

TECHNISCHE UNIVERSITÄT MÜNCHEN

Fakultät für Informatik
Computer Aided Medical Procedures & Augmented Reality / I16

Ultrasound Confidence Maps and Applications in Medical Image Processing

Athanasios Karamalis

Vollständiger Abdruck der von der Fakultät für Informatik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr. R. Westermann

Prüfer der Dissertation: 1. Univ.-Prof. Dr. N. Navab
2. Prof. S.E. Salcudean, Ph.D.
The University of British Columbia,
Vancouver/Kanada

Die Dissertation wurde am 29.01.2013 bei der Technischen Universität München eingereicht und durch die Fakultät für Informatik am 10.06.2013 angenommen.

Abstract

Advancements in medical ultrasound imaging have been driven by the continuous improvement of imaging systems and novel approaches for image processing applications. The modeling and simulation of ultrasound contributes to our understanding of complex sound propagation in human tissue and also to the improvement and development of ultrasound systems. Furthermore, modeling the ultrasound image formation process provides valuable prior information for image processing algorithms.

In terms of simulation this thesis contributes a new framework for fast and realistic ultrasound image simulation. The nonlinear full-wave equation is used for the modeling of sound propagation and solved with a Finite-Difference Time-Domain scheme on GPU hardware to reduce simulation time. The essential image formation pipeline is simulated including pulse transmission, beamforming and signal processing. Realistic images are simulated from computer phantoms and phantoms generated from histology images. In the later case histology and intravascular ultrasound images were acquired from the same cross-section of vessels. Consequently, simulated images based on histology could be compared to real images in the same reference coordinate system.

The main contribution of this thesis is a novel method for estimating confidence in ultrasound images, denoted ultrasound confidence maps. The novelty of the method lies in the formulation of the confidence estimation problem as a random walk on a graph, taking into account essential ultrasound physics and the image formation process. The fast and closed-form solution to this problem comes with an inherent robustness to image acquisition parameters, which is demonstrated by obtaining consistent confidence maps for images acquired with different transducers, systems and settings. In addition, the benefit of using confidence maps for image processing applications is demonstrated by introducing them into algorithms for shadow detection, 3D freehand reconstruction, classification in intravascular ultrasound and mono- and multimodal registration. In all cases algorithms benefited from confidence map information and produced improved results.

In the process of this thesis an ultrasound-guided navigation system was developed for computer-aided orthopedic surgery. For this, a novel feature descriptor is introduced for automatic bone detection in ultrasound, based on the previously introduced confidence maps. The system performed automatic bone surface detection in ultrasound images in real-time and subsequent rigid registration to the bone surface extracted from pre-operative Computer Tomography data set. This would allow fusion of pre-operative imaging and planning data with intra-operative imaging by bringing this information into the same reference coordinate system. The algorithms and system were evaluated on data sets acquired in-vivo and in a cadaver study, demonstrating high accuracy, robustness and the potential for intra-operative application.

Keywords: Ultrasound, Confidence, Simulation, Modeling, Random Walks, Bone Detection, Shadow Detection, Reconstruction, Registration, Medical Image Processing

Zusammenfassung

Fortschritte in der medizinischen Ultraschallbildgebung wurden durch kontinuierliche Verbesserungen der bildgebenden Systeme und durch das Einführen von neuartigen Methoden für die medizinische Bildverarbeitung angetrieben. Die Modellierung und Simulation von Ultraschall trägt nicht nur zum besseren Verständnis der komplexen Schallausbreitung im menschlichen Gewebe bei, sondern auch zur Verbesserung und Weiterentwicklung von Ultraschallsystemen. Zudem bietet die Modellierung der Ultraschallbildgebung wertvolle Informationen für Algorithmen der Bildverarbeitung.

Im Hinblick auf Simulation stellt diese Dissertation eine neue Plattform für schnelle und realistische Simulation von Ultraschallbildern vor. Die Schallausbreitung wurde mit der Westervelt Gleichung modelliert und mit der Finite-Differenzen-Methode auf GPU Architektur gelöst um die Rechenzeit zu reduzieren. Die wesentlichen Bestandteile der Ultraschallbildgebung wurden simuliert einschließlich der Erzeugung von Ultraschallimpulsen, der Formung des Schallstrahls und der Signalverarbeitung. Realistische Bilder wurden mithilfe von Computer Phantomen simuliert die unter anderem auch auf Basis von Histologie Bildern erstellt wurden.

Der wichtigste Beitrag dieser Dissertation ist eine neuartige Methode zur Schätzung der Konfidenz in Ultraschallbildern, welche als “Ultraschallkonfidenzkarte” bezeichnet wird. Die Neuartigkeit der Methode liegt in der Formulierung des Problems als zufälligen Lauf auf einen Graphen, wobei die Grundlagen der Ultraschallphysik und der Bildgebung berücksichtigt werden. Die analytische Lösung bietet inhärente Robustheit in Bezug zur Parametrisierung von Ultraschallsystemen. Dies wurde durch die konsistenten Ergebnisse von Ultraschallkonfidenzkarten auf Bildern die mit unterschiedlichen Proben, Maschinen und Einstellungen aufgenommen wurden dargestellt. Ultraschallkonfidenzkarten bieten wertvolle Information für Algorithmen der Bildverarbeitung, wobei diese in Algorithmen wie Schattenerkennung, 3D Freihand Rekonstruktion, Klassifikation von Gewebe in IVUS Bildern und mono- und multimodale Registrierung eingeführt wurden. In allen Fällen lieferten Ultraschallkonfidenzkarten Information die zur Verbesserung der einzelnen Algorithmen führte.

Im Verlauf dieser Dissertation wurde ein Navigationssystem auf Basis der Ultraschallbildgebung für computergestützte orthopädische Chirurgie entwickelt. Hierfür wurde ein neuartiges Feature für die Erkennung von Knochen in Ultraschallbildern entworfen, welches auf Ultraschallkonfidenzkarten basierte. Das System erkannte dabei vollautomatisch die Knochenoberfläche in Ultraschallbildern und registrierte diese zu pre-operativen Computertomographie Datensätzen. Algorithmen und Gesamtsystem wurden auf Daten evaluiert die in-vivo und in einer Kadaver Studie aufgenommen wurden, wobei Genauigkeit und Robustheit den potentiellen Einsatz in intraoperativen Szenarien unterstützt.

Schlagwörter: Ultraschall, Konfidenz, Simulation, Modellierung, Random Walks, Knochenerkennung, Schattenerkennung, Rekonstruktion, Registrierung, medizinische Bildverarbeitung

Acknowledgments

I would like to start by thanking my advisor Prof. Nassir Navab for giving me this unique opportunity to do my PhD in the wonderful environment of his CAMP chair. I highly appreciate his continuous support throughout the years and the freedom to pursue the ideas that I was excited about. His deep understanding of science and his ever expanding vision for the future of computer aided medical procedures has been a continuous motivation for me. I highly appreciate the experiences that he shared beyond scientific topics and will hopefully remember them for the times to come.

I owe special thanks to Wolfgang Wein. His continuous support, starting with my internship at Siemens, has contributed significantly to my PhD and beyond. In the years throughout my PhD studies he has been a mentor, colleague and a great friend. His vision and everlasting enthusiasms for doing research and development have been a catalyst for my own work. I deeply appreciate all our discussions and late-night coding and paper-writing sessions.

I would further like to thank Amin Katouzian for the great collaboration we had in the last years of my studies. Thanks to him my PhD gained significant momentum at the right time and I deeply appreciate all the opportunities he has given me, taking me on-board exciting projects, introducing me to great people and being a great teammate. Beyond this I also appreciate the time we spent at the kicker table and enjoying good food together.

The time at the CAMP chair has been extraordinary because of the great people that are part of it. I would like to start by thanking my office mates Tassilo Klein, Seyed-Ahmad Ahmadi and Christian Wachinger for the wonderful time at the office, conferences and trips. The laughs we shared will remain unforgettable.

Furthermore, I would like to thank: Stefan Holzer, Jürgen Sotke, Ali Bigdelou, Vasileios Belagiannis, Christoph Hennesperger, Abouzar Eslami, Debdoot Sheet, Markus Müller, Jakob Vogel, José Gardiazabal, Martina Hilla, Mehmet Yigitsoy, Maximilian Baust, Oliver Kutter, Darko Zikic, Tobias Lasser, Tobias Reichl, Olivier Pauly, Martin Horn, Loren Schwarz, Nicolas Brieu, Tobias Blum, Loic Peter, Ben Glocker, Slobodan Ilic, Pascal Fallavollita, Selen Atasoy, Cedric Cagniard, Richard Brosig, Stefan Hinterstoisser, Silvan Kraft, Amit Shah, Mattias Hansson, Peter Maday, Hauke Heibel, Razvan Ioan Ionasec, Victor Castaneda, Christian Schulte zu Berge, Benoit Diotte, Julian Beitzel, Marco Feuerstein, Martin Groher, Alexander Dului, Alexander Ladikos, Arash Taki, Peter Sturm and everybody else I had the pleasure meeting, discussing and working with.

I owe many thanks to my parents Georgia and Dimitrios for all their support. Finally, I would like to express my deepest gratitude to my wife Marilena for her support, patience and sharing with me this adventure. Her everlasting smile and kindness have accompanied me from the beginning up until the end of this exciting journey.

CONTENTS

Abstract	iii
Zusammenfassung	v
Acknowledgments	vii
Thesis Outline	1
1 Introduction	3
1.1 Motivation	3
1.2 Thesis Outline	4
1.3 Contributions	6
I Ultrasound Imaging and Simulation	9
2 Medical Ultrasound Imaging	11
2.1 Medical Ultrasound	11
2.2 Sound Waves	12
2.3 Ultrasound Imaging	16
2.3.1 Acoustic Pulse	16
2.3.2 Beam Forming	17
2.3.3 Ultrasound Transducers	18
2.3.4 Signal Processing Pipeline	19
2.3.5 Ultrasound Transducer Types	20
2.4 Image Characteristics and Artifacts	22
2.4.1 Image Characteristics	22
2.4.2 Imaging Artifacts	23
2.5 Advanced Ultrasound Imaging	25
2.5.1 Doppler Imaging	25
2.5.2 Speckle Reduction	26
2.5.3 Microbubbles	27
2.5.4 Harmonic Imaging	28

CONTENTS

2.5.5	Highly Focused Ultrasound	28
2.5.6	Elastography	29
2.6	Clinical Application Fields	30
2.7	Other Medical Imaging Modalities	31
2.7.1	X-ray Imaging	31
2.7.2	Computer Tomography	32
2.7.3	Magnetic Resonance Imaging	33
2.7.4	Gamma Imaging	33
2.7.5	Histology	34
2.8	Ultrasound and other Modalities	34
3	Simulation	35
3.1	Introduction	35
3.2	Related Work	36
3.3	Nonlinear Full-Wave Model	38
3.4	GPU Computing	39
3.5	Ultrasound Image Simulation	41
3.5.1	Ultrasound Transmission and Reception	41
3.5.2	Radio Frequency Processing	42
3.6	Results	43
3.6.1	Simulation from Synthetic Phantoms	43
3.6.2	Simulation from Histology	45
3.7	Dicussion and Conclusion	46
II	Ultrasound Confidence Maps and Applications	49
4	Ultrasound Confidence Maps	51
4.1	Introduction	51
4.2	Related Work	52
4.3	Methods	55
4.3.1	Incidence and Graph Laplacian Matrix	55
4.3.2	Random Walks for Image Segmentation	56
4.3.3	Random Walks and Circuit Analogy	57
4.3.4	Ultrasound Confidence Map	58
4.3.5	Implementation Considerations	62
4.4	Evaluation	63
4.4.1	Parameters	63
4.4.2	Confidence in Shadow Regions	65
4.5	Confidence Map Applications	66
4.5.1	Shadow Detection	67
4.5.2	Ultrasound Reconstruction	69
4.5.3	Multi-modal Registration	71
4.5.4	Mono-modal registration	73
4.5.5	Confidence Estimation in IVUS	74

4.5.6	Discussion	78
4.5.7	Conclusion	79
4.5.8	Acknowledgments	79
5	Automatic Bone Detection and Registration	81
5.1	Introduction	81
5.2	Related Work	82
5.3	Detection Method	84
5.4	Surface-based Registration	86
5.5	3D Freehand-Ultrasound System Setup	86
5.6	Results and Evaluation	89
5.6.1	Detection Evaluation	90
5.6.2	Registration Evaluation	91
5.7	Conclusion	92
5.8	Acknowledgments	93
III	Conclusion and Appendix	99
6	Conclusion and outlook	101
A	Visualization of 4D Data Sets	105
A.1	Introduction	105
A.2	Previous Work	106
A.2.1	Direct Volume Rendering (DVR)	107
A.2.2	DVR and Vector Geometry	107
A.2.3	DVR and 3D Geometry	108
A.2.4	Motion Transfer Function	109
A.3	Survey	110
A.4	Discussion	111
A.5	Conclusion	113
B	Authored and Co-authored Publications	115
C	Abstracts of Co-author Publications	119
D	List of Abbreviations	123
	List of Figures	124
	References	131

I learned very early the difference between knowing the
name of something and knowing something.

Richard P. Feynman

INTRODUCTION

1.1 Motivation

Medical imaging has revolutionized health care starting with the discovery of X-rays by Wilhelm Conrad Röntgen in 1895. Since then scientists from different fields including mathematics, physics, engineering, medicine and computer science have come together to contribute to medical imaging. Pioneers in the field such as Allan M. Cormack and Godfrey N. Hounsfield (Nobel Prize in Physiology or Medicine 1979) who introduced Computer Tomography (CT) or as Paul C. Lauterbur and Sir Peter Mansfield (Nobel Prize in Physiology or Medicine 2003) who introduced Magnetic Resonance Imaging (MRI) have contributed tremendously to the continuous advancement of the field. Nowadays these advancements are still continuing with new systems being introduced and existing ones being further improved.

Medical ultrasound imaging was early on part of this revolution, offering real-time, non-ionizing and portable imaging capabilities. These characteristics have made ultrasound imaging the preferred choice in clinical scenarios that require non-ionizing imaging (e.g. obstetrics) or intra-operative and fast imaging (e.g. trauma surgery). Ultrasound systems are continuously developed further to provide better image resolution, quality and new imaging capabilities like real-time 4D imaging.

Mathematical modeling and simulation of imaging systems is important for developing new systems and improving existing ones. Mathematical modeling is at the core of developing fundamentally new approaches for imaging, including phase-contrast X-ray imaging [119] and fluorescence tomography [112]. Simulation is the tool for evaluating new models and providing insight into phenomena that would require complex experimental setups to be understood. The mathematical modeling of ultrasound wave propagation together with the development of simulation platforms for these models have contributed significantly to the development of new ultrasound imaging systems and sensors [60]. The research in this field is far from completed and new methods appear that provide further insight on the complex propagation of sound in inhomogeneous media like human tissue.

Computers science has also played an important role for realizing modern medical imaging such as ultrasound, Computer Tomography (CT) and Magnetic Resonance Imag-

ing (MRI). New computer architectures, systems and algorithms provided the necessary computational resources to perform image acquisition and image formation in time frames acceptable for routine clinical application. In addition, image processing, analysis and computer-aided diagnosis (CAD) can enhance image quality, extract valuable diagnostic information from the raw image data and assist the physician in the diagnosis, planning and treatment of patients. Most of the tasks performed semi- or fully-automatically by image processing and analysis algorithms can be performed by physicians. Furthermore, in many cases the accuracy of physicians on these tasks is better than the one of algorithms. However, there are different considerations that motivate the use of algorithms. Physicians have limited time for performing tasks like segmentation and registration of medical images. Moreover, 3-dimensional (3D) or even higher-dimensional data sets require hours of manual processing. The increasing cost of modern health care, the increasing amount of imaging data and the limited time of physicians call for computer-assisted image processing and analysis. An additional consideration is inter- and intra-observer variability [147]. Depending on the experience and training of physicians different image processing and analysis outcomes are expected. This variability can be a limiting factor for standardizing and optimizing treatment options for patients with very similar pathologies. Consequently, the research field of medical image processing and analysis is growing rapidly; with numerous conferences, journals and research groups dedicated to it. Although considerable advancements have been made, many algorithms are still not providing an optimal solution. For example, image registration remains an open challenge - especially deformable registration - even after 30 years of continues research and development in this area [37]. Nevertheless, it should be noted that many algorithms have found their way into clinical products and provide a considerable assistance to physicians. For example, algorithms in radiation therapy reduce significantly the radiation dose to healthy tissue while focusing the dose to diseased ones like cancer [164]. Closing, the search for the optimal algorithmic solutions for different problems continues and is likely to continue in the next decades.

1.2 Thesis Outline

Chapter 1: Introduction

This chapter provides the motivation and an overview of the thesis. In addition, it summarizes the contributions made to ultrasound simulation and modeling for image processing applications.

Part I - Ultrasound Imaging and Simulation

Chapter 2: Medical Ultrasound Imaging

This chapter provides an overview of ultrasound physics, image formation and systems. It discusses different imaging techniques including Doppler, harmonic and contrast agent imaging. Furthermore, the use of therapeutic ultrasound and clinical application scenarios

for ultrasound are discussed. The chapter closes by discussing other medical imaging modalities to provide the context for subsequent chapters.

Chapter 3: Simulation

Initially, this chapter discusses ultrasound simulation frameworks and subsequently introduces a novel simulation framework for fast and realistic image simulation. Implementation considerations for developing the necessary solvers on Graphics Processing Unit (GPU) architecture are discussed. Simulated images are presented using synthetic phantoms and phantoms generated from histology images.

Part II - Ultrasound Confidence Maps and Applications

Chapter 4: Ultrasound Confidence Maps

This chapter discusses the primary contribution of this thesis, namely estimation of confidence in ultrasound images, denoted as ultrasound confidence maps. Initially, related work is discussed together with possible applications. Subsequently, the mathematical framework is presented for graphs, random walks [38] and circuit problems to provide the context for introducing the new method. The method is discussed in detail, its ability to detect uncertainty in attenuated and shadowed regions is demonstrated and its parameter settings are evaluated. Furthermore, the benefit of integrating confidence maps into common ultrasound image processing applications is demonstrated.

Chapter 5: Automatic Bone Detection

This chapter discusses an intra-operative system developed for ultrasound (image)-guide navigation in computer-aided orthopedic surgery. The goal was to develop a framework for fusion of pre-operative Computer Tomography data set with intra-operatively acquired ultrasound images. For this a surface-based registration method was developed to transfer pre-operative data into the intra-operative reference coordinate system. The system and methods were evaluated in a cadaver study on human leg, demonstrating high robustness and accuracy.

Part III - Conclusion and Appendix

Chapter 6: Conclusion and Outlook

This chapter has the conclusion for this thesis and discusses the perspectives for future applications and improvements.

Appendix: 4D Visualization

This chapter introduces concepts for visualization of 4D medical data sets. The goal was to gain insight into approaches for visualizing motion during breathing together with information obtained from deformable registration. Volume rendering was used as the basis for volumetric visualization - as it preserved the anatomical context - and was combined with other visualization elements and approaches. In a user survey the advantages and disadvantages of different visualization approaches were analyzed.

1.3 Contributions

The main objective of this thesis was to improve ultrasound image processing by taking into account the essential physics and image formation. In the process of this thesis different contributions have been made to the modeling and simulation of ultrasound and to image processing and analysis applications. The contributions and their related publications (see also Appendix B) are as follows:

- **Realistic ultrasound image simulation on GPU hardware [69, 73, 86].**
A novel simulation platform was developed for fast ultrasound image simulation on Graphics Processing Unit (GPU) hardware. The goal was to find a balance between realism and computation time. For this purpose the nonlinear full-wave equation was used as the underlying sound propagation model and was solved with a Finite-Difference Time-Domain scheme. The solver was efficiently implemented on GPU hardware, thus, allowing highly parallel computation of wave propagation for fast image simulation. The essential ultrasound imaging pipeline was emulated in the framework including: pulse transmission, beamforming and signal processing. Ultrasound images were simulated from computer-generated phantoms and from phantom generated using histology images. In the later case intravascular ultrasound and histology images were available from the same cross-section of post-mortem coronary arteries. After deformable registration of histology and ultrasound images the real and simulated images could be directly compared for the first time in the same reference coordinate system.
- **Ultrasound confidence maps and applications in medical image processing [67, 66, 74, 132, 133, 71, 72, 73, 82].** The primary contribution of this thesis is a novel method for estimating the confidence in ultrasound images, i.e. the uncertainty in the received signal in attenuated and/or shadowed image regions, which were denoted ultrasound confidence maps. The novelty of the method lies in formulating the problem as a random walk on a graph, taking into account ultrasound transmission, depth-dependent attenuation and beam width. The analytic solution to this random walks problem allows fast computation and provides consistent confidence maps for images acquired with different transducers, systems and settings. Different ultrasound image processing algorithms could benefit from the information provided by confidence maps. Thus, we introduced confidence maps in shadow detection, 3D freehand ultrasound reconstruction, mono- and multi-modal

registration and classification in intravascular ultrasound. In all cases algorithms benefited from confidence information and produced improved algorithmic results, which were demonstrated in a series of qualitative and quantitative experiments.

- **Automatic bone detection and ultrasound-CT registration [9].** In the process of this thesis a new intra-operative ultrasound-guided navigation system was developed for computer-aided orthopedic surgery. For this a new feature descriptor was introduced that was specifically designed for bone surface detection in ultrasound images, which utilized the ultrasound confidence maps. The new feature was integrated in a new method for fully-automatic and real-time detection of bone surfaces in ultrasound images. The goal of the project was to fuse pre-operative CT and intra-operative ultrasound, which was achieved by rigid registration of bone surface extracted from ultrasound and CT. In a cadaver study the necessary Ground-Truth was collected, by implanting trackable markers into the specimen, taking CT acquisition and then ultrasound images with tracking information available for both images and implanted markers. This experimental setup allowed the evaluation of the overall system error including detection, calibration, tracking and registration errors. The results suggest the applicability of such systems in clinical scenarios.
- **Visualization of breathing motion from 4D data sets [106].** In this work different techniques were investigated for the visualization of 4D data sets, with emphasis on visualizing motion of tissue and organs caused by breathing. For this purpose classical visualization techniques were evaluated including: dynamic volume rendering, rendering of segmented organs and rendering of deformation fields with geometric primitives. Additionally, an approach for encoding motion magnitude in the transfer function of volume rendering was developed. In a user survey the advantages and disadvantages of the methods were evaluated. The findings suggest good perception of motion magnitude when it is encoded into the volume rendering transfer functions, but also suggest the need for new visualization techniques that provide consistent perceptual experience of breathing motion to different users.

Part I

Ultrasound Imaging and Simulation

MEDICAL ULTRASOUND IMAGING

In this chapter we discuss the essential physics of ultrasound and the image formation process. After an introduction to the imaging modality we will discuss image characteristics and advanced imaging capabilities (e.g. harmonic imaging). We will close this chapter by discussing different clinical applications of ultrasound and other medical imaging modalities. The major references for this chapter include the classic work of [28] that communicates an insightful understanding of wave physics and beyond. In addition, following essential references were used throughout this chapter [171], [142], [50], [17], [27] and [54]. The subsequent chapters on ultrasound simulation and modeling will build upon the overview provided here.

2.1 Medical Ultrasound

Medical ultrasound imaging is a non-invasive and non-ionizing medical imaging technology. Physicians can image in real-time and dynamically patients anatomy without exposing the patient to harmful radiation that could lead to side-effects; as long as safety features and protocols are in place. The technology operates with sound waves commonly in the frequency range of 3MHz to 10MHz, which is referred to as ultrasound. In comparison, the human hearing of a young adult has a range up to 40kHz, dogs up to approximately 60kHz and bats have a range of up to 150kHz, which allows them to navigate by producing a sound and listening for the echoes. Historically, the Italian biologist Lazzaro Spallanzani already deduced in his studies during the late 18th century that bats use ultrasound for navigation. In the industrial age ultrasound gained significant importance after the sinking of the Titanic and during World War I. Paul Langevin was a pioneer in the field and introduced a first iceberg and submarine detector (sonar) based on ultrasound generated by a piezoelectric transducer in 1917 [101]. Constantin Chilowski had also developed at that time a sonar device; however, the acoustic intensity was still too weak for naval applications [155]. In the 1920s the Russian scientist S.J. Sokolov was the first to explore the non-destructive material analysis capabilities of ultrasound for industrial applications. The first diagnostic application of ultrasound was introduced by Karl Dussik in the 1940s [24] and the first 2D image of living tissue was presented by Wild




Figure 2.1: Illustration of ultrasound imaging working principle.

[162] in 1951. Since these early beginnings ultrasound imaging has continuously evolved; becoming one of the pillars of modern medical imaging.

Diagnostic ultrasound imaging relies on the same fundamental principles as sonar and radar. More specifically, an acoustic pulse is transmitted and is reflected at the boundary of materials/mediums/objects that have different impedance. The time between pulse transmission and receiving an echo is used to calculate the distance from the source to the object, which is also referred to as the pulse-echo principle. Therefore, the distance $D = Tc/2$, where T (s) is the time between transmission and reception of the echo and c (m/s) is the assumed propagation speed in the medium. Furthermore, the amplitude of the received signals gives insight on the impedance difference between the mediums that caused a reflection, i.e., the higher the echo amplitude the higher the impedance difference. An illustration of this working principle is depicted in Fig. 2.1, where different mediums cause reflections that are detected and displayed using varying brightness values, which in turn depend on the echo amplitude. The transducer is the device for emitting and transmitting the acoustic signal. The impedance difference between transducer and medium (tissue) can cause significant reflection of the emitted signal. To avoid this a coupling gel is applied between a transducer and a medium. The gel is free of air and reduces impedance mismatch to human tissue for medical applications. This is only a basic overview of the ultrasound imaging principle. The 2D high-resolution images provided by modern ultrasound systems are the result of multiple processing steps, which we will discuss in more detail. However, before moving on to ultrasound imaging technology it is helpful to first discuss the essential physics of sound waves.

2.2 Sound Waves

Waves are a remarkable phenomenon in physics and are encountered in different contexts including the study of electromagnetic-, sound-, water-waves and oscillating systems. The foundation of our modern understanding of sound waves dates back to Isaac Newton,

Pierre-Simon Laplace and Lord Rayleigh. Sound waves are mechanical and can be understood through Newton's laws. Keeping this in mind we will now discuss in more detail the fundamental laws of sound propagation in mediums.

Sound waves are an oscillation or disturbance in a medium; not related to a particle transportation process. Starting with a simple example from [28], let's assume a stationary object in the air inside an empty room/box. This system is initially in an almost equilibrium state. Now if we move the object rapidly in the air - so that the air is pushed and does not float around the object - then we suppress the air particles in the direction of the motion and create a high pressure region. This pressure difference causes particles from the high pressure region to move to the low pressure region and give their momentum to the particles in that region, which in turn are suppressed and form a high pressure region and so forth; see Fig. 2.2 for an illustration of this process. Essentially, sound waves are the result of pressure inequalities in gases, fluids or solids. The increase in pressure by compression is heating up the medium, whereas, rarefaction of pressure oscillation decreases the temperature. This process leads to kinetic energy being converted to heat, which is one of the main reasons for the attenuation of wave propagation with increasing distance from a source. There are two types of mechanical wave propagation. In transverse waves the particles oscillate perpendicular to the direction of wave propagation, whereas, in longitudinal waves the oscillation has the same direction as the propagation. Transverse waves alter the shear stress of mediums and travel efficiently through solid materials. The variables that are related to sound propagation include the aforementioned pressure and displacement of particles, but also density of medium/material, temperature, frequency, velocity and wavelength. Speed of propagation is related to material properties and most notably to material density and compressibility, i.e., $c_s = (K_s/\rho_s)^{\frac{1}{2}}$, where K is the elastic modulus and ρ is the density. Elasticity and density are in turn affected by temperature. Furthermore, propagation speed is also related to frequency; however, this effect is fairly negligible, especially for air where speed is very nearly independent of frequency. We will discuss these variables and emphasize their relation to effects observed with wave propagation. However, we shall first choose a model as a basis for discussion.

Different mathematical models have been introduced for describing sound wave propagation. A basic derivation of the wave propagation, neglecting phenomena like attenuation, leads to the so-called wave equation that is defined for the 1D case as follows:

$$\frac{\partial^2 u}{\partial t^2} = \frac{1}{c_m^2} \frac{\partial^2 u}{\partial x^2} \quad (2.1)$$

where u is the wave amplitude, x is the spatial variable and t is the temporal variable. The interested reader can refer to [28] for a formal derivation. This equation describes a linear initial value problem [140]. Initially, the system is at equilibrium until a force is introduced (pressure difference) that affects the system over time until it eventually reaches again an equilibrium state (depending on boundary conditions). The speed of the wave propagation is expressed with the parameter c_m , which can differ for various mediums/materials. Although this is a relatively simple model it captures most of the important properties of wave phenomena and is commonly used as the basis for new formulations, including nonlinear models that are discussed in the subsequent chapters.

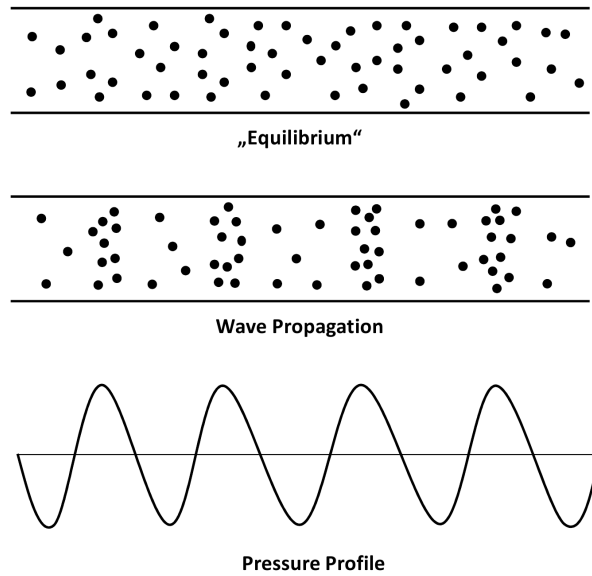


Figure 2.2: Illustration of compression and rarefaction during pressure propagation.

An alternative modeling that is very convenient for an introduction to ultrasound is to simply model sound waves as rays, which is commonly done in optics for modeling propagation of light.

When sound waves traverse from a medium to another one that has a different density and elasticity, then part of the kinetic momentum is reflected and part of it is transmitted into the adjacent medium. The higher the difference the higher the reflection is and the lower the transmitted energy/intensity. In this context acoustic intensity is defined by the sound power per area P , the propagation speed c and the medium density ρ as $I = P^2/(2\rho c)$. Interestingly the human hearing has a nonlinear perception of sound intensity, which is the reason for commonly defining intensity in the logarithmic scale of decibel (dB). In acoustics the quantity for describing different materials is called acoustic impedance and is given by the medium density ρ and the propagation speed in the medium c as $Z = \rho c$; expressed in rayls ($kgm^{-2}s^{-1}$). Table 2.1 has exemplary values for propagation speed and acoustic impedance for different mediums that are commonly encountered in medical ultrasound imaging. Fig. 2.3(a) shows an example of the reflection and transmission process. Apparently, multiple reflections of a transmitted sound wave can occur as it is traveling through different mediums. The relationship between reflection and transmission for two different mediums with acoustic impedance Z_1 and Z_2 is given by the reflection and transmission coefficients as follows:

$$R = \frac{(Z_2 - Z_1)^2}{(Z_1 + Z_2)^2} \quad (2.2)$$

$$T = \frac{4Z_1Z_2}{(Z_1 + Z_2)^2} = (1 - R) \quad (2.3)$$

A difference in acoustic impedance causes also a phenomenon called refraction. More specifically, as a wavefront reaches the boundary of two different mediums part of the

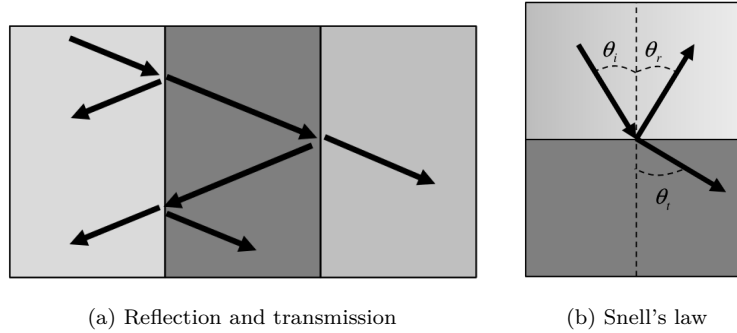


Figure 2.3: Image (a) shows successive transmissions and reflections. Image (b) illustrates Snell's law given an incoming sound beam.

Table 2.1: Speed of sound and density for different mediums

Medium	Speed of sound c (m/s)	Impedance Z (10^6 rayls)
Air	330	0.0004
Water	1480	1.48
Lung	600	0.18
Fat	1460	1.34
Liver	1555	1.65
Kidney	1565	1.63
Muscle	1600	1.71
Skull bone	4080	7.8

wavefront is accelerated/decelerated faster than the other causing a deflection of the wavefront, which is expressed by Snell's law as:

$$\frac{\sin\theta_i}{\sin\theta_t} = \frac{c_1}{c_2} \quad (2.4)$$

where c_1, c_2 are the propagation speeds in the mediums and θ_i is the angle between incident ray and boundary normal and θ_t is the angle between incident ray and opposite boundary normal (pointing to the direction of transmission). Fig 2.3(b) illustrates this process. Consequently, the higher the acoustic impedance mismatch the larger the refraction angle.

Interference refers to the spatial superposition of wavefronts and can be interpreted more intuitively when discussed in terms of phase-difference. One-dimensional waves or oscillations are commonly represented as a sinusoidal function:

$$x(t) = A \cdot \sin(2\pi ft + \phi) \quad (2.5)$$

where x is time-dependent amplitude, A is maximal amplitude, f is the frequency and ϕ is the phase. Consequently, phase is the angle of the sinusoidal function at the beginning of the wave propagation at the source. Now let's take the example of two waves that have the same frequency and are in-phase, i.e., there is no phase difference. In




Figure 2.4: Blue and red wave are constructively interfering, whereas, blue and green are destructively interfering.

this case the compression and rarefaction cycles of the waves would be identical and they would constructively overlap (interfere). On the other hand, if the waves are off-phase, i.e. their phase difference (shift) is equal to 180 degrees, then they would cancel each other out, which is referred to as destructive interference. In Fig. 2.4 an illustration is provided for constructive and destructive interference.

Scattering refers to the redirection of ultrasound incident waves to multiple and arbitrary directions and is commonly observed in presence of small reflectors and irregular material boundaries. The redirected sound waves could then be reflected again and multiple scattering can occur. We already discussed the conversion of kinetic energy to heat, which is also referred to as absorption. Attenuation refers to the combined effect of absorption and scattering, whereas absorption is a greater component of attenuation than scattering [17]. Different materials (tissues) have different scattering and absorption properties. More interestingly, attenuation has a power law dependency to frequency and is expressed as a function of acoustic energy, distance and frequency (dB/cm/MHz).

2.3 Ultrasound Imaging

In this section we will discuss the essential components of an ultrasound imaging system; as illustrated in Fig. 2.5. Modern ultrasound systems are more complex and have advanced features to improve image quality. Nevertheless, our goal here is to provide an overview of the necessary components that are part of all systems, rather than describing vendor or application specific ones.

2.3.1 Acoustic Pulse

As already discussed, ultrasound imaging systems operate by transmitting and receiving ultrasound waves. The amplitude of the emitted acoustic signal is of high importance for identifying objects and regions in tissue. Reflection and attenuation reduce the signal intensity with increasing depth. Higher initial signal amplitude can compensate to some extent this process and effectively increase the possible penetration/imaging depth. Additionally, it can contribute to emphasizing objects with small acoustic impedance mismatch, i.e. with low reflectivity. However, the higher the amplitude of the acoustic pulse

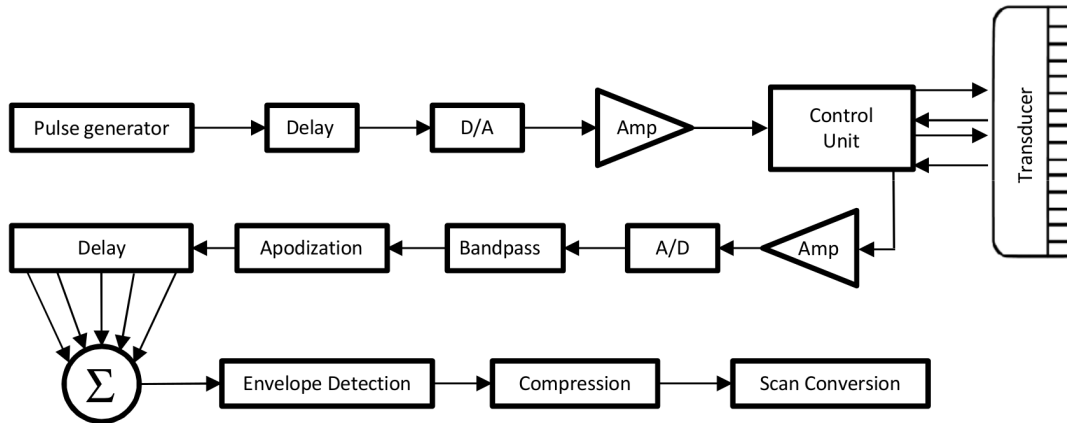


Figure 2.5: Illustration of ultrasound imaging system.

the more heat is transferred into the tissue. This factor cannot be neglected when imaging human tissue. Therefore modern ultrasound systems have build-in safety features that do not allow the amplitude of the emitted acoustic signal to exceed predefined safety thresholds.

Other than the amplitude the signal frequency is highly important. Essentially, the spatial resolution depends on the wavelength of the acoustic signal and the size of reflectors. If the wavelength of the acoustic signal is larger than the reflector then the echo is 'blurred' in the compression and rarefaction phase of the wave. Oppositely, if the signal wavelength is smaller than the reflector then the echo can be distinguished in the acquired signal. In ultrasound imaging a high frequency means better spatial resolution. However, attenuation has a power law dependency to frequency. Consequently, system operators need to adjust the frequency settings depending on the imaging scenario in order to balance image resolution and penetration depth.

Amplitude modulation is applied to the carrier frequency of an ultrasound system - commonly with a Gaussian function - to significantly reduce the sampling frequency that decreases A/D converter dynamic range [17]. Later in the reception phase demodulation of the signal is performed, commonly by taking the envelope of the signal with the Hilbert transform. An illustrative example is shown in Fig. 2.6.

2.3.2 Beam Forming

With the pulse-echo principle we can compute the distance of a reflector in relation to the sound source. However, this does not provide us with the necessary spatial information for generating ultrasound images. To further illustrate this problem lets take a look at Fig. 2.7(a), which shows a sound source that emits a 2D sound wave. The reflectors produce echo waves, Fig. 2.7(b), that are detected at the source. Given that the three reflectors have different impedance characteristics the echo responses can be distinguished and the distance to the reflectors can be computed. However, their spatial position can only be determined at a circle with radius the distance computed from the pulse-echo principle. Positioning the sound sources in a parabolic scheme, see Fig 2.7(c), results

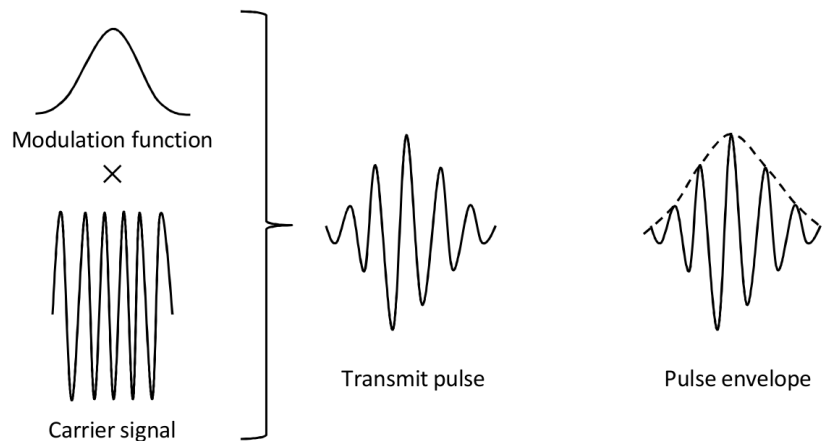


Figure 2.6: Illustration of acoustic pulses. Carrier signal is modulated to form transmit signal. In reception mode the received signal is demodulated by taking it's envelope.

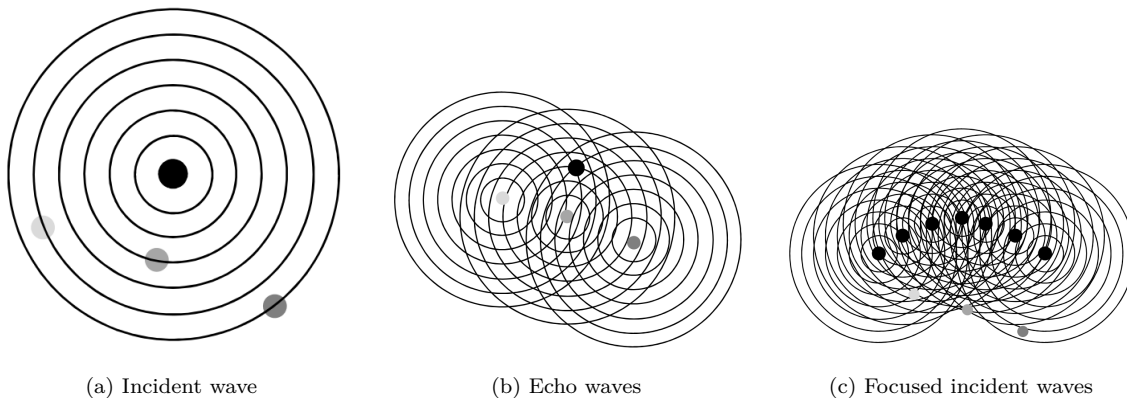


Figure 2.7: Illustration of incident (a) and echo waves (b). Parabolic placement of sound sources to generate focused wavefront (c).

in focusing of the sound energy because of constructive interference. Therefore the echo response from the middle reflector will have a stronger amplitude than the ones coming from the other reflectors. This is the essential principle behind beam-forming; a technique that is also used in radar and sonar applications and aims at concentrating the wave energy in a narrow beam. Any received echo can be associated with a particular beam direction and depth information from the pulse-echo response. Finally, spatial scanning of planes/volumes can be achieved with appropriate beam formation strategies.

2.3.3 Ultrasound Transducers

Ultrasound transducers are devices for transmitting and receiving sound pulses for ultrasound imaging. The active surface - for sound generation and reception - is commonly assembled from small piezoelectric elements. When potential is applied to these elements they contract or expand, thus, generating a sound wave. Furthermore, pressure from sound waves on these elements produces a measurable electric current for detect-




Figure 2.8: Image (a) illustration of triggering scheme, image (b) illustration of sound beam.

ing echoes. Modern systems can control each element individually, which allows beam formation. More specifically, Huygens’s principle states that any complex wavefront can be formed by combining multiple wavefronts. This principle can be used in two ways to achieve beam forming: 1) parabolic assembly of elements and 2) delayed triggering of elements. The first approach was commonly used in early transducer designs, whereas, the second approach is the standard for modern ultrasound systems. Triggering the elements with different delays, as illustrated in Fig. 2.8(a), produces a focused acoustic beam. Along the beam the constructive interference of the individual sound waves reaches maximal amplitude at a point, which is referred to as the focal point (region). Scanline formation refers to the process of scanning a plane/volume by transmitting multiple beams and processing the echo responses. Essentially, an ultrasound image or volume is formed by acquiring the echo responses for each image/volume scanline; see Fig. 2.8(b) for an example of beamforming. Therefore, for the image acquisition process a beam is transmitted, echoes are received and the process is repeated until the entire image plane or volume is acquired. The frame rate of the imaging system depends on the number of beams used, whereas, modern system allow simultaneous acquisition of multiple beams when probe geometry suggests that echo responses from far-apart beams are negligible.

2.3.4 Signal Processing Pipeline

We discussed the essential features and components of an ultrasound imaging system. To complete our picture we will now describe the remaining processing pipeline for generating ultrasound images, as illustrated in Fig. 2.5. The pulse transmission step involves the beamforming (using appropriate delays), a D/A converter step and amplification of the signal. In the reception step the echo signals are initially amplified. The pre-amplifier boosts all acquired echoes to emphasize weak reflectors, but also suppresses high signal responses to protect the A/D converter. The main-amplifier (Gain) is user controlled and allows overall boosting of signal strength to emphasize weak reflectors. Time-Gain-Compensation is a more special amplifier as it allows the user to adjust the amplification

specifically for different depths. This is especially useful when highly reflective structures are present in an image acquisition, which call for a strong amplification after the reflector but limited amplification before it. Lateral-Gain amplifiers are also available for some ultrasound systems in order to control the amplification per scanline.

After the amplification stage and the A/D conversion the signals reach the bandpass filter stage. The nonlinear characteristics of sound propagation results in echo signal responses with a different frequency spectrum than the emitted signal pulse. Commonly only a narrow band around the transmit frequency is used for imaging by applying a bandpass filter, like the Butterworth filter. For each scanline acquisition a group of transducer elements is used to form the beam - see example in Fig. 2.8(b) - but also to listen for the echoes. Before combining the echoes into a single scanline the contribution of the echoes is weighted in the so-called apodization phase. Commonly the further away the elements are from the center element of the active group the less their contribution is to the final scanline, e.g. by using a Hanning window. The delays for the beam forming are taken into account before combining all the measurements into a single scanline; commonly with the Sum-and-Delay algorithm [146]. Alternative approaches are also available [104, 156] for combining the signals of the active element group with the goal of improving image quality. The outputs of these steps are so-called RF (radio-frequency) scanlines.

The aforementioned processing pipeline produces multiple RF scanlines - usually equal to the number of transducer elements - that will form the final ultrasound image. The highly modulated echo signal is processed by taking it's envelope; commonly computed with the Hilbert transform. Dynamic range compression is applied - usually with a logarithmic mapping - to adjust the signal range to the dynamic range of the display device. After these steps the B-mode (brightness-mode) scanlines become available. Depending on the transducer type different interpolation schemes are used between scanlines in the so-called scan-conversion step, i.e., aligning the B-mode scanlines in the image plane by considering the transducer/probe geometry. This essentially concludes the basic pipeline for ultrasound image formation. However, depending on manufacture different pre- and post-processing steps are applied to enhance the image quality. Furthermore, techniques that involve complex beamforming have become the standard in many imaging platforms and will be discussed later on in this chapter.

2.3.5 Ultrasound Transducer Types

Ultrasound imaging is used in different clinical scenarios and different transducer designs improve the scanning capabilities for the anatomy at hand. Linear (Fig. 2.9(a)) and curvilinear (Fig. 2.9(b)) transducers operate similarly as they perform beamforming using small groups of transducer elements, which are triggered appropriately. The difference is - as the name suggests - the geometry of the transducer head. Linear transducers have a smaller field-of-view (FOV) than curvilinear ones and are commonly applied for example in imaging of carotid artery. Curvilinear transducers are commonly used in prenatal, abdominal and renal imaging where a wide FOV is required. Phased array transducers (Fig. 2.9(c)) are commonly applied for echocardiography (ultrasound imaging of the human heart) as they have a very small transducer head, which allows a physician to image the heart




Figure 2.9: Different ultrasound transducers for the Siemens ACUSON S3000 ultrasound system (Siemens press pictures).

through small acoustic windows, e.g., through the ribs. The small contact surface of this transducers together with the wide imaging angle is realized using beam steering, i.e., the elements are triggered in a way that the beam is scanning a fan region. An alternative imaging approach to improve the view on the heart is Transesophageal echocardiogram (TEE). There a small ultrasound transducer is mounted on an endoscope that is inserted into the esophagus of the patient. From the esophagus a clear view on the heart is possible as no reflecting bone anatomy is encountered. The disadvantages of this approach are that a patient needs to prepare for the procedure by fasting, anesthesia might be necessary, a team of physicians is required and that this imaging process takes longer than traditional echocardiography. Transrectal or transvaginal transducers (Fig. 2.9(d)) have been specifically designed for applications in prostate imaging and obstetrics. Intravascular ultrasound (IVUS) imaging is an invasive form of ultrasound imaging as the transducer is placed on the tip of a catheter (see Fig. 2.9(e)) that is inserted into the vascular system of a patient. IVUS is an important tool in diagnostic cardiology as it produces cross-sectional imaging of vessels. IVUS imaging provides valuable information on fat or plaque deposits on vessel walls; however, the invasiveness of this imaging approach is not negligible.

So far we discussed 2D ultrasound imaging transducers. Modern systems and transducers allow 3D and 4D ultrasound imaging. Wobbler transducers, Fig. 2.9(f), are essentially 1D array transducers that are mechanically steered to scan a 3D volume. The mechanical

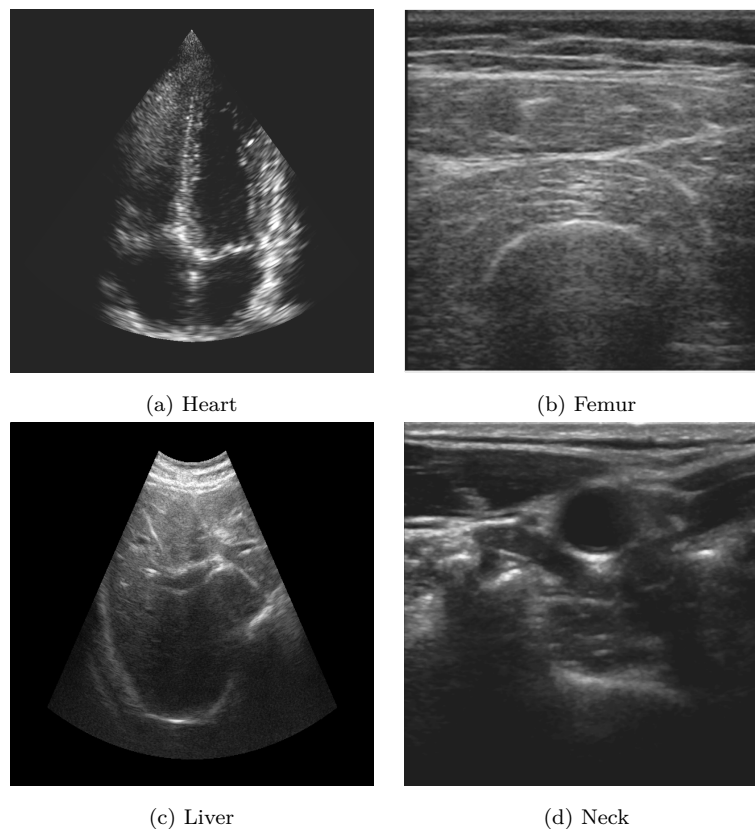


Figure 2.10: Different ultrasound images acquired with different systems from different anatomies.

steering takes time and therefore only low frame rates can be achieved with such systems. Recent high-end systems support 2D matrix transducers. In this case a 2D element matrix is used that allows beam focusing and steering in 3D space. This requires complex electronics and signal processing pipelines and thus only high-end systems of some vendors provide this imaging capabilities. Matrix array transducers allow real-time 4D imaging and are used in echocardiography to study heart motion and valve functionality in 3D.

2.4 Image Characteristics and Artifacts

In this section we will discuss ultrasound images characteristics and artifacts, taking into account the previous discussion on the principles of ultrasound physics and image formation.

2.4.1 Image Characteristics

As a basis for discussion some exemplary ultrasound images are provided in Fig. 2.10. Apparently, large specular reflectors (e.g. organ boundaries) resemble only a subset of the image content. A considerable portion of the image contains a noisy texture, which is referred to as speckle. Speckle is the result of sound scattered by small reflectors or

inhomogeneities in tissue. The echoes from scatterers are not individually detected, but rather their integrated echo response is detected, which is the result of interference effects. Scatterers can be found anywhere in human tissue and on organ boundaries. The image texture/pattern formed by scatterer responses depends on the density and the acoustic impedance of the scatterers. Both of these parameters are related to tissue properties. Thus, tissue types can be inferred from speckle patterns, which allow the distinction between healthy and unhealthy tissue types. The echoes from specular reflectors, like the vessel wall in Fig. 2.10(d), depend on the incident angle and the surface normal. In the cross-sectional image of the vessel wall no reflection is detected on surface normals that are parallel to the beam incident angle. On the other hand, echoes from scatterers do not depend strongly on the incident angle of the ultrasound beam, because sound is scattered by them to different directions.

The simple reflection-transmission model can describe the low-signal region after the highly reflective interfaces in Fig. 2.10(b-d). The strong reflection in Fig. 2.10(b) occurs because of the high acoustic impedance mismatch between tissue and bone (middle of image). In Fig. 2.10(c) the reflection is caused by the liver-lung interface, i.e. the high acoustic mismatch between tissue and air. In Fig. 2.10(d) the esophagus causes a drop of signal in this acquisition of the carotid artery. Table 2.1 provides some exemplary acoustic impedance values for comparison. Strong reflection reduces significantly the transmission and produces a so-called shadow region. Although almost no sound is transmitted into shadow regions the ultrasound images still show signal coming from these areas, as evident for example in Fig. 2.10(c) in the region after the lung-liver interface (bottom of image). These signals arise from the acquisition of echoes from surrounding tissue - which are not part of the shadow region - and are essentially noise.

In all previous image examples compression of soft-tissue occurs at the top of the image; especially evident in the abdominal and femur image. The upper tissue layers consist of skin, fat and muscle tissue, which are compressed and produce strong reflections at the beginning. This region is usually not of interest during ultrasound examinations and Time-Gain-Compensation can significantly reduce its influence on the image content. The heart and the liver image further show that blood is not visible in ultrasound images. Doppler imaging provides a way of visualizing blood flow and will be discussed later on. An important aspect to note is that ultrasound imaging is a dynamic process. The tilting of the transducer, the pressure on the skin and the in- and out-plane appearance of structures and speckle regions in images are important factors during diagnostic processes. This is also one of the reasons why ultrasound imaging requires an experienced operator to harness its full potential and perform accurate diagnosis.

2.4.2 Imaging Artifacts

The term “imaging artifacts” is very general and has been used to describe different ultrasound imaging phenomena in different contexts. For example, in medical image processing literature speckle is often referred to as an artifact or noise. However, speckle contributes most of the information in an ultrasound image and is essential to distinguish tissue types and diseases. The theoretically optimal imaging process would be able to image accurately

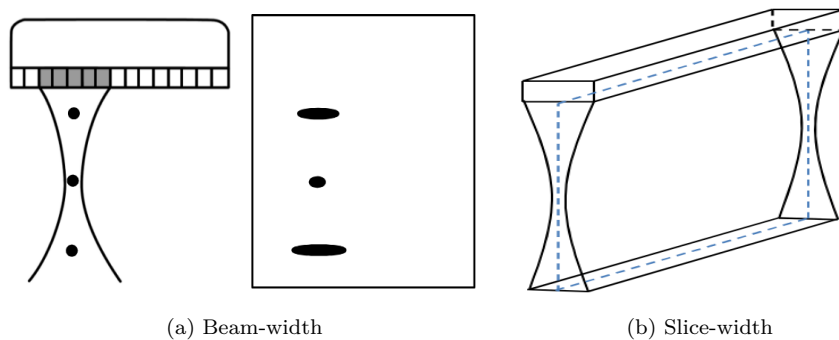


Figure 2.11: (a) Illustration of beam-width and effect on spatial resolution of reflectors, (b) illustration of slice-width.

all tissue types and allow clear distinction between them. However, ultrasound imaging is bound by the physical limitations of the imaging process, i.e. sound propagation in inhomogeneous tissue. We will therefore discuss these limitations on imaging, but also emphasize the clinical value that some of them can provide.

The most basic assumption in ultrasound imaging is the constant speed of sound for the pulse-echo principle. Tissue types have different densities and sound propagates with different velocity; meaning that distances in the images are by default erroneous. Although this error source is fairly negligible it can become significant when high deviations from the assumed speed of sound - usually 1540 m/s for most medical imaging systems - are the case. One example is imaging of breast tissue where the difference in propagation speed can cause errors in spatial information. Furthermore, the acoustic impedance mismatch can cause refraction. This might lead to anatomy displayed off the real ultrasound beam axis in ultrasound images.

Attenuation can also be considered an artifact as it reduces the penetration depth, depending on tissue types and the frequency used. However, highly attenuating regions can also indicate regions of interest, e.g., cysts and tumors have commonly different attenuation characteristics than their surrounding tissue. Hypo- and hyperechoic regions appear darker and brighter respectively compared to surrounding tissue, because of different attenuation properties. Enhancement artifacts are visible after strong hypoechoic regions, as sound was significantly attenuated in adjacent regions.

An artifact that can occur when strong reflectors are present is reverberation. In this case echoes might travel multiple times between reflectors during the acquisition of scanlines. This results in multiple detections of the echo that corresponds to the same anatomy, which is visible in the ultrasound images as a similar echo responses occurring multiple times at different depth intervals. Again, reverberation can be an artifact, but it can also be of clinical importance. Metallic objects like bullets or biopsy needles cause reverberation. These objects appear as a bright “comet” in images - which is the reason why they are referred to as comet-tail artifacts. Mirroring artifacts are very similar to reverberation artifacts; however, the echo commonly travels only between two interfaces, usually reflector and transducer. This produces a mirror image of the anatomy; an effect that can be observed when imaging the liver because of the strongly reflecting lung-liver interface.




Figure 2.12: Example of Doppler imaging of carotid artery (Siemens press picture).

Two artifacts with no clinical value are beam-width and side lobe artifacts, which are the result of imperfect beam formation process. In the ideal case the width of the focused sound beam should be very small and constant for all depths. However, this is not the case and reflectors smaller than the beam-width appear with different sizes along the beam. This is illustrated in Fig. 2.11(a), where multiple equally sized reflectors are imaged with a focused beam. An artifact similar to beam-width is the so-called slice-width. The beam has not only a width in the image plane but also off the image plane, as illustrated in Fig. 2.11(b). Consequently, reflectors that are close to the image plane will become part of the detected echo responses during image formation. Side lobe artifacts are caused by acoustic energy that is transmitted off the main beam axis. This will cause additional echoes that do not belong to structures along the main beam. Grating lobes are similar artifacts, caused by small irregularities when assembling transducer elements into an array, which are mostly avoided in modern ultrasound systems.

2.5 Advanced Ultrasound Imaging

So far we discussed the essential concepts for B-Mode ultrasound imaging. Modern systems provide different imaging capabilities for different application scenarios, which we will discuss in this section.

2.5.1 Doppler Imaging

As the name suggests this imaging technique is based on the Doppler effect, which describes the effect of a change in frequency because of movement of wave source or receiver. In the context of ultrasound imaging the Doppler frequency shift is used to detect the motion of red blood cells, starting with the work of Satomura [125]. The prevailing techniques are continuous wave (CW) and pulsed wave (PW) Doppler imaging. In the first case sound is continuously emitted and received by dedicated transducer elements. The shift in the frequency - caused by the moving blood cells - allows the detection of high blood velocities, even above 1.5 m/s. The disadvantage of this technique is that no depth or region of interest can be selected. This means that all reflections from any blood cells




Figure 2.13: Example of conventional B-mode (left) and spatial compounded image (right). Tumor boundaries and central microcalcifications are better delineated (Reprinted from [30], Copyright(2013), with permission from Elsevier).




Figure 2.14: Example of baseline B-mode (left) and B-mode after injection of contrast agent (right) (Reprinted from [32], Copyright(2013), with permission from Elsevier).

are detected and measured. Pulsed wave (PW) Doppler solves the previous problem by alternating transmission and reception. Combined with focusing techniques this allows the examination of regions of interest in specific depths. The disadvantage of this technique is the limited velocity resolution compared to CW Doppler imaging. Velocities above 1.5-2 m/s cannot be accurately detected. A common approach to visualize the Doppler information is by overlaying a color-coding of the flow direction (in relation to the probe) onto the B-mode image at the selected region of interest, as shown in Fig. 2.12.

2.5.2 Speckle Reduction

Although speckle is of high diagnostic importance it can reduce the Signal-To-Noise-Ratio (SNR) of object and organ boundaries. Different methods have been developed to reduce speckle, including spatial and frequency compounding. In spatial compounding multiple images are acquired from the same region by tilting the image plane with appropriate beamforming. A small tilt produces correlated images that are averaged to reduce speckle effects. An additional benefit of this method is that it allows to partially image structures below small reflectors that cause shadowing, as the tilt of the image acquisition planes allows partial scanning below such reflectors; see for example Fig. 2.13. Frequency compounding relies on acquiring multiple images with different imaging frequencies and




Figure 2.15: Illustration of transducer frequency response (see also [17]).

averaging them to reduce speckle.

The aforementioned techniques require a change in the image acquisition process. Alternatively, numerous techniques have been introduced for post-processing conventional ultrasound images in order to reduce speckle. Common approaches include: anisotropic diffusion filtering [118], speckle reducing anisotropic diffusion [168] and oriented speckle reducing anisotropic diffusion [87]. These methods share the concept of anisotropic diffusion filtering and differ in the way diffusion is controlled for speckle regions and object boundary regions.

2.5.3 Microbubbles

Blood vessels appear as dark regions in ultrasound images surrounded by a bright boundary (the vessel wall). Microbubble contrast agents were introduced to image vascular structures in ultrasound [20, 32]. These are injected into the patient prior or during ultrasound imaging. They are referred to as microbubbles because of their small diameter of about $2\mu\text{m}$, which is smaller than blood cells $6\text{--}8\mu\text{m}$. The bubbles are filled with air or gas and are dissolved in the liver or the pulmonary capillary bed. Ultrasound is reflected at the surface of microbubbles because of the high difference in acoustic impedance compared to their gaseous filling. Consequently, vascular structures are enhanced in the images, which can significantly improve diagnosis for different clinical applications including cancer diagnosis and echocardiography; see Fig. 2.14 for an example. Another interesting property of these contrast agents is their ability to resonate at specific ultrasound frequencies, i.e. microbubbles start acting as sound sources, producing echoes of different frequency than the ultrasound transmit frequency. Microbubbles can also be burst with high pulse amplitude, producing a short strong echo response and enhancing anatomy of interest. This feature is also of high interest for targeted drug delivery. The concept here is to enclose drugs (e.g. for chemotherapy), in the bubble, image the region with ultrasound and burst the bubbles with focused ultrasound only at the target region. This could allow higher drug dose to be delivered at target anatomy without affecting the entire human body [151].



Figure 2.16: Example of conventional B-mode (left) and harmonic imaging (right) (Reprinted from [148], Copyright(2013), with permission from Elsevier).

2.5.4 Harmonic Imaging

The frequency spectrum of the echo response received at the transducer includes frequency components that are higher than the transmit frequency; see Fig. 2.15. Additionally, harmonic frequency responses are detected that are integer multiples of the transmit frequency. These frequencies arise from nonlinear properties of tissues and scatterer response. Harmonic tissue imaging acquires and displays harmonic echo signals to produce clearer ultrasound images than conventional fundamental imaging techniques [98]; see for example Fig. 2.16. Microbubble contrast agents have also nonlinear scattering properties and produce harmonic frequency response, which can be utilized to enhance microbubble presence in images by filtering out the fundamental frequency component of the received echo signal. This is referred to as contrast media harmonic imaging. The technique is particularly useful when only the vascular anatomy is studied.

2.5.5 Highly Focused Ultrasound

We mentioned that some of the kinetic energy of sound propagation is converted to heat and that attenuation has a power law dependency to frequency. Heat transfer is a serious concern for in-vivo imaging of human tissue. Indeed, high frequency and high amplitude combined with a focused beam can cause severe burning of tissue. To prevent this ultrasound systems have build-in safety features to restrict the possible frequency and amplitude that can be set by an operator. Nevertheless, sonographers are trained to further take into account scanning of sensible areas. For example, scanning the human eye with a phased array transducer intended for echocardiography might cause considerable side-effects.

The ability of focused ultrasound to generate enough heat for ablating (burning) tissue might sound dangerous and undesired in clinical applications. However, in some cases this is a highly desired feature. Highly focused ultrasound (HIFU) is regarded as an option for future non-invasive surgery, allowing treatment of tissue through heating and ablation without open surgery [145]; see Fig. 2.17 for an illustration. The benefits are manifold and include less trauma for the patient, reduced hospital stay and only light




Figure 2.17: Illustration of lesion treatment in liver with highly focused ultrasound.

sedation or mild anesthesia. This technique has found application in treatment of tumors [78, 34, 55], uterine fibroids [149] and uterine leiomyomas [144] and is a topic of active research. Furthermore, highly focused ultrasound has been a successful tool for non-invasive treatment of kidney and gallbladder stones with ultrasound shock waves, which is referred to as lithotripsy. Shock waves have sharp and high peak amplitudes followed by an abrupt drop of wave amplitude. For lithotripsy the goal is to introduce shear stress at the focal point that can break down larger stones into smaller pieces, which leave the body through the urinary system.

2.5.6 Elastography

Tissue stiffness depends highly on the anatomical and molecular structure of a region of interest. More importantly, diseased tissues like cancer or cysts have commonly a different stiffness than their surrounding tissue. Ultrasound elastography aims at making use of this property to differentiate between tissues during ultrasound examination. This is a promising feature especially given the fact that tumors, for example in breast or prostate, are sometimes not differentiable from their surrounding tissue. In addition, if the different elastic properties of tissue types could be accurately determined then the contrast of abnormal tissue, i.e. with elastic properties that are different from the surrounding tissue, could be enhanced; see for example Fig. 2.18. The stiffness parameter is a function of elastic modules and can therefore only be measured if a stimulus/pressure is introduced into the tissue. Afterwards the tissue motion is measured that depends on the stimulus and the elastic properties of the tissue [114]. There are different approaches for introducing the necessary tissue compression for elastography [100]. In static compression elastography a controlled compression is introduced with a mechanical aperture and ultrasound is acquired before and after compression. This limits the flexibility of such imaging systems in clinical environments. In quasi-static compression elastography the probe/transducer is used to introduce the compression and is still moving during the acquisition of the images. This approach is very flexible and does not required change in the imaging system,




Figure 2.18: Example of conventional B-mode (left) of carcinoma and corresponding elastogram imaging (right) (Reprinted from [13], Copyright(2013), with permission from Elsevier).

however, the uncontrolled nature of the compression and the inhomogeneity of the pressure field during freehand acquisition make accuracy and repeatability an issue. Sonoelasticity imaging is very similar to static imaging in terms that a vibrating actuator is applied that introduces compression to the tissue. In radiation force elastography the pressure is not introduced by an external mechanical motion, but rather generated by a focused wavefront of high amplitude that is transmitted into the tissue. The compression can be applied in a quasi-static mode or in a continuous vibration mode similar to sonoelasticity imaging. Elastography imaging is still not widespread in clinical routine, but is likely to become part of it given the advancements in this field and the diagnostically important information that could be obtained.

2.6 Clinical Application Fields

The real-time and dynamic character of ultrasound image acquisitions is highly important for diagnosis. Depending on the viewing angle on the target anatomy, the movement of the probe and the pressure on the tissue the image generated from the underlying anatomy changes. This allows experienced sonographers to draw conclusion on the tissue properties. Real-time imaging is important for assessing dynamic anatomy such as heart. Assessment of heart muscles, valves and blood perfusion is essential for diagnosis of cardiac diseases. Doppler imaging is an ultrasound-specific imaging feature that allows real-time measurement of blood flow. Vascular ultrasound imaging, like for carotid artery, makes extensive use of Doppler imaging to assess possible stenosis in vascular anatomy. Prenatal and obstetrics ultrasound imaging is widely known and used for early diagnosis of diseases and possible complication during pregnancy and for follow-up examination. Ultrasound imaging is especially well-suited for this application because of the lack of ionizing radiation. Side-effects from the use of ultrasound for prenatal imaging have not been recorded, as long as safety features and protocols are in place. Abdominal imaging with ultrasound is also very common for assessing functionality of liver, kidney and gallbladder. Breast imaging with ultrasound has become a routine for early detection

of cysts and carcinomas and acts complementary to mammography. Prostate imaging is not only used for diagnosis, but also for cancer treatment options like brachytherapy. Transcranial ultrasound is a relatively new approach for diagnosis of brain diseases like Parkinson [2]. Intravascular ultrasound is an invasive ultrasound imaging technique as the scanning sensor is placed on the tip of a catheter that is inserted into the patient vascular system. It provides cross-section views of vessels and allows distinction of tissue types like plaque, fatty or even vulnerable plaque that is likely to cause acute stenosis or stroke in case of rupture. The high portability of ultrasound systems together with little preparation for patient imaging make it a favorable choice for intra-operative imaging and imaging in emergency situations. Furthermore the comparatively low acquisition and maintenance cost of ultrasound systems have led to widespread use of them in small ambulances. In final consideration, ultrasound imaging is a very important modality in clinical routine and is constantly being improved and extended with new imaging capabilities.

2.7 Other Medical Imaging Modalities

We will briefly discuss the working principles of other imaging modalities to provide the context for the next chapters. Furthermore, we will emphasize the specific advantages and disadvantages of ultrasound compared to other modalities.

2.7.1 X-ray Imaging

X-rays are electro-magnetic waves of high frequencies of about $10^{18} - 10^{20}$ Hz [28] that can be produced in X-ray tubes by hitting a metal target with fast-moving electrons. Furthermore, γ -rays are emitted by radioactive nuclei and have a high frequencies of about $10^{21} - 10^{24}$ Hz. The X-ray tubes consist of a negative (cathode) and positive (anode) electrodes enclosed in vacuum. Heating up the filament frees electrodes that are repelled by the cathode and attracted by the anode and hit the metal target at a speed close to half the light speed, producing X-rays. Both X-rays and γ -rays can and commonly will interact with matter. Transmission refers to the part of the rays “passing through” the material without affecting it. Absorption refers to the process where the energy of a X-ray photon is transferred to the matter, i.e., the photon disappears at this point. Scattering is also an important effect where photons partially transfer their energy to matter and are diverted from their incident direction. For X-ray imaging photons are emitted and detected after passing through a medium. Different tissue types and anatomy have different absorption and scattering characteristics. This property is essential for X-ray imaging as the patient is placed between a X-ray source and a detector (commonly semiconducting diode). Therefore the amount of photon attenuation (absorption and scattering) is essentially what is represented in X-ray images. In Fig. 2.19(a) an exemplary X-ray image is shown, where bright regions account for high attenuation and dark ones for low attenuation. The contrast or color-coding of the image can be chosen arbitrary. Nevertheless, different conventions have prevailed and are used in clinical practice to provide a common ground for discussion. In fluoroscopy imaging X-ray images are generated in real-time

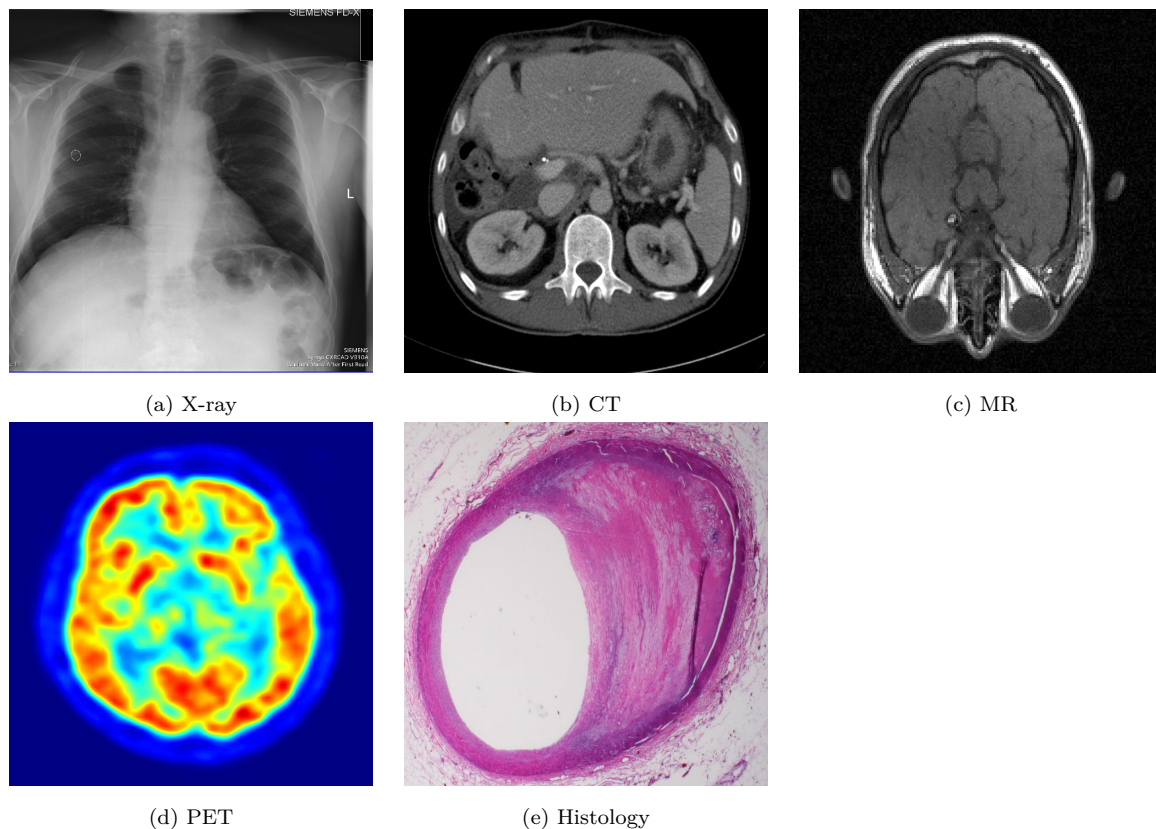


Figure 2.19: Image from different modalities (a) X-ray of chest (Siemens press picture), (b) CT of abdomen, (c) MRI of brain, (d), PET of brain, (e) histology of vessel.

with a phosphor screen that produces a brightness response proportional to the received photon energy.

2.7.2 Computer Tomography

X-ray images of patient anatomy are the projection of the attenuated X-rays. Computer tomography imaging makes use of the mathematical concept introduced by Johann Radon, which states that two- and three-dimensional objects can be reconstructed from an infinite set of projections of the objects. For this a finite set of X-ray projections are acquired around the patient in a CT scanner to reconstruct an axial image of the anatomy. In Fig. 2.19(b) a CT slice from abdominal region is shown that images the liver, kidneys, bones and other anatomy. The acquisition process can be performed for multiple axial slices that together produce a coherent 3D volume acquisition of patient anatomy. Each voxel (3D pixel) corresponds to the average attenuation of X-rays that passed through this voxel from all 2D X-ray projections. The most common reconstruction technique for CT is filtered back-projection in which the attenuation values are projected back to the sources used for acquiring the projections. Different filter kernels can be used for the back-projection process and high-pass filters are commonly applied to reduce star-like artifacts. Different CT scanner generations were introduced including moving-

tubes/ moving-detectors, moving-tubes/ stationary-detectors, parallel beam and fan beam systems. The latest generations of scanners are commonly helical scanners and allow the continuous movement of the patient bed through the scanner as the X-ray tube is rotating. This significantly reduces acquisition time and allows imaging of anatomy with only one breath hold.

2.7.3 Magnetic Resonance Imaging

Every proton has a magnetic momentum that is defined by the external magnetic field and the gyromagnetic ratio of the proton. In Magnetic Resonance Imaging (MRI) the system generates a quasi static strong magnetic field commonly in the range of 1T-3T for clinical applications. The magnetic momentums of protons in the human body align themselves with the field once a patient is placed in the magnetic field of the MRI scanner. The system introduces a high-frequency resonance pulse that flips the spin of the protons. In the so-called relaxation phase the magnetic momentum re-aligns itself with the magnetic field after some time, i.e., the relaxation time. The relaxation causes a change in magnetic field that can induce a measurable current on the receiver coils of a MRI system. Relaxation is observed as a damping of the received signal and different signals are acquired given that different protons exist in the human body. However, the most prominent element is water and therefore mostly the spectral signal from the hydrogen proton response is taken into account for the imaging. In Fig. 2.19(c) a MRI slice of brain is shown. Soft-tissue with high water content is clearly visible, whereas, bone anatomy is barely shown but rather derived from the image context.

2.7.4 Gamma Imaging

In nuclear medicine radiopharmaceuticals are intravenously injected into the patient. The pharmaceutical has the property to concentrate on different organs or tissues of interest and the radionuclide combined with it has the function to signal the position of it by emitting γ -rays. The emission of the γ -ray is a stochastic process, the importance here lies in the detection of the γ -photons by γ -cameras. The cumulative reception of such photons from regions of interest is the diagnostic information provided by gamma imaging systems. There are two different imaging concepts, Single-Photon Emission Computed Tomography (SPECT) and Positron Emission Tomography (PET). Both techniques detect γ -rays produced by radioactive tracers. In SPECT imaging the gamma radiation is measured. In PET imaging the radioisotope of the tracer undergoes a positron emission decay, emitting a positron that interacts with an electron in the patient and annihilates both of them producing γ -rays of opposite direction. These γ -rays can again be detected with γ -cameras and the location of the annihilation point can be determined on a line. An example of a PET image from brain is shown in Fig. 2.19(d), where uptake of radio-pharmaceutical is displayed.

2.7.5 Histology

Histology refers to the study of tissue samples commonly under light or electron microscope. Histopathology is highly important in final diagnosis of diseased tissue at microscopic level, like malignancy of cancer. Histology imaging is performed on ex-vivo tissue that is extracted from patient or from cadavers. Before imaging tissue samples under microscope they need to be prepared. This involves a fixation process, which can be either chemical or cryo-fixation in order to preserve tissue structure. Subsequently, tissues are dehydrated (commonly with ethanol baths), cleaned (to remove the alcohol) and fixed with paraffin wax. These are then embedded in external molds with embedding material and hardened. In the next processing steps the prepared tissue samples are sectioned (frozen samples with microtome), to allow light and electron microscopy. Tissue has a low contrast by itself and staining agents are introduced, like the commonly used Hematoxylin (blue stain of nuclei) and eosin (pink stain for cytoplasm) - also known as H&E stain. An exemplary histology image of coronary artery is shown in Fig. 2.19(e).

2.8 Ultrasound and other Modalities

Ultrasound and MRI imaging are favorable for scenarios where ionizing radiation should not be applied, like for obstetrics. The advantages of ultrasound systems are their mobility and ease of intra-operative use compared to MRI. On the other hand, MRI provides a volumetric representation with good soft-tissue contrast. Ultrasound has limited depth penetrations capabilities and is for example not applicable for imaging lung anatomy, because of the high acoustic impedance mismatch. Fluoroscopy and CT imaging use ionizing radiation, which can have short- and long-term side-effects. Nevertheless, when it comes to imaging bone anatomy both modalities provide superior image contrast - especially CT - compared to ultrasound where only the surface of bone can be imaged. Fluoroscopic imaging has become an essential tool for different intra-operative procedures including cardiac catheterization and stent placement. These cannot be performed with ultrasound imaging because of limited depth penetration and restricted choices for acoustic windows. PET and SPECT imaging are functional imaging modalities with radioactive tracers concentrating on different tissues of interest. Ultrasound contrast agents are further developed and could provide functional imaging in the future, whereas real 3D imaging would still only be possible with modalities that allow volumetric acquisition. It becomes apparent that there is no “best” imaging solution. Moreover, the use of a medical imaging modality depends on the clinical scenario and of course on practical issues including availability, space- and time-restrictions. Instead of choosing only one modality clinicians commonly combine the complementary imaging information from different modalities to assess diseases. Hardware developers are starting to provide systems that perform complementary imaging like PET-CT scanners. Moreover, in this scenario the imaging data from CT improves the reconstruction process of PET imaging as it is used to account for tissue-specific attenuation. Furthermore, combining ultrasound and other modalities has been of high interest and will be discussed in applications of subsequent chapters.

SIMULATION

Computer simulations are an important tool for testing theories and models in different scientific fields. They are essential for modern engineering and allow rapid experimentation without the need for cost and time consuming setups. Ultrasound simulation has played a crucial role in the development of new ultrasound transducers and imaging systems. Furthermore, it has contributed to our understanding of complex propagation of sound in inhomogeneous mediums like human tissue. In this chapter we will discuss ultrasound simulation and introduce a new framework for fast simulation on Graphics Processing Unit (GPU) hardware. The framework uses the nonlinear full-wave equation for modeling sound propagation and implements the essential ultrasound image formation pipeline from pulse generation to B-mode image formation. We will demonstrate the results obtained for simulating B-mode images from computer phantoms and phantoms generated from histology.

3.1 Introduction

The continuous research in the field of ultrasound imaging has provided us with an improved understanding of the relatively complex physics of sound propagation and has led to the development of new imaging technologies. Nevertheless, some phenomena remain a mystery and although we understand the possible reasons that might led to a measurable result, we still lack the mathematical tools to accurately model, simulate and predict different behaviors. For example, HIFU seems to be the future of non-invasive surgery, allowing the ablation of tissue inside the body without open surgery [78]. HIFU treatment allows not only access to target regions without open surgery, but also access to regions for which conventional surgical access is not possible or too dangerous. For example, treatment of brain cancer would highly benefit from HIFU in order to reach a target area without damaging healthy and critical tissue. Although our understanding of ultrasound propagation has significantly improved, the effects of nonlinear sound propagation are still investigated. For HIFU treatment this effectively means that the predicted focal point for treatment might deviate from the real one. Development of appropriate models and simulations for such complex processes could further improve the therapeutic outcome for

patients. In addition, the influence of different imaging settings and tissue properties on ultrasound imaging can be effectively studied with fast and realistic simulation platforms. Finally, ultrasound simulation has been a complementary tool for validating new mathematical models, experimenting with system designs before moving on to manufacturing and for broadening our understanding of ultrasound. As such it has received considerable attention in the research community.

As with any simulation platform the trade-off between realism, accuracy and computational requirements needs to be considered. Simulation platforms that provide results comparable to experimental measurements have been developed and are used throughout ultrasound research and development. However, these platforms may require considerable processing time, which can be problematic for rapid development and experimentation. On the other hand, real-time simulation applications have emerged that perform fast simulation at the cost of realism. The goal of this work was to develop a simulation platform that balances realism and simulation time by using a realistic sound propagation model that is solved in parallel on GPU hardware.

3.2 Related Work

Time-Domain-Diffraction methods allow fast computation of pulse response for short pulses and required only a single convolution operation. However, the necessary integrals could only be defined for simple transducer geometries [142]. More general approaches were introduced that could model arbitrary transducer shapes and apodization schemes. The software package Field II by Jensen [57, 59, 60] is still considered a reference for realistic ultrasound simulation. The framework is based on the computation of the spatial impulse response, where the relative position of transmitter and receiver is taken into account. The approach assumes a homogeneous medium and attenuation and linear wave propagation. Although analytical solutions are available for simple transducer geometries to simulate any geometry, the aperture is divided in small elements, such as triangles or rectangles [58] that are assigned individual apodization weights. The framework requires considerable processing time and resources [60, 129]; however, the increasing availability of high performance computational resources is likely to counterbalance this issue. The Khokhlov-Zabolotskaya-Kuznetsov (KZK) equation models the combined effects of diffraction, nonlinearity, and absorption [98]. It can be solved both in the time and the frequency domain and allows simulation of complex scenarios like harmonic imaging. Furthermore, it has shown good correlation with experimental measurements [56]. A comparison to simulation results obtained with Field II and the KZK equation is presented in [121]. The simulation package Abersim [31] solves the KZK equation with an operator splitting approach, with the diffraction term solved by means of the Angular Spectrum method. Abersim allows linear and nonlinear propagation simulation and attenuation with arbitrary frequencies.

The Westervelt equation, also referred to as the nonlinear full-wave equation, describes wave propagation in a thermoviscous medium [161]. The equation can be efficiently solved with a Finite-Difference Time-Domain scheme and was used in different simulation scenarios, including HIFU simulation [47] and propagation of ultrasound through cranial

bone [19]. In [120] a modified nonlinear full-wave equation is presented that can model arbitrary frequency dependent attenuation. Simulated images are presented; however, details on the Radio-Frequency (RF) processing for the image formation were omitted. Interestingly, a derivation of Perfectly-Matched-Layer boundary conditions was provided. More specifically, different boundary conditions can be defined for solving partial differential equations. The numerical simulation of wave propagation produces reflections at the boundary of the simulation grid that are undesired as they affect the measurements in the target region of interest. The Perfectly-Matched-Layer boundary condition allows absorption of the wave at the boundary and can limit the necessary grid size for the simulation [120]. Simulation and experimental measurements of the KZK and Westervelt equation showed high correlation [56].

Finite difference (FD) and finite element (FE) methods evaluate the wave in a local neighborhood over time. Spectral methods like the K-space method solve the wave propagation problem in the frequency domain [102]. Although they require multiple computations of the Fast-Fourier-Transform (FFT) the overall computational cost and storage requirements compared to FD and FE methods is lower. However, they have limitations in terms of accuracy compared to FD and FE methods [102]. More recently, convolution-based methods have been developed that convolve a precomputed point-spread-function (PSF) for a given beam profile with with a scatter distributions to simulate ultrasound echo response [33, 52]. Compared to spatial impulse response methods, such as Field II, these methods are an order of magnitude faster, allowing real-time simulation. For different settings and beam profiles new PSFs need to be pre-computed. Although this would increase the overall size of the PSF Look-Up-Table this disadvantage has no practical implications. In the work of [23] scatterers were defined for pre-segmented tissue of a CT data set including fat, soft-tissue and bone regions. Acoustic impedance values were assigned to segmented regions and a depth-dependent PSF was utilized for simulation. It is the subject of future research to investigate the accuracy of such simulation approaches and compare them to nonlinear full-wave simulation platforms.

A fast simulation method for generation of ultrasound images was presented in the work of [157]. In this work a similarity measure for registration is evaluated between the real ultrasound images and ones simulated from CT. The realism of the simulated ultrasound was only important in terms of providing the necessary information for multi-modal image registration. A similar approach was introduced in [129], that was based on the same ray-based reflection-transmission model. In addition, a scattering volume was pre-computed from a CT volume to provide realistic scatterer noise. Apodization with a Hanning window improved the perceptual realism of the images as beam-width artifacts could be emulated. Efficient GPU implementation allowed real-time 3D simulation [93, 92] and further made computationally intensive non-rigid registration frameworks [159] possible. These approaches model sound waves as rays, decreasing processing time but also significantly reducing realism of simulation. Ray-based simulation approaches are not able to fully model the complex effects modeled by wave-based approaches. These effects include interference, scattering, diffraction etc. which are common in medical ultrasound propagation and contribute tremendously to the formation of the final ultrasound image. It is also worthwhile mentioning that these methods perform ray-casting, i.e., for each

virtual transducer element a single ray is casted and the reflection and transmission is computed along it. The effect of multiple reflections could be modeled with ray-tracing, i.e., tracing the ray as it is reflected in the medium, however, this would further increase the processing time.

In this work we focus on the realistic and fast simulation of ultrasound images. The goal is to provide a fast simulation platform that can model the entire image formation process from the initial pulse transmission to the final image formation. For this purpose the nonlinear full-wave equation is explicitly solved with a Finite-Difference Time-Domain scheme. In order to achieve fast simulation times the Finite Difference scheme was implemented on the Graphics Processing Unit (GPU), which has already demonstrated its potential for accelerating parallel computations and efficiently solving these schemes [49, 90, 1].

3.3 Nonlinear Full-Wave Model

The propagation of ultrasound waves and their interaction with different media was modeled using the Westervelt Partial Differential Equation (PDE), also referred to as the nonlinear full-wave equation [47, 56]. It describes the propagation of waves and additionally models thermal attenuation and nonlinearity. The reader interested in an accuracy analysis of the Westervelt equation is referred to Huijssen et al. [56], which includes comparisons to an analytical solution, the KZK equation and water tank measurements. The Westervelt equation is given as follows:

$$\nabla^2 p - \frac{1}{c_0^2} \frac{\partial^2 p}{\partial t^2} + \frac{\delta}{c_0^4} \frac{\partial^3 p}{\partial t^3} + \frac{\beta}{\rho_0 c_0^4} \frac{\partial^2 p^2}{\partial t^2} = 0, \quad (3.1)$$

where p [Pa] is the acoustic pressure, c_0 [ms^{-1}] is the propagation speed, ρ_0 [kgm^{-3}] is the ambient density, δ [m^2s^{-1}] is the diffusivity of sound, and β is the coefficient of nonlinearity. Thus, the first two terms are identical to the D'Alembertian operator on p , the third term is the loss term due to thermal conduction and the fourth term describes the nonlinearity.

The Westervelt equation is numerically solved with the Finite Difference method [18, 49, 90, 140]. The basic idea behind this method is to evaluate the PDE equation, more specifically calculate the wave amplitude, on sampling points of a computational grid. The grid can have complex shapes, while we use a regular 2D grid with equidistant sampling points. For the calculation of the wave amplitude the partial derivatives of the equation are substituted with their finite difference representations and the equation is solved for the future time-step. A thorough analysis of higher order Finite Difference schemes for solving the acoustic wave equation is presented by Cohen and Joly [18]. Fourth-order accurate in space and second-order accurate in time schemes have demonstrated good results [18, 47], and are therefore used in our work. The finite differences for equation

(3.1) are given as follows:

$$\frac{\partial p}{\partial t} \approx \frac{p_{i,j}^n - p_{i,j}^{n-1}}{\Delta t}, \quad \frac{\partial^2 p}{\partial t^2} \approx \frac{p_{i,j}^{n+1} - 2p_{i,j}^n + p_{i,j}^{n-1}}{\Delta t^2}, \quad (3.2)$$

$$\frac{\partial^3 p}{\partial t^3} \approx \frac{6p_{i,j}^n - 23p_{i,j}^{n-1} + 34p_{i,j}^{n-2} - 24p_{i,j}^{n-3} + 8p_{i,j}^{n-4} - p_{i,j}^{n-5}}{(2\Delta t)^3}, \quad (3.3)$$

$$\frac{\partial^2 p}{\partial x^2} \approx \frac{-p_{i+2,j}^n + 16p_{i+1,j}^n - 30p_{i,j}^n + 16p_{i-1,j}^n - p_{i-2,j}^n}{12\Delta x^2}, \quad (3.4)$$

$$\frac{\partial^2 p}{\partial y^2} \approx \frac{-p_{i,j+2}^n + 16p_{i,j+1}^n - 30p_{i,j}^n + 16p_{i,j-1}^n - p_{i,j-2}^n}{12\Delta y^2}, \quad (3.5)$$

where i, j are the axial and the lateral indices of the discrete computational grid, n is the time-step, $\Delta x, \Delta y$ are the spatial discretization steps and Δt is the temporal discretization step. Thus, an explicit solution is calculated for each sampling point based on the wave amplitudes at sampling points of the previous time-steps. To model the interaction of the waves with heterogeneous media, different coefficients are used for the speed of sound, the ambient density, the diffusivity of sound and the nonlinearity.

3.4 GPU Computing

The Finite-Difference Time-Domain (FDTD) method for solving the nonlinear full-wave equation requires only local computations, i.e., the wave amplitude at every point for the next time-step can be computed from the amplitude at grid neighborhoods of previous time-steps. This local dependence in the calculations is ideal for efficient implementation of the method on parallel computing architectures. Here we will briefly discuss the difference of Central Processing Unit (CPU) and Graphics Processing Unit architecture and focus on the later one for efficient implementation of the FDTD method.

CPU architectures are designed for general computing and efficient processing of single and sequential tasks, using features like branch-prediction, out-of-order execution and high frequency scaling [95]. The increasing demand for parallel data processing led to the development of multi-core CPU architectures. Crucial aspects for the development of any microprocessor architecture are size, power consumption and heating properties. The features of CPU architectures together with the required cache memory are limiting factors for the development of multi-core systems. On the other hand, GPU architectures were designed and development for parallel processing/rendering of pixels in graphics applications. Currently the latest commodity graphics card from NVIDIA ¹ the GeForce GTX 690 has a total of 3072 cores that can simultaneous process data. This gives GPU architectures a major advantage over CPU for parallel processing. The GPU cores have less shared memory, less features and lower operation frequency than CPU cores. This reduces the die size and number of transistors required for implementing this massive core architecture, at the expense of single-thread performance. Therefore, the efficiency of each computing architecture depends on the application/algorithm and the possible

¹NVIDIA Corporation, CA, USA

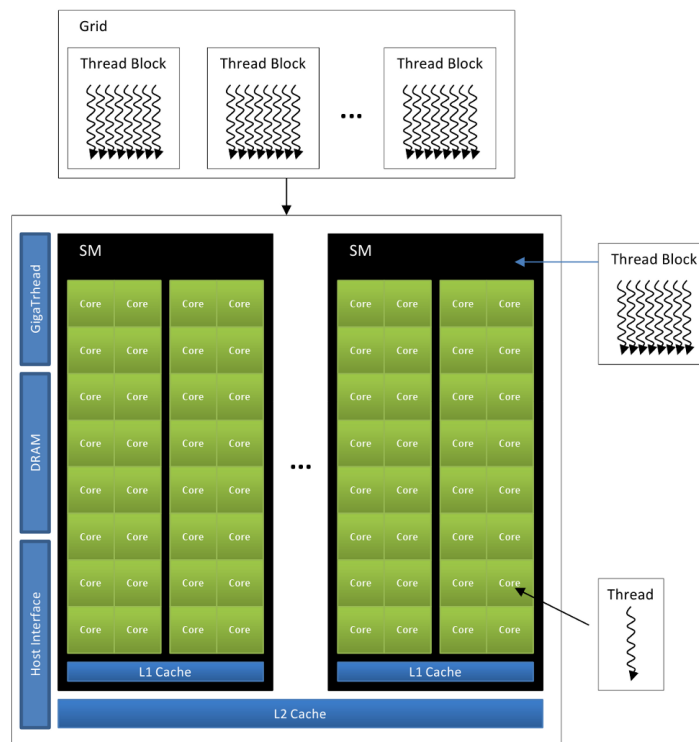


Figure 3.1: Illustration of CUDA programming model and Fermi architecture.

optimization of the implementation [95]. In general, the GPU architecture is favorable for highly parallel algorithms, whereas, the CPU for complex single-thread tasks.

We will briefly discuss the CUDA parallel computing model and Fermi architecture developed by NVIDIA [109] to further illustrate the architecture benefits of GPU for parallel computing. The design concept aims at a scalable parallel architecture, where hundreds or thousands of threads could be executed simultaneously. For this threads are grouped into blocks that share a small and efficient memory space and can communicate and synchronize with each other. Thread-blocks are grouped into grids that share a global, less efficient memory. The blocks need to be independent to provide scalability on GPU hardware that can use different number of cores for execution. In terms of hardware, thread-grids are executed sequentially or parallel by a GPU and thread-blocks by multiple Streaming Multiprocessors (SM). These execute threads in groups of 32 on the latest Fermi architecture [109] and each SM has 32 CUDA cores for the parallel processing. Each core is responsible for the thread execution and includes an arithmetic logic unit (ALU) and floating point unit (FLU). Each SM has an L1 cache (shared memory) and all SM share an L2 cache that is used for the interface with the independent DRAM interfaces. Fig. 3.1 provides an overview illustration of the programming model and architecture. This architecture is scalable in hardware as it allows the manufacturing of boards for different application areas with different number of cores and SM's without the need for changing the processing model.

The architecture of GPU devices was initially utilized by the research community to accelerate time consuming computations that could be parallelized. This was made

possible with the introduction of programmable graphics pipelines [99]. Since then the flexibility and features of the programmable pipelines have been extended and alternative shading languages were introduced. The purpose of these languages was to provide developers with Application-Programming-Interface (API) to define custom graphics processing directly on GPU boards; effectively allowing them to implement new methods for vertex and fragment processing. General-Purpose-Graphics-Processing-Unit (GPGPU) approaches utilized the programmable pipelines and shading languages to perform parallel computation instead of graphics rendering [115]. A typical GPGPU program uses the fragment shader for computation. For this the data input and output can be defined as textures that are up-/downloaded to/from GPU memory. The fragment shader computes the output of a fragment/pixel/texel independently and in parallel on GPU and has access to global memory (textures) and basic input variables (floating point values and matrices). After the execution of the fragment shader the output can be saved in the GPU fragment buffer or into a texture for further processing.

The introduction of C-like languages for GPU programming revolutionized further the GPU computing field. The CUDA programming toolkit introduced by NVIDIA provided developers with a high-level C-like language for GPU computing [108], taking into account the scalable and parallel architecture. Instead of mapping algorithms to the graphics processing pipeline, developers could write code in familiar language environments like C and Fortran. Additionally, the introduction of the OpenCL specification [138] allows heterogeneous parallel computing on different devices that implement the specification. Overall, the number of applications utilizing GPU computing has increased exponentially in the last years. The increase of data and the continuously increasing demand for high-performance parallel computing is likely to act as a catalyst for further developments in this field.

3.5 Ultrasound Image Simulation

In our framework we implemented the basic imaging pipeline to synthesize B-mode ultrasound images. Other than the wave propagation, the image simulation process involves transmission and reception of ultrasound pulses, beamforming and processing of the resulting echoes for forming the final image.

3.5.1 Ultrasound Transmission and Reception

For each simulation run sound waves are emitted at selected points on the computational grid. Various transducer geometries can be simulated by selecting the appropriate points on the grid, with a linear and IVUS transducer modeled in this work. Modifying the wave amplitude at these points introduces a wave disturbance that propagates through the grid with the simulation of consecutive timesteps. The shape of the emitted pulse is of crucial importance. Non-modulated sinusoidal or Gaussian shaped pulses for instance can cause grid disturbances even after the pulse transmission has ceased. We use a 6 cycle sinusoidal pulse modulated by a Gaussian shaped envelope, commonly used by ultrasound

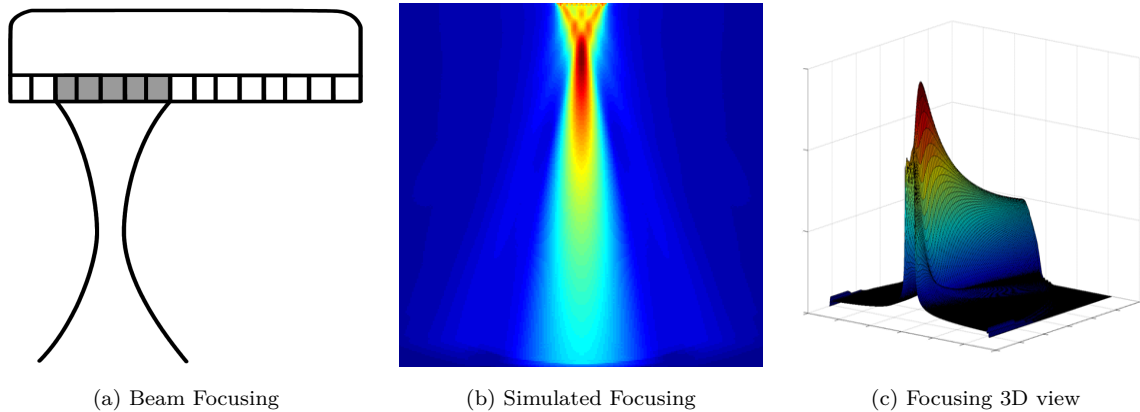


Figure 3.2: Image (a) schematically demonstrates ultrasound beam focusing by triggering a group of elements with different time delays. Image (b) shows the maximum wave amplitudes of a simulation run using a focus scheme for low depths in a medium with uniform speed of sound and image (c) its corresponding 3D visualization.

systems [50]. The pulse is formulated as follows:

$$s(t) = \left(A \cdot \sin\left(\frac{2\pi}{l}t\right) \right) \left(\alpha e^{-\frac{(x - \mu)^2}{2\sigma^2}} \right), \quad (3.6)$$

where $s(t)$ is the pulse amplitude at timestep $t \in [0..l]$, l is the pulse length, A is the maximal pulse amplitude, α is the Gaussian amplitude, μ the mean, and σ is the standard deviation. The echoes are recorded for each timestep at the positions where the grid was perturbed by the pulse transmission. Ultrasound images are formed from multiple scan lines with the total number of scan lines playing an important role for the overall spatial resolution of the final ultrasound image. Each scan line in the ultrasound image corresponds to the echoes received along an ultrasound beam, which brings us to topic of ultrasound beamforming.

Beamforming refers to the constructive/destructive interference of waves emitted by multiple transducer elements. Triggering a group of transducer elements at a time results in high acoustic intensities along the center axis of the group. Fig. 3.2(a) schematically shows the process of focusing using a group of elements and Fig. 3.2(b)-(c) show the simulation run, with our framework, for beam focusing at a low depth. Generating a narrow beam is desirable in ultrasound system development as it improves the spatial resolution of the ultrasound scan. Additionally, triggering the transducer with the appropriate time delays allows to position and steer the beam [50].

Multiple scan lines are acquired for the image formation by transmitting a beam, receiving the echoes and moving the active element group until the entire transducer element surface is covered.

3.5.2 Radio Frequency Processing

The result of the previously described simulation is Radio Frequency (RF) data acquired for each element of an active element group for each scan line. Before forming an image,

the raw data needs to be processed. The RF processing pipeline varies slightly between different ultrasound system vendors, but the basic principles are common and are implemented in this framework.

An ultrasound scan line is formed by combining the RF data acquired at each element of an active element group. For this we apply the Delay and Sum beamforming algorithm [146], which can among others be defined as:

$$d(t) = \sum_{i=0}^N W_i \cdot e_i(t + \tau), \quad (3.7)$$

where $d(t)$ is the final signal response of an ultrasound beam at timestep t , N is the number of active transducer elements, W_i is a weighting function (in our case a Hanning window), $e_i(t)$ is the signal response of each active transducer element i , and τ is the transmission delay expressed in timesteps.

The received signal contains noise that is mainly manifested in the low and high frequency parts of its spectrum [142]. Therefore, the low and high frequency components are removed with a bandpass filter, in our case a Butterworth filter. Afterwards the signal goes through the process of demodulation, which results in a signal that retains its overall pulse response but contains much less high frequency modulations. Demodulation is performed by finding the envelope of the rectified signal and is implemented by taking the absolute of the Hilbert Transform of the signal.

As a pulse traverses through the medium it is attenuated and reflectors at greater depth appear weaker than reflectors at smaller depth. This is compensated by applying a so-called Time-Gain Compensation (TGC), which amplifies echoes based on their reception time (depth). In our implementation the signal is convolved with a simple linear function $f \in [1..n]$, where n is the amplification factor for the maximal depth.

Last but not least, the resulting signal $h(t)$ is compressed by decreasing its dynamic range (ratio of strongest to weakest signal). This is usually done with a logarithmic scaling (also referred to as log-compression):

$$h_c(t) = \log(h(t) + c), \quad (3.8)$$

where c is the compression coefficient. At this point the ultrasound image is formed by combining the processed RF lines into a single image.

3.6 Results

3.6.1 Simulation from Synthetic Phantoms

The proposed simulation framework was utilized for generating 2D ultrasound images out of two synthetic datasets, one showing a fetus, Fig. 3.3(a), and the other one showing multiple anechoic regions embedded in a highly scattering medium, Fig. 3.3(c). The fetus dataset is a modified version of the one presented in Jensen and Munk [61] and the phantom dataset is generated using Rayleigh noise, with similar phantoms being used for testing real ultrasound imaging systems [50]. The mediums had the following




Figure 3.3: Image (a) shows the synthetic fetus dataset and image (c) the synthetic phantom dataset. The intensity values correspond to speed of sound values in the range of 1500-1550 [m/s]. The center of the focal zones are marked with a small triangle on the left side. Image (b) and (d) show the simulated ultrasound images. Baby phantom from [61].

characteristics: $\delta = 4.5 \cdot 10^{-6} [m^2 s^{-1}]$, $\beta = 6$ and $\rho_0 = 1100 [kg m^{-3}]$, which are common for human tissue [167].

The corresponding simulated ultrasound images are shown in Fig. 3.3(b) and 3.3(d). They clearly demonstrate a realistic speckle pattern, interference effects and beam focusing artifacts. The spatial resolution is high at the center of the focal zone and decreases with increasing distance from the focal zone, an effect also observed in real ultrasound imaging. Furthermore, interference of echoes is strongly evident in the anechoic regions of the phantom dataset. The following parameters were used for simulating the presented ultrasound images: 11 transducer elements formed the active group, $\lambda/2$ elements spacing was used, 192 scan lines were processed, and the discretization steps were set to $\Delta x = \Delta y = 5 \cdot 10^{-3} [m]$ and $\Delta t = 5.5 \cdot 10^{-7} [s]$. For the fetus dataset 6000 timesteps were evaluated and for the phantom dataset 8600, because of the increased depth.

The simulation of the ultrasound wave propagation for generating the raw RF data is performed on the GPU using C++, OpenGL, and the GL shading language (GLSL). Implementing the Finite Difference scheme on the GPU is relatively straight forward with explicit and implicit solvers presented in Harris [49] and Krüger et al. [90] respectively. Switching from GLSL to C-like GPU programming languages like CUDA or OpenCL might improve the performance since they offer more elaborate shared memory features.

After simulating the scan lines, the resulting RF data is processed on the CPU, as

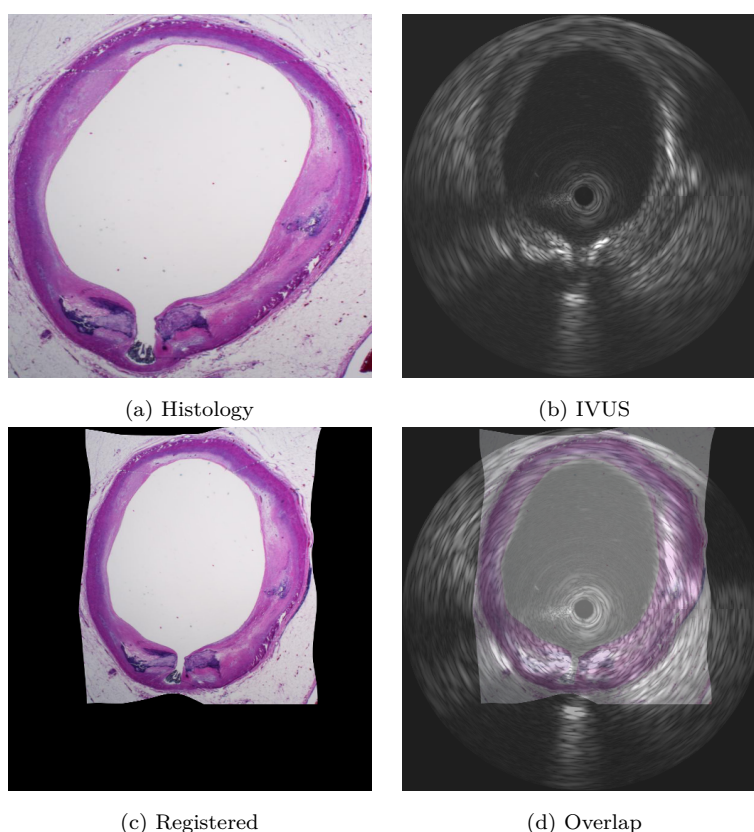


Figure 3.4: Histology image (a) is non-rigidly registered to corresponding IVUS image (b). Registration outcome shown in (c) and (d).

there are no computationally expensive tasks involved. For a 2048^2 grid and 192 scan lines the RF simulation on the GPU required 55 minutes for the fetus dataset and 78 minutes for the phantom dataset, with the image formation on the CPU requiring 19 seconds and 24 seconds respectively. The performance was evaluated on a desktop PC with an Intel Core 2 2.66 GHz with 4 GB RAM and a NVIDIA GeForce GTX 280 with 1GB VRAM. In comparison, the framework presented in Pinton et al. [120] required 32 hours processing time on a 56 PC cluster with 118 GB RAM for a 3D simulation. Details on the processing time of the 2D simulation were omitted, but should be in the same range given the provided parameterizations. Furthermore, one simple 2D image was simulated from a synthetic dataset containing only a single anechoic region.

3.6.2 Simulation from Histology

The aforementioned setup was further utilized to simulate ultrasound from phantoms generated using histology images. In the study of [75] IVUS and histology were acquired post-mortem from dissected coronary arteries from autopsied hearts. The experimental setup allowed accurate acquisition of images from the same arterial cross-sections; for example see Fig. 3.4(a) and (b). Nevertheless, deformations are introduced during the histology preparation process. Consequently, deformable registration of IVUS and his-

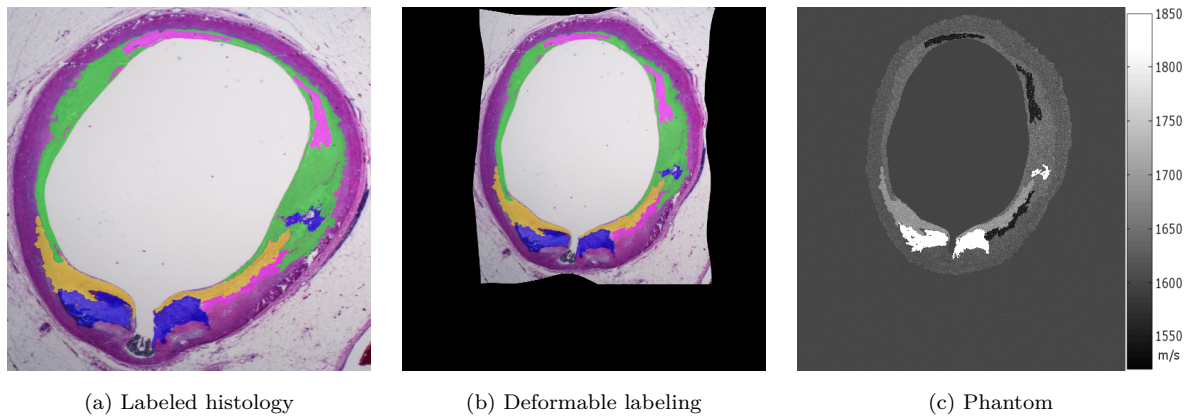


Figure 3.5: Image (a) shows labeled IVUS with regions blue: calcified = 1820m/s, yellow: fibrotic = 1600m/s, green: fibro-lipidic = 1570m/s, pink: lipid= 1520m/s, (b) deformed and labeled histology, (c) phantom with random noise and speed of sound assignments.

tology is required to compare simulated ultrasound from histology with real ultrasound acquired from the same arterial cross-section. In chapter 4 the process for deformable registration will be discussed in detail, whereas, here we will only concern ourselves with the generation of the phantom and the simulation results. The alignment of histology to IVUS, after deformable registration, is shown in Fig. 3.4(c) and (d). The preparation of a simulation phantom is required for the simulation of IVUS from histology. More specifically, acoustic parameters, such as speed of sound, need to be assigned to the different tissue types visible in the histology images. For this purpose a simulation phantom was prepared from an exemplary histology image by manually labeling tissue types and assigning different speed of sound values, as illustrated in Fig. 3.5(a). The labeling could then be mapped to the IVUS image for comparison using the previously computed deformation field; see Fig. 3.5(b). The speed mismatch between tissue types causes reflection of the acoustic signal in real and simulated IVUS. However, a major part of the acoustic response results from scatterers in the tissue. Therefore, random Rayleigh distributed ($\sigma = 1.0$) speed of sound values were assigned to different tissue regions, similarly to [62], to form the simulation phantom. For vessel media speed of sound values were perturbed with ± 50 m/s and for adventitia with ± 7.5 m/s; see Fig. 3.5(c) for the final phantom example and the base speed of sound assignments for tissue types. Finally, the phantom is used in the previously described simulation framework, where a total of 128 scanlines are acquired for the image formation. In Fig. 3.6 original pre-scan-converted IVUS and simulated images are presented.

3.7 Discussion and Conclusion

In this work we presented a framework for fast ultrasound image simulation, covering the imaging pipeline from the initial pulse transmission to the final image formation. Our implementation on GPU hardware simulates realistic ultrasound images in under 80 minutes, avoiding the cumbersome use of PC clusters. The considerably lower simulation




Figure 3.6: Example of simulation results. Image (a) is the real IVUS in polar coordinates (before scan-conversion) and (b) and (c) simulated images using far and near focal regions respectively.

time, compared to other implementations, has practical implications for most simulation related application domains like ultrasound system development. Particularly, our approach has strong implications for future intra-operative simulation of HIFU treatment, as the simulation of a single focal zone is computed in less than 30 seconds. Furthermore, simulating on the GPU allows for an interactive visualization of the wave propagation during the simulation at almost no computational cost, which is of interest to both education and system development applications.

In our current implementation the simulation grid is extended to prevent reflections at the grid boundary from interfering with the region of interest in the simulation grid. Implementing Absorbing Boundary Conditions (ABC), could notably improve the overall performance of the simulation as less grid cells would need to be evaluated. The approach of Perfectly Matched Layers (PML) boundary conditions has been a seminal contribution in this direction [11]. However, derivation of PML conditions for acoustic wave equations is not trivial [170]. In the work of [120] PML conditions for a modified version of the 3D Westervelt equation have been derived and would be of high interest for our future work. This could pave the way for utilizing our framework for 3D ultrasound simulation in acceptable time. Furthermore, GPU hardware and programming capabilities have considerably improved over the years and will also contribute to significantly reduce processing time.

The simulation results obtained using histology images demonstrate a good correlation with real images. Characteristic speckle patterns are evident and reflection at calcified regions is visible. Nevertheless, the simulation of IVUS from histology and the direct comparison to real ultrasound further demonstrates the problems in “reverse-engineering” of the entire image formation process. Numerous parameters influence the image generation including: speckle distributions, speed of sound assignments, tissue parameters, focusing schemes, number of elements, aperture size, apodization, signal processing, frequency, and more. It seems that achieving an accurate simulation of the entire image formation process for a highly complex medium, as encountered in clinical imaging, is an elusive goal. Nevertheless, the progress in simulation models together with high performance computing capabilities that are becoming widely available could allow estimation of this

complex parameter space in the imminent future. This could not only allow the study the effect tissue parameters have on the image generation, but could also lead to new models for describing image formation.

Part II

Ultrasound Confidence Maps and Applications

ULTRASOUND CONFIDENCE MAPS

Advances in ultrasound system development have led to a substantial improvement of image quality and to an increased use of ultrasound in clinical practice. Nevertheless, ultrasound attenuation and shadowing artifacts cannot be entirely avoided and continue to challenge medical image computing algorithms. We introduce a method for estimating a per-pixel confidence in the information depicted by ultrasound images, referred to as an ultrasound confidence map, which emphasizes uncertainty in attenuated and/or shadow regions. Our main novelty is the modeling of the confidence estimation problem as a random walk on a graph, taking into account ultrasound specific constraints. The solution to the random walks equilibrium problem is global and takes the entire image content into account. As a result, confidence maps have an intrinsic robustness to image acquisition parameterization and can be applied to a variety of ultrasound image acquisition setups. We demonstrate the applicability of our confidence maps for ultrasound shadow detection, 3D freehand ultrasound reconstruction, mono- and multi-modal image registration and classification in intravascular ultrasound.

4.1 Introduction

Ultrasound imaging is an essential part of clinical routine offering real-time imaging of patient anatomy. In the last decades advances in ultrasound transducer and system development have led to a substantial improvement of ultrasound image quality and to an ever increasing use in clinical practice. However, ultrasound attenuation and shadow artifacts are still unavoidable, and continue to challenge not only physicians but also engineers in the field of medical image computing. Algorithms for ultrasound image processing are confronted with reduced image contrast and loss of anatomical structures in attenuated and/or shadowed image regions [96]. To overcome these challenges they commonly apply ultrasound specific-constraints and prior knowledge to their underlying methods [111]. Moreover, the accuracy and robustness of ultrasound image processing methods commonly depends on their ability to cope with attenuated and/or shadowed regions. Consequently, approaches that provide knowledge about such regions to other algorithms are of high relevance for different applications, including: ultrasound segmentation, registration, and

reconstruction.

In this work we introduce a novel method for estimating the uncertainty in ultrasound images caused by attenuation and/or shadowing. More specifically, we introduce a per-pixel confidence, i.e., a confidence map, calculated from the ultrasound image content. Our main novelty lies in the confidence estimation process, which is performed within a random walks framework; previously introduced for image segmentation by Grady [38]. The desired confidence map is obtained by answering the following question: What is the probability that a random walk starting from a pixel would be able to reach the virtual transducer elements, given the ultrasound image and ultrasound specific constraints? For this we assume that the likelihood of ultrasound transmission is directly related to the confidence of the image information, i.e., the lower the ultrasound transmission the less the confidence we have in the image information. The ultrasound specific constraints are derived from domain specific knowledge for ultrasound and include: transmission modeling, depth dependent attenuation and ultrasound scanline constraints.

We will further discuss the benefits of a confidence measure in the related work (Section 4.2), together with other approaches that have used measures of confidence, probability, information content, and image quality to improve the outcome of their target application. Subsequently, details of our method will be provided in Section 4.3, including the necessary adjustments to the random walks framework, the modeling assumptions, and implementation details. In Section 4.4 we analyze the free parameters that are involved in the confidence estimation and evaluate the computed estimate for regions of a-priori known confidence. Our confidence maps are not intended to be used independently, but rather to be integrated into other ultrasound image processing methods; providing knowledge of possible attenuated and/or shadowed regions. Consequently, we demonstrate the applicability of our approach for common applications, including: shadow detection, 3D freehand ultrasound reconstruction, mono- and multi-modal image registration and classification in IVUS.

4.2 Related Work

The idea/concept of an ultrasound confidence map has been addressed in the literature with different approaches and in different contexts. Methods that estimate ultrasound attenuation can be used to detect shadowed and/or attenuated areas, thus, providing knowledge about unreliable image regions. The estimation of ultrasound transmission allows algorithms to mask out unreliable regions, like shadows, in order to improve the accuracy and robustness of their underlying methods. More generally, measures of image quality, information content, and probabilistic estimates have been investigated for improving the outcome of various ultrasound image processing methods and will also be discussed in this section.

In ultrasound tissue characterization one tries to extract tissue-related parameters from ultrasound; with the overall goal of differentiating unhealthy from healthy tissue. Estimating the attenuation coefficients of the tissue, i.e. computing a spatial map of attenuation coefficients for a given ultrasound image, has been an elusive goal not only for tissue classification but also for compensating attenuation in the images itself [169].

Various approaches use frequency domain methods to extract the attenuation coefficients. Therefore, a large number of the attenuation estimation methods work directly on the raw Radio-Frequency (RF) data rather than the processed B-mode data. In the work of [91] the spectral shift method is used for which the power spectrum is estimated over a region-of-interest, however, the method is quite sensitive to local spectral noise. In the work of [80] a hybrid method is presented that combines the spectral shift and spectral difference method to overcome the limitations of both methods, which are sensitivity to local noise and to backscatter variation at tissue boundaries. More recently, [169] introduced a joint estimation model for segmenting a region-of-interest (ROI) and estimating the attenuation coefficients. The joint estimation is obtained by alternatively and iteratively solving a cost-function for the backscatter estimation, the contour, and the attenuation. A more comprehensive review on attenuation estimation methods can be found in [110]. In the context of confidence estimation, attenuation estimation methods could provide information on attenuated and shadowed regions [150]. This essentially could provide a measure of confidence as highly attenuated regions, including acoustic shadows, are highly unreliable and should be addressed appropriately by image processing algorithms [51, 110]. Although in recent years attenuation estimation methods have made considerable advances; accurate and reliable attenuation estimation for clinical applications seems an elusive goal [110, 169, 17].

The aforementioned relation of ultrasound transmission to ultrasound B-mode intensities has been utilized for improving the accuracy and results of different image computing algorithms by identifying highly attenuated and/or shadow areas. In [117] shadow areas are identified by sampling the ultrasound scanlines in the B-mode image and thresholding their accumulated intensity. The resulting masks are used for the registration of MR (Magnetic Resonance)-ultrasound of liver images by removing the contribution of the image content in the shadow areas to the registration process. In [96] a similar method was utilized for generating masks of shadow areas. The masks are then used for CT-ultrasound registration of kidney images. The reliability of the shadow detection was improved by thresholding the correlation of the sampled scanline with a heuristic exponential function. In [160] a similar method is applied for detecting shadowing after the tissue-bone interface in ultrasound images of the neck. This information is utilized directly in the cost-function for CT-ultrasound registration. All of the previous approaches integrate the intensities along the ultrasound scanlines. An abrupt change in the accumulated intensities is a strong indicator for the presence of a strong reflector, which is usually followed by a shadowed region. Although this approach improves the registration outcome, it does not provide a spatial map of relative transmission nor of confidence in the data. The high variation of image intensities in ultrasound causes high variation in the accumulated intensities. Therefore, neighboring scanlines exhibit large variation even for tissue of a-priori known homogeneous confidence, like for example after tissue-bone interfaces.

Attenuation and transmission are not the only measures that can be derived from the ultrasound image content. Moreover, ultrasound image analysis approaches have utilized a broad range of different measures for improving their target applications. In the work of [117] multi-modal MR-ultrasound registration of the liver was performed by register-

ing probability maps of MR and ultrasound, which reflect the likelihood of pixels being vessels. The probability maps are generated by estimating the Probability Density Function (PDF) for each modality given pre-segmented datasets. Subsequently, in [116] the same approach was employed for CT-ultrasound registration with the PDF expressing the likelihood of pixels being a tissue-bone interface. The method was evaluated in a cadaver study showing good accuracy and robustness. The downside of such methods is that the PDF is learned from pre-segmented data. Therefore, PDFs are estimated only for a predefined number of classes/objects, whereas the PDF estimation depends on the segmentation parameters and/or user-interaction. Our estimate reflects the confidence in the acquired ultrasound data and not the posterior probability of belonging to predefined classes. In [165] a method is introduced for simultaneous compensation of intensity inhomogeneities and segmentation of ultrasound regions. Moreover, inhomogeneities that are also introduced by the Time-Gain-Compensation (TGC) of the ultrasound system are compensated using a MRF-MAP framework. A uniform compensation is assumed for homogeneous soft-tissue regions that were initially classified. This assumption on homogeneity may introduce errors in the compensation, see [150]. In [44] multi-view ultrasound image reconstruction was performed in a multiscale framework after initial registration of the acquisitions. For the weighting of the different views, feature significance was derived from the phase congruency of the image together with the incidence angle between the ultrasound beam and the local image orientation. Experiments on synthetic and real-time 3D echocardiography acquisitions demonstrated an improvement in image quality. Subsequently, in [43] the registration of multi-view 3D echocardiography is performed with a novel similarity measure that aligns both local phase and orientation. The contribution of the orientation to the registration is weighted by the orientation confidence, which is derived from the local phase, i.e., locations with phase close to zero are assumed to be of low confidence. In [16] the quality and robustness of ultrasonic strain imaging is improved by introducing a quality-guided tracking framework. More specifically, the complex correlation coefficient is used as the measure for comparing pre- and post-deformation RF windows. The generic character of the method allows its use by different displacement estimation methods, e.g., [113], which are commonly required for elasticity imaging methods.

We discussed approaches that extract measures of attenuation, transmission, information content, and image quality from ultrasound image data in order to improve the performance of target applications in different contexts. These assisting measures are primarily derived for application specific tasks and were not evaluated on a broader range of different ultrasound image acquisition setups and/or applications. In this work we introduce a novel approach for estimating confidence, i.e., the uncertainty in attenuated and/or shadowed regions, using a random walks framework. The main advantage of our approach is its generic character. More specifically, the confidence maps can be integrated into various ultrasound image processing methods. These process images acquired from different anatomies using different ultrasound transducers, systems, and settings. Therefore, contrary to previous methods, we offer a general approach for providing knowledge on attenuated and/or shadowed image regions for heterogeneous image acquisition setups and processing scenarios, which can be seamlessly integrated into existing frameworks.




Figure 4.1: Example of graph with corresponding incidence matrix.

4.3 Methods

Random walks for image segmentation, introduced by [38], has become a widely used approach in the computer vision and the medical image segmentation community. Subsequently, the algorithm was applied to several other computer vision and graphics problems including alpha-matting [41], mesh segmentation [172], mesh denoising [141] and stereo matching [135].

In this work we propose the use of the random walk framework for ultrasound confidence estimation. More specifically, the required seeds are automatically placed at the beginning of each scanline to represent virtual transducer elements. Thus, the problem is formulated differently from image segmentation by now computing the probability of a random walk starting from a pixel to reach the virtual transducer elements. The domain specific knowledge for ultrasound is integrated with a simple modeling of the ultrasound image formation process. Specifically, the decrease in ultrasound transmission is derived from the image gradients, depth dependent attenuation is included in the edge-weights of the graph Laplacian, and horizontal/diagonal movements in the graph far from the initial pixel are penalized to model beam-width. Currently our model neglects frequency dependent attenuation and various ultrasound effects like diffraction. Approaches such as in [165] have also made similar assumptions and obtained satisfying results for simultaneous segmentation and compensation of image inhomogeneities. Before moving on to the description of the random walks for image segmentation algorithm and our novel confidence estimation method we first discuss the important matrices involved in defining such equilibrium/balance problems [140].

4.3.1 Incidence and Graph Laplacian Matrix

Graphs are a great tool for modeling different problems including mass-spring systems, circuits, computer networks, trusses, and more [140]. Graphs are commonly represented by matrices, such as the *incidence* matrix [97]. The incidence matrix describes the connection of graph nodes n and edges m with a m by n matrix (a topology matrix); see Fig. 4.1 for an example of a graph together with its incidence matrix. Thus n columns represent the nodes and the non-zero row entries represent the m edges between nodes. For example, row 1 represent edge 1 that connects node 1 to node 3. The signs in the matrix are a matter of convention and can describe the direction of the edges. Another

important matrix for describing graph problems is the *Graph Laplacian*, defined as:

$$L = A^T A = D - W \quad (4.1)$$

For our previous example D and W are defined as:

$$D = \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{bmatrix} \quad W = \begin{bmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{bmatrix} \quad (4.2)$$

The D matrix is referred to as the *degree* matrix and has the number of edges meeting at node, i.e. the matrix is n by n and rows and cols represent the graph nodes. The matrix W is referred to as the *adjacency* matrix. The non-zero matrix entries describe which nodes in the n by n matrix share an edge. The diagonal is zero as nodes are not adjacent to themselves. Commonly we want to place weights on the edges of a graph to model our problem at hand, for example, conductance for circuits. The resulting *Weighted Graph Laplacian* is defined as:

$$L = A^T C A = \begin{bmatrix} -1 & -1 & 0 & 0 \\ 0 & 1 & 0 & -1 \\ 1 & 0 & -1 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} c_1 & 0 & 0 & 0 \\ 0 & c_2 & 0 & 0 \\ 0 & 0 & c_3 & 0 \\ 0 & 0 & 0 & c_4 \end{bmatrix} \begin{bmatrix} -1 & 0 & 1 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 \\ 0 & -1 & 0 & 1 \end{bmatrix} \quad (4.3)$$

where c_1, c_2, c_3, c_4 the weights for the corresponding edges. Both incidence and Laplacian matrix are highly interesting and have various properties. Apparently, the matrices are sparse and the maximum number of non-zeros elements along rows and columns depends on the graph connectivity. What is more interesting is that the nullspace of A , $A^T A$ and $A^T C A$ is not zero, i.e. $Au = [1, 1, \dots, 1]$. This is also evident from the row sums of $A^T A$ and $A^T C A$ which are zero. The rank of A is $r = n - 1$, therefore we need to “ground” a node by removing its corresponding column from the incidence matrix. This makes A invertible and $A^T C A$, invertible and positive definite. Fixing a node to a value is referred to as applying the boundary conditions to the problem and the interpretation of this depends on the problem context. For example, in circuit problems fixing one of the nodes is the condition to introduce a flow-out from the systems, i.e. to get the current flowing (see [140]). In some problem settings its is easier or more intuitive to assemble L first and ground/fix any nodes by removing its corresponding rows and columns from L . The incidence matrix multiplied by u expresses a difference matrix. For example, if u is the potential (voltage) on nodes then Au represent the potential differences between nodes, which causes a current flow. In general the term “potentials” is used for other measures like heat, displacement, etc. If $Au = 0$ we are in a balance situation where no potential differences occur.

4.3.2 Random Walks for Image Segmentation

The random walks algorithm [38] was originally introduced for N -label image segmentation. The image is represented as an undirected weighted graph $G = (V, E)$ with nodes

$v \in V$ representing the image pixels and with edges $e \in E$ where e_{ij} connects nodes v_i and v_j . Each edge e_{ij} is assigned a weight $w_{ij} > 0$ that represents the likelihood of a random walk crossing that edge. The probability of a random walk starting at a pixel to reach first one of the label seeds is computed from the graph Laplacian matrix, which is defined as:

$$L_{ij} = \begin{cases} d_i & \text{if } i = j \\ -w_{ij} & \text{if } v_i \text{ and } v_j \text{ adjacent nodes} \\ 0 & \text{otherwise} \end{cases} \quad (4.4)$$

where $d_i = \sum_j w_{ij}$. Different choices can be made for the weighting function w_{ij} , including the commonly applied Gaussian weighting $w_{ij} = \exp(-\beta(g_i - g_j)^2)$, where g_i is the pixel intensity at node i and β is a free parameter. Alternatively, the Laplacian matrix can also be constructed from the graph incidence matrix A and the diagonal matrix of edge weights C by $L = A^T C A$. Note that L is a sparse, symmetric, and semi-definite matrix. Subsequently, L is re-ordered and decomposed into blocks of marked M and unmarked nodes U :

$$L = \begin{bmatrix} L_M & B \\ B^T & L_U \end{bmatrix} \quad (4.5)$$

The solution to the desired probabilities is obtained by solving the system of linear equations:

$$L_U x_U = -B^T x_M \quad (4.6)$$

where x_U represents the unknown probabilities for the unmarked nodes and x_M the known unit probabilities at the seed nodes. For multiple labels the previous equation becomes:

$$L_U X = -B^T M \quad (4.7)$$

The matrix sizes provide further intuition as:

$$L_{n \times n} = \begin{bmatrix} L_{M(k \times k)} & B_{k \times q} \\ B_{q \times k}^T & L_{U(q \times q)} \end{bmatrix}, \quad L_{U(q \times q)} X_{q \times l} = -B_{q \times k}^T M_{k \times l}$$

where n is the number of nodes, k is the number of marked nodes, q is the number of unmarked nodes, and l is the number of labels. The system is solved for $l - 1$ labels as $\sum_i x_i^l = 1$.

4.3.3 Random Walks and Circuit Analogy

In the same work of [38] the analogy of random walks to circuits was discussed. Here we will further elaborate on this analogy, which helps in interpreting the formulation of ultrasound confidence maps. The aforementioned mathematical framework belongs to the class of discrete equilibrium equations encountered in potential theory and mass-spring systems (see [139]). The basic framework for solving equilibrium equations in the context of circuits is illustrated in Fig. 4.2. In circuit theory we are commonly interested in finding potentials. The incidence matrix A describes the connectivity and multiplied with the potentials at the nodes u produces the potential differences at the edges. External sources

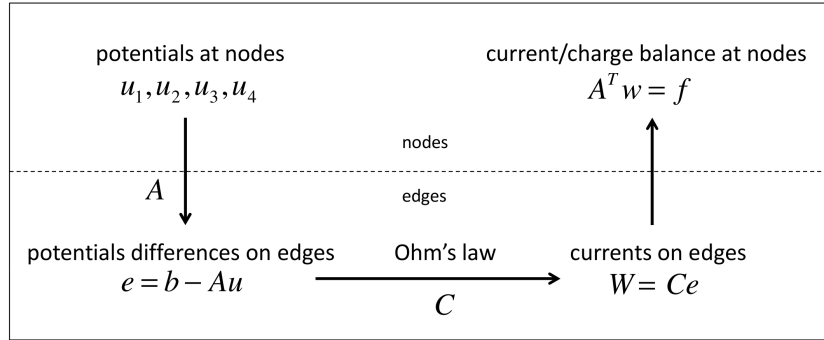


Figure 4.2: Basic framework for solving balancing equations in linear algebra (see [140]), here in the context of circuit problems.

(e.g. batteries) connected to the edges are taken into account $e = b - Au$. Ohm's law tells us that the current is proportional to the potential differences that cause current flow and depends on resistors along the circuit path. Here conductance is used $C = 1/R$ instead of resistance, yielding Ohm's law $W = Ce$. Finally, Kirchoff's Current Law (KCL) states that the incoming current/charge on a node is equal to the outgoing current/charge from the node. Therefore, the last part of the framework provides the balancing term with external current sources f on the right-hand side. Putting the terms of the framework together we end up with the so-called *fundamental equation of equilibrium*:

$$f = A^T w = A^T C e = A^T C (b - Au) \Rightarrow A^T C A u = A^T C b - f \quad (4.8)$$

This is also referred to as the Displacement Method for Finite Element Methods.

In final analysis, the re-ordering and block decomposition of L in the random walks method is not only important to describe the problem in a linear algebra framework, but it is necessary for applying the Dirichlet boundary conditions to the problem. Essentially, L is singular and by fixing/grounding nodes one effectively creates the invertible matrix L_U . Labeling/fixing nodes, i.e., applying the boundary conditions practically means that one removes the rows and columns associated with these nodes from the system of linear equations L , which is an equivalent alternative to re-ordering and decomposing L . Summarizing, in this analogy the graph represents an electric circuit, the nodes electric potentials, the edges resistors, and the weights electrical conductance. The previously described random walks solution gives the potentials at the nodes, subject to the boundary conditions by fixing the electric potential at the boundary/labeled nodes.

4.3.4 Ultrasound Confidence Map

We formulate our confidence estimation problem as follows: What is the probability of a random walk starting from a pixel to reach each of the virtual transducer elements under ultrasound-specific constraints? To answer this question we adjust the original random walks for image segmentation algorithm in order to provide the desired solution, starting with the required seed placement procedure.




Figure 4.3: Illustration of graph setup for confidence maps using an 8-connected lattice. The first row resembles the virtual transducer elements with the probabilities set to unity (shaded dark gray). The last row resembles the 'no signal' region, i.e., the necessary boundary condition, with the probabilities set to zero (shaded bright gray). For all the nodes in-between we compute the probability that a random walk starting from each node would first reach one of the virtual transducer elements.

To model our problem the seeds are automatically placed at the beginning of each scanline, representing the position of the virtual transducer elements; and at the end of each scanline, representing the absorption region after which no signal is practically recorded. An illustration of the seed placement is given in Fig. 4.3. From the formulation of our problem it becomes evident that we are only interested in the probabilities of reaching the seeds at the virtual transducer nodes. The seeds at the absorption region are only required to satisfy the boundary conditions. The formulation of the random walks problem on a graph allows the lattice to represent arbitrary dimension and topology. As a direct consequence all possible transducer geometries can be modeled with this approach by aligning the lattice columns with the scan lines in the axial beam direction.

At this stage solving for the probabilities using the Gaussian function for the edge weights, as originally proposed in [38], would be meaningless for our problem. Basically, the solution obtained would have high probabilities until reaching the strongest gradient in the image, followed by an abrupt drop in probabilities. From a segmentation point of view we would obtain a cut through the most significant gradients in the image, which in ultrasound are usually located at the near-field because of the tissue compression. Therefore, we need to integrate ultrasound specific constraints in order to obtain the desired confidence map. One of the fundamental properties of ultrasound is the attenuation of the signal with increasing depth. Various models exist for ultrasound attenuation, see [17], including variations of the Beer-Lambert Law, which we use for our method to express the depth dependent attenuation. The attenuated signal I is given by

$$I = I_0 \exp(-\alpha d) \quad (4.9)$$

where I_0 is the initial intensity, α the attenuation coefficient, and d the distance from the source. To include this property into our confidence map estimation we adjust the

behavior of the random walks. More specifically, the further away a random walk starts from the virtual transducer the more unlikely it becomes that it will be able to reach one of the virtual transducer elements.

In the ultrasound image formation process a narrow and focused beam of sound is emitted from the transducer to acquire each image scanline [50, 17]. Therefore, it is unlikely that sound outside this narrow beam would notably contribute to the acquisition of each scanline. Nevertheless, the beam width causes structures close to a scanline to contribute to the acquired echo. This property is modeled into the random walks framework as follow: a random walk starting from the node of a scanline should be able to move along the horizontal and diagonal direction of the graph, i.e., approximately perpendicular to the beam/scanline direction. However, the likelihood of horizontal/diagonal walks should decrease with increasing distance to the scanline. A more intuitive explanation of this property is that a random walk starting below a reflector in the image should be able to bend around it, which is the case for sound in reality. However, the extent of the bending should be restricted to model the width of the ultrasound beam that formed the scanline.

These two constraints are integrated into the random walks framework with a new weighting function given by:

$$w_{ij} = \begin{cases} w_{ij}^H & \text{if } i, j \text{ adjacent and } e_{ij} \in E_H \\ w_{ij}^V & \text{if } i, j \text{ adjacent and } e_{ij} \in E_V \\ w_{ij}^D & \text{if } i, j \text{ adjacent and } e_{ij} \in E_D \\ 0 & \text{otherwise} \end{cases} \quad (4.10)$$

$$w_{ij}^H = \exp(-\beta(|c_i - c_j| + \gamma)) \quad (4.11)$$

$$w_{ij}^V = \exp(-\beta(|c_i - c_j|)) \quad (4.12)$$

$$w_{ij}^D = \exp(-\beta(|c_i - c_j| + \sqrt{2} \cdot \gamma)) \quad (4.13)$$

$$c_i = g_i \exp(-\alpha l_i) \quad (4.14)$$

where E_H, E_V, E_D are the edges along the horizontal, vertical, and diagonal graph direction respectively, see Fig. 4.3, with $E_H \oplus E_V \oplus E_D = E$, g_i is the image intensity at node/pixel i , and l_i is the normalized closest distance from node v_i to the virtual transducer elements (seeds). The free parameters (α, β, γ) are described in detail in the following paragraphs.

The Beer-Lambert attenuation is integrated in the random walks weighting function by scaling the pixel intensities as a function of distance to the virtual transducer elements. This effectively decreases the confidence with increasing distance to the virtual transducer elements; effectively modeling the uncertainty in the data because of the depth-dependent attenuation. The solution to the above equilibrium problem results in random walks starting further away from the transducer to have a reduced probability of reaching the virtual transducer elements and vice versa. However, note that the probability still depends on the image content, i.e., the image intensities. This is desired as image structures can cause a high drop in confidence even when they are present close to the transducer elements.




Figure 4.4: The first and third row show B-mode images of a human heart, femur, neck and liver. The second and fourth row show the corresponding confidence maps. Bright values denote high confidence and dark values low confidence. All images were generated with the same α, β parameters, namely $\alpha = 2, \beta = 90$. The γ parameter was optimized for the first image of each sequence/anatomy using Eq. 4.15 and was then kept constant for the remaining images of the sequence.

The γ parameter acts as a penalty on the random walks reducing the probability of a random walk crossing a horizontal/diagonal edge in the graph with increasing horizontal/diagonal distance from the node. The distance is implicitly encoded into this formulation. Basically, the term acts accumulatively for each random walk, which becomes apparent when interpreted in circuit theory. From Ohm's and Kirchoff's law we know that the potential at a node that can be reached from a circuit path of high resistance is lower than the potential at a node that can be reached through a circuit path of low resistance. Therefore, the additive term acts as an accumulative increase in the resistance along possible horizontal/diagonal random walks. Notice that diagonal walks are further penalized by a factor of $\sqrt{2}$ to properly model the greater distance to adjacent nodes compared to horizontal walks.

Ultrasound transmission refers to the sound energy/intensity that is transmitted at the boundary of two mediums of different acoustic impedance. At such boundaries/interfaces part of the energy is reflected, producing the echoes detected by the ultrasound trans-

ducer; and part of the energy is transmitted into the underlying tissue. Reflection and transmission are negatively correlated, i.e., the higher the reflection the lower the transmission. For a more comprehensive overview of ultrasound physics and image formation we refer to [50]. One can assume that high intensities in the ultrasound image correspond to an increased likelihood of low transmission. This would suggest to make the weights w_{ij} directly proportional to the ultrasound image intensities g_i . On the other hand, a change of the image intensity also suggests a change of the ultrasonic transmission. This is because the transmission property of different heterogeneous tissue types is partly suggested by their echogenicity, even in the absence of a single strong reflector (consider the case of a blood vessel within the liver gland). This in turn suggests that we use the image gradient $g_i - g_j$ within the weight term. We have investigated various combinations of using both the intensities and the respective gradient, and concluded that the latter alone yields more stable results.

4.3.5 Implementation Considerations

Similar to the original algorithm we found it useful to normalize the gradients $|c_i - c_j| \forall e_{ij} \in E$ before applying Eq. 6-8. However, instead of a 4-connected lattice we used an 8-connected lattice for all our experiments. The effect of graph connectivity to random walks was already discussed by [40] with regard to segmentation. For our confidence map estimation it makes intuitively more sense to use an 8-connected lattice in order to model more accurately the possible paths of sound. Practically, with an 8-connected lattice the estimated confidence maps are smoother and more homogeneous compared to a 4-connected lattice.

The necessary adjustment for obtaining the confidences map does not alter the favorable properties of the random walks framework. More specifically, we are only changing C in $L = A^T C A$, thus, L remains sparse, symmetric positive-definite as $C > 0$. Consequently, the above system of linear equations can be solved with direct methods such a LU decomposition for 2D ultrasound images.

In terms of performance, GPU implementations of sparse matrix solvers have become widely available [10]. For problem sizes similar to the confidence estimation, i.e., the Laplacian operator discretized as 5-point finite difference stencils on a 1000^2 regular grid, real-time solutions can be obtained on commodity GPU hardware. Furthermore, a fast Poisson multigrid method for solving the inhomogeneous equation was introduced in [39], that computes the solution to the discussed random walks problem almost in real-time on commodity CPUs. In our current non-optimized Matlab implementation, confidence maps for images of size 512×512 were computed in 2.1 sec on an Intel Core i7 3.4 GHz processor. Furthermore, a rectangular and regular lattice is used for every image to form the graph Laplacian. For this purpose, images acquired from convex probes, i.e. all probes except linear ones, are mapped from the Cartesian to the Polar coordinate system with linear interpolation between scanlines. Ideally, the pre-scan converted data of the ultrasound machine should be used for the confidence estimation to minimize interpolation errors. Nevertheless the currently implemented interpolation strategy provides good qualitative and quantitative results, which will be presented in the next sections.

Theoretically, our method could directly be applied to 3D ultrasound, as for example acquired by modern 2D matrix array transducers. From an implementation point of view the required changes are relatively minimal. Similarly, as in the 2D case, nodes are labeled to mark the virtual transducer elements and the absorption region. The only substantial change from 2D to 3D is that random walks can travel along two possible directions perpendicular to the beam, i.e., the elevation and azimuth direction. Therefore, both directions are required to be penalized, similar as in the 2D case, in order to model the beam width. Depending on the transducer properties the beam width could vary along the elevation and azimuth direction, a property that can be modeled by using a different γ values for each direction. The same GPU implementations that were previously discussed would apply to the 3D case [42].

4.4 Evaluation

We provide qualitative results in form of confidence maps obtained from ultrasound images for different transducer geometries and anatomies, as shown in Fig. 4.4. Before we move on to the other results we would first like to discuss the free parameters of our method.

4.4.1 Parameters

The method presented has three free parameters, namely α , β and γ . Starting with the α parameter, it affects the likelihood of vertical random walks, i.e., along $e_{ij} \in E_V$. In the classical Beer-Lambert law it controls the slope of the attenuation as a function of distance. In our random walks formulation it effectively scales the confidence along the vertical graph direction. Fig. 4.5 shows an example of the effect of α on estimating a confidence map for a liver image. For all subsequent experiments we set the parameter to a constant value of $\alpha = 2$, obtaining both good qualitative and quantitative results.




Figure 4.5: The effect of the α parameter. Image (b) and (c) are the confidence maps for image (a) by using different α values and fixing $\beta = 90$ and $\gamma = 0.05$.

The β parameter was adopted from the original random walks algorithms. Its effects on the robustness and accuracy of the segmentation was thoroughly studied in [40]. Similar to the original work we set β to a constant value of $\beta = 90$ for all the experiments. Small deviations from this value did not prove to have any notable effect on the confidence estimation, which is expected based on the analysis of [40]. Fig. 4.6 shows an example

of the effect that β has on estimating a confidence map for a femur image. Differences between the confidence maps are barely noticeable.




Figure 4.6: The effect of the β parameter. Image (b) and (c) are the confidence maps for image (a) by using different β values and fixing $\alpha = 2$ and $\gamma = 0.05$.




Figure 4.7: The effect of the γ parameter. Image (b) and (c) are the confidence maps for image (a) by using different γ values and fixing $\alpha = 2$ and $\beta = 90$.

The γ parameter penalizes horizontal and diagonal random walks in the graph the further away they are from the starting scanline. Selecting the parameter requires a balance between having distinct confidence and having minimal discontinuities in the horizontal direction. Fig. 4.7 shows an example of the effect of γ on estimating a confidence map for a neck image. One could manually adjust the parameter or optimize for the best balance between distinct confidence and low discontinuities in the horizontal direction, which we formulate as

$$\operatorname{argmin}_{\gamma \in [0,1]} SD \left[\left(\left| \frac{\partial I}{\partial y} \right| - \left| \frac{\partial C(\gamma)}{\partial y} \right| \right)^2 \right] \quad (4.15)$$

where I is the ultrasound image, $C(\gamma)$ is the confidence map for a given γ and fixed α, β , and SD is the standard deviation. We choose the range $\gamma \in [0, 1]$ to ensure that the additional penalty from γ will not overpower the horizontal penalty already imposed by the image content, i.e., the Gaussian function, during the optimization process.




Figure 4.8: Boxes in the images indicate regions of high and low confidence; before and after highly attenuating interfaces. Images in column (a) are the original ultrasound of femur and liver, in (b) the corresponding confidence maps, and in (c) the scanline integrated images.

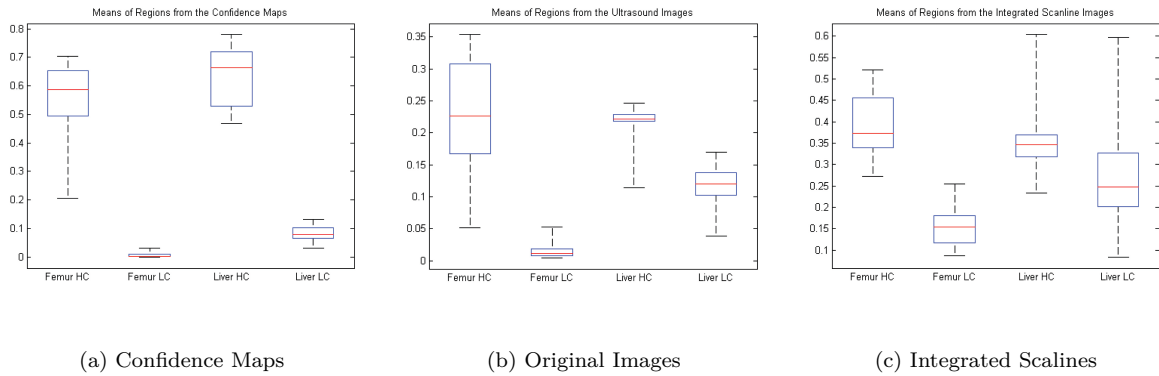


Figure 4.9: Regions of relatively high confidence (HC) and a-priori known low confidence (LC) are defined for 25 human femur and 25 human liver images. Subsequently, the mean values are computed for regions sampled from our confidence maps, the original ultrasound images, and the scanline integrated images. Each box-plot shows the means of the HC and LC regions for every test case.

4.4.2 Confidence in Shadow Regions

Our confidence estimation method emphasizes the uncertainty of attenuated and/or shadowed regions. Consequently, we decided to evaluate our method by assessing its reliability in estimating low confidence for regions that are a-priori known to be of low confidence. More specifically, it is known that almost the entire ultrasound pulse is reflected at the tissue-lung and the tissue-bone interface, resulting in acoustic shadow and unreliable

ultrasound information after the interface, see [50, 17]. We evaluate our method by examining the confidence values before and after such interfaces. The assumption is that our method should indicate a drop in confidence between these regions and an overall very low confidence for the regions after these interfaces. Accordingly, we define regions of high confidence (HC) and low confidence (LC) in ultrasound images of the human femur and the liver. Subsequently, we sample our confidence maps, the original ultrasound images, and the scanline integrated images at the same regions, as shown in Fig. 4.8. We chose scanline integration as it is a simple and commonly applied approach in ultrasound image processing, see Section 4.2, for detecting shadow areas in the images. The scanline integrated image S is given by

$$S(p, s) = S_M - \int_0^p G_\sigma * I(p, s) dx \quad (4.16)$$

where $I(p, s)$ is the image intensity at position p along the s scanline, G_σ a Gaussian kernel of size 5 and $\sigma = 2.0$, and $S_M = \max(S)$. Overall, we examine the means of HC and LC regions for 25 femur and 25 liver images using the original images, the corresponding confidence maps, and the scanline integrated images. The parameters we set to constant values for all experiments, namely $\alpha = 2, \beta = 90, \gamma = 0.05$.

The results of this analysis are presented in Fig. 4.9(a)-(c). The plot in Fig. 4.9(a) shows that our confidence maps reliably detect the low confidence region in both femur and liver datasets. We observe that the confidence in the HC regions of the liver dataset was higher than the one in the femur datasets. This is expected given that larger intensity variations are encountered for femur datasets, because of the fat and muscle tissues before the femur. This becomes clearer when looking at the same statistics sampled from the original ultrasound images, see Fig. 4.9(b). Last but not least, we include the scanline integrated images, see Fig. 4.9(c), to emphasize the problems that algorithms might encounter, see Section 4.2, when simply sampling the scanlines for identifying highly reflecting interfaces and subsequent shadow areas. Before the scanline integration the ultrasound images were filtered with a strong Gaussian kernel to reduce the affect of noise that otherwise would have resulted in even greater intensity variation between the integrated scanlines.

4.5 Confidence Map Applications

Ultrasound confidence maps provide a measure of uncertainty in attenuated and/or shadowed regions in ultrasound. This descriptor is by itself of limit usability, however, when integrated in different image processing applications it can indeed prove beneficial. In this chapter we demonstrate the applicability of confidence maps for different applications including: shadow detection, 3D freehand ultrasound reconstruction, mono-/multi-modal registration and classification in IVUS. We will discuss the working principles of these methods and introduce concepts for integrating confidence in them to improve algorithmic outcome.

4.5.1 Shadow Detection

In this subsection we investigate the applicability of our confidence map concept for shadow detection. Although acoustic shadow is considered an ultrasound imaging specific artifact, it is of diagnostic value and is used for the detection of calcifications, gallstones, and bone structures [50, 51]. Nevertheless, for ultrasound image processing, including: registration and 3D freehand-/multiview ultrasound reconstruction, shadowing might negatively affect the algorithmic outcome [51]. Consequently, the detection of shadow regions in ultrasound images is of high importance.

In the related work, we already discussed the approaches introduced by [117], [96], and [160] to integrated shadow detection in multi-modal ultrasound-CT and ultrasound-MR registration. The work of [51] addresses some of the problems of previous approaches in regard to automatic and robust shadow detection. Discontinuities are detected in the sampled scanlines (line ruptures) and then tested whether shadow occurs after the rupture based on the local intensity statistics. The method works fully-automatically, however, the underlying statistical model assumes low signal and noise in shadow regions. This is certainly a valid assumption for intra-operative brain imaging. However, in other applications, including imaging of bone structures, this assumption is unlikely to hold because of the broad range of strong echo amplitudes in shadow regions (see examples in Fig. 4.10 and [50]).

We decided to evaluate the shadow detection capabilities of the confidence maps with a simple approach that only utilizes the maps and no further information. More specifically, we define shadow regions by merely thresholding the confidence map for each image with:

$$S_i = \begin{cases} 1 & \text{if } C_i \leq T \\ 0 & \text{otherwise} \end{cases} \quad (4.17)$$

where $C_i \in [0..1]$ is the confidence value at pixel/node v_i and $T = 0.25$ is a heuristic threshold for determining the shadow region S . This simple approach was chosen in order to evaluate the direct applicability of the confidence maps for shadow detection. Nevertheless, combining the confidence map and additional image information, e.g. image statistics, is likely to further improve the accuracy and robustness of the detection.

For the evaluation we acquired ultrasound images of the human humerus from 3 healthy volunteers with an Ultrasonix MDP machine in combination with a linear L14-5/38 GPS transducer. The ultrasound system allows the simultaneous acquisition of both pre-processed radio-frequency (RF) data and post-processed B-mode images. For this evaluation the confidence maps were computed on the envelope detected RF images, which is considered to reduce the influence of the post-processing, e.g., logarithmic compression. More specifically, c_i in Eq. 4.14 becomes:

$$c_i^* = |\mathcal{H}(r_i)| \exp(-\alpha l_i) \quad (4.18)$$

where r_i is the RF amplitude at node/sample i and \mathcal{H} is the Hilbert transform. Exemplary detection results are presented in Fig. 4.10. Shadow regions are successfully detected for images that have no or minor reverberation artifacts, see Fig. 4.10(a) and (b) respectively. However, strong reverberation artifacts reduce the accuracy of the detection, Fig. 4.10(c).




Figure 4.10: Images (a)-(c) show the shadow regions, area below dotted green line, for envelope detect RF images (here log-compressed for display) acquired from a human humerus. The corresponding confidence maps are shown in images (d)-(f).

This is expected given that reverberation/mirroring artifacts are not modeled in our framework. Thus, the confidence drops after the last strong reflector present in the shadow area of the image, which is mistakenly interpreted as an anatomical structure although it is a reverberation artifact. Detecting reverberation/mirroring artifacts and including this information in the confidence estimation would further improve the accuracy and robustness of the overall shadow detection.

For a quantitative evaluation we asked two experts (denoted U1 and U2) to manually segment the shadow regions in a total of 280 ultrasound images of the humerus, acquired from 3 healthy volunteers. Half of the images were acquired with 5 MHz and the other half with 10 MHz frequency. Subsequently, the Dice coefficient ($D = 2|X \cap Y| / (|X| + |Y|)$) was computed between manual and confidence based segmentation. The results are presented in Fig. 4.11. The average Dice and standard deviation for U1 was (0.87 ± 0.05) and for U2 (0.84 ± 0.05) . The intra-user variability is evident in the evaluation and has also been previously discussed by [51]; emphasizing the need for fully automatic shadow detection. The results obtained are satisfying given the simplicity of the shadow detection calculation, i.e., binary thresholding of the confidence maps. Similar results were obtained for both acquisitions frequencies; demonstrating the robustness of the confidence estimation. Nevertheless, combining the confidence maps with additional/complementary image information could improve the accuracy and robustness; especially for images corrupted by reverberation/mirroring artifacts.

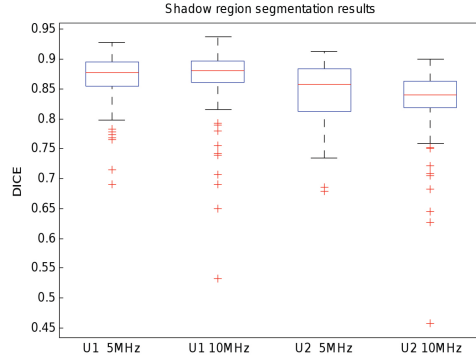


Figure 4.11: Shadow detection results are presented by comparing the manual and automatic detection of shadow regions with the Dice coefficient. Two experts (U1 and U2) segmented for each frequency setting the same 140 images.

4.5.2 Ultrasound Reconstruction

Three-dimensional freehand ultrasound reconstruction usually refers to the process of interpolating a series of 2D freehand-ultrasound slices into a regular 3D Cartesian grid, i.e., a regular voxel volume. The acquisition of 3D freehand-ultrasound follows a general setup. In order to obtain the ultrasound pose information a tracking target is attached to the transducer probe, which among other options can be a configuration of optical markers or an electromagnetic sensor. With this setup ultrasound images can be acquired arbitrary in 3D space. The purpose of this acquisition setup is primarily to extend the field-of-view of ultrasound and to image anatomies from different angles to provide complementary imaging information. Comprehensive overviews on 3D freehand ultrasound reconstruction can be found in [127, 136]. Furthermore, multi-angle ultrasound compounding (MACI), also referred to as spatial compounding, has demonstrated improvement in image quality, i.e., an increase in signal-to-noise ratio (SNR) and reduction of ultrasound speckle, as described in the works of [130, 163, 8].

In our application we acquired 3D-freehand ultrasound of the human femur from different acquisition angles on a volunteer following all ethical guidelines. The acquisition was performed with an Ultrasonix MDP machine in combination with a convex C5-2/60 transducer that was tracked with an NDI Polaris Spectra optical tracking system. For the purpose of the experiment the leg of the volunteer was fixed to ensure consistent acquisitions from different angles. Clearly, this is not an advisable procedure for patients, thus, pre-registration of the different ultrasound views is required before reconstruction in order to produce high-quality results, see [126].

Only the bone surface perpendicular to the ultrasound beam direction is visible because of the view-dependency and shadowing after the tissue-bone interface. Therefore, acquiring the bone from different views allows to acquire the missing bone surface. This complementary information could be of use for applications such as computer-aided orthopedic surgery with ultrasound guidance, as described in the works of [116, 5, 6]. The question that arises is: How to weight the contribution of overlapping or close ultrasound slices for the interpolation/compounding process [127, 43]? This is of particular interest considering the presence of artefacts, i.e. shadow areas that contain low-confidence inten-




Figure 4.12: Image (a) and (b) show a transverse slice of a human femur from a reconstructed volume using the R_G and our newly defined R_C weighting function respectively. Image (c) shows the joint histogram of two volumes, one reconstructed with R_G and the other with R_C .

sity information. The reconstruction is performed on a 3D lattice Ω , where at each discrete voxel $i \in \Omega$ an intensity is reconstructed from a set of measurement tuples $A_i = \{(y_j, d_j)\}$, where each tuple consists of a distance d_j and associated pixel intensity value y_j . Specifically, d_j denotes the Euclidean distance between the voxel center and the pixel location. All tuples in the set fulfill the distance constraint $d_j < D$, where D denotes the radius of a sphere limiting the associated measurements range for each voxel. A common choice for intensity-based reconstruction is to employ a Gaussian weighting function as given by:

$$R_G(A_i) = \frac{\sum_j y_j e^{d_j^2/\sigma^2}}{\sum_j e^{d_j^2/\sigma^2}} \quad (4.19)$$

where $R_G(A_i)$ is the reconstructed voxel intensity at lattice position i using the measurements from the associated tuple set A_i . As a result, the reconstruction is a combination of a smoothed distance-weighted sum of pixel intensities in proximity to the voxel i . However, this results in pixels from shadow areas and tissue areas to contribute alike in the reconstructed voxel intensity. In the optimal case, one would like to reduce the contribution of shadow areas to areas where complementary anatomical information is available. Our newly developed confidence estimation can be directly used for this purpose. More specifically, for each 2D ultrasound B-mode image acquired we estimate a confidence map. Notice, previously for shadow detection (subsection 4.5.1) we estimated confidence maps from the envelope detected RF data, here, we perform the estimation on the post-processed B-mode images using Eq. 4.14. Subsequently, in the reconstruction process we take the confidence into account to weight the intensity information accordingly. More formally, we define a new confidence-based weighting function as:

$$R_C(B_i) = \sum_j y_j \exp\left(-\frac{(c_j - \max(C))^2}{(\max(C))^2}\right) \quad (4.20)$$

where $B_i = \{(y_j, d_j, c_j)\}$ is now a triplet, augmented by the confidence value $c_j \in C$, where $C = \{c_j\}$ denotes the set of confidence values associated with voxel i . Taking the confidence maps into account the new weighting function allows to weight the contribution of different slices based on their confidence; reducing significantly the contribution of low confidence regions.




Figure 4.13: Image (a) and (c) show a sagittal slice of a human femur from a reconstructed volume using the R_G and our newly defined R_C weighting function respectively. Image (b) and (d) show an enlarged version of the regions marked with the white box.

Fig. 4.12 and 4.13 show transverse and sagittal slices from an ultrasound volume of a human femur, reconstructing with a Gaussian (R_G) and our new confidence-based (R_C) weighting function. The improvement in contrast is evident for the reconstruction utilizing our confidence estimate, especially for the bone region. We measure the improvement with the Weber contrast as given by $W = (I - I_b)/I_b$, where I is the mean intensity inside the femur, I_b the background intensity outside the femur, and $I, I_b \in [0..1]$. For the reconstruction using the R_G function the contrast is $W = 0.32$ and using the R_C function it improves to $W = 0.45$. The improvement in contrast is not the result of a linear intensity remapping, but the result of reducing the contribution of low confidence region. This is evident when examining the joint histogram, shown in Fig. 4.12(c), between two volumes reconstructed with R_C and R_G .

4.5.3 Multi-modal Registration

A challenging image processing problem is the automatic spatial alignment (i.e. registration) of pre-operative CT or MRI data to ultrasound imaging acquired during a procedure. The main difficulty lies in the highly different representation of the same anatomy in the respective modalities. In order to allow for image-based registration of CT and ultrasound data, in [157] ultrasonic effects such as large-scale reflections and tissue echogenicity are simulated from CT, in order to make the information from both modalities more comparable. This is combined with a robust similarity measure (LC^2) which is invariant to unknown simulation parameters. This similarity measure was also successfully applied for ultrasound to CT registration of the lumbar spine, using a biomechanically constrained groupwise registration approach [35]. One component of the simulation essentially comprises the CT specific equivalent to integrating over ultrasonic reflections along scanlines. A simulated ultrasound transmission map $T(\vec{x})$ is therefore a by-product of this registration algorithm:

$$T(\vec{x}) = \exp \left(- \int_0^{\lambda_x} \left(\frac{|\nabla\mu(\vec{x}_0 + \lambda\vec{d})|}{2\mu(\vec{x}_0 + \lambda\vec{d})} \right)^2 d\lambda \right) \frac{(\vec{d}^T \nabla\mu(\vec{x}))}{(2\mu(\vec{x}))^2} \frac{|\nabla\mu(\vec{x})|}{(2\mu(\vec{x}))^2} \quad (4.21)$$




Figure 4.14: CT-Ultrasound image pair and their associated transmission and confidence map.

where \vec{x} is the pixel location, μ the function of CT attenuation values, \vec{d} a unit vector denoting the direction of ultrasonic wave propagation, and λ the parameterization of a single scanline. Our confidence map estimated from ultrasound can serve well as the counterpart to such a transmission map. Due to the global solution of an equilibrium system, the estimated confidence emphasizes large-scale structures and is not influenced much by local echogenicity.

We have therefore implemented a new CT-ultrasound similarity measure based on the correlation of a transmission map simulated from CT with the ultrasound confidence maps:

$$Corr(T, C) = \frac{\sum_i (t_i - \bar{t})(c_i - \bar{c})}{\sqrt{\sum_i (t_i - \bar{t})^2 \sum_i (c_i - \bar{c})^2}} \quad (4.22)$$

where t_i , c_i are the individual pixel intensities of the CT transmission and confidence images, respectively, and \bar{t} , \bar{c} their mean values. We compare the registration using this similarity term with the original approach presented in [157]. Fig.4.14 shows typical images from a human femur cadaver, comprising CT, ultrasound, transmission estimated from CT, as well as the confidence map estimated from post-processed B-mode ultrasound. Tracked 3D freehand ultrasound was acquired over the femur, and visually aligned with the CT scan in order to provide a ground truth registration. For the acquisition we used a Siemens X150 ultrasound machine with a convex VF-105 transducer that was tracked with a NDI Aurora magnetic tracking system. For incorporating the confidence similarity, we have investigated three options: 1) confidence similarity (Eq. 4.22) alone; 2) multiplying the original LC^2 similarity measure proposed in [157] with the confidence similarity; 3) executing two successive local optimizations, first with the confidence similarity, then

using the original LC^2 formulation. For assessing the registration performance, the regis-

	trans. (mm)		rot. (deg)		
	t_x	t_y	r_x	r_y	TRE
RMS error					
original	5.2	5.5	5.0	5.9	8.1
confidence	3.5	2.4	2.6	1.7	4.3
product	3.2	3.0	3.1	2.3	3.7
successive	2.4	3.4	2.4	1.6	2.5
standard deviation σ					
original	5.2	4.2	4.5	5.6	6.0
confidence	2.3	2.0	2.6	1.5	3.4
product	3.1	2.5	3.1	2.1	4.8
successive	2.3	2.9	2.4	1.5	4.6

Table 4.1: Standard deviation of registration parameters.

tration was launched 100 times each for the different choices, respectively, from randomly perturbed ($\pm 10\text{mm}/^\circ$) starting estimates of the transformation parameters. The results are shown in Table 4.1, in terms of the target registration error (TRE) of a single point target located in the center of the ultrasound sweep, root-mean-square (RMS) errors of some of the transformation parameters, as well as their standard deviations σ . Not shown are the translation along and rotation around the cranio-caudal axis (t_z , r_z), which contain larger errors for all methods, due to the ambiguity of the long femur bone.

The confidence correlation alone yields small standard deviations and therefore represents a smooth similarity term with large capture range, however, since detailed anatomic structures are neglected the results are somewhat off the ground truth. The multiplicative similarity improves on the absolute transformation errors. Best results however are achieved with a successive optimization. Here, the new confidence-based registration contributes robustness and large capture range, the original LC^2 formulation achieves high accuracy due to the incorporation of simulated ultrasound effects and local soft tissue anatomy.

4.5.4 Mono-modal registration

Mono-modal ultrasound registration is of high interest for multi-view/mosaicing approaches, where an extended field-of-view is assembled through registration of multiple ultrasound volumes [122, 153]. As in multi-modal registration, confidence maps could be used to reduce the effect of local outliers to registration and increase overall capture range. We demonstrate this possible application with a mono-modal 2D registration example of two similar ultrasound images of the liver. In Fig. 4.15(a) and (b) the test images are shown and in (c) and (d) their corresponding confidence maps. We evaluated a standard mono-modal similarity measure - Sum of Squared Difference (SSD) - for different translation and rotation poses of the images. For this only the ultrasound image content, i.e.,




Figure 4.15: Input for mono-modal registration example with confidence maps.

the pixels in the ultrasound fan, are taken into account without the background. Otherwise optimal registration would be given when the fans are aligned, instead of when the anatomy is aligned. Fig. 4.16(a)-(c) shows the change of SSD value for different 2D poses in a range of $[-150..150]$ pixels for translation and $[-120..120]$ degrees for rotation. The evaluation demonstrates a wide capture range and steep gradient for map-map registration with SSD. This suggests that confidence maps could indeed be a valuable feature for mono-modal registration and mosaicing applications. Nevertheless, further evaluation is required and more sophisticated combination of image intensity and confidence map information should be investigated.

4.5.5 Confidence Estimation in IVUS

IVUS provides cross-sectional views of vessels in real-time and is an important tool for assessing cardiovascular diseases like atherosclerosis. Classification of tissues in IVUS images, most notably into fibrotic, lipidic, calcified and necrotic tissue is an important task for assessing vulnerable plaque regions and defining appropriate treatment options [105]. The IVUS grayscale images alone or combined with processing of backscattered RF signals can be employed for further characterization of plaques and more importantly identification of vulnerable ones [137]. As a result, several atherosclerotic tissue classification (TC) algorithms have been developed to assist interventional cardiologists to better characterize plaque constituents and delineate rupture-prone ones [107, 77, 70]. However, the main limitation of commonly used algorithms, such as Virtual-Histology (VH) algorithm [107], is that they systematically classify tissues in hypoechoic regions behind arc




Figure 4.16: Similarity plots of SSD values for different poses: (a) US-US registration (translation), (b) map-map registration (translation), and (c) US-US and map-map registration (rotation).




Figure 4.17: Example from our evaluation dataset showing: (a) IVUS B-Mode, (b) corresponding confidence map, and (c) the corresponding histology.

of calcified plaques [103], where loss of acoustic energy is expected. Similar arguments are also valid for shadowing regions due to tissue-lung interface, where relatively strong signals are recorded but mainly ignored by experts and marked as unreliable information. Moreover, from segmentation perspective, the shadow artifact makes regularization very challenging. For example, a guide-wire close to surface of IVUS transducer element or appearance of side-branch often require special care to provide the final automated detected lumen border in closed smooth elliptical form. Confidences maps could offer the necessary information for tissue-characterization algorithms to improve reliability for detection. Classification in unreliable regions could be improved by taking confidence maps into account as an additional feature or by using confidence information to communicate knowledge of low confidence regions to potential users with application-specific visualization. To evaluate the accuracy and robustness of confidence estimation on IVUS images we used a comprehensive data set of IVUS and histology images. More specifically, the IVUS RF data was acquired from dissected coronary arteries from autopsied hearts 24 hours postmortem. The experimental setup and acquisition protocol is described in [75]. We estimated confidence maps on raw IVUS RF signals, as described in Eq. 4.18, prior to log compression and normalization. Finally, we quantified results using corresponding




Figure 4.18: Illustration of the deformable registration pipeline for mapping labels assigned in histology to confidence maps. Image details: (a) labeled confidence map, (b) IVUS, (c) segmentation mask (dotted-box is the bounding-box for size-normalization), (d) deformation field, (e) labeled histology, (f) histology, and (g) segmentation mask.

histology images. An example of the qualitative evaluation is given in Fig. 4.17. In the confidence map, a bright gray levels indicates high confidence and vice versa. Comparing the map to the B-mode image, we clearly observe a drop of confidence at the IVUS guide-wire, the media-adventitia (MA) border due to hypoechoic property of smooth muscle cells, and at highly reflective interfaces that cause shadowing. Looking at the corresponding histology it becomes clear that these interfaces are at the border of some tissues with especial interest such as calcification and necrotic core that are known factors associated with vulnerable plaques. These regions usually appear as dark areas in IVUS images due to rapid attenuation of signals. The constructed confidence map from RF signals clearly demonstrates that there is not much information and confirms unreliability of tissue characterization results in corresponding regions.

We quantified our results by manually labeling necrosis/calcified and fibrotic/lipidic regions where severe attenuation of ultrasound signals is expected in the former case but not latter, resulting in low (LC) and high confidence (HC) values, respectively. In order to obtain the most accurate and reliable results, the labeling was performed on histology images, where pathological structures were easily distinguishable by histopathologist, and then transformed onto corresponding IVUS image and confidence map. The advantage of the designed experimental setup and acquisition protocol described in [75] was to ensure reliable IVUS-histology matching procedure. Nevertheless, the deformation of inner wall




Figure 4.19: An IVUS grayscale image with arc of calcified plaque at 3-5 o’ clock (a) and corresponding confidence map (b), registered histology (c), PH image before (d) and after (e) weighting with confidence values.

due to formalin fixation that was used prior to histology preparation demanded for a deformable (non-linear) registration between IVUS and corresponding histology image in order to find one-to-one label correspondences. In brief, an intermediate representation of images, which in our case are size-normalized segmentation masks from both IVUS and histology were obtained automatically for the IVUS images and semi-automatically from the histology images using random walks segmentation [38]. Subsequently, deformable registration is computed between the segmentation masks, using the Markov Random Field (MRF) registration framework introduced in [36]¹. The availability of different penalty functions for regularization of the deformation field was decisive for this choice, given additional constraints that needed to be imposed because of missing image information using only binary masks. Finally, the resulting deformation field was utilized to map the HC and LC labels from the histology to the corresponding confidence maps. With this experimental setup the confidence maps could be sampled at HC and LC defined in co-registered histology images. An illustration of the entire registration process and the transfer of labels from histology to corresponding IVUS images is shown in Fig. 4.18. This entire process was repeated for a total of 13 IVUS images and associated labels. The statistics of these sampled regions are presented in Fig. 4.20. The results indicate a good separation between HC and LC regions, which is highly desirable for regularizing in segmentation algorithms and more importantly evaluation of presence of informative regions in regions behind arc of calcified plaques, which has always been a controversial subject in tissue characterization. Our quantitative results confirm that there is not sufficient acoustic information and therefore particular attention must be paid when interpreting tissues in these regions. It is also evident that the confidence values are not distinct for same areas in different data set. This is expected given that sound travels through different mediums with different acoustic impedances to reach the HC and LC region in the different data sets. Essentially, the path to these regions involves different energy reduction for each data set. Although the registration was affirmative, it introduced minor errors due to severe deformation of inner wall in histology images. Consequently, the standard deviation of some sample regions varied proportional to the extend of deformation.

For the first time, we registered IVUS and histology images and illustrated the outcomes. All these, enabled us to search for the most representative regions in ultrasound

¹Framework available at <http://www.mrf-registration.net/>




Figure 4.20: Sampling of confidence maps at HC (blue) and LC (red) regions defined in co-registered histology. Each bar shows the mean and the (+/-) standard deviation of the confidence sampled in the labeled regions.

images or flagged out those that did not carry any informative features. We evaluated our algorithm performance, deploying IVUS RF signals and demonstrated encouraging results. Our results confirmed lack of acoustic energy in regions behind arc of calcified plaques or those with late necrosis. In general, constructed confidence maps could contribute in reducing inter- and intra-variability of experts diagnosis particularly in regions that are challenging to interpret due to uncertainty in image information. A potentially direct application is the weighting of responses from image-based histology, like from IVUS-PH algorithm [75], as shown in Fig. 4.19. This could flag-out regions of low confidence to the physician in order to assist with the decision process. In future, we will use constructed confidence maps to detect guide-wire as well as side-branches to regularize border detection in segmentation algorithms.

4.5.6 Discussion

The introduced confidence maps are not intended to be used independently, but rather in combination with other ultrasound image computing/analysis algorithms; providing an estimate of uncertainty in attenuated and/or shadowed regions. Therefore, we demonstrated the benefit of integrating confidence maps into ultrasound shadow detection, 3D freehand ultrasound reconstruction, ultrasound-ultrasound registration, CT-ultrasound registration and classification in IVUS images. Our goal was not to demonstrate superior segmentation, registration, or reconstruction results compared to other methods in the literature. Instead, we emphasized the generic character of our method for the different applications and acquisition scenarios. Moreover, the global solution to the random walks equilibrium problem comes with an inherent robustness to intensity deviations and noise. As a result, our method was successfully applied on both envelope detected RF and post-processed B-mode images that were acquired from different anatomies using different ultrasound transducers, systems and settings. Although the experiments indicate a good robustness to different ultrasound image acquisition setups, there are many ultrasound

system settings that can be adjusted by a user within a wide range including: focusing, gain, Time-Gain-Compensation (TGC), and spatial compounding [50, 142]. These influence not only the image appearance, but also the outcome of image processing methods to a different extent [165]. Ideally, the systems settings should be included in the confidence map estimation to make the method robust to a wider range of ultrasound system settings.

4.5.7 Conclusion

We presented a generic and novel method for estimating a per-pixel confidence in ultrasound images, which we denoted ultrasound confidence map, that emphasizes the uncertainty in attenuated and/or shadowed regions. The main novelty of our approach is the formulation of the confidence estimation problem within a random walks framework. Domain specific knowledge of ultrasound was integrated into the framework, including: the relation of possible ultrasound transmission to image intensities, depth dependent attenuation and scanline constraints to model the ultrasound beam width. We evaluated the confidence maps on different images and demonstrated the benefit of integrating them into different ultrasound image processing applications.

Future work will focus on integrating ultrasound system settings into the confidence estimation. Of further interest is the extension of the modeling assumptions in order to include the incident angle of the ultrasound beam and the acquisition frequency. We believe that the generic character of our method and the demonstrated applicability to different applications will lead to more sophisticated integration into different ultrasound image processing methods, that could benefit from a confidence estimate.

4.5.8 Acknowledgments

The authors would like to thank Nicolas Brieu and Maximilian Baust for valuable comments, Leo Grady for valuable comments and his help with the revision of this article, Seyed-Ahmad Ahmadi for his help with the ultrasound calibration, and José Gardiazabal for his help with the hardware setup. Further thanks to the anonymous reviewers from Elsevier Medical Image Analysis Journal, whose constructive feedback helped in extending and improving the original manuscript. This work was partly funded by the European projects: "PASSPORT"-Grant Agreement no: 223894 and "ACTIVE"-Grant Agreement no: FP7-ICT-2009-6-270460.

The data for the IVUS applications had been taken in a joint project between the Heffner Biomedical Imaging Lab (Columbia University, New York), Boston Scientific Corp. and Infraredx Inc. We acknowledge the transfer of the data to Technische Universität München through the data transfer agreement MTA39639. We acknowledge the help of Boston Scientific Corporation (Fremont, CA) for providing the 40MHz acquisition system. We also acknowledge Jennifer Lisauskas and Dr. Stephen Sum at InfraRedx (Burlington, MA) who developed the artery fixation cage and their help in collecting IVUS data is acknowledged. The histology samples were processed at CVPPath (Gaithersburg, MD) under supervision of Dr. R. Virmani. Please refer to [76] for a detailed description

of the data collection process.

AUTOMATIC BONE DETECTION AND REGISTRATION

In this chapter we introduce a system for intra-operative ultrasound-CT registration. For this we have developed novel methods for fully automatic bone detection in ultrasound images and automatic registration. Detection of the surface in tracked 3D freehand ultrasound images allows a partial reconstruction of the bone anatomy and registration to pre-operative CT data set. The detection algorithm uses a bone-specific feature descriptor and was evaluated on both in-vivo and ex-vivo acquired ultrasound images. The automatic registration uses a global optimization strategy for surface-based registration to provide a large capture range. The overall system error - including detection, registration, tracking and calibration errors - was evaluated in a cadaver study on a human leg, where bone markers were implanted to provide the necessary Ground-Truth.

5.1 Introduction

Orthopedic diseases are becoming increasingly more common for adults in the Western world; as are the orthopedic interventions performed each year. Minimally invasive surgery is playing a crucial role not only in decreasing the high costs but most importantly in reducing the risk and recovery period for patients. Computer-Aided Orthopedic Surgery (CAOS) provides assistance in the pre-operative planning phase and navigation during surgery.

The acquisition of a pre-operative CT scan is common for patients that will undergo surgery including hip and knee replacement, repair of femoral neck and trochanteric fracture, placement of pedicle screws etc. Intra-operative imaging is usually limited to X-ray and in some cases to fluoroscopic imaging. Surgeons reconstruct a mental image of the patient anatomy during the intervention, taking into account the pre-operative images, the patient position, and available intra-operative imaging. Ultrasound imaging is a viable option for intra-operative imaging. However, the inherent limitations of ultrasound like artifacts, noise, and shadowing reduce its applicability especially when imaging bone anatomy. Fusion of intra-operative ultrasound with pre-operative CT allows the surgeon to make full use of the available diagnostic and planning data and improve the accuracy of the overall procedure. Besides, if a CAOS system utilizes real-time tracking hardware,

fusing anatomic information from CT and tracked ultrasound would allow to transfer the pre-operative CT and planning data into the patient coordinate system during surgery. In the next section we will discuss different applications and methodologies for intra-operative fusion in computer-aided orthopedic surgery.

5.2 Related Work

The focus of earlier work on ultrasound-CT registration of bone anatomy was directed towards feature-based methods. In general, such methods extract the bone surface from CT and ultrasound and register them with algorithms such as the popular Iterative-Closest-Point (ICP) [12]. In the work of [15] the bone surface was manually extracted from ultrasound images and registered to the CT bone surface with ICP. Qualitative results suggested that screw placement was more accurate, however, quantitative results were not presented. In a similar way [7] performed the registration for scaphoid fixation assisted by intra-operative tracked ultrasound. They demonstrated an improvement of the placement accuracy with their system on phantoms. In [4, 5] the manually segmented ultrasound surface was registered to the CT surface and simultaneously the calibration of the freehand ultrasound was refined. This joint optimization of registration and calibration improved the accuracy of the registration compared to previously published approaches.

An alternative branch of ultrasound-CT registration methods pursues the goal of providing patient-specific bone anatomy, combined with intra-operative ultrasound, without the need for a pre-operative CT scan. In general, these methods learn a Statistical Shape Model (SSM) of the bone anatomy from a set of carefully segmented CT training datasets and register the instantiated shape with the bone surface extracted from ultrasound. In [6] a SSM for the femur and pelvis was generated from segmented CT scans by a Principal Component Analysis (PCA) of the principal deformation modes. The bone model was instantiated by non-linearly optimizing the weights of the linear combinations of deformation modes to fit the ultrasound surface as closely as possible. The final instantiated surface is registered to the ultrasound surface with ICP. An average surface misalignment error of about 3.5mm was reported.

Surface-based registration between ultrasound-CT for orthopedic interventions has received attention with the seminal works of Barratt et. al. [4, 6]. The accuracy and robustness of the registration is directly related to the quality of the ultrasound surface detection. Thus the previously discussed approaches strongly depend on the accuracy of the extracted bone surface from the ultrasound images. Manual or semi-automatic segmentation of the bone surface is prone to errors, has limited reproducibility, requires trained users and results in more time for interventions. A plethora of algorithms, including fully-automatic approaches, have been developed for medical image segmentation [143]. Compared to other imaging modalities ultrasound poses further challenges to segmentation given the low contrast, low SNR ratio, and ultrasound specific imaging artifacts. Hence many algorithms have been adjusted for the specific needs of ultrasound segmentation. For a comprehensive review of ultrasound segmentation we refer to [111]. The extraction of the bone surface in ultrasound is not a segmentation task in the classical sense, but rather a detection of only the bone surface. This task is by itself of no imminent diag-

nostic importance. However, as previously discussed, in the context of CAOS it becomes very important, as it allows fast surface-based registration. In [21] the Otsu threshold was used for defining a fuzzy intensity mapping in which the osseous could be segmented automatically by taking into account prior knowledge. A preliminary evaluation on cadaver and patient datasets for segmenting the bone surface of the sacrum demonstrated an alignment error of less than 1mm. In [29] heuristic pre-processing is used in order to enhance edges based on their reflective properties and the shadows beneath, followed by a dynamic programming approach for detection. While there are some absolute thresholds in this method, it worked well on 2D ultrasound images from two cadaver data sets. A theoretically well founded series of works utilizes local phase features for bone detection and subsequent registration for CAOS applications [45, 14]. However, defining local phase filtering parameters that produce consistent results for image series is difficult. In [45] 3D local image phase features were introduced obtained with 3D Log-Gabor filter banks for extracting the bone surface. Bone surfaces and fractured fragments were automatically extracted from 3D ultrasound. Quantitative results on a human sawbone phantom and ex-vivo porcine tibia and fibula showed an error below 1mm. Qualitative results were presented for realistic in-vivo experiments. In the work of [46] an optimization strategy is introduced for 3D log-Gabor filter parameter selection. Nevertheless, selection of speckle in images is required for parameter estimation.

Intensity-based registration approaches offer a viable alternative for ultrasound-CT registration without the need of extracting the bone surface from ultrasound. In [116] an intermediate representation of CT and ultrasound was used for the registration. Features identifying bone structures were defined for both CT and ultrasound giving the probability that a voxel is a bone edge. These probability maps are registered with an intensity-based method and offer improved accuracy. However, a large set of pre-segmented CT and ultrasound volumes is required for defining the features for each specific anatomy to be registered. In [157] an intermediate representation of the CT was created by performing a simulation of ultrasound images from the CT at the position of a virtual transducer probe. Subsequently, the simulated ultrasound was registered with the real ultrasound using a novel similarity measure denoted *LC2*. The method demonstrated high accuracy and robustness for ultrasound-CT registration of abdominal anatomy, i.e., liver and renal datasets. This method was extended from rigid to non-rigid registration in [159]; while the registration time was significantly reduced with a custom GPU implementation. In [35] ultrasound-CT registration of the lumbar spine was performed with the method of [157]. Here each vertebra was treated as a sub-volume and transformed individually. The possible transformations of the sub-volumes were restricted by a biomechanical model. For each iteration of the registration process the sub-volumes were merged into a single volume for which the *LC2* metric was computed. In [79] the SSM approach was combined with the intensity-based registration proposed in [157]. Instead of a SSM of surfaces a Statistical-Appearance-Model of volumetric intensities was generated from CT values. All volumes were registered to each other with a non-rigid approach. The modes of shape variation were extracted with PCA. In each iteration of the optimization a new instance of the CT volume is generated, ultrasound simulated for it and registered to the real ultrasound with the *LC2* metric. A Target Registration Error (TRE) of 3.5mm

was reported, however, the method is computationally very demanding requiring hours of processing time on a modern high-end workstation. In [131] bone surfaces extracted from MR are registered to ultrasound volumes with standard intensity based similarity measures. Taking into account the shadowing artifacts caused by the bone-tissue interface significantly improved the accuracy of the registration.

Subsequently, we introduce a new method for fully-automatic ultrasound bone detection in B-mode images that is robust and provides accurate results for images acquired in-vivo, ex-vivo, and from phantoms in real-time on commodity hardware. The detection method is based on a new bone-specific feature descriptor that uses the previously presented confidence maps. The bone surfaces extracted from ultrasound and CT are rigidly registered with a global optimization strategy; demonstrating a large capture range, high accuracy and robustness.

5.3 Detection Method

We introduce a new feature descriptor designed to produce strong responses at possible bone surface(s) and low responses elsewhere, which we denote *bone confidence localizer*. This intermediate representation serves as a basis for subsequent surface detection. Our feature descriptor requires two inputs: 1) filtered ultrasound image and 2) ultrasound confidence map. An orientation dependent pre-filtering of the ultrasound image I , see Fig. 5.1(a), is performed to enhance possible bone structures by convolving the image along lines with $0, \pm 45^\circ$ direction using the following polynomial kernel:

$$h(x) = 2x^2(x^2 - 2) + 1; \quad x \in \{-1.3 \dots 1.3\} \quad (5.1)$$

The absolute value of the sum of the three direction vectors multiplied with the kernel response if it is positive, is used as filter output, and added to the original ultrasound image intensity, forming the filtered image output F ; see Fig. 5.1(b). This step incorporates the prior knowledge that we are seeking non-vertical ridges in the image, i.e. we expect the reflection both before and behind the bone surface to be weaker than directly on it. The confidence maps have a high confidence drop after tissue-bone interfaces, which are in fact reflecting most of the incident ultrasound wave and transmit only a small part of it (see Fig. 5.1(c)). Bone surface in ultrasound is usually observed as a bright image structure followed by acoustic shadowing. Thus, a strong intensity response followed by a low confidence value is a good indicator for possible bone surface(s). We define our feature descriptor by combining the filtered image and the confidence map as follows:

$$L_{ij} = \sum_{i=m-s}^{m-1} \sum_{n-s}^{n+s} F(i, j) - \sum_{i=m+1}^{m+s} \sum_{n-s}^{n+s} M(i, j), \quad L_{ij} = \begin{cases} L_{ij} & \text{if } L_{ij} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (5.2)$$

where i, j are row and column indices, M is confidence map and $F, M \in [0 \dots 1]$. The discrete kernel size is set constant to $s = 7$ for all subsequent experiments. An example of this feature responses is shown in Fig. 5.1(d).

Initially, for given image I_n the feature response is computed as the average of independently computed feature responses for the image neighborhood I_{n-1}, I_n, I_{n+1} . The reason

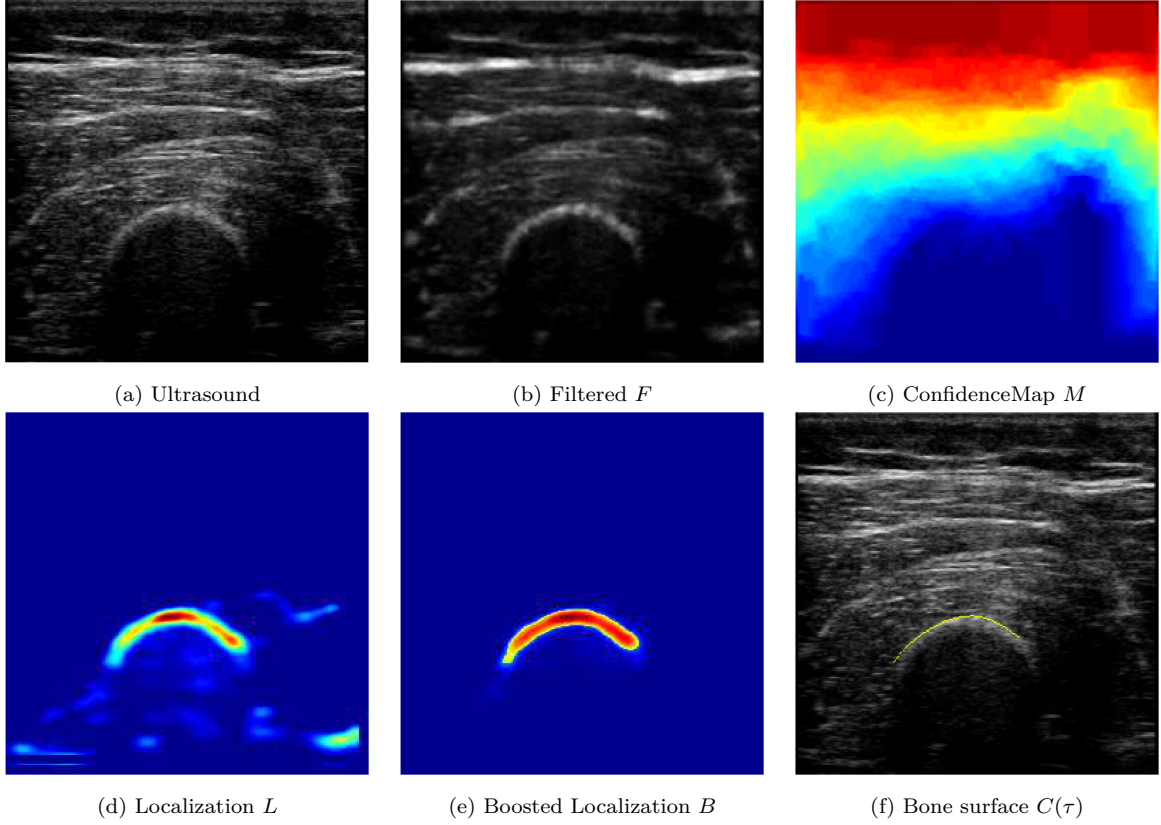


Figure 5.1: Some of the images and intermediate outputs for bone detection.

for this lies in the locality of some ultrasound imaging artifacts, i.e., artifacts might only be present in one of the images of the neighborhood, although image correlation is high because of close proximity in acquisition. Therefore averaging the independent feature response improves the robustness to local artifacts and noise. The search space for bone detection is further constrained by reducing feature responses as a function of normalized unsigned distance from the previously detected contour, as described in the work of [9], by convolving feature response image with:

$$W(x) = \begin{cases} \frac{1-d(x)}{R} & \text{if } x \in \Omega \text{ and } d(x) < R \\ 1 & \text{if } x \in C_S \\ 0 & \text{otherwise} \end{cases}, \quad d(x) = \min_{\hat{x} \in C_S} \|x - \hat{x}\| \quad (5.3)$$

where R is the cut-off distance, $C(\tau) = [x(\tau), y(\tau)]$ the parametrized representation of the contour using a cubic B-spline, and $C_S = \{C(\tau) : \tau \in [0 \dots 1]\}$ the contour set in the image domain Ω . Throughout all experiments the cut-off distance has been set constant to $R = 9mm$; whereas experiments with the range $R \in [9mm..12.5mm]$ provided comparable results.

Responses at pixel locations x are enhanced that are close to the previous detection and similar to the Gaussian distributed intensities of the previous feature responses $L(I_{n-1})$

as follows:

$$S_L = \{L(x) : d(x) < R\} \quad (5.4)$$

$$B(x) = L(x) + \alpha \cdot \exp\left(-\frac{(L(x) - \bar{S}_L)^2}{2M^2}\right) \quad (5.5)$$

where $d(x)$ is the minimal Euclidean distance to previous contour $C(\tau)$, $M = \max(S_L)/2$, $\alpha = M - \frac{\bar{S}_L}{2}$ and $R = 4.5\text{mm}$ for all experiments. An example of this enhanced responses is shown in Fig. 5.1(e). The bone surface is extracted with an approach similar to [85, 9]. In this work statistical thresholding is applied by removing responses that are not in the pixel set $\{x : B(x) < B_\mu + B_\sigma\}$, where B_μ and B_σ are mean and standard deviation of enhanced feature response image respectively. Only the $N = 15$ highest responses per scanline are retained - as proposed in [85] - and the bone region is identified as the largest component. For this morphological closing and opening is performed, to remove outliers, followed by connected component analysis. Fitting of a cubic B-spline results in the final smooth bone surface representation - see Fig. 5.1(f) - required for surface registration.

5.4 Surface-based Registration

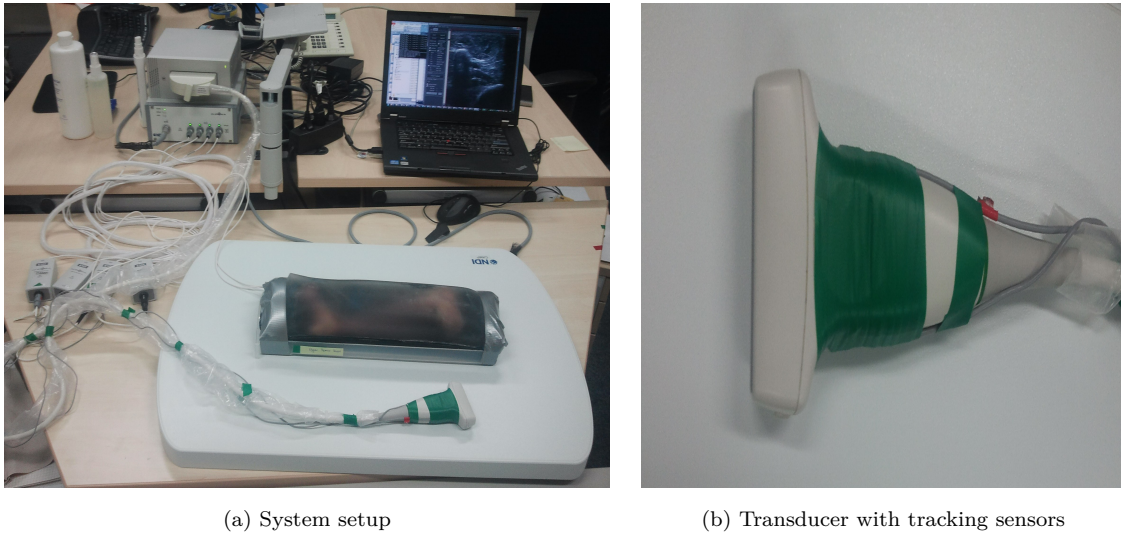
For registration we seek transformation parameters \vec{p} , minimizing a cost function, which represents the distance of all spline control points from the ultrasound bone detection to the CT bone surface:

$$\hat{\vec{p}} = \arg \min_{\vec{p}} \sum_{i=1}^n \sum_{j=1}^{n_i} |D_{bone}(T_i u_{ij})|; \text{ with } T_i = T_{reg} T_{track_i} T_{cal} \quad (5.6)$$

where u_{ij} are the bone surface points in homogeneous coordinates along the fitted splines in each ultrasound image, n_i the number of such points in a certain frame i , n the number of ultrasound frames, T_{track_i} the tracking matrix of frame i (provided by the position sensing system). T_{reg} is the rigid registration matrix, which is parameterized with \vec{p} . Optionally, the calibration matrix T_{cal} (denoting where the tracking sensor/target is with respect to the ultrasound image plane) can be refined by adding some or all of its parameters to \vec{p} . Local optimization strategies in combination with sparse surface reconstruction from ultrasound can produce unfavorable registration results. We address this problem with a global optimization strategy by using Controlled Random Search (CRS) with local mutation [65]. More specifically, during the optimization the registration is repeatedly launched within an exhaustive 2D sampling of the cranio-caudal translation, i.e. along the bone axis (henceforth denoted t_z), and the rotation around it (r_z), since these are the parameters likely to cause ambiguities. The low computational requirements of the cost-function makes this feasible on modern systems.

5.5 3D Freehand-Ultrasound System Setup

Ultrasound image quality depends on multiple factors including imaging system, settings, and patient physiology. High-end ultrasound systems provide advanced image capabilities,



(a) System setup

(b) Transducer with tracking sensors

Figure 5.2: System overview

however, their size can restrict the options for developing new solutions for intra-operative ultrasound-guided orthopedic surgery. We identified a favorable option, namely small systems provided by Teleded Ltd., Vilnius, Lithuania. For the cadaver and volunteer study the Teleded LogicScan 128 EXT-1Z Kit system was used together with a Linear Transducer (HL9.0/40/128Z and HL9.0/60/128Z). The image system is comparatively small (95 x 210 x 215, W x D x H, mm), weighs 3,3 kg and is connected to commodity PC over USB connection. The small size and low price make it a preferable choice for providing a stand-alone system for intra-operative navigation. Although the image quality is not as good as high-end systems, it is more than sufficient for accurate bone detection and intra-operative ultrasound-CT registration - as presented in the subsequent evaluation section. The acquisition of 3D freehand ultrasound is performed by electromagnetic tracking of the ultrasound transducer with a NDI (Northern Digital, Waterloo, Ontario, Canada) Aurora tracking system and the NDI Tabletop Field Generator. For this purpose, two tracking sensors are glued to the ultrasound transducer and the tracking information is sent to PC via USB connection. The bone markers are connected to the same NDI base station as the tracking sensors on the transducer. An image of the system setup (before cadaver study) is shown in Fig. 5.2(a) and of the transducer in Fig. 5.2(b).

Tracking for generation of 3D freehand ultrasound requires accurate positioning sensors that are not interfering with the surgical/clinical workflow. For this purpose Electromagnetic (EM) tracking was chosen as the most viable option as small tracking sensors are not interfering with the ultrasound transducer handling. Furthermore, a direct line of sight to the tracking sensors is not required - as with optical tracking systems - and allows more seamless integration into the clinical workflow. We acquired 3D freehand ultrasound by fixing two EM sensors at the ultrasound transducer. We chose two sensors for redundancy and for avoiding orientation related tracking errors by taking the mean position of both sensors. Furthermore, acquisitions were post-processed by applying Median filtering to reduce the effect of outliers in the tracking data.

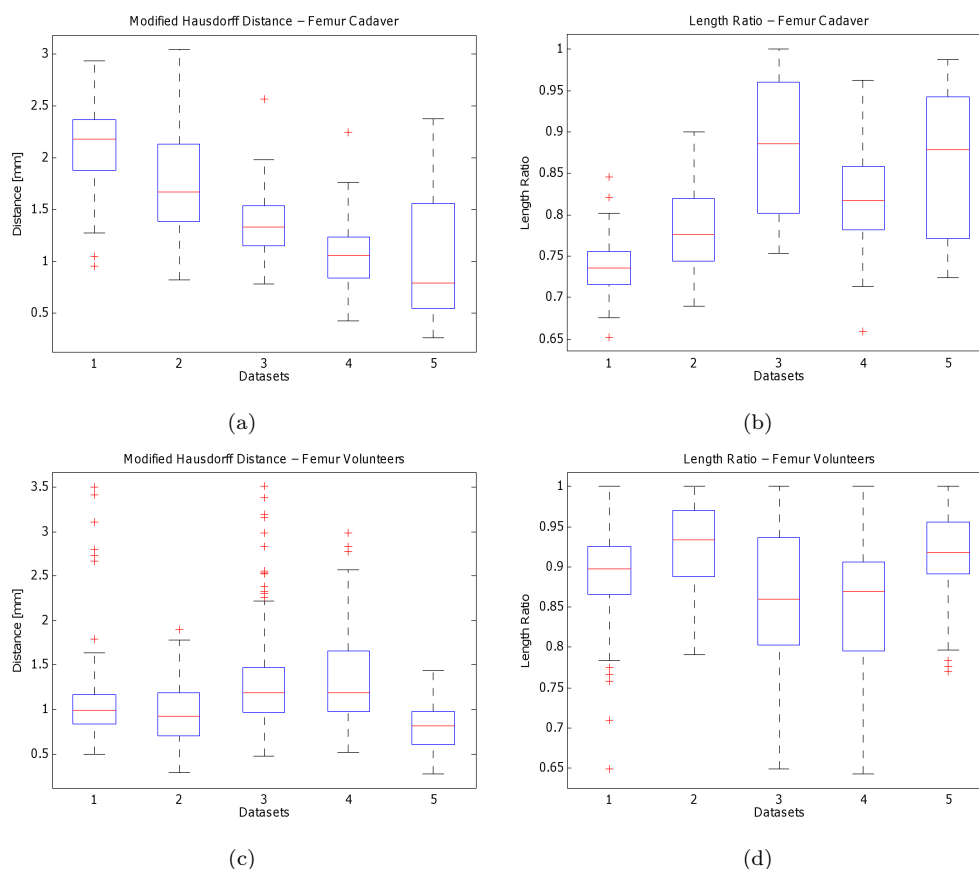


Figure 5.3: Quantitative evaluation of bone detection result for **femur** of five cadaver (images a-b) and five volunteer (images c-d) datasets. The datasets of the cadaver were recorded from the same specimen. The datasets for the volunteers were recorded from different persons with varying physiology. For each dataset 100 images were manually segmented and compared to automatic detection results.

The transformation from the recorded ultrasound image plane to the attached tracking sensors is determined through so-called calibration. Determining this transformation is necessary to establish the complete transformation chain from online acquired ultrasound images to tracking system to co-registered CT data set. For this purpose we used the method introduced in [158] and performed calibration with the EM tracking system and both transducers. For the calibration 3D freehand ultrasound was acquired from a CIRS (Norfolk, VA, USA) 3D abdominal phantom (Model 057), which contains characteristic landmarks for aforementioned registration-based calibration. A CT scan of this multi-modality phantom was acquired and utilized for evaluating the correctness of calibration. Alternatively, other phantoms with similar properties can be used as long as characteristic landmarks are visible from perpendicular acoustic windows and are not displaced by probe pressure.

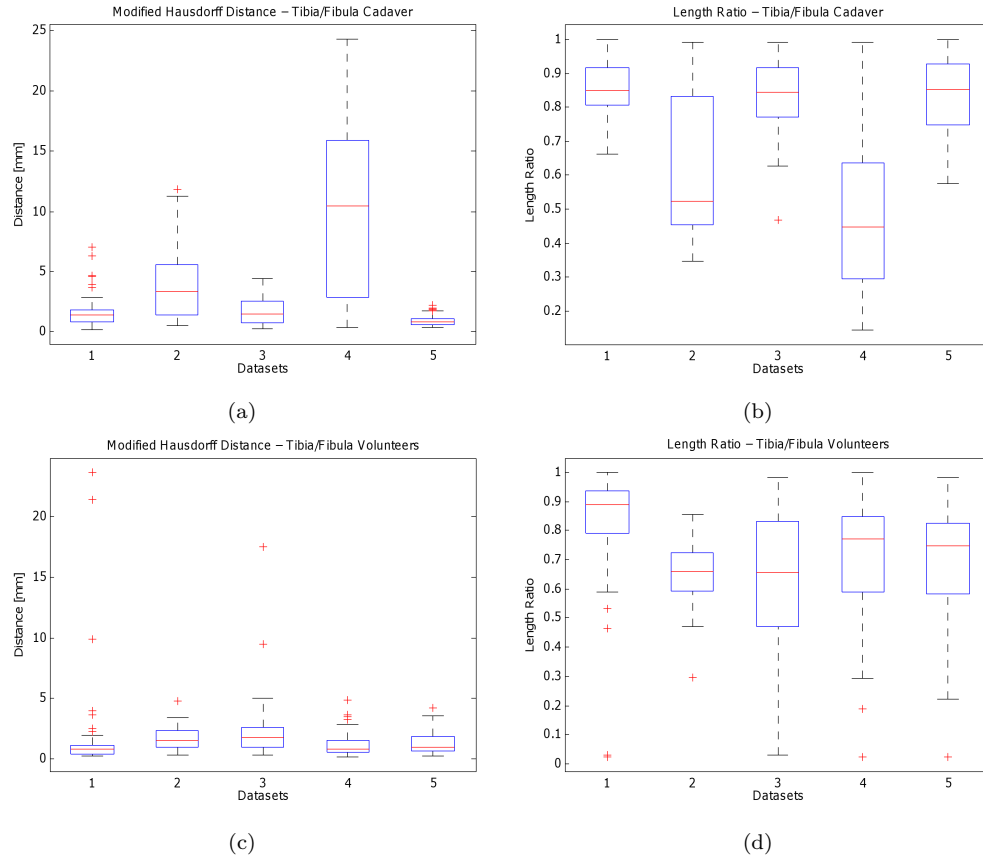


Figure 5.4: Quantitative evaluation of bone detection result for **tibia/fibula** of five cadaver (images a-b) and five volunteer (images c-d) datasets. The datasets of the cadaver were recorded from the same specimen. The datasets for the volunteers were recorded from different persons with varying physiology. For each dataset 100 images were manually segmented and compared to automatic detection results.

5.6 Results and Evaluation

We evaluated automatic bone detection on both in-vivo and ex-vivo data sets. Following all ethical guidelines we evaluated the overall system accuracy and robustness in a cadaver study on a human leg, where markers could be implanted for assessing accurately Fiducial Registration Error (FRE). High-resolution CT was acquired with a Philips Brilliance 64 system, voxel spacing $0.63 \times 0.63 \times 0.5\text{mm}$. Our 3D freehand ultrasound setup used a Teleded LogicScan 128 EXT-1Z Kit ultrasound system with a Teleded HL9.0/40/128Z transducer calibrated to a NDI Aurora electro-magnetic (EM) tracking system. A multitude of partial and complete freehand scans were acquired in various orientations on the femur, tibia and fibula bones. In addition, two EM bone markers with calibrated tracking sensors were implanted into the femur and tibia bone. A second control CT scan was acquired in order to compare the 3D marker locations in it with the positions predicted by the transformation chain of tracking sensor readings and our ultrasound-CT registration. This allows the computation of the overall system error, which includes the detection and registration errors, as well as calibration and tracking errors. A rendering of the CT data set of femur and tibia/fibula - before and after introduction of fractures and

bone markers - is shown in Fig. 5.10. The bone detection was also evaluated for in-vivo image acquisitions from volunteers using the discussed ultrasound system setup with the Telemed HL9.0/60/128Z transducer.

5.6.1 Detection Evaluation

For the quantitative evaluation we followed the following procedure. We manually labeled the bone contour in a subset of the images from different datasets to provide the necessary Ground-Truth and executed automatic bone detection on the same image sets. The contours were represented as point sets $\mathcal{A} = \{x_1 \dots x_m\}$ and $\mathcal{B} = \{x_1 \dots x_n\}$ for manual and automatic delineation respectively. Direct one-to-one comparison of the points sets is not possible, as they do not have a direct mapping and usually have different cardinality. In these circumstances the Hausdorff distance is a common metric for comparing two non-empty point sets. For image processing applications and pattern recognition tasks the modified Hausdorff distance is an appropriate choice to reduce the impact of outliers in the sets [63]. The metric is defined as follows:

$$H(\mathcal{A}, \mathcal{B}) = \max(h(\mathcal{A}, \mathcal{B}), h(\mathcal{B}, \mathcal{A})) \quad (5.7)$$

$$h(\mathcal{A}, \mathcal{B}) = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} \min_{b \in \mathcal{B}} \|a - b\| \quad (5.8)$$

where $\|\cdot\|$ the Euclidean distance. To assess the extend of the surface reconstruction for surface-based registration, the length ratio of the automatically detected and manually delineated contours is compared as follows:

$$R(\mathcal{A}, \mathcal{B}) = \begin{cases} \frac{|\mathcal{A}|}{|\mathcal{B}|} & \text{if } |\mathcal{A}| < |\mathcal{B}| \\ \frac{|\mathcal{B}|}{|\mathcal{A}|} & \text{otherwise} \end{cases} \quad (5.9)$$

Essentially, the modified Hausdorff distance provides a measure of spatial accuracy and robustness. The length ratio resembles the proportionality of the detected surface, i.e., how different is the size of the Ground-Truth contour compared to the automatically detected contour. These metrics were computed for 2000 manually segmented images of the femur and tibia/fibula from ex-vivo cadaver and in-vivo volunteer datasets. More specifically, for every image the contour was manually delineated and compared to the contour defined by the automatic bone detection. Exemplary detection results for femur anatomy are shown in Fig. 5.6 and Fig. 5.7 and for tibia/fibula anatomy in Fig. 5.8 and Fig. 5.9. Images from in-vivo volunteers demonstrate differences in intensity distribution and tissue anatomy; depending on the physiology of the volunteer. For example, images in first row were acquired from normal weight and in second row from overweight volunteer with the same imaging setup and settings. The comprehensive quantitative results are shown in Fig. 5.3 and Fig. 5.4. Please notice that the box-plots (box-and-whisker diagram) contain the measures of Hausdorff distance and length-ratio for 100 images of each datasets - providing a good overview of detection variability within the data sets. The combined statistics for all datasets are presented in Table 5.1.

	Cadaver	Volunteers
Femur - Hausdorff	$1.4664 \pm 0.5884[mm]$	$1.1229 \pm 0.5384[mm]$
Femur - Length	0.8159 ± 0.0826	0.8896 ± 0.0724
Tibia/Fibula - Hausdorff	$3.6295 \pm 4.9553[mm]$	$1.5246 \pm 1.8550[mm]$
Tibia/Fibula - Length	0.7281 ± 0.2185	0.7151 ± 0.1849

Table 5.1: Overall bone detection results from quantitative evaluation (mean \pm standard deviation).




Figure 5.5: Example of US-CT registration result.

The detection on femur anatomy for both in-vivo and ex-vivo data set is accurate, robust and comparable. The detection on tibia/fibula anatomy is good for in-vivo data set, but suboptimal for ex-vivo data set. This outcome is dependent on the choice of the acoustic window, i.e the detection becomes difficult when the bone is very close to the skin, where numerous strong reflections occur. The influence of such outliers on the averaged evaluation result is emphasized by data set no. 4 in Fig. 5.4(a).

5.6.2 Registration Evaluation

We evaluated registration accuracy using the data of the previously described cadaver study. An example of input data and registration result is shown in Fig. 5.5. Table 5.2 contains detailed evaluation for all the sweeps acquired in the study, alongside with the point-based error in mm after automatic registration and the rotational error in degrees. The latter was determined as the angle between the bone marker axis determined from CT, and the axis as given by the tracking data. Additionally, motion of subject (leg) during recording was captured with bone markers and compensated for. On average over all sweeps, a median registration error (evaluation with one or two bone markers) of $3.7mm$ is achieved. For two of the tibia and fibula data sets registration accuracy was low because of incomplete surface detection as bone was very close to the skin. We further performed a randomized trial of 100 test runs with $\pm 45mm / \pm 45^\circ$ from the manually defined initial poses on a subset of data. Table 5.3 summarizes the mean error values with respect to Ground-Truth, percentage of outliers, as well as the standard deviation of all transformation parameters (excluding outliers). The overall variation within the cluster of converged registration poses is fairly small ($1.5mm$ and 0.7°), demonstrating the ro-

no.	Tibia/Fibula						Femur				
	1	2	3	4	5	6	7	8	9	10	11
tran.	12	3.0	6.0	22.7	3.7	0.5	2.6	3.2	3.4	7.4	4.2
rot.	2.7	5.0	6.9	5.5	1.7	2.6	1.8	3.1	9.7	2.9	2.7
motion	0.3	2.3	6.4	7.2	0.3 0.4	2.9 0.5	2.3	11	23.6	8.8	1.3 0.5

Table 5.2: Description of data sets containing bone markers and overall system errors. Rows 'tran.' denote the translation error, 'rot.' the rotational error in degrees and 'motion' their overall motion during the sweep in mm. Two motion entries refer to two simultaneously tracked markers during acquisition.

No.	FRE		outliers	standard deviation					
	trans	rot		t_x	t_y	t_z	r_x	r_y	r_z
2	6.0	1.4	23%	0.0634	0.2466	1.4543	0.1913	0.0290	0.2845
7	2.7	2.5	22%	0.0914	0.0300	0.3918	0.0697	0.0318	0.3306
8	3.6	2.4	10%	0.1274	0.0374	0.3359	0.1393	0.0475	0.7080
9	3.8	9.3	10%	0.0935	0.0368	0.3574	0.0673	0.0113	0.1419

Table 5.3: Detailed random study evaluation results on different data sets.

bustness of the registration framework. For all studies fully automatic registration was achieved with a default stochastic population size of 2000 for the CRS optimization. We chose this value by examining the effect of population size on optimization after applying uniformly randomized registration parameters within the range of $\pm 90mm / \pm 90^\circ$.

5.7 Conclusion

We have presented a system and algorithms to automatically register intra-operatively acquired ultrasound on bones to CT. This will eventually allow to seamlessly use ultrasound as the “glue” connecting pre-operative data with the interventional setup for CAOS, without cumbersome manual interaction. The algorithm automatically detects bone in live ultrasound and reconstructs a surface point-cloud in real-time as the sweep is acquired. A global optimization strategy allows accurate and automatic surface registration with a wide capture range. Feature computation and bone detection require ≈ 70 ms per image and overall registration $\approx 5 - 30$ seconds¹, which is more than acceptable in the context of surgical time constrains. For evaluation data sets were acquired in a cadaver study where CT scans were acquired before and after introducing fracture and implanting trackable bone markers. This provided the necessary Ground-Truth for evaluating the overall system error, which includes registration, tracking and calibration (spatial and temporal) errors. The evaluation of this error is important for assessing the perspective usability of such approaches for computer-aided interventions. Additionally, bone detection was evaluated on data sets acquired from volunteers. In our future work we will focus on improving detection results for tibia/fibula by using additional prior knowledge or a

¹C++ implementation on laptop with Intel i7-2720QM processor

simultaneous detection-registration approach [9].

5.8 Acknowledgments

This work was funded by Depuy Synthes - a Johnson&Johnson company. We would like to thank Adrian Baumgartner from Depuy Synthes for the project management and organizing the cadaver study performed in the scope of this project. Furthermore we would like to thank Seyed-Ahmad Ahmadi for valuable discussions and Julian Beitzel for his work on this project in the scope of his Master thesis.




Figure 5.6: Examples of successful bone detection using **femur dataset from a cadaver**.




Figure 5.7: Examples of successful bone detection using femur dataset from five volunteers.




Figure 5.8: Examples of successful bone detection using **tibia/fibula dataset** from a cadaver.





Figure 5.9: Examples of successful bone detection using tibia/fibula dataset from five volunteers.



(a)



(b)



(c)

Figure 5.10: Renderings of CT data set acquired in cadaver study (a) femur, (b) tibia/fibula and (c) tibia/fibula after introduction of fracture and bone markers. Visualization with Exposure renderer [88].

Part III

Conclusion and Appendix

CONCLUSION AND OUTLOOK

The goal of this dissertation was to provide further insight onto the modeling and simulation of ultrasound and introduce prior knowledge of the image formation into medical image processing applications. The main contribution of this work is a novel method for estimating confidence in ultrasound images, particularly in shadowed and/or attenuated regions, which we denoted ultrasound confidence maps. In addition, this new confidence information was introduced into common ultrasound image processing applications; effectively improving the results of the underlying methods.

Initially, this dissertation focused on the modeling and simulation aspects of ultrasound imaging. For this purpose we discussed the essential physics, image formation, signal processing and advanced imaging capabilities of ultrasound to gain insight on the modality. Subsequently, we introduced a novel framework for fast and realistic ultrasound simulation. For this the nonlinear full-wave equation was used as the underlying model for sound propagation and solved with a parallelized Finite-Difference Time-Domain approach on GPU hardware. The platform simulated the essential ultrasound image formation pipeline including pulse transmission, beamforming and signal processing. We presented realistic simulation results obtained from computer phantoms and phantom generated using histology data sets. In the later case histology images were registered to IVUS images, which allowed direct comparison of simulated and real images. Although the results were promising they also emphasized the problems in “reverse-engineering” of ultrasound imaging through simulation. The parameter space for simulation involves: speckle distributions, speed of sound assignments, tissue parameters, focusing schemes, number of elements, aperture size, signal processing, frequency, and more. Advancements in GPU hardware and new models for accurate simulation of ultrasound propagation could further pave the way for analyzing this complex parameter space in the imminent future. This would further improve our understanding of ultrasound propagation in inhomogeneous mediums, of tissue properties and could led to new imaging methods.

In the second part of this thesis we introduced ultrasound confidence maps; a novel method for estimating relative confidence in ultrasound images particularly in attenuated and/or shadowed regions. The estimation problem was formulated as a random walk on a graph taking into account ultrasound specific constrains. Ultrasound transmission,

depth-dependent attenuation and beam-width were introduced into the random walks framework [38] with new weighting functions, thus, preserving the favorable properties of the original approach. The global and analytic solution to this equilibrium problem comes with an inherent robustness to noise. This allowed consistent computation of confidence maps for images acquired with different transducers, settings and systems. This new relative measure of confidence was not intended to be used independently, but rather in medical image processing applications. Consequently, confidence maps were integrated into different methods, as for example in shadow detection. There simple thresholding of confidence maps already demonstrated good segmentation of shadow regions. In future work confidence maps could act as a feature for accurately detecting shadow regions that contain strong reverberation artifacts. For example, a possible direction could be the combination of speckle classification combined with confidence information. Confidence maps were also used for 3D freehand ultrasound reconstruction of femur anatomy acquired from multiple viewing angles. The goal in this application was to combine the complementary information from multiple ultrasound images into a single Cartesian volume. Shadow areas are reducing the contrast in the reconstruction as they contribute pixels with noise and low SNR to each reconstructed voxel. Therefore contribution of shadowed (low confidence) regions to the reconstructed volume was suppressed by introducing confidence maps into the reconstruction function. In future work further improvements could be expected when confidence maps are combined with alternative reconstruction schemes, e.g. reconstruction from RF data [83]. For multi-modal registration of ultrasound to CT, confidence maps from ultrasound were registered to transmission simulated from CT to effectively increase the capture range and accuracy of intensity-based registration method [157]. Similarly, for mono-modal registration confidence maps of ultrasound images were registered instead of direct image-based registration. Again the results indicated a better capture range and a steeper cost-function, which could be favorable for gradient-based optimization schemes. In future work combining both confidence maps and intensity information into a single cost-function might further improve the registration results. Last but not least, confidence maps computed for IVUS images were evaluated by reliably defining low and high confidence regions in corresponding histology images, and mapping them back to confidence maps with deformable registration. A clear separation between high and low confidence regions was observed, as indicated in the corresponding histology data. In addition, we demonstrated the potential use of confidence maps for visualizing the result of tissue classification obtained from IVUS images, by controlling transparency of tissue labels based on their corresponding confidence values. Our most recent works [132, 133] demonstrated the use of confidence maps as feature for tissue classification algorithms; providing additional information to machine learning approaches.

In the last chapter of this dissertation an ultrasound-guided navigation system was introduced for computer-aided orthopedic surgery. The goal was to perform intra-operative acquisition and registration of ultrasound to pre-operative CT, in order to fuse ultrasound and CT data. This would allow transfer of pre-operative planning and imaging data to operation theater. For this a novel feature descriptor was introduced for bone detection; based on confidence maps. The feature was utilized in an algorithm to perform fully automatic detection of bone surface in real-time on ultrasound images. The detected

bone surface from ultrasound was rigidly registered to the bone surface segmented from corresponding CT data set. To avoid local minima during registration a global optimization strategy was applied by performing exhaustive 2D sampling of the cranio-caudal translation along the long axis of the bone, where most ambiguities during registration occur. Furthermore, the Controlled Random Search (CRS) with local mutation optimization method contributed to the improvement of capture range. The bone detection was evaluated extensively on ultrasound images acquired in-vivo and ex-vivo. The overall system error including calibration, tracking, detection and registration error was evaluated in a cadaver study on a human leg. For this bone markers were implanted, a CT scan was acquired and the ultrasound images were acquired together with the tracking information from the markers. The evaluation with this Ground-Truth data provided a first proof-of-concept for the applicability of such ultrasound-guided navigation system for computer-aided orthopedic surgery. In addition, it supported the use of confidence maps for solving challenging ultrasound image processing problems. In future work improvement of bone detection for anatomy where bone and skin surface are very close could further improve registration result. This could be achieved by integrating further prior knowledge, as for example Statistical-Shape-Models, in the detection method.

Confidence maps were introduced into algorithms for different ultrasound image processing applications; effectively improving algorithmic results. In future work other applications could also benefit from confidence information such as segmentation or tissue classification. Furthermore, confidence maps could further be improved by integrating system specific settings like Time-Gain-Compensation, frequency and focusing into the estimation process. In conclusion, confidence maps offer valuable, versatile and complementary information that has demonstrated its potential when integrated into different ultrasound image processing applications.

VISUALIZATION OF 4D DATA SETS

Volumetric and flow visualization are established fields offering a wide spectrum of techniques for visualizing $3D + t$ datasets. However, visualization of breathing motion in medical data sets is posing new challenges. As opposed to flow visualization we are not directly interested in the underlying flow, but in the deformation caused by the flow. Therefore, visualization of breathing motion focuses on emphasizing the organ motion while preserving the anatomical context provided by the volumetric visualization. In this work we will discuss methods from flow and volume visualization, their applications, and introduce alternative visualization approaches for enhancing the perception of organ motion due to breathing. In a user-survey we will analyze the perceptual benefit of different methods and discuss the future outlooks.

A.1 Introduction

Advancements in medical imaging do not only improve the quality of the images, but also increase the amount of data available to the physicians, a trend that is likely to continue in the future. Visualization has always played a crucial role in presenting the available data in a coherent and intuitive way, assisting physicians and clinical personnel in diagnosis, planning, and treatment. The role of visualization becomes even more important with higher-dimensional datasets like for example high-angular resolution diffusion imaging (HARDI). These complex datasets can only be interpreted with the help of sophisticated visualization techniques [124].

Relatively recent advancements in medical image acquisition allow us to acquire multiple volumetric datasets over time ($3D + t$), which we will also refer to as 4D datasets. These datasets usually come from 4D ultrasound, 4D CT, and 4D MRI. The established fields of volumetric and flow visualization offer a variety of techniques for visualization of $3D + t$ datasets. Visualization of volumetric datasets has constantly improved in the previous years as a result of hardware and algorithmic advancements [25]. The field of flow visualization has a long history and introduced sophisticated methods for the visualization of high-dimensional datasets. For a more comprehensive review and categorization on flow visualization we refer to Laramee et al. [94].

The visualization of organ and tissue motion is introducing new challenges. Although, flow visualization techniques are directly applicable for the visualization of flows in volumetric medical datasets they do demonstrate limited applicability in visualizing organ and breathing specific motion. In this work we are not directly interested in visualizing the flow in $3D + t$ medical image datasets, but in visualizing the effect of the flow, i.e., the tissue and organ motion subject to the flow. Medical education and training is one of the major application areas of such visualizations, which makes it important to preserve the anatomical context in the visualization. In the next sections we will discuss the applicability of methods from volume and flow visualization to our problem. Furthermore, we will introduce an alternative visualization approach that emphasizes motion through the volume rendering transfer function.

A.2 Previous Work

In flow visualization we are mostly concerned with visualizing steady and unsteady flows. Laramee et al. [94] categorized flow visualization into direct, texture/dense, and geometry visualization. Direct visualization is performed with the help of glyphs, most commonly arrows, that follow the direction of the flow and can encode further properties like magnitude. However, these visualizations become rapidly overwhelming with increasing number of glyphs required for the visualization of detailed and complex flows. Texture/dense approaches provide flow information for every pixel. The most popular approach is the Line-Integral-Convolution (LIC) method [94], which convolves a noise texture with the flow. The disadvantage of this method is that it is difficult to distinguish specific characteristics of the flow. The geometry approaches make use of stream, streak, and path lines [123]; and offer a compromise between direct and dense flow visualization by emphasizing important characteristics of a flow.

Ray-casting has been established as the standard for volumetric visualization, especially with the introduction of GPU accelerated ray-casting [89]. Transfer functions are playing the most important role for providing the desired visualization as they allow the implicit segmentation of anatomical structures of interest interactively during rendering. Volume rendering has the benefit of providing visualization directly derived from the volumetric data set; preserving the anatomical context with the use of appropriate transfer functions. For a more comprehensive overview we refer to Engel et al. [25].

The visualization of $3D + t$ medical datasets has already been addressed for different application domains, however, to the best of our knowledge no work is specifically addressing the problem of visualizing organ and breathing specific motion. The closest work we found was from Handels et al. [48] focusing on computing tumor mobility for radiation therapy. The motion of the lung and the tumor were visualized with vectors and color coded meshes. A formal discussion on visualization of organ motion due to breathing was not presented. We believe that the limited number of publications dealing with this subject is related to the relatively recent public availability of $3D + t$ datasets to the visualization community. Fortunately, projects like the POPI model from the Léon Bérard Cancer Center & CREATIS lab (Lyon, France) [152] provide not only the 4D CT datasets, but also the displacement fields with respect to a reference volume, which are




Figure A.1: Direct visualization of motion field with vectors. The vector colors encode the normalized motion magnitude. In image (a) a MPR slice of a volume is displayed together with the motion vectors. In image (b) the vectors are combined with the volume rendering of the dataset (from now on referred to as *vector* visualization). In image (c) only the vector spatially belonging to the liver are displayed.

computed by parametric and non-parametric deformable registration methods.

Visualizing organ motion on a cutting plane of the volumetric dataset can be adequately achieved with standard flow visualization techniques [123, 94]. Figure A.1(a) shows such a rendering. It is evident that rendering only a 2D cutting plane results in losing the anatomical context in the visualization. Furthermore, out-of-plane motion can not be visualized with this approach. For that reason, this work will only discuss 3D visualizations using volume rendering as the technique of choice for providing fast and anatomically preserving volumetric visualizations [25, 89]. The new challenge is how to incorporate the visualization of organ specific motion into volume rendering. In the next sections we will discuss the applicability of well-known visualization techniques for visualizing the effect of breathing motion and introduce a new technique for emphasizing the motion magnitude through the transfer function.

A.2.1 Direct Volume Rendering (DVR)

Ray-casting has been established as the standard for volumetric visualization, especially with the introduction of GPU accelerated ray-casting [89]. For a more comprehensive overview we refer to Engel et al. [25]. The 4D CT POPI dataset - used in this work - is essentially a collection of 3D CT volumes acquired over different breathing stages. Sequential volume rendering of these already allows a decent visualization of the motion. A frame from this visualization is shown in Fig. A.2(a) and A.3(a).

A.2.2 DVR and Vector Geometry

The availability of a displacement field allows the integration of direct flow visualization techniques. These dense displacement fields can be computed with deformable/non-rigid registration techniques. There the goal is to compute a dense deformation field that maps voxels/pixels from one dataset to the corresponding ones in another dataset acquired at a different time. The field of image registration has been studied for almost 30 years and numerous methods have emerged for deformable registration [37]. In this work we use the




Figure A.2: Volume rendering combined with mesh rendering for emphasizing organ specific motion. In image (a) a reference VR is provided (from now on referred to as *direct* visualization), (b) shows a wireframe rendering of the mesh (from now on referred to as *wireframe* visualization), in (c) the mesh is textured with a dense/noisy texture, mimicking cirrhotic liver.

pre-computed displacement fields available for each volume of the POPI dataset [152], which were computed with a free-form deformation approach [22] and a Demons-based approach [128]. Alternatively, deformable registration could be computed with publicly available software packages, including Elastix [81] and Drop [36].

In the field of flow visualization a vector field is often visualized by the use of arrows; encoding both direction and magnitude of the flow [123]. Combining this visualization with DVR allows the examination of the exact deformation for different anatomical structures while preserving the context. A realistic occlusion handling of the vectors does greatly enhance depth perception of the geometry within the volume rendering. Occlusion handling was implemented in our renderer, whereas, efficient techniques for this are further discussed in [25]. A frame from this visualization is shown in Fig. A.1(b). When using this direct vector visualization to display the motion of the complete volume with high detail, the final image is cluttered with too much information so that the user is not able to perceive it properly. Increasing the spacing between the vectors would lead to less elements and clutter, but the information between those elements is not rendered. This problem can be partially addressed by rendering only selective regions of interest like in Fig. A.1(c), where arrows are only rendered for the liver region, which can be determined by segmentation. Throughout this work the random walks image segmentation method was used [38] for semi-automatic segmentation of organs of interest. In general, the direct flow visualization approaches suffer from occlusion of the geometry and cluttering of the scene, thus, preventing organ specific motion perception.

A.2.3 DVR and 3D Geometry

The direct vector visualization is useful if the exact movement of a local element is of interest. However, for medical visualization we are also highly interested in how the internal anatomy is affected by the breathing motion and how organs move in relation to each other. Transfer functions in DVR allow an implicit segmentation of the anatomical regions of interest. However, these are often too fuzzy to recognize small deformations. The relative motion of organs to each other and detailed motion on the surface of or-

gans can be highlighted when using abstraction of these objects in form of 3D polygonal meshes. These can be generated prior to visualization from volume segmentations. If a displacement field is available then the segmentation is performed once and is propagated by the field over time. A visualization with a mesh of the liver is shown in Fig. A.2(b). In order to increase the perception of the mesh/organ distortion we use a noisy texturing as shown in Fig. A.2(c), which has a similar perceptual effects as LIC techniques.

A.2.4 Motion Transfer Function

The transfer function (TF) is playing the most important role for providing the desired visualization. A transfer function maps physical values from the volume to a more appealing value for visualization, using transparency to separate interesting anatomy and color to emphasize different organs and structures. The design of a good transfer function is a field of research by itself and often left to the user through interaction [25].

Here we introduce a new approach of encoding the motion in the transfer function of the volume rendering, that we denote motion transfer function (MotionTF). More specifically, the information from the deformation field is encoded in the TF. Formally, for volume rendering lets define the emissive color $c(t) := c(s(\mathbf{x}(t)))$ and absorption coefficient $k(t) := k(s(\mathbf{x}(t)))$ where $s(\mathbf{x}(t))$ is the scalar value in the volume at position $\mathbf{x}(t)$, t is the distance from the ray origin, and $c : X \rightarrow (R, G, B)$, $k : X \rightarrow A$ are the color and opacity transfer functions defined by the user. The continuous volume rendering integral that determines the total amount of radiant color/energy C is given as:

$$C = \int_0^\infty c(t) \cdot \exp^{-\int_0^\infty k(s)ds} dt \quad (\text{A.1})$$

In the discrete implementation the integral is usually substituted by a Riemann summation. For a more comprehensive overview on volume rendering theory and implementation we refer to [25]. In this work we integrate the information from the deformation field d into the color TF. More specially, we define:

$$c_m(t) = m(\|\nabla d(x(t))\|) \quad (\text{A.2})$$

where $m : X \rightarrow (R, G, B)$ is the transfer function mapping the deformation field magnitude to a color. Substituting the original color TF in Eq. A.1 with the new motion TF (c_m) results in the visualization shown in Fig. A.3(b). A complete sequence of this visualization method is further shown in Fig. A.4. The c_m function can be user defined to emphasize motion depending on its magnitude. The presented visualization method is especially well-suited for assessing the global deformation; a perceptual benefit that is also confirmed in our survey later on. Additionally, a blending of the original and motion TF can be used to preserve further the anatomical context provided from the standard color TF, as shown in figure A.3(c). This is achieved by substituting the original $c(t)$ with $c'(t) = (1 - \alpha)c(t) + \alpha c_m(t)$, were α is user-defined blending weight. The transfer functions should use distinct colors palettes, preferably with different saturations, in order to avoid the blending of similar colors from the color and motion TF.




Figure A.3: In (a) we see a reference volume rendering using a 1D transfer function. In (b) the motion magnitude is encoded directly on the volume with a transfer function (from now on referred to as *MotionTF* visualization). In (c) the reference volume rendering (a) is blended with the motion color volume rendering (b) (from now on referred to as *blended* visualization).

A.3 Survey

We performed a survey with 15 participants that are active researchers in the field of medical image computing, thus, have experience with both medical datasets and visualization. Evaluating the effectiveness of visualization methods remains difficult [64] and we chose to restrict our questions both in number and scope to reduce the bias in the survey [84]. Accordingly, all participants were given videos of five visualization methods on the POPI dataset, which they could view as much as they liked (all images in this work are taken from these videos). Following questions were asked: Which visualization method shows best the: ($Q1$) magnitude of the motion, ($Q2$) direction of the motion, ($Q3$) the interaction/relative motion between anatomies/organs, and ($Q4$) the detailed motion inside the organs. Consequently, the method of choice was recorded for each participant and for each question; summarized in Table A.1.

	Vector	Direct	Wireframe	MotionTF	Blend
$Q1$	0	0	1	14	0
$Q2$	5	2	6	2	0
$Q3$	1	0	12	2	0
$Q4$	2	0	1	7	5

Table A.1: Each cell contains the number of times participants identified a method as the best for a specific question ($Q1$ - $Q4$).

The survey answers for $Q1$ indicate that our method (MotionTF) is very well suited for visualizing the magnitude of the motion in the dataset. This was expected given that the method was designed to specifically encoded the motion magnitude. The answers to the questions regarding the direction ($Q2$) and the detail of the motion inside the organs ($Q4$) do not allow clear conclusions to be drawn. Furthermore, the fact that users had different preferences for these tasks shows the difficulties in defining visualization approaches with a clear and common perceptual effect for all users for medical 4D visualization. Last




Figure A.4: Images show a sequence from the rendering of the 4D CT dataset at different timesteps/breathing stages, using our motion transfer function.

but not least, the wireframe visualization, addressed in $Q3$, proves especially useful for showing organ related motion. The clear delineation with the wireframe geometry proved to be a very important aspect for this result.

A.4 Discussion

To provide interactive framerates we utilized modern GPU programming APIs including: OpenGL, GLSL, and CUDA. Improving the overall quality of the volume rendering could further improve the perception of motion in the datasets, see [25] for high-quality volume rendering techniques. Furthermore, virtual clipping planes can also assist in defining specific regions of interest in the visualization. The presented MotionTF is a specific case of a visualization for breathing motion. Alternative flow information could also be encoded in the function like the direction of the flow, as commonly used in 2D optical flow visualization [3].

The visualizations were tested on a Intel Core 2 Quad CPU with 4 cores running at 2.50Ghz, 4 GB of RAM and a NVIDIA GeForce GTX260 with 896 MB VRAM. Most of the resources are used by our very basic implementation of the volume renderer which is running with around 30 frames per second (FPS) and a resolution of 800x600. Compared to the raycasting, the geometry of the direct vector or mesh visualizations with around 100.000 vertices had no notable impact on the performance (around 1-2 FPS). The MotionTF on the other hand needs two additional texture lookups per sampling point (one for the vector-field and one for the second transfer function), which approximately reduces the performance of our test system by half depending on the sampling rate (see Fig. A.5).

The presented survey focused on the evaluation of basic perceptual questions. We chose videos instead of interactive use of the renderer to ensure the same experience for




Figure A.5: Performance of the MotionTF




Figure A.6: Examples of 4D ultrasound visualization.

all users. A clinical survey towards the evaluation of application specific visualizations with the participation of multiple physicians would be of high interest. Possible applications could involve the assessment of the change in organ motion due to illness. Specific examples include the change of elasticity of the lung due to asbestosis, fibrosis, and in general scar tissue. However, a prerequisite for such an evaluation is the availability of a sufficiently large evaluation dataset, which was not available for this work. The visualization techniques presented are generalizable to other modalities including MRI and ultrasound. However, separate user-studies are required to analyze the perceptual benefit. Especially ultrasound is challenging for visualization because of artifacts and speckle [53], and would require adjustments to transfer function design and blending.

We performed first experiments for visualization of 4D ultrasound acquired from different breathing stages. First results are presented in Fig. A.6. Furthermore, a sequence of the 4D visualization, using the introduced MotionTF and wireframe visualization, is shown in Fig. A.7. Again, semi-automatic segmentation [38] of the liver anatomy was obtained before visualization to generate the required mesh geometry. The acquisition of the dataset was performed with a so-called ultrasound wobbler probe; requiring image-based gating to ensure image consistency in acquired volumes [154]. The deformation field was computed with the method introduced in [166]. These initial results demonstrate that the MotionTF provides good perceptual information of motion magnitude in the dataset. Future research on improving transfer functions for ultrasound volume rendering, combined with the introduced visualization methods, could notably improve the overall perceptual benefit.




Figure A.7: Example of 4D ultrasound visualization sequence using MotionTF method.

A.5 Conclusion

In this work we addressed the specific challenges in the visualization of breathing motion. Furthermore, we presented a new, simple, and effective approach for visualizing the global magnitude of the motion directly in the volume rendering using the transfer function. Flow visualization offers a wide array of methods for visualizing vector fields; however, the perceptual requirements for visualizing breathing motion are different and require alternative visualization approaches. Our preliminary user survey indicates the benefit of specific visualizations for specific perceptual problems and also showed the challenge in defining visualizations with a consistent perceptual experience for multiple users. In our work we also pointed out the lack of research in this specific medical visualization field. We believe that the future increase of publicly available 4D datasets will motivate further research in this field and introduce alternative approaches for handling these new and challenging visualization aspects.

AUTHORED AND CO-AUTHORED PUBLICATIONS

Journal Publications

- Karamalis, A., Wein, W., Klein, T., and Navab, N. Ultrasound confidence maps using random walks. *Medical Image Analysis* 16, 6 (2012), 1101 - 1112
- Katouzian, A., Karamalis, A., Sheet, D., Konofagou, E., Baseri, B., Carlier, S.G., Eslami, A., König, A., Navab, N., and Laine, A. F. Iterative self-organizing atherosclerotic tissue labeling in intravascular ultrasound images and comparison with virtual histology. *IEEE Transactions on Biomedical Engineering* 59, 11 (2012), 3039-3049
- Eslami, A., Karamalis, A., Katouzian, A., and Navab, N. Segmentation by retrieval with guided random walks: Application to left ventricle segmentation in MRI. *Medical Image Analysis* 17, 2 (2012), 236-253
- Sheet, D., Karamalis, A., Eslami, A., Noël, P., Chatterjee, J., Ray, A. K., Laine, A. F., Carlier, S. G., Navab, N., and Katouzian, A. Joint learning of ultrasonic backscattering statistical physics and signal confidence primal for characterizing atherosclerotic plaques using intravascular ultrasound. *Medical Image Analysis* (to appear in print)
- Sheet, D., Karamalis, A., Eslami, A., Noël, P., Virmani, R., Nakano, M., Chatterjee, J., Ray, A. K., Laine, A. F., Carlier, S. G., Navab, N., and Katouzian, A. Hunting for necrosis in the shadows of intravascular ultrasound. *Computerized Medical Imaging and Graphics* (to appear in print)

Peer-reviewed Conference Publications

- Karamalis, A., Wein, W., and Navab, N. Fast ultrasound image simulation using the Westervelt equation. In *Medical Image Computing and Computer Aided Interventions (MICCAI)* (2010), vol. 6361, pp. 243-250

- Karamalis, A., Katouzian, A., Carlier, S. G., and Nassir, N. Confidence estimation in IVUS radio-frequency data with random walks. In 9th IEEE International Symposium on Biomedical Imaging: Nano to Macro (2012), pp. 1068-1071
- Karamalis, A., Wein, W., Kutter, O., and Navab, N. Fast hybrid freehand ultrasound volume reconstruction. In SPIE Medical Imaging (2009), vol. 7261, International Society for Optics and Photonics, pp. 726114-726118
- Klein, T., Hansson, M., Karamalis, A., and Navab, N. Registration of RF ultrasound data using hybrid local binary patterns. In 9th IEEE International Symposium on Biomedical Imaging: Nano to Macro (2012), pp. 1072-1075
- Katouzian, A., Karamalis, A., Laine, A. F., and Navab, N. A systematic approach toward reliable atherosclerotic plaque characterization in IVUS images. In *Bildverarbeitung für die Medizin* (2012), Springer, pp. 21-26
- Ahmadi, S., Baust, M., Karamalis, A., Plate, A., Boetzel, K., Klein, T., and Navab, N. Midbrain segmentation in transcranial 3D ultrasound for Parkinson diagnosis. In *Medical Image Computing and Computer Aided Interventions (MICCAI)* (2011), vol. 6893, pp. 362-369
- Kraft, S., Karamalis, A., Sheet, D., Drecoll, E., Rummeny, E. J., Navab, N., Noël, P. B., and Katouzian, A. Introducing nuclei scatterer patterns into histology-based intravascular ultrasound simulation framework. In SPIE Medical Imaging (2013), vol. 8675, International Society for Optics and Photonics, pp. 86750Y-86750Y-6
- Sheet, D., Karamalis, A., Kraft, S., Noël, P. B., Vag, T., Sadhu, A., Katouzian, A., Navab, N., Chatterjee, J., and Ray, A. K. Random forest learning of ultrasonic statistical physics and object spaces for lesion detection in 2D sonomammography. In SPIE Medical Imaging (2013), vol. 8675, International Society for Optics and Photonics, pp. 867515-867518
- Beitzel, J., Ahmadi, S., Karamalis, A., Wein, W., and Navab, N. Ultrasound bone detection using patient-specific CT prior. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2012), pp. 2664-2667
- Katouzian, A., Karamalis, A., König, C., Carlier, S. G., and Navab, N. Ambiguity in detection of necrosis in IVUS plaque characterization algorithms. In *European Heart Journal* (2012), vol. 33, Oxford University Press, pp. 192-192
- Kutter, O., Karamalis, A., Wein, W., and Navab, N. A GPU-based framework for simulation of medical ultrasound. In SPIE Medical Imaging (2009), vol. 7261, International Society for Optics and Photonics, pp. 726117-9.

Workshop Publications

- Katouzian, A., Karamalis, A., Lisauskas, J., Eslami, A., and Navab, N. IVUS-histology image registration. In *Biomedical Image Registration*, vol. 7359, Springer, 2012, pp. 141-149
- Müller, M., Karamalis, A., and Navab, N. Visualization of organ motion during breathing from 4D datasets. In *Augmented Environments for Computer-Assisted Interventions*, vol. 7815, Springer, 2013, pp. 34-44

ABSTRACTS OF CO-AUTHOR PUBLICATIONS

In this section abstracts of co-authored publications are provided, which were not discussed in the dissertation.

Segmentation By Retrieval With Guided Random Walks: Application To Left Ventricle Segmentation in MRI [26]

A. Eslami, A. Karamalis, A. Katouzian, N. Navab

In this paper, a new segmentation framework with prior knowledge is proposed and applied to the left ventricles in cardiac Cine MRI sequences. We introduce a new formulation of the random walks method, coined as guided random walks, in which prior knowledge is integrated seamlessly. In comparison with existing approaches that incorporate statistical shape models, our method does not extract any principal model of the shape or appearance of the left ventricle. Instead, segmentation is accompanied by retrieving the closest subject in the database that guides the segmentation the best. Using this techniques, rare cases can also effectively exploit prior knowledge from few samples in training set. These cases are usually disregarded in statistical shape models as they are outnumbered by frequent cases (effect of class population). In the worst-case scenario, if there is no matching case in the database to guide the segmentation, performance of the proposed method reaches to the conventional random walks, which is shown to be accurate if sufficient number of seeds is provided. There is a fast solution to the proposed guided random walks by using sparse linear matrix operations and the whole framework can be seamlessly implemented in a parallel architecture. The method has been validated on a comprehensive clinical dataset of 3D+t short axis MR images of 104 subjects from 5 categories (normal, dilated left ventricle, ventricular hypertrophy, recent myocardial infarction, and heart failure). The average segmentation errors were found to be 1.54mm for the endocardium and 1.48mm for the epicardium. The method was validated by measuring different algorithmic and physiologic indices and quantified with manual segmentation ground truths, provided by a cardiologist.

Midbrain Segmentation in Transcranial 3D Ultrasound for Parkinson Diagnosis [2]

A. Ahmadi, M. Baust, A. Karamalis, A. Plate, K. Bötzel, T. Klein, N. Navab

Ultrasound examination of the human brain through the temporal bone window, also called transcranial ultrasound (TC-US), is a completely non-invasive and cost-efficient technique, which has established itself for differential diagnosis of Parkinson's Disease (PD) in the past decade. The method requires spatial analysis of ultrasound hyperechogenicities produced by pathological changes within the Substantia Nigra (SN), which belongs to the basal ganglia within the midbrain. Related work on computer aided PD diagnosis shows the urgent need for an accurate and robust segmentation of the midbrain from 3D TC-US, which is an extremely difficult task due to poor image quality of TC-US. In contrast to 2D segmentations within earlier approaches, we develop the first method for semi-automatic midbrain segmentation from 3D TC-US and demonstrate its potential benefit on a database of 11 diagnosed Parkinson patients and 11 healthy controls.

Registration of RF Ultrasound Data using Hybrid Local Binary Patterns [82]

T. Klein, M. Hansson, A. Karamalis, N. Navab

Registration of ultrasound images is often complicated due to inherent noise. Robust similarity metrics and optimization procedures are required to facilitate medical applicability. In this paper a novel hybrid procedure, incorporating global statistics and local textural features, is proposed for the registration of envelope detected radio frequency ultrasound data. On the global scale this is achieved by Hellinger distance between distribution in images, and on the local scale by a statistics-based extension of Fuzzy Local Binary Patterns (FLBP). The proposed procedure is shown to outperform standard measures such as SSD and NCC, as well as Hellinger distance and histogram matching of standard FLBPs, in rigid registration experiments of envelope detected radio frequency data samples of the human neck.

Random Forest Learning of Ultrasonic Statistical Physics and Object Spaces for Lesion Detection in 2D Sonomammography [134]

D. Sheet, A. Karamalis, S. Kraft, P. B. Noël, T. Vag, A. Sadhu, A. Katouzian, N. Navab, J. Chatterjee, A. K. Ray

Breast cancer is the most common form of cancer in women. Clinicians favor 2D ultrasonography for breast tissue abnormality screening due to high sensitivity and specificity compared to competing technologies. However, inter- and intra-observer variability in visual assessment and reporting of lesions often handicaps its performance. In this work we present a completely automatic system for detection and segmentation of breast lesions

in 2D ultrasound images. We employ random forests for learning of tissue specific primal to discriminate breast lesions from surrounding normal tissues. This enables it to detect lesions of multiple shapes and sizes, as well as discriminate between hypo-echoic lesion from associated posterior acoustic shadowing.

Introducing Nuclei Scatterer Patterns into Histology-based Intravascular Ultrasound Simulation Framework [86]

S. Kraft, A. Karamalis, D. Sheet, E. Drecoll, E. Rummeny, N. Navab, P. B. Noël, A. Katouzian

We introduce a framework to simulate intravascular ultrasound (IVUS) from histological sections. These sections were previously acquired along with real IVUS radiofrequency signals using single-element 40MHz transducer. After labeling and registering the section to the corresponding IVUS image, a virtual phantom was created, incorporating nuclei scatterer patterns. A finite differences simulation of the acoustic signal was performed, resulting in backscattered radiofrequency signals. These were used to process a B-mode image, which in turn was compared to the real IVUS image of the same section. A high image quality with a very promising correlation to the original IVUS images was achieved.

Fast Hybrid Freehand Ultrasound Volume Reconstruction [68]

A. Karamalis, W. Wein, O. Kutter, N. Navab

The volumetric reconstruction of a freehand ultrasound sweep, also called compounding, introduces additional diagnostic value to the ultrasound acquisition by allowing 3D visualization and fast generation of arbitrary MPR(Multi-Planar-Reformatting) slices. Furthermore reconstructing a sweep adds to the general availability of the ultrasound data since volumes are more common to a variety of clinical applications/systems like PACS. Generally there are two reconstruction approaches, namely forward and backward with their respective advantages and disadvantages. In this paper we present a hybrid reconstruction method partially implemented on the GPU that combines the forward and backward approaches to efficiently reconstruct a continuous freehand ultrasound sweep, while ensuring at the same time a high reconstruction quality. The main goal of this work was to significantly decrease the waiting time from sweep acquisition to volume reconstruction in order to make an ultrasound examination more comfortable for both the patient and the sonographer. Testing our algorithm demonstrated a significant performance gain by an average factor of 197 for simple interpolation and 84 for advanced interpolation schemes, reconstructing a 256^3 volume in 0.35 seconds and 0.82 seconds respectively.

A GPU-Based Framework for Simulation of Medical Ultrasound [92]

O. Kutter, A. Karamalis, W. Wein, N. Navab

Simulation of ultrasound (US) images from volumetric medical image data has been shown to be an important tool in medical image analysis. However, there is a trade off between the accuracy of the simulation and its real-time performance. In this paper, we present a framework for acceleration of ultrasound simulation on the graphics processing unit (GPU) of commodity computer hardware. Our framework can accommodate ultrasound modeling with varying degrees of complexity. To demonstrate the flexibility of our proposed method, we have implemented several models of acoustic propagation through 3D volumes. We conducted multiple experiments to evaluate the performance of our method for its application in multi-modal image registration and training. The results demonstrate the high performance of the GPU accelerated simulation outperforming CPU implementations by up to two orders of magnitude and encourage the investigation of even more realistic acoustic models.

LIST OF ABBREVIATIONS

- CAOS:** Computer-Aided Orthopedic Surgery
- CPU:** Central Processing Unit
- CRS:** Controlled Random Search
- CT:** Computer Tomography
- EM:** Electromagnetic
- FDTD:** Finite-Difference Time-Domain
- FG:** Field Generator
- FRE:** Fiducial Registration Error
- GPGPU:** General-Purpose-Graphics-Processing-Unit
- GPU:** Graphics Processing Unit
- HC:** High Confidence
- IVUS:** Intravascular Ultrasound
- KZK:** Khokhlov-Zabolotskaya-Kuznetsov
- LC:** Low Confidence
- LIC:** Line-Integral-Convolution
- MRI:** Magnetic Resonance Imaging
- PDE:** Partial Differential Equation
- PSF:** Point Spread Function
- RF:** Radio-Frequency

RMS: Root-Mean-Square Error

SAM: Statistical-Appearance-Model

SNR: Signal-to-Noise Ratio

SSD: Sum-of-Squared-Differences

SSM: Statistical-Shape-Model

TF: Transfer Function

TRE: Target Registration Error

US: Ultrasound

LIST OF FIGURES

2.1	Illustration of ultrasound imaging working principle.	12
2.2	Illustration of compression and rarefaction during pressure propagation. . .	14
2.3	Image (a) shows successive transmissions and reflections. Image (b) illustrates Snell's law given an incoming sound beam.	15
2.4	Blue and red wave are constructively interfering, whereas, blue and green are destructively interfering.	16
2.5	Illustration of ultrasound imaging system.	17
2.6	Illustration of acoustic pulses. Carrier signal is modulated to form transmit signal. In reception mode the received signal is demodulated by taking it's envelope.	18
2.7	Illustration of incident (a) and echo waves (b). Parabolic placement of sound sources to generate focused wavefront (c).	18
2.8	Image (a) illustration of triggering scheme, image (b) illustration of sound beam.	19
2.9	Different ultrasound transducers for the Siemens ACUSON S3000 ultrasound system (Siemens press pictures).	21
2.10	Different ultrasound images acquired with different systems from different anatomies.	22
2.11	(a) Illustration of beam-width and effect on spatial resolution of reflectors, (b) illustration of slice-width.	24
2.12	Example of Doppler imaging of carotid artery (Siemens press picture). . . .	25
2.13	Example of conventional B-mode (left) and spatial compounded image (right). Tumor boundaries and central microcalcifications are better delineated (Reprinted from [30], Copyright(2013), with permission from Elsevier).	26
2.14	Example of baseline B-mode (left) and B-mode after injection of contrast agent (right) (Reprinted from [32], Copyright(2013), with permission from Elsevier).	26
2.15	Illustration of transducer frequency response (see also [17]).	27
2.16	Example of conventional B-mode (left) and harmonic imaging (right) (Reprinted from [148], Copyright(2013), with permission from Elsevier).	28
2.17	Illustration of lesion treatment in liver with highly focused ultrasound. . . .	29

LIST OF FIGURES

2.18	Example of conventional B-mode (left) of carcinoma and corresponding elastogram imaging (right) (Reprinted from [13], Copyright(2013), with permission from Elsevier).	30
2.19	Image from different modalities (a) X-ray of chest (Siemens press picture), (b) CT of abdomen, (c) MRI of brain, (d), PET of brain, (e) histology of vessel.	32
3.1	Illustration of CUDA programming model and Fermi architecture.	40
3.2	Image (a) schematically demonstrates ultrasound beam focusing by triggering a group of elements with different time delays. Image (b) shows the maximum wave amplitudes of a simulation run using a focus scheme for low depths in a medium with uniform speed of sound and image (c) its corresponding 3D visualization.	42
3.3	Image (a) shows the synthetic fetus dataset and image (c) the synthetic phantom dataset. The intensity values correspond to speed of sound values in the range of 1500-1550 [m/s]. The center of the focal zones are marked with a small triangle on the left side. Image (b) and (d) show the simulated ultrasound images. Baby phantom from [61].	44
3.4	Histology image (a) is non-rigidly registered to corresponding IVUS image (b). Registration outcome shown in (c) and (d).	45
3.5	Image (a) shows labeled IVUS with regions blue: calcified = 1820m/s, yellow: fibrotic = 1600m/s, green: fibro-lipidic = 1570m/s, pink: lipid= 1520m/s, (b) deformed and labeled histology, (c) phantom with random noise and speed of sound assignments.	46
3.6	Example of simulation results. Image (a) is the real IVUS in polar coordinates (before scan-conversion) and (b) and (c) simulated images using far and near focal regions respectively.	47
4.1	Example of graph with corresponding incidence matrix.	55
4.2	Basic framework for solving balancing equations in linear algebra (see [140]), here in the context of circuit problems.	58
4.3	Illustration of graph setup for confidence maps using an 8-connected lattice. The first row resembles the virtual transducer elements with the probabilities set to unity (shaded dark gray). The last row resembles the 'no signal' region , i.e., the necessary boundary condition, with the probabilities set to zero (shaded bright gray). For all the nodes in-between we compute the probability that a random walk starting from each node would first reach one of the virtual transducer elements.	59
4.4	The first and third row show B-mode images of a human heart, femur, neck and liver. The second and fourth row show the corresponding confidence maps. Bright values denote high confidence and dark values low confidence. All images were generated with the same α, β parameters, namely $\alpha = 2, \beta = 90$. The γ parameter was optimized for the first image of each sequence/anatomy using Eq. 4.15 and was then kept constant for the remaining images of the sequence.	61

4.5	The effect of the α parameter. Image (b) and (c) are the confidence maps for image (a) by using different α values and fixing $\beta = 90$ and $\gamma = 0.05$. . .	63
4.6	The effect of the β parameter. Image (b) and (c) are the confidence maps for image (a) by using different β values and fixing $\alpha = 2$ and $\gamma = 0.05$. . .	64
4.7	The effect of the γ parameter. Image (b) and (c) are the confidence maps for image (a) by using different γ values and fixing $\alpha = 2$ and $\beta = 90$	64
4.8	Boxes in the images indicate regions of high and low confidence; before and after highly attenuating interfaces. Images in column (a) are the original ultrasound of femur and liver, in (b) the corresponding confidence maps, and in (c) the scanline integrated images.	65
4.9	Regions of relatively high confidence (HC) and a-priori known low confidence (LC) are defined for 25 human femur and 25 human liver images. Subsequently, the mean values are computed for regions sampled from our confidence maps, the original ultrasound images, and the scanline integrated images. Each box-plot shows the means of the HC and LC regions for every test case.	65
4.10	Images (a)-(c) show the shadow regions, area below dotted green line, for envelope detect RF images (here log-compressed for display) acquired from a human humerus. The corresponding confidence maps are shown in images (d)-(f).	68
4.11	Shadow detection results are presented by comparing the manual and automatic detection of shadow regions with the Dice coefficient. Two experts (U1 and U2) segmented for each frequency setting the same 140 images. . .	69
4.12	Image (a) and (b) show a transverse slice of a human femur from a reconstructed volume using the R_G and our newly defined R_C weighting function respectively. Image (c) shows the joint histogram of two volumes, one reconstructed with R_G and the other with R_C	70
4.13	Image (a) and (c) show a sagittal slice of a human femur from a reconstructed volume using the R_G and our newly defined R_C weighting function respectively. Image (b) and (d) show an enlarged version of the regions marked with the white box.	71
4.14	CT-Ultrasound image pair and their associated transmission and confidence map.	72
4.15	Input for mono-modal registration example with confidence maps.	74
4.16	Similarity plots of SSD values for different poses: (a) US-US registration (translation), (b) map-map registration (translation), and (c) US-US and map-map registration (rotation).	75
4.17	Example from our evaluation dataset showing: (a) IVUS B-Mode, (b) corresponding confidence map, and (c) the corresponding histology.	75
4.18	Illustration of the deformable registration pipeline for mapping labels assigned in histology to confidence maps. Image details: (a) labeled confidence map, (b) IVUS, (c) segmentation mask (dotted-box is the bounding-box for size-normalization), (d) deformation field, (e) labeled histology, (f) histology, and (g) segmentation mask.	76

LIST OF FIGURES

4.19	An IVUS grayscale image with arc of calcified plaque at 3-5 o' clock (a) and corresponding confidence map (b), registered histology (c), PH image before (d) and after (e) weighting with confidence values.	77
4.20	Sampling of confidence maps at HC (blue) and LC (red) regions defined in co-registered histology. Each bar shows the mean and the (+/-) standard deviation of the confidence sampled in the labeled regions.	78
5.1	Some of the images and intermediate outputs for bone detection.	85
5.2	System overview	87
5.3	Quantitative evaluation of bone detection result for femur of five cadaver (images a-b) and five volunteer (images c-d) datasets. The datasets of the cadaver were recorded from the same specimen. The datasets for the volunteers were recorded from different persons with varying physiology. For each dataset 100 images were manually segmented and compared to automatic detection results.	88
5.4	Quantitative evaluation of bone detection result for tibia/fibula of five cadaver (images a-b) and five volunteer (images c-d) datasets. The datasets of the cadaver were recorded from the same specimen. The datasets for the volunteers were recorded from different persons with varying physiology. For each dataset 100 images were manually segmented and compared to automatic detection results.	89
5.5	Example of US-CT registration result.	91
5.6	Examples of successful bone detection using femur dataset from a cadaver	93
5.7	Examples of successful bone detection using femur dataset from five volunteers	94
5.8	Examples of successful bone detection using tibia/fibula dataset from a cadaver	95
5.9	Examples of successful bone detection using tibia/fibula dataset from five volunteers	96
5.10	Renderings of CT data set acquired in cadaver study (a) femur, (b) tibia/fibula and (c) tibia/fibula after introduction of fracture and bone markers. Visualization with Exposure renderer [88].	97
A.1	Direct visualization of motion field with vectors. The vector colors encode the normalized motion magnitude. In image (a) a MPR slice of a volume is displayed together with the motion vectors. In image (b) the vectors are combined with the volume rendering of the dataset (from now on referred to as <i>vector</i> visualization). In image (c) only the vector spatially belonging to the liver are displayed.	107
A.2	Volume rendering combined with mesh rendering for emphasizing organ specific motion. In image (a) a reference VR is provided (from now on referred to as <i>direct</i> visualization), (b) shows a wireframe rendering of the mesh (from now on referred to as <i>wireframe</i> visualization), in (c) the mesh is textured with a dense/noisy texture, mimicking cirrhotic liver.	108

A.3	In (a) we see a reference volume rendering using a 1D transfer function. In (b) the motion magnitude is encoded directly on the volume with a transfer function (from now on referred to as <i>MotionTF</i> visualization). In (c) the reference volume rendering (a) is blended with the motion color volume rendering (c) (from now on referred to as <i>blended</i> visualization).	110
A.4	Images show a sequence from the rendering of the 4D CT dataset at different timesteps/breathing stages, using our motion transfer function. . . .	111
A.5	Performance of the MotionTF	112
A.6	Examples of 4D ultrasound visualization.	112
A.7	Example of 4D ultrasound visualization sequence using MotionTF method.	113

REFERENCES

- [1] ADAMS, S., PAYNE, J., AND BOPPANA, R. Finite difference time domain (FDTD) simulations using graphics processors. In *IEEE DoD High Performance Computing Modernization Program Users Group Conference* (2007), pp. 334–338.
- [2] AHMADI, S., BAUST, M., KARAMALIS, A., PLATE, A., BOETZEL, K., KLEIN, T., AND NAVAB, N. Midbrain segmentation in transcranial 3D ultrasound for Parkinson diagnosis. In *Medical Image Computing and Computer Aided Interventions (MICCAI)* (2011), vol. 6893, pp. 362–369.
- [3] BAKER, S., SCHARSTEIN, D., LEWIS, J., ROTH, S., BLACK, M. J., AND SZELISKI, R. A database and evaluation methodology for optical flow. *International Journal of Computer Vision* 92, 1 (2011), 1–31.
- [4] BARRATT, D., PENNEY, G., CHAN, C., SŁOMCZYKOWSKI, M., CARTER, T., EDWARDS, P., AND HAWKES, D. Self-calibrating ultrasound-to-CT bone registration. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2005), vol. 3749, pp. 605–612.
- [5] BARRATT, D., PENNEY, G., CHAN, C., SŁOMCZYKOWSKI, M., CARTER, T., EDWARDS, P., AND HAWKES, D. Self-calibrating 3D-ultrasound-based bone registration for minimally invasive orthopedic surgery. *IEEE Transactions on Medical Imaging* 25, 3 (2006), 312–323.
- [6] BARRATT, D. C., CHAN, C. S. K., EDWARDS, P. J., PENNEY, G. P., SŁOMCZYKOWSKI, M., CARTER, T. J., AND HAWKES, D. J. Instantiation and registration of statistical shape models of the femur and pelvis using 3D ultrasound imaging. *Medical Image Analysis* 12, 3 (2008), 358–374.
- [7] BEEK, M., ABOLMAESUMI, P., LUENAM, S., ELLIS, R., SELLENS, R., AND PICHORA, D. Validation of a new surgical procedure for percutaneous scaphoid fixation using intra-operative ultrasound. *Medical Image Analysis* 12, 2 (2008), 152–162.
- [8] BEHAR, V., AND NIKOLOV, M. Statistical analysis of image quality in multi-angle compound imaging. In *IEEE International Symposium on Modern Computing* (2006), pp. 197–201.

REFERENCES

- [9] BEITZEL, J., AHMADI, S., KARAMALIS, A., WEIN, W., AND NAVAB, N. Ultrasound bone detection using patient-specific CT prior. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2012), pp. 2664–2667.
- [10] BELL, N., AND GARLAND, M. Implementing sparse matrix-vector multiplication on throughput-oriented processors. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis* (2009), pp. 18:1–18:11.
- [11] BERENGER, J. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics* 114, 2 (1994), 185–200.
- [12] BESL, P., AND MCKAY, N. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 2 (1992), 239–256.
- [13] BHATIA, K. S., CHO, C. C., YUEN, Y.-H., RASALKAR, D. D., KING, A. D., AND AHUJA, A. T. Real-time qualitative ultrasound elastography of cervical lymph nodes in routine clinical practice: Interobserver agreement and correlation with malignancy. *Ultrasound in Medicine & Biology* 36, 12 (2010), 1990 – 1997.
- [14] BROUNSTEIN, A., HACIHALILOGLU, I., GUY, P., HODGSON, A., AND ABUGHAR-BIEH, R. Towards real-time 3D US to CT bone image registration using phase and curvature feature based GMM matching. In *Medical Image Computing and Computer-Assisted Intervention, MICCAI* (2011), vol. 6891, pp. 235–242.
- [15] CARRAT, L., TONETTI, J., MERLOZ, P., AND TROCENZA, J. Percutaneous computer assisted iliosacral screwing: Clinical validation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2000), vol. 1935, pp. 97–140.
- [16] CHEN, L., TREECE, G., LINDOP, J., GEE, A., AND PRAGER, R. A quality-guided displacement tracking algorithm for ultrasonic elasticity imaging. *Medical Image Analysis* 13, 2 (2009), 286–296.
- [17] COBBOLD, R. S. C. *Foundations of Biomedical Ultrasound*. Oxford University Press, 2006.
- [18] COHEN, G., AND JOLY, P. Construction and analysis of fourth-order finite difference schemes for the acoustic wave equation in nonhomogeneous media. *SIAM Journal on Numerical Analysis* 33, 4 (1996), 1266–1302.
- [19] CONNOR, C., AND HYNYNEN, K. Patterns of thermal deposition in the skull during transcranial focused ultrasound surgery. *IEEE Transactions on Biomedical Engineering* 51, 10 (2004), 1693–1706.
- [20] CORREAS, J., BRIDAL, L., LESAVRE, A., MEJEAN, A., CLAUDON, M., AND HELENON, O. Ultrasound contrast agents: properties, principles of action, tolerance, and artifacts. *European Radiology* 11, 8 (2001), 1316–1328.

-
- [21] DAANEN, V., TONETTI, J., AND TROCCAZ, J. A fully automated method for the delineation of osseous interface in ultrasound images. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2004), vol. 3216, pp. 549–557.
- [22] DELHAY, B., CLARYSSE, P., PERA, C., AND MAGNIN, I. A spatio-temporal deformation model for dense motion estimation in periodic cardiac image sequences. In *Statistical Atlases to Personalized Models: Understanding Complex Diseases in Populations and Individuals, Satellite Workshop of MICCAI* (2006), pp. 87–90.
- [23] DILLENSEGER, J., LAGUITTON, S., AND DELABROUSSE, É. Fast simulation of ultrasound images from a CT volume. *Computers in Biology and Medicine* 39, 2 (2009), 180–186.
- [24] DUSSIK, K. Über die möglichkeit, hochfrequente mechanische schwingungen als diagnostisches hilfsmittel zu verwerten. *Zeitschrift für die gesamte Neurologie und Psychiatrie* 174, 1 (1942), 153–168.
- [25] ENGEL, K., HADWIGER, M., KNISS, J., AND REZK-SALAMA, C. *Real-Time Volume Graphics*. AK Peters Ltd, 2006.
- [26] ESLAMI, A., KARAMALIS, A., KATOUIZIAN, A., AND NAVAB, N. Segmentation by retrieval with guided random walks: Application to left ventricle segmentation in MRI. *Medical Image Analysis* 17, 2 (2012), 236–253.
- [27] FARR, R., AND ALLISY-ROBERTS, P. *Physics for Medical Imaging*. Philadelphia: WB Saunder Co., 1997.
- [28] FEYNMAN, R., LEIGHTON, R., SANDS, M., AND TREIMAN, S. The Feynman lectures on physics. *Physics Today* 17 (1963), 45.
- [29] FOROUGHI, P., BOCTOR, E., SWARTZ, M., TAYLOR, R., AND FICHTINGER, G. Ultrasound bone segmentation using dynamic programming. In *IEEE Ultrasonics Symposium* (2007), pp. 2523–2526.
- [30] FORSBERG, F. Ultrasonic biomedical technology; marketing versus clinical reality. *Ultrasonics* 42, 1 (2004), 17–27.
- [31] FRIJLINK, M., KAUPANG, H., VARSLOT, T., AND MASOY, S. Abersim: a simulation program for 3D nonlinear acoustic wave propagation for arbitrary pulses and arbitrary transducer geometries. In *IEEE Ultrasonics Symposium* (2008), pp. 1282–1285.
- [32] FRINKING, P., BOUAKAZ, A., KIRKHORN, J., TEN CATE, F., AND DE JONG, N. Ultrasound contrast imaging: current and new potential methods. *Ultrasound in Medicine & Biology* 26, 6 (2000), 965–975.
- [33] GAO, H., CHOI, H., CLAUS, P., BOONEN, S., JAECQUES, S., VAN LENTHE, G., VAN DER PERRE, G., LAURIKS, W., AND D’HOOGHE, J. A fast convolution-based

REFERENCES

- methodology to simulate 2-D/3-D cardiac ultrasound images. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 56, 2 (2009), 404–409.
- [34] GIANFELICE, D., KHIAT, A., AMARA, M., BELBLIDIA, A., AND BOULANGER, Y. MR imaging-guided focused ultrasound surgery of breast cancer: Correlation of dynamic contrast-enhanced MRI with histopathologic findings. *Breast Cancer Research and Treatment* 82 (2003), 93–101.
- [35] GILL, S., ABOLMAESUMI, P., FICHTINGER, G., BOISVERT, J., PICHORA, D., BORSHNECK, D., AND MOUSAVI, P. Biomechanically constrained groupwise ultrasound to CT registration of the lumbar spine. *Medical Image Analysis* 16, 3 (2012), 662–674.
- [36] GLOCKER, B., KOMODAKIS, N., TZIRITAS, G., NAVAB, N., AND PARAGIOS, N. Dense image registration through MRFs and efficient linear programming. *Medical Image Analysis* 12, 6 (2008), 731–741.
- [37] GLOCKER, B., SOTIRAS, A., KOMODAKIS, N., AND PARAGIOS, N. Deformable medical image registration: Setting the state of the art with discrete methods. *Annual Review of Biomedical Engineering* 13 (2011), 219–244.
- [38] GRADY, L. Random walks for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28, 11 (2006), 1768–1783.
- [39] GRADY, L. A lattice-preserving multigrid method for solving the inhomogeneous poisson equations used in image analysis. In *European Conference on Computer Vision (ECCV)* (2008), vol. 5303, pp. 252–264.
- [40] GRADY, L., AND JOLLY, M. Weights and topology: A study of the effects of graph construction on 3D image segmentation. In *Medical Image Computing and Computer Aided Interventions (MICCAI)* (2008), vol. 5241, pp. 153–161.
- [41] GRADY, L., SCHIWETZ, T., AHARON, S., AND WESTERMANN, R. Random walks for interactive alpha-matting. In *Proceedings of VIIP* (2005), pp. 423–429.
- [42] GRADY, L., SCHIWETZ, T., AHARON, S., AND WESTERMANN, R. Random walks for interactive organ segmentation in two and three dimensions: Implementation and validation. In *Medical Image Computing and Computer-assisted Intervention (MICCAI)* (2005), vol. 3750, pp. 773–780.
- [43] GRAU, V., BECHER, H., AND NOBLE, J. Phase-based registration of multi-view real-time three-dimensional echocardiographic sequences. In *Medical Image Computing and Computer-assisted Intervention (MICCAI)* (2006), vol. 4190, pp. 612–619.
- [44] GRAU, V., AND NOBLE, J. Adaptive multiscale ultrasound compounding using phase information. In *Medical Image Computing and Computer-assisted Intervention (MICCAI)* (2005), vol. 3749, pp. 589–596.

-
- [45] HACIHALILOGLU, I., ABUGHARBIEH, R., HODGSON, A., AND ROHLING, R. Bone segmentation and fracture detection in ultrasound using 3D local phase features. In *Medical Image Computing and Computer Aided Interventions (MICCAI)* (2008), vol. 5241, pp. 287–295.
- [46] HACIHALILOGLU, I., BROUNSTEIN, A., GUY, P., HODGSON, A., AND ABUGHARBIEH, R. 3D ultrasound-CT registration in orthopaedic trauma using GMM registration with optimized particle simulation-based data reduction. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2012), vol. 7511, pp. 82–89.
- [47] HALLAJ, I., AND CLEVELAND, R. FDTD simulation of finite-amplitude pressure and temperature fields for biomedical ultrasound. *The Journal of the Acoustical Society of America* 105 (1999), 7–12.
- [48] HANDELS, H., WERNER, R., SCHMIDT, R., FRENZEL, T., LU, W., LOW, D., AND EHRHARDT, J. 4D medical image computing and visualization of lung tumor mobility in spatio-temporal CT image data. *International Journal of Medical Informatics* 76 (2007), 433–439.
- [49] HARRIS, M. Fast fluid dynamics simulation on the GPU. *GPU Gems 1* (2004), 637–665.
- [50] HEDRICK, W., HYKES, D., AND STARCHMAN, D. *Ultrasound Physics and Instrumentation*. Mosby St. Louis, MI, 2005.
- [51] HELLIER, P., COUPÉ, P., MORANDI, X., AND COLLINS, D. An automatic geometrical and statistical method to detect acoustic shadows in intraoperative ultrasound brain images. *Medical Image Analysis* 14, 2 (2010), 195–204.
- [52] HERGUM, T., LANGELAND, S., REMME, E., AND TORP, H. Fast ultrasound imaging simulation in K-space. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 56, 6 (2009), 1159–1167.
- [53] HONIGMANN, D., RUISZ, J., AND HAIDER, C. Adaptive design of a global opacity transfer function for direct volume rendering of ultrasound data. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)* (2003), pp. 489–496.
- [54] HOSKINS, P. R., MARTIN, K., AND THRUSH, A. *Diagnostic Ultrasound: Physics and Equipment*. Cambridge Medicine. Cambridge University Press, 2010.
- [55] HUBER, P., JENNE, J., RASTERT, R., SIMIANTONAKIS, I., SINN, H., STRITTMATTER, H., VON FOURNIER, D., WANNENMACHER, M., AND DEBUS, J. A new noninvasive approach in breast cancer therapy using magnetic resonance imaging-guided focused ultrasound surgery. *Cancer research* 61, 23 (2001), 8441–8447.

REFERENCES

- [56] HUIJSSEN, J., BOUAKAZ, A., VERWEIJ, M., AND DE JONG, N. Simulations of the nonlinear acoustic pressure field without using the parabolic approximation. In *IEEE Symposium on Ultrasonics* (2003), vol. 2, pp. 1851–1854.
- [57] JENSEN, J. Field: A program for simulating ultrasound systems. In *10th Nordicbaltic Conference on Biomedical Imaging, Vol. 4, Supplement 1, Part 1:351–353* (1996), pp. 351–353.
- [58] JENSEN, J. Simulating arbitrary-geometry ultrasound transducers using triangles. In *IEEE Ultrasonics Symposium* (1996), vol. 2, pp. 885–888.
- [59] JENSEN, J. A new calculation procedure for spatial impulse responses in ultrasound. *The Journal of the Acoustical Society of America* 105, 6 (1999), 3266–3274.
- [60] JENSEN, J. Simulation of advanced ultrasound systems using Field II. In *IEEE International Symposium on Biomedical Imaging: Nano to Macro* (2004), vol. 1, pp. 636–639.
- [61] JENSEN, J., AND MUNK, P. Computer phantoms for simulating ultrasound B-mode and CFM images. *Acoustical Imaging* 23 (1997), 75–80.
- [62] JENSEN, J., AND NIKOLOV, I. Fast simulation of ultrasound images. In *IEEE Ultrasonics Symposium* (2000), vol. 2, pp. 1721–1724.
- [63] JESORSKY, O., KIRCHBERG, K., AND FRISCHHOLZ, R. Robust face detection using the Hausdorff distance. In *Proc. Third International Conference on Audio- and Video-based Biometric Person Authentication* (2001), vol. 2091, pp. 90–95.
- [64] JOHNSON, C. Top scientific visualization research problems. *IEEE Computer Graphics and Applications* 24, 4 (2004), 13–17.
- [65] KAELO, P., AND ALI, M. M. Some variants of the controlled random search algorithm for global optimization. *J. Optim. Theory Appl.* 130, 2 (2006), 253–264.
- [66] KARAMALIS, A., KATOUZIAN, A., CARLIER, S. G., AND NASSIR, N. Confidence estimation in IVUS radio-frequency data with random walks. In *9th IEEE International Symposium on Biomedical Imaging: Nano to Macro* (2012), pp. 1068–1071.
- [67] KARAMALIS, A., WEIN, W., KLEIN, T., AND NAVAB, N. Ultrasound confidence maps using random walks. *Medical Image Analysis* 16, 6 (2012), 1101 – 1112.
- [68] KARAMALIS, A., WEIN, W., KUTTER, O., AND NAVAB, N. Fast hybrid freehand ultrasound volume reconstruction. In *SPIE Medical Imaging* (2009), vol. 7261, International Society for Optics and Photonics, pp. 726114–726118.
- [69] KARAMALIS, A., WEIN, W., AND NAVAB, N. Fast ultrasound image simulation using the Westervelt equation. In *Medical Image Computing and Computer Aided Interventions (MICCAI)* (2010), vol. 6361, pp. 243–250.

-
- [70] KATOUIZIAN, A., BASERI, B., KONOFAGOU, E., AND LAINE, A. Texture-driven coronary artery plaque characterization using wavelet packet signatures. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro* (2008), pp. 197–200.
- [71] KATOUIZIAN, A., KARAMALIS, A., KOENIG, C., CARLIER, S. G., AND NAVAB, N. Ambiguity in detection of necrosis in IVUS plaque characterization algorithms. In *European Heart Journal* (2012), vol. 33, Oxford University Press, pp. 192–192.
- [72] KATOUIZIAN, A., KARAMALIS, A., LAINE, A. F., AND NAVAB, N. A systematic approach toward reliable atherosclerotic plaque characterization in IVUS images. In *Bildverarbeitung für die Medizin* (2012), Springer, pp. 21–26.
- [73] KATOUIZIAN, A., KARAMALIS, A., LISAIKAS, J., ESLAMI, A., AND NAVAB, N. IVUS-histology image registration. In *Biomedical Image Registration*, vol. 7359. Springer, 2012, pp. 141–149.
- [74] KATOUIZIAN, A., KARAMALIS, A., SHEET, D., KONOFAGOU, E., BASERI, B., CARLIER, S. G., ESLAMI, A., KÖNIG, A., NAVAB, N., AND LAINE, A. F. Iterative self-organizing atherosclerotic tissue labeling in intravascular ultrasound images and comparison with virtual histology. *IEEE Transactions on Biomedical Engineering* 59, 11 (2012), 3039–3049.
- [75] KATOUIZIAN, A., AND LAINE, A. F. Methods in atherosclerotic plaque characterization using intravascular ultrasound (IVUS) images and backscattered signals. In *Atherosclerosis Disease Management*. Springer, 2011, pp. 121–152.
- [76] KATOUIZIAN, A., SATHYANARAYANA, S., BASERI, B., KONOFAGOU, E., AND CARLIER, S. Challenges in atherosclerotic plaque characterization with intravascular ultrasound (IVUS): From data collection to classification. *IEEE Transactions on Information Technology in Biomedicine* 12, 3 (2008), 315–327.
- [77] KAWASAKI, M., TAKATSU, H., NODA, T., SANO, K., ITO, Y., HAYAKAWA, K., TSUCHIYA, K., ARAI, M., NISHIGAKI, K., TAKEMURA, G., MINATOGUCHI, S., FUJIWARA, T., AND FUJIWARA, H. In vivo quantitative tissue characterization of human coronary arterial plaques by use of integrated backscatter intravascular ultrasound and comparison with angioscopic findings. *Circulation* 105, 21 (2002), 2487–2492.
- [78] KENNEDY, J. High-intensity focused ultrasound in the treatment of solid tumours. *Nature Reviews Cancer* 5, 4 (2005), 321–327.
- [79] KHALLAGHI, S., MOUSAVI, P., GONG, R. H., GILL, S., BOISVERT, J., FICHTINGER, G., PICHORA, D., BORSCHNECK, D., AND ABOLMAESUMI, P. Registration of a statistical shape model of the lumbar spine to 3D ultrasound images. In *Medical Image Computing and Computer Aided Interventions (MICCAI)* (2010), vol. 6362, pp. 68–75.

REFERENCES

- [80] KIM, H., AND VARGHESE, T. Hybrid spectral domain method for attenuation slope estimation. *Ultrasound in Medicine & Biology* 34, 11 (2008), 1808–1819.
- [81] KLEIN, S., STARING, M., MURPHY, K., VIERGEVER, M., AND PLUIM, J. Elastix: a toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging* 29, 1 (2010), 196–205.
- [82] KLEIN, T., HANSSON, M., KARAMALIS, A., AND NAVAB, N. Registration of RF ultrasound data using hybrid local binary patterns. In *9th IEEE International Symposium on Biomedical Imaging: Nano to Macro* (2012), pp. 1072–1075.
- [83] KLEIN, T., HANSSON, M., AND NAVAB, N. Modeling of multi-view 3D freehand radio frequency ultrasound. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2012), vol. 7510, pp. 422–429.
- [84] KOSARA, R., HEALEY, C., INTERRANTE, V., LAIDLAW, D., AND WARE, C. Thoughts on user studies: Why, how, and when. *IEEE Computer Graphics and Applications* 23, 4 (2003), 20–25.
- [85] KOWAL, J., AMSTUTZ, C., LANGLOTZ, F., TALIB, H., AND BALLESTER, M. G. Automated bone contour detection in ultrasound B-mode images for minimally invasive registration in computer-assisted surgery—an in vitro evaluation. *The Int. Journal of Medical Robotics Computer Assisted Surgery (MRCAS)* 3, 4 (2007), 341–348.
- [86] KRAFT, S., KARAMALIS, A., SHEET, D., DRECOLL, E., RUMMENY, E. J., NAVAB, N., NOËL, P. B., AND KATOUZIAN, A. Introducing nuclei scatterer patterns into histology-based intravascular ultrasound simulation framework. In *SPIE Medical Imaging* (2013), vol. 8675, International Society for Optics and Photonics, pp. 86750Y–86750Y–6.
- [87] KRISSIAN, K., WESTIN, C., KIKINIS, R., AND VOSBURGH, K. Oriented speckle reducing anisotropic diffusion. *IEEE Transactions on Image Processing* 16, 5 (2007), 1412–1424.
- [88] KROES, T., POST, F., AND BOTHA, C. Exposure render: An interactive photo-realistic volume rendering framework. *PLoS ONE* 7, 7 (07 2012).
- [89] KRÜGER, J., AND WESTERMANN, R. Acceleration techniques for GPU-based volume rendering. In *Proceedings of the 14th IEEE Visualization Conference* (2003), pp. 287 – 292.
- [90] KRÜGER, J., AND WESTERMANN, R. Linear Algebra Operators for GPU Implementation of Numerical Algorithms. *GPU Gems 2* (2005), 908 – 916.
- [91] KUC, R., AND LI, H. Reduced-order autoregressive modelling for center-frequency estimation. *Ultrasonic Imaging* 7, 3 (1995), 244–251.

-
- [92] KUTTER, O., KARAMALIS, A., WEIN, W., AND NAVAB, N. A GPU-based framework for simulation of medical ultrasound. In *SPIE Medical Imaging* (2009), vol. 7261, International Society for Optics and Photonics, pp. 726117–9.
- [93] KUTTER, O., SHAMS, R., AND NAVAB, N. Visualization and GPU-accelerated simulation of medical ultrasound from CT images. *Computer Methods and Programs in Biomedicine* 94, 3 (2009), 250–266.
- [94] LARAMEE, R., HAUSER, H., DOLEISCH, H., VROLIJK, B., POST, F., AND WEISKOPF, D. The state of the art in flow visualization: Dense and texture-based techniques. *Computer Graphics Forum* 22, 2 (2004), 203–221.
- [95] LEE, V., KIM, C., CHHUGANI, J., DEISHER, M., KIM, D., NGUYEN, A., SATISH, N., SMELYANSKIY, M., CHENNUPATY, S., HAMMARLUND, P., SINGHAL, R., AND DUBEY, P. Debunking the 100x GPU vs. CPU myth: An evaluation of throughput computing on CPU and GPU. In *ACM SIGARCH Computer Architecture News* (2010), vol. 38, pp. 451–460.
- [96] LEROY, A., MOZER, P., PAYAN, Y., AND TROCCAZ, J. Rigid registration of freehand 3D ultrasound and CT-scan kidney images. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2004), vol. 3216, pp. 837–844.
- [97] LEVITIN, A. *Introduction to the Design & Analysis of Algorithms*. Addison-Wesley Reading, MA, 2003.
- [98] LI, Y., AND ZAGZEBSKI, J. Computer model for harmonic ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 47, 4 (2000), 1000–1013.
- [99] LINDHOLM, E., KILGARD, M., AND MORETON, H. A user-programmable vertex engine. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (2001), pp. 149–158.
- [100] LINDOP, J., TREECE, G., GEE, A., AND PRAGER, R. 3D elastography using freehand ultrasound. *Ultrasound in Medicine & Biology* 32, 4 (2006), 529–546.
- [101] MANBACHI, A., AND COBBOLD, R. Development and application of piezoelectric materials for ultrasound generation and detection. *Ultrasound* 19, 4 (2011), 187–196.
- [102] MAST, T., SOURIAU, L., LIU, D., TABELI, M., NACHMAN, A., AND WAAG, R. A K-space method for large-scale models of wave propagation in tissue. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 48, 2 (2001), 341–354.
- [103] MINTZ, G. S., AND MISSEL, E. What is behind the calcium? The relationship between calcium and necrotic core on virtual histology analyses:reply. *Letters to the Editor, European Heart Journal* 30, 1 (2009), 125–126.

REFERENCES

- [104] MUCCI, R. A comparison of efficient beamforming algorithms. *IEEE Transactions on Acoustics, Speech and Signal Processing* 32, 3 (1984), 548–558.
- [105] MULLER, J., TAWAKOL, A., KATHIRESAN, S., AND NARULA, J. New opportunities for identification and reduction of coronary risk: treatment of vulnerable patients, arteries, and plaques. *Journal of the American College of Cardiology* 47, 8s1 (2006), C2–C6.
- [106] MÜLLER, M., KARAMALIS, A., AND NAVAB, N. Visualization of organ motion during breathing from 4D datasets. In *Augmented Environments for Computer-Assisted Interventions*, vol. 7815. Springer, 2013, pp. 34–44.
- [107] NAIR, A., KUBAN, B., TUZCU, E., SCHOENHAGEN, P., NISSEN, S., AND VINCE, D. Coronary plaque classification with intravascular ultrasound radiofrequency data analysis. *Circulation* 106, 17 (2002), 2200–2206.
- [108] NICKOLLS, J., BUCK, I., GARLAND, M., AND SKADRON, K. Scalable parallel programming with CUDA. *Queue* 6, 2 (2008), 40–53.
- [109] NICKOLLS, J., AND DALLY, W. The GPU computing era. *IEEE Micro* 30, 2 (2010), 56–69.
- [110] NOBLE, J. Ultrasound image segmentation and tissue characterization. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine* 224, 2 (2010), 307–316.
- [111] NOBLE, J., AND BOUKERROUI, D. Ultrasound image segmentation: A survey. *IEEE Transactions on Medical Imaging* 25, 8 (2006), 987–1010.
- [112] NTZIACHRISTOS, V., TUNG, C., BREMER, C., AND WEISSLEDER, R. Fluorescence molecular tomography resolves protease activity in vivo. *Nature medicine* 8, 7 (2002), 757–761.
- [113] O’DONNELL, M., SKOVORADA, A., AND SHAPO, B. Measurement of arterial wall motion using fourier based speckle tracking algorithms. In *IEEE Ultrasonics Symposium* (1991), pp. 1101–1104.
- [114] OPHIR, J., ALAM, S., GARRA, B., KALLEL, F., KONOFAGOU, E., KROUSKO, T., MERRITT, C., RIGHETT, R., SOUCHON, R., SRINIVASAN, S., AND VARGHESE, T. Elastography: imaging the elastic properties of soft tissues with ultrasound. *Journal of Medical Ultrasonics* 29, 4 (2002), 155–171.
- [115] OWENS, J., LUEBKE, D., GOVINDARAJU, N., HARRIS, M., KRÜGER, J., LEFOHN, A., AND PURCELL, T. A survey of general-purpose computation on graphics hardware. In *Computer Graphics Forum* (2007), vol. 26, pp. 80–113.
- [116] PENNEY, G., BARRATT, D., CHAN, C., SŁOMCZYKOWSKI, M., CARTER, T., EDWARDS, P., AND HAWKES, D. Cadaver validation of intensity-based ultrasound to CT registration. *Medical Image Analysis* 10, 3 (2006), 385–395.

-
- [117] PENNEY, G., BLACKALL, J., HAMADY, M., SABHARWAL, T., ADAM, A., AND HAWKES, D. Registration of freehand 3D ultrasound and magnetic resonance liver images. *Medical Image Analysis* 8, 1 (2004), 81–91.
- [118] PERONA, P., AND MALIK, J. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 7 (1990), 629–639.
- [119] PFEIFFER, F., WEITKAMP, T., BUNK, O., AND DAVID, C. Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources. *Nature Physics* 2, 4 (2006), 258–261.
- [120] PINTON, G., DAHL, J., ROSENZWEIG, S., AND TRAHEY, G. A heterogeneous nonlinear attenuating full-wave model of ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 56, 3 (2009), 474–488.
- [121] PINTON, G., AND TRAHEY, G. A comparison of time-domain solutions for the full-wave equation and the parabolic wave equation for a diagnostic ultrasound transducer. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 55, 3 (2008), 730–733.
- [122] POON, T., AND ROHLING, R. Three-dimensional extended field-of-view ultrasound. *Ultrasound in Medicine & Biology* 32, 3 (2006), 357–369.
- [123] POST, F., VROLIJK, B., HAUSER, H., LARAMEE, R., AND DOLEISCH, H. Feature extraction and visualization of flow fields. *Eurographics 2002 State-of-the-Art Reports 1* (2002), 69–100.
- [124] PRCKOVSKA, V., PEETERS, T., VAN ALMSICK, M., ROMENY, B., AND VILANOVA I BARTROLI, A. Fused DTI/HARDI visualization. *IEEE Transactions on Visualization and Computer Graphics* 17, 10 (2011), 1407–1419.
- [125] ROELANDT, J. Seeing the invisible: a short history of cardiac ultrasound. *European Journal of Echocardiography* 1, 1 (2000), 8–11.
- [126] ROHLING, R., GEE, A., AND BERMAN, L. Three-dimensional spatial compounding of ultrasound images. *Medical Image Analysis* 1, 3 (1997), 177–193.
- [127] ROHLING, R., GEE, A., AND BERMAN, L. A comparison of freehand three-dimensional ultrasound reconstruction techniques. *Medical Image Analysis* 3, 4 (1999), 339–359.
- [128] SARRUT, D., BOLDEA, V., MIGUET, S., AND GINESTET, C. Simulation of four-dimensional CT images from deformable registration between inhale and exhale breath-hold CT scans. *Medical Physics* 33, 3 (2006), 605–617.
- [129] SHAMS, R., HARTLEY, R., AND NAVAB, N. Real-time simulation of medical ultrasound from CT images. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2008), vol. 5242, pp. 734–741.

REFERENCES

- [130] SHANKAR, P. Speckle reduction in ultrasound B-scans using weighted averaging in spatial compounding. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 33, 6 (1986), 754–758.
- [131] SHAO, W., WU, R., LING, K., THNG, C., HO, H., CHENG, C., AND NG, W. Evaluation on similarity measures of a surface-to-image registration technique for ultrasound images. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2006), vol. 4191, pp. 742–749.
- [132] SHEET, D., KARAMALIS, A., ESLAMI, A., NOËL, P., CHATTERJEE, J., RAY, A. K., LAINE, A. F., CARLIER, S. G., NAVAB, N., AND KATOZIAN, A. Joint learning of ultrasonic backscattering statistical physics and signal confidence primal for characterizing atherosclerotic plaques using intravascular ultrasound. *Medical Image Analysis (to appear in print)*.
- [133] SHEET, D., KARAMALIS, A., ESLAMI, A., NOËL, P., VIRMANI, R., NAKANO, M., CHATTERJEE, J., RAY, A. K., LAINE, A. F., CARLIER, S. G., NAVAB, N., AND KATOZIAN, A. Hunting for necrosis in the shadows of intravascular ultrasound. *Computerized Medical Imaging and Graphics (to appear in print)*.
- [134] SHEET, D., KARAMALIS, A., KRAFT, S., NOËL, P. B., VAG, T., SADHU, A., KATOZIAN, A., NAVAB, N., CHATTERJEE, J., AND RAY, A. K. Random forest learning of ultrasonic statistical physics and object spaces for lesion detection in 2D sonomammography. In *SPIE Medical Imaging* (2013), vol. 8675, International Society for Optics and Photonics, pp. 867515–867518.
- [135] SHEN, R., CHENG, I., LI, X., AND BASU, A. Stereo matching using random walks. In *19th International Conference on Pattern Recognition (ICPR)* (2008), pp. 1–4.
- [136] SOLBERG, O., LINDSETH, F., TORP, H., BLAKE, R., AND NAGELHUS HERNES, T. Freehand 3D ultrasound reconstruction algorithms—a review. *Ultrasound in Medicine & Biology* 33, 7 (2007), 991–1009.
- [137] STONE, G. W., AND ET AL, M. A prospective natural-history study of coronary atherosclerosis. *New England Journal of Medicine* 364, 3 (2011), 226–235.
- [138] STONE, J., GOHARA, D., AND SHI, G. OpenCL: A parallel programming standard for heterogeneous computing systems. *Computing in Science & Engineering* 12, 3 (2010), 66–72.
- [139] STRANG, G. A framework for equilibrium equations. *SIAM Review* 30, 2 (1988), 283–297.
- [140] STRANG, G. *Computational Science and Engineering*. Wellesley-Cambridge Press, 2007.

-
- [141] SUN, X., ROSIN, P., MARTIN, R., AND LANGBEIN, F. Random walks for feature-preserving mesh denoising. *Computer Aided Geometric Design* 25, 7 (2008), 437–456.
- [142] SZABO, T. *Diagnostic Ultrasound Imaging: Inside Out*. Academic Press, 2004.
- [143] SZELISKI, R. *Computer vision: Algorithms and Applications*. Springer, 2010.
- [144] TEMPANY, C., STEWART, E., MCDANNOLD, N., QUADE, B., JOLESZ, F., AND HYNYNEN, K. MR imaging-guided focused ultrasound surgery of uterine leiomyomas: A feasibility study. *Radiology* 226, 3 (2003), 897–905.
- [145] TER HAAR, G., AND COUSSIOS, C. High intensity focused ultrasound: physical principles and devices. *International Journal of Hyperthermia* 23, 2 (2007), 89–104.
- [146] THOMENIUS, K. Evolution of Ultrasound Beamformers. In *IEEE Ultrasonics Symposium* (1996), pp. 1615–1622.
- [147] TONG, S., CARDINAL, H., MCLOUGHLIN, R., DOWNEY, D., AND FENSTER, A. Intra-and inter-observer variability and reliability of prostate volume measurement via two-dimensional and three-dimensional ultrasound imaging. *Ultrasound in Medicine & Biology* 24, 5 (1998), 673–681.
- [148] TRANQUART, F., GRENIER, N., EDER, V., AND POURCELOT, L. Clinical use of ultrasound tissue harmonic imaging. *Ultrasound in Medicine & Biology* 25, 6 (1999), 889–894.
- [149] TREAT, L., MCDANNOLD, N., VYKHODTSEVA, N., ZHANG, Y., TAM, K., AND HYNYNEN, K. Targeted delivery of doxorubicin to the rat brain at therapeutic levels using MRI-guided focused ultrasound. *International Journal of Cancer* 121, 4 (2007), 901–907.
- [150] TREECE, G., PRAGER, R., AND GEE, A. Ultrasound attenuation measurement in the presence of scatterer variation for reduction of shadowing and enhancement. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 52, 12 (2005), 2346–2360.
- [151] VAN WAMEL, A., BOUAKAZ, A., BERNARD, B., TEN CATE, F., AND DE JONG, N. Radionuclide tumour therapy with ultrasound contrast microbubbles. *Ultrasonics* 42, 1 (2004), 903–906.
- [152] VANDEMEULEBROUCKE, J., SARRUT, D., AND CLARYSSE, P. The POPI-model, a point-validated pixel-based breathing thorax model. In *Conference on the Use of Computers in Radiation Therapy* (2007).
- [153] WACHINGER, C., AND NAVAB, N. Similarity metrics and efficient optimization for simultaneous registration. In *IEEE Conference on Computer Vision and Pattern Recognition* (2009), pp. 667–674.

REFERENCES

- [154] WACHINGER, C., YIGITSOY, M., RIJKHORST, E., AND NAVAB, N. Manifold learning for image-based breathing gating in ultrasound and MRI. *Medical Image Analysis* 16, 4 (2012), 806 – 818.
- [155] WADE, G. Human uses of ultrasound: ancient and modern. *Ultrasonics* 38, 1 (2000), 1–5.
- [156] WANG, Z., LI, J., AND WU, R. Time-delay-and time-reversal-based robust capon beamformers for ultrasound imaging. *IEEE Transactions on Medical Imaging* 24, 10 (2005), 1308–1322.
- [157] WEIN, W., BRUNKE, S., KHAMENE, A., CALLSTROM, M., AND NAVAB, N. Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Medical Image Analysis* 12, 5 (2008), 577–585.
- [158] WEIN, W., AND KHAMENE, A. Image-based method for in-vivo freehand ultrasound calibration. In *SPIE Medical Imaging* (2008), vol. 6920, International Society for Optics and Photonics, pp. 69200K–69200K–7.
- [159] WEIN, W., KUTTER, O., AICHERT, A., ZIKIC, D., KAMEN, A., AND NAVAB, N. Automatic non-linear mapping of pre-procedure CT volumes to 3D ultrasound. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro* (2010), pp. 1225–1228.
- [160] WEIN, W., ROPER, B., AND NAVAB, N. Integrating diagnostic B-mode ultrasonography into CT-based radiation treatment planning. *IEEE Transactions on Medical Imaging* 26, 6 (2007), 866–879.
- [161] WESTERVELT, P. Parametric acoustic array. *The Journal of the Acoustical Society of America* 35, 4 (1963), 535–537.
- [162] WILD, J., AND NEAL, D. Use of high-frequency ultrasonic waves for detecting changes of texture in living tissues. *Lancet* 1, 6656 (1951), 655–657.
- [163] WILHJELM, J., JENSEN, M., BRANDT, T., SAHL, B., MARTINSEN, K., JESPERSEN, S., AND FALK, E. Some imaging strategies in multi-angle spatial compounding. In *IEEE Ultrasonics Symposium* (2000), vol. 2, pp. 1615–1618.
- [164] WU, Q., AND MOHAN, R. Algorithms and functionality of an intensity modulated radiotherapy optimization system. *Medical Physics* 27 (2000), 701–711.
- [165] XIAO, G., BRADY, M., NOBLE, J., AND ZHANG, Y. Segmentation of ultrasound B-mode images with intensity inhomogeneity correction. *IEEE Transactions on Medical Imaging* 21, 1 (2002), 48–57.
- [166] YIGITSOY, M., WACHINGER, C., AND NAVAB, N. Temporal groupwise registration for motion modeling. In *Information Processing in Medical Imaging (IPMI)* (2011), vol. 6801, Springer, pp. 648–659.

-
- [167] YONGCHEN, S., YANWU, D., JIE, T., AND ZHENSHENG, T. Ultrasonic propagation parameters in human tissues. In *IEEE Ultrasonics Symposium* (1986), pp. 905–908.
- [168] YU, Y., AND ACTON, S. Speckle reducing anisotropic diffusion. *IEEE Transactions on Image Processing* 11, 11 (2002), 1260–1270.
- [169] YU, Y., AND WANG, J. Backscatter-contour-attenuation joint estimation model for attenuation compensation in ultrasound imagery. *IEEE Transactions on Image Processing* 19, 10 (2010), 2725–2736.
- [170] YUAN, X., BORUP, D., WISKIN, J., BERGGREN, M., AND JOHNSON, S. Simulation of acoustic wave propagation in dispersive media with relaxation losses by using FDTD method with PML absorbing boundary condition. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 46, 1 (1999), 14–23.
- [171] ZAGZEBSKI, J. *Essentials of Ultrasound Physics*. Mosby, 1996.
- [172] ZHANG, J., ZHENG, J., AND CAI, J. Interactive mesh cutting using constrained random walks. *IEEE Transactions on Visualization and Computer Graphics* 17, 3 (2011), 357–367.

