

REAL-TIME 4D ULTRASOUND RECONSTRUCTION FOR IMAGE-GUIDED  
INTRACARDIAC INTERVENTIONS

(Spine title: Real-Time 4D Ultrasound Reconstruction for Cardiac Therapy)

(Thesis format: Monograph)

by

Danielle F. Pace

Biomedical Engineering Graduate Program

A thesis submitted in partial fulfillment  
of the requirements for the degree of  
Master of Engineering Science

The School of Graduate and Postdoctoral Studies  
The University of Western Ontario  
London, Ontario, Canada

March 2010

© Danielle F. Pace 2010

THE UNIVERSITY OF WESTERN ONTARIO  
SCHOOL OF GRADUATE AND POSTDOCTORAL STUDIES

CERTIFICATE OF EXAMINATION

Supervisor:

Examiners:

---

Dr. Terry Peters

Supervisory Committee:

---

Dr. James Lacefield

---

Dr. Aaron Fenster

---

Dr. Shuo Li

---

Dr. Maria Drangova

---

Dr. Gerard Guiraudon

---

Dr. James White

The thesis by

Danielle F. Pace

entitled

Real-Time 4D Ultrasound Reconstruction for Image-Guided Intracardiac  
Interventions

is accepted in partial fulfillment of the  
requirements for the degree of  
Master of Engineering Science

Date \_\_\_\_\_

\_\_\_\_\_  
Chair of Thesis Examination Board

## Abstract

Image-guided therapy addresses the lack of direct vision associated with minimally-invasive interventions performed on the beating heart, but requires effective intraoperative imaging. Gated 4D ultrasound reconstruction using a tracked 2D probe generates a time-series of 3D images representing the beating heart over the cardiac cycle. These images have a relatively high spatial resolution and wide field of view, and ultrasound is easily integrated into the intraoperative environment. This thesis presents a real-time 4D ultrasound reconstruction system incorporated within an augmented reality environment for surgical guidance, whose incremental visualization reduces common acquisition errors. The resulting 4D ultrasound datasets are intended for visualization or registration to preoperative images. A human factors experiment demonstrates the advantages of real-time ultrasound reconstruction, and accuracy assessments performed both with a dynamic phantom and intraoperatively reveal RMS localization errors of 2.5-2.7 mm, and 0.8 mm, respectively. Finally, clinical applicability is demonstrated by both porcine and patient imaging.

**Keywords:** 4D ultrasound, real-time ultrasound reconstruction, intraoperative imaging, image-guided therapy, minimally-invasive cardiac surgery, echocardiography

## Co-Authorship

Section 3.3 and a portion of Section 3.6 are expanded from:

- D.F. Pace, A.D. Wiles, J. Moore, C. Wedlake, D.G. Gobbi and T.M. Peters. Validation of four-dimensional ultrasound for targeting in minimally-invasive beating-heart surgery. In M.I. Miga and K.H. Wong, editors, *Medical Imaging 2009: Visualization, Image-Guided Procedures, and Modeling*, Proceedings of SPIE, Volume 7261, Page 726115, Orlando, FL, USA, 2009.

Dr. Wiles provided guidance with experimental design and data analysis. John Moore provided support with phantom construction and laboratory equipment. Chris Wedlake provided support with software architecture and is the original author of the ECG-gating classes. Dr. Gobbi is the original author of the real-time 3D ultrasound reconstruction software on which the system presented here is based. Dr. Peters provided project supervision and guidance.

Appendix A is expanded from:

- D.F. Pace, D.G. Gobbi, C. Wedlake, J. Gumprecht, J. Boisvert, J. Tokuda, N. Hata and T.M. Peters. An open-source real-time ultrasound reconstruction system for four-dimensional imaging of moving organs. In *Workshop on Systems and Architectures for Computer Assisted Interventions*, Medical Image Computing and Computer Assisted Interventions, London, UK, 2009. Online at <http://hdl.handle.net/10380/3083>.

The contributions of Dr. Gobbi, Chris Wedlake and Dr. Peters are as above. The application presented here extends software originally written by Dr. Boisvert and further developed by Jan Gumprecht. Dr. Tokuda is the original author of the “4D Imaging” module in 3D Slicer and the OpenIGTLink network communication protocol. Dr. Hata provided project supervision for this work and for that of Jan Gumprecht and Dr. Tokuda.

## Acknowledgments

I would first like to thank my supervisor, Dr. Terry Peters, for your thoughtful guidance and for your continuous encouragement and support. Your example has truly taught me so much about how to be an upstanding independent researcher. Thanks to Dr. Aaron Fenster, Dr. Maria Drangova and Dr. Gerard Guiraudon, who were always available to provide suggestions as members of my advisory committee.

To John Moore and Chris Wedlake, thank you for being such helpful collaborators and patient teachers. Thanks to David Gobbi for giving me lots of great code to start off with, and for your helpful emails at the beginning of this project. Thanks as well to Dr. Roy Eagleson for his many helpful comments.

Dr. Daniel Bainbridge, Dr. Gerard Guiraudon, Dr. Doug Jones, Dr. Bob Kiaii and Dr. Michael Chu, thank you for sharing your clinical knowledge, stimulating many insightful discussions, and for always being willing to lend a helping hand. I would also like to thank Dr. Daniel Bainbridge for his many combined hours spent performing innumerable ultrasound reconstructions over the course of this project. I also very much appreciate the help of the nursing staff at CSTAR during the porcine experiments described in this thesis.

Thanks to everyone in the Peters group, especially Dr. Andrew Wiles, Cristian Linte, Diego Cantor, Carling Cheung, Dan Cho, Dr. Qi Zhang, Jessie Guo, Pencilla Lang, Jennifer Lo and Petar Seslija. Thank you for all of your assistance and suggestions, as well as for all of the great times we've shared together. To Marianne DeGorter and Elizabeth Pile, thank you for always being such wonderful roommates and unwavering friends.

Finally, I would like to sincerely thank my parents Mary and Joe Pace and my sister Emily for your unconditional support. This thesis is dedicated to you.

# Contents

|  |            |
|--|------------|
| <b>Certificate of Examination</b>  | <b>ii</b>  |
| <b>Abstract</b>  | <b>iii</b> |
| <b>Co-Authorship</b>   | <b>iv</b>  |
| <b>Acknowledgements</b>  | <b>v</b>   |
| <b>List of Tables</b>  | <b>ix</b>  |
| <b>List of Figures</b>   | <b>x</b>   |
| <b>List of Abbreviations</b>   | <b>xii</b> |
| <b>1 Introduction: Ultrasound in Minimally-Invasive Image-Guided Cardiac Therapy</b> | <b>1</b>   |
| 1.1 Introduction . . . . .   | 1          |
| 1.2 Minimally-Invasive Cardiac Therapy . . . . .                                     | 2          |
| 1.3 Image-Guided Therapy . . . . .   | 3          |
| 1.4 Ultrasound Imaging . . . . .   | 5          |
| 1.5 Supporting Technologies . . . . .  | 8          |
| 1.5.1 Tracking . . . . .   | 8          |
| 1.5.2 Improving ultrasound imaging . . . . .   | 10         |
| 1.5.3 Patient, cardiac and respiratory motion compensation . . . . .                 | 14         |
| 1.5.4 Image registration . . . . .   | 21         |
| 1.6 Image-Guided Therapy Systems . . . . .   | 29         |
| 1.7 Validation . . . . .   | 35         |
| 1.8 Challenges and Future Directions . . . . .                                       | 38         |
| 1.9 Conclusions . . . . .  | 41         |
| 1.10 Thesis Outline . . . . .  | 42         |

|          |   |            |
|----------|---|------------|
| <b>2</b> | <b>Real-Time 4D Ultrasound Reconstruction</b>                 | <b>44</b>  |
| 2.1      | Introduction . . . . .  | 44         |
| 2.1.1    | 4D ultrasound in image-guided cardiac therapy . . . . .       | 44         |
| 2.1.2    | Ultrasound reconstruction . . . . .                           | 46         |
| 2.1.3    | Motivation . . . . .  | 51         |
| 2.2      | Real-time 4D Ultrasound Reconstruction System . . . . .       | 53         |
| 2.2.1    | System overview . . . . .                                     | 53         |
| 2.2.2    | Data acquisition and synchronization . . . . .                | 57         |
| 2.2.3    | Slice selection . . . . .                                     | 59         |
| 2.2.4    | Slice insertion . . . . .                                     | 61         |
| 2.2.5    | Real-time visualization within an AR environment . . . . .    | 64         |
| 2.2.6    | System implementation . . . . .                               | 67         |
| 2.2.7    | Performance considerations . . . . .                          | 70         |
| 2.3      | Discussion . . . . .  | 72         |
| 2.3.1    | Clinical potential . . . . .                                  | 72         |
| 2.3.2    | Integration into AR environment . . . . .                     | 73         |
| <b>3</b> | <b>Experimental Validation</b>                                | <b>76</b>  |
| 3.1      | Human Factors . . . . .                                       | 76         |
| 3.2      | Porcine Imaging . . . . .                                     | 79         |
| 3.3      | Accuracy Assessment with Dynamic Phantom . . . . .            | 85         |
| 3.4      | Intraoperative Accuracy Assessment . . . . .                  | 99         |
| 3.5      | Clinical Imaging . . . . .                                    | 110        |
| 3.6      | Discussion . . . . .  | 113        |
| 3.6.1    | Human factors benefits of real-time reconstruction . . . . .  | 113        |
| 3.6.2    | Quantitative accuracy assessments . . . . .                   | 115        |
| 3.6.3    | Porcine and human imaging . . . . .                           | 121        |
| <b>4</b> | <b>Future Research Directions and Conclusions</b>             | <b>123</b> |
| 4.1      | Improvements and Future Work . . . . .                        | 123        |
| 4.1.1    | Disadvantages compared to offline US reconstruction . . . . . | 124        |
| 4.1.2    | Reducing reconstruction artifacts . . . . .                   | 125        |
| 4.2      | Additional Applications . . . . .                             | 127        |
| 4.3      | Conclusions . . . . .   | 130        |
| <b>A</b> | <b>Software Overview</b>                                      | <b>132</b> |
| A.1      | Introduction . . . . .  | 132        |
| A.2      | Ultrasound Reconstruction Classes . . . . .                   | 133        |
| A.2.1    | Video capture . . . . .                                       | 133        |
| A.2.2    | Tracking . . . . .  | 135        |
| A.2.3    | Gating . . . . .  | 135        |
| A.2.4    | Ultrasound reconstruction . . . . .                           | 136        |

|       |  |            |
|-------|--|------------|
| A.3   | User Interface and Communication . . . . . | 137        |
| A.3.1 | In the AtamaiViewer . . . . .              | 137        |
| A.3.2 | In 3D Slicer: SynchroGrab4D . . . . .      | 139        |
| A.4   | Conclusions . . . . .                      | 140        |
|       | <b>References</b>                          | <b>141</b> |
|       | <b>Vita</b>                                | <b>172</b> |



# List of Tables

|     |  |     |
|-----|--|-----|
| 2.1 | Tradeoffs in choosing an imaging modality for use in image-guided cardiac therapy . . . . .                                  | 74  |
| 3.1 | Summary statistics for centre reconstruction error in the dynamic phantom experiment . . . . .                               | 96  |
| 3.2 | Summary statistics for distance reconstruction error in the dynamic phantom experiment . . . . .                             | 97  |
| 3.3 | Summary statistics for volume reconstruction error and eccentricity in the dynamic phantom experiment . . . . .              | 98  |
| 3.4 | Summary statistics for localization error in the intraoperative accuracy assessment and analogous laboratory study . . . . . | 107 |

# List of Figures

|      |   |    |
|------|---|----|
| 2.1  | Overview of ultrasound reconstruction . . . . .   | 46 |
| 2.2  | System overview for the real-time 4D ultrasound reconstruction system   | 56 |
| 2.3  | Prospective and retrospective gating . . . . .  | 58 |
| 2.4  | Real-time 4D ultrasound reconstruction in progress . . . . .  | 64 |
| 2.5  | Example visualization of a reconstructed 4D ultrasound dataset within an augmented reality environment . . . . .  | 66 |
| 3.1  | The beating heart phantom used in the human factors experiment . .  | 77 |
| 3.2  | Example output volumes acquired with and without real-time visualization during the human factors experiment, for both 3D and 4D ultrasound reconstruction . . . . .                                      | 78 |
| 3.3  | Example 4D ultrasound reconstruction of a beating porcine heart acquired using the TEE probe with a rotational acquisition under alternating respiration . . . . .  | 81 |
| 3.4  | Example 3D ultrasound reconstruction of an excised porcine heart in a water bath, acquired using the TEE probe with a rotational acquisition  | 82 |
| 3.5  | Example 4D ultrasound reconstruction of a beating porcine heart, acquired using the TEE probe with a fan acquisition under suspended respiration . . . . .  | 83 |
| 3.6  | Example output volumes representing a beating porcine heart at the 0% R-R interval, acquired using the TEE probe with a rotational acquisition under free respiration and suspended respiration . . . . . | 84 |
| 3.7  | The dynamic phantom, the point-source, distance and spherical phantoms, and an example of the light sensor signal used for gating . . . .   | 86 |
| 3.8  | The experimental setup and a schematic illustrating the transformations used to calculate the expected position of the phantom within the ultrasound beam coordinate system . . . . .                     | 88 |
| 3.9  | Assessment of the influence of the dynamic phantom’s motor on the magnetic tracking system . . . . .  | 89 |
| 3.10 | Predicting each phase’s transform matrix for the sensor attached to the dynamic phantom’s platform . . . . .  | 91 |

|      |  |     |
|------|--|-----|
| 3.11 | Example reconstructions of the point-source, distance and spherical phantoms . . . . .   | 94  |
| 3.12 | Example motion artifacts present during 4D ultrasound reconstruction   | 94  |
| 3.13 | Sample ultrasound volumes overlaid with 3D representations of the expected and observed phantom positions for the point-source, distance and spherical phantoms . . . . .                        | 95  |
| 3.14 | Centre reconstruction error, distance reconstruction error, volume reconstruction error and fractional anisotropy for 3D ultrasound reconstruction versus distance from the transducer . . . . . | 100 |
| 3.15 | Centre reconstruction error, distance reconstruction error, volume reconstruction error and fractional anisotropy for 4D ultrasound reconstruction versus phantom motor speed . . . . .          | 101 |
| 3.16 | Components of the intraoperative accuracy assessment study . . . . .   | 103 |
| 3.17 | Laboratory experiment using the tracked pointer tool . . . . .   | 105 |
| 3.18 | Boxplot for the localization errors in the intraoperative accuracy assessment and the analogous laboratory study . . . . .   | 108 |
| 3.19 | Localization error in the intraoperative accuracy assessment and the analogous laboratory study versus distance from the transducer . . . . .  | 108 |
| 3.20 | Example clinical 4D ultrasound reconstruction, acquired using the TEE probe with a rotational acquisition . . . . .  | 111 |
| 3.21 | Example clinical 4D ultrasound reconstruction, acquired using the TEE probe with a pullback acquisition . . . . .  | 112 |
| 4.1  | Augmented reality environments for surgical guidance during spinal facet injections . . . . .  | 129 |
| A.1  | Class diagram of the VTK classes implementing the real-time 4D ultrasound reconstruction algorithm . . . . .   | 134 |
| A.2  | Example 4D ultrasound reconstruction of a beating heart phantom using both the AtamaiViewer and SynchroGrab4D with 3D Slicer . . . . .   | 138 |

# List of Abbreviations

|       |  |
|-------|--|
| 2D    | Two-Dimensional                        |
| 3D    | Three-Dimensional                      |
| 4D    | Four-Dimensional                       |
| AF    | Atrial Fibrillation                    |
| API   | Application Programming Interface      |
| AR    | Augmented Reality                      |
| ASD   | Atrial Septal Defect                   |
| bpm   | Beats per Minute                       |
| CABG  | Coronary Artery Bypass Graft           |
| COPD  | Chronic Obstructive Pulmonary Disease  |
| CPU   | Central Processing Unit                |
| CT    | Computed Tomography                    |
| DOF   | Degrees of Freedom                     |
| ECG   | Electrocardiogram                      |
| EP    | Electrophysiology/Electrophysiological |
| FA    | Fractional Anisotropy                  |
| fps   | Frames Per Second                      |
| GUI   | Graphical User Interface               |
| GPU   | Graphics Processing Unit               |
| ICE   | Intracardiac Echocardiography          |
| IGSTK | Image-Guided Surgery Toolkit           |
| IGT   | Image-Guided Therapy                   |
| IVUS  | Intravascular Ultrasound               |

|            |  |
|------------|--|
| LV         | Left Ventricle                             |
| MI         | Mutual Information                         |
| MRI        | Magnetic Resonance Imaging                 |
| MRML       | Medical Reality Modeling Language          |
| NCC        | Normalized Cross-Correlation               |
| NMI        | Normalized Mutual Information              |
| OR         | Operating Room                             |
| Orthoplane | Orthogonal Slice Plane                     |
| PCA        | Principal Component Analysis               |
| PET        | Positron Emission Tomography               |
| PFO        | Patent Foramen Ovale                       |
| PLL        | Phase-Locked Loop                          |
| PNN        | Pixel Nearest Neighbor                     |
| PTL        | Pixel Trilinear                            |
| PV         | Pulmonary Vein                             |
| RAM        | Random Access Memory                       |
| RF         | Radiofrequency                             |
| RMS        | Root Mean Square                           |
| rpm        | Rotations per Minute                       |
| RT3D       | Real-Time Three-Dimensional                |
| SA         | Spherical Anisotropy                       |
| SAD        | Sum of Absolute Differences                |
| SSD        | Sum of Squared Differences                 |
| SI         | Superior-Inferior                          |
| SLERP      | Spherical Linear Interpolation             |
| SNR        | Signal to Noise Ratio                      |
| SPECT      | Single Photon Emission Computed Tomography |
| TEE        | Transesophageal Echocardiography           |
| TTE        | Transthoracic Echocardiography             |
| US         | Ultrasound                                 |

VSD Ventricular Septal Defect  
VTK Visualization Toolkit

# Chapter 1

## Introduction: Ultrasound in Minimally-Invasive Image-Guided Cardiac Therapy

### 1.1 Introduction

This chapter reviews recent work investigating the use of echocardiography (cardiac ultrasound) within image-guided therapy (IGT) systems developed for minimally-invasive cardiac interventions. Ultrasound (US) is commonly used in clinical practice for preoperative diagnosis and planning and for post-procedural therapy evaluation, but is emerging as an attractive intraoperative imaging modality as well. The review presented below describes how echocardiography has been used to achieve image-based tool tracking, cardiac and respiratory motion compensation, registration between preoperative images and the intraoperative patient, and real-time surgical guidance within sophisticated visualizations incorporating magnetic resonance imaging (MRI), computed tomography (CT) or X-ray along with virtual representations of tracked surgical tools.

## 1.2 Minimally-Invasive Cardiac Therapy

Clinical research and development in minimally-invasive cardiac therapy [1,2] has focused on a wide variety of procedures, including coronary artery bypass grafting (CABG) [3], mitral valve replacement and repair [4], aortic valve replacement and repair [5], ablation of the pulmonary veins (PVs) and/or other targets for treatment of atrial fibrillation (AF) [6] and repair of congenital pathologies such as atrial septal defects (ASDs), ventricular septal defects (VSDs) or patent foramen ovale (PFO) [7].

Anticipated benefits that motivate minimally-invasive cardiac interventions include reduced trauma, improved survival rates, decreased risk of infections or other complications, decreased bleeding, reduced postoperative pain, better cosmetic results and shorter recovery times leading to economic benefits (ex. [8–11]). In particular, performing surgery on the beating heart instead of using cardiopulmonary bypass may decrease subsequent atrial fibrillation [12], inflammatory response [13], and the risk of stroke [14] or other cognitive impairments [15]. Beating heart surgery also allows for post-therapy assessment, which is impossible to perform on the arrested heart. All of these advantages will very often outweigh the surgical learning curves required for clinical adoption of minimally-invasive interventions. An increasing demand for cardiothoracic surgery is forecasted due to an aging population [16]: this patient group has an increased prevalence of comorbidities and would benefit immensely from minimally-invasive approaches to cardiac surgery.

The “invasiveness” of conventional cardiac therapy has two equally important sources: 1) access to the surgical target via a median sternotomy and rib-spreading; and 2) cardiac arrest using cardiopulmonary bypass. Minimally-invasive cardiac therapies reduce the negative effects of one, or ideally both, of these. Less invasive direct access routes include partial sternotomy, thoracotomy or a transapical approach, or alternatively a percutaneous approach can be taken using catheters. There is also a growing focus on robotic systems using port access [17]. The second source of invasiveness can be eliminated by performing interventions on the beating heart.



## 1.3 Image-Guided Therapy

Minimally-invasive cardiac therapy performed on the beating heart presents two major challenges. The first is to design instrumentation that is compatible with the minimally-invasive surgical access and the dynamic cardiac environment. Secondly, surgical guidance must be provided as there is little to no direct vision of surgical targets and tool actuators. The latter is the focus of this chapter.

The first minimally-invasive cardiac techniques, including robot-assisted procedures, relied on thoroscopic video or fluoroscopy for surgical guidance [18–23]. However, the excellent image quality and high frame rate provided by endoscopic video can only be taken advantage of during epicardial procedures or intracardiac procedures performed with cardiopulmonary bypass, because video cannot “see through” the blood-filled intracardiac environment during off-pump procedures. Fluoroscopy is also not ideal as it has poor soft tissue contrast and exposes the patient and clinical staff to harmful ionizing radiation.

More recently, image-guided therapy systems for cardiac surgery have fused data from preoperative imaging, intraoperative imaging and tracking technologies to form sophisticated visualizations for surgical guidance [24]. Fundamentally, IGT systems integrate and display a variety of image, functional and/or spatial information within the same 3D coordinate space registered to the intraoperative patient. Image-guided therapy systems typically include all or a subset of the following steps:

- **Preoperative imaging:** Acquires high quality 3D or 4D (3D+t) images, typically using MRI or CT. These images may be used to generate detailed surface models or to perform preoperative planning. This step may be performed days to weeks before the intervention, but short time intervals are preferred to minimize differences in the heart’s appearance between the preoperative and intraoperative states.
- **Intraoperative imaging:** Provides updates of the intraoperative location of surgical targets and tools during the intervention. For off-pump procedures, this is typically provided by real-time modalities such as ultrasound, fluoroscopy or

endoscopy to capture the rapidly beating heart. Intraoperative C-arm CT is also gaining popularity, but these images are typically used in the same way as preoperative images, with the advantage of imaging the heart in its intraoperative state.

- **Image-to-patient registration:** Finds the transform that transforms the preoperative image so that it can be related to the intraoperative patient space.
- **Tracking:** Localizes surgical tools and/or intraoperative imaging hardware, such as ultrasound probes, in 3D space in real-time.
- **Surgical guidance:** Displays the virtual surgical scene, fusing preoperative and intraoperative images and showing moving representations of tracked surgical tools within the same coordinate system. In the absence of direct vision, clinicians rely on this visualization to perform surgical tasks, such as moving tools towards surgical targets and performing therapy delivery.

IGT systems must satisfy many requirements to be clinically useful. They must address genuine clinical problems while providing clinical outcomes and complication rates that are better than or comparable to those of standard clinical practice (although in some cases a slight reduction in effectiveness may be an acceptable trade-off for minimal invasiveness). IGT systems must also integrate seamlessly into the operating room with minimal disruption to surgical workflow, instrumentation and imaging equipment. Increases in procedure time and manual interaction should be minimized, imposing strict demands on algorithms to be both fully automatic and fast. Finally, IGT systems and their component subsystems must be extensively and thoroughly validated, with robust short-term accuracy and precision and long-term effectiveness and proven patient safety even in non-standard cases. All of these must be achieved with cost-benefit ratios in mind and while convincing the cautious medical community that a new system's benefits warrant procedural changes and surgical learning curves.

## 1.4 Ultrasound Imaging

Ultrasound is very promising for intraoperative imaging as it is the only real-time modality that can image through blood and simultaneously does not induce ionizing radiation or impose restrictions on the intraoperative setup. Echocardiography provides acceptable spatial resolution and good soft tissue contrast and can visualize common surgical targets such as the valves, pulmonary veins, ASDs and PFOs. Ultrasound also provides functional information including blood flow (from Doppler imaging) and ventricular wall motion or synchrony. Finally, ultrasound is low-cost, portable and is already ubiquitous within operating rooms as it is used extensively for intraoperative monitoring, assessment and guidance (as reviewed in [25, 26]). Of course, ultrasound is not without disadvantages, which include its lower signal to noise ratio (SNR) compared to MR or CT, an anisotropic spatial resolution, a small field of view, and the fact that sparse, irregularly sampled data is often all that is available. In addition, there are frequent image artifacts, structures are localized incorrectly in the depth direction if the speed of sound within tissue does not equal that assumed by the ultrasound machine, image quality is dependent on the angle between the US beam and local structures, and acceptable imaging requires good acoustic windows. Despite these disadvantages, ultrasound has been effectively used to perform a variety of IGT tasks including image-based tool tracking, cardiac and respiratory gating and motion compensation, image-to-patient registration, and visualization during surgical guidance. In particular, any deficits in ultrasound image quality can be overcome by presenting real-time echocardiography in the context of image data from other sources.

### 1.4.0.1 Types of echocardiography

Ultrasound transducers vary in size and invasiveness, and may image natively in 2D or 3D. Each has varying standard clinical usage [27] and is suited to different applications during minimally-invasive cardiac interventions:

- **Transthoracic echocardiography (TTE):** These probes are held against the patient’s chest, and are simple to use and completely non-invasive. However,

acoustic windows where the ribs and lungs do not impede imaging are limited, and depth penetration is problematic in obese patients or those with chronic lung disease. Because the probe must remain in contact with the patient’s chest, intraoperative TTE is cumbersome.

- **Transesophageal echocardiography (TEE):** TEE transducers are inserted into the patient’s esophagus or upper stomach to image from directly behind the heart. Most are multiplanar, meaning that the imaging plane can be electronically rotated through  $180^\circ$ . The viewing direction can also be manipulated by translating, flexing and tilting the probe. The proximity of the probe to the heart allows for higher frequency transducers, which increases spatial resolution and overall image quality compared to TTE. During interventions, the transducer is conveniently out of the way inside the patient, but this necessitates general anesthesia and makes imaging more technically challenging.
- **Intracardiac echocardiography (ICE):** Fixed to a steerable catheter, these transducers are navigated directly into the heart via the femoral or jugular vein. In particular, ICE is often used to visualize interventional catheters during percutaneous procedures. A very high imaging frequency provides excellent image quality, and general anesthesia is not required. However, the single-use nature of ICE makes it expensive. In addition, ICE probes are difficult for the clinician to manipulate inside the heart.
- **Reconstructed (“freehand”) echocardiography:** One or a time series of 3D ultrasound images are built from multiple 2D images that cover 3D space, and are acquired by moving the ultrasound probe manually or mechanically [28]. Cardiac gating is required when imaging the beating heart with this approach. Each 2D US image must be accurately localized in 3D space by tracking the ultrasound probe, or, for mechanical probe manipulation, using actuator feedback (although sensorless approaches do exist). Ultrasound reconstruction is flexible and can generate 3D images with good spatial resolution and a large field of view, but can be a lengthy procedure and is subject to artifacts caused

by tracking error, cardiac gating error, or respiratory or patient motion.

- **Real-time 3D (RT3D) echocardiography:** A 2D matrix array transducer with electronic beamsteering natively acquires pyramidal 3D images at 20-30 frames per second with real-time volume rendering [29]. Both TTE and TEE transducers are available and RT3D ICE is on the horizon. Trade-offs between spatial resolution, frame rate and field of view are caused by the finite speed of sound in tissue, so only a small section of the heart can be imaged at any point in time. It is therefore common to stitch together several ECG-gated images acquired over multiple cardiac cycles from different viewing directions with electronic beamsteering. Such “wide-angle” scans are generated without probe tracking and assuming a stationary transducer and patient, and are subject to stitch artifacts at the interfaces between the original RT3D images.

Intravascular ultrasound (IVUS) (mounted on a steerable catheter and moved through an artery) and contrast echocardiography (microbubble contrast agents injected into the bloodstream to enhance imaging) are not commonly used during image-guided intracardiac interventions at this point in time.

#### 1.4.0.2 Ultrasound for surgical guidance

Several groups have examined the use of unembellished ultrasound for surgical guidance during minimally-invasive cardiac interventions in both laboratory and clinical settings. Downing *et al.* [30] described mitral valve suturing under the guidance of multiplanar 2D TEE in pigs. 2D echocardiography has a relatively high frame rate, spatial resolution and field of view, but it is difficult to continuously align the imaging plane with anatomical structures or tools of interest. 3D echocardiography provides a better perception of cardiac structures and surgical tools in 3D space without as much mental effort on behalf of the clinician, and better quantifies volumes, distances and clinical factors such as ejection fraction and wall motion. RT3D US has been shown to outperform 2D US when guiding simulated surgical tasks performed within a water tank with both manual and robot-assisted tool manipulation, and has been used to close ASDs and repair mitral valves on tissue suspended within a water

tank [31–33]. This same group has reported successful beating-heart ASD closure in swine under epicardial RT3D TTE guidance via a median sternotomy [34]. Experience in actual clinical practice has focused on percutaneous techniques. For example, RT3D TTE has been employed during ASD closure [35] and right ventricular endomyocardial biopsy [36], and RT3D TEE has guided the treatment of congenital defects, mitral and aortic valve interventions, and pulmonary vein ablation [37, 38]. 2D ICE is gaining support as an interventional imaging modality and has been used to guide ASD closure [39] and transseptal puncture [40] and PV isolation [41] during cardiac ablations. Knackstedt *et al.* [42] have also proposed basing surgical guidance for ablation procedures on reconstructed 3D ICE images.

## 1.5 Supporting Technologies

This section describes the sub-systems that make up image-guided therapy systems, with a focus on how intraoperative echocardiography impacts their development or facilitates possible solutions.

### 1.5.1 Tracking

Real-time tracking of moving surgical tools and image sources is required for them to be properly displayed relative to each other. If the ultrasound probe itself is tracked, the surgical guidance system’s reference frame is typically centered on the intraoperative patient. If the ultrasound probe is not tracked, the reference frame may instead correspond to that of the ultrasound images or of another imaging source such as an intraoperative C-arm CT scanner. In either case, real-time imaging should be the ultimate source of information that is relied upon during surgical interventions. If this functionality is provided by intraoperative echocardiography, then the degree to which the other virtual components in the IGT system relate to the ultrasound images is the most important relationship to consider when evaluating system accuracy.

### 1.5.1.1 Optical and magnetic tracking systems

Optical and magnetic tracking approaches are currently the most widely used for intraoperative tracking systems. Both work by determining the position and orientation of a sensor that is rigidly attached to the ultrasound probe or surgical instrument to be tracked. In concert with spatial calibration [43], which determines the transformation between this sensor and the ultrasound image's coordinate system, tracking allows each pixel or voxel within the ultrasound images to be spatially located within the three-dimensional tracking space. Optical systems use a stereoscopic camera to localize a rigid system of reflective markers in 3D space. They are very accurate, but the attachments are bulky, and the line of sight required between the reflective markers and the camera is often interrupted within the busy operating room. Their use in minimally-invasive intracardiac surgery is also limited as they cannot be used to track flexible or jointed tools, or instruments inserted into the beating heart such as catheters, TEE probes or ICE probes.

Even though they are less accurate, magnetic tracking systems do not share these limitations, making them more common in image-guided intracardiac therapy. A magnetic field generator generates an inhomogeneous magnetic field, and the spatially-dependent current induced within the small sensors allow them to be localized. Unfortunately, tracking accuracy and precision is decreased by the proximity of ferromagnetic materials, from which surgical tools and OR equipment are commonly made, or of an ultrasound probe [44] or X-ray gantry [45].

### 1.5.1.2 Image-based tracking

Image-based tracking uses techniques from image processing and computer vision to localize surgical tools within streaming 2D or 3D ultrasound images, eliminating the cost and setup restrictions of external tracking systems. Although instruments can be separated from the blood pool quite easily, image-based tracking is difficult because tools and tissue have similar intensity profiles [46] with no clear edges between them, and because of ultrasound's low SNR and imaging artifacts. Despite these challenges, algorithms must still be relatively simple so as to achieve real-time tracking.

Image-based tracking methods can be separated into segmentation-based and marker-based approaches. Both typically assume that surgical tools are rigid cylindrical shafts with distal actuators. Segmentation-based methods include those of Novotny *et al.* [47] for RT3D US and Ortmaier *et al.* [48] for 2D US. Their approaches rely on basic and adaptive thresholding, respectively, and will likely fail whenever tools come in contact with tissue, which was not tested. Linguraru *et al.* [46] used RT3D US to successfully track a tool within the beating porcine heart. Their non real-time implementation used voxel classification from intensity probability distributions, neighborhood analysis, principal component analysis to identify long, thin objects, and a water transform.

Tools may also be supplemented with physical markers to facilitate their localization under ultrasound. Stoll *et al.* [49] added three or four loops along each surgical tool which uniquely determined the tool’s position and orientation when imaged in cross-section. They showed successful tool tracking under 2D US within the laboratory, and a real-time implementation has also been demonstrated within the beating porcine heart for RT3D ultrasound [50].

Despite these promising studies, image-based tool tracking has not been widely adopted in IGT. The technique is limited by the difficulty of continuously keeping surgical instruments within the 2D or 3D ultrasound image, especially when simultaneously trying to image surgical targets. Also, none of the studies described above demonstrated simultaneous tracking of more than one tool, so their effectiveness for this task is unknown.

### 1.5.2 Improving ultrasound imaging

This section reviews strategies developed to increase the image quality and field of view provided by echocardiography, with the aim of simplifying subsequent computational tasks such as image registration and segmentation, and of improving image perception when ultrasound is displayed within surgical guidance systems.



### 1.5.2.1 Speckle reduction

Speckle, which imparts ultrasound’s characteristic “grainy” appearance, arises from interference patterns caused by microscopic scatterers. While speckle is not “noise” in that it represents genuine ultrasonic signal, it does interfere with image processing and so speckle reduction is a common first step. Speckle reduction also improves optimization in intensity-based image registration by smoothing the similarity function and improving the capture range [51–53] (note that intensity quantization can be used to reduce noise and smooth the similarity function as well [51, 54]). Although a variety of sophisticated speckle reduction algorithms exist (ex. [55–58]), IGT researchers often use relatively simple smoothing techniques such as median filtering [48, 51, 52, 59–62], Gaussian filtering [49, 63, 64], “sticks” filtering [65] (used by [66]), and nonlinear anisotropic diffusion [55] (used by [53, 67, 68]). Common procedures such as resampling ultrasound data to Cartesian coordinates or downsampling images to reduce their size have the side-effect of smoothing images and therefore reduce some speckle as well.

### 1.5.2.2 Artifact removal

Even though they are often useful from a diagnostic perspective, artifacts complicate ultrasound processing. Common artifacts include shadowing (dark regions caused by strong reflectors or attenuators), enhancement (bright regions caused by weak attenuators) and reverberations (duplication of the same structure at different depths that results when the ultrasound signal, on return to the transducer, is redirected back into the tissue by strong reflectors or the ultrasound probe itself). Probe pressure artifacts show shallow tissue distortions inflicted by the pressing ultrasound probe, and may be important to consider in TTE (ex. [52]) but are most likely insignificant in TEE. Finally, the ultrasound beam is not infinitely thin and its depth dependent slice thickness can measure several millimetres at depths away from the focal point, and so partial volume effects are common.

A number of ultrasound processing algorithms incorporate an artifact removal step. Voxels representing shadow artifacts have been identified with a simple inten-

sity threshold [60] or, to detect distal shadows in particular, by following each scan line from the bottom to the top and tagging all voxels as shadows until a condition is reached such as minimum intensity [63, 69], minimum intensity variance [70], or minimum correlation to a heuristic function [66]. Probe pressure artifacts have been corrected by simply ignoring a section of data in the near field [63, 69, 70]. Alternatively, probe pressure artifacts and tracking errors can both be reduced in free-hand ultrasound reconstruction by better aligning each 2D US image to its neighbors with non-rigid image registration [62]. Finally, Zhong *et al.* [71] incorporated depth-dependent slice-thickness measurements to automatically adjust manually-segmented endocardial surface contours in 2D ICE.

### 1.5.2.3 Spatial compounding and mosaicing

Spatial compounding combines multiple 2D, reconstructed 3D, or RT3D images acquired from similar viewpoints into a single image to reduce speckle and noise, strengthen edges, and increase the SNR, contrast and spatial resolution. Compounding also reduces the influence of artifacts that appear in only a subset of the images and the dependence of US imaging on the angle between the US beam and local structures [72–75]. Real-time 2D spatial compounding using electronic lateral beam-steering has been integrated into a number of commercial US machines. Ultrasound mosaicing (also known as extended field of view imaging or, for 2D images, panoramic imaging), fuses images acquired from more disparate locations, with the same advantages as spatial compounding plus the additional benefit of extending the field of view [54, 76, 77]. Besides the commonly used RT3D stitching functionality described previously, extended field of view imaging is not particularly common when imaging the heart.

Spatial compounding and mosaicing are accomplished in the same two steps. A “registration” step first aligns the multiple input images so that the composite image does not suffer from excessive blurring [72]. This is implicitly achieved if using a tracked and spatially calibrated ultrasound probe. Any residual misalignments can be overcome using rigid or non-rigid image registration, which may also be used alone

without tracking. Registering all of the images simultaneously using an ultrasound-specific similarity measure has been shown to work best for this task [78]. Image-based registration may be particularly difficult during ultrasound mosaicing, as there may not be an extensive overlap between the images to be fused. Second, a “compounding” step combines the intensities of each image to generate the combined image. The most basic approaches calculate the average [72, 74–76, 79] or maximum [79, 80] intensity, the later of which also reduces the influence of shadows.

Most cardiac applications of these technologies concern multiview spatial compounding, where the same cardiac structures are imaged from widely varying viewpoints or different acoustic windows. The goal is to improve image quality and generally not to vastly extend the field of view. Compared to the original datasets, Ye *et al.* [81] showed that combining two apical and parasternal 4D TTE datasets of the left ventricle (LV) generated more complete derived surface models, and that ejection fraction and volume-time measurements showed better correspondence to those from MRI. This group used feature-based compounding instead of the more common intensity-based compounding: after automatically identifying LV features in each 2D US image (that make up the ultrasound reconstructions) and using the tracking and spatial calibration matrices to transform them into 3D space, a surface was fit to the set of combined features. Both Soler *et al.* [80] and Rajpoot *et al.* [79] compounded two and 3-6 apical RT3D TTE images, respectively, using an intensity-based rigid registration without US probe tracking. This is possible because the original images are quite similar and have a large overlap. Both groups also developed sophisticated compounding techniques. Soler *et al.* present “generalized averaging”, whose behavior depends on whether features are judged to appear in one or both of the images, and “multiview deconvolution”, which considers ultrasound’s anisotropic point spread function. Rajpoot *et al.* treat low and high frequency image components differently with the goal of suppressing noise and maintaining the amount of speckle in the multiview image while combining all of the major features present in each input image. Grau and Noble [82] used a basic manual registration to combine apical and parasternal RT3D US images while introducing a weighted averaging method for the compounding step. Their compounding algorithm prevents strong features existing

only in a subset of the input images from being diluted by the images in which they are absent, based on the information content of each voxel and the ultrasound beam's incidence angle. Later, the same group presented a rigid registration method for use in multiview imaging [53].

### **1.5.3 Patient, cardiac and respiratory motion compensation**

Patient, cardiac and respiratory motion present some of the most significant challenges in image-guided cardiac therapy. Surgical guidance fundamentally relies on the correspondence between the virtual environment and the intraoperative patient, which is invalidated in the face of the rigid and non-rigid transformations induced by these three factors if they are not addressed. Relative motion must also be considered when using imaging techniques that combine multiple images (such as ultrasound reconstruction, spatial compounding and mosaicing) or when performing image registration. Especially when a rigid transform is used, any datasets to be registered should correspond to the same cardiac and respiratory phases.

#### **1.5.3.1 Patient motion**

Any gross patient motion is typically compensated for by working with a coordinate system centred on the patient him/herself, rather than using a fixed intraoperative reference frame. This is easily accomplished by performing all optical or magnetic tracking relative to a reference sensor taped or sutured to a stable region on the patient's torso [83]. This will work as long as the reference sensor and the patient's heart form a rigid system throughout the entire procedure. Unfortunately, this assumption may be invalidated if the reference sensor is attached to loose skin, or because of "heart shift" within the chest wall caused by patient repositioning, respiratory motion, the surgical access or tool manipulation within the heart.

#### **1.5.3.2 Gating and motion compensation**

To compensate for cardiac and respiratory motion, each IGT system must provide solutions for gating and/or motion compensation, both of which can achieve the goal

of making the surgical scene internally consistent. Gating identifies data acquired at the same phase of the cardiac or respiratory cycle, typically by detecting the endpoints of each cycle and interpolating within, assigning one or a range of data to each phase. IGT systems that rely on gating address cardiac or respiratory motion by only integrating data acquired during the same phase. Motion compensation works by applying a transformation to data that “undoes” the transformation caused by motion relative to a certain phase of interest, which is often the end-diastolic cardiac phase or end-expiration respiratory phase because they correspond to minimal motion. This allows data acquired at different phases to be integrated together without motion artifacts. Real-time intraoperative image-to-patient registration may compensate for motion if performed continuously during the intervention, especially if using a non-rigid transform, but the challenges of developing such algorithms have motivated development of alternative approaches as well.

Many gating and motion compensation algorithms using a variety of data sources exist, including methods based on echocardiographic images. They may be used in IGT for offline processes, for example for image-to-patient registration. Intraoperative algorithms must be real-time and fully automatic, but relatively computationally inexpensive so as to preserve resources for other IGT tasks. They must also be robust to changes in cardiac or respiratory cycle duration, rather than determining a mean rate for example. The ideal image-based algorithm will also remain accurate even when instruments intrude into the imaging plane or when the viewing angle changes as the ultrasound probe is moved.

### 1.5.3.3 Cardiac motion

Cardiac gating is most often performed using prospective or retrospective ECG-gating because the ECG signal is easily retrievable, signal analysis is simple and it is a relatively accurate surrogate for the heart’s underlying physical state. Each cardiac cycle is typically delineated using the prominent QRS complex (R-wave), which can be easily identified by signal thresholding or, for retrospective gating, by identifying the minima or maxima in the signal or its slope. Somewhat complicating matters is

variability in the patient’s heart rate, which changes the relative durations of systole and diastole (the diastolic phase’s duration is the more sensitive to patient heart rate of the two). Gating schemes that linearly interpolate within each cardiac cycle do not account for this effect (ex. [53, 59–61, 84]), while schemes that allocate timepoints to each cardiac phase based only on the time passed since the previous R-wave do not even incorporate measurements of the heart rate at all (ex. [85]). A better approach is to ignore those cardiac cycles whose durations differ too much from an expected value (ex. [86]) or to scale the systolic and diastolic phases separately with piecewise linear interpolation (ex. [52, 87, 88]). Although ECG-gating is widely used it is not ideal, as the heart’s electrical patterns do not necessarily reflect its physical motion, a latency may exist between the ECG and ultrasound streams that requires synchronization, and because weak, irregular or unpredictable ECG signals are common in cardiac patients.

A potentially more accurate method is image-based cardiac gating, which uses the echocardiography images that we want to gate themselves or images from a secondary source such as Doppler or M-mode ultrasound imaging. Compared to the ECG signal, these images provide a more comprehensive indication of the heart’s pose, which is likely to be more robust to arrhythmias and more suitable if the patient’s ECG signal is likely to be erratic. Manual identification of end-systolic or end-diastolic frames based on intracardiac chamber volume and/or the cardiac valves (ex. [52, 53]) may be feasible for some offline tasks, but is generally not suitable for use in image-guided therapy. Most automatic image-based cardiac gating methods typically generate a one-dimensional gating signal by exploiting periodicities in image intensity or in the intensity changes between consecutively acquired images. Optima in this signal are then detected in much the same way as the R-wave is detected in ECG signals, and the same heart rate variability considerations apply when interpolating between each detected cardiac cycle.

Development of image-based cardiac gating has concentrated on reducing artifacts in 3D ultrasound reconstructions acquired either of the coronary arteries using IVUS or of the fetal heart, since an ECG signal is often unavailable or unreliable in these cases. Metrics used to bin together IVUS images acquired at the same cardiac phase

have been as simple as average intensity, the mean absolute intensity difference between consecutive images, or normalized cross-correlation [89–91]. More complicated approaches have instead used multidimensional feature vectors [92] or automatically segmented image contours [64]. However, image-based cardiac gating algorithms developed for IVUS are not necessarily extensible to echocardiography because they are typically designed for offline use and may fail when confronted with images acquired from different viewpoints. Gating algorithms designed for fetal echocardiography have analyzed periodic M-mode signals [93] or have correlated B-mode images themselves [94, 95]. Other groups analyzed simultaneous Doppler imaging of the fetal umbilical artery [96] or the fetal heart itself [97]. Somewhat similarly, Rubin *et al.* [98] gated cardiac MR imaging using Doppler images of the great vessels in the neck. Such Doppler-based approaches may be feasible for intraoperative cardiac gating as well.

Image-based cardiac gating methods using the echocardiographic images themselves have been fairly limited, but show great promise as an alternative to ECG-gating during IGT. Treece *et al.* [99] present a simple approach that can be applied to B-mode or Doppler 2D echocardiography. A signal of cardiac phase versus frame index is generated by specifying a few regions of interest showing periodic cardiac motion, and simply counting the number of pixels that fall within a user-specified threshold. This technique was extended by Karadayi *et al.* [100] to eliminate user interaction by instead generating the cardiac phase signal from each image’s intensity centroid along the depth direction. However, the signal processing required to smooth these noisy signals operates on them in their entirety, reducing the approach’s suitability for real-time processing. Sundar *et al.* [101] designed a fully-automatic, real-time image-based cardiac gating method for use during IGT with 2D or 3D images of all modalities. This group used the relationship between the phase shift between two consecutively acquired images and the translation of the moving features within them. The phase shift between each image and its predecessor is integrated over all frequencies to derive a signal of “energy change” versus time that is intensity-independent, resistant to noise and whose maxima or minima can be automatically detected for gating. The energy change signal encapsulates both cardiac and respiratory motion, but cardiac gating information can be extracted with a simple bandpass filter at 1 Hz.

Testing with 3D intracardiac echocardiography showed continuous gating even when the catheter was moved, although further quantitative evaluation of gating accuracy is required.

#### 1.5.3.4 Respiratory motion

The influence of respiratory motion on the heart is significant and highly specific to each individual. The movement of the diaphragm, ribcage and lungs causes cardiac deformations on the order of 1-3 cm [87,102,103], along with changes in stroke volume, heart rate and cardiac output [104]. For most patients, gross cardiac motion on inspiration is most significant in the inferior direction and secondarily towards the anterior [102,103]. Early reports judged that respiratory-induced heart motion could be well-approximated by a global translation along the superior-inferior (SI) axis that is linearly dependent on the motion of the diaphragm in this same direction [102]. However, subject-specific hysteresis causes the cardiac position to be different for the same diaphragmatic position on inhale compared to exhale [105]. In addition, more recent studies revealed that rotational, affine and non-rigid transformations are significant in most cases [87, 103]. The non-rigid transformations deform the cardiac chambers and alter the relationships between intracardiac structures such as the pulmonary veins. Over the respiratory cycle, these deformations can be on the order of 2-4 mm even after rigid motion has been accounted for [106,107].

The simplest way to deal with respiratory motion is to try to eliminate it by performing all critical data collection, including preoperative and intraoperative imaging, and therapy delivery at the same breath-hold position [106]. Unfortunately, voluntary breath-holding is far from reliable as it does not necessarily suspend all diaphragmatic (and therefore cardiac) motion regardless of patient effort [108,109]. The limited time period during which patients can hold their breath (typically less than 30 seconds) also poses severe restrictions, especially in patients with chronic obstructive pulmonary disease (COPD) or congestive heart failure or for the significant percentage who cannot reliably hold their breath at all [110]. It is also nearly impossible to ensure that patients consistently replicate the same breath-hold if more than



one is required. Respiratory suspension in artificially-ventilated anesthetized patients may be simpler than voluntary breath-holding, but even in these cases the potential for respiratory drift is high, the duration of respiratory suspension is limited, and multiple breath-holds may not be consistent.

Image-guided therapy can be performed while the patient breathes freely using respiratory gating and/or motion compensation. Ideally, any preoperative images should also be acquired in the same way because the inhale and exhale positions during free breathing are different than those during breath-holds [111]. However, in practice preoperative images are often acquired at held inspiration.

Some respiratory gating and motion compensation algorithms operate on a surrogate respiratory signal such as the transforms retrieved from a tracked external fiducial marker. Sun *et al.* [85] performed respiratory gating in their fused ICE-CT system by registering only those ICE images acquired when a tracked sensor on the patient’s chest showed little motion along its z-axis. Such algorithms work by assuming that the periodicity of the tracked fiducial’s movement mirrors the periodicity of the respiratory cycle. Motion compensation algorithms that create a patient-specific motion model go one step further by actually correlating the periodic changes in an easily measurable surrogate signal to the motion of an internal target such as the heart or a specific intracardiac structure. During the intervention, the location of the unseen target can then be estimated based on the surrogate signal alone. If the non-rigid cardiac deformations induced by respiratory motion are to be recovered, they must be incorporated into the motion model. Khamene *et al.* [112] create a motion model by simultaneously acquiring high-quality preoperative images of target structures and the position of a rigid system of external markers over multiple respiratory cycles. Atkinson *et al.* [113] adjusted the position of 2D TTE images in the superior-inferior direction for respiratory motion compensation during freehand ultrasound reconstruction. This was accomplished by correlating the anterior-posterior motion of a tracked marker positioned on the patient’s umbilicus to the patient’s diaphragm in tracked 2D ultrasound. This strategy assumes that respiration-induced cardiac motion is dominated by a superior-inferior translation that is linearly correlated with diaphragmatic motion, which as previously described is approximate and not strictly

true.

Although tracked external markers are simple to use, their motion may be only roughly correlated with the heart’s motion throughout the respiratory cycle and they may be obtrusive within the OR. Intraoperative imaging of the right diaphragm with 2D ultrasound may be used to generate a respiratory signal during interventions [114], mirroring the navigation pulse technique used in cardiac and coronary MRI. Such a signal can be input into motion compensation algorithms [115] and has even been used to gate cardiac MRI [116], but continuous diaphragmatic ultrasound imaging may be cumbersome to maintain within the operating room.

Cardiac ultrasound images can themselves serve as a comprehensive and OR-friendly source of information for respiratory gating or motion compensation, although performing real-time analysis is challenging. Image-based approaches eliminate errors associated with surrogate respiratory signals, including errors in measuring the surrogate and when deriving the relationship between the surrogate and the underlying cardiac motion. The gating system proposed by Sundar *et al.* [101] that was described above for use in cardiac gating can also be applied to respiratory gating using a low-pass filter on the gating signal, although this was not explicitly tested by the authors with echocardiography. Two respiratory motion correction methods warping a single preoperative 3D MR image acquired at end-exhale to match the patient’s current respiratory phase were presented by King *et al.* [117]. The preoperative image was first transformed into the space of an end-exhale RT3D image. This end-exhale image was then registered to the current RT3D image by iteratively searching over the space of either 1) the six parameters characterizing a rigid transformation, or 2) a single parameter representing the translation of the right diaphragm, which was then input into a respiratory motion model that returns an affine transformation. The second motion compensation method is superior as it guarantees that the final transform is physiologically realistic and derives the registration transform using more information than available from basic image registration between two RT3D images. However, a real-time implementation of this method is required for it to be useful in IGT.

### 1.5.4 Image registration

Registration aligns corresponding images or other spatial data from different coordinate systems by finding the transformation that best aligns one of the datasets, designated as the moving dataset, to the other, termed the reference dataset. Respiratory motion, skin flexibility and the non-rigidity of the heart within the chest wall precludes the use of external skin markers for registration, as has been proposed for neurosurgery for example, motivating image-based methods. Many excellent articles have reviewed the problem of registering medical images, including general reviews presented in [118–120], a cardiac-specific review [121] and an ultrasound-specific review [122].

Because intraoperative echocardiography either forms the patient coordinate system or is related to it via tracking and spatial calibration, multimodal registration between intraoperative echocardiography and preoperative MR or CT images effectively performs image-to-patient registration. As previously described, monomodal image registration between echocardiography images is also useful in ultrasound compounding and mosaicing, and for cardiac and respiratory gating and motion compensation. Methods have also been developed to register echocardiography with single photon emission computed tomography (SPECT) [61,88] and positron emission tomography (PET) [123], but these nuclear medicine images are not typically used during cardiac interventions. At this point in time, no image-based registration method has been developed to register cardiac ultrasound to 2D X-ray fluoroscopy.

Many demands are made of image registration algorithms proposed for intraoperative use. Although different applications have different requirements, a registration accuracy of 1-5 mm is generally considered to span the range between the best results that are currently achievable and the maximum error that may still be clinically acceptable. Algorithms must also be fast with minimal user interaction. The ultimate goal is to develop a highly accurate algorithm that updates the registration transform in real-time with no user interaction and while remaining robust to varying intraoperative conditions, but this is very difficult.

The following sections describe multimodal and monomodal image registration

algorithms that use echocardiography. Registration transforms may be rigid, rigid + uniform scaling, rigid + nonuniform scaling, affine, or non-linear. However, echocardiography’s relatively low image quality poses problems for non-rigid image registration because the many degrees of freedom required cannot often be estimated with high confidence and the risk of generating spurious transforms is high. Image registration algorithms include feature-based approaches that automatically or manually extract features such as point coordinates, contours or surfaces from the images and use them to drive the registration. Feature-based registration algorithms commonly used with echocardiography can be further split into point-based registration, anatomical feature-based registration and point-to-surface-based registration. Surface-based registration, which minimizes the distances between two surfaces, is not commonly used with echocardiography because of the difficulties associated with segmenting ultrasound images. Feature-based registration algorithms typically rely heavily on manual interaction, as automatic feature detection and/or segmentation algorithms are not yet accurate and robust enough to achieve widespread use. In contrast, intensity-based algorithms directly use the intensities of each image’s pixels or voxels. While both feature-based and intensity-based approaches are commonly used for multimodal registration between echocardiography and MR or CT, intensity-based methods are typically preferred for monomodal echocardiography registration.

#### 1.5.4.1 Point-based registration

Point-based registration aligns two sets of spatial coordinates, which are typically generated by manually localizing image features, by minimizing the distances between corresponding coordinate pairs [124]. Although this approach is the simplest registration approach from a computational perspective, it is often very difficult to use successfully with echocardiography. Manual identification of intracardiac landmarks is relatively simple in preoperative images, but is notoriously difficult to perform accurately in echocardiography, since the heart generally does not contain many distinctive features and because of ultrasound’s relatively poor image quality. For example, Leung *et al.* [125] found that their experienced users localized landmarks in RT3D

echocardiography with a mean intraobserver variability of 3.3-7.0 mm and a mean interobserver variability of 3.8-7.4 mm. Okumura *et al.* [126] avoided this problem by performing a rigid point-based registration in an animal model based on implanted fiducial clips, but this strategy does not extend to clinical use for practical reasons. Point-based registration is therefore often used solely as an initialization step before running a second registration algorithm.

#### 1.5.4.2 Anatomical landmark-based registration

Performing registration based on anatomical landmarks goes a step beyond point-based methods by registering lines, contours or surfaces outlining important anatomical structures. Such algorithms may constitute the final registration step or may be used to initialize an intensity-based algorithm. Registration accuracy will typically be best in the vicinity of the structures driving the registration and may be quite poor elsewhere. However, this may not be important in the context of the clinical procedure if the surgical targets themselves are the anatomical structures used to perform the registration.

Linte *et al.* performed image-to-patient registration using 3D splines outlining the intraoperative mitral and aortic valve annuli, which were manually defined using tracked 2D TEE [84] or a sparse 3D TEE ultrasound reconstruction [127]. These two features were rigidly registered to their counterparts manually traced on preoperative MR or CT images, first by aligning the spline centroids and their normals, and then iteratively minimizing the distances along the spline contours themselves. Ma *et al.* [128] registered RT3D TTE and 3D MR images using a rigid transform by iteratively aligning both the centreline of the aorta and the left ventricle's endocardial surface. Both of these features were automatically detected in the MR images, while approximately three minutes were required to manually identify a few coordinates in the echocardiographic images that were subsequently used to extract the aortic centreline and LV surface. Compared to point-based registration, the grosser features used in anatomical landmark-based registration may be easier to localize, while registration speed is maintained.

#### 1.5.4.3 Point-to-surface-based registration

Many commercial EP systems incorporate a 3D surface model derived from preoperative MR or CT imaging, which is registered to the patient by collecting a set of point coordinates along the intraoperative endocardial surface and minimizing the distance between them and the preoperative surface. This set of intraoperative coordinates may be localized manually using a tracked catheter [129–131] but can also be generated by acquiring a sweep of 2D ICE images with a tracked probe. A set of intraoperative endocardial contours from which the 3D point set is generated is then derived by segmenting each image, typically manually [132], although semi-automatic [133] and automatic [134] algorithms have been described. Point-to-surface registrations are also very user-intensive and time-consuming. First, the preoperative MR or CT image must be segmented, although this is often required regardless to generate a preoperative surface model for visualization as part of the surgical guidance system. More significant is the effort required to collect the intraoperative points with a tracked catheter or by manually segmenting each ICE image. When performing these types of registrations, it is important to also consider the rotational symmetry of many heart chambers, and to either include data from more feature-rich cardiac regions such as the area around the pulmonary veins, or to constrain the registration by simultaneously incorporating a point-based registration with a few coordinates identified in both coordinate systems [131].

#### 1.5.4.4 Intensity-based registration

Intensity-based registration is often the preferred approach for intraoperative use as it does not rely on potentially inaccurate feature localization, requires no user interaction besides an optional manual initialization, and in some cases can be implemented in real-time. Intensity-based registration uses an iterative approach to determine the registration transform. At each iteration, the current transformation estimate is applied to the moving image, the transformed moving image is compared to the reference image using a similarity measure (or, equivalently, a dissimilarity measure), and a search algorithm is applied to explore the transformation's param-

eter space with the aim of maximizing the similarity measure (or minimizing the dissimilarity measure).

Echocardiography’s relatively low SNR, few strong features, small field of view, common artifacts, viewpoint dependence and anisotropic spatial resolution pose challenges for any type of image registration, but has the strongest impact on intensity-based registration. In particular, the optimal transform may not correspond to a global maximum in the similarity measure but only to a local maximum, necessitating accurate registration initialization [51, 60]. The algorithms described below that do estimate capture range report capture ranges between 5-45 mm, although these are calculated with different definitions for what qualifies as an “acceptable” registration. Optimization may also be difficult if there no sharp peak in the similarity measure at the local optimum because few features are present in the echocardiographic images to be registered (ex. [135]). Most common image registration optimization algorithms have been applied for use with echocardiography, including “best-neighbor” [85], gradient ascent/descent [59, 60, 80], Powell optimization [53, 79] and variants of simplex optimization [51, 52, 125, 136]. Optimization is often done in a “coarse-to-fine” framework [53, 69, 79, 85, 125, 136–139] to increase speed and reduce the impact of small features on the initial transform estimation. However, these multiresolution approaches can be sometimes be problematic if the first low-resolution registration has a small capture range or converges on the incorrect local optimum.

Registration between ultrasound and MR or CT is difficult for any organ because ultrasound primarily images the interfaces between different tissues, while MR and CT images show relatively homogeneous regions for each tissue type. Strategies to overcome this include preprocessing the MR/CT image to increase its resemblance to ultrasound by highlighting boundaries [66], incorporating the MR/CT image’s gradient magnitude in the similarity measure [140], or preferring transforms that align edges in the MR/CT image with high intensities in the ultrasound image [70]. Registrations between ultrasound and MR/CT have also been based on aligning a derived probability map of a specific property such as tissue type [63], registering a simulated ultrasound image from CT instead of the original CT [70], or incorporating a simulation of ultrasound from CT into the similarity measure by adding terms for

ultrasound reflection and echogeneity [141] or reflection alone [136]. More straightforward registrations that do not require extensive preprocessing often use the mutual information (MI) similarity measure [142–144] because it assumes only a statistical correspondence between intensities in the two image types.

Algorithms developed to register MR or CT with 2D echocardiography include that of Sun *et al.* [85], who rigidly registered a set of tracked 2D ICE images with an intraoperative C-arm 3D CT image, both acquired with ECG-gating. Intended for use during EP procedures, their algorithm maximized the normalized cross-correlation (NCC) between each 2D ICE image and the CT image’s thresholded gradient magnitude. Each registration required approximately one minute to converge. Huang *et al.* [59] proposed a fast rigid registration algorithm between tracked 2D echocardiography and a time series of preoperative CT images. The speed and accuracy with which mutual information could be maximized between each 2D US image and its corresponding 3D CT image (as identified with ECG-gating) was increased by using a very close registration initialization based on aligning a sparse 3D US reconstruction. Their system is intended for updating the pose of a preoperative images displayed within an augmented reality environment in real-time throughout the procedure. Finally, Zhang *et al.* accommodated non-rigid deformations between cardiac ultrasound and MR images by performing affine registrations between 2D MR and 2D US [137], and poly-affine registrations between 2D MR and 2D US [138] and 2D+t MR and 2D+t US [139]. In the latter, a temporal alignment first created a single 2D US image for each 2D MR image by first automatically identifying the maximum contraction and end-diastolic ultrasound image in each sequence, and then merging the two ultrasound images whose cardiac phases were closest to the MR image’s by interpolating the motion field between them. Spatial alignment between each MR-US pair was based on maximizing the mutual information between their phase images, which emphasizes the alignment of local image structures without assuming a clear relationship between the US and MR intensities [145]. A poly-affine transformation was found by iteratively creating control points in feature-rich areas that simultaneously suffered from poor alignment, and generating a local affine transformation for each surrounding block.



A number of authors have also taken advantage of the 3D information provided by RT3D US for rigid image-to-patient registration. Huang *et al.* [60] presented a rapid (approximately 1 Hz) registration method between tracked RT3D TTE and preoperative 4D CT that closely follows their tracked 2D US method described above. However, mutual information was calculated only on a pseudorandom subset of voxels to increase speed, and a single rigid transformation matrix was assumed to be valid over the entire cardiac cycle. Real-time rigid registration between untracked RT3D ICE and an end-diastolic 3D C-arm CT image was presented by Wein *et al.* [136]. Each registration took approximately 0.6 seconds with a GPU implementation and image-based cardiac gating. A single 3D ICE image could be used, or, to increase the amount of ultrasound data available for registration, an untracked sweep of 3D ICE images were simultaneously registered to an ultrasound image simulated from CT and to each other (by maximizing the mean Pearson’s correlation coefficient and minimizing the mean sum of absolute differences, respectively). Additional work is required for this technique because their algorithm’s accuracy compared to manual registration is currently quite poor. King *et al.* [69] used a unique similarity measure to rigidly register one or more tracked RT3D TTE images with a preoperative surface model segmented from 3D CT, and found an improvement in the registration when two ultrasound images from different acoustic windows were registered at the same time. At each registration iteration, they used the current estimation of the registration transform to calculate the statistical likelihood that the ultrasound image(s) arose from the surface model, using a simple model of ultrasound physics incorporating concepts of tissue scattering strength, differences in acoustic impedance at tissue boundaries, the ultrasound beam’s slice thickness, and the angle of incidence between the US beam’s normal and that of the surface model. Non-rigid registration between cardiac ultrasound and MR or CT images constitute future work in the field. Although Zhang *et al.* [139] claim that their 2D+t non-rigid registration algorithm is extensible to 3D+t sequences, they show results only between 2D MR images and 2D short-axis and long-axis US images.

Most intensity-based monomodal echocardiography image registration algorithms estimate rigid transformations, or affine transformations at their most complex. This

is not as limiting as it may originally seem, because many applications register cardiac-gated images acquired from the same patient and within a short time interval. As previously mentioned, non-rigid transformations become difficult to estimate using echocardiography. For example, Shekhar *et al.* [51] found an increase in registration error and a decrease in capture range when attempting to recover known deformations with increasing transform complexity from rigid to affine. However, previously developed elastic monomodal ultrasound registration algorithms originally demonstrated for use in the breast [146] and liver [147] may be extended to echocardiography in future.

Along with simpler transforms, simpler similarity measures become feasible in monomodal image registration because stronger assumptions can be made regarding the relationships between the intensities of the echocardiographic images to be registered. For example, the sum of absolute differences (SAD) [125, 136], sum of squared differences (SSD) [125, 148] and normalized cross-correlation (NCC) [79, 125] similarity measures have been used in addition to mutual information (MI) [51, 52] and normalized mutual information (NMI) [125].

Many monomodal registration algorithms are concerned with registering RT3D images to enable multiview image compounding. Rajpoot *et al.* [79] rigidly fuse apical RT3D images within a multiresolution, multistart framework by maximizing normalized cross-correlation with Powell optimization, while Soler *et al.* [80] accomplish a similar task using gradient descent to optimize a robust similarity measure that minimizes the influence of ultrasound artifacts and noise. Rather than registering their untracked RT3D TTE apical and parasternal images directly, Grau *et al.* [53] derived a similarity measure based on local phase and orientation. This should overcome the intensity differences between apical and parasternal images caused by differences in artifacts and ultrasound attenuation, and indeed the authors showed improvement over instead performing registration by maximizing mutual information or NCC.

A registration algorithm optimizing mutual information to align two 3D TTE images with a rigid, rigid+uniform scaling, rigid+nonuniform scaling or affine transform using the Nelder-Mead simplex optimizer has also been developed [51] and applied to the problem of aligning a pre-stress and a post-stress time series of 3D TTE images

in 3D stress echocardiography [52]. For the same task, Leung *et al.* [125] did not register the pre- and post-stress images directly but instead used a sparse pre-stress image generated by combining 3-5 2D standard views extracted from the original. In addition to increasing registration speed, this reduces the influence of artifacts within less important image planes, which were deemed to cause the substantial errors that they sometimes found when registering the RT3D image pairs in their native form. Finally, Herlambang *et al.* [148] generated a time series of high-quality images over the cardiac cycle from a single 3D MR image. The MR image was first rigidly registered to a RT3D US image acquired at the same cardiac phase (by maximizing mutual information). It was then deformed a second time by applying a time series of non-linear B-spline transforms determined by monomodal registrations between this RT3D image and each of a series of RT3D US images acquired over the cardiac cycle (by maximizing SSD).

## 1.6 Image-Guided Therapy Systems

A variety of image-guided therapy systems have been developed for minimally-invasive cardiac therapy that incorporate echocardiography to accomplish surgical guidance and the supporting technologies described above. These include systems fusing echocardiography with MR, CT and/or X-ray fluoroscopy. The clinical applicability of many of these systems has been demonstrated in animal studies and even in human subjects.

### 1.6.0.5 Echocardiography and endoscopic video

Ultrasound has been superimposed onto live 2D or stereoscopic video for surgical guidance during a variety of clinical applications, including needle biopsy [149] and laparoscopic liver interventions [150]. However, endoscopy cannot typically be used during intracardiac beating-heart therapy because it cannot image through blood. Recently, specialized endoscopy cameras designed for use during beating-heart interventions have been developed, and work by either clearing blood from the camera's

vicinity by flushing saline in front of it [151,152] or by enclosing the endoscope within a specialized port into which blood cannot enter [153]. However, such endoscopes have not experienced widespread use. Leven *et al.* [154] integrated flexible laparoscopic 2D ultrasound imaging with the DaVinci robot-assisted surgical system, enabling both 2D and reconstructed 3D ultrasound imagery to be overlaid on stereoscopic endoscopy. For robot-assisted epicardial interventions such as CABG, such a combined visualization may be a significant improvement during surgical guidance compared to the typical reliance on endoscopic video alone, as discussed in [155].

#### 1.6.0.6 Echocardiography and tracked tools

Relatively simple surgical guidance systems visualizing tracked real-time 2D US alongside virtual representations of tracked surgical tools, such as needles or anastomosis (fastening) devices, have been proposed for beating-heart intracardiac interventions [46, 47, 50] and even for prenatal cardiac interventions to guide needles through the maternal abdomen and into the fetal heart [156]. In the robotic system proposed by Hastenteufel *et al.* [157] for minimally-invasive ablation for treatment of atrial fibrillation, models of both their highly-flexible multi-electrode catheter and of the left atrial endocardial surface were created using reconstructed 3D TEE and used to validate the catheter’s pose before ablation. Although overlaying virtual representations of tracked surgical tools onto echocardiography should enhance image interpretability during surgical tasks, greater improvements can be provided by fusing echocardiography with information from other imaging modalities.

#### 1.6.0.7 Echocardiography and MR or CT

This image combination is ideal in many ways, fusing the high quality and large field of view provided by MR or CT with the real-time imaging capabilities of intraoperative ultrasound. Surgical guidance systems that follow this approach typically display patient-specific cardiac surface models segmented from MR or CT imaging, which is most often performed preoperatively, along with texture-mapped 2D or volume rendered 3D echocardiography and virtual representations of tracked surgical

tools. Linte *et al.* [84] propose dividing the tool-to-target manipulation task into two steps, each of which relies primarily on one of the two image sources. Initial navigation towards the surgical target’s vicinity can be accomplished using the 3D context provided by the cardiac model. Afterwards, detailed manipulations of the surgical target should rely primarily on intraoperative echocardiography, which provides real-time feedback of the position and orientation of the target and tool, and also visualizes any discrepancies between the patient’s intraoperative heart and its depiction within the MR/CT image or derived model [158].

Example IGT systems include that described by Linte *et al.* [84], who use feature-based image registration to integrate magnetically-tracked multiplanar 2D TEE with a dynamic patient-specific cardiac model that is segmented from four-dimensional preoperative MR or CT and is synchronized with the intraoperative beating heart using the patient’s ECG signal. Their system is applicable for a variety of procedures performed within the beating heart and has been used to successfully perform minimally-invasive mitral valve replacement and atrial septal defect closure in swine [127]. The visualization provided by the system described by Rettmann *et al.* and Robb *et al.* [129, 159] is similar, with a focus on cardiac ablation procedures. Both systems can also display intraoperative electroanatomical mapping formation (ex. activation times) onto the preoperative surface model for use during EP procedures [160]. Sun *et al.* [85] and Wein *et al.* [136] both fused intracardiac echocardiography with an intraoperative C-arm CT scan in pigs, using cardiac gating to integrate image information acquired at a single cardiac phase only. Sun *et al.* applied a fully automatic intensity-based registration algorithm to align the CT image with a set of magnetically-tracked 2D ICE images, and is the only surgical guidance system fusing MR/CT with echocardiography that integrates intraoperative respiratory gating. Wein *et al.* used untracked RT3D ICE, which was displayed with volume rendering after real-time fusion with C-arm CT data, using an intensity-based registration algorithm and image-based cardiac gating at end-diastole. Duan *et al.* [124] also discuss the situation where only a single static preoperative image is available, and used a stationary, untracked RT3D TTE probe so that the relationship between the two image coordinate systems remained constant. Image-to-patient registration was based

on a single ultrasound image, and their surgical guidance visualization displayed the CT image with overlaid surface models of the left ventricle that were automatically segmented from streaming RT3D TTE.

#### 1.6.0.8 Echocardiography and X-ray

Fluoroscopy and echocardiography are both ubiquitous intraoperative real-time imaging modalities with complementary imaging properties. Fluoroscopy can localize structures in 3D space with from two images acquired from different angles or from rotational X-ray and provides excellent visualization of catheters, other surgical tools, and the vasculature (following the application of contrast agent). In contrast, echocardiography provides better soft tissue contrast and can also be used to collect functional information, for example left ventricular wall motion or synchrony.

Image registration between ultrasound and X-ray is difficult because the intensity profiles of each of these modalities are very different and because the anatomical structures that they image well do not overlap. Therefore, X-ray and echocardiography are usually fused by separately relating each imaging system's coordinate system to a third coordinate system using spatial calibration and/or tracking. Rasche *et al.* [161] describe a planning and guidance system for biventricular pacemaker implantation that is somewhat unusual in that echocardiography serves as a peri-operative (rather than an intraoperative) imaging modality. In the peri-operative step, a 3D surface model of the left ventricle is first generated from optically-tracked RT3D TTE and annotated with functional information. Localizing five external markers on the patient's chest both in the peri-operative optical tracking space, and relative to the intraoperative X-ray coordinate system, gives sufficient information for the LV surface model to be fused with a 3D coronary vessel surface model derived from intraoperative X-ray imaging modeling or projection imaging. Although the results of their accuracy assessment were generally positive, their reliance on tracked chest sensors for registration is not ideal.

Both Jain *et al.* [162] and Gutierrez *et al.* [45] enable real-time image fusion between fluoroscopy and RT3D echocardiography (RT3D TEE and RT3D TTE, respec-

tively) by localizing coordinates within the ultrasound and X-ray images relative to a tracked sensor on the patient bed. This is achieved for ultrasound via magnetic tracking and spatial calibration, and for the X-ray system by rigidly registering landmarks on a calibration phantom localized in X-ray images and with a tracked pointer tool relative to the magnetic tracking system, followed by application of the X-ray system's pose and its intrinsic calibration matrix. However, such a system would require the clinician to manually manipulate the ultrasound probe under continuous fluoroscopy exposure, which poses significant radiation risks. This concern is addressed by the robotic system of Ma *et al.* [163], in which an RT3D ultrasound probe is manipulated by a slave robot controlled by a remote master controller. Registration between the X-ray and RT3D echocardiography images was shown for two patient cases and proceeds in much the same way as in the two previously described papers, and the resulting visualization platform presents an overlay of volume rendered 3D ultrasound or a derived surface model onto 2D X-ray images.

#### 1.6.0.9 Clinical integration: electrophysiological procedures

Clinical integration of image-guided intracardiac interventions is most widespread in electrophysiological procedures. Original computer-assisted EP systems generated a 3D electroanatomical model of the heart (including both spatial and EP mapping information) from a set of 3D coordinates collected by touching a magnetically-tracked catheter to the endocardial heart wall [164, 165]. During the intervention, the surface model was visualized along with the tracked ablative catheter's tooltip and annotations showing previous ablation points. These systems reduced radiation exposure compared to guidance based on intraoperative fluoroscopy alone, but produced very low resolution surface models that were subject to cardiac wall deformation from the pressure of the tracked catheter, sometimes had missing data for hard to reach intracardiac anatomy, and were time-consuming to generate.

Subsequent improvements focused on how to generate the cardiac surface model and on integrating new types of intraoperative data. Surface models generated using a tracked catheter can be replaced by higher-resolution patient-specific models

segmented from a preoperative MR or CT image [130, 131]. These surface models show several anatomical structures that were not clearly visible previously, including the number, location and branching pattern of the pulmonary veins which are often the surgical targets during radiofrequency (RF) ablation. However, intraoperative collection of endocardial surface points is still required in this strategy to register the preoperative surface to the intraoperative patient, thus retaining many of the disadvantages associated with tracked catheters described above. In addition, the accuracy with which the preoperative surface model could be registered to the intraoperative patient depends on the severity of any non-rigid deformations between them.

Recently, intraoperative 2D ICE has been proposed as an alternative to preoperative MR/CT as an imaging source from which endocardial surface models can be generated. ICE-derived surface models have been used within surgical guidance systems used in patients to perform pulmonary vein and linear ablations [126] and left ventricular tachycardia ablations [133]. A set of spatially-localized 2D ICE images acquired by sweeping a magnetically-tracked 2D ICE probe to view the left atrium, left ventricle, pulmonary veins or any other structures of interest, typically with ECG-gating and under respiratory breath-hold or quiet breathing. The endocardial surface contour in each image is then segmented, and is subsequently transformed into the reference catheter's coordinate system using magnetic tracking and spatial calibration information. Combining the contours provided by each 2D ICE image creates a 3D endocardial surface model. This technique is much faster than using a tracked catheter, and because 2D ICE imaging does not deform the heart wall it is likely to be more accurate as well. However, endocardial surface segmentation may be complicated by artifacts in the ICE images, including shadows, and typically requires significant user interaction. Although ICE-derived surface models have a lower spatial resolution than those segmented from preoperative MR or CT, they are generated in the operating room and therefore have the significant advantage of representing the heart in its actual intraoperative state.

Finally, intraoperative endocardial surface data derived from 2D ICE have been used to integrate preoperative MR/CT surface models with the intraoperative patient [134, 135]. This approach takes advantage of ICE's ability to collect endocardial



surface points without deforming the cardiac wall, which should increase registration accuracy while still integrating a high-quality surface model. However, challenges related to ICE segmentation and the possible deformations between the preoperative image and the intraoperative morphology of the heart still apply. Zhong *et al.* [134] presented a fully automatic algorithm to segment endocardial surface points from 2D ICE, and also described a phantom study demonstrating a dramatic increase in registration success enabled by the increased number of endocardial coordinates that can be collected with echocardiography, compared to the small number that can be manually localized with a tracked catheter because of time constraints. Real-time imaging with 2D ICE can also be integrated into the surgical guidance system along with the cardiac model and the representation of the ablation catheter. Such a system has been used by den Uijl *et al.* [132] while delivering therapy to seventeen AF patients.

## 1.7 Validation

The process of validation verifies that an IGT system meets its requirements with respect to accuracy (both trueness and precision), robustness, task completion time, clinical outcomes, or one of many other application-specific measures [166]. It is important to note that when validating IGT systems, like any other kind of “mission-critical” software, worst-case performance is often a more important indicator than average performance and so outliers cannot be ignored [167]. One must also be wary when comparing the results of accuracy assessments performed by different groups, as validation methodologies vary widely and their results are also highly dependent on the specific equipment and human operators employed for testing. Finally, there are significant differences between patients and healthy volunteers with respect to anatomy, ultrasound image quality, cardiac and respiratory rates, ECG signal quality and physical limitations such as breath-holding, which means that testing should be done with the patient population for which a system is developed whenever possible.

### 1.7.0.10 Quantitative subsystem validation

Before evaluating the accuracy of an IGT system during actual clinical tasks, each of its component subsystems, including strategies for imaging, tracking, spatial calibration, gating, motion correction and registration, are often evaluated individually. This allows multiple approaches under consideration to be compared, and also enables a portion of the overall system error to be attributed to each subsystem, allowing efforts for improvement to be focused on those that contribute the most error. Validation based solely from a qualitative standpoint must be followed by a comprehensive quantitative study that calculates trueness and precision statistics in order to ensure that requirements are truly being met.

Initial validation is often performed under controlled conditions within the laboratory using phantoms (objects of known structure). Phantoms used to validate systems intended for use during cardiac interventions include bead [161], cross-wire [154, 163] and geometrical shape [71, 136] phantoms, anatomically-unrealistic cardiac chamber phantoms [84, 129], cardiac chamber phantoms created from a high-quality patient image with rapid prototyping [130, 134], expandable balloon-type phantoms [115, 148], and/or anatomically-realistic stationary [59, 135, 160] or beating [59, 60, 124, 158, 160] artificial hearts. It is important for laboratory studies to replicate as many clinical conditions as possible, including using a phantom that is anatomically realistic, has a dynamic motion pattern similar to that of the beating heart, and has tissue-mimicking imaging properties. Unfortunately, no laboratory setup will exactly replicate the intraoperative setting, and any such deviations will contribute to either over-estimating *or* under-estimating the accuracy of the system under consideration.

Quantitative validation procedures are numerous, and so a review of all of the possibilities is outside the scope of this chapter, and the reader is referred to [166] for a more comprehensive discussion. Most quantitative assessments compare the result of applying an algorithm on phantom, animal or human data to a “gold standard” expected value [166]. If the gold standard is not known absolutely, any errors sustained in its determination will increase the unreliability of the evaluation results. Unfortunately, this is often the case for systems that work on echocardiography. Their

results are often compared to the results of manual interaction, and ultrasound’s relatively poor image quality makes performing tasks such as image registration and segmentation ambiguous even for experts.

#### 1.7.0.11 User studies

User studies simulating surgical tasks within the laboratory often constitute an intermediate pre-clinical step before testing an IGT system on animals or human subjects. The task may be dramatically simplified: for example the user, preferably a clinician, may perform surgical “drills” within a water tank while relying on the surgical guidance system rather than direct vision [31–33]. Increasing in realism, experiments may ask users to guide tools towards small targets within an anatomically-realistic beating heart phantom [158] or to perform a more complicated procedure such as suturing, mitral valve repair or ASD closure on excised tissue [30, 32, 33] or on an intact excised porcine heart [84]. Targeting tasks typically give accuracy and precision statistics (i.e. quantifying the distance between the final tool position and the target), while the results of more clinically-oriented tasks are typically binary (i.e. was the simulated intervention considered successful, or not?).

#### 1.7.0.12 Clinical validation

The final test of an IGT system is to demonstrate its performance during interventions on animal models and, finally, in actual human subjects. Animal studies [30, 34, 85, 126, 127, 136, 156, 160] allow the effectiveness of an IGT system to be demonstrated within a clinical environment. Considerations involving human factors and surgical workflow can also be addressed, and necessary improvements that may not have been apparent during laboratory testing can be identified. In many acute animal experiments, the animal is sacrificed and its heart opened to verify correct therapy delivery, which may be indicated by ablation lines or the location of implanted clips or prosthetic devices. Okumura *et al.* [126] describe a particularly comprehensive assessment of an ICE-guided electrophysiological system in dogs. They implanted fiducial clips on the endocardial heart surface before preoperative

imaging, attempted to target these fiducial points with the ablation catheter during the procedure, and post-operatively calculated the distance between each actual ablation point and the implanted clip.

Once success has been demonstrated on animals, final validation should be performed on human subjects. This addresses residual differences between humans and animals with respect to anatomy, imaging properties, and cardiac and respiratory rates, and can also be used to examine long-term outcomes. Because of the relative novelty of image-guided therapy for minimally-invasive cardiac interventions, most systems have not yet progressed to the point where they can be evaluated in human subjects. As previously described, an exception is the use of tracked 2D ICE during EP procedures [126, 132, 133]. The final step before widespread clinical acceptance and regulatory approval for any IGT system would ideally be a randomized prospective clinical trial [168] comparing surgical accuracy, procedure duration, complication and conversion rates, and long-term effectiveness and freedom from repeat procedures with that of standard clinical practice.

## 1.8 Challenges and Future Directions

### 1.8.0.13 Visualization and human factors

The question of how to best present data within surgical guidance systems is an important one to address [169]. If not done carefully, the many fused images and geometrical models can be visually confusing, leading to lowered scene interpretability, slowed surgical decision making and a decreased ability to accurately perform surgical tasks. This is especially important when designing systems for cardiac interventions, as the many moving objects within the virtual scene may further compound these perception problems. The use of real-time ultrasound imaging also presents challenges because it often requires more mental effort to interpret than other modalities. 3D ultrasound in particular is difficult to display, and the most common visualization technique used, namely volume rendering, often does not display internal structures well without extensive interactive reslicing, while alternatives such as display on or-

thogonal slice planes shows only a few slices through the 3D dataset. Future work should focus on developing specialized display strategies for use with ultrasound data, and on examining the human factors associated with image-guided therapy systems used during cardiac interventions. This may involve studies that vary the information that is shown in the surgical scene, the control of the virtual camera (i.e. surgeon-control vs. delegated-control vs. fixed view [170]), the display hardware used (i.e. wall-mounted displays vs. head-mount display vs. tomographic reflection [171]) and whether data is visualized stereoscopically [172] or not.

#### **1.8.0.14 Echocardiographic imaging**

Many advances have been recently made in ultrasound imaging, including miniaturization in the development of ICE, IVUS and RT3D TEE transducers, and in image quality improvement including the development of coded pulses, tissue harmonic imaging, and adaptive image enhancement techniques [173, 174]. Any future improvements in ultrasonic imaging that increase SNR, remove speckle, reduce artifacts, increase spatial resolution or facilitate image display via post-processing will benefit image-guided therapy, as higher image quality greatly facilitates rapid image interpretation during surgical guidance and also makes processing tasks such as image registration and segmentation easier. Of particular interest for intracardiac interventions is the advent of novel RT3D ICE transducers and the continuing development of RT3D TEE. The spatial resolution, frame rate and field of view of RT3D ultrasound is fundamentally limited by the finite speed of sound in tissue. However, future strategies that improve any of these factors in RT3D imaging would be highly beneficial, as would the development of more robust techniques for freehand ultrasound reconstruction or extended field of view RT3D imaging.

#### **1.8.0.15 Algorithms**

Continued work on each of the supporting technologies for IGT described above is expected to continue. In particular, further development of the following algorithms may have the largest impact on image-guided cardiac therapy: real-time rigid

image-to-patient registration, non-rigid image-to-patient registration, extended field of view RT3D imaging without stitch artifacts, real-time ultrasound segmentation and feature tracking, and image-based cardiac gating methods that are robust to a moving ultrasound probe and instruments appearing within the field of view. Algorithms that are tailored for use with ultrasound, for example by considering typical imaging artifacts and the finite thickness of typically US beams, will also be beneficial where appropriate. Many algorithms still require significant user input, and further automation would increase clinical penetration of IGT strategies by minimizing manual interaction within the operating room. Finally, increased processing speed is desirable for many algorithms, and in some cases may be accomplished by a specially-constructed hardware architecture [175], a parallel software implementation [176] or, as is more common for IGT algorithms operating on echocardiography, using GPUs (ex. [50, 77, 101, 136, 148]).

#### **1.8.0.16 Instrumentation**

Minimally-invasive cardiac surgery offers challenges in tool design, even when image-guided therapy is not used, for example requiring specialized collapsible valves or patches, fastening devices and laparoscopic instruments. However, employing tracked echocardiography as a mainstay of surgical guidance presents additional challenges. Traditional tools are made of metal, and so their exact position and orientation is not easily determined under ultrasound, they become less clearly visible when angled parallel to the ultrasound beam, they cause strong reflections and shadow artifacts under ultrasound imaging and they have an intensity distribution similar to tissue [46, 177–179]. Metallic instruments also reduce the accuracy of magnetic tracking systems. Future tools constructed from plastics or with ultrasound-specific coatings are required to reduce these negative effects. In addition, ultrasound probes themselves may require modification to facilitate image-guided intracardiac therapy. ICE transducers may be integrated with ablative devices within the same catheter [180], or magnetic tracking sensors may need to be integrated within the probe housing of transesophageal transducers [181].

### **1.8.0.17 Open systems**

Constant communication between imaging, tracking, robotic and data processing systems is required during image-guided therapy. However, at this point in time, access to ultrasound images and data is often restricted. Although few open alternatives do exist [182–184], this issue continues to hamper research and development in the field. Researchers typically do not have access to streaming ultrasound images, the raw radiofrequency (RF) data or critical image formation parameters, and also cannot control image acquisition remotely. Communication standards and open interfaces will be required for clinical deployment of IGT systems in future and would also ease research and development at the present time. Just as open-source libraries have accelerated software development in image-guided therapy, open communication standards should accelerate the development of practical IGT system implementations that can be used clinically.

### **1.8.0.18 Validation**

Finally, standards in validation methodology are currently required so that assessments evaluating methods presented by different groups can be compared [167]. Most methods developed for image-guided cardiac are validated in the laboratory or in animal models at this point in time, but as their development continues more systems will reach the stage where human testing and clinical trials will become necessary.

## **1.9 Conclusions**

Minimally-invasive cardiac procedures performed on the beating heart have the potential to dramatically improve upon standard clinical practice in cardiothoracic surgery but present unique challenges, not least of which is the problem of surgical guidance in the absence of direct vision. Development of image-guided therapy systems to address this clinical application is ongoing, and can benefit greatly from the real-time imaging provided by echocardiography. Intraoperative ultrasound has been used to accomplish many critical tasks required by IGT systems, including registra-

tion between preoperative images and the intraoperative patient, image-based tool tracking, and cardiac and respiratory motion compensation. In addition, echocardiography is extremely useful for visualization within surgical guidance systems by providing real-time intraoperative imaging of surgical tools and targets, especially when fused with complementary image and/or geometrical data derived from MR, CT or X-ray data. Clinical adoption of image-guided therapy systems that employ echocardiography is growing, and is expected to increase in the future alongside developments in ultrasonic imaging, image processing, and display technologies.

## 1.10 Thesis Outline

The objective of this thesis is to present and validate a real-time 4D ultrasound reconstruction algorithm and its practical implementation, intended for intraoperative imaging of the beating heart during image-guided intracardiac therapy. Fusing high-quality preoperative images to intraoperative images acquired using “OR-friendly” modalities is arguably the best approach to image-guided therapy, and gated free-hand 4D ultrasound reconstruction is an ideal intraoperative imaging technique in many respects, providing flexible and customizable volumetric imaging with a wide field of view and high spatial resolution without radiation dose, using widely-available equipment while remaining safe, inexpensive and compatible with standard OR equipment. This thesis examines the use of real-time reconstruction to ameliorate freehand 4D US reconstruction’s typically challenging acquisition procedure, with the goal of improving reconstruction quality to make the images better-suited for subsequent registration to preoperative images and/or direct visualization.

Chapter 1 provided a literature review of the use of ultrasound in minimally-invasive cardiac therapy, focusing particularly on its use within image-guided interventional systems. This chapter provides the necessary background knowledge to appreciate the motivation for a real-time 4D US reconstruction intraoperative imaging system and to understand its applications as it fits within the broader context of image-guided interventions performed on the beating heart.

The remainder of this thesis proceeds as follows. Chapter 2 begins with a more



focused review of 3D and 4D ultrasound reconstruction, and goes on to describe the real-time 4D ultrasound reconstruction system developed here. Chapter 3 details the system's evaluation, including human factors experiments that illustrate the tangible benefits of providing incremental visualization during 4D ultrasound reconstruction, assessments of reconstruction accuracy (including a study performed using a dynamic phantom and a second performed intraoperatively), and representative imaging performed on porcine subjects and on patients. Chapter 4 completes the thesis by summarizing its main points, outlining potential improvements and avenues for future research, and describing additional applications for real-time ultrasound reconstruction in image-guided therapy. Finally, Appendix A is intended for future users of the ultrasound reconstruction software and for programmers wishing to extend it, and provides details of the software classes and of two implementations provided within the AtamaiViewer and 3D Slicer applications.

# Chapter 2

## Real-Time 4D Ultrasound Reconstruction

### 2.1 Introduction

#### 2.1.1 4D ultrasound in image-guided cardiac therapy

Whether used for visualization or for registration to preoperative images, intraoperative imaging is an integral part of most image-guided therapy (IGT) systems. In particular, effective intraoperative imaging is essential when facilitating minimally-invasive therapies performed within the beating heart, as they present a unique combination of challenges. Any surgical guidance system must replace direct vision while the surgeon targets highly-deformable structures moving quickly within an environment that is itself in motion. Further complicating matters for direct access procedures is the non-linear deformation that the heart undergoes compared to preoperative images. However, the dynamic nature of beating heart surgery makes the use of intraoperative imaging modalities critical while simultaneously complicating their use. The selected intraoperative imaging approach should provide good delineation of intracardiac structures in three-dimensions and throughout the cardiac cycle, while remaining unobtrusive within the operating room (OR).

Of all the available imaging modalities, ultrasound shows particular promise for

intraoperative cardiac imaging, as it provides real-time imaging with good soft-tissue contrast without causing a major intrusion into the OR or subjecting the patient and clinical staff to radiation dose. During surgical guidance, any difficulties in ultrasound image interpretation caused by artifacts or the inherent low quality of ultrasound images can often be ameliorated via fusion with preoperative images or models as described in Chapter 1.

Traditional ultrasound scanners image in two dimensions at a high frame rate, but effective IGT for beating heart surgery would benefit greatly from volumetric intraoperative ultrasound, since it is difficult to manipulate 2D probes to keep the imaging plane aligned with anatomical structures and/or surgical tools, impeding their use for surgical guidance. In addition, access to 2D image information alone is limiting for both visualization and non-rigid image registration, especially considering the dynamic 3D intracardiac environment.

4D echocardiography generates a time series of 3D US volumes representing the beating heart over the cardiac cycle, and can be accomplished using either RT3D ultrasound or gated ultrasound reconstruction. Originally developed in the early 1990s [185,186], RT3D echocardiography, particularly RT3D transesophageal echocardiography, is currently receiving substantial attention but has several disadvantages that should inspire caution when proposing to rely on it alone for intraoperative imaging. First, the finite speed of sound in tissue results in a compromised field of view, spatial resolution and frame rate. Commercial systems can extend the field of view by stitching together multiple 3D volumes acquired over several cardiac cycles, however this approach still fundamentally combines several low resolution images. RT3D TEE is still not ubiquitous in cardiac ORs, streaming 3D image data out of proprietary machines is difficult, and maintaining patient safety despite frequent probe overheating means that live 3D imaging is limited to 2-3 minute intervals.

The reconstruction of ultrasound volumes from 2D images uses a tracked 2D US probe to acquire multiple images from a variety of poses, which are subsequently composited together into one or more 3D volumes. Reconstructions of this nature can maintain the wide field of view and relatively high spatial resolution of 2D US imaging while using commonly available 2D US equipment and maintaining the 2D scanning




Fig. 2.1: Overview of ultrasound reconstruction.

technique with which clinicians are familiar. Unlike RT3D ultrasound, ultrasound reconstruction systems do not image natively in three dimensions, but nevertheless provide 4D image information encapsulating the intraoperative heart that can be used for non-rigid registration to preoperative images or for direct visualization.

### 2.1.2 Ultrasound reconstruction

The goal of ultrasound reconstruction is to generate a 3D or 4D ( $3D + t$ ) ultrasound dataset using multiple 2D ultrasound images. Each image is associated with a transformation describing the ultrasound probe's position and orientation (with respect to some coordinate reference frame) during its acquisition, which is used to insert the image into an output volume with the correct pose (Figure 2.1). 4D ultrasound reconstruction can be used to image rhythmically moving structures, such as the beating heart or organs like the abdominal organs that are influenced by respiratory motion. In 4D US reconstruction, a gating scheme is needed to assign specific 2D US images to their corresponding output volumes. Once they have been generated, reconstructed ultrasound volumes can be displayed on orthogonal or arbitrarily-oriented slice planes (a method known as multiplanar reformatting), with volume rendering, or with surface rendering following image segmentation.

The development of 3D ultrasound systems was initially motivated by the ability to perform accurate distance and volume quantification, reduce user variability by

minimizing the need to mentally reconstruct 3D spatial relationships, improve ultrasound image quality through compounding [72], and flexibly adjust the viewing plane regardless of practical constraints on probe positioning. Ultrasound reconstruction systems have been developed for intraoperative imaging during a variety of clinical procedures. For example, reconstructed 3D ultrasound has been used to correct for intraoperative brain shift [187] and to define targets for both robotic liver biopsy [188] and robotic prostate brachytherapy [189]. In addition, augmented reality surgical guidance systems incorporating 3D ultrasound reconstruction have been developed for laparoscopic liver surgery [150] and for needle biopsy [149]. In both of these systems, 3D US volumes imaging structures of interest are reconstructed, volume rendered, and overlaid onto live video data. However, to date reconstructed ultrasound has not yet been integrated into an augmented reality environment for cardiac therapy.

Both Fenster and Downey [28] and Gee *et al.* [190] give a comprehensive overview of research and development in ultrasound reconstruction (Gee *et al.* in particular provide practical advice to the system designer). Multiple 2D US images are typically acquired by moving the US imaging plane in a linear translation, along a fan trajectory, or around a rotational axis, although a completely unstructured motion can also be used. The image plane may be manipulated using a mechanical system, which can be an external device that moves the entire probe or may be built into the transducer itself. Mechanical probes are very precise and the transformation associated with each 2D US image can be easily determined based on the state of the mechanical device itself. Of particular interest are fast rotating mechanical probes that perform the task of ultrasound reconstruction in real-time [191, 192]. A second alternative is to have the clinician maneuver the ultrasound probe manually in a technique referred to as freehand ultrasound reconstruction. This approach is more flexible than the mechanical approach and allows the clinician to manipulate the transducer to optimal 2D US views of the target structures. To associate a transform with each 2D US image, some older systems simply required the user to follow a predefined trajectory, although this is not ideal as any deviation from this path will result in reconstruction errors. Several groups have investigated image-based tracking, which attempts to

determine the spatial relationships between 2D US images based on their intensity information alone, for example using speckle decorrelation, although this has proven to be a difficult task [193]. The most popular freehand reconstruction method is to track the ultrasound probe using a tracking system, allowing the probe’s movement to be captured during the scan. A variety of tracking systems have been used, including articulated arms, acoustic and, more recently, magnetic and optical tracking systems [28].

Both spatial and temporal calibration must be performed when using a tracking system for ultrasound reconstruction, as reviewed in [43]. Since the tracking system localizes the sensor mounted on the ultrasound probe and not the ultrasound fan origin itself, spatial calibration is required to determine the transformation between the ultrasound tracking sensor’s coordinate system and that of the input 2D US images themselves. This allows input image pixels to be correctly localized within the output volume. Spatial calibration is usually performed using a stationary phantom of known construction. First, the phantom is localized relative to the tracking system’s coordinate system, several 2D US images of the phantom are acquired, and features in the images are related to the corresponding points in the phantom’s reference frame. Using several such examples the spatial calibration matrix can then be determined. Many spatial calibration procedures exist, and example studies comparing different methods include [194, 195]. Secondly, temporal calibration is required to match each 2D US image to its corresponding transform matrix from the tracking system. Image and tracking data are captured at different rates and with an approximately constant lag time between them. This latency is the parameter estimated by temporal calibration. In general, temporal calibration is performed by concurrently collecting and timestamping a series of 2D US images and tracking data. Correlating measurements derived from each data stream allows the mean latency between the two timestamp sequences to be calculated. During ultrasound reconstruction, the transformation for each 2D US image can then be pulled from the streaming transform matrices using their timestamping information and this known lag.

Once the 2D US images and transform matrices are acquired, a reconstruction algorithm is typically used to resample the 2D US image pixels onto an output voxel

grid. An exception is Stradx [196, 197], an open-source system that relies solely on the original set of 2D US images and their associated transformation data to perform interactive reslicing and volume rendering. However, most ultrasound reconstruction systems do employ a regular output voxel grid for ease of subsequent image processing. Solberg *et al.* [198] review previously developed reconstruction algorithms, grouping them into three categories. Voxel-based (“backward”) methods iterate over the output voxels and fill them with intensities from one or more input pixels, using either a nearest-neighbor approach or an interpolation function. In contrast, pixel-based (“forward”) methods iterate over the pixels of each of the 2D US images and contribute each of their intensities towards one or more output voxels. Although a nearest-neighbor method can be used here as well, a three-dimensional spherical or ellipsoidal kernel function is often employed to spread each input pixel onto multiple output voxels such that those voxels closest to the input pixel are more greatly influenced than those farther away. A post-reconstruction hole-filling step is typically used to fill any small gaps in the output volumes that were not directly hit by an input pixel. Pixel-based algorithms are often preferable over a voxel-based approach as the latter provide no speed advantages when the 2D US image pose is arbitrary (as in freehand reconstruction), become slower with increasing output volume size, and are not well-suited for real-time implementation. Finally, function-based methods fit a function to the input pixel intensities that is subsequently evaluated at the output voxel grid coordinates. Function-based methods may produce very high-quality reconstructions but are computationally expensive.

Many 4D ultrasound reconstruction systems for imaging the beating heart have been developed since the first systems were created in the 1970s [199], most using ECG-gating to bin together 2D images corresponding to the same timepoint in the cardiac cycle. Representative work includes the early VHS video-based system of Barry *et al.* [200], the use of raw digital ultrasound data to improve spatial and temporal resolution by Berg *et al.* [201] and the 4D imaging system developed by Meairs *et al.* [202] to reconstruct both B-mode and Doppler cardiovascular images. More recently, Knackstedt *et al.* [42] used a stepper motor to rotate an untracked 2D intracardiac echocardiography transducer for intraoperative visualization during

electrophysiological interventions.

There are many potential error sources in ECG-gated 4D ultrasound reconstruction that combine to result in localization errors and stitch artifacts. These error sources can be itemized as follows:

- **2D ultrasound imaging:** finite 2D US spatial resolution; frame grabbing reduces frame rate and image quality compared to the use of raw US data; discrepancy between actual speed of sound in tissue and assumed speed; artifacts hamper image interpretation
- **Tracking:** intrinsic tracking bias and imprecision; ferromagnetic objects may disrupt the magnetic field in magnetic tracking, thus introducing measurement errors
- **ECG-gating:** imperfect R-wave detection; heart rate changes may bin together images acquired during different physiological cardiac phases; ECG signal does not necessarily represent the heart's actual state
- **Synchronization:** temporal calibration error; timestamping error
- **Gross heart movement:** patient movement; patient respiration and/or unequal lung volume between subsequent breath-holds; probe pressure effects
- **Reconstruction:** spatial calibration error; finite output volume resolution; blurring from interpolation onto the output voxel grid

As described in Chapter 3, it is therefore important to perform validation studies for any ultrasound reconstruction system to determine its accuracy.

A common complaint limiting the clinical integration of ultrasound reconstruction is the need for offline reconstruction following 2D image acquisition. While current offline reconstruction techniques can be extremely fast (ex. [203]), image acquisition errors are not detected until after reconstruction, at which time the acquisition must be repeated. Common acquisition errors include a failure to capture all of the structures of interest, or moving the probe too quickly, causing gaps in the output



volume(s). Real-time ultrasound reconstruction (also known as incremental or interactive reconstruction) solves this problem by inserting 2D US images into the output volume(s) as they are acquired, thus providing the dual advantages of (1) being able to visualize the reconstructed volume(s) during data collection and hence provide the clinician with valuable feedback to prevent or correct acquisition errors, and (2) eliminating the post-acquisition reconstruction time altogether.

Past development of real-time ultrasound reconstruction systems has focused on 3D imaging of stationary organs. This includes the Stradx system described above. In 1992, Ohbuchi *et al.* presented one of the first incremental 3D ultrasound reconstruction systems, displaying results using an asymmetric 3D Gaussian kernel and volume rendering [204]. To simplify the real-time reconstruction task, Edwards *et al.* used a simple pixel-replacement distribution scheme [205], while Welch *et al.* inserted the input images in blocks [206]. Finally, our group has presented a real-time 3D US reconstruction system for intraoperative brain shift compensation [207]. This system implemented four kernel-based resampling schemes within a multithreaded software architecture, achieving interactive reconstruction rates while displaying the output on three adjustable orthogonal slice planes (“orthoplanes”).

### 2.1.3 Motivation

Although image acquisition is more technically challenging for 4D ultrasound reconstruction than for 3D reconstruction, potentially making real-time feedback even more valuable, work on real-time 4D US reconstruction is practically non-existent. Although Stradx can perform image-based gating [99], their method is a post-acquisition technique. Developing a real-time freehand 4D ultrasound reconstruction system should optimize image acquisition and therefore improve the quality of the reconstructed volumes. I propose that such improvements would facilitate subsequent image visualization and registration to preoperative images, further increasing the approach’s applicability for intraoperative imaging of the beating heart during image-guided therapy.

This chapter presents a real-time 4D ultrasound reconstruction system for intra-

operative imaging of the beating heart during minimally-invasive intracardiac interventions. This system is a hybrid freehand-mechanical system that uses the patient's ECG signal for gating. The reconstruction software provides real-time visualization using standard computer hardware and is integrated within an augmented reality environment designed for intracardiac therapy delivery [84]. Clinicians using the real-time visualization provided by the incremental 4D ultrasound reconstruction system should be able to effectively prevent and correct acquisition errors, and hence reconstruct improved 4D image datasets.

Compared to RT3D ultrasound, reconstructed 4D ultrasound has a wider field of view and higher spatial resolution. This is especially true considering that the intraoperative setting considered here prohibits the use of transthoracic echocardiography and requires transesophageal echocardiography instead. In addition, ultrasound reconstruction uses standard 2D ultrasound probes that are more widespread clinically at this point in time. Finally, it is often difficult, if not impossible, to extract streaming RT3D images due to manufacturer restrictions, while the 2D ultrasound images input into ultrasound reconstruction algorithms can be easily accessed with frame grabbing. Compared to untracked mechanical reconstruction systems using motorized probes, freehand ultrasound reconstruction offers enhanced flexibility with respect to the ability to optimize each of the 2D ultrasound views making up the output volumes. Also, freehand ultrasound reconstruction can generate output data with irregular shapes that is customized to the individual patient, in order to ensure that all structures of interest are imaged.

The reconstructed 4D ultrasound datasets generated in this thesis are intended for both image-to-patient registration and intraoperative visualization within our group's augmented reality environment for minimally-invasive interventions performed on the beating heart. Reconstructed 4D ultrasound data offers a global 3D view of the intraoperative heart while existing natively within the AR environment's coordinate system. Therefore, it forms an intermediate of sorts between registered preoperative image data (which provides 3D context but may suffer from residual registration errors) and real-time 2D or RT3D echocardiography (which does not require registration into the augmented reality environment but provides only a small amount of

intraoperative image data at any point in time). The large field of view and high spatial resolution is an especially useful combination for image-to-patient registration that is not provided by other ultrasound image approaches. However, feature-based or intensity-based image-to-patient registration may instead be based on a sparse dataset of tracked 2D ultrasound images. Although a direct comparison between this approach and to registration based on reconstructed ultrasound has not yet been performed for cardiac applications, the denser image information provided by ultrasound reconstruction may improve registration accuracy, especially if non-rigid registration is attempted. When applied to intraoperative visualization, 4D ultrasound reconstruction gives a gross 3D view of the heart in its intraoperative state. As with registered preoperative images, this can be used to provide a 3D context to the virtual scene, but has the additional advantage of capturing the heart in its intraoperative state. RT3D TEE remains useful for visualization during final detailed manipulation of surgical targets, although RT3D TEE has a tendency to overheat and therefore cannot be used for extended imaging to maintain patient safety.

The remainder of this chapter focuses on the real-time implementation of the 4D ultrasound reconstruction algorithm, while Chapter 3 discusses its clinical integration and validation.

## **2.2 Real-time 4D Ultrasound Reconstruction System**

### **2.2.1 System overview**

The core of the real-time 4D ultrasound reconstruction algorithm is implemented in C++ using the Visualization Toolkit (VTK) [208], while the user interface and interactive visualization tasks are performed by the python-based AtamaiViewer software package ([www.atamai.com](http://www.atamai.com)). The AtamaiViewer supports common IGT tasks including tool tracking, registration and visualization of both images and models, and is the platform with which our group’s augmented reality intracardiac IGT system [84]

was developed. I have also provided the source code to the research community for a command-line implementation [209] that can communicate with any application supporting the OpenIGTLink network transfer protocol [210]. This includes 3D Slicer, a widely-used open-source application for medical image processing and image-guided therapy ([www.slicer.org](http://www.slicer.org)). Please see Appendix A for more details.

The reconstruction algorithm described below supports tracked multiplanar ultrasound probes, which contain a small motor that rotates the imaging plane from 0-180° around the central axis parallel to the imaging plane. This conveys a significant advantage to the system: unlike the majority of mechanical US reconstruction systems, which use an untracked transducer, the system can provide the reliability and precision of motorized manipulation in combination with the flexibility of the freehand approach. With the system, clinicians can manipulate the probe manually using a translational or fan-like motion, mechanically with the rotational approach, or using any combination of the above. Even when performing a strictly rotational acquisition, tracking the US transducer allows the system to accommodate for slight probe movements.

The system employs a pixel-based (“forward”) reconstruction algorithm, iterating through the pixels of each input 2D US image and transforming them into output voxel coordinates before blending their intensities into the output volume. The 3D output volumes are axis-aligned with the AR environment’s world coordinate system (which is that of the tracking system or of the tracked reference sensor if one is used), and so the transformation from discrete 2D input pixel coordinates  $p_{in} = \begin{bmatrix} x_{in} & y_{in} & 0 & 1 \end{bmatrix}^T$  to discrete 3D output voxel coordinates  $v_{out} = \begin{bmatrix} x_{out} & y_{out} & z_{out} & 1 \end{bmatrix}^T$  is given by:

$$v_{out} = S_{out}T_{US}C_{US}R_{\theta}S_{in}p_{in} \quad (2.1)$$

Here,  $S_{in}$  converts from input pixel coordinates to input image coordinates. In the implementation used here, the 2D US images lie on the  $\mathbf{xy}$ -plane, centred around the  $\mathbf{y}$ -axis and with the bottom edge lying on the  $\mathbf{x}$ -axis. If using a multiplanar probe,  $R_{\theta}$  applies the rotation for the current image plane rotation  $\theta$  ( $R_{\theta} = I$  otherwise).  $C_{US}$  denotes the spatial calibration matrix that gives the constant transformation from input image coordinates to the tracking sensor affixed to the US transducer. Tracking

the ultrasound probe gives the rigid transformation  $T_{US}$  into world coordinates, and finally  $S_{out}$  converts from world coordinates into output voxel coordinates.

In particular, if the spacings of the input image and output volumes are given by  $s_{in}$  and  $s_{out}$ , respectively, and their origins are  $o_{in}$  and  $o_{out}$ , respectively, then  $R_\theta$ ,  $S_{in}$  and  $S_{out}$  are defined as followed:

$$R_\theta = \begin{bmatrix} \cos(\theta) & 0 & \sin(\theta) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta) & 0 & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2.2)$$

$$S_{in} = \begin{bmatrix} s_{in_x} & 0 & 0 & o_{in_x} \\ 0 & s_{in_y} & 0 & o_{in_y} \\ 0 & 0 & s_{in_z} & o_{in_z} \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2.3)$$

$$S_{out} = \begin{bmatrix} 1/s_{out_x} & 0 & 0 & (-1 \times o_{out_x})/s_{out_x} \\ 0 & 1/s_{out_y} & 0 & (-1 \times o_{out_y})/s_{out_y} \\ 0 & 0 & 1/s_{out_z} & (-1 \times o_{out_z})/s_{out_z} \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2.4)$$

This real-time 4D ultrasound reconstruction system extends the 3D system of Gobbi *et al.* [207,211] and performs the following tasks to reconstruct a time series of 3D ultrasound volumes while simultaneously visualizing the results (see Figure 2.2):

- **Data Acquisition:** 2D images from the ultrasound scanner, transform matrices from the tracking system, and ECG information are streamed into the system.
- **Slice Selection:** Not all of the 2D US images will correspond to a timepoint in the cardiac cycle that is represented by an output volume. Retrospective or prospective ECG-gating is used to pick a 2D US image that does, and then determine its associated transform matrix.

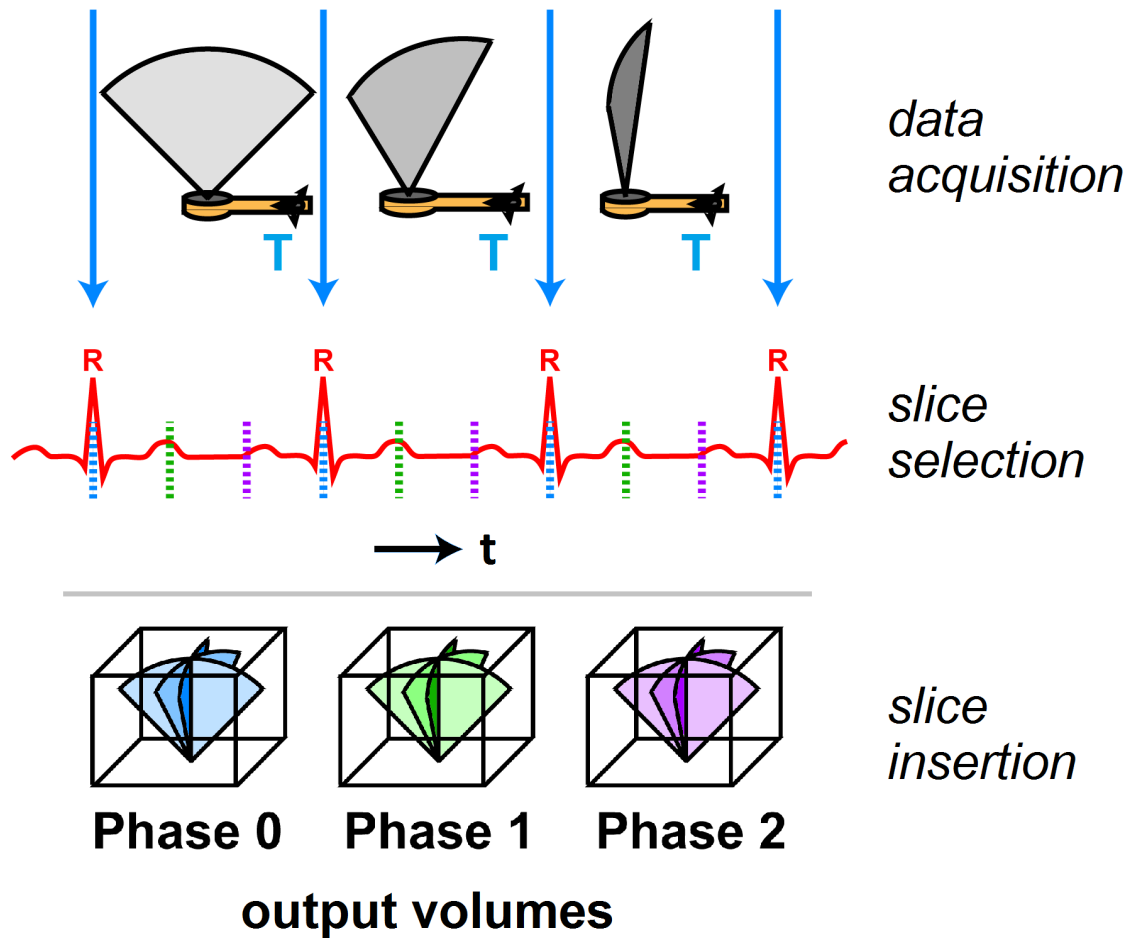


Fig. 2.2: System overview for the real-time 4D ultrasound reconstruction system. The system streams 2D ultrasound images, transform matrices and ECG data, selects 2D US images corresponding to an output volume and inserts them into the output volumes, all while providing real-time visualization of the results for a representative output volume.

- **Slice Insertion:** The 2D input pixels are transformed into output voxel coordinates, and are injected into the appropriate 3D output volume using a three-dimensional splatting kernel.
- **Real-time Visualization:** A representative 3D output volume is visualized on three interactively-controlled orthoplanes, which are updated many times per second.

Since a pixel-based approach is used, upon the completion of the reconstruction a hole filling step is performed to interpolate small gaps in the output volumes. The remainder of this section further details the steps listed above. Because a multi-threaded software architecture is employed to perform these tasks concurrently, the related temporal synchronization and shared data protection considerations are also discussed. Finally, I describe the cardiac imaging equipment and the spatial and temporal calibration procedures, and discuss the performance considerations for real-time 4D ultrasound reconstruction.

### 2.2.2 Data acquisition and synchronization

Three separate threads handle streaming 2D US images, transform matrices from the tracking system and ECG information. All three forms of input data are associated with timestamps to synchronize the three threads. First, a frame grabber is used to capture  $640 \times 480$  2D US images at 30 frames per second (fps). Each frame is associated with a timestamp from the system clock, and both are stored within memory in a circular buffer holding 1000 frames. Both intensity and colour frames can be acquired, enabling reconstruction of both B-mode and colour Doppler ultrasound, although B-mode ultrasound is used exclusively here. To produce high-quality reconstructions, it is important for the clinician to perfect the ultrasound gain, contrast, depth and focus settings to optimize the individual 2D US images. Second, both transform matrices and their timestamps are received from the tracking system over a serial connection. The tracking rate of our group's magnetic tracking system depends on the number of tracked tools and is either 45 Hz (five or fewer tracked

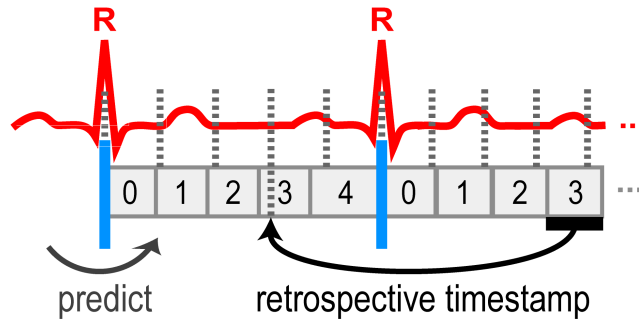


Fig. 2.3: Prospective and retrospective gating. (a) At any point in time, prospective gating predicts the current cardiac cycle phase (boxed integers), which will usually differ from the true cardiac cycle phase (dashed lines); (b) Retrospective gating can accurately determine the correct timepoints at which phases began in the previous cardiac cycle.

coils) or 20 Hz (greater than five tracked coils). Since the tracking timestamps are not generated using hardware interrupts, a software implementation of a phase-locked loop (PLL) circuit is used to reduce their temporal jitter [211]. The transform matrices and their timestamps are stored within a second circular buffer that typically holds 1000 elements, but may hold 54000 (approximately 20 minutes of tracking data at 45 Hz) when post-acquisition analysis of tracking data is required. For each transform matrix, the tracking system also provides a flag if a tracked tool is outside of the system’s field of view or cannot be localized correctly. Finally, every 5 ms the patient’s ECG signal is sampled and saved in memory along with its corresponding system timestamp.

The R-waves of the ECG signal are automatically detected based on a user-defined threshold and are used to delineate each cardiac cycle. Each cardiac cycle can be evenly split into  $N$  phases, of which a subset  $P \subseteq \{0, 1, \dots, N - 1\}$  are designated as being of interest. The reconstructed 4D ultrasound dataset will contain  $N_P = |P|$  output volumes, each corresponding to the beginning of one of these phases.

To aid in slice selection, the ECG thread maintains estimates of the current heart rate and cardiac phase. If an R-wave is detected at time  $t_{R_i}$  (in seconds), the heart rate estimate (in beats per minute) can be updated using the timestamp  $t_{R_{i-1}}$  of the



previously-detected R-wave:

$$HR = 60 / (t_{R_i} - t_{R_{i-1}}) \quad (2.5)$$

During the subsequent cardiac cycle, for each iteration of the ECG thread's main loop occurring at system time  $t_C$ , prospective gating can estimate the current cardiac phase as follows:

$$p = \left\lfloor \left( \frac{t_C - t_{R_i}}{t_{R_i} - t_{R_{i-1}}} \times N \right) \bmod N \right\rfloor \quad (2.6)$$

However, any changes in the patient's heart rate will render estimates based on the previous cardiac cycle inaccurate (Figure 2.3). Since the exact duration of the current cardiac cycle is unknown, an alternative is to use retrospective gating. For any phase  $p$  ( $0 \leq p < N$ ), the timestamp  $t_r$  ( $t_{R_{i-1}} \leq t_r < t_{R_i}$ ) at which it began in the previous cardiac cycle can be calculated exactly:

$$t_r = t_{R_{i-1}} + \left( \frac{t_{R_i} - t_{R_{i-1}}}{N} \times p \right) \quad (2.7)$$

Although the three data acquisition threads run at different rates and may be temporally offset with respect to each other, in the slice selection step each 2D US image must be grouped with its corresponding transform matrix and cardiac phase information. Temporal synchronization is accomplished using the timestamps associated with each piece of data. It is assumed that the lag in receiving the 2D US images with respect to the ECG signal is negligible, and since both of these streams are timestamped using the same system clock, they are already synchronized. However, there is an approximately constant time delay  $t_{lag}$  between the 2D US image stream and the tracking stream, which is estimated using the temporal calibration procedure described below. Since multiple threads access the same underlying data structures, mutex locks are used whenever reading from or writing to the 2D US image and tracking buffers and when updating the ECG data.

### 2.2.3 Slice selection

For each cardiac cycle, one 2D US image is inserted into each of the  $N_P$  output volumes. A fourth reconstruction thread monitors the current status of the ECG

thread and uses retrospective or prospective gating to find the 2D US images whose timestamps best correspond to the beginning of the phases designated in  $P$ . As described above, retrospective gating is more accurate and so is generally preferred.

Changes in heart rate alter the relative durations of systole and diastole, so the user can elect to reduce motion artifacts by enforcing a limited range of acceptable heart rates. Any 2D US images acquired when the heart rate differs too much from an expected value ( $HR_{exp}$ ) are rejected, as defined by  $\%dev_{HR}$ , the maximum allowed percentage deviation from the expected heart rate. To do this the actual cardiac cycle duration for each 2D US image must be known, and so the selected images are inserted into an output volume one cardiac cycle after they were acquired when this option is turned on. Prior to the reconstruction,  $HR_{exp}$  is measured by calculating the mean heart rate over a 5 s interval while simultaneously ensuring that no measured heart rate is beyond  $HR_{exp} \pm \%dev_{HR}$  by repeating the measurement until the condition is true.

To stagger the computationally demanding task of inserting  $N_P$  2D US images per cardiac cycle, one image is inserted into its output volume whenever the ECG thread predicts the beginning of a phase of interest using prospective gating (equation (2.6)). If the known duration of the previous cardiac cycle is within range, then retrospective gating can be used to find the 2D US image to be inserted and its corresponding transform matrix:

1. Use equation (2.7) to calculate  $t_r$ , the timestamp for the beginning of phase  $p$  in the previous cardiac cycle.
2. Use a binary search to find the buffered 2D US image whose timestamp  $t_{US}$  is closest to  $t_r$ .
3. Calculate the tracking timestamp:  $t_t = t_{US} - t_{lag}$ .
4. Use a binary search to find the two buffered transform matrices whose timestamps straddle  $t_t$ , and interpolate the transform matrix  $T_{US}$  using linear interpolation for the position and spherical linear interpolation (SLERP) [212] for the orientation.

If the user elects to use prospective, rather than retrospective, gating, the procedure differs only when selecting the 2D US images. Instead of following steps 1) and 2) above, whenever the ECG thread predicts the beginning of a phase of interest the system selects the 2D US image whose timestamp is closest to that of the ECG thread’s most recent measurement. This will be either the most recently acquired 2D US image or the upcoming image. Even when using prospective gating, if the user elects to check the heart rate against the expected value then each selected 2D US image is saved in memory for insertion during the next cardiac cycle.

#### 2.2.4 Slice insertion

Multiple threads (equal in number to the number of available processors) are used to insert each selected 2D US image into its output volume. This is possible because the slice insertion task operates on a pixel-by-pixel basis, and can therefore be easily parallelized by assigning a horizontal block to each thread (i.e. if there are two processors, one thread operates on the top half of the US image and the other operates on the bottom half). Each thread transforms those input pixels falling within the US fan into output voxel coordinates using equation (2.1). The extraction of the US fan from each of the images is performed using the known fan origin position, the left and right fan angles and the current fan depth. These values are determined by the current depth setting (in cm) and whether the fan is flipped up or down, both of which are automatically detected at the beginning of the reconstruction and are assumed to remain constant throughout. If a multiplanar probe is used,  $R_\theta$  is determined for each 2D US image by interrogating the pixels in the image that display the current image plane rotation. The 2D US image is rejected if the difference between its image plane rotation and that identified for the previous image is greater than  $20^\circ$ , as this may indicate a problem with the rotation detection algorithm.

Each input pixel will very rarely map exactly onto an output voxel, and so a mechanism is needed to spread the pixel’s intensity onto one or more neighboring voxels. In addition, multiple input pixels from different 2D US images will sometimes intersect with the same output voxel, in which case it is often desirable to average

their intensities in some way. This is achieved using the method described by [207], where each input pixel is splatted into the output volume using a discrete kernel. This reconstruction method is similar to the kernel-based algorithms presented in [200,202,204], but uses a small, non-configurable kernel and a simpler linear weighting scheme to reduce the time required for real-time slice insertion.

Gobbi *et al.* [207] describe two interpolation techniques, pixel nearest neighbor (PNN) and pixel trilinear (PTL) interpolation, and two blending techniques, alpha blending and compounding. The software implementation proposed here implements all four combinations, but since PTL interpolation with compounding was used exclusively for all experiments this method is described in full below while [207] is referred to for the others. When using PTL interpolation with compounding, at the end of the reconstruction each output voxel will contain the weighted average of the intensities from all the input pixels that intersected with it. Although it is the most computationally expensive method, PTL interpolation with compounding can still execute sufficiently quickly for use in real-time reconstruction, and of the four alternatives produces the most accurate and least noisy output volumes [211].

Adding each input pixel incrementally to an output volume involves an interpolation step followed by a blending step. First, the interpolation step determines the relative influence that the input pixel’s intensity  $I_{\text{pixel}}$  has on the eight “hit” output voxels ( $v_k$ ) that surround the transformed input pixel coordinate  $v_{\text{out}}$ . This is achieved by applying a  $2 \times 2 \times 2$  kernel containing the trilinear interpolation coefficients  $b_k$  that relate the distance between  $v_{\text{out}}$  and each  $v_k$  ( $\sum b_k = 1$ ).

In the compounding step, the eight output voxel intensities  $I_{k\text{voxel}}$  are updated: for each of the eight output voxels, the product  $b_k I_{\text{pixel}}$  is combined with the similarly-weighted intensities of any input pixels that previously intersected with the output voxel. For each output volume, an accumulation buffer with the same dimensions is stored and holds the accumulated kernel weights  $a_k$ . To blend a new input pixel into the output volume, the eight output voxel intensities and the eight accumulated kernel weights are updated as follows:

$$I_{k\text{voxel}} := \frac{b_k I_{\text{pixel}}}{b_k + a_k} + \frac{a_k I_{k\text{voxel}}}{b_k + a_k} \quad (2.8)$$

$$a_k := a_k + b_k \quad (2.9)$$

In (2.8), the left operand represents the contribution of the new input pixel, while the right represents the combined contribution of all input pixels that previously intersected with the output voxel. The first time an output voxel is hit, (2.8) reduces to  $I_{k\text{voxel}} := I_{\text{pixel}}$  because each  $a_k$  is initialized to zero.

In addition to the intensity information stored at each output voxel, a second scalar (which does not play a part in the reconstruction algorithm) representing whether the output voxel has been intersected by a 2D US image is also stored. Since the output voxel intensities are initialized to zero, after performing the reconstruction it allows output voxels with zero intensity that were never hit by a 2D US image to be distinguished from those that genuinely represent very hypoechoic ultrasound data. In any subsequent image processing of the reconstructed 4D US data, care must be taken not to process this additional scalar as if it represented actual ultrasound data.

To reconstruct a dense 4D US dataset, the clinician must move the ultrasound transducer by a small increment approximately once per cardiac cycle. In this way, a 2D US image is inserted for each phase of interest before the probe is moved to the next location. Clearly, this is a more technically challenging acquisition compared to that used for 3D US reconstruction, where the clinician can move the probe in one smooth motion without introducing significant gaps into the reconstructed volume.

After the clinician has reconstructed a 4D dataset, a hole-filling step is required to fill in any small holes in the output volumes. Since a  $2 \times 2 \times 2$  kernel is used, a gap will occur whenever the distance between 2D US images is more than two output voxel widths (approximately 1.3 mm). Small holes may occur whether using a freehand or rotational acquisition. For example, if performing a rotational acquisition at  $1^\circ$  increments with a US depth setting of 12 cm, two voxel widths measures 1.39 mm while the distance between the most lateral pixels of consecutive US fans is 1.48 mm. The hole-filling procedure described in [211] is applied to each output volume using as many threads as there are processors, which is very fast. This step is applied only to those empty output voxels whose  $3 \times 3 \times 3$  or  $5 \times 5 \times 5$  neighborhoods are already more than half-filled, preventing large gaps from being filled based on poorly-supporting



Fig. 2.4: Real-time 4D ultrasound reconstruction in progress. A monitor positioned close to the clinician visualizes the AtamaiViewer output, which shows the partially reconstructed output volumes being filled in in as the clinician manipulates the ultrasound probe.

data so that they can be subsequently identified. The empty voxels are isolated using the second scalar component described above. If an empty voxel’s  $3 \times 3 \times 3$  neighborhood is more than half-filled, then its intensity will be the average of those filled neighbors. Otherwise, if an empty voxel’s  $5 \times 5 \times 5$  neighborhood is more than half-filled, it is similarly filled in by averaging the intensities of its filled neighbors, but those intensities coming from the voxel’s  $3 \times 3 \times 3$  neighborhood are weighted four times more heavily.

### 2.2.5 Real-time visualization within an AR environment

Performing 4D ultrasound reconstruction in real-time allows an interactive visualization of the incremental results to be displayed to the clinician during the acquisition. Figure 2.4 shows an example real-time 4D US reconstruction in progress. The AtamaiViewer application thread texture maps the current contents of a representative output volume onto three interactive orthoplanes. Note that the visualization will show 2D US images being inserted one heart beat “late” whenever retrospective gat-

ing is used or whenever prospective gating is used with an expected heart rate. Since the goal is to enable the clinician to monitor the reconstruction's progress and correct acquisition errors, displaying the results on orthoplanes allows for a more detailed view of the individual output voxels than would be possible using other visualization techniques, such as volume rendering. The operator of the reconstruction software typically views the reconstruction results as well, and brings any potential problems to the clinician's attention. During the reconstruction, the operator can smoothly slice through the visualized output volume, adjust the camera position (i.e. pan and zoom), and toggle the display between the different 3D volumes in the 4D dataset. Depending on the user's mouse movement and the current state of the tracked tools, the AtamaiViewer's rendering mechanism renders the scene at a variable rate which is generally greater than or equal to the 30 fps required for smooth motion. The AtamaiViewer thread attempts to monitor the tracked tools every 20 ms and triggers a render whenever a tool either moves compared to its previous measurement or is added, goes out of bounds, or cannot be tracked. The AtamaiViewer also triggers a render whenever the mouse is in motion.

Since a tracked, spatially-calibrated ultrasound transducer is used to perform the 4D US reconstruction, it is possible to reconstruct and display the 4D US dataset within an augmented reality environment (Figure 2.5). The ultrasound volumes are automatically in the same coordinate system as the 2D ultrasound probe and any other tracked tools, while preoperative images and/or derived surface models can be integrated using a registration procedure (such as that described in [84]). After reconstruction and hole-filling, the output volumes can be displayed on orthoplanes as before, or alternatively the AtamaiViewer can perform volume rendering. The result is somewhat similar to the AR systems developed for laparoscopic liver surgery [150] and needle biopsy [149] described above. However, this cardiac-specific system offers the advantages of 1) being able to perform the reconstruction in real-time; and 2) reconstructing a 4D dataset rather than a single 3D volume.




Fig. 2.5: Example visualization of a reconstructed 4D ultrasound dataset registered to a dynamic cardiac surface model, using a beating heart phantom (The Chamberlain Group, Great Barrington, MA, USA). Within the augmented reality environment, the cardiac surface model beats in synchrony with the patient's heart rate and the user can toggle the output volume from within the 4D dataset that is visualized.



## 2.2.6 System implementation

### 2.2.6.1 Hardware

I developed an ultrasound reconstruction module within the AtamaiViewer that can perform both real-time 3D and 4D ultrasound reconstruction. The ultrasound probe can be tracked using a variety of optical and magnetic tracking systems. Alternatively, a tracking simulator can be used for testing. Prior to the reconstruction, the user specifies the output volumes' dimensions and spacing, while their origin is dictated by a 2D US "key frame" that should be located approximately halfway along the clinician's planned trajectory. In the "centring" procedure that determines the output volumes' origin, the corners of a box with the specified output dimensions is aligned with the key frame. The bounding box in the world coordinate system is then used to initialize the output volumes, typically making the dimensions of the output volumes larger than those specified by the user. During the reconstruction, the user can stop and restart the acquisition multiple times. Additionally, the user can choose to save the selected 2D US images and/or their timestamps to a file (this is also performed in real-time, and so caution must be taken when saving the 2D US images as the additional processing required can cause the software to drop 2D US images during the reconstruction). Following the reconstruction, the software saves the output volumes, the contents of the transform matrix buffers for all tracked tools, a summary file of the reconstruction parameters, and a file keeping track of the acquisition time.

For patient safety, the equipment used for laboratory and animal work was kept separate from a second set used for humans. For phantom and porcine studies, a 32-bit workstation equipped with two dual-core Xeon CPUs (3.2 GHz), 2 GB RAM and a Matrox Meteor frame grabber (Montreal, QC, Canada) was used. Human reconstructions were performed using a 32-bit workstation with two quad-core Xeon CPUs (3.2 GHz), 4 GB RAM and a Matrox Morphis frame grabber. Ultrasound imaging was provided by a Philips SONOS 7500 ultrasound machine (Andover, MA, USA) and a 4-7 MHz adult multiplanar 2D TEE probe (M/N:T6210, Philips). The TEE transducer's imaging plane can be rotated at 1° increments using button controls

on the ultrasound machine. I also experimented with transthoracic imaging using a Philips X4 2-4 MHz RT3D probe operating exclusively in 2D mode. Since magnetic tracking systems must be used with transesophageal probes as there is no line of sight to an optical camera, the Aurora<sup>®</sup> magnetic tracking system (Northern Digital Inc., Waterloo, ON, Canada) was used in this work. For phantom and porcine work a 6-degrees of freedom (DOF) sensor was fixed onto the TEE probe's exterior (Northern Digital Inc.), while for human imaging a second TEE transducer with the magnetic tracking sensor integrated directly into the beryllium copper body of the probe was employed [181]. The integrated sensor's tracking accuracy is somewhat reduced compared to an externally mounted sensor, but its use was deemed necessary for patient safety. The TTE probe was also tracked using an externally mounted 6-DOF sensor (Traxtal, Toronto, ON, Canada). All tracking was performed relative to a 6-DOF reference sensor (phantom/humans: Northern Digital Inc.; porcine: Traxtal). Standard leads were used to acquire the ECG signal, which was imported into the workstation over a parallel (phantom/porcine) or USB (humans) port. In the first configuration, the ECG signal was sent to hardware that performed signal thresholding with an adjustable threshold followed by analog-to-digital conversion, outputting a 5 V pulse at the computer's parallel port at each R-wave. These pulses were detected by the ECG-gating software to delineate between subsequent cardiac cycles. For humans, the ECG signal was passed through an analog-to-digital converter and sent directly to the computer over USB, where the ECG thread itself performed thresholding using a threshold specified interactively by the user.

### 2.2.6.2 Calibration

Spatial calibration for each ultrasound probe was performed using a Z-bar phantom [213]. As previously described, the goal of spatial calibration is to determine the transformation between points in 2D US image coordinates and their correspondents relative to the tracked sensor mounted on the ultrasound probe. In this procedure, multiple 2D US images are acquired of the Z-bar phantom from a variety of viewpoints. The centroid of each string cross-section is determined in each image and the

ratio between the three centroids of each Z-triplet is used to determine where the central points lie in the phantom’s coordinate system. These points are then transformed into the ultrasound sensor’s coordinate system, first into the tracking system’s coordinate system using divots on the phantom localized using a tracked pointer tool, and second by applying the transform matrices from the tracked ultrasound probe. The final result is a system of equations relating points in the ultrasound sensor’s coordinate system to points in the 2D US images via the spatial calibration matrix  $C_{US}$ , which can then be solved for using least-squares minimization. The accuracy of this spatial calibration procedure has been assessed for tracked 2D US imaging in a previous laboratory study, which found an RMS error for our group’s TEE probe of 2.4 mm [214].

Temporal calibration was achieved using a procedure based on imaging a string in cross-section to determine the average lag between 2D US images and their corresponding transform matrices from the magnetic tracking system. The method described in [211] was followed, where a series of 2D US images and 3D positional measurements are acquired by slowly moving a tracked TEE probe from side-to-side while keeping the imaging plane perpendicular to the string. Data are acquired over approximately 5-10 seconds, being mindful to keep the string in the ultrasound focal zone and using a plastic block to maintain the linearity of the probe’s motion. A custom-developed Matlab script was used to perform the computations for temporal calibration. The centroid of the string’s cross section is determined for the first image and is then automatically tracked for the remainder of the image sequence by using the coordinate from the previous image to seed a thresholded centroid calculation. Applying a principal component analysis (PCA) both to the resulting series of 2D image coordinates and to the 3D coordinates for the ultrasound sensor yields two one-dimensional signals with similar profile. After normalizing the signals to zero mean and unit variance, the tracking signal is shifted in 3 ms increments (minimum -130 ms, maximum 130 ms) to find the lag that minimizes the least-squares difference between the two waveforms. The temporal calibration procedure was repeated six times, each time while tracking a total of three 6-DOF sensors. On average, the ultrasound image data lagged behind the tracking data by approximately 32 ms ( $t_{lag} = 0.032$ ), although

the results of these six trials showed a temporal lag ranging from approximately -20 ms to +90 ms.

We are constrained to magnetic tracking because of our intracardiac application, but the majority of groups developing ultrasound reconstruction techniques use optical tracking systems because of their greater accuracy. An additional advantage of optical tracking is its speed in tool localization, as it is virtually guaranteed that the transform matrices for each tracked object will be determined within the time period associated with the tracking frequency. The resulting constant frame rate in optical tracking systems makes the traditional assumption of a constant lag between the image and tracking data very reasonable. In comparison, magnetic tracking systems have a more variable tracking rate that depends on the position and orientation of each tracked tool, explaining the large range in the temporal calibration results. Although the assumption of a constant lag is not ideal, for magnetic tracking systems, developing a system that integrates a variable lag between the image and tracking data is not trivial. Note that as probe movement must be very slow for 4D US reconstruction, in particular compared to 3D US reconstruction, slight errors in temporal calibration will not manifest as large errors in the reconstructed volumes. Therefore a slightly inaccurate temporal calibration is likely not one of the more significant error sources in the system.

### **2.2.7 Performance considerations**

Practical constraints on available memory and processing speed will impose limitations on any real-time 4D ultrasound reconstruction system, particularly concerning the number of output volumes, their dimensions and those of the 2D US images, and the time difference between reconstructed phases.

The amount of available memory limits the number of output volumes and their size because the output volumes and the accumulation buffers (if compounding is used as implemented here) must be allocated prior to the reconstruction. In the preliminary experiments described in the next chapter, the area over which the clinician would like to manipulate the probe was uncertain, and so large output volumes, often

containing between 70 and 90 million elements each, were used. Therefore, only two output volumes for each 4D dataset were normally reconstructed. In many cases, smaller output volumes can be used to increase their number. Using our laboratory workstation, the largest number of output volumes (up to a maximum of 20) that could be allocated for different output volume dimensions without memory allocation problems was determined using trial and error. The results were as summarized below for varying output volume size:

- $320 \times 240 \times 320$ : maximum 5 (compounding) or 6 (alpha blending)
- $240 \times 180 \times 240$ : maximum 16 (compounding) or 20 (alpha blending)
- $160 \times 120 \times 160$ : maximum 20 (compounding) or 20 (alpha blending)

While the available workstation limited the possible number of output volumes, this could easily be increased by using a machine with more random access memory (RAM) and a 64-bit architecture.

The second potential limitation is processor speed. To prevent dropped frames, the time difference between subsequent phases of interest must be large enough to perform a slice insertion. The speed with which 2D US images can be inserted depends on the processor, the dimensions of the input images, and the type of interpolation and blending functions used. When using our laboratory workstation for real-time 3D US reconstruction, approximately twenty  $640 \times 480$  2D US images can be inserted per second when using pixel trilinear interpolation with compounding. The improvement over the 12 fps reported for  $320 \times 240$  images for the software on which the present system is based [207] is caused by the more advanced processor employed here. For real-time 4D ultrasound reconstruction, our group is typically interested in reconstructing output volumes at evenly spaced intervals, and since memory restricts us to reconstructing few output volumes the time between phases of interest is relatively long and therefore the slice insertion rate is not a limiting factor.

## 2.3 Discussion

### 2.3.1 Clinical potential

As of yet, no image-guided surgical system has overcome all of the challenges associated with minimally-invasive intracardiac therapy performed without cardiopulmonary bypass. Further research and development in this area is encouraged by the potential advantages of such a system, including reduced patient trauma, shorter post-operative recovery times and a decreased incidence of side effects including infection, negative cognitive effects and post-operative stroke. The real-time 4D ultrasound reconstruction system described here has the potential to provide effective yet practical 4D intraoperative imaging as part of our group's larger project of developing augmented reality-based surgical guidance technologies for minimally-invasive interventions performed within the beating heart, including valve replacement, atrial septal defect closure, and endocardial EP measurement and ablation for the treatment of atrial fibrillation.

To the best of my knowledge, the work presented here represents both the first real-time 4D ultrasound reconstruction system and the first integration of reconstructed ultrasound volumes into an AR environment for use during image-guided cardiac therapy. The system maintains freehand ultrasound reconstruction's advantages of flexibility and the ability to optimize 2D US views, but additionally should reduce typical acquisition errors via real-time updates of the reconstruction results displayed to the clinician. Compared to offline reconstruction systems, the predicted improvement in the output volumes' reconstruction quality represents a significant advantage for their subsequent use within a surgical guidance system. Performing real-time reconstruction also presents the output volumes for use immediately following the acquisition process, which may increase clinical penetration of ultrasound reconstruction by eliminating the wait times required by more common offline approaches.

### 2.3.2 Integration into AR environment

A second major advantage of the present system is that a mechanism is provided to display reconstructed 4D ultrasound datasets within our group’s augmented reality environment for intracardiac therapy [84] alongside tracked 2D TEE, virtual representations of tracked surgical tools and a patient-specific beating-heart surface model derived from preoperative 4D MR or CT imaging. In the original AR environment, the dynamic cardiac model and tracked multiplanar 2D TEE probe serve complementary purposes: the cardiac model provides a global context but may be misregistered to the surgical scene, while tracked 2D TEE provides the real-time imaging required for final detailed navigation onto surgical targets and therefore comprises the “ground truth” coordinate frame to which all virtual objects should ideally be registered, but in turn provides only a thin 2D slice of image data. This AR environment has been used to perform a preliminary mitral valve implantation and atrial septal defect closure in swine [127] and its users can successfully compensate for errors sustained in the registration process transforming the cardiac model into the AR environment’s coordinate system [158]. Previously reported accuracy assessments involved having users target epicardial landmarks on a stationary heart phantom using a tracked pointer tool [215] or employ a tracked catheter to target endocardial landmarks within a beating heart phantom [216]. These studies yielded RMS errors of 2.83 mm and 1.10 mm, respectively.

However, the current workflow integrating the patient-specific cardiac model into the AR environment (which is based on a rigid registration aligning outlines of the aortic and mitral valves specified manually on the preoperative MR/CT and on intraoperative tracked 2D TEE images) does not compensate for any non-rigid deformation between the preoperative images and the intraoperative heart. Such deformations are expected and may be caused by varying patient positioning, intracardiac tool manipulation during therapy delivery or, in the case of our group’s direct surgical access procedure, by the minithoracotomy or the attachment of the Universal Cardiac Introducer [217]. Additionally, clinicians have reported difficulty in using 2D TEE for surgical guidance, specifically in continuously aligning the imaging plane with

Table 2.1: Tradeoffs in choosing an imaging modality for use in image-guided cardiac therapy (++++: Excellent; +++: Good; ++: Somewhat good; +: Poor; (-): Functionality not available)

|                          | Imaging Modality  |                     |                         |                         |
|--------------------------|-------------------|---------------------|-------------------------|-------------------------|
|                          | Tracked<br>2D TEE | Tracked<br>RT3D TEE | Reconstructed<br>4D TEE | Preoperative<br>(MR/CT) |
| Real-time frame rate     | ++++              | +++                 | +                       | (-)                     |
| Spatial resolution       | +++               | ++                  | +++                     | ++++                    |
| Field of view            | ++                | +                   | +++                     | ++++                    |
| 3D context               | +                 | ++                  | +++                     | ++++                    |
| Insensitive to artifacts | ++++              | ++++                | ++                      | ++++                    |
| Natively in AR coords.   | ++++              | ++++                | ++++                    | (-)                     |

structures of interest and in maintaining their spatial orientation when non-standard rotational increments of the multiplanar probe are presented to them.

Adding the ability to perform real-time 4D ultrasound reconstruction for intraoperative imaging within our group’s AR environment represents a significant advantage over the previous reliance on tracked 2D ultrasound, as 4D ultrasound reconstruction provides a significantly greater amount of image data for visualization and registration while maintaining ultrasound’s traditional advantages with respect to ease of OR integration, portability and cost-effectiveness. A reconstructed 4D US dataset cannot completely replace real-time methodologies such as 2D US, RT3D US or fluoroscopy for visualization, but could instead provide a more global 3D context and wider field of view without the potential registration inaccuracies inherent in integrating preoperative images (see Table 2.1 for a visualization of the advantages and disadvantages of the imaging modalities that are of interest here). For example, Knackstedt *et al.* [42] have proposed using reconstructed 3D ICE volumes for surgical guidance during EP catheter ablation procedures.

The time series of reconstructed ultrasound volumes could also serve as the basis for algorithms registering the dynamic surface model derived from preoperative MR or CT images to the intraoperative patient. As previously mentioned, the feature-based registration algorithm currently employed by our group is partially based on outlining the mitral and aortic valves under tracked 2D TEE. The task of manually



delineating these landmarks may be made simpler and more robust by using a reconstructed 4D ultrasound dataset instead. Alternatively, a reconstructed 4D US dataset may be used for image-based rigid registration to the preoperative images, or may even be input into a non-rigid registration algorithm elastically warping the preoperative surface model to match the patient’s intraoperative heart. The availability of the latter algorithm would vastly improve the AR environment by increasing the fidelity with which the dynamic cardiac model represents the patient’s intraoperative heart, although its implementation would be difficult considering the real-time requirements involved in surgical guidance and the vast difference between the two imaging modalities involved. To the best of my knowledge, no algorithm registering reconstructed cardiac ultrasound datasets to MR or CT currently exists, although there are examples using RT3D US instead [60].

Tracked RT3D TEE could also be used to perform these visualization and registration tasks. Ultrasound reconstruction has the advantage of a wider field of view and higher spatial resolution, albeit at the expense of stitch artifacts caused by the extended acquisition time (Table 2.1). RT3D TEE’s smaller field of view, near real-time frame rate and limited imaging time (for patient safety from overheating) make it more suitable for use in final detailed navigation once the surgical target has been approached, and could be integrated into the AR environment in future. RT3D images are routinely stitched together to extend their field of view by acquiring a few RT3D images over several cardiac cycles [76]. This process generates volumes with fewer artifacts than does reconstructed US, but still suffers from the relatively low spatial resolution of each composite subvolume.

Since the presented real-time 4D ultrasound reconstruction system is aimed for use in image-guided therapy rather than for diagnosis alone, it is particularly important to thoroughly evaluate its accuracy. For surgical guidance, intraoperative image data must strictly represent the patient’s anatomy because direct vision of the surgical targets is removed during minimally-invasive interventions. For successful therapy delivery, it is therefore essential that the relationships between entities in the “visible” augmented reality environment that the surgeon relies on closely match those between the “invisible” structures within the patient that are actually being manipulated.

# Chapter 3

## Experimental Validation

This chapter describes the studies performed to evaluate the real-time 4D ultrasound reconstruction system described in Chapter 2, including a human factors experiment that compares 3D and 4D ultrasound reconstructions acquired with and without real-time visualization, experimental porcine imaging demonstrating clinical feasibility of the approach, two accuracy assessment studies performed both within the laboratory and within the operating room and our initial experience with human subjects.

### 3.1 Human Factors

The first goal was to confirm that real-time visualization truly confers an advantage during 4D ultrasound reconstruction. For this study the beating heart phantom shown in Figure 3.1 was used, which consists of a two-chambered heart that is made of a hydrogel material which mimics soft tissue under ultrasound imaging suspended in a distilled water bath (Shelley Medical Imaging Technologies, London, ON, Canada). The phantom is controlled by a servomotor-driven actuator and beats at 17 beats per minute (bpm) while outputting a simulated ECG voltage waveform. A skilled echocardiographer performed several freehand reconstructions by translating a magnetically tracked TTE probe linearly from above the phantom. In some ways, this acquisition task is more difficult than imaging in human subjects because of the com-




Fig. 3.1: The beating heart phantom used in the human factors experiment (courtesy [www.simutec.com](http://www.simutec.com)).

paratively slower heart rate and because the probe must be steadily manipulated without having a surface to rest it on.

3D and 4D reconstructions of the stationary and beating phantom, respectively, were performed under each of two conditions: 1) with real-time visualization, where the clinician could view both the ultrasound monitor and the incremental reconstruction visualization provided by the real-time reconstruction system; and 2) without real-time visualization, where the clinician had access to the ultrasound monitor alone and viewed the output volume(s) only after the reconstruction's completion. To reduce training effects, these two visualization conditions were alternated for both the 3D and 4D acquisition segments of the experimental protocol. Because of its configuration, no attempt was made to shield the beating heart phantom itself from the clinician's view, as this would obstruct the clinician from freely scanning the phantom. Acquisition time increased when real-time visualization was provided compared to when it was not, on average by a factor of 2.4 for 3D US reconstruction and 1.8 for 4D US reconstruction. This is a result of the extra time required to review the reconstruction results and to re-acquire any gapped areas.

Figure 3.2 shows the beneficial effect of real-time visualization on 3D and 4D ultrasound reconstruction, particularly with respect to the number of gaps in the output volume(s). The output volumes are shown after hole filling, meaning that any remaining gaps represent significant acquisition errors. For 3D ultrasound reconstruction,




Fig. 3.2: Example output volumes acquired with and without real-time visualization during the human factors experiment, for both 3D and 4D ultrasound reconstruction. The output volumes are displayed after hole-filling in the equivalent of a four-chamber view. All images represent the clinician's second trial under each imaging paradigm with the exception of the 4D ultrasound reconstruction with real-time visualization, which shows the first attempt. (US depth setting = 14 cm,  $s_{in} = 0.40$  mm,  $s_{out} = 0.81$  mm; for 4D US reconstruction  $HR_{exp} = 17$  bpm,  $\%dev_{HR} = 1.5\%$ ,  $N = 2$ ,  $P = \{0, 1\}$ , images shown represent the 50% R-R interval).

a few gaps resulted when the ultrasound monitor alone was shown to the clinician, while none were present in the output volumes reconstructed using the visualization from the real-time reconstruction system. However, the relative simplicity of the 3D acquisition task makes the difference slight. In contrast, there was a considerable difference between the two visualization conditions for 4D ultrasound reconstruction. Even after training, 4D ultrasound datasets reconstructed when the clinician did not view the incremental reconstruction results showed a large number of gaps, which are dramatically reduced even in the first dataset reconstructed using the real-time reconstruction visualization. This study reveals the significant advantages of the real-time 4D US reconstruction system, as extremely controlled probe manipulation is required to reconstruct usable 4D ultrasound datasets but is difficult to execute without visual guidance.

## 3.2 Porcine Imaging

Two swine were imaged to qualitatively evaluate the real-time 4D ultrasound reconstruction system in a more clinically-relevant setting. Here, the goals were to determine intraoperative feasibility and to evaluate the quality of the output volumes under different acquisition techniques and as a function of respiratory motion.

All animals were treated according to the policies of the Animal Use Subcommittee of The University of Western Ontario Council on Animal Care. Respiration was controlled by a mechanical ventilator, which for these acute studies could be safely suspended for up to 6-8 minutes. To mimic more closely the allowable ventilation in human subjects, 4D ultrasound reconstruction was also performed by suspending and resuming respiration in alternating 30 s “imaging” and 30 s “rest” blocks (“alternating respiration”). The field generator associated with the magnetic tracking system used to track the TEE probe, reference sensor and other tools was placed within the OR mattress directly below the pig’s heart to maintain a suitable tracking volume while minimizing distortions in the magnetic tracking system’s field caused by ferromagnetic OR equipment. In both studies, a 6-DOF sensor (Traxtal, Toronto, ON, Canada) was sutured to the pig’s back to act as the reference sensor, and in the first porcine study

an additional 6-DOF sensor (Traxtal) was sutured onto the pig’s chest as part of a parallel study. In the first porcine experiment, the reference sensor and the chest sensor were tracked continuously while only one of the sensors attached to the TEE and TTE probes was plugged into the tracking system at any point in time. There was no chest sensor in the second experiment and the reference, TEE and TTE sensors were tracked throughout.

The multiplanar TEE probe was used to acquire images used to reconstruct the majority of the 4D US datasets, each of which was generated using retrospective gating and comprised two output volumes. In the first porcine study the output volumes corresponded to the 0% and 50% R-R intervals, while in the second study the mitral valve was imaged under M-mode echocardiography to determine that the 25% and 80% R-R intervals corresponded approximately to mid-systole and mid-diastole, respectively. To do this, the ECG signal was acquired while simultaneously performing M-mode echocardiography through the mitral valve, which can identify both mid-systole and end-systole. This allows one to measure the duration of the cardiac cycle ( $t_{\text{cycle}}$ ), the approximate time difference between each R-wave and mid-systole ( $t_{\text{mid}}$ ) and the approximate time difference between each R-wave and end-systole ( $t_{\text{end}}$ ). The % R-R interval corresponding to mid-systole is  $(t_{\text{mid}}/t_{\text{cycle}} \times 100\%)$  and the % R-R interval corresponding to mid-diasole is  $\left(\left(\frac{t_{\text{cycle}}}{2} - \frac{t_{\text{end}}}{2} + t_{\text{end}}\right)/t_{\text{cycle}} \times 100\%\right)$ . Real-time visualization of the reconstruction results was provided to the echocardiographer via a screen mounted above the ultrasound machine’s monitor, allowing the clinician to view both simultaneously.

To maximize the quality of the reconstructed 4D ultrasound dataset, an acquisition trajectory must provide a series of good 2D US views while allowing the clinician to manipulate the US imaging plane slowly and with a high degree of control. When imaging swine, keeping the transducer relatively stationary and rotating the imaging plane at  $1^\circ$  increments provided the best delineation of cardiac structures, including all four cardiac chambers, the left atrial appendage and the mitral and aortic valves (Figure 3.3). As is typical for such gated systems, the 4D ultrasound reconstructions suffer from motion artifacts that follow the acquisition trajectory (in this example, artifacts radiate from the rotational axis). Comparing the stitch artifacts in these




Fig. 3.3: Example 4D ultrasound reconstruction of a beating porcine heart acquired using the TEE probe with a rotational acquisition under alternating respiration. The two output volumes represent the (a) 0% R-R interval and (b) 50% R-R interval. (Fig 1: US depth setting = 11 cm,  $s_{in} = 0.32$  mm,  $s_{out} = 0.64$  mm, output dimensions =  $472 \times 502 \times 370$ ,  $HR_{exp} = 107$  bpm,  $\%dev_{HR} = 10\%$ ,  $N = 2$ ,  $P = \{0, 1\}$ ).




Fig. 3.4: Example 3D ultrasound reconstruction of an excised porcine heart in a water bath, acquired using the TEE probe with a rotational acquisition. (US depth setting = 11cm,  $s_{in}$  = 0.32 mm,  $s_{out}$  = 0.64 mm, output dimensions =  $320 \times 240 \times 320$ ).

output volumes to those in a similarly acquired 3D US reconstruction of an excised porcine heart (Figure 3.4) emphasizes the degree to which artifacts are caused by the additional error sources associated with imaging a moving organ.

Compared to the rotational technique where the transducer is relatively stationary, the increased tracking error sustained when performing ultrasound reconstruction by moving the 2D US probe generally introduces more artifacts in the output volumes. Nevertheless, a manual fan acquisition can successfully visualize the left atrium, left ventricle and mitral valve in pigs (Figure 3.5). Linearly translating the probe (“pull-back” acquisition) was not as successful in pigs as can be expected in humans because the air-filled right bronchus extends between the esophagus and the heart, impeding ultrasound imaging. Finally, freehand epicardial imaging using the TTE probe was attempted. In this technique, the clinician translated the probe directly over the beating heart with access provided by a left minithoracotomy. However, this method produced inferior output volumes, primarily because of difficulties in controlling the probe’s movement while maintaining contact with the rapidly beating heart. In addition, this form of imaging places the sensor attached to the TTE probe at the edge of the field generator’s tracking volume, leading to increased tracking error.



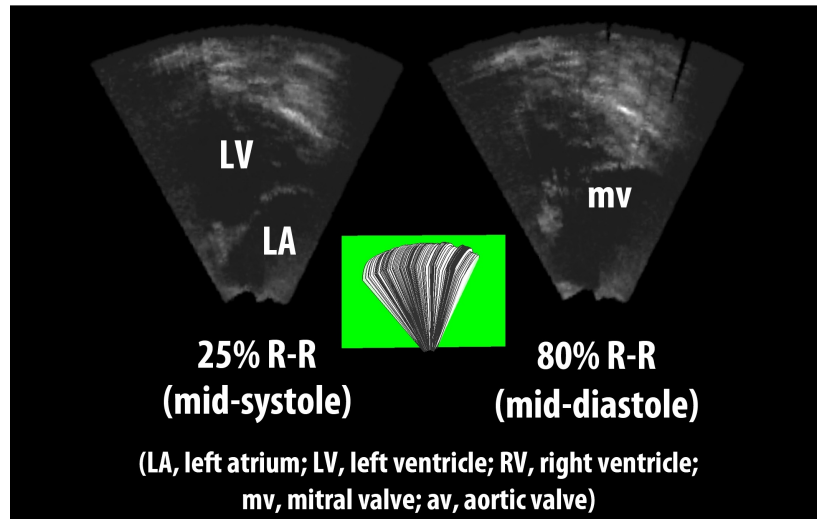


Fig. 3.5: Example 4D ultrasound reconstruction of a beating porcine heart, acquired using the TEE probe with a fan acquisition under suspended respiration. The two output volumes represent the (a) 25% R-R interval (approximately mid-systole) and (b) 80% R-R interval (approximately mid-diastole). (Fig 2: US depth setting = 12 cm,  $s_{in} = 0.35$  mm,  $s_{out} = 0.69$  mm, output dimensions =  $411 \times 464 \times 351$ ,  $HR_{exp} = 96$  bpm,  $\%dev_{HR} = 10\%$ ,  $N = 20$ ,  $P = \{5, 16\}$ ).

I examined the effect of respiratory motion on reconstruction quality by performing reconstructions under free breathing, suspended respiration and alternating respiration (Figure 3.6). Respiratory motion is indeed a significant error source, as reconstructions acquired while the pig was breathing freely show more visible stitch artifact compared to those acquired under suspended or alternating respiration. Suspending respiration for an extended period of time is not possible in human subjects, but the alternating respiration strategy produced qualitatively similar images and can be safely performed in humans.

During the porcine studies, the clinician was observed to use the real-time 4D ultrasound reconstruction system to identify and successfully correct acquisition errors. In addition, suboptimal acquisitions could quickly be recognized as such and abandoned, while with an offline reconstruction system one would need to complete the entire acquisition and wait for the reconstruction results. For the TEE reconstructions, the minimum and maximum acquisition times were three and eight minutes, respectively. However, for each type of acquisition there existed one or more tri-




Fig. 3.6: Example output volumes representing a beating porcine heart at the 0% R-R interval, acquired using the TEE probe with a rotational acquisition under (a) free respiration, (b) suspended respiration. Figure 3.3a shows the equivalent for alternating respiration. (Fig 1: US depth setting = 11 cm,  $s_{in} = 0.32$  mm,  $s_{out} = 0.64$  mm, output dimensions =  $473 \times 492 \times 387$  for (a) and  $473 \times 502 \times 369$  for (b),  $HR_{exp} = 107$  bpm,  $\%dev_{HR} = 2.5\%$  for (a) and  $10\%$  for (b),  $N = 2$ ,  $P = \{0, 1\}$ ).

als where the acquisition time was 3-4 minutes, the shortest time possible given the constraints of the gating scheme used. Given enough experience with the system, it should be possible to perform all types of acquisition in this time frame. Since 4D ultrasound reconstruction is more difficult in swine than in humans due to a faster heart rate and anatomical differences that degrade 2D US imaging, the successful 4D US reconstructions in pigs described here constitutes a strong validation of the system.

### 3.3 Accuracy Assessment with Dynamic Phantom

While some quantitative validations of 4D ultrasound reconstruction systems aimed for cardiac imaging compare measurements derived from beating heart reconstructions to those based on 2D US [42] or MRI [81], many groups assess the accuracy of their systems by imaging static objects [200–202], thus ignoring the many error sources associated with imaging a beating organ. All imaging systems that are intended for interventional guidance are subject to the constraints on procedural accuracy that must be satisfied for positive clinical outcomes, and hence their performance must be fully studied before they are used therapeutically.

Here, a comprehensive accuracy assessment of the 4D ultrasound reconstruction system using a dynamic phantom is presented, an approach also taken in [218]. This study used a precursor of the final reconstruction system, which was implemented as a standalone application, incorporated prospective gating only, and produced output volumes that were in the same coordinate system as (and axis-aligned with) the ultrasound key frame rather than the reference sensor’s coordinate system. The dynamic phantom was imaged while it was either stationary (3D US reconstruction) or moving (4D US reconstruction), and accuracy metrics were calculated by comparing manual segmentations of the resulting images to gold standards derived from micro-CT and tracking information. This protocol is similar to that described by Wiles *et al.* [214], incorporating measurements of both trueness and precision, and is extended here to determine the system’s ability to localize and delineate moving objects. In particular, the assessment characterizes the quality of the imaging parameters, magnetic track-

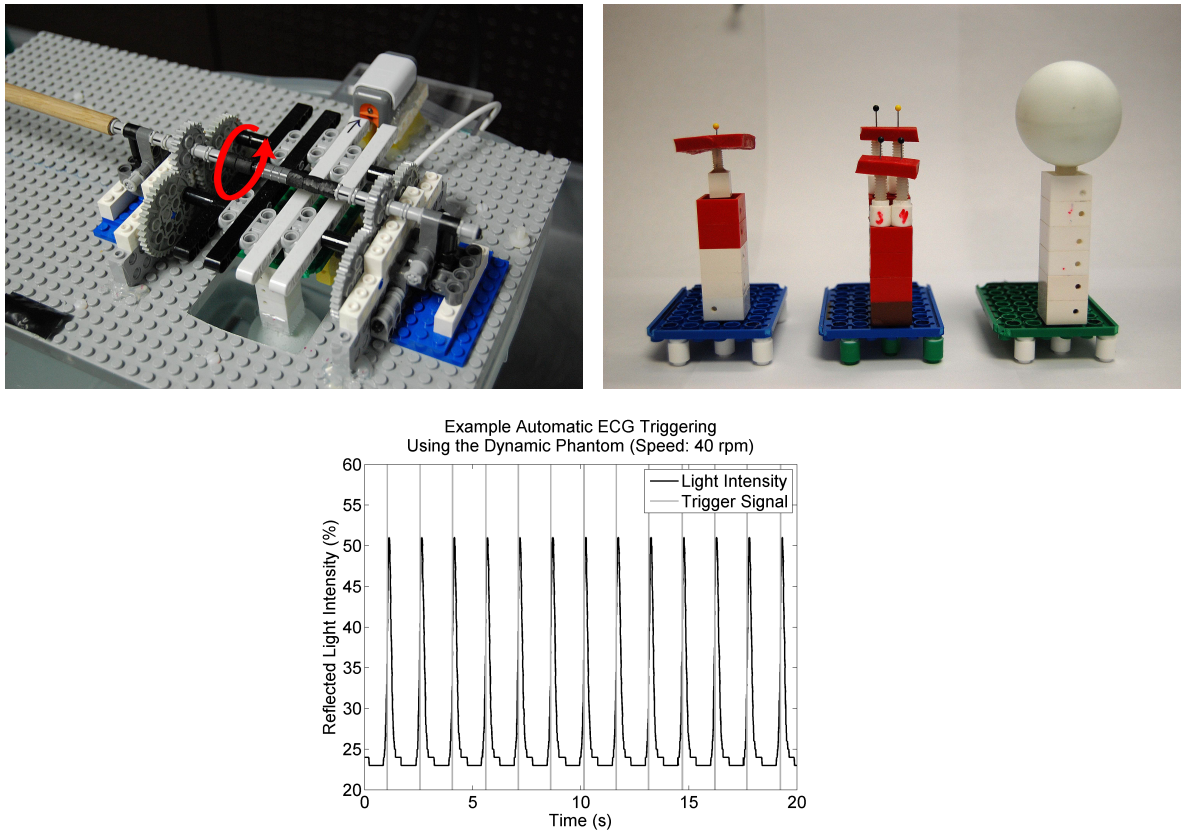


Fig. 3.7: (a) The dynamic phantom, annotated to show the direction of phantom movement; (b) The point-source, distance and spherical phantoms; (c) Example light sensor signal, with automatically-detected ECG trigger signals that indicate the beginning of a new cycle.

ing, spatial calibration, temporal calibration and gating in combination, and could be used to evaluate any form of 4D ultrasound, such as RT3D US, when imaging moving targets.

To assess the accuracy with which the system performs 3D and 4D ultrasound reconstruction, I built a dynamic phantom (Figure 3.7a) comprising a simple robot which moves a base to which one of the following three modular phantoms can be attached (Figure 3.7b): 1) a point-source phantom, composed of a 2.45 mm diameter hard plastic polycarbonate sphere mounted upon a stainless steel pin above a small piece of sound absorbing material (Sorbothane Inc., Ohio, USA); 2) a distance phantom, comprising of four similar polycarbonate spheres arranged in a ramp config-

uration above two pieces of sound absorbing material (inspired by the wedge phantom presented by Barratt *et al.* [219]); and 3) a spherical phantom, consisting of a 40.25 mm diameter hollow table tennis ball fixed upon a nylon screw. To emulate the moving anatomy within the beating heart, the custom-built robot (LEGO® Mindstorms NXT™, The LEGO Group, Billund, Denmark) was programmed to move the phantom around a circular path of diameter 2.4 cm at a constant speed using an open-source robot control library [220]. The dynamic phantom’s movement was not perfectly smooth and the duration of each cycle varied slightly, thus mimicing “imperfections” in the beating heart. A 6-DOF sensor (Traxtal, Toronto, ON, Canada) was connected to the phantom so that it could be localized by the magnetic tracking system. To approximate the speed of sound in tissue, all experiments were performed within a bath of 7% glycerol in water solution by mass at room temperature.

Within the operating room, the current phase of the patient’s heart within the cardiac cycle is determined by analysis of the ECG signal. To mimic the ECG signal in the dynamic phantom, the light sensor of the LEGO® Mindstorms NXT™ kit was positioned so that the emitted LED light reflected off of one of the moving white beams on the phantom’s base. This results in a signal from the sensor’s phototransistor that varied predictably as the phantom moved. The beginning of a new cycle was automatically detected when the difference between the minimum and maximum light signal within a predefined time interval exceeded an empirically-determined threshold that depended on motor speed (Figure 3.7c).

The experimental setup is shown in Figure 3.8a. The TEE probe could be secured at varying heights and orientations to promote imaging of the dynamic phantom from below. Both the sensor attached to the ultrasound probe and to the phantom were tracked relative to a 6-DOF reference sensor (Northern Digital Inc., Waterloo, ON, Canada).

Because motors can influence the accuracy of magnetic tracking systems, primarily causing jitter in the returned transformations, the robot’s motor was positioned as far away from the experimental setup as possible. To examine any effect on the tracking precision, the sensor attached to the stationary US transducer was tracked over approximately four minutes while the motor was turned off, run without a phan-




Fig. 3.8: (a) The experimental setup; (b) Schematic illustrating the transformations used to calculate the expected position of the phantom within the ultrasound beam coordinate system. In the diagram, the phantom intersects the  $90^\circ$  plane of the multiplanar transducer, which lies along the  $\mathbf{yz}$ -plane of the ultrasound beam coordinate system.

tom attached, and run with the spherical phantom attached. Trials were acquired at 20, 30, 40, 50 and 60 rpm, resulting in a point cloud for each trial representing the measured 3D coordinates for the TEE sensor. Precision was quantified by calculating the histogram over the Euclidean distances from each coordinate in the point cloud to the point cloud's centroid (Figure 3.9). The results show that the motor itself did not have a noteworthy influence on the tracking precision, as the histograms generated when the motor ran without a phantom attached followed those corresponding to when the motor was turned off. However, there was a difference at low motor speeds (20/30 rpm) when the spherical phantom was attached.

Multiple 3D and 4D acquisitions were performed for each phantom using the calibration matrix resulting from a single spatial calibration. Twenty-five static 3D ultrasound volumes were reconstructed using the rotational acquisition approach at  $1^\circ$  increments for each phantom at varying depths (i.e. distance from the transducer) and probe orientations. For each phantom, two dynamic 4D US datasets comprising five 3D volumes were reconstructed for two transducer positions and at five different speeds (20, 30, 40, 50 and 60 rpm, corresponding to 2.5, 3.8, 5.0, 6.3 and 7.5 cm/s, respectively). Once again, the rotational acquisition approach was used, with the

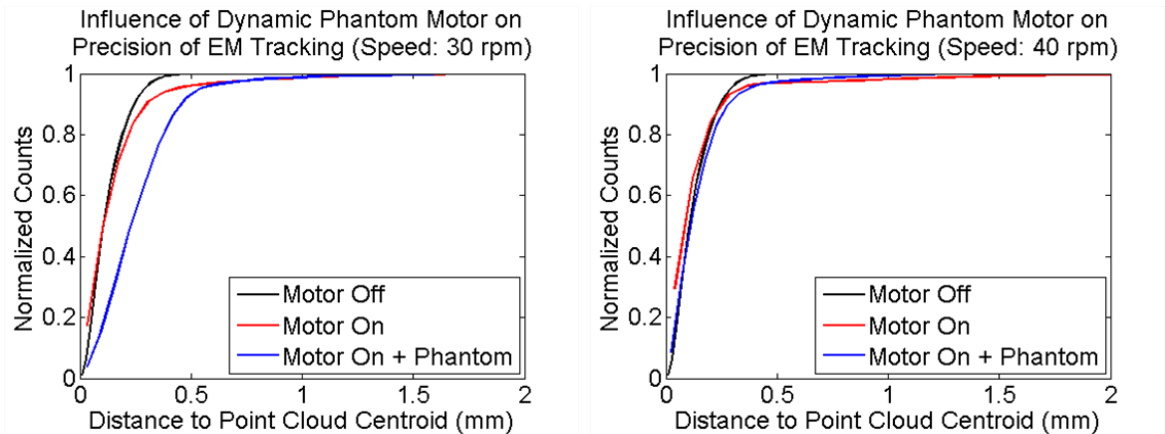


Fig. 3.9: Assessment of the influence of the dynamic phantom’s motor on the magnetic tracking system. Histograms were calculated over the Euclidean distances between each coordinate measured for the sensor attached to the stationary TEE probe and the point cloud’s centroid. Trials were acquired while the motor was turned off (“Motor Off”), turned on with no phantom attached (“Motor On”) and turned on with the spherical phantom attached (“Motor On + Phantom”). (a) shows a representative histogram for slow motor speed (20/30 rpm) while (b) shows a representative histogram for fast motor speeds (40/50/60 rpm).

imaging plane rotated by  $1^\circ$  approximately once every two cycles (US depth setting = 12 cm,  $s_{in} = 0.35$  mm; for 3D US reconstruction  $s_{out} = 0.35$  mm, output volumes were resampled to  $s_{out} = 0.69$  mm before subsequent analysis to reduce file size; for 4D US reconstruction  $s_{out} = 0.69$  mm,  $HR_{exp} = 20/30/40/50/60$  bpm,  $\%dev_{HR} = 1.5\%$ ,  $N = 5$ ,  $P = \{0, 1, 2, 3, 4\}$ ). The axes of the output volume coordinate system were defined by the ultrasound beam at  $0^\circ$  rotation, such that the  $\mathbf{x}$ -axis points in the lateral direction, the  $\mathbf{y}$ -axis points in the axial (depth) direction, and the  $\mathbf{z}$ -axis points in the elevation direction (see Figure 3.8b). During 4D ultrasound reconstruction, the ultrasound transducer was positioned so that the plane of the phantom’s circular movement lay approximately parallel to the  $\mathbf{yz}$  plane. Following each reconstruction, the transform matrices acquired for the tracked US probe, phantom and reference sensors and their corresponding timestamps were saved for future use.

To determine the reconstruction accuracy, the expected phantom position for each image was calculated based on the phantom’s structure and the position of the attached sensor, and was then compared to an observed position derived from manual

segmentation of the reconstructed output volumes. The process of deriving the expected phantom position within the output volume coordinate system proceeded as follows and is depicted in Figure 3.8b.

- **Micro-CT:** A micro-CT image (eXplore Locus Ultra, GE Healthcare Biosciences, London, ON, Canada) was acquired of each phantom and a set of 3D coordinates outlining each phantom’s sphere(s) was manually generated. The centre  $p_{mCT}$  and radius  $r_{mCT}$  of each sphere was then found using a least-squares sphere-fit algorithm [221].
- **Point-based registration:** Each phantom was equipped with ten divots, which were manually identified in the micro-CT image and localized relative to the sensor mounted on the phantom using a tracked pointer tool. The transformation  $T_{div}$  between micro-CT coordinates and those relative to the phantom’s attached sensor was then found using a point-based linear registration.
- **Magnetic tracking:** The transformations  $T_{ptm}$  and  $T_{US}$  result from tracking the sensors mounted on the phantom and ultrasound probe, respectively, relative to the reference sensor. For each of the two sensors, the series of transformations acquired over the course of each reconstruction must be aggregated to determine the  $T_{ptm}$  and  $T_{US}$  used to calculate the expected phantom position. When the transformations were approximately constant throughout the acquisition (i.e.  $T_{US}$  during 3D and 4D US reconstructions and  $T_{ptm}$  during 3D US reconstructions) the transformation that gives rise to the median coordinate, after removing outliers, was used. Outlier removal was achieved by removing those points within the point cloud whose distance to the median coordinate was outside an empirically-determined percentile. For 4D US reconstructions, it was assumed that the motor operates at a constant speed and  $T_{ptm}$  was calculated for each phase by:
  1. Finding the transformation that gave rise to the median coordinate, after removing outliers, of all those transforms identified by the reconstruction software as corresponding to phase 0



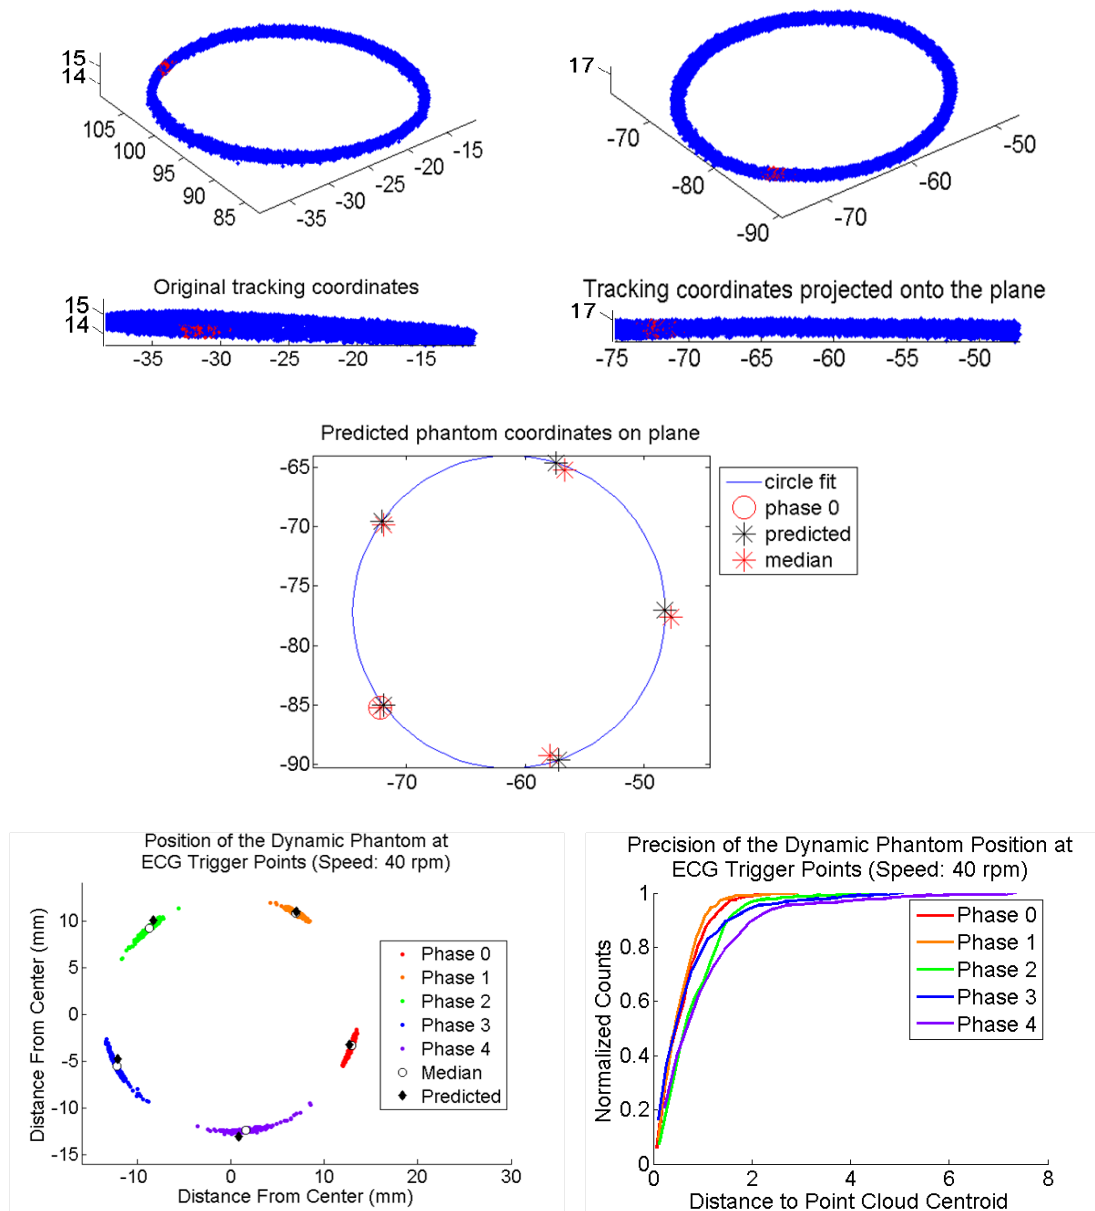


Fig. 3.10: Predicting each phase's transform matrix for the sensor attached to the dynamic phantom's platform. (a) Example tracked coordinates for the phantom sensor, where red points correspond to the beginning of phase 0; (b) Results of applying PCA to the points in (a); (c) and (d) The evenly-distributed coordinates for each phase corresponding to the predicted  $T_{US}$  matrices, and the median coordinates for the point clouds associated with each phase; (e) Each phase's histogram over the Euclidean distances between each associated point and the point cloud's centroid.

2. Projecting all coordinates onto the 3D space defined by the two principal eigenvectors found by applying a principal component analysis (PCA) to all of the tracking data (Figures 3.10a and 3.10b)
3. Subdividing the best-fit circle of the  $(x,y)$  coordinates into  $N$  evenly spaced sections starting at the projection of the point for phase 0 found in step 1 (Figure 3.10c)
4. Projecting the coordinates associated with each phase back into the original 3D space.

The coordinates acquired for the tracked phantom whose timestamps were associated with the beginning of each phase (i.e. those corresponding to an inserted 2D US image) are shown in Figure 3.10d for an example 4D US reconstruction. This figure also shows the median coordinate for each phase’s point cloud and the coordinates predicted for each phase using the procedure described above. In addition, Figure 3.10e shows the histograms for each phase’s point cloud calculated over the Euclidean distances from each point to the point cloud’s centroid. Although the accuracy with which  $T_{ptm}$  can be predicted for 4D US reconstruction decreases as the phantom’s cycle progresses, the median coordinate for each phase’s point cloud cannot be used as it is derived from the very gating mechanism being validated. The RMS Euclidean distance between the predicted coordinates and the median coordinates for each phase was approximately 0.75 mm for all three phantoms, with a maximum of 1.2 mm over all trials.

- **Calculating the expected position:** Finally, the spatial calibration matrix  $C_{US}$  gives the transformation from the 2D US images to the sensor mounted on the ultrasound probe, so the expected position of the phantom  $p_{exp}$  within the output volume coordinate system is given by

$$p_{exp} = (C_{US})^{-1}(T_{US})^{-1}T_{ptm}T_{div}P_{mCT} \quad (3.1)$$

The acquired output volumes were characterized by a single user, who marked the centre of the bright spots representing the spheres on the point-source and distance

phantoms and multiple points on the edges of the spherical phantom’s outline closest to the transducer’s fan origin. The user dealt with the motion artifacts present in the 4D US reconstruction output volumes by attempting to choose the centre of the cloud of bright pixels for the point-source and distance phantoms, and by defining a smooth curve for the spherical phantom. The single coordinates representing each sphere for the point-source and distance phantoms were taken directly as the observed phantom coordinates  $p_{obs}$  for comparison to the expected phantom positions, while the points outlining the spherical phantom were input to an ellipsoid-fitting algorithm [222]. Note that this algorithm’s assumption that the length of the ellipsoid’s minor axis is at least half that of its major axis will be true for all images of the spherical object. Finally, the method described by Schneider and Eberly [223] was used to extract the centre ( $p_{obs}$ ), the three radii lengths ( $r_{max}$ ,  $r_{med}$  and  $r_{min}$ ), and the three eigenvalues ( $\lambda_1$ ,  $\lambda_2$  and  $\lambda_3$ ) corresponding to the equation of the best-fit ellipsoid defining the image of the spherical phantom.

Example 3D and 4D US reconstructions of the three phantoms are shown in Figure 3.11. Qualitatively, the output volumes generated by the real-time reconstruction system show good correspondence to the actual phantom structures. As expected, the motion artifacts in the 4D ultrasound reconstructions become more prominent with increasing phantom speed. Figure 3.12 depicts further examples of these artifacts. To quantify the error sources present in 3D and 4D ultrasound reconstruction, the expected and observed phantom characterizations were compared with respect to localization, distance, volume and shape. Example comparisons are depicted in three dimensions in Figure 3.13.

For each phantom, the localization accuracy of the real-time reconstruction system was assessed by comparing the expected and observed phantom centres. Table 3.1 lists summary statistics computed for the Euclidean centre reconstruction error  $CRE = \|p_{obs} - p_{exp}\|$  and for the errors along each of the three axes in the output volume system. For the 3D US reconstructions, the RMS Euclidean centre reconstruction error ranged between 1.2-1.7 mm, with the spherical phantom giving both the smallest mean error and the least variability. The RMS Euclidean centre reconstruction error for 4D US reconstruction was similar for all three phantoms and ranged between




Fig. 3.11: Example reconstructions of the point-source (top), distance (middle) and spherical (bottom) phantoms. Column 1 shows 3D reconstructions of a static phantom, while columns 2-4 show the phase 0 volumes from ultrasound reconstructions of a dynamic phantom moving at 20, 40 and 60 rpm, respectively.




Fig. 3.12: Example motion artifacts present during 4D ultrasound reconstruction (spherical phantom at 50 rpm).




Fig. 3.13: Sample ultrasound volumes overlaid with 3D representations of the expected (light blue) and observed (red) phantom positions for the (a) point-source, (b) distance and (c) spherical phantoms.

Table 3.1: Summary statistics for centre reconstruction error in the dynamic phantom experiment over  $N$  trials. The X, Y and Z directions correspond to the lateral, axial and elevation directions, respectively, of the key frame used to determine the output volumes’ origin (which in this study corresponded almost exactly to the 2D US image acquired at  $0^\circ$ , as shown in Figure 3.8b.)

| Direction | Centre Reconstruction Error (mm) |       |         |     |                         |       |           |     |
|-----------|----------------------------------|-------|---------|-----|-------------------------|-------|-----------|-----|
|           | Static Phantom (3D US)           |       |         |     | Dynamic Phantom (4D US) |       |           |     |
|           | RMS                              | Mean  | Std Dev | $N$ | RMS                     | Mean  | Std. Dev. | $N$ |
|           | Point-Source Phantom             |       |         |     |                         |       |           |     |
| Euclidean | 1.54                             | 1.46  | 0.50    | 25  | 2.68                    | 2.54  | 0.87      | 50  |
| X         | 0.67                             | 0.25  | 0.64    | 25  | 0.77                    | -0.33 | 0.70      | 50  |
| Y         | 1.17                             | 0.99  | 0.64    | 25  | 1.73                    | 1.18  | 1.28      | 50  |
| Z         | 0.74                             | 0.02  | 0.75    | 25  | 1.90                    | -0.68 | 1.79      | 50  |
|           | Distance Phantom                 |       |         |     |                         |       |           |     |
| Euclidean | 1.72                             | 1.59  | 0.65    | 100 | 2.53                    | 2.39  | 0.81      | 200 |
| X         | 0.93                             | 0.38  | 0.86    | 100 | 0.51                    | -0.10 | 0.50      | 200 |
| Y         | 0.99                             | 0.82  | 0.57    | 100 | 1.62                    | 0.91  | 1.34      | 200 |
| Z         | 1.04                             | 0.61  | 0.85    | 100 | 1.87                    | 0.07  | 1.88      | 200 |
|           | Spherical Phantom                |       |         |     |                         |       |           |     |
| Euclidean | 1.22                             | 1.18  | 0.30    | 25  | 2.71                    | 2.54  | 0.95      | 50  |
| X         | 0.69                             | -0.37 | 0.59    | 25  | 1.24                    | -1.12 | 0.54      | 50  |
| Y         | 0.71                             | 0.64  | 0.30    | 25  | 1.62                    | 1.23  | 1.07      | 50  |
| Z         | 0.71                             | -0.36 | 0.63    | 25  | 1.78                    | -0.57 | 1.70      | 50  |

2.5-2.7 mm. Compared to the 3D US reconstructions, both the magnitude and the variability of the error increased; for all three phantoms this was attributed first to increases in the mean and variability of the error in the  $\mathbf{z}$ -direction followed by increases in the  $\mathbf{y}$ -direction. For both 3D and 4D ultrasound imaging and for all three phantoms, there was also a bias towards overestimating the distance in the depth ( $\mathbf{y}$ -) direction between the observed phantom centre and the ultrasound fan origin.

To evaluate the accuracy with which distances are preserved, the distance reconstruction error  $DRE = \|p_{1obs} - p_{2obs}\| - \|p_{1exp} - p_{2exp}\|$  was calculated for the two pairs of diagonally-opposite spheres on the distance phantom (Table 3.2). Although the expected Euclidean distance between each pair of spheres is the same for all image volumes, the expected distances in the  $\mathbf{x}$ ,  $\mathbf{y}$  and  $\mathbf{z}$  directions depend on the pose of the ultrasound transducer relative to the phantom and changes as the ultrasound

Table 3.2: Summary statistics for distance reconstruction error in the dynamic phantom experiment over  $N$  trials (mean expected Euclidean distance = 21.0 mm). The use of percentages and the reasons for varying  $N$  are explained in the main body of text. The X, Y and Z directions are as explained in Table 3.1.

| Direction | Distance Reconstruction Error (mm) |        |         |     |                         |        |           |     |
|-----------|------------------------------------|--------|---------|-----|-------------------------|--------|-----------|-----|
|           | Static Phantom (3D US)             |        |         |     | Dynamic Phantom (4D US) |        |           |     |
|           | RMS                                | Mean   | Std Dev | $N$ | RMS                     | Mean   | Std. Dev. | $N$ |
| Euclidean | 0.54                               | 0.32   | 0.44    | 50  | 0.97                    | 0.36   | 0.91      | 100 |
| Euclidean | 2.70%                              | 1.62%  | 2.19%   | 50  | 4.83%                   | 1.80%  | 4.50%     | 100 |
| X (%)     | 5.84%                              | 3.65%  | 4.67%   | 21  | 7.51%                   | 5.28%  | 5.40%     | 50  |
| Y (%)     | 4.49%                              | 3.04%  | 3.35%   | 42  | 7.12%                   | 2.54%  | 6.69%     | 100 |
| Z (%)     | 4.42%                              | -1.72% | 4.17%   | 23  | 8.94%                   | -0.39% | 8.98%     | 100 |

transducer is repositioned on each trial. Since a single expected distance for each axis does not exist, these distance reconstruction errors are reported as percent differences  $\%_{DRE} = 100(\|p_{1obs} - p_{2obs}\| - \|p_{1exp} - p_{2exp}\|) / (\|p_{1exp} - p_{2exp}\|)$ . Furthermore,  $\%_{DRE}$  along a particular axis was calculated only when the expected distance along the axis was  $\geq 10$  mm, since small variations in the manual point-identification process manifest as substantial distance errors when the expected distance is very small. The RMS Euclidean distance reconstruction error for 3D US reconstruction was 0.54 mm and increased to 0.97 mm for 4D US reconstruction. This increase was attributed mostly to an increase in variability, which was most prominent in the  $\mathbf{z}$ -direction. In both 3D and 4D US reconstruction, there was a bias towards overestimating the Euclidean distance between objects, and there was no striking difference between the RMS distance reconstruction errors in the  $\mathbf{x}$ -,  $\mathbf{y}$ - and  $\mathbf{z}$  directions.

The spherical phantom was used to determine the volume reconstruction error  $VRE = v_{obs} - v_{exp}$  and the major, medium and minor radii length errors ( $r_{maj} - r_{mCT}$ ), ( $r_{med} - r_{mCT}$ ) and ( $r_{min} - r_{mCT}$ ), by comparing the known structure of the table tennis ball to the parameters of the best-fit ellipsoids derived from the image data (Table 3.3). The volume was overestimated in every image volume analyzed, with RMS volume reconstruction errors for 3D and 4D ultrasound reconstruction corresponding to percent differences of 18.24% and 16.75%, respectively.

Finally, object shape was evaluated in the 3D and 4D ultrasound reconstructions

Table 3.3: Summary statistics for volume reconstruction error and eccentricity in the dynamic phantom experiment over  $N$  trials (expected volume = 34167.95 mm<sup>3</sup>, expected radius = 19.59 mm). The X, Y and Z directions are as explained in Table 3.1.

| Metric | Volume Reconstruction Error (mm <sup>3</sup> ) and Radii Length Error (mm) |       |         |     |                         |       |           |     |
|--------|--|-------|---------|-----|-------------------------|-------|-----------|-----|
|        | Static Phantom (3D US)   |       |         |     | Dynamic Phantom (4D US) |       |           |     |
|        | RMS  | Mean  | Std Dev | $N$ | RMS                     | Mean  | Std. Dev. | $N$ |
| Volume | 5741   | 5675  | 881     | 25  | 5271                    | 5146  | 1152      | 50  |
| Volume | 18.2%  | 18.0% | 2.8%    | 25  | 16.8%                   | 16.4% | 3.7%      | 50  |
| Major  | 1.47   | 1.46  | 0.20    | 25  | 1.57                    | 1.54  | 0.28      | 50  |
| Medium | 1.11   | 1.10  | 0.21    | 25  | 1.06                    | 1.02  | 0.27      | 50  |
| Minor  | 0.80   | 0.79  | 0.16    | 25  | 0.60                    | 0.49  | 0.35      | 50  |

| Metric      | Medium/Minor Axis Ratios (unitless), Fractional Anisotropy (unitless) and Spherical Anisotropy (unitless) |       |         |     |                         |       |           |     |
|-------------|---|-------|---------|-----|-------------------------|-------|-----------|-----|
|             | Static Phantom (3D US)  |       |         |     | Dynamic Phantom (4D US) |       |           |     |
|             | RMS   | Mean  | Std Dev | $N$ | RMS                     | Mean  | Std. Dev. | $N$ |
| Medium Axis | 0.98  | 0.98  | 0.009   | 25  | 0.98                    | 0.98  | 0.016     | 50  |
| Minor Axis  | 0.97  | 0.97  | 0.007   | 25  | 0.95                    | 0.95  | 0.018     | 50  |
| FA          | 0.034   | 0.033 | 0.007   | 25  | 0.057                   | 0.053 | 0.021     | 50  |
| SA          | 0.064   | 0.063 | 0.014   | 25  | 0.102                   | 0.096 | 0.035     | 50  |

by examining the eccentricity of the ellipsoids characterizing the images of the spherical phantom. First, the medium axis ratio ( $r_{med}/r_{maj}$ ) and the minor axis ratio ( $r_{min}/r_{maj}$ ) were computed: in the ideal case all three radii lengths are equal and these ratios are equal to one. Two metrics were also borrowed from the evaluation of diffusion anisotropy in diffusion tensor MRI imaging, specifically fractional anisotropy (FA) [224] and a spherical anisotropy metric  $SA = 1 - (\lambda_3/\lambda_1)$  [225]. Ideally, anisotropy is low for a perfect sphere and these metrics are equal to zero. The eccentricity results, listed in Table 3.3, showed a consistently excellent conservation of shape, with RMS medium and minor axis ratios of 0.98 and 0.97, respectively, for the images of the spherical phantom derived from 3D US reconstruction. The results from 4D US reconstruction are also very good, although slightly worse than those from 3D US reconstruction: the variability for both ratios increases compared to 3D ultrasound imaging, and although the RMS medium axis ratio is unchanged, the RMS minor axis ratio decreases to 0.95. The RMS fractional and spherical anisotropies were



also small for both 3D and 4D ultrasound reconstruction, signifying a near-spherical object, and once again do increase slightly when comparing 4D to 3D ultrasound reconstruction.

Plots of the 3D ultrasound reconstruction errors versus depth are shown in Figure 3.14. The centre reconstruction errors for the point-source and distance phantoms showed a strong depth dependency, which is not as apparent for the spherical phantom. For all three phantoms, plotting the centre reconstruction error in the  $y$ -direction shows the most significant depth dependency of all three axes, and is particularly clear for the point-source and distance phantoms (plot not shown). We also found that volume reconstruction error had an inverse relationship with distance from the transducer, and that neither distance reconstruction error nor eccentricity were depth-dependent.

Strong conclusions regarding the dependence of 4D US reconstruction accuracy on the speed of the moving object should not be drawn since there were only ten trials for each phantom for each motor speed tested, but the distributions of the 4D ultrasound reconstruction errors are nevertheless shown in Figure 3.15 as a function of speed. In general, the magnitude of the centre reconstruction error had a tendency to increase with speed. The RMS values for distance reconstruction error and fractional anisotropy also increased with speed, although volume reconstruction error showed no consistent relationship.

### 3.4 Intraoperative Accuracy Assessment

The results from the dynamic phantom experiment are promising, yet intraoperative reconstruction errors may differ from what can be measured in the laboratory. In particular, there may be differences related to ultrasound imaging properties within the two environments, magnetic tracking accuracy, the quality of ECG gating, heart rate variability and patient movement. I began to address these factors, examining the first two in particular, by performing an intraoperative accuracy assessment. In this study, I imaged a stationary pointer tool suspended within the beating heart of the second pig and compared the tooltip location in the reconstructed volumes to

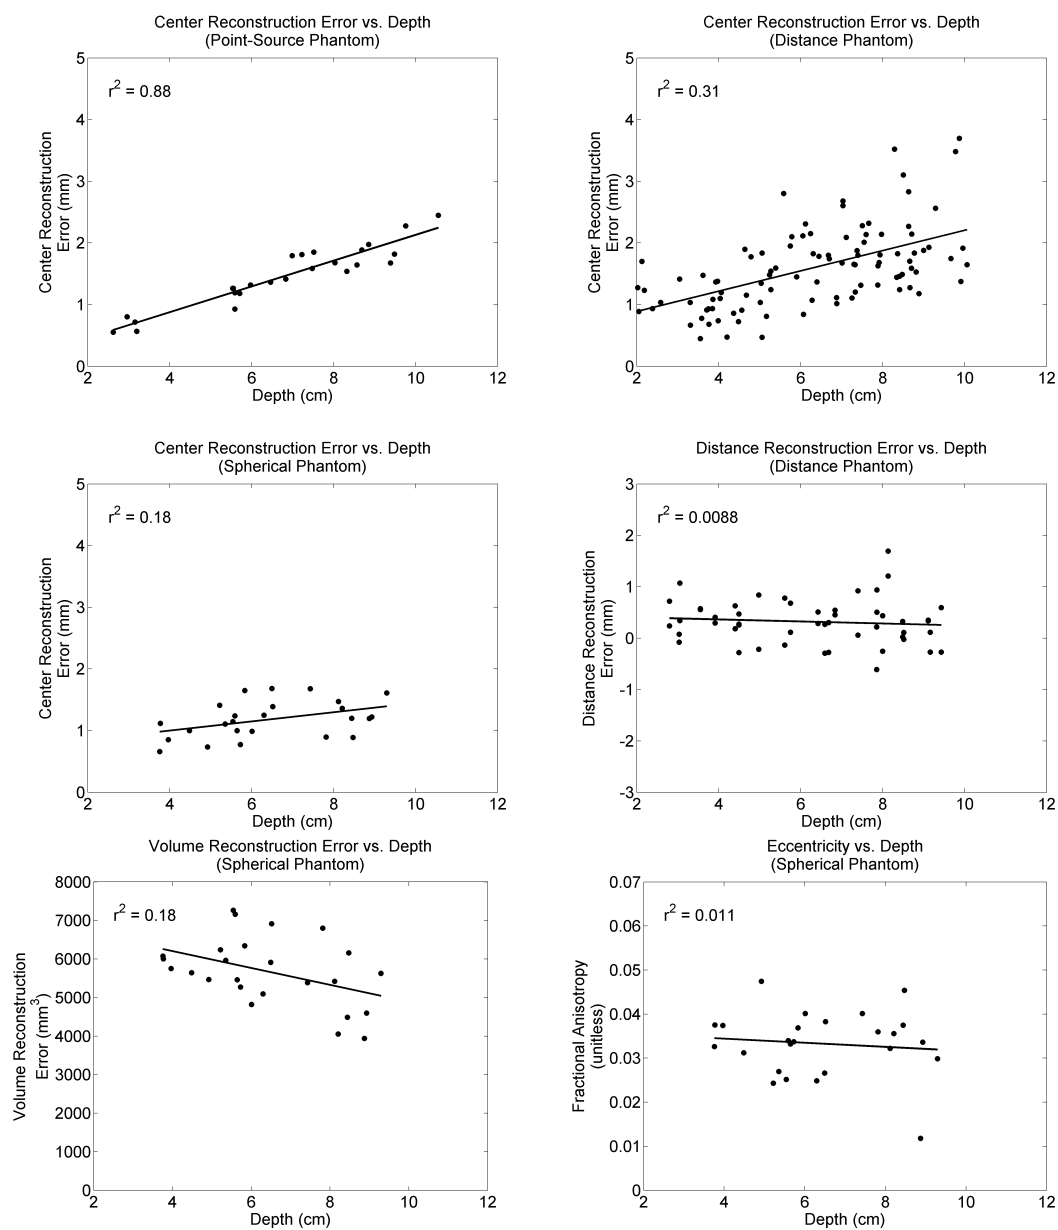


Fig. 3.14: Centre reconstruction error, distance reconstruction error, volume reconstruction error and fractional anisotropy for 3D ultrasound reconstruction versus distance from the ultrasound transducer.

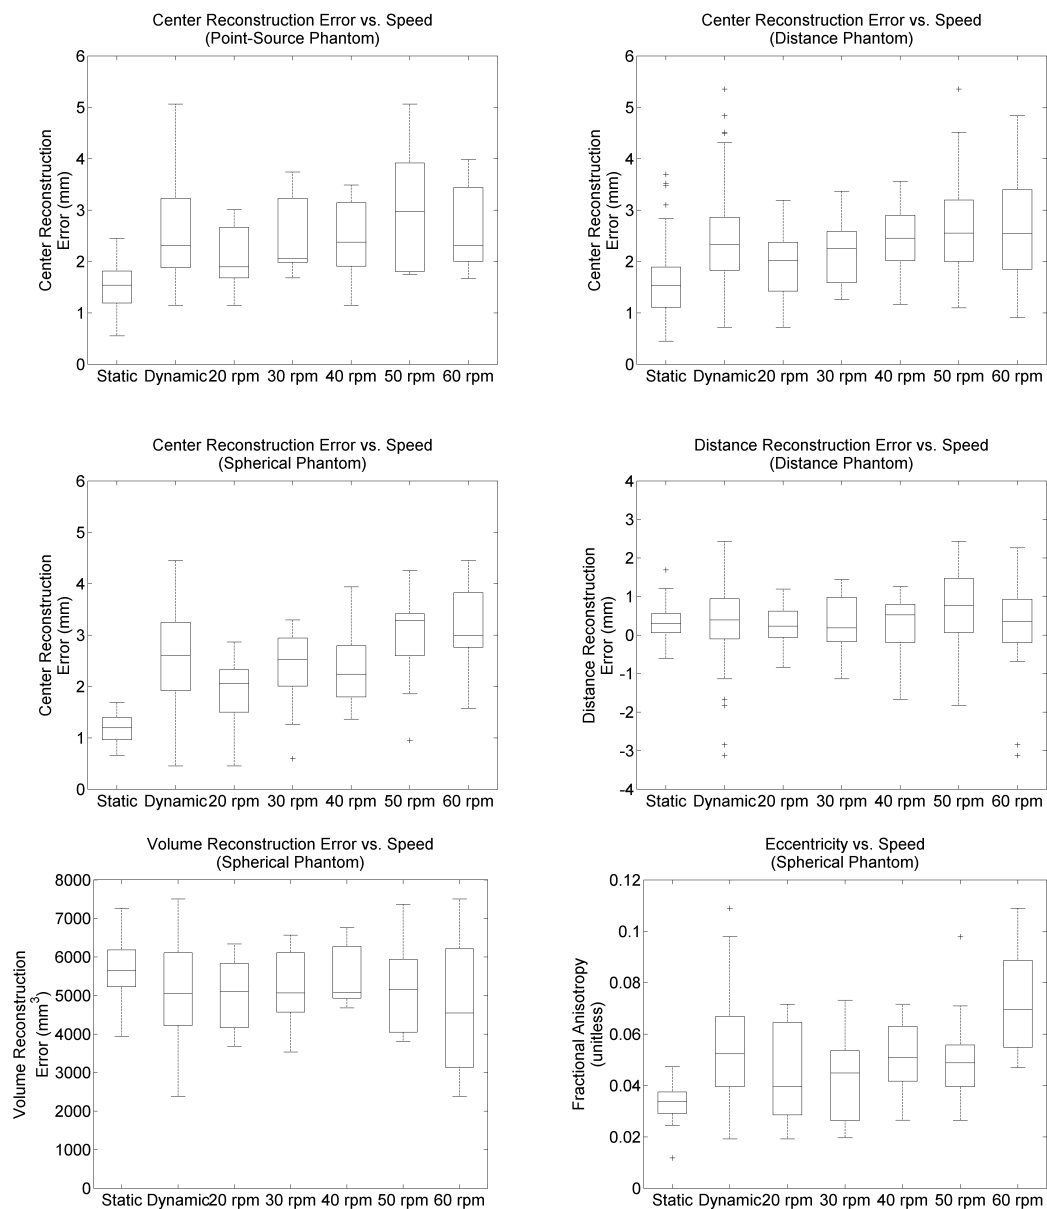


Fig. 3.15: Centre reconstruction error, distance reconstruction error, volume reconstruction error and fractional anisotropy for 4D ultrasound reconstruction versus phantom motor speed. Complete distributions for 3D ultrasound (“static”) and 4D ultrasound (“dynamic”) reconstruction errors are shown for comparison.

a gold standard derived from tracking information and the tool’s characterization. The errors measured intraoperatively were subsequently compared to those from an analogous laboratory study.

For this experiment, a magnetically-tracked painted aluminum pointer tool was fitted with two rubber rings to aid identification of the tooltip under ultrasound (Figure 3.16a). After epicardial imaging of the second pig, a miniature Universal Cardiac Introducer [217] was sutured to the left atrial appendage via the left minithoracotomy, providing direct access inside the beating heart. The pointer tool was suspended inside the left atrium within view of 2D TEE imaging and clamped to prevent tool movement (Figures 3.16b and 3.16c). Four 4D ultrasound datasets were reconstructed using different tool placements within the left atrium, where the distance from the tooltip to the US fan origin ranged between 4.1-5.7 cm. The stationary tool was imaged using a rotational ultrasound reconstruction while the animal’s respiration was suspended (Figure 3.16d), using retrospective gating to image at mid-systole (25% R-R) and mid-diastole (80% R-R) (US depth setting = 9 cm,  $s_{in} = 0.30$  mm,  $s_{out} = 0.59$  mm,  $HR_{exp} = 80-85$  bpm,  $\%dev_{HR} = 15\%$ ,  $N = 20$ ,  $P = \{5, 16\}$ ). Four 6-DOF sensors were magnetically tracked during each reconstruction: the reference sensor and the sensors attached to the pointer tool and the TEE and TTE probes. As in the dynamic phantom experiment, the contents of the tracker buffers for all tracked objects were saved after each reconstruction, as were the timestamps for each inserted 2D US image.

Intraoperative accuracy was assessed by comparing the observed tooltip location for each output volume to its expected position. The observed coordinate in each output volume was identified semi-automatically by calculating the centroid of all voxels belonging to the blob representing the tooltip. These voxels were defined as those whose intensities were greater than a specified threshold and that were connected to a seed point without exceeding a maximum distance from it. The seed point, threshold and maximum distance were manually defined for each output volume; note that the centroid coordinate was not usually very sensitive to these values. The identified tooltip coordinates  $p_{ref}$  are relative to the reference sensor, and were subsequently transformed into the ultrasound beam coordinate system, as described




Fig. 3.16: Components of the intraoperative accuracy assessment study. (a) The pointer tool equipped with two rubber rings; (b) The OR scene showing the pointer tool within the animal's beating heart; (c) An image of the pointer tool under 2D US; (d) The virtual view of an ultrasound reconstruction of the pointer tool within the AR environment; (e) The setup for the laboratory experiment.

for the dynamic phantom experiments:

$$p_{obs} = (C_{US})^{-1} (T_{US})^{-1} p_{ref} \quad (3.2)$$

As previously described,  $T_{US}$  represents the transforms returned from the tracking system for the sensor attached to the ultrasound probe, and  $C_{US}$  is the spatial calibration matrix. The gold standard tooltip coordinate was calculated by transforming the tooltip  $p_{tip} = [0, 0, 0]^T$  in the pointer tool's coordinate system, where the origin represents the tooltip and the positive  $\mathbf{z}$ -axis passes through the tool axis, into the ultrasound beam coordinate system:

$$p_{exp} = (C_{US})^{-1} (T_{US})^{-1} T_{tool} C_{tool} p_{tip} \quad (3.3)$$

Here,  $C_{tool}$  is the transformation between the tooltip and the tool's attached magnetic tracking sensor as determined by a pivot calibration procedure, and  $T_{tool}$  represents the transform matrices resulting from tracking the pointer tool relative to the reference sensor. Both  $p_{obs}$  and  $p_{exp}$  were calculated using every transform matrix acquired for the relevant tracked objects throughout the reconstruction, and were then set to the centroid of the resulting point cloud after removing outliers. Outlier removal was performed by removing those points within the point cloud whose distance to the median coordinate was outside the 90th percentile. Finally, the localization error was calculated as the Euclidean distance between  $p_{obs}$  and  $p_{exp}$ , as were the errors along each of the three axes relative to the ultrasound beam at  $0^\circ$  rotation.

An analogous experiment was performed within the laboratory for comparison to the intraoperative case. This experiment also evaluated the accuracy of rotational versus fan acquisitions for 3D US reconstruction and of tracked 2D ultrasound versus 3D US reconstruction. Unlike the intraoperative experiment, 3D US reconstruction was used because the laboratory setup was entirely static, however the laboratory results can be directly compared to the intraoperative results as the pointer tool was stationary in both cases. To mimic the speed of sound in tissue, all imaging was performed within a 10.6% glycerol solution by mass at room temperature (Figure 3.16e). Note that in the interval between the dynamic phantom experiment and the intraoperative accuracy experiment, an independent study within our laboratory showed




Fig. 3.17: Laboratory experiment using the tracked pointer tool. (a) Using the AR environment to position the 2D US probe to intersect with the tooltip; (b) Laboratory setup for rotational TEE acquisition; (c) Laboratory setup for fan TEE acquisition (2D US image plane is approximately perpendicular to probe motion).

that the speed of sound within a 10.6% glycerol solution is closer to 1540 m/s (the average speed of sound in tissue) than that within a 7% glycerol solution. The same four tools were magnetically tracked as in the intraoperative experiment: the reference sensor, the pointer tool, the TEE probe and the TTE probe. Although the TTE probe was not used in the laboratory experiment, it was plugged into the magnetic tracking system to maintain the same number of tracked tools as in the intraoperative case. The pointer tool was supplemented with a piece of sound absorbing material (Sorbothane Inc., Ohio, USA) to absorb any ultrasound waves reflecting off the tank walls, and imaging was performed using the 2D TEE transducer with the same spatial calibration matrix as in the intraoperative experiment. For each of twenty tool placements, a 2D US image was captured after using the augmented reality environment to position the probe and rotate the imaging plane so that it intersected with the pointer tool's tooltip (Figure 3.17a). Then, two 3D ultrasound reconstructions were performed as shown in Figures 3.17b and 3.17c, first using a rotational acquisition (at  $1^\circ$  increments) and second by manipulating the probe in a fan motion while the imaging plane was set at  $90^\circ$  (distance from tooltip to US fan origin ranged between 2.9-8.9 cm, US depth setting = 12 cm,  $s_{in} = 0.35$  mm,  $s_{out} = 0.69$  mm). As in the intraoperative study, the tooltip was semi-automatically identified in both the 2D US

images and the reconstructed 3D volumes. The resulting observed coordinates and the expected coordinates were then transformed into the ultrasound beam coordinate system for comparison.

For both the intraoperative and laboratory experiments, it was important to validate the assumption that the pointer tool remained stationary throughout the reconstructions. For each intraoperative trial, the expected tooltip coordinate was recalculated relative to the reference sensor ( $p_{exp} = T_{tool}C_{tool}p_{tip}$ ) twice, each time determining  $T_{tool}$  using only those transform matrices for the pointer tool corresponding to either the first or the second phase of interest. If the pointer tool moved rhythmically with the beating heart, then there would be a large distance between these two points. In actuality, this distance was less than 0.06 mm for all trials, indicating that the tool was not influenced by the cardiac cycle. In addition, for each intraoperative and laboratory trial a point cloud was created for the expected tooltip coordinate relative to the reference sensor by calculating it using every  $T_{tool}$  acquired from the tracking system. On average, the distance from each coordinate to the point cloud centroid was  $0.21 \pm 0.14$  mm for the intraoperative trials and  $0.05 \pm 0.02$  mm for the laboratory trials. The compactness of these point clouds demonstrates that pointer tool movement did not significantly corrupt the accuracy assessment for both the intraoperative and laboratory experiments.

Table 3.4 and Figure 3.18 show the results from the intraoperative and laboratory studies using the tracked pointer tool. The RMS error for the intraoperative datasets was very low, measuring 0.82 mm, and arises primarily from errors in the  $\mathbf{z}$ -direction followed by errors along the  $\mathbf{x}$ -axis. There was no meaningful difference in error between the output volumes reconstructed at mid-systole compared to those reconstructed at mid-diastole, which can be expected since the pointer tool was stationary. The error associated with performing 3D US reconstruction within the laboratory using the same rotational acquisition had an RMS of 1.55 mm and was significantly greater than the intraoperative error, as tested by two-sample, one-tailed t-tests without assuming equal population variances:  $p = 0.0114$  for rotational laboratory vs. intraoperative mid-systole and  $p = 0.0038$  for rotational laboratory vs. intraoperative mid-diastole (`ttest2`, Matlab, The MathWorks Inc., Natick, MA,



Table 3.4: Summary statistics for localization error in the intraoperative accuracy assessment and analogous laboratory study over  $N$  trials. The X, Y and Z directions correspond to the lateral, axial and elevation directions, respectively, of the  $0^\circ$  2D US image plane when the probe is at the mean position acquired over the reconstruction.

| Setting        | Image Type                                 | Direction | RMS  | Mean  | Std. Dev. | $N$ |
|----------------|--|-----------|------|-------|-----------|-----|
| Intraoperative | Mid-Systole<br>(Rotational)                | Euclidean | 0.82 | 0.74  | 0.40      | 4   |
|                |  | X         | 0.49 | 0.35  | 0.40      | 4   |
|                |  | Y         | 0.27 | 0.07  | 0.30      | 4   |
|                |  | Z         | 0.59 | -0.53 | 0.31      | 4   |
|                | Mid-Diastole<br>(Rotational)               | Euclidean | 0.81 | 0.77  | 0.30      | 4   |
|                |  | X         | 0.40 | 0.32  | 0.27      | 4   |
|                |  | Y         | 0.24 | 0.10  | 0.25      | 4   |
|                |  | Z         | 0.67 | -0.64 | 0.21      | 4   |
|                | Combined<br>Intraoperative<br>(Rotational) | Euclidean | 0.82 | 0.76  | 0.33      | 8   |
|                |  | X         | 0.45 | 0.34  | 0.32      | 8   |
|                |  | Y         | 0.25 | 0.08  | 0.26      | 8   |
|                |  | Z         | 0.63 | -0.59 | 0.25      | 8   |
| Laboratory     | Rotational 3D US                           | Euclidean | 1.55 | 1.49  | 0.44      | 20  |
|                |  | X         | 0.53 | 0.41  | 0.35      | 20  |
|                |  | Y         | 0.80 | -0.75 | 0.30      | 20  |
|                |  | Z         | 1.21 | -1.11 | 0.50      | 20  |
|                | Fan 3D US                                  | Euclidean | 1.62 | 1.56  | 0.45      | 20  |
|                |  | X         | 0.74 | 0.19  | 0.73      | 20  |
|                |  | Y         | 0.59 | -0.49 | 0.33      | 20  |
|                |  | Z         | 1.32 | -1.09 | 0.75      | 20  |
|                | Combined                                   | Euclidean | 1.58 | 1.52  | 0.44      | 40  |
|                |  | X         | 0.65 | 0.30  | 0.58      | 40  |
|                |  | Y         | 0.70 | -0.62 | 0.34      | 40  |
|                |  | Z         | 1.26 | -1.10 | 0.63      | 40  |
| Laboratory     | Tracked 2D US                              | Euclidean | 1.52 | 1.49  | 0.30      | 20  |
|                |  | X         | 0.48 | 0.08  | 0.49      | 20  |
|                |  | Y         | 1.02 | -1.01 | 0.17      | 20  |
|                |  | Z         | 1.01 | -0.81 | 0.62      | 20  |

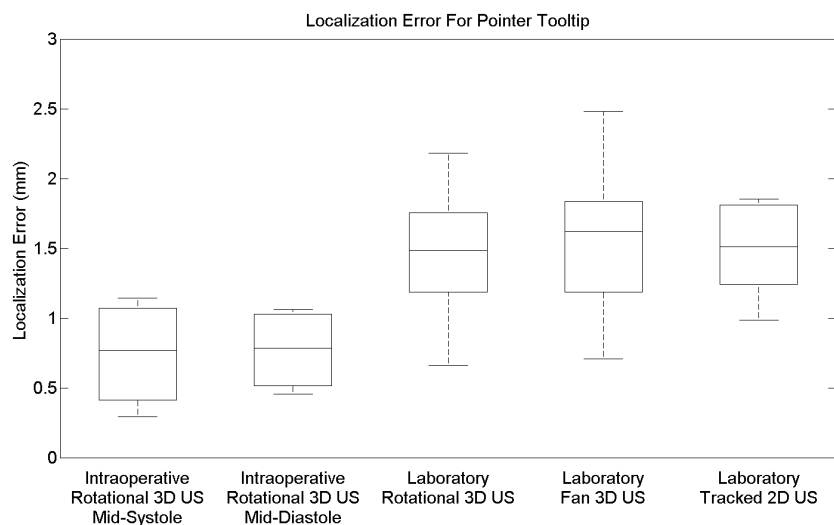


Fig. 3.18: Boxplot for the localization errors in the intraoperative accuracy assessment and the analogous laboratory study.

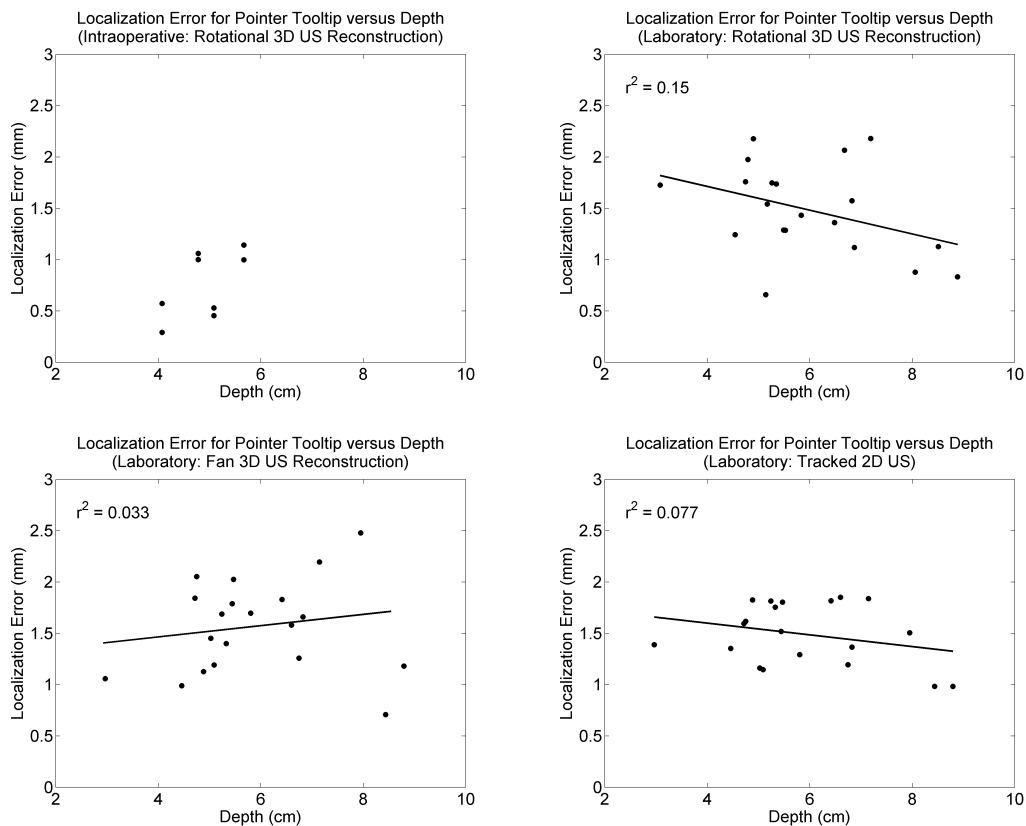


Fig. 3.19: Localization error in the intraoperative accuracy assessment and the analogous laboratory study versus distance from the ultrasound transducer.

USA). However, the RMS rotational laboratory error does parallel the laboratory results for 3D US reconstruction of the dynamic phantom (RMS centre reconstruction error = 1.22/1.54/1.72 mm). Compared to the intraoperative trials, there was an increase both in the mean error, mostly caused by larger errors in the  $\mathbf{y}$ - (depth) and  $\mathbf{z}$ -directions, and in the error variance, which can primarily be attributed to a decreased precision in localizing the tooltip along the  $\mathbf{z}$ -axis.

In the laboratory, comparisons were also made between rotational and fan acquisitions for 3D US reconstruction and between 3D US reconstruction and tracked 2D ultrasound. First, the laboratory results show slightly decreased mean accuracy for the fan technique compared to the rotational technique, mostly caused by increased variance in the  $\mathbf{x}$ -direction. Visually, ultrasound volumes acquired using the rotational technique were significantly less blurred than those generated using a fan acquisition, and could for example distinguish the rubber rings from the tool's metallic shaft, which the fan acquisitions could not. There was no notable difference between the mean errors associated with 3D US reconstruction and with tracked 2D ultrasound, although the errors for tracked 2D ultrasound were slightly less variable. For tracked 2D ultrasound, errors were primarily along the  $\mathbf{y}$ - and  $\mathbf{z}$ -axes. For both intraoperative and laboratory reconstructions, the tooltip position along the  $\mathbf{z}$ -axis was typically underestimated and contributed the most to the overall error, while the tooltip position along the  $\mathbf{x}$ -axis was usually overestimated. Error in the  $\mathbf{y}$ -direction was minimal intraoperatively, but suffered from underestimation in the laboratory.

Finally, Figure 3.19 plots the tooltip localization error versus depth. Trends can be extracted for the laboratory studies, but there are too few intraoperative trials to do so with any confidence. Error decreases with depth for both 3D US reconstruction with a rotational approach and tracked 2D US, which can be attributed clearly to error reductions in the  $\mathbf{y}$ -direction and somewhat to decreases along the  $\mathbf{z}$ -direction (plots not shown). For 3D US reconstruction using a fan acquisition, error increases with increasing depth, primarily caused by larger errors along the  $\mathbf{x}$ -axis (plot not shown).

### 3.5 Clinical Imaging

Finally, a preliminary study with human subjects was performed to evaluate the real-time 4D ultrasound reconstruction system within a true clinical environment. Reconstructions were acquired both before and after left lung collapse and insufflation in two patients who were undergoing routine TEE imaging as part of a robotic coronary artery bypass graft (CABG) procedure. All imaging was performed according to the policies of The University of Western Ontario’s Research Ethics Board. Both patients were undergoing perioperative procedures during ultrasound reconstruction resulting in significant patient motion, and therefore this clinical component represents the most severe test of the ultrasound reconstruction system.

The clinical setup mirrored that previously described for porcine imaging. For patient safety, imaging was performed using the integrated TEE transducer and with respiratory suspension limited to 2-3 minute blocks as dictated by the clinician performing the ultrasound reconstruction. The magnetic tracking system’s field generator was mounted on the underside of the OR table and tracking was performed relative to a 6-DOF patient sensor (Northern Digital Inc., Waterloo, ON, Canada) securely taped to the patient’s back. Two 6-DOF sensors were therefore tracked during these studies: the reference sensor and the sensor within the integrated TEE probe used for humans. Retrospective gating was used to reconstruct two output volumes per 4D dataset corresponding to mid-systole and mid-diastole, once again based on measurements derived from imaging the mitral valve under M-mode ultrasound. ECG-gating was complicated by patient arrhythmias and interference induced by the magnetic tracking system in the non-shielded ECG leads. To simplify R-wave detection under these conditions, the “pacemaker” functionality of the anesthesiology equipment was used to adjust the patient’s ECG signal to produce a more typical output. Practical constraints prevented us from mounting a monitor providing real-time visualization of the reconstruction’s progress close to the OR table, and so the clinician viewed the workstation’s monitor located some distance away. As before, the technician operating the reconstruction software also viewed the reconstruction’s progress and gave verbal feedback to the clinician.




Fig. 3.20: Example clinical 4D ultrasound reconstruction at the 15% (left) and 65% (right) R-R intervals, acquired using the TEE probe with a rotational acquisition before left lung collapse. (Patient 1: US depth setting = 15 cm,  $s_{in} = 0.43$  mm,  $s_{out} = 0.81$  mm, output dimensions =  $433 \times 354 \times 478$ ,  $HR_{exp} = 48$  bpm,  $\%dev_{HR} = 10\%$ ,  $N = 20$ ,  $P = \{3, 13\}$ ).





Fig. 3.21: Example clinical 4D ultrasound reconstruction at the 15% (left) and 70% (right) R-R intervals, acquired using the TEE probe with a pullback acquisition after left lung collapse and insufflation. (Patient 2: US depth setting = 14 cm,  $s_{in} = 0.40$  mm,  $s_{out} = 0.81$  mm, output dimensions =  $437 \times 388 \times 398$ ,  $HR_{exp} = 48$  bpm,  $\%dev_{HR} = 27\%$  as erratic heart rate prevented reconstructing with strict HR limits,  $N = 20$ ,  $P = \{3, 14\}$ ).

4D ultrasound reconstruction was performed both using a rotational acquisition (at 2° increments to reduce acquisition time) and using the freehand pullback (linear) method (Figures 3.20 and 3.21). Freehand transthoracic imaging from between the patient’s ribs was also attempted, but this technique was very cumbersome and produced output volumes of low quality. The initial clinical imaging was relatively successful, but the human subject output volumes contained more stitch artifact compared to the reconstructed porcine images. This was due to significant patient motion, increased time pressure (all reconstructions were finished in 1.75 to 2.75 minutes), decreased gating accuracy when relying on the erratic ECG signals of patients with arrhythmias, and reduced tracking accuracy caused by the integrated probe and potentially also by more substantial magnetic interference within the patient OR.

## 3.6 Discussion

### 3.6.1 Human factors benefits of real-time reconstruction

Although several real-time systems have been developed [196,204–207], user studies actually evaluating the perceived benefits of providing real-time visualization during ultrasound reconstruction have not been previously reported. Such human factors experiments would ideally evaluate clinicians as they performed the actual task of imaging human subjects, however time limitations within the operating room made this impractical and a phantom experiment was performed instead.

In a similar study, Brekke *et al.* [226] examined the benefits of real-time visualization as applied to incremental volume rendering during extended field of view RT3D TTE, a task which is usually performed by commercial ultrasound systems without any such visualization. This group compared both image quality and acquisition time for “wide-angle” RT3D volumes acquired with and without real-time visualization, measuring the former by calculating the correlation between the two facing 2D planes at each interface between stitched RT3D volumes.

My human factors study showed only minor improvements in the 3D output volumes reconstructed when real-time visualization was shown to the clinician compared

to when it was not. Similarly, Brekke *et al.* found that no significant improvement in image quality resulted from providing real-time visualization. Both of these results can be attributed to the fact that these acquisition tasks are sufficiently simple for users to complete them satisfactorily without any visual aid.

In contrast, 4D ultrasound reconstruction is technically challenging and benefited greatly from real-time visualization. Both my human factors results and observations of clinicians working with the system in the OR indicate that a significant reduction in acquisition errors follows from providing an incremental visualization of the reconstruction results. In particular, guidance was especially valuable in preventing large gaps in the output volumes. Typical algorithms would either leave such gaps blank or would interpolate them using data from filled voxels regardless of their proximity, which does not necessarily regenerate the correct intensities within the gap regions. In contrast, 4D US datasets reconstructed under real-time visualization generally do not contain such missing or unreliable data, and will therefore be more suitable for subsequent visualization or registration during image-guided therapy.

The observed increase in acquisition time associated with real-time visualization can be attributed to the time taken to review reconstructions and to correct acquisition errors. Although the chance of significant patient motion or cardiac cycle variability does increase with acquisition duration, this risk is likely to be an acceptable trade-off for most applications for the benefits of more densely sampled output volumes. Brekke *et al.* found no significant difference in mean acquisition time between the two visualization conditions for their RT3D system, although they measured acquisition time in a more clinical setting as the amount of time taken to perform patient scans deemed to be acceptable by the echocardiographer upon subsequent review. In the human factors study presented here, achieving an “acceptable” 4D ultrasound reconstruction of the beating heart phantom was practically impossible in the absence of real-time visualization, and so a similar evaluation was not attempted. It is anticipated that the real-time visualization condition would prove superior for 4D ultrasound reconstruction had this alternative scheme for measuring acquisition time been followed.



### 3.6.2 Quantitative accuracy assessments

Two quantitative validation studies evaluated the real-time 4D ultrasound reconstruction system: the first used a dynamic phantom to study the accuracy of 3D and 4D ultrasound reconstruction with respect to target localization and conservation of distance, volume and shape, while the second assessed target localization of a static object under both intraoperative and laboratory conditions. Both evaluations used magnetic tracking information to calculate the gold standard and are therefore subject to the implicit assumption that tracking error is minimal. This assumption is not true in practice, but is somewhat unavoidable when performing such assessments as few alternative strategies exist for determining gold standards.

#### 3.6.2.1 Dynamic phantom accuracy assessment

The RMS localization errors from the dynamic phantom experiments for 3D US reconstruction were between 1.2-1.7 mm, while the 4D ultrasound reconstruction RMS localization errors were between 2.5-2.7 mm. Both these 3D and 4D US reconstruction errors are reasonable within the clinical context of minimally-invasive therapy performed within the beating heart, especially considering that the system accuracy is limited by that of the magnetic tracking system (specified by the manufacturer to have RMS accuracy of 0.6 mm in position and 0.4° in orientation). Distance reconstruction errors were also small, indicating that the relationships between objects are well-preserved despite any errors in their gross localization. The relatively large volume reconstruction errors do indicate that caution should be taken when performing intracardiac volume measurements. Because the radii length errors were passable for the spherical phantom, the large volume reconstruction errors arise from the cubed relationship between radii length and volume. However, the large volume reconstruction errors may be partially caused by the procedure used to manually outline the spherical phantom within the output volumes. This is supported by the fact that the radii length errors for the spherical phantom (RMS: 3D 0.80-1.47 mm / 4D 0.60-1.57 mm) were generally larger than the Euclidean distance reconstruction errors for the distance phantom (RMS: 3D 0.54 mm / 4D 0.97 mm) even though the gold standard

radii and distances are similar. Finally, the eccentricity statistics characterizing the conservation of shape were excellent for both 3D and 4D US reconstruction. Since the stitch artifacts within the 4D ultrasound reconstructions of the dynamic phantom appear similar to those present when imaging the beating porcine or human heart, the quantitative results presented for 4D US reconstruction may be a good approximation for the errors present clinically.

The accuracy statistics for 3D US reconstruction of the dynamic phantom are better than those presented in our group's previous assessment of tracked 2D ultrasound [214], which found an RMS localization error of 2.4 mm, although the same adult transesophageal probe and spatial calibration procedure and a similar phantom design were used. This may illustrate the advantages of spatial compounding, but a corresponding improvement of 3D US reconstruction compared to tracked 2D imaging was not replicated in the laboratory experiments performed using the tracked pointer tool (see below). Since the previous tracked 2D US study, the gain and compression settings for ultrasound imaging were optimized to sharpen the edges defining the phantom boundaries, and the phantom's characterization was improved by using a micro-CT image rather than digitizing the phantom surface with an optically-tracked pointer tool.

The dynamic phantom assessment showed that 4D ultrasound reconstruction is generally associated with an increase in error magnitude and variability compared to 3D ultrasound reconstruction, which can be attributed to several error sources. A significant problem is presented by the limited capture rate provided by the frame grabber, which causes an inadequate sampling of moving structures especially at high speed. A timestamp corresponding to the beginning of each phase of interest can be calculated with relatively high accuracy using retrospective gating, but there is a maximum discrepancy of approximately 15 ms between this time and the closest timestamp from a 2D US image in the video stream. In this worst-case scenario, the phantom will move approximately 0.4 mm at 20 rpm and 1.26 mm at 60 rpm within this interval. As discussed in [227], stitch artifacts result when the 2D US images making up each output volume do not image the heart at the exact same cardiac phase. A second source of error in the dynamic phantom experiments derives

from the use of prospective gating. Although a strict 1.5% allowed deviation from the expected heart rate was applied, prospective gating is not as accurate as retrospective gating in the common situation where a given cardiac cycle does not have the same duration as its predecessor. Finally, the motion artifacts present in the 4D ultrasound datasets increased the difficulty of performing manual segmentation, compared to that for output volumes generated using 3D ultrasound reconstruction. The stitch artifacts within the output volumes derived from 4D ultrasound reconstruction make developing an automatic segmentation method very difficult, and so manual segmentation was used for both the 3D and 4D ultrasound datasets for the sake of consistency.

Several accuracy measures showed a directional dependence when analyzed along the  $\mathbf{x}$ -,  $\mathbf{y}$ - and  $\mathbf{z}$ -axes of each output volume's coordinate system. The decreased accuracy of 4D US reconstruction compared to 3D US reconstruction was generally attributed to increased errors along the  $\mathbf{y}$ - and  $\mathbf{z}$ -axes, which are the two directions along which phantom movement was greatest and therefore the directions in which motion artifacts will be most significant. Localization error also increased with distance from the ultrasound transducer, partially because of increasing separation between adjacent 2D US images with increasing depth for rotational reconstructions. In addition, partial volume effects caused by increased beam width are stronger away from the beam focus, causing blur in 2D US imaging and reducing the precision with which objects can be identified. Partial volume effects are also a likely contributor to the observed decrease in volume reconstruction error as a function of depth. Since the outline of the sphere was more blurry at increased distance from the ultrasound transducer, the strategy of segmenting the edge closest to the fan origin will decrease the observed sphere size, reducing the degree to which it was typically overestimated. The observed tendency to overestimate distances in the distance phantom and radii length in the spherical phantom is consistent with our group's previous finding that area is overestimated under 2D ultrasound [214]. For the spherical object, this may be caused by the technique employed for manual segmentation or by the imaging properties of the table tennis ball, but is also consistent with a slower speed of sound within a 7% glycerol solution compared to the speed of sound assumed by the ultra-

sound machine. This hypothesis is also supported by the bias towards overestimating the distance between objects and the US probe along the depth direction, and was later confirmed in a separate study within our group, which recommended a 10.6% glycerol solution as used in the later studies presented here.

### **3.6.2.2 Intraoperative accuracy assessment**

The second accuracy assessment aimed at evaluating US reconstruction accuracy under clinically relevant conditions, in this case by imaging a stationary pointer tool within the beating heart. The intraoperative experiment designed here focused on examining environmental factors such as US imaging quality (including the associated speed of sound considerations) and tracking error, not specifically on the artifacts associated with imaging a moving object. In future, a dynamic assessment should be performed within the beating heart using tracked moving targets, although the practical challenges involved in designing such an experiment are not insignificant.

The pointer tool assessment gave RMS localization errors of approximately 0.8 mm within an intraoperative setting and 1.6 mm in the laboratory. The latter measurements are similar to the 1.2-1.7 mm RMS errors previously determined for performing 3D US reconstructions of the dynamic phantom within the laboratory. Of note, the Euclidean localization error for 3D US reconstructions of the point-source phantom ( $1.46 \pm 0.50$ , RMS 1.54) is almost identical to that for similarly performing a rotational 3D US reconstruction of the tracked pointer tool within the laboratory ( $1.49 \pm 0.44$ , RMS 1.55). However, direct comparisons between the results from the dynamic and pointer tool studies are impaired by differences in the objects' imaging properties under ultrasound and in the methods used for their characterization.

It was originally hypothesized that tracking error sustained within the operating room caused by nearby ferromagnetic equipment would increase the error for intraoperative ultrasound reconstruction compared to reconstructions performed within the laboratory. The observation that accuracy was in fact better within the operating room was therefore unexpected, although both RMS errors are within a reasonable range for clinical use. The duration of the laboratory experiment was not constrained

as it was in the operating room, and so the more numerous laboratory trials ( $N = 20$  versus  $N = 4$ ) explored a wider variety of pointer tool configurations than could be tested intraoperatively. This may have exposed error-prone configurations that were not tested within the operating room. Additional sources of error sustained within the laboratory include stronger ultrasound reflections of the pointer tool that made reliable identification of the tooltip more difficult, as well as the increased US depth setting used. The very low errors realized along the depth ( $\mathbf{y}$ -) direction for the intraoperative trials indicate that the speed of sound through the porcine heart was well matched to that assumed by the ultrasound machine while the tooltip location in the depth direction was typically underestimated in the laboratory trials, potentially indicating an accelerated speed of sound in the 10.6% glycerol solution.

Finally, the direction of greatest error for the fan reconstructions performed within the laboratory (the  $\mathbf{x}$ - direction) corresponds to the direction of probe movement and to the ultrasound beam's elevation direction (when the multiplanar TEE probe is set to  $90^\circ$  as done here), both of which contributed to increased error. Partial volume effects caused by the ultrasound beam's increased slice thickness with depth also explains the observed depth dependence for error along the  $\mathbf{x}$ - direction in the laboratory fan reconstructions.

As described by Rohling *et al.* [72], spatial compounding in 3D ultrasound reconstruction can reduce ultrasound speckle and increase the signal to noise ratio, although the errors associated with positioning each individual 2D US image within the output volume cause blurring. However, when attempting to manually localize an object such as the tooltip or an intracardiac structure, using a blurred 3D US image may be preferable to an individual 2D US image whose imaging plane may not intersect directly with the object of interest, or whose associated transform matrix may be corrupted by tracking error. However, spatial compounding did not have a considerable effect in the laboratory experiments performed using the tracked pointer tool, as there was no notable difference between the errors associated with 3D ultrasound reconstruction and with tracked 2D ultrasound. The results shown here indicate that, at least within a laboratory setting, 3D US reconstruction provides accuracy similar to that of tracked 2D ultrasound while generating a much larger amount of image

data, with inaccuracies dominated by errors in the spatial calibration, magnetic tracking and the fidelity with which objects can be identified under ultrasound imaging. However, this claim should be re-evaluated within a more clinically-realistic setting than a laboratory experiment can provide.

### 3.6.2.3 Comparison to results from assessments by other groups

The results from the 3D ultrasound reconstruction assessments can be examined relative to previous evaluations performed by other groups to compare the accuracy of the spatial and temporal calibration strategies described here. However, the usefulness of such comparisons should not be exaggerated because of considerable differences in experimental methodology, phantom design and the ultrasound transducer and imaging settings used. For the evaluations described below, please note that optical tracking, which has a greater accuracy than magnetic tracking, was used in [194, 195, 197, 228, 229]. In addition, an automatic segmentation method was used in [194, 197, 200], reducing the imprecision associated with manual segmentation as performed for the dynamic phantom experiments.

The mean localization errors reported for 3D reconstructions of the dynamic phantom and the tracked pointer tool ranged between 0.8-1.6 mm, and are in range of the mean errors reported by other researchers, specifically 0.8-1.2 mm (Lindseth *et al.* [194], FPA/FLA transducers and diagonal/pyramidal phantoms), 1.2 mm (Blackall *et al.* [228]) and 0.6 mm (Muratore and Galloway [229]). These results are also better than the means of 2.3-3.2 mm reported by Pagoulatos *et al.* [230].

The distance errors calculated for the distance phantom ( $0.3 \pm 0.4$  mm, or  $1.6 \pm 2.2\%$ ) show a slight bias towards overestimating distance compared to the  $0.03 \pm 0.7\%$  given by Blackall *et al.* [228] and the  $-0.5 \pm 3.4\%$  measured by Barratt *et al.* [219]. The errors from the experiments presented in this thesis are very similar to the  $0.1-0.3 \pm 0.3-0.5$  mm reported by Lindseth *et al.* [194] and are generally better than those found by Rousseau *et al.* [195] ( $2.9-6.9\% \pm 1.5-3.4\%$ ) and Pagoulatos *et al.* [230] ( $-0.1-0.3 \pm 1.3-2.9$  mm, when the true intertarget distance was between 1-9 cm).

The RMS volume error measured here for 3D US reconstruction of the spherical phantom was approximately 18%. This is much larger than the errors measured by other groups by scanning a water-filled balloon and calculating the gold standard volume based on its mass. For example, Berg *et al.* [201] found volume errors ranging between -1.4% and 2.6% and Barry *et al.* [200] reported an RMS volume error of 1.1%. However, it is somewhat closer to the mean 92-105% of the gold standard volume measured by Rousseau *et al.* [195] based on manual segmentations of their egg-shaped phantom. Finally, few ultrasound reconstruction systems have been evaluated with respect to conservation of shape. Although Rousseau *et al.* [195] do evaluate shape, their metric is very different from that presented here and so a direct comparison would be meaningless.

### 3.6.3 Porcine and human imaging

The results of the porcine and human subject imaging sessions demonstrate the promise of real-time 4D ultrasound reconstruction for intraoperative imaging during image-guided cardiac therapy. Despite the fast heart rate and lower-quality 2D ultrasound imaging in swine compared to humans, the porcine 4D ultrasound datasets clearly show several important intracardiac structures and represent the reconstruction quality that can be reasonably expected from the real-time ultrasound reconstruction system at this point in its development.

The patient studies represent our group's first experience with ultrasound reconstruction in humans. Although patient imaging was somewhat successful, the reduction in reconstruction quality compared to that of the porcine images emphasizes the additional challenges that must be overcome in translational research to develop systems suitable for patient use, as well as the need to validate IGT systems with patients in addition to animals and/or healthy volunteers. The quality of patient imaging would be improved substantially by addressing the cardiac gating error arising from erratic patient ECG signals, reducing the additional tracking error sustained when using the integrated TEE probe, and preventing patient motion and rushed imaging by performing the ultrasound reconstruction within an environment

where it was prioritized as a formal part of the surgical workflow. The advantages of real-time 4D ultrasound reconstruction are highlighted under such challenging conditions. More specifically, the real-time visualization allowed the clinician to fill in gaps caused by patient motion or by improper cardiac cycle gating, which can cause a series of 2D US images to be rejected if the measured heart rate falls beyond the acceptable range.

In general, the rotational acquisition approach was preferred in both swine and human subjects, agreeing with the qualitative results from the pointer tool laboratory assessment. This can be attributed to the reduced tracking error associated with a stationary ultrasound probe compared to one that is moving, albeit slowly. Even for a purely rotational acquisition, real-time visual feedback of the reconstruction results is useful for gap prevention because the timing of each US imaging plane rotation is a manual process in the current system. Although the risk of gaps would be significantly reduced if the imaging plane was rotated automatically at the beginning of each cardiac cycle detected by ECG-gating, real-time reconstruction would still be useful to monitor probe positioning and to eliminate the time required for offline ultrasound reconstruction. It is also important to note that the entire heart cannot be captured within the cone-shaped volume generated by a rotational acquisition, and so some structures, especially those distal to the ultrasound probe such as the aortic valve, may be better captured using a focused linear or fan reconstruction.



# Chapter 4

## Future Research Directions and Conclusions

### 4.1 Improvements and Future Work

When developing the 4D ultrasound reconstruction system for intraoperative imaging during image-guided intracardiac interventions, the initial focus was on a real-time implementation guiding clinicians to reconstruct improved output volumes and on validating the quality of 3D and 4D ultrasound reconstructions quantitatively and within clinically realistic environments. As elaborated upon in Chapter 2's discussion, the resulting time series of reconstructed ultrasound volumes can serve various purposes in image-guided therapy. Investigating the use of reconstructed 4D ultrasound datasets for both visualization and registration within our group's augmented reality environment for minimally invasive beating-heart interventions constitutes two substantial items for future work, as are studies evaluating the efficacy of reconstructed 4D US datasets for these two tasks compared with tracked 2D US (as used currently by our group) or RT3D TEE.

### 4.1.1 Disadvantages compared to offline US reconstruction

Although I show that real-time ultrasound reconstruction generates superior 4D ultrasound datasets in the studies presented here, there are some disadvantages compared to offline ultrasound reconstruction approaches. However, most of these relate to the memory and processing speed required for real-time reconstruction, both of which can be improved by using a better machine or by future anticipated improvements in computer hardware. Ultrasound reconstruction also lends itself to a GPU implementation [203], which could vastly speed slice insertion and therefore address some of the following considerations.

As previously described, limitations in memory and processing speed restrict the number and size of the output volumes making up 4D ultrasound datasets reconstructed in real-time. Although the former has obvious implications for registration to preoperative 4D MR/CT datasets, registration based on a single timepoint in the cardiac cycle may be acceptable, for rigid registration at least, assuming that the heart's motion remains relatively periodic throughout each cardiac cycle [59–61].

Secondly, one must specify the output volume sizes and the phases of interest in advance when using a real-time ultrasound reconstruction system. This is not true for offline ultrasound reconstruction algorithms because all acquired 2D US images are retained and so the extent of each output volume can be appropriately calculated and an output volume can be generated for any desired cardiac phase. However, the phases of interest are generally known beforehand for our group's clinical application, either from prior knowledge of appropriate phases to visualize or according to those corresponding to the preoperative image(s) available for registration. Secondly, appropriate output volume sizes were easily estimated using the ultrasound reconstruction software's interactive parameter settings. Output volume expansion can be performed on the fly in real-time ultrasound reconstruction systems if an input 2D US image extends beyond the allocated output volume extent, but this uses additional computational resources and was not deemed to be necessary here.

The requirement for fast slice insertion means that an expensive reconstruction kernel cannot necessarily be used in real-time US reconstruction, even though such

kernels may reduce output volume artifacts [231]. More sophisticated kernels include those that consider the ultrasound imaging plane’s depth-dependent slice thickness or incorporate models of ultrasound’s point spread function [198], carrying obvious advantages compared to the assumption employed here that the imaging plane is infinitely thin. A possible extension to the real-time 4D US reconstruction system is to follow the suggestion proposed by Solberg *et al.* [198] and use the quick interpolation method described here for real-time visualization during image acquisition and afterwards perform a second offline reconstruction using a more advanced reconstruction kernel. Although this strategy may improve reconstruction quality, it does prohibit the use of the reconstructed datasets directly after their acquisition.

#### 4.1.2 Reducing reconstruction artifacts

Major sources of error in the current system causing reconstruction artifacts include the limited US video frame rate, tracking errors in the integrated probe used for human imaging and errors associated with ECG-gating. The 2D US images comprising each output volume could be input into an artifact reduction algorithm based on image registration that better aligns the images with respect to each other, similarly to the algorithm described by Treece *et al.* [62] to correct for probe pressure artifacts, but first each error source should be minimized as much as possible.

The limited 2D US frame rate available from the frame grabber constitutes a significant source of error, making it impossible to extract 2D US images corresponding exactly to the timestamps predicted by cardiac gating [227]. This can be rectified by reconstructing with raw ultrasound data instead, as demonstrated by Berg *et al.* [201], however for our research group and many others this is currently impossible as access to these data are restricted. A second potential improvement would be to trigger video acquisitions based on prospective cardiac gating instead of using a constant video digitizing rate [227]. However, this scheme relies on accurate prospective cardiac gating which cannot be guaranteed given variations in patient heart rate.

Attempts are currently underway in our laboratory to improve the tracking accuracy of the integrated ultrasound probe used clinically by adjusting the position of the

internal magnetic tracking sensor to reduce the influence of the surrounding copper beryllium metal [181]. Tracking imprecision can also be reduced for the tracked 2D ultrasound transducer (or for any other tracked tool as well) by applying a sliding averaging window to the transformations retrieved from the tracking system.

Finally, although its use is exceedingly common in commercial and research ultrasound systems, ECG-gating poses major problems for use in patients with arrhythmias or presenting with erratic ECG signals. Incorrect gating causes artifacts in the reconstructed volumes as many of the 2D US images may not correspond to the same temporal snapshot of the heart. A first step towards improving ECG gating is to interpolate the gating timestamps for systole and diastole separately, as the duration of these cardiac phases scales non-linearly with cardiac cycle duration [52]. Also, current induced within the ECG leads from the magnetic tracking system posed problems during human imaging and so shielded ECG leads should be used in future. Another alternative is image-based cardiac gating [101, 197], which has the potential to rectify the problems of ECG-gating by providing a truer representation of the heart's configuration throughout the cardiac cycle than the ECG signal. However, the use of image-based gating for real-time 4D ultrasound reconstruction is complicated by the requirement for real-time processing and especially by the continuous changes in the US imaging plane's viewpoint.

We found that suspending respiration was ideal during ultrasound reconstruction when possible, while the alternating respiration paradigm produced similar results in the porcine imaging experiments and is more suitable when imaging humans. Suspending respiration for up to 2-3 minutes is fairly standard clinical practice during several surgical procedures. However, as described in Chapter 1 many ultrasound-based systems integrate respiratory gating [85, 101, 113] or respiratory motion compensation [117], which could be added to the real-time 4D US reconstruction system to limit the reliance on respiratory suspension.

Finally, compounding several 3D output volumes corresponding to the same time-point in the cardiac cycle could reduce the appearance of artifacts, improve the signal to noise ratio and/or extend the field of view, as has been achieved for freehand ultrasound reconstruction [54, 81] and for RT3D US [53, 78]. From the experiments

described here, rotational imaging often gives superior results when imaging some cardiac chambers and the mitral valve, while distal structures such as the aortic valve may be better imaged using a focused translational or fan approach. Compounding may provide an ideal mechanism to combine two (or more) such acquisitions. Similarly, several rotational acquisitions from different viewpoints within the esophagus could be combined, as could reconstructions using a transesophageal and a transthoracic transducer. Since output volumes are always reconstructed in the AR environment's coordinate system and because probe pressure artifacts are minimal for the TEE probe, these compounded datasets could be created without image registration by simply averaging the intensities of voxels corresponding to the same 3D coordinate (while of course employing an appropriate interpolation scheme).

Finally, when comparing the reconstructed ultrasound images presented here to displays of RT3D US images, it should be noted that RT3D US images are subject to a heavy amount of post-processing and smoothing before they are displayed. This improves visual comprehension for diagnosis, but it is unknown how much these displays deviate from faithfully representing intracardiac structures as required in image-guided therapy. The reconstructed ultrasound volumes presented here are subject to artifacts but are depicted in their native form. Subsequent smoothing, whether by a Gaussian, median or anisotropic diffusion filter, could be employed in future to reduce the appearance of artifacts provided that such processing does not interfere with the images' intended use.

## 4.2 Additional Applications

The real-time 4D ultrasound reconstruction system is also applicable for any application requiring the acquisition of many tracked 2D US images, whether such a sparse 2D US dataset is intended for ultrasound reconstruction or for direct use in other tasks such as registration to preoperative images as in [63, 66, 70, 85, 141, 232]. Providing real-time feedback visualizing the set of 2D US images as they are acquired may be especially valuable if the acquisition process is difficult, for example when acquiring sets of respiratory-gated 2D US images of the liver [63].

Although not specifically examined in this paper, one of the most promising applications of the real-time 4D ultrasound reconstruction system may be in 2D ICE imaging during electrophysiological catheter ablation procedures. Tracked 2D ICE is currently gaining attention as a source of endocardial surface contours useful in registering the patient with a 3D surface model derived from a preoperative MR or CT scan, without relying on the time-consuming process of localizing 3D endocardial surface coordinates with a tracked catheter [132, 134]. Such algorithms typically do not involve an ultrasound reconstruction step, but a real-time ultrasound reconstruction system could be used to give the clinician visual feedback of the acquired 2D ICE images while he or she manually manipulates the tracked 2D ICE transducer to collect a variety of images from different viewpoints. This visual feedback would encourage even and comprehensive sampling of the intraoperative heart, potentially increasing the final reconstruction accuracy. The real-time 4D ultrasound reconstruction system's current computational limitation with respect to the number of output volumes that can be reconstructed is less significant in this ICE application, as typically only one cardiac phase is of interest.

3D ultrasound reconstruction also shows promise for use during image-guided spinal therapy. Our group's augmented reality environment, incorporating a patient-specific preoperative surface model, tracked 2D ultrasound and virtual representations of tracked surgical tools and originally designed to guide intracardiac interventions, has recently been applied to the task of guiding spinal facet injections (Figure 4.1a). A study performed by our group [233] demonstrated that clinicians could accurately perform the needle insertion task using the augmented reality guidance system in phantoms and in a cadaver. Displaying a 3D ultrasound reconstruction of the spinal phantom with volume rendering (Figure 4.1b) also provides a nice visualization of the bony anatomy. A surgical guidance system that uses a volume rendered 3D ultrasound reconstruction for gross context instead of a patient-specific surface model (Figure 4.1c) may therefore give acceptable clinical outcomes without the need for preoperative CT imaging. As proposed in [234], reconstructed 3D ultrasound images may also be used to register the patient to their preoperative CT image if a surface model is desired instead.

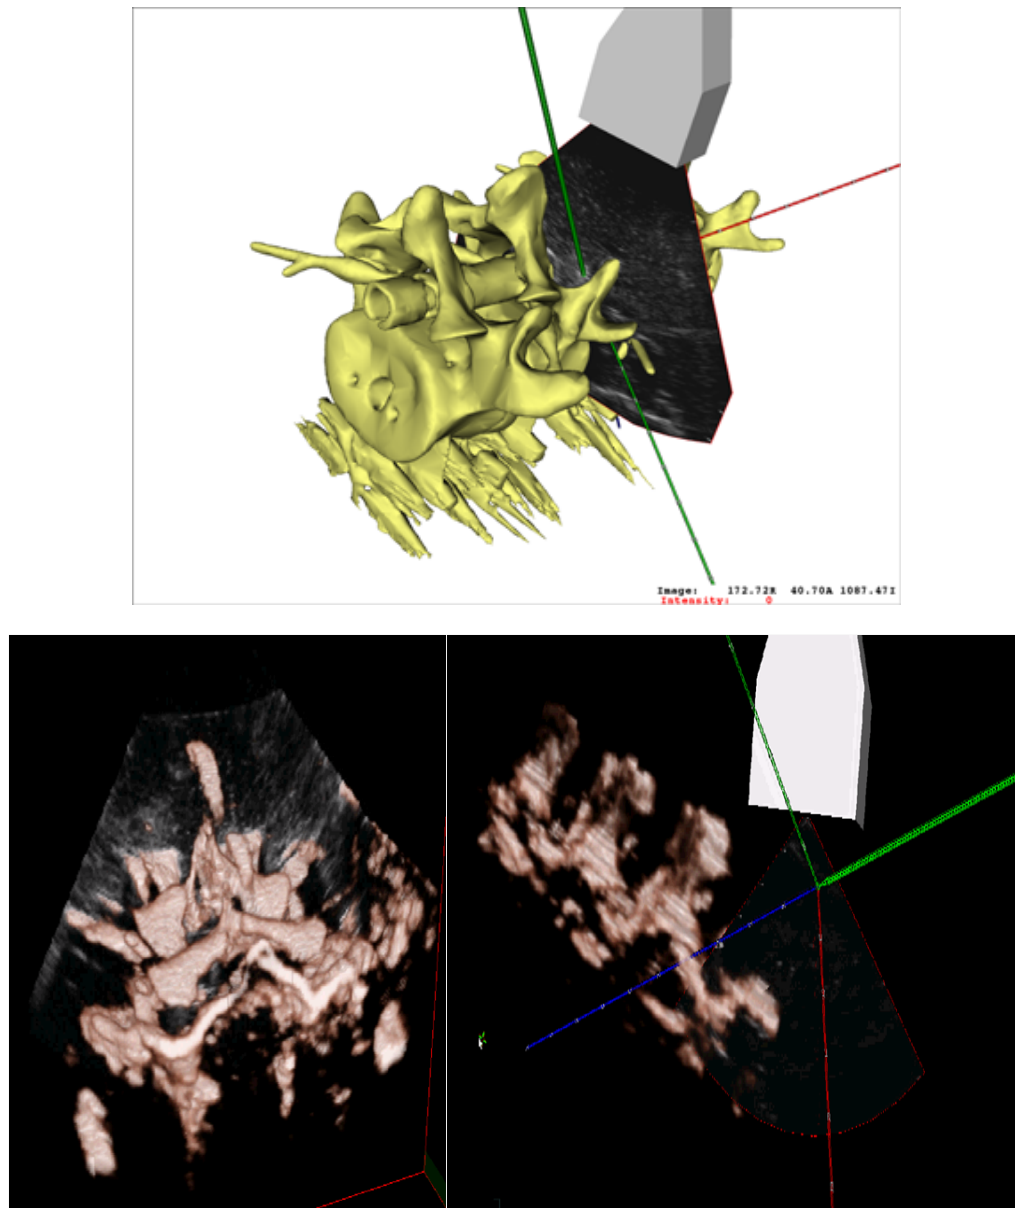


Fig. 4.1: Augmented reality environments for surgical guidance during spinal facet injections [233] (a) Original AR environment incorporating a patient-specific surface model derived from preoperative CT, tracked 2D US and a virtual representation of the tracked needle (b) Volume rendered 3D ultrasound reconstruction of the spine phantom (c) Alternative AR environment incorporating a volume rendered 3D ultrasound reconstruction, tracked 2D US and a virtual representation of the tracked needle. Images courtesy John Moore.

Finally, 4D ultrasound reconstruction can be used as an imaging modality for general research use, especially when OR-compatibility is required. For example, our group is currently involved in a project aimed at improving algorithms for optimal port placement determination in robot-assisted cardiac interventions [235]. Such algorithms typically determine the optimal port placements with respect to factors such as access to the surgical target(s), maneuverability and collision prevention based on a single preoperative end-diastolic 3D CT image. However, surgical plans made on such preoperative images do not extend seamlessly into the operating room because of cardiac transformations caused by different steps of the robotic procedure, including lung collapse and chest insufflation. Our group is currently characterizing these intraoperative deformations, partially by acquiring 4D ultrasound reconstructions of the patients' hearts at each of these steps in the surgical workflow. The final aim of this project is to develop a mechanism of altering the port placement plan to better reflect anticipated changes in the intraoperative anatomy.

### 4.3 Conclusions

Intraoperative imaging forms an integral component of image-guided therapy. This is particularly true during minimally-invasive interventions performed within the beating heart, where significant discrepancies between preoperative images and the actual intraoperative heart are introduced by rapid cardiac motion and non-rigid deformations caused by the surgical access. Ultrasound is well suited to intraoperative imaging as it is easily integrated into the operating room, safe and inexpensive, while providing acceptable spatial resolution and soft tissue contrast. 4D ultrasound reconstruction generates a time series of 3D ultrasound images representing the beating heart over the cardiac cycle and, unlike 2D ultrasound and real-time 3D ultrasound, simultaneously provides a wide field of view, 3D context and high spatial resolution. This form of imaging has potential for use in visualization or image-to-patient registration, but is subject to a highly user-dependent and error prone acquisition procedure.

This thesis presents a real-time 4D ultrasound reconstruction system that provides



real-time visual feedback to clinicians during the ultrasound reconstruction procedure to prevent acquisition errors. Real-time 4D ultrasound reconstruction of the beating heart is enabled by an efficiently-calculated reconstruction kernel, magnetic tracking and ECG-gating, and uses a multithreaded software architecture to reconstruct those frame-grabbed 2D US images acquired at the same point in the cardiac cycle into each corresponding 3D output volume of the 4D ultrasound dataset. By tracking a multiplanar ultrasound probe, the clinician can manipulate the imaging plane manually or by using a rotational approach, thus combining the benefits of freehand and mechanical strategies for ultrasound reconstruction. 4D ultrasound reconstruction also proceeds within an augmented reality environment developed to guide minimally-invasive beating-heart interventions, and so the resulting 4D ultrasound datasets can be visualized alongside a dynamic patient-specific cardiac model, tracked 2D ultrasound and virtual representations of tracked surgical tools.

A human factors experiment performed on a beating heart phantom confirmed the hypothesis that providing real-time visualization improves the quality of reconstructed 4D ultrasound datasets. A clinically acceptable reconstruction accuracy was confirmed in two studies, the first assessing accuracy with a dynamic phantom and giving RMS localization errors between 2.5-2.7 mm for 4D ultrasound reconstruction, and the second determining that a tracked pointer tool rigidly suspended within a beating porcine heart could be localized with an RMS error of approximately 0.8 mm. Experiments analogous to this second intraoperative accuracy assessment were also performed within the laboratory. Finally, the real-time 4D ultrasound reconstruction system was tested in a more clinically realistic setting by acquiring representative images in both swine and clinical patients.

Real-time ultrasound reconstruction algorithms eliminate the post-acquisition reconstruction time that is required by traditional offline systems and currently limits clinical penetration of ultrasound reconstruction imaging. More importantly, the demonstrated quality improvement and clinically relevant accuracy of output volumes reconstructed with the real-time 4D ultrasound reconstruction system make them well suited for subsequent use in image registration and/or visualization during intracardiac image-guided therapy.

# Appendix A

## Software Overview

### A.1 Introduction

This appendix briefly describes the software implementation of the real-time 4D ultrasound reconstruction system described in this thesis. The core reconstruction algorithm and the supporting frame grabbing, tracking and gating components are implemented using the Visualization Toolkit (VTK) [208], a comprehensive and widely-used open-source toolkit for computer graphics and data visualization. Of particular importance for real-time ultrasound reconstruction is the VTK pipeline linking image processing algorithms with both their input(s) and the graphical entities displaying their output(s), as well as the classes supporting multithreading and mutex locks. Software has been developed to provide the user interface for reconstruction parameter specification and to enable dynamic interaction with the output volume(s) within the AtamaiViewer ([www.atamai.com](http://www.atamai.com)) and for any application supporting the OpenIGTLink network transfer protocol [210], such as 3D Slicer [236] ([www.slicer.org](http://www.slicer.org)).

Both the VTK classes and the implementation designed for use with 3D Slicer have been provided as open-source software (called “SynchroGrab4D”) to the research community [209]. Although 3D and 4D ultrasound reconstruction has been integrated into commercial ultrasound systems for quite some time, the inaccessibility of the resulting raw image data makes these systems unsuitable for research, as opposed to

clinical, use. Some highly-specialized open-source software can perform both non-gated and gated ultrasound reconstruction, including the well-known Stradx system [99, 196], but usually cannot perform subsequent image registration, segmentation, advanced visualization (such as volume rendering) or data transfer for integration into surgical guidance systems. In contrast, 3D Slicer is an actively growing open-source application that bundles an extensive collection of medical image processing and image-guided therapy modules. Yet despite current research interest, 4D imaging and image processing algorithms are currently limited in 3D Slicer and in other open-source medical imaging packages.

## A.2 Ultrasound Reconstruction Classes

The core classes are written in C++ using VTK 5.2. Figure A.1 shows a complete class diagram indicating the relationships between the different classes. As described in Chapter 2, in the real-time 4D US reconstruction system separate threads are used to (1) grab and timestamp 2D US frames; (2) retrieve and timestamp transform matrices from the tracking system; (3) perform automatic phase detection from the incoming ECG signal; and (4) perform ultrasound reconstruction by updating pixels within the output volume(s). Additional documentation can be found in the comments within each of the classes described below.

### A.2.1 Video capture

2D US frame grabbing, timestamping and buffering is based on the `vtkVideoSource` VTK classes, which have been refactored to provide a clearer representation for a single frame, the circular buffer of frames and the frame grabbing mechanism (`vtkVideoFrame2`, `vtkVideoBuffer2` and `vtkVideoSource2`, respectively). Although VTK was used here, please note that frame grabbing has recently been added to the Image-Guided Surgery Toolkit (IGSTK) as well [237]. The base class, `vtkVideoSource2`, outputs noise images and is useful for testing, while the derived classes currently support Video-for-Windows and Matrox video cards (Montreal,

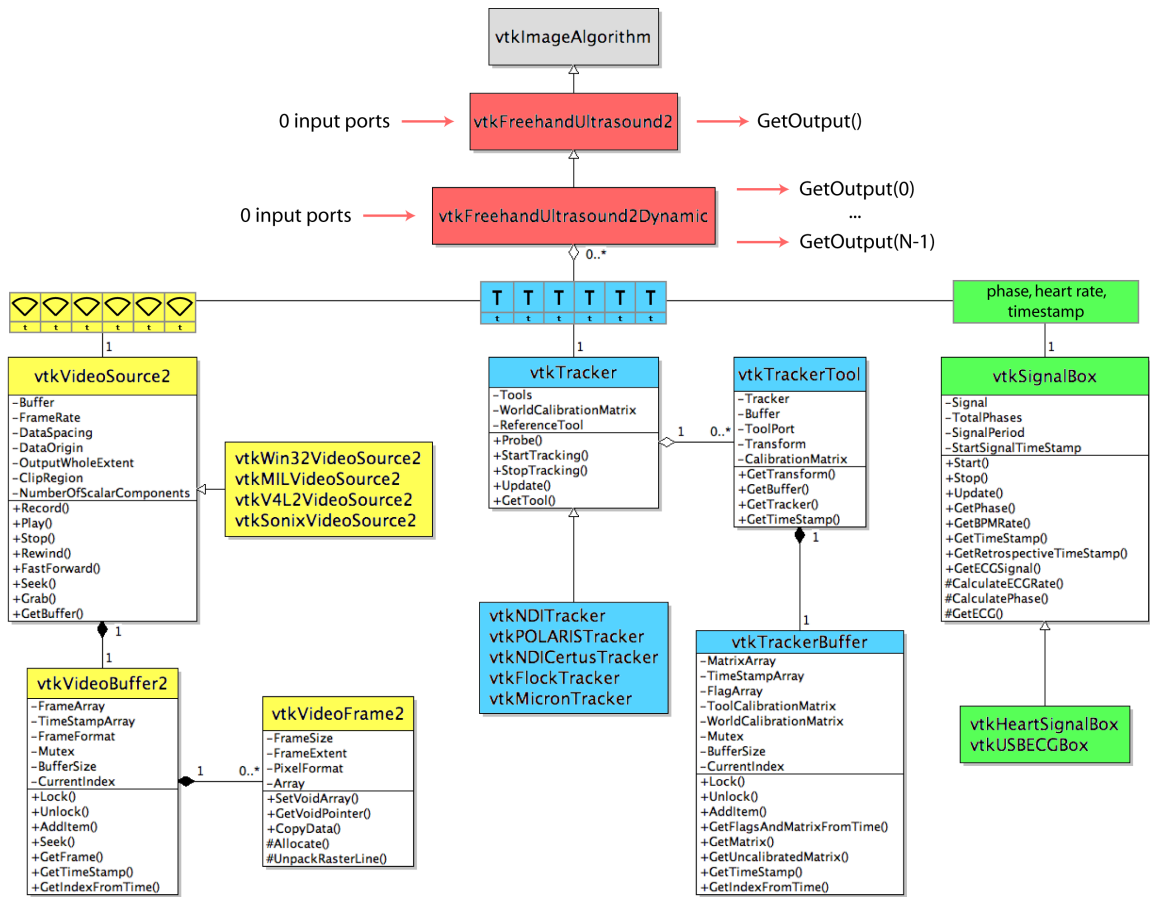


Fig. A.1: Class diagram of the VTK classes implementing the real-time 4D ultrasound reconstruction algorithm.

QC, Canada). Plans exist to add support for Linux drivers and for the open-interface Sonix RP ultrasound system (Ultrasonix, Vancouver, BC, Canada [238]). Although the original VTK frame grabbing classes always output the most recently acquired frame, in the new classes a “desired” timestamp can be specified so that the filter output will instead be the buffered frame whose timestamp is closest to the desired timestamp.

### A.2.2 Tracking

The tracking classes are provided by Atamai Inc. (Calgary, AB, Canada) and support the following tracking systems: Northern Digital Inc.’s Aurora<sup>®</sup>, Polaris<sup>®</sup> and Certus<sup>®</sup> (Waterloo, ON, Canada), Ascension’s Flock of Birds<sup>®</sup> (Burlington, VT, USA) and Claron Technology’s Micron Tracker<sup>®</sup> (Toronto, ON, Canada). A tracking simulator is also available (`vtkFakeTracker`) for testing when a tracking system is not available. Each tracked object is represented by a `vtkTrackerTool`, which holds the current transform matrix and a `vtkTrackerBuffer` storing recently acquired transform matrices, timestamps and status flags in a circular buffer. The interface to each tracking system’s Application Programming Interface (API) is provided by a `vtkTracker`-derived class, which can handle multiple tracked tools. The tracking classes also allow for the specification of a world calibration matrix, a spatial calibration matrix for each tool, and a reference tool.

### A.2.3 Gating

An integral part of the reconstruction software is the gating subsystem, which currently operates on ECG signals but could be easily modified to perform other forms of gating. The base class, `vtkSignalBox`, is an ECG-gating simulator useful for testing purposes with a user-specified cardiac cycle period, and allows the user to set the number of phases within each cardiac cycle. This class also implements prospective gating for estimates of the current cardiac phase and heart rate, and retrospective gating to determine the true time at which each phase began in the previous cycle. The ECG classes provide this numerical output in real-time only, and

unlike the frame grabbing and tracking subsystems, do not store an internal circular buffer. The first derived class, `vtkHeartSignalBox`, gates based on a 5V pulse arriving at the machine’s parallel port, having been originally designed for use with a beating-heart phantom (The Chamberlain Group, Great Barrington, MA, USA) that outputs a voltage pulse at the beginning of each cardiac cycle. The second derived class, `vtkECGUSBBox`, reads a patient ECG signal (amplified from millivolts to volts) over a USB connection. Automatic R-wave detection can then be accomplished by the software by applying a user-specified threshold. Our group has used both `vtkHeartSignalBox` and `vtkECGUSBBox` to interface with clinical anesthesiology equipment within various operating rooms at our institution.

#### A.2.4 Ultrasound reconstruction

Real-time ultrasound reconstruction is performed by two classes: `vtkFreehandUltrasound2` implements real-time 3D ultrasound reconstruction and is based on software originally described in [207, 211], while the derived class `vtkFreehandUltrasound2Dynamic` incorporates prospective and retrospective gating for real-time 4D ultrasound reconstruction. Additional functions used by the reconstruction classes can be found in `vtkFreehandUltrasound2Helper.h`, including the main reconstruction loop. Both `vtkFreehandUltrasound2` and `vtkFreehandUltrasound2Dynamic` derive from `vtkImageAlgorithm`, fitting directly into the VTK 5 pipeline as image filters. A typical `vtkImageAlgorithm` processes its input data only when requested to by a downstream object in the pipeline. In contrast, the ultrasound reconstruction filters operate continuously to detect 2D US images as they arrive and to perform slice selection and slice insertion, sleeping whenever the most recent 2D US image has not changed since the last iteration, or if the ultrasound probe has gone out of view of the tracking system. In the most basic configuration, `vtkFreehandUltrasound2Dynamic` has  $N$  image outputs depicting the beginning of the  $N$  phases of the associated `vtkSignalBox`. However, a subset of phases can be selected to allow greater flexibility in the timepoints of the cardiac cycle that are represented by the output volumes. In addition, this class contains

the functions to impose an expected heart rate in order to reject any 2D US images corresponding to unacceptably long or short cardiac cycles.

## **A.3 User Interface and Communication**

### **A.3.1 In the AtamaiViewer**

The “US Reconstruction” module in the AtamaiViewer enables real-time 3D and 4D ultrasound reconstruction (Figure A.2a) and is written in python using wrapped VTK classes. The US reconstruction module interfaces with the AtamaiViewer’s “Surgical Tracker” module handling tracked surgical tools and the ECG modules for ECG-gating. Using the Surgical Tracker module, the user first configures the tracked tool(s), video source(s), calibration matrices and, if desired, the reference tool. The user must also set one or more callbacks specifying which of the tracked tools are ultrasound transducers to be used for ultrasound reconstruction. The number of phases is specified using one of the AtamaiViewer’s ECG modules, which can also be used to interactively determine the threshold used for automatic R-wave detection in patient imaging.

Once frame grabbing, tracking and gating have been configured, the US Reconstruction module handles the GUI (Graphical User Interface) allowing the user to set the reconstruction parameters, to perform the centring procedure dictating the ultrasound key frame used to determine the origin of the output volume(s), to start and stop the real-time ultrasound reconstruction, and subsequently to save the output volumes and related files. By default, the spacing of the output volume(s) is set to double that of the video source used to grab 2D US images, while their extent is set to match the video source’s dimensions. The user can toggle between the output volumes to select the one to be displayed in real-time. This volume is added to the AtamaiViewer’s orthoplane display, which is updated whenever the AtamaiViewer detects the user’s mouse movement or whenever a tracked tool either moves or goes out of view.

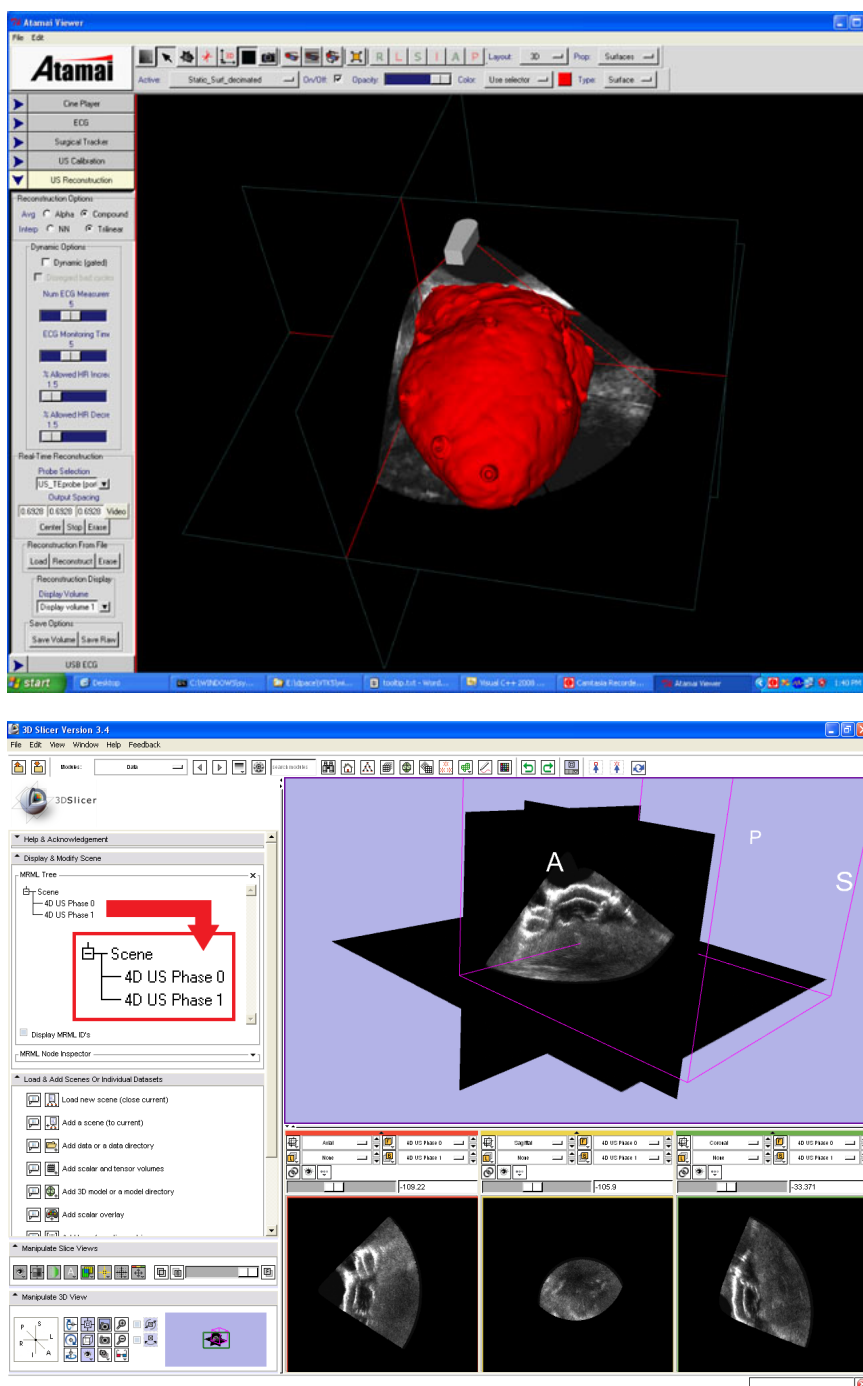


Fig. A.2: Example 4D ultrasound reconstruction of a beating heart phantom (a) using the US Reconstruction module in the AtamaiViewer; (b) using SynchroGrab4D with 3D Slicer.



### A.3.2 In 3D Slicer: SynchroGrab4D

Unlike the AtamaiViewer implementation, SynchroGrab4D is not integrated within any particular image-guided therapy system, and is instead designed as a command-line 3D/4D ultrasound reconstruction application that sends reconstructed output volumes to an IGT system over a network connection. Although I have focused on visualizing the reconstruction results within 3D Slicer 3.4 (Figure A.2b), SynchroGrab4D can be used with any application implementing the OpenIGTLink network protocol. The software is written in C++, is cross-platform and is based on a previously-described open-source 3D ultrasound reconstruction system [238] that used an earlier variant of the 3D US reconstruction VTK classes described above. In addition to the ultrasound reconstruction threads, three threads (1) buffer the output volume(s) to be transferred; (2) send OpenIGTLink messages containing the image data for the output volume(s); and (3) send OpenIGTLink messages containing the transform matrices for the tracked ultrasound probe retrieved from the tracking system (optional).

Command-line options allow the user to specify a calibration file containing the reconstruction parameters, the type of video grabbing, tracking and gating hardware to be used, and the server IP and port number. During real-time US reconstruction, the output volume(s) can be sent at a user-specified rate over the network connection for interactive visualization within an IGT system. SynchroGrab4D is currently implemented such that the output volume(s) are transferred in their entirety, which does restrict the rate at which they can be updated within 3D Slicer’s virtual scene. Alternatively, the user can choose to delay the image transfer until after the completion of the reconstruction if the transfer speed is insufficient. The transform matrices from the tracking system can also be sent using the OpenIGTLink protocol to a second port, therefore allowing the US probe’s position and orientation to be visualized by the IGT system as well.

3D Slicer’s “OpenIGTLink IF” module can manage the incoming output volumes; for 4D US reconstruction their names in 3D Slicer’s Medical Reality Modeling Language (MRML) data tree are annotated with the phase that they correspond to. If

the user elects to send the output volume(s) over the network connection in real-time during the reconstruction, then 3D Slicer’s user interface can be used to specify the output volume to be displayed and the incremental reconstruction results can be visualized. These real-time updates allows the operator to ensure that the reconstructed volume(s) contain all of the required image data. A relatively new “4D Imaging” module that visualizes 4D image datasets is also under active development by the 3D Slicer community. The reconstructed 3D volumes making up the four-dimensional ultrasound dataset are added to a “4D bundle” within 3D Slicer along with their relative timestamps, and the 4D Imaging module can then display the volumes as a cine loop to visualize the motion of the organ of interest over time.

## A.4 Conclusions

Within the operating room, intraoperatively reconstructed 4D ultrasound datasets can be used as part of the registration process between preoperative images and the intraoperative patient and for surgical guidance within augmented reality systems. Since the two real-time US reconstruction programs described here are implemented for use with the AtamaiViewer and 3D Slicer, any of these tasks can be performed using the variety of tools within these two suites. In addition, the frame grabbing, tracking and gating classes have been released to the open-source community and are of general interest for research use, as their modular design makes them easily extensible to work with currently unsupported hardware. Many popular tracking systems are currently supported, and a variety of ultrasound machines can be used by specifying the appropriate parameters for extracting the US fan from the 2D US frames. Although the automatic rotation detection for multiplanar US probes is currently customized for frame grabbing from a Philips Sonos 7500 or 2500 scanner with a Matrox Meteor or Morphis card, this piece of code can be easily modified for use with different ultrasound machines. Finally, it would also be a simple matter to integrate additional gating schemes into the system by sub-classing the base gating class, whether operating on ECG signals or for other forms of gating such as respiratory gating.

# References

- [1] M.J. Mack. Minimally invasive and robotic surgery. *The Journal of the American Medical Association*, 285:568–572, 2001.
- [2] A.P. Kypson. Recent trends in minimally invasive cardiac surgery. *Cardiology*, 107:147–158, 2007.
- [3] V.A. Subramanian, D.F. Loulmet, and N.C. Patel. Minimally invasive coronary artery bypass grafting. *Seminars in Thoracic and Cardiovascular Surgery*, 19:281–288, 2007.
- [4] Y.J. Woo, J. Seeburger, and F.W. Mohr. Minimally invasive valve surgery. *Seminars in Thoracic and Cardiovascular Surgery*, 19:289–298, 2007.
- [5] J.G. Webb and S. Lichtenstein. Transcatheter percutaneous and transapical aortic valve replacement. *Seminars in Thoracic and Cardiovascular Surgery*, 19:304–310, 2007.
- [6] A.E. Saltman. Minimally invasive surgery for atrial fibrillation. *Seminars in Thoracic and Cardiovascular Surgery*, 19:33–38, 2007.
- [7] P.J. del Nido. Minimal incision congenital cardiac surgery. *Seminars in Thoracic and Cardiovascular Surgery*, 19:319–324, 2007.
- [8] D.N. Wijeyesundera, W.S. Beattie, G. Djaiani, V. Rao, M.A. Borger, K. Karkouti, and R.J. Cusimano. Off-pump coronary artery surgery for reducing mortality and morbidity: Meta-analysis of randomized and observational studies. *Journal of the American College of Cardiology*, 46:872–882, 2005.

- [9] E.A. Grossi, A.C. Galloway, G.H. Ribakove, P.K. Zakow, C.C. Derivaux, F.G. Baumann, D. Schwesinger, and S.B. Colvin. Impact of minimally invasive valvular heart surgery: A case-control study. *The Annals of Thoracic Surgery*, 71:807–810, 2001.
- [10] J.A. Morgan, J.C. Peacock, T. Kohmoto, M.J. Garrido, B.M. Schanzer, A.R. Kherani, D.W. Vigilance, F.H. Cheema, S. Kaplan, C.R. Smith, M.C. Oz, and M. Argenziano. Robotic techniques improve quality of life in patients undergoing atrial septal defect repair. *The Annals of Thoracic Surgery*, 77:1328–1333, 2004.
- [11] J.A. Morgan, B.A. Thornton, J.C. Peacock, K.W. Hollingsworth, C.R. Smith, M.C. Oz, and M. Argenziano. Does robotic technology make minimally invasive cardiac surgery too expensive? A hospital cost analysis of robotic and conventional techniques. *Journal of Cardiac Surgery*, 20:246–251, 2005.
- [12] R. Ascione, M. Caputo, G. Calori, C.T. Lloyd, M.J. Underwood, and G.D. Angelini. Predictors of atrial fibrillation after conventional and beating heart coronary surgery: A prospective, randomized study. *Circulation*, 102:1530–1535, 2000.
- [13] R. Ascione, C.T. Lloyd, M.J. Underwood, A.A. Lotto, A.A. Pitsis, and G.D. Angelini. Inflammatory response after coronary revascularization with or without cardiopulmonary bypass. *The Annals of Thoracic Surgery*, 69:1198–1204, 2000.
- [14] S.C. Stamou, K.A. Jablonski, A.J. Pfister, P.C. Hill, M.K.C. Dullum, A.S. Bafi, S.W. Boyce, K.R. Petro, and P.J. Corso. Stroke after conventional versus minimally invasive coronary artery bypass. *The Annals of Thoracic Surgery*, 74:394–399, 2002.
- [15] D. Van Dijk, E.W.L. Jansen, R. Hijman, A.P. Nierich, J.C. Diephuis, K.G.M. Moons, J.R. Lahpor, C. Borst, A.M.A. Keizer, H.M. Nathoe, D.E. Grobbee, P.P.T. De Jaegere, and C.J. Kalkman. Cognitive outcome after off-pump and

- on-pump coronary artery bypass graft surgery. *The Journal of the American Medical Association*, 287:1405–1412, 2002.
- [16] D.A. Etzioni, J.H. Liu, M.A. Maggard, and C.Y. Ko. The aging population and its impact on the surgery workforce. *Annals of Surgery*, 238:170–177, 2003.
- [17] Y.J. Woo. Robotic cardiac surgery. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 2:225–232, 2006.
- [18] F.J. Benetti, C. Ballester, G. Sani, P. Doonstra, and J. Grandjean. Video assisted coronary bypass surgery. *Journal of Cardiac Surgery*, 10:620–625, 1995.
- [19] W.R. Chitwood, C.L. Wixon, J.R. Elbeery, J.F. Moran, W.H.H. Chapman, and R.M. Lust. Video-assisted minimally invasive mitral valve surgery. *The Journal of Thoracic and Cardiovascular Surgery*, 114:773–782, 1997.
- [20] A. Cribier, H. Eltchaninoff, A. Bash, N. Borenstein, C. Tron, F. Bauer, G. Derumeaux, F. Anselme, F. Laborde, and M.B. Leon. Percutaneous transcatheter implantation of an aortic valve prosthesis for calcific aortic stenosis. *Circulation*, 106:3006–3008, 2002.
- [21] A.E. Saltman, L.S. Rosenthal, N.A. Francalancia, and S.J. Lahey. A completely endoscopic approach to microwave ablation for atrial fibrillation. *The Heart Surgery Forum*, 6:E38–E41, 2003.
- [22] A. Carpentier, D. Loulmet, B. Aupeple, J.P. Kieffer, D. Tournay, P. Guibourt, A. Fiemeyer, D. Meleard, P. Richomme, and C. Cardon. First computer assisted open heart surgery: First case operated on with success. *Comptes Rendus de l'Academie des Sciences - Series III*, 321:437–442, 1998.
- [23] D. Loulmet, A. Carpentier, N. d'Attellis, A. Berrebi, C. Cardon, O. Ponzio, B. Aupeple, and J.Y.M. Relland. Endoscopic coronary artery bypass grafting with the aid of robotic assisted instruments. *The Journal of Thoracic and Cardiovascular Surgery*, 118:4–10, 1999.

- [24] T.M. Peters. Image-guidance for surgical procedures. *Physics in Medicine and Biology*, 51(14):R505–R540, 2006.
- [25] J.S. Shanewise, R. Zaffer, and R.P. Martin. Intraoperative echocardiography and minimally invasive cardiac surgery. *Echocardiography*, 19:579–582, 2002.
- [26] D. Bainbridge I. Iglesias and J. Murkin. Intraoperative echocardiography: Support for decision making in cardiac surgery. *Seminars in Cardiothoracic and Vascular Anesthesia*, 8:25–35, 2004.
- [27] C.M. Otto. *The Practice of Clinical Echocardiography*. Saunders Elsevier, 3rd edition, 2007.
- [28] A. Fenster and D.B. Downey. Three-dimensional ultrasound imaging. *Annual Reviews in Biomedical Engineering*, 2:457–475, 2000.
- [29] I.S. Salgo. 3D echocardiographic visualization for intracardiac beating heart surgery and intervention. *The Journal of Thoracic and Cardiovascular Surgery*, 19:325–329, 2007.
- [30] S.W. Downing, W.A. Herzog, J.S. McLaughlin, and T.P. Gilbert. Beating-heart mitral valve surgery: Preliminary model and methodology. *The Journal of Thoracic and Cardiovascular Surgery*, 123:1141–1146, 2002.
- [31] 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 3-dimensional ultrasound for guiding surgical tasks. *Computer Aided Surgery*, 8:82–90, 2003.
- [32] Y. Suematsu, G.R. Marx, J.A. Stoll, P.E. DuPont, R.O. Cleveland, R.D. Howe, J.K. Triedman, T. Mihaljevic, B.N. Mora, B.J. Savord, I.S. Salgo, and P.J. del Nido. Three-dimensional echocardiography-guided beating-heart surgery without cardiopulmonary bypass: A feasibility study. *The Journal of Thoracic and Cardiovascular Surgery*, 128:579–587, 2004.
- [33] Y. Suematsu, B. Kiaii, D.T. Bainbridge, P.J. del Nido, and R.J. Novick. Robotic-assisted closure of atrial septal defect under real-time three-dimensional

- echo guide: in vitro study. *European Journal of Cardiothoracic Surgery*, 32(4):573–576, 2007.
- [34] Y. Suematsu, J.F. Martinez, B.K. Wolf, G.R. Marx, J.A. Stoll, P.E. DuPont, R.D. Howe, J.K. Triedman, and P.J. del Nido. Three-dimensional echo-guided beating heart surgery without cardiopulmonary bypass: Atrial septal defect closure in a swine model. *The Journal of Thoracic and Cardiovascular Surgery*, 130:1348–1357, 2005.
- [35] R. McKendrick and C.Y. Owada. Real-time 3D echocardiography-guided transcatheter device closure of atrial septal defects. *Catheterization and Cardiovascular Interventions*, 65:442–446, 2005.
- [36] M. Scheurer, V. Bandisode, P. Ruff, A. Atz, and G. Shirali. Early experience with real-time three-dimensional echocardiographic guidance of right ventricular biopsy in children. *Echocardiography*, 23:45–49, 2006.
- [37] J. Balzer, H. Kuhl, T. Rassaf, R. Hoffmann, P. Schauerte, M. Kelm, and A. Franke. Real-time transesophageal three-dimensional echocardiography for guidance of percutaneous cardiac interventions: First experience. *Clinical Research in Cardiology*, 97:565–574, 2008.
- [38] G. Perk, R.M. Lang, M.A. Garcia-Fernandez, J. Lodato, L. Sugeng, J. Lopez, B.P. Knight, D. Messika-Zeitoun, S. Shah, J. Slater, E. Brochet, M. Varkey, Z. Hijazi, N. Marino, C. Ruiz, and I. Kronzon. Use of real time three-dimensional transesophageal echocardiography in intracardiac catheter based interventions. *Journal of the American Society of Echocardiography*, 22:865–882, 2009.
- [39] M.J. Mullen, B.F. Dias, F. Walker, S.C. Siu, L.N. Benson, and P.R. McLaughlin. Intracardiac echocardiography guided device closure of atrial septal defects. *Journal of the American College of Cardiology*, 41:285–292, 2003.

- [40] E.G. Daoud, S.J. Kalbfleisch, and J.D. Hummel. Intracardiac echocardiography to guide transseptal left heart catheterization for radiofrequency catheter ablation. *Journal of Cardiovascular Electrophysiology*, 10:358–363, 1999.
- [41] J.M. Mangrum, J.P. Mounsey, L.C. Kok, J.P. DiMarco, and D.E. Haines. Intracardiac echocardiography-guided, anatomically based radiofrequency ablation of focal atrial fibrillation originating from pulmonary veins. *Journal of the American College of Cardiology*, 39:1964–1972, 2002.
- [42] C. Knackstedt, A. Franke, K. Mischke, M. Zarse, F. Gramley, T. Schimpf, J. Plisiene, G. Muehlenbruch, E. Spuentrup, S. Ernst, S. Willems, P. Krichhof, and P. Schauerte. Semi-automated 3-dimensional intracardiac echocardiography: Development and initial clinical experience of a new system to guide ablation procedures. *Heart Rhythm*, 3:1453–1459, 2006.
- [43] L. Mercier, T. Lango, F. Lindseth, and L.D. Collins. A review of calibration techniques for freehand 3-D ultrasound systems. *Ultrasound in Medicine & Biology*, 31:449–471, 2005.
- [44] M. Hastenteufel, M. Vetter, H.P. Meinzer, and I. Wolf. Effect of 3D ultrasound probes on the accuracy of electromagnetic tracking systems. *Ultrasound in Medicine & Biology*, 32:1359–1368, 2006.
- [45] L.F. Gutierrez, G. Shechter, D. Stanto, S. Dalal, D. Elgort, R. Manzke, R.C. Chan, and L. Zagorchev. Multimodality image guidance system integrating X-ray fluoroscopy and ultrasound image streams with electromagnetic tracking. In *Medical Imaging: Visualization and Image-Guided Procedures*, volume 6509 of *SPIE 2007*, page 6509K, 2007.
- [46] M.G. Linguraru, N.V. Vasilyev, P.J. del Nido, and R.D. Howe. Statistical segmentation of surgical instruments in 3-D ultrasound images. *Ultrasound in Medicine & Biology*, 33:1428–1437, 2007.
- [47] P.M. Novotny, J.W. Cannon, and R.D. Howe. Tool localization in 3D ultrasound



- images. In *Medical Image Computing and Computer Assisted Interventions*, volume 2879, pages 969–970. Lecture Notes in Computer Science, 2003.
- [48] T. Ortmaier, M.A. Vitrani, G. Morel, and S. Pinault. Robust real-time instrument tracking in ultrasound images for visual servoing. In *IEEE International Conference on Robotics and Automation*, pages 2167–2172, 2005.
- [49] J. Stoll and P. Dupont. Passive markers for ultrasound tracking of surgical instruments. In *Medical Image Computing and Computer Assisted Interventions*, volume 3750, pages 41–48. Lecture Notes in Computer Science, 2005.
- [50] P.M. Novotny, J.A. Stoll, N.V. Vasilyev, P.J. del Nido, P.E. Dupont, T.E. Zickler, and R.D. Howe. GPU based real-time instrument tracking with three-dimensional ultrasound. *Medical Image Analysis*, 11:458–464, 2007.
- [51] R. Shekhar and V. Zagrodsky. Mutual information-based rigid and nonrigid registration of ultrasound volumes. *IEEE Transactions on Medical Imaging*, 21:9–22, 2002.
- [52] R. Shekhar, V. Zagrodsky, M.J. Garcia, and J.D. Thomas. Registration of real-time 3-D ultrasound images of the heart for novel 3-D stress echocardiography. *IEEE Transactions on Medical Imaging*, 23:1141–1149, 2004.
- [53] V. Grau, H. Becher, and J.A. Noble. Registration of multiview real-time 3-D echocardiographic sequences. *IEEE Transactions on Medical Imaging*, 26:1154–1165, 2007.
- [54] A.H. Gee, G.M. Treece, R.W. Prager, C.J.C. Cash, and L. Berman. Rapid registration for wide field of view freehand three-dimensional ultrasound. *IEEE Transactions on Medical Imaging*, 22:1344–1357, 2003.
- [55] K.Z. Abd-Elmoniem, A.M. Youssef, and Y.M. Kadah. Real-time speckle reduction and coherence enhancement in ultrasound imaging via nonlinear anisotropic diffusion. *IEEE Transactions on Biomedical Engineering*, 49:997–1014, 2002.

- [56] A. Achim, A. Bezerianos, and P. Tsakalides. Novel bayesian multiscale method for speckle removal in medical ultrasound images. *IEEE Transactions on Medical Imaging*, 20:772–783, 2001.
- [57] X. Zong, A.F. Laine, and E.A. Geiser. Speckle reduction and contrast enhancement of echocardiograms via multiscale nonlinear processing. *IEEE Transactions on Medical Imaging*, 17:532–540, 1998.
- [58] N. Azzabou and N. Paragios. Spatio-temporal speckle reduction in ultrasound sequences. In *Medical Image Computing and Computer Assisted Interventions*, volume 5241, pages 951–958. Lecture Notes in Computer Science, 2008.
- [59] X. Huang, G. Guiraudon, D.L. Jones, D. Bainbridge, J. Ren, and T.M. Peters. Dynamic 2D ultrasound and 3D CT registration of the beating heart. *IEEE Transactions on Medical Imaging*, 28:1179–1189, 2009.
- [60] X. Huang, J. Ren, G. Guiraudon, D. Boughner, and T.M. Peters. Rapid dynamic image registration of the beating heart for diagnosis and surgical navigation. *IEEE Transactions on Medical Imaging*, 28:1802–1814, 2009.
- [61] V. Walimbe, V. Zagrodsky, S. Raja, W.A. Jaber, F.P. DiFilippo, M.J. Garcia, R.C. Brunken, J.D. Thomas, and R. Shekhar. Mutual information-based multimodality registration of cardiac ultrasound and SPECT images: A preliminary investigation. *The International Journal of Cardiovascular Imaging*, 19:483–494, 2003.
- [62] G.M. Treece, R.W. Prager, A.H. Gee, and L. Berman. Correction of probe pressure artifacts in freehand 3D ultrasound. *Medical Image Analysis*, 6:199–214, 2002.
- [63] G.P. Penney, J.M. Blackall, M.S. Hamady, T. Sabharwal, A. Adam, and D.J. Hawkes. Registration of freehand 3D ultrasound and magnetic resonance liver images. *Medical Image Analysis*, 8:81–91, 2004.

- [64] S.K. Nadkarni, D. Boughner, and A. Fenster. Image-based cardiac gating for three-dimensional intravascular ultrasound imaging. *Ultrasound in Medicine & Biology*, 31:53–63, 2005.
- [65] R.N. Czerwinski, D.L. Jones, and W.D. O’Brien. Detection of lines and boundaries in speckle images - applications to medical ultrasound. *IEEE Transactions on Medical Imaging*, 18:126–136, 1999.
- [66] A. Leroy, P. Mozer, Y. Payan, and J. Troccaz. Rigid registration of freehand 3D ultrasound and CT-scan kidney images. In *Medical Image Computing and Computer Assisted Interventions*, volume 3216, pages 837–844. Lecture Notes in Computer Science, 2004.
- [67] J. Montagnat, M. Sermesant, H. Delingette, G. Malandain, and N. Ayache. Anisotropic filtering for model-based segmentation of 4D cylindrical echocardiographic images. *Pattern Recognition Letters*, 24:815–828, 2003.
- [68] C. Bergmeir and N. Subramanian. Operator guidance in 2D echocardiography via 3D model to image registration. In *Medical Imaging: Ultrasonic Imaging and Signal Processing*, volume 7265 of *SPIE 2009*, page 726518, 2009.
- [69] A.P. King, Y.L. Ma, C. Yao, C. Jansen, R. Razavi, K.S. Rhode, and C.P. Penney. Image-to-physical registration for image-guided interventions using 3-D ultrasound and an ultrasound imaging model. In *Information Processing in Medical Imaging*, volume 5636, pages 188–201. Lecture Notes in Computer Science, 2009.
- [70] W. Wein, B. Roper, and N. Navab. Integrating diagnostic B-mode ultrasonography into CT-based radiation treatment planning. *IEEE Transactions on Medical Imaging*, 26:866–879, 2007.
- [71] H. Zhong, T. Kanade, and D. Schwartzman. Image thickness correction for navigation with 3D intra-cardiac ultrasound catheter. In *Medical Image Computing and Computer Assisted Interventions*, volume 5242, pages 485–492. Lecture Notes in Computer Science, 2008.

- [72] R. Rohling, A. Gee, and L. Berman. Three-dimensional spatial compounding of ultrasound images. *Medical Image Analysis*, 1:177–193, 1997.
- [73] R.R. Entrekin, B.A. Porter, H.H. Sillesen, A.D. Wong, P.L. Cooperberg, and C.H. Fix. Real-time spatial compound imaging: Application to breast, vascular and musculoskeletal ultrasound. *Seminars in Ultrasound, CT, and MRI*, 22:50–64, 2001.
- [74] A.R. Groves and R.N. Rohling. Two-dimensional spatial compounding with warping. *Ultrasound in Medicine & Biology*, 30:929–942, 2004.
- [75] J.F. Krucker, C.R. Meyer, G.L. LeCarpentier, J.B. Fowlkes, and P.L. Carson. 3D spatial compounding of ultrasound images using image-based nonrigid registration. *Ultrasound in Medicine & Biology*, 26:1475–1488, 2000.
- [76] T.C. Poon and R.N. Rohling. Three-dimensional extended field-of-view ultrasound. *Ultrasound in Medicine & Biology*, 32:357–369, 2006.
- [77] O. Kutter, W. Wein, and N. Navab. Multi-modal registration based ultrasound mosaicing. In *Medical Image Computing and Computer Assisted Interventions*, volume 5761, pages 763–770. Lecture Notes in Computer Science, 2009.
- [78] C. Wachinger, W. Wein, and N. Navab. Registration strategies and similarity measures for three-dimensional ultrasound mosaicing. *Academic Radiology*, 15:1404–1415, 2008.
- [79] K. Rajpoot, J.A. Noble, V. Grau, C. Szmigielski, and H. Becher. Multiview RT3D echocardiography image fusion. In *Functional Imaging and Modeling of the Heart*, volume 5528. Lecture Notes in Computer Science, 2009.
- [80] P. Soler, O. Gerard, P. Allain, E. Saloux, E. Angelini, and I. Bloch. Comparison of fusion techniques for 3D+T echocardiography acquisitions from different acoustic windows. *IEEE Computers in Cardiology*, 32:141–144, 2005.

- [81] X. Ye, J.A. Noble, and D. Atkinson. 3-D freehand echocardiography for automatic left ventricle reconstruction and analysis based on multiple acoustic windows. *IEEE Transactions on Medical Imaging*, 21:1051–1058, 2002.
- [82] V. Grau and J.A. Noble. Adaptive multiscale ultrasound compounding using phase information. In *Medical Image Computing and Computer Assisted Interventions*, volume 3749, pages 589–596. Lecture Notes in Computer Science, 2005.
- [83] M.L. Chuang, M.G. Hibberd, R.A. Beaudin, M.G. Mooney, M.F. Riley, J.T. Fearnside, and P.S. Douglas. Patient motion compensation during transthoracic 3-D echocardiography. *Ultrasound in Medicine & Biology*, 27:203–209, 2001.
- [84] C.A. Linte, J. Moore, A.D. Wiles, C. Wedlake, and T.M. Peters. Virtual reality-enhanced ultrasound guidance: A novel technique for intracardiac interventions. *Computer Aided Surgery*, 13:82–94, 2008.
- [85] Y. Sun, S. Kadoury, Y. Li, M. John, J. Resnick, G. Plambeck, R. Liao, F. Sauer, and C. Xu. Image guidance of intracardiac ultrasound with fusion of pre-operative images. In *Medical Image Computing and Computer Assisted Interventions*, volume 4791, pages 60–67. Lecture Notes in Computer Science, 2007.
- [86] D.F. Pace, A.D. Wiles, J. Moore, C. Wedlake, D.G. Gobbi, and T.M. Peters. Validation of four-dimensional ultrasound for targeting in minimally-invasive beating-heart surgery. In *Medical Imaging: Visualization, Image-Guided Procedures, and Modeling*, 2009.
- [87] G. Shechter, C. Ozturk, J.R. Resar, and E.R. McVeigh. Respiratory motion of the heart from free breathig coronary angiograms. *IEEE Transactions on Medical Imaging*, 23:1046–1056, 2004.
- [88] V. Walimbe, W.A. Jaber, M.J. Garcia, and R. Shekhar. Multimodality cardiac stress testing: Combining real-time 3-dimensional echocardiography and

- myocardial perfusion SPECT. *The Journal of Nuclear Medicine*, 50:226–230, 2009.
- [89] H. Zhu, K.D. Oakeson, and M.H. Friedman. Retrieval of cardiac phase from IVUS sequences. In *Medical Imaging: Ultrasonic Imaging and Signal Processing*, volume 5035 of *SPIE 2003*, pages 135–146, 2003.
- [90] D. Guo and P. Richardson. Detection of cardiac cycle from intracoronary ultrasound. *Ultrasound in Medicine & Biology*, 32:345–356, 2006.
- [91] S.M. O’Malley, J.F. Granada, S. Carlier, M. Naghavi, and I.A. Kakadiaris. Image-based gating of intravascular ultrasound pullback sequences. *IEEE Transactions on Information Technology in Biomedicine*, 12:299–306, 2008.
- [92] S.A. de Winter, R. Hamers, M. Degertekin, K. Tanabe, P.A. Lemos, P.W. Seruys, J.R.T.C. Roelandt, and N. Bruining. A novel retrospective gating method for intracoronary ultrasound images based on image properties. *Computers in Cardiology*, 30:13–16, 2003.
- [93] J. Deng, J.E. Gardener, C.H. Rodeck, and W.R. Lees. Fetal echocardiography in three and four dimensions. *Ultrasound in Medicine & Biology*, 22:979–986, 1996.
- [94] G.R. Devore, P. Falkensammer, M.S. Sklansky, and L.D. Platt. Spatio-temporal image correlation (STIC): New technology for evaluation of the fetal heart. *Ultrasound in Obstetrics and Gynecology*, 22:380–387, 2003.
- [95] P. Falkensammer and H. Brandl. Ultrasound technology update: 4D fetal echocardiography: Spatio-temporal image correlation (STIC) for fetal heart acquisition. Report, GE Medical Systems, Kretz Ultrasound, 2003.
- [96] J. Deng, R. Yates, A.G. Birkett, C.F. Ruff, A.D. Linney, W.R. Lees, M.A. Hanson, and C.H. Rodeck. Online motion-gated dynamic three-dimensional echocardiography in the fetus - preliminary results. *Ultrasound in Medicine & Biology*, 27:43–50, 2001.

- [97] S. Brekke, E. Tegnander, H.G. Torp, and S.H. Eik-nes. Tissue Doppler gated (TDOG) dynamic three-dimensional ultrasound imaging of the fetal heart. *Ultrasound in Obstetrics and Gynecology*, 24:192–198, 2004.
- [98] J.M. Rubin, J.B. Fowlkes, M.R. Prince, R.T. Rhee, and T.L. Chenevert. Doppler US gating of cardiac MR imaging. *Academic Radiology*, 7:1116–1122, 2000.
- [99] G.M. Treece, R.W. Prager, A.H. Gee, C.J.C. Cash, and L. Berman. Grey-scale gating for freehand 3D ultrasound. In *IEEE International Symposium on Biomedical Imaging*, pages 993–996, 2002.
- [100] K. Karadayi, T. Hayashi, and Y. Kim. Automatic image-based gating for 4D ultrasound. In *28th Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 2388–2391, 2006.
- [101] H. Sundar, A. Khamene, L. Yatziv, and C. Xu. Automatic image-based cardiac and respiratory cycle synchronization and gating of image sequences. In *Medical Image Computing and Computer Assisted Interventions*, volume 5762, pages 381–388. Lecture Notes in Computer Science, 2009.
- [102] Y. Wang, S.J. Riederer, and R.L. Ehman. Respiratory motion of the heart: Kinematics and the implications for the spatial resolution in coronary imaging. *Magnetic Resonance in Medicine*, 33:713–719, 1995.
- [103] K. McLeish, D.L.G. Hill, D. Atkinson, J.M. Blackall, and R. Razavi. A study of the motion and deformation of the heart due to respiration. *IEEE Transactions on Medical Imaging*, 21:1142–1150, 2002.
- [104] K. Toska and Eriksen. Respiration-synchronous fluctuations in stroke volume, heart rate and arterial pressure in humans. *The Journal of Physiology*, 472:501–512, 1993.
- [105] K. Nehrke, P. Bornert, D. Manke, and J.C. Bock. Free-breathing cardiac MR

- imaging: Study of implications of respiratory motion - initial results. *Radiology*, 220:810–815, 2001.
- [106] P.A. Noseworthy, Z.J. Malchano, J. Ahmed, G. Holmvang, J.N. Ruskin, and V.Y. Reddy. The impact of respiration on left atrial and pulmonary venous anatomy: Implications for image-guided intervention. *Heart Rhythm*, 2:1173–1178, 2005.
- [107] J. Ector, S. De Buck, D. Loeckx, W. Coudyzer, F. Maes, S. Dymarkowski, J. Bogaert, and H. Heidbuchel. Changes in left atrial anatomy due to respiration: Impact on three-dimensional image integration during atrial fibrillation ablation. *Journal of Cardiovascular Electrophysiology*, 19:828–834, 2008.
- [108] C. Jahnke, I. Paetsch, S. Achenbach, B. Schnackenburg, R. Gebker, E. Fleck, and E. Nagel. Coronary MR imaging: Breath-hold capability and patterns, coronary artery rest periods and  $\beta$ -blocker use. *Radiology*, 239:71–78, 2006.
- [109] A.E. Holland, J.W. Goldfarb, and R.R. Edelman. Diaphragmatic and cardiac motion during suspended breathing: Preliminary experience and implications for breath-hold MR imaging. *Radiology*, 209:483–489, 1998.
- [110] S.B. Gay, C.L. Siström, C.A. Holder, and P.M. Suratt. Breath-holding capability of adults: Implications for spiral computed tomography, fast-acquisition magnetic resonance imaging, and angiography. *Investigative Radiology*, 29:848–851, 1994.
- [111] S. Hunjan, G. Starkschall, I. Rosen, K. Prado, N. Tolani, and P. Balter. Comparison of breath-hold and free-breathing positions of an external fiducial by analysis of respiratory traces. *Journal of Applied Clinical Medical Physics*, 9:34–42, 2008.
- [112] A. Khamene, J.K. Warzelhan, S. Vogt, D. Elgort, C. Chefd’Hotel, J.L. Duerk, J. Lewin, F.K. Wacker, and F. Sauer. Characterization of internal organ motion



- using skin marker positions. In *Medical Image Computing and Computer Assisted Interventions*, volume 3217, pages 526–533. Lecture Notes in Computer Science, 2004.
- [113] D. Atkinson, M. Burcher, J. Declerek, and J.A. Noble. Respiratory motion compensation for 3-D freehand echocardiography. *Ultrasound in Medicine & Biology*, 27:1615–1620, 2001.
- [114] Q. Xu and R.J. Hamilton. A novel respiratory detection method based on automated analysis of diaphragm video. *Medical Physics*, 33:916–921, 2006.
- [115] H. Timinger, S. Krueger, K. Dietmayer, and J. Borgert. Motion compensated coronary interventional navigation of diaphragm tracking and elastic motion models. *Physics in Medicine and Biology*, 50:491–503, 2005.
- [116] D.A. Feinberg, D. Geise, D.A. Bongers, S. Ramanna, M. Zaitsev, M. Markl, and M. Gunther. Hybrid ultrasound MRI for improved cardiac imaging and real-time respiration control. *Magnetic Resonance in Medicine*, 63:290–296, 2010.
- [117] A.P. King, C. Jansen, K.S. Rhode, D. Caulfield, R.S. Razavi, and G.P. Penney. Respiratory motion correction for image-guided cardiac interventions using 3-D echocardiography. *Medical Image Analysis*, 14:21–29, 2010.
- [118] D.L.G. Hill, P.G. Batchelor, M. Holden, and D.J. Hawkes. Medical image registration. *Physics in Medicine and Biology*, 46:R1–R45, 2001.
- [119] B. Zitova and J. Flusser. Image registration methods: A survey. *Image and Vision Computing*, 21:977–1000, 2003.
- [120] W.R. Crum, T. Hartkens, and D.L.G. Hill. Non-rigid image registration: Theory and practice. *British Journal of Radiology*, 77:S140–S153, 2004.
- [121] T. Makela, P. Clarysse, O. Sipila, N. Pauna, Q.C. Pham, T. Katila, and I.E. Magnin. A review of cardiac image registration methods. *IEEE Transactions on Medical Imaging*, 21:1011–1021, 2002.

- [122] C. Ewertzen, K.R. Nielsen, B. Hesse, and M.B. Nielsen. Image fusion of diagnostic ultrasound with other modalities. *Current Medical Imaging Reviews*, 5:150–155, 2009.
- [123] A. Savi, M.C. Gilardi, G. Rizzo, M. Pepi, C. Landoni, C. Rossetti, G. Lucignani, A. Bartorelli, and F. Fazio. Spatial registration of echocardiographic and positron emission tomographic heart studies. *European Journal of Nuclear Medicine*, 22:243–247, 1995.
- [124] Q. Duan, G. Shechter, L.F. Gutierrez, D. Stanton, L. Zagorchev, A.F. Laine, and D.R. Elgort. Augmenting CT cardiac roadmaps with segmented streaming ultrasound. In *Medical Imaging: Visualization and Image-Guided Procedures*, volume 6509 of *SPIE 2007*, page 65090V, 2007.
- [125] K.Y.E. Leung, M. van Stralen, A. Nemes, M.M. Voormolen, G. van Burken, M.L. Geleijnse, F.J. ten Cate, J.H.C. Reiber, N. de Jong, A. F.W. van der Steen, and J.G. Bosch. Sparse registration for three-dimensional stress echocardiography. *IEEE Transactions on Medical Imaging*, 27:1568–1579, 2008.
- [126] Y. Okumura, B.D. Henz, S.B. Johnson, J. Bunch, C.J. O’Brien, D.O. Hodge, A. Altman, A. Govari, and D.L. Packer. Three dimensional ultrasound for image-guided mapping and intervention: Methods, quantitative validation and clinical feasibility of a novel multimodality image mapping system. *Circulation: Arrhythmia and Electrophysiology*, 1:110–119, 2008.
- [127] C.A. Linte, J. Moore, C. Wedlake, D. Bainbridge, G.M. Guiraudon, D.L. Jones, and T.M. Peters. Inside the beating heart: An in vivo feasibility study on fusing pre- and intra-operative imaging for minimally-invasive therapy. *International Journal of Computer Assisted Radiology and Surgery*, 4:113–123, 2009.
- [128] Y.L. Ma, G.P. Penney, C.A. Rinaldi, M. Cooklin, R. Razavi, and K.S. Rhode. Echocardiography to magnetic resonance image registration for use in image-guided cardiac catheterization procedures. *Physics in Medicine and Biology*, 54:5039–5055, 2009.

- [129] M.E. Rettmann, D.R. Holmes III, Y. Su, B.M. Cameron, J.J. Camp, D.L. Packer, and R.A. Robb. An integrated system for real-time image guided cardiac catheter ablation. In *Medicine Meets Virtual Reality*, volume 14, pages 455–460. IOS Press, 2006.
- [130] Z.J. Malchano, P. Neuzil, R.C. Cury, G. Holmvang, J. Weichet, E.J. Schmidt, J.N. Ruskin, and V.Y. Reddy. Integration of cardiac CT/MR imaging with three-dimensional electroanatomical mapping to guide catheter manipulation in the left atrium: Implications for catheter ablation of atrial fibrillation. *Journal of Cardiovascular Electrophysiology*, 17:1221–1229, 2006.
- [131] L.F. Tops, J.J. Bax, K. Zeppenfeld, M.R.M. Jongbloed, H.J. Lamb, E.E. van der Vall, and M.J. Schalij. Fusion of multislice computed tomography imaging with three-dimensional electroanatomical mapping to guide radiofrequency catheter ablation procedures. *Heart Rhythm*, 2:1076–1081, 2005.
- [132] D.W. den Uijl, L.F. Tops, J.M. Tolosana, J.D. Schuijf, S.A.I.P. Trines, K. Zeppenfeld, J.J. Bax, and M.J. Schalij. Real-time integration of intracardiac echocardiography and multislice computed tomography to guide radiofrequency catheter ablation for atrial fibrillation. *Heart Rhythm*, 5:1403–1410, 2008.
- [133] Y. Khaykin, A. Skanes, B. Whaley, C. Hill, M. Beardsall, C. Seabrook, Z. Wulffhart, R. Oosthuizen, L. Gula, and A. Verma. Real-time integration of 2D intracardiac echocardiography and 3D electroanatomical mapping to guide ventricular tachycardia ablation. *Heart Rhythm*, 5:1396–1402, 2008.
- [134] H. Zhong, T. Kanade, and D. Schwartzman. “Virtual touch”: An efficient registration method for catheter navigation in left atrium. In *Medical Image Computing and Computer Assisted Interventions*, volume 4190, pages 437–444. Lecture Notes in Computer Science, 2006.
- [135] A.D. Koolwal, F. Barbagli, C.R. Carlson, and D.H. Liang. An incremental method for registering electroanatomical mapping data to surface mesh models

- of the left atrium. In *Medical Image Computing and Computer Assisted Interventions*, volume 5242, pages 847–854. Lecture Notes in Computer Science, 2008.
- [136] W. Wein, E. Camus, M. John, M. Diallo, C. Duong, A. Al-Ahmad, R. Fahrig, A. Khamene, and C. Xu. Towards guidance of electrophysiological procedures with real-time 3D intracardiac echocardiography fusion to C-arm CT. In *Medical Image Computing and Computer Assisted Interventions*, volume 5761, pages 9–16. Lecture Notes in Computer Science, 2009.
- [137] W. Zhang, J.A. Noble, and J.M. Brady. Real time 3-D ultrasound to MR cardiovascular image registration using a phase-based approach. In *IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pages 666–669, 2006.
- [138] W. Zhang, J.A. Noble, and J.M. Brady. Adaptive non-rigid registration of real time 3D ultrasound to cardiovascular MR images. In *Image Processing in Medical Imaging*, volume 4584, pages 50–61. Lecture Notes in Computer Science, 2007.
- [139] W. Zhang, J.A. Noble, and J.M. Brady. Spatio-temporal registration of real time 3D ultrasound to cardiovascular MR sequences. In *Medical Image Computing and Computer Assisted Interventions*, volume 4791, pages 343–350. Lecture Notes in Computer Science, 2007.
- [140] A. Roche, X. Pennec, G. Malandain, and N. Ayache. Rigid registration of 3-D ultrasound with MR images: A new approach combining intensity and gradient information. *IEEE Transactions on Medical Imaging*, 20:1038–1049, 2001.
- [141] W. Wein, S. Brunke, A. Khamene, M.R. Callstrom, and N. Navab. Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Medical Image Analysis*, 12:577–585, 2008.
- [142] J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever. Mutual-information-based

- registration of medical images: A survey. *IEEE Transactions on Medical Imaging*, 22:986–1004, 2003.
- [143] F. Maes, A. Collignon, D. Vandermeulen, G. Marchall, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16:187–198, 1997.
- [144] W.M. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multimodal volume registration by maximization of mutual information. *Medical Image Analysis*, 1:35–51, 1996.
- [145] M. Mellor and M. Brady. Phase mutual information as a similarity measure for registration. *Medical Image Analysis*, 9:330–343, 2005.
- [146] J.F. Krucker, G.L. LeCarpentier, J.B. Fowlkes, and P.L. Carson. Rapid elastic image registration for 3-D ultrasound. *IEEE Transactions on Medical Imaging*, 21:1384–1394, 2002.
- [147] P. Foroughi, P. Abolmaesumi, and K. Hashtrudi-Zaad. Intra-subject elastic registration of 3D ultrasound images. *Medical Image Analysis*, 10:713–725, 2006.
- [148] N. Herlambang, H. Liao, K. Matsumiya, K. Masamune, and T. Dohi. Real-time autostereoscopic visualization of registration-generated 4D MR image of beating heart. In *Medical Imaging and Augmented Reality*, volume 5128, pages 349–358. Lecture Notes in Computer Science, 2008.
- [149] A. Khamene, S. Vogt, F. Azar, T. Seilhorst, F. Sauer, and H. Niemann. Local 3D reconstruction and augmented reality visualization of free-hand ultrasound for needle biopsy procedures. In *Medical Image Computing and Computer Assisted Interventions*, volume 2879, pages 344–355. Lecture Notes in Computer Science, 2003.
- [150] K. Konishi, M. Nakamoto, Y. Kakeji, K. Tanoue, H. Kawanaka, S. Yamaguchi, S. Ieiri, Y. Sato, Y. Maehara, S. Tamura, and M. Hashizume. A real-time nav-

- igation system for laparoscopic surgery based on three-dimensional ultrasound using magneto-optic hybrid tracking configuration. *International Journal of Computer Assisted Radiology and Surgery*, 2:1–10, 2007.
- [151] N. Suzuki, A. Hattori, S. Suzuki, and Y. Otake. Development of a surgical robot system for endovascular surgery with augmented reality function. In *Medicine Meets Virtual Reality*, volume 15, pages 460–463. IOS Press, 2007.
- [152] K. Masamune, T. Horiuchi, M. Mizutani, H. Yamashita, H. Tsukihara, N. Motomura, S. Takamoto, H. Liao, and T. Dohi. Novel endoscope system with plasma flushing for off-pump cardiac surgery. In *Medical Image Computing and Computer Assisted Interventions*, volume 5761, pages 451–458. Lecture Notes in Computer Science, 2009.
- [153] N.V. Vasilyev, J.F. Martinez, F.P. Freudenthal, Y. Suematsu, G.R. Marx, and P.J. del Nido. Three-dimensional echo and videocardioscopy-guided atrial septal defect closure. *Annals of Thoracic Surgery*, 82:1322–1326, 2006.
- [154] J. Leven, D. Burschka, R. Kumar, G. Zhang, S. Blumenkranz, X.D. Dai, M. Awad, G.D. Hager, M. Marohn, C. Choti, C. Hasser, and R.H. Taylor. DaVinci Canvas: A telerobotic surgical system with integrated, robot-assisted, laparoscopic ultrasound capability. In *Medical Image Computing and Computer Assisted Interventions*, volume 3749, pages 811–818. Lecture Notes in Computer Science, 2005.
- [155] S.L. Lee, M. Lerotic, V. Vitiello, S. Giannarou, K.W. Kwok, M. Visentini-Scarzanella, and G.Z. Yang. From medical images to minimally invasive intervention: Computer assistance for robotic surgery. *Computerized Medical Imaging and Graphics*, 34:33–45, 2010.
- [156] S.P. Emery, J. Kreutzer, F.R. Sherman, K.L. Fujimoto, B. Jaramaz, C. Nikou, K. Tobita, and B.R. Keller. Computer-assisted navigation applied to fetal cardiac intervention. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 3:187–198, 2007.

- [157] M. Hastenteufel, S. Yang, C. Christoph, M. Vetter, H.P. Meinzer, and I. Wolf. Image-based guidance for minimally invasive surgical atrial fibrillation ablation. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 2(1):60–69, 2006.
- [158] C.A. Linte, J. Moore, A.D. Wiles, C. Wedlake, and T.M. Peters. Targeting accuracy under model-to-subject misalignments in model-guided cardiac surgery. In *Medical Image Computing and Computer Assisted Interventions*, volume 5761, pages 361–368. Lecture Notes in Computer Science, 2009.
- [159] R.A. Robb. Medical imaging and virtual reality: A personal perspective. *Virtual Reality*, 12:235–257, 2008.
- [160] K. Wilson, G. Guiraudon, D.L. Jones, and T.M. Peters. Mapping of cardiac electrophysiology onto a dynamic patient-specific heart model. *IEEE Transactions on Medical Imaging*, 28:1870–1880, 2009.
- [161] V. Rasche, M. Mansour, V. Reddy, J.P. Singh, A. Qureshi, R. Manzke, S. Sokka, and J. Ruskin. Fusion of three-dimensional X-ray angiography and three-dimensional echocardiography. *International Journal of Computer Assisted Radiology and Surgery*, 2:293–303, 2008.
- [162] A. Jain, L. Gutierrez, and D. Stanton. 3D TEE registration with X-ray fluoroscopy for interventional cardiac applications. In *Functional Imaging and Modeling of the Heart*, volume 5528, pages 321–329. Lecture Notes in Computer Science, 2009.
- [163] Y.L. Ma, G.P. Penney, D. Box, B. Frissen, G. de Fockert, C. Yao, A. King, G. Gao, C.A. Rinaldi, R. Razavi, and K.S. Rhode. Using a robotic arm for echocardiography to X-ray image registration during cardiac catheterization procedures. In *Workshop on Cardiovascular Interventional Imaging and Biophysical Modelling*, volume Medical Image Computing and Computer Assisted Interventions, 2009.

- [164] L. Gepstein and S.J. Evans. Electroanatomical mapping of the heart: Basic concepts and implications for the treatment of cardiac arrhythmias. *Pacing and Clinical Electrophysiology*, 21:1268–1278, 1998.
- [165] L. Gepstein, G. Hayam, and S.A. Ben-Haim. A novel method for nonfluoroscopic catheter-based electroanatomical mapping of the heart. *Circulation*, 95:1611–1622, 1997.
- [166] P. Jannin, J.M. Fitzpatrick, D.J. Hawkes, X. Pennec, R. Shahidi, and M.W. Vannier. Validation of medical image processing in image-guided therapy. *IEEE Transactions on Medical Imaging*, 21:1445–1449, 2002.
- [167] S. DiMaio, T. Kapur, K. Cleary, S. Aylward, P. Kazanzides, K. Vosburgh, R. Ellis, J. Duncan, K. Farahani, H. Lemke, T. Peters, W. Lorensen, D. Gobbi, J. Haller, L. Clarke, S. Pizer, R. Taylor, R. Galloway Jr., G. Fichtinger, N. Hata, K. Lawson, C. Tempany, R. Kikinis, and F. Jolesz. Challenges in image-guided therapy system design. *NeuroImage*, 37:S144–S151, 2007.
- [168] M. Janda and B. Buch. The challenges of clinical validation of emerging technologies: Computer-assisted devices for surgery. *The Journal of Bone and Joint Surgery*, 91:17–21, 2009.
- [169] T. Sielhorst, M. Feuerstein, and N. Navab. Advanced medical displays: A literature review of augmented reality. *Journal of Display Technology*, 4:451–467, 2008.
- [170] J. Lo, J. Moore, C. Wedlake, G. Guiraudon, R. Eagleson, and T. Peters. Surgeon-controlled visualization techniques for virtual reality-guided cardiac surgery. In *Medicine Meets Virtual Reality*, volume 17, pages 162–167. IOS Press, 2009.
- [171] G.D. Stetten and V.S. Chib. Overlaying ultrasonographic images on direct vision. *Journal of Ultrasound in Medicine*, 20:235–240, 2001.



- [172] N.Y. Vasilyev, P.M. Novotny, J.F. Martinez, H. Loyola, I.S. Salgo, R.D. Howe, and P.J. del Nido. Stereoscopic vision display technology in real-time three-dimensional echocardiography-guided intracardiac beating-heart surgery. *The Journal of Thoracic and Cardiovascular Surgery*, 135:1334–1341, 2008.
- [173] R.T. O’Brien and S.P. Holmes. Recent advances in ultrasound technology. *Clinical Techniques in Small Animal Practice*, 22:93–103, 2007.
- [174] F. Fosberg. Ultrasonic biomedical technology; marketing versus clinical reality. *Ultrasonics*, 42:17–27, 2004.
- [175] C.R. Castro-Pareja, J.M. Jagadeesh, and R. Shekhar. FAIR: A hardware architecture for real-time 3-D image registration. *IEEE Transactions on Information Technology in Biomedicine*, 7:426–434, 2003.
- [176] T. Rohlfing and C.R. Maurer Jr. Nonrigid image registration in shared-memory multiprocessor environments with application to brains, breasts, and bees. *IEEE Transactions on Information Technology in Biomedicine*, 7:16–25, 2003.
- [177] J. Huang, P.E. Dupont, A. Undurti, J.K. Triedman, and R.O. Cleveland. Producing diffuse ultrasound reflections from medical instruments using a quadratic residue diffuser. *Ultrasound in Medicine & Biology*, 32:721–727, 2006.
- [178] K. Nichols, L.B. Wright, T. Spencer, and W.C. Culp. Changes in ultrasonographic echogenicity and visibility of needles with changes in angles of insonation. *Journal of Vascular and Interventional Radiology*, 14:1553–1557, 2003.
- [179] J. Huang, J.K. Triedman, N.V. Vasilyev, Y. Suematsu, R.O. Cleveland, and P.E. Dupont. Imaging artifacts of medical instruments in ultrasound-guided interventions. *Journal of Ultrasound in Medicine*, 26:1303–1322, 2007.
- [180] K.L. Gentry and S.W. Smith. Integrated catheter for 3-D intracardiac echocardiography and ultrasound ablation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51:799–807, 2004.

- [181] J.T. Moore, A.D. Wiles, C. Wedlake, D. Bainbridge, B. Kiaii, A.L. Trejos, R. Patel, and T.M. Peters. Integration of trans-esophageal echocardiography with magnetic tracking technology for cardiac interventions. In *Medical Imaging: Visualization, Image-Guided Procedures, and Modeling*, volume 7625 of *SPIE 2010*, page 7625Y, 2010.
- [182] T. Wilson, J. Zagzebski, T. Varghese, Q. Chen, and M. Rao. The Ultrasonix 500RP: A commercial ultrasound research interface. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 53:1772–1782, 2006.
- [183] R. Rohling, W. Fung, and P. Lajevardi. PUPIL: Programmable Ultrasound Platform and Interface Library. In *Medical Image Computing and Computer Assisted Interventions*, volume 2879, pages 424–431. Lecture Notes in Computer Science, 2003.
- [184] V. Shamdasani, U. Bae, S. Sikdar, Y.M. Yoo, K. Karadayi, R. Managuli, and Y. Kim. Research interface on a programmable ultrasound scanner. *Ultrasonics*, 48:159–168, 2008.
- [185] S.W. Smith, H.G. Pavy Jr., and O.T. von Ramm. High-speed ultrasound volumetric imaging system - Part I: Transducer design and beam steering. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 38:100–108, 1991.
- [186] O.T. von Ramm, S.W. Smith, and H.G. Pavy Jr. High-speed ultrasound volumetric imaging system - Part II: Parallel processing and image display. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 38:109–115, 1991.
- [187] I.A. Rasmussen Jr., F. Lindseth, O.M. Rygh, E.M. Berntsen, T. Selbekk, J. Xu, T.A. Nagelhus Hernes, E. Harg, A. Haberg, and G. Unsgaard. Functional neuronavigation combined with intra-operative 3D ultrasound: Initial experiences during surgical resections close to eloquent brain areas and future directions in

- automatic brain shift compensation of preoperative data. *Acta Neurochirurgica*, 149:365–378, 2007.
- [188] E.M. Boctor, M.A. Choti, E.C. Burdette, and R.J. Webster III. Three-dimensional ultrasound-guided robotic needle placement: An experimental evaluation. *International Journal of Medical Robotics and Computer Assisted Surgery*, 4:180–191, 2008.
- [189] Z. Wei, L. Gardi, C. Edirisinghe, D. Downey, and A. Fenster. Three-dimensional ultrasound guidance and robot assistance for prostate brachytherapy. In T.M. Peters and K. Cleary, editors, *Image-Guided Interventions: Technology and Applications*, pages 429–460. Springer US, 2008.
- [190] A. Gee, R. Prager, G. Treece, and L. Berman. Engineering a freehand 3D ultrasound system. *Pattern Recognition Letters*, 24:757–777, 2003.
- [191] K.K. Djoa, N. De Jong, F.C. van Egmond, J.D. Kasprzak, W.B. Vletter, C.T. Lancee, A.F. van der Steen, N. Bom, and J.R.T.C. Roelandt. A fast rotating scanning unit for real-time three-dimensional echo data acquisition. *Ultrasound in Medicine & Biology*, 26:863–869, 2000.
- [192] J. Blancher, C. Leger, and L.D. Nguyen. Time-varying 3-D echocardiography using a fast-rotating probe. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51:634–639, 2004.
- [193] R.J. Housden, A.H. Gee, G.M. Treece, and R.W. Prager. Sensorless reconstruction of unconstrained freehand 3D ultrasound data. *Ultrasound in Medicine & Biology*, 33:408–419, 2007.
- [194] F. Lindseth, G.A. Tangen, T. Lango, and J. Bang. Probe calibration for freehand 3-D ultrasound. *Ultrasound in Medicine & Biology*, 29:1607–1623, 2003.
- [195] F. Rousseau, P. Hellier, M.M.J. Letteboer, W.J. Niessen, and C. Barillot. Quantitative evaluation of three calibration methods for 3-D freehand ultrasound. *IEEE Transactions on Medical Imaging*, 25:1492–1501, 2006.

- [196] R.W. Prager, A. Gee, and L. Berman. Stradx: Real-time acquisition and visualization of freehand three-dimensional ultrasound. *Medical Image Analysis*, 3:129–140, 1998.
- [197] G.M. Treece, A.H. Gee, R.W. Prager, C.J.C. Cash, and L.H. Berman. High-definition freehand 3-D ultrasound. *Ultrasound in Medicine & Biology*, 29:529–546, 2003.
- [198] Ole Vegard Solberg, Frank Lindseth, Hans Torp, Richard E. Blake, and Toril A. Nagelhus Hernes. Freehand 3D ultrasound reconstruction algorithms—a review. *Ultrasound in Medicine & Biology*, 33(7):991–1009, 2007.
- [199] A. Salustri and J.R.T.C. Roelandt. Ultrasonic three-dimensional reconstruction of the heart. *Ultrasound in Medicine & Biology*, 21:281–293, 1995.
- [200] C.D. Barry, C.P. Allott, N.W. John, P.M. Mellor, P.A. Arundel, D.S. Thomson, and J.C. Waterton. Three-dimensional freehand ultrasound: Image reconstruction and volume analysis. *Ultrasound in Medicine & Biology*, 23:1209–1224, 1997.
- [201] S. Berg, H. Torp, D. Martens, E. Steen, S. Samstad, I. Hoivik, and B. Olstad. Dynamic three-dimensional freehand echocardiography using raw digital ultrasound data. *Ultrasound in Medicine & Biology*, 25:745–753, 1999.
- [202] S. Meairs, J. Beyer, and M. Hennerici. Reconstruction and visualization of irregularly sampled three- and four- dimensional ultrasound data for cerebrovascular applications. *Ultrasound in Medicine & Biology*, 26:263–272, 2000.
- [203] A. Karamalis, W. Wein, O. Kutter, and N. Navab. Fast hybrid freehand ultrasound volume reconstruction. In *Medical Imaging: Visualization, Image-guided Procedures and Modeling*, volume 7261 of *SPIE 2009*, page 726114, 2009.
- [204] R. Ohbuchi, D. Chen, and H. Fuchs. Incremental volume reconstruction and rendering for 3D ultrasound imaging. In *Visualization in Biomedical Computing*, volume 1808 of *SPIE 1992*, pages 312–323, 1992.

- [205] W.S. Edwards, C. Deforge, and Y. Kim. Interactive three-dimensional ultrasound using a programmable multimedia processor. *International Journal of Imaging Systems and Technology*, 9:442–454, 1998.
- [206] J.N. Welch, J.A. Johnson, M.R. Bax, R. Badr, and R. Shahidi. A real-time freehand 3D ultrasound system for image-guided surgery. In *IEEE Ultrasonics Symposium*, pages 1601–1604, 2000.
- [207] D.G. Gobbi and T.M. Peters. Interactive intra-operative 3D ultrasound reconstruction and visualization. In *Medical Image Computing and Computer Assisted Interventions*, volume 2489, pages 156–163. Lecture Notes in Computer Science, 2002.
- [208] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit: An object-oriented approach to 3D graphics*. Kitware Inc., 4th edition, 2006.
- [209] D.F. Pace, D.G. Gobbi, C. Wedlake, J. Gumprecht, J. Boisvert, J. Tokuda, N. Hata, and T.M. Peters. An open-source real-time ultrasound reconstruction system for four-dimensional imaging of moving organs. In *Workshop on Systems and Architectures for CAI*, volume Medical Image Computing and Computer Assisted Interventions, 2009.
- [210] J. Tokuda, G.S. Fischer, X. Papademetris, Z. Yaniv, L. Ibanez, P. Cheng, H. Lui, J. Blevins, J. Arata, A. Golby, T. Kapur, S. Pieper, E.C. Burdette, G. Fichtinger, C.M. Tempany, and N. Hata. OpenIGTLink: An open network protocol for image-guided therapy environment. *International Journal of Medical Robotics and Computer Assisted Surgery*, 5:423–434, 2009.
- [211] D.G. Gobbi. *Brain deformation correction using interactive 3D ultrasound imaging*. PhD thesis, Department of Medical Biophysics, The University of Western Ontario, London, ON, Canada, 2003.
- [212] K. Shoemake. Animating rotations with quaternion curves. In *ACM SIGGRAPH Computer Graphics*, volume 19, pages 245–254, 1985.

- [213] D.G. Gobbi, R.M. Comeau, and T.M. Peters. Ultrasound probe tracking for real-time ultrasound/MRI overlay and visualization of brain shift. In *Medical Image Computing and Computer Assisted Interventions*, volume 1679, pages 920–927. Lecture Notes in Computer Science, 1999.
- [214] A.D. Wiles, J. Moore, C.A. Linte, C. Wedlake, A. Ahmad, and T.M. Peters. Object identification accuracy under ultrasound enhanced virtual reality for minimally invasive cardiac surgery. In *Medical Imaging: Visualization, Image-Guided Procedures, and Modeling*, volume 6918 of *SPIE 2008*, page 39180E, 2008.
- [215] C.A. Linte, A. Wiles, J. Moore, C. Wedlake, and T.M. Peters. Virtual reality-enhanced ultrasound guidance for atrial ablation: in vitro epicardial study. In *Medical Image Computing and Computer Assisted Interventions*, volume 5242, pages 644–651. Lecture Notes in Computer Science, 2008.
- [216] C.A. Linte, J. Moore, A. Wiles, J. Lo, C. Wedlake, and T.M. Peters. In vitro cardiac catheter navigation via augmented reality surgical guidance. In *Medical Imaging: Visualization, Image-Guided Procedures, and Modeling*, volume 7261 of *SPIE 2009*, page 72610O1, 2009.
- [217] G.M. Guiraudon, D.L. Jones, D. Bainbridge, and T.M. Peters. Mitral valve implantation using off-pump closed beating intracardiac surgery: A feasibility study. *Interactive Cardiovascular and Thoracic Surgery*, 6(5):603–607, 2007.
- [218] D. Debrun, F. Therain, L.D. Nguyen, C.P. Leger, J.J.N. Visser, and E. Busemann-Sokole. Volume measurements in nuclear gated SPECT and 4D echocardiography: Validation using a dynamic cardiac phantom. *International Journal of Cardiovascular Imaging*, 21:239–247, 2005.
- [219] D.C. Barratt, A.H. Davies, A.D. Hughes, S.A. Thom, and K.N. Humphries. Accuracy of an electromagnetic three-dimensional ultrasound system for carotid artery imaging. *Ultrasound in Medicine & Biology*, 27:1421–1425, 2001.

- [220] D.F. Pace, R. Kikinis, and N. Hata. An accessible, hands-on tutorial system for image-guided therapy and medical robotics using a robot and open source software. In *Workshop on Open Source and Open Data*, volume Medical Image Computing and Computer Assisted Interventions, 2007.
- [221] A.B. Forbes. Least-squares best-fit geometric elements. Report NPL Report DITC 140/89, National Physical Laboratory, 1989.
- [222] Q. Li and J.G. Griffiths. Least squares ellipsoid specific fitting. In *Geometric Modeling and Processing*, pages 335–340. IEEE Computer Society, 2004.
- [223] P. Schneider and D.H. Eberly. *Geometric tools for computer graphics*. Morgan Kaufmann, 2002.
- [224] D. Le Bihan, J.-F. Mangin, C. Poupon, C.A. Clark, S. Pappata, N. Molko, and H. Chabriat. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging*, 13:534–546, 2001.
- [225] C.F. Westin, S.E. Maier, B. Khidhir, P. Everett, F.A. Jolesz, and R. Kikinis. Image processing for diffusion tensor magnetic resonance imaging. In *Medical Image Computing and Computer-Assisted Interventions*, volume 1679, pages 441–452. Lecture Notes in Computer Science, 1999.
- [226] S. Brekke, S.I. Rabben, A. Stoylen, A. Haugen, G.U. Haugen, E.N. Steen, and H. Torp. Volume stitching in three-dimensional echocardiography: Distortion analysis and extension to real time. *Ultrasound in Medicine & Biology*, 33:782–796, 2007.
- [227] S.K. Nadkarni, D.K. Boughner, M. Drangova, and A. Fenster. In vitro simulation and quantification of temporal jitter artifacts in ECG-gated dynamic three-dimensional echocardiography. *Ultrasound in Medicine & Biology*, 27:211–222, 2001.
- [228] J.M. Blackall, D. Rueckert, C.R. Maurer, G.P. Penney, D.L.G. Hill, and D.J. Hawkes. An image registration approach to automated calibration for freehand

- 3D ultrasound. In *Medical Image Computing and Computer Assisted Interventions*, volume 1935, pages 462–471. Lecture Notes in Computer Science, 2000.
- [229] D.M. Muratore and R.L. Galloway. Beam calibration without a phantom for creating a 3-D freehand ultrasound system. *Ultrasound in Medicine & Biology*, 27:1557–1566, 2001.
- [230] N. Pagoulatos, D.R. Haynor, and Y. Kim. A fast calibration method for 3-D tracking of ultrasound images using a spatial localizer. *Ultrasound in Medicine & Biology*, 27:1219–1229, 2001.
- [231] R. Rohling, A. Gee, and L. Berman. A comparison of freehand three-dimensional ultrasound reconstruction techniques. *Medical Image Analysis*, 3:339–359, 1999.
- [232] R. Brooks, D.L. Collins, X. Morandi, and T. Arbel. Deformable ultrasound registration without reconstruction. In *Medical Image Computing and Computer-Assisted Interventions*, volume 5242, pages 1023–1031. Lecture Notes in Computer Science, 2008.
- [233] J. Moore, C. Clarke, D. Bainbridge, C. Wedlake, A. Wiles, D. Pace, and T. Peters. Image guidance for spinal facet injections using tracked ultrasound. In *Medical Image Computing and Computer Assisted Interventions*, volume 5761, pages 516–523. Lecture Notes in Computer Science, 2009.
- [234] S. Gill, P. Mousavi, G. Fichtinger, E. Chen, J. Boisvert, D. Pichora, and P. Abolmaesumi. Biomechanically constrained groupwise US to CT registration of the lumbar spine. In *Medical Image Computing and Computer Assisted Interventions*, volume 5761, pages 803–810. Lecture Notes in Computer Science, 2009.
- [235] A.L. Trejos, R.V. Patel, I. Ross, and B. Kiaii. Optimizing port placement for robot-assisted minimally invasive cardiac surgery. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 3:355–364, 2007.



- [236] S. Pieper, M. Halle, and R. Kikinis. 3D Slicer. In *IEEE International Symposium on Biomedical Engineering*, pages 632–635, 2004.
- [237] K. Cleary, P. Cheng, A. Enquobahrie, and Z. Yaniv, editors. *IGSTK Image-Guided Surgery Toolkit: An open source C++ software library*. Signature Book Printing, 4.2 edition, 2009.
- [238] J. Boisvert, D. Gobbi, S. Vikal, R. Rohling, G. Fichtinger, and P. Abolmaesumi. An open-source solution for interactive acquisition, processing and transfer of interventional ultrasound images. In *Workshop on Systems and Architectures for Computer Assisted Interventions*, volume Medical Image Computing and Computer Assisted Interventions, 2008.

## Vita

**Name:** Danielle F. Pace

**Post-secondary Education and Degrees:** M.E.Sc. Biomedical Engineering Candidate  
The University of Western Ontario, London, Ontario  
2007-2010

B.Cmp.H. Biomedical Computing  
Queen's University, Kingston, Ontario  
2003-2007

**Honours and Awards:** Canada Graduate Scholarship M  
Natural Sciences and Engineering Research Council of Canada  
2007-2009

Western Graduate Research Scholarship  
The University of Western Ontario  
2007-2009

“People’s Choice” Advanced Undergraduate Project Award  
School of Computing, Queen’s University  
2007

Best Undergraduate Contribution Award  
Canadian Student Conference on Biomedical Computing  
2006

Undergraduate Scholarship  
HSBC Bank Malta  
2003-2007

Principal’s Scholarship  
Queen’s University  
2003-2004

**Related Work Experience:** Teaching Assistant  
The University of Western Ontario, London, Ontario  
2007, 2009

Training Specialist  
National Center for Image-Guided Therapy  
2008

Research Assistant  
Surgical Planning Laboratory, Boston, Massachusetts  
2007

Research Assistant, Department of Physiology  
Queen's University, Kingston, Ontario  
2005-2006

### **Conference Proceedings and Presentations:**

D.F. Pace, D.G. Gobbi, C. Wedlake, J. Gumprecht, J. Boisvert, J. Tokuda, N. Hata and T.M. Peters. "An open-source real-time ultrasound reconstruction system for four-dimensional imaging of moving organs". Workshop on Systems and Architectures for Computer Assisted Interventions, Medical Image Computing and Computer Assisted Intervention (MICCAI), London, UK, September 24, 2009.

J. Moore, C. Clarke, D. Bainbridge, C. Wedlake, A. Wiles, D. Pace and T. Peters. "Image guidance for spinal facet injections using tracked ultrasound". Medical Image Computing and Computer Assisted Intervention (MICCAI), London, UK, September 20-24, 2009.

D.F. Pace, A.D. Wiles, J. Moore, C. Wedlake, D.G. Gobbi and T.M. Peters. "Validation of four-dimensional ultrasound for targeting in minimally-invasive beating-heart surgery". SPIE Medical Imaging, Orlando, FL, USA, February 7-12, 2009.

J. Jomier, L. Ibanez, A. Enquobahrie, D.F. Pace and K. Cleary. "An open-source testing framework for tracking devices using Lego Mindstorms". SPIE Medical Imaging, Orlando, FL, USA, February 7-12, 2009.

D.F. Pace, R. Kikinis and N. Hata. "An accessible, hands-on tutorial system for image-guided therapy and medical robotics using a robot and open source software". Workshop on Open Source and Open Data, Medical Image Computing and Computer Assisted Intervention (MICCAI), Brisbane, Australia, November 2, 2007.

D. F. Pace, T. Bui and P.K. Rose. "Computational estimates of the effect of asynchronous synaptic activity on fluctuations in the membrane potential of motoneurons". Society for Neuroscience, Atlanta, GA, USA, October 14-18, 2006.