

SURGICAL NAVIGATION USING TRACKED ULTRASOUND

By

Thomas Steven Pheiffer

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Biomedical Engineering

August, 2014

Nashville, Tennessee

Approved:

Michael I. Miga, Ph.D.

Brett C. Byram, Ph.D.

Benoit M. Dawant, Ph.D.

Robert L. Galloway Jr., Ph.D.

Reid C. Thompson, M.D.

ACKNOWLEDGEMENTS

There are many people without whom I would not have been able to finish this work. First of all I would like to thank my advisor and mentor Dr. Michael Miga, for all of his invaluable help and guidance throughout my time at Vanderbilt, not only with regard to scientific research but also career development. My appreciation also goes out to the rest of my committee members, including Dr. Benoit Dawant, Dr. Bob Galloway, Dr. Reid Thompson, and Dr. Brett Byram for their helpful comments and suggestions.

I also need to thank the members of the Biomedical Modeling Laboratory and SNARL lab groups, past and present, who helped at all stages of this work and without whom my life would have been much more difficult. In particular I would like to thank Dr. Jao Ou, Dr. David Kwartowitz, Dr. Prashanth Dumpuri, Dr. Logan Clements, Dr. Kay Sun, and Dr. Amber Simpson for their mentorship during various times of this process. My thanks also go out to Yifei Wu, Rebekah Conley, and Dr. Mike Delisi for their assistance with experiments and their camaraderie over the past several years.

Most of all, I would like to thank my family and friends for their love and support: my mom and dad, for instilling the value of education and always encouraging me to fulfill my aspirations; my brothers – Bill and Dave – for all the good times while we were growing up; Lindsay, for being at my side the last five years as we each attempted to maintain our sanity during grad school; the various members of the BME coffee break crew, past and present, who gave me something to look forward to each morning; and lastly, my friend Dr. Dean Paras, who was taken from us far too early. I owe Dean so much for his friendship and his influence on my life while I had the privilege to know him at Vanderbilt.

Finally, I would like to acknowledge the funding source which supported me and this project over the past several years: National Institutes of Health grant R01 NS049251 of the National Institute for Neurological Disorders and Stroke.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	ii
LIST OF TABLES	viii
LIST OF FIGURES	x
LIST OF ABBREVIATIONS	xvi
 Chapter	
I. INTRODUCTION	1
I.1 Specific Aims	1
I.1.1 AIM 1: Characterize the accuracy of a novel laser range scanner	2
I.1.2 AIM 2: Develop a patient-specific compression correction method for tracked ultrasound	2
I.1.3 AIM 3: Develop a generalized real-time compression correction method for tracked ultrasound	3
I.2 Dissertation Overview	3
II. BACKGROUND	4
II.1 Introduction to Image-Guided Interventions	4
II.2 Ultrasound Imaging	5
II.3 Tracking Systems	10
II.4 Registration Methods	11
II.5 Model-Updated IGS	13
III. METHODOLOGY	15
III.1 Image Segmentation	15
III.1.1 Preoperative Tomogram Segmentation	15
III.1.2 Intraoperative Ultrasound Segmentation	16
III.2 Model Construction	17
III.3 Tracked Ultrasound	19
III.3.1 Tracked Ultrasound Hardware	19
III.3.2 Tracked Ultrasound Software	21
III.3.3 Spatial Calibration	23
III.3.4 Temporal Calibration	27
III.3.5 Registration	30
III.3.6 Strain Imaging	30
III.3.7 Phantom Construction	37

IV.	DESIGN AND EVALUATION OF AN OPTICALLY-TRACKED SINGLE-CCD LASER RANGE SCANNER ...	40
	IV.1 Introduction and Significance of Study	40
	IV.2 Abstract.....	40
	IV.3 Introduction	41
	IV.4 Materials and Methods.....	44
	IV.4.1 Design and Development.....	44
	IV.4.2 Laser Range Scanner Evaluation	48
	IV.5 Results.....	52
	IV.6 Discussion	54
	IV.7 Conclusions	55
V.	MODEL-BASED CORRECTION OF TISSUE COMPRESSION FOR TRACKED ULTRASOUND IN SOFT- TISSUE IMAGE-GUIDED SURGERY	56
	V.1 Introduction and Significance of Study	56
	V.2 Abstract.....	57
	V.3 Introduction	57
	V.4 Materials and Methods.....	59
	V.4.1 Phantom Construction	59
	V.4.2 Patient Model from Preoperative Image Volume.....	60
	V.4.3 Intraoperative Data Collection.....	61
	V.4.4 Proposed Compression Error Correction	62
	V.4.5 Rigid Correction.....	64
	V.4.6 Model-Based Correction	64
	V.4.7 Experimental Validation.....	67
	V.4.7.1 Simulations of Proposed Correction	67
	V.4.7.2 Phantom Experiments.....	70
	V.4.7.3 Clinical Case.....	73
	V.5 Results.....	73
	V.5.1 Simulations.....	73
	V.5.2 Phantom Experiments.....	74
	V.5.3 Clinical Case.....	79
	V.6 Discussion	82
	V.6.1 Simulations.....	82
	V.6.2 Phantom Experiments.....	83
	V.6.3 Clinical Case.....	84
	V.6.4 Limitations.....	85
	V.7 Conclusions	88
VI.	TOWARD A GENERIC REAL-TIME COMPRESSION CORRECTION FRAMEWORK FOR TRACKED ULTRASOUND.....	89
	VI.1 Introduction and Significance of Study	89
	VI.2 Abstract.....	90
	VI.3 Introduction	90
	VI.4 Methods.....	92
	VI.4.1 Phantom Construction	92

VI.4.2 Preoperative Imaging and Patient Model	93
VI.4.3 Intraoperative Data Collection	94
VI.4.4 Compression Correction	94
VI.4.4.1 Patient-Specific Correction	95
VI.4.4.2 Proposed Generic Correction	97
VI.4.5 Experimental Validation	101
VI.4.5.1 Simulations	101
VI.4.5.2 Phantom Experiments	102
VI.4.5.3 Clinical Case	103
VI.4.5.4 Computational Efficiency	104
VI.5 Results	104
VI.5.1 Simulations	104
VI.5.2 Phantom Experiments	106
VI.5.3 Clinical Case	109
VI.5.4 Computational Efficiency	111
VI.6 Discussion	112
VI.6.1 Simulations	112
VI.6.2 Phantom Experiments	113
VI.6.3 Clinical Case	114
VI.6.4 Computational Efficiency	114
VI.6.5 Additional Comments	115
VI.7 Conclusions	117
VII. OTHER STUDIES TOWARDS USING ULTRASOUND IMAGE CONTENT	119
VII.A Geometric Reconstruction Using Tracked Ultrasound Strain Imaging	119
VII.A.1 Introduction and Significance of Study	119
VII.A.2 Abstract	120
VII.A.3 Purpose	120
VII.A.4 Novel Contributions to be Presented	121
VII.A.5 Methods	122
VII.A.6 Results	125
VII.A.7 Conclusions	128
VII.B Intraoperative Segmentation of Tracked Ultrasound Images	129
VII.B.1 Introduction and Significance of Study	129
VII.B.2 Abstract	129
VII.B.3 Introduction	129
VII.B.4 Methods	131
VII.B.4.1 Preoperative Patient Model	131
VII.B.4.2 Intraoperative Tracked Ultrasound	131
VII.B.4.3 Soft-Tissue Deformation Correction	132
VII.B.4.4 Segmentation Procedure	132
VII.B.4.5 Experiments	136
VII.B.5 Results	137
VII.B.6 Discussion	142
VII.B.7 Conclusions	144
VIII. SUMMARY AND FUTURE DIRECTIONS	145

REFERENCES	154
APPENDIX	171
A. Monitoring surgical resection of tumors with ultrasound strain imaging.....	171
A.1 Introduction and Significance of Study.....	171
A.2 Background	171
A.3 Aims.....	172
A.4 Methods.....	172
A.5 Results.....	172
A.6 Conclusions	173
B. Automatic Generation of Boundary Conditions Using Demons Non-Rigid Image Registration for Use in 3D Modality-Independent Elastography.....	175
B.1 Introduction and Significance of Study	175
B.2 Abstract	176
B.3 Introduction	176
B.4 Methodology.....	180
B.4.1 Automatic Generation of Boundary Conditions.....	180
B.4.2 Simulations.....	182
B.4.3 Phantom Experiment 1	183
B.4.4 Phantom Experiment 2	184
B.5 Results	186
B.5.1 Simulations.....	186
B.5.2 Phantom Experiment 1	189
B.5.3 Phantom Experiment 2	190
B.6 Discussion.....	193
B.6.1 Simulations.....	193
B.6.2 Phantom Experiment 1	196
B.6.3 Phantom Experiment 2	197
B.7 Conclusions	200

LIST OF TABLES

Table	Page
1. Recent examples of LRS integration into image-guided procedures.....	43
2. Comparison of the arithmetic and RMS tracking error over all 30 scans (total) and also over the separate groups of 10 scans taken while tracking each face (Face 1 is the rear right face, Face 2 is the top face, and Face 3 is the rear left face). The errors were calculated by comparing the centroid points of the phantom discs against points collected with a tracked pen probe.	54
3. Simulation results to assess the assumption of compression purely in the depth direction of the ultrasound image plane. The two angles are the pitch and yaw of the ultrasound probe about which the displacement vectors were rotated when assigning Dirichlet boundary conditions. The mean and standard deviation of the boundary node error and the mean centroid distance were calculated in comparing each corrected simulation tumor with the original uncompressed tumor.	74
4. Results of phantom elasticity test to compare the effect on model-correction from the assumption of 1:1 tumor-to-bulk stiffness ratio against the known 9:1 stiffness ratio from material testing.	78
5. Approximate execution time for each step in the patient-specific and the generic correction pipelines. Times were determined using a single thread of an Intel Core2 Quad CPU at 2.4 GHz.	111
6. Results of geometric reconstruction of the phantom tumor volumes using CT, tracked B-mode images, and tracked strain images.....	126
7. Height-to-width ratios for the phantom tumor volumes from CT, B-mode reconstructions, and strain imaging reconstructions. Height and width were measured approximately through the tumor centroids.....	127
8. The level set weighting parameters used for each of the segmentation trials (six trials for both the B-mode and strain image). The parameters in Trial 1 were chosen empirically to give a satisfactory segmentation, and then following trials modified one weighting factor at a time (bolded).	137
9. Comparison of the semi-automatic contours generated for the B-mode image with various weighting parameters, compared to the manual segmentation.	140
10. Comparison of the semi-automatic contours generated for the strain image with various weighting parameters, compared to the manual segmentation.	142
11. Comparison of boundary condition mapping error and MIE reconstruction results between the four methods for the simulations (*[249], **[211]).	188

12. Comparison of boundary condition mapping error between the four methods for the two phantom experiments (*[249], **Based on work in [250])	190
13. MIE-reconstructed elasticity contrast ratios for phantoms 2&3 and gel material testing data (*[211])	192

LIST OF FIGURES

Figure	Page
1. Time evolution of a HIFU focus region showing close-ups of the bright hyperechoic spot produced during ablation (b-h). Reprinted from [23], with permission from Elsevier.	6
2. A B-mode image of a rabbit kidney before (left) and during (right) histotripsy treatment, showing the hyperechoic cavitation zone. Reprinted from [27], with permission from Elsevier.	7
3. A thermal lesion resulting from radiofrequency ablation of a metastatic liver mass, imaged with B-mode (a) and ARFI (b), with verification of the tumor borders in postoperative CT (c&d). Reprinted from [44], with permission from Elsevier.	8
4. Tracking hardware setup. The passive target attached to the probe (left) is tracked with respect to a reference target by the Polaris tracking system (right).	21
5. Illustration of the two discrete transformations needed in a freehand 3D ultrasound system. The calibration transformation must be computed by the user before using the system and maps image space (left) to probe sensor space (center), whereas the tracking transformation is computed automatically by the tracking system and maps probe space to patient space (right).	23
6. Spatial calibration setup. A tracked probe is inserted into the ultrasound beam (left) until the tip produces a bright dot in the image (right). This is done repeatedly to establish the transformation between image space and sensor space.	25
7. Illustration of the synchronization problem in tracked ultrasound. Interpolation of tracking data needs to be performed in order to match each video frame with its corresponding tracked pose (dotted line).	27
8. Temporal calibration setup. The tracked transducer (left) is oscillated up and down while imaging a flat surface (right), which moves correspondingly in the video.	29
9. Siemens commercial elastography software. A gel phantom containing a hard inclusion was imaged using B-mode (left) and strain imaging (right).	31
10. Example of unprocessed RF data (left) and processed B-mode image (right) from a gel phantom containing hard inclusion.	33
11. Pre-compression (solid line) and post-compression (dashed line) axial RF signals from two frames in a single window. These curves represent approximately the same signal separated by a phase difference.	34
12. Cross correlation function for one window of RF data matched in the axial direction. Here, the maximum discrete value occurs at a lag of 2 samples.	35

13. Phase angle of the analytic CC function at each lag value. Each zero crossing corresponds to a maximum CC value. The exact zero crossing may be estimated by approximating the phase angle function as a line using three points (shown above).....	36
14. B-mode image (left) and strain image overlay (right). Both images were computed from raw RF data acquired from the research interface on the Acuson Antares machine.....	37
15. Mechanical tester (right) and workstation (left) which are used to obtain material properties of phantom samples to validate strain imaging.....	39
16. The novel LRS, showing the single CCD design from the front (top left picture) and the tracking marker configuration from the top of the scanner (top right picture). The bottom diagram shows the improved functionality of the single CCD system (right) in comparison with a dual-CCD system (left).....	46
17. Geometric accuracy test setup (left) showing the nine positions of the precision phantom (right).....	49
18. Orientations used in tracking visibility test. For reference, a pitch of 0° and a yaw of 0° denotes the orientation in which the top of the LRS is facing toward the camera, whereas a pitch of 90 and yaw of 180 denotes a horizontal orientation facing away from the camera.	50
19. Experimental setup used for the tracking accuracy test, showing the fixed phantom and LRS (left) and the Polaris Spectra optical tracking system (right).....	51
20. Bitmap view of the phantom captured by the LRS (top left) and textured point cloud rotated slightly to show the 3D geometry of the data (top right). The bottom left shows an example of intraoperative data collection with the laser line sweeping across a brain surface. The corresponding LRS point cloud reconstructed from this scan is shown on the bottom right.	53
21. Experimental setup with a gel phantom attached to a base with fiducial markers (a), and the digital phantom and tumor surfaces segmented from the CT image volume (b).....	60
22. Polaris tracking system (a), and the laser range scanner with passively-tracked targets (b). A passively-tracked target was also used to track the ultrasound probe (c), and the laser range scanner was used to construct a digital representation of the probe surface (d).....	62
23. Procedure for the rigid correction and model-based correction. The rigid correction was performed by calculating the distances between the top of each A-line and the co-registered tomographic surface (a) and then translating the A-lines upward in the depth direction by those distances (b). The model-based correction was performed by calculating the distances from the tomogram mesh nodes to a digital representation of the ultrasound probe surface, and designating the distances as boundary conditions to a FEM model (c). The model was used to solve for tissue	

displacements throughout the mesh (d) and then reversed to deform the ultrasound image to its pre-compressed state (e).	63
24. Initial boundary conditions were assigned manually (a) to designate the bottom fixed, and the top as stress free (b). These boundary conditions were then modified for each model-based correction by changing the appropriate nodes to Dirichlet displacement conditions (c) based on the position of the ultrasound probe surface, and then the model was solved to produce an approximation of the deformed tissue (d).	66
25. Simulation methodology. The probe surface was inserted into the mesh to simulate its position during ultrasound imaging of the tumor (a). Nine sets of boundary conditions were then generated to simulate various possible probe-tissue contact scenarios by rotating the displacement vectors about angles θ_1 and θ_2 defined with respect to the probe orientation (b). The solid arrows descending from the brain to the surface represent the assumption of compression in the depth direction, whereas the dashed arrows are an example of how the displacement vectors were systematically rotated to test that assumption (here $\theta_2 = 30^\circ$). Each set of boundary conditions was used to drive the forward model and deform the mesh and tumor (c). The proposed rigid and model-based corrections (using the assumption of compression perfectly parallel with the depth direction) were then applied to each deformed tumor and compared to the original uncompressed tumor (d).	68
26. Examples of B-mode (a-c) and strain image contours (d-f). (a) and (d) show the uncorrected images in 3D space with the segmented tumor contour shown in red and the tomogram tumor volume shown in green. (b) and (e) show the alignment after the rigid correction method. (c) and (f) show the alignment after the model-based correction method.	76
27. Alignment error results for the B-mode (a & b) and strain imaging (c & d) modalities for the organ-like phantom ($n = 178$ for B-mode, and $n = 83$ for strain). The position of tumor borders in each modality was evaluated in terms of Modified Hausdorff Distance to the co-aligned CT borders (a & c), as well as the distance between the centroid of the ultrasound tumor with the co-planar CT tumor border (b & d). The edges of the boxes are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered as outliers.	77
28. Results of the clinical case using tracked B-mode to localize a tumor border. The coronal and axial views are each shown for the uncorrected (a & b), rigid correction (c & d) and model correction (e & f) with the segmented tumor contour shown in red and the tomogram tumor volume shown in green. The segmented tumor surface in each B-mode slice is outlined for comparison with the MR tumor surface.	80
29. Alignment error results for the clinical case ($n = 118$ B-mode images). The position of tumor borders was evaluated in terms of Modified Hausdorff Distance to the co-aligned MR borders (a), as well as the distance between the centroid of the ultrasound tumor with the co-planar MR tumor border (b). The edges of the boxes	

are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered as outliers.	81
30. Example of a mesh used for the patient-specific model correction method (a), constructed from preoperative imaging and aligned to the ultrasound data using intraoperative registration methods, and an example of a block mesh for the generic model correction method (b), which is pre-aligned to the ultrasound data by performing a calibration to the ultrasound probe.	95
31. Example of the meshes used in the simulation studies with sizes of 10 x 10 x 5 cm (a), 10 x 10 x 10 cm (b) and 10 x 10 x 15 cm (c). In this case a simulated 20 mm diameter tumor is shown at the equivalent 3 cm depth in each mesh.	101
32. Effects of mesh resolution on three sizes of a tumor after a simulated 10 mm surface compression. The tumor node error is defined relative to the result of the solution of a mesh with 2 mm edge length resolution.	105
33. Effects of block mesh depth, tumor stiffness, and tumor size upon the model-predicted tumor border deformation under varying amounts of surface compression.	106
34. Example of the patient specific-model correction (a-d) and generic model correction (e-h) applied to a B-mode image slice. (a) and (e) show the co-registered LRS point cloud on the respective mesh, (b) and (f) show the tracked probe surface and the misalignment between the ultrasound tumor border with the CT tumor, (c) and (g) show the deformed mesh after creation of boundary conditions, and (d) and (h) show the undeformed mesh and corrected ultrasound image.	107
35. Alignment error results for the B-mode (a & b) and strain imaging (c & d) modalities for the organ-like phantom (n = 178 for B-mode, and n = 83 for strain). The position of tumor borders in each modality was evaluated in terms of Modified Hausdorff Distance to the co-aligned CT borders (a & c), as well as the distance between the centroid of the ultrasound tumor with the co-planar CT tumor border (b & d). The edges of the boxes are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered as outliers.	109
36. Alignment error results for the clinical case (n = 118 B-mode images). The position of tumor borders was evaluated in terms of Modified Hausdorff Distance to the co-aligned MR borders (a), as well as the distance between the centroid of the ultrasound tumor with the co-planar MR tumor border (b). The edges of the boxes are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered as outliers.	110
37. The tracked ultrasound system consisted of a Polaris optical tracking system (a) and a passive target attached to the ultrasound probe (b).	123
38. Calibration of the tracked ultrasound system. The ultrasound beam is approximated to be planar, and is sampled with a tracked pen probe (a) to produce bright corresponding points in the ultrasound images (b). A calibration transformation is computed which best matches the image points to the physical probe points.	123

39. Embedded tumor surface generated from CT (a) and corresponding lesion shown in two roughly cross-sectional tracked strain images (b) and as a volume reconstructed from strain contours (c).	124
40. An example of the CT tumor dimensions (a), the corresponding strain tumor dimensions (b), and the B-mode tumor dimensions (c).	126
41. Ultrasound image processing and segmentation pipeline using the shape model derived from a co-registered tomogram segmentation.	135
42. Creation of statistical shape model from co-registered tomogram target for the ultrasound segmentation pipeline.	136
43. Example of the statistical shape model created from intersection of a tracked ultrasound slice with a co-registered MR tumor, showing the mean image (left) and three standard deviations of the first and second modes (middle and right).	138
44. A B-mode image of the brain tumor (left), and the tracked ultrasound image in 3D space with the co-registered MR tumor (right).	138
45. Illustration of the segmentation variability in the B-mode image according to level set weighting parameters. A manual user-drawn segmentation is shown in (a). (b) shows the result of the semi-automatic segmentation using empirically selected weighting parameters to give a contour similar to the manual contour. The contours given in (c-g) show some examples of altering the parameters chosen for (b): (c) is the result of doubling the advection weight; (d) is the result of doubling the curvature weight; (e) is the result of doubling the propagation weight; (f) is the result of doubling the shape model weight; and (g) is the result of eliminating the shape model term from the level set equation.	139
46. A strain image of the brain tumor (left), and the tracked ultrasound image in 3D space with the co-registered MR tumor (right).	140
47. Illustration of the segmentation variability in the strain image according to level set weighting parameters. A manual user-drawn segmentation is shown in (a). (b) shows the result of the semi-automatic segmentation using empirically selected weighting parameters to give a contour similar to the manual contour. The contours given in (c-g) show some examples of altering the parameters chosen for (b): (c) is the result of doubling the advection weight; (d) is the result of doubling the curvature weight; (e) is the result of doubling the propagation weight; (f) is the result of doubling the shape model weight; and (g) is the result of eliminating the shape model term from the level set equation.	141
48. Examples of the compression correction applied in a clinical brain case (a&b), a liver phantom (c&d) and a clinical breast case (e&f from [209]). In each pair of images, the first image shows a slice from the preoperative tomogram, and the second shows an overlay of the co-registered ultrasound image after compression correction.	146

49. Example of MIE applied to a CT breast volume (a) with cross-section showing tumor-to-normal elasticity contrast (b), and an MR breast volume (c) with cross-section showing similar contrast (d). [215]	148
50. An example of tracked ultrasound data acquired during a breast examination of a patient with benign cysts (a), showing the registration of LRS cloud to MR surface (b). The swabbed ultrasound cyst contours are spread over a wide volume due to dragging of the tissue (c), which was confirmed in the probe tracking vector history (d).....	149
51. Example of an ultrasound cyst surface (red) misaligned with the corresponding MR surface, with the prior swab trajectory visualized as a blue arrow (a). The known vector of misalignment between the ultrasound and MR centroids was computed and the closest vector in the previous three seconds of tracking data was used to translate the ultrasound cyst to an earlier position in the swab.....	151
52. Example of a stereovision reconstruction on a cortical surface phantom, showing the left and right camera views and the resulting 3D point cloud [219].	152
53. Effects of tumor resection on strain imaging in a cryogel phantom. Boxes (a), (c), (e), and (g) show increasing amounts of tumor resection, from no resection to total resection. Boxes (b), (d), (f), and (h) show the corresponding strain images of the tumor remnant.	173
54. Representative slices from the two data sets used for the simulations. Slices (a) and (b) show the pre- and post-deformed CT set, whereas slices (c) and (d) show the pre- and post-deformed MR set, respectively.	187
55. TRE distribution (in mm) across the surfaces of the CT mesh (a) and the MR mesh (b) for the demons-based boundary conditions compared to the known conditions.....	187
56. Objective function maps for the CT simulation (a) and the MR simulation (b). The objective function value calculated by the optimization framework is plotted on the ordinate axis against selected elasticity contrast ratios (tumor-to-normal) as affected by the boundary conditions. Shown are the objective maps of the demons case (solid lines) and the known boundary conditions as the control (dashed lines). The ordinate is scaled in both cases.....	189
57. Representative source (a) and target (b) geometry of the anthropomorphic phantoms, with center slices. Phantom tumors as shown here were present in Phantoms 2&3 only.....	191
58. Objective function maps for Phantom 2 (a) and the Phantom 3 (b). The objective function value calculated by the optimization framework is plotted on the ordinate axis against selected elasticity contrast ratios (tumor-to-normal) as affected by the boundary conditions. Shown are the objective maps of the demons case (solid lines) and the known boundary conditions as the control (dashed lines). The ordinate is scaled in both cases.	193

LIST OF ABBREVIATIONS AND NOMENCLATURE

Abbreviation	Full Name
1. ARFI	acoustic radiation force impulse
2. CC	cross correlation
3. CCD	charge-coupled device
4. CT	computed tomography
5. FEM	finite element method
6. FLE	fiducial localization error
7. FRE	fiducial registration error
8. FTC	freeze-thaw cycle
9. ICP	iterative closest point
10. IGS	image-guided surgery
11. IRED	infrared light emitting diode
12. LRS	laser range scanner
13. MHD	modified Hausdorff distance
14. MIE	modality-independent elastography
15. MRE	magnetic resonance elastography
16. MRI	magnetic resonance imaging
17. OR	operating room
18. PCA	principal component analysis
19. PDE	partial differential equation
20. PVA	polyvinyl alcohol
21. RBF	radial basis function
22. RF	radiofrequency

23. RMS	root mean square
24. SRE	surface registration error
25. SVD	singular value decomposition
26. SWEI	shear wave elasticity imaging
27. TPS	thin plate spline
28. TRE	target registration error
29. US	ultrasound
30. USE	ultrasound elastography

CHAPTER I

INTRODUCTION

I.1 Specific Aims

Ultrasound provides several contrast mechanisms, such as conventional B-mode and strain imaging, which have the potential to provide complementary information beyond that which is offered by preoperative volumetric imaging modalities such as computed tomography (CT) or magnetic resonance (MR) imaging [1]. Given the critical importance of accurate margin delineation for surgical targets such as tumors or blood vessels, ultrasound can be a useful interventional tool within a comprehensive image-guidance framework by supplementing information regarding the extent of diseased tissue. Unfortunately, ultrasound suffers from several drawbacks compared to other imaging modalities. One of the most common criticisms of ultrasound centers around the challenge of image interpretation, which can be difficult without other supporting information. In addition, the contact between ultrasound probe and tissue can result in nonrigid tissue deformation which distorts positional and size measurements of subsurface targets. There are emerging ultrasound technologies which have yet to be widely realized in an interventional role due in part to the issues mentioned above, despite the fact that ultrasound is already a common fixture within the operating room. These technologies have the potential to impact a wide array of procedures by offering novel information to the clinician, provided that intraoperative ultrasound can be made more quantitative with respect to spatial localization of subsurface objects.

The primary objective of this dissertation was to advance ultrasound as an interventional image-guidance platform by creating methods to enhance the spatial context of ultrasound data intraoperatively. The following tasks were accomplished in order to meet this goal: 1) a tracked ultrasound system was developed with capability for both B-mode and strain imaging, 2) a high-

accuracy laser range scanner was characterized and used to register tracked ultrasound data to other imaging modalities, 3) methods were developed to correct for nonrigid tissue deformation during ultrasound imaging, and 4) these methods were evaluated with simulations, phantoms, and clinical data. The results and applications of this work are important to the development of a comprehensive approach to ultrasound-guided interventions. Thus, while the goals of this project lie within the domain of a much larger field, the scope of this dissertation was constrained by the formulation of the following specific aims:

I.1.1 AIM 1: Characterize the accuracy of a novel laser range scanner

The first aim of this project was to describe the design and performance of a laser range scanner (LRS). The LRS was developed with the capability to create a tracked geometric point cloud with color information automatically aligned to cloud, and it was evaluated with respect to its geometric and tracking accuracy. This device was necessary to support the next two aims as a high accuracy measurement tool capable of creating surface geometry in a common 3D coordinate system as the tracked ultrasound data. The data from the LRS was used for both surface-based image-to-physical registrations and for tissue deformation measurements.

I.1.2 AIM 2: Develop a patient-specific compression correction method for tracked ultrasound

The second aim of this project was to develop a method to correct the positional and geometric information in tracked ultrasound in cases involving significant tissue compression exerted by the ultrasound probe. This was accomplished by registering intraoperative tracked ultrasound data to a tissue model created from preoperative imaging of the patient, and utilizing that registration to estimate the amount of tissue compression induced during ultrasound imaging. This estimation of the compression was then used to drive a biomechanical tissue model and create a full displacement field prediction throughout the tissue volume, which was then used to correct the ultrasound data.

I.1.3 AIM 3: Develop a generalized real-time compression correction method for tracked ultrasound

The third aim of this project was to extend the patient-specific correction to a generic framework which does not rely on the patient-specific model as in Aim 2. The patient-specific method was limited to clinical interventions which routinely utilize volumetric preoperative imaging which closely matches the intraoperative state of the patient. This aim modified the previous approach so that a generic block of tissue was modeled on the tip of the tracked ultrasound probe and used to estimate tissue displacement. The generic block correction method only required relatively sparse intraoperative measurements of tissue compression depth in order to generate a model correction. In addition, the workflow modifications by this method enabled the model solution to be calculated at a much faster speed compared to the patient-specific correction, and thus preserved the rapid acquisition characteristic of ultrasound.

I.2 Dissertation Overview

This dissertation begins in Chapter II with an introduction to image-guided interventions, including ultrasound imaging, tracking systems, registration methods, and model-updated image guidance. A comprehensive description of the methodology used throughout this dissertation is given in Chapter III, including image segmentation methods, model creation, the creation and calibration of the tracked ultrasound system, and ultrasound strain imaging. The following chapters expand upon the aforementioned specific aims of the project. Chapter IV describes the methods and results of the characterization of the novel LRS. Chapter V describes the development of the patient-specific model correction for tissue compression during ultrasound imaging, with results in phantoms and clinical data. In Chapter VI, the methods and results of the generic compression correction method are described. Chapter VII presents additional studies in which the ultrasound image information is exploited to detect surgical targets. Finally, Chapter VIII presents the overall conclusions of the work and suggests future areas of investigation.

CHAPTER II

BACKGROUND

The following section provides a brief overview of the state of image-guided interventions and the most common technologies which enable these procedures. Special note will be made of groups which have used intraoperative tracked ultrasound as a navigational tool, as the Specific Aims of this project rely on the development and use of a tracked ultrasound system. This chapter is intended to provide a broader context to the contributions of this dissertation in the field of image-guided interventions as a whole.

II.1 Introduction to Image-Guided Interventions

Image-guided interventions are defined as medical procedures which leverage computer displays of imaging data to assist the physician in localizing a surgical target in a spatially and temporally quantitative manner [2, 3]. This is achieved by merging information from imaging data with the physical working space in the operating room and thus aligning the coordinate systems of image and physical space. A modern image-guidance technology platform most commonly consists of the following components: 1) imaging data, which is often preoperative tomographic volumes but also includes a variety of intraoperative modalities, 2) intraoperative tools which are tracked with a localizing system, 3) an image-to-physical registration method, and 4) a method of visualizing the co-registered information in an interactive display.

The following sections will describe some of these technological components and current clinical applications. Although many anatomical targets in the human body are the subject of image-guided interventions, this dissertation will focus on a subset of these. Many of the examples will be centered on neurosurgery, as image-guidance has been standard of care for several decades in

those procedures. Special attention will also be given to emerging ultrasound technologies which may provide substantial benefit to current clinical practices when integrated with the contributions made in this dissertation. It should therefore be noted that many of the following concepts and technologies are also being employed to aid interventional procedures in other fields, such as in radiation therapy and radiosurgery, as well as in other anatomy such as the kidney, liver, and prostate to name a few [4-6].

II.2 Ultrasound Imaging

Intraoperative ultrasound is an interventional imaging modality capable of subsurface measurements. It has the benefit of being a low-cost and safe alternative to both computed tomography (CT) and magnetic resonance (MR) imaging, and is a tool readily available in almost all operating rooms (ORs). Conventional ultrasound is an attractive imaging modality because it is relatively inexpensive, noninvasive, nonionizing, and can provide quick real time imaging of nearly any human tissue. It is widely used for a variety of medical applications, including intraoperative imaging for treatment of breast [7-9], liver [10-12], prostate [13], thyroid [14], and gynecological targets [15]. Despite its appeal, however, it does generally suffer the limitation of providing only 2D images as opposed to CT and MR 3D volumes. The clinical availability of 3D ultrasound probes is increasing, but conventional 2D probes remain the standard intraoperative technology. It is widely recognized that the 2D format and relatively low signal-to-noise ratio of ultrasound images makes them more difficult to interpret in isolation. Because of this limitation, the utility of intraoperative ultrasound has also been investigated almost in parallel with stereotactic procedures in attempts to place US images in a spatial context [16-19]. For example, intraoperative ultrasound has been used to provide some measure of brain shift in neurosurgery, as well as to measure the difference in tumor volumes predicted by preoperative CT and ultrasound [20, 21].

In addition to conventional open surgeries, ultrasound guidance has also been deployed in a variety of noninvasive procedures. A procedure of particular note is the use of high intensity focused ultrasound (HIFU) to ablate surgical targets such as solid tumors by converting mechanical energy into heat and also through inertial cavitation effects [22, 23]. A recurrent challenge in the HIFU community is accurate monitoring of the ablation zone with noninvasive imaging, which is typically done with either MR or ultrasound itself. Ultrasound has been used to guide HIFU ablation of tumors in the breast [24], liver [25], uterus [26], and other organs. One of the difficulties usually encountered is accurate spatial identification of the target and ablation zone borders. Figure 1 shows an example of a HIFU ablation zone evolving with time during a procedure.

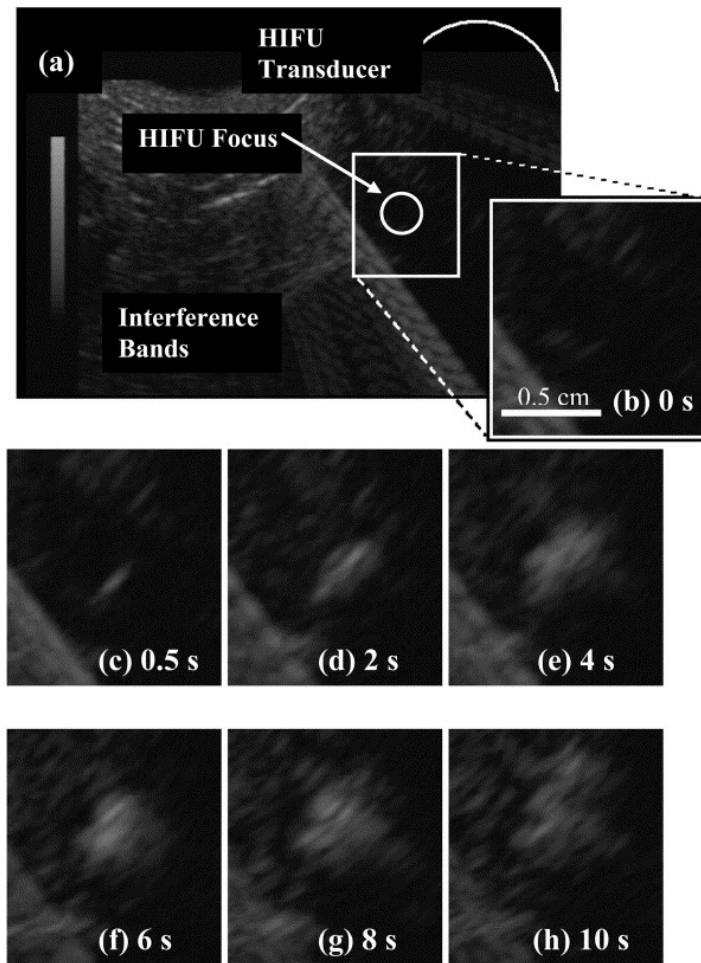


Figure 1. Time evolution of a HIFU focus region showing close-ups of the bright hyperechoic spot produced during ablation (b-h). Reprinted from [23], with permission from Elsevier.

A related technique involves using short pulses of focused ultrasound energy to induce cavitation destruction of the tissue in a noninvasive and nonthermal manner. This effect is known as histotripsy, and it has been presented as a more precisely targeted method of destroying tissue by avoiding the problems associated with thermal ablation, namely the inhomogeneous heating of tissue due to tissue heterogeneity, perfusion, and charring [27]. Ultrasound histotripsy has shown potential for use in thrombolysis [28] and has also been proposed for targeted ablation of small incidental lesions in various organ systems such as the kidney [29] and prostate [30]. Figure 2 shows an example of a targeted ablation in an animal model using histotripsy.




Figure 2. A B-mode image of a rabbit kidney before (left) and during (right) histotripsy treatment, showing the hyperechoic cavitation zone. Reprinted from [27], with permission from Elsevier.

Acoustic radiation force impulse (ARFI) imaging and shear wave elasticity imaging (SWEI) are ultrasound-based technologies which have emerged as noninvasive methods of measuring tissue stiffness [31-33]. Traditionally, imaging strategies for measuring tissue stiffness have relied on external tissue excitation methods or physiological motion in order to induce mechanical deformation of the tissue which can be related to material properties [34]. The acoustic radiation force based methods avoid the need for external excitation, relying instead on the interactions of the tissue and acoustic beam to estimate stiffness. ARFI has been commercially available for the last

several years and has been presented in the clinical literature for several applications. Thus far it has primarily been used in a diagnostic role, such as for assessing solid hepatic lesions [35], liver fibrosis [36-38], myocardial stiffness [39], breast lesions [40, 41], and peripheral nerves [42]. There has also been recognition of the potential usefulness of the elasticity information offered by ARFI in a guidance sense, such as for needle biopsies in the prostate [43] and monitoring the target zone in radiofrequency ablations in the heart and the liver [44, 45] as shown in Figure 3.




Figure 3. A thermal lesion resulting from radiofrequency ablation of a metastatic liver mass, imaged with B-mode (a) and ARFI (b), with verification of the tumor borders in postoperative CT (c&d). Reprinted from [44], with permission from Elsevier.

A technology which has also been emerging is freehand tracked ultrasound, which has several advantages over conventional ultrasound. It visualizes tissue in an intuitive 3D format, allows for the acquisition of 2D slice views in arbitrary 3D orientations, and facilitates the placement of ultrasound image slices in a meaningful spatial and temporal context. Freehand tracked ultrasound systems generally consist of a tracked sensor attached to the ultrasound probe which is used to determine its 3D position and orientation by an external tracking system. This information is used to calculate the 3D position of each object in the ultrasound images, and can also be used to reconstruct a 3D ultrasound volume by compounding a series of tracked 2D images [46]. By placing the images within an external coordinate system, it becomes possible to directly align ultrasound images to the patient, as well as to other modalities. As an example, tracking information has been used to assist in automatic co-registration of ultrasound with MR for correction of brain deformation [47]. Similar work has been done to register ultrasound with CT images of the spine in navigated orthopedic surgery [48]. The fusion of multiple imaging modalities in a unified image-guided surgery (IGS) platform can produce a greater understanding of the content in both modalities for the surgeon, and it has also been shown to lead to changes to the therapeutic strategy compared to standalone imaging information [49, 50]. Groups have demonstrated freehand 3D ultrasound to monitor brain shift, for example, and also incorporate image warping to register intraoperative data to preoperative images [51-56]. Current commercial ultrasound systems designed for IGS integration include the SonoWand (SonoWand AS, Trondheim, Norway), Brainlab (Brainlab AG, Munich, Germany) [49, 57] and PercuNav systems (Philips, Amsterdam, Netherlands). While these imaging systems do provide a quantitative measurement of intraoperative tissue movement, methods of compensating for tissue deformation in real-time for use in surgical guidance have not yet reached maturation.

II.3 Tracking Systems

There are a variety of spatial localization technologies for tracking surgical tools in physical space. Early on in the evolution of image-guidance systems, geometric localizers such as articulated arms were used to track points in 3D space, but these devices were ultimately deemed too unwieldy for routine use [58]. The two most common tracking modalities currently in clinical use are optical and electromagnetic tracking. Optical tracking entails the use of an optical recording camera which employs triangulation to localize rigid objects affixed with either actively emitting LEDs or passively reflecting spheres arranged in a known geometric configuration. Optical systems are capable of highly accurate spatial localization in relatively large work volumes, but suffer from the limitation of requiring line-of-sight from the tracking camera and the tracked rigid body. This leads to the additional implication that it is not typically possible to track the tip of nonrigid flexible tools using optical methods, as the optical trackers need to remain outside the body. Nevertheless, the submillimetric accuracy of optical systems has led them to become the standard tracking modality in practice [59]. Electromagnetic tracking devices work on the principle of detecting alterations in a static magnetic field induced by the movement of a sensor through that field. This method offers the advantage of avoiding the issues associated with the line-of-sight requirements of optical tracking between the tracking system and the tracked object. Electromagnetic sensors have become the preferred tracking technology for the tips of nonrigid instruments inserted into the body, such as flexible endoscopes in orbital procedures [60] or flexible ultrasound probes. The primary drawback of electromagnetic tracking systems, however, is their susceptibility to localization inaccuracy when the magnetic field is distorted by objects such as metal tools in the working volume of the system. The working volume for most electromagnetic tracking systems is also much smaller than for optical systems.

While any tracking technology could in theory be used to track an ultrasound probe, the most typical tracking mechanisms have been electromagnetic and optical as discussed above. The

three most common magnetic tracking devices used for freehand 3D ultrasound have been: 1) the Polhemus AC electromagnetic tracking systems (Polhemus Inc., Colchester, USA) [61, 62]; 2) the Ascension DC electromagnetic systems (Ascension Technology, Burlington, USA) [63-66]; and 3) the NDI Aurora (Northern Digital Inc., Waterloo, Canada) system [67]. The most common optical tracking devices used for freehand 3D ultrasound have been: 1) the NDI Polaris optical tracker [52, 68-70]; and 2) the NDI Optotrak optical tracker [71-76].

II.4 Registration Methods

The primary function of any IGS platform is to provide correspondence between intraoperative data and a wealth of preoperative images [3]. This correspondence is accomplished by aligning the patient in the OR and the patient's preoperative images within a common coordinate system. This process of alignment is known as registration.

Registration was initially accomplished by the surgeon using intuition and training to align anatomical landmarks to get a sense of where to operate. Quantitative registration in neurosurgery was popularized with the use stereotactic frames [77, 78], which helped to formalize the problem of registering image space (the coordinate system in the preoperative images) and physical space (a similar coordinate system within a real world localization system, such as the stereotactic frame). The advent of tomographic imaging modalities led to widespread exploration of IGS methods using stereotaxy [79-82].

Frameless stereotaxy has become a popular method of providing correspondence in recent years as a more convenient alternative to full stereotactic frames. These are devices which localize the physical space of the patient in alternative ways, such as with acoustical [83], robotic [84], magnetic [85], or optical [86] mechanisms as discussed in the previous section. These systems, particularly optical tracking, have been shown to provide the same level of localization accuracy as stereotactic frames, and have become the current standard in IGS procedures.

The process of registering without frames essentially consists of matching landmarks in the images to corresponding physical landmarks on the patient. Provided at least three points, a transformation can be computed which maps all points in one space to the other [87]. These points are sometimes chosen to be immobile anatomical landmarks such as the bridge of the nose in craniofacial surgery [88]. Difficulty in consistently localizing these types of landmarks, however, has led to the common use of synthetic markers which are attached to the skin [89, 90] or rigidly implanted in bone [91] and provide a much more accurate and precise marker localization across both spaces.

In addition to point-based registration methods, surface-based methods are also commonly used. Alignment in these methods is usually achieved by first digitizing some portion of the patient's skin or organ surface. The digitized surface is then registered to the corresponding surface in image space using an iterative closest point (ICP) algorithm described in [92]. For example, methods of digitizing the patient surface in neurosurgery commonly consist of swabbing the craniofacial region with a tracked stylus. Recent methods have included laser range scanning (LRS) technology, which creates a geometric reconstruction of the target surface [93-95] by sweeping a line of laser light onto the object of interest. The surface is digitized by capturing the shape of the laser line with a digital camera and using triangulation to form a point cloud. Calibration is done to determine how points detected by the digital camera are mapped to the physical location of the laser line. LRS systems are attractive for assisting image-guidance because they can provide relatively fast and accurate sampling of the entire exposed surface of the organ. Alignment can be facilitated by tracking a conventional LRS in 3D space via optical targets attached to the exterior enclosure. Some of these devices also collect color information from the field and apply it to the surface as a textured bitmap. This represents additional information which can help guide the registration, such as by matching vessels or other landmarks on the textured surface with the contrast-enhanced MR image [94, 96-98].

Much progress has been made in the last 30 years to provide registration between these data, but there remain obstacles to achieving accurate correspondence between the physical state of the patient and high-detail tomograms throughout the entire procedure. In particular, the integrity of IGS procedures continues to be challenged by the lack of standard mechanisms to account for intraoperative soft tissue deformations, such as the well-described phenomenon of brain shift in neurosurgery, for example. Deformation of the organ compromises the correspondence between the patient and images if there are no means of compensating for the change throughout the course of the procedure. Soft tissue deformation is particularly prohibitive in translating IGS practices to anatomy outside of the cranium such as the abdomen, due to the lack of rigid structures to restrict tissue movement. In addition, uncertainty regarding anatomical landmarks in the preoperative images themselves can undermine the fidelity of IGS. This dissertation highlights the current challenges associated with IGS procedures and presents the development of novel tools to help address these challenges.

II.5 Model-Updated IGS

While intraoperative imaging modalities can provide some level of information regarding the surgical status of the patient, any data which is not acquired intraoperatively needs to be updated to reflect the realities of the operating room. A complementary approach to intraoperative imaging is to computationally model the changes to the preoperative state of the patient using intraoperative data acquisition methods. Generally, the behavior of the organ of interest is assumed to conform to a patient-specific soft tissue biomechanical model. The model can be used to predict the soft tissue motion of the organ under loading conditions measured in the OR. One domain in which this approach has seen promising results is in neurosurgery to address brain shift, and there has been a great deal of research done to investigate ways to accurately model brain deformation [99-105]. One of the most accurate brain shift models, used by Miga *et al.*, is based on Biot's theory

of consolidation mechanics and approximates the brain's behavior as a sponge [106]. The levels of cerebrospinal fluid (CSF) in the brain affect brain motion in this model, as lower levels of CSF will result in the brain sagging in the direction of gravity. This model has been shown in the cited studies as accurately mimicking the brain's observed behavior, and has demonstrated the ability to predict 70-80% of controlled shift in a porcine brain model.

In order for the model to accurately predict movement of the soft tissue, it requires boundary conditions to drive and constrain the biomechanical equations which characterize the behavior of the organ. These boundary conditions are measured via a variety of devices and techniques. Workflow and resource demands often result in this data being sparse in information and extent within the surgery. Sources of data may include intraoperative imaging modalities (such as MR, CT, or US), a tracked stylus, or surface acquisition methods such as laser range scanners. Each of these methods may be used to provide patient-specific boundary conditions for the mathematical model of the tissue and thus present customized guidance to the surgeon. The speed of surface acquisition methods make them preferred for model-based guidance, but the ability to quickly localize subsurface landmarks with ultrasound also has considerable merit. A combined approach may provide the necessary data while keeping cost manageable.

CHAPTER III

METHODOLOGY

The following chapters represent a series of studies which were designed to evaluate the innovations arising over the course of this project. Each chapter focuses on components of a generalized image guidance platform which have been enhanced in this dissertation. The focus of this chapter is to provide an overview of the entire image guidance framework and the methods used. The primary steps involved are: image segmentation, biomechanical model creation and designation of boundary conditions, ultrasound calibration and tracking, image-to-physical registration, and the fundamentals of ultrasound strain imaging. This thesis contributes novel contributions to the areas of boundary condition designation, and image-to-physical registration. The remainder of this chapter is devoted to outlining this framework step-by-step, with indications to where the new contributions will be made.

III.1 Image Segmentation

With respect to the general image guidance framework, image segmentation is a necessary component of the workflow for both preoperative tomographic imaging and intraoperative ultrasound imaging. The following subsections describe the techniques used in this dissertation to accomplish this task.

III.1.1 Preoperative Tomogram Segmentation

Preoperatively acquired images for surgical navigation are usually high resolution CT or MR images which offer detailed anatomical information. The structures of interest must be extracted from unnecessary information in order to construct a geometric model, which will be covered in the next section. There is a rich variety of segmentation techniques in the literature for CT and MR

volumes, including methods specific to the imaging modality, organs, or tissue types. These methods can also be classified as manual, semi-automatic, or automatic. Although there are automatic segmentation techniques in current use for procedures such as MR-guided neurosurgery [107], the work done in this dissertation utilizes primarily manual and semi-automatic segmentation. As far as manual segmentation, commercial packages such as Analyze 9.0 (Mayo Clinic, Rochester, MN) and freely available software such as ITK-SNAP [108] and VTK/ITK (Kitware Inc., Clifton Park, NY) can be used to draw contours with line segments and/or splines to outline the organ of interest in each image slice and create a segmented image mask. With respect to semi-automatic segmentation, the packages mentioned above offer several different methods to facilitate supervised segmentation of structures. The primary method in this dissertation consisted of the simple intensity thresholding tools in Analyze 9.0, in which seed points are placed in the image volume and voxels with intensity values within a defined range and connected to the seed points are assigned to the image mask. The second method utilized consisted of the contour evolution tools in ITK-SNAP, in which the segmentation border propagates based upon a partial differential equation (PDE) framework using either edges or intensity values in the image. The contour is seeded in a similar manner as the simple thresholding technique, but can allow for a more refined segmentation by adjusting the various PDE weighting parameters such as propagation, advection toward edges, or curvature of the contour.

III.1.2 Intraoperative Ultrasound Segmentation

The task of ultrasound segmentation presents different challenges than CT or MR image segmentation due to the characteristic noise and artifacts which are often prevalent in the images. The presence of speckle, attenuation, and shadowing in the data often complicate segmentation techniques traditionally used in other imaging modalities. In addition, the contrast between various tissue types of interest can be quite low in B-mode. These issues have motivated interest in segmentation algorithms which leverage the unique characteristics of ultrasound imaging or

information from other imaging modalities to assist in the segmentation problem at hand. Much effort has been made toward tissue segmentation for diagnostic purposes, but image-guided interventions are emerging as an important application in which image segmentation may have a strong clinical impact. A review of the ultrasound segmentation literature from the last decade is given by [109, 110]. Generally speaking, the most prevalent segmentation strategies primarily target B-mode images rather than the raw radiofrequency (RF) or envelope-detected signals, as most clinical ultrasound machines do not provide access to those forms of data. Although several machines now exist which offer research interfaces to the raw data, this dissertation deals with segmentation of B-mode images. The primary method used in this dissertation was a Livewire technique [111] implemented using VTK image processing and visualization filters in custom-made software. In addition, a novel segmentation framework for intraoperative tracked ultrasound will be presented in the second half of Chapter VII.

III.2 Model Construction

Subsequent to segmentation of the organ of interest from a CT or MR image volume, a finite element mesh is constructed to use as the numerical domain for a biomechanical model. First, the segmented image mask is used to create a 3D surface description using the marching cubes algorithm [112]. This surface is then smoothed with a Laplacian filter, and then it is used as the input for SPESH, a mesh generation program, to create a mesh with tetrahedral elements [113]. The edge length of the tetrahedrons can be specified by the user, and typically an edge length of approximately 5 mm was utilized for each mesh in this dissertation, resulting in approximately 10,000 nodes and 50,000 elements for a liver phantom. Each element is assigned material property values according to prior knowledge about the tissue, and according to the particular biomechanical continuum model.

The finite element mesh described above is designed for numerical integration techniques, which are used in conjunction with a carefully selected model to predict how tissue will deform when subjected to various boundary conditions reflecting the intraoperative state. In this dissertation, the linear elastic model was chosen to simulate tissue biomechanics within the general tracked ultrasound framework because of its computational simplicity [114]. The governing PDE consists of the standard 3D Navier-Cauchy equations for a displacement field u :

$$\frac{E}{2(1+\nu)}\nabla^2 u + \frac{E}{2(1+\nu)(1-2\nu)}\nabla(\nabla \cdot u) + F = 0 \quad (1)$$

where E is Young's modulus, ν is Poisson's ratio, and F is an applied body force distribution. Linear basis functions are defined on the tetrahedral elements to perform the Galerkin weighted residual method and obtain a linear system of equations for a mesh with n nodes:

$$[K]\{\mathbf{d}\} = \{\mathbf{f}\} \quad (2)$$

$$\mathbf{d} = [u_{1x} \quad u_{1y} \quad u_{1z} \dots u_{nx} \quad u_{ny} \quad u_{nz}]^T \quad (3)$$

where K is a $3n \times 3n$ global stiffness matrix, \mathbf{d} is the vector of nodal displacements, and \mathbf{f} is the vector containing applied body forces and surface traction at each node. The assignment of boundary conditions is accomplished using a custom-made graphical user interface. Surface nodes are selected using either a plane or box widget, and then are assigned either displacement or stress conditions by the user. In addition, the user designates whether the assigned conditions are with respect to the global Cartesian coordinate system or the local normal-tangential coordinate system of each individual surface node. Boundary conditions are applied to the appropriate surface nodes by modification of the corresponding equation rows in (2), giving a final system of equations which is solved for the 3D nodal displacements throughout the mesh that satisfy static equilibrium for the boundary conditions. A previously developed remote submission system can be used to take the mesh, boundary conditions, and material properties as inputs, and return the displacements as output [115].

III.3 Tracked Ultrasound

The following subsections describe all of the hardware, software, and methods that were used to support the goals of this research. These primarily include the equipment and software used to track ultrasound images, the calibration procedures needed to ensure accurate tracking, the registration methods which were used to align tracked ultrasound images and tomograms with the patient, the methods used for generating strain images, and the methods used to construct phantoms for preliminary studies.

III.3.1 Tracked Ultrasound Hardware

The task of acquiring tracked ultrasound images must be accomplished by capturing two data streams: 1) ultrasound video frames and 2) tracking data. These two sources of data require two hardware setups connected to a single host PC, which then synchronizes the data. This section provides a brief overview of the hardware used in this dissertation.

A tracked ultrasound system requires synchronization of the tracking information with real-time ultrasound video. There are two primary ways of acquiring images from the ultrasound machine. The first and most popular method is to stream the analog video output of the ultrasound machine (S-video, composite video, coaxial, etc.) to a dedicated frame-grabber card on the computer [52, 116, 117]. This solution is simple and usable on virtually any ultrasound machine, provided that it outputs video in some form. However, it is associated with an implied loss of accuracy due to the repeated signal processing necessary to convert the onboard digital image data on the ultrasound machine to output analog video data, and then to digital data again once it is captured by the frame-grabber. The second method is to directly stream the digital images (pre- or post-processed) from the ultrasound machine to the computer, usually through a network cable [63, 69, 118]. However, this method requires the ultrasound machine to be specifically built with this capability in mind, and it is not currently a typical feature of most clinical ultrasound models. There have recently been attempts by manufacturers to fulfill the demand for direct real-time

access to digital data, such as with Ultrasonix units (Ultrasonix Medical Corporation, Burnaby, Canada) [119].

The ultrasound machine used in this dissertation was an Acuson Antares unit produced by Siemens (Siemens Inc., Munich, Germany). It comes with a standard array of video output ports, including VGA, S-video, and RS-170. In order for the host PC to acquire video output from any of these ports, it needs to be equipped with a frame-grabbing device. Several considerations were taken into account when choosing the capture device, such as price, video formatting, applications programming interface (API) availability, and form factor. The device chosen for this project was a Matrox Morphis Dual frame grabber card (Matrox Imaging, Dorval, Canada). This card is capable of capturing from NTSC, PAL, RS-170 and CCIR video sources and so could easily be adapted to capturing video from other devices as well. It has a relatively small form factor and interfaces via a standard PCI card slot in the host PC, which enabled simultaneous use with other devices requiring a footprint on the motherboard (such as the LRS). It was connected directly to the video output of the ultrasound machine to capture RS-170 format video in real-time.

The second hardware component of the tracked ultrasound system was the tracking instrumentation. In this project, a passive optical tracking system was used as it is the current standard-of-care in neurosurgery cases at Vanderbilt University Medical Center (see Figure 4). The tracking device used was an NDI Polaris Spectra (Northern Digital Inc., Waterloo, Canada). The target attached to the ultrasound probe was also produced by NDI, and was pre-calibrated by the manufacturer to be tracked by the Polaris. The tracking system was connected to the PC through a USB port to acquire the tracked target pose in real-time.




Figure 4. Tracking hardware setup. The passive target attached to the probe (left) is tracked with respect to a reference target by the Polaris tracking system (right).

III.3.2 Tracked Ultrasound Software

A tracked ultrasound system requires real time streaming of ultrasound images and tracked transducer poses. While the hardware framework described above did provide the two data streams, the host PC needed to synchronize them. The software used for this synchronization was adapted from the Synchrograb collection of open-source software for real time 3D ultrasound reconstruction [120, 121]. This software was primarily based upon the Visualization Toolkit (VTK) libraries (Kitware Inc., Clifton Park, New York), which is a collection of high-level functions for facilitating visualization and processing of images and 3D graphics. The Synchrograb software was roughly organized into the following categories: 1) video acquisition, 2) tracking acquisition, and 3) synchronization.

The acquisition of video was accomplished by utilizing the API of the Matrox frame grabber card, which was known as the Matrox Imaging Library (MIL) and was produced by Matrox. The MIL suite of functions came in the form of dynamic link libraries (DLLs) and C++ header files. The MIL functions could be called from an external program to control the physical card and request images be sent to the program as they were acquired from the ultrasound machine. Synchrograb contained an array of classes which wrapped the MIL functions into VTK-style classes for ease of use and

consistency. These classes grabbed the video frames and stored them in a rotating buffer along with their timestamps in real time. All frames that were collected could be optionally output as individual bitmap image files at the end of data acquisition.

The acquisition of tracking information was accomplished by utilizing the API of the Polaris tracking system. The API was produced by NDI and provided methods for communicating with the machine. The functions for controlling the system were wrapped into VTK-style classes in the same manner as the MIL functions discussed above. These classes acquired the tracking matrices and timestamps from the tracking system and stored them in a rotating buffer in real-time. The entire tracking buffer could be optionally output as a text file for postprocessing.

The synchronization of video frames and tracking matrices was accomplished by timestamp matching. It was assumed that the timestamps which were recorded for the two streams of data were based on the same relative time defined by the host PC, as enforced by a temporal calibration performed prior to data collection. Even with this assumption, each ultrasound frame was not necessarily perfectly matched with a corresponding tracking matrix, due to the difference in sampling rate between the two data sources. Since in general the tracking data was sampled at a much higher rate than the video data, the synchronization was approached from the view that each video frame should be matched to an interpolated tracking matrix, rather than vice versa. It was not appropriate to simply average the elements of 4×4 transformation matrices, however, as this could have unintended consequences to the orthonormality of the resulting matrix and produce an unsatisfactory result. A commonly used method for finding an intermediate transformation matrix is known as spherical linear interpolation (Slerp). Slerp is a method for interpolating the quaternion representation of a transformation matrix to a path through 3D rotations with uniform angular velocity around a fixed rotation axis [122]. This method has been shown to generate motion between two quaternions that is smooth and natural. In the synchronization code, each tracking matrix was converted to its quaternion form and the two nearest quaternions to each

frame were interpolated using Slerp to find the corresponding quaternion, which was then converted back to the conventional 4 x 4 matrix form.

III.3.3 Spatial Calibration

A spatial calibration procedure is required in order to transform ultrasound image coordinates into patient coordinates. The rigid body file for the target attached to the ultrasound probe defines the geometry of the target such that it can be localized by the tracking system. However, the relationship between the ultrasound image slice and the attached target needs to be established in order to fully connect ultrasound image space with physical patient space. The tracking system is used to track the pose of the sensor rigidly attached to the ultrasound probe, but this does not directly track the image plane itself. Thus, an additional transformation has to be computed in the form of a calibration which maps the coordinates of image pixels to the coordinate system defined by the attached sensor. The transformations required to express pixel coordinates in terms of patient coordinates are illustrated in Figure 5.

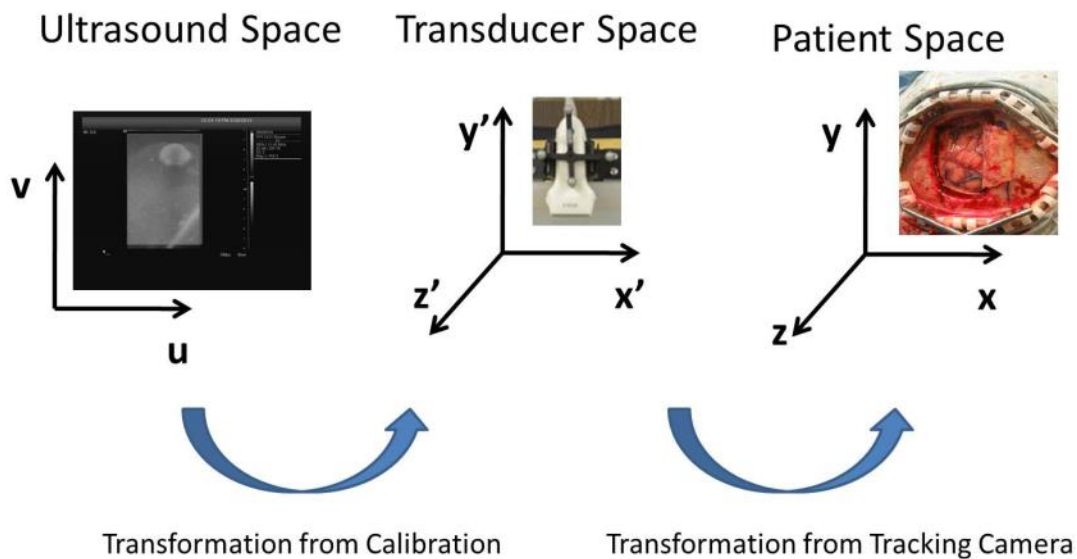


Figure 5. Illustration of the two discrete transformations needed in a freehand 3D ultrasound system. The calibration transformation must be computed by the user before using the system and maps image space (left) to probe sensor space (center), whereas the tracking transformation is computed automatically by the tracking system and maps probe space to patient space (right).

There have been a wide variety of tracked ultrasound calibration procedures proposed in the literature, and it remains an area of active research. A thorough review of freehand 3D ultrasound calibration methods is presented in [123]. The basic idea is the same for each method, which is to establish a consistent relationship between pixels in the ultrasound images and the sensor mounted to the transducer. This relationship can be expressed in equation form as:

$$\begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} = T_w \cdot T_c \cdot \begin{pmatrix} u \\ v \\ 0 \\ 1 \end{pmatrix} \quad (4)$$

where u and v are pixel coordinates in image space, and x , y , and z are coordinates in patient space. T_c and T_w are the 4x4 transformation matrices mapping from image space to sensor space (from the calibration) and from sensor space to patient space (from the tracking system), respectively. Creation of T_c is commonly done by scanning a phantom of known geometry such that a set of features is identifiable in both image space and in sensor space. Phantom designs initially consisted of single-point targets in a coupling medium, in the form of small spherical objects or a single wire [116, 124]. A calibration method developed by [74] is unique in that no phantom is used, but rather the tip of a tracked stylus is imaged by the ultrasound probe. Later groups developed multi-target and cross-wire phantoms to aid in alignment during calibration [64, 73, 117, 125]. The design which has probably become most common is the N-wire phantom [65, 69], which is inspired by stereotactic head frames [126].

The calibration is performed by first localizing phantom features in physical space, usually with a tracked tool. Corresponding features are then localized in image space, either manually or utilizing automatic edge detection [127], depending on the phantom used. The calibration solution (creation of T_c) is essentially the rigid registration of these point sets by minimization of fiducial registration error:

$$FRE^2 = \sum_{j=1}^N \|b_j - sRa_j - t\|^2 \quad (5)$$

where R , t , and s are the rotational, translational, and scaling components of T_c , FRE is the fiducial registration error, a_j is the set of image points, and b_j is the set of physical points. This problem is usually solved in a least-squares minimization sense using singular value decomposition (SVD) or Horn's method [128, 129].

In this dissertation the method of Muratore *et al.* was used, as it did not require the construction of any special phantoms and was shown to have similar accuracy to conventional N-wire calibration phantoms [74]. In this calibration procedure, the tracked transducer was secured to a fixed stand, and the ultrasound beam was projected into a water bath consisting of 9.5% ethanol so that the speed of sound was approximately 1540 m/s [130]. A tracked stylus was then used to sample the ultrasound beam by inserting the probe into the water bath until the tip showed as a bright dot in the image as illustrated in Figure 6.




Figure 6. Spatial calibration setup. A tracked probe is inserted into the ultrasound beam (left) until the tip produces a bright dot in the image (right). This is done repeatedly to establish the transformation between image space and sensor space.

The ultrasound image was recorded, as well as the 3D position of the stylus tip in the sensor coordinate system. This process was repeated a number of times, with the points distributed

equally throughout the extent of the ultrasound image. Typically, about 30 points were gathered in this manner, 25 of which were used to generate the calibration and the remainder were used to test its accuracy. The stylus points were collected under the assumption that the ultrasound beam was approximately planar, and that the points gathered should comprise a plane in 3D space. The stylus points were therefore fit to a plane by calculating a least-squares estimate of the coefficients of the general equation for a plane:

$$Ax + By + Cz + D = 0 \quad (6)$$

Once the coefficients (A , B , C , and D) were known, the stylus points were projected onto the plane along the vector normal to the plane. The bright points in the ultrasound images which corresponded to the 3D probe points were manually selected and given 3D coordinates of the form $(u, v, 0)$, where u and v are the 2D pixel coordinates of the point. The ultrasound image points were then fitted to the plane-projected probe points by a 3D affine registration using a standard SVD algorithm [87]. The result of this registration was a 4x4 matrix which described the necessary transformation of ultrasound image pixel coordinates to transducer coordinates. The points in transducer coordinates were then automatically transformed into the physical coordinate system by the tracking camera using the NDI calibration file for that target.

During the tracked ultrasound calibration, a rough indicator of the noise in the calibration data was the average distance of the probe points to the plane created from (6). This average distance was used as a metric to help judge whether the collected points were in agreement with the assumption of planarity. A large distance suggested that the stylus did not properly sample the ultrasound beam, most likely because the point of the stylus was not adequately close to the center of the ultrasound beam. This indicated that the calibration data should be discarded and recollected.

An even more relevant and direct metric of the quality of the calibration was the calculation of a calibration target registration error (TRE) using points gathered during the calibration which

were not actually used to generate the plane and 4 x 4 matrix, T_c . These points were selected in the ultrasound images and then transformed into physical coordinates using the calibration and tracking system. The transformed image points were then directly compared to the corresponding tracked stylus points as the gold standard to calculate TRE. The primary contributions to this error were the calibration error and the inherent error of the tracking system. Given that passive tracking error can be on the order of 1.5 mm, it was reasonable to obtain a tracked ultrasound TRE in the range of 1.5 to 2.5 mm.

III.3.4 Temporal Calibration

In order to track each frame of a live video feed from the ultrasound machine, the video was synchronized with the tracking information from the tracking system in the form of a temporal calibration. The tracking system and the ultrasound machine both sent their respective data streams to the host computer, but each was sampled at a different frequency and had a different transfer time from the device to the PC. This problem is illustrated in Figure 7.




Figure 7. Illustration of the synchronization problem in tracked ultrasound. Interpolation of tracking data needs to be performed in order to match each video frame with its corresponding tracked pose (dotted line).

The practical implication of this was that the images arrived at the PC in a slightly delayed fashion compared to their corresponding tracking data. The standard approach to addressing this asynchrony is to determine the temporal offset between the data streams experimentally and also

to interpolate the measured pose values to find poses corresponding to the image timestamps [70, 131]. Once the offset t_o is calculated (generally on the order of milliseconds), it is integrated into the guidance platform by simply associating each image frame with the pose data timestamped t_o milliseconds prior. An alternative solution is to use an external trigger to interrogate the tracking system whenever an image frame is acquired, but this requires more specialized hardware [61]. An important phenomenon that affects this process, however, is the difference in transfer time to the PC for the video frames and for the tracking information. In general, it can be assumed that there is a relatively constant time lag between the two data sets due to the size differences between the two (an image is much larger than a 4x4 transformation matrix), and the differences in transfer method (coaxial cable into a frame grabber card for the video versus USB for tracking).

The temporal calibration method used for this project was a variant which relies on estimating the time lag difference between two analogous signals (video and tracking) using cross-correlation. This method was essentially the same as the method which will be described in a later section for finding a displacement field using two sets of RF signals. The primary challenge in this application was the creation of the two analogous signals to be compared. The oscillatory signals here were created by moving the tracked ultrasound probe up and down repeatedly while imaging a flat surface as shown in Figure 8. The height of the line in the image thus oscillated in the same manner as the tracking data, with a simple phase difference between the two signals corresponding to the time delay which needed to be corrected.




Figure 8. Temporal calibration setup. The tracked transducer (left) is oscillated up and down while imaging a flat surface (right), which moves correspondingly in the video.

The position of the probe was recorded in units of millimeters, while the motion of the line in the ultrasound image was recorded in units of pixels. Each of the signals was normalized from a range of 0 to 1 based on the minimum and maximum values in the axis of motion in order to better visualize their correspondence. Each signal was also resampled using a cubic spline interpolation in order to enforce equal spacing between sample points. The original data was unequally spaced because the sampling rates of the video and the tracking stream were different, and also because some frames were not tracked properly when the user inadvertently blocked line of sight from the tracking system to the target. The two signals were then compared in a cross-correlation framework using a Hilbert transform approach to estimate the time lag between the two. This approach yielded a time lag estimate of approximately 11 ms between the tracking data and the video frames. To compensate for this discrepancy as the tracking data and video frames were acquired, they were each stored in buffers on the PC along with their timestamps. As each new

frame was entered into the buffer, it was matched to the tracking matrix from 11 ms prior in the tracking buffer. This simple approach was sufficient to address the issue of temporal synchronization of the two data streams.

III.3.5 Registration

Given that the ultrasound probe was calibrated, alignment with preoperative tomograms was accomplished with standard image-to-physical registration techniques. A Polaris tracking system was the primary intraoperative tracking equipment used in this dissertation. A passively tracked stylus was also utilized in order to digitize points in physical space. It is currently standard of care at Vanderbilt University Medical Center (VUMC) for the neurosurgeon to use the StealthStation stylus to perform the initial image-to-physical registration by swabbing facial landmarks and allowing the Medtronic software to iteratively align the swab points to the preoperative images. Registration for this project was accomplished in essentially the same manner with a stylus via points using Horn's method [129], swabs using iterative closest points [92], or a combination of both [132].

III.3.6 Strain Imaging

One of the original goals of this project was to investigate the utility of tracked strain imaging as an intraoperative guidance tool. Early on during the initial utilization of tracked strain imaging in clinical cases, other obstacles were identified which needed to be resolved and which ultimately became the primary focus of this dissertation, namely correcting for tissue compression effects which are often induced while creating conventional axial strain images. For the sake of background, this subsection will provide an overview of strain imaging.

The simplest method for implementing a strain imaging capability was to use a commercially available elastography module, such as the software produced by Siemens for the Antares ultrasound unit. This module was eSie Touch Elasticity Imaging and was installed separately on the ultrasound machine as an additional imaging option. It was designed to be used

for breast imaging with the VFX13-5 transducer, which has a frequency bandwidth of 5 to 13 MHz. Although it was intended for breast imaging, it was easily translatable to other anatomy like the brain, so long as the objects to be imaged were no deeper than the transducer's display depth of 6 cm and the large probe could be placed on the tissue of interest. The strain images produced by this software were displayed as real-time video in the same manner as conventional B-mode images and thus were captured and synchronized to the tracking data in exactly the same manner described previously. An example strain image acquired from a gel phantom containing a stiff inclusion is shown in Figure 9 along with the corresponding B-mode image.




Figure 9. Siemens commercial elastography software. A gel phantom containing a hard inclusion was imaged using B-mode (left) and strain imaging (right).

The Siemens elastography software had the advantage of providing real time strain images, but had the disadvantage of not allowing access to the raw data used to generate the images, and did not provide any quantitative measurement of calculated tissue displacements or relative strain values. These values would be useful in validation of the commercial strain images, as well as additional analysis of the underlying tissue behavior. The Acuson Antares ultrasound machine did,

however, also have a separate ultrasound research interface (URI) called Axis Direct which provided access to raw beamformed radiofrequency (RF) data. The unprocessed RF data could be collected during a normal imaging procedure by manually triggering the URI software on the ultrasound machine, which then saved RF data files to the hard drive. The files were then transferrable to a PC to be processed by the user. In this framework, it was extremely difficult to synchronize the raw RF data files with the external tracking system. Therefore there was a need for both the commercial strain imaging package and the URI, in order to get the benefits of tracking data and quantitative strain imaging, respectively. Although the raw RF data was not ultimately used in this dissertation due to the synchronization issue, an overview of strain image creation from RF signals will now be described to provide a general understanding of the process.

The raw RF data can be converted to strain images using a variety of algorithms [133, 134]. An ultrasound elastography algorithm used by Solbekk *et al.* was been implemented in Matlab (Mathworks Inc., Natick, MA) due to its simplicity and demonstrated efficacy in brain tumor imaging [135, 136]. This method was used for generating only axial strain images, as the axial resolution is greater than the lateral resolution in ultrasound images. The general procedure for generating a strain image begins with the acquisition of at least two frames of RF data while dynamically compressing the tissue of interest. The RF data is recorded as an array of voltage values generated by the piezoelectric elements in the transducer as the acoustic waves reflected by the tissue are recorded over time. A non-uniform distribution of scatters in the interrogated medium gives unique RF signatures throughout an image. An example of an RF frame and its equivalent processed B-mode image is shown in Figure 10 (acquired from a linear array transducer).




Figure 10. Example of unprocessed RF data (left) and processed B-mode image (right) from a gel phantom containing hard inclusion.

Assuming that the deformation applied to the tissue was very small, the basic shape of the post-compression echo, for a given window at a specific depth, will not have changed significantly in shape when compared with the pre-compression echo. However, there will be a phase difference between the two echo signals due to the difference in acoustic travel time arising from the change in distance to the probe brought about by compression. An illustration of this phase shift is shown for a single axial trace in Figure 11.

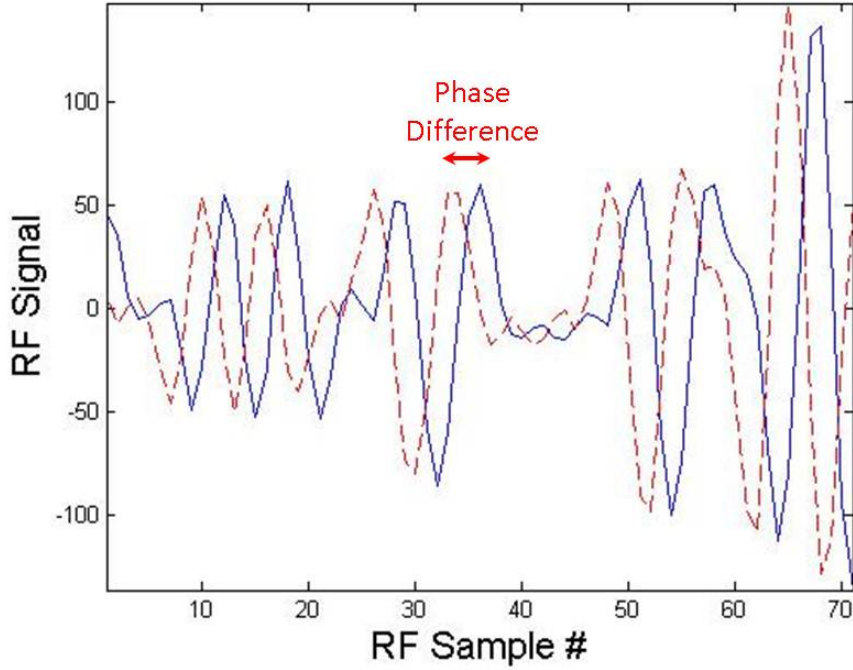


Figure 11. Pre-compression (solid line) and post-compression (dashed line) axial RF signals from two frames in a single window. These curves represent approximately the same signal separated by a phase difference.

Local tissue displacements may be estimated by exploiting the phase shifted signal described above. A computational framework for matching the two signals based on cross-correlation (CC) can calculate the lag between them, and thus the displacement which would have resulted in the phase difference. For an RF signal $r(m,n,k)$ at sample m of trace n in frame k , the cross-correlation at lag q between frames k and $k+1$ is found from:

$$\gamma(m,n;q,0) = \sum_{m'=-M}^M \sum_{n'=-N}^N r(m+m',n+n',k) \cdot r(m+m'+q,n+n',k+1) \quad (7)$$

In the algorithm used here, the cross-correlation function is only calculated along the axial traces, with zero lag in the lateral (n) direction. The correlation function needs a certain window size in order to generate accurate results. A window that is too large, however, will reduce the ability to detect smaller local displacements. The applied window size was generally selected based on visual

inspection of the resulting strain image. Given a sufficiently high frame rate, the tissue movement between consecutive RF frames was small enough that the maximum value of the cross-correlation function in each window occurred within a few samples of zero lag. An example of the cross-correlation function for one window is shown in Figure 12.




Figure 12. Cross correlation function for one window of RF data matched in the axial direction. Here, the maximum discrete value occurs at a lag of 2 samples.

In Figure 12, the maximum CC value appears at a time lag of 2 samples, which would suggest that the tissue displaced a distance of 2 samples. However, the CC equation only computes values discretely located at sample intervals and does not provide subsample estimates of the CC function. Further processing must be done in order to determine the true maximum value of the underlying function. One method of achieving this is to exploit a characteristic of the Hilbert transform, which is defined as:

$$\tilde{h}(t) = \frac{1}{\pi} \int \frac{h(t - \xi)}{\xi} d\xi \quad (8)$$

The Hilbert transform is used to create the analytic form of the CC function:

$$\hat{\gamma}(m, n, k) = \gamma(m, n, k) - j \cdot [\gamma(m, n, k) * \tilde{h}(m)] \quad (9)$$

where $*$ is the convolution operator. The analytic form of a signal created in this fashion is known to have some useful properties. The most relevant is that the phase of the analytic signal crosses zero at maximum values of the original signal [137]. The significance of this is that the true maximum value of the CC function may be estimated at subsample resolution via simple linear interpolation of the analytic phase values:

$$dt(m, n) = \left(\frac{-2 \cdot \angle \hat{\gamma}(m, n; q_{max}, 0)}{\angle \hat{\gamma}(m, n; q_{max} + 1, 0) - \angle \hat{\gamma}(m, n; q_{max} - 1, 0)} + q_{max} \right) T_{samp} \quad (10)$$

where dt is the estimated lag between the two RF signals, \angle is the phase operator, q_{max} is the discrete maximum lag of the CC function, and T_{samp} is the sampling time of the RF data. The equation above is derived from a simple line equation using a center-difference approach for the slope. At each zero crossing the phase angle function becomes approximately linear, and so the angles corresponding to q_{max} , $q_{max} + 1$, and $q_{max} - 1$ can be used to determine the subsample lag estimation. This interpolation procedure is illustrated in Figure 13.




Figure 13. Phase angle of the analytic CC function at each lag value. Each zero crossing corresponds to a maximum CC value. The exact zero crossing may be estimated by approximating the phase angle function as a line using three points (shown above).

The CC window slides in the axial direction down each RF trace and produces a field of time lag values using the procedure described above. Utilizing the assumption that the shift in signal is the direct result of tissue displacement, these lag values are converted to axial strain by differentiation:

$$\varepsilon(m, n) = \frac{dt(m + 1, n) - dt(m, n)}{T_{samp}} \quad (11)$$

An example of a completed strain image using this method is shown in Figure 14. A gel phantom containing a hard inclusion was sampled with the URI and RF data was collected to reconstruct both a B-mode image and strain image.




Figure 14. B-mode image (left) and strain image overlay (right). Both images were computed from raw RF data acquired from the research interface on the Acuson Antares machine.

III.3.7 Phantom Construction

This dissertation made use of phantom models prior to collection of *in vivo* data in order to validate the tracked ultrasound methods developed in later chapters. A variety of tissue-mimicking materials could be utilized to replicate characteristics of biological tissue, but many of these substances require the use of potentially hazardous reagents. Polyvinyl alcohol (PVA), however, is non-toxic and biocompatible. It is also known to have the interesting property of becoming

increasingly stiffer with the application of repeated freeze-thaw cycles (FTC), which makes it suited for elastography studies [138-140].

The creation of a PVA cryogel phantom begins by preparing the basic polymer mixture. A 7% w/v suspension of hydrolyzed PVA powder (Flinn Scientific, Batavia, IL) in cold water is first heated to 80°C. Glycerol (Fisher Scientific, Pittsburgh, PA) is then added at 10% by volume until the mixture is clear and fluid. The container is then covered tightly to minimize dehydration, and the mixture is cooled to room temperature while mixed by a magnetic stir plate. The final creation of the cryogel phantom occurs after full polymerization, achieved by the application of sequential FTCs in which the material is brought to -37°C over the course of 12 hours and then gradually returned to approximately 20°C over another 12 hours.

The phantom design used for most of the studies in this dissertation consisted of bulk PVA mixture poured into an organ-like mold, which was then allowed to go through one FTC. During this FTC, a tumor-like object was suspended in the mixture using wires, which could be removed easily by simply pulling them out of the final phantom. The tumor used varied depending on the requirements of the study. One type which lent itself to elastography studies was to use a small amount of PVA mixed with contrast agent in a tumor-like mold which was allowed to go through one or more FTCs by itself. When it was suspended in the bulk PVA mixture, it went through an additional FTC, and thus the tumor became stiffer than the surrounding phantom material. The contrast agent facilitated segmentation of the tumor in tomograms to establish a ground truth location, while the difference in material properties caused it to show up on ultrasound strain images.

In order to evaluate the performance of ultrasound elastography for detecting differences in phantom material properties, validation was performed by independent mechanical tests. To achieve this, a sample from the bulk phantom and from the phantom tumor were put aside during phantom construction. The samples were placed in polystyrene cell culture plates (Corning Inc.,

Corning, NY) and subjected to the same number of FTCs as the main phantom materials. This process resulted in cylindrical samples with diameter and height both about 15 mm, which was then subjected to compression testing using an ElectroForce 3100 material tester (Bose, Eden Prairie, MN) shown in Figure 15.



Figure 15. Mechanical tester (right) and workstation (left) which are used to obtain material properties of phantom samples to validate strain imaging.

Each sample was subjected to several cycles of a load rate by a transducer for small strains. The instrument was used to determine average elastic modulus values using the slope of the stress-strain curves of the steady-state loading phases. The ratio of moduli between the two materials was then used to validate the ultrasound strain images from the Siemens elasticity software.

CHAPTER IV

DESIGN AND EVALUATION OF AN OPTICALLY-TRACKED SINGLE-CCD LASER RANGE SCANNER

IV.1 Introduction and Significance of Study

This chapter will give an overview of the design and evaluation of a novel tracked laser range scanner device which supported the tracked ultrasound research presented in subsequent chapters. The primary intent of this device was to act as a relatively quick and accurate surface digitizer for intraoperative procedures. The dense geometric point clouds produced by the LRS are useful for surface-based registration methods, and also for measuring tissue deformation for model-updated image guidance. In this work, the LRS was described and subjected to a series of performance tests to establish its clinical efficacy. The geometric point cloud accuracy was determined using phantoms to be submillimetric, and the tracking accuracy of the system was found to be similar to other passive optical tracking tools. The significance of this study was to establish a gold standard registration and surface measurement tool to be used in Chapters V and VI. In addition, this work also led to the use of this device in studies targeting brain, liver, and breast surgery [141-144]. This work was published in Medical Physics in 2012.

Appearing in: Thomas S. Pheiffer, Amber L. Simpson, Brian Lennon, Reid C. Thompson, Michael I. Miga. Design and evaluation of an optically-tracked single-CCD laser range scanner. Medical Physics, Volume 39 (2), Pages 636 – 642, February 2012.

IV.2 Abstract

Purpose: Acquisition of laser range scans of an organ surface has the potential to efficiently provide measurements of geometric changes to soft tissue during a surgical procedure. A laser range

scanner design is reported here which has been developed to drive intraoperative updates to conventional image-guided neurosurgery systems.

Methods: The scanner is optically tracked in the operating room with a multi-face passive target. The novel design incorporates both the capture of surface geometry (via laser illumination) and color information (via visible light collection) through a single lens onto the same charge-coupled device (CCD). The accuracy of the geometric data was evaluated by scanning a high-precision phantom and comparing relative distances between landmarks in the scans with the corresponding ground truth (known) distances. The range-of-motion of the scanner with respect to the optical camera was determined by placing the scanner in common operating room configurations while sampling the visibility of the reflective spheres. The tracking accuracy was then analyzed by fixing the scanner and phantom in place, perturbing the optical camera around the scene, and observing variability in scan locations with respect to a tracked pen-probe ground truth as the camera tracked the same scene from different positions.

Results: The geometric accuracy test produced a mean error and standard deviation of 0.25 ± 0.40 mm with an RMS error of 0.47 mm. The tracking tests showed that the scanner could be tracked at virtually all desired orientations required in the OR set up, with an overall tracking error and standard deviation of 2.2 ± 1.0 mm with an RMS error of 2.4 mm. There was no discernible difference between any of the three faces on the LRS with regard to tracking accuracy.

Conclusions: A single-lens laser range scanner design was successfully developed and implemented with sufficient scanning and tracking accuracy for image-guided surgery.

IV.3 Introduction

An ongoing problem in the field of image-guided neurosurgery is the measurement and compensation of intraoperative brain shift. It is well understood that there is often significant movement of brain tissue between the time of preoperative imaging and the time of resection of

soft tissue [145, 146]. As the most common strategy for image-guidance relies solely on registering preoperative tomograms with the physical intraoperative coordinate frame, brain shift reduces navigational accuracy [147, 148]. Efforts to address the problem of brain shift have included methods to characterize intraoperative tissue deformation. Intraoperative imaging modalities are often utilized to provide updates to the pre-operative surgical plan derived from higher-resolution magnetic resonance (MR) or computed tomography (CT) images [149-151]. There has also been a movement toward using intraoperative ultrasound for shift measurement, as in the SonoWand (Trondheim, Norway) and BrainLab (Munich, Germany) systems [49, 57]. While these imaging systems do provide a quantitative measurement of brain movement, methods of compensating for shift in real-time for use in surgical guidance have not yet reached maturation.

Movement of the cortical surface is an attractive metric for brain shift, as it is readily observed and can provide intuition on the positions of internal structures of the brain. Any method which can capture and digitize the intraoperative surface of the patient could be used to provide quantitative measurements of shift. Once the surface has been acquired, it can be used to drive a number of shift compensation strategies. These strategies can include rigid or nonrigid registration of the surface to preoperative imaging to provide a corrective transformation to the guidance system [95, 152, 153]. Another approach is to use the acquired surface to drive a biomechanical model of the brain, which provides displacement updates throughout the imaged tissue [104, 154]. Sources of data may include intraoperative imaging modalities (such as intraoperative MR, CT, or ultrasound) or surface acquisition methods such as lasers range scanners (LRS). Each of these methods may be used to provide patient-specific boundary conditions for the mathematical model of the brain and thus present customized guidance to the surgeon. Regardless of the method used, real-time guidance requires that data acquisition be both fast and accurate.

LRS systems are traditionally used for geometric measurement of objects for which tactile means of measurement are either undesirable or infeasible [155-157]. As soft tissue deforms when

contacted, an LRS lends itself very well to surgical applications as a way of measuring geometry. Laser range scanners have been used for surface capture in a variety of procedures such as orthodontics, cranio-maxillofacial surgery, liver surgery, and neurosurgery. Table 1 summarizes a list of publications that examine the integration of various LRS devices into image-guided procedures.

Table 1. Recent examples of LRS integration into image-guided procedures.

Year	Author	Procedure	Application
2000	Commer <i>et al.</i> [158]	orthodontics	tooth position tracking
2003	Audette <i>et al.</i> [93]	neurosurgery	registration, brain deformation tracking
2003	Cash <i>et al.</i> [159]	liver surgery	registration of liver surface
2003	Marmulla <i>et al.</i> [160]	cranio-maxillofacial surgery	face registration
2003	Meehan <i>et al.</i> [161]	cranio-maxillofacial surgery	facial tissue deformation tracking
2003	Miga <i>et al.</i> [94]	neurosurgery	cortical surface registration
2005	Cash <i>et al.</i> [162]	liver surgery	liver deformation tracking
2005	Sinha <i>et al.</i> [97]	neurosurgery	cortical surface deformation tracking
2006	Sinha <i>et al.</i> [98]	neurosurgery	cortical surface registration
2008	Cao <i>et al.</i> [96]	neurosurgery	comparison of registration methods
2009	Ding <i>et al.</i> [163]	neurosurgery	semiautomatic LRS cloud registration
2009	Shamir <i>et al.</i> [95]	neurosurgery	face registration
2010	Dumpuri <i>et al.</i> [164]	liver surgery	liver deformation compensation

Conventional LRS devices work by sweeping a line of laser light onto the object of interest, and the surface is digitized by capturing the shape of the laser line with a digital camera and using triangulation to form a point cloud. Calibration is done to determine how points detected by the digital camera are mapped to the physical location of the laser line. The digital camera may also be used to collect texture information from the surface and map it onto the geometry to form a textured point cloud [97, 98, 159]. LRS systems are attractive for assisting image-guidance because they can provide relatively fast and accurate sampling of the entire exposed surface of the brain. Sun *et al.* have also used stereopsis via operating microscopes to capture the brain surface to address the problem of deformation [156]. This intraoperative information can be used both to align image-to-physical space as well as to track deformations. Alignment can be facilitated by tracking a conventional LRS in 3D space via optical targets attached to the exterior enclosure. In

addition to assisting with image-to-physical alignment, the role of an LRS in brain shift compensation is well defined by its ability to quickly acquire a series of scans over the course of surgery in order to track deformation. Work has also been done to use the texture associated with the point clouds to nonrigidly register a series of LRS scans, thus providing measurements of brain shift [97, 163]. Although the accuracy of LRS data has been encouraging, efforts to improve LRS-driven model-updated systems have highlighted aspects of conventional LRS design which could be altered to increase system fidelity and ease of use.

We present two fundamental contributions in this paper: 1) a tracked single-CCD LRS design and 2) an accuracy assessment of the new device. LRS devices which provide field-of-view colored point clouds are usually constructed from a two-lens design in which one lens captures geometric information from the laser line, and the other lens captures color information via a digital camera. The use of separate lenses unfortunately makes it necessary to create an additional calibration to map the 2D color information onto the 3D scanner point cloud. We present here a solution to this problem in the form of a single-lens system design. The novel LRS design was implemented and evaluated with the intent to use in cortical surface tracking; however, the LRS could be used to characterize any anatomy with sufficient surgical access.

IV.4 Materials and Methods

The following sections describe the two production phases of the new LRS: 1) the design and development decisions which composed the final system, and 2) an analysis of its scanning and tracking accuracy.

IV.4.1 Design and Development

This work presents the results of a collaborative effort to design a new laser range scanner to capture both geometric and field-of-view color information without the need for two lenses. Working in conjunction with engineers at Pathfinder Therapeutics Inc. (Nashville, TN, USA), we

developed a single-lens solution which is unique in that existing commercial systems such as 3D Digital (Sandy Hook, CT, USA) or ShapeGrabber (Ottawa, ON, Canada) products capture both geometric and color information with two lenses and two charge-coupled devices (CCDs), or do not collect color information at all. The older designs not only carry additional cost, but also require the overlay of color information onto the 3D point cloud. This process is another source of error, as each lens imparts a unique geometric distortion on the captured scene, and each lens also has a different line-of-sight to the target. One solution considered was to capture the field with a single lens and feed the geometry and color to two CCDs via a beam splitter. Ultimately it was decided that this option was less attractive in terms of cost, size, and complexity compared to a single-CCD approach.

The novel single-CCD solution here utilizes a Basler Pilot camera (Basler Vision Technologies, Ahrensburg, Germany) running at 1920x1080 at 32 fps. This camera is part of a family of cameras with uniform physical dimensions and electrical interfaces, which enables other camera models to be swapped out to meet varying scanning accuracy or speed requirements. To capture the geometry of the field, a standard red laser with a wavelength of 635 nm and a uniform line generator was selected. This wavelength was selected because of the wide availability of diode modules as well as its known reflectivity on the organs of interest (primarily brain and liver). One drawback to using a red laser is that the Bayer color filter pattern (which filters pixels to record color as either red, green, or blue before interpolation generates the final image) used on the CCD only assigns one out of every four pixels to capture red light, which effectively reduces the resolution of the scanner. Since the Bayer filter pattern assigns two out of every four pixels to green light, there was some consideration to using a green laser. However, this would result in reduced contrast of the laser on the background image in some of our intended applications, such as scanning the liver surface, as a red-brown object would tend to absorb green light. The laser line is swept across the field-of-view using a mirror attached to a standard galvanometer. The

galvanometer chosen can rotate over a 40° arc with approximately 15 bit precision and a settling time of about 0.1 ms. Using a video frame rate of 32 Hz, the maximum exposure length is 31.25 ms. The galvanometer is allowed to settle to its next resting position during the small window of time when the CCD is transferring data out to the frame buffer and is not actively collecting photons.

In order to maximize the scanning speed, the full frame rate of the CCD is used. At 8 bits per pixel, the CCD outputs data at a rate of 531 Mb/s. Conventionally, this high data rate would lead to a digital signal processor (DSP) based processing solution such that the point cloud could be calculated in the scanner and then transmitted to a host PC upon scan completion. However, by leveraging modern CPUs and high-speed communications links, the raw video frames are transferred to the host PC via gigabit ethernet for processing and calculation of the final textured point cloud.

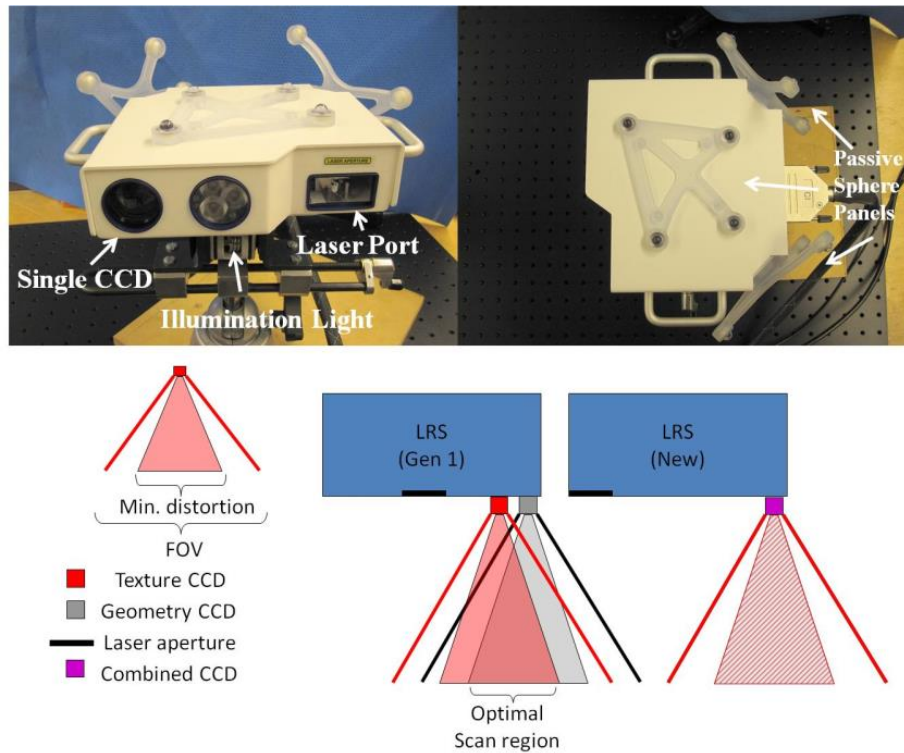


Figure 16. The novel LRS, showing the single CCD design from the front (top left picture) and the tracking marker configuration from the top of the scanner (top right picture). The bottom diagram shows the improved functionality of the single CCD system (right) in comparison with a dual-CCD system (left).

The external enclosure of the LRS was designed to be as small as possible to accommodate the following hardware: camera, lens, white-light illuminator, galvanometer, galvanometer driver board, motherboard containing the microcontroller with support circuitry, and passive tracking targets (see Figure 16). Within the enclosure, an internal structure was created to hold the camera, laser, and galvanometer perfectly rigid with respect to each other, as even slight changes in their relative positions would invalidate the scanner calibration. Although the calibration process and fixture are proprietary in nature, it can be stated that it is a semi-automated procedure in which the scanner is trained to measure distance, determine various optical parameters specific to the hardware used, and correct for geometric distortions.

The tracking electronics were originally designed to be compatible with the NDI Certus position sensor (Northern Digital Inc., Waterloo, ON, Canada) for active optical tracking. This first design included active tracking infrared emitting diodes (IREDS) that were the same height as the scanner enclosure. However, initial testing indicated that the error in triangulating the position of these diodes was too great, which led to the design and construction of an alternate configuration of IREDS. The attachment design of the marker housings was chosen to be modular, such that they could be changed easily depending on the application without needing to modify the scanner enclosure itself. The initial marker geometry on the LRS was replaced due to preliminary problems with marker visibility in the operating room. The active marker housings were replaced with reflective spheres compatible with the NDI Polaris Spectra position sensor for passive tracking. Specifically, a passive target was added to the top face of the LRS to increase the number of viable poses in the tracking volume, as it was not always possible to position the LRS within the confines of normal OR workspace such that at least one of the rear targets was visible to the tracking system. Other tracking systems could also be used, such as the NDI Polaris Vicra, but the relatively large work volume of the Polaris Spectra allows for greater flexibility in positioning the equipment, as it is not always possible to position the tracking system close to the patient, in our specific

application. The passive configuration is currently preferred in our work due to the ease of use of wireless tools and the existence of passive tracked surgical instruments and reference rigid bodies in the StealthStation (Fridley, MN, USA) workflow, currently used by our clinical colleagues at Vanderbilt University Medical Center. While the accuracy of active tracking was attractive, much consideration was given to the tradeoff between achievable accuracy and ease of integration (since the introduction of wired tools was intrusive to surgical workflow), and it was decided that passive tracking provided sufficient accuracy and minimized disruption in the OR.

IV.4.2 Laser Range Scanner Evaluation

The LRS accuracy was characterized over the course of two tests. The first test was designed to evaluate the accuracy of the geometric range scans. A multi-level platform phantom (see Figure 17) was scanned by a coordinate measurement machine such that the distances between disc centers were known to within a tolerance of 0.05 mm. The phantom was first used to determine the effective work volume of the scanner by the ability of the LRS to construct a geometric point cloud as it swept the laser across the phantom surface. Once the work volume was defined the geometric accuracy test was performed by mounting the LRS horizontally on an optical breadboard, facing the phantom. The LRS was held completely stationary, the phantom was moved systematically throughout the work volume, and the LRS scanned the phantom multiple times at each position. Nine positions in the work volume were used, consisting of three positions on each of three planes (see Figure 17) such that at least six of the discs were visible in any scan (non-central discs were occasionally outside of the LRS work volume due to field-of-view limitations). Ten scans were taken at each position for a total of 90 scans. From the acquired point clouds, the geometric centroids of each visible disc at each position were calculated. Then the relative distances between centroids were compared to the known disc distances.




Figure 17. Geometric accuracy test setup (left) showing the nine positions of the precision phantom (right).

The second test was to characterize the ability of the LRS to be tracked with respect to a global coordinate system defined by the tracking system and a reference target. The first part of this test was to observe the tracking behavior of the scanner. A rigid body file describing the LRS passive sphere configuration was generated by characterizing the LRS as a passive three-face tool within the NDI software. Each face was composed of four of the twelve markers and was divided into the planes formed by the top panel and two posterior panels, respectively (see Figure 16). The visibility of the spheres was tested by placing the camera and the LRS in “typical” operating room configurations. The camera was mounted horizontally at a height of approximately 2 m, whereas the LRS was mounted at a height of 1 m at a horizontal distance of 1.5 m directly in front of the tracking system. The relative positions of the tracking system and LRS were kept constant with respect to each other while the orientation of the LRS was incremented in its pitch (ψ) and yaw (θ) to simulate plausible orientations in the operating room (see Figure 18). The pitch was varied between 0, 45, and 90° with respect to the floor, and at each pitch angle the yaw was incremented by 30° through a full 360° rotation. The number of spheres tracked at each orientation was recorded using software provided by the manufacturer of the tracking system.




Figure 18. Orientations used in tracking visibility test. For reference, a pitch of 0° and a yaw of 0° denotes the orientation in which the top of the LRS is facing toward the camera, whereas a pitch of 90° and yaw of 180° denotes a horizontal orientation facing away from the camera.

The second tracking test was designed to observe the robustness of the rigid body file description for the passive targets attached to the LRS. A calibration was first performed to determine the transformation placing the raw point cloud into the coordinate frame of the tracked LRS rigid body [159]. The calibration is performed by scanning the block phantom described above, and then calculating the geometric centroids of the discs in the point cloud. Using the tracked LRS rigid body as the reference coordinate system, the locations of the discs are also digitized with a tracked pen probe. The scan centroid points are then fitted to the probe points with a standard least-squares method to produce a 4×4 calibration matrix which transforms scan points into the space of the LRS. The navigation software then automatically transforms the point cloud into the space of the reference target as the LRS is tracked. This means that all scans of the patient will be in a common coordinate frame.




Figure 19. Experimental setup used for the tracking accuracy test, showing the fixed phantom and LRS (left) and the Polaris Spectra optical tracking system (right).

After the scanner was calibrated, the LRS was positioned horizontally facing the block phantom such that the phantom was in the center of the LRS work volume. The LRS, block phantom, and reference target were fixed in place as shown in Figure 19. The Polaris camera was then moved between 30 positions distributed approximately 360° around the LRS such that the camera tracked each of the three faces of the LRS for ten of the scans. A scan of the phantom was acquired for each position of the camera, and the disc centroids were calculated in the coordinate frame of the reference target. In addition, the phantom discs were digitized with a tracked pen probe each time the phantom was scanned. These points were considered the gold standard positions for the discs, and the point cloud centroids were compared against them. While this gold standard was simple and convenient to create, it did inherently add error to the test, as there was tracking error associated with tracking the pen probe itself. There was also error in digitizing the discs, as placing the tip of the probe in the disc centers was a manual process. A more robust gold standard would entail a precision grid spanning the work volume of the tracking system throughout which the phantom and LRS setup could be stepped such that its position relative to the camera was known

with higher precision than achievable with passive optical tracking. However, the pen probe method used above was deemed to be more practical for this study. The 30 scans were analyzed as a group and in the three subsets corresponding to the different faces to determine variability in tracking the LRS.

IV.5 Results

The geometric accuracy test determined a mean error and standard deviation of 0.25 ± 0.40 mm, with an RMS error of 0.47 mm for the set of 90 scans acquired. The 95% confidence interval for this error was 0.24 to 0.27 mm. The maximum error encountered in this dataset was 1.6 mm.

The face visibility test indicated that in all of the tested LRS orientations except for one (in which the LRS was positioned vertically with its top face pointing away from the camera, *i.e.* a pitch of 0° and yaw of 180°) that the camera was able to track at least one of the faces. It should be noted that the NDI software (and navigation systems in general) only tracks a single face of a multi-face tool at a time. As each face on the LRS contains four markers, four is the maximum number of usable markers at any particular position or orientation.




Figure 20. Bitmap view of the phantom captured by the LRS (top left) and textured point cloud rotated slightly to show the 3D geometry of the data (top right). The bottom left shows an example of intraoperative data collection with the laser line sweeping across a brain surface. The corresponding LRS point cloud reconstructed from this scan is shown on the bottom right.

The second part of the tracking test resulted in a set of 30 scans such as the example in Figure 20. The nine disc centroids in each scan were determined and compared to the corresponding points collected by the pen probe. The results of this comparison are shown in Table 2, which shows the error across all 30 scans, as well as the error among just the 10 scans acquired while tracking each of the respective faces on the LRS. The mean overall error (across all scans) was 2.2 ± 1.0 mm, with an RMS tracking error of 2.4 mm. The 95% confidence interval for this error across all 30 scans was 2.1 to 2.4 mm. When the data was examined per face, it was found that Face 1 (rear right face) had a mean error of 2.1 ± 1.2 mm with an RMS error of 2.4 mm. Face 2 (top face)

had a mean error of 2.5 ± 1.1 mm with an RMS error of 2.7 mm. Face 3 (rear left face) had a mean error of 2.3 ± 1.0 mm with an RMS error of 2.5 mm.

Table 2. Comparison of the arithmetic and RMS tracking error over all 30 scans (total) and also over the separate groups of 10 scans taken while tracking each face (Face 1 is the rear right face, Face 2 is the top face, and Face 3 is the rear left face). The errors were calculated by comparing the centroid points of the phantom discs against points collected with a tracked pen probe.

	Mean Tracking Error (mm)	RMS Tracking Error (mm)
Total	2.2 ± 1.0	2.4
Face 1	2.1 ± 1.2	2.4
Face 2	2.5 ± 1.1	2.7
Face 3	2.3 ± 1.0	2.5

IV.6 Discussion

The results of the geometric accuracy test show that the average geometric scanning error is about half a millimeter with sub-millimetric standard deviation, which is acceptable for the intended applications. A previous generation LRS using a dual-CCD design was reported to have a scanning accuracy of 0.3 mm at best, and its performance degraded outside of the center of the work volume at least in part due to the computational error in aligning the texture and geometric information from their respective CCDs [97]. It is possible to increase the resolution of the point cloud through the scanner API by collecting more range points, at the cost of scanning speed.

The face visibility test showed that four markers were visible on average to the tracking system at virtually all of the tested positions, which provided enough markers to compute the LRS position and orientation in the reference frame. It was important to conduct this test in order to determine practical positioning limitations of the LRS with respect to the optical tracking system in the operating room. Equipment logistics are often beyond the control of research engineers in the operating room due to the requirements of normal surgical workflow, which necessitates flexibility in the positions in which the LRS can be tracked.

The second part of the tracking test showed that the error in LRS point cloud locations in the reference frame is on the order of normal optical tracking error for passive systems. The

tracking accuracy of the previous generation scanner used by Sinha *et al.* was similar at 1.0 ± 0.5 mm [97]. However, the LRS in that case utilized actively-emitting IREDS rather than the passively tracked solution used for the new LRS. The data in Table 2 show the mean error in point cloud locations in the reference frame for all 30 scans and for each subset of 10 scans per tracked face. It is apparent that the error is approximately equivalent across all three faces, which implies that the accuracy of scan tracking is insensitive to the orientation in which the LRS is viewed by the optical tracking system (for configurations used in the OR).

IV.7 Conclusions

The design goal of a single-CCD LRS capable of capturing both geometric and color information was met in terms of possessing sub-millimetric scanning accuracy and tracking accuracy that is typical of passive tracking systems (on the order of 2 mm). It was evaluated with regard to its scanning accuracy and tracking ability using a precision phantom and found to be appropriate for image-guided procedures. While the overall error in the system is approximately 2 mm (primarily contributed by the tracking of the LRS), this performance is similar to the 3D Digital LRS used by Dumpuri *et al.* to provide a TRE of 2 to 4 mm in liver phantom targets, and also used by Cao *et al.* to provide a TRE of about 2 mm for cortical targets [96, 164]. We have integrated the LRS into our guidance software and are currently evaluating its contribution to our shift correction system.

Acknowledgements: This work is funded by the National Institutes of Health: grant R01 NS049251 of the National Institute for Neurological Disorders and Stroke, and R01 CA162477 of the National Cancer Institute. Dr. Miga is a co-founder and holds ownership interest in Pathfinder Therapeutics Inc. which is currently licensing rights to approaches in image-guided liver surgery and has incorporated the reported laser range scanner within their product line.

CHAPTER V

MODEL-BASED CORRECTION OF TISSUE COMPRESSION FOR TRACKED ULTRASOUND IN SOFT-TISSUE IMAGE-GUIDED SURGERY

V.1 Introduction and Significance of Study

This study is the first contribution toward the improvement of tracked ultrasound as an intraoperative guidance platform. It arose as a response to the problem of registration error induced by tissue compression during ultrasound imaging, especially strain imaging. It is well known that routine ultrasound imaging can cause deformation on the order of 1 cm, which is a clinically unacceptable error in many surgical procedures. This phenomenon was detected early in the deployment of tracked ultrasound in tandem with the LRS device from Chapter IV. The tracked LRS point clouds and the tracked ultrasound data are automatically recorded in a common coordinate space, and in the absence of compression error it would be expected that the top of the ultrasound image would align with the LRS surface. In practice, however, it was observed that the ultrasound data tended to be located well below the LRS point cloud due to the pressure exerted by the user. This chapter presents a strategy for correcting this error by using knowledge of the pose of the ultrasound probe surface within the tissue of interest. An initial image-to-physical registration of the tracked ultrasound to a patient-specific finite element model must be done in order to calculate this pose. After the registration, the pose of the ultrasound probe within the tissue is used to assign boundary conditions to the finite element tissue model. The solution of the model is then reversed to provide an estimate of the tissue in the uncompressed state. This strategy was found to be capable of reducing errors of approximately 1 cm to a clinically useful 2 to 3 mm. This work was published in *Ultrasound in Medicine and Biology* in 2014.

Appearing in: Thomas S. Pheiffer, Reid C. Thompson, Daniel C. Rucker, Amber L. Simpson, Michael I. Miga. Model-based correction of tissue compression for tracked ultrasound in soft-tissue image-guided surgery. *Ultrasound in Medicine and Biology* 40 788-803 (2014).

V.2 Abstract

Acquisition of ultrasound data negatively impacts image registration accuracy during image-guided therapy due to tissue compression by the probe. We present a novel compression correction method which models subsurface tissue displacement resulting from application of a tracked probe to the tissue surface. Patient landmarks are first used to register the probe pose to preoperative imaging. The ultrasound probe geometry is used to provide boundary conditions to a biomechanical model of the tissue. The deformation field solution of the model is inverted to non-rigidly transform the ultrasound images to an estimation of the tissue geometry prior to compression. Experimental results with gel phantoms demonstrated that the proposed method reduced the tumor margin Modified Hausdorff Distance (MHD) from 5.0 ± 1.6 mm to 1.9 ± 0.6 mm, and reduced tumor centroid alignment error from 7.6 ± 2.6 mm to 2.0 ± 0.9 mm. The method was applied to a clinical case, and reduced the tumor margin MHD error from 5.4 ± 0.1 mm to 2.6 ± 0.1 mm, and the centroid alignment error from 7.2 ± 0.2 mm to 3.5 ± 0.4 mm.

V.3 Introduction

Ultrasound is commonly used as an intraoperative imaging modality to monitor surgical targets such as tumors. The need to maintain acoustic coupling between the probe and tissue often results in significant compression of the target by the user. This is especially a concern when using ultrasound strain imaging, in which a certain level of pre-compression of the tissue may be necessary. However, this tissue deformation affects the geometry of the scanned objects and the

resulting images. Soft tissue can undergo surface compression on the order of 1 cm during routine freehand imaging [165, 166]. This leads to incorrect estimates of the size and location of landmarks within the ultrasound images.

Compressional effects from the probe are especially apparent in image-guided procedures, which align intraoperative data with preoperative tomographic images. In these procedures, it is important that data collected during the surgery is accurately registered to high-resolution computed tomography (CT) or magnetic resonance (MR) image volumes for optimal guidance. Typically this is done by digitizing physical landmarks on the patient with a tracked instrument, selecting the corresponding landmarks in the tomograms, and computing a rigid transformation which best aligns the two coordinate spaces. Although there are a variety of methods to track and calibrate an ultrasound probe such that each image slice is recorded with a known pose in physical space [71, 72, 74, 123, 167, 168], the usefulness of tracked ultrasound relies on an accurate registration. Registration accuracy is compromised by non-rigid tissue deformation such as that which occurs with manipulation of the ultrasound probe. The goal of this work is to improve the usefulness of tracked ultrasound in image-guided procedures by improving this registration.

There are several approaches in the literature which have sought to address the problem of tissue deformation exerted by an ultrasound probe. One method is to create a digital representation of the surface and then use a combination of Bayesian theory and prior knowledge of the surgical scene to create a deformation which matches the observed ultrasound data [169], but this approach did not incorporate a physical model of tissue which could be used to provide more realistic priors. Another approach is to acquire B-mode or raw radiofrequency (RF) data from the ultrasound and use non-rigid image-based registration and positional tracking to correct for deformation [166, 170], but this approach requires a series of ultrasound images in order to provide sequential estimates of compression correction. There has also been work done to model tissue compression using data from a force transducer attached to the ultrasound probe along with a position sensor to

drive a tissue model [171, 172]. Our proposed method is similar to this method, but eliminates the need for a force measurement apparatus on the probe by using measured 3D surface displacements, rather than force, to drive the model. Our method utilizes just the tracking system which is routinely used in surgical procedures such as image-guided neurosurgery. To our knowledge, there has not been an attempt to model the tissue deformation from the physical probe surface itself in the correction. This work presents a compression correction method which measures and compensates for this effect using a biomechanical tissue model with validation in simulations, phantoms, and a preliminary clinical case.

V.4 Materials and Methods

We present our compression correction method as one component in the context of a patient-specific data pipeline for image-guided therapy. Prior to correction, we perform several data acquisition and processing steps. The procedures described below were used in all phantom experiments and were similar for the acquisition and analysis of clinical data.

V.4.1 Phantom Construction

Two compliant phantoms were each constructed by mixing 7% by mass polyvinyl alcohol (PVA) in water, 10% by volume glycerol, and heating to 80 °C to ensure saturation [138, 140]. For each phantom, a smaller amount of PVA was treated with barium sulfate powder for CT contrast and poured into a separate mold to act as the tumor target. The tumor was subjected to five freeze-thaw cycles in which it was frozen at -40 °C for 12 hours and then thawed for 12 additional hours, in order to produce a stiffer material. The tumor was then suspended in the bulk phantom mixture and the phantom underwent one freeze-thaw cycle to produce a tissue-like phantom containing a stiff tumor. The volumes of the tumor and bulk phantom mixtures were 3.2 cm³ and 720 cm³, respectively. The stiffness properties for the bulk tissue and tumor were tested using small samples with an ElectroForce 3100 instrument (Bose, Eden Prairie, MN). One of the phantoms was

constructed in a small cup-like container covered in fiducial markers and was used for the baseline accuracy test described in the Phantom Experiments section. The second phantom was fixed to a rigid base which contained 8 evenly distributed fiducial markers used in the image-to-physical registration and was used to test the compression correction method.

V.4.2 Patient Model from Preoperative Image Volume

CT image volumes of the phantoms were acquired using a clinical CT machine. This data simulated a typical preoperative tomogram acquisition, and was defined in the experiment as the baseline undeformed state against which our corrected ultrasound data would be compared. All volumes were 512 x 512 x 422 with 0.6 mm isotropic voxels. The phantom structures were segmented using intensity thresholding tools within Analyze 9.0 (Mayo Clinic, Rochester, MN). Isosurfaces were generated from the bulk phantom and tumor segmentations via the marching cubes algorithm, and were smoothed using a Laplacian filter. A tetrahedral mesh was generated from the segmentation surfaces using custom-built mesh generation methods [113]. One phantom and mesh are shown in Figure 21.




Figure 21. Experimental setup with a gel phantom attached to a base with fiducial markers (a), and the digital phantom and tumor surfaces segmented from the CT image volume (b).

V.4.3 Intraoperative Data Collection

All ultrasound images were acquired with an Acuson Antares ultrasound machine (Siemens Inc., Munich, Germany), using a VFX13-5 linear array probe with a 6 cm depth setting at 10 MHz. The ultrasound unit was also capable of producing strain images via the eSie Touch elasticity software. For the compression correction experiment, B-mode images were collected as well as strain images, and both types of images were analyzed in order to evaluate the effect of correction on target locations in ultrasound images having different contrast mechanisms. Ultrasound data was tracked in 3D space by synchronizing the ultrasound video and tracking data using software based on the Visualization Toolkit (VTK) on a host PC [120, 121]. The video was captured by a Matrox Morphis Dual card (Matrox Imaging, Dorval, Canada), which recorded the analog video output of the ultrasound machine in real-time. A passive optical tracking rigid body (Northern Digital, Waterloo, ON, Canada) was fixed to the ultrasound probe as shown in Figure 22. The pose of the rigid body was measured by a Polaris Spectra (Northern Digital Inc., Waterloo, Canada) optical tracking system. The tracked ultrasound system was calibrated using a method which relies on imaging a tracked tool in the ultrasound plane [74]. Once the tracked ultrasound system was calibrated, all pixels in each image were associated with a 3D pose.




Figure 22. Polaris tracking system (a), and the laser range scanner with passively-tracked targets (b). A passively-tracked target was also used to track the ultrasound probe (c), and the laser range scanner was used to construct a digital representation of the probe surface (d).

V.4.4 Proposed Compression Error Correction

The goal of our project was to reduce the error associated with rigid registration of the tracked ultrasound data to preoperative volumetric image data by additionally correcting for tissue compression. The general strategy was to use probe tracking information in combination with co-registered tomograms to calculate the distance into the tissue that the probe was pressed, and then use that distance to correct the tracked ultrasound image pose. We chose to first evaluate rigid translations of the ultrasound pixels based on this calculated compression distance as a simple, yet naive correction. This method would have the benefit of very low computational expense, and may be sufficient for aligning subsurface targets to within a surgically relevant threshold. However, this rigid approach would not account for non-rigid compressional effects throughout the ultrasound image itself. Our proposed correction method instead used the compression distance as input to a

more sophisticated tissue-mechanics model-based approach, which non-rigidly deformed the ultrasound image in a physically realistic fashion. Both the naive rigid correction and our model-based correction are outlined in Figure 23 and described below, and both methods were performed for comparison in terms of alignment error reduction during simulations, phantom experiments, and a clinical case described in the following section.




Figure 23. Procedure for the rigid correction and model-based correction. The rigid correction was performed by calculating the distances between the top of each A-line and the co-registered tomographic surface (a) and then translating the A-lines upward in the depth direction by those distances (b). The model-based correction was performed by calculating the distances from the tomogram mesh nodes to a digital representation of the ultrasound probe surface, and designating the distances as boundary conditions to a FEM model (c). The model was used to solve for tissue displacements throughout the mesh (d) and then reversed to deform the ultrasound image to its pre-compressed state (e).

V.4.5 Rigid Correction

The rigid correction procedure stems from the notion that, in the absence of registration error from tissue compression, the near-field row of each A-line in a tracked ultrasound image should ideally be aligned with the organ surface from the co-registered preoperative tomogram. In the presence of compression error, our tracked ultrasound capabilities would predict that the near-field pixels in the image would typically be some distance below the organ surface. Thus, the correction consists of calculating the distances between the unperturbed surface and the near-field ultrasound image pixels representing the interface between transducer and organ, and then rigidly translating each A-line of the image by that distance opposite the direction of compression to render the ultrasound image in the appropriate preoperative image space as shown in Figure 23a-b. The translation of each A-line is calculated by assuming the direction of compression occurred purely in the depth direction of the ultrasound image, and vectors are created for each pixel in the top row of the image with origin at that pixel and direction opposite of compression. The points of intersection on the surface are determined for each vector, and the distances from the pixels in the top row to these points are calculated to provide the translational correction vectors for the each A-line in the image.

V.4.6 Model-Based Correction

The model-based correction procedure consists of modeling the volumetric tissue displacement exerted by the ultrasound probe. Knowledge of the exterior geometry of the probe is a necessary component of this model. A laser range scanner (LRS) was used to scan the face of the ultrasound probe in order to create a 3D point cloud representing the probe surface as shown in Figure 22d. The accuracy of the LRS device was previously characterized and known to have geometric RMS accuracy of approximately 0.5 mm [173]. The LRS was tracked by the same tracking system as the ultrasound probe during the scan, and thus the created point cloud had a known pose based on the pose of the probe. As a consequence of this relationship and the tracking of the

ultrasound probe during data acquisition, the position and orientation of the digital probe surface was known for every ultrasound image collected. The LRS point cloud of the probe geometry was converted to a smooth surface using radial basis functions (RBF), and this surface was used for all subsequent calculations (FastRBF, FarField Technology Ltd., Christchurch, New Zealand). It should be noted that the LRS scan of the probe surface was only necessary in the absence of 3D geometry files from the ultrasound probe manufacturer, which could be used to provide an equivalent geometry.

The model-based correction method is illustrated in Figure 23c-e. The general approach is to use the pose of the probe surface to create boundary conditions for a forward FEM model solution to predict tissue deformation. Due to the standard image-to-physical registration that is done in image-guided interventions, the digital probe surface is placed in the same coordinate space as the patient-specific FEM mesh. Provided a good initial registration between the patient and preoperative imaging, we assume that the probe surface is located within the mesh at some distance beneath the surface depending on the magnitude of compression exerted by the user. The distances from the probe surface opposite the direction of compression to the mesh surface are computed automatically in a similar fashion as in the rigid correction method. The computed distances are then assigned as Dirichlet boundary conditions to the model. This is accomplished by first generating initial mesh boundary conditions which describe the tissue at rest, without the influence of the ultrasound probe. For the phantom experiments, the far-field of the mesh was set as fixed due to the phantom base enclosure and the near-field and surrounding regions of the mesh were set to stress-free as shown in Figure 24a-b. The top of the phantom was considered stress-free because that portion of the phantom was left exposed to the atmosphere and was not subjected to any external forces. The ultrasound probe was applied to the top of the phantom in the stress-free region. The set of initial boundary conditions are then altered to reflect the position of the probe in

the tissue. Thus, for each ultrasound image the appropriate mesh nodes were changed from stress-free to Dirichlet with the probe displacement values as shown in Figure 24c.

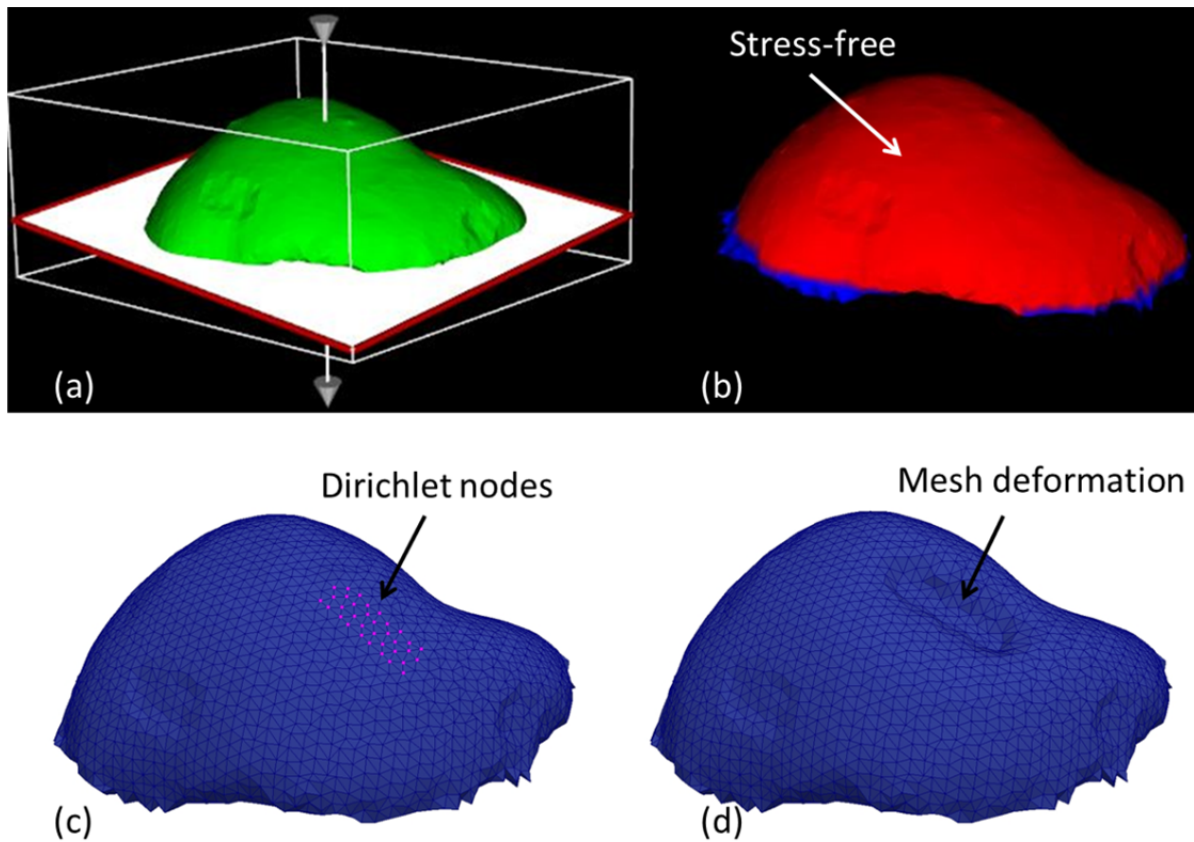


Figure 24. Initial boundary conditions were assigned manually (a) to designate the bottom fixed, and the top as stress free (b). These boundary conditions were then modified for each model-based correction by changing the appropriate nodes to Dirichlet displacement conditions (c) based on the position of the ultrasound probe surface, and then the model was solved to produce an approximation of the deformed tissue (d).

After the boundary conditions are generated, they are used to drive a FEM biomechanical model to solve for 3D displacements throughout the tissue. A linear elastic model was used in the phantom experiments, as the displacements were relatively small and the phantoms were largely homogenous. The model consists of the standard 3D Navier-Cauchy equations for the displacement field:

$$\frac{E}{2(1+\nu)}\nabla^2\mathbf{u}+\frac{E}{2(1+\nu)(1-2\nu)}\nabla(\nabla\cdot\mathbf{u})+F=0 \quad (12)$$

where E is Young's modulus, ν is Poisson's ratio, \mathbf{u} is the 3D displacement vector at a point in the tissue, and F is the applied body force distribution. The partial differential equation is solved using the Galerkin weighted residual technique with linear basis functions. The system of equations that solves for the displacement vectors at every node in the mesh can be written as:

$$[K]\{\mathbf{u}\} = \{\mathbf{f}\} \quad (13)$$

where K is the global stiffness matrix, \mathbf{u} is the vector of nodal displacements, and \mathbf{f} contains the contributions of any applied body forces or surface movement at each node. This system of equations is solved for the nodal displacements which satisfy static equilibrium for the supplied boundary conditions.

These nodal displacement vectors are used to deform the mesh nodes to the compression state exerted by the probe as shown in Figure 24d, giving a deformed mesh in which each node is associated with a displacement vector. We can then interpolate displacements corresponding with the pixels associated with the co-localized ultrasound image slice and apply the reversed field to simulate the uncompressed state, i.e. the reverse displacement vectors are used to non-rigidly undeform the ultrasound image to an estimation of the tissue in its uncompressed state.

V.4.7 Experimental Validation

V.4.7.1 Simulations of Proposed Correction

The proposed model-based correction method primarily relied on the assumption that the ultrasound probe would be applied with pressure directly normal to the tissue surface during data collection. The limits of this assumption were tested using a simulation dataset and applying the rigid and model-based corrections for a set of tissue compression scenarios in which the trajectory of the probe into the tissue was varied.

A patient-specific model was constructed from the preoperative MR volume of a low-grade glioma patient at Vanderbilt University Medical Center who previously gave informed consent with approval of our Institutional Review Board. The brain surface and contrast-enhanced tumor were segmented and a patient-specific model was constructed as a tetrahedral mesh. For simplicity, the brain and tumor were approximated as homogenous linear elastic materials with Young's modulus of 2100 Pa and Poisson's ratio of 0.45, taken from the values used by Dumpuri *et al.* for white and gray matter. There are few consistent values reported in the literature for brain material properties *in vivo*, but the values reported here fall within a range of values proposed by various groups [174, 175].




Figure 25. Simulation methodology. The probe surface was inserted into the mesh to simulate its position during ultrasound imaging of the tumor (a). Nine sets of boundary conditions were then generated to simulate various possible probe-tissue contact scenarios by rotating the displacement vectors about angles θ_1 and θ_2 defined with respect to the probe orientation (b). The solid arrows descending from the brain to the surface represent the assumption of compression in the depth direction, whereas the dashed arrows are an example of how the displacement vectors were systematically rotated to test that assumption (here $\theta_2 = 30^\circ$). Each set of boundary conditions was used to drive the forward model and deform the mesh and tumor (c). The proposed rigid and model-based corrections (using the assumption of compression perfectly parallel with the depth direction) were then applied to each deformed tumor and compared to the original uncompressed tumor (d).

First, the digital probe surface was manually inserted into the mesh over the tumor to simulate misalignment between the tracked ultrasound and co-registered MR volume due to tissue compression, as shown in Figure 25a. The compression distance from the brain surface down to the probe surface was approximately 7.5 mm, which is within the 1 cm compression magnitude often observed in routine imaging [165, 166]. An initial set of boundary conditions was created for the brain mesh. A patch of boundary nodes corresponding to a hypothetical craniotomy region above the tumor was designated as stress-free, the brain stem was set as fixed, and then the rest of the brain surface nodes were designated to have zero normal displacement, with stress-free tangential components to allow slip along the skull. This slip condition has been observed and used in modeling brain shift previously [176-178]. Then, vector distances from the brain surface nodes to the probe surface were calculated and used as Dirichlet boundary conditions for a forward solve of the linear elastic model. In order to determine the impact of our assumption of purely axial probe compression, the orientation of these vectors was systematically altered across the pitch and yaw angles (θ_1 and θ_2 in Figure 25b) of the probe surface from -30 to 30° in steps of 15° . For each of these vector alterations, the probe surface remained stationary, and the only change was the selection of brain surface nodes which were designated to displace toward the probe with Dirichlet conditions. These scenarios were intended to illustrate the primary limitation of our approach, which is a lack of known correspondence between the tissue surface and the ultrasound probe. While tracking the probe gave knowledge of its location within the tissue, its trajectory into the tissue to arrive at that location was not necessarily known due to tissue contact uncertainty. These forward boundary conditions were used to deform the brain and tumor to simulate nine possible physical conditions (one for each of the angle rotations) during an ultrasound data acquisition.

Finally, the proposed compression corrections were applied to each of the deformed tissue models generated from the boundary condition variants from above, and resulted in nine corrected

tumor locations. The corrections incorporated the assumption that the probe was compressed purely in the depth direction of the image plane, even for the simulations in which the compression trajectory was at an angle not parallel to the depth direction. The average distance from tissue surface to probe surface was used in the rigid correction method to translate the entire tumor in the direction of the mesh surface. The model-based procedure was used as described previously in order to create an inverted displacement field and to correct the tumor. Each of the simulated tumor corrections was then compared to the gold standard uncompressed tumor from the original patient model prior to deformation, in terms of both the boundary node distances between the corrected and gold standard tumor and the distance between the tumor centroids.

V.4.7.2 Phantom Experiments

We performed three validation tests of our method using phantoms. The first was to image a phantom through a layer of water with no contact between probe and phantom, in order to assess the baseline alignment accuracy of the system in the absence of compression. Second, in order to test our method we performed an experiment with an anthropomorphic phantom containing a stiff, contrast-enhanced lesion. We compared the tracked ultrasound tumor borders and centroid locations with the equivalent data from co-registered CT as a gold standard before and after compression correction. Third, we also evaluated the effect of tumor elasticity on our model-based correction by using the actual tumor-to-phantom stiffness ratio from material testing, and compared this to the correction using our assumption of tissue homogeneity.

The two phantoms were first constructed as described previously. The phantom fiducial markers were localized in physical space with a tracked stylus, and the corresponding marker positions in the CT volumes were also recorded. The image-to-physical registration was computed using a standard point-based method. The registration produced a transformation matrix which was then automatically applied to all tracked ultrasound images to align them with the CT data.

After the registration was calculated for each phantom, the first phantom which was constructed in a cup-like container was covered with a thin layer of water and then imaged with freehand tracked ultrasound for a total of 116 B-mode images in several sweeps while avoiding direct contact with the phantom surface. With respect to this contained phantom, the purpose was not to assess our compression technique with realistic geometries but instead to quantitatively assess geometric accuracy of a reconstructed target based on tracked ultrasound images without the presence of deformation. Each captured image was stored along with the concatenation of calibration, tracking, and registration transformations describing the 3D location of each slice, and then was compared to the co-aligned CT surface with metrics described below.

The second phantom was then imaged normally with full contact between probe and phantom with the more realistic organ geometry. The tumor was first fully outlined with B-mode using several slow sweeps in approximately 1 mm increments. A total of 178 B-mode images were collected. Then, the elasticity software was used to create strain images of the tumor. A total of 83 strain images were collected. The tracking and registration transformations were applied to the LRS probe surface for each image in order to generate boundary conditions as described previously. Following the ultrasound data collection from this phantom, the rigid and model-based correction methods were applied to each ultrasound image. The result was a collection of uncorrected and corrected images. Each type of correction was evaluated by comparing each population of images to the baseline CT images in terms of tumor geometry. In each B-mode and strain image, the tumor borders were segmented semi-automatically using an implementation of the Livewire technique. For each ultrasound image, the CT volume was re-sliced to provide a co-planar CT slice according to the 3D location and orientation of the co-registered ultrasound slice. The CT tumor borders were segmented using intensity thresholding in Analyze 9.0. The tumor borders segmented from ultrasound images and CT were then compared to each other in terms of the Modified Hausdorff

Distance (MHD) between the two contours, as well as the average distance between the contour centroids [179]. The MHD value is defined for two sets of points A and B as follows:

$$d(A, B) = \frac{1}{N_a} \sum_{a \in A} \min_{b \in B} (\|a - b\|) \quad (14)$$

$$\text{MHD} = \max(d(A, B), d(B, A)) \quad (15)$$

where $d(A, B)$ is the mean closest point distance in the direction of A to B , a and b are respectively points in sets A and B , and N_a is the number of points in set A . Thus, the MHD value is created by calculating the mean closest point distance from A to B , and then B to A , and then choosing the maximum of those two values. This metric was chosen as it was less prone to underestimating the error between tumor borders than calculating just a mean closest point distance from the ultrasound to CT contours. The MHD and centroid distance comparison was made for every slice before and after correction and were the primary metrics used to evaluate the improvement offered by our proposed model-based correction algorithm. In addition, the tumor volume measured during phantom construction was recorded for comparison with the CT tumor volume, the volume measured by tracked B-mode in the baseline accuracy test, and the volume of the model-corrected B-mode and strain data.

The third phantom experiment was performed to test the effect of tumor elasticity on the model-based correction. As mentioned previously, one of the simplifications of the method is the assignment of homogenous material properties to the patient specific FEM mesh, which implies that the tumor-to-bulk stiffness ratio in terms of Young's modulus is assumed to be 1:1. The impact of this assumption was tested by instead using the material testing data done on the phantom materials during the correction, which was a tumor-to-bulk stiffness ratio of 9:1. The 1:1 and 9:1 model-based corrections for the 178 B-mode and 83 strain images were then compared to each other in terms of MHD and centroid distances. PVA is known to be nearly incompressible, so Poisson's ratio was set to 0.49 for all corrections.

V.4.7.3 Clinical Case

In addition to the phantom study, we also investigated the feasibility of our method on a preliminary clinical case. As in the case of the simulation dataset, informed written patient consent was obtained for this work. A preoperative MR volume was used to construct the patient model. The tumor in this case was determined to be a meningioma, located on the left side of the brain just beneath the surface. Material properties for the linear elastic model were those used in the simulation, based on the average brain tissue properties used by Dumpuri *et al.* [180]. Intraoperative tracked ultrasound images were aligned to the patient model by first collecting a scan of the patient face with a tracked LRS and then registering the point cloud to the corresponding MR surface using the iterative closest point (ICP) algorithm [96]. After the craniotomy was completed, tracked B-mode images were obtained of the tumor. The rigid and model-based correction methods were then applied to the ultrasound data and were evaluated in terms of MHD between the co-aligned B-mode and MR tumor borders, as well as the distance between the co-planar tumor centroid locations.

V.5 Results

V.5.1 Simulations

The results of the correction simulations are shown in Table 3. These simulations were designed to validate the assumption of compression applied purely in the depth direction. The corrected tumor volumes were compared to the original uncompressed tumor volume in terms of node positional error around the tumor boundary, as well as the distances between tumor centroids. Table 3 reports in every case that the model correction outperformed the rigid-based method regardless of application inaccuracy. Boundary errors ranged between 2.7 and 4.3 mm for the rigid method and below 2.8 mm for all model-corrected. Similarly, centroidal errors ranged 2.4-3.8 mm and less than 2.6 mm respectively. On average over all cases, the model correction

improved localization over the rigid method by 55% and 50% for the mean boundary node and centroid error, respectively.

Table 3. Simulation results to assess the assumption of compression purely in the depth direction of the ultrasound image plane. The two angles are the pitch and yaw of the ultrasound probe about which the displacement vectors were rotated when assigning Dirichlet boundary conditions. The mean and standard deviation of the boundary node error and the mean centroid distance were calculated in comparing each corrected simulation tumor with the original uncompressed tumor.

θ_1 (°)	θ_2 (°)	Rigid Correction		Model Correction	
		Mean and SD of Boundary Node Error (mm)	Centroid Error (mm)	Mean and SD of Boundary Node Error (mm)	Centroid Error (mm)
0	0	2.7 ± 1.5	2.4	0.0 ± 0.0	0.0
15	0	3.0 ± 1.3	2.5	1.1 ± 0.4	1.2
-15	0	3.3 ± 1.3	2.8	1.2 ± 0.5	1.1
0	15	2.9 ± 1.4	2.4	1.0 ± 0.4	1.0
0	-15	3.2 ± 1.2	2.7	1.0 ± 0.4	1.0
30	0	4.0 ± 1.0	3.6	2.5 ± 0.8	2.3
-30	0	4.3 ± 1.1	3.8	2.8 ± 1.2	2.6
0	30	3.5 ± 1.0	2.9	2.2 ± 0.8	2.1
0	-30	4.0 ± 0.9	3.4	2.3 ± 1.0	2.2

V.5.2 Phantom Experiments

The baseline accuracy test using the cup phantom provided an assessment of the best alignment of tracked ultrasound and CT which could be achieved using only tracking and point-based registration. Using the 116 images from this test, the MHD between the B-mode and CT tumor contours was 1.2 ± 0.4 mm, and the average centroid error was 1.7 ± 0.6 mm. Additionally, the tumor volume enclosed by the ultrasound contours was found to be approximately 3.0 cm^3 , whereas the volume given by the CT segmentation was 3.3 cm^3 . For comparison, the volume of PVA mixture used for the tumor during construction was 3.2 cm^3 .

The results of the phantom experiments are shown in Figure 26 and Figure 27. We showed previously in Figure 24 an example of the FEM mesh of the phantom organ in its original undeformed state, as well as in its deformed state after the probe compression has been modeled

for a particular ultrasound slice. We note qualitatively that the displacements displayed by the model match very well with those observed in the physical phantom during data acquisition. Figure 26 shows examples of B-mode and strain images before and after compression correction using the rigid method and model-based method. The segmented tumor contour from each ultrasound image was outlined, and the CT tumor surface was overlaid for comparison. The MHD between the ultrasound borders and CT borders was 5.0 ± 1.6 mm for B-mode and 5.6 ± 1.1 mm for strain images prior to correction. The quantitative compression correction results in Figure 27 display the MHD and co-planar-centroid distances as error metrics in comparing the ultrasound tumor borders with the co-registered CT tumor borders. A Wilcoxon signed rank test was also computed for the null hypothesis that the median difference between the error metrics was zero. The alternative was that the median was not zero, with the implication that the correction method offered a statistically significant improvement to the data. The number of samples for each test was the number of corrected ultrasound images. In results not presented here, we also assessed the effect of discarding image slices that were within 1.5 mm (approximately the tracking system accuracy) of the previous slice in the image stream to enforce uniqueness of each observation in the statistical analysis, but no differences were observed in the resulting p values compared to using every image, and so we included all images for the Wilcoxon tests (recall $n = 178$ B-mode, and $n = 83$ strain).




Figure 26. Examples of B-mode (a-c) and strain image contours (d-f). (a) and (d) show the uncorrected images in 3D space with the segmented tumor contour shown in red and the tomogram tumor volume shown in green. (b) and (e) show the alignment after the rigid correction method. (c) and (f) show the alignment after the model-based correction method.

With respect to contour boundary metrics, after the rigid correction, the MHD for the B-mode images dropped to approximately 2.8 ± 0.9 mm, which was a significant improvement ($p <$

0.01). The rigid correction also significantly improved the alignment of the contours in the strain images using this metric with an error of 3.3 ± 0.6 mm ($p < 0.01$). In the case of the model correction, the MHD was reduced to approximately 1.9 ± 0.6 mm for B-mode and 2.0 ± 0.5 mm for strain images. The model-based correction was found to make a significant improvement to the data compared to both the uncorrected and rigid-corrected data ($p < 0.01$).

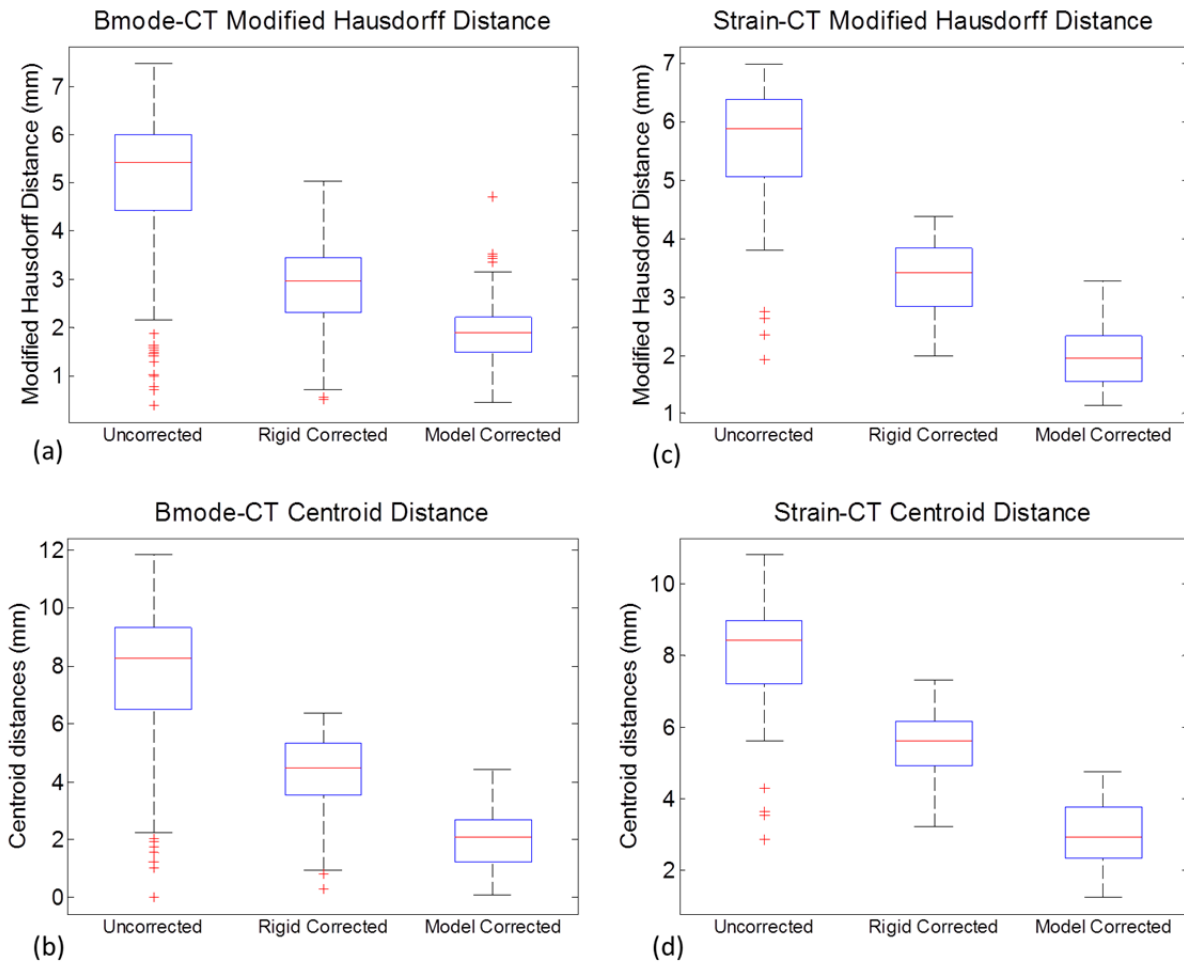


Figure 27. Alignment error results for the B-mode (a & b) and strain imaging (c & d) modalities for the organ-like phantom ($n = 178$ for B-mode, and $n = 83$ for strain). The position of tumor borders in each modality was evaluated in terms of Modified Hausdorff Distance to the co-aligned CT borders (a & c), as well as the distance between the centroid of the ultrasound tumor with the co-planar CT tumor border (b & d). The edges of the boxes are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered as outliers.

With respect to centroid metrics, the mean uncorrected centroid distance was approximately 7.6 ± 2.6 mm for B-mode and 8.0 ± 1.6 mm for strain images, which conformed to typical deformation values of up to 1 cm reported in soft anatomy by others [165, 166]. After application of the rigid correction, the mean centroid distance was reduced to 4.3 ± 1.3 mm in B-mode images, and was only reduced to 5.4 ± 0.9 mm in strain images but was still significant according to the Wilcoxon test (both improved with $p < 0.01$). After application of the model-based correction, the mean centroid error was reduced to 2.0 ± 0.9 mm for B-mode and 3.0 ± 0.9 mm for strain images and was a significant improvement over the uncorrected and rigid-corrected data (both $p < 0.01$).

The tumor volume enclosed by the rigid-corrected and model-corrected ultrasound B-mode contours was found to be approximately 3.2 cm^3 and 3.8 cm^3 for this phantom, respectively, and the volumes enclosed in the rigid-corrected and model-corrected strain images were 2.5 cm^3 and 2.9 cm^3 . Recall that in the baseline accuracy test using the cup phantom, the tumor volume (made from the same mold) was measured with tracked B-mode as 3.0 cm^3 , and the tumor mold was measured as 3.2 cm^3 .

The last phantom experiment was the comparison of model-correction with the assumption of a tumor-to-bulk stiffness ratio of 1:1 against the known 9:1 ratio. These results are compiled in Table 4 and show that there was no measureable difference in how much the tumor border was corrected when the actual 9:1 ratio was used.

Table 4. Results of phantom elasticity test to compare the effect on model-correction from the assumption of 1:1 tumor-to-bulk stiffness ratio against the known 9:1 stiffness ratio from material testing.

	1:1 Stiffness Ratio		9:1 Stiffness Ratio	
	MHD Error (mm)	Centroid Error (mm)	MHD Error (mm)	Centroid Error (mm)
B-mode (n = 178)	1.9 ± 0.6	2.0 ± 0.9	1.8 ± 0.8	2.2 ± 1.0
Strain (n = 83)	2.0 ± 0.5	3.0 ± 0.9	1.9 ± 0.5	3.1 ± 0.9

V.5.3 Clinical Case

The results of the clinical case are shown in Figure 28 and Figure 29. Figure 28 shows an example of a B-mode image before and after compression correction using the rigid method and model-based method. The segmented tumor border in each ultrasound image was outlined for comparison with the co-aligned MR tumor surface. Quantitative compression correction results are shown in Figure 29 in the same manner as the phantom studies, with the error metrics being the MHD and co-planar centroid distances. The MHD error before any correction was approximately 5.4 ± 0.1 mm, and the average centroid error was approximately 7.2 ± 0.2 mm. A Wilcoxon signed rank test was also computed to assess the significance of the improvement offered by each method. The rigid and model-based corrections both offered substantial improvement compared to the uncorrected alignment (both $p < 0.01$) across the two error metrics. The rigid correction resulted in a reduced MHD error of 2.3 ± 0.1 mm and centroid distance error of 4.6 ± 0.5 mm, whereas the model-based correction resulted in an MHD error of 2.6 ± 0.1 mm and centroid distance error of 3.5 ± 0.4 mm. Due to the difficulty in maneuvering the large ultrasound probe in the craniotomy of this patient, only a few angles of insonation were achieved and so a meaningful ultrasound tumor volume could not be constructed.




Figure 28. Results of the clinical case using tracked B-mode to localize a tumor border. The coronal and axial views are each shown for the uncorrected (a & b), rigid correction (c & d) and model correction (e & f) with the segmented tumor contour shown in red and the tomogram tumor volume shown in green. The segmented tumor surface in each B-mode slice is outlined for comparison with the MR tumor surface.

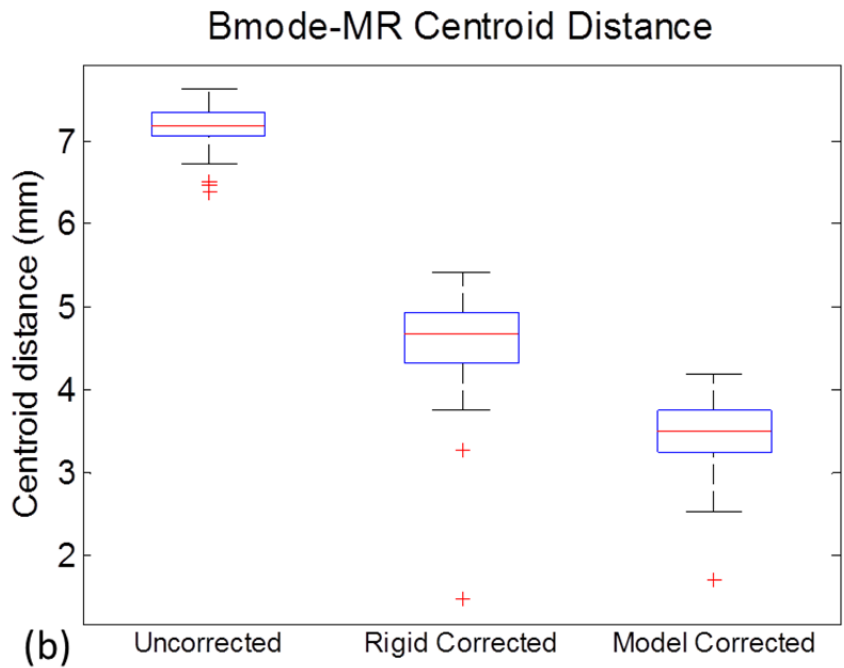
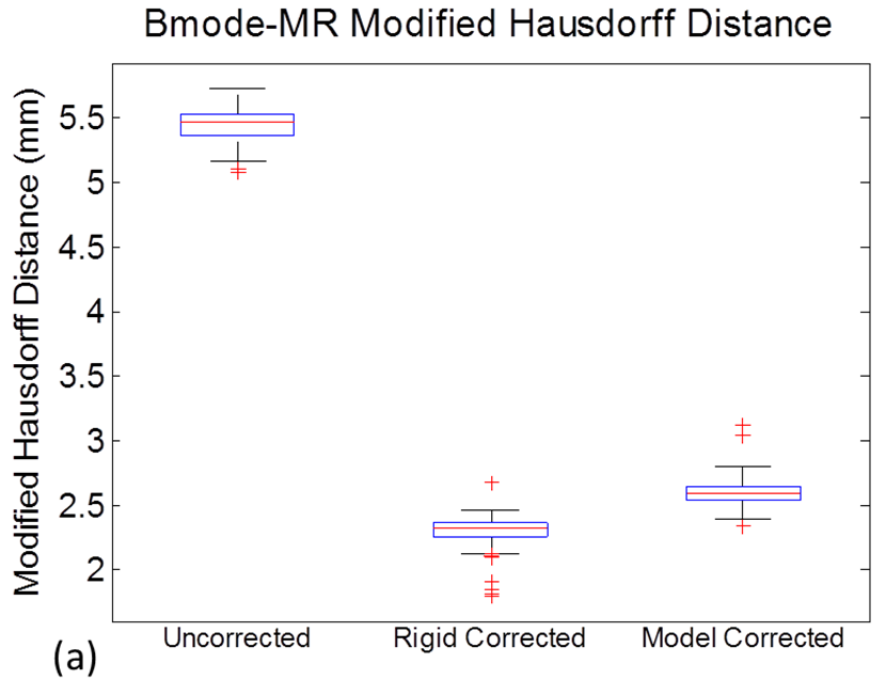


Figure 29. Alignment error results for the clinical case ($n = 118$ B-mode images). The position of tumor borders was evaluated in terms of Modified Hausdorff Distance to the co-aligned MR borders (a), as well as the distance between the centroid of the ultrasound tumor with the co-planar MR tumor border (b). The edges of the boxes are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered as outliers.

V.6 Discussion

V.6.1 Simulations

The results of the compression correction simulations in Table 3 show the effect of deviating from the assumption that the user would always apply compression purely in the depth direction of the ultrasound image. As expected, when the displacement vectors were actually parallel with depth (θ_1 and θ_2 of 0°), the model-based correction method essentially resulted in a corrected tumor which was equivalent to the original uncompressed tumor. The rigid correction method resulted in several millimeters of error in terms of both the tumor boundary and centroid location. However, when the displacement vectors used to create the simulated deformation were rotated to approach the probe surface from a different set of surface nodes (see Figure 25), the average boundary node error after both corrections grew. The performance variations due to different angles can be attributed to the assumption made in assigning boundary conditions to the FEM model, which is that the insertion trajectory of the probe into the tissue is perfectly parallel with the depth direction of the ultrasound images. The results in Table 3 indicate the sensitivity of the method to physical deviations from that assumption. Essentially, when the insertion trajectory is not perfectly perpendicular to the tissue surface, then the method creates inappropriate boundary conditions, which leads to model solutions which do not accurately reflect the actual tissue movement. These numbers represent a large error addition to an image-guidance platform, especially in procedures such as neurosurgery. It could be concluded that it is important that the user apply the probe primarily in the depth direction of the image plane in order for this compression correction method to be effective. This may not be a large limitation in practice, however, due to the typical presence of slip between the probe and tissue surface from ultrasound gel or other fluids. This would conceivably make it unlikely for the ultrasound probe to apply enough friction to the tissue surface necessary for the “dragging” effect which is described by the

scenarios simulated above, in which the probe trajectory into the tissue is off-axis from the depth direction.

V.6.2 Phantom Experiments

The results of the baseline accuracy test show that the tracked ultrasound system is capable of aligning the ultrasound and CT tumor borders with error between 1 and 2 mm when tissue compression is not a factor. This represents the best outcome that could potentially be expected from our correction. In addition, the tumor volume measured by tracked B-mode was quite similar to both the volume measured from the tumor mold and the CT volume.

The results of the compression correction methods shown in Figure 27 clearly demonstrate the improvement of the model-based method to the alignment between ultrasound and co-registered tomograms in the phantom experiments. The MHD error metric shows a decrease in misalignment after application of the rigid method, and shows a further decrease when the model-based method is used instead. The varying magnitude in improvement between B-mode and strain image when using the rigid method highlights the differences in contrast mechanism between the two modalities as well as the lack of tumor shape change, which both contribute to the delineation of mock tumor borders in this case. However, after application of the model-based compression correction, both modalities show similar MHD error values. This indicates that shape change of the inclusion is significant and needs to be taken into account when correcting guidance systems. One interesting effect of the correction was the increase in B-mode tumor volume after model-correction to 3.8 cm^3 , which is a slight overestimate compared to the tumor mold and CT volumes and is likely due to the stretching effect of the model displacements. The model-corrected strain volume was actually an underestimate of the volume, but this was primarily related to the sparser manner in which those images were collected in order to avoid out-of-plane movement. The discrepancy in rigid-corrected tumor volumes compared to the uncorrected volume was due to the

movement of the tumor contours in non-uniform directions determined by the freehand insonation angles for the slices.

The centroid distance error metric showed an even clearer improvement after applying the model correction compared to the uncorrected ultrasound images. A similar trend was again noted with respect to the rigid correction in that a modest improvement of both error metrics was observed, and model-based correction showed a large correction in both modalities. It can be seen that the rigid correction gave a slightly larger improvement to the B-mode images than the strain images. This again points to the significance of contrast and shape change factors in our study. Overall, the results demonstrated that while the rigid method occasionally led to a modest improvement, the model-based correction consistently led to significant reduction in alignment error.

V.6.3 Clinical Case

The results of the clinical case showed a clear improvement in alignment between the tracked B-mode tumor and MR tumor borders, as shown in Figure 28 and Figure 29. Interestingly, the rigid correction was slightly more effective than expected in improving the alignment given the results of the previous phantom experiments, which may be related to the location and structure of the tumor in this particular case. The model-based correction did offer a large improvement to alignment, especially along the region of the tumor closest to the surface as shown by Figure 28d and Figure 28f. One important note about this case was that the preoperative brain surface was not updated to account for general intraoperative brain shift which may have occurred prior to applying our ultrasound compression correction methods. Typically our group employs a model-based shift correction driven by sparse surface data from a LRS or other instrumentation in order to update the preoperative imaging to reflect intraoperative sag or swelling of the brain [96, 141, 181]. However, qualitatively there was little brain shift observed at the time of ultrasound acquisition for this patient before resection of the dural membrane, so no LRS data was obtained so

as to limit the amount of intraoperative data collection for this case. It is possible that there was some quantifiable amount of brain shift before the ultrasound acquisition, which would directly affect our correction method and could be reflected in the error metrics reported in Figure 29. In the future, the preoperative brain surface would ideally be corrected for brain shift prior to collection and correction of ultrasound data.

V.6.4 Limitations

It is worthwhile to note that although the model-based correction demonstrated a marked improvement over the uncorrected alignment in both the phantoms and clinical case, the gold standard used in this comparison (the co-registered tomogram borders) had some bulk-to-lesion segmentation variability. It was also likely that, due to our realistic guidance environment, several sources of error propagated through to the final analysis. For example, the optical tracking system used in this work had inherent error which was imparted to each measurement made with a tracked device and was a primary contributor to the overall error. Propagating uncertainty in optical tracking is an open problem. We have recently proposed a solution to this problem [182] and are working on an empirical evaluation [183] with encouraging results. The tracking error of our tracking system has been reported to have an RMS error of approximately 0.2 mm when tracking a passive rigid body, although typically the tip of a stylus-like tool which is subject to lever-arm effects can be localized at 1 to 2 mm [184, 185]. This influences fiducial localization error (FLE), which affects the accuracy of the image-to-physical registration between tomogram and tracked ultrasound data. While FLE can be difficult to directly quantify, the mean fiducial registration error (FRE) in our phantom experiments was calculated as 0.4 ± 0.3 mm for the phantoms. It should be noted this error only represents the accuracy with which the fiducial markers in the tomograms could be matched to the physical fiducials localized with a pen probe, and does not necessarily imply that the target registration error (TRE) throughout the phantoms was the same value. In addition, there was error associated with the tracked ultrasound calibration,

which was found to be approximately 0.4 mm using the method of Muratore and Galloway (2001). It should be noted that the errors discussed above are somewhat interrelated and not necessarily additive, but in our experience a reasonable estimate of the error involved in using a passively tracked ultrasound system ranges from 1.5 to 2.5 mm.

There were also several assumptions and simplifications made in the development of the rigid and model-based corrections. Regarding the rigid correction, it would be intuitively more accurate to stretch the A-lines to approximate the deformation rather than simply translate them. However, without a known reference in the far field, it is difficult to choose an appropriate bounding condition since there is no guarantee of having a known fixed object in view. In the case of the model-based correction, the first assumption was that the user would apply the probe purely in the depth direction for each image acquisition. This simplified the creation of boundary conditions for the model, but this was not always accurate. It was challenging to avoid lateral and out-of-plane probe movement during freehand imaging. This type of movement would invalidate the assignment of Dirichlet boundary conditions based upon the position of the digital probe surface, because the correspondence between the probe surface and the tomogram surface could potentially be lost in the absence of perfect slip at the boundary. The other assumption was the assignment of material properties to the FEM mesh. Overall, material properties are largely irrelevant in our approach because the model is driven by only Dirichlet conditions and thus only large stiffness ratios between tissue regions might have an effect. Absolute values for Young's modulus do not affect the displacement solution in the scenarios described in this work, and Poisson's ratio affects it only slightly. Possible future incorporation of force measurements into the model, however, would indicate a need for accurate material properties. In our phantom experiments, there was prior knowledge of the phantom material properties from mechanical testing data, and using the known properties was not shown to improve the correction. The phantoms were also highly homogenous compared to real tissue. In translating this method to the

clinic, there is the challenge of assigning appropriate property values throughout the mesh. It is well-recognized that determining patient-specific material properties *in vivo* is difficult, so one solution which has been employed is to use values from the literature, as well as those deduced through optimization experiments [180]. Lastly, one simplification which was made was the use of a simple linear elastic model to describe tissue movement. More sophisticated governing equations that better capture soft-tissue mechanics could be employed to estimate displacements, such as Biot's consolidation theory in the case of brain tissue [101, 186]. However, with added complexity comes added computational and integration burden that can affect adoption, yet despite these sources of model error, the improvement we found using this simple model was considerable. The model-based method offered consistent alignment error reductions of between 4 to 6 mm, i.e. the final alignment error was approximately 2 mm, which has clinical utility in the case of brain, liver, and breast surgery. These results suggest that our approach is a meaningful improvement to the utility of tracked ultrasound in image-guided surgery. In addition, anatomy with well-defined constitutive frameworks may give opportunity to improve this performance.

Other improvements can be made with respect to workflow. Currently, the correction cannot be run at real-time frame rates given that the model solution takes several seconds to calculate on an Intel i7 processor. However, updated ultrasound images can be provided in less than ten seconds, which is not an unreasonable burden on normal surgical workflow. The system could potentially be made real-time with the proper hardware and parallelization methods. Our method also made no use of the actual ultrasound images themselves to assist in realignment. One possible addition would be the utilization of either raw RF or B-mode pixel values to generate displacement fields by comparing frames in close proximity to one another and performing deformable registration as other groups have proposed [166, 170]. The inclusion of these smaller displacements could help address small inaccuracies due to the tracking system, while still making use of the biomechanical model to provide an estimate of bulk displacements.

V.7 Conclusions

In this work we proposed and validated two very practical methods for correcting alignment error due to tissue compression exerted by an ultrasound probe within the context of image-guided therapy. In one method, a simple rigid correction was applied. In a second more sophisticated method, patient-specific models were used to estimate physical tissue deformation as the direct result of pressing the tracked probe to the tissue surface. These model solutions were then used to transform the ultrasound images to an undeformed state for assessing pathology. The method was validated in simulations, phantoms, and a preliminary clinical case and showed that alignment of freehand tracked ultrasound with co-registered tomographic images was improved to within clinically useful margins. Experimental results indicate that integration of this correction method into conventional image-guided therapeutic platforms could assist the clinicians in decision-making by providing more accurate intraoperative data.

Acknowledgements: This work was supported in part by the National Institutes of Health award R01 NS049251 of the National Institute for Neurological Disorders and Stroke, and by the National Institute of Health award R01 CA162477 from the National Cancer Institute.

CHAPTER VI

TOWARD A GENERIC REAL-TIME COMPRESSION CORRECTION FRAMEWORK FOR TRACKED ULTRASOUND

VI.1 Introduction and Significance of Study

This study is concerned with the improvement of tracked ultrasound towards a real-time intraoperative guidance platform. The method presented previously in Chapter V used a patient-specific model to address the problem of registration error induced by tissue compression during ultrasound imaging. There were two areas of improvement that were identified with the patient-specific method which are addressed in this chapter. The first was the reliance of the patient-specific method on a full volumetric image volume from preoperative imaging in order to construct the patient model, which limits the method to procedures which routinely utilize such image volumes. The second area was the speed of the method, which was not real-time due to the need to generate new boundary conditions and solve the entire patient-specific model for each ultrasound slice. This chapter presents a novel strategy for performing an equivalent compression correction by modeling a generic block of tissue calibrated to the tip of the tracked ultrasound probe. The generic model may be precomputed and calibrated to any ultrasound system without the need for preoperative imaging of the patient, and only requires a sparse intraoperative measurement of compression depth in order to calculate the model solution. In addition, this new correction framework requires significantly less computational expense, making possible a nearly real-time compression correction. The generic method was evaluated with simulations, phantoms, and a clinical case, and was found to provide correction results which were similar to the patient-specific method. This work is in preparation for submission to the International Journal of Computer Assisted Radiology and Surgery.

VI.2 Abstract

Tissue compression during ultrasound imaging leads to error in the location and geometry of subsurface targets during soft tissue interventions. We present a novel compression correction method which models a generic block of tissue and its subsurface tissue displacements resulting from application of a probe to the tissue surface. The block model is calibrated to the tip of any tracked ultrasound probe. Intraoperatively digitization of the tissue surface is used to measure the depth of compression and provide boundary conditions to the biomechanical model of the tissue. The tissue displacement field solution of the model is inverted to non-rigidly transform the ultrasound images to an estimation of the tissue geometry prior to compression. This method was compared to a previously developed method using a patient-specific mesh rather than a generic block mesh. Experimental results with gel phantoms demonstrated that the proposed generic method reduced the tumor margin Modified Hausdorff Distance (MHD) from 5.0 ± 1.6 mm to 2.1 ± 0.7 mm, and reduced tumor centroid alignment error from 7.6 ± 2.6 mm to 2.6 ± 1.1 mm. The method was applied to a clinical case, and reduced the tumor margin MHD error from 5.4 ± 0.1 mm to 2.9 ± 0.1 mm, and the centroid alignment error from 7.2 ± 0.2 mm to 3.8 ± 0.4 mm.

VI.3 Introduction

Ultrasound is ubiquitous as an interventional imaging modality, and is commonly used to assess the location and geometry of disease intraoperatively. An inherent problem with this role is the shape distortion of visualized tissue structures introduced by pressured exerted on the tissue by the ultrasound probe itself during imaging. It is widely recognized that relatively large tissue compression can occur in soft tissue anatomy, e.g. the liver or breast. As a result, compression can obfuscate geometrical and locational measurements of subsurface targets such as tumors. This is particularly a problem for image-guided interventions which rely upon tracked ultrasound to provide intraoperative spatial measurements of structures taken during an intervention and then

compared to their co-registered preoperative imaging data counterparts. Nonrigid tissue compression is a primary cause of misalignment and shape distortion with these other sources of information. As image-guided navigation strategies in soft tissue environments continue to be developed, methods of correcting the tissue deformation from routine ultrasound imaging are necessary in order to ensure that all of these data are in a consistent spatial arrangement.

There are several methods described in the literature for performing compression correction. A common approach is to utilize the intensity information in the ultrasound images to perform a nonrigid intensity-based registration with positional tracking of compressed images over a range of compression states [166, 170]. One drawback of this method is that it requires a stream of ultrasound images, and intensity based registration for ultrasound is a challenging task in practice. For example, in Treece *et al.* they demonstrated a method to correct for compression using correlation of a stream of radiofrequency (RF) or amplitude frames, and although the method performed well in a phantom dataset, the authors noted its reliance on good image quality as well as the possibility of correction drift when compression estimates are accumulated across a large sequence of images. Another method of correction is to use a mechanical model of the tissue in order to estimate the subsurface tissue displacements caused by the interaction of the probe with the tissue surface. One group proposed using a force measurement apparatus to provide force boundary conditions to a tissue model [171, 187], although force boundary conditions require some prior estimate of absolute material properties for the tissue. We recently proposed an alternative method which utilizes a biomechanical model based correction which is driven by displacement boundary conditions provided by the position of a tracked ultrasound probe within a co-registered patient-specific organ surface from preoperative tomograms [188]. This method was shown to reduce ultrasound compressional error of nearly 1 cm to approximately 2 to 3 mm.

There is a subset of image-guided procedures for which preoperative tomographic image volumes are not commonly acquired, or the volumes are acquired with the patient in a much

different state than the operative state. This can be the case in liver surgery, for example, in which there is often significant manipulation of the organ by the surgeon leading up to the surgical presentation of the tissue. The previously developed method for ultrasound compression correction, however, relies on registration of intraoperative tracked ultrasound to a geometrically accurate patient model constructed from preoperative imaging. There is therefore a need for a method of compression compensation which does not rely on a preoperative model. In addition, there are also implications for utilizing subsurface information to perform image-to-physical registration. Provided with at least some form of intraoperative measurement of compression, be it from contact triggering or knowledge of the tissue surface with respect to the probe position, subsurface structures could be uncompressed to give true shapes in physical space. The true subsurface shapes could be used in combination with surface information to compute a combined image-to-physical registration, such as with a surface point cloud from a laser measurement device and subsurface structures like blood vessels [189]. With these possibilities in mind, the first goal of this work was to create a compression correction method which utilizes a generic correction method which is independent of tomographic imaging and requires no registration to a preoperative surface. The second goal was to compare the new method with the method previously described by Pheiffer et al. by deployment in phantoms and clinical data.

VI.4 Methods

The compression compensation method is one step within a pipeline for image-guidance using tracked ultrasound. A brief description of the pipeline is given below to provide context for the new method and to indicate how it differs from the previous correction method.

VI.4.1 Phantom Construction

A compliant gel phantom was constructed by mixing 7% by mass polyvinyl alcohol (PVA) in water with 10% by volume glycerol. A small amount was poured into a tumor mold and subjected

to four freeze-thaw cycles, in which the gel was frozen at -40 °C for 12 hours and then thawed at room temperature for 12 additional hours. The first freeze-thaw cycle produces a gel with a tissue-like consistency, and each additional cycle results in an increasingly stiffer material. The phantom tumor was then suspended by wire in a larger anthropomorphic liver mold with PVA mixture and subjected to one additional freeze-thaw cycle. This resulted in a soft tissue phantom containing a stiffer tumor. The completed phantom was fixed to a rigid base containing fiducials which were used to initialize the image-to-physical registration.

VI.4.2 Preoperative Imaging and Patient Model

Image-guided interventions often begin with acquisition of high resolution CT or MR image volumes prior to the procedure. The patient-specific compression correction method utilized a patient model created from these images. However, this data was unnecessary for the proposed generic method. Tomograms were acquired in this study only in order to compare the patient-specific and generic correction methods. CT image volumes were acquired for the phantoms using a clinical CT machine at 512 x 512 x 422 with 0.6 mm isotropic voxels. The bulk phantom and tumor were segmented using intensity thresholding in Analyze 9.0 (Mayo Clinic, Rochester, MN). Isosurfaces were generated from the segmentations using the marching cubes algorithm and smoothed with a Laplacian filter. A patient-specific finite element mesh with tetrahedral elements was created from the smoothed phantom isosurface using custom-built mesh generation software [113].

The clinical dataset consisted of a meningioma patient at Vanderbilt Medical Center. Informed written consent was obtained from the patient prior to the study with the approval of our Institutional Review Board. The preoperative MR volume was segmented to produce brain and tumor surfaces, which were used to create a patient specific model in the same manner as the phantom data.

VI.4.3 Intraoperative Data Collection

An Acuson Antares ultrasound machine (Siemens Inc., Munich, Germany) with a VFX13-5 linear array probe at 10 MHz was used to acquire all ultrasound images in this study. The machine was used to collect both B-mode and strain images with the eSie Touch elasticity software from the manufacturer in order to illustrate the general applicability of the correction method to all forms of ultrasound data. The ultrasound images were tracked in 3D space by synchronizing each image with the pose detected by a Polaris Spectra optical tracking system (Northern Digital, Waterloo, ON, Canada) for a passive rigid body attached to the ultrasound probe. The tracked ultrasound probe was calibrated using the method described by Muratore et al. such that all pixels in each image were associated with a 3D pose.

In addition to the ultrasound data collected above, the other intraoperative tools used in this study were a tracked pointer and laser range scanner (LRS). The pointer was used to digitize point fiducials such as beads on the phantom base and craniofacial landmarks on the patient. These points were used to initialize a surface-based registration of the dense point clouds from the LRS to the preoperative patient-specific model in the case of the phantoms. There was no LRS cloud available in the case of this specific patient, and so a random sampling of the MR model surface was used to simulate LRS data in that case.

VI.4.4 Compression Correction

The goal of this work was to reduce the registration error arising from soft tissue deformation exerted by an ultrasound probe. We recently presented a method in [188] utilizing both probe tracking information in combination with a co-registered patient model in order to estimate the compression depth of the probe into the tissue during insonation, and then to use that depth to correct the tracked ultrasound image poses using a biomechanical model-based approach. The novel method which we now propose does not use a patient-specific model derived from preoperative imaging, but instead uses a generic model to drive the correction as shown in Figure

30. In order to thoroughly compare the two methods, a summary of the patient-specific method is described in the next subsection, followed by a description of the novel generic method.

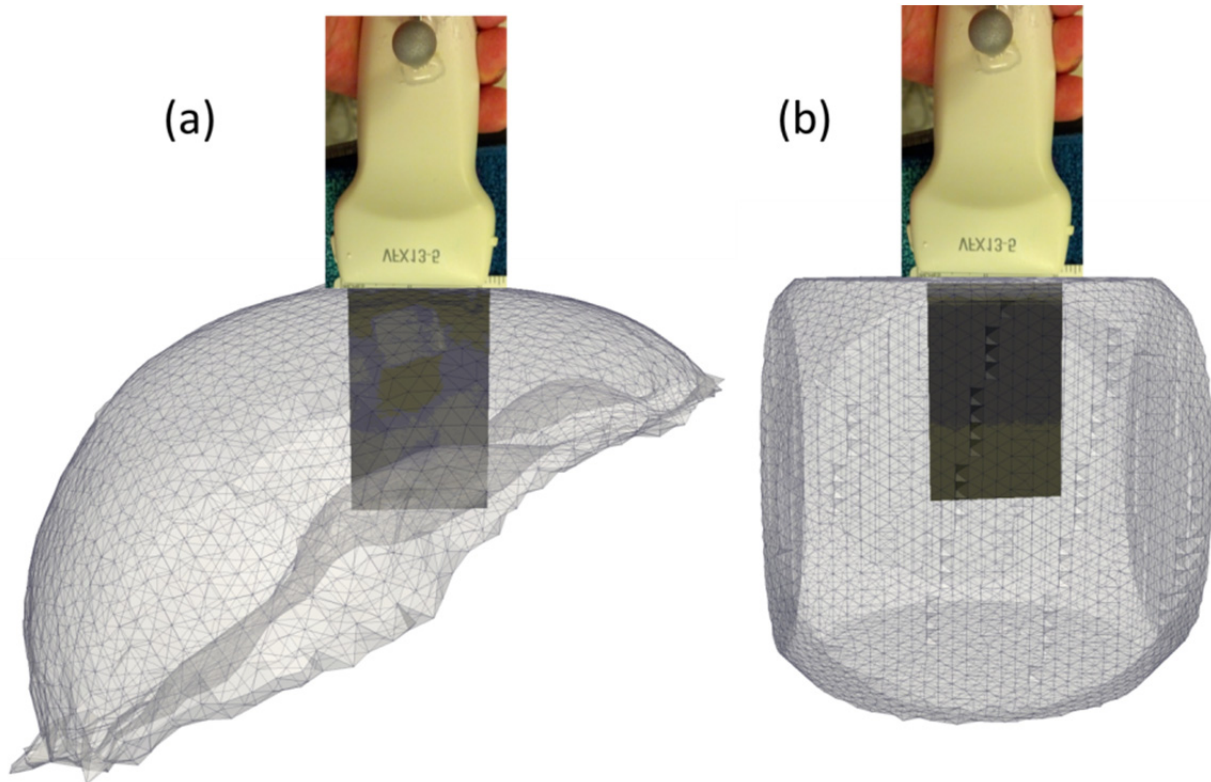


Figure 30. Example of a mesh used for the patient-specific model correction method (a), constructed from preoperative imaging and aligned to the ultrasound data using intraoperative registration methods, and an example of a block mesh for the generic model correction method (b), which is pre-aligned to the ultrasound data by performing a calibration to the ultrasound probe.

VI.4.4.1 Patient-Specific Correction

The patient-specific compression correction method presented in [188] made several key assumptions in order to compute a reasonable correction. The primary assumption was that intraoperative ultrasound data could be aligned to the preoperative imaging with an initial rigid registration utilizing surface features. The second assumption was that the tissue presentation in terms of mechanical state in the preoperative imaging was similar to the intraoperative presentation, in the absence of tissue manipulation by the ultrasound probe. These two

assumptions led to the next assumption, which was that in the presence of tissue compression by the ultrasound probe, the tip of the probe would be some distance below the surface of the co-registered patient model depending on the magnitude of compression.

The strategy which was proposed was to use the position of the probe geometry within the tissue model in order to estimate the tissue deformation resulting from application of the probe to the tissue surface for every ultrasound slice to be corrected. The geometry of the probe tip was constructed by scanning it with the LRS, giving a digital probe surface which was tracked in physical space using the tracked target attached to the probe. The next assumption was that during imaging the user would press the probe only in the depth direction of the ultrasound plane into the tissue, with no lateral or out-of-plane movement through the tissue which could cause a dragging effect. This assumption was made to simplify the next step of the correction, which was to assign boundary conditions to the biomechanical tissue model using the pose of the probe geometry. Assuming purely depth compression, the model surface nodes directly above the digital probe surface were assigned Dirichlet boundary conditions corresponding to the compression vectors calculated from the patient surface to the probe surface. The rest of the patient model was assigned a set of initial boundary conditions based upon the surgical plan and prior knowledge of the patient presentation. In the case of the liver phantom, the inferior surface was set to zero displacement and the superior surface was set to zero stress, because the bottom of the phantom was rigidly fixed to a rigid base whereas the top of the phantom was left open to the atmosphere. In the case of the clinical brain case, the mesh nodes corresponding to the craniotomy region were set as stress free, the base of the brain was set as fixed, and the rest of the brain nodes were set to have zero displacement in the normal direction but stress free in the tangential directions in order to allow for slip along the skull.

After the assignment of boundary conditions, the model was solved for 3D displacements over the entire mesh to estimate the probe-deformed state of the tissue. The model used in the

previous method and in the method proposed in this work consists of the standard 3D Navier-Cauchy equations for the displacement field:

$$\frac{E}{2(1+\nu)}\nabla^2\mathbf{u}+\frac{E}{2(1+\nu)(1-2\nu)}\nabla(\nabla\cdot\mathbf{u})+F=0 \quad (16)$$

where E is Young's modulus, ν is Poisson's ratio, \mathbf{u} is the 3D displacement vector at a point in the tissue, and F is the applied body force distribution. The partial differential equation is solved within a finite element method framework using the Galerkin weighted residual technique with linear basis functions. The system of equations that solves for the displacement vectors at every node in the mesh can be written as:

$$[K]\{\mathbf{u}\}=\{\mathbf{f}\} \quad (17)$$

where K is the global stiffness matrix, \mathbf{u} is the vector of nodal displacements, and \mathbf{f} contains the contributions of any applied body forces or surface movement at each node. For each ultrasound image to be corrected, this system of equations is constructed and solved for the nodal displacements which satisfy static equilibrium for the supplied boundary conditions. These displacements are then reversed and interpolated onto the tracked ultrasound data, which was then deformed with this 3D displacement field to an estimate of its state in the absence of compression.

VI.4.4.2 Proposed Generic Correction

The first difference between the generic correction and the patient-specific correction is that instead of a patient-specific mesh constructed from preoperative imaging and registered to intraoperative space, the generic method instead uses a preconstructed block mesh which is calibrated to follow the tip of the tracked ultrasound probe (see Figure 30). The most important consequence of this framework is that the generic method only requires a sparse intraoperative measurement of tissue compression in order to provide a model correction, rather than a

registration to preoperative imaging. In addition, a precomputed mesh offers certain computational advantages which will be described later.

The block mesh calibration procedure simply requires the alignment of the top of the ultrasound image with the center of one side of the mesh, and of the image plane itself with the plane through the center of the block. The pose of the generic block mesh thus is defined by the same tracking information which defines the pose of the ultrasound image. The general strategy is to acquire intraoperative measurements of the undeformed tissue surface using an LRS or tracked pointer, and use that surface to estimate the depth to which the tissue was compressed by the ultrasound probe. This depth is then used to assign Dirichlet boundary conditions to the block mesh in a similar manner as the patient-specific correction, although with a slight difference. The initial boundary conditions for the mesh in this method were assigned such that the inferior face of the block was fixed, and the superior and side surfaces were stress-free. This was somewhat arbitrary, as one downside to this approach is a lack of knowledge about what part of the patient anatomy the block mesh would essentially represent, depending on the angle of insonation.

After assignment of boundary conditions, the model is solved for 3D displacements in the block of tissue, and the displacements are reversed and interpolated onto the ultrasound data to perform the correction. The model construction is governed by the same constitutive equations given by (16) and (17). However, there are several advantages that the generic correction offers compared to the patient-specific model. The global stiffness matrix, K , needs to be reconstructed whenever the type of boundary condition for any boundary node is changed, such as from a displacement to a force condition or vice versa. In the patient-specific model correction, boundary nodes are often reassigned different types of boundary conditions depending on the position of the probe within the mesh (as described in the patient-specific correction section). This implies that the K matrix and boundary conditions in f often need to be recomputed and solved for u , which is a relatively expensive process. With respect to the generic correction, however, the type of boundary

condition assigned to each boundary node will always remain the same, as each correction proceeds by merely altering the magnitude of the displacement boundary conditions on the top of the block mesh. Thus, it is possible to pre-compute K and reuse it for each correction whenever f is updated in a simple matrix multiplication:

$$\{\mathbf{u}\} = [K]^{-1}\{\mathbf{f}\} \quad (18)$$

Another property of the generic method offers a further computational speedup. In order to correct the ultrasound data, only the model solution at a plane of the mesh which corresponds to the ultrasound image plane is actually needed. The computations solving for the rest of the mesh nodes are therefore superfluous, which makes it desirable to remove those nodes from the system of equations. This can be accomplished through the method of condensation in order to result in a smaller system of equations which can be solved rapidly [190]. The first step in this process is to carefully arrange the ordering of the mesh node indices to ensure that the first N equations belong to the nodes lying on the ultrasound plane, as well as any nodes on the top surface which are assigned varying amounts of compression boundary conditions. Assuming this ordering, the equation from (17) can be rewritten as a block matrix system where the subscripts p and a indicate the plane nodes and all other nodes, respectively:

$$\begin{bmatrix} K_{pp} & K_{pa} \\ K_{ap} & K_{aa} \end{bmatrix} \begin{Bmatrix} \mathbf{u}_p \\ \mathbf{u}_a \end{Bmatrix} = \begin{Bmatrix} \mathbf{f}_p \\ \mathbf{f}_a \end{Bmatrix} \quad (19)$$

The block matrix system in (19) can be rearranged to a form involving only the displacement solution of the plane nodes, \mathbf{u}_p :

$$[\check{K}_{pp}]\{\mathbf{u}_p\} = \{\check{\mathbf{f}}_p\} \quad (20)$$

where

$$[\check{K}_{pp}] = [K_{pp}] - [K_{pa}][K_{aa}]^{-1}[K_{ap}] \quad (21)$$

$$\{\check{\mathbf{f}}_p\} = \{\mathbf{f}_p\} - [K_{pa}][K_{aa}]^{-1}\{\mathbf{f}_a\} \quad (22)$$

The modified stiffness matrix given in (21) represents a transfer of the displacements from all non-plane nodes to the plane nodes which are the primary concern, and maintains the volumetric nature of the model. Using this stiffness matrix offers significant computational benefit because it is a fraction of the size of the full K matrix. It can be similarly pre-computed and stored for very fast solutions of the \mathbf{u}_p vector of plane node displacements. In addition, given the careful ordering of the node indices explained above and the assignment of initial boundary conditions, it will also be the case that all nodes in the \mathbf{f}_a vector will always be assigned either zero stress or zero displacement boundary conditions. Given (22), this implies that changes in the compression depth during imaging will result in simple reassignment of the values in $\check{\mathbf{f}}_p$:

$$\{\check{\mathbf{f}}_p\}|_{\{\mathbf{f}_a\}=\mathbf{0}} = \{\mathbf{f}_p\} \quad (23)$$

Given the pre-computation of the modified stiffness matrix in (21) and the speed of assigning new values in (23), the generic method offers a very large speed increase compared to the patient-specific method and can potentially be performed at near real-time frame rates. Both correction methods were implemented in MATLAB on an Intel Core 2 Quad CPU at 2.4 GHz with 4 GB of RAM.

To reiterate, the primary hypothesis of this work was that a correction using a generic block mesh calibrated to the ultrasound probe could significantly reduce geometric compression error, and could possibly perform as well as the patient-specific method without the need for a preoperative mesh. In the next section, two simulation studies were performed to evaluate the behavior of the generic correction under various conditions. Both correction methods were then performed for comparison in terms of alignment error reduction during phantom experiments and a clinical case described in the following section, with a brief analysis of the computational expense required by each method.

VI.4.5 Experimental Validation

VI.4.5.1 Simulations

Several simulations were performed to examine the sensitivity of the generic correction method to various factors. The first simulation performed was to analyze the effect of the finite element mesh resolution on the model correction. This simulation consisted of constructing equivalent 10 x 10 x 5 cm block meshes with a tetrahedral element edge length ranging from 2 to 10 mm. The number of nodes in the meshes ranged from 53,018 to 634 between the 2 and 10 mm resolution, respectively. Three simulated tumors were created with diameters of 10, 20, and 40 mm and placed separately in an instance of each mesh. A 10 mm surface compression was then simulated for each mesh and the model solution was interpolated to the tumor nodes for comparison of the effects of the mesh resolution on the correction strategy. The comparison was performed by utilizing the most finely resolved mesh (2 mm edge length) as the ground truth solution, with each subsequent model solution from the coarser meshes being compared to the ground truth solution in terms of the difference in final tumor position.




Figure 31. Example of the meshes used in the simulation studies with sizes of 10 x 10 x 5 cm (a), 10 x 10 x 10 cm (b) and 10 x 10 x 15 cm (c). In this case a simulated 20 mm diameter tumor is shown at the equivalent 3 cm depth in each mesh.

The second simulation performed was aimed at determining the potential effects of the following variables upon the correction: 1) the block mesh size, 2) the tumor stiffness, and 3) the tumor size. Three block meshes were created with dimensions of 10 x 10 x 5 cm, 10 x 10 x 10 cm, and 10 x 10 x 15 cm constructed with 5 mm edge length. Three simulated tumors were created with diameters of 10, 20, and 40 mm and each placed at a 3 cm depth (half the maximum depth of the linear array probe used for the phantom and clinical data in this work) in an instance of each block mesh described above. To illustrate, the 20 mm tumor is shown in the meshes of different sizes in Figure 31. The tumors were assigned stiffness values of 1:1, 10:1, and 30:1 compared to the rest of the tissue block, resulting in 27 meshes (three mesh sizes, three tumor sizes, and three stiffness ratios). Each mesh was then subjected to surface compression ranging from 0 to 10 mm. For each state of compression, the model-deformed tumor surfaces were compared to the uncompressed tumor surfaces to illustrate the effect of tumor stiffness and size on the model solution, which would in turn affect the correction.

VI.4.5.2 Phantom Experiments

A liver phantom was constructed as described previously. The phantom fiducial markers were localized with a tracked pointer and an initial rigid point-based registration to the CT images was performed [128]. An LRS scan of the liver surface was acquired and an iterative closest point (ICP) registration was performed of the tracked LRS point cloud to the CT surface in order to refine the registration [92]. This alignment was used to perform the patient-specific correction method and served as the gold standard validation for the proposed generic correction method. Tracked B-mode and strain images were acquired of the embedded tumor, and the transformation matrix from the ICP registration was used to automatically align all tracked ultrasound images with the CT data. A total of 178 B-mode and 83 strain images were collected of the tumor. The tracking and registration transformations were then applied to the digital probe surface in order to generate boundary conditions for the two correction methods as described previously.

After the generation of boundary conditions, the patient-specific and generic correction methods were applied to each ultrasound image. This resulted in a collection of uncorrected, patient-specific corrected, and generic corrected images. With respect to the generic correction, a 10 x 10 x 10 cm block mesh was used. For both methods, the meshes were assigned a tumor Young's modulus ratio of 1:1 with Poisson's ratio at 0.49 because PVA is known to be nearly incompressible. Each population of images was compared to the baseline CT images in terms of tumor geometry in order to evaluate the corrections. The tumor borders in each B-mode and strain image were segmented semi-automatically using the Livewire technique, and for each ultrasound image the CT volume was re-sliced to provide a co-planar CT slice and tumor contour corresponding to the 3D pose of the co-registered ultrasound slice. The tumor borders from ultrasound and CT were then compared in terms of Modified Hausdorff Distance (MHD) and centroid distance between the two contours [179]. The MHD and centroid error metrics were computed prior to and after each correction, and were the primary metrics in determining the efficacy of the methods.

VI.4.5.3 Clinical Case

The patient-specific and the generic correction methods were deployed in a preliminary clinical case in addition to the phantom study. With respect to the generic correction, a 10 x 10 x 10 cm block of tissue was used, and a preoperative MR image volume was used to construct the patient-specific brain model. The tumor in this case was a meningioma located superficially on the left side. In this case the tumor and brain were assigned a 1:1 stiffness ratio and Poisson's ratio of 0.45 [180]. Alignment of the intraoperative tracked ultrasound images to the MR was performed by scanning the face of the patient with the LRS and performing an ICP registration between the LRS face point cloud and the MR patient model. Tracked B-mode images were obtained immediately after the craniotomy. Both correction methods were then applied to the ultrasound data and

compared to the co-aligned MR tumor borders in terms of the co-planar contour MHD and centroid error.

VI.4.5.4 Computational Efficiency

In order to provide an idea of the gain in speed offered by the generic correction framework, the time the time to compute the generic model solution for an image slice in the phantom B-mode dataset was recorded in the case of a 10 x 10 x 5 cm block mesh with 5 mm edge length consisting of 4,042 nodes and 19,672 tetrahedral elements. Furthermore, only the subset of the mesh which corresponds to the ultrasound plane was used, which reduced the number of nodes and elements used in the generic correction to 697 and 2,698, respectively. The analogous correction using the patient-specific correction was computed using a mesh with 5 mm edge length consisting of 10,989 nodes and 55,165 tetrahedral elements. The difference in the numbers of nodes and elements in the meshes having similar edge length corresponds to the difference between the volume of the full patient-specific organ versus the volume of the block of tissue in the generic method.

VI.5 Results

VI.5.1 Simulations

The results of the mesh resolution simulations are shown below in Figure 32. This figure displays how the model solutions at varying mesh resolutions changes compared to the solution to the high resolution mesh using a 2 mm element edge length. The general trend in each case was that as the mesh becomes coarser, the interpolated model solution deviates from the solution obtained from the more finely resolved mesh, especially above an edge length of 7 mm.

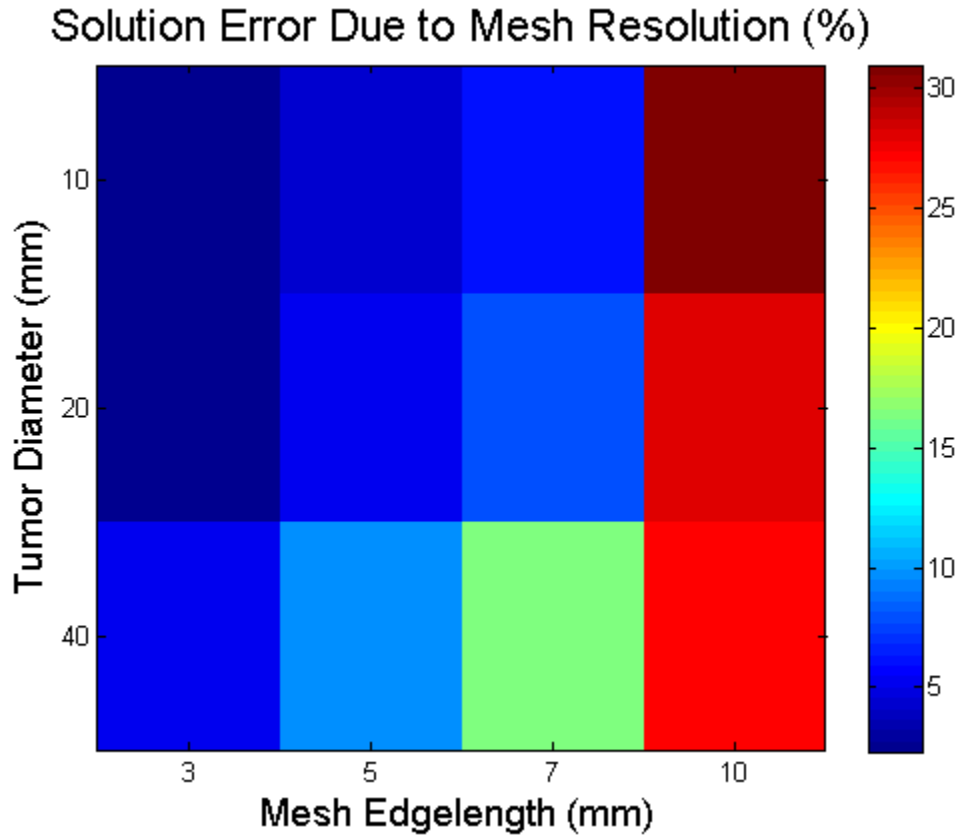


Figure 32. Effects of mesh resolution on three sizes of a tumor after a simulated 10 mm surface compression. The tumor node error is defined relative to the result of the solution of a mesh with 2 mm edge length resolution.

The results of the second simulation testing the effects of mesh size, tumor stiffness, and tumor size is shown below in Figure 33. Each graph shows that as the applied surface compression increases, the tumor boundary nodes displace correspondingly. The magnitude of tumor displacement varies slightly, however, by the size of the block mesh and even more so by the size of the tumor and its stiffness compared to the rest of the tissue block.

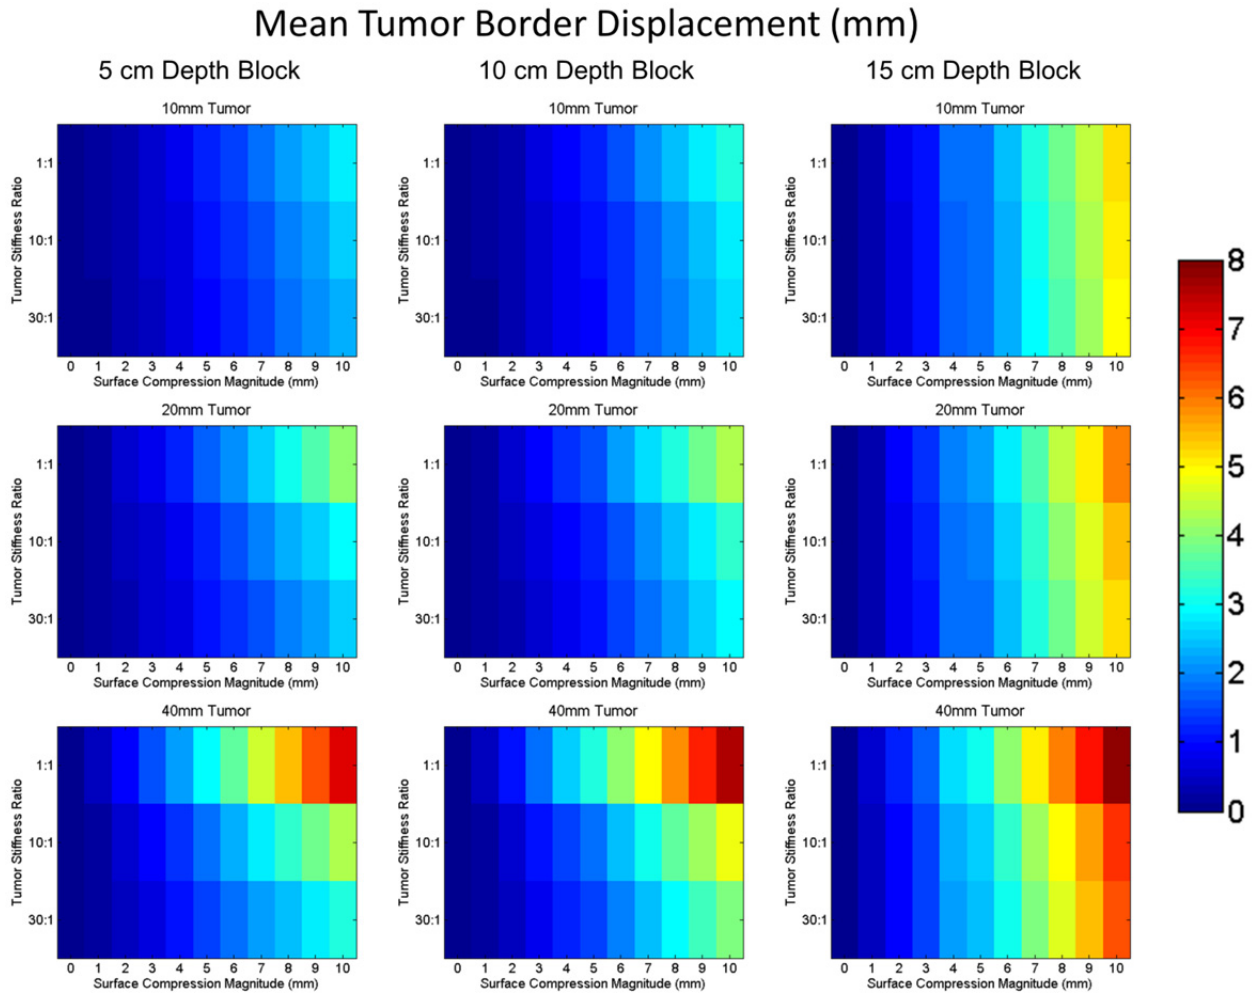


Figure 33. Effects of block mesh depth, tumor stiffness, and tumor size upon the model-predicted tumor border deformation under varying amounts of surface compression.

VI.5.2 Phantom Experiments

The patient-specific and generic model corrections were deployed in the liver phantom, and an example of the correction process applied to a tracked ultrasound slice is shown in Figure 34. Qualitatively, there was a clear improvement to the alignment between ultrasound and co-registered tomograms in the phantom experiments. In addition, the ultrasound contours corrected with the generic model method were very geometrically similar to the ultrasound contours corrected with the patient-specific method.




Figure 34. Example of the patient specific-model correction (a-d) and generic model correction (e-h) applied to a B-mode image slice. (a) and (e) show the co-registered LRS point cloud on the respective mesh, (b) and (f) show the tracked probe surface and the misalignment between the ultrasound tumor border with the CT tumor, (c) and (g) show the deformed mesh after creation of boundary conditions, and (d) and (h) show the undeformed mesh and corrected ultrasound image.

The quantitative results of the phantom experiments in Figure 35 show the MHD and coplanar-centroid distances as error metrics in comparing the ultrasound tumor borders with the co-registered CT borders, for both B-mode and strain images. The B-mode MHD values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 5.0 ± 1.6 mm, 1.9 ± 0.6 mm, and 2.1 ± 0.7 mm, respectively. A Wilcoxon signed rank test was computed for the null hypothesis that the median difference between the error metrics was zero. It was found that there was a statistically significant difference between each of the image populations using this metric ($p < 0.01$). The B-mode centroid error values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 7.6 ± 2.6 mm, 2.0 ± 0.9 mm, and 2.6 ± 1.1 mm, respectively. The Wilcoxon test again found the three image populations to be significantly different from one another based on this metric ($p < 0.01$).

With respect to the strain images, the strain MHD values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 5.6 ± 1.1 mm, 2.0 ± 0.5 mm, and 2.2 ± 0.5 mm, respectively. The Wilcoxon test found all three image populations to be statistically different using the MHD metric ($p < 0.01$). The strain centroid error values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 8.0 ± 1.6 mm, 3.0 ± 0.9 mm, and 3.3 ± 1.1 mm, respectively. The Wilcoxon test also found all three image populations to be statistically different using the centroid error metric ($p < 0.01$).

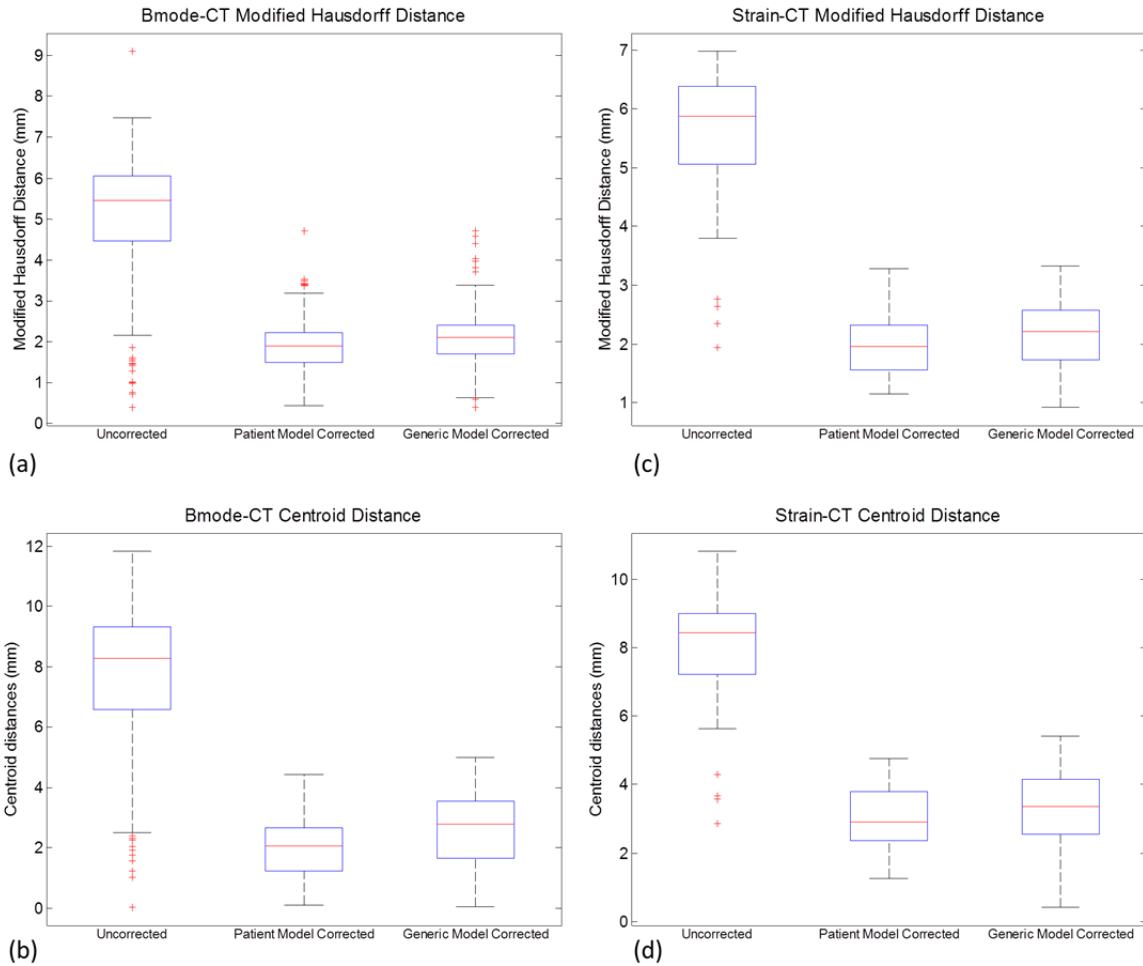


Figure 35. Alignment error results for the B-mode (a & b) and strain imaging (c & d) modalities for the organ-like phantom ($n = 178$ for B-mode, and $n = 83$ for strain). The position of tumor borders in each modality was evaluated in terms of Modified Hausdorff Distance to the co-aligned CT borders (a & c), as well as the distance between the centroid of the ultrasound tumor with the co-planar CT tumor border (b & d). The edges of the boxes are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered as outliers.

VI.5.3 Clinical Case

The quantitative results of the clinical case in Figure 36 show the MHD and co-planar-centroid distances as error metrics in comparing the ultrasound tumor borders with the co-registered CT borders, for both B-mode and strain images. The B-mode MHD values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 5.4 ± 0.1 mm, 2.6 ± 0.1 mm, and 2.9 ± 0.1 mm, respectively. A Wilcoxon signed rank test was computed for the

null hypothesis that the median difference between the error metrics was zero. It was found that there was a statistically significant difference between each of the image populations using this metric ($p < 0.01$). The B-mode centroid error values for the uncorrected, patient-specific corrected, and generic corrected tumor borders were 7.2 ± 0.2 mm, 3.5 ± 0.4 mm, and 3.8 ± 0.4 mm, respectively. The Wilcoxon test again found the three image populations to be significantly different from one another based on this metric ($p < 0.01$).

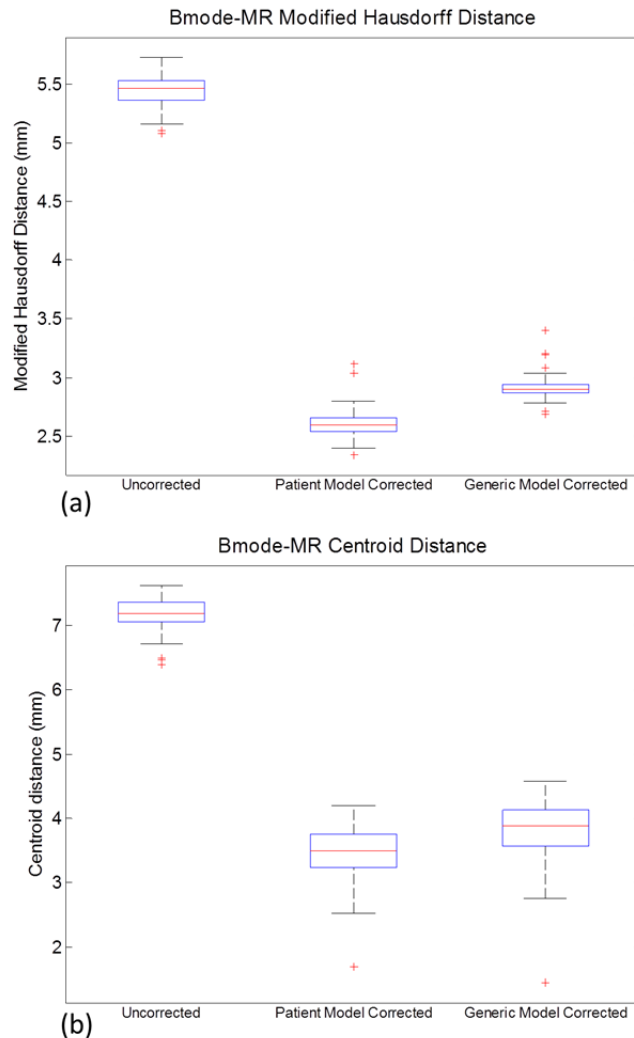


Figure 36. Alignment error results for the clinical case ($n = 118$ B-mode images). The position of tumor borders was evaluated in terms of Modified Hausdorff Distance to the co-aligned MR borders (a), as well as the distance between the centroid of the ultrasound tumor with the co-planar MR tumor border (b). The edges of the boxes are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered as outliers.

VI.5.4 Computational Efficiency

A breakdown of the various computational costs in terms of execution time for the patient-specific correction and generic correction is given in Table 5.

Table 5. Approximate execution time for each step in the patient-specific and the generic correction pipelines. Times were determined using a single thread of an Intel Core2 Quad CPU at 2.4 GHz.

		Patient-Specific Correction	Generic Correction
Preoperative Phase	Image Segmentation	30 min	-
	Mesh Construction	5 min	2 min
	Mesh Calibration	N/A	2 min
	Stiffness Matrix Pre-construction	N/A	1 min
	TOTAL	35 min	5 min
Intraoperative Phase	Boundary Condition Determination	50 ms	50 ms
	Stiffness Matrix Construction	40 s	N/A
	Model Solve	12 s	10 ms
	Ultrasound Tumor Interpolation	0.4 s	20 ms
	TOTAL	52.5 s	80 ms

In the case of the patient-specific method, the mesh is created from preoperative imaging, which typically requires at least 30 min assuming that some manual oversight of the image segmentation is required. Creation of the mesh from the segmentation mask takes at least 5

minutes. In terms of actual intraoperative expense, the construction of the 32967×32967 stiffness matrix K and the solution of equation (17) were conducted together in approximately 50 s, and this represented the vast majority of the total intraoperative computation time of 52.5 s.

In the case of the generic method, the modified stiffness matrix given by (21) was pre-computed for the block mesh prior to collection of ultrasound data. There were only 697 nodes in the ultrasound plane region of the mesh, and so the condensed stiffness matrix in this case contained 2091×2091 entries. For each model correction, the vector given by (22) was modified with the detected compression vectors and (20) was solved for the plane node displacements. The inverse of the modified stiffness matrix in (21) was stored and the solution time for the model was approximately 10 ms. The overall intraoperative computation time was approximately 80 ms.

VI.6 Discussion

VI.6.1 Simulations

The mesh resolution sensitivity study showed in Figure 32 that there is less than 10% difference between mesh solutions when the element edge length is at least below approximately 7 mm. Given that there is motivation to use a mesh with as few nodes as possible due to the time needed to solve the linear system of equations, the primary conclusion of this simulation is that an element edge length of approximately 5 mm is an appropriate tradeoff between speed and accuracy.

The second simulation study showed how the block mesh size, tumor size, and tumor stiffness affect the model solution. The first observation is that the size of the block mesh did not affect the solution at the tumor nodes until the size was much larger than the depth at which the tumor was located (recall the tumor was placed at 3 cm to be in the center of a 6 cm US image). The solutions at the tumor nodes were similar when utilizing block sizes of $10 \times 10 \times 5$ cm and $10 \times 10 \times 10$ cm, but at the $10 \times 10 \times 15$ cm size the tumor solutions tended to become less variable across

different tumor sizes and stiffness. One trend to note from Figure 33 is that there is very little impact on the model solution from the size of a tumor or its stiffness ratio at low levels of surface compression. However, as the surface compression becomes quite large, there is a divergence in the solutions on the basis of both tumor size (illustrated by the displacement magnitudes in each graph) and the stiffness ratio (the vertical axis on each graph). Larger tumor size resulted in greater overall tumor boundary deformation, which was expected because a larger tumor diameter implies that a greater proportion of tumor nodes were closer to the surface deformation, since all three tumors were placed at the same tissue depth. It is especially worth observing that the importance of tumor stiffness increased with increasing tumor size. In the case of the 10 mm diameter tumor at the maximum surface displacement of 10 mm, the difference in mean tumor displacements when using the 1:1 and 30:1 stiffness ratios varied by approximately 1 mm. However, in the case of the 40 mm diameter tumor, the difference in mean tumor displacements when using the 1:1 and 30:1 stiffness ratios varied by approximately 3 mm for the same surface displacement magnitude. These simulations indicate that although in many cases the tumor geometry and material properties do not greatly impact the model solution, these variables can become important when the target is a large tumor with a much different stiffness from the surrounding normal tissue.

VI.6.2 Phantom Experiments

The results of the compression correction methods shown in Figure 35 clearly demonstrate the improvement offered by both the patient-specific and generic methods to the alignment between ultrasound and co-registered tomograms in the phantom experiments. The MHD error metric showed a significant decrease in misalignment after application of both methods, but with a greater decrease for the patient-specific method. The centroid distance error metric showed an even clearer improvement after applying the two model-based corrections compared to the uncorrected ultrasound images. A similar trend was noted with respect to the corrections, in that

the patient-specific method resulted in a greater decrease in error overall compared to the generic model correction. However, although the correction was significantly better for the patient-specific method according to the Wilcoxon test, the mean difference between the resultant errors for the corrections was submillimetric for both the B-mode and strain images. This result is important because it indicates that the generic model correction performs nearly as well as the patient-specific method, making it a reasonable alternative for cases in which a patient-specific model may not be available.

VI.6.3 Clinical Case

The results of the clinical case in Figure 36 showed a clear improvement in alignment between the tracked B-mode and MR tumor borders after both of the correction methods. The same trend from the phantom dataset was noted in this case, which was that the patient-specific model correction yielded a slightly greater reduction in error than the generic model correction. However, as with the phantom data, the difference in the mean error for both metrics was submillimetric in comparing the two corrections. This reinforces the idea that the generic model correction could be used to perform a comparable compression correction in the absence of a patient-specific model from preoperative imaging.

VI.6.4 Computational Efficiency

It was found that the patient-specific method on average needed approximately 50 s to provide a compression correction update to each individual ultrasound frame during freehand movement of the probe. This long computation time was primarily a consequence of the need to reassign boundary condition types to surface nodes in the patient-specific mesh as the probe was moved around the tissue, thus necessitating a full reconstruction of the stiffness matrix K for each correction. The stiffness matrix for the patient-specific mesh was much larger than the mesh in the generic correction due to the greater number of nodes needed to represent the full patient organ, thus leading to a longer solution time as well. This correction time removes one of the primary

advantages of ultrasound as an interventional imaging modality, which is its real-time data acquisition. There is clearly motivation to provide both a corrected image while maintaining a high frame rate.

By contrast, the generic correction method was shown to provide a model solution in approximately 10 ms using the condensation method to only solve for the mesh nodes in the immediate vicinity of the ultrasound plane. This essentially represents the removal of the primary computational bottleneck from the patient-specific correction, which was the 50 s required for construction and solution of the stiffness matrix for a large organ-shaped mesh. The rest of the intraoperative steps consisted of determining boundary conditions prior to the model solution and then interpolating the model solution to the ultrasound data. The other steps combine with the model solution to give a total intraoperative correction time of approximately 80 ms, which is nearly real-time at 12.5 frames per second. In addition, this work was implemented on only a single CPU, and could possibly be further improved by the use of GPU programming. This work demonstrated that although there is a slight reduction of the accuracy of the solution provided by the generic correction versus the patient-specific method, the considerable computational benefits provided by the former makes it attractive for maintaining a near real-time workflow during ultrasound imaging.

VI.6.5 Additional Comments

The proposed generic model correction shares many of the same limitations of the patient-specific correction enumerated in [188]. For example, the generic correction is still subject to several sources of propagating error in the image guidance workflow. It heavily relies upon the optical tracking system, which imparts an inherent error to each measurement made with the device, including the surface digitization using a tracked pointer or LRS, as well as the tracked ultrasound data. It also retains the assumption from the patient-specific method that the user applies the probe purely in the depth direction for each image acquisition. This again simplified the

creation of boundary conditions for the model, which is a challenge shared by both methods. In addition to the accuracy of the boundary conditions, the geometry of the mesh itself was likely the primary cause for the difference in error observed between the generic model correction and the patient-specific correction. A block of tissue is clearly a very simplistic representation of most anatomical structures on which this method would be used. The size of the block mesh also needs to be chosen before the correction can occur. Although a 10 x 10 x 10 cm cube was used for each generic model correction for the phantoms and clinical case in this study, it would be possible to pre-construct block meshes of various sizes based on prior knowledge of the anatomy of interest, which could be selected during the procedure depending on the tissue depth. It would be fairly trivial to exchange various preconstructed block meshes of different depths intraoperatively. The most computationally expensive operation during the procedure would be computing the compression depth using the LRS or pointer point cloud, and then performing the interpolation of the solution to the ultrasound data. It would conceivably be possible to perform these operations at a real-time frame rate if efficiently implemented and with the support of GPU programming.

Another assumption retained in the generic correction method was the assignment of material properties to the finite element mesh. Accurate intraoperative measurement of tissue mechanical properties is very challenging in practice. The approach taken in this work was to assume the mesh was composed of a single homogenous tissue type. Under this assumption, the biomechanical model solution would become less accurate with increasing contrast in tissue stiffness between various tissue types within the tissue of interest. It should be noted that only the relative stiffness values would affect the solution of the model in either the patient-specific or the generic model correction, as only Dirichlet boundary conditions drive the solution. Absolute values for Young's modulus would only affect the solution if force conditions were integrated into the approach in the future.

The overall result of this work is that information in tracked ultrasound data can be corrected in near real-time, provided that a measure of tissue compression is available intraoperatively. The immediate benefits are obviously in providing the clinician with more accurate size and position measurements of subsurface targets. This is important in a wide variety of procedures and anatomy, such as determining resection or ablation margins. Additionally, there are implications for more speculative work using subsurface information in other ways. One example is the image registration method presented by Lange *et al.* in which they proposed using subsurface targets such as blood vessel centerlines or bifurcations to align intraoperative space with a preoperative model [191]. The localization accuracy of such subsurface target locations could certainly affect registration accuracy, and correction of compression error would impact the result of these registrations. An analysis of the effects of integrating corrected and uncorrected ultrasound data in a registration framework would be interesting, and we leave that work to a future study.

VI.7 Conclusions

In this work we proposed and validated a novel method for correcting tissue compression error exerted by an ultrasound probe. In this method, a generic tissue model was used to estimate physical tissue deformation as a result of pressing the tracked probe into the tissue surface. The model solution was used to transform the ultrasound images to an undeformed state. The generic model correction was compared to a previously developed correction method which required a patient-specific model constructed from preoperative imaging. These two correction methods were validated in phantoms and a clinical case, and showed that alignment of freehand tracked ultrasound with co-registered tomographic images was significantly improved compared to the uncorrected data. The experimental results indicate that the generic model correction method could assist in providing more accurate intraoperative data to clinicians in procedures in which

patient-specific models may not be available from preoperative imaging, and has future implications in registration strategies utilizing subsurface features.

Acknowledgments: This work was supported in part by the National Institutes of Health award R01 NS049251 of the National Institute for Neurological Disorders and Stroke, and by the National Institutes of Health award R01 CA162477 from the National Cancer Institute.

CHAPTER VII

OTHER STUDIES TOWARDS USING ULTRASOUND IMAGE CONTENT

VII.A Geometric Reconstruction Using Tracked Ultrasound Strain Imaging

VII.A.1 Introduction and Significance of Study

One of the compelling aspects of using ultrasound as an intraoperative guidance tool is that it can offer imaging information which complements the information offered by preoperative CT or MR, due to the difference in underlying contrast mechanisms. While conventional B-mode ultrasound has been long established as an interventional imaging modality, other ultrasound-based modalities have yet to be widely realized in a surgical navigation role. Strain imaging, ARFI, and SWEI have emerged in the ultrasound literature as ways to estimate subsurface tissue stiffness, which is well recognized in the medical community as being a biomarker for many diseases, especially surgical targets such as cancerous lesions. Given that tissue stiffness is a different contrast mechanism than what is used in B-mode imaging, there are open questions with respect to the accuracy with which target borders can be identified in 3D space using tracked strain imaging. In this work, ultrasound strain imaging was integrated into a tracked surgical navigation framework and was analyzed with respect to the accuracy of a 3D strain volume reconstructed from multiple tracked 2D strain images, compared with a co-registered CT volume. This study was presented as a conference paper at SPIE Medical Imaging by Dr. Michael I. Miga in 2013.

Appearing in: T. S. Pheiffer, A. L. Simpson, J. E. Ondrake, and M. I. Miga. Geometric reconstruction using tracked ultrasound strain imaging. In: Image-Guided Procedures, Robotic Interventions, and Modeling, SPIE Medical Imaging Orlando, FL, 2013.

VII.A.2 Abstract

The accurate identification of tumor margins during neurosurgery is a primary concern for the surgeon in order to maximize resection of malignant tissue while preserving normal function. The use of preoperative imaging for guidance is standard of care, but tumor margins are not always clear even when contrast agents are used, and so margins are often determined intraoperatively by visual and tactile feedback. Ultrasound strain imaging creates a quantitative representation of tissue stiffness which can be used in real-time. The information offered by strain imaging can be placed within a conventional image-guidance workflow by tracking the ultrasound probe and calibrating the image plane, which facilitates interpretation of the data by placing it within a common coordinate space with preoperative imaging. Tumor geometry in strain imaging is then directly comparable to the geometry in preoperative imaging. This paper presents a tracked ultrasound strain imaging system capable of co-registering with preoperative tomograms and also of reconstructing a 3D surface using the border of the strain lesion. In a preliminary study using four phantoms with subsurface tumors, tracked strain imaging was registered to preoperative image volumes and then tumor surfaces were reconstructed using contours extracted from strain image slices. The volumes of the phantom tumors reconstructed from tracked strain imaging were approximately between 1.5 to 2.4 cm³, which was similar to the CT volumes of 1.0 to 2.3 cm³. Future work will be done to robustly characterize the reconstruction accuracy of the system.

VII.A.3 Purpose

Image-guided neurosurgical procedures largely rely on the premise that the tumor can be clearly demarcated in preoperative tomograms, but this is not always possible with gliomas. The tumor boundaries seen in image volumes are often not in agreement with histological examination of surrounding tissue. It has been shown that in both contrast-enhanced computed tomography (CT) [192] and magnetic resonance (MR) scans [193], cancer cells can infiltrate far beyond the region of high contrast uptake which defines the tumor boundary in these imaging modalities. It is

widely recognized in the medical community that most tumors commonly have significantly different mechanical properties than the surrounding soft tissue. The biological basis for this effect is due to changes in tissue composition, such as varied expression of collagen and greater numbers of fibroblasts [194, 195]. The contrast in mechanical properties is an important tactile cue in guiding the surgeon during resection. Technology which leverages this mechanism, such as ultrasound strain imaging, can potentially assist in identification of cancerous tissue during surgery.

Strain imaging employs a combination of image/signal processing and measurements of the physical deformation of tissue to create a relative representation of the mechanical strength of structures inside the organ of interest [196, 197]. The central premise of this work is that strain imaging can detect margins that are otherwise undetectable in contrast-enhanced tomograms, and that these margins should complement conventional contrast mechanisms, given the link between biomechanical changes and pathology. It would be significantly helpful to establish a clear understanding of the accuracy with which a strain imaging system can detect lesion boundaries. This work presents an analysis of the ability of strain imaging to characterize a target shape, by placing all of the strain images within a consistent spatial and temporal context via 3D tracking.

VII.A.4 Novel Contributions to be Presented

In the recent literature, several groups have begun to investigate the use of intraoperative ultrasound strain imaging in neurosurgery [136, 198, 199]. Chakraborty *et al.* utilized a tracking system to track the tip of an ultrasound probe in order to help position the probe over a lesion before calculating a strain image. However, they did not present any quantitative data describing the accuracy with which their strain imaging system could localize a target. Another group used tracking of an ultrasound probe to help select RF frames for block-matching such that the movement between frames is primarily in the axial direction [200]. The tracking information in that case was used to improve the generation of quality elastograms, rather than to provide surgical

guidance. To date, work to assess the ability of strain imaging to accurately localize surgical targets in a guidance setting remains relatively scarce.

The novel aspect of this paper is the merging of strain imaging with a tracking system to create a surgical guidance method which offers information not routinely available with any other intraoperative tool. The strain image plane was calibrated to the ultrasound probe such that strain images could be generated from the surgical field with automatic alignment to preoperative tomograms through tracking and registration. Co-registered strain images and tomograms were then directly compared to each other. A phantom study was performed to assess the ability of strain imaging to reconstruct the shape of a subsurface target, using contrast-enhanced CT and tracked B-mode image segmentation of the target for comparison.

VII.A.5 Methods

The ultrasound machine used in this study was an Acuson Antares with a VFX13-5 linear array transducer at 10 MHz and depth setting of 6 cm. The ultrasound was capable of standard B-mode imaging, as well as strain imaging through the commercial eSie Touch elasticity software module (Siemens Inc., Munich, Germany). Ultrasound data was tracked in 3D space by synchronizing the ultrasound video and tracking data using software based primarily on the Visualization Toolkit (VTK) on a host computer [120, 121]. The ultrasound video was captured by a Matrox Morphis Dual card (Matrox Imaging, Dorval, Canada), which recorded the analog video output of the ultrasound machine in real time. The ultrasound probe was tracked with a NDI Polaris (Northern Digital Inc., Waterloo, Canada) via the attachment of a passive optical target to the probe as shown in Figure 37.



Figure 37. The tracked ultrasound system consisted of a Polaris optical tracking system (a) and a passive target attached to the ultrasound probe (b).




Figure 38. Calibration of the tracked ultrasound system. The ultrasound beam is approximated to be planar, and is sampled with a tracked pen probe (a) to produce bright corresponding points in the ultrasound images (b). A calibration transformation is computed which best matches the image points to the physical probe points.

The ultrasound images were calibrated for tracking by the method of Muratore *et al.* as shown in Figure 38, in which the ultrasound beam was approximated as a plane and was sampled with a tracked stylus to generate a calibration transformation which mapped from image space to physical space for each pixel in the images [74]. The combination of calibration and tracking

matrices transformed each ultrasound image to a physical pose based on the pose of the ultrasound probe. The tracked ultrasound data could then be registered to preoperative images in the same manner as other intraoperative information, such as data from a tracked pen probe or laser range scanner [173].

A phantom experiment was performed to assess the ability of tracked strain imaging to localize a subsurface target shape within a co-registered coordinate system with preoperative tomograms. To perform this study, four compliant polyvinyl alcohol (PVA) gel phantoms were constructed in an anthropomorphic organ-shaped mold, using a polyester sphere mock tumor doped with barium sulfate for CT contrast, and embedded glass beads in the surface to provide landmarks for registration [140]. The phantoms were fixed to a wooden base along with a tracked reference target. The phantoms were then each imaged with an xCAT ENT mobile CT scanner (Xoran Technologies, Ann Arbor, MI) at 640 x 640 x 356 with 0.4 mm isotropic voxels. The CT volumes were registered to physical space with a point-based registration using Horn's method and the surface beads localized with a tracked stylus and manually selected in the image volumes [129]. The bulk phantom and tumor were each segmented from the surrounding structures with Analyze 9.0 (Mayo Clinic, Rochester, MN) and the resulting segmentations were used to create digital surfaces with a standard marching cubes algorithm [112]. An example of one phantom with the CT tumor and co-registered ultrasound strain images is shown in Figure 39.




Figure 39. Embedded tumor surface generated from CT (a) and corresponding lesion shown in two roughly cross-sectional tracked strain images (b) and as a volume reconstructed from strain contours (c).

The phantoms were then imaged freehand using the tracked ultrasound probe at several positions above the subsurface target, with care taken to ensure that the probe was applied only in the axial direction for quality strain image generation. B-mode images were collected in a series of several swabs over the tumors. Strain images were collected incrementally with the probe at several angles, to minimize undesirable out-of-plane motion during strain image formulation. Each ultrasound acquisition lasted approximately 10 seconds. Lesion borders in each ultrasound image slice were segmented semi-automatically using a VTK implementation of the livewire technique based on Dijkstra's algorithm [201]. The segmentation resulted in a contour of points corresponding to the lesion border in B-mode and strain images. These points were then transformed to physical space using the ultrasound calibration from above. When a target was imaged from multiple points of view, tracking was used to place each border contour in its proper pose with respect to each other contour. Thus, the tracked strain images generated a point cloud which described the target shape based on the observed relative tissue stiffness. A volumetric representation of the point cloud was created by performing a 3D Delaunay triangulation of the contour points to produce a tetrahedral mesh [202]. The volume of the tumor mesh from B-mode and strain imaging was then compared to the volume of the CT tumor mesh.

VII.A.6 Results

The result of the point-based registration is shown in Figure 39. This figure shows the alignment of CT and tracked ultrasound images, as well as the ability of the strain imaging to localize the phantom tumor. Note that the borders of the strain image lesion were very well defined. The location and orientation of the strain images qualitatively displayed good agreement with the CT tumor as a result of the point-based registration done to align physical and CT space. The ability to accurately align ultrasound tumor borders with borders from preoperative imaging is a very important component of a tracked strain imaging system, and this aspect will be explored more

thoroughly in a future work. The primary focus of this work was to assess the capability of a series of tracked strain image contours to generate a volume.

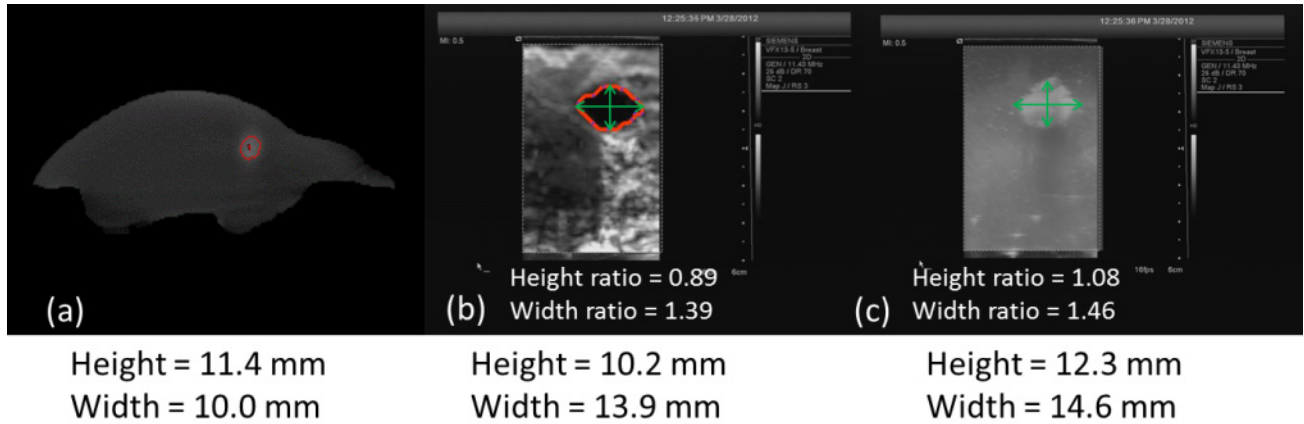


Figure 40. An example of the CT tumor dimensions (a), the corresponding strain tumor dimensions (b), and the B-mode tumor dimensions (c).

In order to reconstruct the shape of the phantom tumors, the borders of the lesions were extracted from contours segmented from tracked B-mode and strain images acquired from various positions above the target. Figure 40 shows an example of the tumor geometry in CT, strain imaging, and B-mode. For a quantitative comparison, the volume of the tumors in each phantom measured by each modality is shown below in Table 6.

Table 6. Results of geometric reconstruction of the phantom tumor volumes using CT, tracked B-mode images, and tracked strain images.

Phantom #	CT Volume (cm ³)	B-mode Volume (cm ³)	Strain Volume (cm ³)
1	1.17	1.50	1.70
2	2.25	1.77	1.50
3	1.98	1.59	2.39
4	1.01	1.95	1.90

The volume of the physical tumors during phantom construction was intentionally variable, and this was reflected in the segmented CT tumor volumes, which ranged from approximately 1.0 to 2.3 cm³. The volumes of the tumors reconstructed from tracked ultrasound contours were slightly different from the corresponding CT volumes, but were on the same order of magnitude and ranged from 1.5 to 2.4 cm³. In addition, the ultrasound-based volumes were typically more similar to each other than to the CT volume. This behavior was probably due to a combination of factors, particularly compression of the tissue by the ultrasound probe, as well as the scarcity of ultrasound border data compared to CT.

Table 7. Height-to-width ratios for the phantom tumor volumes from CT, B-mode reconstructions, and strain imaging reconstructions. Height and width were measured approximately through the tumor centroids.

Phantom #	CT Height/Width	B-mode Height/Width	Strain Height/Width
1	1.14	0.734	0.842
2	1.02	1.11	0.643
3	1.13	1.21	0.728
4	1.67	1.33	0.931

Height-to-width ratios for the phantom tumors were measured through the tumor centroids along axes defined by the CT slicing, and these results are shown in Table 7. The height-to-width ratio of the tumor in the strain-reconstructed tumor volumes showed a tendency to be smaller than either the CT or B-mode ratios. This could be partially attributed to compressional effects from the probe, as a certain level of pre-compression of the tissue was sometimes needed in order to generate proper tissue displacements for calculation of strain images. It should also be noted that although the CT in this case was considered the gold standard, the CT contrast enhancement of the tumor was imperfect due to uneven distribution of the contrast agent. This was verified using the B-mode images such as the example in Figure 40, which displayed a larger cross-sectional area than

the CT through the tumor centroid. This uncertainty could be reduced or eliminated in the future by using a tumor manufactured with known dimensions with a machined mold.

VII.A.7 Conclusions

We conclude that combining tracking data with ultrasound strain imaging is a feasible method of reconstructing the shape of a subsurface target. Given its unique contrast mechanism, tracked strain imaging shows promise as a complementary intraoperative data source for image-guided procedures. Future work will include further testing of the robustness of shape reconstruction in phantoms, as well as the development of novel methods which leverage the tracking system to correct for tissue displacement during strain imaging.

Acknowledgements: This work was supported by the National Institutes of Health–National Institute for Neurological Disorders and Stroke under Grant R01 NS049251-01A1, and by the National Institutes of Health award R01CA162477 from the National Cancer Institute. The authors also gratefully acknowledge Dr. Brett Byram (Duke University) for assistance and training with ultrasound strain imaging.

VII.B Intraoperative Segmentation of Tracked Ultrasound Images

VII.B.1 Introduction and Significance of Study

One of the current limitations with ultrasound as an interventional imaging modality is the difficulty in image interpretation and structure identification. This study presents an initial effort toward expanding the utility of intraoperative ultrasound by creating a tool to assist in interpretation of image features. This tool took the form of a semi-automatic image segmentation approach designed to identify the margins of subsurface surgical targets such as tumors. The segmentation algorithm was formulated specifically to take advantage of the tracked ultrasound and biomechanical model correction pipeline previously developed in order to delineate target borders with an appropriate spatial and temporal context. The following sections describe the segmentation framework and give preliminary results in clinical data, with a further discussion regarding the limitations of the proposed technique.

VII.B.2 Abstract

Identification of objects is an important and challenging task in interventional ultrasound imaging. This work proposes an algorithm to semi-automatically segment targets in ultrasound images using both intensity information and co-registration with preoperative imaging modalities to help guide the segmentation. A level set formulation was used as the overall framework, which included terms for advection, curvature, propagation, and a shape model constructed from the preoperative image volume. The proposed segmentation method was shown to produce contours which matched manually created contours to within 0.71 and 1.29 mm in B-mode and strain image datasets, respectively.

VII.B.3 Introduction

Ultrasound image segmentation is a challenging task due to the characteristic noise and artifacts which are often prevalent in the images. The presence of speckle, attenuation, and shadowing in the data often complicate segmentation techniques traditionally used in other

imaging modalities. In addition, the contrast between various tissue types of interest can be quite low in B-mode. These issues have motivated interest in segmentation algorithms which leverage the unique characteristics of ultrasound imaging or information from other imaging modalities to assist in the segmentation problem at hand. Much effort has been made toward tissue segmentation for diagnostic purposes, but image-guided interventions are emerging as an important application in which image segmentation may have a strong clinical impact.

A review of the ultrasound segmentation literature from the last decade is given by [109, 110]. Generally speaking, the most prevalent segmentation strategies primarily target B-mode images rather than the raw radiofrequency (RF) signals, as most clinical ultrasound machines do not provide convenient access to those forms of data. Although several machines now exist which offer research interfaces to the raw data, this work deals with segmentation of B-mode images. The two most popular frameworks for B-mode segmentation have been the ‘snakes’ and level set propagation approaches, in large part due to the relative ease and speed of implementation. Regardless of the chosen framework, it is almost always necessary to constrain the segmentation in some way due to the noisy characteristics of ultrasound. Common constraints include the B-mode intensity distribution, gradient, phase, and texture. In addition, prior knowledge of general shape can be used to guide segmentation using the image intensity-based constraints above.

The primary contribution of this work is the investigation of an intraoperative ultrasound tumor segmentation technique which utilizes optical tracking, tissue deformation modeling, and tomogram co-registration to provide shape guidance. A level set framework was chosen for the ability to evolve contours that change topology, such as might be the case with tumors having complex shapes, without the need to parameterize these objects. The governing level set equation was formulated with terms for controlling edge advection, overall contour propagation, and contour curvature. In addition, a term for enforcing a general shape to the contour propagation was

included utilizing novel shape prior information from co-registered model-updated tomogram segmentations.

VII.B.4 Methods

The proposed segmentation method is one part of a general intraoperative guidance pipeline. In the following sections we describe the preoperative data collection and preprocessing steps necessary for the pipeline, followed by the intraoperative data collection and the segmentation procedure itself.

VII.B.4.1 Preoperative Patient Model

Prior to the image-guided intervention, the patient is imaged with magnetic resonance imaging (MRI). The MR volume in this study was acquired with dimensions 256 x 256 x 216 at 1 x 1 x 1.2 mm voxels. The organ and tumor were both segmented from these images and triangulated surfaces are created using the marching cubes algorithm [112]. Some amount of manual interaction is often necessary for the tumor segmentation, due to blurred edges occurring due to diffuse infiltration of the lesion into normal tissue. After the organ and tumor surfaces were created, custom software was used to create a tetrahedral mesh of the organ geometry [113].

VII.B.4.2 Intraoperative Tracked Ultrasound

The ultrasound machine used in this study was an Acuson Antares (Siemens, Munich, Germany) with a VFX13-5 linear array probe at 10 MHz. The pose of the ultrasound probe was detected by a Polaris Spectra optical tracking camera (Northern Digital Inc., Waterloo, Canada). A passively-tracked optical target was rigidly attached to the probe and calibrated using the method of Muratore *et al.* [74]. The probe tracking data and ultrasound video were collected and synchronized on a host computer [120, 121], and the ultrasound images were captured with a Morphis Dual frame-grab card (Matrox Imaging, Dorval, Canada). Each image was stored along with its pose in 3D space given by the tracking data.

VII.B.4.3 Soft-Tissue Deformation Correction

Our group currently employs model-based correction strategies to address misalignment between intraoperative and preoperative data due to non-rigid tissue movement during brain [94, 96-98, 101, 163, 180, 203, 204] and liver [159, 162, 164] surgeries. The general approach is to measure surface displacements from the organ of interest, and apply these as boundary conditions to a biomechanical model of the tissue. Most recently, a method was proposed which utilizes a patient-specific finite element model of tissue to predict subsurface target movement due to tissue compression exerted by the probe during ultrasound imaging [188]. This method was found to improve the alignment of compressed ultrasound targets with co-aligned tomogram targets to within 2 to 3 mm of error. In this work, this correction method was employed to ensure that the shape priors from co-registered tomograms (discussed further in the next section) were positionally and geometrically consistent with the ultrasound data during the segmentation procedure.

VII.B.4.4 Segmentation Procedure

Our approach utilizes level sets with geodesic active contours [205]. The signed distance function, ψ , was chosen as the embedding function for the zero level set [206]. The movement of the contour is controlled by the generic level-set equation:

$$\frac{d}{dt}\psi = -\alpha A(x) \cdot \nabla\psi - \beta P(x)|\nabla\psi| + \gamma Z(x)\kappa|\nabla\psi| \quad (24)$$

where A is an advection term, P is an expansion term, and Z is a spatial modifier term for the mean curvature κ [207]. Scalar constants α , β , and γ are used to weight the relative contributions of the three terms to the curve evolution. The contour is extracted at any time as the zero level-set

$$\Gamma((x), t) = \{\psi(x, t) = 0\} \quad (25)$$

The function ψ is solved at each time step in a finite difference scheme to create an image field of function values. The best approximation of the object surface is then to calculate the image positions corresponding to zero crossings in function values.

The geodesic active contours equation in (24) may be extended with an additional shape prior term. We used the formulation of this term presented by Leventon *et al.*:

$$\lambda(\psi^*(x) - \psi(x)) \quad (26)$$

where ψ^* is the best estimate of the final curve as determined by a maximum *a posteriori* approach and λ is a scalar weight constant [208]. The estimate of the final curve can be calculated by

$$\langle \alpha_{MAP}, r_{MAP} \rangle = \underset{\alpha, r}{\operatorname{argmax}} P(\alpha, r | \psi, \nabla I) \quad (27)$$

which seeks to maximize the probability of a set of shape parameters, α , and rigid pose parameters, r , (both defined later) given the surface ψ at some point in time and the gradient of the image, ∇I . Equation (27) can be expanded using Bayes' Rule:

$$P(\alpha, r | \psi, \nabla I) = \frac{P(\psi | \alpha, r) P(\nabla I | \alpha, r, \psi) P(\alpha) P(r)}{P(\psi, \nabla I)} \quad (28)$$

where the denominator can be discarded as it does not depend on shape or pose. The first term in the numerator is modeled as a Laplacian density function over $V_{outside}$, the volume of the current curve ψ which lies outside the estimated final curve ψ^* :

$$P(\psi | \alpha, r) = \exp(-V_{outside}) \quad (29)$$

The second term describes the probability of observing a given image gradient given the current and final curves. It can be assumed that when the curve correctly outlines the boundary of an object, its relationship with $|\nabla I|$ should be Gaussian. Thus this term is modeled as a Laplacian of the goodness of fit of the Gaussian $h(\psi^*)$ to the samples $(\psi^*, |\nabla I|)$:

$$P(\nabla I | \alpha, r, \psi) = \exp\left(-|h(\psi^*) - |\nabla I||^2\right) \quad (30)$$

The third term is the probability distribution of the shape prior parameters. The typical strategy for computing a prior on shape variation is to build a shape model given a set of training images. Given

n training images, the training set $R = \{u_1, u_2, \dots, u_n\}$ consists of the signed distance maps for each image. The mean surface, μ , is computed as the mean of the signed distance maps, $\mu = \frac{1}{n} \sum u_i$. Principle component analysis (PCA) is used to compute the shape variance. First, the mean shape is subtracted from each u_i to create mean-offset maps, which are then each placed as column vectors in an $N \times n$ -dimensional matrix M , where N is the number of samples (pixels) in each distance map. Instead of performing the eigen decomposition on the large $N \times N$ covariance matrix $M^* M^T$, we decompose the much smaller $n \times n$ inner product matrix $M^* M$. The resulting eigen vectors, E , are then multiplied by the matrix M to get the principle component images, U . The object ψ can be estimated by the first k principle components as a k -dimensional vector of shape parameters, α :

$$\alpha = U_k^T (\psi - \mu) \quad (31)$$

The shape prior probability term is thus modeled as a Gaussian with shape variance Σ_k :

$$P(\alpha) = \frac{1}{\sqrt{(2\pi)^k |\Sigma_k|}} \exp\left(-\frac{1}{2} \alpha^T \Sigma_k^{-1} \alpha\right) \quad (32)$$

where Σ_k is a diagonal matrix containing the first k eigenvalues.

The last term, $P(r)$, is the probability of observing a set of pose parameters. We currently do not assume any pose more likely than another, and so merely assume a uniform distribution over these parameters. The approximate signed distance map to the shape prior can be computed as

$$\psi^*(x) = \mu(T(x)) + \sum_i^k \alpha[i] * \sigma[i] * U_i(T(x)) \quad (33)$$

where $\sigma[i]$ are the square root of the eigenvalues, and $T(x)$ is a transform which defines the pose of the shape, which we chose to be a rigid transform function with parameters, r , of rotation and translation. The shape prior term in (26) is thus constructed at each iteration of the level set evolution by referring to the PCA description of the object shape and then using an optimizer to solve for a set of shape and pose parameters which maximize the posterior probability in (28). The

segmentation algorithm described above was implemented in C++ primarily using ITK (Kitware Inc., Clifton Park, NY) and the segmentation pipeline is shown in Figure 41.

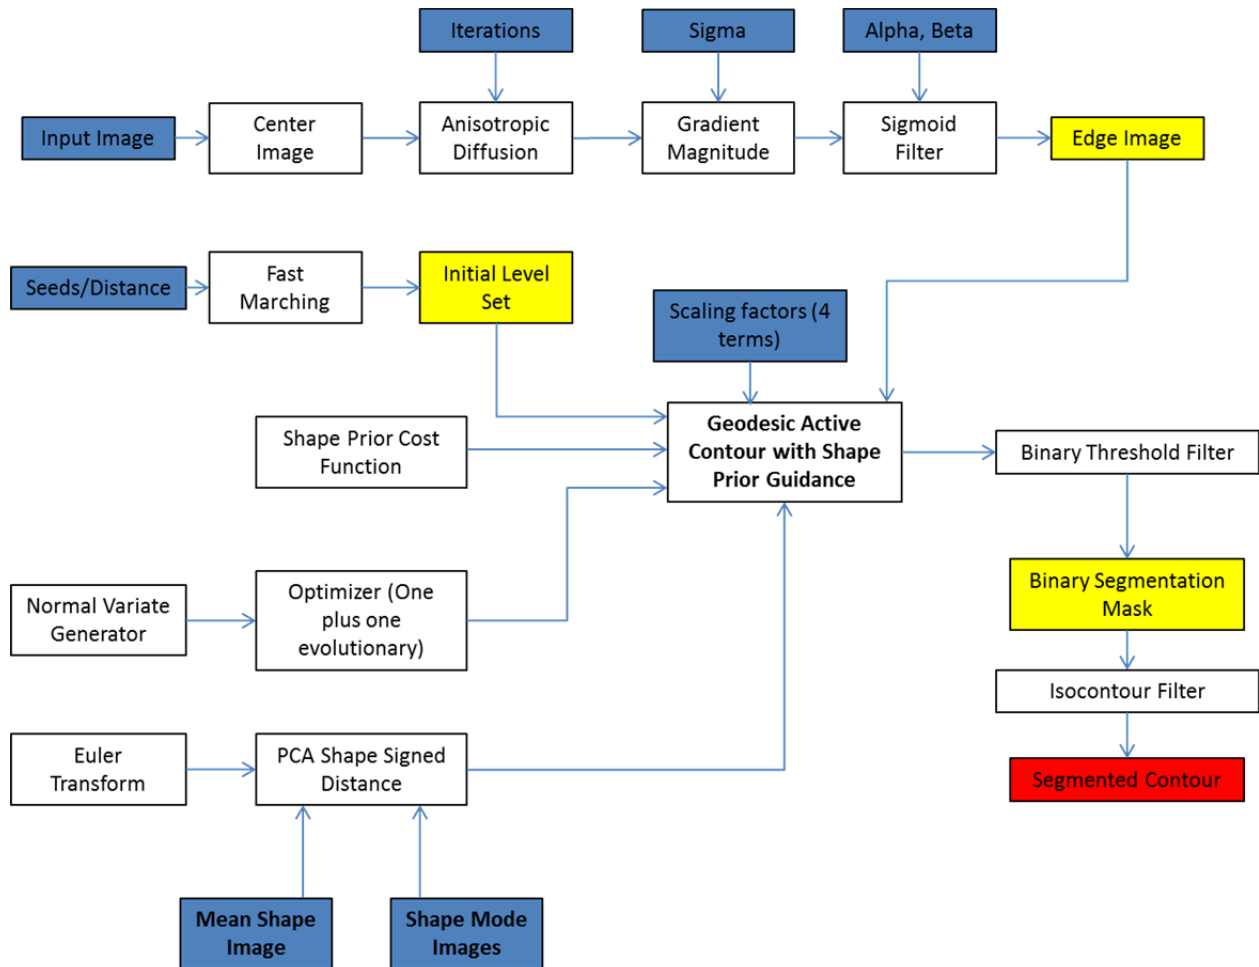


Figure 41. Ultrasound image processing and segmentation pipeline using the shape model derived from a co-registered tomogram segmentation.

The segmentation is performed after soft tissue movement has been corrected with our model-based approach. Each ultrasound image is thus associated with a specific model deformation derived from the position of the ultrasound probe in the tissue. The statistical shape model is then created from intersections of the ultrasound beam with the co-registered tomogram target segmentation as shown in Figure 42. Typically it is desirable to create a shape model from a very large dataset, and since this is not available in practice for irregular patient-specific structures such

as tumors, the approach here was to perturb the ultrasound beam in the elevational directions several millimeters in each direction and record the varying intersected target borders. This was done for several planes in either direction in order to gather at least five shape images with which to create the shape model.

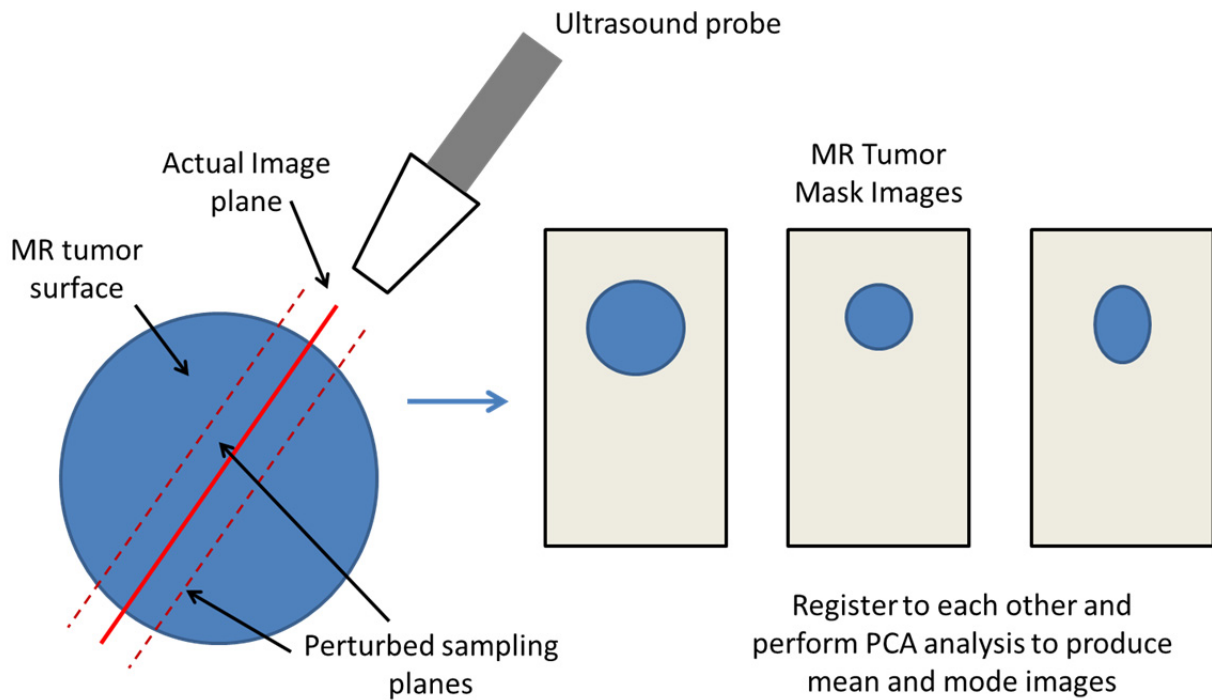


Figure 42. Creation of statistical shape model from co-registered tomogram target for the ultrasound segmentation pipeline.

VII.B.4.5 Experiments

The segmentation method was applied to a B-mode image and strain image from a brain tumor resection patient at Vanderbilt Medical Center. Prior to the study, informed written consent was obtained from the patient with the approval of our Institutional Review Board. For each image, the tumor was first segmented manually, and then the tumor was segmented using the semi-automatic approach over a series of six trials. In the first trial, the four weighting parameters for the level set equation were chosen empirically to give a reasonable segmentation. In each of the remaining trials, the weighting parameters were altered in the manner given in Table 8. These trials

were not intended to be a fully comprehensive sensitivity study of the weighting parameters, but rather provide an idea of at least the variability of the resulting segmentation contour based upon selection of these parameters.

Table 8. The level set weighting parameters used for each of the segmentation trials (six trials for both the B-mode and strain image). The parameters in Trial 1 were chosen empirically to give a satisfactory segmentation, and then following trials modified one weighting factor at a time (bolded).

Level Set Parameter	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6
Advection	6	12	6	6	6	6
Propagation	1.5	1.5	1.5	3	1.5	1.5
Curvature	1	1	2	1	1	1
Shape	0.2	0.2	0.2	0.2	0.4	0

VII.B.5 Results

Figure 43 shows an example of the shape model constructed from the intersection of a tracked ultrasound image with the co-registered MR tumor borders. The mean image and first two modes are shown and give an idea of how the shape given by the MR changed with respect to the intersection with the ultrasound plane.




Figure 43. Example of the statistical shape model created from intersection of a tracked ultrasound slice with a co-registered MR tumor, showing the mean image (left) and three standard deviations of the first and second modes (middle and right).

Figure 44 shows a B-mode image slice of a tumor collected prior to surgical resection of the brain tumor, and also shows the registration between the tracked ultrasound data and preoperative MR. This alignment was used to construct the shape model from the MR tumor as described above.




Figure 44. A B-mode image of the brain tumor (left), and the tracked ultrasound image in 3D space with the co-registered MR tumor (right).

The manual B-mode segmentation and the results of the segmentation trials using varying weighting parameters are visualized in Figure 45.




Figure 45. Illustration of the segmentation variability in the B-mode image according to level set weighting parameters. A manual user-drawn segmentation is shown in (a). (b) shows the result of the semi-automatic segmentation using empirically selected weighting parameters to give a contour similar to the manual contour. The contours given in (c-g) show some examples of altering the parameters chosen for (b): (c) is the result of doubling the advection weight; (d) is the result of doubling the curvature weight; (e) is the result of doubling the propagation weight; (f) is the result of doubling the shape model weight; and (g) is the result of eliminating the shape model term from the level set equation.

The comparison of each B-mode semi-automatic segmentation contour with the manually created segmentation contour in terms of MHD value is given in Table 9.

Table 9. Comparison of the semi-automatic contours generated for the B-mode image with various weighting parameters, compared to the manual segmentation.

	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6
MHD (mm)	0.71	16.89	0.98	13.70	2.10	8.70

Figure 46 shows a strain image slice of the same tumor shown in Figure 44, and also shows the registration of the tracked strain image with the preoperative MR.




Figure 46. A strain image of the brain tumor (left), and the tracked ultrasound image in 3D space with the co-registered MR tumor (right).

The manual strain image segmentation and the results of the segmentation trials using varying weighting parameters are visualized in Figure 47.




Figure 47. Illustration of the segmentation variability in the strain image according to level set weighting parameters. A manual user-drawn segmentation is shown in (a). (b) shows the result of the semi-automatic segmentation using empirically selected weighting parameters to give a contour similar to the manual contour. The contours given in (c-g) show some examples of altering the parameters chosen for (b): (c) is the result of doubling the advection weight; (d) is the result of doubling the curvature weight; (e) is the result of doubling the propagation weight; (f) is the result of doubling the shape model weight; and (g) is the result of eliminating the shape model term from the level set equation.

The comparison of each strain image semi-automatic segmentation contour with the manually created segmentation contour is given in Table 10.

Table 10. Comparison of the semi-automatic contours generated for the strain image with various weighting parameters, compared to the manual segmentation.

	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6
MHD (mm)	1.29	2.75	2.52	7.59	5.07	26.46

VII.B.6 Discussion

The results of the B-mode segmentation shown in Figure 45 demonstrate the effects of altering the level set weighting parameters for the advection, propagation, curvature, and shape prior terms. Figure 45b shows that when the parameters are chosen carefully, the semi-automatic method can produce a final contour which is very similar to the manually created contour shown in Figure 45a. However, Figure 45c-g show that modifying any of the four weighting parameters has significant effects on the segmentation. Figure 45c shows that when the advection weight is increased, the contour is more likely to adhere to very strong edges which can be outside or inside the bounds of the desired object, as shown by the lower left and lower right edges of the resulting contour, respectively. Figure 45d shows that when the curvature weight is increased, the contour is overall forced to have fewer regions of high curvature, such as on the left portion of the contour, and some finer details may be lost as a consequence. Figure 45e shows that when the propagation weight is increased, the contour expands well beyond the edges of the tumor. Figure 45f shows that when the shape model weighting is increased, the contour is forced to conform more to the shape of the co-registered MR tumor border (not shown). Finally, Figure 45g shows that when the shape term is eliminated completely, the level set propagates and tends to adhere to image edges which do not necessarily correspond to the tumor. The MHD values comparing the semi-automatic segmentation contours to the manual contours in Table 9 further demonstrate the variability of the segmentation performance. In the case of Trial 1, in which the segmentation parameters were empirically chosen by the user, the contour MHD error value was only 0.71 mm, indicating that the overall shape and size of the semi-automatic contour was quite similar to the manually created

contour. However, over the other trials, the MHD errors ranged from 0.98 to 16.89 mm depending on the weighting parameters which were altered. This range is obviously much larger than would be acceptable for the intended purposes of accurately detecting the size and position of intraoperative targets, indicating the importance of the level set weights in this method.

The same overall trends were observed in Figure 47 with respect to the segmentation of the tumor in the strain image. The primary finding was that alterations of any of the weighting parameters can lead to drastic differences in the final segmentation contour. The differences observed between Figure 47 and Figure 45 are primarily due to the difference in intensity profiles between the strain and B-mode image, which also points to the difference in contrast mechanism between the two modalities. The strain image MHD values in Table 10 also show a trend similar to the B-mode image. The Trial 1 MHD error was the lowest, with a value of 1.29 mm. Over the other five trials, the MHD error ranged from 2.52 to 26.46 mm depending on the selection of level set parameters.

The reliance of the proposed segmentation algorithm upon the choice of weighting parameters does present a significant obstacle toward adoption within a general intraoperative ultrasound framework. There is generally no way to predict the optimal weights to produce an acceptable result unless the image intensity profile is well understood beforehand, which is a challenging requirement for intraoperative ultrasound due to the freehand nature of the acquisitions. The simplest technical solution is to acquire the ultrasound images and then manually alter the weights until a satisfactory contour is observed. However, this would be a large disruption to normal surgical workflow for most procedures, since the semi-automatic segmentation using shape priors takes approximately 10 to 20 seconds on one thread of an Intel Core 2 Quad CPU at 2.4 GHz. Future work should most likely focus on ways to improve the speed of the segmentation, or identify interventions in which very consistent ultrasound images are routinely produced which could take advantage of *a priori* knowledge of useful weighting values for the level set equation.

VII.B.7 Conclusions

This study presented a level set method for intraoperatively segmenting ultrasound images using the image intensity information in the ultrasound data in combination with co-registered preoperative segmentations from tomographic image volumes. The method was shown to produce contours which were similar to manually created contours, but required careful selection of weighting parameters in the level set equation. The dependence of the segmentation on these parameters presents an obstacle to the broad adoption of this method for image-guided interventions, as the time to determine proper parameters intraoperatively is a significant workflow hinderance. A possible avenue of future research is to determine whether ultrasound images of certain anatomy are consistent enough that certain parameter presets could be used without the need for intraoperative determination of the proper parameters.

CHAPTER VIII

SUMMARY AND FUTURE DIRECTIONS

This dissertation presented studies that evaluated devices and methodology in support of a tracked ultrasound system for accurate intraoperative image guidance. Chapter IV presented the design and characterization of a novel laser range scanner to be used as a highly accurate surface digitization tool. This device was a key component in an intraoperative pipeline by enabling surface-based image-to-physical registrations, and measuring soft tissue deformations. The tracked data from this device led to the identification of a fundamental problem of tracked ultrasound systems, which is the registration error introduced by soft tissue deformation exerted by the ultrasound probe on the organ of interest. Chapter V introduced a novel method of correcting this registration error by utilizing a biomechanical model approach. This method utilized the rigid registration from the LRS or tracked pointer in order to align the tracked ultrasound probe to the preoperative patient model. The position of the probe within the model was then used to create boundary conditions for a forward model solution, given a set of assumptions regarding the handling of the probe by the user, and the solution was used to correct for the tissue deformation. That chapter demonstrated that this method is capable of reducing alignment error due to surface compression from up to 1 cm down to a clinically relevant range of 2 to 3 mm after correction. The compression correction was demonstrated in that chapter in a brain tumor case, but is applicable in a variety of surgical domains, as shown by the examples below in Figure 48.




Figure 48. Examples of the compression correction applied in a clinical brain case (a&b), a liver phantom (c&d) and a clinical breast case (e&f from [209]). In each pair of images, the first image shows a slice from the preoperative tomogram, and the second shows an overlay of the co-registered ultrasound image after compression correction.

Chapter VI then presented a more generalized method for compression correction which can be extended to any tracked ultrasound system without the need for preoperative imaging, and which also has the potential for real-time corrections. The generic model correction method was compared to the previously developed patient-specific correction method and was found to improve alignment with co-registered tomograms with only submillimetric differences between the two methods. Finally, Chapter VII presented other studies utilizing ultrasound image information to designate features intraoperatively.

There are many avenues for further improvement to the methodology in this dissertation. For instance, the compression compensation methods presented in Chapters V and VI relied on an assumption of homogenous material properties in the imaged tissue in order to simplify the problem of modeling the deformation. The reality of surgical interventions often involves heterogeneous regions of tissue due to the presence of disease or various other normal tissue types in close proximity. Large differences in material properties between tissue types can lead to significant error in the biomechanical model solution if the assumption of homogeneity is kept. The

ability to accurately designate material properties in the model intraoperatively would be a valuable addition to the field. A particularly attractive approach would be to utilize ultrasound strain imaging to provide at least relative stiffness values in real-time (see Appendix A). The compression correction methods in Chapters V and VI used only Dirichlet boundary conditions to drive the solution, implying that stiffness ratios are sufficient to solve the model and that absolute material properties would not be necessary. Strain imaging would be both able to satisfy this input requirement and would also utilize the same tracked ultrasound framework in order to place the data in a consistent spatial context with all other data. Preliminary work has been done to co-register tracked strain imaging with other surgical data [198, 210] during image-guided interventions, but so far strain imaging has not been extensively used to estimate material properties intraoperatively to improve guidance within the framework of a predictive biomechanical tissue model. One potential obstacle to this approach would be the potential need for extended intraoperative imaging time with the ultrasound machine to acquire the strain images in addition to normal B-mode images.

An alternative strategy would be to compute material properties from preoperative information, and then align this preoperative data to the intraoperative space using standard registration techniques. Previous work has shown that modality-independent elastography (MIE) can be used on typical preoperative CT or MR images to produce estimates of tissue stiffness throughout the image volume [211-214]. For example, Figure 49 shows stiffness contrast of a tumor located in the center of two breast volumes reconstructed using MIE.




Figure 49. Example of MIE applied to a CT breast volume (a) with cross-section showing tumor-to-normal elasticity contrast (b), and an MR breast volume (c) with cross-section showing similar contrast (d). [215]

Recent advancements to MIE workflow, such as the automation of boundary conditions in the tissue model (see Appendix B), indicate that MIE could potentially be performed prior to surgery with minimal disruption to normal surgical preparation or workflow [215]. Most recently, MIE is being validated in preclinical breast cancer models [216, 217]. Integrated knowledge of material properties in some form during surgery, whether from MIE or another source, would greatly complement the advancements presented in this work toward a more accurate and streamlined ultrasound guidance system.

In addition to the issue of material properties, one of the primary caveats of the compression correction methods presented in this dissertation was that the user must press the probe directly in the depth direction of the ultrasound beam. This assumption was made because in practice it is a very challenging task to measure the intraoperative probe-tissue interactions in

other compression scenarios such as dragging. Although it was shown in the preceding studies that it is feasible to account for purely depth-wise compression, there are ultrasound-guided procedures in which it is much more difficult for the user to carefully control the tissue interaction. For example, image-guided breast surgery is a procedure which is recently emerging as a target for navigational methodology including tracked ultrasound. One major hurdle in translating established image-guidance techniques to this anatomy is the greater potential for large tissue movements during imaging. This is illustrated below in data obtained from a healthy volunteer with benign cysts. In this dataset, an image-to-physical registration was performed to a supine MR volume using skin fiducials. The tracked ultrasound data was acquired by an experienced surgeon in one continuous swab while attempting to localize every cyst. In this case, the surgeon was not given special instruction with respect to manipulation of the tissue with the ultrasound probe, and so this dataset is representative of normal target localization in the breast domain. Each ultrasound slice in which a cyst was visible was segmented and the contour was rendered in 3D space.




Figure 50. An example of tracked ultrasound data acquired during a breast examination of a patient with benign cysts (a), showing the registration of LRS cloud to MR surface (b). The swabbed ultrasound cyst contours are spread over a wide volume due to dragging of the tissue (c), which was confirmed in the probe tracking vector history (d).

Figure 50c shows a large amount of tissue dragging due to the swabbing search for each cyst. It can be seen that the majority of the ultrasound contours are spread out in broad swaths which are not spatially consistent with the MR cyst segmentations. The misalignment in this case was not only just in the depth direction of the ultrasound images, but rather it was in the lateral/elevational directions as well. Analysis of the tracking data during the tissue swab confirmed large lateral components to the probe movement. The probe trajectory vectors from the tracking data tended to be very spatially consistent with the direction of misalignment between the ultrasound and MR data, as shown in Figure 50d. In general the direction in which the probe was moving immediately prior to an ultrasound slice corresponded to the misalignment, indicating a dragging effect exerted on the tissue. Obviously this phenomenon confounds the compression correction methods previously presented in this dissertation, which did not integrate any way to compensate for lateral movement.

The data in Figure 50 above was obtained with no special instruction to the breast surgeon, and so one preliminary attempt at resolving this issue was to take extra precautions to avoid dragging effects during the swab procedure. The same patient was imaged again, but when the largest cyst was located during the swab, the surgeon lifted the probe off of the breast surface in an attempt to let the tissue settle from the dragging manipulation back into the rest state. Then the probe was brought back perpendicularly to the tissue and the cyst was imaged from several different angles. The resulting ultrasound surface and its unmodified registration with the MR (based on a rigid breast surface registration, as before) is shown below in Figure 51a.

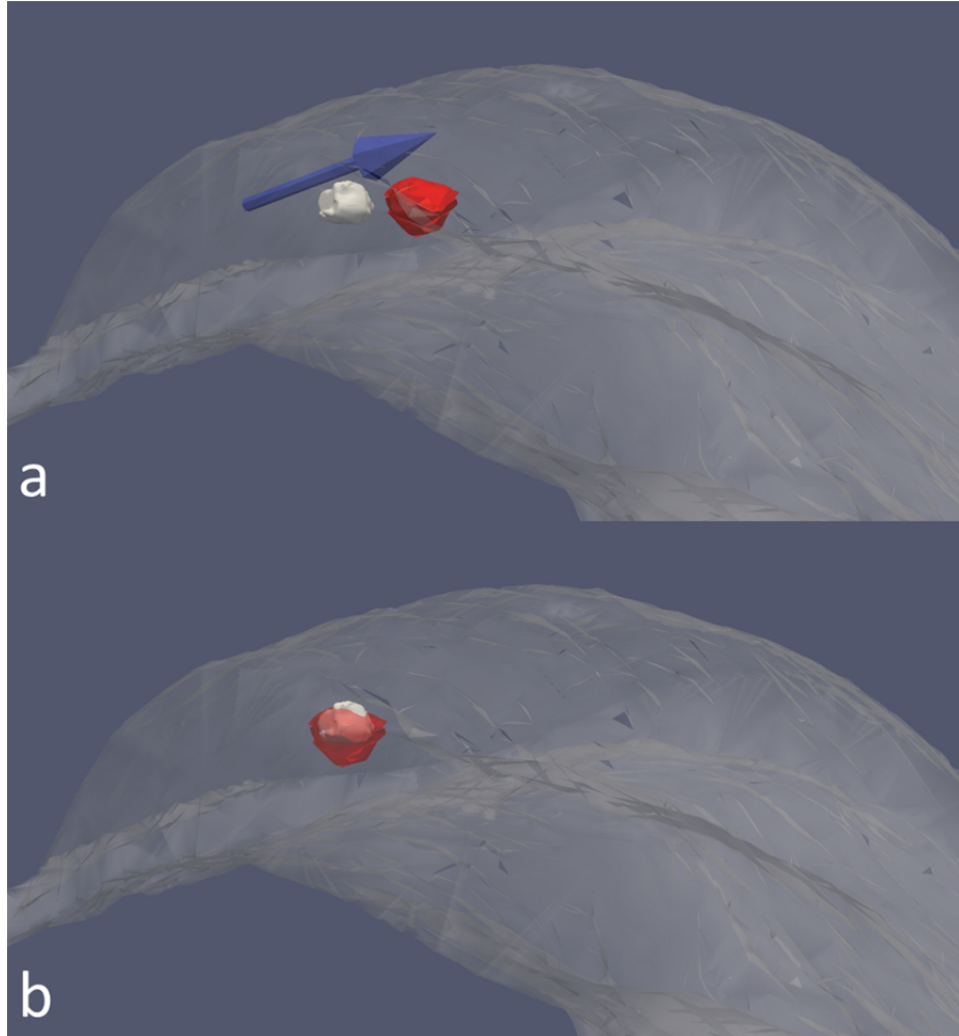


Figure 51. Example of an ultrasound cyst surface (red) misaligned with the corresponding MR surface, with the prior swab trajectory visualized as a blue arrow (a). The known vector of misalignment between the ultrasound and MR centroids was computed and the closest vector in the previous three seconds of tracking data was used to translate the ultrasound cyst to an earlier position in the swab.

Figure 51 illustrates that even after the precautions taken with respect to careful swabbing there can still be lingering effects from tissue swabbing, although likely exaggerated in this case because the benign breast cysts were noted by the surgeon as being highly mobile targets. Examination of the probe tracking data showed again that immediately prior to locating the cyst and lifting the probe off of the surface, the swab trajectory closely corresponded to the direction of misalignment between MR and ultrasound. In order to further bolster the hypothesis that tissue

dragging was the cause of the cyst misalignment, a crude correction was devised in which the known misalignment vector between the MR and ultrasound centroids was calculated, and then the tracking data from three seconds prior was searched for the swab vector which most closely matched the centroid vector. The best tracking vector was then used to naively translate the ultrasound cyst surface backward along the swab path, which resulted in the alignment shown by Figure 51b. While this approach would not be robust enough to deploy in a general sense, it does serve to illustrate a scenario in which a more sophisticated correction method is needed to reconstruct all of the tissue movement.

One possible approach to this problem may be to record surface movement during ultrasound imaging with another device, such as an LRS. The movement of surface landmarks visible in the textured point cloud data could be used to more accurately formulate boundary conditions for the model-based correction. However, the LRS is not an ideal solution because it is intrusive to workflow and cannot provide real-time surface acquisitions. A technology which is emerging as an alternative to the laser-based measurement, however, is stereo-pair surface reconstruction [156, 218]. In this method, two cameras in a stereo configuration are calibrated to construct a 3D point cloud of a target surface, as shown below in Figure 52.




Figure 52. Example of a stereovision reconstruction on a cortical surface phantom, showing the left and right camera views and the resulting 3D point cloud [219].

It was shown by Kumar *et al.* that a stereovision surgical microscope system can be used to create point clouds with accuracy in the 0.46 – 1.5 mm range [219]. While this is not quite as good as what can be achieved with the LRS, the stereovision approach can be performed in real-time. In

addition, this approach can potentially be performed continuously throughout a procedure to provide surface measurements with little interruption to normal workflow, as long as the cameras are pointed at the surgical field. Integration of this data into the model corrections developed in this dissertation, along with the possibility of using more sophisticated friction models for the probe-tissue interactions, could further advance the broader utility of ultrasound as a quantitative interventional localization tool.

REFERENCES

1. Regner, D.M., et al., *Breast lesions: evaluation with US strain imaging--clinical experience of multiple observers*. Radiology, 2006. **238**(2): p. 425-37.
2. Cleary, K. and T.M. Peters, *Image-Guided Interventions: Technology Review and Clinical Applications*. Annual Review of Biomedical Engineering, Vol 12, 2010. **12**: p. 119-142.
3. Galloway, R.L., *The process and development of image-guided procedures*. Annual Review of Biomedical Engineering, 2001. **3**: p. 83-108.
4. Cash, D.M., et al., *Concepts and preliminary data toward the realization of image-guided liver surgery*. Journal of Gastrointestinal Surgery, 2007. **11**(7): p. 844-859.
5. Benincasa, A.B., et al., *Feasibility study for image-guided kidney surgery: Assessment of required intraoperative surface for accurate physical to image space registrations*. Medical Physics, 2008. **35**(9): p. 4251-4261.
6. Tokuda, J., et al., *Integrated navigation and control software system for MRI-guided robotic prostate interventions*. Computerized Medical Imaging and Graphics, 2010. **34**(1): p. 3-8.
7. Bennett, I.C., J. Greenslade, and H. Chiam, *Intraoperative ultrasound-guided excision of nonpalpable breast lesions*. World Journal of Surgery, 2005. **29**(3): p. 369-374.
8. Davis, K.M., et al., *Intraoperative Ultrasound Can Decrease the Re-excision Lumpectomy Rate in Patients with Palpable Breast Cancers*. American Surgeon, 2011. **77**(6): p. 720-725.
9. Henry-Tillman, R., et al., *Intraoperative ultrasound and other techniques to achieve negative margins*. Seminars in Surgical Oncology, 2001. **20**(3): p. 206-213.
10. Leen, E., et al., *Potential value of contrast-enhanced intraoperative ultrasonography during partial hepatectomy for metastases - An essential investigation before resection?* Annals of Surgery, 2006. **243**(2): p. 236-240.
11. Jakimowicz, J.J., *Intraoperative ultrasonography in open and laparoscopic abdominal surgery: an overview*. Surgical Endoscopy and Other Interventional Techniques, 2006. **20**: p. 425-435.
12. Fioule, B., et al., *Additional value of contrast enhanced intraoperative ultrasound for colorectal liver metastases*. European Journal of Radiology, 2008. **67**(1): p. 169-176.
13. Polo, A., et al., *Review of intraoperative imaging and planning techniques in permanent seed prostate brachytherapy*. Radiotherapy and Oncology, 2010. **94**(1): p. 12-23.
14. Sippel, R.S., et al., *Localization of Recurrent Thyroid Cancer Using Intraoperative Ultrasound-Guided Dye Injection*. World Journal of Surgery, 2009. **33**(3): p. 434-439.
15. Criniti, A. and P.C. Lin, *Applications of intraoperative ultrasound in gynecological surgery*. Current Opinion in Obstetrics & Gynecology, 2005. **17**(4): p. 339-342.

16. Berger, M.S., *Ultrasound-guided stereotaxic biopsy using a new apparatus*. J Neurosurg, 1986. **65**(4): p. 550-4.
17. Chandler, W.F., et al., *Intraoperative use of real-time ultrasonography in neurosurgery*. J Neurosurg, 1982. **57**(2): p. 157-63.
18. Gooding, G.A., et al., *Neurosurgical sonography: intraoperative and postoperative imaging of the brain*. AJNR Am J Neuroradiol, 1984. **5**(5): p. 521-5.
19. Koivukangas, J. and P.J. Kelly, *Application of ultrasound imaging to stereotactic brain tumor surgery*. Ann Clin Res, 1986. **18 Suppl 47**: p. 25-32.
20. LeRoux, P.D., et al., *Correlation of intraoperative ultrasound tumor volumes and margins with preoperative computerized tomography scans. An intraoperative method to enhance tumor resection*. J Neurosurg, 1989. **71**(5 Pt 1): p. 691-8.
21. LeRoux, P.D., et al., *A comparison between preoperative magnetic resonance and intraoperative ultrasound tumor volumes and margins*. J Clin Ultrasound, 1994. **22**(1): p. 29-36.
22. Kennedy, J.E., *High-intensity focused ultrasound in the treatment of solid tumours*. Nature Reviews Cancer, 2005. **5**(4): p. 321-327.
23. Rabkin, B.A., V. Zderic, and S. Vaezy, *Hyperecho in ultrasound images of HIFU therapy: Involvement of cavitation*. Ultrasound in Medicine and Biology, 2005. **31**(7): p. 947-956.
24. Wu, F., et al., *"Wide local ablation" of localized breast cancer using high intensity focused ultrasound*. Journal of Surgical Oncology, 2007. **96**(2): p. 130-136.
25. Leslie, T., et al., *High-intensity focused ultrasound treatment of liver tumours: post-treatment MRI correlates well with intra-operative estimates of treatment volume*. British Journal of Radiology, 2012. **85**(1018): p. 1363-1370.
26. Zhou, X.D., et al., *Therapeutic response assessment of high intensity focused ultrasound therapy for uterine fibroid: Utility of contrast-enhanced ultrasonography*. European Journal of Radiology, 2007. **62**(2): p. 289-294.
27. Roberts, W.W., et al., *Pulsed cavitation ultrasound: A noninvasive technology for controlled tissue ablation (histotripsy) in the rabbit kidney*. Journal of Urology, 2006. **175**(2): p. 734-738.
28. Maxwell, A.D., et al., *Noninvasive Thrombolysis Using Pulsed Ultrasound Cavitation Therapy - Histotripsy*. Ultrasound in Medicine and Biology, 2009. **35**(12): p. 1982-1994.
29. Kieran, K., et al., *Refining histotripsy: Defining the parameter space for the creation of nonthermal lesions with high intensity, pulsed focused ultrasound of the in vitro kidney*. Journal of Urology, 2007. **178**(2): p. 672-676.
30. Schade, G.R., et al., *Prostate histotripsy: evaluation of prostatic urethral treatment parameters in a canine model*. Bju International, 2014. **113**(3): p. 498-503.

31. Nightingale, K., *Acoustic Radiation Force Impulse (ARFI) Imaging: A Review*. Current Medical Imaging Reviews, 2011. **7**(4): p. 328-339.
32. Nightingale, K., et al., *Acoustic radiation force impulse imaging: In vivo demonstration of clinical feasibility*. Ultrasound in Medicine and Biology, 2002. **28**(2): p. 227-235.
33. Sarvazyan, A.P., et al., *Shear wave elasticity imaging: A new ultrasonic technology of medical diagnostics*. Ultrasound in Medicine and Biology, 1998. **24**(9): p. 1419-1435.
34. Greenleaf, J.F., M. Fatemi, and M. Insana, *Selected methods for imaging elastic properties of biological tissues*. Annual Review of Biomedical Engineering, 2003. **5**: p. 57-78.
35. Cho, S.H., et al., *Acoustic Radiation Force Impulse Elastography for the Evaluation of Focal Solid Hepatic Lesions: Preliminary Findings*. Ultrasound in Medicine and Biology, 2010. **36**(2): p. 202-208.
36. Muller, M., et al., *Quantitative Viscoelasticity Mapping of Human Liver Using Supersonic Shear Imaging: Preliminary in Vivo Feasibility Study*. Ultrasound in Medicine and Biology, 2009. **35**(2): p. 219-229.
37. Palmeri, M.L., et al., *Noninvasive evaluation of hepatic fibrosis using acoustic radiation force-based shear stiffness in patients with nonalcoholic fatty liver disease*. Journal of Hepatology, 2011. **55**(3): p. 666-672.
38. Friedrich-Rust, M., et al., *Liver Fibrosis in Viral Hepatitis: Noninvasive Assessment with Acoustic Radiation Force Impulse Imaging versus Transient Elastography*. Radiology, 2009. **252**(2): p. 595-604.
39. Hsu, S.J., et al., *In Vivo assessment of myocardial stiffness with acoustic radiation force impulse imaging*. Ultrasound in Medicine and Biology, 2007. **33**(11): p. 1706-1719.
40. Meng, W., et al., *Preliminary Results of Acoustic Radiation Force Impulse (Arfi) Ultrasound Imaging of Breast Lesions*. Ultrasound in Medicine and Biology, 2011. **37**(9): p. 1436-1443.
41. Tozaki, M., S. Isobe, and E. Fukuma, *Preliminary study of ultrasonographic tissue quantification of the breast using the acoustic radiation force impulse (ARFI) technology*. European Journal of Radiology, 2011. **80**(2): p. E182-E187.
42. Palmeri, M.L., et al., *On the Feasibility of Imaging Peripheral Nerves Using Acoustic Radiation Force Impulse Imaging*. Ultrasonic Imaging, 2009. **31**(3): p. 172-182.
43. Zhai, L., et al., *Acoustic Radiation Force Impulse Imaging of Human Prostates: Initial in Vivo Demonstration*. Ultrasound in Medicine and Biology, 2012. **38**(1): p. 50-61.
44. Fahey, B.J., et al., *In Vivo Guidance and Assessment of Liver Radio-Frequency Ablation with Acoustic Radiation Force Elastography*. Ultrasound in Medicine and Biology, 2008. **34**(10): p. 1590-1603.

45. Fahey, B.J., et al., *Acoustic radiation force impulse imaging of myocardial radiofrequency ablation: Initial in vivo results*. *Ieee Transactions on Ultrasonics Ferroelectrics and Frequency Control*, 2005. **52**(4): p. 631-641.
46. Solberg, O.V., et al., *Freehand 3D ultrasound reconstruction algorithms - A review*. *Ultrasound in Medicine and Biology*, 2007. **33**(7): p. 991-1009.
47. Arbel, T., et al., *Automatic non-linear MRI-ultrasound registration for the correction of intra-operative brain deformations*. *Comput Aided Surg*, 2004. **9**(4): p. 123-36.
48. Brendel, B., et al., *Registration of 3D CT and ultrasound datasets of the spine using bone structures*. *Comput Aided Surg*, 2002. **7**(3): p. 146-55.
49. Lindseth, F., et al., *Multimodal image fusion in ultrasound-based neuronavigation: improving overview and interpretation by integrating preoperative MRI with intraoperative 3D ultrasound*. *Comput Aided Surg*, 2003. **8**(2): p. 49-69.
50. Berntsen, E.M., et al., *Functional Magnetic Resonance Imaging and Diffusion Tensor Tractography Incorporated Into an Intraoperative 3-Dimensional Ultrasound-Based Neuronavigation System: Impact on Therapeutic Strategies, Extent of Resection, and Clinical Outcome*. *Neurosurgery*, 2010. **67**(2): p. 251-264.
51. Comeau, R.M., A. Fenster, and T.M. Peters, *Intraoperative US in interactive image-guided neurosurgery*. *Radiographics*, 1998. **18**(4): p. 1019-27.
52. Comeau, R.M., et al., *Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery*. *Med Phys*, 2000. **27**(4): p. 787-800.
53. Gobbi, D.G., R.M. Comeau, and T.M. Peters, *Ultrasound probe tracking for real-time ultrasound/MRI overlay and visualization of brain shift*. *Medical Image Computing and Computer-Assisted Intervention, MICCAI'99, Proceedings*, 1999. **1679**: p. 920-927.
54. Bucholz, R.D., et al., *The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device*. *CVRMED-MRCAS'97*, 1997. **1205**: p. 459-466.
55. Letteboer, M.M., et al., *Brain shift estimation in image-guided neurosurgery using 3-D ultrasound*. *IEEE Trans Biomed Eng*, 2005. **52**(2): p. 268-76.
56. Mercier, L., et al., *New prototype neuronavigation system based on preoperative imaging and intraoperative freehand ultrasound: system description and validation*. *Int J Comput Assist Radiol Surg*, 2011. **6**(4): p. 507-22.
57. Gronningsaeter, A., et al., *SonoWand, an ultrasound-based neuronavigation system*. *Neurosurgery*, 2000. **47**(6): p. 1373-9; discussion 1379-80.
58. Galloway, R.L., R.J. Maciunas, and C.A. Edwards, *Interactive Image-Guided Neurosurgery*. *IEEE Transactions on Biomedical Engineering*, 1992. **39**(12): p. 1226-1231.

59. Peters, T. and K. Cleary, *Image-guided interventions: technology and applications*. 2008: Springer.
60. Atuegwu, N.C. and R.L. Galloway, *Volumetric characterization of the Aurora magnetic tracker system for image-guided transorbital endoscopic procedures*. *Physics in Medicine and Biology*, 2008. **53**(16): p. 4355-4368.
61. Barry, C.D., et al., *Three-dimensional freehand ultrasound: image reconstruction and volume analysis*. *Ultrasound Med Biol*, 1997. **23**(8): p. 1209-24.
62. Prager, R.W., et al., *Rapid calibration for 3-D freehand ultrasound*. *Ultrasound Med Biol*, 1998. **24**(6): p. 855-69.
63. Berg, S., et al., *Dynamic three-dimensional freehand echocardiography using raw digital ultrasound data*. *Ultrasound Med Biol*, 1999. **25**(5): p. 745-53.
64. Leotta, D.F., *An efficient calibration method for freehand 3-D ultrasound imaging systems*. *Ultrasound Med Biol*, 2004. **30**(7): p. 999-1008.
65. Pagoulatos, N., D.R. Haynor, and Y. Kim, *A fast calibration method for 3-D tracking of ultrasound images using a spatial localizer*. *Ultrasound Med Biol*, 2001. **27**(9): p. 1219-29.
66. Rousseau, F., P. Hellier, and C. Barillot, *Robust and automatic calibration method for 3D freehand ultrasound*. *Medical Image Computing and Computer-Assisted Intervention - Miccai 2003, Pt 2*, 2003. **2879**: p. 440-448.
67. Krucker, J., et al., *An electro-magnetically tracked laparoscopic ultrasound for multi-modality minimally invasive surgery*. *International Congress Series*, 2005. **1281**(0): p. 746-751.
68. Bouchet, L.G., et al., *Calibration of three-dimensional ultrasound images for image-guided radiation therapy*. *Phys Med Biol*, 2001. **46**(2): p. 559-77.
69. Lindseth, F., et al., *Probe calibration for freehand 3-D ultrasound*. *Ultrasound Med Biol*, 2003. **29**(11): p. 1607-23.
70. Treece, G.M., et al., *High-definition freehand 3-D ultrasound*. *Ultrasound Med Biol*, 2003. **29**(4): p. 529-46.
71. Boctor, E.M., et al., *Bootstrapped ultrasound calibration*. *Stud Health Technol Inform*, 2006. **119**: p. 61-6.
72. Blackall, J.M., et al., *An image registration approach to automated calibration for freehand 3D ultrasound*. *Medical Image Computing and Computer-Assisted Intervention - Miccai 2000*, 2000. **1935**: p. 462-471.
73. Kowal, J., et al., *On the development and comparative evaluation of an ultrasound B-mode probe calibration method*. *Comput Aided Surg*, 2003. **8**(3): p. 107-19.

74. Muratore, D.M. and R.L. Galloway, *Beam calibration without a phantom for creating a 3-D freehand ultrasound system*. *Ultrasound in Medicine and Biology*, 2001. **27**(11): p. 1557-1566.
75. Sato, Y., et al., *Image guidance of breast cancer surgery using 3-D ultrasound images and augmented reality visualization*. *Ieee Transactions on Medical Imaging*, 1998. **17**(5): p. 681-693.
76. Zhang, W.Y., R.N. Rohling, and D.K. Pai, *Surface extraction with a three-dimensional freehand ultrasound system*. *Ultrasound Med Biol*, 2004. **30**(11): p. 1461-73.
77. Leksell, L., *The stereotaxic method and radiosurgery of the brain*. *Acta Chir Scand*, 1951. **102**(4): p. 316-9.
78. Spiegel, E.A., et al., *Stereotaxic Apparatus for Operations on the Human Brain*. *Science*, 1947. **106**(2754): p. 349-50.
79. Birg, W. and F. Munding, *Direct target point determination for stereotactic brain operations from CT data and the calculation of setting parameters for polar-coordinate stereotactic devices*. *Appl Neurophysiol*, 1982. **45**(4-5): p. 387-95.
80. Goerss, S., et al., *A computed tomographic stereotactic adaptation system*. *Neurosurgery*, 1982. **10**(3): p. 375-9.
81. Leksell, L. and B. Jernberg, *Stereotaxis and tomography. A technical note*. *Acta Neurochir (Wien)*, 1980. **52**(1-2): p. 1-7.
82. Talairach, J. and G. Szikla, *Application of stereotactic concepts to the surgery of epilepsy*. *Acta Neurochir Suppl (Wien)*, 1980. **30**: p. 35-54.
83. Roberts, D.W., et al., *A Frameless Stereotaxic Integration of Computerized Tomographic Imaging and the Operating Microscope*. *Journal of Neurosurgery*, 1986. **65**(4): p. 545-549.
84. Kosugi, Y., et al., *An articulated neurosurgical navigation system using MRI and CT images*. *IEEE Trans Biomed Eng*, 1988. **35**(2): p. 147-52.
85. Manwaring, K.H., M.L. Manwaring, and S.D. Moss, *Magnetic field guided endoscopic dissection through a burr hole may avoid more invasive craniotomies. A preliminary report*. *Acta Neurochir Suppl*, 1994. **61**: p. 34-9.
86. Zamorano, L.J., et al., *Interactive intraoperative localization using an infrared-based system*. *Neurol Res*, 1993. **15**(5): p. 290-8.
87. Hajnal, J.V., D.J. Hawkes, and D.L.G. Hill, *Medical image registration*. 2001, Boca Raton: CRC Press. 382 p., [8] p. of plates.
88. Watanabe, E., et al., *Three-dimensional digitizer (neuronavigator): new equipment for computed tomography-guided stereotaxic surgery*. *Surg Neurol*, 1987. **27**(6): p. 543-7.

89. Barnett, G.H., D.W. Miller, and J. Weisenberger, *Frameless stereotaxy with scalp-applied fiducial markers for brain biopsy procedures: experience in 218 cases*. J Neurosurg, 1999. **91**(4): p. 569-76.
90. Vinas, F.C., et al., *Application accuracy study of a semipermanent fiducial system for frameless stereotaxis*. Comput Aided Surg, 1997. **2**(5): p. 257-63.
91. Maurer, C.R., Jr., et al., *Registration of head volume images using implantable fiducial markers*. IEEE Trans Med Imaging, 1997. **16**(4): p. 447-62.
92. Besl, P.J. and N.D. McKay, *A Method for Registration of 3-D Shapes*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1992. **14**(2): p. 239-256.
93. Audette, M.A., et al., *An integrated range-sensing, segmentation and registration framework for the characterization of intra-surgical brain deformations in image-guided surgery*. Comput. Vis. Image Underst., 2003. **89**(2-3): p. 226-251.
94. Miga, M.I., et al., *Cortical surface registration for image-guided neurosurgery using laser-range scanning*. IEEE Trans Med Imaging, 2003. **22**(8): p. 973-85.
95. Shamir, R.R., et al., *Surface-based facial scan registration in neuronavigation procedures: a clinical study* *Clinical article*. Journal of Neurosurgery, 2009. **111**(6): p. 1201-1206.
96. Cao, A., et al., *Laser range scanning for image-guided neurosurgery: investigation of image-to-physical space registrations*. Med Phys, 2008. **35**(4): p. 1593-605.
97. Sinha, T.K., et al., *A method to track cortical surface deformations using a laser range scanner*. IEEE Trans Med Imaging, 2005. **24**(6): p. 767-81.
98. Sinha, T.K., et al., *Intraoperative cortical surface characterization using laser range scanning: preliminary results*. Neurosurgery, 2006. **59**(4 Suppl 2): p. ONS368-76; discussion ONS376-7.
99. Miller, K., *Constitutive model of brain tissue suitable for finite element analysis of surgical procedures*. J Biomech, 1999. **32**(5): p. 531-7.
100. Miller, K. and K. Chinzei, *Constitutive modelling of brain tissue: experiment and theory*. J Biomech, 1997. **30**(11-12): p. 1115-21.
101. Paulsen, K.D., et al., *A computational model for tracking subsurface tissue deformation during stereotactic neurosurgery*. IEEE Trans Biomed Eng, 1999. **46**(2): p. 213-25.
102. Miga, M.I., et al., *In Vivo Analysis of Heterogeneous Brain Deformation Computations for Model-Updated Image Guidance*. Comput Methods Biomech Biomed Engin, 2000. **3**(2): p. 129-146.
103. Skrinjar, O., A. Nabavi, and J. Duncan, *Model-driven brain shift compensation*. Med Image Anal, 2002. **6**(4): p. 361-73.
104. Skrinjar, O.M. and J.S. Duncan, *Real time 3D brain shift compensation*. Information Processing in Medical Imaging, Proceedings, 1999. **1613**: p. 42-55.

105. Edwards, P.J., et al., *A three-component deformation model for image-guided surgery*. Med Image Anal, 1998. **2**(4): p. 355-67.
106. Biot, M.A., *General Theory of Three-Dimensional Consolidation*. Journal of Applied Physics, 1941. **12**(2): p. 155-164.
107. D'Haese, P.F., et al., *Atlas-based segmentation of the brain for 3-dimensional treatment planning in children with infratentorial ependymoma*. Medical Image Computing and Computer-Assisted Intervention - MICCAI 2003, Pt 2, 2003. **2879**: p. 627-634.
108. Yushkevich, P.A., et al., *User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability*. Neuroimage, 2006. **31**(3): p. 1116-1128.
109. Noble, J.A., *Ultrasound image segmentation and tissue characterization*. Proceedings of the Institution of Mechanical Engineers Part H-Journal of Engineering in Medicine, 2010. **224**(H2): p. 307-316.
110. Noble, J.A. and D. Boukerroui, *Ultrasound image segmentation: a survey*. IEEE Trans Med Imaging, 2006. **25**(8): p. 987-1010.
111. Mortensen, E.N. and W.A. Barrett. *Intelligent scissors for image composition*. in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. 1995: ACM.
112. Lorensen, W.E. and H.E. Cline, *Marching cubes: A high resolution 3D surface construction algorithm*. SIGGRAPH Comput. Graph., 1987. **21**(4): p. 163-169.
113. Sullivan, J.M., G. Charron, and K.D. Paulsen, *A three-dimensional mesh generator for arbitrary multiple material domains*. Finite Elements in Analysis and Design, 1997. **25**(3-4): p. 219-241.
114. Lynch, D.R., *Numerical partial differential equations for environmental scientists and engineers: a first practical course*. 2005: Springer.
115. Weis, J.A., *A novel finite element inverse analysis to assess bone fracture healing*, in *Biomedical Engineering*. 2011, Vanderbilt University: [Nashville, Tenn.]. p. 1 online resource (xiv, 202 p.).
116. Detmer, P.R., et al., *3D ultrasonic image feature localization based on magnetic scanhead tracking: in vitro calibration and validation*. Ultrasound Med Biol, 1994. **20**(9): p. 923-36.
117. Meairs, S., J. Beyer, and M. Hennerici, *Reconstruction and visualization of irregularly sampled three- and four-dimensional ultrasound data for cerebrovascular applications*. Ultrasound Med Biol, 2000. **26**(2): p. 263-72.
118. Barratt, D.C., et al., *Accuracy of an electromagnetic three-dimensional ultrasound system for carotid artery imaging*. Ultrasound Med Biol, 2001. **27**(10): p. 1421-5.
119. Lang, A., et al., *Multi-modal registration of speckle-tracked freehand 3D ultrasound to CT in the lumbar spine*. Med Image Anal, 2012. **16**(3): p. 675-86.

120. Boisvert, J., et al., *An Open-Source Solution for Interactive Acquisition, Processing and Transfer of Interventional Ultrasound Images*. The MIDAS Journal - Systems and Architectures for Computer Assisted Interventions, 2008.
121. Pace, D., et al., *An open-source real-time ultrasound reconstruction system for four-dimensional imaging of moving organs*. Med Image Comput Comput Assist Interv. MICCAI 2009, 2009.
122. Shoemake, K., *Animating rotation with quaternion curves*. SIGGRAPH Comput. Graph., 1985. **19**(3): p. 245-254.
123. Mercier, L., et al., *A review of calibration techniques for freehand 3-D ultrasound systems*. Ultrasound Med Biol, 2005. **31**(4): p. 449-71.
124. Trobaugh, J.W., D.J. Trobaugh, and W.D. Richard, *Three-dimensional imaging with stereotactic ultrasonography*. Comput Med Imaging Graph, 1994. **18**(5): p. 315-23.
125. Trobaugh, J.W., et al., *Frameless stereotactic ultrasonography: method and applications*. Comput Med Imaging Graph, 1994. **18**(4): p. 235-46.
126. Brown, R.A., *A stereotactic head frame for use with CT body scanners*. Invest Radiol, 1979. **14**(4): p. 300-4.
127. Canny, J., *A computational approach to edge detection*. IEEE Trans Pattern Anal Mach Intell, 1986. **8**(6): p. 679-98.
128. Arun, K.S., T.S. Huang, and S.D. Blostein, *Least-Squares Fitting of 2 3-D Point Sets*. Ieee Transactions on Pattern Analysis and Machine Intelligence, 1987. **9**(5): p. 699-700.
129. Horn, B.K.P., H.M. Hilden, and S. Negahdaripour, *Closed-Form Solution of Absolute Orientation Using Orthonormal Matrices*. Journal of the Optical Society of America a-Optics Image Science and Vision, 1988. **5**(7): p. 1127-1135.
130. Martin, K. and D. Spinks, *Measurement of the speed of sound in ethanol/water mixtures*. Ultrasound Med Biol, 2001. **27**(2): p. 289-91.
131. Prager, R.W., A. Gee, and L. Berman, *Stradx: real-time acquisition and visualization of freehand three-dimensional ultrasound*. Med Image Anal, 1999. **3**(2): p. 129-40.
132. Clements, L.W., et al., *Robust surface registration using salient anatomical features for image-guided liver surgery: Algorithm and validation*. Medical Physics, 2008. **35**(6): p. 2528-2540.
133. Konofagou, E. and J. Ophir, *A new elastographic method for estimation and imaging of lateral displacements, lateral strains, corrected axial strains and Poisson's ratios in tissues*. Ultrasound Med Biol, 1998. **24**(8): p. 1183-99.
134. Chen, L., et al., *A hybrid displacement estimation method for ultrasonic elasticity imaging*. IEEE Trans Ultrason Ferroelectr Freq Control, 2010. **57**(4): p. 866-82.

135. Selbekk, T., J. Bang, and G. Unsgaard, *Strain processing of intraoperative ultrasound images of brain tumours: initial results*. *Ultrasound Med Biol*, 2005. **31**(1): p. 45-51.
136. Selbekk, T., et al., *Tissue Motion and Strain in the Human Brain Assessed by Intraoperative Ultrasound in Glioma Patients*. *Ultrasound in Medicine and Biology*, 2010. **36**(1): p. 2-10.
137. Cabot, R.C., *A Note on the Application of the Hilbert Transform to Time-Delay Estimation*. *Ieee Transactions on Acoustics Speech and Signal Processing*, 1981. **29**(3): p. 607-609.
138. Surry, K.J., et al., *Poly(vinyl alcohol) cryogel phantoms for use in ultrasound and MR imaging*. *Phys Med Biol*, 2004. **49**(24): p. 5529-46.
139. Chu, K.C. and B.K. Rutt, *Polyvinyl alcohol cryogel: an ideal phantom material for MR studies of arterial flow and elasticity*. *Magn Reson Med*, 1997. **37**(2): p. 314-9.
140. Fromageau, J., et al., *Estimation of polyvinyl alcohol cryogel mechanical properties with four ultrasound elastography methods and comparison with gold standard testings*. *IEEE Trans Ultrason Ferroelectr Freq Control*, 2007. **54**(3): p. 498-509.
141. Simpson, A.L., et al., *Comparison Study of Intraoperative Surface Acquisition Methods for Surgical Navigation*. *Ieee Transactions on Biomedical Engineering*, 2013. **60**(4): p. 1090-1099.
142. Rucker, D.C., et al., *A Mechanics-Based Nonrigid Registration Method for Liver Surgery Using Sparse Intraoperative Data*. *Medical Imaging, IEEE Transactions on*, 2013. **33**(1): p. 147-158.
143. Chen, I., et al., *Integrating Retraction Modeling Into an Atlas-Based Framework for Brain Shift Prediction*. *Ieee Transactions on Biomedical Engineering*, 2013. **60**(12): p. 3494-3504.
144. Shannon, M.J., et al. *Initial study of breast tissue retraction toward image guided breast surgery*. in *SPIE Medical Imaging*. 2012: SPIE.
145. Nabavi, A., et al., *Serial intraoperative magnetic resonance imaging of brain shift*. *Neurosurgery*, 2001. **48**(4): p. 787-97; discussion 797-8.
146. Roberts, D.W., et al., *Intraoperative brain shift and deformation: a quantitative analysis of cortical displacement in 28 cases*. *Neurosurgery*, 1998. **43**(4): p. 749-58; discussion 758-60.
147. Hill, D.L.G., et al., *Measurement of intraoperative brain surface deformation under a craniotomy*. *Neurosurgery*, 1998. **43**(3): p. 514-526.
148. Nauta, H.J., *Error assessment during "image guided" and "imaging interactive" stereotactic surgery*. *Comput Med Imaging Graph*, 1994. **18**(4): p. 279-87.
149. Black, P.M., et al., *Development and implementation of intraoperative magnetic resonance imaging and its neurosurgical applications*. *Neurosurgery*, 1997. **41**(4): p. 831-42; discussion 842-5.
150. Butler, W.E., et al., *A mobile computed tomographic scanner with intraoperative and intensive care unit applications*. *Neurosurgery*, 1998. **42**(6): p. 1304-10; discussion 1310-1.

151. Lunsford, L.D. and A.J. Martinez, *Stereotactic exploration of the brain in the era of computed tomography*. Surg Neurol, 1984. **22**(3): p. 222-30.
152. Henderson, J.M., K.R. Smith, and R.D. Bucholz, *An Accurate and Ergonomic Method of Registration for Image-Guided Neurosurgery*. Computerized Medical Imaging and Graphics, 1994. **18**(4): p. 273-277.
153. Raabe, A., et al., *Laser surface scanning for patient registration in intracranial image-guided surgery*. Neurosurgery, 2002. **50**(4): p. 797-801; discussion 802-3.
154. Roberts, D.W., et al., *Intraoperatively updated neuroimaging using brain modeling and sparse data*. Neurosurgery, 1999. **45**(5): p. 1199-1206.
155. Isheil, A., et al., *Systematic error correction of a 3D laser scanning measurement device*. Optics and Lasers in Engineering. **49**(1): p. 16-24.
156. Sun, H., et al., *Cortical surface tracking using a stereoscopic operating microscope*. Neurosurgery, 2005. **56**(1 Suppl): p. 86-97; discussion 86-97.
157. Van Gestel, N., et al., *A performance evaluation test for laser line scanners on CMMs*. Optics and Lasers in Engineering. **47**(3-4): p. 336-342.
158. Commer, P., et al., *Construction and testing of a computer-based intraoral laser scanner for determining tooth positions*. Medical Engineering & Physics, 2000. **22**(9): p. 625-635.
159. Cash, D.M., et al., *Incorporation of a laser range scanner into image-guided liver surgery: surface acquisition, registration, and tracking*. Med Phys, 2003. **30**(7): p. 1671-82.
160. Marmulla, R., et al., *Soft tissue scanning for patient registration in image-guided surgery*. Comput Aided Surg, 2003. **8**(2): p. 70-81.
161. Meehan, M., M. Teschner, and S. Girod, *Three-dimensional simulation and prediction of craniofacial surgery*. Orthod Craniofac Res, 2003. **6 Suppl 1**: p. 102-7.
162. Cash, D.M., et al., *Compensating for intraoperative soft-tissue deformations using incomplete surface data and finite elements*. IEEE Trans Med Imaging, 2005. **24**(11): p. 1479-91.
163. Ding, S.I., et al., *Semiautomatic Registration of Pre- and Postbrain Tumor Resection Laser Range Data: Method and Validation*. Ieee Transactions on Biomedical Engineering, 2009. **56**(3): p. 770-780.
164. Dumpuri, P., et al., *Model-updated image-guided liver surgery: preliminary results using surface characterization*. Prog Biophys Mol Biol, 2010. **103**(2-3): p. 197-207.
165. Artignan, X., et al., *Online ultrasound image guidance for radiotherapy of prostate cancer: impact of image acquisition on prostate displacement*. Int J Radiat Oncol Biol Phys, 2004. **59**(2): p. 595-601.

166. Xiao, G., et al., *Nonrigid registration of 3-D free-hand ultrasound images of the breast*. IEEE Trans Med Imaging, 2002. **21**(4): p. 405-12.
167. Hsu, P.W., et al., *Real-time freehand 3D ultrasound calibration*. Ultrasound Med Biol, 2008. **34**(2): p. 239-51.
168. Hsu, P.W., et al., *Comparison of freehand 3-D ultrasound calibration techniques using a stylus*. Ultrasound Med Biol, 2008. **34**(10): p. 1610-21.
169. King, A.P., et al., *Bayesian estimation of intra-operative deformation for image-guided surgery using 3-D ultrasound*. Medical Image Computing and Computer-Assisted Intervention - Miccai 2000, 2000. **1935**: p. 588-597.
170. Treece, G.M., A.H. Gee, and R.W. Prager, *RF and amplitude-based probe pressure correction for 3D ultrasound*. Ultrasound Med Biol, 2005. **31**(4): p. 493-503.
171. Burcher, M.R., H. Lianghao, and J.A. Noble. *Deformation correction in ultrasound images using contact force measurements*. in *Mathematical Methods in Biomedical Image Analysis, 2001. MMBIA 2001. IEEE Workshop on*. 2001.
172. Sun, S.-Y., B.W. Anthony, and M.W. Gilbertson. *Trajectory-based deformation correction in ultrasound images*. in *SPIE Medical Imaging 2010: Ultrasonic Imaging, Tomography, and Therapy*. 2010. San Diego, CA: SPIE.
173. Pfeiffer, T.S., et al., *Design and evaluation of an optically-tracked single-CCD laser range scanner*. Med Phys, 2012. **39**(2): p. 636-42.
174. Bassar, P.J., *Interstitial pressure, volume, and flow during infusion into brain tissue*. Microvasc Res, 1992. **44**(2): p. 143-65.
175. Nagashima, T., T. Shirakuni, and S.I. Rapoport, *A two-dimensional, finite element analysis of vasogenic brain edema*. Neurologia medico-chirurgica, 1990. **30**(1): p. 1.
176. Coffey, A.M., et al., *Toward a preoperative planning tool for brain tumor resection therapies*. International Journal of Computer Assisted Radiology and Surgery, 2013. **8**(1): p. 87-97.
177. Chen, I., et al., *Intraoperative Brain Shift Compensation: Accounting for Dural Septa*. IEEE Transactions on Biomedical Engineering, 2011. **58**(3): p. 499-508.
178. Miga, M.I., et al., *Model-updated image guidance: initial clinical experiences with gravity-induced brain deformation*. IEEE Trans Med Imaging, 1999. **18**(10): p. 866-74.
179. Dubuisson, M.P. and A.K. Jain. *A modified Hausdorff distance for object matching*. in *Pattern Recognition, 1994. Vol. 1 - Conference A: Computer Vision & Image Processing., Proceedings of the 12th IAPR International Conference on*. 1994.
180. Dumpuri, P., et al., *An atlas-based method to compensate for brain shift: preliminary results*. Med Image Anal, 2007. **11**(2): p. 128-45.

181. Ding, S.Y., et al., *Tracking of Vessels in Intra-Operative Microscope Video Sequences for Cortical Displacement Estimation*. IEEE Transactions on Biomedical Engineering, 2011. **58**(7): p. 1985-1993.
182. Simpson, A.L., et al., *Uncertainty propagation and analysis of image-guided surgery*. Proceedings of SPIE medical imaging: visualization, image-guided procedures, and modeling, 2011. **7964**: p. 79640H-7.
183. Simpson, A.L., et al. *A framework for measuring TRE at the tip of an optically tracked pointing stylus*. in *SPIE Medical Imaging*. 2013: International Society for Optics and Photonics.
184. Wiles, A.D., D.G. Thompson, and D.D. Frantz, *Accuracy assessment and interpretation for optical tracking systems*. 2004: p. 421-432.
185. Glossop, N.D., *Advantages of optical compared with electromagnetic tracking*. J Bone Joint Surg Am, 2009. **91 Suppl 1**: p. 23-8.
186. Nagashima, T., T. Shirakuni, and S.I. Rapoport, *A two-dimensional, finite element analysis of vasogenic brain edema*. Neurol Med Chir (Tokyo), 1990. **30**(1): p. 1-9.
187. Sun, S.-Y., *Deformation correction in ultrasound imaging in an elastography framework*, in *Department of Electrical Engineering and Computer Science*. 2010, Massachusetts Institute of Technology: Cambridge. p. 80 p.
188. Pfeiffer, T.S., et al., *Model-Based Correction of Tissue Compression for Tracked Ultrasound in Soft Tissue Image-Guided Surgery*. Ultrasound Med Biol, 2014.
189. Lange, T., et al., *3D ultrasound-CT registration of the liver using combined landmark-intensity information*. International Journal of Computer Assisted Radiology and Surgery, 2009. **4**(1): p. 79-88.
190. Bro-Nielsen, M., *Finite element modeling in surgery simulation*. Proceedings of the IEEE, 1998. **86**(3): p. 490-503.
191. Lange, T., et al., *Augmenting intraoperative 3D ultrasound with preoperative models for navigation in liver surgery*. Medical Image Computing and Computer-Assisted Intervention - Miccai 2004, Pt 2, Proceedings, 2004. **3217**: p. 534-541.
192. Burger, P.C., et al., *Topographic anatomy and CT correlations in the untreated glioblastoma multiforme*. J Neurosurg, 1988. **68**(5): p. 698-704.
193. Van Den Hauwe, L., et al., *Postmortem Mri of the Brain with Neuropathological Correlation*. Neuroradiology, 1995. **37**(5): p. 343-349.
194. Burns-Cox, N., et al., *Changes in collagen metabolism in prostate cancer: a host response that may alter progression*. J Urol, 2001. **166**(5): p. 1698-701.
195. Lee, H., et al., *Phagocytosis of collagen by fibroblasts and invasive cancer cells is mediated by MT1-MMP*. Biochem Soc Trans, 2007. **35**(Pt 4): p. 704-6.

196. Doyley, M.M., et al., *A freehand elastographic imaging approach for clinical breast imaging: system development and performance evaluation*. *Ultrasound Med Biol*, 2001. **27**(10): p. 1347-57.
197. Konofagou, E.E., T. Harrigan, and J. Ophir, *Shear strain estimation and lesion mobility assessment in elastography*. *Ultrasonics*, 2000. **38**(1-8): p. 400-4.
198. Chakraborty, A., J.C. Bamber, and N.L. Dorward, *Preliminary investigation into the use of ultrasound elastography during brain tumour resection*. *Ultrasound*, 2012. **20**(1): p. 33-40.
199. Scholz, M., et al., *Current status of intraoperative real-time vibrography in neurosurgery*. *Ultraschall Med*, 2007. **28**(5): p. 493-7.
200. Rivaz, H., et al., *Tracked regularized ultrasound elastography for targeting breast radiotherapy*. *Med Image Comput Comput Assist Interv*, 2009. **12**(Pt 1): p. 507-15.
201. Dijkstra, E.W., *A note on two problems in connexion with graphs*. *Numerische Mathematik*, 1959. **1**(1): p. 269-271.
202. Okabe, A., B.N. Boots, and K.O. Sugihara, *Spatial tessellations : concepts and applications of Voronoi diagrams*. *Wiley series in probability and mathematical statistics*. 1992, Chichester, England ; New York: Wiley. x, 532 p.
203. Miga, M.I., et al., *Updated neuroimaging using intraoperative brain modeling and sparse data*. *Stereotact Funct Neurosurg*, 1999. **72**(2-4): p. 103-6.
204. Ding, S., et al., *Estimation of intra-operative brain shift using a tracked laser range scanner*. *Conf Proc IEEE Eng Med Biol Soc*, 2007. **2007**: p. 848-51.
205. Caselles, V., R. Kimmel, and G. Sapiro, *Geodesic active contours*. *International Journal of Computer Vision*, 1997. **22**(1): p. 61-79.
206. Sethian, J.A., *Level set methods : evolving interfaces in geometry, fluid mechanics, computer vision, and materials science*. *Cambridge monographs on applied and computational mathematics 3*. 1996, Cambridge: Cambridge University Press. xviii, 218 p.
207. Ibanez, L., et al., *The ITK Software Guide*. Second ed. 2005: Kitware, Inc.
208. Leventon, M.E., W.E.L. Grimson, and O. Faugeras. *Statistical shape influence in geodesic active contours*. in *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*. 2000.
209. Conley, R.H., et al. *Image to physical space registration of supine breast MRI for image guided breast surgery*. in *SPIE Medical Imaging*. 2014.
210. Chakraborty, A., et al., *Intra-operative ultrasound elastography and registered magnetic resonance imaging of brain tumors: a feasibility study*. *Ultrasound*, 2006. **14**(1): p. 43-49.

211. Ou, J.J., et al., *Evaluation of 3D modality-independent elastography for breast imaging: a simulation study*. Phys Med Biol, 2008. **53**(1): p. 147-63.
212. Miga, M.I., M.P. Rothney, and J.J. Ou, *Modality independent elastography (MIE): Potential applications in dermoscopy*. Medical Physics, 2005. **32**(5): p. 1308-1320.
213. Washington, C.W. and M.I. Miga, *Modality independent elastography (MIE): A new approach to elasticity imaging*. IEEE Transactions on Medical Imaging, 2004. **23**(9): p. 1117-1128.
214. Miga, M.I., *A new approach to elastography using mutual information and finite elements*. Physics in Medicine and Biology, 2003. **48**(4): p. 467-480.
215. Pfeiffer, T.S., et al., *Automatic Generation of Boundary Conditions Using Demons Nonrigid Image Registration for Use in 3-D Modality-Independent Elastography*. IEEE Transactions on Biomedical Engineering, 2011. **58**(9): p. 2607-2616.
216. Weis, J.A., et al. *Validation and reproducibility assessment of modality independent elastography in a pre-clinical model of breast cancer*. in *SPIE Medical Imaging*. 2014: International Society for Optics and Photonics.
217. Kim, D.K., et al. *Utilizing a reference material for assessing absolute tumor mechanical properties in modality independent elastography*. in *SPIE Medical Imaging*. 2014: International Society for Optics and Photonics.
218. DeLorenzo, C., et al., *Image-Guided Intraoperative Cortical Deformation Recovery Using Game Theory: Application to Neocortical Epilepsy Surgery*. IEEE Transactions on Medical Imaging, 2010. **29**(2): p. 322-338.
219. Kumar, A.N., et al. *Phantom-based comparison of the accuracy of point clouds extracted from stereo cameras and laser range scanner*. in *SPIE Medical Imaging*. 2013: International Society for Optics and Photonics.
220. Bilgen, M., *Target detectability in acoustic elastography*. IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control, 1999. **46**(5): p. 1128-1133.
221. Dooley, M.M., P.M. Meaney, and J.C. Bamber, *Evaluation of an iterative reconstruction method for quantitative elastography*. Phys Med Biol, 2000. **45**(6): p. 1521-40.
222. Ophir, J., et al., *Elastography: a quantitative method for imaging the elasticity of biological tissues*. Ultrason Imaging, 1991. **13**(2): p. 111-34.
223. Muthupillai, R., et al., *Magnetic resonance elastography by direct visualization of propagating acoustic strain waves*. Science, 1995. **269**(5232): p. 1854-7.
224. Miga, M.I., *A new approach to elastography using mutual information and finite elements*. Phys Med Biol, 2003. **48**(4): p. 467-80.
225. Shi, P. and H. Liu, *Stochastic finite element framework for simultaneous estimation of cardiac kinematic functions and material parameters*. Med Image Anal, 2003. **7**(4): p. 445-64.

226. Goshtasby, A., *Registration of Images with Geometric Distortions*. IEEE Transactions on Geoscience and Remote Sensing, 1988. **26**(1): p. 60-64.
227. Ong, R.E., J.J. Ou, and M.I. Miga, *Non-rigid registration of breast surfaces using the laplace and diffusion equations*. Biomed Eng Online, 2010. **9**(1): p. 8.
228. Papademetris, X., et al., *Estimation of 3-D left ventricular deformation from medical images using biomechanical models*. IEEE Trans Med Imaging, 2002. **21**(7): p. 786-800.
229. Tanner, C., et al., *Factors influencing the accuracy of biomechanical breast models*. Med Phys, 2006. **33**(6): p. 1758-69.
230. Rueckert, D., et al., *Nonrigid registration using free-form deformations: Application to breast MR images*. IEEE Transactions on Medical Imaging, 1999. **18**(8): p. 712-721.
231. Crum, W.R., C. Tanner, and D.J. Hawkes, *Anisotropic multi-scale fluid registration: evaluation in magnetic resonance breast imaging*. Phys Med Biol, 2005. **50**(21): p. 5153-74.
232. Froh, M.S., et al., *Piecewise-quadrilateral registration by optical flow--applications in contrast-enhanced MR imaging of the breast*. Med Image Comput Comput Assist Interv, 2006. **9**(Pt 2): p. 686-93.
233. Thirion, J.P., *Image matching as a diffusion process: an analogy with Maxwell's demons*. Med Image Anal, 1998. **2**(3): p. 243-60.
234. Pennec, X., P. Cachier, and N. Ayache, *Understanding the "Demon's Algorithm": 3D non-rigid registration by gradient descent*, in MICCAI'99, C. Taylor and A. Colchester, Editors. 1999, Springer-Verlag Berlin Heidelberg: Cambridge, UK. p. 597-606.
235. Cachier, P., et al., *Iconic feature based nonrigid registration: the PASHA algorithm*. Computer Vision and Image Understanding, 2003. **89**(2-3): p. 272-298.
236. Modersitzki, J., *Numerical methods for image registration*. Numerical mathematics and scientific computation. 2004, Oxford ; New York: Oxford University Press. x, 199 p.
237. Vercauteren, T., et al., *Non-parametric diffeomorphic image registration with the demons algorithm*. Med Image Comput Comput Assist Interv, 2007. **10**(Pt 2): p. 319-26.
238. Vercauteren, T., et al., *Symmetric log-domain diffeomorphic Registration: a demons-based approach*. Med Image Comput Comput Assist Interv, 2008. **11**(Pt 1): p. 754-61.
239. Vercauteren, T., et al., *Diffeomorphic demons: Efficient non-parametric image registration*. Neuroimage, 2009. **45**(1): p. S61-S72.
240. Cahill, N.D., J.A. Noble, and D.J. Hawkes, *A Demons algorithm for image registration with locally adaptive regularization*. Med Image Comput Comput Assist Interv, 2009. **12**(Pt 1): p. 574-81.

241. Mansi, T., et al., *LogDemons revisited: consistent regularisation and incompressibility constraint for soft tissue tracking in medical images*. Med Image Comput Comput Assist Interv. **13**(Pt 2): p. 652-9.
242. Miga, M.I., M.P. Rothney, and J.J. Ou, *Modality independent elastography (MIE): potential applications in dermoscopy*. Med Phys, 2005. **32**(5): p. 1308-20.
243. Joachimowicz, N., C. Pichot, and J.P. Hugonin, *Inverse Scattering - an Iterative Numerical-Method for Electromagnetic Imaging*. IEEE Transactions on Antennas and Propagation, 1991. **39**(12): p. 1742-1752.
244. Yoo, T.S., et al., *Engineering and algorithm design for an image processing Api: a technical report on ITK--the Insight Toolkit*. Stud Health Technol Inform, 2002. **85**: p. 586-92.
245. Broit, C., *Optimal registration of deformed images*. 1981, Graduate School of Arts and Sciences, University of Pennsylvania, 1981.: [Philadelphia],. p. vi, 151 l.
246. Cachier, P. and N. Ayache, *Isotropic energies, filters and splines for vector field regularization*. Journal of Mathematical Imaging and Vision, 2004. **20**(3): p. 251-265.
247. Krouskop, T.A., et al., *Elastic moduli of breast and prostate tissues under compression*. Ultrason Imaging, 1998. **20**(4): p. 260-74.
248. Ou, J.J., *Development of modality-independent elastography as a method of breast cancer detection*, in *Department of Biomedical Engineering*. 2008, Vanderbilt University: Nashville, TN.
249. Ong, R.E., J.J. Ou, and M.I. Miga, *Non-rigid registration of breast surfaces using the laplace and diffusion equations*. Biomedical Engineering Online, 2010. **9**.
250. Ou, J.J., *Development of modality-independent elastography as a method of breast cancer detection*. 2008. p. xiv, 169 p. (1 file).

APPENDIX

A. Monitoring surgical resection of tumors with ultrasound strain imaging

A.1 Introduction and Significance of Study

There are several ultrasound-based imaging modalities, such as strain imaging, ARFI, and SWEI, which offer measurements of tissue stiffness. Although the clinical importance of stiffness as a biomarker for disease extent is well recognized, these modalities have not thus far been widely utilized in a guidance role for surgical resection of interventional targets. This study presents an initial feasibility assessment of the ability of strain imaging to consistently localize a phantom lesion during resection even with the presence of a resection cavity and irrigation with fluids. This work was presented at the International Tissue Elasticity Conference in 2012.

Appearing in: T. S. Pheiffer, B. C. Byram, M. I. Miga. Monitoring surgical resection of tumors with ultrasound strain imaging. In: Clinical and Animal Applications, ITEC Deauville, France, 2012.

A.2 Background

Resection of tumors is often performed by the surgeon using tactile sensory information to distinguish between normal and abnormal tissue. Ultrasound strain imaging has potential to supplement conventional guidance methods with quantitative information about tissue stiffness at depth. It has been suggested that strain imaging may be capable of distinguishing tumor from normal tissue during surgery [136, 198, 199]. With respect to diagnostic lesion inspection, localization with strain imaging of a potential surgical target is well understood. In this work, we assess the efficacy of this modality to monitor resection.

A.3 Aims

The aim of this work is to demonstrate the feasibility of using ultrasound strain imaging to monitor a tumor remnant during surgical resection.

A.4 Methods

A phantom was constructed of tissue-mimicking polyvinyl alcohol gel with graphite scatterers and a fabric sphere to serve as the target lesion. The tumor was incrementally resected in three stages, with larger amounts excised at each stage until complete removal was accomplished. Strain imaging was performed of the tumor remnant at each stage of resection, and the tumor cavity was irrigated with water to eliminate air pockets introduced by the resection process. An Acuson Antares ultrasound machine (Siemens, Munich, Germany) was used with a VFX13-5 probe at a frequency of 11.4 MHz. Strain images were generated using the eSie Touch Elasticity Imaging software on the ultrasound machine. All imaging was conducted freehand with the probe in approximately the same location lateral to the resection site.

A.5 Results

The strain images show the tumor mass with clear contrast against the bulk phantom material prior to resection. After resection of approximately one third of the tumor volume, the lesion still appears in the strain images with a corresponding reduction in image slice cross sectional area. After resection of another third of the tumor, the tumor mass still clearly appears in the strain images but with a more noticeable decrease in cross sectional area. Following the complete resection of the entire tumor mass, strain imaging no longer showed a region of low strain in the resection cavity.

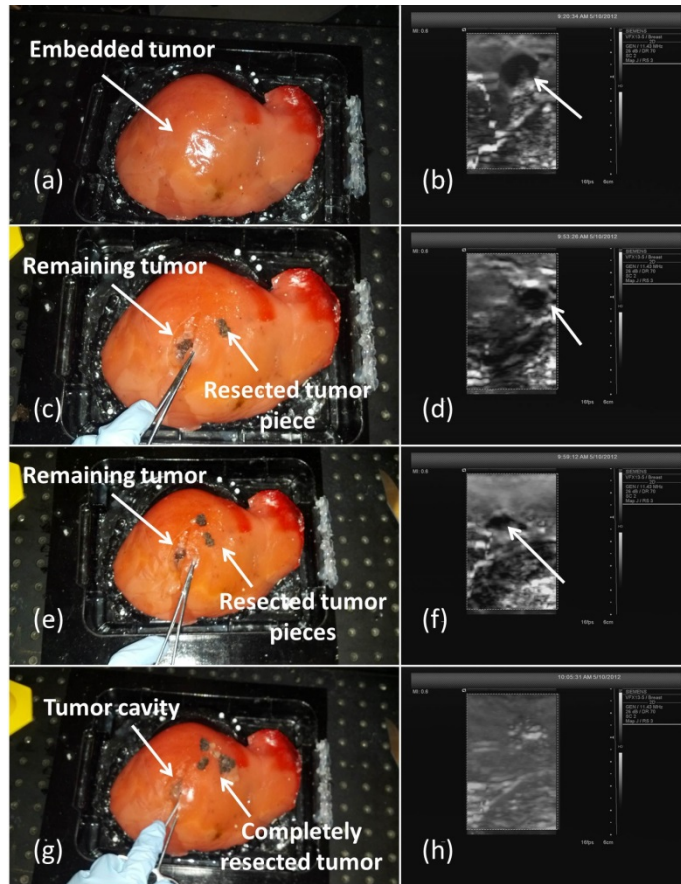


Figure 53. Effects of tumor resection on strain imaging in a cryogel phantom. Boxes (a), (c), (e), and (g) show increasing amounts of tumor resection, from no resection to total resection. Boxes (b), (d), (f), and (h) show the corresponding strain images of the tumor remnant.

A.6 Conclusions

The phantom study performed shows the feasibility of using ultrasound strain imaging as a tool for monitoring surgical resection of lesions. It was shown that the excision of tumor mass and the creation of a resection cavity did not obstruct the creation of strain images. In addition, the lesion area in the images decreased in correlation with decreasing remnant tumor volume. Ultrasound strain imaging shows promise as a surgical localization method and awaits further studies within the clinical environment.

Acknowledgements: This work is funded by the National Institutes of Health: grant R01 NS049251 of the National Institute for Neurological Disorders and Stroke.

B. Automatic Generation of Boundary Conditions Using Demons Non-Rigid Image Registration for Use in 3D Modality-Independent Elastography

B.1 Introduction and Significance of Study

This study demonstrates a novel approach to solving a workflow problem in modality-independent elastography (MIE), which utilizes a finite element model approach to estimating material properties in tissue. The general framework of MIE is to acquire pre- and post-deformation images of tissue subjected to a mechanical loading condition, and then to optimize the material property distribution in the tissue model until the model-prediction of the subsurface tissue movement matches the observed tissue movement in the images. A necessary input to the finite element model, however, are boundary conditions which accurately reflect the physical loading conditions applied to the tissue. These boundary conditions are difficult to measure in practice, and previously required a large amount of effort to generate using surface fiducials. In this work, the boundary condition generation step was automated through the use of the demons deformable registration algorithm to detect boundary displacements using the image volumes themselves. This brings the MIE method closer to clinical deployment, in which the ready availability of tissue material properties has implications in a variety of domains, including the model-based ultrasound corrections presented in this dissertation. This work was published in IEEE Transactions on Biomedical Engineering in 2011.

Appearing in: T. S. Pheiffer, J. J. Ou, R. E. Ong and M. I. Miga. Automatic Generation of Boundary Conditions Using Demons Nonrigid Image Registration for Use in 3-D Modality-Independent Elastography. IEEE Transactions on Biomedical Engineering 58, 2607-2616 (2011).

B.2 Abstract

Modality-independent elastography (MIE) is a method of elastography that reconstructs the elastic properties of tissue using images acquired under different loading conditions and a biomechanical model. Boundary conditions are a critical input to the algorithm, and are often determined by time-consuming point correspondence methods requiring manual user input. This study presents a novel method of automatically generating boundary conditions by non-rigidly registering two image sets with a demons diffusion-based registration algorithm. The use of this method was successfully performed *in silico* using magnetic resonance and X-ray computed tomography image data with known boundary conditions. These preliminary results produced boundary conditions with accuracy of up to 80% compared to the known conditions. Demons-based boundary conditions were utilized within a 3D MIE reconstruction to determine an elasticity contrast ratio between tumor and normal tissue. Two phantom experiments were then conducted to further test the accuracy of the demons boundary conditions and the MIE reconstruction arising from the use of these conditions. Preliminary results show a reasonable characterization of the material properties on this first attempt and a significant improvement in the automation level and viability of the method.

B.3 Introduction

An imaging methodology that utilizes the mechanical properties of tissue is known as elastography. Elastography employs a combination of image processing and measurements of the physical deformation of the tissue to create a representation of the mechanical strength of structures inside an organ [220, 221]. The overall principle behind elastography for use in cancer imaging is that regional changes in tissue architecture resulting from the manifestation of disease result in detectable changes in mechanical properties. For example, breast cancers have been widely recognized in the medical community as much firmer to the touch than the surrounding soft tissue. The biological basis for this effect is due to changes in tissue composition, such as varied

expression of collagen and greater numbers of fibroblasts [194, 195]. The exploitation of a contrast mechanism based on elastic properties may have considerable potential to characterize disease states.

Several kinds of elastography exist, such as ultrasound elastography (USE) and magnetic resonance elastography (MRE) which have already shown promise in diagnosing solid lesions in breast tissue and other physiological locations. The first introduction of USE demonstrated that images from A-line ultrasound could provide axial strain estimates [222]. Elastography has also been applied within the MR imaging domain whereby motion-sensitized gradient sequences were used to visualize and quantify strain wave propagation in media [223]. A relatively new method known as modality-independent elastography (MIE) has recently shown potential for supplementing other imaging modalities such as MR and CT for detection of solid tumors in soft tissue [224]. MIE has the benefit of being flexible with regard to its inputs, and unlike USE and MRE, it is not reliant on a particular imaging modality. MIE involves imaging a tissue of interest before and after compression, and then applying a finite element (FE) soft-tissue model within a nonlinear optimization framework in order to determine the elastic properties of the tissue. The group of Shi *et al.* also has used FE biomechanical models to estimate material properties of the heart using *a priori* information [225]. A requirement of the MIE method is that appropriate boundary conditions be designated for use in the biomechanical model. Generation of accurate boundary conditions is problematic because soft-tissue organs are non-rigid structures, which invalidates the use of standard rigid registration techniques. Techniques which have addressed this issue in the past have required a significant amount of user interaction. The goal of this work is to develop and validate a method of generating boundary conditions automatically by registering organ surfaces before and after mechanical loading. While breast-like phantoms are used as a demonstration of the method, the intention is to make a more broadly applicable method that may be applied to other organs, such as the brain, liver, or lung.

The previous gold standard in generating boundary conditions for MIE has been feature-based registration methods [211]. Conventionally this entails employing point correspondence methods facilitated by attached fiducials and assisted by thin-plate spline (TPS) interpolation [226] to create boundary conditions that non-rigidly map the pre-deformed organ surface to the post-deformed organ surface. This registration process requires the tedious task of applying and subsequently localizing numerous surface markers within the image space, determining point correspondence, creating a thin-plate spline interpolation, and finally calculating a set of Dirichlet boundary conditions for use in the MIE method. Initial attempts to reduce the complexity and level of user interaction have focused on the use of two energy minimization techniques [227]. These techniques relied upon partial differential equation (PDE) solutions of Laplace's equation or the diffusion equation, respectively, across the surface of the organ geometry in the pre- and post-deformed states. Like-valued isocontours from the solutions on each surface (i.e. pre-deformed, and post-deformed) act as 'virtual' fiducials to assist in correspondence using a symmetric closest point approach [228]. Dirichlet boundary conditions are generated after the assigned correspondence is determined and this completes the required input for the MIE algorithm. Both methodologies required conditions to be manually specified to various regions of the mesh. While the results presented by [227] indicated better performance via the Laplacian method, the diffusion method did not require the difficult task of assigning a boundary condition to the chest wall in both pre- and post-deformed mesh domains. These methods, as well as the TPS method, will be compared to the intensity-based approach in this paper.

While the above PDE-based methods represented an improvement in automation over the TPS method for generating boundary conditions for the MIE algorithm, the ideal boundary condition method would be both fully automated and require no fiducials. There is precedent for using non-rigid image registration to generate boundary conditions for FE tissue volume models, as presented by Tanner *et al.* [229]. There is a wide variety of non-rigid registration algorithms based

on intensity, such as the B-spline registrations of Rueckert [230] or various diffusive flow methods [231, 232]. This study presents an approach for automatically generating boundary conditions through the use of a non-rigid image registration algorithm called demons diffusion. The demons method was arbitrarily chosen based on its popularity and ease of implementation. The demons method was first proposed by Thirion [233], and is well-understood to have a strong mathematical foundation as in the works of Pennec *et al.* [234], Cachier *et al.* [235], and Modersitzki [236]. The basic premise presented by Thirion in [233] is to use an optical flow model governed by the idea of Maxwell's demons to drive the registration. In this model, the intensity of a moving object is considered to be constant with time, which implies that some level of correspondence can be achieved between deformed and undeformed images, as long as the intensity profiles are very similar. The object boundaries in one image are characterized as semi-permeable membranes, and the other image is allowed to diffuse through these membranes based upon the optical flow equation:

$$D(X) \cdot \nabla f(X) = -(m(X) - f(X)) \quad (34)$$

where $f(X)$ is the fixed target image, $m(X)$ is the source image being deformed for the registration, and $D(X)$ is the displacement field mapping the source to the target image through an instantaneous optical flow. Thirion reformulated (34) to an algorithmic iterative form as follows:

$$D^N(X) = D^{N-1}(X) - \frac{\left(m(X + D^{N-1}(X)) - f(X) \right) \nabla f(X)}{\|\nabla f\|^2 + \left(m(X + D^{N-1}(X)) - f(X) \right)^2} \quad (35)$$

The displacement field obtained from (35) is smoothed with a Gaussian filter between each iteration in order to regularize the registration. The popularity of the demons algorithm has also helped it remain an active area of research. Vercauteren *et al.* recently introduced symmetric diffeomorphic demons [237-239], and improvements to the registration regularization continue to be made by Cahill *et al.* [240], Mansi *et al.* [241] and other groups. In this paper, the demons algorithm was used to perform image matching of pre- and post-deformation images and tested

against a controlled *in silico* simulation with known boundary conditions. The generated boundary conditions were also used to perform an MIE elasticity reconstruction to evaluate its effectiveness in determining the elasticity contrast of a previously characterized system. The simulation study was followed by two phantom experiments to further stress the abilities of this new approach.

B.4 Methodology

B.4.1 Automatic Generation of Boundary Conditions

As described in previous work, the MIE algorithm is comprised of three major components: 1) a biomechanical FE model of soft-tissue deformation based on material properties, 2) a similarity metric with which to compare images, and 3) an optimization routine to update the material properties in the model [242].

The process of generating an elasticity reconstruction begins with the acquisition of an image of the organ. A mechanical load is then applied to the tissue, and the organ is imaged again. These pre- and post-deformation images comprise the primary input to the MIE algorithm, and are referred to as the source and target images, respectively. The organ boundary is then segmented manually in the pre-deformed source image and its surface geometry is extracted using the marching cubes algorithm, which allows a finite element mesh of tetrahedrons to be created from the surface information. The mesh is partitioned into 'regions' to which elasticity properties are assigned, which defines the resolution of the elastographic reconstruction. The biomechanical model used for the reconstruction is a linear elastic model, which holds that the strain experienced is proportional to the applied stress. We further assume that the materials of the FE mesh are isotropic and nearly incompressible in nature. Although breast tissue is known to not be linearly elastic, the system may be approximated as linear elastic with sufficiently small strains (all strains in this work are less than 15%). In work not presented here, it was found that a Poisson's ratio of 0.485 was optimal for use in MIE, and was used for all of the following experiments. The other critical material property in the model is Young's modulus, E , which is solved for by the MIE

framework. The ability of the biomechanical model to accurately deform the mesh of the candidate tissue volume is dependent on the accuracy of its boundary conditions. Once boundary conditions have been designated, the model is run and the FEM displacement solution for all the nodes in the mesh is obtained. The displacements are then used to deform the original pre-deformation image, which is then compared with the known post-deformation target image to generate an image similarity measurement. A non-linear optimization framework is used to update the material properties of the mesh based on the modeled deformation. The optimization is the minimization of the objective function:

$$\psi = |S_{TRUE} - S_{EST}|^2 \quad (36)$$

where S_{TRUE} is the similarity value achieved when comparing the target image to itself and S_{EST} is the similarity between the target and model-deformed source images. Differentiating (36) with respect to the elasticity distribution and setting the resulting expression equal to zero gives a series of nonlinear equations which is solved using the Levenberg-Marquardt method:

$$[J^T J + \alpha I] \{\Delta E\} = [J^T] \{S_{TRUE} - S_{EST}\} \quad (37)$$

where J is the Jacobian matrix whose size is determined by the number of material property regions, ΔE is the vector of updates to the material property distribution defined by the regions, and α is an empirical regularization parameter determined by the methods of Joachimowicz [243]. Modulus values in the mesh are updated by ΔE until an error tolerance on the relative objective function error evaluation is reached, at which point the reconstructed elastographic image is created from the current distribution of E values in the mesh regions.

The implementation of the demons algorithm used in this work to generate boundary conditions for the above model was based on the Insight Toolkit (ITK) [207, 244], and was derived from the original demons registration presented by Thirion. This included the use of simple Gaussian smoothing of the deformation field as the regularization of the registration. It should be noted that there are a multitude of regularization schemes in the literature, including those that

imitate linear elasticity [245], or elastic-like vector filters [246], which could potentially improve the results presented here. The two parameters of the registration were the number of iterations (single-resolution registrations) and the standard deviation of the Gaussian smoothing kernel. The number of iterations required was chosen separately for each data set such that the updates to the deformation field were observed to become very small by the end of the registration. In results not presented here, a brief sensitivity analysis was performed on the sigma value for the registration of the simulations, and it was found that over a range of 0 to 3 voxels for sigma, that 1.5 resulted in the most accurate deformation field. As the smoothing became stronger than 1.5, we noticed that the error became greater around the area of the depressions, as the kernel began to excessively smooth the depression displacements. Therefore, this number was used in all of the reported experiments. When the registration is executed, it produces displacements at the centroid of every voxel. The displacement vectors are then interpolated onto the nodal coordinates of the FE mesh using a cubic 3D interpolation. The displacements which are assigned to boundary nodes are thus designated as Type I boundary conditions for the biomechanical model.

B.4.2 Simulations

In order to evaluate the demons method of generating boundary conditions for MIE as described above, a controlled experiment was conducted by obtaining a CT and an MR image volume of human breast tissue and registering them to target images created by simulated mechanical loads. The two image sets (CT and MR) of normal tumor-free human breast tissue were obtained from the UC-Davis Department of Radiology and the Vanderbilt University Institute of Imaging Science, respectively, for use in this work. The surface of each tissue volume was segmented from the surrounding structures in the images with ANALYZE 8.1 (Mayo Clinic, Rochester, MN) and the resulting segmentation was used to create a 3D FE mesh using a tetrahedral mesh generation algorithm [113]. For both the CT set and the MR set, a 2-cm spherical tumor was synthetically implanted in the center of the respective mesh and assigned an elasticity

value six times higher than the surrounding material, which is consistent with breast cancer elasticity contrasts in the literature [247]. This contrast ratio of 6:1 was thus considered to be the goal for reconstruction in both cases.

Each finite element mesh was deformed by applying a depression to one side of the tissue volume. The displacements predicted by the model were then used to deform the CT and MR source images to provide simulated target images. Using the pre- and post-deformed image volumes, the demons registration could be executed and compared to the known displacements responsible for the simulated tissue deformations. In addition, the surface displacements could be used to test the accuracy and fidelity of the 3D MIE reconstructions conducted with demons-based boundary conditions. The registration for both simulations utilized 2,500 iterations with a σ of 1.5 voxels.

B.4.3 Phantom Experiment 1

After demonstrating the efficacy of the demons method in this highly controlled *in silico* simulation study, the next step was to apply the same tests to real-world data with realistic amounts of noise and uncertainty. To this end, phantom images were acquired to evaluate the ability of the demons method to produce accurate boundary conditions when compared to the current gold standard method.

As described in [227], the phantom used in this study (hereafter referred to as Phantom 1) was created from an 8% w/v solution of polyvinyl alcohol (Flinn Scientific, Batavia, IL) in an anthropomorphic breast mold. To provide intrinsic fiducial markers, 34 1-mm stainless steel beads were distributed over the phantom directly under its surface. It should be noted that, except for the beads, there was little to provide intensity heterogeneity within this phantom. A mechanical load was applied to the phantom in a custom-built acrylic chamber via a neoprene sphygmomanometer air bladder (W.A. Baum, Copiague, NY) positioned on the side of the phantom.

The phantom was subjected to three levels of compression by inflation of the air bladder: no compression, inflation with 50% of the maximum bladder pressure, and full inflation of the bladder.

At each state of compression, CT images were acquired with dimensions 512 x 512 x 174, and 0.54 x 0.54 x 1 mm voxel size. The images were then segmented and triangular meshes were created from the surface geometry of the phantom. From the surface meshes, the fiducial bead centroid positions were localized and then used in a TPS interpolation to provide the gold standard boundary conditions for two scenarios: 1) deforming from the uncompressed state to the 50% compression state, and 2) deforming from the uncompressed state to the 100% compression state. In generating the TPS boundary conditions, 33 of the beads were used in calculating the interpolation, while the last fiducial was used to evaluate the target registration error (TRE). In an effort to evaluate the error over the entire surface, the TPS registration was conducted 34 times, each time using a different fiducial for the TRE calculation. The final TRE for the TPS gold standard was the average of these repetitions. The demons method was then used independently to generate boundary conditions mapping from the pre- to the post-deformed surface of the phantom for the two scenarios, and compared to the control TPS result, as well as the previous semi-automated methods (Laplace equation and diffusion methods). The registration in both scenarios utilized 120,000 iterations and σ of 1.5 voxels.

B.4.4 Phantom Experiment 2

Following the evaluation of the performance of the demons method in generating boundary conditions in the above phantom study, a second phantom experiment was designed to test the performance of demons-based boundary conditions in the context of a full MIE reconstruction. Two more phantoms (hereafter referred to as Phantom 2 and Phantom 3) were constructed of polyvinyl alcohol cryogel (PVA-C) to test the accuracy of the reconstruction when validated with material testing data. As described by [248], the two new phantoms were created in a manner similar to the first phantom. However, these phantoms each included a 25-mm diameter phantom tumor composed of a stiffer gel than the bulk gel. Barium sulfate was mixed with the tumor gel and was randomly added in streams to the bulk gel to provide contrast for the CT images. Similar to the first

phantom study, polytetrafluoroethylene spherical beads (McMaster-Carr, Atlanta, GA) with a 1.6-mm diameter were distributed just under the surface of the phantoms in order to facilitate a TPS interpolation to act as the gold standard boundary conditions. Phantom 2 received 35 beads, while Phantom 3 received 32 beads. The TRE for the TPS registration was calculated as before. To provide validation for MIE reconstructions in Phantoms 2&3, independent mechanical tests were performed on samples of the two gel elasticity constituents of the phantom. A sample from each gel (tumor and normal) was set aside for this testing during fabrication. Each was subjected to compression testing using an ElectroForce 3100 material tester (Bose, Eden Prairie, MN). The instrument was programmed to provide fixed displacements to the cryogels when the samples were mounted on a platform over a 22.5 N load cell. Each sample was subjected to five cycles of a load rate of 0.15 mm/s and then held for 300 s for strains of 2, 5, 10, and 15% in compliance with small deformation theory. Average elastic modulus values for the two gels were obtained from the slope of the stress-strain curves of the steady-state loading phases.

The phantoms were imaged in the previously described air bladder chamber using a CT scanner (Philips Medical, Bothell, WA). The Phantom 2 CT images (pre- and post-deformation) were reconstructed with dimensions of 512x512x143 and voxel spacing of 0.27 x 0.27 x 0.8 mm, while the Phantom 3 CT images were reconstructed with dimensions of 512x512x139 and voxel spacing of 0.26 x 0.26 x 0.8 mm. The pre-deformed source image surfaces were then used to create tetrahedral meshes. The Phantom 2 mesh was constructed of 30,900 nodes and 166,509 elements, while the Phantom 3 mesh was constructed of 33,930 nodes and 183,609 elements. The TPS boundary conditions were generated using the implanted beads as control points for a thin-plate spline interpolation between the pre- and post-deformation surfaces for each phantom set. The PDE-based and demons methods were then utilized to independently generate boundary conditions for the two phantoms. The demons registration was set to run for 30,000 iterations with a σ of 1.5 voxels.

The accuracy of the demons-based boundary conditions was evaluated by comparing the gold standard TRE of the TPS method, the TRE of the PDE-based methods, and the TRE of the points when used in the demons method. The appropriateness of demons-based boundary conditions was then tested by employing them in a MIE reconstruction comparing elastic modulus values to independent measurements. To constrain the problem, only two regions of material properties were designated in the mesh: the tumor and the bulk normal gel. *A priori* knowledge of the location of the tumor was also used by segmenting the tumor margins from the normal gel beforehand in order to assign the material types to their corresponding elements in the FE model. The results of the MIE reconstruction using demons-based boundary conditions were also compared to the results of the reconstruction when using TPS boundary conditions, and those derived from the PDE methods. The Poisson's ratio used in the model for both experiments was 0.485 to approximate an incompressible tissue-mimicking material.

B.5 Results

B.5.1 Simulations

The CT and MR image source images were acquired and then deformed with the set of known boundary conditions as shown in Figure 54. The deformations applied in both cases were approximately Gaussian in distribution across the depressions. The maximum displacement experienced by the CT set was approximately 13 mm whereas the maximum experienced by the MR set was approximately 12 mm.




Figure 54. Representative slices from the two data sets used for the simulations. Slices (a) and (b) show the pre- and post-deformed CT set, whereas slices (c) and (d) show the pre- and post-deformed MR set, respectively.

The demons method was then used to register the source images to their respective target images and automatically generate boundary conditions for the source meshes. The TRE calculated from the boundary nodes was then calculated, and is visualized in Figure 55.

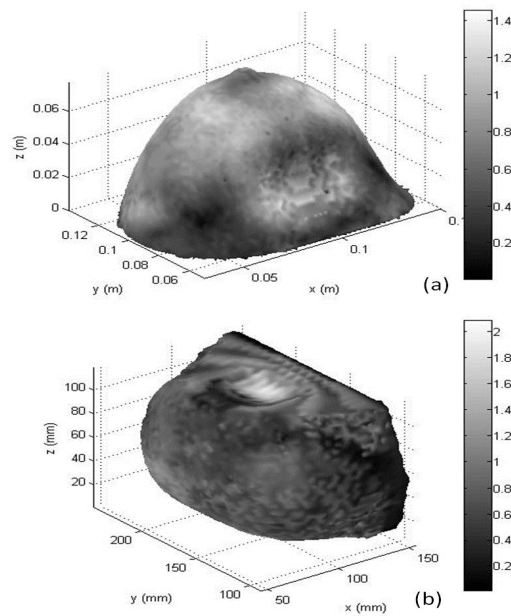


Figure 55. TRE distribution (in mm) across the surfaces of the CT mesh (a) and the MR mesh (b) for the demons-based boundary conditions compared to the known conditions.

The light surfaces of the mesh correspond to areas that experienced greater error when compared to the known boundary conditions. Averaging over all the nodes on the boundary, the CT

set experienced a mean error of 0.6 mm \pm 0.3 mm with a maximum error of 1.5 mm, which represents an average difference of about 17% between the magnitude of the TRE vectors and the magnitudes of the known displacement vectors. The MR set experienced a mean error of 0.5 mm \pm 0.3 mm with a maximum error of 1.9 mm, which represents a mean difference of about 23%. The demons-based boundary conditions were then utilized in an MIE reconstruction in an attempt to recapture the known 6:1 contrast in the simulations. The tumor-to-normal elasticity contrast calculated by the MIE algorithm was 3.63:1 for the CT set, and was 5.46:1 for the MR set. The results of the boundary condition accuracy and the resulting contrast ratios are shown in Table 11, as well as a comparison with the results of the three other boundary condition methods.

Table 11. Comparison of boundary condition mapping error and MIE reconstruction results between the four methods for the simulations (*[249], **[211]).

	Boundary Condition Mapping Error		MIE Reconstruction Results (x:1)	
	CT Mean TRE (max) mm	MR Mean TRE (max) mm	CT Elasticity Contrast Ratio	MR Elasticity Contrast Ratio
TPS (40 pts.)*	0.30 (2.6) *	0.033 (0.6)*	5.66**	6.26**
Laplace*	0.53 (2.6)*	0.48 (2.5)*	5.02**	673**
Diffusion*	1.5 (8)*	0.61 (2.9)*	17.5**	348**
Demons	0.60 (1.5)	0.50 (1.9)	3.63	5.46

Figure 56 illustrates the relationship between elasticity contrast ratios (tumor-to-normal) and the associated objective function values in the MIE optimization routine. The minima in the objective function space correspond to elasticity contrast values which resulted in an optimally

deformed image. Shown in the figure are the objective function values of the deformations using the known boundary conditions (as the control) and the demons boundary conditions.

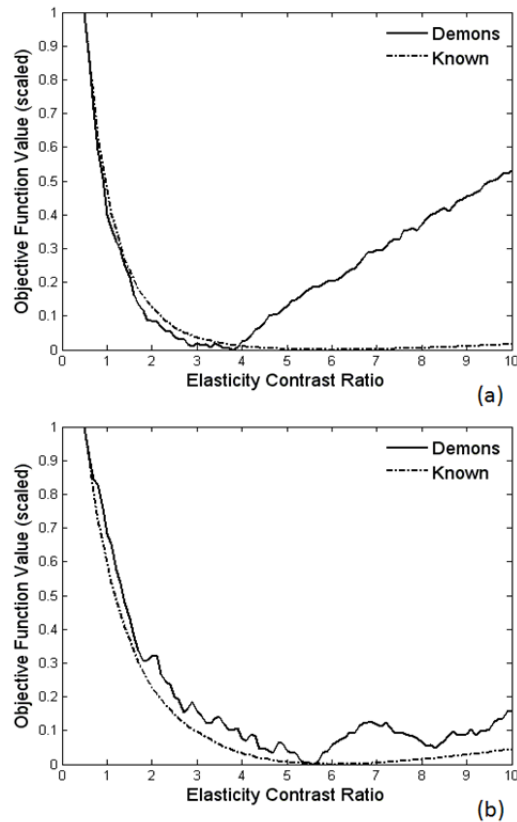


Figure 56. Objective function maps for the CT simulation (a) and the MR simulation (b). The objective function value calculated by the optimization framework is plotted on the ordinate axis against selected elasticity contrast ratios (tumor-to-normal) as affected by the boundary conditions. Shown are the objective maps of the demons case (solid lines) and the known boundary conditions as the control (dashed lines). The ordinate is scaled in both cases.

B.5.2 Phantom Experiment 1

In the first phantom experiment CT images of Phantom 1 were acquired at no compression, 50% compression, and 100% compression and segmented from the compression chamber. The demons method was used to generate Type I boundary conditions to map from the uncompressed state to the 50% state, and another set of boundary conditions to map from the uncompressed state to the 100% state. The implanted beads on the surface of the phantom were used to calculate the TRE of this surface registration in both cases. The average TRE for 50% compression when using

the demons boundary conditions was approximately 3.3 mm \pm 1.32 mm, with a maximum TRE of 6.1 mm. The average TRE for 100% compression was approximately 6.8 mm \pm 3.2 mm, which a maximum of 14.2 mm. The Phantom 1 results are directly compared in Table 12 to the gold standard TPS result and the results of the previous semi-automated methods, as well as to analogous results from Phantom 2 and Phantom 3.

Table 12. Comparison of boundary condition mapping error between the four methods for the two phantom experiments (*[249], **Based on work in [250])

	Boundary Condition Mapping Error			
	Phantom 1		Phantom 2	Phantom 3
	50%	100%	Single Compression Mean TRE (max) mm	Single Compression Mean TRE (max) mm
	Compression Mean TRE (max) mm	Compression Mean TRE (max) mm		
TPS	1.1 (3.4)*	1.7 (5.1)*	1.4 (7.08)**	1.24 (4.9)**
Laplace	3.4 (8.6)*	6.3 (15.3)*	4.22 (7.26)	2.24 (4.74)
Diffusion	2.7 (6.9)*	5.7 (13.6)*	4.11 (6.57)	2.35 (6.36)
Demons	3.3 (6.1)	6.8 (14.2)	1.55 (4.92)	1.85 (4.34)

B.5.3 Phantom Experiment 2

In the second phantom experiment, CT images of Phantom 2 and Phantom 3 were acquired and segmented from the compression chamber for pre- and post-depression. Phantom 2 is shown in Figure 57 as an example. The embedded tumor in Phantom 2 was about 12 mm from the surface. The tumor in Phantom 3 was located further from the site of depression, at about 26 mm from the surface. Qualitatively, the streams of barium sulfate which were distributed throughout the gel

provided an increase in the image texture of these phantom images over the Phantom 1 images, which lacked this texture enhancement.




Figure 57. Representative source (a) and target (b) geometry of the anthropomorphic phantoms, with center slices. Phantom tumors as shown here were present in Phantoms 2&3 only.

The demons method was applied to both phantoms to acquire Type I boundary conditions for each mesh. The TRE of the demons-based conditions was evaluated by comparing to the known point correspondence of the implanted surface beads. The average demons-based TRE for Phantom 2 was calculated to be approximately $1.6 \text{ mm} \pm 1.0 \text{ mm}$, with a maximum experienced TRE of 4.9 mm. For Phantom 3, the average TRE was $1.9 \text{ mm} \pm 1.2 \text{ mm}$, with a maximum TRE of 4.3 mm. These values are directly compared in Table 12 to the performance of the gold standard TPS interpolation method and two previous semi-automated methods, as well as the Phantom 1 results. As the results showed that the PDE-based methods were not notably more accurate for Phantoms 2 and 3 than the TPS or demons methods, only the demons method and TPS method were used in MIE reconstructions for comparison. The material testing data resulted in an average contrast ratio of 4.10:1 for the gels. The demons-based boundary conditions were then used in an MIE reconstruction for each phantom. The tumor-to-normal elasticity contrast for Phantom 2 was calculated by the MIE algorithm to be 4.70:1.

The elasticity contrast for Phantom 3 was calculated to be 2.46:1. In Table 13, these values are compared to the contrast ratios that were calculated by MIE using the gold standard TPS

boundary conditions, and to the material testing data as validation for the accuracy of the MIE method.

Table 13. MIE-reconstructed elasticity contrast ratios for phantoms 2&3 and gel material testing data (*[211])

	Phantom 2 Reconstructed Contrast Ratio	Phantom 3 Reconstructed Contrast Ratio	Material Tester Contrast Ratio*
TPS*	3.81	3.06	4.10
Demons	4.70	2.46	

Figure 58 illustrates the relationship between elasticity contrast ratios (tumor-to-normal) and the associated objective function values in the MIE optimization routine. Shown in the figure are the objective function values of the deformations using the TPS boundary conditions (as the control) and the demons boundary conditions.

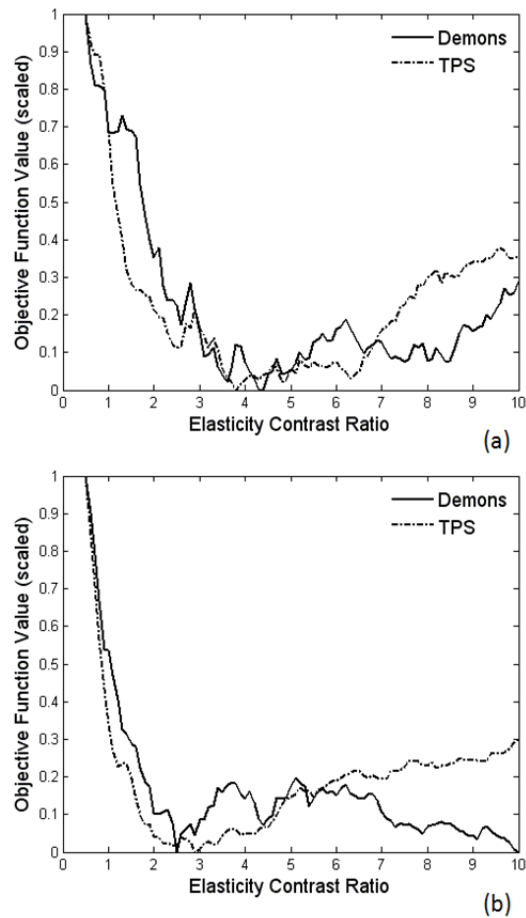


Figure 58. Objective function maps for Phantom 2 (a) and the Phantom 3 (b). The objective function value calculated by the optimization framework is plotted on the ordinate axis against selected elasticity contrast ratios (tumor-to-normal) as affected by the boundary conditions. Shown are the objective maps of the demons case (solid lines) and the known boundary conditions as the control (dashed lines). The ordinate is scaled in both cases.

B.6 Discussion

B.6.1 Simulations

The demons-based boundary conditions resulted in deformed meshes for the simulation experiment which were qualitatively very close in appearance to the known target meshes for both the CT and MR data sets. Quantitatively, the average difference between the demons conditions and the known conditions was about 20% for both sets, which was an encouraging indication of the ability of the demons methods to automatically provide boundary conditions which would have

adequate accuracy for use in MIE. In Figure 55 it can be seen that the largest errors were spread across the regions of high curvature around the tip of the tissue volume and in the dip of the artificial depression for the CT set, while in the MR set the errors were mostly localized to the depression area.

The accuracy of the demons-based boundary conditions for the simulations were compared to the results of past methods in Table 11. Unsurprisingly, the TPS method remained the most accurate of the four methods when considering the average boundary condition error. The demons method performed about as well as the Laplace method, and clearly outperformed the diffusion method for the CT set in terms of the average error. However, the demons method performed favorably compared to all of the other methods in terms of maximum TRE for that set, as its maximum error was well below those of the other methods. In terms of average surface TRE, the demons method was also comparable to the PDE-based method for the MR set as well. However, with the exception of the TPS method, the demons boundary conditions again compared favorably against the other methods in terms of the maximum error experienced on the boundary.

The results of the boundary condition accuracy experiment were encouraging and indicated that demons-based boundary conditions were a feasible solution to the MIE boundary condition problem. The results of the MIE reconstruction for the CT and MR simulation sets were shown in Table 11 and compared to the results of reconstructions which utilized boundary conditions generated from the other three methods. Unsurprisingly, the table shows that the TPS boundary conditions, which were the most accurate of the four, resulted in elasticity contrast ratios for both sets that were closer to the known ratio of 6:1 than any of the other boundary conditions. For the application of the demons registration-based boundary conditions to the CT data set, the elasticity reconstruction with spatial *a priori* knowledge of the tumor converged to a contrast ratio of 3.63:1. Similarly, the MR data resulted in a contrast ratio of 5.46:1. Compared to the known designated material contrast of 6:1, there is clearly a discrepancy in these reconstruction behaviors that needs

to be investigated. The difference, particularly between the different modalities of input data, is likely due to a combination of factors including mesh geometry and image quality. In addition, the distance of the tumor from the area of greatest displacement likely affects the accuracy of the reconstruction since the displacements of nodes are expected to decrease the further they are located away from the depression. These simulations did not investigate the effect of tumor distance on the reconstruction. Notably, the diffusion method resulted in a much higher contrast ratio for the CT set than the demons method, while the Laplace method resulted in a contrast ratio that was closer to 6:1 but was an underestimation rather than an overestimation of the true value. The ability of the demons-based conditions to provide a contrast that was more accurate than the diffusion method for the CT simulation was encouraging. Even more suggestive was the behavior of the MR reconstruction. The Laplace and diffusion boundary conditions introduced instabilities into the MIE algorithm, which resulted in contrast estimates that were unreasonably higher than the true value. The demons-based conditions allowed the algorithm to provide a contrast estimate which was closer to the known value.

Introducing the inexact demons boundary conditions to the model had a noticeable effect on the objective function profile, as shown in Figure 56 by shifting the minimum objective function value to a different optimal elastic contrast ratio for both the CT and the MR simulation. The shift was much more pronounced for the CT simulation, for which the new optimal objective function value corresponded to a contrast ratio of about 3.80:1 instead of 6:1 as predicted by the known boundary conditions. Additionally, the convexity of the function was altered significantly, with very little variation in the objective function for contrast ratios in the immediate vicinity of the global minimum. The MR simulation also experienced a shift in the optimal objective function when demons boundary conditions were used instead of the known conditions, with a new optimal contrast of about 5.50:1. This represented only a slight decrease from the desired 6:1 prediction. The objective function values arise from the image similarity metric, which again suggests that the

difference in objective maps between the two simulations is influenced by the image texture characteristics. It is also clear that the addition of inaccuracies within the boundary conditions alters the nature of the objective function by injecting local minima and undesirable variations, which may necessitate a filtering approach to ensure global minima are found.

B.6.2 Phantom Experiment 1

While the efficacy of the automated demons method was shown by the simulations to be comparable to the semi-automated Laplace method and somewhat better than the diffusion method, the simulations were in several ways performed under optimal conditions. The image volumes qualitatively had a great deal of heterogeneity and texture on which the demons registration could act, and with which the MIE optimization routine could use to help accurately update material property assignments. There was also an absolute truth with which to compare, in the form of known boundary conditions. The first phantom experiment sought to provide additional challenge to the demons method in its ability to generate reasonably accurate boundary conditions.

In the first phantom experiment, the results of the demons method were compared to the results of the three other methods in Table 12 for the two compression states applied to Phantom 1. The table shows that the demons algorithm performed about as well in relation to the other PDE methods as it did in the simulation experiment. Note that Phantom 1 had very little image heterogeneity and would indicate that with a lack of image intensity contrast that the demons-based registration is at least no worse than that achieved by the PDE methods. The gold standard TPS method gave the lowest error. As seen in Table 12, the errors given by all of the methods increased when a larger deformation was applied to Phantom 1. The demons boundary conditions became slightly worse in relation to the other methods at the increased level of compression, which suggests that the number of iterations used by the demons algorithm may need to be increased to

accommodate larger differences between pre- and post-deformation images, or that the algorithm may be somewhat more sensitive to the lack of image intensity heterogeneity.

In moving from simulation data to “real-world” phantom data, the errors experienced by all four of the methods increased significantly. The Phantom 1 image data was different from the simulation data in several key respects. For example, the target image volume of Phantom 1 represents a completely new acquisition, whereas in the simulation work post-deformed image sets were generated from the pre-deformed set. This discrepancy in target image acquisition introduces some uncertainty to the determination of source-to-target correspondence. Another major change from the simulation experiment was the markedly smaller presence of texture in the images due to the homogeneity of the gel. More specifically, the TRE performance varied among the Phantom 1, Phantoms 2&3, and simulation results which are listed respectively in terms of increasing image texture. Qualitatively observing the results across Table 11 and Table 12, the trend of decreasing TRE with increasing texture for the demons-based approach can be observed.

B.6.3 Phantom Experiment 2

It was shown in the first phantom experiment that the demons method could produce reasonably accurate boundary conditions compared to the semi-automated Laplace and diffusion methods. The second phantom experiment introduced another set of real-world data, but the images from this experiment had more texture in the form of barium sulfate as a contrast agent, which was intended to allow the demons registration to provide more accurate boundary conditions as needed by the MIE algorithm. In addition, the presence of the stiff tumor allowed for a test of the MIE algorithm’s ability to distinguish elasticity contrast in a phantom while using demons-based boundary conditions. This experiment was thus the first in which demons-based boundary conditions were used in an MIE reconstruction for which the true boundary conditions were not absolutely known.

The surface errors calculated from the fiducial point correspondence for the TPS, Laplace, diffusion and demons methods were compared in Table 12 for Phantom 2 and Phantom 3. Unsurprisingly, the TPS method performed better with respect to mean accuracy. Notably, the maximum error experienced by the demons method was less than that of the TPS method, which was similar to the result of the CT simulation study. The two PDE-based methods presented error which was similar in scope to their Phantom 1 results. Overall, the demons method performed considerably better on these two phantom sets than it did on Phantom 1, and notably outperformed the Laplacian and diffusion methods. This is most likely due to the increase in image texture which can be qualitatively observed from visual inspection of the images. Given that clinical images tend to have even more image texture and geometric heterogeneity than found in these phantom images, further investigation into the efficacy of the demons method seems merited.

The utilization of the demons boundary conditions in MIE reconstructions in the second phantom experiment successfully resulted in realistic tumor-to-normal modulus contrast ratios for both phantoms. Due to the observation that the demons method resulted in boundary conditions with comparable (and sometimes superior) accuracy to the Laplace and diffusion methods, only the TPS and demons boundary conditions were utilized in these reconstructions. The results for the TPS- and demons-based MIE reconstructions were compared to each other in Table 13 as well as to the material tester results. As the table shows, the elasticity contrast ratios for each phantom when using TPS boundary conditions were reconstructed to values with 14-40% difference from the material testing data average. The reconstructions using demons boundary conditions resulted in contrast ratios which were very similar to those of the TPS-based reconstructions, with only a slight drop in contrast. This suggests that the demons boundary conditions were sufficiently accurate for the MIE algorithm to provide a reasonable estimate of the actual gel contrast.

Compared to the control TPS boundary conditions, the demons conditions had a noticeable effect by shifting the minimum objective function value to a different optimal elastic contrast ratio

for both phantoms, as shown in Figure 57. Additionally, the convexity of the function was altered slightly for each. The global minimum of the Phantom 2 objective function was located at an approximate contrast ratio of 4.20:1, which was more similar to the material testing average of 4.10:1 than the case in which TPS boundary conditions were used. The actual contrast ratio to which the MIE reconstruction converged was 4.70:1, which was located on the slope of a local minimum. This behavior was most likely a result of the regularization parameters used in the Levenberg-Marquardt optimization. In the case of Phantom 3, the global minimum was about 2.50:1, which was the approximate value to which the algorithm converged. In this case, the global minimum decreased slightly when using demons instead of TPS conditions.

Observations of Figure 56 and Figure 58 indicate the change in algorithm performance with respect to simulation and physical data. While the nature of a simulation-to-real transition may be responsible for the increased error in reconstruction, there are several other likely factors involved. Over-constraint of the problem is a possible candidate with the incorporation of the spatial prior. The MIE method works by sampling similarity regionally, i.e. the method breaks up evaluation into many similarity zones (usually over 100) distributed spatially over the domain. The method tries to improve the similarity among all the zones with the use of only two parameters in this case (the elasticity of the background and tumor). This constraint within this type of problem can lead to this type of oscillatory behavior. Another possible reason is the inaccuracy in boundary condition determination due to the dramatic difference in image heterogeneity between simulation and real data. This is supported by the change in TRE. Related to this, it is interesting to note the difference between CT and MR reconstruction for the simulation work associated with Figure 56 and in light of Table 11. The first observation can be made by comparing the control objective function map across CT and MR simulation sets in Figure 56. Both simulation sets had a contrast ratio of 6:1, with the only difference being the level of intensity heterogeneity, and potential different tissue-volume/tumor geometries/locations. The CT control had a shallower minimum which may affect

the reconstruction. When adding to this observation the objective function maps associated with the demon-based boundary condition it would seem that the CT reconstruction may perform better due to its convexity; but when observing how the minimum has been shifted, and the shape of the control that has no error in boundary conditions, it can be seen that in fact the MR demons-based objective function maps more closely to its control which is reflected in the elasticity contrast ratio.

B.7 Conclusions

The simulations and phantom experiments conducted in this work indicate that while TPS interpolation remains the most accurate method used thus far in MIE for generating boundary conditions, the demons method shows promise in situations where fiducial point correspondence data may not be available. In addition, when transitioning from simulation to real data, the discrepancy in performance between TPS and the demons-based boundary condition mapping becomes less (at least in cases where image intensity contrast within the domain is available). Furthermore, while the higher accuracy of the TPS method is desirable, the much higher level of manual user interaction and numerous fiducials needed for the method make clear the desire for alternative methods of boundary condition generation. The demons method proposed represents a fully automated approach.

While the results are encouraging, the challenge of predicting (prior to workflow initiation) how well a pre-post deformation image set will fare prior to execution of the demons registration and MIE optimization routine still remains. Since the demons registration algorithm possesses diffusive behavior based upon intensity contours as described in [233], it is obvious that the images require a certain level of texture and intensity heterogeneity in order to provide these membranes a meaningful registration. This is one of the likely causes of the varying performance of the demons method in generating accurate boundary conditions among the experiments presented in this work. In work not presented here, our group observed similar error performance when using a more recent diffeomorphic log-domain version of the demons registration algorithm. While the

registration error was virtually the same as the original demons implementation, it is possible that with noisier images, the diffeomorphic nature of the more recent demons may help minimize phenomena such as collapsed mesh elements due to overlapping displacements. Development of a feasibility metric which can predict the success of applying the MIE algorithm to a given image set is a needed next step for the project.

In addition to a threshold criterion to evaluate the potential for a successful reconstruction, the need to generate more realistic phantoms with controllable stiffness properties is also necessary. The breast has a complex image signature even within CT and the reproduction of those patterns coupled with controllable elasticity properties is very challenging. While obstacles remain, the results presented here demonstrate the potential of treating elastographic reconstructions using non-rigid image registration approaches and that the possibility of full automation is also within reach.

Acknowledgments: This work was supported by the National Institutes of Health–National Institute for Neurological Disorders and Stroke under Grant R01 NS049251. The authors thank the Vanderbilt Medical Center CT department for imaging the phantoms, as well as Dr. Thomas Yankeelov and Dr. John Boone (University of California-Davis Medical Center, Dept. of Radiology) for the human breast MR and CT data sets, respectively.