta_1 t_0) & \sin(\theta_1 t_1) & \cdots & \sin(\theta_1 t_{N-1}) \\ \vdots & \vdots & & \vdots \\ \sin(\theta_m t_0) & \sin(\theta_m t_1) & \cdots & \sin(\theta_m t_{N-1}) \\ \cos(\theta_1 t_0) & \cos(\theta_1 t_1) & \cdots & \cos(\theta_1 t_{N-1}) \\ \vdots & \vdots & & \vdots \\ \cos(\theta_m t_0) & \cos(\theta_m t_1) & \cdots & \cos(\theta_m t_{N-1}) \end{bmatrix}, \quad (2.15)$$

and $\mathbf{x} = [c_0, a_1, \dots, a_m, b_1, \dots, b_m]^T$. The state coefficients \mathbf{x} can be estimated using least squares. Our implementation makes use of singular value decomposition to solve this linear system. Each coordinate (translation and rotation) of the bending tip and bending base is modeled independently.

Once the model parameters are determined, the controller is activated, meaning the observations from the EM sensors on the catheter cannot be used to update the respiration model any more. The model tends drift across time due to variations in the respiratory cycle, therefore the respiration frequency is measured using an EM tracker attached to the patient's chest, and the corresponding phase shift is added to the model to keep it aligned with the current respiration cycle.

2.4 Robot and System Design

The kinematic algorithms were implemented and tested on a robotic prototype for manipulating an ultrasound catheter. The design of the robot (Fig. 2.11) connects four actuators with the four degrees of freedom of the catheter handle. A system of coaxial rotating rings mates to the bending knobs and the catheter handle to actuate pitch, yaw, and roll (see Fig. 2.13). The rotational transmission rests on a linear stage which is translated by a lead screw. The robot was designed such that an ultrasound catheter can be quickly fastened into (and removed from) the robot during clinical procedures, as demonstrated in Fig. 2.11. This is a safety




Figure 2.13: Cross-sectional view showing knob interactions with the catheter. Parts B, E, H are bearings. Parts B, C, D, E enable yaw; parts E, F, G, H enable pitch; and parts A, B, H, I, J enable roll.

feature that allows clinicians to take over manual control of the ICE catheter in case of a system failure. Two scissor lifts adjust the inclination of the robot for aligning the catheter with the introducer to minimize buckling of the catheter shaft.

Each of the four DoF of the catheter is actuated by a 25W brushless DC motor (EC-max22, Maxon, Switzerland) with a 53:1 gear reduction. Each motor is driven by a motor controller (EPOS2 24/2) that runs an internal servo loop at 1 kHz. The error between the commanded and measured joint angle motion is negligible [52,53]. Fig. 2.14 shows a diagram of the devices that enable autonomous steering of ICE catheters.

The host PC uses the CANopen protocol to communicate with the motor controllers through a PCIe-based active CAN interface card (CAN-IB200/PCIe, IXXAT, Bedford, NH, USA). We switched from a USB connection in our earlier prototypes to a CAN connection in this version due to the high latency of USB communication (3 ms vs. 460 μ s) per axis. This switch required no change in the control code; the high-level Service Data Object (SDO)

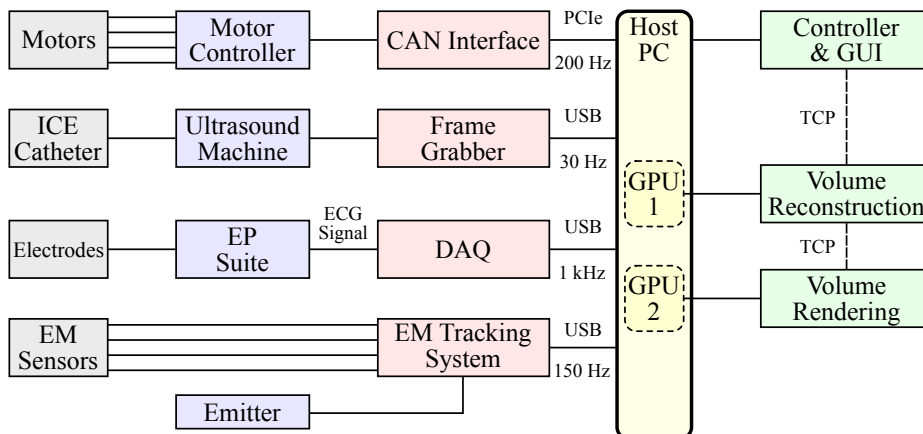


Figure 2.14: Device diagram showing the communication layers between the physical devices and the software framework components.

communication uses the same API calls for both USB and CAN communication modes.

An S-video frame grabber (USB3HDCAP, StarTech, London, Ontario, Canada) captures ultrasound images at 30 frames per second. The frame grabber is capable of faster acquisition at 60 fps, however the ultrasound imaging system used in our studies (ACUSON Sequoia 512, Siemens Healthcare, Erlangen, Germany) can only transmit S-video signals at 30 fps. Depending on the imaging settings, the ICE catheter acquisition rate can go up to 120 fps.

ECG signals are acquired from a pacing catheter positioned in the coronary sinus. The pacing catheter is connected to an electrophysiology (EP) system (LABSYSTEM, Boston Scientific, Marlborough, MA, USA) for signal conditioning. A digital data acquisition system (U6, LabJack, Lakewood, CO, USA) captures the ECG signals from the EP system at a 1 kHz sampling rate.

The motion of the ICE catheter, the target catheter, and the chest wall are tracked at 150 Hz using four 6-DOF electromagnetic (EM) sensors (trakSTAR, Ascension Technology/NDI, Canada). The EM trackers have a spatial resolution of 0.5 mm and angular resolution of 0.1° , and RMS accuracy of 1.4 mm and 0.5° .

The host machine houses an Intel i7 4790K quad-core overclocked CPU running at 4.4 GHz with 32GB DDR3 RAM and two NVIDIA TITAN GPUs with 6GB RAM each. The data generated by the system, system events, and user actions are stored on an SSD drive,

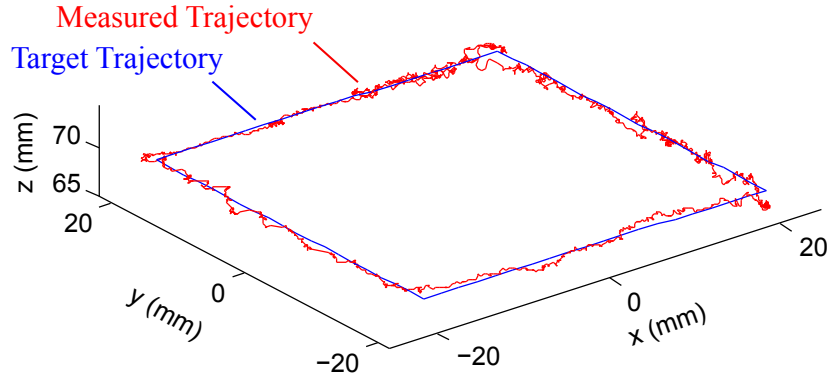


Figure 2.15: The catheter navigates a 4 cm square trajectory of 80 set points while pointing the ultrasound imager in the same direction for all points.

which reduces data I/O overhead.

2.5 Experiments

2.5.1 Bench Top Motion

In preparation for *in vivo* studies, the performance of the controller was tested in a simulated vasculature environment on the bench top. The catheter was introduced through a 14 Fr introducer into a smooth-walled plastic tube (Teflon FEP, 1.3 cm inner diameter, 62 cm length). The distal bending section of the catheter was able to bend freely outside the distal end of the tube.

The first bench top motion study measured the ability of the robotic system to navigate the catheter tip through space along a 4 cm square trajectory of 80 setpoints. The catheter was navigated through multiple trajectory orientations and sizes. The ultrasound imager was adjusted to point in the same heading for all trials. An example trajectory is shown in Fig. 2.15, where the blue line is the target trajectory and the red line is the measured trajectory of the catheter tip. The catheter tip reached the desired set points with mean absolute errors of 0.92 mm (standard deviation $\sigma = 0.12$ mm) and 1.01° ($\sigma = 1.05^\circ$). The base tracker moved up to 8.4 mm away from the initial point due to catheter actuation. The average time to navigate to each set point (dependent on the allowable error threshold and the




Figure 2.16: Bench top tests demonstrating the ability of the robotic system to maintain the catheter tip at the target pose despite disturbance step and target step inputs.

proportional controller gain) ranged from 1.5 s to 3.7 s, which is well within a 5 s specification for breathing motion.

The next motion study isolated ultrasound imager rotation while maintaining a fixed position in space. The ultrasound imager was commanded to rotate 2° a total of 20 times with a specified allowable angular error threshold, while the position of the catheter tip was maintained within the 1 mm position error threshold (measured mean absolute error 0.83 mm, $\sigma = 0.3$ mm). The average time to converge on the target imager heading ranged from 2.61 s ($\sigma = 2.67$ s) with error threshold 0.25° to 1.81 s ($\sigma = 1.15$ s) with error threshold 1° .

The last bench top study examined the performance of the system in response to a step disturbance. This is important because interaction with adjacent catheters can displace the

shaft, and respiratory motion causes the heart and the supporting vasculature to shift. A step input disturbance was created by moving the vasculature phantom at the base of the catheter bending section, thereby displacing and slightly rotating the base of the catheter. The controller senses the disturbance and makes the necessary adjustments to converge back to the original target position and ultrasound imager direction. An example disturbance response test is shown in Fig. 2.16. The system was disturbed by moving the plastic tube before data collection. At $t = 0$ s the controller was activated and the catheter converged to the target pose. Data collection began at $t = 5$ s when the disturbance was removed (Fig. 2.16 (top)), causing a disturbance step input to the base of the bending section. This disturbance moved the catheter tip to the wrong pose (Fig. 2.16 (middle, bottom)). The controller then navigated the catheter tip back to the target pose. At $t = 17$ s the user commanded a $\Delta\psi = 2^\circ$ imager rotation (a target step input), changing the desired imager angle from 0° to 2° and maintaining the position at 0 mm from the target. The tip converged to the target pose with mean errors 0.5 mm ($\sigma = 0.3$ mm) and 0.18° ($\sigma = 0.12^\circ$).

2.5.2 Initial in vivo studies

In vivo animal testing was performed on a live porcine model due to similarities between human and porcine cardiac anatomy. The *in vivo* protocol was approved by the Institutional Animal Care and Use Committee (IACUC). The animal received humane care in accordance with the 1996 *Guide for the Care and Use of Laboratory Animals*, recommended by the US National Institutes of Health. An ICE catheter with two EM sensors was introduced through the 14 Fr introducer in the femoral vein to the right atrium (RA). A 6 Fr quadpolar catheter was outfitted with one EM sensor near the tip to serve as the imaging target. The target catheter was introduced through an 11 Fr introducer in the femoral artery to the left ventricle (LV).

In previous studies the introducer seal caused numerous problems: (1) Friction between the introducer seal and the catheter body caused nonlinear effects on steering motion. Mechanical solutions were built into the robot to improve joint motion efficiency. (2) During *in vivo*




Figure 2.17: *Introducer seal is bypassed using a brass tube filled with silicone.*

studies the EM tracker cables applied outward pressure opening the rubber seal, causing continuous blood leakage from the femoral vein. A saline drip was set up to apply gentle pressure against blood leakage. This replaced the blood leakage with a continuous saline leakage from the introducer seal.

Previous introduction methods were replaced with a new strategy. A brass tube (diameter 6 mm, length > 2 cm) was used to completely prop open the seal (Fig. 2.17), thereby eliminating the rubber-to-catheter friction from the introducer seal. Some friction from other sources still existed, but the negative effects of friction on steering motion were greatly reduced. The brass tube was fit into place by first sliding it over the catheter tip to rest at the proximal handle. After the ultrasound catheter was introduced to the RA, the tube was then forced into the seal. This process was borrowed from existing sheath and guidewire catheterization techniques. Silicone grease was packed into the tube around the ultrasound catheter body and EM sensor wires. This prevented blood from leaking through the introducer and eliminated the need for a saline drip. Silicone grease was also used in bench top experiments. This process was designed specifically for non-survival animal procedures in which an EM sensor must be attached to the exterior of the ultrasound catheter. This process can be changed or eliminated for an ultrasound catheter navigated by built-in EM sensors housed within the catheter tip, which are currently not accessible to us for proprietary reasons.

The subject's heart rate was paced at 120 bpm. The ventilator pumped 800 cc of air into

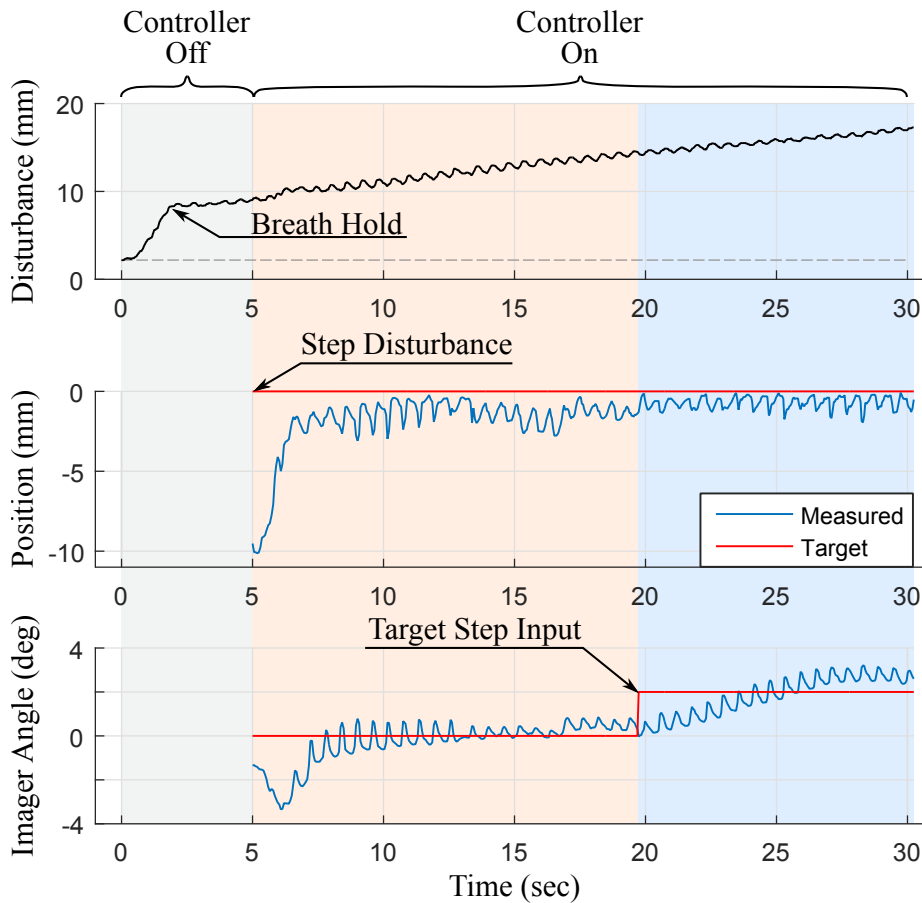


Figure 2.18: *In vivo* tests demonstrating the ability of the robotic system to maintain the catheter tip at the target pose despite disturbance step, disturbance ramp, and target step inputs.

the subject's lungs every 7.5 sec. To create a repeatable step disturbance input, respiration was temporarily paused on inhalation. While the respirator is paused, the pressure inside the lungs gradually increases (due to oxygen input, diaphragm relaxation, and other physiological effects), causing the heart to displace slowly with constant velocity as a ramp input.

The disturbance to the base of the catheter bending section is shown in Fig. 2.18 (*top*) where low frequency changes are due to respiratory motion and high frequency changes are due to cardiac motion. Before data collection, the controller was de-activated and the subject inhaled a breath, causing the catheter to displace 7 mm from the desired pose. Then the subject's breath was held for 30 seconds, causing the ramp input disturbance (shown in Fig. 2.18 (*top*) as the increasing line of constant slope). At $t = 5$ s the controller was

activated and overcame the step disturbance input from respiration to converge at the target pose. At $t = 20$ s the system was given a 2° step command of ultrasound imager adjustment. From $t = 5$ s to $t = 30$ s the ramp input continuously caused base displacement with an approximately constant velocity while the controller maintained the target tip pose. During the ramp input, the tip converged to the target pose and remained at the target pose with mean errors 1.1 mm ($\sigma = 0.7$ mm) and 0.44° ($\sigma = 0.31^\circ$). The controller was also able to avoid divergence from the target pose while cardiac motion caused higher frequency (1.67 Hz) disturbance.

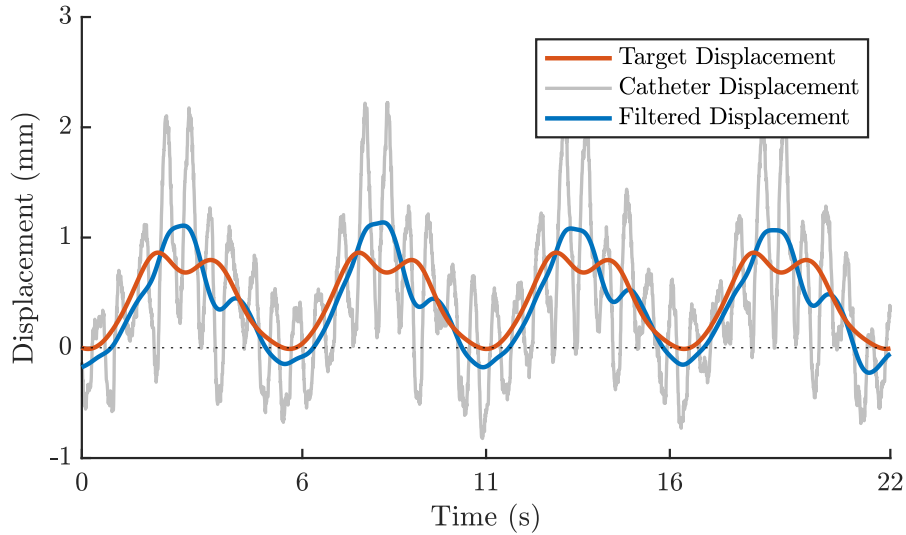


Figure 2.19: *In vivo respiration compensation trajectory.*

2.5.3 Instrument Tracking with Respiration Compensation

The goal of this test was to maintain the ICE catheter imaging plane aligned with a target ablation catheter by using respiratory motion compensation. The ablation catheter tip was in the LV. The ICE catheter was teleoperated to a region of the workspace and rotated to point the imaging plane roughly towards the target. Breathing models were initialized and then the robot was activated to begin tracking. The system maintained the same position of the catheter tip with respect to the breathing motion and continuously rotated the ultrasound




Figure 2.20: *In vivo* respiration compensation error histogram.




Figure 2.21: *Automatically pointing the ultrasound imager at a manually-steered target catheter in vivo.* (A) Robot points ultrasound imager at target. (B) Clinician moves target out of view. (C) Robot reorients ultrasound imager to continue pointing at target.

imager to align with the target.

This test was repeated multiple times. An example data set from a one-minute long trial is shown in Fig. 2.19. Based on the respiration model, the ICE catheter is commanded a target displacement (shown in orange) which would keep it in a fixed location with respect to the cardiac tissue. The ICE catheter response is shown in light gray. This response is low-pass filtered and shown in blue for easier comparison with the target displacement. Fig. 2.20 shows a histogram of the error between the target and achieved (unfiltered) displacement. The (unfiltered) position was maintained with a 0.53 mm root mean squared error (0.20 mm for filtered). Ultrasound images acquired during an ablation task is shown in Fig. 2.21 which demonstrates the utility and effectiveness of the instrument tracking function. Fig. 2.22 shows an ultrasound image from another ablation task in the left ventricle where the ablation site and the catheter tip are clearly seen.




Figure 2.22: *In vivo* ablation of the left ventricular wall.

2.6 Discussion

This chapter presented a method for robustly and accurately navigating flexible manipulators through unconstrained environments with external disturbances. The control strategy was designed to calculate the configuration space parameters of the catheter based on the current tip pose and base rotation that were measured using two EM sensors. This enabled navigation to be robust to friction and backlash nonlinearities as well as external disturbances to the catheter body.

Experimental results from bench top and *in vivo* studies demonstrate the robustness of the robotic system in steering the catheter tip to desired poses in the presence of unmodeled disturbances. Disturbances due to respiration were modeled using a fourth-order Fourier series, and the model period was updated using an EM sensor mounted on the chest. Using this model, accurate tracking of an ablation catheter was demonstrated during *in vivo* trials. Instrument tracking enabled automatically re-orienting the ICE catheter imaging plane towards a target

moving at respiratory speeds. Such effective instrument tracking can potentially reduce or completely eliminate the use of X-ray fluoroscopy in certain cardiac interventions. The methods presented here are also applicable towards compensating for higher frequency cardiac motions, however this was not demonstrated due to the mechanical limitation of the ICE catheter.

While the performance of the system is satisfactory, convergence time and error can be further improved through deadzone compensation, or by using catheters that are better built for robotic steering. The brass tube that was used to eliminate friction due to the introducer seal can be eliminated by integrating the EM sensors into the catheter body.

Chapter 3

Enabling Real-Time Ultrasound Reconstruction for Advanced Procedure Guidance

In the previous chapter we presented a robotic system that can autonomously steer cardiac catheters, demonstrating automatic positioning of the ultrasound imager and adjusting the direction of the imaging plane in the presence of physiological disturbances. This chapter describes how these capabilities are utilized to acquire, generate, and display wide field-of-view volumetric ultrasound images of cardiac anatomy. These volumes can help reduce the cognitive burden on the clinician, obviating the need to steer the ICE catheter while mentally piecing together a series of 2D ultrasound and fluoroscopy images. I developed a versatile software framework to enable concurrent execution of autonomous catheter steering, volume reconstruction, and rendering on a single workstation. We tested the system capabilities in two *in vivo* live animal trials.

3.1 Introduction

The system presented here is capable of autonomously acquiring 2D ICE images from within a beating heart, reconstructing 3D ultrasound volumes from these images in real-time (and 4D volumes in post processing), and displaying these high resolution volumes through real-time volume rendering, providing a clear view of cardiac structures and other interventional tools inside the heart (Fig. 3.1).

A review of freehand 3D ultrasound reconstruction algorithms can be found in [54]. Real-time 3D registration of ultrasound volumes was demonstrated in [55]. A wavelet-based fusion method was used in [56] to compound ultrasound volumes that were acquired from different ultrasound probe locations. Real-time 4D ‘mosaicing’ of 3D ultrasound images of moving targets is demonstrated in [57], where ECG-gating is used to split the ultrasound volumes into bins. 3D reconstruction of 2D ultrasound slices was demonstrated in a benchtop setting in [58] and *in vivo* in [59]. Researchers performed 3D reconstruction of ICE images to extend the field of view of ICE catheters in [60, 61]. A SLAM-like 3D reconstruction method was adopted in [62] to reconstruct 2D ICE images of a stationary phantom into 3D volumes in real-time during benchtop studies. The CARTO electroanatomical mapping system (Biosense Webster, Irvine, CA, USA) creates a 3D map of the heart based on position information from a mapping catheter [63]. The CartoSound module overlays ultrasound images acquired from an ICE catheter on this map [64], however the ultrasound image itself is not used to generate or update the map.

Acquiring, building, and rendering 4D (3D + time) ultrasound volumes of a beating heart is a computationally-challenging task, which requires the reliable operation of many diverse, intricate parts. In this chapter we continue to address the challenges that were outlined in Chapter 2:

1. Implementing robust and efficient algorithms to enable reliable and real-time execution,
2. Handling data input/output to/from multiple devices asynchronously at different sampling rates,




Figure 3.1: *The volume rendering graphical user interface enables real-time user interaction with a 4D (3D + time) ultrasound image of the anatomy. The ultrasound volume is reconstructed from a series of 2D ultrasound images that are acquired at different imaging angles throughout the cardiac cycle.*

3. Creating an intuitive and informative user-interface that enables quick debugging of the system, and
4. Reconstructing and rendering 4D ultrasound volumes.

3.2 Software Framework

Building robotic systems for surgical and interventional applications often involves integrating many devices and complex behavior. Over the years, a number of research teams developed software frameworks to enable robust, real-time operation of surgical robots. The CISST/SAW libraries is one such framework that is widely used [65]. It is a collection of libraries designed to ease the development of computer assisted intervention systems [66]. However, as programming languages evolve, and with the release of highly optimized libraries, these frameworks quickly become obsolete. These advances ultimately make developing new frameworks faster and reliable.

Our initial system development during the proof-of-concept stage of our system relied on

the Microsoft Foundation Class (MFC) library for the graphical user interface (GUI) elements. As we moved towards animal trials and added more functionality to our system (e.g. real-time reconstruction and visualization), we elected to create a new framework that could handle the computational challenges of the problem and maintain reliability through hour of testing during *in vivo* trials.

Our main goals in developing the software framework were to ensure:

- *High performance*: The algorithms, data structures, and libraries used should be selected to minimize computation time.
- *Minimal memory management*: Objects that go out of scope should automatically deallocate memory. This reduces the coding efforts and the likelihood of memory leaks, as well as improving code readability.
- *A responsive UI*: The graphical user interface should feel responsive and have an intuitive layout. It should support visualization tools such as real-time plotting and OpenGL for 3D scene rendering where the user can interact with the rendered volume (rotation, zoom, thresholding, etc.).
- *Modularity*: The class structure should allow individual submodules to be easily adopted in different projects, minimizing coding efforts in the future. This also means that the code should be cross-platform.

We elected to switch the development platform to Qt cross-platform application framework (the Qt Company, Espoo, Finland) in the interest of minimizing development time while maximizing functionality. Qt has many advantages: it is cross-platform, has built-in OpenGL support, it is well documented, has many useful GUI element prototypes that can be easily extended, and it is built on top of C++. It is widely used in industrial applications, including medical devices, therefore it is robust and well-tested. The signals and slots mechanism in Qt allows communication between objects, even across threads, obviating the need for messy callback functions. Unlike callback functions, signals and slots ensure type safety, since function signature mismatches will be caught at compile time. Signal and slot connections

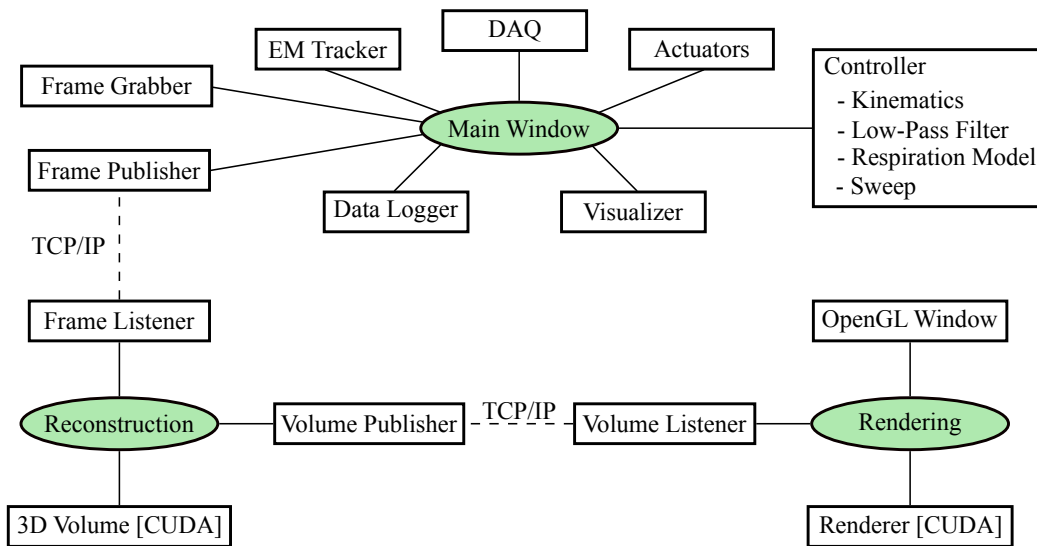


Figure 3.2: Class diagram for the overall system. User interface classes are indicated in green. Inter-process communication (e.g. Frame Publisher to Listener) is handled using TCP/IP connections.

across threads are established in a queued fashion, ensuring thread-safety. The signals and slots mechanism makes it possible to ensure modularity, since classes do not need access to each other’s methods. The signals and slots mechanism is similar to Lightweight Communications and Marshalling (LCM) and Boost Signals; in fact, Qt allows substituting LCM and Boost in place of the built-in signals and slots.

Fig. 3.2 shows a class diagram for our framework. For each device in the system, a worker class handles data I/O and device operation through function calls to the device application programming interface (API). Each device also has a ‘widget’ class associated with it, which implements a GUI element for user interaction with the device. These classes are modular, and therefore can easily be used in other applications. Each worker class lives in a separate thread to allow for concurrent code execution and communication with each of the devices. The GUI elements are always responsive even if the workers are busy. Data is transferred across classes and threads through the signals and slots mechanism.

The OpenCV library is used to acquire images from the frame grabber. In order to minimize data transfer overhead (deep copy), we use shared pointers to pass the image data to other threads. Saving the images to disk is one of the slowest operations due to the limited

data transfer bandwidth, therefore the Data Logger keeps a queue of shared pointers to the images that need to be saved. Once each image is saved, the shared pointer automatically goes out of scope, and the pointer count is decremented. Once the pointer count reaches zero, the image memory is automatically deallocated. Use of smart pointers greatly reduces the likelihood of memory leaks, code complexity, and potential for future bugs.

The Data Logger handles logging events, errors, and data from each device. The `QFile` and `QDataStream` classes are used for asynchronous file I/O. Data logging is also executed on a separate thread.

User interface elements are activated or deactivated based on a state machine model of UI operation. User, device, and system events transition the state, and these updates are propagated from the workers to the widgets through the use of the Qt signals and slots architecture.

The Eigen linear algebra library [67] is used for highly optimized computations in the controller, such as coordinate transformations, filtering, and singular value decomposition. Eigen utilizes SIMD instruction sets to vectorize operations. It also performs automatic loop unrolling to speed up computation at compile time.

During respiration model generation, the EM readings are low-pass filtered to filter out the cardiac motion. `QtConcurrent` is used for simultaneous low-pass filtering of multiple EM sensor readings to speed up the control loop.

The `QCustomPlot` library enables real-time plotting of signals, such as ECG and respiration. This functionality enables easy debugging and system monitoring for the engineer, and provides useful information to the clinician.

Communication between the main GUI and the reconstruction and rendering widgets is handled through TCP/IP connections, implemented using `QTcpSocket` and `QTcpServer` classes.

The ultrasound volume reconstruction is a compute-intensive but highly parallelizable operation. Thus, we offloaded this computation to an NVIDIA graphics processing unit (GPU). We wrote custom CUDA kernels to process millions of pixels from the ultrasound

images at once. We also use a GPU for real-time rendering of the ultrasound volumes using raycasting.

3.3 Ultrasound-based Procedure Guidance

The real-time, high frame rate images taken from within the heart using the ICE catheter can improve procedure guidance during cardiac interventions. Data and results presented in this section were acquired during *in vivo* studies with porcine animal models. The *in vivo* protocol was approved by the Institutional Animal Care and Use Committee (IACUC). The animals received humane care in accordance with the 1996 Guide for the Care and Use of Laboratory Animals, recommended by the US National Institutes of Health.

3.3.1 Ultrasound Volume Reconstruction

2D ultrasound images offer limited visibility of cardiac structures. In addition, piecing together a mental image from a series of 2D images is a cognitive burden on clinicians. Using our system, we can combine the 2D ultrasound images into a 4D (3D + time) volume to provide advanced procedure guidance to clinicians.

Data Acquisition

Data acquisition for volume reconstruction is automatically controlled by our system. The clinician positions the ICE catheter in a region of interest, then enters the angular spacing between slices, and the total desired sector angle. The system then controls the ICE catheter such that at least two cardiac cycles are observed at each sweep location, where the imaging plane is actively controlled to remain at a fixed heading while compensating for respiration. Then the system proceeds to the next imaging angle. Instead of performing a dense sampling where no interpolation is needed, we sample at set angular intervals, which reduces the sampling time. Less time is needed to reconstruct a volume, especially considering that at each sampling location we need to acquire samples for a few heart beats.

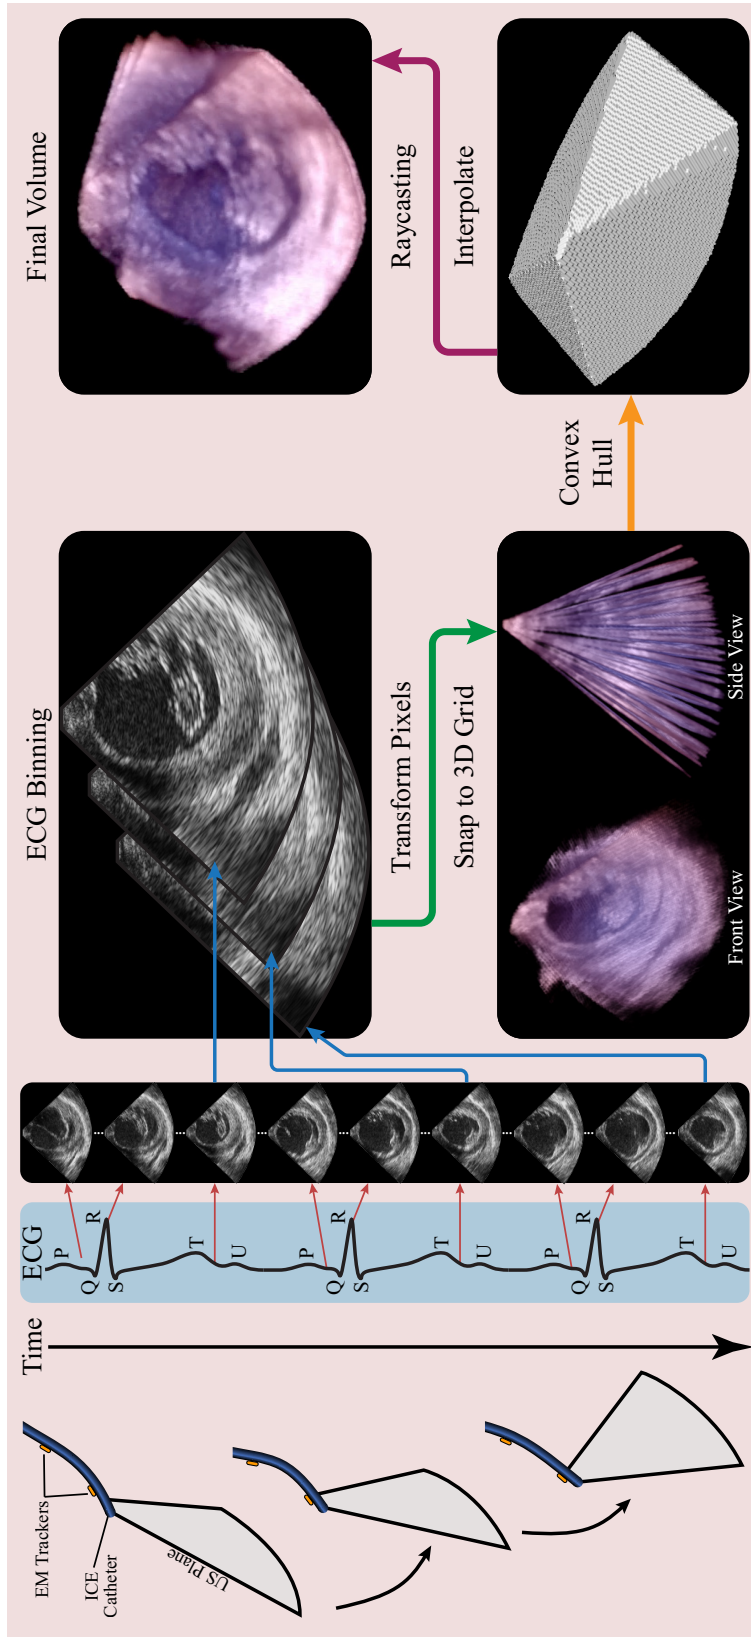


Figure 3.3: 3D reconstruction pipeline. Ultrasound images are binned according to their cardiac phase. Pixels are transformed in space based on the EM sensor readings and placed in a volumetric grid. The convex hull of these voxels defines the interpolation boundary, which is required to inpaint the gaps between the slices.

During acquisition, each image is labeled with its pose based on the EM sensor readings. In addition, based on the ECG signal, the image is also labeled with the percent cardiac phase it is acquired from (Fig. 3.3). During reconstruction, images with the same cardiac phase (within $\pm\frac{1}{2}$ bin size) are binned together. The number of bins depends on the temporal resolution desired by the clinician. In our studies, we elected to divide one cardiac cycle into fifteen bins. Our animal models had heart rates in the 100 to 120 bpm range, and ultrasound images were acquired at 30 fps, therefore we expected to acquire a minimum of 15 frames per cardiac cycle.

Reconstruction

After ECG binning, the Frame Publisher sends these images, the ECG label, and the EM data to the Reconstruction Widget over a TCP/IP connection. We separated the Controller, Reconstruction, and Rendering widgets into different processes and provided TCP/IP connections between them. This isolates each process from any errors or crashes that occur in the others, which is an important safeguard in safety-critical systems such as surgical robots. Even though these processes currently run on a single host machine, the TCP/IP connection allows these processes to be easily executed on different machines, which is a common practice in commercial systems.

The Reconstruction Widget transforms each pixel in the ultrasound images to a common coordinate frame. This computation is performed efficiently on the CPU using SIMD vectorization. The pixel intensities and coordinates are then transferred to the GPU. The pixels are assigned to voxels in a 3D grid. In cases when multiple pixels are assigned to a single voxel, the data collision is resolved based on maximum intensity.

Each pixel in the ultrasound image is transformed to a common coordinate frame based on the EM sensor readings.

$${}^{world}T_{US} = {}^{world}T_{EM} \cdot {}^{EM}T_{US} \quad (3.1)$$

Depending on the angular spacing between imaging angles prescribed by the clinician,

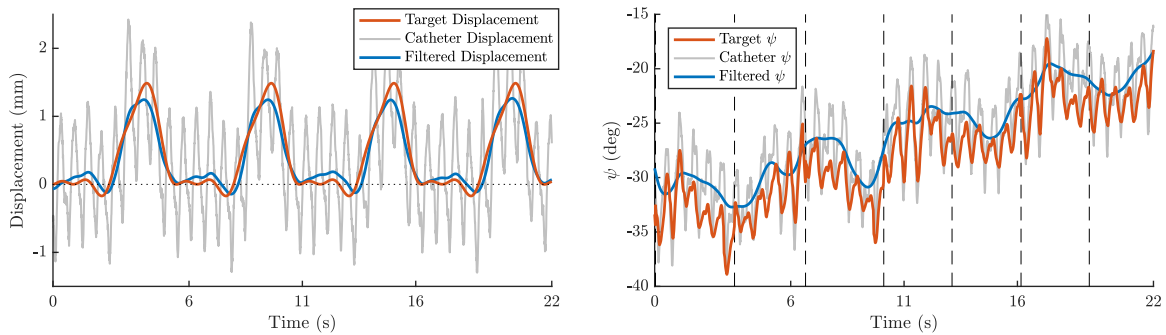


Figure 3.4: Catheter tip position and imaging plane angle during image acquisition for 4D reconstruction. The heart rate was 100 bpm and the respiration period was 5.4 s. The vertical dashed lines on the right indicate the instances when the catheter imaging plane converged to the target imaging angle. The angle $\Delta\psi$ between targets was 2° .

there can be gaps between the imaging planes, as shown in Fig. 3.3. These can be filled in using interpolation. We first find the convex hull of the filled voxels to determine the interpolation boundaries using the QHull library [68]. Afterwards, a simple nearest neighbor interpolation is performed on the GPU using a marching progression to fill the gaps in the volume. The interpolated volumetric grid is then saved to disk and transferred to the Rendering Widget over TCP/IP for visualization.

3.3.2 Rendering

The Rendering Widget has an OpenGL-based interface that allows a user to interact with the ultrasound volume, such as rotating or scaling it. The rendering algorithm is based on raycasting, where 10^6 ‘rays’ are cast from the observer’s position into the volume and the pixel information is accumulated along the ray based on a transfer function. The user can also control the parameters of this transfer function to enhance the detail of different regions. The false color is based on the depth of the voxels being rendered. Raycasting is also performed on the GPU as it is computationally intensive but highly parallelizable. We achieved 60 Hz update rates using this technique. Ultrasound volumes shown in Fig. 3.3, 3.5, and 3.6 were rendered using the Rendering Widget.




Figure 3.5: Fast intraoperative 3D reconstruction result (systole). The ICE catheter was positioned in the right atrium and imaged the left side of the heart. This volume rendered image is a view from the left atrium looking towards the bottom of the heart. The mitral valve is visible in the near field, and the left ventricle (positioned behind the mitral valve) can be seen through the aortic opening.

3.4 Results

The goal of the robotic system *in vivo* was to quickly navigate the catheter tip and ultrasound imager to a target pose, collect images of target structures, and then navigate to the next target pose. The controller was designed to actively converge towards the target pose. Ultrasound image acquisition was triggered once the catheter reached the target tip pose within a specified error threshold. During the *in vivo* study images for 4D reconstruction were collected at increments of 2° steps. Ultrasound images were collected through multiple heart beats and then the catheter was moved to the next target pose with a step input to the target angle. 30 ultrasound slices were collected in 96 seconds with motion compensation, during which the

mean absolute error was 1.57° ($\sigma = 1.06^\circ$) and 0.55 mm ($\sigma = 0.33 \text{ mm}$).

The desired and achieved catheter tip positions and imaging plane angles are shown in Fig. 3.4. The time series on the left shows the commanded target displacement of the ICE catheter (shown in orange), which is computed based on the respiration model. The controller tries to keep the ICE catheter tip in a fixed location with respect to the cardiac tissue. The ICE catheter tip response is shown in light gray. This response is low-pass filtered and shown in blue for easier comparison with the target displacement. Based on the catheter displacement (gray) we can see that the catheter is cyclically displaced due to cardiac motion at 1.67 Hz, which matches the 100 bpm heart rate of the subject. In addition, the filtered data (blue) shows an external cyclical disturbance due to respiration at roughly 0.2 Hz, which agrees with the respiration cycle period of 5.42 s that was detected by the respiration model (orange).

The plot on the right in Fig. 3.4 shows the catheter imaging plane angle. The vertical dashed lines indicate the instances when the catheter imaging plane converged to the target imaging angle. Once converged, the catheter was commanded to maintain heading and acquire ultrasound images for 1 s, and then to proceed to the next desired imaging angle. The angle $\Delta\psi$ between targets was 2° .

The acquired ultrasound images are reconstructed during the procedure into 3D or 4D volumes to enhance visualization of anatomical structures and working instruments. A reconstructed volume of the left ventricle and the mitral valve is shown in Fig. 3.5.

Table 3.1 shows the execution time of each operation in the 3D reconstruction pipeline. Discretization, data gridding, and interpolation all scale sub-linearly with the input and output size, clearly demonstrating the value of parallel execution on the GPU. These timings can be further reduced through optimized memory access. The slowest operation in the pipeline is finding the convex hull, which is currently executed on the CPU, and scales almost linearly with the data size. However the current performance still allows fast intraoperative reconstruction.

Our reconstruction pipeline on the GPU is limited to 3D at the moment. In post-processing, we extended this to 4D, where we used a more accurate, GPU-accelerated inpainting algorithm




Figure 3.6: Intraoperative 4D reconstruction results. (top) 2D ultrasound images of the mitral valve acquired throughout one cardiac cycle using an ICE catheter during an *in vivo* trial. (bottom) 4D (3D + time) reconstruction of the mitral valve and the left ventricle, generated in post processing using 30 spatial slices for each of the 15 temporal bins, with 2° spacing between slices.

based on the discrete cosine transform [69] to improve the interpolation quality. A subset of the reconstructed volumes are shown in Fig. 3.6. Top row shows the original 2D ultrasound images acquired by the ICE catheter. The bottom row shows the reconstructed volumes from a set of 30 slices with 2° spacing. The motion of the mitral valve and the ventricular wall is clearly visible. This reconstruction can be completed under three minutes.

Table 3.1: Execution time for various stages of the reconstruction pipeline with different input and output sizes.

Input (# pixels)	1,444,080	2,166,120	2,888,160
Output (# voxels)	128x64x128	128x64x128	128x128x128
Operation (device)	Execution Time (ms)		
Discretize (GPU)	0.88	1.12	1.17
Grid Data (GPU)	0.59	0.66	0.78
Convex Hull (CPU)	1,254.03	1,799.52	3,979.89
Interpolate (GPU)	168.94	171.97	243.11
Transfer (x2)	1.92	3.06	4.70
Total (s)	1.43	1.98	4.23

3.5 Discussion

In this chapter we presented methods for robotic acquisition and reconstruction of 3D and 4D ultrasound volumes from 2D ICE images. We presented results from *in vivo* live animal studies to demonstrate the performance and validity of our approach. We successfully generated a 3D volume of the left heart in real-time with our system. The robot control, volume reconstruction, and rendering algorithms can all run on a single host computer, demonstrating the computational efficiency of the system. The computationally-intensive but highly parallelizable steps of the reconstruction pipeline were performed on the GPU to accelerate reconstruction speeds. Our robotic system and volume reconstruction techniques can provide autonomous procedure guidance in cardiac interventions, potentially reducing the use of fluoroscopy and improving the efficacy of the procedures. Situational awareness in the clinic can be further improved by overlaying instrument location in the rendered ultrasound

volumes.

Throughout the development of the system, each part was tested rigorously to make sure that the overall system would survive the challenges of *in vivo* trials. Many robotic surgery research platforms do not make the transition from the benchtop to *in vivo* studies due to the complexities involved in preparing a system to face many unanticipated problems.

The EM sensors have a limited accuracy, therefore there is often an alignment error between the spatially-registered 2D images that are using during volume reconstruction. This can lead to blurry volumetric images and inaccurate representation of anatomy. This problem can be alleviated by extracting salient features from the 2D images and imposing a continuity constraint on these features (like a rubber band) between slices in 3D space to improve the alignment.

Chapter 4

Gaussian Process Regression for Ultrasound Scanline Conversion

In image-based procedure guidance, the resolution and accuracy of the medical image sets a fundamental limit on the precision of the guidance system. Therefore improving medical image quality has a direct impact on increasing the accuracy of feedback provided to clinicians and surgeons.

As described in Chapter 1, the piezoelectric crystal array inside an ultrasound probe acquires radio frequency (RF) data along scanlines. This scanline data is then processed, undergoing time gain compensation, low-pass filtering, envelope detection, log compression, and many other steps. Finally, the processed scanlines (A-mode) are interpolated such that a dense 2D image or 3D volume (B-mode) can be displayed to the clinician.

The real-time, high frame rate display requirements of ultrasound imaging has limited the type of interpolation that is used to generate B-mode images. Bilinear or cubic interpolation is often used to generate the final US image. These simple, local interpolation approaches ignore important data from more distant neighboring samples, leading to inaccuracies in interpolation. Therefore a more accurate interpolation method can improve the quality of image-based procedure guidance.

In a probabilistic framework, the data along the scanlines can be considered as samples

drawn from an underlying probability distribution – then the goal of the interpolation step is to approximate the values of unobserved pixels from this distribution. Statistical methods can better infer the underlying distribution, leading to more accurate interpolation.

In this chapter, we use Gaussian process regression to improve the accuracy of ultrasound scanline conversion. In \mathcal{GP} regression, a Gaussian distribution function is fit to each observation point. A covariance function is used to compute spatial correlation between observations. The parameters of this function can be tuned to capture different characteristic length scales that are specific to each dimension of the data. This adaptability is useful for ultrasound, since the observations along a scanline are dense, whereas the number of scanlines is small.

Using ultrasound scanlines acquired from two different ultrasound scanners during *in vivo* trials, we evaluate the performance of different interpolation methods. Using several quality metrics, we show that \mathcal{GP} regression outperforms the nearest neighbor, bilinear, and cubic spline interpolation methods. In addition to improved image quality, the variance estimate of \mathcal{GP} regression can be used to improve the accuracy of further image processing steps, such as registration.

4.1 Background

4.1.1 Previous Work on Improving Ultrasound Image Quality

Fast interpolation methods, such as bilinear and cubic interpolation, are often the de facto interpolation methods for B-mode ultrasound image generation. Most work on ultrasound image enhancement has focused on speckle reduction [70] and denoising [71]. A spatio-temporal smoothing filter was developed in [72] to improve scan conversion accuracy. Spline interpolation for 3D ultrasound volume compounding was explored in [73]. Kriging (another name for \mathcal{GP} regression in geostatistics) was used for 3D ultrasound reconstruction from multiple scan converted 2D slices in [74]. A comprehensive review of interpolation for 3D ultrasound compounding can be found in [75].

\mathcal{GP} regression has been used with other imaging modalities: single image super-resolution

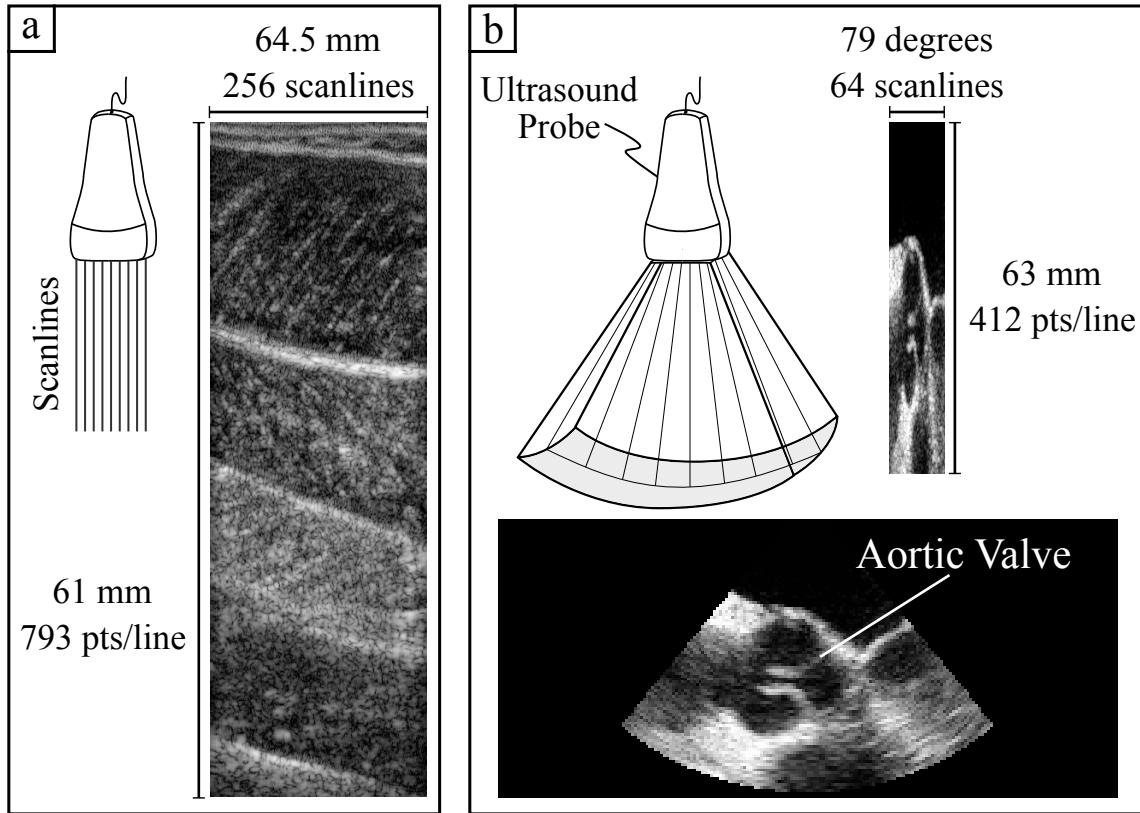


Figure 4.1: *In vivo* scanline data. (a) An ultrasound image of a posterior human lower leg was acquired using a probe with 256 parallel scanlines, with 793 points along each scanline. The gastrocnemius and soleus muscles are clearly visible in the image. (b) An ultrasound image of a human aortic valve was acquired using a 3D TEE probe with 64×48 divergent scanlines, with 412 points along each scanline. The image in the upper right corner is the scanline data, plotted in polar coordinates. The image on the bottom is scan converted to Cartesian coordinates using bilinear interpolation.

was explored in [76] for natural images; [77] explored MRI super-resolution; [78] used the variance estimates from \mathcal{GP} regression to improve MRI image registration. To the best of our knowledge, \mathcal{GP} regression has not been applied to ultrasound scanline interpolation before.

4.1.2 Statistical Models for Ultrasound Signals

It is important to understand the process that generates the ultrasound radio frequency signal in order to choose appropriate statistical models and parameters for the problem. Utilizing prior knowledge about the data distribution will lead to a more accurate regression.

A simple model for ultrasound signals is the Rayleigh distribution [79]. Assuming the




Figure 4.2: *The PDF, CDF, and 1-CDF of the K-distribution.*

imaging media has a large number of scatterers that are smaller than the wavelength of the ultrasound signal, and the magnitude and phase of the scattered ultrasound waves are independent and have a normal distribution, then their joint probability will be a Rayleigh distribution,

$$f(x) = \frac{x}{\sigma^2} \exp\left(-\frac{x^2}{2\sigma^2}\right) \quad (4.1)$$

where σ is the scale parameter.

The Rayleigh distribution assumptions only hold for diffuse scattering. However biological tissue usually also exhibits specular scattering due to the presence of larger scatterers, which leads to a deviation from the Rayleigh distribution. The K-distribution is proposed as a generalized model for scattering in [80, 81], which has the form

$$f(x) = \frac{x^{\nu+(n/2)-1} b^{\nu+(n/2)}}{2^{\nu+(n/2)-2} \Gamma(\nu) \Gamma\left(\frac{n}{2}\right)} K_{\nu-(n/2)}(bx), \quad (4.2)$$

where $\nu > -1$ is a shape parameter that describes the homogeneity of the tissue [82], $b = \frac{\sqrt{2n\nu}}{\rho}$ is the scaling parameter, n is the number of dimensions, ρ is the characteristic length scale, $\Gamma(\cdot)$ is the Gamma function, and $K_\nu(\cdot)$ is an order- ν modified Bessel function of the second kind. In the limit $\nu \rightarrow \infty$, the K-distribution approaches the Rayleigh distribution. In ultrasound imaging, we set $n = 2$, resulting in the following probability density function (PDF),

$$f(x) = 4 \frac{\sqrt{\nu}}{\rho \Gamma(\nu)} \left(\frac{\sqrt{\nu}}{\rho} x\right)^\nu K_{\nu-1}\left(\frac{2\sqrt{\nu}}{\rho} x\right) \quad (4.3)$$

where $x > 0$. The PDF of the K-distribution for different values of ν can be seen in Fig. 4.2.

Ultrasound data collected from 72 subjects in [83] show that the K-distribution fits the

data better compared to the Rayleigh distribution. In fact, the K-distribution is widely used to describe scattering phenomena with non-Rayleigh behavior, from sonar reverberation modeling in oceanic engineering to laser scattering in atmospheric turbulence [84].

We can utilize this prior on the ultrasound signal distribution through the use of non-parametric Bayesian regression to improve interpolation quality.

4.2 Methods

4.2.1 Gaussian Process Regression

A Gaussian process (\mathcal{GP}) is a collection of random variables, any finite number of which have a joint Gaussian distribution [85]. A Gaussian process is completely specified by its mean and covariance functions $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$. A zero mean function is commonly used in practice. Observations are considered as random variables that are drawn from a multivariate normal distribution

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}_* \end{bmatrix} = \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} \mathbf{K} + \sigma_n^2 \mathbf{I} & \mathbf{K}_* \\ \mathbf{K}_*^T & \mathbf{K}_{**} \end{bmatrix} \right) \quad (4.4)$$

where \mathbf{y} is a vector of scanline intensities, \mathbf{y}_* is a vector of pixel intensities at the interpolation locations, and σ_n^2 is the observation noise. The covariance matrix is a positive semi-definite symmetric matrix, where \mathbf{K} is the covariance between observations, \mathbf{K}_* is the covariance between observations and interpolation points, and \mathbf{K}_{**} is the covariance between interpolation points. The elements of the covariance matrix are computed using a covariance function (Section 4.2.4) $K_{ij} = k(x_i, x_j)$, $K_{*ij} = k(x_i, x_{*j})$, and $K_{**ij} = k(x_{*i}, x_{*j})$. \mathbf{y}_* and its variance can be estimated as

$$\bar{\mathbf{y}}_* = \mathbf{K}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} \quad (4.5)$$

$$\mathbb{V}[\mathbf{y}_*] = \mathbf{K}_{**} - \mathbf{K}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{K}_* \quad (4.6)$$

From Eq. (4.5) we can see that \mathbf{y}_* is in a reproducing kernel Hilbert space (RKHS) with reproducing kernel $k(\mathbf{x}, \mathbf{x}')$ [85].




Figure 4.3: (left) The interpolation points (colored) for generating a B-mode image lie on a grid, however the observations that contribute to each pixel lie along an arc (light blue). (right) Mapping these interpolation points to polar coordinates resolves this warping of the interpolant, but results in a non-uniform spacing between interpolation points. Scan conversion is often performed in polar coordinates, then the interpolated points are mapped back into Cartesian coordinates.

4.2.2 Scanline Conversion - Parallel and Divergent Scanlines

In ultrasound probes, the scanlines can be either parallel to each other or spread out radially from the source (divergent). With parallel scanlines, the data points are in Cartesian coordinates, and it is straightforward to interpolate and display them on a user’s screen (Fig. 4.1a).

In contrast, with divergent scanlines, the scanline data is in polar coordinates, but the clinician is often interested in the Cartesian space representation for ease of spatial reasoning (Fig. 4.1b). Fig. 4.3 shows the arrangement of scanlines in both polar and Cartesian coordinates. The desired interpolation points (colored) lie on the image grid in Cartesian space. Notice that the density of observations in Cartesian space gets sparser as the distance from the ultrasound transducer (r) increases. In practice, when using linear interpolators, scan conversion is performed in polar coordinates because the observations lie on a grid. The interpolated intensity values can be then simply displayed in Cartesian coordinates [72]. However this method ignores the increasing distance between observations due to the mapping between polar and Cartesian space. In order to achieve better accuracy, the interpolation method should account for this change in distance, capturing the underlying process better. In Section 4.2.4 we will discuss how this can be achieved with \mathcal{GP} regression by using different kernel lengths.

4.2.3 Validation - Measuring Interpolation Accuracy

In order to validate our claim that \mathcal{GP} regression leads to higher interpolation accuracy, we compare the performance of commonly used interpolation methods (nearest neighbor, bilinear, and cubic spline interpolation) against the performance of \mathcal{GP} regression. We conducted leave-N-out studies, where only N_s scanlines were kept out of the total N_t scanlines. The missing data was then filled in using interpolation. We computed the mean of absolute errors (MAE) and the peak signal-to-noise ratio (PSNR) between the ground truth and the interpolated images.

The MAE is defined as

$$\text{MAE} = \frac{1}{I} \sum_i^I |Z_i^* - \hat{Z}_i|, \quad (4.7)$$

where Z_i^* are the intensities of the ground truth scanlines, \hat{Z}_i are the intensities of the interpolated scanlines, and I is the total number of observations in the ultrasound image.

Assuming that the scanline intensities are normalized to $\mathbb{R} : [0, 1]$, the PSNR is defined as

$$\text{PSNR} = 10 \log_{10} \left(\frac{1}{\text{MSE}} \right), \quad (4.8)$$

where

$$\text{MSE} = \frac{1}{I} \sum_i^I (Z_i^* - \hat{Z}_i)^2. \quad (4.9)$$

\mathcal{GP} regression should ideally have the highest PSNR and lowest MAE and MSE among the interpolation methods tested.

4.2.4 Covariance Functions

Covariance functions (i.e. kernels) capture the spatial dependence of observations. There are various kernels that can be used to compute the covariance matrix. A comprehensive list can be found in [85]. In this work, we consider two covariance functions: the squared exponential and the Matérn.

Kernels usually have a parameter that controls the characteristic length scale of observation locations. It is possible to set a different length scale for each dimension of the data, which is




Figure 4.4: *Leave-N-out studies.* The dotted lines indicate ultrasound scanlines. The scanlines indicated in red are removed from the dataset. Nearest neighbor, bilinear, and cubic spline interpolation, and GP regression is used to estimate the pixel intensities at these locations. The estimates are compared to the ground truth values in the full dataset.

useful when dealing with a non-isometric sampling density, such as in the case of ultrasound, especially in the case of divergent scanlines.

Squared Exponential

The squared exponential (SE) covariance function (also called the radial basis function (RBF) kernel) is an infinitely differentiable function, which results in a very smooth process. The SE covariance function with separate length scales for each dimension is expressed as,

$$k(\mathbf{x}_i, \mathbf{x}_j | \theta) = \sigma^2 \exp \left[- \sum_{n=1}^2 \frac{(x_{i,n} - x_{j,n})^2}{2l_n^2} \right], \quad (4.10)$$

where θ are the function parameters (σ and l_n), σ is the signal variance, l_n is the characteristic length scale for dimension n , and $x_{i,n}$ is the n^{th} component of the i^{th} observation.

In our case, we will denote the length scale along the scanlines as l_r , and normal to the scanlines as l_θ .

Matérn

The Matérn autocovariance function has the form

$$k(x_i, x_k | \theta) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{l} r \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{l} r \right), \quad (4.11)$$

where r is the distance between x_i and x_j , and l is the characteristic length scale, and ν is the smoothness parameter. In D dimensions, the distance metric is expressed as

$$r(\mathbf{x}_i, \mathbf{x}_j) = \sqrt{\sum_{n=0}^D \frac{(x_{i,n} - x_{j,n})^2}{l_n^2}}. \quad (4.12)$$

According to [86] the strong smoothness assumptions of the SE covariance function are unrealistic for modeling many physical processes, and therefore the Matérn class is recommended [85]. The Matérn covariance function is $\lceil \nu \rceil - 1$ times differentiable. In general, ν is chosen to be a half-integer, which reduces Eq. (4.11) to an exponential multiplied by a polynomial and the spectral density becomes rational [86]. When $\nu = 1/2$, the Matérn covariance function reduces to the exponential covariance function, and as $\nu \rightarrow \infty$ the kernel converges to the SE covariance function. In literature, ν is usually set to 3/2 or 5/2 [85].

Now let us reconsider the PDF of the K-distribution (Eq. (4.3)), restated here for convenience,

$$f(x) = 4 \frac{\sqrt{\nu}}{\rho \Gamma(\nu)} \left(\frac{\sqrt{\nu}}{\rho} x \right)^\nu K_{\nu-1} \left(\frac{2\sqrt{\nu}}{\rho} x \right)$$

where $x > 0$. This function and the Matérn covariance function given in Eq. (4.11) look similar. In fact, the cumulative density function (CDF) of the K-distribution is,

$$F(x) = \int_0^x f(u) du = 1 - \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{2\sqrt{\nu}}{\rho} x \right)^\nu K_\nu \left(\frac{2\sqrt{\nu}}{\rho} x \right). \quad (4.13)$$

Setting $\rho = \sqrt{2}\rho^*$,

$$1 - F(x) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{\rho^*} x \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{\rho^*} x \right), \quad (4.14)$$

which gives us the Matérn covariance function with unit variance ($\sigma^2 = 1$).

In our experiments, we observed that the Matérn covariance function with $\nu = 3/2$ leads to better interpolation compared to the SE kernel.




Figure 4.5: *Incorrect choice of the kernel length can lead to reconstruction errors (blurry images or missing data).*

4.2.5 Finding the Optimal Length Scale

In general, the kernel parameters are determined by minimizing the log marginal likelihood of the \mathcal{GP} ,

$$\log p(y | X) = -\frac{1}{2} \mathbf{y}^\top (K + \sigma_n^2 I)^{-1} \mathbf{y} - \frac{1}{2} \log |K + \sigma_n^2 I| - \frac{N}{2} \log 2\pi, \quad (4.15)$$

where N is the number of observations. However, since we have ground truth data, we instead opted to optimize the kernel parameters through a constrained optimization problem with the MSE metric as the cost function, described as

$$\begin{aligned} \min_{l_n} \quad & \frac{1}{I} \sum_i (Z_i^* - \hat{Z}_i^{\mathcal{GP}})^2 \\ \text{s.t.} \quad & l_n > 0, \quad n = \alpha, r \end{aligned} \quad (4.16)$$

where Z_i^* is the ground truth pixel intensity and $\hat{Z}_i^{\mathcal{GP}}$ is the \mathcal{GP} estimate. In our optimization studies, we noticed that the landscape is convex, however the gradients near the optimum are small.

Choosing a larger than optimal length scale result in blurring, while shorter length scales lead to unfilled regions between the scanlines (Fig. 4.5), both of which result in a decrease in the PSNR.

The kernel parameters correspond to physical parameters of the underlying process. Therefore we hypothesized that we can use our knowledge about the distances between data points to determine the kernel parameters. The distance between observations along the scanlines, Δr , and the distance between the scanlines, $|r\Delta\alpha|$, are constants that are determined

by the ultrasound probe and imaging settings.

In the case of divergent scanlines, our confidence along the rays r is constant, however as the scanlines diverge, we have less observations per unit area. Therefore our confidence between scanlines (in the α axis) should decrease. This can be achieved by scaling l_α by $1/r$, which reduces the characteristic length scale in the α -axis, thus increasing the uncertainty of the pixels in between the scanlines.

Inspection of the optimized kernel lengths led to the discovery of following relationships for the kernel lengths:

- For parallel scanlines,

$$l_\alpha = \frac{1}{\sqrt{3}} \Delta\alpha \quad (4.17)$$

$$l_r = k_N \Delta r, \quad (4.18)$$

where $\Delta\alpha$ and Δr are the (standardized) distances between observations (i.e. voxel size), k_N is the reduction factor for leave- N -out studies. For example, if only $N/4$ scanlines are used in the interpolation, then $k_N = 4$.

- For divergent scanlines,

$$l_{\alpha|\bar{r}_i} = \frac{1}{\sqrt{3}} \frac{|\bar{r}|}{2\bar{r}_i} \Delta\alpha \quad (4.19)$$

$$l_r = k_N \Delta r, \quad (4.20)$$

where $|\bar{r}|$ is the overall (standardized) scanline length.

4.2.6 Reducing Computational Cost - Patched \mathcal{GP} Regression

\mathcal{GP} regression has $O(N^3)$ memory and $O(N^2)$ time complexity (where N is the number of pixels) [85], which is perhaps the main reason why it has not been widely adopted in clinical imaging. Medical images can contain millions of pixels or voxels, making it computationally challenging to perform \mathcal{GP} regression in real time.

The main bottleneck in the algorithm is in inverting the covariance matrix. The covariance

function gives us a measure of the spatial dependence between observations. With certain covariance functions (such as the ones used in this work), the covariance function quickly decays as the distance between observation pairs increases. Therefore, outside of a local region, the contributions from other observations can practically be ignored. We utilize this feature to reduce the computational demands of \mathcal{GP} regression. We split the scanline data into overlapping patches and evaluate \mathcal{GP} regression separately for each patch. We chose a window size of 15 observations, since larger window sizes only had a small impact on our quality metrics. Note that this procedure can be massively parallelized on a GPU or FPGA to process the entire image in a single pass.

Also note that the computation time and costs can be greatly reduced by pre-computing the matrix $\mathbf{K}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1}$ in Eq. (4.5) and \mathbf{K}_{**} in (4.6). For a given probe and imaging depth, the interpolation locations are not going to change, therefore these matrices can be precomputed and stored in memory to avoid performing the costly Cholesky decomposition in real time. This reduces the computational complexity to just $O(NM)$ for the matrix-vector multiplication in Eq. (4.5), where $M < N$ is the number of observations.

4.3 Experiments and Results

We acquired *in vivo* scanline data from two different ultrasound imaging systems. 2D images of the lower leg of a healthy male subject were acquired at 61 mm imaging depth and 6 MHz frequency using a linear-array ultrasound probe with parallel scanlines (LV8-4L65S-3, TELEMED Ltd, Vilnius, Lithuania). The 256 scanlines with 793 observations along each scanline are shown in Fig. 4.1(a). In another experiment, 3D divergent scanline data of a human heart was acquired at 63 mm imaging depth and 7 MHz frequency using a phased-array 3D transesophageal echocardiography (TEE) probe (X7-2t transducer connected to an iE33 ultrasound imaging system, Philips Healthcare, Andover, MA, USA). A slice from the 3D dataset with 64 scanlines and 412 observations along each scanline are shown in Fig. 4.1(b). These studies were approved by the Harvard University Institutional Review Board (IRB) and the Boston Children’s Hospital IRB.




Figure 4.6: Scan converted B-mode images of the parallel scanline dataset. Leftmost column shows the ground truth data with 256 scanlines. Rightmost column shows the normalized variance matrix generated during GP regression, where darker regions indicate higher confidence in the estimates.




Figure 4.7: Interpolator performance for scan conversion of parallel scanlines. (a) absolute PSNR scores (higher is better), (b) relative PSNR scores (computed by subtracting the lowest scores from the rest, higher is better), and (c) MAE scores (lower is better).

The observation noise σ_n^2 was set to 2×10^{-3} through a hyperparameter optimization of the \mathcal{GP} . We used the imaging processing toolbox and the statistics and machine learning toolbox in MATLAB (MathWorks, Natick, MA, USA) in our calculations.

4.3.1 Parallel Scanlines

Fig. 4.6 shows the interpolated B-mode images for three different leave-N-out cross-validation studies. Leftmost column shows the ground truth data. The variance map on the right shows the normalized values of the variance reported by \mathcal{GP} regression from Eq. (4.6), where darker regions indicate smaller variance (i.e. higher confidence). We can see that the amount of uncertainty increases as the number of scanlines is reduced. Fig. 4.7 shows (a) the absolute PSNR scores, (b) the relative PSNR scores (computed by subtracting the lowest scores from the rest), and (c) the MAE scores. Nearest neighbor interpolation has the largest error as expected, and \mathcal{GP} regression has the lowest error and highest PSNR.

As the number of observations (scanlines) is reduced, we see a wider separation in performance between \mathcal{GP} regression and the other interpolation methods. Looking at the relative PSNR scores, we see that the scores for the bilinear and spline methods have a negative slope that tend towards the performance of the nearest neighbor interpolation. However \mathcal{GP} regression maintains a constant separation from the baseline (nearest neighbor).




Figure 4.8: Scan converted B-mode images of the divergent scanline dataset, presented in polar coordinates. Leftmost column shows the ground truth data with 64 scanlines. Rightmost column shows the normalized variance matrix generated during GP regression.




Figure 4.9: Interpolator performance for scan conversion of divergent scanlines. (a) absolute PSNR scores (higher is better), (b) relative PSNR scores (computed by subtracting the lowest scores from the rest, higher is better), and (c) MAE scores (lower is better).

4.3.2 Diverging Scanlines

Fig. 4.8 shows the interpolated B-mode images for three different leave-N-out cross-validation studies in polar coordinates. Leftmost column shows the ground truth data. The variance map on the right shows the normalized values of the variance reported by \mathcal{GP} regression from Eq. (4.6). We can see that the amount of uncertainty increases as the number of scanlines is reduced. The variance also increases along the scanlines due to the $1/r$ scaling of the l_α parameter. Fig. 4.9 shows (a) the absolute PSNR scores, (b) the relative PSNR scores (computed by subtracting the lowest scores from the rest), and (c) the MAE scores. Nearest neighbor interpolation has the largest error as expected, and \mathcal{GP} regression has the lowest error and highest PSNR. \mathcal{GP} regression exhibits a larger increase in performance for divergent scanlines compared to the linear scanlines.

We observe an even wider separation (compared to parallel scanline results) in performance between \mathcal{GP} regression and the other interpolation methods as the number of observations (scanlines) is reduced. Looking at the relative PSNR scores, we see that the scores for the bilinear and spline methods again have a negative slope that tend towards the performance of the nearest neighbor interpolation. Differently from the parallel case, the relative PSNR score for \mathcal{GP} regression has a peak around 13 scanlines, then rapidly decreases near the performance of bilinear interpolation. However, we should notice that this occurs when the number of scanlines is less than 15, which is the window size we chose to reduce the computation time. This suggests that we simply do not have enough data points to accurately infer the underlying




Figure 4.10: Scan converted B-mode images of the divergent scanline dataset, presented in Cartesian coordinates. Leftmost column shows the ground truth data with 64 scanlines. Rightmost column shows the normalized variance map generated during GP regression.

distribution of the data, thus \mathcal{GP} regression performance drops rapidly.

Fig. 4.10 shows the interpolated B-mode images for three different leave-N-out cross-validation studies in Cartesian coordinates. Leftmost column shows the ground truth data. The variance map on the right shows the normalized values of the variance reported by \mathcal{GP} regression from Eq. (4.6). We can see that the amount of uncertainty increases as the number of scanlines is reduced. The variance also increases along the scanlines due to the $1/r$ scaling of the l_α parameter. At $N = 13$ and $N = 8$, the B-mode images generated by \mathcal{GP} regression are visibly clearer compared to the other interpolation methods. Bilinear interpolation exhibits the laterally elongated features that is characteristic to ultrasound images seen in commercial ultrasound systems.

4.4 Discussion

In this chapter we investigated the use of Gaussian process regression for ultrasound scanline interpolation. We analyzed *in vivo* ultrasound data acquired using both parallel and divergent scanline probes to validate the performance of \mathcal{GP} regression. Using leave-N-out cross-validation, we qualitatively and quantitatively showed that \mathcal{GP} regression leads to better B-mode conversion than other interpolation methods, indicated by the higher PSNR and lower MAE scores. We also showed that the covariance matrix can be precomputed, greatly reducing the computational complexity of the problem. In our implementation, the average time to estimate the value of one data point was 0.3ms. It should be possible to enable real-time execution of \mathcal{GP} regression through an optimized implementation.

We demonstrated that \mathcal{GP} regression leads to better scanline conversion using real clinical data that was acquired from different organ systems. The images we collected span a wide range of clinical cases. The leg dataset exhibits high frequency texture and has a dense set of scanlines, whereas the cardiac dataset comprises uniform patches of dark (blood) and bright (muscle) regions and has a sparse set of scanlines.

One open question in our study is whether the kernel length relationships that we identified hold for different anatomy. Kernel lengths should be optimized for ultrasound images acquired

from different regions of the body, and these optimal lengths should be compared with the relationships we identified.

Choosing a covariance function prescribes a prior on the spatial variation of the data. Even though the Matérn_{3/2} function resulted in lower errors compared to the squared exponential kernel, use of other covariance functions should be investigated as well. Furthermore, a gamma process may recover the underlying distribution better than a Gaussian process.

In clinical practice, better quantification leads to more informative clinical decision making, and quantification depends heavily on the image quality. In the case of parallel scanlines, visual inspection of the interpolated images does not reveal significant improvements in image quality. This may raise concerns over the utility of using \mathcal{GP} regression. However the PSNR scores indicate a marked improvement. This, in unison with the utility of using the variance matrix generated by the \mathcal{GP} , indicates that using \mathcal{GP} regression can lead to improvements in the performance of image processing steps that are further down the pipeline by avoiding compounding of errors.

Chapter 5

High Dynamic Range Ultrasound Imaging

High dynamic range (HDR) imaging is a popular computational photography technique that has found its way into every modern smartphone and camera. In HDR imaging, images acquired at different exposures are combined to increase the luminance range of the final image, thereby extending the limited dynamic range of the camera.

Ultrasound imaging suffers from a limited dynamic range as well; at higher power levels, the hyperechogenic tissue is overexposed, while at lower power levels hypoechogenic tissue details are not visible. In this chapter, we apply HDR techniques to ultrasound imaging, where we combine ultrasound images acquired at different power levels to improve the level of detail visible in the final image. We present results for images acquired from *ex vivo* and *in vivo* tissue and demonstrate that HDR-US enables visualizing both hyper- and hypoechogenic tissue at once in a single image. The performance of five tone mapping operators is quantitatively evaluated to determine the most suitable mapping for HDR-US imaging. Our results clearly demonstrate that HDR-US imaging can improve the utility of ultrasound in image-based diagnosis and procedure guidance.

5.1 Motivation

Ultrasound imaging is an invaluable tool for medical diagnosis and intraoperative guidance. Images are acquired in real time at high frame rates, with no radiation exposure, and at much lower expense than CT and MRI. However, ultrasound imaging suffers from a limited dynamic range: at higher power levels, rigid surgical tools and hyperechogenic tissue in the imaging field get overexposed (Fig. 5.1c), while at lower power levels hypoechogenic tissue details are lost (Fig. 5.1b).

There is extensive literature on image processing methods for improving ultrasound image quality. Some popular techniques are speckle and noise reduction through spatial and frequency-domain filters [71], detail enhancement through spatial and frequency compounding [70, 87], and contrast enhancement using histogram equalization or tissue harmonic imaging. See [88] for a review of ultrasound image enhancement techniques. However, these methods do not address the issue of limited dynamic range.

Similar to ultrasound imaging, optical cameras have a limited dynamic range as well. Imaging shadows requires a long exposure, while imaging highlights requires a short one. High dynamic range (HDR) imaging is a method which extends the dynamic range of optical cameras by combining images taken at multiple exposures [89]. In the last few years, this method has become ubiquitous in smartphones and digital cameras.

In this chapter, we show how HDR imaging techniques can be applied to ultrasound images to improve their dynamic range, enabling the imaging of media over a wide range of reflection and attenuation characteristics. We also compare several tone mapping methods, which allow the HDR-US images to be displayed on regular (*i.e.*, low dynamic range) screens, while preserving the visibility of features from across the entire dynamic range. This leads to better visualization for clinicians and can improve the performance of subsequent image processing operations. Therefore the extended dynamic range of HDR ultrasound (HDR-US) can allow for better diagnosis and procedure guidance.

Extending the dynamic range is of interest to other medical imaging modalities as well. A similar treatment was applied to MR imaging in [90], but there was no investigation of tone




Figure 5.1: (a) *Experimental setup:* A phased linear array ultrasound probe (Philips 15-6L) images a chicken thigh at 4 cm depth while an ablation catheter (ThermoCool SF, Biosense Webster, Diamond Bar, CA) is pressed against the tissue. (b) *B-mode ultrasound image acquired at power level -18.0 dB.* The catheter body and the skin surface exhibit strong reflection, showing a clear image of the contact between the tissue and the instrument, while the deeper tissue and the bone surface have poor detail. (c) *Ultrasound image acquired at power level 0.0 dB.* The bone surface and the hypodermis details are pronounced, whereas the instrument and skin surface details are lost due to overexposure as well as ringing from the instrument. Scale on the right shows the estimated drop in signal strength due to attenuation and reflection through the catheter and tissue layers.

mapping in this study.

5.2 Methods

5.2.1 Ultrasound Image Formation

We begin by considering the process that generates the ultrasound image. The intensity of a pixel in the B-mode ultrasound image is proportional to the pressure of the reflected sound signal. The pressure of a sound wave that is scattered by a target away from the transducer can be expressed as

$$p_s(r, t) = A_0 \frac{R e^{-\mu_a(d+r)}}{r} F\left(t - \frac{d+r}{c}\right) \quad (5.1)$$

where r is the radius from the target, t is time, A_0 is the amplitude, R is the reflection coefficient, μ_a is the amplitude attenuation factor, d is the distance between the target and the transducer, $F(\cdot)$ is the acoustic envelope, and c is the speed of sound [91]. Increasing the amplitude increases the pressure. The only constant intrinsic to the ultrasound transducer in this equation is the amplitude; the rest of the terms are specific to the media being imaged.

Table 5.1: *Acoustic properties of relevant media [91].*

Property	Water	Fat	Muscle	Bone	Aluminum
Z [kg/m ² / s ¹] × 10 ⁶	1.5	1.35	1.7	5.5	17
a [dB/cm ¹]	0.0022	0.63	1.8	20	0.018
b	2	1	1	1	1

The relative amplitude loss of the returning signal due to attenuation and reflection losses can be calculated as

$$\frac{A_z}{A_0} = \sqrt{R_I} e^{-\mu_a z} \quad (5.2)$$

where R_I is the intensity reflectivity, and z is the total distance traveled by the signal. Assuming a normal incidence at the boundary between two media, the reflected power loss in dB can be calculated as

$$20 \log_{10} \frac{A_z}{A_0} = 20 \log_{10} \left| \frac{Z_2 - Z_1}{Z_2 + Z_1} \right| e^{-\mu_a z} \quad (5.3)$$

where Z_1 and Z_2 are the acoustic impedance of media 1 and 2. The amplitude attenuation factor is a function of the signal frequency f and media-specific values a and b , written as $\mu_a = \frac{a f^b}{20 \log_{10} e}$. Fig. 5.1c shows the estimated drop in signal strength due to attenuation and reflection through the catheter and tissue layers. The acoustic impedance and the attenuation coefficients were taken from [91].

5.2.2 High Dynamic Range Formulation

In photography, the intensity Z of a pixel in a photograph is a nonlinear function of the exposure X , expressed as $Z = f(X)$. Malik and Debevec [92] derived a method to recover the response curve of a camera from multiple exposures. Exposure can be defined as $X = E\Delta t$, where E is the sensor irradiance, and Δt is the exposure time. Assuming f is monotonic, we can write $f^{-1}(Z) = E\Delta t$. Taking the logarithm of both sides, we have $g(Z) = \log f^{-1}(Z) = \log E + \log \Delta t$. The response curve g and the irradiances E_i can be then recovered by

minimizing the loss function

$$\mathcal{L} = \sum_i^N \sum_j^P [g(Z_{ij}) - \log E_i - \log \Delta t_j]^2 + \lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} g''(z)^2, \quad (5.4)$$

where N is the number of pixels in each image, P is the number of photographs, and λ is a regularization parameter that controls the smoothness of g [92]. The response curve is then used to map the pixel intensities in each low dynamic range (LDR) image to a relative radiance value, by rearranging the expression for $g(Z)$ given above to $\log E_i = g(Z_{ij}) - \log \Delta t_j$ [92]. In practice, radiance is computed using a weighted average of all exposures from the input images, which increases robustness and reduces noise. Once the response curve is computed, it can be used with any image taken with the system and does not need to be recomputed, given other imaging settings remain constant.

A similar response curve is recoverable for ultrasound probes. In Eq. 5.1, the only term that depends on the ultrasound transducer is the signal amplitude, A_0 . Grouping all other terms in U , which describes the scattering and attenuation properties of the medium, Eq. 5.1 can be written as $p_s = UA_0$. Assuming that the mapping f from the pressure of the scattered wave to the pixel intensity Z in the ultrasound image is monotonic, we can write $f^{-1}(Z) = p_s = UA_0$. This has the same form as $f^{-1}(Z) = E\Delta t$.

The reflection equation (Eq. 5.2) is also linear in the amplitude, but it depends on the frequency as well. Assuming that the ultrasound probe is held stationary, kept at the same imaging depth and frequency, and the target did not move during imaging, a change in source




Figure 5.2: Response function recovered from the fifteen ultrasound images acquired at different power settings.

signal amplitude should result in a linear change in the received signal amplitude. This is similar to keeping the camera stationary with the same aperture and focus while changing the exposure duration.

Ultrasound machines generally express acoustic power in decibels (dB). This is not an absolute measure of power, but rather a logarithmic ratio between two amplitudes. We can convert from power to a relative amplitude A_r using

$$A_r = 10^{(Power/20dB)}. \quad (5.5)$$

This quantity can be used in the HDR calculations. Fig. 5.2 shows the response curve recovered using fifteen ultrasound images acquired at different power levels. The curve has a form that is very similar to that of digital cameras computed using natural photos (see [92] for examples). The nonlinearity in this function is due to nonlinearities in the transducer behavior, as well as the log compression applied during B-mode image formation in the ultrasound machine.

The recovered response curve is then used to compute $\log U_i = g(Z_{ij}) - \log A_j$, which can be interpreted as mapping B-mode image pixel intensities to the relative scattering and attenuation characteristics of the tissue.

5.2.3 Tone Mapping – Display of HDR-US Images

Once the HDR ultrasound image is computed, there remains the problem of displaying the information to the user on a low dynamic range (LDR) screen. The human ocular system can only resolve up to 30 dB, and conventional monitors are limited to two orders of magnitude of dynamic range. Mapping an HDR image to the LDR is called tone mapping, where the goal is to preserve details and color appearance from the HDR image, which would be lost in a simple linear compression of pixel intensities. There is extensive work on tone mapping for HDR images; a detailed treatment is presented in [93]. Note that this step is only necessary if the HDR-US images will be shown to the clinician for diagnosis and procedure guidance. If the images are instead going to be used in an image processing pipeline, the rich HDR information can be retained to improve performance.

Tone mapping operators (TMOs) are often distinguished into *global* and *local* operators. Global operators find a single compression function that changes the luminance of all pixels in the image [93]. In contrast, local operators compress the luminance of each pixel using a mapping that varies spatially based on the neighboring pixels [94]. Therefore local TMOs preserve detail better, whereas global TMOs can lead to a loss of contrast. However local TMOs can cause halo artifacts around sharp features due to the Gaussian filter that is commonly used to compute the local parameters. The use of edge-preserving smoothing operators, such as the bilateral filter, can help prevent halos.

The simplest tone mapping method is normalization, where the HDR pixel intensities are linearly scaled down to the LDR range. However, this operation leads to a loss of detail, especially in the darker regions. In this work, we focus on five common tone mapping operators: (1) contrast-limited adaptive histogram equalization (CLAHE) [95], (2) a global luminance compression TMO [96], (3) a local luminance compression TMO that computes the local parameters based on a difference of Gaussians (DoG) approach [96], (4) a local luminance compression TMO based on a bilateral filter [93], and lastly (5) a two-scale decomposition method using bilateral filtering [94].

In CLAHE, the image is divided into regions, and the histograms of these regions are redistributed to have a wider, more uniform distribution, effectively increasing the contrast.

The luminance compression approach suggested by Reinhard in [96] has several implementations. The global implementation computes the geometric average of each pixel’s luminance to steer the compression [93]. The DoG approach computes a difference of Gaussians to determine the appropriate filter size to be used in determining the local compression parameter. The bilateral filter approach replaces the DoG with an edge-preserving smoothing operator, which is better suited to preserving sharp contrasts and can be faster to compute than the DoG.

The decomposition approach by Durand in [94] preserves sharp gradients by separating the HDR image into a base layer (low frequency components) and a detail layer (high frequency components), and only reducing the contrast of the base layer, which preserves details in the

final LDR image.

5.2.4 Quality Metrics

Quantifying the improvement (or degradation) in image quality of a tone-mapped HDR image over the original images is not straightforward since there is no ground truth. Quality metrics, such as mean-squared error (MSE), apply a direct comparison between pixel intensities. However, comparing the tone-mapped HDR image to any one of the input images necessarily leads to a mismatch between intensities, resulting in lower scores.

Instead, the analysis technique we adopt in our work is to use a sliding window to compute the similarity between the tone-mapped HDR image, and each of the original LDR ultrasound images that are used to generate the HDR image. Using only a small patch instead of the entire image makes it possible to determine local correspondences between the HDR and LDR images. This analysis shows how much of the dynamic range detail is preserved when the HDR-US image is tone mapped for display on an LDR medium: if the tone-mapped HDR-US image only has high agreement with a small subset of the original LDR images (*i.e.*, uses a narrow power level range), this means that the tone mapping failed to preserve the HDR information. Through this analysis, we aimed to determine the TMO that maintains the most detail.

The similarity metric we used is the peak signal-to-noise ratio (PSNR), defined as

$$\text{PSNR}_j = 10 \log_{10} \left(255^2 / \text{MSE}_j \right), \quad (5.6)$$

in which

$$\text{MSE}_j = \frac{1}{N} \sum_i^N \left(Z_i^{\text{HDR}} - Z_i^{\text{LDR}_j} \right)^2, \quad (5.7)$$

where N is the number of pixels in the sliding window, j is the index of the reference image, Z_i^{HDR} is the intensity of pixel i in the tone-mapped HDR image, and $Z_i^{\text{LDR}_j}$ is the intensity of pixel i in the reference image j . Computing the $\arg \max$ of PSNR_j yields the index (and therefore the power setting) of the reference image that has the best correspondence with the tone-mapped HDR image. Computing this index at each sliding window location yields




Figure 5.3: (a–d) Four of the fifteen ultrasound images acquired at various power levels, and their histograms. In low power settings, the catheter body and the skin surface have good contrast and detail; however, the deeper structures are missing. At the higher power settings, the bone surface and the deeper tissue have good detail, but the catheter and the skin are overexposed and have lost fine detail. (e) Tone-mapped high dynamic range ultrasound image. The deep tissue structures, as well as the catheter and skin details, are easily visible.

a contour map (Fig. 5.6), which shows a power level landscape of the best corresponding reference image. The TMO that results in an HDR image with the widest set of reference images will be chosen as the TMO that is most suitable for ultrasound imaging.

5.3 Experiments and Results

We conducted an *ex vivo* and an *in vivo* experiment to test our methods. In both experiments, the ultrasound images were acquired using a Philips SONOS 7500 ultrasound machine (Philips Healthcare, Andover, MA, USA) with a 128-element linear phased array probe (Philips 15-6L), which has an operating frequency range of 6–15 MHz. The TGC controls were set to their lowest value. The grayscale B-mode ultrasound images were transferred to a PC using an S-video frame grabber (USB3HDCAP, StarTech, London, Ontario, Canada).

5.3.1 *Ex Vivo*

In our *ex vivo* experiments, we placed a chicken thigh into a water tank and pressed an ablation catheter (ThermoCool SF, Biosense Webster, Diamond Bar, CA) against the tissue surface. The bottom of the tank was covered with an ultrasound absorbing mat to prevent ringing due to reflections from the tank surface. See Fig. 5.1a for a picture of the experimental setup.




Figure 5.4: (a) Ultrasound image acquired at 0.0 dB. (b) Tone mapped HDR-US image. (c) Regions discussed in the text. Highlighted in green is the ablation catheter body, purple is the ringing from the catheter, red is the fat layer, cyan is the bone, and yellow is the deep tissue structure. (d) HDR-US image intensities before tone mapping, shown in false color to convey the entire dynamic range. This image contains over three orders of magnitude of dynamic range, from 0.09 to 140.

Fifteen B-mode ultrasound images were acquired at 4 cm depth and 6 MHz frequency, with power levels linearly spaced between -26.7 to 0.0 dB (mechanical index of 0.1 to 1.0, there was no return below -26.7 dB). A subset of these images can be seen in Fig. 5.3. The power settings were dialed in manually using the ‘power’ control knob on the ultrasound machine. The response function for these images, calculated using Debevec and Malik’s method [92], can be seen in Fig. 5.2. The HDR-US image computed from the original fifteen LDR images is shown in Fig. 5.4d, which contains over three orders of magnitude of dynamic range. The tone-mapped HDR image is shown in Fig. 5.4b and demonstrates the utility of HDR in combining the low- and high-power details in a single image. Fig. 5.4c shows regions of interest.

The HDR Toolbox [89] was used to perform the HDR and tone mapping operations in MATLAB (MathWorks, Natick, MA, USA). Fig. 5.5 shows the output of five different tone mapping algorithms: (a) contrast-limited adaptive histogram equalization (CLAHE), (b) global Reinhard, (c) local Reinhard with DoG, (d) local Reinhard with bilateral filter, and (e) Durand’s bilateral separation. For all TMOs, we used the default parameters. The Durand and Reinhard TMOs automatically adjust their parameters based on the input images. A gamma correction value of 2.2 was also applied, except for the CLAHE-based TMO.

A qualitative assessment shows that the adaptive histogram equalization amplifies the




Figure 5.5: Output of five different tone mapping algorithms: (a) CLAHE, (b) global, (c) local Reinhard with DoG, (d) local Reinhard with bilateral filter, and (e) Durand's bilateral separation.




Figure 5.6: Similarity with baseline images for (a) CLAHE, (b) global, (c) local Reinhard with DoG, (d) local Reinhard with bilateral filter, and (e) Durand's bilateral separation. The blue rectangle in the top left corner indicates the size of the sliding window (40×40 pixels) used for computing the local PSNR.

noise in the deeper regions and the ringing around the catheter body. The DoG TMO results in saturated regions along the catheter and skin (Fig. 5.5c) that lead to a loss of detail. Of the remaining, the Reinhard bilateral TMO maintains finer texture along the catheter and the skin compared to the global TMO. However the Durand TMO has the best ringing suppression at the catheter tip and sharpest detail along the skin overall.

In our quantitative analysis, we calculated the PSNR for the output of each TMO, which is a measure of the similarity between the reference image and the tone-mapped HDR-US image, as described in Section 5.2.4. We iterated over each of the fifteen ultrasound images as the reference input in the PSNR calculations. A sliding window size of 40×40 pixels was chosen heuristically. For each patch, we calculated its $\arg \max_j \text{PSNR}_j$ to find the reference image power level that corresponded the most to the tone-mapped HDR-US image. We display

this information in the form of a contour plot in Fig. 5.6. The colors in the contour plot correspond to the power levels of the reference images.

We expect the area around the catheter-tissue interface to have the highest PSNR score at lower power levels, since these areas have better contrast and detail at the lower power levels. Furthermore, when the region of interest focuses on the deep tissue structures and bone surface, the PSNR should peak near or at the highest power reference image, since these structures are only visible in the high power images.

Based on Fig. 5.6, the PSNR contours at the catheter-tissue interface helps us rule out the adaptive histogram and Reinhard DoG TMOs as viable candidates. For these two TMOs, the catheter-tissue interface shows a higher correspondence with the high-power images; however in reality, these structures are overexposed and have lost detail at the higher power levels. Therefore, these two TMOs fail to preserve the detail and contrast in these regions. Of the remaining three, the Reinhard global and the Reinhard bilateral TMOs have almost identical performance, suggesting that the high frequency texture in ultrasound images results in a large kernel size, which reduces the compressive power of the bilateral approach to that of the global. The Durand TMO has the widest range of power levels, indicating that this TMO preserves the widest dynamic range among all, and is therefore best suited for use in our application. This is expected, since the Durand TMO only compresses the low frequency components of the image and preserves the high frequency details in the image.

5.3.2 *In Vivo*

In a second experiment, we imaged the lower leg of a healthy human male subject. Fifteen B-mode ultrasound images were acquired in the transverse plane at 6 MHz using the same probe, imaging depth, and power settings as before. The study was approved by the Harvard University Institutional Review Board (IRB). Fig. 5.7 shows ultrasound images taken at three power levels, as well as the false color HDR-US image displaying the full dynamic range, and its tone-mapped LDR version. The tibialis anterior muscle and the reflection from the surface of the tibia are hard to discern in the low-power image (Fig. 5.7a), whereas the muscle




Figure 5.7: HDR-US image of a human lower leg acquired in the transverse plane. The tibialis anterior muscle and the surface of the tibia can be seen. Images at fifteen power levels were acquired, three of which are shown in (a) -15.9 dB, (b) -6.3 dB and (c) 0 dB. (d) HDR image tone-mapped using the Durand TMO. (e) HDR-US image intensities prior to tone mapping, shown in false-color with values ranging from 0.1 to 49 , almost three orders of magnitude of dynamic range.

fascicles and the skin surface are saturated in the high-power images (Fig. 5.7b-c). Note that the ultrasound image in Fig. 5.7b is already saturated near the skin surface, while the deep tissue structures are still not visible. This demonstrates that capturing all details of the anatomy is not possible at a single power level.

The HDR-US image in Fig. 5.7d is tone-mapped using the Durand TMO. This image shows a clear improvement in the contrast and visibility of the deep tissue structures and the muscle fascicles while preventing saturation in the skin and bone surface. Fig. 5.7e shows the

HDR-US image prior to tone-mapping, which contains almost three orders of magnitude of dynamic range.

5.3.3 Number of Images Required for HDR-US Computations

An important point to consider in HDR computations is the number of images that is sufficient for capturing the dynamic range in the imaging field. For a given dynamic range, minimizing the number of images used in HDR reconstruction will allow for faster acquisition, resulting in higher frame rates and improved robustness to tissue and transducer motion. On the other hand, increasing the number of images can reduce noise and increase image quality.

In order to understand how the number of images included in the HDR computations affects HDR quality, we conducted an exhaustive leave- p -out cross-validation. We calculated the mean squared error (MSE) between the “ground truth” (*i.e.*, the HDR image computed using all available images) and the HDR image computed using p fewer images.

In addition to the number of images, the spacing between power levels is also crucial to recovering the widest dynamic range available. Images that are close to each other in terms of their dynamic range will not contribute much towards extending the dynamic range. In our cross-validation studies, we calculated the MSE over all possible image combinations (32,767) to obtain bounds on best and worst MSE possible for each value of p . Fig. 5.8 shows these values both for the *ex vivo* and *in vivo* datasets, which exhibit almost identical trends. Using the popular structural similarity (SSIM) index [97] instead of the MSE score leads to similar trends.

We hypothesized that uniform sampling is a good strategy for choosing which images to include in the HDR reconstruction for a given number of images. For example, if three LDR images are to be used, we would pick -26.7 dB, -12 dB, and 0 dB. Our study results validate this hypothesis; the MSE scores for the uniform sampling strategy (shown in black in Fig. 5.8) quickly converge to the optimal MSE score when four out of fifteen images are used.

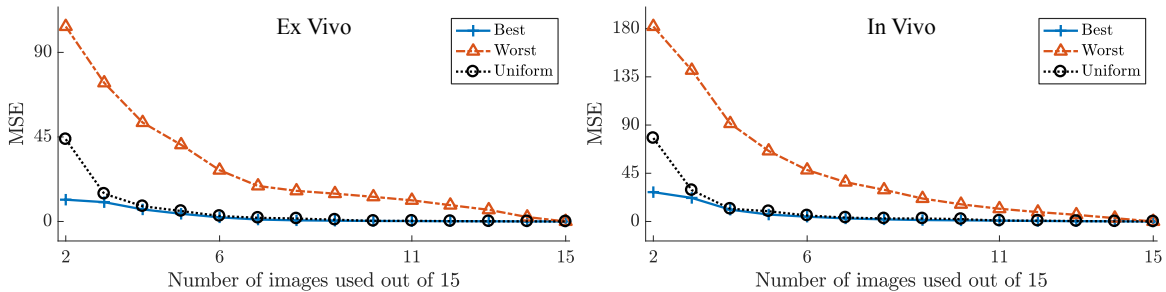


Figure 5.8: Leave- p -out cross-validation results for *ex vivo* and *in vivo* studies. The best and worst mean squared error (MSE) scores provide lower and upper bounds on the deviation from the “ground truth” HDR image constructed using fifteen power levels. The uniform sampling strategy quickly converges to the best (*i.e.*, lowest) MSE score.

5.4 Discussion

Our results demonstrate that it is possible to enhance the amount of detail that is visible in an ultrasound image by combining images that were acquired at different power levels. This enables capturing a wide dynamic range, from hypoechogenic tissue details to highly reflective instruments, and presenting them to the user in a single image. We tested the performance of five TMOs through qualitative and quantitative analysis. The peak signal-to-noise ratio studies clearly show that the Durand TMO preserves the most detail during tone mapping. The leave- p -out cross-validation results indicate that four uniformly-sampled images are sufficient for capturing the full dynamic range of the specific imaging targets used in our studies. However, the minimum number of images required to construct an HDR-US image depends on the dynamic range of the imaging target. See [92] for a discussion on the number of images required to recover the response curve and the radiance map.

We did not conduct an exhaustive study of TMOs, and the parameters for the TMOs we tested could be optimized to further improve the level of detail in the tone-mapped LDR images. Additionally, the TMO can be a user-selected post-processing option. Just as clinicians can currently control the amount of log compression applied to the image using a knob or slider on the ultrasound machine, once HDR-US is implemented clinically, the users can switch between different TMOs.

One limitation of our study was in terms of acquisition time. We currently do not have

a way of automatically adjusting the power settings on the ultrasound machine in real-time. Thus, acquiring an HDR stack of muscle activity, such as the heart, is not possible (although ECG-gating can be used to gate cardiac imaging). However, ultrasound machine manufacturers, such as BK Ultrasound (Peabody, MA, USA), have already integrated real-time frame-by-frame power level adjustments into their system application programming interfaces (APIs). Implementation of HDR-US imaging in such systems would enable a set of HDR images to be acquired in just four consecutive frames, which would result in an acquisition time of 50 ms for a probe with an 80 Hz acquisition rate. Images at different power levels can be acquired fast enough such that tissue or probe motion causes minimal misalignment. These images can then be aligned using standard algorithms, such as SIFT or Ward’s method [98]. The HDR computations, which are already optimized for mobile computing platforms, can easily be integrated into the ultrasound data processing pipeline, including 3DUS imaging.

In our studies, we were limited to using the S-video output for image capture, which introduces noise in the image. Once the HDR computations are integrated into the image formation pipeline, using the clean digital signals will lead to higher-quality HDR images.

One common issue with imaging rigid instruments with ultrasound is the ringing and shadowing artifacts from the tool. Here we did not explicitly address these artifacts; however, the tone-mapping step will automatically prevent the instruments from getting saturated, and the HDR computations will reduce the prominence of these artifacts. See [99, 100] for in-depth analyses of ultrasound artifacts arising from metallic instruments, and methods to reduce such artifacts.

Readers familiar with ultrasound imaging may question whether time gain compensation (TGC) can compensate for signal loss due to attenuation, thereby obviating the need for HDR-US. It should be noted that the utility of TGC is limited to cases when the attenuation media is approximately homogeneous in the lateral direction (*i.e.*, at a given depth, the dynamic range of the media is narrow). Otherwise, adjusting the TGC will not prevent structures from getting overexposed or lost.

5.5 Conclusions

In this chapter, we have shown that standard HDR imaging techniques can be used to improve the dynamic range of ultrasound images. Based on acoustic equations, we argued that the power settings can be used as an analog to exposure time in HDR calculations. We evaluated five tone mapping operators based on the peak signal-to-noise ratio. The results based on images acquired from *ex vivo* and *in vivo* tissue showed that the Durand TMO produces LDR images with the widest perceivable dynamic range. Extending the dynamic range of ultrasound imaging can improve the performance of further image processing steps. The benefit of HDR-US will be the greatest in procedure guidance, where we face the broadest dynamic range, due to the presence of highly reflective surgical tools and deep tissue structures.

Chapter 6

Quantifying Tissue Dynamics In Biomechanics Studies Through Ultrasound Imaging

This chapter describes the development of a non-invasive system for characterizing human tissue dynamics. The system combines medical ultrasound imaging with an optical tracking system and a vertical exciter that can impart whole-body vibrations on seated subjects. Tissue motion is extracted from the ultrasound images, and in combination with the optical tracker data, the frequency response of the tissue is calculated based on the commanded vibration of the seat and the resultant motion of the tissue. Dynamics of abdominal organs and the upper leg are characterized using the developed system. The ultrasound imaging method presented here will provide insight into the dynamics of soft tissues, as well as their boundary conditions with surrounding organ systems. The identified characteristics can be utilized for surgical planning and simulation as well as for validating a finite element model that is being developed to predict vehicular ride comfort in the early stages of automobile design. I closely collaborated with Daisuke Yamada on this work.

6.1 Motivation

Understanding the *in vivo* dynamics of soft tissue, including its interaction with adjacent tissues, is a key problem for many fields of study. An accurate understanding of the dynamics can improve the fidelity of computational models of the human body. For instance, in surgical planning and simulation, a computational model of the patient based on medical images is used to evaluate surgical procedures [101]. The static configuration of internal organs has been modeled in previous work, but the dynamic response of soft tissue, including contact forces between adjacent tissues, was not well characterized. In the automobile engineering field, computational models of passengers have been studied to predict vehicular ride comfort in automobile design [102]. Distribution of contact pressure between the seat cushion and the passenger’s anatomy have been well validated. However, conventional modeling studies did not focus on the dynamic response of soft tissues inside the passenger’s body, thus the relationship between dynamics and subjective comfort are yet to be clarified. The reason for this gap between tissue dynamics modeling and comfort estimation can be attributed to the lack of robust methods to observe internal tissue dynamics and to validate tissue dynamics models that have already been developed. Therefore, validating models of the human body dynamics requires developing a measurement method that can characterize the *in vivo* dynamics of soft tissue.

Ultrasound is a promising imaging modality for observing the *in vivo* dynamic motion and deformation of soft tissue due to its compact form, low cost, high sampling rate, and noninvasive nature. Ultrasound imaging has been previously used in biomechanics studies to estimate the muscle and tendon dynamics during walking [28–30]. Such work has shown how deformation of aponeuroses and connecting fascicles impact the biomechanics of gait.

The work presented here combines an ultrasound imaging system with a whole-body vibration exciter. Using this system, the dynamic characteristics of soft tissues can be identified in terms of a frequency response function by associating the observed response with the input vibration. An optical tracker is used to compensate for the vibration induced motion of the ultrasound probe. In addition to the design of the system, initial results of




Figure 6.1: *System setup.*

abdominal organs and upper leg are presented.

6.2 Experimental Setup

The system consists of (1) an ultrasound imaging device, (2) a vibration exciter, and (3) an optical tracker, as shown in Fig. 6.1. The ultrasound imaging device (Sonos 7500, Philips Healthcare, Andover, MA, USA) was used to acquire images of soft tissue for recording the dynamic response. A transesophageal echocardiography (TEE) ultrasound probe (Omni III, Philips Healthcare, Andover, MA, USA) was affixed to the surface of the subject's body. The TEE probe was chosen for its low-profile, reducing moments of inertia that can cause the probe to move during excitation, as well as its steerable ultrasound imaging plane. An active seat suspension system (Bose Ride System, Bose Corporation, Framingham, MA, USA), which was originally developed to reduce passenger vibration by counteracting external disturbances, was repurposed as a vibration exciter. Its powerful electromagnetic actuator lends itself well to generating input vibrations that can mimic road conditions. An optical tracker (fusionTrack500, AtracSys GmbH, Switzerland) was used to track the spatial position

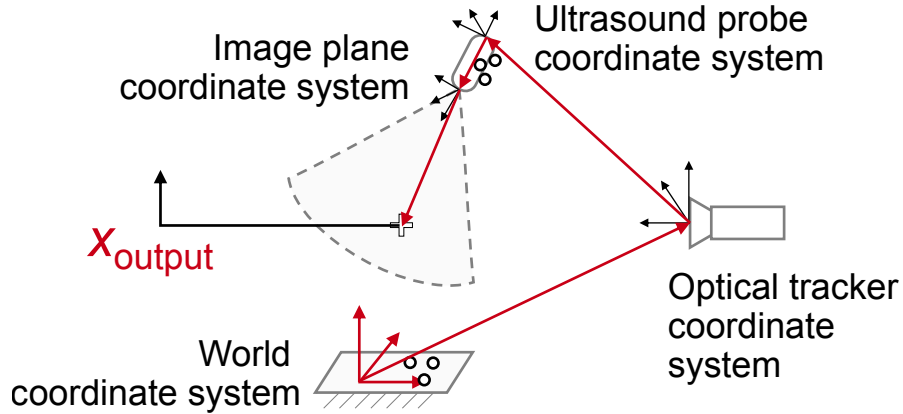


Figure 6.2: *Coordinate frames of targets of interest. The optical tracker records the pose of reflective markers that are attached to the ultrasound probe, the subject, and the rigid base frame of the vibration exciter (defined as the world coordinate frame).*

of reflective markers at a sampling rate of 333 Hz and accuracy of 90 μm . Passive motion of the ultrasound probe, the responses at the surface of the subject’s skin, and the vertical seat motion x_{seat} were tracked.

After attaching the reflective markers, ultrasound images and trajectories of the reflective markers were recorded while the subject was exposed to whole-body vibrations. The experimental subject maintained their posture throughout the experiment. In order to temporally align the ultrasound images with the optical tracker data, the ultrasound probe was given a step input both at the beginning and at the end of the experiment.

An image processing pipeline was developed to extract the motion of the stomach from the ultrasound images. First, features in the ultrasound image were automatically extracted using the Shi-Tomasi minimum-eigenvalue algorithm [103]. Second, the features were automatically tracked by using the Kanade-Lucas-Tomasi feature tracking algorithm [104], and displacement fields were calculated. Velocity fields were also calculated (Fig. 6.3) using the Lucas-Kanade optical flow algorithm [105]. Feature trajectories extracted from ultrasound images were temporally aligned with optical tracker data. The feature trajectories were then transformed to world coordinates using the measured poses of the ultrasound probe, thus compensating for the passive motion of the ultrasound probe. The vertical component of the transformed feature

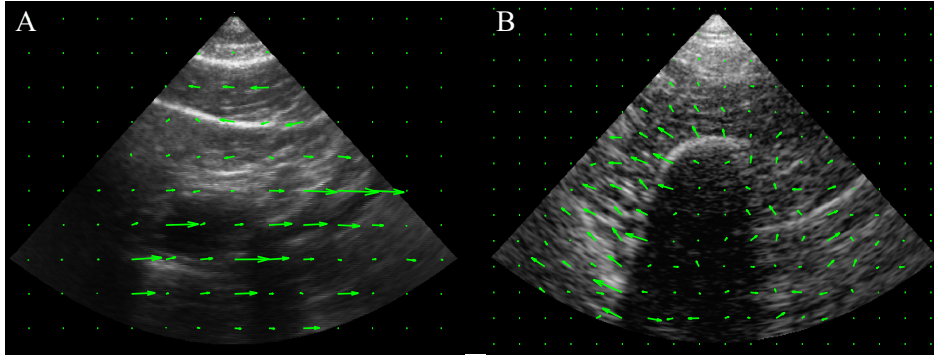


Figure 6.3: Velocity fields extracted from ultrasound images of (a) abdominal organs, and (b) the upper leg. Images were acquired at resonance frequencies of 4 Hz for the stomach and 8 Hz for the femur.

trajectories is defined as x_{output} . Relative position and orientation between the ultrasound imaging plane and the reflective markers attached to the ultrasound probe were previously calibrated using the algorithm described in [106]. Finally, the frequency response function (FRF) of soft tissue, $A(\omega)$, was calculated as a transfer function defined as

$$A(\omega) = \frac{X_{output}(\omega)}{X_{seat}(\omega)} \quad (6.1)$$

where X is the Fourier transform of x , and ω is the frequency. Analysis was limited to only the vertical direction in these experiments; we aim to study the full three dimensional deformation of tissue in future studies.

Ultrasound images of both the abdominal organs in the sagittal plane and the upper leg in the transverse plane were acquired to identify dynamics using the developed system. Two kinds of input signals, sinusoidal and frequency sweep, were introduced. Frequency of the sinusoidal signal ranged from 1 Hz to 10 Hz at 1 Hz intervals. The frequency of the frequency-sweep continuously ranged from 1 Hz to 10 Hz across a 20 s time span. Two acceleration amplitudes, 1 m/s^2 r.m.s. and 2 m/s^2 r.m.s., were employed for each signal type. However, the combination of 2 m/s^2 r.m.s. and 1 Hz was omitted due to the travel limits of the vibration exciter. The experimental procedure and conditions were approved by the Harvard University Committee on the Use of Human Subjects.

The experimental setup was controlled using a Qt-based user interface. The Frame Grabber

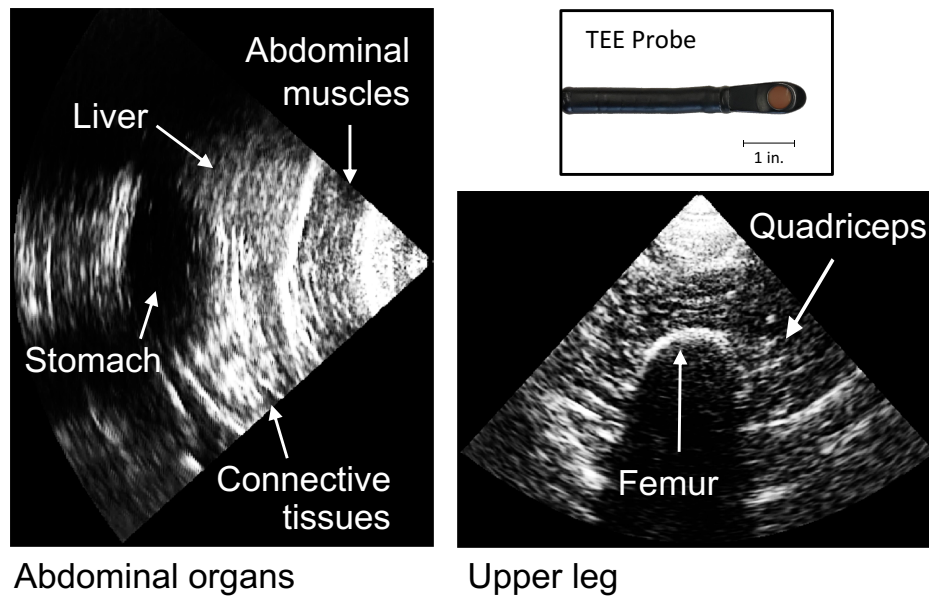


Figure 6.4: *Ultrasound images of the femur and the stomach were acquired using a TEE probe. The tip of the TEE probe is shown at the top.*

and Data Logger classes were inherited from the system developed in Chapter 3.

6.3 Results

Fig. 6.4 (a) shows a raw ultrasound image of the abdominal organs. Stomach, liver, abdominal muscles, and the surrounding connective tissue - such as omentum - can be seen in the image. The abdominal muscles were recognized as a thick layer separated from other internal organs with a bright curve. The stomach was observed as a black hollow area surrounded by a bright ellipsoidal perimeter. Fig. 6.4 (b) shows the upper leg, where the femur and surrounding quadriceps can be seen in the image. Since the ultrasound is mostly reflected at the surface of the bone, only the top edge of the femur was observed as a bright arc.

6.3.1 Frequency Response

Fig. 6.5 shows FRFs for (a) the head, (b) the sternum, and (c) the knee. The responses at the surface of the passenger's skin were observed using the optical tracker during the frequency-sweep. FRFs for (d) the stomach, and (e) the femur are also shown in Fig. 6.5.




Figure 6.5: Frequency response of internal and external anatomy to sinusoidal inputs at 1 m/s^2 and 2 m/s^2 acceleration.

These responses were extracted from the ultrasound images acquired with the sinusoidal excitation. The solid lines represent the response to 1 m/s^2 , and the dotted lines represent those to 2 m/s^2 . The curves were interpolated using a cubic spline algorithm.

6.3.2 3D Reconstruction of the Stomach

Using the techniques developed in Chapter 3, the shape of the stomach was reconstructed from 37 ultrasound images acquired using the TEE probe. A cross-section through this 3D ultrasound volume is shown in Fig. 6.6, where the stomach cavity is shaded in green.

6.4 Discussion

Peaks of magnitude and phase delays shown in Fig. 6.5 suggest that the passenger-seat coupled system has resonance characteristics. However, the resonance frequencies vary with location on the body. Since the experimental subject reported that 4 Hz of imparted vibration was the most uncomfortable, it is implied that the resonances at the head, the abdominal muscle, and the stomach can induce the subjective feeling of discomfort. Correlations between




Figure 6.6: 3D reconstruction of the stomach at 4 Hz sinusoidal input. 37 ultrasound images were acquired from 0 to 180 degrees at 5 degree intervals. The stomach cavity is shaded in green.

x_{seat} and all x_{output} were high, suggesting that the linear analyses are feasible. However, as shown in Fig. 6.5, the FRFs depend on the amplitude of the input signal, indicating that the passenger-seat coupled system is nonlinear. In external responses shown in Fig. 6.5 (a), (b), and (c), higher amplitude causes the resonance frequency and the peak levels of $|A(\omega)|$ to be lower. In contrast, no consistent tendency due to the amplitude is observed in the internal responses shown in Fig. 6.5 (d) and (e). Further experiments are required to reveal the amplitude dependency.

Cyclic vertical vibrations of soft tissue were clearly observed in the ultrasound images. We observed that the direction of velocity of the stomach is shifted from that of the abdominal muscle. Passive motion of the ultrasound probe caused by the imparted vibration can be one reason of the phase difference. However, the difference in dynamics between the abdominal muscle and the stomach can be considered to be dominant to the phase delay between the velocities, because the phase delays at 4 Hz have a difference of approximately $\pi/8$. Moreover, the quadriceps respond not only in the vertical but also the lateral direction, though only a vertical input was applied. It is revealed that multi-dimensional vibration modes can be observed inside the subject's body.

6.5 Conclusion

A tissue dynamics characterization system comprising an ultrasound imaging system, a vibration exciter, and an optical tracker is presented. The system can identify *in vivo* soft tissue dynamics as a frequency response function. Dynamics of soft tissue in abdominal organs and upper leg were characterized using the developed system, and shown to be different from those observed at the skin surface. 4D (3D + time) reconstruction of resonance mode shapes will be conducted in future work. The identified shapes of soft tissue, including boundaries with adjacent tissues, will be utilized for estimating contact forces with inverse finite element analysis.

Motivated by the small size of the TEE probe and the high frequency image acquisition capability ultrasound imaging, we created a new tool for biomechanics research, which enables researchers to study and characterize the dynamics of internal organs and muscles. We aim to combine the ultrasound images acquired during our studies with MR images of our subjects to create accurate whole-body finite element models. These models will then be validated using the analysis methods presented here.

Chapter 7

Conclusions and Future Work

In this work we aimed to present new methods and systems to improve the acquisition and processing of ultrasound images for clinical applications, with the goal of advancing diagnostic capabilities and enhancing robotic and image-guided interventions.

7.1 Contributions

Contributions of this work can be grouped into the four areas below:

7.1.1 Autonomous Steering of Flexible Instruments

We developed a robotic system for autonomous steering of ICE catheters to enable real-time tracking of cardiac structures and other instruments (e.g. ablation catheters) during cardiac interventions. I developed a computationally efficient framework to enable real-time control of the robot. Results from *in vivo* animal trials demonstrated accurate autonomous steering of ICE catheters inside a beating heart despite unmodeled disturbances due to respiratory motion and cyclical disturbances from cardiac motion. This enables real-time monitoring of the interaction between instruments and the cardiac tissue, potentially leading to a decrease in the use of X-ray fluoroscopy, and improving the quality of ablations.

7.1.2 Real-Time 3D Ultrasound Image Generation and Visualization for Procedure Guidance

Using the aforementioned robotic catheter system, I developed a framework to enable real-time reconstruction of 3D ultrasound volumes from the 2D images acquired by the ICE catheter. I implemented the computationally-demanding volume reconstruction and rendering algorithms in CUDA. We conducted two *in vivo* trials, in which we generated 3D views of the heart in real time. This work improves visualization of cardiac structures during interventions, potentially improving the efficacy of ablation procedures, thus leading to a decrease in the recurrence of arrhythmia after ablation and in the use of X-ray fluoroscopy.

7.1.3 Improving Ultrasound Image Quality Through Bayesian Methods and Enhanced Dynamic Range

We applied computational photography and machine learning techniques to improve ultrasound image quality. We used Gaussian process regression to increase interpolation accuracy in converting ultrasound scanlines to B-mode images. We analyzed *in vivo* ultrasound data acquired using both parallel and divergent scanline probes to validate the performance of \mathcal{GP} regression. Using leave-N-out cross-validation, we qualitatively and quantitatively showed that \mathcal{GP} regression leads to better B-mode conversion than other interpolation methods, indicated by the higher values of the quality metrics.

We also developed methods to acquire and generate high dynamic range ultrasound images, similar to how smartphones and digital cameras can combine pictures taken at different exposures to improve the visibility of different structures. We have shown that standard HDR imaging techniques can be used to improve the dynamic range of ultrasound images. Based on acoustic equations, we argued that the power settings can be used as an analog to exposure time in HDR calculations. We evaluated five tone mapping operators based on the peak signal-to-noise ratio. The results based on images acquired from *ex vivo* and *in vivo* tissue showed that the Durand TMO produces LDR images with the widest perceivable dynamic range. Extending the dynamic range of ultrasound imaging can improve

the performance of further image processing steps. The benefit of HDR-US will be the greatest in procedure guidance, where we face the broadest dynamic range, due to the presence of highly reflective surgical tools and deep tissue structures.

7.1.4 Ultrasound Image Acquisition and Processing for Biomechanics Studies

Thinking about uses for ultrasound imaging outside the confines of surgical robotics, we proposed to use ultrasound imaging to acquire real-time, high frame rate information about the state of muscles and other tissue inside the human body in a non-invasive fashion. Together with collaborators, I developed a platform for characterizing the dynamic response of human muscles and internal organs, such as the stomach, and conducted IRB-approved studies with human subjects. Ultrasound images of the stomach are compounded in 3D space to reconstruct the shape of the stomach. The data collected in these studies will help improve the fidelity of finite element models of the human body.

7.2 Future Directions

In general, as researchers we should make sure that our projects are clinically relevant, and that the benefits of the systems we develop are worth the costs. There are many medical robots that never make it into an operating room, and I firmly believe that we can do better if we do our due diligence in finding the clinical needs, including our clinical collaborators in design decisions, and making clinical impact our primary goal.

An interesting investigation would be on whether the HDR-US imaging techniques and \mathcal{GP} regression methods developed in this thesis can improve the performance of further image processing steps such as registration or segmentation. The HDR-US imaging techniques and \mathcal{GP} regression methods can be combined to further increase the level of detail visible in ultrasound images.

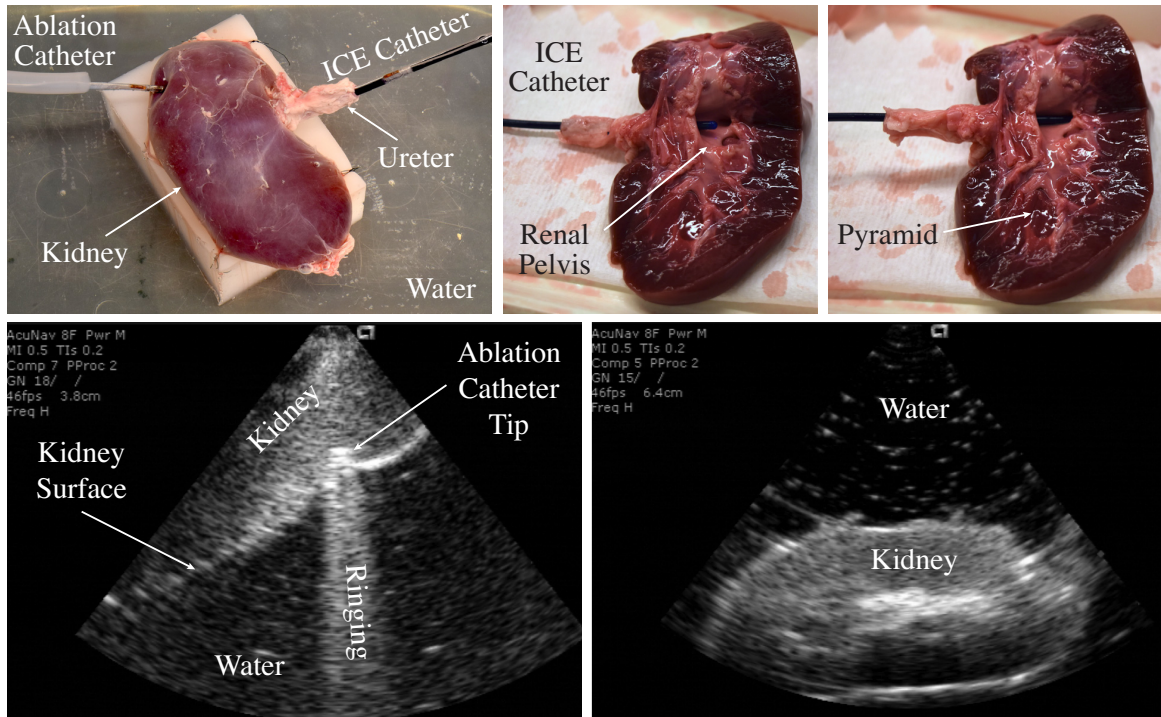


Figure 7.1: *Ex vivo* ultrasound study of a kidney.

7.2.1 Applications to Other Organs

The methods and systems described in this thesis are not limited to use in cardiac interventions, but can be adopted to other organ systems as well. The autonomous catheter steering system we developed can enable new procedures through ultrasound-based image guidance. For example, during partial nephrectomy, a portion of the kidney is surgically resected to remove tumors. Ultrasound imaging can be used during these procedures to identify tumor margins and in real time provide guidance to the surgeon. We performed an *ex vivo* ultrasound imaging study to determine whether an ICE catheter can be inserted inside the renal pelvis through the ureter to image the kidney from within. Fig. 7.1 shows an excised porcine kidney placed in a water tank. An ablation catheter was pressed against the outer surface of the kidney to as a surrogate for surgical instruments. The tissue-catheter interface was imaged using an ICE catheter, which was positioned inside the renal pelvis. The tissue deformation due to contact with ablation catheter can be clearly seen in the ultrasound image, as well as the

ringing artifact due to the metallic ablation catheter tip. After the experiment, the kidney was dissected to expose the renal pelvis to demonstrate how the ICE catheter can be steered to change the imaging plane direction. Results from this study indicate that an ICE catheter can be used to monitor tissue-instrument interaction during procedures such as nephrectomy, and can be used to detect tumor margins and provide guidance to the surgeon.

7.2.2 Catheter Steering

Another area of interest is improving controller performance for continuum manipulators. There are often unmodeled parameters that reduce accuracy and stability, but these parameters could be learned using adaptive algorithms and other machine learning methods. Modeling or learning catheter dynamics can allow for more accurate control and motion compensation.

7.2.3 4D Reconstruction, Strain Imaging, Image-Guided Steering

Extending the real-time 3D volume reconstruction to enable 4D reconstruction is algorithmically straightforward. Our current implementation interpolates the volume when all of the slices have been acquired. Changing this such that the volume is updated after each slice could potentially speed up the computation, because a more efficient convex hull algorithm can be implemented in this case.

The volume renderer can be upgraded to add shading to the volume based on surface normals. This technique could help differentiate cardiac structure better.

One exciting extension of this work is to perform strain imaging on the cardiac tissue. Using the high resolution 4D volumes generated from the ICE catheter, it should be possible to measure the contraction of the cardiac muscle, which could help identify ischemic tissue regions to guide targeted therapy delivery.

In addition, the volumetric data can be used to improve electromagnetic sensor readings through a sensor fusion approach, and allow for tissue tracking in addition to instrument tracking during cardiac interventions.

7.2.4 GANs for Scanline Conversion

Generative adversarial networks (GANs) have received a great deal of attention recently due to their generative nature [107]. GANs can learn complex distributions, therefore we are no longer restricted to models like Gaussian processes or Gaussian mixture models. This makes them interesting for use in ultrasound scanline conversion, since a GAN can potentially better learn the underlying distribution. Similar to how we perform kernel optimization in Chapter 4, a GAN can be trained to receive a subset of ultrasound scanlines, minimize the MSE error during training, and then generate a full resolution image. GANs have been used to create super-resolution images in [108].

7.2.5 Compensation for Brain Shift

In brain surgery, MRI-based image guidance is crucial. However as brain tissue is removed (and other factors such as swelling, gravity, etc) the brain tissue shift and the MRI image is no longer accurate. Currently, ultrasound imaging is used to deform the preoperative MR image to the current state, however this can only be done once an hour, since the ultrasound probe is large and imaging the brain disrupts the surgical field. ICE catheters are FDA approved for clinical use, and they can be used to simply sit on the surface of the brain and continuously acquire ultrasound images during surgery. This will allow the MR image to be updated much more frequently. Continuous ultrasound monitoring of brain shift can therefore improve procedure guidance, preserving the maximum amount of healthy brain tissue.

7.2.6 3D Printed Ultrasound Phantoms

Tissue phantoms are valuable medical imaging tools that are critical in the development and validation of new imaging technologies and analysis methods, as well as training residents and clinicians. However current techniques prohibit the manufacturing of phantoms which can replicate real tissue characteristics as well as anatomy, limiting the realism and utility of these phantoms. Manufacturing cost-effective phantoms that have both the desired mechanical and optical properties is still an elusive goal.




Figure 7.2: (left) Top view of a phantom with fibrillar structures. (right) Ultrasound image of the phantom acquired using a 15 MHz probe along the dashed red line.

Phantoms are useful for a multitude of applications:

1. Training residents who are not familiar with the imaging technology,
2. Ground truth data for testing and validating medical image processing algorithms, such as image segmentation, or elastography.
3. Ground truth data for testing and calibrating imaging systems (e.g. ultrasound probes).

Building tissue phantoms for medical imaging is currently a time consuming task, with very limited feature complexity and repeatability. Often, 3D printed molds are used to cast internal features, which are then embedded in a surrounding matrix. The curing time for these parts is on the order of tens of hours, leading to long manufacturing times. Phantoms do not have the look and feel of real tissue [109–112].

In collaboration with James Weaver, PhD, Christoph Bader, and Dominik Kolb, we developed 3D printing methods that can be used to directly print arbitrary geometries using multiple materials at once. This allows the printing of anatomical features at very fine resolutions 30 μm with excellent repeatability in a few hours without requiring any molding or assembly. The phantoms are ready for imaging as soon as they are done printing. A similar process was independently reported in [113].

These 3D printed ultrasound phantoms can be used for training new residents and clinicians in ultrasound imaging and palpation, validating image processing algorithms, and calibrating

ultrasound probes and systems. For use in training, the phantoms can be printed using clear or clear material, or completely opaque such that trainees cannot see the features inside the phantom.

Future work should focus on identifying impedance, reflectivity and transmittivity indices, attenuation constant, speed of sound, and Young's modulus.

7.2.7 Future of Ultrasound

A 3D ICE catheter with a $24^\circ \times 90^\circ$ imaging cone has been recently released (AcuNav V, Siemens Healthcare GmbH, Erlangen, Germany) and can provide upwards of 40 volumes per second. The limited size of the imaging cone still does not provide a full view of most cardiac structures, however the 3D images can be combined more accurately using 3D registration techniques developed in [55].

An exciting future awaits ultrasound imaging, where ultrasound will enable the development of a real-world medical tricorder from Star Trek. Butterfly Network, Inc. is a startup company working to revolutionize ultrasound imaging by using silicon-based ultrasound transducers instead of piezoelectric crystals [114]. Combined with miniaturized electronics and by harnessing the computing power of mobile devices (e.g. iPhones), this company is aiming to sell self-contained ultrasound probes for less than \$2,000. In addition, machine learning techniques will enable inexperienced clinicians (and perhaps home users in the future) to get accurate diagnosis.

It is quite likely that the first extraterrestrial settlers will take ultrasound probes with them to their colonies. The low cost, small size, and low power requirements make them ideal for interplanetary travel. Combined with machine learning and high intensity focused ultrasound (HIFU), these systems of the future can both autonomously diagnose and treat diseases. The work presented here improves the quality and utility of ultrasound imaging for future generations to come, both terrestrial and extraterrestrial.

References

- [1] *Image-Guided Interventions [Fact sheet]*, U.S. National Institutes of Health, 10 2010. [Online]. Available: <https://report.nih.gov/nihfactsheets/ViewFactSheet.aspx?csid=98>
- [2] W. G. Bradley, “History of medical imaging,” *Proceedings of the American Philosophical Society*, vol. 152, no. 3, pp. 349–361, 2008. [Online]. Available: <http://www.jstor.org/stable/40541591>
- [3] S. Singh and A. Goyal, “The origin of echocardiography: a tribute to inge edler,” *Texas Heart Institute Journal*, vol. 34, no. 4, p. 431, 2007.
- [4] R. Damadian, L. Minkoff, M. Goldsmith, M. Stanford, and J. Koutcher, “Field focusing nuclear magnetic resonance (fonar): Visualization of a tumor in a live animal,” *Science*, vol. 194, no. 4272, pp. 1430–1432, 1976. [Online]. Available: <http://www.jstor.org/stable/1744021>
- [5] W. Hinshaw, P. Bottomley, and G. Holland, “Radiographic thin-section image of the human wrist by nuclear magnetic resonance,” *Nature*, vol. 270, no. 5639, p. 722, 1977.
- [6] T. Meaney, “Magnetic resonance without nuclear.” *Radiology*, vol. 150, no. 1, pp. 277–277, 1984.
- [7] K. Nakajima, J. W. Milsom, and B. Böhm, *History of Laparoscopic Surgery*. New York, NY: Springer New York, 2006, pp. 1–9.
- [8] I. S. Inc., “Intuitive surgical - investors faq,” <http://phx.corporate-ir.net/phoenix.zhtml?c=122359&p=irol-faq>, 2018.
- [9] R. L. Mueller and T. A. Sanborn, “The history of interventional cardiology: Cardiac catheterization, angioplasty, and related interventions,” *American Heart Journal*, vol. 129, no. 1, pp. 146 – 172, 1995. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0002870395900551>
- [10] N. Foundation, “The nobel prize in physiology or medicine,” 1956. [Online]. Available: http://www.nobelprize.org/nobel_prizes/medicine/laureates/1956/
- [11] M. Moscucci, *Grossman & Baim’s Cardiac Catheterization, Angiography, and Intervention*. Lippincott Williams & Wilkins, 2013.
- [12] W. Kilby, J. R. Dooley, G. Kuduvalli, S. Sayeh, and J. C. R. Maurer, “The cyberknife® robotic radiosurgery system in 2010,” *Technology in Cancer Research & Treatment*, vol. 9, no. 5, pp. 433–452, 2010, pMID: 20815415.

- [13] Z. M. Hijazi, K. Shivkumar, and D. J. Sahn, “Intracardiac echocardiography during interventional and electrophysiological cardiac catheterization,” *Circulation*, vol. 119, no. 4, pp. 587–596, 2009.
- [14] J. M. Cooper and L. M. Epstein, “Use of intracardiac echocardiography to guide ablation of atrial fibrillation,” *Circulation*, vol. 104, no. 25, pp. 3010–3013, 2001.
- [15] S. G. Dravid, B. Hope, and J. J. McKinnie, “Intracardiac echocardiography in electrophysiology: a review of current applications in practice,” *Echocardiography*, vol. 25, no. 10, pp. 1172–1175, 2008.
- [16] L. M. Epstein, M. A. Mitchell, T. W. Smith, and D. E. Haines, “Comparative study of fluoroscopy and intracardiac echocardiographic guidance for the creation of linear atrial lesions,” *Circulation*, vol. 98, no. 17, pp. 1796–1801, 1998.
- [17] J.-F. Ren and F. E. Marchlinski, “Utility of intracardiac echocardiography in left heart ablation for tachyarrhythmias,” *Echocardiography*, vol. 24, no. 5, pp. 533–540, 2007.
- [18] N. F. Marrouche, D. O. Martin, O. Wazni, A. M. Gillinov, A. Klein, M. Bhargava, E. Saad, D. Bash, H. Yamada, W. Jaber *et al.*, “Phased-array intracardiac echocardiography monitoring during pulmonary vein isolation in patients with atrial fibrillation impact on outcome and complications,” *Circulation*, vol. 107, no. 21, pp. 2710–2716, 2003.
- [19] 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*, vol. 8, no. 2, pp. 82–90, 2003.
- [20] R. J. Siegel, H. Luo, and S. Biner, “Transcatheter valve repair/implantation,” *The International Journal of Cardiovascular Imaging*, vol. 27, no. 8, pp. 1165–1177, Dec 2011.
- [21] I. Hitachi Aloka Medical America, “Robot assisted partial nephrectomy,” <http://www.hitachi-aloka.com/assets/pdf/HAMA-RAPN-Surgical-Brochure-2015.pdf>, 2015.
- [22] G. Fichtinger, J. Fiene, C. W. Kennedy, G. Kronreif, I. I. Iordachita, D. Y. Song, E. Clif Burdette, and P. Kazanzides, “Robotic assistance for ultrasound guided prostate brachytherapy,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007*, N. Ayache, S. Ourselin, and A. Maeder, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 119–127.
- [23] T. K. Adebar, A. E. Fletcher, and A. M. Okamura, “3-d ultrasound-guided robotic needle steering in biological tissue,” *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 12, pp. 2899–2910, 2014.
- [24] J. Hong, T. Dohi, M. Hashizume, K. Konishi, and N. Hata, “An ultrasound-driven needle-insertion robot for percutaneous cholecystostomy,” *Physics in Medicine & Biology*, vol. 49, no. 3, p. 441, 2004.

- [25] L. VascuLogic, “Venouspro robotic venipuncture,” <https://vasculogic.com/index.html>, 2018.
- [26] S. G. Yuen, N. V. Vasilyev, J. Pedro, and R. D. Howe, “Robotic tissue tracking for beating heart mitral valve surgery,” *Medical image analysis*, vol. 17, no. 8, pp. 1236–1242, 2013.
- [27] S. B. Kesner and R. D. Howe, “Robotic catheter cardiac ablation combining ultrasound guidance and force control,” *The International Journal of Robotics Research*, vol. 33, no. 4, pp. 631–644, 2014.
- [28] M. Rana, G. Hamarneh, and J. Wakeling, “Automated tracking of muscle fascicle orientation in b-mode ultrasound images,” *Journal of Biomechanics*, vol. 42, pp. 2068–2073, 2009.
- [29] N. Cronin, C. Carty, R. Barrett, and G. Lichtwark, “Automatic tracking of medial gastrocnemius fascicle length during human locomotion,” *Journal of Applied Physiology*, vol. 111, pp. 1491–1496, 2011.
- [30] N. Cronin and G. Lichtwark, “The use of ultrasound to study muscle–tendon function in human posture and locomotion,” *Gait and Posture*, vol. 37, no. 3, pp. 305–312, 2013.
- [31] S. S. Chugh, R. Havmoeller, K. Narayanan, D. Singh, M. Rienstra, E. J. Benjamin, R. F. Gillum, Y.-H. Kim, J. H. McAnulty, Z.-J. Zheng, M. H. Forouzanfar, M. Naghavi, G. A. Mensah, M. Ezzati, and C. J. Murray, “Worldwide epidemiology of atrial fibrillation,” *Circulation*, vol. 129, no. 8, pp. 837–847, 2014.
- [32] R. Cappato, H. Calkins, S.-A. Chen, W. Davies, Y. Iesaka, J. Kalman, Y.-H. Kim, G. Klein, A. Natale, D. Packer, A. Skanes, F. Ambrogi, and E. Biganzoli, “Updated worldwide survey on the methods, efficacy, and safety of catheter ablation for human atrial fibrillation—clinical perspective,” *Circulation: Arrhythmia and Electrophysiology*, vol. 3, no. 1, pp. 32–38, 2010.
- [33] E. N. Marieb, *Human Anatomy & Physiology*, 9th ed. Boston: Pearson, 2013.
- [34] Corindus, “Corindus, inc. *CorPath Robotic PCI*,” <http://www.corindus.com/>, 2016.
- [35] C. Robotics, “Catheter robotics, inc. *Amigo Remote Catheter System*,” <http://www.catheterrobotics.com/>, 2016.
- [36] Stereotaxis, “Stereotaxis *V-Drive Robotic Navigation System*,” <http://www.stereotaxis.com/products/vdrive/>, 2016.
- [37] —, “Stereotaxis *Niobe ES*,” <http://www.stereotaxis.com/products/niobe/>, 2016.
- [38] Hansen, “Hansen Medical, Inc. *Sensei X Robotic Catheter System*,” <http://hansenmedical.com>, 2016.
- [39] R. S. Penning, J. Jung, J. A. Borgstadt, N. J. Ferrier, and M. R. Zinn, “Towards closed loop control of a continuum robotic manipulator for medical applications,” in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 4822–4827.

- [40] I. Robert J. Webster and B. A. Jones, “Design and kinematic modeling of constant curvature continuum robots: A review,” *The International Journal of Robotics Research*, vol. 29, no. 13, pp. 1661–1683, 2010.
- [41] M. Khoshnam, M. Azizian, and R. V. Patel, “Modeling of a steerable catheter based on beam theory,” in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 4681–4686.
- [42] Y. Ganji and F. Janabi-Sharifi, “Catheter kinematics for intracardiac navigation,” *Biomedical Engineering, IEEE Transactions on*, vol. 56, no. 3, pp. 621–632, 2009.
- [43] D. B. Camarillo, C. R. Carlson, and J. K. Salisbury, “Configuration tracking for continuum manipulators with coupled tendon drive,” *Robotics, IEEE Trans. on*, vol. 25, no. 4, pp. 798–808, 2009.
- [44] M. C. Yip and D. B. Camarillo, “Model-less hybrid position/force control: a minimalist approach for continuum manipulators in unknown, constrained environments,” *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 844–851, 2016.
- [45] G. J. Vrooijink, T. Ellenbroek, P. Breedveld, J. G. Grandjean, and S. Misra, “A preliminary study on using a robotically-actuated delivery sheath (rads) for transapical aortic valve implantation,” in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 4380–4386.
- [46] P. Loschak, L. Brattain, and R. Howe, “Automated pointing of cardiac imaging catheters,” in *Robotics and Automation (ICRA), 2013 IEEE Int’l Conf. on*, May 2013, pp. 5794–5799.
- [47] S. Kesner and R. Howe, “Design and control of motion compensation cardiac catheters,” in *Robotics and Automation (ICRA), 2010 IEEE Int’l Conf. on*, May 2010, pp. 1059–1065.
- [48] R. S. Penning, J. Jung, M. R. Zinn, and N. J. Ferrier, “An Evaluation of Closed-loop Control Options for Continuum Manipulators,” in *Robotics and Automation (ICRA), 2012 IEEE Int’l Conf. on*, May 2012, pp. 5392–5397.
- [49] Y. Ganji, F. Janabi-Sharifi, and A. N. Cheema, “Robot-assisted catheter manipulation for intracardiac navigation,” *Int’l J of Comp. Asst. Radiology and Surgery*, vol. 4, no. 4, pp. 307–315, 2009.
- [50] S. Seung, P. Liu, S. Park, J.-O. Park, and S. Y. Ko, “Single-port robotic manipulator system for brain tumor removal surgery: Siromans,” *Mechatronics*, vol. 26, pp. 16 – 28, 2015.
- [51] P. M. Loschak, A. Degirmenci, and R. D. Howe, “Predictive Filtering In Motion Compensation With Steerable Cardiac Catheters,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, May 2017, pp. 4830–4836.
- [52] P. M. Loschak, Y. Tenzer, A. Degirmenci, and R. D. Howe, “A 4-DOF Robot for Positioning Ultrasound Imaging Catheters,” in *Proc. of the ASME 2015 Int’l Design Eng. Tech. Conf. & Comp. and Info. in Eng. Conf. IDETC/CIE*, August 2015.

- [53] P. M. Loschak, A. Degirmenci, Y. Tenzer, C. M. Tschabrunn, E. Anter, and R. D. Howe, “A Four Degree of Freedom Robot for Positioning Ultrasound Imaging Catheters,” *ASME J. Mechanisms Robotics*, vol. 8, no. 5, pp. 051 016–051 025, May 2016.
- [54] R. Ohbuchi, D. Chen, and H. Fuchs, “Incremental volume reconstruction and rendering for 3-d ultrasound imaging,” in *Visualization in Biomedical Computing’92*, vol. 1808. International Society for Optics and Photonics, 1992, pp. 312–324.
- [55] R. J. Schneider, D. P. Perrin, N. V. Vasilyev, G. R. Marx, J. Pedro, and R. D. Howe, “Real-time image-based rigid registration of three-dimensional ultrasound,” *Medical image analysis*, vol. 16, no. 2, pp. 402–414, 2012.
- [56] 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 and Biology*, vol. 37, no. 7, pp. 1056–1072, 2011.
- [57] L. J. Brattain and R. D. Howe, “Real-Time 4D Ultrasound Mosaicing and Visualization,” in *Medical Image Computing and Computer-Assisted Interventions – MICCAI 2011*. Springer, 2011, pp. 105–112.
- [58] L. J. Brattain, P. M. Loschak, C. M. Tschabrunn, E. Anter, and R. D. Howe, “Instrument tracking and visualization for ultrasound catheter guided procedures,” in *Workshop on Augmented Environments for Computer-Assisted Interventions*. Springer, 2014, pp. 41–50.
- [59] 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 and Biology*, vol. 33, no. 7, pp. 991–1009, 2007.
- [60] C. Knackstedt, A. Franke, K. Mischke, M. Zarse, F. Gramley, T. Schimpf, J. Plisiene, G. Muehlenbruch, E. Spuentrup, S. Ernst, S. Willems, P. Kirchhof, and P. Schauerte, “Semi-automated 3-dimensional intracardiac echocardiography: Development and initial clinical experience of a new system to guide ablation procedures,” *Heart Rhythm*, vol. 3, no. 12, pp. 1453 – 1459, 2006.
- [61] A. B. Koolwal, F. Barbagli, C. R. Carlson, and D. H. Liang, “A probabilistic framework for freehand 3d ultrasound reconstruction applied to catheter ablation guidance in the left atrium,” *International Journal of Computer Assisted Radiology and Surgery*, vol. 4, no. 5, pp. 425–437, Sep 2009.
- [62] —, “A fast slam approach to freehand 3-d ultrasound reconstruction for catheter ablation guidance in the left atrium,” *Ultrasound in Medicine and Biology*, vol. 37, no. 12, pp. 2037–2054, 2011.
- [63] I. Biosense Webster. (2016) Biosense Webster, Inc. *CARTO System*. [Online]. Available: <https://www.biosensewebster.com/products/carto-3.aspx>

- [64] D. L. Packer, S. B. Johnson, M. W. Kolasa, T. J. Bunch, B. D. Henz, and Y. Okumura, “New generation of electro-anatomic mapping: full intracardiac ultrasound image integration,” *Europace*, vol. 10, no. suppl_3, pp. iii35–iii41, 2008.
- [65] B. Vagvolgyi, S. DiMaio, A. Deguet, P. Kazanzides, R. Kumar, C. Hasser, and R. Taylor, “The surgical assistant workstation,” in *Proc MICCAI Workshop: Systems and Architectures for Computer Assisted Interventions*, 2008, pp. 1–8.
- [66] P. Kazanzides, A. Deguet, A. Kapoor, O. Sadowsky, A. LaMora, and R. Taylor, “Development of open source software for computer-assisted intervention systems,” in *MICCAI Workshop on Open-Source Software*, (online at <http://hdl.handle.net/1926/46>), 2005.
- [67] G. Guennebaud, B. Jacob *et al.*, “Eigen v3,” [http:// eigen.tuxfamily.org](http://eigen.tuxfamily.org), 2010.
- [68] C. B. Barber, D. P. Dobkin, and H. Huhdanpaa, “The quickhull algorithm for convex hulls,” *ACM Trans. Math. Softw.*, vol. 22, no. 4, pp. 469–483, Dec. 1996.
- [69] D. Garcia, “Robust smoothing of gridded data in one and higher dimensions with missing values,” *Computational Statistics & Data Analysis*, vol. 54, no. 4, pp. 1167 – 1178, 2010.
- [70] P. Coupé, P. Hellier, C. Kervrann, and C. Barillot, “Nonlocal Means-Based Speckle Filtering for Ultrasound Images,” *IEEE Transactions on Image Processing*, vol. 18, no. 10, pp. 2221–2229, 2009.
- [71] Y. Erez, Y. Y. Schechner, and D. Adam, “Ultrasound image denoising by spatially varying frequency compounding,” in *Pattern Recognition: 28th DAGM Symposium, Berlin, Germany, September 12-14, 2006*. Springer Berlin Heidelberg, Sep. 2006, pp. 1–10.
- [72] I. Herlin and N. Ayache, “A new methodology to analyze time sequences of ultrasound images,” INRIA, Research Report RR-1390, 1991. [Online]. Available: <https://hal.inria.fr/inria-00075171>
- [73] R. Rohling, A. Gee, L. Berman, and G. Treece, “Radial Basis Function Interpolation for Freehand 3D Ultrasound,” in *Information Processing in Medical Imaging: 16th International Conference, IPMI’99 Visegrád, Hungary, June 28 – July 2, 1999*. Springer-Verlag, Jun. 1999.
- [74] M. R. Styzt and R. W. Parrott, “Using kriging for 3D medical imaging,” *Computerized Medical Imaging and Graphics*, vol. 17, no. 6, pp. 421–442, Nov. 1993.
- [75] O. V. Solberg, F. Lindseth, H. Torp, R. E. Blake, and T. A. Nagelhus Hernes, “Freehand 3D Ultrasound Reconstruction Algorithms—A Review,” *Ultrasound in Medicine and Biology*, vol. 33, no. 7, pp. 991–1009, Jan. 2007.
- [76] H. He and W.-C. Siu, “Single image super-resolution using Gaussian process regression,” in *CVPR ’11: Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2011.

- [77] H. D. V. Cardona, A. F. López-Lopera, Á. Á. Orozco, M. A. Álvarez, J. A. H. Tamames, and N. Malpica, “Gaussian Processes for Slice-Based Super-Resolution MR Images,” *ISVC*, pp. 692–701, 2015.
- [78] C. Wachinger, P. Golland, M. Reuter, and W. Wells, “Gaussian Process Interpolation for Uncertainty Estimation in Image Registration,” in *MICCAI 2014, Part I*, P. Golland, N. Hata, C. Barillot, J. Hornegger, and R. Howe, Eds. Boston, MA: Springer, 2014, pp. 267–274.
- [79] R. F. Wagner, M. F. Insana, and D. G. Brown, “Statistical properties of radio-frequency and envelope-detected signals with applications to medical ultrasound,” *J. Opt. Soc. Am. A*, vol. 4, no. 5, pp. 910–922, May 1987.
- [80] E. Jakeman and P. N. Pusey, “Significance of k distributions in scattering experiments,” *Phys. Rev. Lett.*, vol. 40, pp. 546–550, Feb 1978.
- [81] E. Jakeman and R. J. A. Tough, “Generalized k distribution: a statistical model for weak scattering,” *J. Opt. Soc. Am. A*, vol. 4, no. 9, pp. 1764–1772, Sep 1987.
- [82] R. Molthen, P. Shankar, and J. Reid, “Characterization of ultrasonic b-scans using non-rayleigh statistics,” *Ultrasound in Medicine & Biology*, vol. 21, no. 2, pp. 161 – 170, 1995.
- [83] R. Molthen, P. Shankar, J. Reid, F. Forsberg, E. Halpern, C. Piccoli, and B. Goldberg, “Comparisons of the rayleigh and k -distribution models using in vivo breast and liver tissue,” *Ultrasound in Medicine & Biology*, vol. 24, no. 1, pp. 93 – 100, 1998.
- [84] D. A. Abraham and A. P. Lyons, “Reliable methods for estimating the k -distribution shape parameter,” *IEEE Journal of Oceanic Engineering*, vol. 35, no. 2, pp. 288–302, April 2010.
- [85] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*. Cambridge, MA: MIT Press, 2006.
- [86] M. L. Stein, *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media, 2012.
- [87] G. Cincotti, G. Loi, and M. Pappalardo, “Frequency decomposition and compounding of ultrasound medical images with wavelet packets,” *IEEE Transactions on Medical Imaging*, vol. 20, no. 8, pp. 764–771, Aug 2001.
- [88] A. Perperidis, “Postprocessing Approaches for the Improvement of Cardiac Ultrasound B-Mode Images: A Review,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 3, pp. 470–485, March 2016.
- [89] F. Banterle, A. Artusi, K. Debattista, and A. Chalmers, *Advanced High Dynamic Range Imaging: Theory and Practice*. Natick, MA, USA: AK Peters (CRC Press), 2011.
- [90] A. H. Hung, T. Liang, P. A. Sukerkar, and T. J. Meade, “High dynamic range processing for magnetic resonance imaging,” *PLOS ONE*, vol. 8, no. 11, pp. 1–11, 11 2013.

- [91] J. L. Prince and J. Links, *Medical Imaging Signals and Systems (2nd Edition)*, 2nd ed. Pearson, 3 2014.
- [92] P. E. Debevec and J. Malik, “Recovering high dynamic range radiance maps from photographs,” in *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’97. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997, pp. 369–378.
- [93] E. Reinhard, T. Kunkel, Y. Marion, J. Brouillat, R. Cozot, and K. Bouatouch, “Image display algorithms for high- and low-dynamic-range display devices,” *Journal of the Society for Information Display*, vol. 15, no. 12, pp. 997–1014, 2007.
- [94] F. Durand and J. Dorsey, “Fast bilateral filtering for the display of high-dynamic-range images,” in *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’02. New York, NY, USA: ACM, 2002, pp. 257–266.
- [95] S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. T. H. Romeny, and J. B. Zimmerman, “Adaptive histogram equalization and its variations,” *Comput. Vision Graph. Image Process.*, vol. 39, no. 3, pp. 355–368, Sep. 1987.
- [96] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, “Photographic tone reproduction for digital images,” in *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’02. New York, NY, USA: ACM, 2002, pp. 267–276.
- [97] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, April 2004.
- [98] G. Ward, “Fast, robust image registration for compositing high dynamic range photographs from hand-held exposures,” *Journal of Graphics Tools*, vol. 8, no. 2, pp. 17–30, 2003.
- [99] 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*, vol. 26, no. 10, pp. 1303–1322, 2007.
- [100] H. Ren, B. Anuraj, and P. E. Dupont, “Varying ultrasound power level to distinguish surgical instruments and tissue,” *Medical & Biological Engineering & Computing*, Aug 2017.
- [101] M. A. Scherer and D. A. Geller, “New preoperative images, surgical planning, and navigation,” in *Imaging and Visualization in The Modern Operating Room*. Springer, 2015, pp. 205–214.
- [102] A. Siefert, “Occupant comfort—a mixture of joint angles, seat pressure and tissue loads,” SAE Technical Paper, Tech. Rep., 2016.

- [103] J. Shi *et al.*, “Good features to track,” in *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR’94., 1994 IEEE Computer Society Conference on.* IEEE, 1994, pp. 593–600.
- [104] C. Tomasi and T. Kanade, “Detection and tracking of point features,” 1991.
- [105] B. D. Lucas, T. Kanade *et al.*, “An iterative image registration technique with an application to stereo vision,” 1981.
- [106] A. Lasso, T. Heffter, A. Rankin, C. Pinter, T. Ungi, and G. Fichtinger, “Plus: open-source toolkit for ultrasound-guided intervention systems,” *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 10, pp. 2527–2537, 2014.
- [107] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [108] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang *et al.*, “Photo-realistic single image super-resolution using a generative adversarial network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4681–4690.
- [109] R. O. Bude and R. S. Adler, “An easily made, low-cost, tissue-like ultrasound phantom material,” *Journal of Clinical Ultrasound*, vol. 23, no. 4, pp. 271–273, 1995.
- [110] J. L. Kendall and J. P. Faragher, “Ultrasound-guided central venous access: a homemade phantom for simulation,” *Canadian Journal of Emergency Medicine*, vol. 9, no. 5, p. 371–373, 2007.
- [111] L. Nolting, P. Hunt, T. Cook, and B. Douglas, “An inexpensive and easy ultrasound phantom,” *Journal of Ultrasound in Medicine*, vol. 35, no. 4, pp. 819–822, 2016.
- [112] H. J. Won, N. Kim, G. B. Kim, J. B. Seo, and H. Kim, “Validation of a ct-guided intervention robot for biopsy and radiofrequency ablation: experimental study with an abdominal phantom,” *Diagnostic and Interventional Radiology*, vol. 23, no. 3, pp. 233–237, 2017.
- [113] J.-R. Jacquet, F. Levassort, F. Ossant, and J.-M. Grégoire, “3d printed phantom for high frequency ultrasound imaging,” in *Ultrasonics Symposium (IUS), 2015 IEEE International.* IEEE, 2015, pp. 1–4.
- [114] B. N. Inc., “Meet Butterfly iQ,” <https://www.butterflynetwork.com/>, 2018.