Computer Aided Medical Procedures
Prof. Dr. Nassir Navab

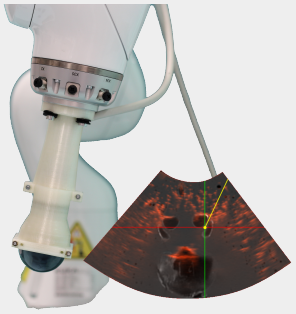

Dissertation

Advanced Ultrasound Imaging Techniques for Computer Assisted Interventions

Oliver Zettinig

Fakultät für Informatik
Technische Universität München

Technische Universität München
Fakultät für Informatik
Lehrstuhl für Informatikanwendungen in der Medizin & Augmented Reality

Advanced Ultrasound Imaging Techniques for Computer Assisted Interventions

Oliver Zettinig

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:		Prof. Dr. Nils Thuerey
Prüfer der Dissertation:	1.	Prof. Dr. Nassir Navab
	2.	Prof. Dr. Terry Peters University of Western Ontario, Kanada

Die Dissertation wurde am 22.03.2017 bei der Technischen Universität München eingereicht und durch die Fakultät für Informatik am 02.08.2017 angenommen.

Oliver Zettinig
Advanced Ultrasound Imaging Techniques for Computer Assisted Interventions
Dissertation, Version 1.1

Technische Universität München
Fakultät für Informatik
Lehrstuhl für Informatikanwendungen in der Medizin & Augmented Reality
Boltzmannstraße 3
85748 Garching bei München

Abstract

Ultrasound imaging is commonplace in clinical routine and has become the standard of care for a plethora of diagnostic scenarios. Due to issues such as limited image quality or obstructed visibility of anatomy, the exclusive use of this modality for interventional guidance purposes has, however, not yet reached a comparable level of maturity. This thesis addresses several challenges associated with ultrasound imaging by proposing advanced techniques for interventional use. Although applied to a broad spectrum of clinical fields and anatomies, their underlying methodology is generic and can be transferred to other medical scenarios.

First, a framework for multi-modal prostate biopsy guidance is introduced, allowing urologists to accurately target suspicious lesions by combining trans-rectal ultrasound information with complementary functional tomographic data. The crucial part of this fusion consisting of deformable image registration is solved by two novel algorithms based on automatically segmented prostate surfaces or a preconditioned intensity similarity metric and a statistical deformation model.

To overcome the challenges of manual acquisitions, in particular navigation to and maintenance of appropriate location and suitable acoustic window, robotic solutions are studied. Based on multi-modal image registration, a visual servoing control scheme for neurosurgical navigation is introduced. While compensating for target anatomy movements in real-time, it allows for automatic needle guide alignment for accurate manual insertions. The suitability of such systems for reliable robotic acquisitions even in absence of planning data is demonstrated by applying the developed methods, including image quality optimizations using confidence maps, for automated abdominal aortic aneurysm screenings.

Through Doppler modes, ultrasound physics uniquely allows fast analysis of blood flow dynamics, albeit limited to 2D projections. This thesis introduces a novel technique to recover 3D velocity information in combination with a temporal flow profile using measurements from multiple directions. Due to the importance of accurate and linearly independent sampling, the advantages of robotic acquisition schemes can be hereby fully exploited.

Results of phantom experiments, volunteer studies and clinical patient evaluations, all in close collaboration with medical partners, demonstrate the great potential benefit of advanced ultrasound imaging techniques in interventional settings in terms of both efficacy and efficiency.

Zusammenfassung

Ultraschall-Bildgebung ist aus der klinischen Routine nicht mehr wegzudenken und hat sich als Untersuchungsstandard in vielen diagnostischen Bereichen etabliert. Aufgrund der begrenzeten Bildqualität und der eingeschränkten Sichtbarkeit anatomischer Strukturen bleibt die Verbreitung als alleiniges Bildgebungsverfahren für interventionelle Navigation allerdings deutlich zurück. Diese Dissertation befasst sich mit vielfältigen Herausforderungen medizinischer Ultraschall-Bildgebung und präsentiert erweiterte Verfahren für den interventionellen Einsatz. Trotz der Anwendung auf einige ausgewählte klinische Bereiche und Anatomien bleiben die zugrundeliegenden Methoden generisch und leicht auf andere medizinische Szenarien übertragbar.

Zunächst wird ein System für die multi-modale Führung von Prostatafusionsbiopsien vorgestellt, das es Urologen erlaubt, verdächtige Läsionen zielgerichtet unter Berücksichtigung von transrektalem Ultraschall und komplementären, funktionellen Schnittbildern zu biopsieren. Der hierfür wesentliche Rechenschritt, die elastische Bildregistrierung, wird durch zwei neuartige Algorithmen realisiert, die entweder auf automatisch segmentierten Prostataoberflächen oder auf einer vorkonditionierten Ähnlichkeitsmetrik in Kombination mit einem statistischen Deformationsmodell beruhen.

Im Hinblick auf die Herausforderungen manueller Bildaquise, insbesondere Navigation zu einer geeigneten Position und Beibehaltung des akustischen Fensters, werden robotische Verfahren untersucht. Aufbauend auf multi-modaler Bildregistrierung wird ein Visual Servoing-Schema für die neurochirurgische Navigation präsentiert. Neben der Kompensation von Bewegungen der Zielanatomie in Echtzeit ist es in der Lage, eine Nadelführung für präzise, manuelle Punktionen automatisch auszurichten. Die Einsetzbarkeit derartiger Systeme auch ohne patientenspezifische Planungsdaten wird anhand eines autonomen Screeningssystems für abdominale Aortenaneurysmata demonstriert, zu dessen wichtigsten Komponenten eine Bildqualitätsoptimierung mithilfe von Confidence Maps zählt.

Im Rahmen von Doppler-Aufnahmen erlaubt Ultraschallbildgebung auch einzigartige Analysen der Blutflussdynamik, wenn auch nur in der Form von zweidimensionelen Projektionen. Diese Dissertation stellt ein neues Verfahren zur gleichzeitigen, dreidimensionalen Rekonstruktion eines Blutflussgeschwindigkeitsfeldes und eines temporalen Flussprofiles aus Messungen verschiedener Richtungen vor. Wegen der hohen Anforderungen in Bezug auf präzise, linear unabhängige Abtastung können auch hier die Vorteile robotischer Aufnahmeschemata ausgenutzt werden.

Die präsentierten Ergebnisse von Phantomexperimenten, Freiwilligenevaluationen und klinischen Patientenstudien, die alle in enger Zusammenarbeit mit klinschen Partnern durchgeführt wurden, demonstrieren das große Potential von erweiterten Verfahren der Ultraschall-Bildgebung im interventionellen Einsatz hinsichtlich Effizienz und Wirksamkeit.

Acknowledgments

First and foremost, I would like to express by deep gratitude to Prof. Nassir Navab. I could not have appreciated the freedom he granted me in pursuing my own ideas more. Vividly remembering many inspiring discussions, I want to thank him for his honest encouragement, enduring guidance and contagious enthusiasm all along the way.

Likewise, I am extremely thankful to my colleagues at CAMP for their great support in all scientific and non-scientific matters, especially to my advisor Dr. Benjamin Frisch for his cordial yet demanding supervision and his invaluable ability to regularly add a different perspective. I also want to thank Dr. Maximilian Baust, Dr. Tobias Lasser and Dr. Christoph Hennersperger for their great support and guidance, and I already miss the both fruitful and humorous discussions with Salvatore Virga, Julia Rackerseder, Marco Esposito, Beatrice Lentes, Rüdiger Göbl, Amit Shah, Fausto Milletarì, Jabob Weiss, and Pierre Chatelain at IFL and in Garching. A very special thank you is dedicated to Martina Hilla for her protection from bureaucratic calamities.

Working on interventional ultrasound imaging, my research has been highly dependent on clinical collaborators, in particular at Klinikum Rechts der Isar. It is my wish to warmly thank Dr. Tobias Maurer, especially for serving as my interdisciplinary mentor, and the team of the urology department for their precious time, continuous confidence and tedious manual annotation sessions. I also want to thank neurosurgeon Dr. Yu-Mi Ryang, nuclear radiologist Dr. Matthias Eiber, and vascular surgeon Dr. Karin Pfister (Universitätsklinikum Regensburg) for their support and am grateful for the fruitful collaboration.

During the last three years, I spent some time in Baltimore, where I had the wonderful chance to work with Dr. Bernhard Fuerst, Risto Kojcev, Javad Fotouhi, and Sing Chun Lee, whom I want to thank not only for providing valuable insights but also for the great time off-campus. A strong connection has also grown to the team of ImFusion, especially to Dr. Wolfgang Wein, long-term companion Dr. Christian Schulte zu Berge, as well as Mehrdad Salehi and Markus Müller.

Finally, I would like to express my gratitude to my parents Gabriela and Raimund, my sister Désirée, and my closest friends and colleagues. I am thankful for their great support and perpetual confidence. This most notably applies to Maria, for whose unconditional, emotional support I am deeply grateful. Ευχαριστώ που ολοκληρώνεις την ζωή μου.

Contents

Part I

Introduction

1 Introduction

1.1 Motivation

Since the 1940s, when ultrasound (US) waves were first used to examine interior parts of the human body [215], sonography has seen a tremendous technological development. In contrast to X-ray diagnostics, which was only capable of producing integral projection images, B-mode ultrasound allowed to visualize cross-sections of the body, almost two decades before the advent of Computed Tomography (CT) in 1971 and Magnetic Resonance Imaging (MRI) in 1973. Due to its lack of ionizing radiation, independence of potentially nephrotoxic contrast agents and the relatively high mobility of US scanners, ultrasound quickly gained popularity in many medical disciplines. Further breakthroughs such as real-time performance and the Doppler-based color and power imaging modes laid the foundation of modern practice and helped to establish ultrasound as the modality of choice for a broad variety of clinical scenarios. Today, sonographers can choose between a multitude of inexpensive systems and hand-held transducers with high spatial and temporal resolution, not only enabling fast diagnostic on-site scans but also image-based guidance for a limited set of interventions.

Nevertheless, the nature of currently available ultrasound systems still poses significant challenges for their application in clinical routine. On the one hand, US images inherently incorporate a high amount of noise, suffer from limited anatomical visibility and may contain a variety of artifacts due to the underlying physical principles of image formation. On the other hand, both image quality itself and the subsequent interpretation of the manually acquired data is highly dependent on well-experienced clinical staff. Because tissue reflectance varies with the direction of the incoming ultrasonic beam and limited anatomical context has to be regularly compensated with high cognitive load, the manual navigation of the transducer to an appropriate location and the maintenance of a suitable acoustic window are cumbersome and require extensive training.

While these challenges are commonly considered manageable for diagnostic scenarios, they often prohibit the replacement of other, more harmful or more expensive modalities by sonographic techniques in interventional settings. Previous approaches to reconstruct 3D US volumes from multiple individual frames, possibly acquired autonomously by a robot, paved the way for more advanced guidance methods, often coupled with state-of-the-art segmentation and registration techniques. Individually, the required components for sophisticated interventional US guidance have seen great leaps of technological evolution in recent years, especially in terms of hardware improvements and algorithmic advancement. Yet, their integration and translation from laboratory conditions into mature solutions for everyday clinical usage regularly fails due to the high complexity involved.

Technologically well-grounded in this environment, this thesis aims at bridging the final gap toward regular usage on humans by addressing current challenges with innovative methodologies. In particular, novel interventional multi-modal registration algorithms, original control laws for robotic navigation using visual servoing, and a practical method to reconstruct 3D velocity fields over time from Doppler information are proposed.

1.2 Outline

This thesis is structured in three parts. In **Part I**, **chapter 1**, basic concepts of ultrasound imaging, including 3D compounding, as well as image registration are introduced, serving as a basis for the subsequent chapters.

Part II contains the scientific contributions of this thesis:

- **Chapter 2** introduces a framework for multi-modal prostate biopsy guidance, allowing urologists to accurately target suspicious lesions by combining trans-rectal ultrasound information with complementary functional tomographic data. Key elements are two novel elastic registration algorithms based on automatically segmented prostate surfaces or a preconditioned intensity similarity metric and a statistical deformation model.

- **Chapter 3** deals with image-based visual servoing concepts for robotic ultrasound. Based on multi-modal image registration, a control scheme for neurosurgical needle navigation is introduced that compensates for target anatomy movements in real-time. The developed method, reliably working even without patient-specific planning data, also allows for autonomous abdominal aortic aneurysm screenings.

- **Chapter 4** presents a novel technique to recover 3D velocity information in combination with a temporal flow profile using arbitrarily sampled Doppler measurements from multiple directions. Due to the importance of accurate and linearly independent sampling, the advantages of robotic acquisition schemes can be hereby fully exploited.

In each of these chapters, the presented methodology is not only evaluated using phantoms but also validated in-vivo on human patients (chapter 2) or healthy human volunteers (chapter 3, 4). Substantial parts of this thesis have already been published, and the respective publications are clearly indicated at the beginning of each chapter. Although the presented work can be considered my own if not explicitly declared otherwise, the usage of the first-person plural form indicates that many efforts where only possible as a team.

Finally, **Part III**, **chapter 5**, concludes the thesis and outlines potential directions of future work. The appendix (**Part IV**) contains lists of publications, figures, tables, and the bibliography.

1.3 Essentials of Ultrasound Imaging

Sound waves, from Latin *sonus*, are mechanical, predominantly longitudinal compression and displacement waves. The ultrasound spectrum (from Latin *ultra*, i.e. *beyond*) covers frequencies above the human audible limit of around 20 kHz. This is not to be confused with the term *supersonic* (from Latin *super*, i.e. *above*), which refers to velocities beyond the speed of sound. In contrast, the word *echo*, originally from Greek ἦχος (*echos*, i.e. *sound*), has a slightly different meaning and refers to the reflection of sound waves reaching a listener after some delay. The reason behind this seemingly different meaning can be found in ancient Greek mythology, where *Echo* was the name of a cursed mountain nymph who could not speak freely but only repeat the last words spoken to her. In the domain of ultrasound imaging, both Latin and Greek terms are used interchangeably today. While the term (ultra-)sonography, a Latin-Greek hybrid from γράφειν (*graphein*, i.e. *write*), refers to ultrasound imaging in general, the purely Greek name of the modality has often prevailed in particular fields, leading to a rich nomenclature of methods and techniques. Ultrasonic examinations of the heart, for instance, are commonly referred to as echocardiography.

Unlike other imaging modalities with more linear evolution, ultrasound methodology has seen a long, serpentine development toward modern clinical applications, and a vast amount of ingenious mathematicians, physicists and physicians have contributed over time. This section aims at summarizing ultrasound history and explaining the essential principles of sonographic acquisition, image formation and 3D processing.

1.3.1 Brief History

The fundamental properties of sound waves had already been studied and described by ancient Greek philosophers, but the English scientist Robert Hooke (1635–1703) was likely the first one to foresee their potential for diagnostic purposes, when he wrote [225]:

> "*It may be possible to discover the motion of the internal parts of bodies, whether animal, vegetable, or mineral, by the sound they make; that one may discover the works performed in the several offices and shops of a man's body, and thereby discover what instrument or engine is out of order, what works are going on at several times, and lie still at others, and the like. I could proceed further, but methinks I can hardly forbear to blush when I consider how the most part of men will look upon this: but, yet again, I have this encouragement, not to think all these things utterly impossible.*"

Throughout the 18th and 19th centuries, deep theoretical understanding of non-linear (acoustic) wave propagation in various materials was gained by mathematicians and physicists such as Leonhard Euler, Pierre-Simon Laplace, Jean-Baptiste d'Alembert, Hermann von Helmholtz, Gustav R. Kirchhoff, Siméon D. Poisson, Bernhard Riemann, Christian Doppler, and Lord John Rayleigh. The latter published a two-volume milestone compendium in 1877 titled *The Theory of Sound*, which is still referred to today [174].

Nonetheless, the emergence of medical ultrasound imaging is strongly linked to naval requirements in early 20$^{\text{th}}$ century. Shortly after the Titanic catastrophe, several scientists, including Lewis F. Richardson, a British meteorologist, suggested to utilize underwater sound wave reflections to detect icebergs. Roughly at the same time, Reginald Fessenden, a Canadian inventor, already practically demonstrated echo ranging capabilities for icebergs, albeit direction-blind due to the high wave length. During World War I, military efforts led to submarine detection devices pioneered by Paul Langevin, Robert Boyle, and Constantin Chilowsky, whose success is based on the integration of recent discoveries such as (reverse) piezo-electricity and triode amplifier tubes. Together, they coined the term *hydrophone* for underwater ultrasound sensing [215]. The technology quickly proved useful for depth sounding, too. Hereby, the naval term *to sound*, from Old English *sund* (the power, capacity, or act of swimming), ironically only refers to techniques measuring the depth of water at sea (traditionally with weights at the end of a long rope), and is etymologically unrelated to Latin *sonus*.

In the immediate years preceding World War II, multiple nations secretly and independently developed pulse-echo ranging systems using electromagnetic waves, which were later termed *radar* (radio detection and ranging). The important contribution of directional sweeps around 360°, combined with plan position indicators, was in return translated back to the supersonic domain and inspired the development of *sonar* (sound navigation and ranging). As a result, underwater ranging devices capable of observing more than one particular direction soon superseded previous setups [215].

The domain of medical imaging was only opened when pulse-echo methods were applied to probe parts of the human body. The *reflectoscope*, originally invented by Floyd Firestone in 1940 for the localization of metal defects, was applied to several anatomies roughly at the same time by teams in the United States, Sweden and Japan. The device visualized the amplitude of reflected sound waves over time on an oscilloscope, a 1D acquisition technique later called *A-mode* (for amplitude). The earliest efforts toward brightness, i.e. *B-mode* scans date back to 1942, when Karl Dussik used light bulbs glowing proportionally to the strength of the received signal. Despite his difficulties in transmitting ultrasound trough the skull, his work greatly inspired others. Douglass Howry was one of the first ones to study tomographic, cross-sectional images obtained by rotating a single element transducer around a patient in a water bath. Similarly, John J. Wild and John M. Reid developed one of the first hand-held scanners, again relying on a single transducer with controlled motion [240]. In either case, a display would accumulate time traces of individual scan rays, eventually forming a brightness image of strong reflectors in the tissue, especially bones. It was soon understood, that because of the remarkably constant speed of sound in biological tissues of approximately $c = 1,540$ m/s, these images could faithfully represent tissue geometry [215]. In addition, such tomographic cross-sections were capable of providing soft-tissue information that are not obtainable with X-ray projections.

Major breakthroughs in subsequent years were due to two seemingly unrelated developments. First, the invention of transistors and the fast pace of integrated circuit progress greatly facilitated electronic gray-scale acquisition and scan conversion to cathode ray tube formats, thus allowing a higher dynamic range than previously used storage oscilloscopes or conventional film by means of logarithmic compression and amplification [119]. Second, translational ef-

forts regarding electronic scanning and focused phase-array technology from electromagnetic applications marked a clear milestone. Image acquisition with a single element transducer necessitated its time-consuming, mechanical translation to cover the desired area of interest. Until the introduction of real-time *cinematographers*, ultrasound scans were therefore static and required the patient to remain motionless during the acquisition. While the first commercial real-time system, Siemens Vidoson (1965), still relied on a single, rotating element in combination with a parabolic mirror for parallel beam alignment, faster acquisitions can be achieved using arrays of transducer elements fired in turn. Combined, these two paradigm shifts led to modern, cart-integrated, phased-array systems such as the Hewlett Packard 70020A presented in 1981 [215]. With its 64 elements operating at 2.5, 3.0 or 5.0 MHz, it was capable of acquiring 121 scanlines at a frame rate of 30 Hz.

Ultrasound imaging quickly proved useful for blood velocity measurements using the Doppler effect, i.e. the change in frequency due to relative motion between source and observer, which was already well understood at the time. Originally based on phase shift measurements, the group of Henry P. Kalmus [110] used continuous wave (CW) ultrasound to measure flow velocities in fluids using pairs of transducers in 1954. A few years later, Shigeo Satomura and his colleagues confirmed the applicability of ultrasound to investigate blood flow in-vivo, suggesting that the method could be useful to diagnose vascular diseases. It took, however, several more years until the underlying physiological principles had been understood, and that the measurements were successful because of the scattering of red blood cells. Because the frequency shift generated by blood flow falls in the audio range, Doppler signals can not only be visualized on displays but also directly heard.

In the late 1960s, the advantages of pulsed wave (PW) methods for flow measurements became evident, because in contrast to CW-based signals, the depth of detected velocities could be recovered as well. In addition, only a single transducer was required for the estimation. As later exhaustively investigated, not the frequency shift caused by small scatterers itself but rather their movement between consecutive pulses contributed to the measured signal here [41]. Donald Baker, Vern Simmons, and Peter N. T. Wells pioneered range-gated Doppler to isolate different targets and laid the foundation for combined anatomical (B-mode) and functional imaging of cardiovascular structures. Color coding of flow measurements, later termed *color-coded duplex* (CCD), allowed for a direct mapping between anatomy and measured blood velocity, deepening the understanding of vascular diagnosis. Following early attempts toward fused visualizations, Marco A. Brandestini and Fred K. Forster [26] were the first ones to superimpose color-coded flow data onto B-mode images in 1978, which drastically increased the intuitiveness of ultrasound displays and remains a commonly used imaging mode on modern scanners.

Since the 1980s, when the fundamental technology had reached a sufficient maturity and had become widely established in clinical practice, ultrasound research has greatly diversified. Continuous incremental improvements in transducer design and technology have increased the overall image quality and opened new medical domains. 2D matrix arrays were introduced by Olaf T. von Ramm and Stephen Smith in 1987, allowing real-time scanning of an entire volume without physically moving the transducer, albeit with drastically reduced frame rates [200]. Based on the discovery of the advantageous effects of microbubbles injected into the blood stream, also novel contrast agents designed to enhance the ultrasound sensitivity gained more

importance [215]. A similar impact on sonographic diagnosis had the advent of elastography, a term used for methods measuring the elastic or viscoelastic properties of tissue. Conventional approaches in fact measured strain and relied on externally induced compression, regularly produced in a mechanical way by vibration. Tissue stiffness was then inferred by comparing images before and after. In contrast, shear wave elasticity imaging as proposed by Armen P. Sarvazyan [186] utilizes focused ultrasound (acoustic radiation force) to induce transverse shear waves through the tissue. The recently introduced technique of plain wave imaging has not only skyrocketed the achievable ultrasound frame rates by two orders of magnitude to the kHz range (*ultrafast* imaging), but also facilitated quantification in Doppler and shear wave elastography imaging [13]. Finally, with high-intensity focused ultrasound (HIFU) at hand, sonographic techniques have left the diagnostic realm and are now also used for ablation of a variety of cancer types as well as several neurological disorders, with modern systems by far exceeding the capabilities of lithotripsy devices traditionally used to break kidney stones [215].

Today, almost three centuries after Hooke's groundbreaking visions, medical ultrasound imaging has emerged to one of the most widely used diagnostic modalities, second to X-ray only, and the most widely used modality in certain fields such as gynecology, obstetrics and urology [155, 197]. As such, ultrasound is commonly considered a disruptive technology in the medical imaging field [115].

1.3.2 Ultrasound Wave Physics

Wave Properties

Two elementary types of acoustic waves propagating through some medium, schematically depicted in Fig. 1.1, can be distinguished based on the direction of particle motion:

i) **Longitudinal** waves, also denoted compression waves, are characterized by sinusoidal back-and-forth motions of particles, inducing propagating areas of local pressure disturbances, which are called compression and rarefaction. Both the local displacement of particles and the local density change over time as the wave travels through the medium. The transmission speed of these wave disturbances, commonly denoted c, is called the speed of sound and depends on the material stiffness. Since such material properties are regularly unknown a priori, assumptions regarding the speed of sound are common, and a constant value of $c = 1,540$ m/s is chosen as soft tissue compromise in most modern systems.

ii) **Transversal** waves, also called shear waves, are characterized by particle movements orthogonal to the wave propagation direction. Instead of areas of compression and rarefaction, motion in shear direction is present. Also in this case, a material properties define the amplitude of maximum displacement, especially the density and the shear modulus.

Although the latter type is gaining increasing importance because of advanced modalities such as shear wave elastography [186], a common simplification in understanding ultrasound physics is to assume that waves in the body propagate like waves *within* fluids due to the

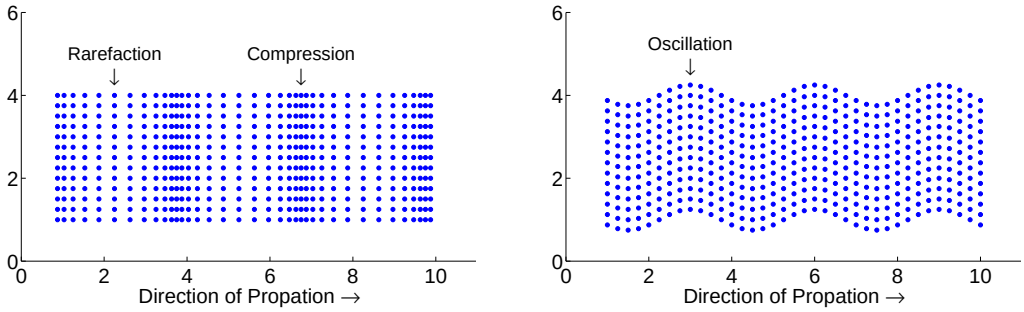

(a) Longitudinal wave **(b)** Transversal wave

Fig. 1.1. Overview of the two elementary wave types for an isotropic medium. Longitudinal waves exhibit areas of compression and rarefaction which propagate through the medium. In transversal waves, particles oscillate orthogonal to the direction of wave transmission.

high water content of biological tissues. This assumption may be sufficient for many cases of US imaging and relates to a predominantly longitudinal propagation of waves. Note the difference to waves observable at the *surface* of fluids, for instance as occurring on the sea, which incorporate both longitudinal and transversal components. Special care is also required for particular anatomies. The human heart, for instance, is a notable exception because its muscle fibers, arranged in sheets, form an orthotropic elastic solid [253], which has significant implications with respect to acoustic wave propagation.

As already mentioned above, the ultrasonic spectrum starts at the human audible limit of around 20 kHz and reaches beyond 1 GHz. Depending on the application, however, frequencies typically range from 1 to 60 MHz for medical imaging purposes, with most common ultrasounds scans being acquired with 20 MHz or less. The relationship between frequency f and wavelength λ, i.e. the distance between consecutive areas of compression or rarefaction, is given by:

$$\lambda = \frac{c}{f}. \tag{1.1}$$

For instance for the frequencies used in this thesis, this results to wavelengths of 0.47 mm for abdominal US (3.3 MHz, see chapter 3) or 0.19 mm for vascular US (8 MHz, see chapter 4).

The wavelength is tightly coupled to the spatial resolution of an US system, i.e. the minimum distance between two reflectors that can still be distinguished. Because US excitation is performed using multiple ($n = 2, 3$) pulses, axial resolution, i.e. along a scanline, is limited by the length of the entire pulse L_p:

$$R_a = \frac{L_p}{2} = \frac{n\,\lambda}{2}. \tag{1.2}$$

In lateral direction, i.e. perpendicular to the scanlines, the resolution is heavily dependent on the shape of the US beam and influenced by the focal length F and the active aperture of the transducer D [95]:

$$R_l = \sqrt{2}\lambda\frac{F}{D}. \tag{1.3}$$

In the general case, the fundamental principle of mass conservation states that the mass of fluid flowing into a partial volume is equal to the mass leaving it. This is stated by the continuity equation

$$\frac{\partial \rho}{\partial t} + \operatorname{div}(\rho \mathbf{v}) = 0, \tag{1.4}$$

where $\mathbf{v}(\mathbf{x}, t) = \partial \mathbf{u}(\mathbf{x})/\partial t$ is the velocity of a particle at location $\mathbf{x}$ at time t, expressed in terms of its displacement $\mathbf{u}$. ρ denotes the density and div the divergence operator defined as $\operatorname{div} \mathbf{a} = \nabla \cdot \mathbf{a}$ for some differentiable vector field $\mathbf{a}$. These velocities can be related to the induced changes of pressure $p(\mathbf{x}, t)$ with Euler's equation of motion [95]:

$$\rho \left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla)\mathbf{v} \right) = -\nabla p. \tag{1.5}$$

Since this relation does not account for viscosity, additional terms need to be added. In fact, Euler's equation is a special case of the well-known Navier-Stokes equation

$$\rho \left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla)\mathbf{v} \right) = -\nabla p + \mu \Delta \mathbf{v} + \left(\xi + \frac{1}{3}\mu \right) \nabla (\mathbf{v} \cdot \nabla). \tag{1.6}$$

Hereby, $\Delta = \nabla^2$, μ is the dynamic shear viscosity and $\xi = \lambda + \frac{2}{3}\mu$ the bulk viscosity with λ referring to the first Lamé parameter [41]. Equation 1.6 thus accounts for shear stress, which has a significant impact on wave propagation in viscous liquids. In combination with Eq. 1.4, it allows to fully describe longitudinal waves in homogeneous media. For convenience, it is possible, cf. [215], to derive $\mathbf{v}$ from a velocity potential ϕ such that $\mathbf{v} = -\nabla \phi$. This leads to the wave equation

$$\Delta \phi + \kappa \left(\xi + \frac{4}{3}\mu \right) \frac{\partial}{\partial t} \Delta \phi = \kappa \rho_0 \frac{\partial^2 \phi}{\partial t^2}, \tag{1.7}$$

where κ refers to the adiabatic compressibility, i.e. the relative change in volume or density due to a pressure change (ρ_0 is the equilibrium density):

$$\kappa = \frac{1}{\rho_0} \frac{\partial \rho}{\partial p} = \frac{1}{\rho_0} \frac{1}{c^2}. \tag{1.8}$$

For the assumption of idealized inviscid fluids, Eq. 1.7 simplifies to:

$$\Delta \phi = \kappa \rho_0 \frac{\partial^2 \phi}{\partial t^2} = \frac{1}{c^2} \frac{\partial^2 \phi}{\partial t^2}, \tag{1.9}$$

which is known as the Helmholtz equation in the frequency domain:

$$\Delta \Phi = -k^2 \Phi. \tag{1.10}$$

Hereby, k refers to the wave number $k = \omega/c$ with the angular frequency $\omega = 2\pi f$, and Φ is the Fourier transform of ϕ [215]. In this scenario, an intuitive relationship between particle velocity $\mathbf{v}$, and pressure p as well as its gradient can be found:

$$p = \rho_0 \frac{\partial \phi}{\partial t}, \tag{1.11}$$

$$\nabla p = -\rho_0 \frac{\partial \mathbf{v}}{\partial t}. \tag{1.12}$$

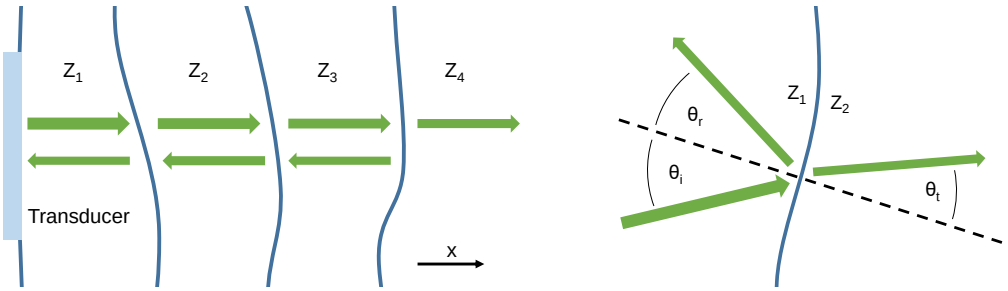

(a) Transmission across several interfaces **(b)** Snell's law

Fig. 1.2. The impedance difference between several boundaries as well as the incident angle with respect to an impedance boundary define the amplitude and direction of reflection and transmission.

The general solution in 3D to Eq. 1.9 can be formulated as:

$$\phi(\mathbf{x},t) = g\left(t - \frac{\mathbf{x}\cdot\mathbf{r}}{c}\right) + h\left(t + \frac{\mathbf{x}\cdot\mathbf{r}}{c}\right). \tag{1.13}$$

Vector $\mathbf{r}$ is the direction of the US wave, and g and h represent forward and backward traveling waves. The terms $t \pm (\mathbf{x}\cdot\mathbf{r})/c$ can be interpreted as phase, with the ratio computing the travel time of the wave due to the speed of sound. An important specific solution is given by the time harmonic,

$$\phi = \phi_0\left(e^{i(\omega t - \mathbf{k}\cdot\mathbf{r})} + e^{i(\omega t + \mathbf{k}\cdot\mathbf{r})}\right), \tag{1.14}$$

where vector $\mathbf{k}$ contains the three projections $[k_1, k_2, k_3]$ with $k^2 = \sum_1^3 k_i^2$ [215]. In practice, the real components of these complex exponentials would be considered. The pressure of the forward-traveling wave, for instance, can be computed as follows:

$$p(t) = p_0\left(\mathrm{Re}\left\{e^{i(\omega t - \mathbf{k}\cdot\mathbf{r})}\right\} + \cos(\omega t + \mathbf{k}\cdot\mathbf{r})\right). \tag{1.15}$$

As pointed out by Szabo [215], the presented wave equations can also be formulated for p and v instead of ϕ.

Impedance, Reflection and Transmission

An essential quantity for US imaging is the specific acoustic or characteristic impedance defined as the product of characteristic material properties i) density ρ_m and ii) medium-specific speed of sound c_m:

$$Z = \frac{p}{v} = \rho_m\, c_m, \tag{1.16}$$

where p and v correspond to the pressure of a forward-traveling particle wave only and the magnitude of the particle velocity therein, respectively. The unit of Z is Rayls, where 1 Rayl = 1 kg/m^{2}s. For instance, the characteristic impedance of water at 20° C is $Z = 1.48 \cdot 10^6$ Rayls [50] ($c = 1,481$ m/s, $\rho_0 = 998$ kg/m^3, see Tab. 1.1). The magnitude of the sound wave intensity, i.e. the energy carried by the wave per unit area, is then defined as follows:

$$I = p\, v = p^2 \frac{1}{Z}. \tag{1.17}$$

Fig. 1.2a illustrates the transmission of an ultrasound wave across several interfaces of different impedances. At the first interface, i.e. the boundary between Z_1 and Z_2, the wave bounces off and gets partially reflected. The solution to the wave equation in 1D (see Eq. 1.14) will have the following form:

$$p(x,t) = p_0 \left(e^{i(\omega t - kx)} + R\, e^{i(\omega t - kx)} \right). \tag{1.18}$$

In this scenario, R is a reflection factor scaling the amplitude of the backward-traveling wave. Note that the sign of kx is inverted compared to Eq. 1.14 because by convention, R will be negative for inversions of the incident wave. Whereas different conventions to represent the reflection factor have been reported in literature [41, 168], the one used in [215] is presented here.

The pressure decrease at the boundary defined by Eq. 1.18 then follows:

$$p_2 = p_0(1 + R), \tag{1.19}$$

and the particle velocity is given by:

$$v_2 = \frac{p_0(1 - R)}{Z_1}. \tag{1.20}$$

Using Eq. 1.16 and a few algebraic transformation, an intuitive relation for the reflection factor can be found:

$$R = \frac{Z_2 - Z_1}{Z_2 + Z_1}. \tag{1.21}$$

Because not the entire wave will be reflected, the remaining energy is passed along to the second medium with the transmission factor $T = 1 + R$, which can be determined as follows:

$$T = \frac{2Z_2}{Z_2 + Z_1}. \tag{1.22}$$

In analyzing the implications of these two factors, it is evident that reflections will only occur when there is a difference in impedance, i.e. $Z_1 \neq Z_2$. Open boundaries, i.e. $Z_2 = 0$, will cause a full inversion of the incident wave ($R = -1$). Incident and reflected waves will then cancel out so that $T = 0$. This scenario can be experienced in boundaries with air, for instance when transmitting US through the lung ($T \approx 0.5 \cdot 10^{-3}$). For the rather theoretical case of $Z_2 = \infty$, the incident wave will be reflected back directly ($R = +1$), both incident and reflected waves will add in phase ($T = 2$), but no energy is transmitted because $v_2 = p_2/Z_2 = 0$. Note that because the wave intensity is proportional to the squared amplitude (see Eq. 1.17), the squared coefficients R^2 and T^2 are used to determine the amount of reflection and transmission for the wave intensity.

It is, however, rarely the case that the acoustic wave hits impedance boundaries perpendicular to their surface. In any case, pressures and particle velocities remain continuous at any boundary. According to Snell's law, depicted in Fig. 1.2b, the reflection angle $\theta_r = \theta_i$ is identical to the incidence angle, the for the transmission angle θ_t, the ratio between the different speeds of sound in the two media needs to be considered:

$$\frac{\sin\theta_i}{\sin\theta_t} = \frac{c_1}{c_2}. \tag{1.23}$$

This relation can be used to determine the reflection and transmission factors in the general case:

$$R = \frac{Z_2 \cos\theta_i - Z_1 \cos\theta_t}{Z_2 \cos\theta_i + Z_1 \cos\theta_t}, \tag{1.24}$$

$$T = \frac{2Z_2 \cos\theta_i}{Z_2 \cos\theta_i + Z_1 \cos\theta_t}. \tag{1.25}$$

Note that the direction of the US wave changes at tissue interfaces for $\theta_i \neq \theta_t \neq 0$, which is commonly to as refraction.

Attenuation as Absorption and Scattering

The considerations above did not involve any energy losses and assumed that the acoustic wave does not additionally interact with its medium. In reality, however, energy is lost due to several phenomena including weak local heating (absorption), scattering, diffraction, and refraction [215], of which the first two are most important for practical considerations. Absorption, on the one hand, occurs as a result of the compressional displacement of particles, causing a dissipation of energy in the form of heat. Hence, the amount of absorption exhibited by a medium is dependent on its material properties. Scattering, on the other hand, is produced by diffuse reflections at objects smaller then the wavelength. The resulting interference waves, extending in different directions, continue to give rise to further reflections themselves. In the end, a recursive speckle pattern that appears to be noise can be observed in the US image. Although recent findings in the field of radar interference suggest that such patterns may in fact be deterministic, further investigations will be necessary to confirm this hypothesis in the field of acoustic waves [202]. Nevertheless, the presence of speckle is commonly used in a variety of applications, including tissue classification and image registration [196, 231]. Together, both mechanisms are commonly summarized as frequency-dependent ultrasound attenuation and are modeled using a multiplicative loss term:

$$p(x,t) = p_0 \, e^{i(\omega t - kx)} \, e^{-\alpha x}, \tag{1.26}$$

where α denotes an attenuation factor according to the Beer-Lambert-Law and is often measured in decibels (dB) per centimeter [215]. Note that the ratio r in dB can be computed with both amplitudes and intensities:

$$r = 20 \log_{10} \frac{p}{p_0} = 10 \log_{10} \frac{I}{I_0}. \tag{1.27}$$

The frequency dependence of α is often modeled using an approximation of the form:

$$\alpha(f) = \alpha_0 + \alpha_1 |f|^y, \tag{1.28}$$

where usually $\alpha_0 = 0$, α_1 a tissue-dependent scaling factor and the power exponent $y \in [0.9, 2]$. In Tab. 1.1, commonly used tissue properties constants are listed. Except for bone, which is hard to penetrate due to the enormous impedance difference, all listed tissue types include an exponential increase in attenuation with respect to the wave frequency. This explains the common trade-off between spatial resolution and penetration depth in choosing an appropriate frequency: Higher frequencies will offer superior resolution close to the surface but fail to reach farther into the tissue, and vice versa.

Tissue	c [m/s]	α_1 [dB/(MHz^y cm)]	y	ρ_0 [kg/m^3]	Z [MegaRayls]
Water (20° C)	1,481	$2.7{\cdot}10^{-3}$	2.00	998	1.482
Blood	1,584	0.14	1.21	1,060	1.679
Bone	3,198	3.54	0.90	1,990	6.364
Brain	1,562	0.58	1.30	1,035	1.617
Fat	1,430	0.60	1.00	928	1.327
Heart	1,554	0.52	1.00	1,060	1.647
Kidney	1,560	10.00	2.00	1,050	1.638
Liver	1,578	0.45	1.05	1,050	1.657
Muscle	1,580	0.57	1.00	1,041	1.645

Tab. 1.1. Tissue and material properties with respect to acoustic waves. *Data from [50].*

1.3.3 Image Formation and Modes

Transducer Geometry

Contrary to historical US systems, modern transducers contain an arrangement of numerous piezo-electric elements, usually between 128 and 256. By aligning them in an array, it becomes possible to utilize multiple elements at the same time to acquire a single scanline. This facilitates focusing the US beam, effectively increasing the spatial resolution in lateral direction.

A variety of probe geometries, shapes and sizes has been developed to account for different demands related to the anatomy and the route of access. Nonetheless, there is a limited set of fundamental design types that dominate the spectrum. Commonly used transducers can be categorized with the following criteria [41]:

i) The **geometric alignment** of piezo-electric elements either follows a linear or a curvilinear design. *Linear* arrays incorporate a rectangular array of elements, which allows for insonification with parallel US beams. Because of their high resolution, they are frequently used in vascular, orthopedic and rheumatological applications. *Curvilinear* or convex arrays feature elements aligned on a circular arc of a given radius and operate with US beams virtually originating from an origin point within the probe. Such a configuration increases the field of view as especially appreciated in abdominal imaging.

ii) The **number of scanlines** acquired in each frame either corresponds with the number of transducer elements (*sequential imaging*), or a sophisticated sequence of pulses allows to steer US beams to multiple directions using the same elements (*phased array*). The footprint of sequentially used arrays is larger than in phased arrays, which allow for imaging through narrow acoustic windows. In addition, elements aligned in a rectangle allow for a dynamic selection of the imaging plane or even 3D volumentric imaging. Common applications of phased arrays include cardiac and transcranial imaging.

iii) In terms of **anatomical access**, *extracorporeal* US probes are most commonly found in practice and allow for imaging the human body from outside. In miniaturized form,

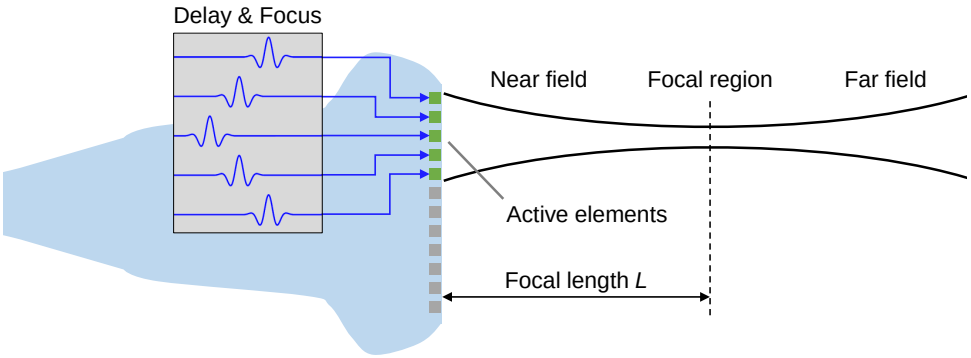

Fig. 1.3. Gaussian pulses, usually 2-3 half cycles long, are transmitted with a particular delay for each transducer element to focus the beam in the desired focal region. *Inspired by [95].*

and equipped with sterile wraps, similar designs are also used intra-operatively. In contrast, *trans-rectal* or *trans-vaginal* US systems utilize natural orifices to achieve a higher proximity to the region of interest. A special case is *intra-vascular* ultrasound (IVUS), where a small array is attached to the distal end of a catheter and remotely inserted into a blood vessel.

Beamforming

Regardless of the used mode of acquisition, US image formation always follows the same principle. To generate a pulse for insonification, a carrier signal with the desired frequency f is electronically convolved with a chosen selection function. Regularly, Gaussian functions according to the desired spatial pulse width are used for this purpose, controlling the axial resolution as required for a particular application. For the acquisition of one scanline, a set of active transducer elements is then selected, each transmitting the pulse into the tissue. As schematically shown in Fig. 1.3, small delays τ_i between different elements i allow for a focusing of the ultrasound energy:

$$\tau_i = \frac{L - \sqrt{\Delta s_i^2 + D^2}}{c} + \tau_0, \tag{1.29}$$

where L is the focal distance, Δs_i denotes the distance between element i and the central element, and τ_0 is a constant delay required by the electronic circuits to ensure that $\tau_i > 0$. This technique is called *time-delay focus* and optimizes the resolution in lateral direction. In fact, Eq. 1.29 simulates the effect of an acoustic lens compensating for quadratic diffraction properties [215]. Note that focusing is dependent on a chosen focal depth and that the optimal resolution is only achieved in a narrow region around it. In the case of phased arrays, additional time delays proportional to the distance to the central active element are used for beam steering, essentially implementing a lateral offset of the focal point. Modern US systems also apply a technique called *apodization* for an improved control of the beam properties and scale the pressure amplitude such that the central element transmits with the highest, and the outer-most ones with the least intensity. The interested reader is referred to [215] for a detailed mathematical discussion on advanced beamforming techniques.

B-Mode and M-Mode Image Formation

After transmission of the pulses, the transducer elements are used as receivers, listening to the reflected waves as a result of the tissue interactions described above. As stated in the previous section, the core concept of B-mode imaging is to encode an US sample's intensity as brightness.

Let $g(t)$ be the acquired raw signal. Note that as multiple elements are used to transmit the US beam, also $g(t)$ depends on readings from multiple elements, which are aligned (inverse principle for beamforming) and summed for further processing. Two different kinds of amplification form the first part of the signal processing pipeline:

$$h(t) = B_0 \, B_1(t) \, g(t). \tag{1.30}$$

Hereby, B_0 is a global gain, and $B_1(t)$ incorporates local amplification factors depending on the depth the signal originated from, i.e. the time of arrival at the transducer. The latter scaling is known as *time-gain-compensation* (TGC) and can be commonly found on the user interfaces of US scanners for step-wise adjustment of image gain, allowing to compensate for diminished contrast with increasing depth due to attenuation.

For demodulation and separation from the carrier signal with $\omega = 2\pi f$, it is helpful to represent $h(t)$ again in complex notation:

$$z(t) = h(t) + i\,\mathcal{H}\{h\}(t) = a(t)\, e^{i(\omega t + \varphi(t))}, \tag{1.31}$$

with amplitude $a(t)$ and phase $\varphi(t)$. $\mathcal{H}$ denotes the Hilbert transform

$$\mathcal{H}\{h\}(t) = \int_{-\infty}^{\infty} \frac{h(x)}{t-x} dx. \tag{1.32}$$

The absolute value of this signal is finally used for envelope detection and B-mode display, mathematically corresponding to the demodulation of the acquired signal:

$$A(t) = |z(t)|. \tag{1.33}$$

In the simplest scenario, this envelope signal $A(t)$ is then directly used for image generation, logarithmically compressing the signal to the available gray values of an image with N bit color depth:

$$I(t) = \frac{\ln A(t) - \ln \min_t A(t)}{\ln \max_t A(t) - \ln \min_t A(t)} \left(2^N - 1\right). \tag{1.34}$$

Note that the ratio in Eq. 1.34 essentially performs a windowing to the interval $[0, 1]$, and its denominator defines the dynamic range of the acquisition:

$$D = \ln \frac{\max_t A(t)}{\min_t A(t)}. \tag{1.35}$$

The resulting intensities I represent one individual scanline. In the case of a linear array, they can directly be used to populate the corresponding column's pixels in the final image.

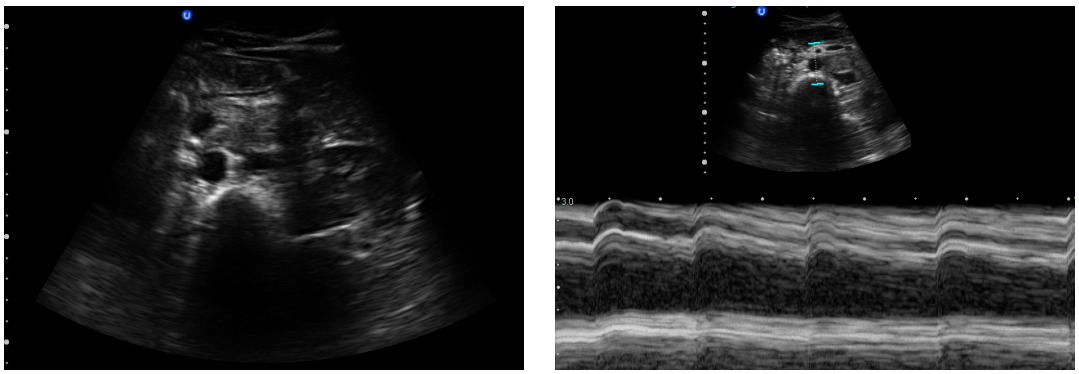

(a) B-mode **(b)** Combined B-/M-mode

Fig. 1.4. Exemplary brightness images of a healthy human abdomen (female, 28), acquired with a curvilinear probe and clearly visualizing the aorta. While B-mode frames assemble numerous scanlines in 2D, thus generating a spatial view, M-mode presents parts of a single scanline as selected by the sonographer (*blue bars*) over time and thus facilitates a temporal view.

In practice, image formation may be slightly more complex. First, the complex signals $z(t)$ might be useful to retain, for instance for speckle reduction [215]. Second, the envelope signal $A(t)$ is rarely fed to the logarithmic compression stage without further filtering, for instance to reduce the level of noise. Detailed algorithms and filter designs for these two additional steps are commonly considered critical knowledge by US system vendors and often remain confidential. Third, non-linear transducer geometries eventually require a so-called *scan conversion*, i.e. an interpolation from scanlines to pixels in order to generate rectangular 2D images. An exemplary B-mode image is depicted in Fig. 1.4a.

M-mode imaging is a related concept that follows the same brightness paradigm as in B-mode visualization. However, instead of assembling multiple scanlines corresponding to different spatial locations to a 2D image, one single scanline is followed *over time*. This facilitates the assessment of periodically moving structures, especially due to respiratory or pulsatile motion. For spatial reference, B-mode images are usually visualized as well, allowing the sonographer to interactively select the desired scanline for the temporal view. An exemplary M-mode visualization is shown in Fig. 1.4b.

Doppler Flow Image Formation

As already outlined in Sec. 1.3.1, different techniques to exploit the Doppler effect have been developed in the past decades. One distinguishes continuous wave (CW) and pulsed wave (PW) Doppler modes, whose principle of acquisition is quite different. In CW operation, on the one hand, the US system emits a continuous wave and directly measures the frequency shift Δf in spectrum of the reflections as a result of a moving scatterer, for instance blood particles, according to the well-known Doppler effect:

$$\Delta f = -\frac{2v}{c} f \cos\theta. \tag{1.36}$$

Hereby, θ is the Doppler angle, i.e. the angle between the US beam and the scatterer's direction of movement. Note that in contrast to Sec. 1.3.2, v is now defined as the velocity of the scatterer. Although there is no upper limit to the detectable frequency shifts, the biggest

disadvantage of CW signals is their lack of spatial reference. Because reflected frequency changes caused by moving scatterers in different depths are indistinguishable, CW modes only allow for cumulative observation along an entire scanline selected by the user. Similar to an M-mode visualization, modern scanners present the cumulative frequency or velocity spectrum over time.

On the other hand, PW operation neglects the frequency change and estimates the speed of a scatterer from its observed movement between short consecutive pulses. Usually, *ensembles* of eight to twelve such pulses are emitted for a single scanline with pulse repetition frequency f_{PRF}, which is regularly in the range of 5-20 kHz. This amounts to pulse repetition intervals $T_{PRF} = 1/f_{PRF}$ between 50 and 200 μs. If a scatterer has moved between two pulses, its traveled distance is consequently given by

$$\Delta x = v \cos\theta \, T_{PRF}, \tag{1.37}$$

which relates to the induced time delay Δt of the reflected pulse as follows:

$$\Delta t = \frac{2\Delta x}{c}. \tag{1.38}$$

Instead of directly estimating Δt, the phase shift between pulses is determined. We recall Eq. 1.31 on the complex notation of the received signal and assume that a reflected pulse is a scaled and phase-shifted version of an emitted one [58]:

$$z(t) = a(\xi(t - t_r - \Delta t)) \; e^{i\omega\xi(t-t_r)} \; e^{i\varphi}, \tag{1.39}$$

where $t_r = 2r/c$ is a pulse's round trip time to depth r and back, and $\xi = 1 + 2v\cos\theta/c$ a time compression factor according to Eq. 1.36. Note that while the signal amplitude is corrected for the combined effect of round trip duration and scatterer motion, the induced changed in phase change can be decoupled. The frequency of the Doppler phase function φ can finally be estimated as discrete derivative [58], for instance:

$$\Delta f = \frac{1}{2\pi} \frac{\varphi_k - \varphi_{k-1}}{T_{PRF}}, \tag{1.40}$$

where index k relates to the k-th pulse of an ensemble. To determine the sign of the detected velocity according to Eq. 1.36, in-phase and quadrature components of the complex signal $z(t)$ need to be inspected.

The pulse repetition frequency is of crucial importance, because it directly influences the interval between minimum and maximum possibly distinguished velocities:

$$[v_{min}, v_{max}] = \frac{c \, f_{PRF}}{f \, \cos\theta} \left[\frac{1}{2n}, \frac{1}{4}\right]. \tag{1.41}$$

Both of these limits are related to the Nyquist-Shannon sampling theorem. The limit on v_{min} can be intuitively understood since at least a scatterer movement by one period needs to be observed within a sequence of n pulses. The upper limit on v_{max} states that the scatterer must not move by more than a half period. If either of these criteria is violated, *aliasing* effects will misestimate the observed particle velocities by at least one full period, which might also invert the obtained reading. In practice, it is not always straightforward to find an appropriate

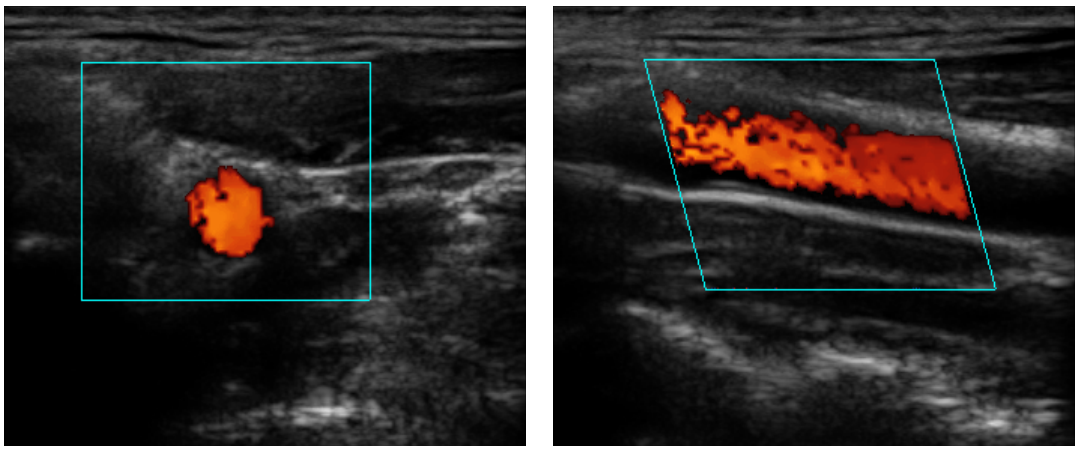

(a) Cross-sectional **(b)** Longitudinal

Fig. 1.5. Exemplary color-coded duplex images of a healthy human carotid artery (male, 28), acquired with a linear probe in duplex mode during systole. The *cyan rectangle* indicates the Doppler window, where flow measurements are acquired and superimposed (*red* indicates flow toward the transducer, *blue* would signify flow away from it). Note that for the longitudinal scan (**b**), beamsteering was used to reduced the angle between the vessel and the US beam.

balance between measurable absolute velocities (low f_{PRF}) and the resolution in terms of velocity (high f_{PRF}).

All considerations above concerning PW Doppler acquisition refer to one particular sample location. In modern US scanners, a similar visualization as in CW operation is offered to the sonographer, only that an additional gate position along the scanline can be defined. The presented frequency or velocity spectrum over time then refers to the flow through this gate, which can also be rotated and hence aligned with the direction of a vessel to compensate for the $\cos\theta$ terms.

However, the concept can easily be extended to sample an entire area, i.e. using several scanlines with multiple samples each similar to B-mode formation. The acquired velocity estimates can then be rendered using false colors, with a common convention of red referring to flow toward and blue away from the transducer (remember "BART": blue away, red toward). Several names for this imaging mode can be found in literature, including color-flow mapping (CFM) and color-Doppler imaging (CDI). For anatomical reference, CFM frames are almost always jointly visualized with B-mode information, which is commonly referred to as color-coded duplex (CCD) mode. To maintain an acceptable frame rate, Doppler sampling is regularly reduced to a small user-selected window. Exemplary duplex images are presented in Fig. 1.5.

1.3.4 3D Processing and Compounding

Not surprisingly, the benefit of 3D ultrasound data has been understood early, and 3D reconstructions, i.e. spatial compounding, were sporadically performed already in the 1980s [172]. One of the simplest forms of tracking a transducer in space is a mechanical arm with position sensors at all joints. Geiser et al. [75] used such a device to estimate left ventricular wall mo-

tion in 1982. Later, electromagnetic and optical tracking systems shaped the most commonly used techniques for so-called *freehand* 3D ultrasound [74, 203].

Even without compounding, tracked 2D US frames within a registered 3D environment can greatly support a variety of interventions. Luo et al. [257] demonstrated that both delivery and deployment of aortic stent-grafts is feasible using tracked ultrasound. In particular for cardiovascular applications, gating (synchronization) techniques facilitate acquisitions independent of pulse phase and enable 3D compounding for a specific stage in the cardiac cycle. Both ECG gating [27] and the usage of pulse-oximetry sensors [98] has been shown to be effective in this regard, for instance removing artifacts due to vessel pulsation. In the following, the necessary components to embed 2D US frames in a 3D environment are briefly discussed.

Types of 3D Localization

Three-dimensional ultrasound processing is a prerequisite for all contributions in this thesis and in general necessitates a form of 3D localization. This can be achieved in the following three ways:

i) By rigidly attaching a marker (*target*) of a tracking system to the US probe, the tracking stream can be used to relate multiple frames to each other. As mentioned above, optical and electromagnetic tracking systems are most widespread in US-based navigational and interventional applications, and mechanical arms continue to fill niches in certain applications [206]. The prostate biopsy system presented in chapter 2 relies on optical tracking of a trans-rectal convex probe, while the freehand acquisitions for 3D blood flow reconstructions (chapter 4) were performed using a linear probe and an electromagnetic tracking system.

ii) Similarly, active robotic manipulators can be utilized for the same purpose. Not only are robots regularly more accurate than conventional tracking systems, they also facilitate 3D acquisitions and guidance applications by means of automation. Robotic concepts are intensively studied in chapter 3, and the precision the robotic tracking also proved beneficial in chapter 4. A thorough overview of robotized US systems can be found in Sec. 3.2.3.

iii) Some transducers incorporate a motorized linear or curvilinear array, which allows for 3D acquisitions without moving the probe [61]. Because the motors are equipped with encoders, the geometric relation between subsequent frames, or even scanlines, can be inferred. In fact, the robot-attached transducer used in chapter 3 is a motorized curvilinear probe, allowing for 3D US acquisitions while keeping a needle aligned to a planned target position.

Ultrasound Calibration

Regardless of the type of 3D localization, a calibration is required to learn the transformation ${}^{f}\mathrm{T}_{t}$ between the US image coordinate system F_{f} and the frame F_{t} obtained from the tracking system, robot controller, or motor encoder, as shown in Fig. 1.6. Since a rigid attachment of

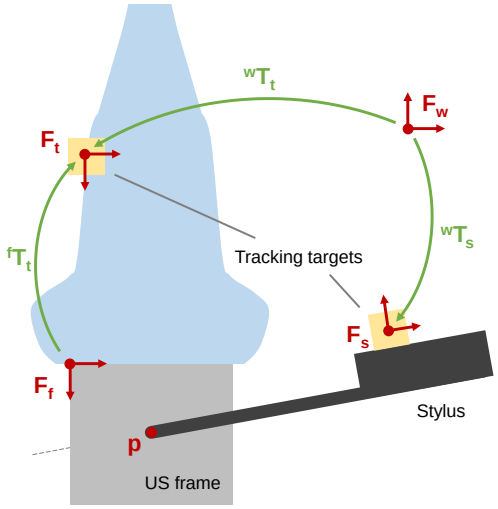

Fig. 1.6. Ultrasound calibration consists of estimating the transformation ${}^{f}\mathrm{T}_{t}$ between image frame F_{f} and tracking target (or alternatively robot end-effector) frame F_{t}. Only then, US frames can be embedded in some world coordinate frame F_{w}, which is often defined as tracking system origin or robot base. ${}^{f}\mathrm{T}_{t}$ can be estimated by collecting samples of corresponding coordinates of some points $\mathbf{p}$, for instance by moving the tip of a tracked stylus (${}^{w}\mathrm{T}_{s}$) to various positions in the US image.

the tracking target onto the probe is assumed, ${}^{f}\mathrm{T}_{t}$ can be written in homogeneous notation as:

$$ {}^{f}\mathrm{T}_{t} = \begin{bmatrix} \mathrm{R} & \mathbf{t} \\ \mathbf{0}^{\intercal} & 1 \end{bmatrix} \cdot \begin{bmatrix} s_u & 0 & 0 & 0 \\ 0 & s_v & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{1.42} $$

where the rigid transformation with the 3×3 rotation matrix R and the translation vector $\mathbf{t}$ constitutes six degrees of freedom, and factors s_u and s_v scale from the US image space in pixels to real distances, e.g. millimeters. Together, this amounts to eight unknown parameters in the general case. In practice, however, s_u and s_v can often be either directly obtained from the US scanner interface in open access systems or inferred based on the configured imaging depth and the pixel resolution of the acquired B-mode images. Hence, only rotations and translations are considered for the calibration procedure in many algorithms [102].

To solve for the unknowns, a set of N constraints are usually collected in an error function of the form

$$ f = \sum_{i=1}^{N} \| \mathbf{p}_i^{w} - {}^{w}\mathrm{T}_{t}{}^{-1} \cdot {}^{f}\mathrm{T}_{t} \cdot \mathbf{q}_i^{f} \|, \tag{1.43} $$

where $\mathbf{p}_i^{w}$ and $\mathbf{q}_i^{f}$ both refer to the same, i-th point represented in world (F_{w}) and US image space (F_{f}), respectively. Note that the tracking information ${}^{w}\mathrm{T}_{t}$ usually differs for each correspondence i as well. In their extensive review, Hsu et al. [102] summarize different manual and semi-automatic methods to compile a list of such constraints in freehand scenarios.

In the simplest case, a pointing stylus is also equipped with a tracking target as illustrated in Fig. 1.6. Its own calibration, denoted ${}^{0}\mathrm{T}_{s}$, i.e. the Euclidean offset from the tip to its own tracking target frame F_{s} can easily be found with a so-called pivot calibration [102]. Thereafter, as proposed by Muratore et al. [146], the stylus is positioned in such a way that the tip is visible in the B-mode image, which will require some back-and-forth motion to ensure that the point visible in the image is not an arbitrary point on the shaft. The location of the tip can then be manually selected or automatically segmented in the US frame ($\mathbf{q}^{f}$), and its location in world space is given by:

$$\mathbf{p}^{w} = {}^{w}\mathrm{T}_{s}{}^{-1} \cdot {}^{0}\mathrm{T}_{s} \cdot [0,0,0,1]^{\intercal}. \tag{1.44}$$

The error function of Eq. 1.43 is finally minimized in an iterative fashion, for instance using a Levenberg-Marquardt solver. A similar technique is used in chapter 2 for spatial calibration of a trans-rectal US probe.

A detailed overview of related calibration techniques, including solid and (N-)wire phantoms, is provided in [102]. In recent years, novel approaches have been proposed to facilitate spatial calibration. Incorporating prior knowledge of a wire phantom, Boctor et al. [20] presented a closed form solution for the transformation ${}^{f}\mathrm{T}_{t}$. Further considerations yielded an automated segmentation in the US image [181], rendered expensively 3D-printed phantoms obsolete [201], and exploited robotic manipulators for improved automated sampling [121].

Finally, Wein et al. [237] have suggested a technique entirely based on image registration, which is also employed in this thesis. Instead of mapping between US frames and the world coordinate system, various US frames are directly related with each other. If points $\mathbf{q}^{f}$ in image space correspond between frames i and j, the following relation holds:

$$\mathbf{q}_{j}^{f} = {}^{f}\mathrm{T}_{t}{}^{-1} \cdot {}^{w}\mathrm{T}_{t}(j) \cdot {}^{w}\mathrm{T}_{t}{}^{-1}(i) \cdot {}^{f}\mathrm{T}_{t} \cdot \mathbf{q}_{i}^{f}. \tag{1.45}$$

To estimate the components of ${}^{f}\mathrm{T}_{t}$, two roughly orthogonal sweeps A and B of the same region are acquired, yielding a plethora of potential correspondences. Using 3D compounding as summarized below and image registration as described in Sec. 1.4, the goodness of fit can then be directly measured and optimized using an image similarity metric S:

$${}^{f}\mathrm{T}_{t} = \arg\max S(A,B). \tag{1.46}$$

Compounding Strategies

Unless a matrix probe is available, compounding techniques are necessary to reconstruct 3D US volumes from multiple 2D frame. Only in very specific cases, where all acquired frames are parallel and equidistant, straightforward stacking of pixels is possible similar to the slices in tomographic modalities. Otherwise, an interpolation strategy has to be chosen to handle the mapping between differently oriented and scaled image grids.

Following the classification of Solberg et al. [203], direct compounding methods can be categorized in *forward* and *backward* warping techniques, as illustrated in Fig. 1.7. **Forward warping** methods are pixel-based, as the individual samples of input US frames are traversed and projected onto the grid of the output volume and their intensity is inserted at the obtained

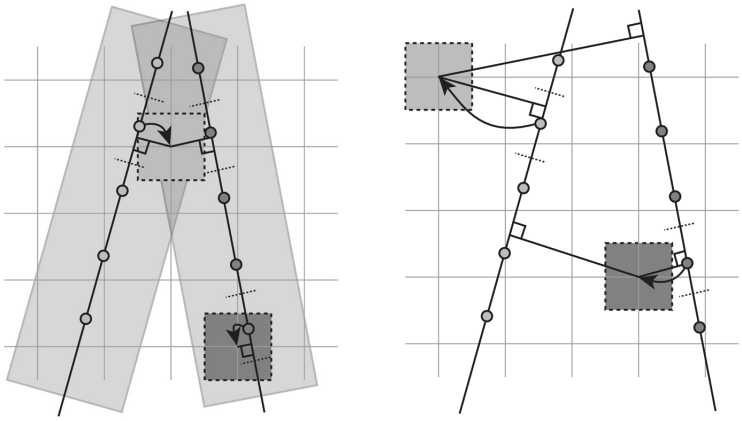

(a) Forward warping (b) Backward warping

Fig. 1.7. **a)** Forward warping methods traverse all samples and project them into the target grid. **b)** Backward warping approaches traverse the target grid instead, and search for appropriate samples to populate it with. *Reprint from [203], with kind permission of Elsevier.*

locations. As shown in Fig. 1.7a, multiple pixels might reach the same voxel. Various strategies, including minimum distance selection, have been proposed to solve this ambiguity. While forward warping is computationally cheap, holes due to a lack of samples at every single voxel are quite common [203].

In contrast, **backward warping** techniques, also called voxel-based methods, traverse the grid of the output volume and search for respective pixels of US frames to populate it with. As this might also include pixels farther away than the voxel spacing, as shown in Fig. 1.7b, the occurrence of holes is reduced. Since again multiple pixels might be in the vicinity, a weighted sum based on Euclidean distance is commonly employed.

Although it has been shown that superior reconstruction qualities can be achieved with backward warping [239], the advantage of pixel-based methods is that processing and visualization of the 3D volume being acquired can start immediately with the first US frame. Depending on the application, fast hybrid approaches as presented by Karamalis et al. [113] might constitute an acceptable trade-off between interactivity and image quality. Another frequent limitation of compounding schemes is the strong dependence of US intensities on the insonification direction. In sweeps with overlapping frames acquired from different angles, naive implementations are prone to yield significant inhomogeneity artifacts. To address these issues, techniques have been developed to either compound directional clusters independently first and then fuse their potentially contradictory information [193], or to model the direction dependence from the very beginning in the form of a tensor field [96].

1.4 Ultrasound Image Registration

Image registration generally refers to the process of aligning two or more images, establishing a common coordinate system and thus achieving correspondence between the information contained within the images. For many modern applications of computer-assisted diagnosis

and therapy, it has evolved into a key technology and continues to be actively investigated in research [183].

The objective of image registration can be achieved in two distinct ways, either by extracting and aligning a set of features from the input images (*feature-based registration*), or by employing the image intensities directly (*intensity-based registration*). Both of these approaches are briefly summarized in the following sections to introduce the fundamental concepts which the contributions of this thesis are based upon. In doing so, they focus on the case that at least one of the input datasets is an ultrasound image.

1.4.1 Overview of Feature-based Registration

In the simplest scenario, sets of corresponding points are known, one for each image. For instance, both the manually annotated landmarks points and mesh vertices of automatically generated segmentations are considered such features in chapter 2. Let $\mathrm{P} = [\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n]$ and $\mathrm{Q} = [\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_n]$ again be sets of corresponding 3D points. An optimal rigid transformation with 3×3 rotation matrix R^* and translation $\mathbf{t}^*$ minimizes the quadratic projection error:

$$(\mathrm{R}^*, \mathbf{t}^*) = \arg\min \sum_{i=1}^{n} (\mathrm{R}\mathbf{p}_i + \mathbf{t} - \mathbf{q}_i)^2. \tag{1.47}$$

Umeyama [226] proposed an elegant and widely used method to uniquely solve for R and $\mathbf{t}$. First the point sets are demeaned, i.e. points $\mathbf{p}'_i = \mathbf{p}_i - \overline{\mathbf{p}}$, with $\overline{\mathbf{p}}$ denoting the mean of point set P, are collected in P', and likewise for Q'. Next, a singular value decomposition (SVD) of the covariance matrix is computed:

$$\mathrm{U}\mathrm{D}\mathrm{V}^\intercal = \Sigma_{PQ} = \frac{1}{n}\mathrm{Q}'\mathrm{P}'^\intercal. \tag{1.48}$$

Assuming that $\det \Sigma_{PQ} > 0$, the final transformation is then given as:

$$\mathrm{R}^* = \mathrm{U}\mathrm{V}^\intercal, \qquad \mathbf{t}^* = \overline{\mathbf{q}} - \mathrm{R}^*\overline{\mathbf{p}}. \tag{1.49}$$

However, correspondence is not known a-priori in many applications, for instance because the features are automatically detected based on edges visible in the image. Also point clouds, e.g. as acquired with RGB-D cameras, are unstructured. The **Iterative Closest Point** (ICP) algorithm was presented by Besl and McKay [16] to tackle this problem. Essentially, the method alternates between an estimation of correspondence in a nearest-neighbor fashion given the sets' current positions, and an update of these positions according to Eq. 1.49. With a sufficiently accurate initialization, the algorithm iteratively converges to an optimal transformation.

A plethora of varieties have been proposed to improve ICP with respect to accuracy and performance, including better suited error metrics such as point-to-plane, stable point sampling, efficient correspondence search using binary trees, anisotropic noise models, and multi-resolution schemes [167]. Unlike ICP, where correspondence is defined in a binary way, Robust Point Matching (RPM) [82] uses a normalized exponential function to define "soft" correspondence.

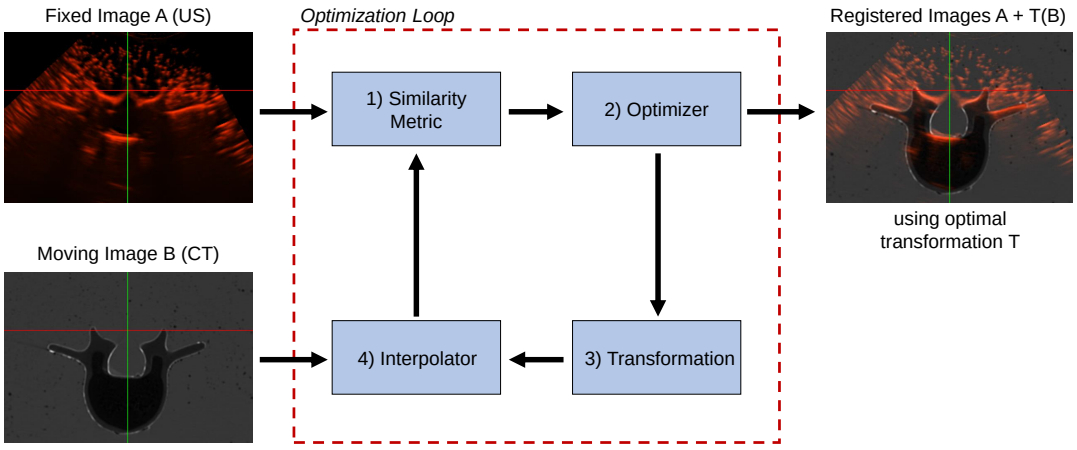

Fig. 1.8. A general image registration framework consists of four essential components: 1) A similarity metric quantifies how well two images are aligned, which is used by 2) an optimizer that tunes the parameters of a 3) transformation model. Finally, 4) an interpolator is required to obtain a representation of the moving image in the fixed image grid for comparison.

In many applications, the estimated transformation is not necessarily rigid but rather affine or deformable. Despite the development of non-rigid ICP variants [167], probabilistic approaches have been successfully adopted to the area of registration and are most commonly used today. Myronenko and Song [148] presented the **Coherent Point Drift** algorithm, which is based on Gaussian mixture models and enforces a smooth (coherent) movement between points. The method is agnostic with respect to the used transformation model, but is originally intended to cover the non-rigid case, with rigid or affine transformation simply incorporating additional constraints. Because the contribution of fast surface-based prostate registration is based on the CPD algorithm, a detailed description can be found in Sec. 2.3.3.

1.4.2 General Intensity-based Registration Framework

The advantage of intensity-based registration algorithms is that no feature extraction is necessary and the image intensities are directly used in the iterative optimization scheme. A general intensity-based registration framework as depicted in Fig. 1.8 aims at finding an optimal transformation T^* to align the fixed reference image A with the moving image B:

$$T^* = \arg\max_T S\left[\mathrm{A}, T(\mathrm{B})\right] \tag{1.50}$$

At least four essential components are required for an iterative optimization:

1) The similarity metric S quantifies how well the fixed image A and a transformed version of the moving image $T(\mathrm{B})$ are aligned with each other. An overview of commonly used similarity metrics is provided in Sec. 1.4.3 below, including metrics dedicated for the registration of US images with CT or MRI datasets.

2) The optimizer considers the output of S to iteratively find a set of transformation parameters that maximizes the similarity. If derivatives $\partial S/\partial T$ can be easily obtained or at least numerically approximated, gradient-based methods such as conjugate gradient,

Gauss-Newton, or Levenberg-Marquardt as a combination of these two can be employed. The most common gradient-free methods are heuristic, for instance Nelder-Mead and NEWOA [180]. If the similarity metric can be assumed to locally resemble a multi-variate quadratic function, BOBYQA [169] may constitute an acceptable trade-off between speed and accuracy [70].

3) The transformation model determines the number of degrees of freedom to optimize. As previously stated, there are six degrees of freedom in the rigid case, i.e.

$$\mathbf{x}' = \begin{bmatrix} \mathrm{R} & \mathbf{t} \\ \mathbf{0}^\intercal & 1 \end{bmatrix} \mathbf{x}, \tag{1.51}$$

where the rotation R is regularly parametrized with Euler angles to allow for optimization. Affine transformations additionally include three scaling parameters. A broad variety of methods can be found in the literature to capture non-linear deformation. Apart from physically expressed models, which are defined by partial differential equations, interpolation schemes based on basis functions are frequently used to limit the amount of parameters to a manageable range. While approaches like free-form deformations (FFD) employ a set of control points being translated throughout optimization, more advanced, statistical techniques incorporate prior knowledge [207]. As an example, the eigenvalues of a statistical deformation model are optimized in Sec. 2.3.5 to register prostate MRI and US images.

4) Finally, an interpolation method is required so that $S(\mathrm{A}, T(\mathrm{B}))$ can be numerically evaluated because the voxel positions of $T(\mathrm{B})$ generally do not coincide with those of image A. Computationally the least expensive approach, denoted nearest-neighbor interpolation, is achieved by simply rounding the coordinates to the next integer values followed by a mere look-up. Better results are achieved using linear or higher order polynomial interpolation, which eliminates most aliasing artifacts [113]. Note that since most similarity metrics can be easily efficiently evaluated on Graphical Processing Units (GPUs), shader implementations often directly exploit built-in interpolation capabilities without the need to explicitly generate $\mathrm{B}' = T(\mathrm{B})$ in memory.

A detailed, general review on intensity-based medical image registration can be found in [183]. For the application to the challenging prostate anatomy, a thorough overview of the contemporary literature is provided in Sec. 2.2.2.

1.4.3 Similarity Metrics for Intensity-based Registration

The purpose of a similarity metric $S(\mathrm{A}, \mathrm{B}')$ is to quantify how well a fixed image A and a transformed moving image $\mathrm{B}' = T(\mathrm{B})$ correspond in the overlapping area. This is commonly achieved by comparing individual corresponding pixels or voxels of the two images using a suitable metric f, and summarizing the result by integration over the entire overlapping image domain Ω:

$$S(\mathrm{A}, \mathrm{B}') = \frac{1}{|\Omega|} \int_\Omega f(\mathbf{x}) \, d\mathbf{x}. \tag{1.52}$$

In the simplest case, f is defined as the squared intensity difference, i.e.

$$f_{\text{SSD}}(\mathbf{x}) = \left(\mathrm{A}(\mathbf{x}) - \mathrm{B}'(\mathbf{x})\right)^2, \tag{1.53}$$

resulting in the **Sum of Squared Differences** (SSD) similarity metric. Because it is sensitive to large intensity differences, for instance for image pairs with and without contrast agent, the **Sum of Absolute Differences** (SAD) is also commonly used:

$$f_{\text{SAD}}(\mathbf{x}) = |\mathrm{A}(\mathbf{x}) - \mathrm{B}'(\mathbf{x})|, \tag{1.54}$$

The metrics defined above assume that corresponding structures in both images share identical intensity values, which is regularly not the case. Even if both images are derived from the same modality, device settings can result in different illumination. The commonly used **Normalized Cross-Correlation** (NCC) similarity metric compensates for different illumination levels using the mean μ and standard deviation σ of both images' intensities:

$$f_{\text{NCC}}(\mathbf{x}) = \frac{(\mathrm{A}(\mathbf{x}) - \mu_A)(\mathrm{B}'(\mathbf{x}) - \mu_B)}{\sigma_A \sigma_B}. \tag{1.55}$$

Yet, NCC assumes a positive linear correlation between intensities in both images, i.e. that bright structures in one image will corresponding to (relatively) bright ones in the other image, and that the relative scale between both is constant throughout the entire intensity space. In the case of multi-modal registration problems, for instance in CT/MRI or MRI/ultrasound combinations, these assumptions generally do not hold. Therefore, a great variety of similarity metrics have been specifically proposed to cope with such non-linear intensity correlations, ranging from hybrid intensity/gradient metrics to information theoretic approaches relying on joint intensity histograms. Because a thorough analysis of the available metrics would by far exceed the scope of this thesis, the most commonly employed ones shall be briefly introduced.

Among the pixel-wise multi-modal similarity metrics is **Linear Correlation of Linear Combination** (LC^2) [70, 238], which was specifically developed for MRI/ultrasound registration and is regularly used in this thesis (see Sec. 2.3.5, Sec. 3.3.5). In short, LC^2 correlates ultrasound intensities of a fixed image A_{US} with a linear combination of moving MR image intensities (B_{MR}), their gradients, and a constant:

$$z(\mathbf{x}) = \alpha \mathrm{B}'_{\text{MR}}(\mathbf{x}) + \beta \left\|\nabla \mathrm{B}'_{\text{MR}}(\mathbf{x})\right\| + \gamma. \tag{1.56}$$

This combination is motivated by the physical principles of ultrasound image generation. Intensities may correlate well with other modalities in rather homogeneous areas, for instance within the liver, but at sharp edges in CT or MRI, i.e. areas with a high gradient, strong ultrasound reflections will produce high contrast. As a result, US intensities may locally either correlate to MR intensities or their gradients, suggesting a patch-based consideration. Note that the subscripts *MR* and *US* in Eq. 1.56ff indicate that the images are not interchangeable.

For a small neighborhood $\Phi(\mathbf{x})$ centered around a voxel $\mathbf{x}$, coefficients $\mathbf{c} = (\alpha, \beta, \gamma)$ are determined by minimizing the difference between A_{US} and the linear combination z in a least squares sense:

$$\hat{\mathbf{c}} = (\mathrm{M}^\intercal \mathrm{M})^{-1} \mathrm{M}^\intercal \mathrm{U}, \tag{1.57}$$

with

$$\mathrm{M} = \begin{pmatrix} \mathrm{B}'_{\mathrm{MR}}(\mathbf{x}_1) & \|\nabla \mathrm{B}'_{\mathrm{MR}}(\mathbf{x}_1)\| & 1 \\ \vdots & \vdots & \vdots \\ \mathrm{B}'_{\mathrm{MR}}(\mathbf{x}_m) & \|\nabla \mathrm{B}'_{\mathrm{MR}}(\mathbf{x}_m)\| & 1 \end{pmatrix}, \tag{1.58}$$

$$\mathrm{U} = \begin{pmatrix} \mathrm{A}_{\mathrm{US}}(\mathbf{x}_1) \\ \vdots \\ \mathrm{A}_{\mathrm{US}}(\mathbf{x}_m) \end{pmatrix}, \tag{1.59}$$

where $m = |\Phi(\mathbf{x})|$ is the number of voxels in the neighborhood. The local similarity function is then defined as

$$f_{\mathrm{LC}^2}(\mathbf{x}) = 1 - \frac{\sum_{x_i \in \Phi(\mathbf{x})} \left[\mathrm{A}_{\mathrm{US}}(\mathbf{x}_i) - z(\mathbf{x}_i)\right]^2}{|\Phi(\mathbf{x})| \ \mathrm{var}(\mathrm{A}_{\mathrm{US}}(\Phi(\mathbf{x})))}. \tag{1.60}$$

A common alternative to LC^2 is the **Modality Independent Neighborhood Descriptor** (MIND) as proposed by Heinrich et al. [94]. The approach reduces multi-modal registration to the mono-modal case by defining an intermediate representation computed from the input images, followed by registration using SSD. For each image I, neighborhood descriptors $\mathbf{L}_{\mathrm{MIND}}$ can be computed as follows:

$$L_{\mathrm{MIND}}(\mathrm{I}, \mathbf{x}, \mathbf{r}) = \frac{1}{n} \exp\left(-\frac{D_p(\mathrm{I}, \mathbf{x}, \mathbf{x}+\mathbf{r})}{\mathrm{var}(\mathrm{I}(\mathbf{x}))}\right), \tag{1.61}$$

where vectors r are within some search region R (dense sampling, sampling every 45°, or six-neighborhood), so that every voxel $\mathbf{x}$ is represented by a $|R|$-dimensional vector $\mathbf{L}$. Distances D_p evaluate a patch-based self-similarity

$$D_p(\mathrm{I}, \mathbf{x}_1, \mathbf{x}_2) = \sum_{\mathbf{p} \in \Phi} \left(I(\mathbf{x}_1 + \mathbf{p}) - I(\mathbf{x}_2 + \mathbf{p})\right)^2, \tag{1.62}$$

with patches Φ of size $(2p+1)$ in each dimension and centered at $\mathbf{x}_1$ and $\mathbf{x}_2$. Eventually, the similarity metric is then defined as

$$f_{\mathrm{MIND}}(\mathbf{x}) = \frac{1}{|R|} \sum_{\mathbf{r} \in R} |L_{\mathrm{MIND}}(\mathrm{A}, \mathbf{x}, \mathbf{r}) - L_{\mathrm{MIND}}(\mathrm{B}, \mathbf{x}, \mathbf{r})|. \tag{1.63}$$

Finally, there are similarity measures not directly relying on pixel/voxel information but establishing a statistic relation between sets of images instead. **Mutual information** (MI) was

introduced to image registration by Viola and Wells [229] and is based on the information theory formulation of Shannon entropy H:

$$H(\mathrm{I}) = -\sum_i p_{\mathrm{I}}(i) \log p_{\mathrm{I}}(i), \tag{1.64}$$

where $p_{\mathrm{I}}(i)$, estimated from a histogram, denotes the probability distribution function measuring the probability of a voxel of image I having the intensity i. The concept can be extended to sets of images using the joint probability distribution function $p_{\mathrm{AB}}(i, j)$. After minimizing the joint entropy between two images, i.e

$$H(\mathrm{A}, \mathrm{B}) = -\sum_i \sum_j p_{\mathrm{AB}}(i, j) \log p_{\mathrm{AB}}(i, j), \tag{1.65}$$

an image can be assumed to be optimally aligned. To make the method more robust, mutual information defined as follows is maximized instead:

$$S_{\mathrm{MI}}(\mathrm{A}, \mathrm{B}') = H(\mathrm{A}) + H(\mathrm{B}') - H(\mathrm{A}, \mathrm{B}') = \sum_i \sum_j p_{\mathrm{AB}'}(i, j) \log \frac{p_{\mathrm{AB}'}(i, j)}{p_{\mathrm{A}}(i) p_{\mathrm{B}'}(j)}. \tag{1.66}$$

Various extensions and modifications have since been proposed, including normalized MI to cope with varying overlap, and conditional MI, which encodes more spatial information than the otherwise globally acting MI method [94]. Due to the very different nature of image generation in US compared to tomographic modalities, and the accompanied mutually exclusive visibility of certain structures, MI-based approaches have been shown to perform poorly in multi-modal scenarios with ultrasound [70] and are therefore not used in this thesis.

Part II

Contributions

2 Multi-Modal Prostate Biopsy Guidance

2.1 Overview and Publications

This chapter presents my contributions toward multi-modal prostate biopsy guidance based on trans-rectal ultrasound (TRUS) imaging. Ultimately, the goal of the methods presented in this chapter is to support urologists in performing targeted TRUS-based biopsies by accurately augmenting the live TRUS image with additional modalities such as MRI and PET. While the main challenge of this task is to perform a fast, precise, and deformable image registration between TRUS and MRI, auxiliary work is necessary for ultrasound acquisition, 3D compounding, segmentation, and deformation model generation.

After outlining the clinical background on prostate cancer, its diagnosis and prior art on image-guided biopsies (Sec. 2.2), this chapter describes in Sec. 2.3 the developed methodology for multi-modal fusion biopsy guidance. It starts by presenting how 3D TRUS volumes are acquired once the patient is positioned (Sec. 2.3.1). Two distinct, fully automatic methods have been developed to perform deformable image registration between TRUS and MRI prostate images:

(1) The **surface-based registration** method as explained in Sec. 2.3.3 utilizes a modified version of the Coherent Point Drift algorithm [148] to elastically align surface meshes of prostate segmentations in both modalities. While MRI segmentations can be performed in advance, the method relies on a fast, automatic TRUS segmentation algorithm, which is based on Hough forests (Sec. 2.3.2).

(2) The **intensity-based registration** method as explained in Sec. 2.3.5 is based on a Statistical Deformation Model derived using the automatic segmentations from (A), see Sec. 2.3.4. A novel, lesion-specific, anisotropic preconditioning of the LC^2 similarity metric [70] is a key feature, as it emphasizes correct registration outcomes at crucial locations along expected directions.

Finally, once a registration is established, multi-modal biopsy guidance as described in Sec. 2.3.6 supports the urologist in targeting the identified lesions. Materials and experiments are described in Sec. 2.4. Each component of the presented guidance system is individually validated in Sec. 2.5, and thoroughly discussed in Sec. 2.6.

Substantial parts of this chapter have already been published and are quoted verbatim. Using method (1) as registration algorithm, methodology and results of the entire pipeline from

TRUS acquisition to biopsy guidance have been published in the following article (especially sections 2.2.2, 2.3.1-3, 2.3.6, 2.4.1-2, 2.5.1-3, 2.5.6, and 2.6):

[256] **Oliver Zettinig**, Amit Shah, Christoph Hennersperger, Matthias Eiber, Christine Kroll, Hubert Kübler, Tobias Maurer, Fausto Milletarì, Julia Rackerseder, Christian Schulte zu Berge, Enno Storz, Benjamin Frisch, and Nassir Navab. "Multimodal image-guided prostate fusion biopsy based on automatic deformable registration". In: *International Journal of Computer Assisted Radiology and Surgery* 10.12 (Dec. 2015), pp. 1997–2007. DOI: 10.1007/s11548-015-1233-y

All relevant additions in terms of method (2) are to be published in the following article (especially sections 2.2.2, 2.3.4-5, 2.5.4-5, and 2.6):

[255] **Oliver Zettinig**, Julia Rackerseder, Beatrice Lentes, Tobias Maurer, Kay Westenfelder, Matthias Eiber, Benjamin Frisch, and Nassir Navab. "Preconditioned intensity-based prostate registration using statistical deformation models". In: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*. IEEE, Apr. 2017, pp. 853–857. DOI: 10.1109/ISBI.2017.7950651

2.2 Introduction

2.2.1 Clinical Background

Prostate cancer, i.e. carcinoma of the prostate (PCa), is the leading cause of cancer in male adults, with an estimated 21% of newly detected cancer cases contributing to this diagnosis [199]. However, only around a third of these carcinomas are responsible for death, so that prostate cancer currently contributes to approximately 8% of cancer-related deaths in the USA. Numbers in Western Europe have shown to be similar [143]. The probability of developing invasive PCa is marginal in men under 50 years of age (0.3%) and increases steadily with age to a total of 14% over the entire lifetime.

Early detection, grading and staging of PCa is vital for risk stratification and therapy selection. The European Association of Urology (EAU) first suggests in its guidelines (Update March 2015, [143]) early PSA tests for men over 50 years of age, men over 45 years of age with PCa family history, and African-Americans. Hereby, PSA refers to the prostate-specific antigen, a protein produced by prostate cells, which blood serum levels are measured in ng/ml. Although PSA is organ-, but not cancer-specific and can also be elevated in benign prostatic hypertrophy or prostatitis, it is a better PCa predictor than digital rectal examinations.

The current gold standard for the diagnosis of prostate cancer is a systematic 10 to 12 core random biopsy under trans-rectal ultrasound (TRUS) guidance [18]. Hereby, the term *random*

Fig. 2.1. Example of a TRUS image of the prostate. As usual, the image allows anatomical guidance but shows poor contrast within the gland.

refers to the invisibility of the majority of cancer lesions in the TRUS images. An exemplary TRUS image of a patient with confirmed prostate cancer is shown in Fig. 2.1, illustrating that while ultrasound can be used for anatomical guidance, targeted selection of biopsy sampling locations is not possible. To tackle this blindness to lesions, schemes have been developed to systematically sample a standardized set of locations throughout the gland, regularly focusing on the peripheral zone where the prevalence is increased [63]. Nevertheless, TRUS continues to exhibit a poor detection rate of suspicious areas, leading to a high rate of false negative results [224]. It is not uncommon for patients with elevated PSA blood levels to undergo several all-negative biopsy sessions until eventually a positive sample is found that can be used for staging and therapy selection.

Among other systems, the Gleason system has been widely adapted for histopathological grading of biopsy samples [55]. Figure 2.2 shows the patterns that are used to classify the biopsy specimen. A primary grade is assigned to the dominant pattern of the tumor. Secondary and – if applicable – tertiary grades describe the next-most frequent patterns. The final Gleason score is reported as a sum of primary and secondary (or tertiary, if higher) grades, e.g. 2+1=3. This elaborate way of reporting is essential, because 3+4=7 cancers exhibit a better prognosis than the 4+3=7 ones.

Depending on the clinical findings, including tumor staging using the well-established TNM system (tumor, lymph nodes, metastasis), one of the following treatment options is selected [143]: Watchful waiting / active surveillance, radical prostatectomy, radiotherapy, cryotherapy, high-focused ultrasound ablation, or androgen suppression.

2.2.2 Image Registration for Prostate Cancer Diagnosis

To tackle the above identified challenges associated with purely TRUS-guided biopsies, multi-modal approaches with additional, pre-interventional imaging data have been adopted in clinical practice. Recent studies report that combining TRUS-guided biopsy workflows with multi-parametric magnet resonance imaging (MRI) is more accurate in detecting cancerous lesions with a significant Gleason score than the exclusive usage of ultrasound [71, 132,

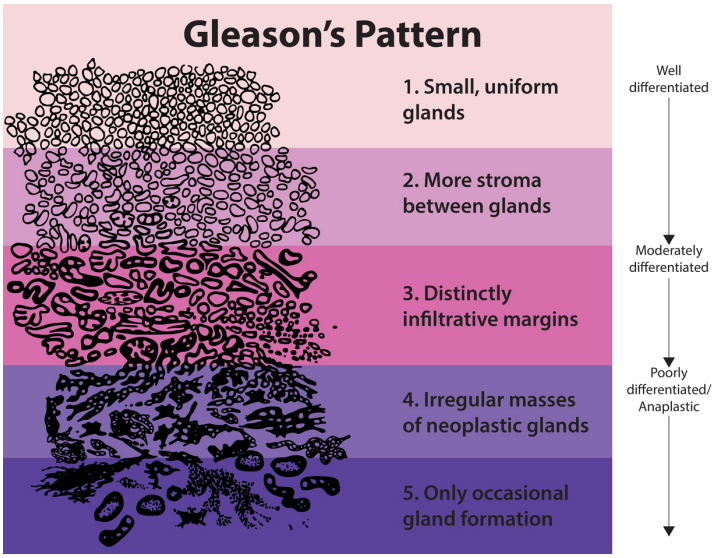

Fig. 2.2. Gleason's cell patterns for prostate grading. *From NIH (*`https://training.seer.cancer.gov`*), public domain.*

154]. Because the excellent anatomical soft-tissue contrast MRI is regularly not sufficient to identify suspicious lesions, fusion with functional imaging data, in particular positron emission tomography (PET), has been investigated. PET/TRUS fusion for prostate biopsy guidance has initially generated only moderate interest due to the low specificity of available radiotracers like ^{11}C-acetate, ^{11}C-choline and ^{18}F-FDG [224]. With the introduction of ^{68}Ga labeled ligands of Prostate-Specific Membrane Antigen (PSMA), exhibiting almost exclusive expression in the prostate and increased expression in prostate cancer [53], PET/TRUS and PET/MRI/TRUS fusion however gains increasing attention [134, 135]. Its higher cost is leveraged by the possibility to avoid serial biopsies with unclear outcome. Prior to our work [256], the usage of PSMA-PET/MRI has not yet been implemented into a TRUS-based biopsy guidance system.

Recent studies have shown that temporally enhanced ultrasound can improve the detection of tissue-specific prostate cancer patterns [151]. Applying machine learning techniques to time-series of ultrasound data allows differentiation between benign and malignant tissues [10]. While such approaches may greatly benefit tissue classification directly in TRUS, they may not be able to render additional modalities such as MRI and PET obsolete in the near future due to the underlying shortcomings of ultrasound itself.

Main challenges for the development of a multi-modal fusion image-guided prostate biopsy framework are time and space constraints during the procedure. In particular, cognitive fusion of these imaging modalities, i.e. the exclusively mental alignment of images presented side-by-side, is error prone, highly depends on the ability of the urologist to interpret MRI or PET images, and thus remains marginally useful [49]. Therefore, computer-aided approaches are of great interest. Recent studies suggest that overall, accurate MRI/TRUS registration outperforms cognitive fusion [198]. Even if such cognitive alignment performed by experts with great experience might lead to comparable results, the introduction of device-based fusion systems by general urologists is beneficial in terms of learning curve, "making the results quoted by experts more achievable", as Sonn et al. pointed out [206]. This trend

drives the ongoing development of image fusion-guided biopsy systems. Algorithms for image segmentation and registration need, however, to be fast and sufficiently accurate to be applicable in daily clinical routine.

Currently, a variety of commercial solutions on the market offer fusion biopsy functionality to some degree, employing different tracking systems to determine the position of their US transducer in space and registration algorithms to perform manual or semi-automatic rigid or elastic registration between 3D TRUS and MRI. They are usually rather bulky, require several lengthy interactions by the urologist – in particular for the segmentation of the prostate – and are mostly limited in their overall accuracy. The currently available commercial systems offering MRI/TRUS fusion, are summarized by Sperling et al. [209] in an excellent review.

The actual registration itself has been the topic of many prior studies. Recalling from Sec. 1.4, existing approaches to perform registration between 3D TRUS and MRI can be mainly distinguished along the following two dimensions:

i) **Domain**: **Point-based** (fiducial-based) or **surface-based** (mesh-based) algorithms rely on geometrical features either annotated manually (e.g. fiducials, boundary delineation) or using a feature extraction or segmentation algorithm. Even though they inherently neglect inhomogeneous deformations within the prostate gland and yield linear or spline-interpolated image alignment, they are predominantly used in approved clinical systems. In contrast, **intensity-based** approaches employ a similarity metric, which measures the pixel/voxel-wise quality of an alignment over the entire image domain, to find an optimal transformation. Because of the challenging anatomy and the multi-modal nature of the problem, they highly depend on good initialization and a proper transformation model.

ii) **Transformation model**: While a linear **rigid** or **affine** registration constitutes only few degrees of freedom (easier to solve), it cannot capture the deformations evident in the change of pose between MRI and TRUS acquisition and the deflections induced by the ultrasound transducer. **Elastic** registration can be modeled in various ways, including free-form deformations (FFD) and statistical/prior-based approaches [207].

In terms of geometrical features as basis of the registration, extensive prior art is available. Xu et al. [244] and Kaplan et al. [111] register MRI and TRUS by respectively finding rigid and affine transformations based on fiducials. Reynier et al. [178] register the point clouds from TRUS and MRI surfaces using first rigid and then elastic transformations, however without modeling the real organ deformations. Prostate surface-based registration models using thin plate spline basis functions are reported by Cool et al. [42] and Mitra et al. [140], both requiring manual interaction. In the work by Narayanan et al. [153] on MRI to TRUS deformable surface registration, an adaptive focus deformable model of a prostate phantom is used. Sparks et al. [208] as well as Hu et al. [103] have presented probabilistic and statistical shape and motion models of the prostate in patient datasets. The authors of these methods achieved low registration errors but also face time constraints for an integration into the clinical routine due to the manual prostate segmentation.

For surface-based techniques, a segmentation of the prostate in at least one modality is required. In many commercial systems, manual or semi-automatic contouring is the method of choice to obtain such segmentations, often leading to a large inter-observer variation [198]. A fully automatic segmentation method with clinically acceptable accuracy within permissible time limits is thus highly desired when relying on non-intensity-based registration approaches. In this context, Ghose et al. [77] review the extensive research on semi-automatic and automatic segmentation of the prostate from TRUS and MRI images. The main approaches in delineating prostate boundaries are contour-, shape- or region-based, and can be distinguished in supervised or unsupervised classification, as well as combinations of these. A recent study by Qui et al. [243] has reported promising segmentation results by using convex optimization with axial symmetry enforcement for 3D TRUS and MRI prostate images. However, the use of axial symmetry is a disputable assumption, especially in diseased prostate conditions. With the advent of deep learning techniques in field of prostate segmentation, Dice score accuracies beyond 90% come into reach [246].

The limitation of geometrical constraints capturing the true organ deformation can only partially be overcome by performing a deformable intensity-based registration between both images [207]. Focusing on geometric constraints of the transformation model, proposed methods are either based on control points and an interpolation scheme [80] or on a dense, voxel-wise formulation of the deformation field [242]. In either case, non-linear deformation models regularly include a high number of parameters, with known challenges in overcoming local minima, on physically reasonable regularization. Sun et al. [213] proposed a purely intensity-based approach to derive a deformation field, requiring structures to be visible in both images for optimal results. The approach uses the MIND similarity descriptor and performs convex optimization to achieve image alignment.

The availability of sufficiently large annotated datasets of various anatomies has often been exploited to generate statistical models of shape, texture and deformation as priors for deformable registration, effectively reducing the dimensionality of the optimization problem while at the same time enforcing physically meaningful deformations [207]. This also includes methods using image registration to create an atlas [182], or to map ground truth segmentation onto unseen images [78]. A learning technique is used for dimensionality reduction, for instance in the case of principal component analysis (PCA) allowing to optimize an unseen dataset's representation in PCA space. While the work of Onofrey et al. [160] allows for a MRI/TRUS registration using a population-based statistical deformation model, the algorithm is point-based and does not consider image intensities.

A popular choice to generate patient-specific statistical models of prostate motion is biomechanical simulation [60, 103, 234] based on prior segmentations in MRI. Such approaches regularly require a large quantity of finite element simulations for a wide range of TRUS probe positions etc., which might be prohibitive in clinical environments. Recently, Hu et al. [104] have shown that predictions of prostate deformation can be performed using population FEM data, i.e. without patient-specific modeling, with similar accuracy. Yet, MRI segmentations are again required for personalization of unseen datasets.

One of the contributions of this chapter is a novel, fully automatic MRI/TRUS registration scheme combining a statistical deformation model (SDM) generated from a population of

clinically observed prostate deformations with an intensity-based image registration algorithm that does not require a segmentation of unseen datasets. This is different from the work by Tahmasebi et al. [216], where an SDM is employed to estimate eigenmodes using a set of known landmark-based deformations, which are then used to extrapolate a likely whole-gland deformation without considering image intensities. Instead of merely incorporating the SDM as regularizer to penalize unlikely, i.e. physically unrealistic deformations as in [120], we directly optimize for eigenmode coefficients, greatly reducing the dimensionality of the registration problem.

2.2.3 Contributions

Based upon our early prototype [194], which solely relies on rigid MRI/TRUS registration using four manually selected fiducial landmarks, a comprehensive TRUS-based multi-modal prostate biopsy guidance framework was developed in the course of this thesis. Building on top of the experience gained with the prototype, feedback from urologists could be rapidly translated into the development process, leading to a guidance system currently in regular, weekly clinical use in the urology department of Klinikum Rechts der Isar. As of February 2017, more than 300 patients have undergone prostate biopsy procedures using the presented framework, with early clinical results already published [212] and further clinical trials in progress.

The methodological development can be described in three stages, which are illustrated in Fig. 2.3 depicting the clinical protocol:

0) The mentioned prototype [194] established the foundations of 2D B-mode TRUS acquisition and 3D compounding, as well as the possibility to visualize MRI and PET slices (MPR, multi-planar reconstructions) corresponding to the current US frame once a linear transformation between TRUS and MRI is known. In this prototype, a rigid registration using fiducial landmarks is employed. An optical tracking system is utilized for both 3D acquisition and live targeted biopsy guidance.

1) The framework currently in regular clinical use was presented in [256] and introduces surface-based deformable segmentation. In particular, a Hough forest approach is employed to automatically detect the prostate boundary in TRUS, which in return is used for an automatic registration based on the Coherent Point Drift algorithm [148]. Its integration into the regular clinical workflow is possible as the system requires minimal user interaction and performs US acquisition, segmentation and registration in less than five minutes – the time allocated for the local anaesthetic to take effect.

2) Finally, an intensity-based approach has been developed specifically for the challenging prostate anatomy and presented in [255], aiming at ultimately substituting the surface-based technique. First, we present a statistical deformation model (SDM) for the prostate between MR and TRUS images acquired and automatically segmented using the system described above. Second, a novel, multi-modal deformable registration scheme based on the derived SDM is introduced. Combining the LC^2 similarity metric [70] with a Mahalanobis distance-inspired preconditioning, our method achieves optimal alignment

at crucial locations within the organ, intended for accurate biopsy targeting in clinical routine.

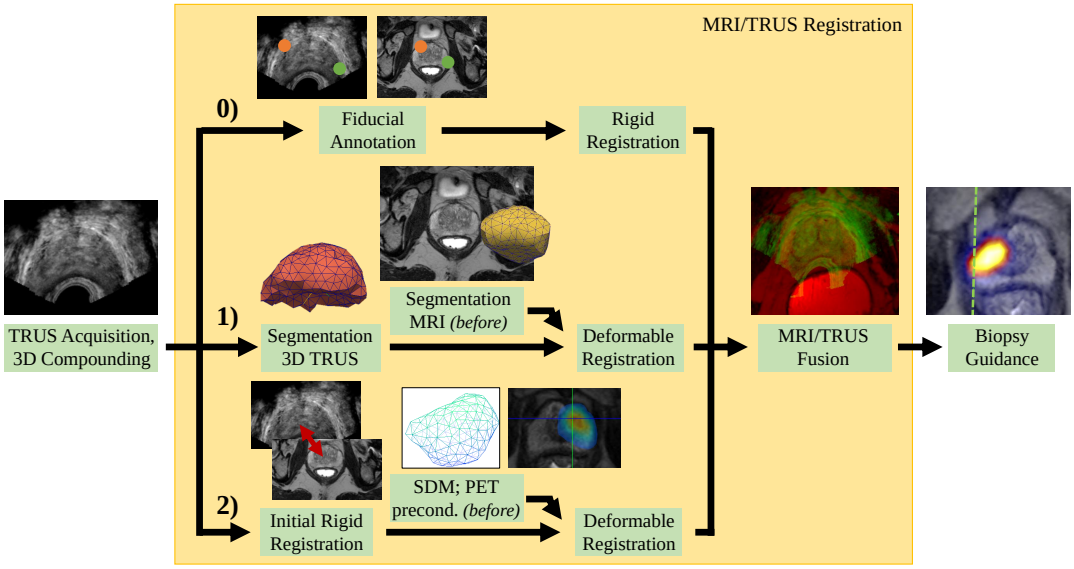

Fig. 2.3. Overview of the clinical protocol for multi-modal prostate biopsy guidance. Three different algorithms were implemented for MRI/TRUS registration: 0) Rigid based on fiducial landmarks, 1) Deformable surface-based using segmentations, and 2) Deformable intensity-based using a statistical deformation model (SDM) and similarity preconditioning using PET information. Apart from the MRI segmentation and SDM model generation, which are obtained in advance, all steps can be performed within a time frame of around five minutes to fit into the time requirements of conventional, only US-guided biopsy procedures.

For both novel registration approaches, results of a quantitative evaluation of our automatic system on phantom and patient datasets are reported and compared against expert annotations serving as ground truth data. While the underlying algorithms used for segmentation and surface-based registration have already been presented in the literature, our work is the first to employ these algorithms with necessary modifications on the challenging prostate anatomy and present results of their performance, not only on image data but in their real application on phantoms and patients. To the best of our knowledge, our framework remains to be the first TRUS-based multimodal prostate biopsy guidance system using elastically registered PSMA-PET/MRI datasets implemented in the routine clinical workflow.

2.3 Methods

The proposed multimodal prostate fusion biopsy system is embedded into the clinical routine as illustrated in Fig. 2.3. After the patient has assumed lithotomy position on the examination chair (a supine position with the legs separated, flexed, and supported in raised stirrups), 2D TRUS images of the prostate are acquired and compounded into a 3D volume. Section 2.3.1 describes the components of the system as well as the image acquisition and compounding process. Immediately after TRUS acquisition, the urologist initiates local anesthesia, which takes a few minutes to take full effect. In this period, one of the registration methods is used to align the patient's (PET/)MRI to the just acquired TRUS volume. The proposed algorithms to achieve surface-based (Sec, 2.3.3) or intensity-based registration (Sec. 2.3.5) rely on

prerequisites, which are presented in Sec. 2.3.2 (automatic segmentation) and Sec. 2.3.4 (statistical deformation model). Finally, Sec. 2.3.6 describes how the registered images are used for live prostate biopsy guidance.

2.3.1 Image Acquisition and Compounding

Ultrasound acquisition is based on standard components used in urological practice, including a conventional 2D front fire trans-rectal probe. Such systems are routinely used for US-guided prostate biopsies by attaching calibrated needle holders to the probe and following a needle guide on the US images [194]. In this work, the US-based guidance is not altered but augmented with MRI and PET information.

Compounding of 3D volumes is achieved using an optical tracking system as schematically illustrated in Fig. 2.4. One tracking target (*probe*) is rigidly attached to the transducer opposite of the needle guide. A second tracking target (*reference*) is attached to the patient chair in a known, reproducible way (${}^{\text{reference}}\text{T}_{\text{chair}}$). This allows the compounded 3D volumes to be embedded in the chair coordinate system, which is mapped to the 2D TRUS slices as follows:

$$ {}^{\text{frame}}\text{T}_{\text{chair}} = {}^{\text{reference}}\text{T}_{\text{chair}} \cdot ({}^{\text{reference}}\text{T}_{\text{world}})^{-1} \cdot {}^{\text{probe}}\text{T}_{\text{world}} \cdot {}^{\text{frame}}\text{T}_{\text{probe}} \tag{2.1} $$

This serves two purposes: First, the 3D volumes will be independent from the tracking camera

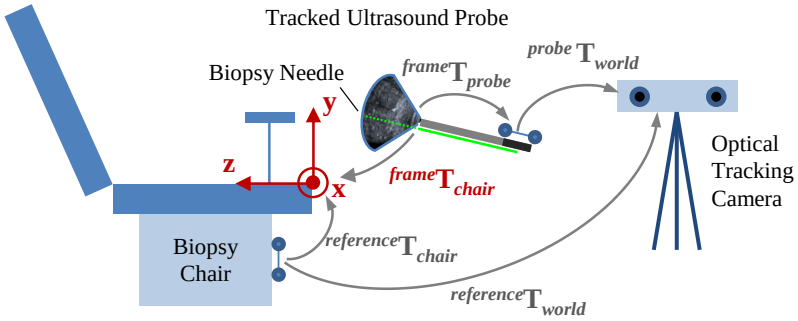

Fig. 2.4. Schematic system overview for prostate fusion biopsy along with the required coordinate systems used to reconstruct a 3D TRUS volume relative to the chair to make the initialization for MRI registration straightforward. *Reprint from [256] with permission of Springer.*

position, allowing to move it freely during the intervention. Because needle access to various lesions throughout the prostate requires transducer rotations of more than 180°, changes of the camera position are in fact rather likely during biopsy sessions to avoid the common line-of-sight problem. Second, the resulting volumes will be homogeneous across patients and facilitate the initialization of the subsequent registration step by means of a simple translation. Note that the patient is strictly told to remain static during the entire intervention.

For US compounding, i.e. 3D volume reconstruction, linear interpolation and hole filling with the Gaussian accumulation technique is employed, based on the approaches in [24] and [81]. Best performance of our system was achieved with an isotropic voxel spacing of 0.5 mm. Larger spacings incorporate more averaging artifacts, while smaller spacings only increase the segmentation time without any further improvement in its quality.

2.3.2 Automatic Hough Forest-based Segmentation

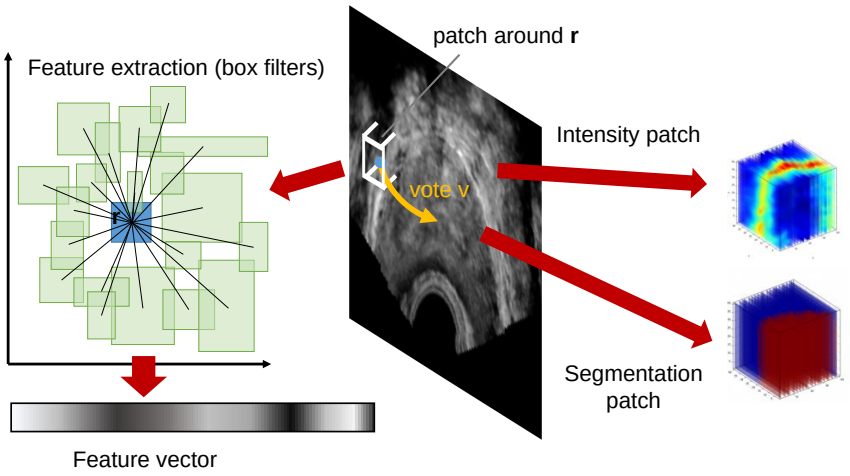

Fig. 2.5. Around each sampling point $\mathbf{r}$, a high dimensional feature vector is extracted using the entire pattern of intensities in a local neighborhood (*blue box*) as well as mean intensities from box filters (*green boxes*). Vectors $\mathbf{v}$ from boundary sampling points to the prostate center of mass are essential for the Hough voting scheme. Intensity and segmentation patches are furthermore stored for reprojection weighting.

A fully automatic prostate segmentation in compounded 3D TRUS volumes is obtained using Hough forests, a method coupling the classification performance of random forests with object localization capabilities. In our method, we further extend these with a strategy to extract a segmentation contour once the position of the prostate in the image is obtained through a voting strategy. The algorithm is similar to the one presented by Rematas et al. in [176], and is based on the technique presented in [139] already applied to the segmentation of 3D echocardiographic data.

Our Hough forest is an ensemble of Hough trees, trained on a set of TRUS volumes of various size and common spacing using segmentation contours, which were manually annotated by experts. During training, a grid of sampling points over each annotated volume is defined. As illustrated in Fig. 2.5, around each sampling point $\mathbf{r}$, a 1,000-dimensional feature vector corresponding to the entire pattern of intensities in a pixel patch of the volume and to the mean intensities of random box filters is extracted. We chose high dimensional feature vectors since we rely on the capabilities of the Hough forest to select the features that are most discriminative for the problem at hand [175].

The sampling points belonging to a narrow region around the prostate boundary, denoted as foreground (*fg*), are also associated with a vote $\mathbf{v}$, in the form of a displacement vector from the sampling point itself to the center of mass of the prostate. Additionally, the original position $\mathbf{r}$ in the training volume of each point of the foreground is stored and supplied to the training algorithm.

Training a Hough tree amounts to recursively splitting the training set T in the splitting nodes until one of the termination criteria is met, in which case the recursion is stopped and a leaf is instantiated. In our framework, the termination criteria are defined as a depth of 18 levels, or alternatively, a population smaller than 30 data points reaching a particular node. The splits are chosen such that either the information gain is maximized (Eq. 2.2) or the vote

scattering (Eq. 2.3) is minimized. In the first equation, the splitting threshold θ is chosen such that the difference between the Shannon entropy of the splitting node $H(T)$ and the weighted Shannon entropies $H(T_l)$ and $H(T_r)$ of each of the child nodes is maximized. While this ensures good classification capabilities, Eq. 2.3 provides a criterion to choose the splitting threshold using only the votes $\mathbf{v}$ associated with data points $\mathbf{d}$ belonging to the foreground. The votes reaching each child node are compared in terms of squared distance to the mean vote $\bar{\mathbf{v}}$ for that child node. This measure is minimized to ensure minimal vote scattering after the split.

$$I(T,\theta) = H(T) - \sum_{k=\{l,r\}} \frac{|x \in T_k|}{|x \in T|} H(T_k) \tag{2.2}$$

$$V(T,\theta) = \sum_{k=\{l,r\}} \sum_{i|\mathbf{d}_i \in \text{fg}} (\mathbf{v}_i^k - \bar{\mathbf{v}}_k)^2 \tag{2.3}$$

Apart from the class posterior distribution, when a leaf is instantiated, we store one vote $\mathbf{v}_i$ and the original position of the data point $\mathbf{r}_i$ for each data point $\mathbf{d}_i$ belonging to the foreground and reaching that leaf.

During **testing**, an image that was not employed for training is sampled using a regular grid. The features of each data point are used to traverse each tree until reaching a leaf node. If a data point is classified as foreground by a sufficient number of trees in the forest, all the votes contained in the leaf nodes reached by this data point are weighted by the classification confidence and accumulated in a volume having the same dimensionality as the image. Additionally, as the votes are being cast, a pixel-wise reprojection list of votes is filled with information about the provenance of each vote and the position of the corresponding training data point in the ground truth images. The position where the maximum number of votes was attained represents the position of the prostate's center of mass in the test image.

Using the information stored in the reprojection list in the immediate neighborhood of the maximum vote, we are then able to propagate portions of ground truth contours from the training images onto the newly acquired test image. Additionally, we weight each contribution by the truncated normalized cross correlation (NCC) between the intensities of the currently considered patch around $\mathbf{r}_i$ in the test image and the patch around $\mathbf{r}_t$ of the training image (see Eq. 1.55 in Sec. 1.4.3).

After the reprojection is complete, a normalization step is performed in order to obtain a probabilistic segmentation, which can thereafter be thresholded at 0.5. We perform an automatic connected component analysis as post-processing, resulting in a precise delineation of the boundary of the prostate. The binary segmentation image is finally approximated to a triangular surface mesh Γ_{US} following the algorithm outlined in [23].

2.3.3 Deformable Surface-based Registration

The previously generated surface meshes from the TRUS and MRI segmentations are now used to deformably register both images. Such a registration can be understood as a two-step process: Firstly, only the surface meshes are deformably registered, for instance by computing displacements for each vertex of one mesh such that a distance d between both meshes is

minimized. Secondly, these sparse displacements are used to interpolate a dense deformation field over the entire image domain in order to warp the MRI image onto the TRUS image.

In this work, we achieve the first step, i.e. the mesh registration, by employing a modified version of the Coherent Point Drift (CPD) algorithm [148], which is a fully deformable point registration method without any geometrical constraints regarding the shapes to be registered. Additionally, it also shows robustness in terms of outliers as it forces the points to move coherently. The algorithm registers two sets of points, which are both modeled as centroids of Gaussian mixture models (GMMs). The first set of points $\mathrm{X} \in \mathbb{R}^{n\times 3}$ comprises the vertices of the TRUS mesh Γ_{US} and is considered static. The vertices of the MRI mesh Γ_{MR} are collected in the second, moving set of points $\mathrm{Y} \in \mathbb{R}^{m\times 3}$, which is fitted to X during the registration procedure using expectation maximization (EM). m and n denote the number of points in the two sets, respectively. For numerical robustness, both point sets are demeaned and normalized with respect to translation and scale. The goal of the algorithm is now to find a non-rigid transformation $\mathcal{T}$ such that the deformed point set $\mathrm{Y}' = \mathcal{T}(\mathrm{Y}, v) = \mathrm{Y} + v(\mathrm{Y})$ is given as the initial positions plus an optimal displacement vector field v. Assuming that for a proper mesh registration the two Gaussian mixtures will be statistically similar, mesh vertex correspondences will be derived, expressed as the maximum of the GMM posterior probability for any given data point. Therefore, the L_2 distance is well suited as a cost function for the parameter optimization:

$$d_{L_2} = \int \left[\mathrm{gmm}(\mathcal{T}(\mathrm{Y}, v)) - \mathrm{gmm}(X)\right]^2 dx. \tag{2.4}$$

In this notation, $\mathrm{gmm}(\mathcal{P}) = \sum \alpha \mathcal{N}(\mu, \Sigma)$ denotes the Gaussian mixture density of point set $\mathcal{P}$ with weights α. As parametrization for $\mathcal{T}$, thin-plate splines (TPS) transformations are used, which can be decomposed into an affine as well as a non-linear part parametrized by a set of warping coefficients w. As such, the relation between the moving and the original MRI mesh vertices can be expressed as:

$$\mathrm{Y}' - \mathrm{Y} = \left[\mathrm{Y}|\mathbf{1}\right] \mathrm{A}^\intercal + \mathrm{U}\mathrm{w}. \tag{2.5}$$

Hereby, $\mathbf{1}$ is a column vector of size m. The matrix $\mathrm{A} \in \mathbb{R}^{3\times 4} = [\mathrm{B}|\mathbf{t}]$ models the affine part of the motion with translation $\mathbf{t}$ as well as rotation and scaling B. The basis matrix $\mathrm{U} \in \mathbb{R}^{m\times m}$ for the local, non-linear transformations is expressed using radial basis functions in 3D, i.e. $\mathrm{U} = \{u_{i,j}\}$, $u_{i,j} = \|\mathrm{Y}'_i - \mathrm{Y}_j\|$. To ensure that the non-linear deformation becomes zero at infinity, new parameters $\mathrm{v} \in \mathbb{R}^{(m-4)\times 3}$ are introduced such that $\mathrm{w} = \mathrm{N}\mathrm{v}$, where N denotes the left null space of $[\mathrm{Y}|\mathbf{1}]$. It has been shown that efficient, gradient-based numerical optimization techniques are suited to solve for A and w [17], using the following discretized cost function $D_{L_2}(\mathrm{Y}')$ and its derivatives:

$$\begin{aligned} D_{L_2} &= \sum_{i=1}^{m}\sum_{j=1}^{n} p_{i,j} + \frac{\lambda}{2}\,\mathrm{tr}(\mathrm{w}^\intercal \mathrm{K}\mathrm{w}), \\ \frac{\partial D_{L_2}}{\partial \mathrm{A}} &= \left[\mathrm{Y}|\mathbf{1}\right]^\intercal \mathrm{G}, \\ \frac{\partial D_{L_2}}{\partial \mathrm{w}} &= \mathrm{U}^\intercal \mathrm{G} + \lambda \mathrm{K}\mathrm{w}. \end{aligned} \tag{2.6}$$

The matrix $\mathrm{P} \in \mathbb{R}^{m \times n} = \{p_{i,j}\}$ contains the GMM posterior probabilities:

$$p_{i,j} = \exp\left[-\frac{\|\mathrm{Y}_i - \mathrm{X}_j\|}{2\sigma^2}\right]. \tag{2.7}$$

For the optimization, the derivative matrix $\mathrm{G} \in \mathbb{R}^{m \times 3} = \partial D_{L_2}/\partial \mathrm{Y}'$ can be obtained using P. The regularization term with kernel matrix $\mathrm{K} \in \mathbb{R}^{m \times m} = \{k_{i,j}\}$, $k_{i,j} = \|\mathrm{Y}_i - \mathrm{Y}_j\|$ and weighting parameter λ impose penalties on the bending of the TPS deformation field.

As the parameters v are initialized with 0, the algorithm will intuitively first establish an affine transformation, which could cause significant rotation or shearing of the prostate meshes, and only then fine-tune the local non-rigid TPS components. However, minimal rotation is to be expected because the chair coordinate system (see Fig. 2.4) is defined to match the orientation of the coordinate system of the MRI, and the translation should also be small because both point sets were demeaned before the registration. Therefore, we constrain the affine part to the identity transformation in the first k iterations of the optimization (identity rotation $\mathrm{I} \in \mathbb{R}^{3 \times 3}$ and $\mathbf{0}$ translation):

$$\mathrm{A}^{i \leq k} = \left[\mathrm{I}|\mathbf{0}\right]. \tag{2.8}$$

After the point set registration is performed, the second step of the mesh registration consists of interpolating a dense displacement field for the entire MRI image using the transformed MRI mesh points Y'. One advantage of our interpolation scheme is that the TPS parameters w already obtained by the deformable registration can be employed to warp an arbitrary point from the original MRI space into the TRUS image. Incorporating Y and w as RGB textures, the warping can easily be performed directly on the graphics processing unit (GPU) in the fragment shader. Note that such a mapping can be performed in both directions, i.e. from Y to Y' and vice versa, but the mapping is not bijective in the general case [56]. Because MRI and PET images are acquired simultaneously and reconstructed in the same coordinate system, all three modalities are now registered.

2.3.4 Statistical Deformation Model Generation

The generation of the SDM is based on a dataset of N corresponding pre-interventional MR images I_{MR} and interventional TRUS images I_{US} acquired respectively in the supine and lithotomy positions. Triangular surface meshes of the prostate are created from available binary segmentations in both images (manually segmented Γ_{MR}, automatically segmented Γ_{US} using Hough forest, see Sec. 2.3.2), and demeaned so that their center of gravity is at the origin. As shown in Fig. 2.6, their vertices are elastically registered with the Coherent Point Drift (CPD) algorithm [148] to obtain a warped MR mesh Γ'_{MR} as in Sec. 2.3.3. Due to the arbitrary vertex numbering in the N meshes, point correspondence across patients is established by intersecting M angularly equidistant rays starting at the origin with meshes Γ_{MR} for all patients. Thus, new vertex positions in MR space $\mathbf{p}_{i,j} \in \Gamma_{MR_i}$ and – using barycentric interpolation – in US (i.e. warped MR) space $\mathbf{p}'_{i,j} \in \Gamma'_{MR_i}$ are obtained for patients

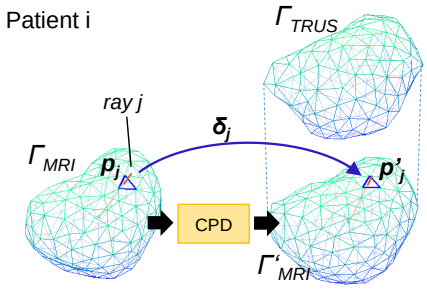

Fig. 2.6. Determination of prostate shape and deformation vectors corresponding across patients using ray casting technique. *Reprint from [255], ©2017 IEEE.*

$i = 1, .., N$ and rays $j = 1, .., M$. The point set $\hat{P}$ of the mean prostate shape in MR, and the mean deformations $\hat{D}$ are then defined as

$$\hat{P} = \{\hat{\mathbf{p}}_j = \frac{1}{N}\sum_{i=1}^{N} \mathbf{p}_{i,j}\}, \tag{2.9}$$

$$\hat{D} = \{\hat{\boldsymbol{\delta}}_j = \frac{1}{N}\sum_{i=1}^{N} \boldsymbol{\delta}_{i,j}\}, \tag{2.10}$$

with $\boldsymbol{\delta}_{i,j} = \mathbf{p}'_{i,j} - \mathbf{p}_{i,j}$ (see Fig. 2.6). As in [182], the SDM is created using PCA. In brief, the deformations are demeaned ($\boldsymbol{\delta}'_{i,j} = \boldsymbol{\delta}_{i,j} - \hat{\boldsymbol{\delta}}_j$) and vectorized into matrix $\Delta \in \mathbb{R}^{N\times 3M}$. An Eigen analysis of cov(Δ) yields the sorted and devectorized eigenvectors ϕ_k and corresponding eigenvalues λ_k.

2.3.5 Preconditioned Intensity-based Registration using Statistical Deformation Model

Recalling from Sec. 1.4, image registration aims at finding an optimal transformation T^* between reference and moving images, in our case TRUS and MRI, using a similarity metric S and an iterative solver, in this work BOBYQA [169]:

$$T^* = \arg\max_T S\left[\mathrm{I}_{US}, T(\mathrm{I}_{MR})\right] \tag{2.11}$$

Hereby, the similarity metric $S(\mathrm{I}_{US}, \mathrm{I}'_{MR})$ measures how well fixed image I_{US} and transformed moving image $\mathrm{I}'_{MR} = T(\mathrm{I}_{MR})$ correspond. A common way to define S is the integration of a metric f over the entire overlapping image domain Ω (see Sec. 1.4.3):

$$S(\mathrm{I}_{US}, \mathrm{I}'_{MR}) = \int_\Omega f(\mathbf{x})\, d\mathbf{x}. \tag{2.12}$$

In this work, the multi-modal LC^2 similarity metric presented in [70] is employed due to its excellent behavior for MRI/ultrasound registration.

Similarity Metric Preconditioning

The aim of the proposed preconditioning is to emphasize the optimization of image alignment at crucial locations $\mathbf{t}$ for a given clinical application. For prostate biopsy guidance, suffering from a high false negative rate, such locations could be suspicious lesions present in MRI and/or PET. In this work, locations $\mathbf{t}$ are automatically identified by the position of maximum PSMA expression in PET images, which are already registered to the corresponding MR images by acquisition (combined PET/MR scanner). Note that for other clinical applications, or in absence of nuclear imaging, crucial locations could also be defined in a different manner, for instance by manual annotation by an expert.

We propose to modify the metric of the otherwise Euclidean space for similarity integration as follows:

$$S(\mathrm{I}_{US}, \mathrm{I}'_{MR}) = \int_{\Omega} f(\mathbf{x}) \quad \underbrace{\zeta(\mathbf{x})\, d\mathbf{x}}_{\text{metric change}} \quad . \tag{2.13}$$

In contrast to Eq. 2.12, the function $\zeta(\mathbf{x}) \in [0; 1]$ modifies the "density" of the image space, effectively emphasizing the registration on areas where ζ is close to 1, and removing influence of areas where ζ is close to 0. A simple, isotropic (i.e. direction-independent) preconditioning around $\mathbf{t}$ can now be achieved using the logistic function and the regular Euclidean norm:

$$\zeta_{iso}(\mathbf{x}) = 1 - \left(1 + e^{-k\cdot(\|\mathbf{x}-\mathbf{t}\|-d_0)}\right)^{-1}, \tag{2.14}$$

where parameters k and d_0 control logistic slope and the curve's inflection point, respectively.

Zikic et al. [258] used location-independent gradient normalization to improve mutual information-based registration. The idea of our approach is to not only focus the registration on the region around $\mathbf{t}$ but also predominantly along the expected directions of deformation at $\mathbf{t}$. To this end, we estimate the deformations $\mathbf{d}_k$ at this point using the first three SDM eigenmodes $\sqrt{\lambda_k} \cdot \boldsymbol{\phi}_k$, for $k = 1, 2, 3$, and thin-plate splines (TPS) for interpolation between control points. Note that because the TPS interpolation does not guarantee a linear mapping, the deformations' covariance matrix $\Sigma = \mathrm{cov}([\mathbf{d}_1\, \mathbf{d}_2\, \mathbf{d}_3])$ is not necessarily diagonal. Denoting

$$\|\mathbf{x} - \mathbf{t}\|_{\Sigma} = \sqrt{(\mathbf{x} - \mathbf{t})^{\intercal} \Sigma^{-1} (\mathbf{x} - \mathbf{t})} \tag{2.15}$$

the Mahalanobis distance with respect to location $\mathbf{t}$, we propose the anisotropic preconditioning function:

$$\zeta_{aniso}(\mathbf{x}) = 1 - \left(1 + e^{-k\cdot\left(\|\mathbf{x}-\mathbf{t}\|_{\Sigma}-d_0\right)}\right)^{-1}. \tag{2.16}$$

An example of this function is shown in Fig. 2.15f.

Automatic Deformable Registration Workflow

The workflow to perform automatic deformable registration first requires both MR and ultrasound images to be rigidly aligned. A reasonably good rotatory initialization is achieved using an optical reference target placed consistently for all patients during US acquisition (see Sec. 2.4.1 for details on setup and protocol). In terms of translation, both images are first aligned to the center of the image, followed by a rigid registration using LC^2.

For the essential part of the proposed registration scheme, we formulate a deformation field using a linear combination of the SDM eigenmodes. For each control point $\hat{\mathbf{p}}_j$, i.e. vertex of the mean prostate shape, the corresponding deformation is defined as follows:

$$\mathbf{d}_j = (1 + \theta_0) \cdot \hat{\boldsymbol{\delta}}_j + \sum_{k=1}^{L} \theta_k \cdot \sqrt{\lambda_k} \cdot \boldsymbol{\phi}_k, \tag{2.17}$$

where L is the number of used eigenmodes, and $\boldsymbol{\theta} = (\theta_0, ..., \theta_k)$ the vector of optimization parameters. For registration, we now directly optimize for an optimal parametrization

$$\boldsymbol{\theta}^* = \arg\max_{\boldsymbol{\theta}} S\left[\mathrm{I}_{US}, T_d(\mathrm{I}_{MR}, \boldsymbol{\theta})\right]. \tag{2.18}$$

Hereby, $T_d(x, \boldsymbol{\theta}) = x + D(x, \boldsymbol{\theta})$ denotes the elastic transformation with the dense deformation field D created using thin-plate splines (TPS) as in [256] for efficient image warping on graphics processors. Note that θ_0 allows to scale the mean deformation, allowing to cope with various probe pressures during US acquisition.

2.3.6 Interventional Visualization and Biopsy Guidance

Finally, the mapping parameters obtained in the previous step are utilized to guide the urologist during the biopsy procedure. Assuming the patient has not moved in the meantime, the transformation ${}^{\text{frame}}\mathrm{T}_{\text{chair}}$ (see Fig. 2.4) containing the real-time tracking information of the US probe is used to interactively identify which part of the 3D TRUS volume is currently visible on the US machine. Then, corresponding MPRs (multi-planar reconstructions) of the MRI and PET images according to the deformable registration are rendered and presented to the urologist. A virtual biopsy needle guide, previously calibrated, indicates an approximate needle insertion path, as illustrated in Fig. 2.8. Two separate views are employed for successful navigation: While the sole MRI image provides anatomical context, a fused representation of gray-scale MRI and colored PET allows for precise targeting of suspicious lesions. Target biopsies are taken by manoeuvring the US probe such that the virtual biopsy guide aligns with the target site.

2.4 Materials and Experiments

In this section, all necessary details to re-create the presented system setup are explained in Sec. 2.4.1. Next, the phantom generation procedure is outlined in Sec. 2.4.2, before an overview of patient datasets and the experimental protocols is provided in Sec. 2.4.3.

2.4.1 System Setup and Implementation Details

The final system setup as established in the university hospital's urology department is depicted in Fig. 2.7. It comprises the following components:

- A conventional Hitachi AVIUS® (Hitachi Ltd., Tokyo, Japan) US system with a 2D front fire trans-rectal probe, Model EUP-V53W, is employed for image acquisition. Due to

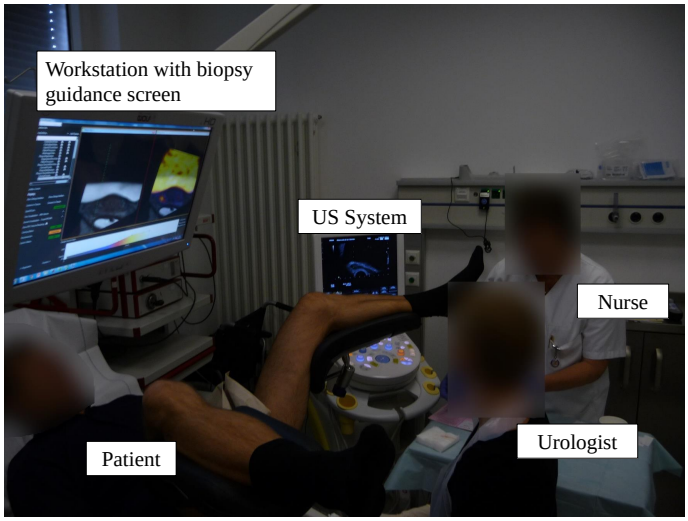

Fig. 2.7. System setup as used in the urology clinic, where a large screen is employed for multi-modal biopsy guidance. *Reprint from [256] with permission of Springer.*

lack of access to RF data from the US machine, a StarTech frame grabber (StarTech.com Ltd., London, ON, Canada) is used to acquire high resolution (1280×1024 pixel) digital images.

- The US probe is tracked by an NDI Polaris® (Northern Digital Inc., Waterloo, ON, Canada) optical tracking system. Two tracking targets are attached to the back of the US transducer, opposite of the needle guide, and the patient chair, respectively.

- For the image processing of all volumes used for this work, a workstation with 2 Intel Xeon® processors running at 2.13 GHz with 32 GB RAM and an NVIDIA GeForce® GTX TITAN Black graphics card was used. The system was recently replaced by an Asus ROG notebook with an Intel Core i7 running at 2.5 GHz, 16 GB RAM and an NVIDIA GeForce® GTX 980M for increased mobility and less overhead in the intervention room.

In terms of software libraries and frameworks, we relied on readily existing solutions. The open source software SlicerIGT[1] proved useful to perform spatial calibration of the US image to the optical tracking (${}^{\text{frame}}\text{T}_{\text{probe}}$). Furthermore, we use the freely available PLUS framework[2] [124] for temporal calibration, tracked freehand US acquisition and for the compounding of 3D TRUS images [194].

With the exception of the intensity-based registration method, which has not been introduced to clinical usage yet, the workflow illustrated in Fig. 2.3 was implemented in CAMPVis, an open source game-engine inspired research framework[3] [192]. The relevant workflow plugins and pipelines are also publicly available. Our implementation of the Hough forest algorithm is based on the Sherwood C++ library [44]. For subsequent registration, we finally employ the CGAL library[4] to triangulate the binary volumetric segmentation image and create a

[1] Available online: `http://www.slicerigt.org`
[2] Available online: `https://app.assembla.com/spaces/plus/wiki`
[3] Available online: `http://campar.in.tum.de/Main/CAMPVis`
[4] Computational Geometry Algorithms Library, available online: `http://www.cgal.org`

surface mesh. Our CPD implementation for surface-based registration is heavily inspired by the version of Jian and Vemuri[5] [17].

The statistical deformation model (SDM) was computed using MATLAB, R2013b (MathWorks, Inc., Natick, MA, USA). The intensity-based registration methodology as described in Sec. 2.3.5 was implemented within custom plugins for ImFusion Suite 1.2.20 (ImFusion GmbH, Munich, Germany), an extensible GPU-based framework for medical images.

Fig. 2.8 shows a screenshot of the image-based guidance system during a biopsy procedure, including the two MRI and PET/MRI views and the virtual needle guide lines.

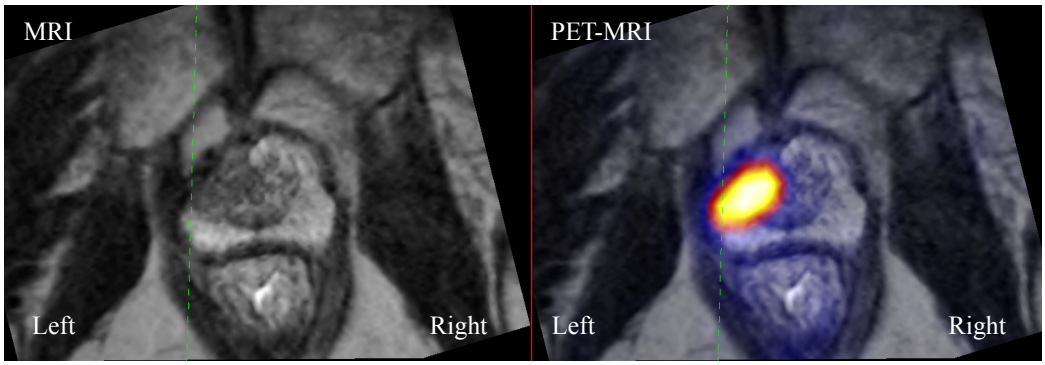

Fig. 2.8. The presented prostate biopsy guidance system provides the urologist with MRI and PET/MRI views corresponding to the current position of the TRUS transducer. The PET/MRI view shows a hotspot in a suspicious region in left apical prostate zone where the target biopsies were taken. Note that because the urologist turns the transducer by 180° to take biopsies from the patient's left side (thus flipping the TRUS image), also the MRI and PET/MRI slices become flipped in our visualization. *Reprint from [256] with permission of Springer.*

2.4.2 Prostate Phantom Construction

Based on the previous work of Dang et al. [45] for multi-modality phantoms, we used agar and gelatin to generate several prostate phantoms with suitable tissue contrast in both MRI and ultrasound. First, egg-shaped (roughly 50×30×30 mm) prostate glands consisting of 6 weight percent (wt%) gelatin, 3 wt% agar, and blue food coloring were cast. Into each gland, three lesions with a diameter of around 5 mm were positioned at varying locations during cooling. To ensure appropriate contrast between normal prostate tissue and lesions in MRI, the latter were made with 12 wt% gelatin, 8 wt% agar and red food coloring. The different colors of gland and lesions allow to visually distinguish the tissue types in biopsy samples. The urethra was mimicked by an air-filled plastic tube. For each prostate, an artificial rectum (modeled during casting by a glass cylinder of around 35 mm in diameter) was surrounded with a very elastic compound of 8 wt% gelatin and 1 wt% agar. Finally, the remaining phantom, now surrounding the prostate gland and maintaining a distance between rectal wall and gland of around 20 mm, was filled with a softer compound with only 1.8 wt% gelatin and 0.5 wt% agar. For efficiency, three prostate glands were positioned next to each other in one big plastic container and jointly MR-scanned, as illustrated in Fig. 2.9.

[5] Available online: `https://code.google.com/p/gmmreg/`

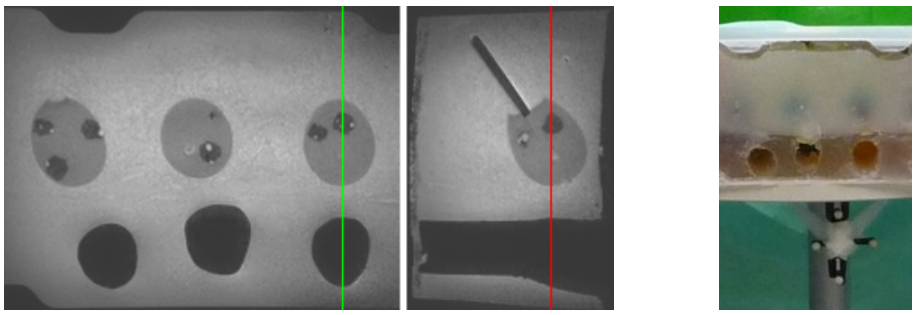

(a) (b)

Fig. 2.9. **a)** Axial (left) and sagittal (right) MRI slices of agar-gelatin phantoms with lesions well visible. The vertical lines indicate the slice in the respective other view. **b)** Setup for phantom biopsies with reference tracking target. *Reprint from [256] with permission of Springer.*

2.4.3 Patient Dataset and Protocol

In total, a collection of 50 patient datasets of prostate (PET/)MRI from various scanners and 3D TRUS were available for experiments. For parts of this collection, additional features such as manual TRUS segmentations were obtained from medical experts. An overview of the datasets and the additionally available information is provided in Tab. 2.1. All patients underwent prostate biopsies without pathological finding in the past and were referred for a follow-up biopsy to our clinic.

For all patients, the biopsy guidance framework presented in this thesis was employed, albeit with different registration algorithms (cf. workflow in Fig. 2.3). Regardless of the employed method, our system was used to acquire 3D TRUS volumes and guide targeted biopsies to suspicious areas identified in PET or MRI. In addition to the 12 regular biopsy cores spread throughout the prostate following the standard protocol, two targeted biopsy samples were taken from these areas.

The first part of the dataset collection (*Part 1*) was acquired in the beginning of the project, using rigid registration with a varying set of manually placed landmarks [194]. Careful manual TRUS segmentations were conducted by a medical expert, so that the Hough Forest, see Sec. 2.3.2, could be trained for subsequent clinical usage.

For the datasets accumulated in *Part 2*, the automatic segmentation algorithm (Sec. 2.3.2) and the deformable surface-based registration approach (Sec. 2.3.3) were already used in clinical practice. As a prerequisite, the prostate was carefully manually segmented in MRI in advance, to allow for automatic, deformable MRI-TRUS registration without user interaction during the intervention. After the acquisition of a 3D TRUS volume, the prostate gland was then automatically segmented and registered. For retrospective evaluation of the segmentation and registration performance, additional manual TRUS segmentations and sets of four corresponding anatomical landmarks in both modalities were collected. A preliminary clinical evaluation is possible in cases with available histology data (*Part 2a*), confirming that successful guidance toward previously identified suspicious lesions is possible with the developed system.

Dataset	N	Registration method	Manual MRI Segmentation	Automatic TRUS Segmentation	Manual TRUS Segmentation	# Manual Landmarks	PSMA-PET	Histology	Usage / Evaluation of
Total	50		✓	✓			*		SDM generation
⌞ Part 1	23	R	✓	✓	✓		*		Hough Forest training
⌞ Part 2	13	S	✓	✓	✓	4	*		Segmentation, surface-based registration
⌞ Part 2a	5	S	✓	✓	✓	4	✓	✓	Histology evaluation
⌞ Part 3	10	S	✓	✓	✓	6	✓		Intensity-based registration

Tab. 2.1. Patient dataset collection for prostate biopsy guidance evaluation. For various parts, additional features were obtained from medical experts for method validation. *R* denotes rigid landmark-based registration, and *S* the deformable surface-based method. Asterisks (*) indicate partial availability.

Once in total 50 patients had undergone fusion biopsies, manual MR segmentations as well as automatic Hough Forest-based segmentations were performed retrospectively for the cases where such data was still missing, in particular for the ones from *Part 1*. The complete collection of datasets was used for the generation of the statistical deformation model as outlined in Sec. 2.3.4.

Finally, the most recently acquired datasets of *Part 3* were used to retrospectively validate the intensity-based registration approach. To this end, manual TRUS segmentations and six structured point correspondences in both modalities (four at the prostate boundary, two at structures within the organ), were manually annotated by an expert. For all cases, PSMA-PET images acquired with the in-house PET/MR scanner were available. This enabled the automatic identification of locations $\mathbf{t}$ for preconditioning as determined by the position of maximum PSMA expression in PET images. For each patient, we asked for one landmark point to be placed in the vicinity of $\mathbf{t}$, ideally not farther away than 5 mm. Note that all landmarks were chosen based on the visibility of certain structures, including but not limited to urethra and calcifications. Although similar, the landmarks might thus not necessarily correspond across patients.

The reason why some of the datasets were not reused for other experiments, e.g., datasets of *Part 2* for the validation of the intensity-based approach, is related to the protocol of manual landmark selection. In general, the provision of landmarks occurred during actual biopsy sessions while the urologist was waiting for the local anesthesia to take effect. This did not only reduce the necessity to have time-consuming follow-up sessions devoted to data annotation but also increased the confidence of the urologist in the selected landmarks as both MR and TRUS images were still fresh in memory. As a result, it was more efficient to prepare a new dataset collection with updated landmark placement instructions than revising the landmark set of previously biopsied patients.

2.5 Results

In order to evaluate the performance of the presented methods, we first examine the applicability of the biospsy guidance system on agar-gelatin phantoms mimicking the conditions in prostate biopsy interventions as closely as possible (Sec. 2.5.1). Next, we analyze the performance of the individual components of our system, i.e. segmentation (Sec. 2.5.2), surface-based (Sec. 2.5.3) and intensity-based registration (Sec. 2.5.4-5). Finally, the image-based guidance capabilities in real patients is validated in Sec. 2.5.6 using histology data.

2.5.1 Phantom Experiments

TRUS images were acquired following the procedure described in Sec. 2.3.1 for all three phantoms two days after their creation (see Sec. 2.4.2). This time frame allowed diffusion between normal prostate tissue and lesions, making them almost invisible in ultrasound and thus mimicking real biopsy conditions well (cf. Fig. 2.10a). Due to the different intensity distributions compared to TRUS images of real patients, a semi-automatic segmentation approach [85] was employed for both TRUS and MRI, before the images were elastically registered as outlined in Sec. 2.3.3. Finally, an experienced urologist and a non-medical domain expert used the presented image-based guidance system to perform target biopsies. We considered a biopsy intervention on a particular phantom successful if all three lesions were hit at least once (visually identified using the red food color, cf. Fig. 2.10b), which was the case for all phantoms. Table 2.2 reports the biopsy samples in detail. Only in 2 out of 14 samples, the lesion was missed.

2.5.2 Retrospective Evaluation of Prostate Segmentation

Datasets of *Part 1* (see Sec. 2.4.3) were employed to train a Hough Forest, which was then used to automatically segment the prostate glands of the patients of *Part 2* during the biopsy session. For ground truth segmentations of both training and – retrospectively – testing images, manual expert annotations were used. As previously reported, e.g. by Reynier et al. [178], the

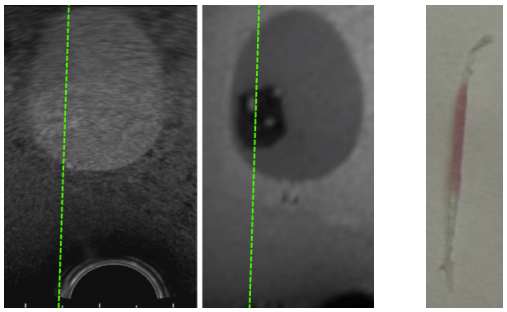

Fig. 2.10. **a)** Live TRUS image with invisible lesion and image-based MR guidance view during biopsy session. **b)** Positive core biopsy sample, lesion material visible due to red food color. *Reprint from [256] with permission of Springer.*

Phantom	User	Lesion 1	Lesion 2	Lesion 3	Total
1	Urologist	1/1	1/1	1/1	3/3 (100%)
2	Non-med. expert	2/2	2/2	1/2	5/6 (83%)
3	Urologist	2/2	1/1	1/2	4/5 (80%)

Tab. 2.2. Results of prostate phantom biopsy samples (number of positive cores / total number of cores). All lesions were positively sampled at least once. *Reprint from [256] with permission of Springer.*

Patient	Segmentation Surf. Dist.			TRE Rigid		TRE Deform.	
	Median	Mean	Std.Dev.	Mean	Max	Mean	Max
1	0.56	0.98	1.09	1.14	1.45	0.94	1.45
2	0.77	1.03	1.00	1.88	3.18	1.82	2.17
3	0.41	0.62	0.67	1.92	2.96	0.88	1.12
4	0.67	0.79	0.70	2.58	4.66	2.42	4.14
5	2.16	2.69	2.37	3.62	5.62	1.74	2.70
6	0.49	0.75	0.83	3.68	5.11	1.80	2.85
7	0.86	1.14	1.13	2.38	2.79	2.80	3.73
8	0.56	1.08	1.29	2.40	3.06	1.61	2.35
9	0.41	0.66	0.76	2.19	3.96	3.55	6.56
10	0.56	0.90	1.02	2.70	3.18	2.27	2.75
11	0.62	0.81	0.77	1.82	2.00	1.79	2.35
12	0.67	0.96	0.95	4.99	7.80	2.08	2.65
13	0.77	1.13	1.18	2.60	3.60	2.27	3.64
Average	**0.73**	**1.04**	**1.06**	**2.60**		**2.00**	

Tab. 2.3. Segmentation and registration results for all patients. Segmentation results are reported as surface distances between ground truth and Hough forest-based segmentation [mm]. Registration results compare target registration errors (TRE) [mm] for rigid and automatic deformable registration, computed with four landmarks each. *Reprint from [256] with permission of Springer.*

segmentation of the basal and the apical parts of the prostate in US images is very challenging. Even to the human eye, delineating the boundaries in these regions is often ambivalent and will result in high intra- and inter-observer variability. In order to maintain comparability to ground truth segmentations, the prostate is subdivided into the three subregions *base*, *mid-gland* and *apex* following the convention described in [243], i.e. according to 30%, 40% and 30% of the base-apex axis, respectively, and focus our evaluation on the *mid-gland* region only.

As similarity metric, the surface distance between ground truth and the automatically obtained segmentations is used (cf. Tab. 2.3). The method achieved an average surface distance of 1.04 ± 1.06 mm (mean $\pm$ standard deviation). Across all records, the maximum distance between both surfaces was on average 5.28 mm. Segmentation results for six representative patient datasets are illustrated in Fig. 2.11, showing the triangulated segmentation surfaces as well as their color-coded distance to the ground truth annotation. We also computed a Dice similarity score, which was on average 87.81 ± 2.92 for the *mid-gland* region. On our workstation, the training of the Hough forest using 23 datasets took approximately 105

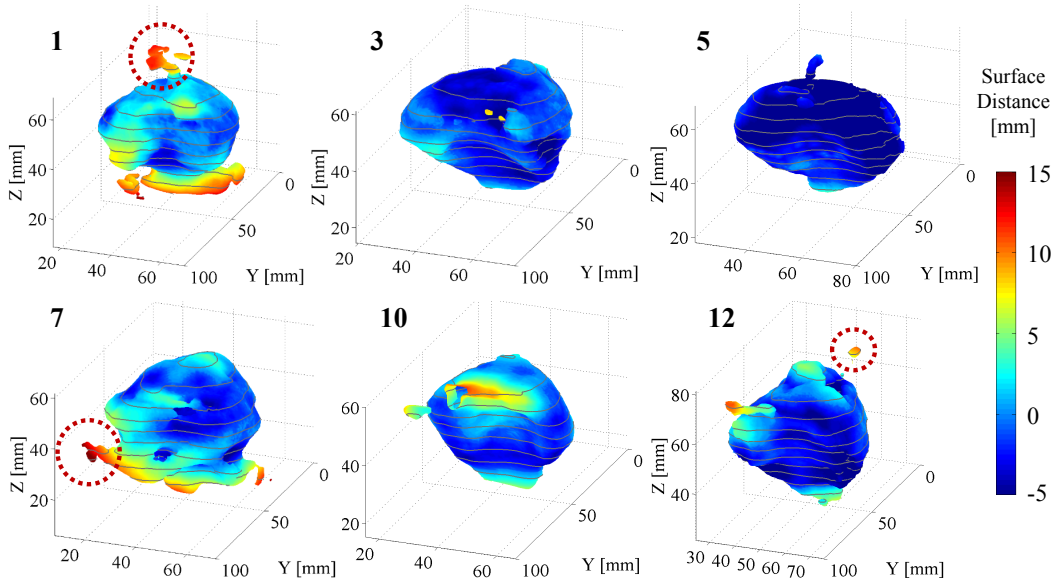

Fig. 2.11. Automatic segmentation results for six representative patient datasets. The color of the segmented surface indicates the distances to the ground truth segmentation according to the color bar on the right hand side. While segmentation artifacts extended far beyond the prostate in some datasets (dashed circles), the majority of the gland is segmented reasonably well. The patient numbers coincide with the ones in Tab. 2.3. *Reprint from [256] with permission of Springer.*

minutes, and segmentation of the unseen dataset – without any optimization – around 3 minutes depending on the size of the image.

2.5.3 Retrospective Evaluation of Surface-based Registration

We evaluated the quality of the automatic registration, obtained using the automatic segmentation result from the previous step, by comparing it with rigid landmark-based registration. As described in Sec. 2.4.3, four corresponding pairs of anatomical landmarks were carefully selected in MRI and TRUS for each patient, and rigid transformation matrices were obtained using the Umeyama method [226]. The experts performing these annotations were instructed to select landmarks from all parts of the gland such that the overall match would be optimal. In Fig. 2.12, rigid and deformable axial slice registration results are presented. In addition, we computed the target registration errors (TRE) for both methods and reported mean and maximum for each patient in Tab. 2.3. Overall, the presented method was able to improve the TRE in all but two cases, from an average of 2.60 to 2.00 mm. In all cases, including these two, the TRE did not exceed 3.55 mm.

2.5.4 Statistical Deformation Model

On the basis of all 50 patient datasets, an SDM was generated as described in Sec. 2.3.4, using $M = 81$ rays and the automatic TRUS segmentations only. Fig. 2.13 shows the mean deformation, and the first three eigenmodes are depicted in Fig. 2.14. While all three induce a volume change, compression is most prominent in the second mode (95% variation between

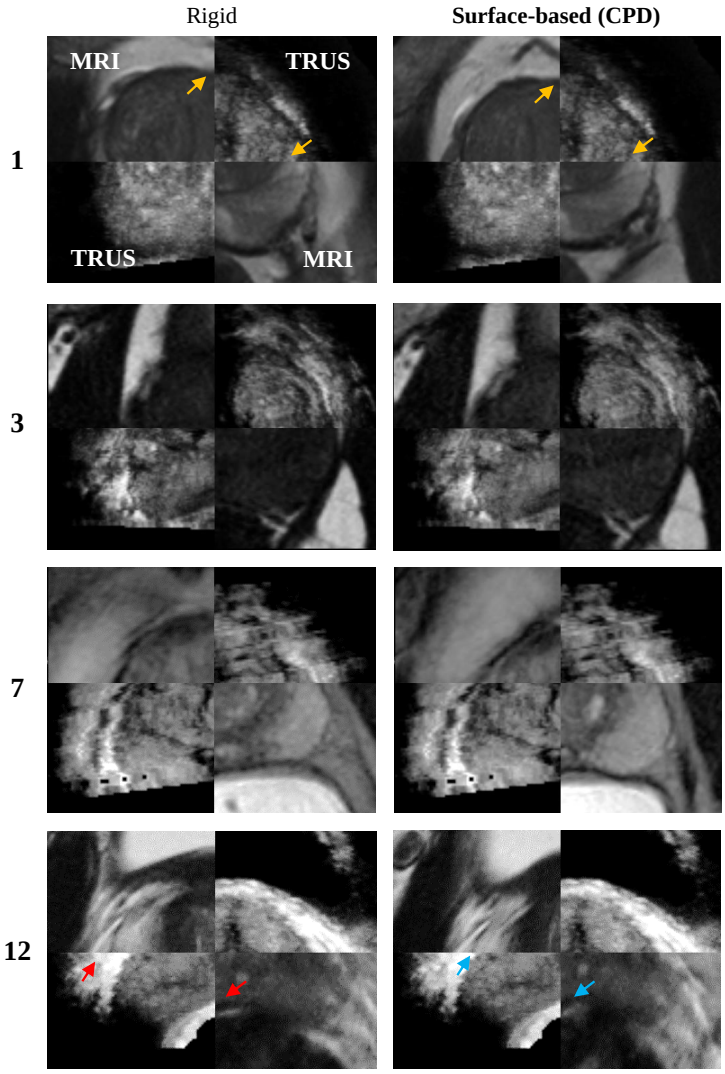

Fig. 2.12. Axial slice comparison between landmark-based rigid registration (*left column*) and deformable surface-based registration (*right column*) for four representative patients. In general, the deformable registration shows a good alignment between both images and captures the non-linear deformation better than the rigid deformation model. The patient numbers coincide with the ones in Tab. 2.3. *Reprint from [256] with permission of Springer.*

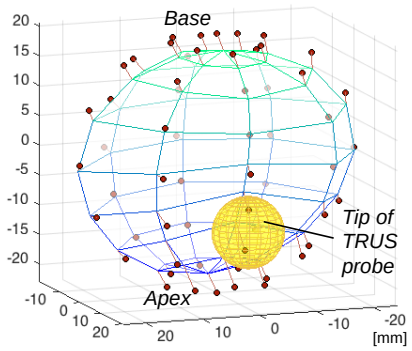

Fig. 2.13. Mean prostate shape **before** (*red vertices*) and after (*mesh*) SDM mean deformation (*red lines*), showing the compression induced by ultrasound transducer (*yellow sphere*). *Reprint from [255], ©2017 IEEE.*

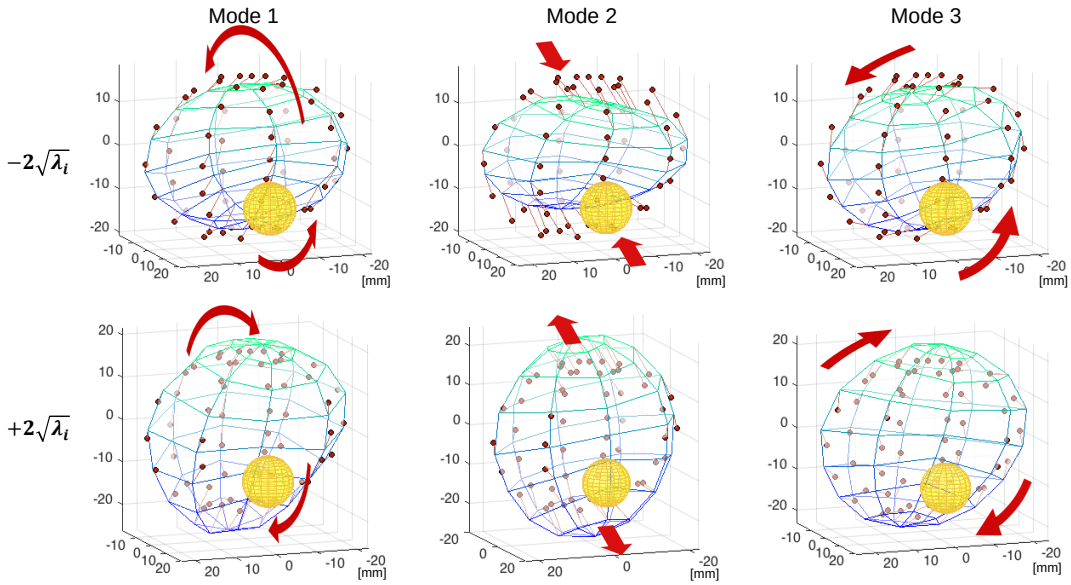

Fig. 2.14. Deformation caused by the first three SDM eigenmodes (*red lines*). For reference, the red vertices show the mean prostate **after** average deformation. Red arrows indicate the main direction of deformation (rotation around the left-right axis, compression in the direction of transducer insertion, and rotation around the cranio-caudal axis, respectively). *Reprint from [255], ©2017 IEEE.*

$\pm 2\sqrt{\lambda}$). Already the first 9 eigenmodes were found to explain 92% of the variation in the dataset.

2.5.5 Retrospective Evaluation of Intensity-based Registration

The intensity-based approach was validated on the 10 fully annotated datasets (*Part 3*, see Sec. 2.4.3) using the target registration error (TRE) of the landmark points. For each of these experiments, the patient dataset under investigation was excluded from the generation of the SDM (leave-one-out cross validation). On these fully annotated datasets, the Dice score between automatic and manual ultrasound segmentation in the mid-gland region was found to be 86.2±3.5, which is similar to the ones reported in Sec. 2.5.2 and in [60].

Table 2.4 lists TREs for rigid registration, surface-based registration as well as the intensity-based method in three different variants. For each method, not only the average TRE for all landmarks is provided, but also the error for the single landmark close to the lesion $\mathbf{t}$ identified in PET (see Sec. 2.4.3). The rigid registration (Tab. 2.4a) was computed purely based on the four boundary landmarks using the Umeyama method. The average TRE evaluated on all six landmarks was 3.20±1.33 mm, but yielded errors of up to 16.2 mm for individual landmarks, in particular for the two not included in the computation of the transformation. In Tab. 2.4b, we report TREs for the surface-based registration method (Sec. 2.3.3). The average TRE for all landmarks amounted to 4.40±1.95 mm.

Finally, the proposed intensity-based registration method is evaluated. We optimized for the first $L = 9$ eigenmodes, together with θ_0 leading to a total of 10 parameters. Empirically identified parameters $d_0 = 40$ and $k = 0.001$ were used for preconditioning, providing an

	Experiment	TRE for	1	2	3	4	5	6	7	8	9	10	$\mu \pm \sigma$
a)	Rigid	All LMs	3.63	6.17	1.63	3.18	4.74	2.59	2.41	1.99	3.56	2.14	3.20±1.33
		Lesion	7.13	16.2	1.99	3.91	4.42	1.47	3.16	1.16	4.91	3.44	4.78±4.16
b)	Surface-based	All LMs	3.09	2.34	1.35	8.40	3.01	4.65	4.50	6.26	5.46	4.91	4.40±1.95
		Lesion	0.32	2.03	2.48	7.96	1.60	1.65	4.33	2.64	2.36	5.99	3.14±2.19
c)	Without	All LMs	2.85	2.65	4.26	2.31	4.75	3.40	2.83	1.97	2.60	3.35	3.10±0.82
	Precond.	Lesion	2.16	3.25	1.20	1.48	1.84	2.37	1.61	0.79	2.13	1.73	**1.86±0.64**
	Isotropic	All LMs	1.88	3.64	2.87	2.39	4.71	3.07	3.03	2.08	6.26	3.48	3.34±1.25
	Precond.	Lesion	2.28	2.37	1.27	1.31	1.54	2.15	1.20	1.14	1.74	1.87	**1.69±0.44**
	Anisotropic	All LMs	2.71	3.90	2.96	2.36	5.03	3.91	2.85	2.68	5.42	2.79	3.46±1.01
	Precond.	Lesions	0.84	2.47	1.31	1.24	1.08	2.39	1.18	0.48	2.06	1.09	**1.41±0.63**

Tab. 2.4. Average target registration errors (TRE) in mm for **a)** rigid registration purely based on four boundary landmarks (LM), and **b)** surface-based registration. **c)** Effect of isotropic and anisotropic preconditioning onto TRE for landmark placed close to **t** (lesion). The patient numbers coincide with the ones in Fig. 2.15. *Reprint from [255], ©2017 IEEE.*

appropriate balance between the weighting around the lesion and the capture range of the similarity metric. Representative examples of results are illustrated in Fig. 2.15a-d. The effect of the preconditioning is shown in Fig. 2.15e-h. In Tab. 2.4c, TREs are reported for SDM-based registration without preconditioning, isotropic preconditioning (ζ_{iso}), and anisotropic preconditioning (ζ_{aniso}). TREs of the lesion-specific landmarks were lowest with 1.41 mm for anisotropic preconditioning. The median TREs were 1.79, 1.64, and 1.21 mm, respectively. Using the ground truth annotations, prostate surface meshes were created for MRI and TRUS. After our registration pipeline, the average Hausdorff distance between the two meshes was 1.84±0.59 mm.

The TREs of the lesion-specific landmarks are also reported in Fig. 2.16 for all evaluated methods. Paired Kolmogorov-Smirnov tests indicate that all SDM-based methods performed significantly better than rigid registration ($p < 0.05$), and that the proposed SDM registration with anisotropic preconditioning performs significantly better than surface-based registration. Our optimized GPU implementation yielded an average runtime for the entire registration pipeline of 17.5 seconds, roughly 6× faster than the method described in [213].

2.5.6 Clinical Histology Evaluation

Since a detailed medical discussion of the findings for all patients would exceed the scope of this thesis, results for five patients are presented where available histological evaluations indicated pathological findings (*Part 2a*, see Sec. 2.4.3). A full evaluation discussing the identification of suspicious lesions on PET/MRI and the detection rates of the biopsy guidance system will be the subject of a medical publication currently in preparation.

First, we report the PSA (Prostate-Specific Antigen) blood level prior to the biopsy and the location of the suspicious lesion identified in PET/MRI in Tab. 2.5. In addition, the number of positive and total biopsy cores for random and targeted biopsies, respectively, and the

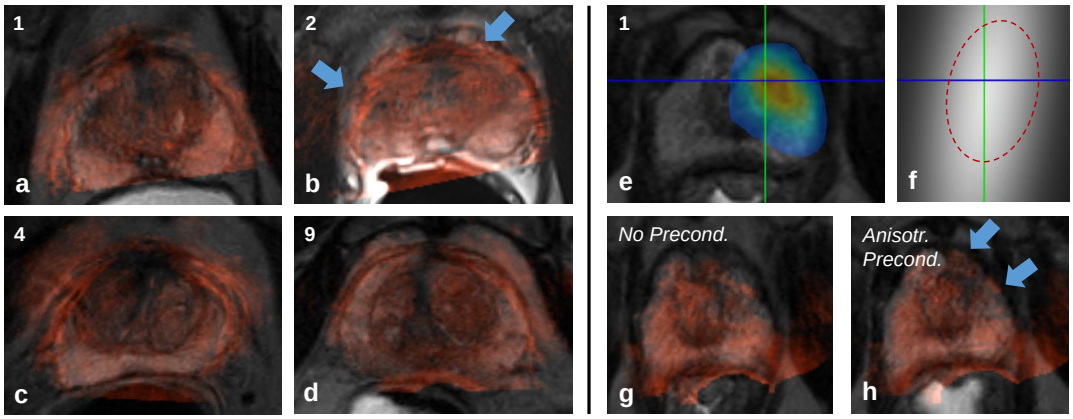

Fig. 2.15. **a-d)** Registration result for four example datasets shown as axial slice overlays of the TRUS image (*red*) onto warped MR (*greyscale*). For Pat. 2, the registration converged to a local minimum (*blue arrows*). **e)** MRI with PET overlay for one patient in the apical region, with selection lesion **t**. **f)** Preconditioning weight map ζ_{aniso}, with estimated main directions of deformation (*ellipse*), allowing to improve registration without preconditioning (**g**) compared to the proposed anisotropic preconditioning (**h**, *blue arrows*). The patient numbers coincide with the ones in Tab. 2.4. *Reprint from [255], ©2017 IEEE.*

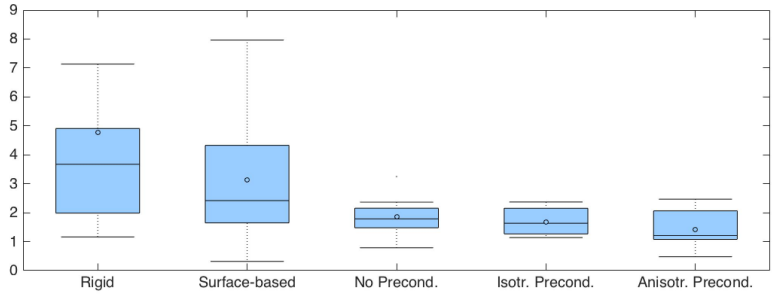

Fig. 2.16. Target registration errors (TRE) in mm for the landmarks placed close to **t** (lesion). The intensity-based method outperforms rigid and surface-based registration, and the anisotropic preconditioning can improve the image alignment around critical structures. *Reprint from [255], ©2017 IEEE.*

final Gleason score are reported. In all cases, at least one of the targeted biopsy samples was positive, except for patient 4, where the only positive randomized biopsy core showed pathological findings in a different area than the targeted one.

2.6 Discussion

In this chapter, a multi-modal image-guided biopsy framework was presented, which combines PET/MR images with interventional TRUS. Novel algorithms to solve for the challenging registration between MRI and TRUS using either surface-based or intensity-based techniques were proposed.

Prostate experiments early validated the suitability of the developed methodology because all lesions were successfully sampled at least once. The few individual samples that missed the lesion were possibly due to a different pressure of the transducer on the prostate compared

Patient	PSA [ng/ml]	Suspicious Lesion	Random Biopsies	Target Biopsies	Gleason Score
1	9.8	left apical	3/12	1/1	3+4=7
4	10.0	right median basal	1/11	0/1	3+3=6
7	6.0	right central lateral	5/12	1/1	3+4=7
10	14.5	right median central	3/12	1/3	3+3=6
12	7.5	right central lateral	6/12	3/4	3+4=7

Tab. 2.5. Histology results for five cases are reported as number of positive cores / total number of cores for random and targeted biopsies, respectively. For all but one patient, at least one targeted sample was positive, indicating that the system is able to map suspicious lesions. *Reprint from [256] with permission of Springer.*

to the initial TRUS image acquisition, leading to a distorted TRUS-MRI registration. This influenced the guidelines for the TRUS acquisition protocol for subsequent experiments.

The accuracy of the surface-based registration method is limited by the accuracy of the automatic, Hough forest-based segmentation algorithm. Here, achieved average surface distances of around 1 mm are within the clinically acceptable accuracy [18]. The algorithm produced in some cases artifacts stretching beyond the ground truth segmentation, having a clear impact on the surface distance (see Fig. 2.11, patients 1, 7, and 12). However, it has been previously shown that the registration algorithm is robust in terms of outliers and able to cope with such artifacts as long as the majority of the surface is appropriately well captured [148]. The presented Dice scores for all datasets indicate potential for further improvements, also by using an extended and more diverse training dataset in the future. Hopefully, this will allow a better capture range outside of the *mid-gland* region, which is commonly considered today [243].

Surface-based registration results overall support the suitability of the method for the intended usage in clinical routine. In some of the investigated cases, little deformation was evident, resulting in low registration errors for both methods and good agreement of the fused images, for instance in patient 1 (Fig. 2.12, *yellow arrows*). The benefit of the proposed method is clearly visible in cases with severe deformations, for instance in patient 12. Here, rigid landmark registration errors of up to eight millimeters were obtained, causing significant misalignment of both images (*red arrows*). With the deformable registration approach, a much better agreement between the surfaces was achieved (*blue arrows*), decreasing the mean TRE to less than half. Because the registration result was slightly worse in two cases, urologists retain the option to choose between the rigid and elastic surface-based registration methods in the workflow implemented in clinical routine today. This way, they can override the automatic default algorithm and adjust the registration in case of obvious misalignment. In the future, more reliable segmentation techniques might reduce such occurrences, as discussed above.

Closely related to these issues are the presented developments on intensity-based elastic registration, aiming to reduce the dependence on segmentation algorithms in the first place. The generated statistical deformation model confirmed our hypothesis that despite the variation in prostate size, the expected deformations during a biopsy session compared to the patient's pose in the MRI scanner are quite homogeneous. We can therefore assume that an SDM is

able to capture the patient-specific deformations reasonably well, as is thus a suitable tool to parametrize a transformation model for registration schemes.

Using landmark points within the prostate and in the vicinity of lesions, we were able to confirm that rigid registration may lead to errors of more than 1.5 cm, which is inappropriate for this clinical application. Surface-based registration errors were better than the rigid fusion in many cases, but the method produced severe misalignments in the few cases where it failed, mostly due to inaccurate automatic segmentation. In either case, results on interior landmarks indicate that purely feature-based approaches are not able to capture the deformations within the gland accurately. In contrast, better results were obtained using intensity-based registration, even without preconditioning.

The obtained results do not only show that an SDM can be used to elastically register unseen patient MR and TRUS images, but also that inaccuracies in the segmentation process, as required for SDM generation using a large dataset, can be overcome by the optimizer. Errors obtained using the proposed method were in the range of the experiments conducted in [60, 213]. As expected, anisotropic preconditioning performs better than isotropic preconditioning or no preconditioning at all, and significantly improves registration accuracy compared to the other methods investigated. The average landmarks error also indicates that improvements toward the critical lesion for interventional guidance might affect the registration accuracy at other locations. For targeted prostate biopsies, where urologists predominantly focus on the suspicious lesions in the scope of multi-modal image guidance, this might be an acceptable trade-off. In terms of computational aspects, the proposed algorithm is well suited for the tight time requirements of clinical routine.

The clinical histology evaluation finally validates the effectiveness of the proposed multi-modal biopsy guidance framework under real conditions. The presented system was able to successfully map suspicious regions from PET/MRI to the interventional TRUS image, and facilitated the diagnosis of prostate cancer.

Altogether, the presented guidance system was eagerly accepted in the urology department of Klinikum Rechts der Isar. Rapid introduction into weekly clinical use was facilitated by the fact that additional steps of the system did not interrupt the established clinical procedure as the urologist could prepare the local anesthesia in the meantime. As mentioned in Sec. 2.2.3, the system has already been used in biopsy sessions of more than 300 patients, and it will continue to support urologists in performing the intervention in the course of ongoing studies.

3 Visual Servoing-based Robotic Ultrasound

3.1 Overview and Publications

This chapter presents the contributions toward robotically assisted interventional ultrasound. Intelligent guidance applications based on ultrasound for interventional scenarios do not only require appropriate robot control laws for scans with sufficient quality but also necessitate certain levels of automation to facilitate navigation and execution of pre-interventional plans by the surgeon. After describing novel aspects of automated 3D ultrasound acquisitions, this chapter therefore focuses on innovative concepts of ultrasound-based visual servoing, i.e. control laws where features extracted from sonographic images directly serve as input variables. Hereby, this thesis emphasizes the suitability of continuous re-registration using image intensities for automatic motion compensation and tool alignment.

While the presented methods are generic and can be potentially applied to a wide range of clinical tasks, two important procedures were chosen to demonstrate the capabilities of the proposed methods. First, needle guidance for facet joint injections is investigated, potentially improving the accuracy of manual targeting while at the same time reducing X-ray exposure for patient and staff to a minimum. Second, the achievement of a fully automatic screening system for abdominal aortic aneurysms confirms that robotic ultrasound is useful for vascular monitoring. This opens the field of minimally invasive procedures with crucial blood vessel involvement in the future. Together, both applications exemplify how the described concepts can be adapted to and implemented for intervention-specific scenarios, and allow to medically validate the achievable performance of the proposed framework within realistic environments.

This chapter is organized as follows: In Sec. 3.2, the clinical and methodological background on the two chosen fields motivates the application of robotic ultrasound and summarizes prior art as well as this work's contributions. Thereafter, Sec. 3.3 establishes the necessary foundation for automated, constant-force ultrasound acquisitions and describes the developed visual servoing techniques in detail. Depending on the targeted clinical application, the respective components are then validated within the scope of phantom or human volunteer experiments, which are explained in detail in Sec. 3.4. Results are reported in Sec. 3.5 and thoroughly discussed in Sec. 3.6.

Substantial parts of this chapter have already been published and are quoted verbatim. Relevant parts regarding the neurosurgical application of facet joint needle insertions are incorporated in the following article (especially sections 3.2.3, 3.3.1-3, 3.3.5-6, 3.4.1-4, 3.5.1-2, 3.6):

[249] **Oliver Zettinig***, Benjamin Frisch*, Salvatore Virga, Marco Esposito, Anna Rienmüller, Bernhard Meyer, Christoph Hennersperger, Yu-Mi Ryang*, and Nassir Navab*. "3D ultrasound registration-based visual servoing for neurosurgical navigation". In: *International Journal of Computer Assisted Radiology and Surgery* 12.9 (Sept. 2017), pp. 1607–1619. DOI: 10.1007/s11548-017-1536-2

* *These authors contributed equally to this work, respectively.*

The article above is based on the following, already published work, where preliminary results on needle guidance were presented:

[250] **Oliver Zettinig**, Bernhard Fuerst, Risto Kojcev, Marco Esposito, Mehrdad Salehi, Wolfgang Wein, Julia Rackerseder, Edoardo Sinibaldi, Benjamin Frisch, and Nassir Navab. "Toward real-time 3D ultrasound registration-based visual servoing for interventional navigation". In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2016, pp. 945–950. DOI: 10.1109/ICRA.2016.7487226

All relevant additions to the robot control architecture for automatic aneurysms screening are published in the following article (especially sections 3.2.3, 3.3.4, 3.3.7-9, 3.4.5, 3.5.3, 3.6):

[230] Salvatore Virga*, **Oliver Zettinig***, Marco Esposito, Karin Pfister, Benjamin Frisch, Thomas Neff, Nassir Navab, and Christoph Hennersperger. "Automatic force-compliant robotic ultrasound screening of abdominal aortic aneurysms". In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Oct. 2016, pp. 508–513. DOI: 10.1109/IROS.2016.7759101

* *These authors contributed equally to this work.*

3.2 Introduction

3.2.1 Clinical Background

For multiple medical indications in the scope of neurosurgery, for instance needle injections into various joints [65, 205], or general vascular conditions such as suspected stenoses or peripheral artery disease, ultrasound is already an established first-line imaging modality [156, 241]. This section outlines the clinical background for the two chosen applications and describes how ultrasound imaging relates to the gold standard procedure today.

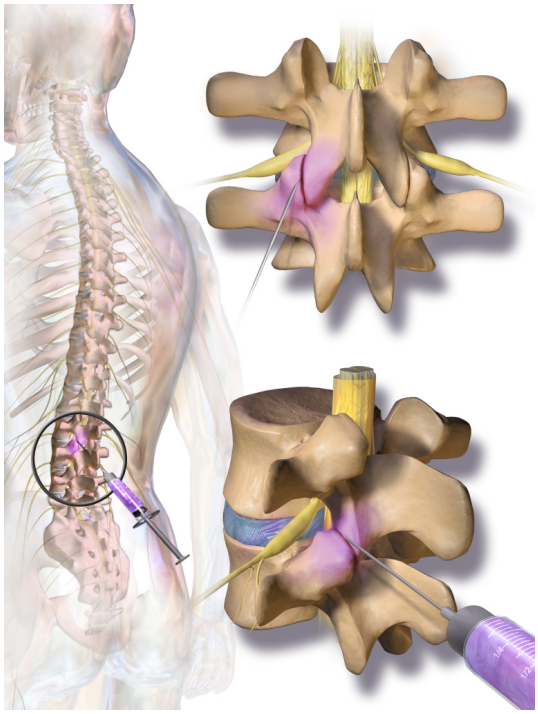

Fig. 3.1. Schematic illustration of a facet joint injection in the lumbar spine. *From [19], CC-BY-3.0.*

Neurosurgery: Lumbar Spine Facet Joint Injections

The first application focuses on facet joint injections in the lumbar spine, which is illustrated in Fig. 3.1. Chronic low back pain originating from this part of the spine affects a large part of the adult population, with a reported overall one-year prevalence of 38% [130]. In general, adults over the age of 45 are more likely to experience back pain than younger ones, with a slight drop in oldest cohorts. In terms of genders, women seem to be affected more often than men. Following the classification of Cassidy et al. [32], chronic Grade I-IV low back pain is distinguished based on severeness of pain and caused disability: Grade I describes low-intensity/low-disability low back pain (49%), Grade II high-intensity/low-disability low back pain (12%), and Grades III and IV describes high-intensity/high-disability low back pain (together 11%). The rest (28%) did not encounter chronic pain [130]. In the last 15 years, prevalence has been reported to steadily increase. In the vast majority of affected individuals (80% to 90%), attacks resolve in up to six weeks, but frequent relapses lead to the experience of multiple episodes in roughly a third of patients. Around 5% to 10% of patients develop persistent low back pain. A variety of comorbid, psychological, occupational, lifestyle and social demographic factors have been associated to the disease [130].

Low back pain may be caused by a large number of underlying complications, including but not limited to muscular or skeletal issues such as sprains or strains, osteoarthritis, degeneration of the discs between the vertebrae, spinal disc herniation, vertebra fractures, or, rarely, an infection or tumor of the spine [31]. Using double diagnostic nerve blocks, in up to 45% of chronic low back pain patients, the facet joint as shown in Fig. 3.1 can be identified as the source of pain [129]. Hereby, a nerve block refers to the infiltration of the facet joint nerve with local anesthetic drugs such as lidocaine or bupivacaine. Administered by means of needle injections, they are either performed diagnostically to determine if the joint is a source of pain,

or therapeutically for chronic spinal pain management [33]. In both cases, X-ray fluoroscopy guidance is recommended and considered gold standard [248].

Several studies have reported that US-guided injections show no significant difference in diagnostic accuracy and therapeutic efficacy compared to guidance using fluoroscopy or CT, yet substantially reducing the radiation exposure of both patient and clinicians [9, 65, 205]. Nevertheless, the limited field of view and difficulties in interpreting US images are common obstacles to guarantee accurate needle placement [248].

The manual navigation of the US transducer to the appropriate location and maintenance of a suitable acoustic window are linked to additional challenges. While the cumbersome procedure induces a slow learning curve [46, 144], work-related musculoskeletal disorders have been associated with medical sonographers [59], causing chronic pain to the examiners and impeding the quality of US-based interventions over time. Both of these issues related to the freehand nature of previous US-guided facet joint injections suggest more automated solutions [127].

Vascular Surgery: Abdominal Aortic Aneurysms

Another target area with a high potential benefit from more automated ultrasound scanning procedures is the abdominal aortic aneurysm (AAA), a localized dilation (ballooning) of one of the major vessels in the human body, as illustrated in Fig. 3.2. More specifically, AAA is a "segmental, full-thickness dilatation of the abdominal aorta exceeding the normal vessel diameter by 50%, although an aneurysm diameter of 3.0 cm is commonly regarded as the threshold" [204]. A marked variability in epidemiological reporting has been seen in recent literature, with prevalence varying between 1.7% and 12.7%, only partly depending on the investigated age group and the country of study origin [211]. The major risk of an AAA is

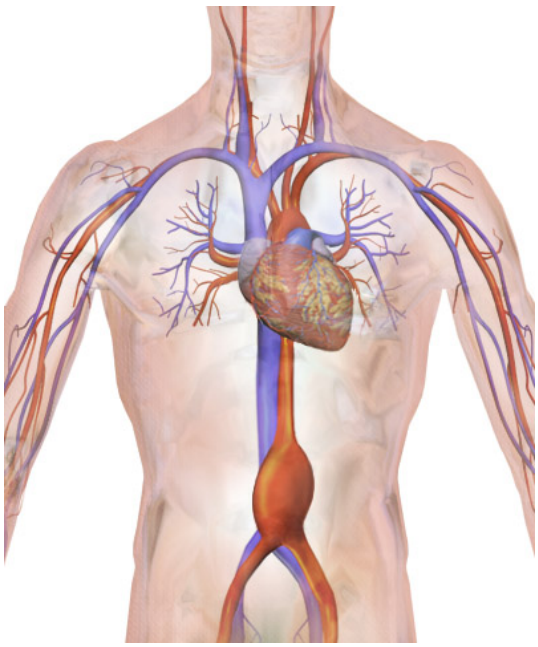

Fig. 3.2. Schematic illustration of an abdominal aortic aneurysm (AAA). *From [19], CC-BY-3.0.*

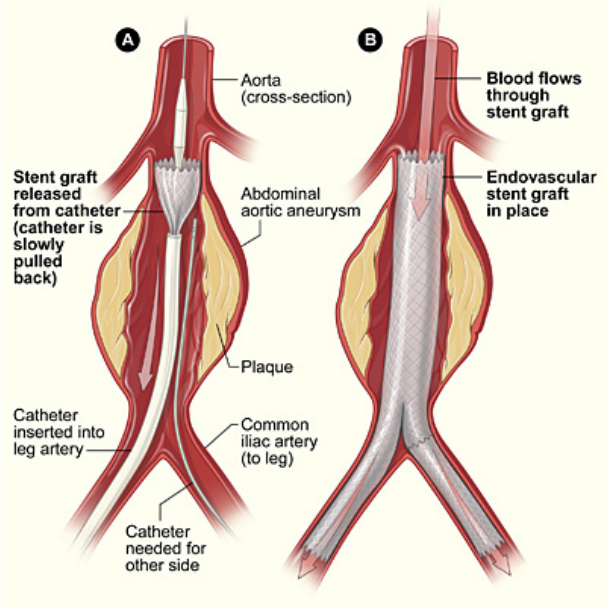

Fig. 3.3. Schematic illustration of an endovascular aortic repair (EVAR) procedure. *From NIH (*`https://www.nhlbi.nih.gov/health/health-topics/topics/arm/`*), public domain.*

the rupture of the aneurysm, which is regularly lethal with mortality rates up to 95% [211]. Of those patients who reach a hospital in time, survival rates remain high between 50% and 70% [204]. The probability for rupture depends on the size, shape and stress of the aneurysm, with a substantially increased risk for diameters above 6 cm [28]. Risk factors include the ones commonly associated with cardiovascular diseases, including smoking, hypertension, and hypercholesterolemia.

Since the benefit of medical treatment with a variety of drugs could not be confirmed, treatment necessitates a surgical intervention with the goal to prevent rupture. Hereby, two techniques for aneurysm repair can be distinguished. Historically, the vessels above and below the aneurysm were controlled in open surgery, and a synthetic graft was inserted into the opened aneurysm sac. Today, more than 75% of patients undergo percutaneous endovascular aortic repair (EVAR) procedures, where one or multiple grafts are inserted through the femoral and iliac arteries (see Fig. 3.3). While both procedures lead to similar long-term mortality rates, the less invasive endovascular technique offers the advantages of significantly reduced hospitalization times and the duration until full recovery [204]. However, the EVAR currently requires a high level of fluoroscopic X-ray guidance.

Ultrasound is already employed as a standard diagnostic tool for the imaging of the aorta, and US-based staging is widely accepted in clinical practice [187]. However, challenges with respect to inter-operator variability and standardized measurement approaches still impair the implementation of national or international sonography-based screening programs [11, 47, 187]. In contrast to conventional clinical 2D ultrasound, 3D US showed an improved localization of aneurysm-shape and endoleak after EVAR compared to 2D US, using contrast-enhanced imaging [165].

Based on the extraction of quantitative values from these data, a 3D ultrasound scanning system with reproducible and constant acquisition quality could facilitate the establishment of screening frameworks. While the patient benefit of an early detection and regular checks for aortic aneurysms is obvious in such a program, discussions about the overall cost-effectiveness of potential screening activities remain [188]. The full automation of the acquisition using robotic technology would reduce personnel costs, a major factor for the data acquisition, while achieving full comparability of acquired data.

3.2.2 Visual Servoing as Robot Control Law

Proceeding the subsequent sections, the fundamental concepts of visual servoing shall be outlined. This introduction follows the excellent two-part tutorial by Chaumette and Hutchinson [39, 40], where more detailed descriptions can be found.

Single Task Visual Servoing

Visual servoing generally refers to a closed control loop that involves a robotic manipulator and a visual sensor such as a camera or a US probe. It aims at minimizing a visual error $\mathbf{e}$ defined as follows:

$$\mathbf{e}(t) = \mathbf{s}(t) - \mathbf{s}^*. \tag{3.1}$$

Hereby, $\mathbf{s}(t) \in \mathbb{R}^k$ denotes a k-dimensional visual feature vector which depends on a set of measurements obtained from the visual sensor, potentially including any prior knowledge, e.g. calibration information. The desired visual features are expressed as $\mathbf{s}^*$. The choice of $\mathbf{s}$, and thus $\mathbf{s}^*$, is highly dependent on the intended task.

To define a simple velocity controller, the relationship between time variation of $\mathbf{s}(t)$ and the velocity of the robot is required. For now, it is helpful to define the velocity as screw, i.e. $\mathbf{v} = (v, \omega) \in \mathbb{R}^6$ with v referring to the linear and ω to the angular velocity components. In addition, let $\mathbf{x}(\mathbf{q})$ be the Cartesian position of the robot end-effector given a vector of joint positions $\mathbf{q}$ (angles for revolute joint) at a given point in time. We can then define:

$$\dot{\mathbf{s}} = \frac{\partial \mathbf{s}}{\partial \mathbf{x}} \mathbf{v} = \mathrm{L}_s \mathbf{v}, \tag{3.2}$$

where $\mathrm{L}_s \in \mathbb{R}^{k \times 6}$ is called the interaction matrix or feature Jacobian matrix. From Eq. 3.1, it is possible to obtain the time variation of the visual error as:

$$\dot{\mathbf{e}} = \mathrm{L}_s \mathbf{v}, \tag{3.3}$$

Considering $\mathbf{v}$ as output, L_s needs to be inverted. In practice, it is commonly neither possible to directly invert L_s nor to compute the Moore-Penrose pseudo-inverse $\mathrm{L}_s^+ = (\mathrm{L}_s^\intercal \mathrm{L}_s)^{-1} \mathrm{L}_s^\intercal$ [39]. In these cases, approximations of L_s^+ are used for the final control law:

$$\mathbf{v} = -\lambda \mathrm{L}_s^+ \mathbf{e}, \tag{3.4}$$

where $\lambda > 0$ is a control gain causing an exponential decoupled decrease of the error, i.e. $\dot{\mathbf{e}} = -\lambda \mathbf{e}$. Conveniently, local asymptotic stability is ensured as long as

$$\mathrm{L}_s \mathrm{L}_s^+ > 0 \tag{3.5}$$

is satisfied. This allows strong assumptions in the estimation of L_s or L_s^+.

In-plane Position Control using a Single Point

A straightforward control law can be defined for geometric visual servoing using a single point $\mathbf{p} = (u, v)^\intercal$ in the 2D ultrasound frame, for instance a wire in a water bath. We assume that for its identification within the image, a fast, suitable algorithm is available. The visual error can then be defined as the displacement from its desired location in the US frame $\mathbf{p}^* = (u^*, v^*)^\intercal$:

$$\mathbf{e} = \begin{bmatrix} u - u^* \\ v - v^* \end{bmatrix}. \tag{3.6}$$

It is furthermore assumed that the coordinate frame used to control the robot is the end-effector frame and that its y- and z-axes are aligned with the US frame so that $\mathbf{p}$ can be expressed in 3D coordinates as:

$$\begin{bmatrix} 0 \\ u \\ v \end{bmatrix} = a \cdot \begin{bmatrix} 0 \\ p_y \\ p_z \end{bmatrix}, \tag{3.7}$$

where a is an image scaling factor. As originally presented in [184], the following interaction matrix can then be derived:

$$\mathrm{L}_s = \begin{bmatrix} 0 & -a & 0 & v & 0 & 0 \\ 0 & 0 & -a & -u & 0 & 0 \end{bmatrix}. \tag{3.8}$$

By applying the control law of Eq. 3.4, the point $\mathbf{p}$ can be brought to its desired location. Extensions to 3D can be implemented analogously. Special care is necessary when multiple tasks are defined. In this case, the *redundancy formalism* as described by Samson et al. [185] is commonly employed. By defining a secondary task in the null-space of a primary one, the system will only aim at minimizing the secondary visual error as long as the primary task is not disturbed.

In this chapter, only position control is considered, i.e. the algorithms presented in Sec. 3.3 will estimate desired poses, i.e. desired positions and orientations, of the robot end-effector in some external (world) coordinate system. The necessary velocities to reach these poses are then indirectly estimated as described above.

3.2.3 From Ultrasound-Assisted Guidance to Robotic Visual Servoing

In the past decades, numerous methods for more operator-independent ultrasound acquisitions and guidance techniques based thereupon have been presented. This section attempts to summarize the most important milestones in this direction by categorizing with respect to the employed degree of automation.

Non-Robotized Approaches

As pointed out in Sec. 1.3.4, 3D freehand techniques have been widely adopted to enhance the possibilities of planar 2D sonography. Some of these developments were dedicated to certain interventional neurosurgical applications. The simplest form of ultrasound-guided needle insertion is to superimpose the needle path, which is calibrated to a guide rigidly attached to the transducer, on the live 2D image. With such a setup, Tran et al. [222] first achieved a real time 2D guidance framework not in need of a second operator. Brudfors et al. [29] later successfully extended the system to allow for 3D guidance as required for navigational tasks, especially on the challenging spine anatomy.

Since such approaches do not alleviate limited target visibility and poor contrast issues, tracking systems were utilized to perform freehand 3D acquisitions, as mentioned above, that can be registered to pre-operative imaging data such as CT [245]. Assuming that the patient has not moved, features such as vertebrae contours can be then highlighted on the live US image along with the insertion path of a tool, which is also coupled with a tracking target. Moore et al. [142] were the first ones to present both the live US frame and the tracked needle accurately within a 3D augmented reality view along with a high-resolution geometric model of the spine to the physician. In a preliminary study, the proposed system was found to not only reduce vertebra level confusion (needle inserted correctly but into the wrong joint) but also increase needle placement accuracy significantly. In a similar approach, Ungi et al. [227] recorded tracked US snapshots, which are visualized in a virtual 3D environment and directly serve as basis for guidance of a needle that, again, also was equipped with a tracking target.

When pre-interventional images are not available, registration can also be established with statistical spine atlases, learned from pre-operative MR or CT images over a range of the population. Rasoulian et al. [173] have shown that such atlases allow for a sufficiently accurate model-based registration to provide enough context for successful needle insertions. A further extension superseded the need for a tracking system by falling back to an attached needle guide and employing a motorized transducer, directly enabling 3D US acquisitions [29]. In return for a greatly reduced frame rate on the insertion plane, this system provides updated anatomy highlighting also if either transducer or patient move.

However, none of the aforementioned works eliminate the need for manually navigating to the site of interest, finding a suitable acoustic window, and maintaining sufficient image quality with the handheld US probe. In not purely diagnostic scenarios, the observation of the tool, e.g. a needle, in the US frame has furthermore been described as cumbersome [248].

Robotic Ultrasound Systems

To overcome at least some of the mentioned limitations, robotic assistance for US acquisitions has been studied in a broad variety of previous works. Taylor and Stoianovici [219] introduced the terms CAD/CAM systems and surgical assistants to distinguish two general classes of medical robotic systems:

i) **CAD/CAM systems**, named in analogy with computer-integrated design and manufacturing, emphasizes a paradigm of an integrated surgical workflow incorporating planning, execution and follow-up. In this context, where a set of pre- and intraoperative imaging and detailed planning information are regularly available, key steps in the procedure are executed with the help of a robot, reaching levels of precision otherwise unachievable with conventional techniques. A notable example is the ROBODOC system for use in total hip arthroplasty [218].

ii) **Intraoperative surgical assistants**, on the other hand, aim at providing surgeons with a new set of tools that extend their capabilities in treating patients. Many of such systems are designed for minimally invasive environments, greatly increasing both accessibility of anatomy and dexterity of surgical tools. The most prominent example of this category is the da Vinci® Surgical System (Intuitive Surgical, Inc., Sunnyvale, CA, USA), featuring cable-driven laparoscopic manipulators and a master console for tele-operation [89]. Ultrasound guidance in combination with surgical assistants has been implemented in numerous works, most notably for TRUS-based surgeries on the prostate [170], and for pick-up transducers that can be gripped and manipulated by da Vinci graspers [30].

In both cases, US imaging can contribute greatly toward improved clinical outcomes. While in the former class of systems, the registration between patient and planning data might be facilitated, raising in turn targeting accuracy, real-time imaging and the visualization of hidden structures may cause tremendous benefit in the latter case [170]. Yet, a key factor distinguishing these two categories is the degree of automation. By design, CAD/CAM system are intended to perform at least parts of a desired diagnostic or interventional procedure autonomously, requiring a certain level of awareness of the task and the system's surrounding environment [97]. In contrast, surgical assistant frameworks generally emphasize their purely supportive nature for human surgeons, which might be limited to situation-independent aid such as hand tremor suppression [219].

Therefore, the remainder of this section focuses on CAD/CAM systems. We previously proposed to further classify into *Automatic Robotic Support Systems* (ARSS), which concerns the lumbar spine needle guidance application, and *Automatic Data Acquisition Systems* (ADAS), which covers the aneurysm screening application [97].

Automatic Robotic Support Systems provide automatic support for a particular, well-defined task. As outlined in the excellent review by Priester et al. [170], a significant portion of such systems tackle accurate needle placement as required in many clinical interventions, most notably biopsies, central venous access, local anesthesia, brachytherapy and thermal ablation. In the case of biopsy assistance, almost all identified studies rely on the same general procedure. In the first step, 2D or 3D US images of the target region of interest are acquired. Then, a target point (and an insertion point, if applicable) is identified in the patient coordinate system

either by manual annotation or using a registration to a pre-interventional plan. Finally, the needle is aligned as intended and ultimately inserted.

In an early work, Megali et al. [136] marked a target in a custom-made phantom using 3D US and automatically aligned a robotic needle holder, achieving positioning errors of around 2 mm. Kettenbach et al. [114] registered their robot with seven degrees of freedom to a hand-held US probe and showed, that accuracies in the millimeter-scale are also possible with only 2D images. By combining optical tracking, 3D compounded US, and a robotized needle holder, Freschi et al. [66] were able to implement a cooperative control mode, where the physician could trigger trajectory corrections by manipulating a 3D virtual scene. In their originally presented setup, Boctor et al. [21] used two robots for US-guided liver ablation, one holding a US transducer, the other one the thermal needle. Although reaching promising errors, the first arm was later replaced in favor of a freehand 3D US setup for improved usability [22].

A subset of reported efforts specialized on particular anatomies such as breast and prostate. Intending to avoid tumor misplacements between imaging and needle insertion for breast biopsies, Mallapragada et al. [128] applied controlled external forces using a robotic system based on imaging data. Results of phantom experiments indicated that repositioning instead of needle steering is a suitable way to guarantee accurate needle placement. Other works on breast focus more on resolving the anatomy-specific difficulties, for instance by designing dedicated automated scanning equipment for the pendant breast in a water tank, facilitating lesion visibility for robotic biopsies [157].

Prostate biopsies, regularly performed under TRUS guidance (see Sec. 2.2.1), also were the topic of robotic support system developments [170]. Schneider et al. [189] proposed a partially motorized (three degrees of freedom) robotic needle insertion probe that used novel parametric curves as needle path. Instead of transrectal needle entry, Ho et al. [101] chose transperineal access with multiple entry points to avoid obstruction and internal damage. The transperineal route also opens the application to robotic brachytherapy, i.e. the automated implantation of radioactive sources into the prostate as means of cancer treatment [62, 158]. Because efforts in this area are rather specific to the urological realm, the reader is referred to [170] for an extensive review.

In the field of spine interventions, several robotic systems have been presented [15]. However, the majority of these do not rely on ultrasonic guidance but either use X-ray fluoroscopy or utilize optical tracking systems to maintain relative alignment with the target. A prominent example of these systems is the SpineAssist/ Renaissance®system (Mazor Robotics Ltd., Caesarea, Israel), a small hexapod with six degrees of freedom allowing alignment according to planned trajectories in pre-operative CT images. While the accuracy of such systems per se is undisputed, complications in robot attachment, tool canula stability and fluoroscopy-based (re-)registration, e.g. due to imaging artifacts after partial implantation, have prevented general acceptance and recommendation in clinical guidelines [179].

In contrast, Weber et al. [235] presented an intra-operative navigation framework that depends on 2D ultrasound images. The setup consisted of a parallel robot equipped with a linear stage for tool or needle insertion, and a static holder for the US transducer. By means of feature-

based 2D-3D registration after edge detection, alignment between US and pre-operative CT or MR images (and planned trajectories defined therein) was established. Calibration between the robot and the other coordinate systems was achieved using an optical or electromagnetic tracking system. While preliminary screw insertion results revealed an overall insufficient accuracy in many cases, the study confirmed that robotic ultrasound-based guidance is possible for spine applications.

Automatic Data Acquisition Systems originated from tele-operated systems for US-based diagnosis [97, 170]. In particular, a lack of experienced sonographers in rural areas motivated master-slave systems for remote scans. An early system presented by Salcudean et al. [184] featured a robot with a parallel linkage structure resulting in six degrees of freedom, which had counterbalances and backdrivable joints to satisfy safety concerns. In later extensions, visual servoing approaches for acquisition started to emerge in the community. Abolmaesumi et al. [4] tracked the location of an artery in the US image and adjusted the transducer's position accordingly to keep it centered. Robot control was thus shared between the human operator and the visual adjustment algorithm, each being responsible for different degrees of freedom.

Robotic systems have also been used to facilitate ultrasound tomography [52]. Aalamifar et al. [1] used a setup with two opposing US transducers, of which the first one was hand-held and optically tracked, and the other one automatically aligned by a robotic manipulator. The concept was later extended to work with a hand-held TRUS probe coupled with an abdominal transducer held by the robot. This way, the authors were able to generate quantitative speed-of-sound maps of the prostate [2]. Because of the difficulty in alignment and anatomical limits in possible depths of US transmission, similar setups are however limited to imaging of smaller structures such as the prostate or the breast.

The approaches described above operate either directly under guidance of a human operator or semi-automatically. In any case, manual setup, alignment, safe approach of the patient, and definition of an appropriate working environment has relied on human interaction. As a first step, Onogi et al. [161] were able to autonomously acquire B-mode scans of a constrained volume of interest with their pneumatically steered probe holder. A more autonomous system was proposed by Merouche et al. [137]. Their robotic system provides a compounded 3D volume of parts of the lower limb arterial tree. To follow a manually selected vessel, tracking algorithms were implemented to follow the artery under investigation.

Recently, first systems for fully autonomous data acquisitions reached wider interest. In a study focusing on liver screening, Mustafa et al. [147] utilized an RGB camera to detect the region of interest on the patient and followed a predefined protocol of various scanning directions to automatically acquire a set of meaningful 2D images. Nevertheless, their system does not implement 3D reconstructions and is dependent on classified body features which might lead to misaligned 2D imaging planes, limiting the applicability for diagnostic use. Focusing on a different aspect, Graumann et al. [87] proposed a trajectory planning algorithm for automatic 3D acquisitions to cover predefined volumes of interest. The authors employed an RGB-D camera to register the subject with previously acquired tomographic imaging data, which might not necessarily be available in screening applications.

Visual Servoing using Ultrasound

To achieve their objectives, several of the aforementioned robotic systems implemented some kind of visual servoing technique, i.e. they extracted features from live US images and defined control laws thereupon. Looking on prior art where visual servoing forms an essential component of the presented methodology, one can distinguish between **eye-to-hand** and **eye-in-hand** approaches [39].

Eye-to-hand techniques are characterized by a fixed sensor which observes an externally actuated robot end-effector or a tool attached to it. Most prominently, this setup refers to an ultrasound probe observing the insertion of a robotized needle into tissue. In the case of straight needles, such as in a system proposed by Boctor et al. [22], visual servoing is mostly used for initial alignment of the needle. The more challenging case includes steerable needles inserted deep into biological tissue. Adebar et al. [7] used high-frequency vibration to detect the needle using Doppler modes in 3D US. Krupa et al. [122] reached sub-millimeter targeting accuracy using two orthogonal views. The work was later extended to support duty-cycling using full 3D volumes [36]. Abayazid et al. [3] showed that also moving targets can be aimed at using US guidance. Visual servoing is not limited to only detecting the needle but can also incorporate tissue tracking. Nadeau et al. [150] proposed a method for simultaneous tracking of both anatomy and surgical instruments. While not directly the focus of this chapter, these works are relevant because the developed control schemes can commonly be translated to eye-in-hand scenarios. The interested reader can find an extensive review on US-based visual servoing for needle steering in [35].

Eye-in-hand approaches use a sensor that is directly attached to robot. In this case, visual servoing regularly consists of optimizing the robot pose toward a desired view of an object of interest. Abolmaesumi et al. [5] were among the first ones to incorporate vessel detection algorithms directly into robot control laws. Nakadate et al. [152] were able to extend carotid artery vessel tracking to also cope with out-of-plane motion. By exploiting characteristic speckle patterns in US images, Krupa et al. [123] explored motion compensation both in-plane and out-of-plane. Organ motion compensation was also studied by Nadeau et al. [149]. For the first time directly working with image intensities rather than extracted features, the authors proved the applicability to compensate for periodic motions. However, intensity-based 3D-to-3D volume registration for the purpose of US transducer visual servoing has not yet been achieved.

Apart from B-mode images, visual servoing based on Doppler information has also been investigated. Fröhlich et al. [69] proposed a robotic system to aid the resection of the internal mammary artery as required for certain coronary artery bypass procedures. Exploiting color Doppler information, the system identified, tracked and marked the artery automatically for subsequent removal. More advanced image processing recently involved uncertainty information in the form of confidence maps as introduced in [112]. As a tool, confidence maps did not only allow for automated maintenance of US quality [38] but also adjustment of the insonification angle for optimal acoustic coupling [37]. While these methods achieve overall higher repeatability in acquisition and guidance, they are also characterized by a lack of automation beyond a narrowly confined task.

3.2.4 Contributions

In this chapter, a comprehensive, autonomous robotic ultrasound framework based on visual servoing schemata is introduced. Combining a robotic manipulator equipped with internal torque sensors with a motorized 3D ultrasound probe, the employed hardware setup enabled the development of application-specific control laws, which were validated on both phantoms and real human anatomy.

Our previous, preliminary work [97] serves as starting point in terms of robot control and autonomous US image acquisition. In this work, a generic control architecture is built on this foundation, fusing potentially conflicting force and position control laws, both of which are necessary to acquire high quality US images. Several methodological improvements, including the performance of 3D ultrasound compounding and registration, enable real-time behavior of the system and greatly improve needle placement accuracy. In addition, a complete re-design of the transducer mount and needle holder now comply with clinical safety protocols.

Two clinical applications with a high potential benefit of robotic assistance with ultrasound were identified, one each in the scope of ARSS and ADAS. The main contribution for guidance of lumbar spine needle insertions is a novel visual servoing control law. A continuous re-registration of live 3D ultrasound images with an interventional plan, which is based on pre-interventional CT or MR images, allows for accurate guidance of manual needle insertions. Since the transformation obtained by the registration algorithm is directly fed to the robotic servo-controller, target anatomy movements are quickly compensated by re-adjusting the US transducer position. As a result, the presented system does not only relieve the physician from manually maneuvering the US probe while injecting the facet joint, but also provides continuous guidance regardless of target motion.

For this application, results of an extensive gel phantom analysis of the visual servoing system behavior are reported for target anatomy tracking under various conditions. By including smooth movements with different velocities and along different degrees of freedom, the limitations of the presented method are estimated. A set of human volunteer experiments demonstrates the capabilities of the proposed visual servoing methodology under realistic conditions. Finally, the proposed framework was validated within a realistic neurosurgical operating environment to demonstrate clinical feasibility in terms of needle placement accuracy. First results of manual needle insertions performed by an expert surgeon into an ultrasound-realistic lumbar spine gel phantom are reported and compared against the gold standard.

The other investigated clinical application concerns the fully autonomous acquisition of abdominal 3D US images to facilitate AAA screening in clinical routine. Designed to cope with a high anatomical variety in the general population, the proposed system aims to adapt the performed US trajectory to the individual patient. Similar to the approach in our earlier work [97], an RGB-D camera is used to register the patient with respect to the system setup. However, a deformable registration in combination with a generic patient atlas is employed to account for various body sizes and shapes. In this way, patients for whom no tomographic imaging data is available can undergo the proposed screening protocol. Similarly, due to the great variety of possible patient conditions, there is a need for an adaptive parametrization of the employed contact force of the US transducer onto the skin. The proposed system utilizes

confidence maps [112] to automatically determine the optimal force for the US scan. An additional innovative control law, again based on confidence maps, furthermore allows to initially estimate and continuously adjust the out-of-plane rotation of the US transducer during the acquisition. This enables the optimization of the image quality at specified penetration depths, for this application at the location of the aorta, compensating for shadowing artifacts due to the presence of bowel gas.

Presented results in this chapter demonstrate the effectiveness of the proposed algorithms to optimize for contact force and transducer alignment, improving the quality of the resulting US scans compared to naive acquisitions. In its current form, the presented system aims at providing US scans for manual diameter measurements. Therefore, the system is validated on healthy volunteers by comparing measurements from both automatic and manual US acquisitions, which demonstrated the feasibility of the presented system for AAA screening.

3.3 Methods

In this section, the modular components to perform autonomous robotic ultrasound acquisitions and to allow more sophisticated robotically assisted clinical applications are presented. It is organized as follows: First, the general control architecture is presented in Sec. 3.3.1, followed by a description of the force-based control law employed in both presented applications (Sec. 3.3.2). Both scenarios also share the techniques to acquire ultrasound images, which are described in Sec. 3.3.3. The calculation of confidence maps is detailed in Sec. 3.3.4. Thereafter, the application-specific parts are explained, focusing on the respective visual servoing control laws. For neurosurgical needle insertions, the implemented clinical workflow (Sec. 3.3.5) is introduced and the necessary registration-based visual servoing law presented (Sec. 3.3.6). Finally, the implemented workflow, the patient registration method and the necessary control laws for fully autonomous abdominal aneurysm screening are described in Sec. 3.3.7-9.

3.3.1 System Design and Generic Control Architecture

The presented robotic ultrasound system consists of a manipulator equipped with internal torque sensors at all joints, allowing to estimate the force at its end-effector. A convex, motorized US transducer, suited for abdominal and spine scans, is rigidly attached to the end-effector. An exemplary system setup is depicted in Fig. 3.4a. Please note that the presented methodology is independent from the used hardware, as long as a means to measure end-effector forces is available. The full specifications of the employed hardware is listed in Sec. 3.4.1.

Two distinct control laws are jointly responsible for robot movements, as illustrated in the control fusion architecture in Fig. 3.5. On the one hand, a compliant direct force controller maintains a constant contact force onto the patient, ensuring sufficient image quality yet allowing the patient or the operator to manually move the robot away, if it were needed. On the other hand, an application-specific visual servoing controller is responsible for modifying the robot trajectory to maintain sonographic visibility, imaging quality, or needle alignment.

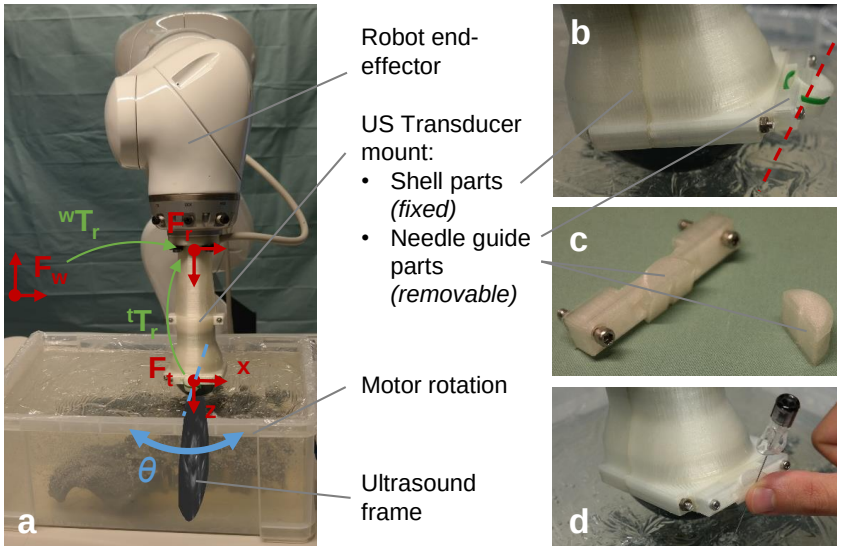

Fig. 3.4. **a)** System design including robot, mount, motorized US transducer (θ denotes the motor angle) and needle guide. The green transformations map between world (F_w, can coincide with robot base), robot end-effector (F_r) and US transducer (F_t) coordinate frames. **b-c)** The mount consists of four parts: Two fixed shell parts and two removable needle guide parts. **d)** When the outer needle guide shell is removed, accurate and safe guidance is possible using a finger. *Reprint from [249] with permission of Springer.*

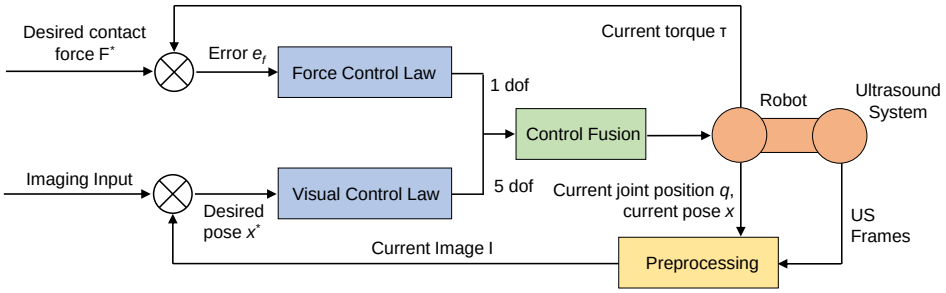

Fig. 3.5. The generic, dual control architecture combines force control with visual servoing in a parallel fashion. *Reprint from [249], modified, with permission of Springer.*

3.3.2 Force Control Scheme

For optimal ultrasonic acquisitions, a constant contact force onto the surface is desirable. As already demonstrated in prior art [37, 69], a compliance control scheme is adopted to regulate the robotic manipulator, which allows to relate the displacement of the end-effector to the forces acting on it. Considering the tool center point (TCP) frame of the end-effector, it is possible to control the force applied along a constrained task direction, while the position of the end-effector is controlled along the unconstrained task directions. In our setup, the constrained direction is the z-axis of frame F_t, which is the orthogonal vector to the contact surface. For a manipulator with compliant joints and rigid links, the generalized end-effector stiffness matrix $\mathrm{K}_{TCP} \in \mathbb{R}^{m \times m}$ can be expressed in task space as in [48, 210]:

$$F_{TCP} = \mathrm{K}_{TCP}\,(x - x_0) = \mathrm{K}_{TCP}\,\Delta x, \tag{3.9}$$

where $F_{TCP} \in \mathbb{R}^m$ is the external force at the robot end-effector, x_0 and $x \in \mathbb{R}^m$ are the initial and the current Cartesian position of the robot, respectively, and $\Delta x \in \mathbb{R}^m$ is the position displacement. For real-time behavior, torque sensors in all robotic joints are utilized to compute the external forces using both the Jacobian and the known inverse dynamics system of the manipulator. A constant force along the constrained direction is obtained for planar contact surfaces with the selection of the desired force F^* and a low stiffness K^* as values of the respective components along the constrained axis in F_{TCP} and K_{TCP}. A high stiffness along the remaining components allows for a classic position control scheme in the other directions. That is, while the given desired force F^* and stiffness K^* determine the end-effector position along its z-axis, the remaining degrees of freedom can be controlled by the proposed visual control scheme, as shown in Fig. 3.5. Due to the compliance of the system, the robot can be manually moved away from the contact surface at any time, either by the patient or the physician. During needle insertion, this is useful to make slight adjustments, if necessary.

3.3.3 2D and 3D Ultrasound Image Acquisition

Depending on the enabled visual servoing control law, which in turn depends on the application and the current clinical stage of the procedure, either 2D or 3D ultrasound images are used. Note that while in chapter 4, Doppler acquisitions are utilized to reconstruct 3D blood velocity fields, this chapter is limited to series of B-mode frames $I_i \in \mathbb{R}^2$.

For every single frame, the full chain of transformations into the world coordinate frame F_w is considered (see Fig. 3.4a):

$$ {}^{us}\mathrm{T}_w = ({}^{w}\mathrm{T}_r)^{-1} \cdot {}^{t}\mathrm{T}_r \cdot \mathrm{T}_\theta \cdot \mathrm{Z}, \tag{3.10} $$

where Z is the intrinsic calibration (see Sec. 1.3.4) determined using the known pixel spacing and the pixel coordinates of the first central US ray sample as provided by the US system, and T_θ a rotation matrix determined by the current motor position of the US transducer. The extrinsic calibration ${}^{t}\mathrm{T}_r$ is initialized based on the CAD design of the mount and fine-tuned using orthogonal robotic sweeps similar to [237]. Finally, ${}^{w}\mathrm{T}_r$ is the current end-effector pose of the robot.

In the case that 3D volumes are required by the visual servoing controller, individual 2D B-mode frames are compounded using a GPU-based backward warping strategy as in [113]. As illustrated in Fig. 3.6, various scenarios including movements of both the robotic arm and the US transducer motor need to be considered. In previous works [250], we assumed the static case as in Fig. 3.6b and only used the transformations $\mathrm{T}_\theta \cdot \mathrm{K}_c$ for 3D compounding. 3D sweeps created in this way were thus originated in the end-effector frame F_t and were ultimately, as a whole, moved to the world frame F_w according to one single robot pose. To allow for undistorted 3D volumes even if the robot is moving during acquisition, the full transformation as in Eq. 3.10 is now considered for compounding.

Special care is required to avoid holes in the 3D volume when robot and motor are moving in the same direction (see Fig. 3.6c), effectively limiting the robot's velocity v_{max} for a desired 3D resolution. If the sweeping motion is directed against the manipulator's trajectory, as sche-

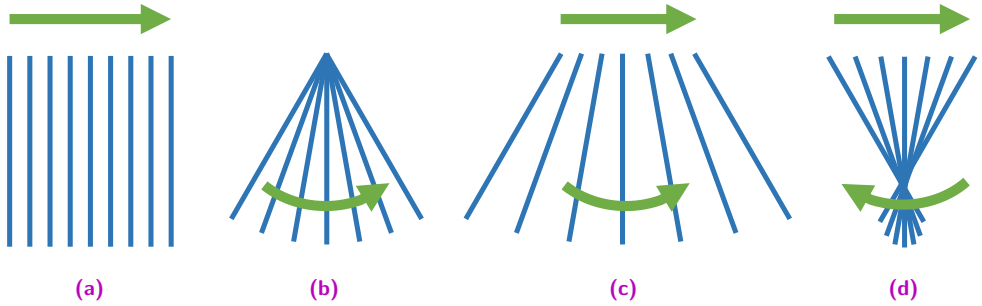

Fig. 3.6. Geometry of US frames due to fused motion of motorized US transducer and robotic arm. **a)** Transducer motor off, only linear motion by robotic arm. **b)** Robotic arm fixed, sweep using the transducer motor ("wobbeling"). **c,d)** Challenging geometries for 3D compounding, if both transducer motor and robot are moving simultaneously.

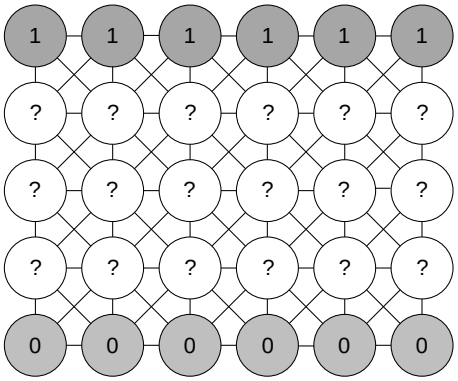

Fig. 3.7. Schematic illustration of the graph used for confidence map computation. Top and bottom rows represent source (1) and sink nodes (0), respectively. Nodes in between are connected to their 8-neighborhood, and the edge weights are defined as a function of direction and image intensity gradients.

matically depicted in Fig. 3.6d, imaging artifacts in the overlapping parts of several US frames are to be expected because ultrasound signals are dependent on the direction of insonification, causing B-mode intensity variation of the same physical location between frames. This effect can be mitigated by selecting the most appropriate frames during compounding in the case of overlap, solving the compounding in a energy-optimal [99] or tensor-based way [96], or by assuming that the predominant (orthogonal) frames outweigh the ones acquired from less steep angles and applying a median filter [113].

For practical reasons, no correction of distorted US frames due to motor and robot movement is applied in this work, i.e. US frames are assumed to be of planar geometry, even if there was out-of-plane motion between the acquisition of the first and the last US ray.

3.3.4 Confidence Map Computation

For an estimation of the ultrasound quality, we employ confidence maps as introduced by Karamalis et al. [112]. While this technique only models ultrasound reflection physics in biological tissue to a very limited extend, per-pixel attenuation information has successfully been used to approximate uncertainty in B-mode images [191].

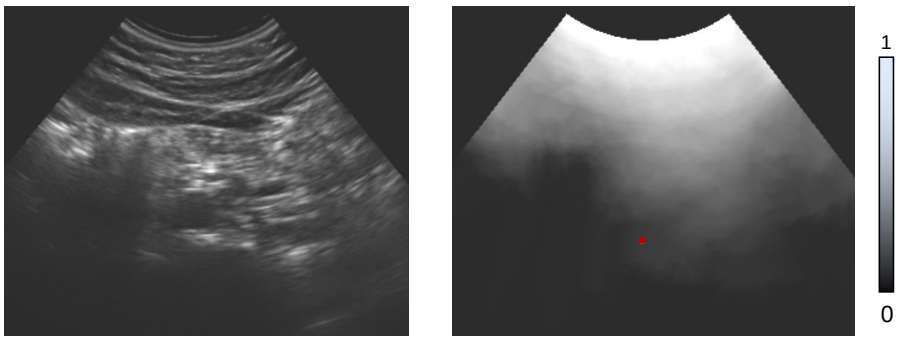

Fig. 3.8. **a)** Exemplary B-mode frame of the aorta, and **b)** corresponding confidence map. The expected center of the aorta is shown as *red* dot. *Reprint from [230], modified, ©2016 IEEE.*

As illustrated in Fig. 3.7, an 8-connected graph is constructed between the pixels of the B-mode image, with source (confidence: 1) and sink (confidence: 0) nodes at the transducer elements and the bottom of the image, respectively. The edge weights are directly assigned based on their directions and the US intensity gradients between pixels. The core idea of confidence maps is now to compute the probability of a random walker starting from a particular pixel to rather first reach the (virtual) transducer elements at the top than the sink vertices at the bottom. In essence, the problem can be therefore formulated as an equilibrium diffusion solution using US-specific constraints. Because the edges of the graph are of undirected nature, it is possible to express the problem effectively in the form of a linear system

$$\mathrm{L}\mathbf{y} = \mathbf{b}, \tag{3.11}$$

where L is a symmetric, sparse graph Laplacian matrix encoding the weights between graph vertices, for which a confidence value should be computed. $\mathbf{b}$ contains the Dirichlet boundary conditions, including source and sink vertices. By solving the system, $\mathbf{y}$ yields the confidence values in vectorized form. After rearrangement and mapping to the dimensions of the input B-mode image I_i, a confidence map $C_i \in \mathbb{R}^2 \rightarrow [0, 1]$ is obtained.

For any given US frame, we then denote with the feature

$$\zeta_R = \frac{1}{|R|} \sum_{(x,y)\in R} C(x, y), \tag{3.12}$$

the average confidence for some rectangular region R in the corresponding B-mode frame I. The averaging hereby copes with the inherent level of noise in confidence maps [99]. In Fig. 3.8, exemplary corresponding B-mode and confidence map images are shown.

3.3.5 Needle Guidance Workflow

After presenting the generic components of the robotic ultrasound system in the previous sections, the application-specific parts of this chapter are described. In Fig. 3.9, an exemplary clinical workflow for lumbar facet joint needle insertions is depicted.

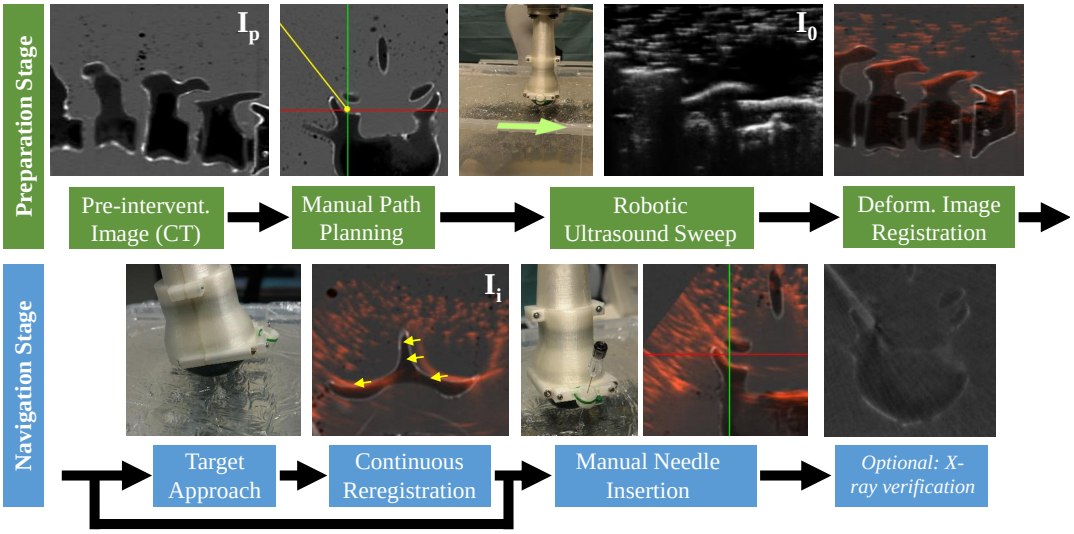

Fig. 3.9. Workflow for facet joint needle insertion using the proposed ultrasound-based visual servoing guidance framework. In the **preparation stage** (*green*), a multi-modal set of images is acquired and registered to establish an interventional plan in the current patient coordinate system. Thereafter, the automatic navigation system takes over during the **navigation stage** (*blue*), maintaining registration over time and enabling precise manual needle insertion. *Reprint from [249] with permission of Springer.*

For needle insertions using the proposed system, adequate imaging of the patient is required for manual path planning (I_p). Ideally, pre-interventional CT (as in the figure) or MR images are available, which is commonly (but not necessarily) the case in today's practice to confirm the treatment indication. Yet, 3D ultrasound volumes as acquired during the robotic sweep would be sufficient as long as the joints can be clearly visualized. Along with the needle path, additional annotations such as vertebrae labels might be helpful for subsequent orientation.

Preparation Stage

Once the patient is positioned in the intervention room and sufficient ultrasound gel is applied, the robot can be manually steered towards the patient's sacrum in gravity compensation mode. It will then automatically approach the skin until a predefined scanning force is reached, and automatically perform an initial US sweep I_0 in caudocranial direction. During the acquisition of this sweep, the motor of the transducer is turned off ($\theta = 0$) so that parallel frames are acquired. The sweep is compounded to a 3D image I_0 and registered to the pre-interventional image I_p by estimating a non-linear transformation $\tilde{T}_p$:

$$\tilde{T}_p = \arg\max_{T} S_{\mathrm{LC}^2}(\mathrm{I}_0, T(\mathrm{I}_p)), \tag{3.13}$$

where S_{LC^2} is the multi-modal LC^2 similarity metric [70], efficiently evaluated on the GPU (see Sec. 1.4.3 for further details). Initialization of T, which is modeled using free-form deformations [141], is facilitated by the assumption that the robotic scan starts at the sacrum and continues in cranial direction. This allows reliable mapping of vertebrae in I_0 and I_p and helps to avoid level misalignments. All manual annotations can be embedded in the world coordinate frame F_w using the final transformation $\tilde{T}_p$, including the planned needle path. The optimization of the parameters is performed using BOBYQA [169].

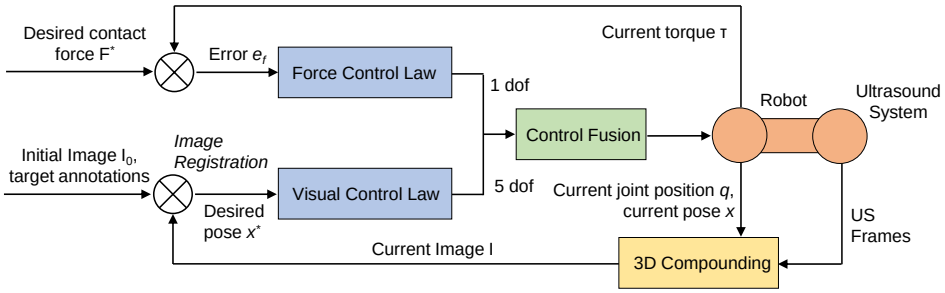

Fig. 3.10. The dual control architecture combines force control with registration-based visual servoing in a parallel fashion. Compared to the generic architecture in Fig.3.5, preprocessing includes a 3D compounding operation, and the visual error computation is performed using 3D intensity-based image registration. *Reprint from [249], modified, with permission of Springer.*

Navigation Stage

At this point, the system is prepared to start the navigation stage and approach the target position by aligning the needle guide attached to the US transducer (see Fig. 3.4b) with the planned needle path. The updated control architecture is shown in Fig. 3.10. The objective of the visual control law is now to update the desired pose of the US transducer by means of 3D image registration. This allows to follow and compensate for target anatomy movements, and to keep the needle guide aligned with a pre-interventional plan.

For continuous re-registration, sweeps to acquire images I_i are obtained using the transducer motor. As described below, the proposed visual servoing scheme performs a continuous re-registration of the target region in 3D and compensates for patient motion (note that the patient is awake throughout the entire intervention). Eventually, the surgeon manually inserts the needle under ultrasonic guidance. In this final stage, the motor is turned off again to enable high frame rate visual feedback to the surgeon. The lack of 3D imaging information then also prohibits further automatic robot pose updates, avoiding harm to the patient once the needle has been inserted. If desired, a multi-planar reconstruction of the registered pre-interventional image can be visualized as well for better contrast. Ultimately, since the needle guide can be opened and the robot manually removed at any time, traditional X-ray verification of the final needle position is easily possible before drug injection. Note that this optional step is intended for clinical verification studies but not the final, radiation-free interventional procedure.

3.3.6 Registration-based Visual Control Schemes

Motion Compensation Constraints

Instead of defining and optimizing a visual error functional as outlined in Sec. 3.2.2, a transformation $\hat{T}_i$ to align two 3D ultrasound images is estimated using registration and then used to update the desired robot pose x_i^* via standard position control:

$$\hat{\mathrm{T}}_i = \begin{pmatrix} \mathrm{R}(\alpha,\beta,\gamma) & (t_x,t_y,t_z)^\intercal \\ \mathbf{0} & 1 \end{pmatrix} = \arg\max_T S_{\mathrm{NCC}}(\mathrm{I}_i, T(\mathrm{I}_0)). \tag{3.14}$$

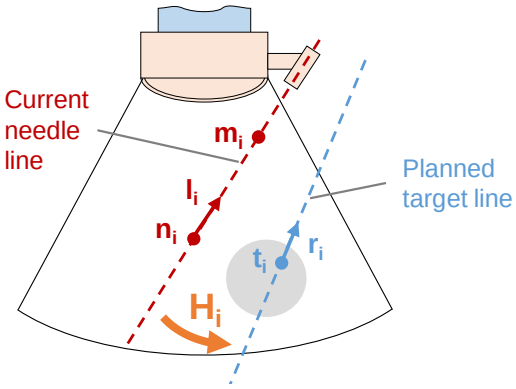

Fig. 3.11. To allow for robotic needle guidance, the visual control law needs to bring the current needle guide line (*red*, $\mathbf{l_i}$) and the planned insertion path (*blue*, $\mathbf{r_i}$), two skew lines, into alignment. *Reprint from [250], ©2016 IEEE.*

Note that throughout the remainder of this chapter, linear transformations and vectors are expressed in computer vision notation, i.e. using 4×4 homogeneous matrices and 4×1 vectors. In this work, we assume locally rigid movements [236], constituting six degrees of freedom. R denotes a 3×3 rotation matrix parametrized with Euler angles α, β, γ. The fixed image is now the most recent 3D compounding I_i, and the moving one the initial US sweep I_0 with all annotations. Normalized cross correlation (see Sec. 1.4.3) is employed for the similarity S_{NCC} as both volumes are derived from the same modality [236]. The transformation is initialized with the previous registration result $\hat{\mathrm{T}}_{i-1}$ for fast convergence. In addition, the degrees of freedom of the optimization procedure are restrained, since pose updates in the direction of the z-axis of frame F_t will be ignored anyway by the force controller. The exclusion of further degrees of freedom is also possible, e.g. if only translations are expected.

If there has been no change of the target anatomy, both images will be almost identical, and the similarity function S_{NCC} at a local maximum with all transformation parameters set to 0. Small movements, however, can be recovered efficiently within a few iterations of the optimizer, directly allowing for their compensation with the obtained transformation $\hat{\mathrm{T}}_i$.

In scenarios without a needle guide, where the system is purely used to track a moving target, the visual control law determining a new, desired Cartesian pose x_i^* of the US transducer is formulated as follows:

$$\mathrm{x}_i^* = \begin{bmatrix} \mathrm{R}(\alpha^*, \beta^*, \gamma^*) & \mathbf{t}^* \\ \mathbf{0} & 1 \end{bmatrix} = ({}^t\mathrm{T}_r)^{-1} \cdot (\hat{\mathrm{T}}_i)^{-1} \cdot \mathrm{P}_0, \tag{3.15}$$

where P_0 is a manually annotated US transducer target pose defined either directly in the initial sweep I_0 or transformed from the pre-interventional image: $\mathrm{P}_0 = \tilde{T}_p(\mathrm{P}_p)$.

Needle Alignment Constraints

If a needle guide needs to be aligned with a predefined insertion path, additional constraints based on two further manual annotations have to be considered. Let $\mathbf{t_0}$ be a needle target point and $\mathbf{r_0}$ a suitable insertion path, both defined in I_0 similarly to P_0. Furthermore, consider

a set of two points in the transducer frame F_t, $\mathbf{m_0}$ and $\mathbf{n_0}$, to indicate two points the needle would hit if inserted into the guide. The two lines defined in this way are schematically illustrated in Fig 3.11. Obtained once by calibration in water using manual annotation of two arbitrary needle points, they can be expressed in the world frame using the known chain of transformations. All are carried along in each iteration as the target might have moved:

$$\begin{bmatrix} \mathbf{t_i} & \mathbf{r_i} \end{bmatrix} = (\hat{\mathrm{T}}_i)^{-1} \cdot \begin{bmatrix} \mathbf{t_0} & \mathbf{r_0} \end{bmatrix}, \tag{3.16}$$

$$\begin{bmatrix} \mathbf{m_i} & \mathbf{n_i} \end{bmatrix} = (\hat{\mathrm{T}}_i)^{-1} \cdot ({}^{w}\mathrm{T}_r)^{-1} \cdot {}^{t}\mathrm{T}_r \cdot \begin{bmatrix} \mathbf{m_0} & \mathbf{n_0} \end{bmatrix}. \tag{3.17}$$

The normalized needle guide direction can then be defined as $\mathbf{l_i} = (\mathbf{m_i} - \mathbf{n_i}) / \|\mathbf{m_i} - \mathbf{n_i}\|$ [250]. The required correction to orient skew $\mathbf{r_i}$ and $\mathbf{l_i}$ in a parallel fashion is a rotation

$$\mathrm{C}_i = \mathrm{R}(\mathbf{r_i} \times \mathbf{l_i}, \mathrm{acos}\langle \mathbf{r_i}, \mathbf{l_i} \rangle), \tag{3.18}$$

where R converts the axis-angle representation to a 3×3 matrix. An additional offset is finally necessary to let both lines coincide:

$$\mathbf{u_i} = (\mathbf{m_i} - \mathbf{t_i}) - \mathbf{r_i} \cdot \langle \mathbf{r_i}, \mathbf{m_i} - \mathbf{t_i} \rangle \tag{3.19}$$

Together, these two adjustments are embedded into the matrix

$$\mathrm{H}_i = \begin{pmatrix} \mathrm{C}_i & \mathbf{u_i} \\ 0 & 1 \end{pmatrix}, \tag{3.20}$$

which forms the basis for the modified visual control law

$$\mathrm{x}_i^* = ({}^{t}\mathrm{T}_r)^{-1} \cdot \mathrm{H}_i \cdot (\hat{\mathrm{T}}_i)^{-1} \cdot \mathrm{P}_0. \tag{3.21}$$

The behavior induced by this combined control law serves two purposes: (1) The target anatomy (and thus also the defined injection target), are tracked and will be followed as they move. (2) The calibrated needle guide will be aligned with the defined needle insertion path.

3.3.7 Aneurysm Screening Workflow

In the remaining methodological sections of this chapter, the application-specific parts for abdominal aortic aneurysm screening are presented. For this application, the goal is to automatically acquire an abdominal US scan of the aorta with optimal quality, so that clinically meaningful measurements of aortic diameters can be obtained afterwards. To achieve full automation, an RGB-D camera is used to automatically detect the patient position, and to register a suitable, application-specific trajectory onto the patient, which can then be executed by the robot.

By means of the intended system setup as shown in Fig. 3.12, the proposed workflow can be outlined: First, a point cloud of the patient is acquired using an RGB-D camera and registered

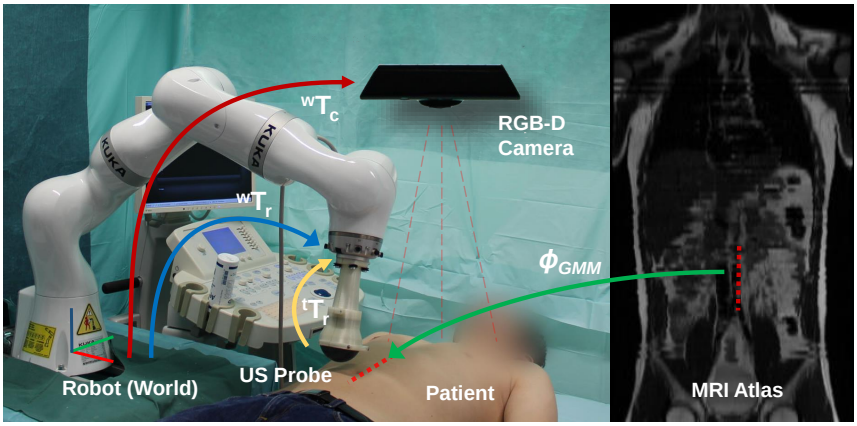

Fig. 3.12. System setup showing the robot, the US transducer mounted to its end-effector, the RGB-D camera on the ceiling, the patient, a coronal slice of the MRI atlas, and all required transformations. *Reprint from [230], ©2016 IEEE.*

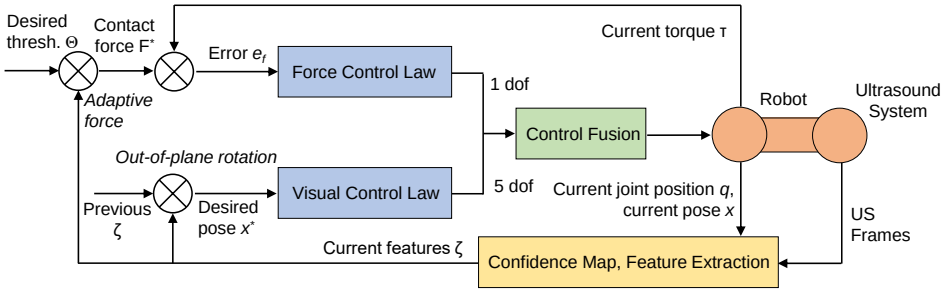

Fig. 3.13. The dual control architecture combines force control with confidence-based visual servoing in a parallel fashion. Compared to the generic architecture in Fig.3.5, preprocessing includes a 2D confidence map computation and feature (ζ) extraction for every B-mode frame. These features do not only influence the visual servoing component by triggering adjustment of the out-of-plane rotation but also the desired contact force F_d. *Reprint from [249], modified, with permission of Springer.*

to a suitable atlas. This allows for automatic planning of a robotic acquisition trajectory of the abdominal aortic region. The robot will then approach the patient's surface until the transducer touches the skin, and use an ultrasound-based adaptive force algorithm to estimate an appropriate contact force used throughout the acquisition. An automatic, rotatory sweep completes the necessary preparations before the actual acquisition is performed, identifying the optimal angle of the transducer. Eventually, a visual servoing technique is employed to continually optimize the out-of-plane rotation online during the acquisition.

The updated control architecture is shown in Fig. 3.13. The two newly introduced control schemes for adaptive force estimation and the optimization of the out-of-plane rotation are explained in Sec. 3.3.9.

3.3.8 Patient Registration and Trajectory Transfer

The aim of a global patient-to-world registration is to gather knowledge about the current patient position with respect to the system setup, so that accurate and safe motions of the robotic arm can be achieved. We exploit the sensing information of a RBG-D camera to

transfer a generic trajectory aimed at the aorta of an atlas onto the current patient. For camera-to-robot calibration, we employ a technique based on augmented reality markers similar to the approach in [97] to determine the transformation ${}^{W}\mathrm{T}_{C}$ (cf. Fig. 3.12).

Given the wide target population of a screening program, it is not feasible to assume the availability of individual tomographic images for each patient. Hence, we propose to leverage a statistical MRI atlas based on physical and anatomical characteristics, such that the anatomy of each patient can be taken into account for the trajectory planning [177]. A surface point cloud P_{MRI} extracted from a selected MRI atlas image can be elastically registered to the live point cloud P_{RGBD} obtained from the RGB-D camera. We make use of an implementation of the Coherent Point Drift (CPD) algorithm [148], a probabilistic non-rigid registration method that fits a Gaussian Mixture Model (GMM) to the moving point set. The algorithm is also used in Sec. 2.3.3 to register the surfaces of MRI and TRUS prostate segmentations. The reader is referred there for a detailed mathematical description. In short, the GMM is initialized using the target points P_{RGBD} and a coherent velocity is enforced to its motion so that a smooth non-linear transformation $\Phi_{GMM} : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ can be computed using spline interpolation. Both point sets are subsampled by a factor f_s for this process, allowing for an acceptable trade-off between the fitting accuracy and the computational performance.

By a projection of the start- and endpoint ($\mathbf{p}_s$,$\mathbf{p}_e \in \mathbb{R}^4$) of the aortic region of interest from the atlas to its surface, the robotic trajectory on the patient surface ($\mathbf{p}'_s$,$\mathbf{p}'_e$) is obtained by transfering these points to the world coordinate system

$$\mathbf{p}'_s = ({}^{W}\mathrm{T}_{C})^{-1} \cdot \Phi_{GMM}\left[f_{NN}(\mathbf{p}_s, \mathbf{e}_z)\right], \tag{3.22}$$

where $\mathbf{e}_z = (0, 0, 1, 0)$ is the vertical unit vector, and $f_{NN}(\mathbf{p}, \mathbf{n})$ computes the nearest element of the point cloud P_{MRI} to the ray $\mathbf{p} + \lambda\mathbf{n}, \lambda \in \mathbb{R}_0$ (analogous for $\mathbf{p}'_e$).

3.3.9 Visual Control Schemes for Aneurysm Screening

Adaptive Force Estimation

Although the area of general robotic force control has been extensively discussed throughout the last decades, the choice of an appropriate force F_d for a particular medical scenario depended on a manual parametrization up to now. While too little pressure will compromise good acoustic coupling and sufficient image quality during US acquisitions, excessive force might overly deform the anatomy or even harm the patient. In the view of a fully autonomous robotic system for US screening, an optimal force value cannot be known a priori but has to be estimated online to cope with a variety of patients constitution and tissue density. Therefore, we propose an online adaptive force estimation based on confidence values presented in Sec. 3.3.4. During initialization, we vertically approach the start pose $\mathbf{p}'_s$ until skin contact ($F^*_0 = 0$). Next, the desired force exerted onto the tissue is increased iteratively by F_{step} until a mean confidence threshold Θ is reached (H is the Heaviside step function):

$$F^*_{i+1} = F^*_i + F_{step} \cdot H\left(\Theta - \zeta_i\right). \tag{3.23}$$

Optimization of Out-of-Plane Rotation

The Euler angles $(\alpha^*, \beta^*, \gamma^*)$ of the pose x^* can be interpreted as out-of-plane rotation, in-plane rotation, and rotation around the transducer axis, respectively. Contrary to the target anatomies investigated in [37] and lumbar spine regions, axial aortic scans benefit only marginally from in-plane or transducer axis rotation: On the one hand, turning the US probe around its axis does not avoid acoustic obstacles between the transducer and the aorta. On the other hand, an in-plane rotation during axial scans constitutes lateral tilting, quickly translates the aorta away from the image center, and is not considered helpful in clinical routine [165]. Thus, we define $\beta^* = \gamma^* = 0$ for all experiments and concentrate on the more challenging out-of-plane rotation α for image quality optimization.

Initially, a sweep is acquired with angles between $[\alpha_{min}; \alpha_{max}]$, where the confidence feature ζ_α is recorded for each rotatory pose. The optimal out-of-plane rotation α_0 for the start point of the sweep is then defined as the angle that maximizes confidence at the aorta: $\alpha_0^* = \arg\max_\alpha \ \zeta_\alpha$.

Throughout the acquisition, we aim to maintain an optimal echoing pose. As the direction to tilt the probe out of its plane cannot be directly inferred from 2D frames, we propose to compute the following parameters for each frame. First, the binary parameter κ indicates a drop in confidence below the average of the preceding M frames

$$\kappa = H\left(\zeta_i - \frac{1}{M}\sum_{j=i-M-1}^{i-1} \zeta_j\right). \tag{3.24}$$

It is used to determine whether the current probe orientation provides sufficient image quality. Second, the parameter $\sigma \in \{-1, 1\}$, $\sigma_0 = 1$, states the direction the probe should tilt to:

$$\sigma_i = \begin{cases} -\sigma_{i-1} & \text{if } \zeta_i < \zeta_{i-M} \ \wedge \left|\sum_{j=i-M-1}^{i-1} \sigma_j\right| = M \\ \sigma_{i-1} & \text{else .} \end{cases} \tag{3.25}$$

As a result, σ will change sign only if there has not been a change in the previous M iterations, and the current confidence dropped below the one M iterations ago. Altogether, the desired out-of-plane rotation can be computed by combining these factors

$$\alpha_i^* = \alpha_{i-1}^* + \kappa \cdot \sigma \cdot \alpha_{step}. \tag{3.26}$$

Updates of α are thus not continuous but are handled by the position controller in a smooth fashion as in [97].

3.4 Materials and Experiments

3.4.1 System Setup

As shown in Fig. 3.4a and Fig. 3.12, a KUKA® LBR iiwa 7 R800 robot (KUKA Roboter GmbH, Augsburg, Germany) was employed for all experiments. The manipulator is equipped with torque sensors in all joints, allowing for impedance control applications. A KUKA Sunrise Connectivity SmartServo application[1] was utilized to connect to ROS (Robot Operating System[2]) on a client workstation for reporting the current robot status with 400 Hz, including the current end-effector pose ${}^{w}\mathrm{T}_{r}$, and executing incoming commands [230]. On the hand-held control pad, custom functionality was added to allow quick changes to gravity compensation mode for manual manipulator guidance. The robot is CE-certified and designed for human interaction for safe usage in collaborative scenarios compliant with ISO 10218-1. Desired force and stiffness in tool direction were set to $F^{*} = 5$ N and $K^{*} = 0$ N/m, respectively. The stiffness in all other directions was fixed to $1,000$ N/m.

An Ultrasonix® Sonix RP system with a motorized curvilinear transducer (model 4DC7-3/40 Convex 4D) was used for B-mode acquisition (BK Ultrasound, Analogic Corp., Peabody, MA, USA). The 2D acquisition rate of the curvilinear array was set to 30 Hz. The motor speed was defined so that sweeps of $\pm 15^{\circ}$ are covered at 4 Hz. The frequency, depth and gain were set to 3.3 MHz, 70 mm, and 50%, respectively. A modified component in the publicly available PLUS library 2.2.0 [124], based on the Ultrasonix Porta SDK 5.75, maintained a bi-directional OpenIGTLink [220] Ethernet connection to the client workstation for a) streaming incoming B-mode frames and b) receiving user commands to start and stop the motor as desired.

A custom 3D-printed mount consisting of four parts was used to attach the US probe with the manipulator and allowed for needle guidance, as shown in Fig. 3.4b. It was designed to meet the safety and hygienic requirements of clinical environments. Two symmetric shell parts surround the transducer and enable the fixation to the end-effector using four screws, forming a permanent fixation of the two devices to avoid frequent calibration procedures. In contrast, a semicircular needle guide with desired inclination could be individually attached for an intervention (see Fig. 3.4c-d). Its open design allows an operator to guide the needle by hand but remove the compliant robot at any time, even after the needle is already partially inserted. Finally, an optional second semicircular shell makes it possible to completely surround the needle for highest precision guidance. All parts in contact with the needle are quickly removable and can be manufactured either for single-use or for re-use after sterilization.

The essential part of the proposed visual servoing control schemes is the image processing component, for which a set of custom plugins for ImFusion Suite 1.2.43 (ImFusion GmbH, Munich, Germany), an extensible GPU-based framework for medical images, was used. Running on the mentioned client workstation (Intel® Core i7-4770K processor at 3.5 GHz, 32 GB RAM, NVIDIA® GeForce GTX 970 graphics card), they incorporated OpenIGTLink and ROS interfaces and implemented the visual control laws described in the methods section.

[1] https://github.com/SalvoVirga/iiwa_stack
[2] http://www.ros.org/

For both applications, US frames were compounded on the GPU to 3D volumes with an isometric resolution of 0.3 mm. In the case that 3D US sweeps were acquired by means of the transducer motor, e.g. for needle guidance applications, compounding added a latency of around 30 ms, before subsequent processing could start. However, compounding tasks ran in parallel to any continuous registration tasks, which also utilized the graphics processor for the evaluation of the NCC similarity metric. For continuous re-registration, the registration task was alloted a fixed time budget of 230 ms, and the result was discarded if the optimizer had not converged until then. This helped to avoid a pile-up of 3D volumes (acquired at 4 Hz) to be processed.

Finally, for the autonomous screening application, spatial and depth information were acquired using a Kinect camera (Microsoft Corporation, Redmond, WA, USA) placed above the patient.

3.4.2 Experimental Phantoms and Human Volunteers

Phantoms were specifically created for all conducted experiments, offering realistic contrast in both ultrasound and CT. As tissue mimicking material, Ceraflex N530 transparent gel (Th. C. Tromm GmbH, Cologne, Germany) was used. Lab experiments to validate the visual servoing control schemes were performed on a geometric phantom consisting of 4 rubber spheres (diameters $2\times$ 27 mm, $2\times$ 35 mm) embedded in a $180\times160\times90$ mm box filled with gel (see Fig. 3.14).

For needle insertions experiments, a realistic and radiopaque spine phantom was made. A lumbar spine model (vertebra L1 to sacrum; Sawbones, Pacific Research Laboratories Inc., Vashon Island, WA, USA) was embedded in a gelatinous box, approximately the size of an adult human abdomen ($380\times240\times150$ mm). To avoid an unrealistically homogeneous gel filling, the gel was poured quickly and cooled to 5 °C immediately after. This allowed for the retention of air bubbles, generating challenging speckle noise in the US images.

For both applications, experiments on human volunteers were conducted. In the case of motion compensation for neurosurgery usage, one healthy volunteer (male, age 28) could be registered. Robotic US scans of his lumbar spine (L2-sacrum) were used for the experiments. The presented aneurysm screening method was evaluated on five different healthy volunteers (2 female, 3 male, age 24–27), scanning the aorta from slightly inferior of the rib cage in downward direction roughly until the navel (scanning time approx. 2 minutes per subject). The atlas for this application consisted of a T2-weighted MR image (resolution $1.2\times1.2\times6$ mm) of one healthy individual (age 26, male), which was deemed sufficient for this volunteer study due to the similar anatomical condition.

3.4.3 Motion Compensation Performance

In order to evaluate the tracking and motion compensation capabilities of our system, the geometric phantom was attached to a second robot, KUKA LWR 4+ (KUKA Roboter GmbH, Augsburg, Germany), as shown in Fig. 3.14a. After manual positioning of the ultrasound

Fig. 3.14. **a)** Experimental setup for servoing validation experiments. A second robot held a geometrical gel phantom and moved horizontally. **b)** Center US frame and **c)** coronal slice of the compounded 3D US volume. In the latter, the *blue* volume is the current volume, and the *red* one the initial sweep the algorithm continuously registers against. *Reprint from [249] with permission of Springer.*

equipped robot such that the central US frame roughly contained both smaller spheres' centers, a horizontal sweep of ± 60 mm was acquired for subsequent registration (Fig. 3.14c, *red* volume).

Three sets of experiments were conducted to measure the influence of various parameters on the tracking performance. At each point in time t, the spatial lag

$$\Delta s(t) = \|\Delta p_{iiwa} - \Delta p_{4+}\| \tag{3.27}$$

is measured as the difference between the relative movements of both robots. Depending on the experiment, $\Delta p = p(t) - p(t_0)$ either extracts the translation on the horizontal plane ($p_t = [t_x \; t_y]^\intercal$) or the rotation around the z-axis ($p_\gamma = [\gamma]$) of the end-effector poses, which were recorded at 50 Hz. In addition to the initial temporal lag Δt_0, i.e. the time between the onset of both robots' movements, we defined the time until convergence Δt_c as the time from the end of the 4+ movement until $t_e : \Delta s(t) < 0.05$ mm $\forall t > t_e$. The final positioning error after convergence was then defined as $e_p = \Delta s(t_e)$. Because only such relative movements were considered and the robots were aligned by design, we abstained from performing an additional robot-to-robot calibration.

In the first set of experiments (*A*), the lower robot was smoothly moved orthogonally to the central US image plane, i.e. along the x-axis of F_t, with different velocities v between 12 and 32 mm/s by 280 mm. The degrees of freedom for the registration were restricted to t_x and t_y to avoid that the registration exceeded its time budget. In the second set of experiments (*B*), the phantom was rotated by $45°$ with different angular velocities ω between 3.6 and $14.3°/s$. In a last set of experiments (*C*), the phantom was translated with constant, intermediate velocity of $v = 15$ mm/s in different directions ϕ by 150 mm to estimate whether motions orthogonal to the central US frame yield higher errors than the ones parallel to it. Finally, we performed a drift experiment over a time span of 4 minutes with several translations within a 200×100 mm box with $v = 12.1$ mm/s.

In a similar manner, the tracking performance was validated on the real anatomy of the mentioned human volunteer. As shown in Fig. 3.18a, the subject was positioned in prone

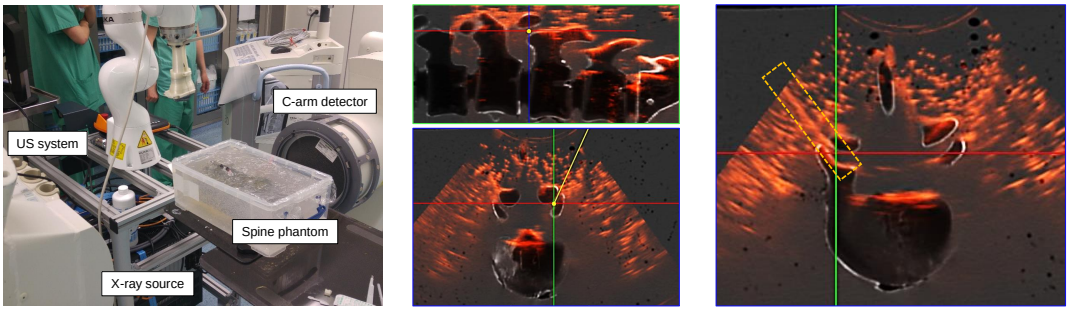

Fig. 3.15. **a)** Experimental setup in neurosurgery OR. **b)** Planned needle path shown in sagittal (*top*) and axial (*bottom*) slices of pre-interventional CT (*grayscale*) registered to initial US sweep (*red*). **c)** The needle can be clearly visualized in US during and after insertion (*yellow box*). *Reprint from [249] with permission of Springer.*

position on a wheeled bed, steerable in all directions, and instructed to perform regular, shallow breathing. Instead of a second robot performing translations, the bed was manually moved with 5-10 mm/s perpendicular and orthogonal to the central US frame. An NDI® Polaris Vicra (Northern Digital Inc., Waterloo, ON, Canada) was used to measure ground truth positions. The tracking target was fixed to the subject, as close as possible to the lumbar spine. A hand-eye calibration as in [223] was performed to facilitate reporting of positioning errors e_p on the horizontal plane as defined above.

3.4.4 Needle Insertion Accuracy

To validate the navigation capabilities of the proposed framework, needle insertion experiments were conducted in a realistic neurosurgical operating environment by two expert spine surgeons with standard 22G needles of 88 mm length. The phantom was positioned on an operating table in prone position as shown in Fig. 3.15a. The workflow as described in the methods section was fully carried out four times, each starting with a manual path annotation by one surgeon in a preoperatively acquired navigation CT image, as illustrated in Fig. 3.15b (*bottom*). After robotic target approach and registration-based refinement, four needles were inserted under live US guidance by the same surgeon who annotated the plan. For the first two, the closed needle guide (Fig. 3.4b) was used, for the last two, the open one in combination with a finger as depicted in Fig. 3.4d.

To compare with the gold standard, four additional needles were inserted by a surgeon without robotic assistance but under X-ray fluoroscopy guidance, for which a Siemens Arcadis® Orbic 3D iso-C-Arm system (Siemens Healthineers GmbH, Erlangen, Germany) was used. The C-arm was also used to acquire 3D fluoroscopy scans (CBCT), as shown in Fig. 3.19, to allow quantification of the achieved placement accuracies in both cases.

Two error measures were used. First, the total Euclidean error $e_d = \|t_0 - t_v\|$ reports the distance between the planned needle target t_0 and the manually annotated tip of the needle in the CBCT image, denoted t_v and reported in frame I_0 after registration to the already deformed ($\tilde{T}_p$) pre-interventional image. Second, since the presented navigation system has no influence

on the manual insertion depth, the orthogonal error $e_\perp = \|t_0 - t_v - \langle \mathbf{s_0}, t_0 - t_v \rangle\|$ reports the perpendicular component of e_d, i.e. the distance of the needle tip to the planned insertion line.

3.4.5 Aortic Aneurysm Screening Performance

For all five subjects, patient registration as described in Sec. 3.3.8 and several US acquisitions were performed. In a first sweep, a minimum force for US screening of 5 N was applied statically. To evaluate the adaptive force estimation, the algorithm outlined in Sec. 3.3.9 can then be executed for a second sweep with optimal contact force, but without optimization of the out-of-plane rotation, i.e. $\alpha^* = 0$. Finally, a third sweep allowed for validation of the proposed transducer angle control law. In a single sweep (only one volunteer), we tested the capabilities of the robot controller to compensate for motion in real-time to maintain a constant force. Therefore, we asked the volunteer to perform one deep chest inhalation and then breath abdominally throughout the acquisition.

The dependence between exerted force on the tissue and confidence was modeled using two regression models: A linear and a logarithmic model were examined to describe the relationship:

$$\zeta_{lin}(F) \approx a \cdot F + b \qquad \zeta_{log}(F) \approx a \cdot \log F + b \tag{3.28}$$

In all five volunteers, the aortic diameter was measured by a medical expert in the compounded US volumes (last sweep) in sagittal and axial slices according to the guidelines in [187]. For comparison, a medical expert blind to formerly mentioned measurements performed a standard US scan on each volunteer according to the clinical protocol as in [165]. In particular, inner diameters d_{in} (without walls), outer diameters d_{out} (with walls) and leading-edge diameters d_{le} (with closer wall only) were measured.

For all experiments, the following set of parameters was used: $f_s = 0.01$, $F_{step} = 2$ N, $F_{max} = 25$ N, $\alpha_{step} = 2°$, $M = 4$ and the region R comprised an area of 10×10 px. The robot moved with $v = 5$ mm/s during sweeps, and Θ was empirically set to 0.2.

3.5 Results

3.5.1 Motion Compensation Performance

Table 3.1 lists mean and standard deviation for the mentioned metrics for all experiments. For easy comparison, positions of both robots during all experiments of set *A* are collectively visualized in Fig. 3.16a. Only for the fastest motion with $v = 32$ mm/s, tracking was lost. In all other cases, the continuous movements were reliably detected and compensated by the visual servoing controller. As shown in Fig. 3.16c (*red curve*), we found a strong linear correlation between v and the average lag $\Delta\bar{s}$ ($R^2 = 0.98$), which is shown in Fig. 3.16b. Below the limit of the US resolution in transducer sweep direction (ca. 1 mm at the depth of the spheres), the final error e_p increased approximately in a linear fashion with higher velocities (overall

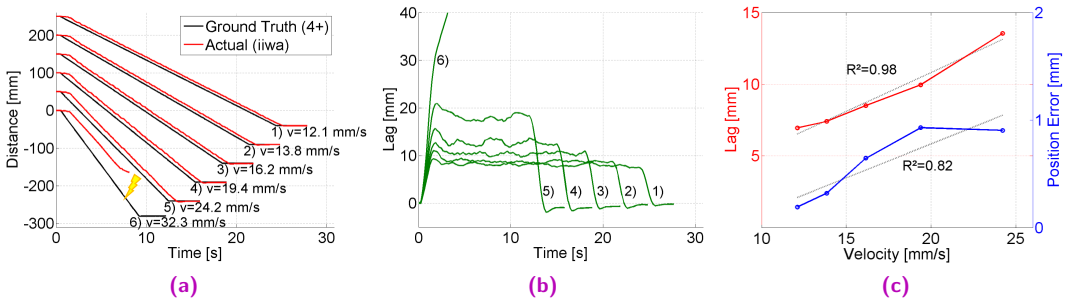

Fig. 3.16. Experiments *A*. **a-b)** Tracking was successful in all experiments except the one with fastest motion (32 mm/s), with spatial lags of less than 20 mm, almost no overshoot and fast convergence. **c)** Strong linear relation between the velocity and the spatial lag (*red, left ordinate*) as well as final position error (*blue, right ordinate*). *Reprint from [249] with permission of Springer.*

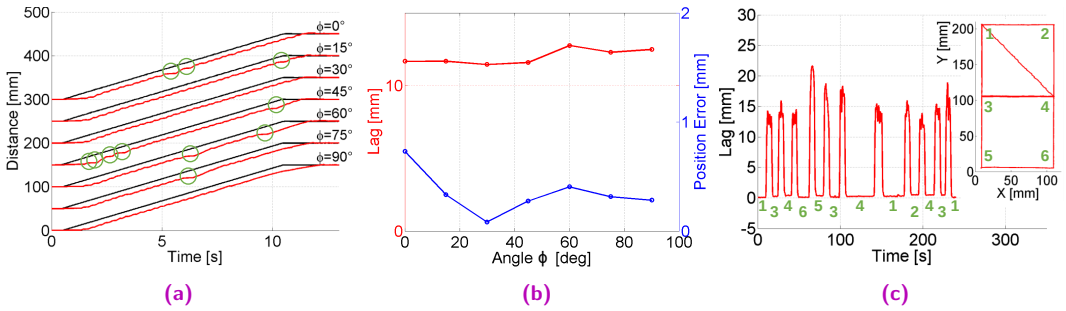

Fig. 3.17. **a)** Experiments *C*. Trajectories of both robots show successful tracking in all experiments even though the continuous registration occasionally exceeded its time budget (*circles*); **b)** both the spatial lag (*red, left ordinate*) and the final position error (*blue, right ordinate*) were independent of the translation angle ϕ. **c)** Spatial lag during several consecutive translations between corner points (labeled in *green*) of the box shown in the inlet suggests the absence of a drift over time. *Reprint from [249] with permission of Springer.*

$R^2 = 0.82$, see Fig. 3.16b, *blue curve*). Similar results were obtained for the experiment set *B*, with no clear relation between ω and error e_p. For experiments *C*, the positions of both robots during all experiments, projected into the respective vertical plane, are collectively visualized in Fig. 3.17a. The registration exceeded its time budget 11 times (≈3.3%), not updating the desired robot pose, (*see green circles*) but recovered quickly in all cases. No clear relation between ϕ and $\Delta\bar{s}$ (see Fig. 3.17b, *red curve*) was found. The positioning error e_p was also independent of ϕ (see Fig. 3.17b, *blue curve*), as long as the movement was not exclusively orthogonal to the central US plane (highest error $e_p = 0.73$ mm for $\phi = 0$). In all 17 experiments (*A-C*), we found Δt_0 and Δt_c to be independent of the examined parameters. In both cases, the random time until the completion of the next US sweep due to the significantly lower update rate compared to the robot was the determining factor. Finally, as depicted in Fig. 3.17c, the robot returned in the drift experiment to its initial position after 11 translations with an error of $e_p = 0.05$ mm, suggesting the absence of any drift over time. The average spatial lag after convergence of each section was 0.26 ± 0.15 mm.

The tracking performance for the human subject experiment is visualized in Fig. 3.18b, first showing translations perpendicular to the central US plane (*top*), followed by ones parallel to it (*bottom*). The spatial lag $\Delta\bar{s}$ during movements was on average slightly lower for the

Set	Parameter	N	$\Delta\bar{s}$	Δt_0 [s]	Δt_c [s]	e_p
A: Translation	v	6	9.27 ± 2.65 mm	0.46 ± 0.18	0.96 ± 0.18	0.59 ± 0.33 mm
B: Rotation	ω	4	$5.85° \pm 4.2°$	1.05 ± 0.14	1.36 ± 0.43	$0.61° \pm 0.45°$
C: Translation	ϕ	7	12.00 ± 0.51 mm	0.70 ± 0.11	1.16 ± 0.25	0.35 ± 0.19 mm

Tab. 3.1. Results of the three sets of visual servoing experiments, reporting the number of experiments per set N, average spatial lag $\Delta\bar{s}$, time until initial movement Δt_0, time until convergence Δt_c and the final positioning error e_p. *Reprint from [249] with permission of Springer.*

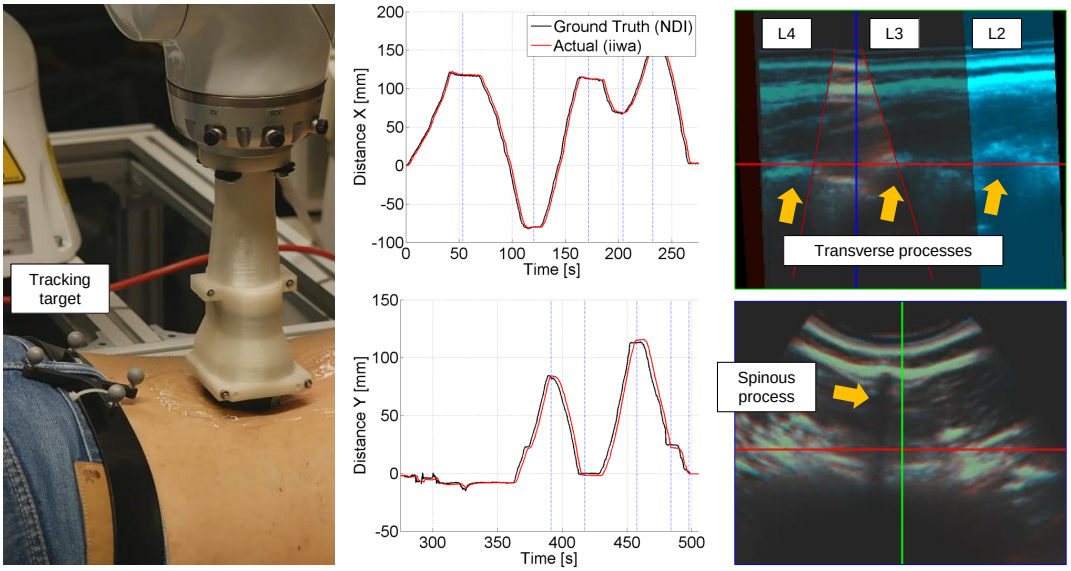

(a) (b) (c)

Fig. 3.18. **a)** Experimental setup for human volunteer experiments with optical tracking target. **b)** Trajectories of robot and body tracking target over 500 seconds, showing excellent visual servoing behavior. Translations along the x- and y-axes of F_t were performed consecutively. **c)** Exemplary initial US sweep (*blue*) with overlay of one motorized sweep after registration (*red*). Anatomical landmarks annotated for reference. *Reprint from [249] with permission of Springer.*

former (8.26 vs. 9.32 mm). In total, 11 translations were performed. The positioning error e_p after convergence of each segment, indicated as blue lines in Fig. 3.18b, was on average 1.45 ± 0.78 mm. The maximum error was 2.72 mm. Exemplary US images are shown in Fig. 3.18c, including both an initial sweep (*blue*) and one motorized sweep (*red*). Several anatomical structures as inherently used for alignment during continuous re-registration are highlighted for better understanding. As long as the subject maintained low tidal volumes, no detrimental respiratory disturbances in the horizontal plane were encountered. Similar to [126], maximum displacements in vertical direction were around 2 mm.

3.5.2 Needle Insertion Accuracy

All eight insertions into the facet joint were deemed technically successful by the expert spine surgeon (needle tip in the gap between the articular processes [144]), with average orthogonal positioning errors for robotic and conventional way of guidance of 1.36 ± 0.33 mm

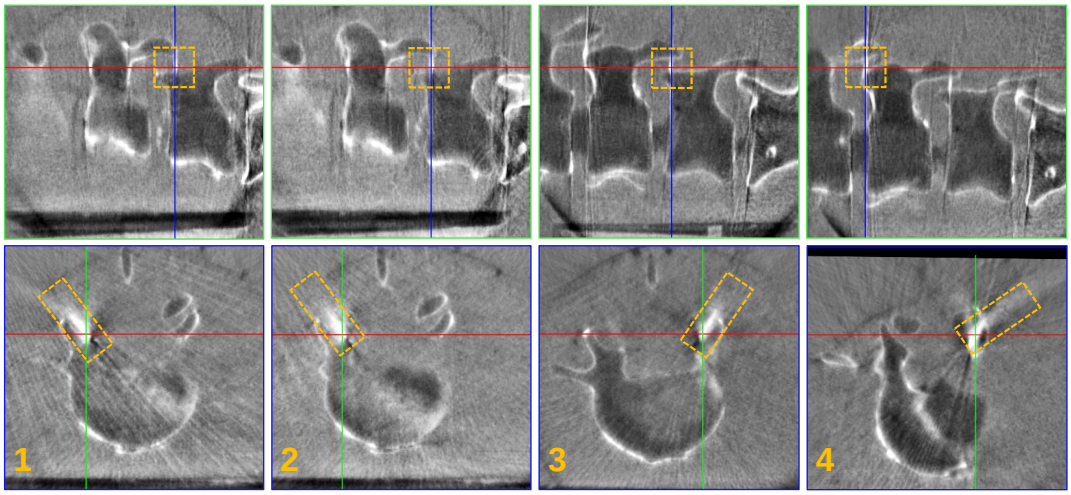

Fig. 3.19. Sagittal (*top*) and axial (*bottom*) slices of the verification cone-beam CT images after needle insertion, showing successful targeting of the facet joint in all cases (*yellow boxes*). *Reprint from [249] with permission of Springer.*

Needle	Robotic Assistance (Ultrasound)				Conventional (X-ray)		
	Site	Guide	e_d [mm]	$e_\perp$ [mm]	Site	e_d [mm]	$e_\perp$ [mm]
1	L4-5 R	closed	1.57	1.55	L4-5 R	4.29	3.22
2	L4-5 R	closed	1.46	1.20	L3-4 R	3.96	2.89
3	L4-5 L	open	1.83	1.71	L4-5 L	1.28	1.27
4	L3-4 L	open	3.60	0.97	L3-4 L	1.93	1.34
Average			2.12 ± 1.00	1.36 ± 0.33		2.86 ± 1.49	2.18 ± 1.02

Tab. 3.2. Results of the needle insertion experiments, comparing accuracies achieved with and without robotic assistance. *Reprint from [249] with permission of Springer.*

and 2.18 ± 1.02 mm, respectively. All results are reported in Tab. 3.2, and CT slices of the phantom after robot assisted insertions in Fig. 3.19. An exemplary registration result between pre-interventional CT and the initial US sweep I_0 is visible in Fig. 3.15b. Figure 3.15c shows an exemplary live US guidance view, which was used by the expert neurosurgeon to insert the needle. The average time from manual path planning to manual needle tip annotation in the CBCT image was 25 minutes for the robotic case, with routine established quickly so that the last iteration only took 11 minutes. As the conventional needle insertions were performed by an expert, the total applied X-ray dosage for navigation only amounted to 5.93 cGycm^2. For each of the CBCT acquisitions, 15.11 cGycm^2 were measured (rotation by $190°$ while acquiring 100 images in 60 seconds).

3.5.3 Aortic Aneurysm Screening Performance

Similar to [97], the Hausdorff distance between P_{RGBD} and the warped atlas surface mesh was on average 3.7 mm (maximum 9.8 mm), robustly allowing the visualization of the

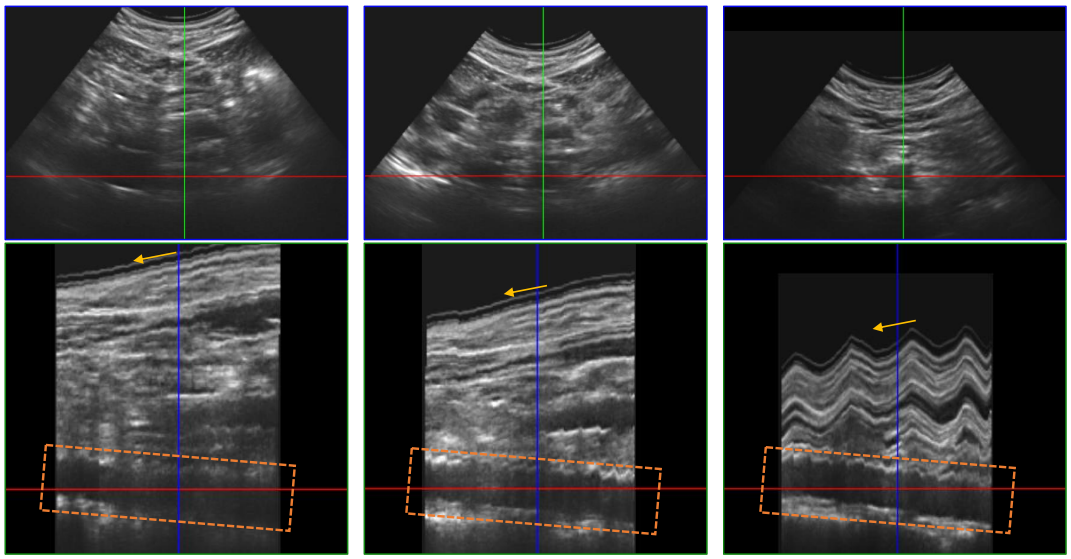

(a) $F = 5$ N, breath-hold **(b)** $F = 20$ N, breath-hold **(c)** $F = 20$ N, breathing

Fig. 3.20. Results of force adaptation experiments for one subject, including axial (*top*) and sagittal (*bottom*) slices of the compounded US volumes. Without proper force onto the tissue (**a**), the US quality at the aorta, in particular its posterior wall, is poor (*orange box*). Quality improves with the adaptive force estimation (**b**). Results also show that the force controller can successfully account for breathing motion and allows for steady aortic acquisitions, even if the US image of tissue directly underneath the skin becomes unusable (**c**). Yellow arrows indicate the scan direction. *Reprint from [230], modified, ©2016 IEEE.*

Model	a	b	MSE
Linear (ζ_{lin})	0.0081	0.0245	$6.8499 \cdot 10^{-4}$
Logarithmic (ζ_{log})	0.0727	-0.0452	$7.2828 \cdot 10^{-4}$

Tab. 3.3. Regression model parameters after fitting to describe relationship between exerted force on the tissue and confidence, including mean squared error (MSE). *See text for details.*

aorta (expected diameter < 4 cm) in the US frame (width ca. 10 cm) without further compensation.

In Fig. 3.20, a comparison between minimal force sweeps (*left*), sweeps with optimized contact force (*center*) and in addition during normal respiration (*right*) is shown. The vertical positions of the transducer and the exerted forces onto the tissue corresponding to these three sweeps are visualized in Fig. 3.21. In total, the estimated force F^* was 14.8±6.4 N for all volunteers. Results show that the force controller maintained the desired force with an average error of 0.17±0.24 N.

The dependence between exerted force on the tissue and confidence is visualized in Fig. 3.22. Table 3.3 presents the results of the regression model fit, assuming that forces were measured in N and confidence values as positive scalars in the interval $[0; 1]$. The fitted models are also shown in Fig. 3.22.

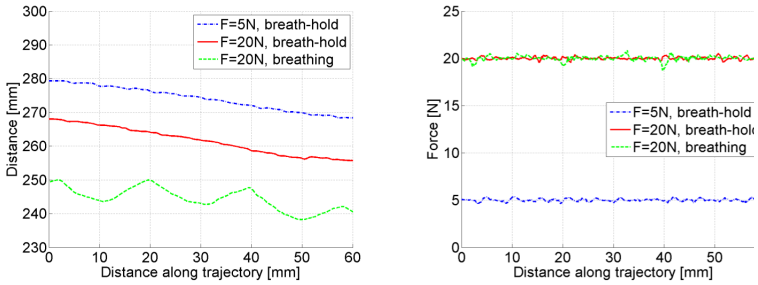

Fig. 3.21. **a)** Vertical position of transducer and **b)** corresponding exerted force during force adaptation experiments for one subject (see Fig. 3.20), showing the accuracy of the force controller and the elasticity of the tissue. *Reprint from [230], modified, ©2016 IEEE.*

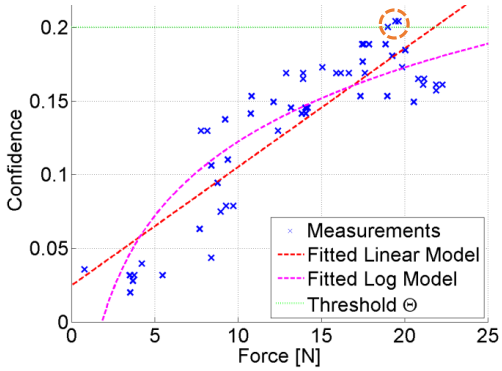

Fig. 3.22. Confidence values at the depth of the aorta during initial force estimation for one representative subject, showing a dependency on the exerted force. The final force of 20 N was reached after the confidence exceeded $\Theta = 0.2$ (*circled*). *Reprint from [230], modified, ©2016 IEEE.*

The difference between static rotation ($\alpha^* = 0$) and the proposed out-of-plane rotation estimation is depicted in Fig. 3.23 for a representative case. In particular in the beginning of the sweep, shadowing artifacts made the aorta not detectable. In contrast, the initial rotation estimation for a second sweep determined an optimal angle of $\alpha^* = 9.2°$, which led to a significantly increased confidence within the first 30 mm of the sweep. For the remaining trajectory, our controller gradually lowered the out-of-plane rotation and maintained high confidence. For all subjects, the optimal initial rotation was found to be $\alpha_0^* = 3.2 \pm 8.0°$.

Table 3.4 reports the results obtained in both scenarios. In Fig. 3.24, the selection of appropriate B-mode frames is shown, and Fig. 3.25 illustrates how manual measurements of aortic diameters were then obtained therein. On average, the error between manual US scan measurement and the ones performed in the robotically acquired volumes was 0.5±0.3 mm.

3.6 Discussion

The phantom motion compensation experiments indicated that visual servoing-based tracking of moving anatomies is possible with velocities of up to 25 mm/s. This should exceed the expected velocities in neurosurgical scenarios, which are dominated by respiratory mo-

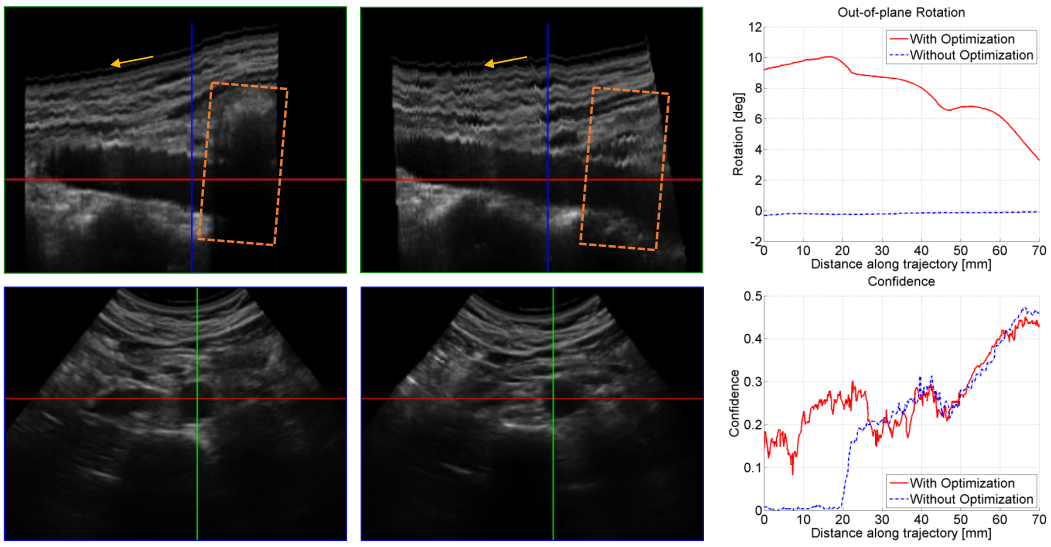

(a) Without optimization **(b)** With optimization **(c)** Comparison

Fig. 3.23. Comparison between static out-of-plane rotation and optimization based on confidence for one representative subject. The proposed method avoids the shadowing in the beginning of the trajectory (*orange box*). Throughout the acquisition, optimization of the rotation (**c**, *top*) maintains high confidence (**c**, *bottom*). Yellow arrows indicate the scan direction. *Reprint from [230], ©2016 IEEE.*

	Robotic Sweep		Manual Scan		Error	
	S [mm]	ED [mm]	S [mm]	ED [mm]	S [mm]	ED [mm]
d_{in}	13.50	11.68	13.38	11.58	0.32±0.13	0.54±0.22
d_{out}	16.68	15.00	16.68	14.84	0.48±0.28	0.56±0.34
d_{le}	15.08	13.34	15.00	13.08	0.52±0.38	0.46±0.15

Tab. 3.4. Average aortic diameter measurement results [mm] for systole (S) and end-diastole (ED) over all five subjects, comparing measurements obtained from sweeps acquired with the robot with conventional (manual) ones. *Reprint from [230], ©2016 IEEE.*

vements [126]. In combination with the spatial lag achievable by the system, higher velocities would lead to the target leaving the capture range of the similarity metric [70] or the field of view of the US transducer completely. The obtained results thus do not only show that tracking with submillimeter accuracy is possible as long as sufficient 3D image contrast is available to guide the image registration, but also that there is no drift over time. Both characteristics of the presented system are evenly important for interventions necessitating image guidance over longer time spans. Despite the small field of view and the limited 3D US frame rate, in turn reducing the spatial resolution in sweep direction, overall smooth movements with an oscillation-free convergence behavior (no overshoots) have been achieved. The average time until convergence reached a clinically acceptable level.

As expected, the system was able to successfully track human vertebrae under realistic conditions in volunteer experiments. Movements orthogonal to the central US frame were slightly lower than when the subject moved parallel to it. This is counterintuitive due to

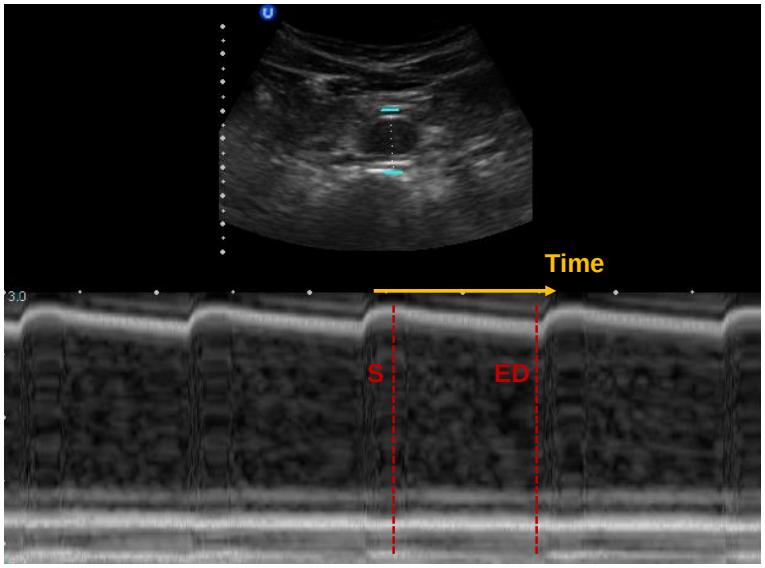

Fig. 3.24. Result of manual measurement as in clinical routine. **a)** M-mode frame showing cross-section of the aorta over time, allowing for the selection of B-mode frames corresponding to systole (*S*) and end-diastole (*ED*) for manual diameter measurement.*Reprint from [230], modified, ©2016 IEEE.*

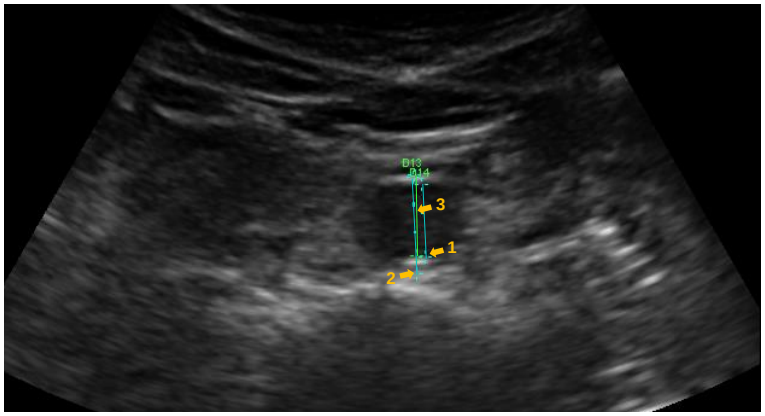

Fig. 3.25. B-mode frame of the aorta, showing how inner (*1*), outer (*2*), and leading-edge diameters (*3*) were manually determined. *Reprint from [230], modified, ©2016 IEEE.*

the lower resolution in this direction, and can possibly be explained by the better coverage with the initial sweep. Although the final positioning errors were higher than in phantom experiments, the achieved accuracy is well below the size of the target anatomy for facet joint injections [72].

Although the registration optimizer occasionally exceeded its fixed time budget, any temporarily increased spatial lag could be reliably compensated, yielding successful tracking outcomes in all experiments. Nevertheless, the issue could potentially be resolved by using the optimal transformation parameters at the end of the allotted time slot as initialization for the subsequent registration task. A Kalman filter [109] could furthermore improve the lag of the system and help avoiding the propagation of wrong local minima captures. In addition, the application of motion prediction algorithms as in [108] might be a starting point to better

cope with repetitive motions such as breathing with higher amplitude than in the conducted experiments, which was not investigated in this work.

One major limiting factor of the presented system is the update rate of the motorized US transducer. While the usage of 4D matrix probes or very high frame rate technology [171] will be suited to improve the rate at which 3D volumes can be acquired, the bottleneck of 3D image registration will remain, in particular if many degrees of freedom are considered for optimization. Adaptive strategies to for example reduce the image size (downsampling) during rapid movements until target approach might mitigate this limitation in future studies.

Promising needle placement results were obtained in the spine phantom experiments using the proposed needle guidance system. For the first time, a robotic ultrasound-based navigation system has been employed to allow manual needle insertions into the narrow facet joint with clinically sufficient accuracy [72]. No differences in needle placement accuracy could be identified concerning the type of needle guide (open vs. closed). In comparison to the proposed system, our X-ray guidance experiments demonstrated the placement variability of the conventional approach. Yet, the obtained errors should be interpreted with care, as the planned target positions were chosen to be further within the facet joint for better US-guided aiming, while the conventionally inserted needles approached the joints in a slightly steeper angle. The experimental setup in a realistic neurosurgical environment proved the clinical applicability of the system. In particular, the sharp decrease in procedure time throughout the conducted experiments underlined its minimal overhead and usability without extensive training. The compatibility with existing C-arm systems in terms of space and workflow requirements directly open the possibility of a subsequent clinical study, comparing the accuracy, efficacy and safety of automated US-based needle guidance with the gold standard of fluoroscopic facet joint injection, where X-ray confirmation of accurate placement will be required.

Regarding the aortic aneurysm screening, results demonstrated that confidence maps are a suitable method to derive meaningful visual servoing features and that the proposed control laws could improve the quality of 3D US acquisitions compared to naive sweeps. The employed adaptive force controller could successfully adapt to the physique of the scanned volunteers. To our surprise, a steady acquisition of the almost incompressible aorta is possible even while breathing. This might be useful in screening programs with elderly patients who may not be able to comply with breath-hold instructions for the entire scan duration. The initial estimation of the out-of-plane rotation proved very useful in avoiding shadowing artifacts due to bowel gas. However, the continuous improvement throughout the remainder of the acquisition was marginal in the majority of scanned individuals. Nevertheless, there was no confidence drop in any of the subjects compared to the static rotation scan, which suggests to continue considering such control laws in the future.

The obtained errors between aortic diameter measurements with the presented system and measurements following the clinical standard protocol were clearly below one millimeter and thus negligible for diagnostic purposes [72]. The preliminary study has proven that safe, fully autonomous robotic US acquisitions are feasible within clinical environments and that the quality of the obtained images are sufficient for embedding such a system in screening programs.

4 Toward 3D Doppler Vascular Imaging

4.1 Overview and Publications

This chapter presents the contributions of this thesis concerning the 3D reconstruction of blood velocities measured using Doppler ultrasound. Sonographic quantification of blood flow is not trivial because of two reasons: First, Doppler US only measures the projection of the true velocity in the beam direction and is thus blind to velocity components orthogonal to it. Second, due to pulsation of the vessel blood flow, acquisitions over time, for instance from multiple directions, cannot directly be related to each other and require some form of temporal modeling. In this chapter, novel techniques to reconstruct both a 3D velocity field and a temporal flow profile are presented. As a first step toward application of the proposed methods in clinical routine, their suitability for vascular screening is demonstrated, paving the way for more interventional scenarios with blood vessel involvement in the future.

This chapter is organized as follows: Sec. 4.2 outlines the clinical fields where Doppler sonography is frequently used, presents prior art on 3D blood flow reconstruction and summarizes this chapter's contributions. Thereafter, Sec. 4.3 establishes the mathematical foundation to reconstruct 3D velocity fields over time, before the experiments conducted on in-silico phantoms and human volunteers are explained in Sec. 4.4. Results are reported in Sec. 3.5 and thoroughly discussed in Sec. 3.6.

Substantial parts of this chapter have already been published in the following article and are quoted verbatim (especially the non-robotic sections 4.2.2, 4.3.1, 4.4, 4.5.1-2, and 4.6):

[251] **Oliver Zettinig***, Christoph Hennersperger*, Christian Schulte zu Berge, Maximilian Baust, and Nassir Navab. "3D Velocity Field and Flow Profile Reconstruction from Arbitrarily Sampled Doppler Ultrasound Data". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Vol. 8674 LNCS. PART 2. Springer Verlag, 2014, pp. 611–618. DOI: 10.1007/978-3-319-10470-6_76

* *These authors contributed equally to this worky.*

A more extensive journal article concerning the entire content of this chapter is currently in preparation. Apart from the team that contributed to the publication listed above, also Thiemo Taube shall be mentioned here. Under my supervision, he concentrated on 3D reconstructions of robotic Doppler measurements in the course of his Master's thesis [217].

4.2 Introduction

4.2.1 Clinical Background

Since the advent of Doppler sonography, in particular since the breakthrough of color-flow mapping, this imaging modality has gained significant popularity and is used in a broad variety of clinical examinations today [57]. As a generic tool for blood flow assessment, Doppler US is suitable for diagnosis and observation in numerous parts of the human body. Without claiming to be exhaustive, the following list summarizes the fields where Doppler US is most commonly used in today's practice:

i) **Echocardiography** is often concentrating on the function of the heart as opposed to purely anatomical information. As a result, Doppler modes play an essential role in the assessment of valvular regurgitation, abnormal communication conditions between the two sides, (septal) wall motion, and the calculation of cardiac output [162]. Apart from classical transthoracic imaging protocols, transesophageal echocardiography (TEE) has gained popularity due to the better anatomical access, especially in combination with 3D matrix probes. In recent years, advanced understanding in intra-ventricular blood dynamics have triggered a more in-depth assessment of vortex structures, which redirect the blood in form of jets toward the outflow tracks [164].

ii) Following examinations of the heart, **vascular diseases** form the second-largest area of Doppler sonography [57]. In particular, the imaging modality is widely utilized to measure stenosis of blood vessels, for instance as a result of atherosclerosis, or thrombosis, often occurring in the deep veins [86]. In particular due to the systemic nature of atherosclerosis [100], non-invasive screening using US of easily accessible vessels allows appropriate risk classification and patient stratification. In the case of the carotid artery, this immediately concerns conditions such as transient ischemic attacks and strokes, but indirectly also indicates an increased risk for coronary problems and heart attacks. Surveillance is especially important in patients with known cardiovascular diseases and in temporal proximity to surgeries. Even though there is an ongoing discussion whether carotid Doppler examinations can reduce the risk of perioperative stroke after major cardiac procedures, selective measures are recommended for patient management [6, 133].

iii) In the **kidney**, sonographers often determine the *renal resistive index* using Doppler techniques [228]. The assessment is useful in the detection and management of renal artery stenosis, chronic kidney disease, and chronic renal allograft rejection. Apart from intrarenal perfusion, evidence has furthermore linked the index also to systemic hemodynamics including subclinical atherosclerosis and primary hypertension. In a pilot study, also novice sonographers reached good reliability in their assessment, making Doppler US an easy-to-learn method for routine usage [190].

iv) Ultrasound is also used in the **liver** to diagnose abnormal hepatic blood flow conditions such as reversed portal venous blood flow as a result of portal hypertension [107]. While

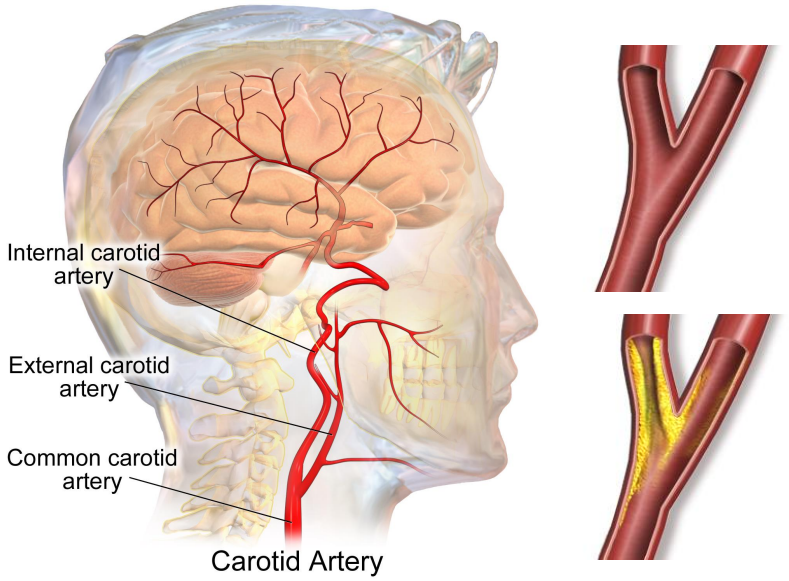

Fig. 4.1. Schematic overview of the common, external and internal carotid artery, indicating their regions of blood supply. On the right, a healthy artery is compared to a stenosis as usually caused by atherosclerosis. *From [19], CC-BY-3.0, modified.*

no additional value of Doppler US could be confirmed for liver fibrosis [14], the method has also been proven useful in diagnosis and management of cirrhosis [25].

v) In a variety of other use cases, Doppler imaging has been found useful. The vascularization of breast tumors as identified with power Doppler, for instance, has been confirmed to be a predictor of tumor grades, which is likely to be applicable to other forms of malignancies as well [34]. Other areas include the measurement of hypervascularization in rheumatology to detect inflammations, certain gynecology and urology applications, for instance to measure endometrial vascularity or assess testicular torsion. Also dynamic analysis of non-blood fluids such as semen in cases of ejaculatory dysfunction has been reported [90].

As a first application, this chapter focuses on the quantification of blood flow in the carotid artery as shown in Fig. 4.1. As a result of atherosclerosis, a thickening of the inner vessel wall by build-up of fibrofatty plaque, the lumen of the vessel is reduced. The most dangerous complications of such a stenosis is the sudden rupture of soft plaque, causing the formation of a thrombus that is able to block the blood flow completely, with devastating consequences to the cells fed by it. Depending on the affected vessel, these events present as myocardial infarction (coronary artery) or as ischemic stroke (carotid artery).

Stroke alone is one of the leading causes of death, with an approximately 1.1 million inhabitants of Europe suffering a stroke each year, and a one-month case-fatality rate between 13 and 35% [12]. Of these cases, roughly 80% account for ischemic strokes. In the USA, every 40 seconds, on average, someone experiences as stroke, and someone dies of one approximately every 4 minutes [145]. Up to 50% of surviving patients also suffer from poor outcomes, including cognitive disorders and depression.

In clinical routine, Doppler imaging is clearly not the only modality capable to assess blood flow. Because conventional Doppler sonography suffers from limited anatomical accessibility as well as high inter-observer variability, alternatives are in frequent use today, including MRI, digital subtraction angiography (DSA) and its tomographic counterpart (CTA) [163]. Conventional MR angiography (MRA) as well as the X-ray based techniques are hereby limited to morphological data. More advanced studies of blood dynamics can be obtained by more recently presented 4D MRI protocols. In a study on the flow characteristics in the carotid artery bifurcation, authors confirmed that flow-sensitive 4D MRI yields similar pulsatility index and resistance index measurements as Doppler examination [91]. Yet, these modalities, in particular the tomographic ones, are by far not as widely available as US imaging and regularly require longer acquisition times and yield inferior frame rates. In addition, all of these methods expose the patient to nephrotoxic contrast agents (Gadolinium for MRI, iodine for X-ray contrast). Finally, DSA and CTA also involve exposure to ionizing radiation.

4.2.2 2D and 3D Doppler Reconstruction

Doppler US suffers from one particular physical limitation. Because it can only measure the projection of the true velocity vector along the echo beam direction, it is thus blind to the flow orthogonal to it. As a consequence, examiners need experience in order to make qualitative assertions regarding the three-dimensional blood flow, which also explains the recent endeavors to reconstruct three-dimensional flow fields from multiple Doppler acquisitions to overcome this limitation.

Fox [64] was the first one to utilize simultaneous measurements from multiple directions to recover 3D flow information already in 1978. His experimental setup included two US transmitters, in cross-beam arrangement, and a single US receiver element. After the derivation of a close-form solution to compute calibrated 3D velocity vectors, the method was successfully validated using a revolving turntable. This design inspired a variety of related methods with intersecting US rays, also utilizing beam steering techniques on linear transducers instead of relying on various single US elements [51]. Hussain et al. [105] later extended this concept, termed *Vector Doppler*, to 3D by surrounding a single transmission element with multiple receivers in a hexagonal architecture, or by rotating a linear array around central transmission elements [106]. Pihl et al. [166] directly used a matrix probe and switched between different beam configurations. In contrast, Giarre et al. [79] investigated the use of subsequent measurements at different locations and introduced laser Doppler anemometry (LDA, [125]) as ground truth method for validation. Such techniques have successfully been applied under laboratory conditions, for instance to optimize spiral flow patterns in vascular prostheses [118], they have not yet translated to clinical routine.

A fundamentally different approach to obtain velocity information of more than one dimension, speckle tracking, emerged at the end of the 1980s. Trahey et al. [221] proposed to rely on B-mode imaging data and to optimize for patch-wise correlation in order to approximate 2D velocities. Essentially implementing a variant of optical flow estimation, many other groups presented approaches to track the movement of speckle generated by interference of of the backscattered ultrasonic waves [93]. Swillens el al. [214] compared speckle tracking and vector Doppler techniques in a simulation study and found both methods may be suitable

for clinical use. However, they concluded that while the former might be a more practical approach, the latter can provide superior axial velocity measurements.

Plane-wave imaging promises to alleviate some of the shortcomings in terms of robustness of the vector Doppler systems outlined above. Ekroll et al. [54] introduced an in-vivo US system using tilted plane waves and parallel receive beamforming, reaching more than 50 frames per second in a duplex configuration. In a broad extension using very high frame rate technology [171], Correia et al. [43] recently presented 4D ultra-fast flow imaging with matrix probes, tilted plane wave transmission in 3D and multiple sub-apertures for receiving. The authors indicate, however, the difficult choice and the number of the transmitted tilted angles for specific clinical applications and the small field-of-view as current limitations. In addition, ultra-fast US technology is still not widely available, even for research purposes.

A number of works approach the problem of measuring 3D blood flow from a different direction. Garcia et al. [73] proposed to compute a time-resolved 2D velocity field from conventional color Doppler data by incorporating ventricular wall motion. Under planar flow assumption, the wall position constraints thus replace the otherwise required second Doppler direction. Although the approach allowed for successfully reconstruction of cardiac vortex patterns to some extend, it is limited by its 2D nature and the additional constraints. Arigovindan et al. [8] introduced B-spline grids to regularize the reconstruction problem for spatially non-uniformly sampled Doppler data from multiple directions. Based on their work, Gómez et al. [83] presented a method to quantify 3D blood flow from multiple registered B-Mode and Doppler volumes acquired by a 2D matrix array probe. However, the method requires pulse phase consistency, i.e. samples acquired at (or interpolated for) the same instant in time, and fails if temporally arbitrary Doppler samples are used. Usage in adults is furthermore impaired due to the limited field of view of the employed matrix probes. An impressive extension of their method added time as forth dimension for the reconstruction problem, and also incorporated ventricular wall motion constraints again [84]. Just as the algorithm presented in this chapter, their work also assumes velocity consistence across consecutive cardiac cycles, but in addition, also a constant heart rate is intrinsically modeled. Since the wall motion estimated was based on segmentations, which were carried out manually by medical experts, and separate 3D sequences had to be manually registered using landmarks, direct clinical implementation is not straightforward.

4.2.3 Contributions

This chapter introduces a novel technique to reconstruct, at the same time, a continuous three-dimensional velocity field of blood flow and a continuous flow profile over time. Employing B-spline basis functions, the method can cope with arbitrarily sampled Doppler data in space and time.

The flow profile is inspired by the work of Waechter et al. [232], who employ a waveform model to regularize the estimation of blood flow parameters from rotational angiography. The presented model assumes a that the flow direction at any point in the reconstructed volume is constant over time, and only its magnitude is changing due to pulsation. Such a model has been shown to be applicable with negligible errors unless bifurcations, high-grade stenoses, or

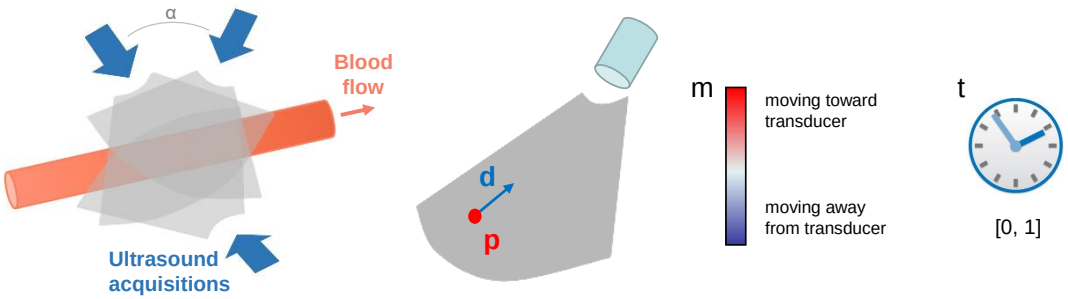

Fig. 4.2. **a)** Three ultrasound sweeps from different directions are necessary to reconstruct the velocity field in a blood vessel. *Reprint from [251], with permission of Springer.* **b)** 4-tuples $(\mathbf{p}, \mathbf{d}, m, t)$ indicating position, beam direction, scalar Doppler measurement and normalized time, serve as input.

aneurysms are investigated [232]. The posed assumptions do not, however, enforce that the blood flow is organizing in parallel sheets, so that also more complex patter can be potentially reconstructed.

Due to the temporal regularization provided by the waveform model, the proposed method is gating-free in the sense that no data selection according to a gating signal is required and that the reconstruction is performed over the entire dataset at once. Therefore, it supports temporally arbitrary Doppler samples from any pulse phase as input. Moreover, the proposed approach is also compounding-free as it is based on the individual Doppler samples in the form of raw in-phase and quadrature or color-mapping data. Thus, the method also allows spatially arbitrary, non-uniform data points, without prior compounding of US volumes and associated interpolation artifacts.

The proposed method is validated using in-silico phantoms and pre-clinically on healthy volunteer datasets, which were either acquired manually using a tracking system or using a robotic manipulator. The presented results successfully demonstrate the feasibility to reconstruct 3D blood flow information over time using sets of measurements from multiple directions.

4.3 Methods

After introducing the concept of B-spline grids to model arbitrary signals, Sec. 4.3.1 first derives a formulation to represent a 3D velocity field and then extends the model by a temporal flow profile. Based on a description how to numerically represent the model's components, Sec. 4.3.2 outlines how the B-spline coefficients are eventually determined in a least-squares sense.

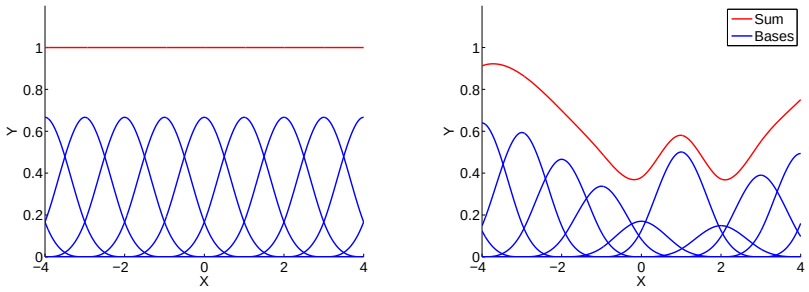

Fig. 4.3. Principle of modeling signals using a linear combination (sum in *red*) of B-spline basis functions β (*blue*). **a)** Uniform weights c result in a straight line. **b)** Up to the resolution of the grid, arbitrary signals can be represented with appropriate weights c.

4.3.1 Combined Spatial and Temporal B-spline Model

As shown in Fig. 4.2a, the reconstruction method relies on Doppler samples acquired from multiple directions as input. In particular, 4-tuples of $(\mathbf{p}, \mathbf{d}, m, t)$ are collected for every sample, where $\mathbf{p} = [p_x, p_y, p_z]^\intercal$ denotes the sample position in Cartesian space, $\mathbf{d}$ the corresponding normalized echo beam direction, m the measured Doppler signal, and $t \in [0, 1]$ the normalized pulse phase time (see Fig. 4.2b for reference).

Three-Dimensional B-spline Grid for each Velocity Component

For the static case as in [83], a smooth 3D velocity field $\mathbf{v}(\mathbf{p})$ can be described using B-spline basis functions, as illustrated in Fig. 4.3, by modeling each of the three Cartesian components individually:

$$\mathbf{v}(\mathbf{p}) = [v_x(\mathbf{p}), v_y(\mathbf{p}), v_z(\mathbf{p})]^\intercal, \tag{4.1}$$

The components of $\mathbf{v}(\mathbf{p})$ are then defined as:

$$v_x(\mathbf{p}) = \sum_{i=-w}^{w} \sum_{j=-w}^{w} \sum_{k=-w}^{w} c^x_{i,j,k}\, S_{i,j,k}(\mathbf{p}), \tag{4.2}$$

$$v_y(\mathbf{p}) = \sum_{i=-w}^{w} \sum_{j=-w}^{w} \sum_{k=-w}^{w} c^y_{i,j,k}\, S_{i,j,k}(\mathbf{p}), \tag{4.3}$$

$$v_z(\mathbf{p}) = \sum_{i=-w}^{w} \sum_{j=-w}^{w} \sum_{k=-w}^{w} c^z_{i,j,k}\, S_{i,j,k}(\mathbf{p}), \tag{4.4}$$

where

$$S_{i,j,k}(\mathbf{p}) = \beta^n_{s,i}(p_x)\beta^n_{s,j}(p_y)\beta^n_{s,k}(p_z). \tag{4.5}$$

The indices i, j, and k indicate a location in the regular, three-dimensional B-spline grid with spacing s, which defines the resolution of the recovered vector field. The used patch size is defined as $2w + 1$. $\beta^n_{s,i}(p_x) = \beta^n(i/s - p_x)$ denotes the scaled and translated B-spline basis

function of degree n at grid position i, likewise for $\beta^n_{s,j}(p_y)$ and $\beta^n_{s,k}(p_z)$. Finally, $c^x_{i,j,k}$, $c^y_{i,j,k}$, and $c^z_{i,j,k}$ are the corresponding B-spline coefficients to be found.

The original approach in [83] relies on an input dataset with N samples sharing the same pulse phase, from which the reconstruction is performed by minimizing the error between the measured Doppler value denoted by m_r and the projection of the recovered velocity at sample position $\mathbf{p}_r$ onto the beam direction $\mathbf{d}_r$

$$J_{proj}(\mathbf{v}) = \sum_{r=1}^{N} \|\mathbf{d}_r \cdot \mathbf{v}(\mathbf{p}_r) - m_r\|_2^2 \,. \tag{4.6}$$

However, this formulation has to be adjusted for different pulse phases because the blood flow at every point within an artery varies over time. Thus, uncorrected Doppler measurements lead to inconsistent samples, prohibiting the reconstruction of the velocity field at a given instant. In this work, we introduce a flow profile function

$$\varphi(t)\colon [0,1] \to [0,1] \tag{4.7}$$

mapping from normalized pulse phase times to scaling factors compensating the varying flow patterns. An exemplary flow profile function, which is also used as ground truth for phantom experiments in Sec. 4.4.1 is visualized in Fig. 4.4. The actual velocity vector at any point in time is then given by

$$\mathbf{v}(\mathbf{p}, t) = \varphi(t)\, \mathbf{v}_{max}(\mathbf{p}), \tag{4.8}$$

assuming a pulsatile laminar flow. While allowing varying velocity magnitudes over time, our formulation keeps the flow direction constant at any position in the reconstructed volume. Unless bifurcations, high-grade stenoses, or aneurysms are investigated, this model has been shown to introduce only negligible errors [232]. It should be noted, however, that this formulation does not enforce the velocity directions to be organized in parallel sheets. In order to compute φ, we employ a B-spline parametrization as well:

$$\varphi(t) = \sum_{\xi=1}^{T} c^t_\xi \beta^n_{\tau,\xi}(t), \tag{4.9}$$

where $\beta^n_{\tau,\xi}(t)$ is a periodic B-spline basis function of degree n and spacing τ evaluated at grid point ξ, and c^t_ξ denote the corresponding temporal B-spline coefficients. As a result, we propose the cost function

$$J_{proj}(\mathbf{v}, \varphi) = \sum_{r=1}^{N} \|\varphi(t_r)\, \mathbf{d}_r \cdot \mathbf{v}_{max}(\mathbf{p}_r) - m_r\|_2^2 \qquad \text{s.t.} \quad \max_t |\varphi(t)| = 1 \tag{4.10}$$

The constraint in Eq. 4.10 is added to the system to ensure that $\mathbf{v}_{max}$ is correctly scaled.

Numerical Representation

In order to solve for B-spline coefficients modeling a given dataset, a numerical representation of involved quantities is necessary. For assembling a linear system of equations, the indices i, j, and k of the 3D B-spline grid first need to be vectorized:

$$\text{vec}(i,j,k) \in [1, M] = [(i+w)(2w+1) + j + w] \cdot (2w+1) + k + w + 1, \tag{4.11}$$

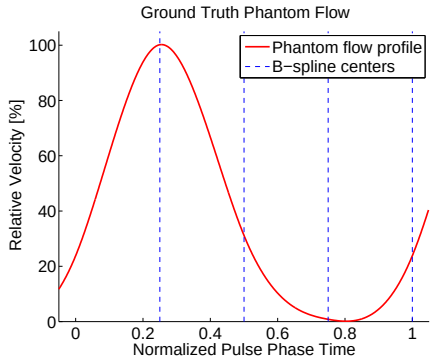

Fig. 4.4. Ground truth flow profile function φ defined using periodic B-spline coefficients.

where i, j, and k are in the interval $[-w, w]$, again yielding a patch size of $2w+1$. Thus, $M = (2w+1)^3$ coefficients are considered in total. The spatial B-spline coefficients are gathered in column vector $\mathbf{C}_s \in \mathbb{R}^{3M} = [\mathbf{C}_s^x \mathbf{C}_s^y \mathbf{C}_s^z]^\intercal$ with elements:

$$\{\mathbf{C}_s^x\}_{1,\mathrm{vec}(i,j,k)} = c_{i,j,k}^x, \tag{4.12}$$

and analogously for $\mathbf{C}_s^y$ and $\mathbf{C}_s^z$. Similarly, the temporal B-spline coefficients are collected in column vector $\mathbf{C}_t \in \mathbb{R}^T$ so that

$$\{\mathbf{C}_t\}_{\xi,1} = c_\xi^t. \tag{4.13}$$

Using the vectorization scheme of Eq. 4.11, the evaluation of the B-spline functions is encoded by sampling matrices $\mathrm{S}_s \in \mathbb{R}^{3N \times 3M}$

$$\mathrm{S}_s = \begin{bmatrix} \mathrm{S}_p & 0 & 0 \\ 0 & \mathrm{S}_p & 0 \\ 0 & 0 & \mathrm{S}_p \end{bmatrix} \tag{4.14}$$

with elements

$$\{\mathrm{S}_p\}_{r,\mathrm{vec}(i,j,k)} = \beta_{s,i}^n(p_{r,x})\beta_{s,j}^n(p_{r,y})\beta_{s,k}^n(p_{r,z}) \tag{4.15}$$

and $\mathrm{S}_t \in \mathbb{R}^{N \times T}$ with elements

$$\{\mathrm{S}_t\}_{r,\xi} = \beta_{\tau,\xi}^n(t_r). \tag{4.16}$$

This allows to represent $\varphi = \mathrm{S}_t \mathbf{C}_t$, and $\mathbf{v}_{max} = [\mathbf{v}_{max}^x \mathbf{v}_{max}^y \mathbf{v}_{max}^z]^\intercal = \mathrm{S}_s \mathbf{C}_s$. The dot product is realized using the direction matrix $\mathrm{D} = [\mathrm{D}^x \mathrm{D}^y \mathrm{D}^z]$ containing diagonal matrices with the three Cartesian components of the beam directions, respectively: $\{\mathrm{D}^x\}_{r,r} = d_{x,r}$, likewise for D^y and D^z. With $\mathbf{m} = \{m_r\}$ referring to the vector of Doppler measurements and $\otimes$ to the component-wise multiplication, the cost function reads

$$J_{proj}(\mathbf{C}_s, \mathbf{C}_t) = \left\| \left(\underbrace{\mathrm{S}_t}_{(N\times T)} \underbrace{\mathbf{C}_t}_{(T\times 1)} \right) \otimes \left(\underbrace{\mathrm{D}}_{(N\times 3N)} \underbrace{\mathrm{S}_s}_{(3N\times 3M)} \underbrace{\mathbf{C}_s}_{(3M\times 1)} \right) - \underbrace{\mathbf{m}}_{(N\times 1)} \right\|_2^2. \tag{4.17}$$

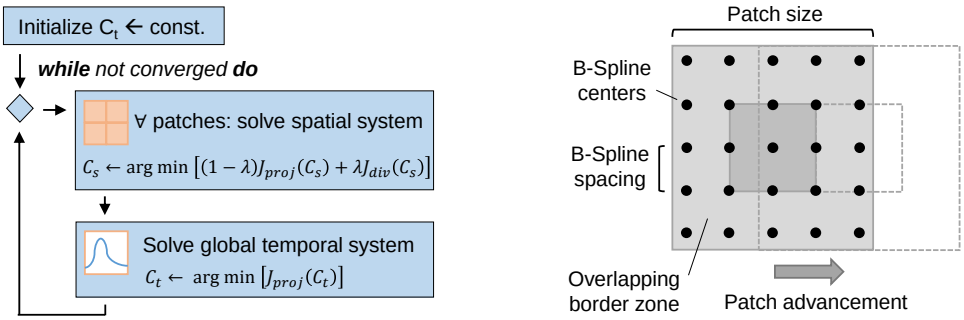

Fig. 4.5. **a)** Algorithm to solve spatial velocity field and temporal flow profile function. **b)** Illustration of the spatial patch-wise reconstruction scheme. After reconstruction, only the estimated B-spline coefficients $\mathbf{C}_s$ in the non-overlapping core zone are stored. *Reprint from [251], with permission of Springer.*

For smoothness, a spatial regularization term is added to the final cost function:

$$J(\mathbf{C}_s, \mathbf{C}_t) = (1-\lambda)\; J_{proj}(\mathbf{C}_s, \mathbf{C}_t) + \lambda\; J_{div}(\mathbf{C}_s). \tag{4.18}$$

As shown in [83], the locally acting divergence term $J_{div} = \|\nabla \cdot \mathbf{v}_{max}\|_2^2$, derived using the B-spline formulation, enforces the incompressibility of the blood flow. It is computed, in matrix form, as follows:

$$J_{div}(\mathbf{C}_s) = \dot{\mathrm{S}}_s\; \mathbf{C}_s, \tag{4.19}$$

whereby the matrix $\dot{\mathrm{S}}_s$ is assembled as $\dot{\mathrm{S}}_s = \left[\dot{\mathrm{S}}_s^x \dot{\mathrm{S}}_s^y \dot{\mathrm{S}}_s^z\right]$ with elements

$$\{\dot{\mathrm{S}}_s^x\}_{r,\mathrm{vec}(i,j,k)} = \dot{\beta}_{s,i}^n(p_{r,x})\beta_{s,j}^n(p_{r,y})\beta_{s,k}^n(p_{r,z}), \tag{4.20}$$

and analogous formulations for $\dot{\mathrm{S}}_y$ and $\dot{\mathrm{S}}_z$ with the respective B-spline function replaced by its derivative $\dot{\beta}$. It should be noted that the flow profile function φ inherently acts as temporal regularization itself, rendering further regularization terms unnecessary.

4.3.2 3D Velocity Reconstruction over Time

Unfortunately, solving for both $\mathbf{v}_{max}$ and φ simultaneously is not possible. Instead, spatial and temporal coefficients are obtained in an alternating fashion as reported in Fig. 4.5a. For the first iteration, the flow profile is initialized to be constant, i.e. $\varphi = 1$.

Spatial Reconstruction

The spatial solution directly follows the method in [83] and assumes a given global flow profile function φ to be constant. Exploiting the local support of B-splines, the velocity field $\mathbf{v}_{max}$ can then be reconstructed directly in a patch-wise fashion. As illustrated in Fig. 4.5b, patches mutually overlap for increased stability, and only the core, non-overlapping coefficients in vector $\mathbf{C}_s$ (dark gray in the figure) are stored for assembling the final output.

For better readability, we introduce:

$$\mathrm{Z} = \Big(\underbrace{(\mathrm{S}_t \mathbf{C}_t)}_{(N\times 1)} \underbrace{[1,1,\ldots,1]}_{(1\times 3M)} \Big) \otimes \underbrace{(\mathrm{DS}_s)}_{(N\times 3M)} . \tag{4.21}$$

Note that this expression essentially performs a component-wise multiplication of vector $\mathrm{S}_t\mathbf{C}_t$ with all columns of matrix DS_s. Minimization is performed by derivation of the cost function and equating to zero:

$$\frac{\partial}{\partial \mathbf{C}_s} J = 0 \qquad \rightarrow \qquad \mathbf{C}_s^\intercal \left((1-\lambda)\mathrm{Z}^\intercal\mathrm{Z} + \lambda \dot{\mathrm{S}}_s^\intercal \dot{\mathrm{S}}_s \right) = (1-\lambda)\mathrm{Z}^\intercal \mathbf{m}. \tag{4.22}$$

The resulting linear system of equations can be solved with an LU decomposition scheme with partial pivoting, directly yielding the spatial coefficients $\mathbf{C}_s$ for each patch.

Temporal Reconstruction

While we used a Levenberg-Marquardt solver for our early efforts in [251] to solve for coefficients $\mathbf{C}_t$, it can be shown that also the temporal system holds a closed-form solution [217].

Because of its global nature, all samples in every patch $q \in Q$ (Q denotes the set of all patches) need to be considered in solving for φ. Again, we introduce for better readability:

$$\mathrm{Y}_q = \underbrace{\mathrm{S}_{tq}}_{(N_q\times T)} \otimes \Big(\underbrace{(\mathrm{D}_q \mathrm{S}_{sq} \mathbf{C}_{sq})}_{(N_q\times 1)} \underbrace{[1,1,\ldots,1]}_{(1\times T)} \Big). \tag{4.23}$$

Note that this expression essentially performs a component-wise multiplication of the patch-specific vector $\mathrm{D}_q\mathrm{S}_{sq}\mathbf{C}_{sq}$ with all columns of the patch-specific matrix S_{tq}. For the minimization, the spatial regularization can be ignored because it does not depend on $\mathbf{C}_t$. It is thus sufficient to optimize for J_{proj}:

$$\frac{\partial}{\partial \mathbf{C}_t} J_{proj} = 0 \qquad \rightarrow \qquad \gamma\, \mathbf{C}_t^\intercal \sum_{q\in Q} \mathrm{Y}_q^\intercal \mathrm{Y}_q = \sum_{q\in Q} \mathrm{Y}_q^\intercal \mathbf{m}_q. \tag{4.24}$$

In this cost function, an unknown normalization factor γ is supposed to ensure the constraint of φ not exceeding 1. Since γ only reciprocally scales the final coefficients $\mathbf{C}_t$, it is possible to ignore it for solving the resulting linear system of equations, similar to the spatial reconstruction described above. Afterwards, the obtained coefficients are normalized by the maximum velocity scaling factor of all patches:

$$\gamma = \max_q \max_t |\mathrm{S}_{tq}\mathbf{C}_t|. \tag{4.25}$$

As shown in Fig. 4.5a, the spatial and temporal steps are repeated in an alternating fashion until convergence. Eventually, the algorithm yields two continuously defined results: the reconstructed maximum velocity field $\mathbf{v}_{max}$ and the blood flow profile φ.

4.4 Materials and Experiments

4.4.1 In-silico Phantom Evaluation

For evaluating our method, three virtual in-silico phantom datasets were created, each spanning 34×34×34 mm. The first two phantoms, denoted *Linear* and *Curved*, contained a cylinder and a torus with a vessel radius of $R = 6$ mm, respectively. A quadratic Poiseuille flow as in [92] with a maximum velocity of 1 m/s in the centerline was applied, see Fig. 4.6. Assuming that the linear vessel passes the origin, and the unit vector $\mathbf{b}$ defines the direction of the vessel, we can compute the distance of an arbitrary point to the centerline as:

$$r(\mathbf{p}) = \|\mathbf{p} - \mathbf{b}\,(\mathbf{p} \cdot \mathbf{b})\|_2. \tag{4.26}$$

A similar expression is possible for the torus-shaped vessel. Knowing the distance from the centerline allows for a concise velocity definition for the first two phantoms (see Fig. 4.6):

$$\mathbf{v}_{max,ph}(\mathbf{p}) = \begin{cases} a\left[1 - \left(\frac{r(\mathbf{p})}{R}\right)^2\right]\mathbf{b} & \text{if } r(\mathbf{p}) \leq R, \\ 0 & \text{otherwise.} \end{cases} \tag{4.27}$$

The velocity for the third phantom, denoted *Barrel Roll*, was defined as

$$\mathbf{v}_{max,ph}(\mathbf{p}) = a[\mathbf{p} - \mathbf{b}\,(\mathbf{b} \cdot \mathbf{p})] \times \mathbf{b}, \tag{4.28}$$

where $\mathbf{b}$ now defines the direction of the roll. In both cases, a represents a scaling factor to ensure $\|\mathbf{v}_{max,ph}\| \leq 1$. In all cases, a phantom flow profile function φ_{ph} as shown in Fig. 4.4

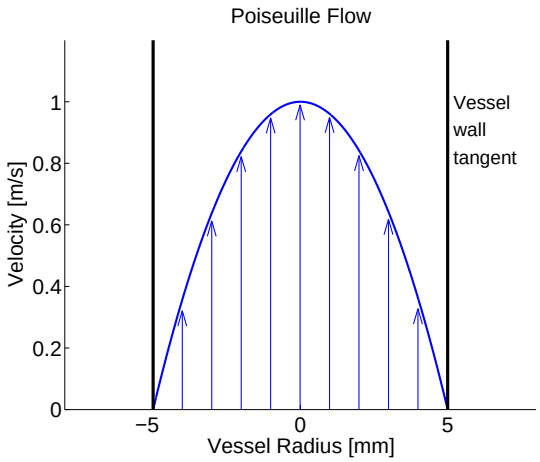

Fig. 4.6. Velocity distribution for cross-sections through *Linear* and *Curved* phantoms, exhibiting quadratic Poiseuille flow.

was used to construct a time-varying flow profile. Finally, synthetic Doppler sweeps from three directions $\mathbf{d}$ were generated with in total N measurements m computed as follows:

$$m = \varphi_{ph}(t \sim [0,1])\, \mathbf{d} \cdot \mathbf{v}_{max,ph}(\mathbf{p}) + \eta_s + \eta_v. \tag{4.29}$$

The pairwise angle between the sweeps, denoted opening angle α as illustrated in Fig. 4.2a, was varied between $10°$ and $90°$. For the linear phantom, the mean direction $\bar{\mathbf{d}}$ was set to be $45°$ tilted against the vessel direction. The sample positions $\mathbf{p}$ were chosen from a regular grid with 1 mm spacing. Hereby, different grid origin positions were used for each direction $\mathbf{d}$ to avoid sample overlap. Pulse phase times t to evaluate φ_{ph} were randomly drawn from a uniform distribution.

Two different Gaussian noise models were considered: On the one hand, noise with standard deviation σ was added to the samples such that $\eta_s = \mathcal{N}(0, \sigma)$ simulated inaccuracies of the Doppler measurement itself. On the other hand, noise was added to the velocity vector under investigation such that

$$\eta_v = \mathcal{N}(0, \sigma)\, \frac{\mathbf{v}}{\|\mathbf{v}\|} \cdot \mathbf{d}. \tag{4.30}$$

This modeled inaccuracies of the tracking stream or badly synchronized temporal data.

The reconstruction was performed using 6×6×6 mm patches, with 2 mm overlap, for which only ρN samples were used to mimic sparse sampling (data density $\rho \in [0.6, 1]$). The spatial spacing of the cubic ($n = 3$) B-splines was fixed at $s = 1.5$ mm (see Sec. 4.3.1) Due to the high runtime complexity of the alternating algorithm (Fig. 4.5a), a coarse temporal spacing of $\tau = 0.25$ was used. For all experiments, moderate regularization ($\lambda = 0.1$) was employed.

Three measures were used to analyze reconstruction errors of computed velocities $\mathbf{v}_{max,c}$ and the estimated flow profile φ_c. While the mean cosine similarity

$$s_{cos} = \frac{1}{N} \sum_{r=1}^{N} \frac{\mathbf{v}_{max,c}(\mathbf{p}_r)}{\|\mathbf{v}_{max,c}(\mathbf{p}_r)\|_2} \cdot \frac{\mathbf{v}_{max,ph}(\mathbf{p}_r)}{\|\mathbf{v}_{max,ph}(\mathbf{p}_r)\|_2} \tag{4.31}$$

and the mean Euclidean norm

$$e_d = \frac{1}{N} \sum_{r=1}^{N} \|\mathbf{v}_{max,c}(\mathbf{p}_r) - \mathbf{v}_{max,ph}(\mathbf{p}_r)\|_2 \tag{4.32}$$

quantify errors of the velocity field, the root mean squared (RMS) error

$$e_{rms} = \sqrt{\frac{1}{N} \sum_{r=1}^{N} [\varphi_c(t_r) - \varphi_{ph}(t_r)]^2} \tag{4.33}$$

measures how well the flow profile was recovered. A reconstruction of the velocity field was considered successful if $s_{cos} \geq 0.85$ and $e_d \leq 0.15$ mm.

4.4.2 Real Case Acquisition and Preprocessing Protocol

The method was evaluated on the carotid artery of ten human subjects aged 25 to 31. For the first five of these, a manual screening protocol was implemented. Once the capabilities of the proposed method to successfully reconstruct blood velocity fields and flow profiles in real human data, the remaining five volunteers underwent sets of robotic US acquisitions for improved imaging quality. After preprocessing, the full reconstruction pipeline as described in Sec. 4.3 was executed for all subjects. This section details the respective experimental system setups, including the necessary preprocessing steps to establish an input dataset for the reconstruction algorithm in the form of 4-tuples $(\mathbf{p}, \mathbf{d}, m, t)$.

Freehand Screening Protocol

An open access ultrasound system Aurotech® MANUS (Aurotech Ultrasound AS, Tydal, Norway) was used with a linear array probe (128 elements, single element width 0.27 mm, focal depth 30 mm) operating at 8 MHz. This system facilitates the duplex acquisition of B-mode images and raw Doppler in-phase and quadrature samples. For 3D positioning, a Ascension® TrakStar electromagnetic (EM) tracking system (Ascension Technology Corp., Shelburne, VT, USA) was used. Timing data with respect to the pulse phase was recorded with a Medlab® P-OX100 dedicated pulse-oximetry sensor (Medlab GmbH, Stutensee, Germany). We performed the calibration of all devices and the Doppler signal processing as described in [98]. The same reconstruction parameters as in Sec. 4.4.1 were used.

Preprocessing of the pulse signal was straightforward by mapping the peak-to-peak distances [247] to a linear signal in the interval $[0, 1]$. In contrast, significant preprocessing was necessary to align the obtained US imaging data spatially. To correct for the limited precision of the EM tracking system, the Doppler magnitudes of each sweep were manually thresholded to obtain a point cloud of the blood vessel. The three point clouds were then registered using the Iterative Closest Point (ICP) algorithm (see Sec. 1.4.1) in a pairwise manner, resulting in small translational corrections.

Each subject was positioned on the edge of an examination bed in supine position, with the head turned toward the center of the bed as much as conveniently possible. Three US sweeps of the carotid artery, mutually around $\alpha = 60°$ apart, were acquired. The volunteers were instructed to breathe normally but to refrain from swallowing or performing any other movements in the neck area.

Robotic Screening Protocol

For robotic acquisitions, the same manipulator and robot control infrastructure as in Sec. 3.4.1 was used. Likewise, the methods to acquire US in 3D follow the principles in Sec. 3.3.3, including the calibration using orthogonal robotic sweeps similar to the approach presented in [237]. However, a different open access US system was used. We acquired both Doppler data (flow estimates after auto-correlation of in-phase and quadrature samples) and B-mode images in duplex mode with a Cephasonics® cQuest Cicada system (Cephasonics, Santa Clara, CA, USA) with a linear CPLA12875 probe (128 elements, focal depth 20 mm) at 7.6 MHz. Using the cQuest Ultrasound SDK, acquired US data were directly imported into ImFusion

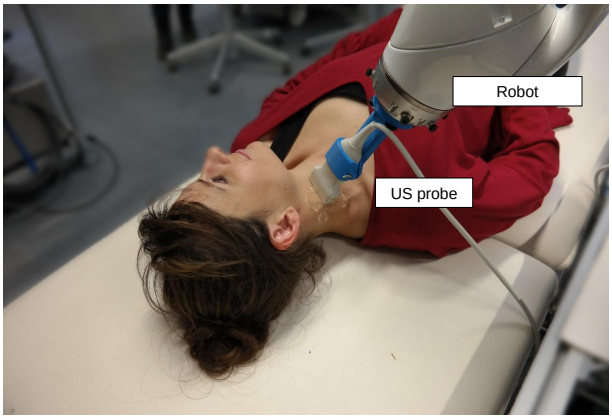

Fig. 4.7. Experimental setup for robotic Doppler acquisitions of the carotid artery. Note that for some subjects, a small pillow was additionally used to stabilize the head.

Suite (see Sec. 3.4.1). Apart from the increased temporal resolution ($\tau = 0.1$), the same reconstruction parameters as in Sec. 4.4.1 were used.

For robotic acquisitions, no pulse-oximetry sensor was available. Instead, a surrogate signal $u(t)$ was defined as the maximum absolute Doppler signal of each frame $F(t)$ in a particular sweep:

$$u(t) = \max |\{m_i\}| \qquad \text{for } m_i \in F(t) \tag{4.34}$$

As above, mapping the peak-to-peak distances of $u(t)$ to the interval $[0, 1]$ yielded normalized time points for each sample as required for the reconstruction algorithm.

While the tracking accuracy of the robotic system was by far superior to the EM tracking used for manual acquisitions, the carotid artery was still slightly pushed by the force exerted onto the tissue. To compensate for the resulting translations in different directions, conventional pair-wise 3D image registrations of the compounded B-mode volumes were performed (see Sec. 1.4 for more details). In this way, small translational corrections similar to the ones in the freehand experiments were obtained.

As illustrated in Fig. 4.7, each subject was positioned on an examination bed as outlined above, including the same instructions in terms of breathing and movements. Depending on the subject, a small pillow was used to stabilize the head. In each case, the robot was manually steered to the proximity of the carotid artery using gravity-compensation mode, before a force controller with a desired force $F^* = 2$ N was enabled, similar to the workflow in Sec. 3.3.5. Following manual exploration of the anatomic area of interest, three US sweeps were automatically acquired, mutually around $\alpha = 45°$ apart. The scanning time was approximately 15 minutes per subject.

4.4.3 Implementation Details

For reconstructions of blood velocities and flow profiles, implementations in MATLAB, R2013b (MathWorks, Inc., Natick, MA, USA) and in C++ were used. For the latter, the Eigen library[1] facilitated numerical computations. Depending on the size of the input dataset and the amount of spatial and temporal B-spline basis functions, reconstructions on our workstation (Intel® Core i7-4770K processor at 3.5 GHz, 32 GB RAM) lasted up to 15 minutes. Incorporating various computational optimizations, reconstruction times could be reduced to 2–5 minutes for confined imaging volumes.

4.5 Results

4.5.1 Model Robustness Evaluation

Figure 4.8 illustrates the reconstruction results for all phantoms with the ideal parameter configuration $\alpha = 90°$, $\eta_s = \eta_v = 0$, and $\rho = 1$.

For the estimation of the flow profile with respect to the measurement angle separation, the correlation r between e_{rms} and the acquisition parameters was investigated. As illustrated

[1] Available online: `http://eigen.tuxfamily.org`

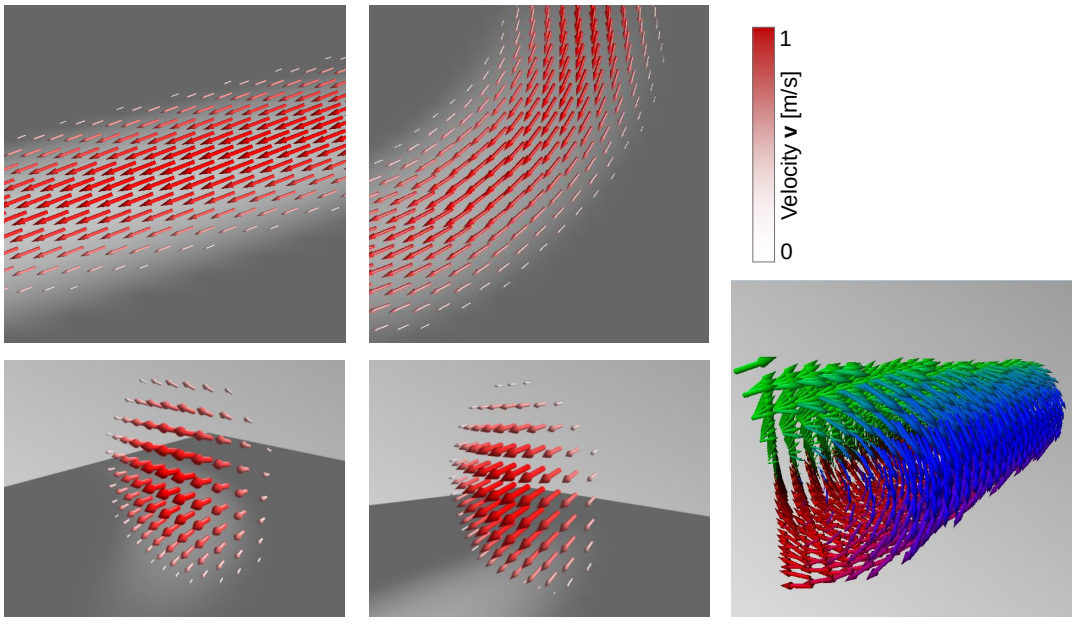

(a) *Linear* (b) *Curved* (c) *Barrel Roll*

Fig. 4.8. Results of phantom experiment reconstructions closely reflect expected results. Vector fields for vessel phantoms (**a**, **b**), as visible in longitudinal (*top*) and cross-sectional slices (*bottom*), follow vessel direction indicated in *white* in the background and exhibit quadratic velocity distributions. Expected rotatory symmetry in the *Barrel Roll* phantom (**c**) can be easily observed using direction-based color-mapping. *Partial reprint from [251], with permission of Springer.*

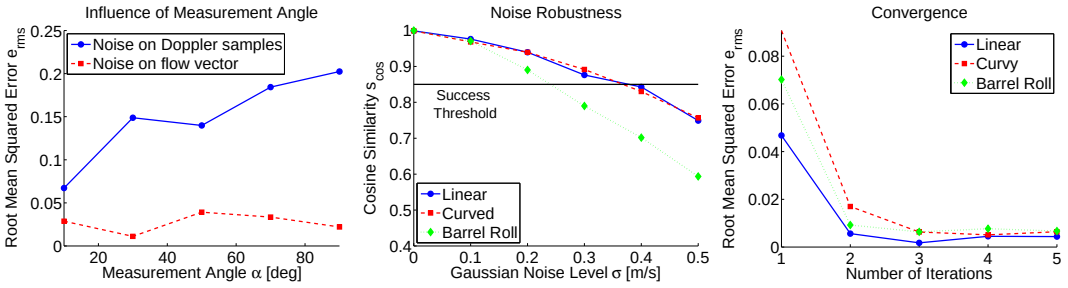

(a) (b) (c)

Fig. 4.9. **a)** The counter-intuitive relation between α and the root mean squared error is only present if noise is directly applied to the Doppler samples, regardless of the direction of acquisition. **b)** Gaussian noise of up to $\sigma \leq 0.35$ m/s (*Linear* and *Curved* phantoms) can be added to the Doppler samples before the cosine similarity s_{cos} drops below the success threshold. **c)** The alternating solver converges quickly within few iterations. *Partial reprint from [251], with permission of Springer.*

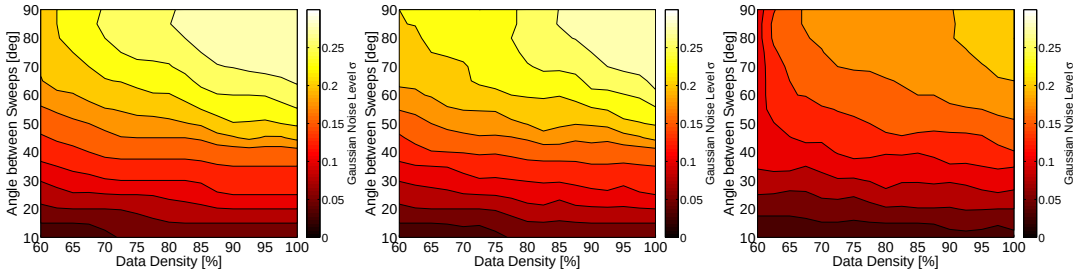

(a) *Linear* (b) *Curved* (c) *Barrel Roll*

Fig. 4.10. Maximum tolerable Gaussian noise level σ such that reconstructions fulfill $s_{cos} > 0.85$ and $e_d < 0.15$. More noise (up to $\sigma = 0.25$ m/s) can be added to samples with higher data density and higher angles between sweeps to still obtain successful reconstructions. *Reprint from [251], with permission of Springer.*

in Fig. 4.9a, increasing opening angles α led to higher errors ($r_{rms,\alpha} = +0.18$) when only noise on the Doppler measurements η_s was applied. An application of noise only onto the underlying velocity vector η_v did not show any significant relationship between angle and error ($r_{rms,\alpha} = -0.03$). Because the former noise model (η_s) facilitates evaluating the limitations of the method, we set $\eta_v = 0$ for all subsequent experiments.

Next, the influence of noise on the velocity field reconstruction was investigated, following the success criteria defined in Sec. 4.4.1. As shown in Fig. 4.9b, Gaussian noise of up to $\sigma \leq 0.35$ m/s could be added to the Doppler samples under ideal measurement conditions of $\alpha = 90°$ and $\rho = 1$, before the cosine similarity s_{cos} dropped below the success threshold. Similar results were obtained for the Euclidean norm e_d. The effect of data density and angle between the measurements for the same noise range was also jointly analyzed. Figure 4.10 illustrates the maximum tolerable noise level σ for various acquisition parameters. Finally, Fig. 4.9c shows the convergence behavior of the alternating solver.

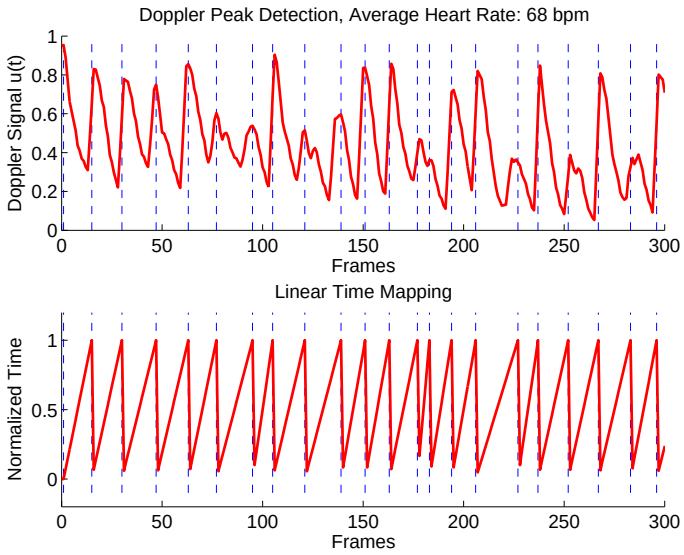

Fig. 4.11. A surrogate signal $u(t)$ based on the frame-wise maximum Doppler amplitude was used if the pulse-oximetry sensor was not available (*top*). For all experiments, regardless whether pulse-oximetry or surrogate signals were used, a linear peak-to-peak mapping yielded normalized times in the interval $[0, 1]$ (*bottom*). In both graphs, *blue lines* indicate detected pulse peaks.

4.5.2 Freehand and Robotic Volunteer Evaluation

For a representative subject, the peak-to-peak linear mapping for obtaining normalized pulse phase times is visualized in Fig. 4.11 (*bottom*). Since no pulse-oximetry sensor was available during the scan of the presented patient, peak detection was performed on the surrogate Doppler signal $u(t)$ as described in Sec. 4.4.2 (*top*), yielding an average heart rate of 68 beats per minute (bpm).

The obtained 3D velocity field and flow profile reconstruction results are presented for all ten subjects. Results for volunteers scanned according to the freehand protocol (*F1-5*) are shown in Fig. 4.14. In these cases, a pulse-oximetry sensor was available. Results for subjects scanned according to the robotic protocol (*R1-5*) are shown in Fig. 4.15. Note that in all cases, compounding of B-mode frames and rendering of a suitable cross-section was only performed for reference and did not influence the blood flow reconstruction.

For two representative freehand volunteers (subjects *F1* and *F2*), the velocity field is visualized in Fig. 4.12 with longitudinal B-mode image slices for reference. Similarly, one representative case (*R5*) of the robotically scanned subjects is shown in Fig. 4.13. Despite the overall qualitatively good reconstructions, errors at the vessel boundaries were regularly evident, regardless of the acquisition technique.

For a preliminary evaluation, all ten subject cases were discussed with an expert vascular surgeon. The obtained velocity magnitudes were considered reasonable for all scanned volunteers, and visualizations of the 3D blood velocity field were deemed overall meaningful.

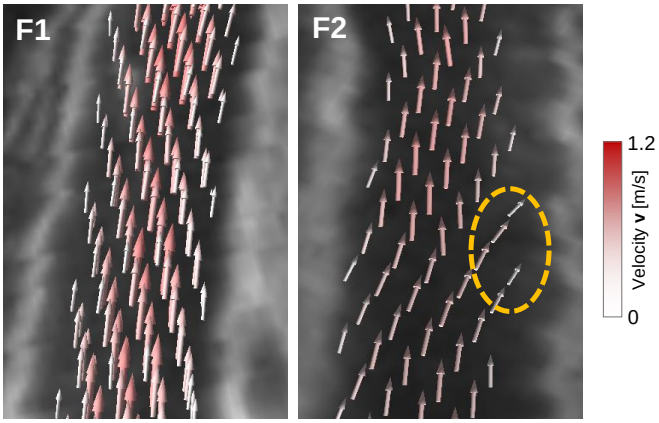

Fig. 4.12. Carotid artery reconstruction results for two representative freehand subjects, illustrating that the blood flow pattern was qualitatively well captured but errors at the vessel borders are evident in some datasets (*yellow ellipse*). Overlay of velocity field on longitudinal B-mode image slice only for reference. *Reprint from [251], modified, with permission of Springer.*

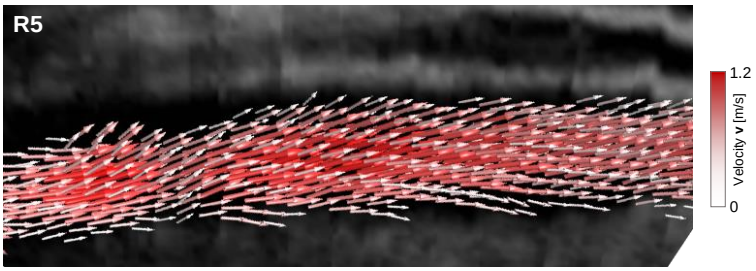

Fig. 4.13. Carotid artery reconstruction results for one representative robotic subject, illustrating that also longer artery sections can be qualitatively well captured. Overlay of velocity field on longitudinal B-mode image slice only for reference.

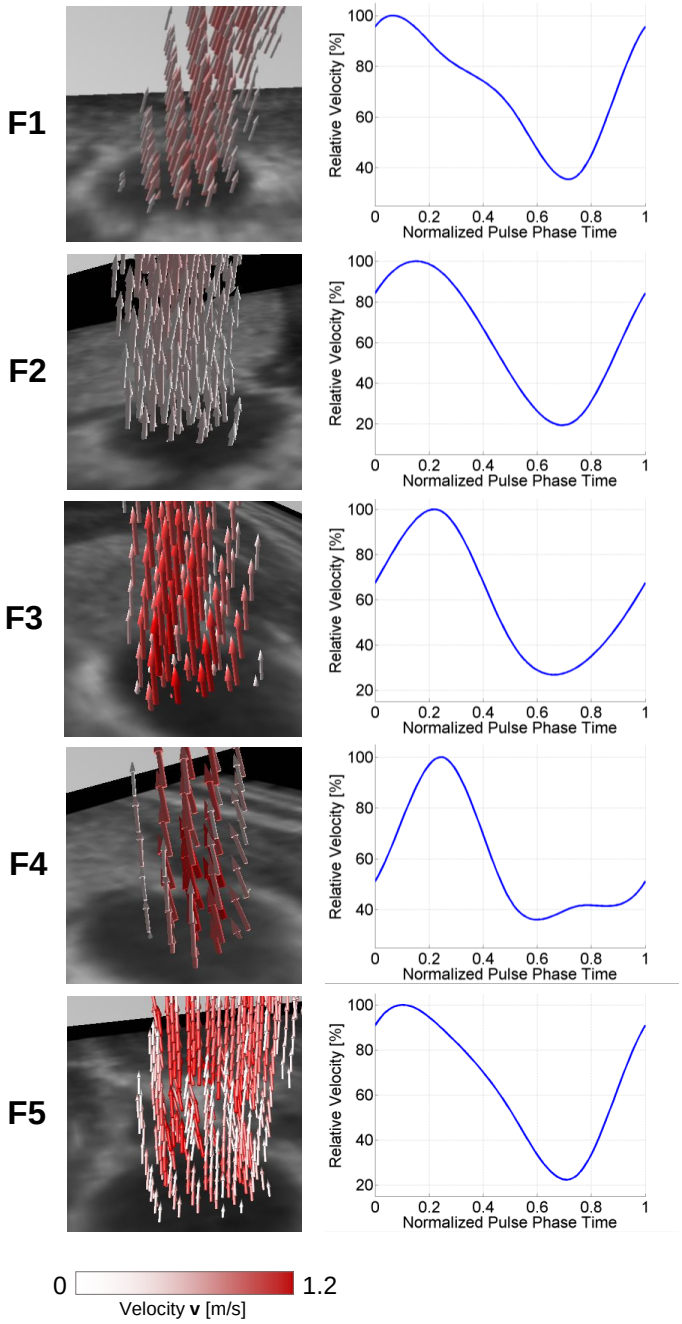

Fig. 4.14. Carotid artery 3D velocity field and flow profile reconstruction results for all five subjects of **freehand experiments**. Overlay of velocity fields on cross-sectional B-mode image slices only for reference. The estimated flow profiles of all five subjects, reconstructed using 4 B-spline basis functions, reveal the pulsatility of the artery. *Partial reprint from [251], with permission of Springer.*

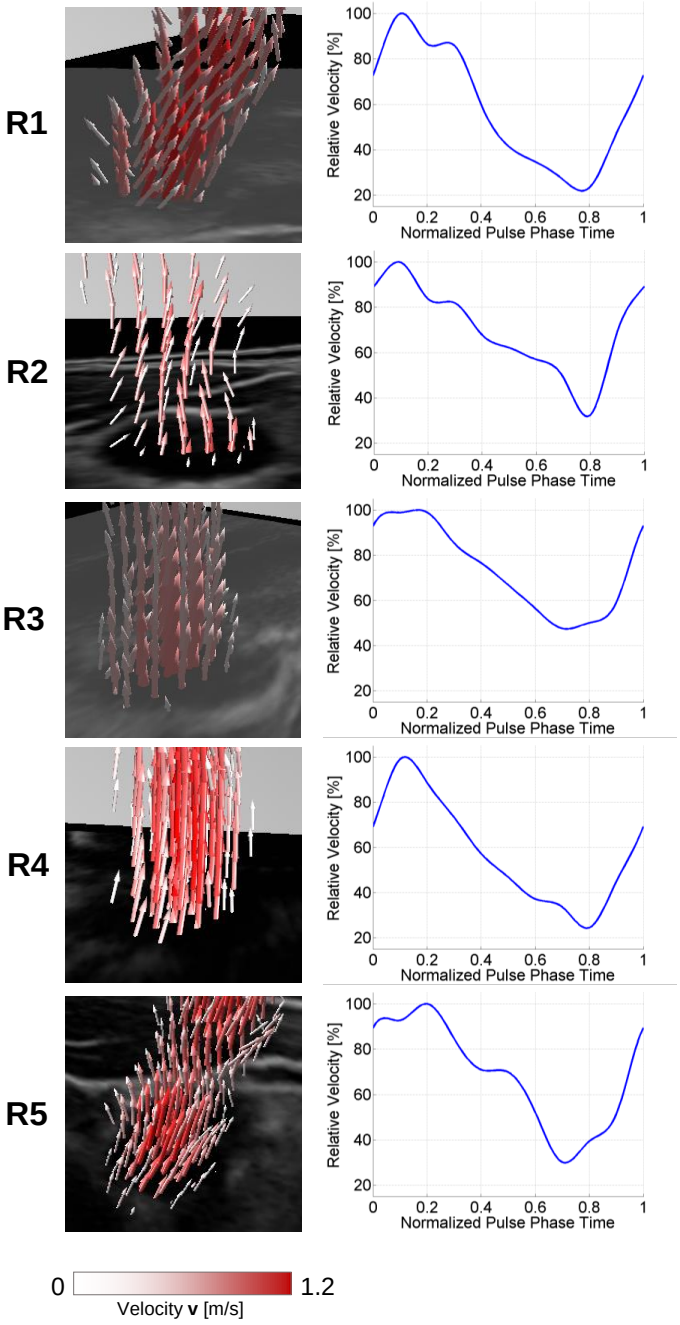

Fig. 4.15. Carotid artery 3D velocity field and flow profile reconstruction results for all five subjects of **robotic experiments**. Overlay of velocity fields on cross-sectional B-mode image slices only for reference. Similar to the freehand experiments, the estimated flow profiles reveal the pulsatility of the artery. Note that more (*10*) B-spline basis functions were used to model the flow profile compared to the freehand experiments shown in Fig. 4.14.

4.6 Discussion

In this chapter, a novel approach to quantify varying blood velocities over time from spatially and temporally arbitrarily sampled Doppler ultrasound measurements was presented. For the first time, to the best of our knowledge, we were able to simultaneously reconstruct a continuous three-dimensional blood velocity field and a periodic temporal flow profile function by coupling a patch-wise B-spline formulation of blood velocity with a waveform model.

The proposed algorithm was first validated using extensive in-silico phantom experiments with two distinct Gaussian noise models. In general, the obtained reconstructions with data sampled from both ideal and noisy environments closely reflected the expected results. Counter-intuitively, noise on the Doppler measurements increased the errors as the three acquisition directions were spread farther apart. This can be explained by the fact that such noise is more likely to cancel out if all samples are roughly obtained from one direction ($\alpha \approx 0°$).

In contrast, tracking inaccuracies and pressure-induced translations of the artery will cause the other investigated noise model (noise on the underlying velocity vectors) to be dominant in a clinical scenario. This is because of the relatively high accuracy of raw Doppler measurements [233], i.e. before correction according to the cosine law, in modern ultrasound systems when optimal acquisition parameters are used. However, experimental results indicated that different acquisition geometries did not influence the reconstruction error when only noise on the velocity vectors was applied. Since the accuracy of clinical measurements has been broadly shown to be very dependent on the insonification angle relative to the vessel [233], this is a unique advantage of reconstruction techniques such as the one presented in this chapter.

The obtained results furthermore suggest a high robustness in terms of the permissible level of noise, lack of acquisition angle separation and data sparsity. The method was most sensitive to noise for the vortex patterns of the *Barrel Roll* phantom and provided best results for tubular vessel structures expected in real scenarios.

Despite the coarse temporal grid spacing, realistic and qualitatively accurate velocity fields and flow profiles could be successfully reconstructed for all five volunteer subjects of the freehand study. Overall, promising agreement between vector fields and vessel anatomies as seen in the B-mode images has been achieved. Similarly, the estimated flow profiles captured the pulsatility of the arteries well. However, in some datasets, reconstruction errors at the vessel borders were evident. Phantom experiments suggested an unsatisfactory registration of the three sweeps and thus missing vector information at the borders to be the cause of this phenomenon.

The original intent of using a robotic manipulator to acquire three sweeps for 3D reconstruction was to mitigate these registration issues and have less distortion due to the pressure induced by the US transducer. Yet, a sufficiently precise tracking system as offered by the robot, and minimal contact force did not make the preprocessing in this regard obsolete. For the five additional volunteer subjects, who were studied using the robot, pair-wise image registration turned out to be essential for acceptable reconstruction results. Similar to the freehand

evaluation, overall promising 3D blood velocity fields and flow profiles could then be obtained. However, apparent reconstruction errors at the vessel boundaries still remained because the non-linear distortions caused by the US probe could not fully be compensated. Nevertheless, early qualitative clinical feedback on the reconstructed velocity magnitude levels and 3D flow visualizations was overall positive.

The increased temporal grid resolution used in the robotic experiments yielded more pronounced flow profile characteristics that were otherwise smoothed out. While the profile inherently becomes more noisy due to the decreased regularization, this could potentially be useful in diagnostic scenarios. A thorough analysis of the temporal regularization as a result of the number of temporal B-spline basis functions might lead to an optimal trade-off between smoothness and expressiveness of the flow profile.

Because all scanned volunteers were healthy, young subjects, no obstructions or otherwise pathological conditions of the carotid artery were detected at all. Further evaluation, potentially with laser Doppler anemometry measurements as ground truth, is necessary to assess the effect of the spatially relaxed laminarity assumption in this regard. Alternatively, a clinical study including a pertinent patient cohort might also shed light on the suitability of the proposed reconstruction methodology in severe stenoses and related conditions.

Apart from a potentially expedient diagnostic value in clinical routine, the presented approach may have further important applications, for instance as initialization of computational fluid dynamics models by the recovered flow profile. Robotic Doppler measurements might also prove useful in more interventional scenarios with blood vessel involvement in the future, in particular for accurate quantification of blood velocities in central parts of the vessel, i.e. when marginal flow direction and magnitudes at the boundaries can be neglected.

Part III

Conclusion

5 Conclusion

In this thesis, several advanced ultrasound imaging techniques for computer assisted interventions were presented, each addressing current challenges encountered in clinical routine. This chapter will summarize the contributions of this work, all of which have been successfully validated in vivo, and provide an outlook to potential future work.

5.1 Multi-Modal Prostate Biopsy Guidance

In chapter 2, a multi-modal image-guided biopsy framework was presented, combining pre-interventional PET and MR images with interventional trans-rectal ultrasound (TRUS). Hereby, the main challenge consists of performing a fast, precise, deformable image registration between MRI and TRUS, allowing a mapping of suspicious lesions identified in advance to the interventional scenario. A successful registration ultimately supports the urologist in performing targeted TRUS-based biopsies by accurately augmenting the live TRUS image with information from additional modalities.

To this end, two novel, fully automatic deformable registration methods were presented. The proposed surface-based registration algorithm requires a segmentation of the prostate in both modalities, which are then elastically aligned using a modified version of the Coherent Point Drift algorithm. Since the accuracy of the method as well as its applicability in the clinical environment is highly dependent on the TRUS segmentations, a fast, automatic segmentation algorithm based on Hough forests was developed. In contrast, the proposed intensity-based registration algorithm does not require segmentations and utilizes a statistical deformation model to reduce the parameter space of the deformable, purely intensity-based image registration. To facilitate correct registration outcomes at crucial locations along expected deformation directions, a novel, lesion-specific, anisotropic preconditioning of the LC^2 similarity metric is presented.

Both methods were extensively validated on phantom and patient data. The latter was only possible because the developed guidance system was eagerly accepted in the urology department of Klinikum Rechts der Isar and rapidly integrated into all prostate biopsy interventions where pre-interventional PET/MR imaging data was available. The obtained results on the performance of both methods overall support their suitability for the intended usage in clinical routine, facilitating guidance with respect to targeted biopsy sampling. While the obtained registration errors were overall within the clinically acceptable accuracy range, measures to manually adjust the automatic registration in case of misalignment ensure the generally safe applicability of the system. As clinical histology evaluations have validated the effectiveness of the presented framework under real conditions, the system has already been used in biopsy sessions of more than 300 patients. After the publication of preliminary clinical results, the

system will continue to support urologists in performing the intervention in the course of ongoing studies.

From a technological point of view, various improvements of both methods could potentially reduce the registration errors in the future. More sophisticated segmentation algorithms based on advanced machine learning techniques [88, 138, 159], which were not the focus of this work, could increase the robustness of the surface-based registration, in turn improving the statistical deformation model. A combination of both surface- and intensity-based approaches, for instance by using the former as initialization for an intensity-based refinement step, still awaits implementation. Clearly, a broader clinical validation of the presented algorithms on a larger collection of patient datasets will be helpful in assessing to which extend successful mapping of suspicious regions from functional imaging data to the interventional TRUS image is possible, providing a lower bound on the size of the suspicious lesion for reliable targeting. Beyond that, biopsy guidance systems as the one presented in this thesis play an important role in the assessment of and guideline development for novel PET tracers such as ^{68}Ga-PSMA, whose positive influence on cancer detection, localization or exclusion still remains to be confirmed [135].

5.2 Visual Servoing-based Robotic Ultrasound

Chapter 3 concerns novel techniques toward robotically assisted interventional ultrasound. Image-based guidance applications utilizing robotic manipulators can help to overcome some of the challenges of manual acquisitions, in particular navigation to and maintenance of an appropriate location and a suitable acoustic window. Not only adequate robot control laws for US scans with sufficient quality but also certain levels of automation are necessary within intraoperative guidance applications to facilitate navigation and execution of pre-interventional plans by the surgeon.

Based on innovative visual servoing schemata, i.e. control architectures where features derived from image information directly serve as input variables, a comprehensive, autonomous robotic US framework was proposed. It facilitated the development of application-specific control laws for two clinical scenarios with a high potential benefit of robotic assistance with ultrasound. First, a continuous re-registration of live 3D ultrasound images with an interventional plan allowed for an accurate guidance of manual lumbar spine needle insertions. Hereby, target anatomy movements were quickly compensated by re-adjusting the US transducer position. Second, a fully autonomous acquisition of high-quality, abdominal 3D US images was achieved to facilitate aortic aneurysm screening in clinical routine. Using US confidence maps, automatic optimizations of the image quality with respect to the predefined target anatomy ensured the suitability of acquired 3D US volumes for diagnosis.

In both cases, the presented methods were not only validated on artificial phantoms but also in vivo on human volunteers, confirming their robustness in real anatomy. The robotic system was able to successfully track human vertebrae under realistic conditions, converging within clinically acceptable time spans in a smooth, oscillation-free manner. The conducted experiments have also indicated that the proposed adaptive controllers are also able to

successfully adapt to the physique of the scanned volunteers and increasingly avoid bowel gas in the acquisition of abdominal sweeps.

The presented visual servoing methodology is generic and can possibly be translated to other clinical applications. Potential future improvements include the increased exploitation of US systems natively acquiring 3D volumes with high frame rates [171]. When coupled with more sophisticated strategies to mitigate the bottleneck of 3D image registration and confidence map computation, anatomy tracking and image quality optimization may reach near real-time performance. In turn, this might allow for robotic guidance applications in currently inaccessible clinical scenarios due to the required capabilities of motion compensation.

In the long term, ultrasound-based approaches could lead to an increased effectiveness of surgical navigation with reduced exposure to radiation compared to today's practice. Due to the reliable, clinically sufficient [72] needle insertion accuracy achieved on phantoms in a real neurosurgical operating theater, an initial clinical trial to demonstrate the efficacy of facet joint injections using robotic US guidance for chronic spinal pain management is already in preparation, with currently pending approval of the institutional review board.

5.3 Toward 3D Doppler Vascular Imaging

In chapter 4, an innovative method for the reconstruction of 3D blood velocity fields and flow profile functions from Doppler ultrasound samples was presented. The approach addresses two common challenges associated with sonographic quantification of blood flow, namely the projective nature of Doppler velocity measurements, and the inconsistency of 3D acquisitions over time due to the pulsation of the blood flow within a vessel. As a first step toward application of the proposed method in clinical routine, its suitability for vascular screening is demonstrated, paving the way for more interventional scenarios with blood vessel involvement in the future.

The presented method uses B-spline basis functions to model both a spatial, 3D velocity field describing the peak (systolic) flow patterns within a vessel, and a temporal flow profile function incorporating the patient-specific pulsatility over the heart cycle. By assuming that the flow direction at any point in the reconstructed volume is constant over time, and only its magnitude is changing due to pulsation, an efficient iterative reconstruction scheme is able to cope with arbitrarily sampled Doppler data in space and time.

In numerous in-silico phantom experiments, the method was validated and its robustness with respect to different noise models investigated. Especially for tubular vessel structures as commonly expected in clinical scenarios, a high robustness in terms of the permissible level of noise, lack of acquisition angle separation and data sparsity was determined. Using both freehand and robotic scanning protocols, in vivo validation on human volunteers was conducted. Overall, promising agreement between blood velocity vector fields and vessel anatomies as seen in the also acquired and independently compounded B-mode images has been achieved. However, reconstruction errors at the vessel borders were evident in some cases as a result of an insufficient distribution of input samples due to an imperfect registration between acquisition directions.

Possible future extensions of the presented approach include a detailed parameter evaluation, a thorough quantitative validation, a more local definition of the flow profile for handling bifurcations, and an improvement of ultrasound sweep registration toward higher velocity field resolution. In addition, further evaluation is necessary to assess the effect of the spatially relaxed laminarity assumption. A more elaborate journal publication, including repeated experiments with ground truth flow measurements from experienced physicians, is currently in preparation.

Part IV

Appendix

List of Publications and Patents

A

A.1 List of Authored and Co-authored Publications

2017

[249] **Oliver Zettinig***, Benjamin Frisch*, Salvatore Virga, Marco Esposito, Anna Rienmüller, Bernhard Meyer, Christoph Hennersperger, Yu-Mi Ryang*, and Nassir Navab*. "3D ultrasound registration-based visual servoing for neurosurgical navigation". In: *International Journal of Computer Assisted Radiology and Surgery* 12.9 (Sept. 2017), pp. 1607–1619. DOI: 10.1007/s11548-017-1536-2

[255] **Oliver Zettinig**, Julia Rackerseder, Beatrice Lentes, Tobias Maurer, Kay Westenfelder, Matthias Eiber, Benjamin Frisch, and Nassir Navab. "Preconditioned intensity-based prostate registration using statistical deformation models". In: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*. IEEE, Apr. 2017, pp. 853–857. DOI: 10.1109/ISBI.2017.7950651

[117] Risto Kojcev, Ashkan Khakzar, Bernhard Fuerst, **Oliver Zettinig**, Carole Fahkry, Robert DeJong, Jeremy Richmon, Russell Taylor, Edoardo Sinibaldi, and Nassir Navab. "On the reproducibility of expert-operated and robotic ultrasound acquisitions". In: *International Journal of Computer Assisted Radiology and Surgery* 12.6 (June 2017), pp. 1003–1011. DOI: 10.1007/s11548-017-1561-1

[97] Christoph Hennersperger*, Bernhard Fuerst*, Salvatore Virga*, **Oliver Zettinig**, Benjamin Frisch, Thomas Neff, and Nassir Navab. "Towards MRI-Based Autonomous Robotic US Acquisitions: A First Feasibility Study". In: *IEEE Transactions on Medical Imaging* 36.2 (2017), pp. 538–548. DOI: 10.1109/TMI.2016.2620723

2016

[250] **Oliver Zettinig**, Bernhard Fuerst, Risto Kojcev, Marco Esposito, Mehrdad Salehi, Wolfgang Wein, Julia Rackerseder, Edoardo Sinibaldi, Benjamin Frisch, and Nassir Navab. "Toward real-time 3D ultrasound registration-based visual servoing for interventional navigation". In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2016, pp. 945–950. DOI: 10.1109/ICRA.2016.7487226

[230] Salvatore Virga*, **Oliver Zettinig***, Marco Esposito, Karin Pfister, Benjamin Frisch, Thomas Neff, Nassir Navab, and Christoph Hennersperger. "Automatic force-compliant robotic ultrasound screening of abdominal aortic aneurysms". In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Oct. 2016, pp. 508–513. DOI: 10.1109/IROS.2016.7759101

[116] Risto Kojcev, Bernhard Fuerst, **Oliver Zettinig**, Javad Fotouhi, Sing Chun Lee, Benjamin Frisch, Russell Taylor, Edoardo Sinibaldi, and Nassir Navab. "Dual-robot ultrasound-guided needle placement: closing the planning-imaging-action loop". In: *International Journal of Computer Assisted Radiology and Surgery* 11.6 (June 2016), pp. 1173–1181. DOI: 10.1007/s11548-016-1408-1

[68] Benjamin Frisch, **Oliver Zettinig**, Bernhard Fuerst, Salvatore Virga, Christoph Hennersperger, and Nassir Navab. "Collaborative Robotic Ultrasound: Towards Clinical Application". In: *Radiological Society of North America (RSNA) Annual Meeting*. 2016, SSM12–04. URL: http://archive.rsna.org/2016/16005047.html

2015

[256] **Oliver Zettinig**, Amit Shah, Christoph Hennersperger, Matthias Eiber, Christine Kroll, Hubert Kübler, Tobias Maurer, Fausto Milletarì, Julia Rackerseder, Christian Schulte zu Berge, Enno Storz, Benjamin Frisch, and Nassir Navab. "Multimodal image-guided prostate fusion biopsy based on automatic deformable registration". In: *International Journal of Computer Assisted Radiology and Surgery* 10.12 (Dec. 2015), pp. 1997–2007. DOI: 10.1007/s11548-015-1233-y

[212] Enno Storz, Amit Shah, **Oliver Zettinig**, Matthias Eiber, Hans-Jürgen Wester, Hubert Kübler, Jürgen E Gschwend, Markus Schwaiger, Benjamin Frisch, and Tobias Maurer. "PSMA-PET/MRI-guided transrectal fusion biopsy for the detection of prostate cancer". In: *European Urology Supplements* 14.2 (Apr. 2015), e217. DOI: 10.1016/S1569-9056(15)60217-2

[67] Benjamin Frisch, Enno Storz, **Oliver Zettinig**, Amit Shah, Hubert Kübler, Nassir Navab, Hans-Jürgen Wester, Markus Schwaiger, Matthias Eiber, and Tobias Maurer. "PET/MRI/TRUS image fusion guided prostate biopsy: Development of a research platform and initial clinical results". In: *Journal of Nuclear Medicine* 56 (2015). URL: http://jnm.snmjournals.org/content/56/supplement_3/510

[195] Amit Shah, **Oliver Zettinig**, Enno Storz, Tobias Maurer, Matthias Eiber, Nassir Navab, and Benjamin Frisch. "Challenges in Multimodal Image-guided Targeted Prostate Biopsy". In: *Hamlyn Symposium on Medical Robotics, London, UK*. 2015, pp. 21–22. URL: http://hamlyn.doc.ic.ac.uk/hsmr/sites/default/files/HSMR-2015-Proceedings-FINAL.pdf

2014

[254] **Oliver Zettinig**, Tommaso Mansi, Dominik Neumann, Bogdan Georgescu, Saikiran Rapaka, Philipp Seegerer, Elham Kayvanpour, Farbod Sedaghat-Hamedani, Ali Amr, Jan Haas, Henning Steen, Hugo Katus, Benjamin Meder, Nassir Navab, Ali Kamen, and Dorin Comaniciu. "Data-driven estimation of cardiac electrical diffusivity from 12-lead ECG signals". In: *Medical Image Analysis* 18.8 (Dec. 2014), pp. 1361–1376. DOI: 10.1016/j.media.2014.04.011

[251] **Oliver Zettinig***, Christoph Hennersperger*, Christian Schulte zu Berge, Maximilian Baust, and Nassir Navab. "3D Velocity Field and Flow Profile Reconstruction from Arbitrarily Sampled Doppler Ultrasound Data". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Vol. 8674 LNCS. PART 2. Springer Verlag, 2014, pp. 611–618. DOI: 10.1007/978-3-319-10470-6_76

[194] Amit Shah, **Oliver Zettinig**, Tobias Maurer, Cristina Precup, Christian Schulte zu Berge, Jakob Weiss, Benjamin Frisch, and Nassir Navab. "An Open Source Multimodal Image-Guided Prostate Biopsy Framework". In: *Clinical Image-Based Procedures. Translational Research in Medical Imaging*. Vol. 8680. Springer LNCS, 2014, pp. 1–8. DOI: 10.1007/978-3-319-13909-8_1

[131] Tommaso Mansi, Roy Beinart, **Oliver Zettinig**, Saikiran Rapaka, Bogdan Georgescu, Ali Kamen, Yoav Dori, M. Muz Zviman, Daniel A. Herzka, Henry R. Halperin, and Dorin Comaniciu. "A Framework for the Pre-clinical Validation of LBM-EP for the Planning and Guidance of Ventricular Tachycardia Ablation". In: *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges*. Vol. 8330 LNCS. Springer Verlag, 2014, pp. 253–261. DOI: 10.1007/978-3-642-54268-8_30

2013

[252] **Oliver Zettinig**, Tommaso Mansi, Bogdan Georgescu, Elham Kayvanpour, Farbod Sedaghat-Hamedani, Ali Amr, Jan Haas, Henning Steen, Benjamin Meder, Hugo Katus, Nassir Navab, Ali Kamen, and Dorin Comaniciu. "Fast Data-Driven Calibration of a Cardiac Electrophysiology Model from Images and ECG". in: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*. Vol. 8149. Springer LNCS, 2013, pp. 1–8. DOI: 10.1007/978-3-642-40811-3_1
MICCAI Young Scientist Award 2013

[253] **Oliver Zettinig**, Tommaso Mansi, Bogdan Georgescu, Saikiran Rapaka, Ali Kamen, Jan Haas, Karen S Frese, Farbod Sedaghat-Hamedani, Elham Kayvanpour, Ali Amr, Stefan Hardt, Derliz Mereles, Henning Steen, Andreas Keller, Hugo A Katus, Benjamin Meder, Nassir Navab, and Dorin Comaniciu. "From Medical Images to Fast Computational Models of Heart Electromechanics: An Integrated Framework towards Clinical Use". In: *Functional Imaging and Modeling of the Heart*. Vol. 7945. Springer LNCS, 2013, pp. 249–258. DOI: 10.1007/978-3-642-38899-6_30

[76] Bogdan Georgescu, Saikiran Rapaka, Tommaso Mansi, **Oliver Zettinig**, Ali Kamen, and Dorin Comaniciu. "Towards Real-Time Cardiac Electrophysiology Computations Using GP-GPU Lattice-Boltzmann Method". In: *MICCAI Workshop on High Performance Computing for Biomedical Image Analysis–HPC-MICCAI*. 2013

*Authors marked with * contributed equally to the respective work.*

A.2 List of Patent Applications

[P1] Tommaso Mansi, **Oliver Zettinig**, Bogdan Georgescu, Ali Kamen, Dorin Comaniciu, Saikiran Rapaka, "Method and System for Interactive Computation of Cardiac Electromechanics," June 2015, US Patent App. 14/655,083.

B

Abstracts of Publications not Discussed in this Thesis

Data-driven Estimation of Cardiac Electrical Diffusivity from 12-lead ECG Signals

O. Zettinig, T. Mansi, D. Neumann, B. Georgescu, S. Rapaka, P. Seegerer, E. Kayvanpour, F. Sedaghat-Hamedani, A. Amr, J. Haas, H. Steen, H. Katus, B. Meder, N. Navab, A. Kamen, D. Comaniciu

Diagnosis and treatment of dilated cardiomyopathy (DCM) is challenging due to a large variety of causes and disease stages. Computational models of cardiac electrophysiology (EP) can be used to improve the assessment and prognosis of DCM, plan therapies and predict their outcome, but require personalization. In this work, we present a data-driven approach to estimate the electrical diffusivity parameter of an EP model from standard 12-lead electrocardiograms (ECG). An efficient forward model based on a mono-domain, phenomenological Lattice-Boltzmann model of cardiac EP, and a boundary element-based mapping of potentials to the body surface is employed. The electrical diffusivity of myocardium, left ventricle and right ventricle endocardium is then estimated using polynomial regression which takes as input the QRS duration and electrical axis. After validating the forward model, we computed 9,500 EP simulations on 19 different DCM patients in just under three seconds each to learn the regression model. Using this database, we quantify the intrinsic uncertainty of electrical diffusion for given ECG features and show in a leave-one-patient-out cross-validation that the regression method is able to predict myocardium diffusion within the uncertainty range. Finally, our approach is tested on the 19 cases using their clinical ECG. 84% of them could be personalized using our method, yielding mean prediction errors of 18.7 ms for the QRS duration and 6.5° for the electrical axis, both values being within clinical acceptability. By providing an estimate of diffusion parameters from readily available clinical data, our data-driven approach could therefore constitute a first calibration step toward a more complete personalization of cardiac EP.

Medical Image Analysis, vol. 18, iss. 8, pp. 1361-1376, 2014. [254]

A Framework for the Pre-clinical Validation of LBM-EP for the Planning and Guidance of Ventricular Tachycardia Ablation

T. Mansi, R. Beinart, O. Zettinig, S. Rapaka, B. Georgescu, A. Kamen, Y. Dori, M. M. Zviman, D. A. Herzka, H. R. Halperin, D. Comaniciu

This manuscript presents a framework for the pre-clinical validation of LBM-EP, a fast cardiac electrophysiology model based on the lattice-Boltzmann method (LBM). The overarching goal is to assess whether the model is able to predict ventricular tachycardia (VT) induction given lead location and stimulation protocol. First, the random-walk algorithm is used to interactively segment the heart ventricles from delayed-enhancement magnetic resonance images (DE-MRI). Scar and border zone are visually delineated using image thresholding. Then, a detailed anatomical model is generated, comprising fiber architecture and spatial distribution of action potential duration. That information is rasterized to a Cartesian grid, and the cardiac potentials are computed. The framework is illustrated on one swine data, for which two different pacing protocols at four different sites were tested. Each of the protocols were then virtually tested by computing seven seconds of heart beat. Model predictions in terms of VT induction were compared with what was observed in the animal. Our parallel implementation on graphics processing units required a total computation time of about two minutes at an isotropic grid resolution of 0.8 mm (21s at a resolution of 1.5 mm), thus enabling interactive VT testing.

Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges, Springer LNCS, 2014, vol. 8330, pp. 253-261. [131]

From Medical Images to Fast Computational Models of Heart Electromechanics: an Integrated Framework towards Clinical Use

O. Zettinig, T. Mansi, B. Georgescu, S. Rapaka, A. Kamen, J. Haas, K. S. Frese, F. Sedaghat-Hamedani, E. Kayvanpour, A. Amr, S. Hardt, D. Mereles, H. Steen, A. Keller, H. A. Katus, B. Meder, N. Navab, and D. Comaniciu

With the recent advancements in computational power, realistic modeling of heart function within a clinical environment has come into reach. Yet, current modeling frameworks either lack overall completeness or computational performance, and their integration with clinical imaging and data is still tedious. In this paper, we propose an integrated framework to model heart anatomy, electrophysiology, biomechanics and hemodynamics from clinical and imaging data, which is fast enough to be embedded in clinical setting. More precisely, we introduce an efficient GPU implementation of the orthotropic Holzapfel-Ogden model of myocardium tissue and couple it with a near real-time cardiac electrophysiology model, efficient lumped models of cardiac hemodynamics and data-driven techniques for cardiac anatomy estimation. Benchmark experiments conducted on patient data showed that the computation of a whole-heart cycle with a mesh resolution of 64k elements is possible in about one minute on a standard desktop machine (Intel Xeon 2.4GHz, NVIDIA GeForce GTX 580). Our experiments demonstrate that our approach is able to compute electrophysiology up to $21\times$ faster and biomechanics up to $22\times$ faster than prior CPU-based approaches. This speed-up breaks ground towards clinical therapy planning.

Functional Imaging and Modeling of the Heart, Springer LNCS, 2013, vol. 7945, pp. 249-258. [253]

Fast Data-driven Calibration of a Cardiac Electrophysiology Model from Images and ECG

O. Zettinig, T. Mansi, B. Georgescu, S. Rapaka, A. Kamen, J. Haas, K. S. Frese, F. Sedaghat-Hamedani, E. Kayvanpour, A. Amr, S. Hardt, D. Mereles, H. Steen, A. Keller, H. A. Katus, B. Meder, N. Navab, and D. Comaniciu

Recent advances in computational electrophysiology (EP) models make them attractive for clinical use. We propose a novel data-driven approach to calibrate an EP model from standard 12-lead electrocardiograms (ECG), which are in contrast to invasive or dense body surface measurements widely available in clinical routine. With focus on cardiac depolarization, we first propose an efficient forward model of ECG by coupling a mono-domain, Lattice-Boltzmann model of cardiac EP to a boundary element formulation of body surface potentials. We then estimate a polynomial regression to predict myocardium, left ventricle and right ventricle endocardium electrical diffusion from QRS duration and ECG electrical axis. Training was performed on 4,200 ECG simulations, calculated in $\approx 3\,s$ each, using different diffusion parameters on 13 patient geometries. This allowed quantifying diffusion uncertainty for given ECG parameters due to the ill-posed nature of the ECG problem. We show that our method is able to predict myocardium diffusion within the uncertainty range, yielding a prediction error of less than $5\,ms$ for QRS duration and $2°$ for electrical axis. Prediction results compared favorably with those obtained with a standard optimization procedure, while being 60 times faster. Our data-driven model can thus constitute an efficient preliminary step prior to more refined EP personalization.

Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013, Springer LNCS, 2013, vol. 8149, pp. 1-8. [252]

Towards Real-Time Cardiac Electrophysiology Computations Using GP-GPU Lattice-Boltzmann Method

B. Georgescu, S. Rapaka, T. Mansi, O. Zettinig, A. Kamen, and D. Comaniciu

With recent advances in numerical methods and experimental validation, cardiac electrophysiology models can become surrogate tools for improved diagnostics and therapy planning. However, day-to-day clinical applications require models that are accurate and detailed enough to capture the main pathological patterns, but at the same time fast, with near real-time computation time. In particular, the models should be computed in a reasonable amount of time to enable personalization and on-line therapy guidance. Towards this goal, we present in this manuscript a novel algorithm adapted to graphics processing units (GPU) that enables near real-time cardiac electrophysiology computation with state-of-the-art cellular models. Our method relies on LBM-EP, a Lattice-Boltzmann method, which is naturally scalable to massively parallel architectures. Tested on a synthetic case and on a patient geometry, our experiments demonstrate the high scalability of the algorithm, reaching $10\times$ speed up with respect to the CPU implementation of the algorithm.

MICCAI Workshop on High Performance Computing for Biomedical Image Analysis–HPC-MICCAI, 2013. [76]

Bibliography

[1] F. Aalamifar, R. Khurana, A. Cheng, X. Guo, I. Iordachita, and E. M. Boctor. "Enabling technologies for robot assisted ultrasound tomography". In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 13.1 (Mar. 2017), e1746. DOI: 10.1002/rcs.1746 (cit. on p. 73).

[2] F. Aalamifar, R. Seifabadi, M. Bernardo, A. H. Negussie, B. Turkbey, M. Merino, P. Pinto, A. Rahmim, B. J. Wood, and E. M. Boctor. "Ultrasound Tomosynthesis: A New Paradigm for Quantitative Imaging of the Prostate". In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*. 2016, pp. 577–584. DOI: 10.1007/978-3-319-46720-7_67 (cit. on p. 73).

[3] M. Abayazid, G. J. Vrooijink, S. Patil, R. Alterovitz, and S. Misra. "Experimental evaluation of ultrasound-guided 3D needle steering in biological tissue". In: *International Journal of Computer Assisted Radiology and Surgery* 9.6 (Nov. 2014), pp. 931–939. DOI: 10.1007/s11548-014-0987-y (cit. on p. 74).

[4] P. Abolmaesumi, S. Salcudean, Wen-Hong Zhu, M. Sirouspour, and S. DiMaio. "Image-guided control of a robot for medical ultrasound". In: *IEEE Transactions on Robotics and Automation* 18.1 (2002), pp. 11–23. DOI: 10.1109/70.988970 (cit. on p. 73).

[5] P. Abolmaesumi, S. Salcudean, and W. Zhu. "Visual servoing for robot-assisted diagnostic ultrasound". In: *Proceedings of the 22nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. Vol. 4. IEEE. IEEE, 2000, pp. 2532–2535. DOI: 10.1109/IEMBS.2000.901348 (cit. on p. 74).

[6] B. C. Adams, R. M. Clark, C. Paap, and J. M. Goff. "There is No Benefit to Universal Carotid Artery Duplex Screening before a Major Cardiac Surgical Procedure". In: *Annals of Vascular Surgery* 28.1 (Jan. 2014), pp. 93–101. DOI: 10.1016/j.avsg.2013.06.018 (cit. on p. 102).

[7] T. K. Adebar, A. E. Fletcher, and A. M. Okamura. "3-D Ultrasound-Guided Robotic Needle Steering in Biological Tissue". In: *IEEE Transactions on Biomedical Engineering* 61.12 (Dec. 2014), pp. 2899–2910. DOI: 10.1109/TBME.2014.2334309 (cit. on p. 74).

[8] M. Arigovindan, M. Suhling, C. Jansen, P. Hunziker, and M. Unser. "Full Motion and Flow Field Recovery From Echo Doppler Data". In: *IEEE Transactions on Medical Imaging* 26.1 (Jan. 2007), pp. 31–45. DOI: 10.1109/TMI.2006.884201 (cit. on p. 105).

[9] I. B. Atci, N. Ucler, O. Ayden, S. Albayrak, H. Bitlisli, S. Kilic, and H. B. Altinsoy. "The Comparison of Pain Management Efficiency of Ultrasonography-guided Facet Joint Injection With Fluoroscopy-guided Injection in Lower Lumbar Facet Syndrome". In: *Neurosurgery Quarterly* 26.3 (Aug. 2016), pp. 246–250. DOI: 10.1097/WNQ.0000000000000172 (cit. on p. 66).

[10] S. Azizi, F. Imani, J. T. Kwak, A. Tahmasebi, S. Xu, P. Yan, J. Kruecker, B. Turkbey, P. Choyke, P. Pinto, B. Wood, P. Mousavi, and P. Abolmaesumi. "Classifying Cancer Grades Using Temporal Ultrasound for Transrectal Prostate Biopsy". In: *19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part I*. Springer International Publishing, 2016, pp. 653–661. DOI: 10.1007/978-3-319-46720-7_76 (cit. on p. 36).

[11] L. Beales, S. Wolstenhulme, J. A. Evans, R. West, and D. J. A. Scott. “Reproducibility of ultrasound measurement of the abdominal aorta”. In: *British Journal of Surgery* 98.11 (Nov. 2011), pp. 1517–1525. DOI: 10.1002/bjs.7628 (cit. on p. 67).

[12] Y. Béjot, H. Bailly, J. Durier, and M. Giroud. “Epidemiology of stroke in Europe and trends for the 21st century”. In: *La Presse Médicale* 45.12 (Dec. 2016), e391–e398. DOI: 10.1016/j.lpm.2016.10.003 (cit. on p. 103).

[13] J. Bercoff, G. Montaldo, T. Loupas, D. Savery, F. Mézière, M. Fink, and M. Tanter. “Ultrafast compound doppler imaging: providing full blood flow characterization”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 58.1 (Jan. 2011), pp. 134–147. DOI: 10.1109/TUFFC.2011.1780 (cit. on p. 8).

[14] T. Bernatik, D. Strobel, E. G. Hahn, and D. Becker. “Doppler measurements: a surrogate marker of liver fibrosis?” In: *European Journal of Gastroenterology & Hepatology* 14.4 (Apr. 2002), pp. 383–387. DOI: 10.1097/00042737-200204000-00008 (cit. on p. 103).

[15] A. Bertelsen, J. Melo, E. Sánchez, and D. Borro. “A review of surgical robots for spinal interventions”. In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 9.4 (Dec. 2013), pp. 407–422. DOI: 10.1002/rcs.1469 (cit. on p. 72).

[16] P. Besl and N. D. McKay. “A method for registration of 3-D shapes”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14.2 (Feb. 1992), pp. 239–256. DOI: 10.1109/34.121791 (cit. on p. 24).

[17] Bing Jian and B. C. Vemuri. “Robust Point Set Registration Using Gaussian Mixture Models”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33.8 (Aug. 2011), pp. 1633–1645. DOI: 10.1109/TPAMI.2010.223 (cit. on pp. 44, 50).

[18] M. A. Bjurlin, J. S. Wysock, and S. S. Taneja. “Optimization of Prostate Biopsy: Review of Technique and Complications”. In: *Urologic Clinics of North America* 41.2 (May 2014), pp. 299–313. DOI: 10.1016/j.ucl.2014.01.011 (cit. on pp. 34, 60).

[19] Blausen.com-staff. “Medical gallery of Blausen Medical 2014”. In: *WikiJournal of Medicine* 1.2 (2014). DOI: 10.15347/wjm/2014.010 (cit. on pp. 65, 66, 103).

[20] E. Boctor, A. Viswanathan, M. Choti, R. Taylor, G. Fichtinger, and G. Hager. “A novel closed form solution for ultrasound calibration”. In: *2004 2nd IEEE International Symposium on Biomedical Imaging: Macro to Nano (IEEE Cat No. 04EX821)*. Vol. 2. IEEE, 2004, pp. 527–530. DOI: 10.1109/ISBI.2004.1398591 (cit. on p. 22).

[21] E. Boctor, G. Fischer, M. Choti, G. Fichtinger, and R. Taylor. “A dual-armed robotic system for intraoperative ultrasound guided hepatic ablative therapy: a prospective study”. In: *IEEE International Conference on Robotics and Automation (ICRA), 2004*. Vol. 3. IEEE. IEEE, 2004, pp. 2517–2522. DOI: 10.1109/ROBOT.2004.1307440 (cit. on p. 72).

[22] E. M. Boctor, M. A. Choti, E. C. Burdette, and R. J. Webster III. “Three-dimensional ultrasound-guided robotic needle placement: an experimental evaluation”. In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 4.2 (June 2008), pp. 180–191. DOI: 10.1002/rcs.184 (cit. on pp. 72, 74).

[23] J.-D. Boissonnat and S. Oudot. “Provably good sampling and meshing of surfaces”. In: *Graphical Models* 67.5 (Sept. 2005), pp. 405–451. DOI: 10.1016/j.gmod.2005.01.004 (cit. on p. 43).

[24] J. Boisvert, D. Gobbi, S. Vikal, R. Rohling, G. Fichtinger, and P. Abolmaesumi. “An open-source solution for interactive acquisition, processing and transfer of interventional ultrasound images”. In: *The MIDAS Journal-Systems and Architectures for Computer Assisted Interventions* (2008), p. 70 (cit. on p. 41).

[25] L. Bolondi, S. Li Bassi, S. Gaiani, G. Zironi, G. Benzi, V. Santi, and L. Barbara. “Liver cirrhosis: changes of Doppler waveform of hepatic veins.” In: *Radiology* 178.2 (Feb. 1991), pp. 513–516. DOI: 10.1148/radiology.178.2.1987617 (cit. on p. 103).

[26] M. Brandestini and F. Forster. “Blood flow imaging using a discrete-time frequency meter”. In: *Ultrasonics Symposium, IEEE*. 1978, pp. 348–352. URL: http://faculty.washington.edu/forster/forsterpubs/brandestini78.pdf (cit. on p. 7).

[27] K. Bredahl, N. Eldrup, C. Meyer, J. Eiberg, and H. Sillesen. “Reproducibility of ECG-gated Ultrasound Diameter Assessment of Small Abdominal Aortic Aneurysms”. In: *European Journal of Vascular and Endovascular Surgery* 45.3 (Mar. 2013), pp. 235–240. DOI: 10.1016/j.ejvs.2012.12.010 (cit. on p. 20).

[28] D. C. Brewster, J. L. Cronenwett, J. W. Hallett, K. Johnston, W. C. Krupski, and J. S. Matsumura. “Guidelines for the treatment of abdominal aortic aneurysms: Report of a subcommittee of the Joint Council of the American Association for Vascular Surgery and Society for Vascular Surgery”. In: *Journal of Vascular Surgery* 37.5 (May 2003), pp. 1106–1117. DOI: 10.1067/mva.2003.363 (cit. on p. 67).

[29] M. Brudfors, A. Seitel, A. Rasoulian, A. Lasso, V. A. Lessoway, J. Osborn, A. Maki, R. N. Rohling, and P. Abolmaesumi. “Towards real-time, tracker-less 3D ultrasound guidance for spine anaesthesia”. In: *International Journal of Computer Assisted Radiology and Surgery* 10.6 (June 2015), pp. 855–865. DOI: 10.1007/s11548-015-1206-1 (cit. on p. 70).

[30] R. P. Budde, R. Meijer, P. F. Bakker, C. Borst, and P. F. Gründeman. “Endoscopic localization and assessment of coronary arteries by 13 MHz epicardial ultrasound”. In: *The Annals of Thoracic Surgery* 77.5 (May 2004), pp. 1586–1592. DOI: 10.1016/j.athoracsur.2003.10.051 (cit. on p. 71).

[31] B. A. Casazza. “Diagnosis and treatment of acute low back pain.” In: *American Family Physician* 85.4 (Feb. 2012), pp. 343–50. URL: http://www.ncbi.nlm.nih.gov/pubmed/22335313 (cit. on p. 65).

[32] J. D. Cassidy, L. J. Carroll, and P. Côté. “The Saskatchewan health and back pain survey. The prevalence of low back pain and related disability in Saskatchewan adults.” In: *Spine* 23.17 (Sept. 1998), pp. 1860–1867. URL: http://www.ncbi.nlm.nih.gov/pubmed/9762743 (cit. on p. 65).

[33] P. Center and L. Manchikanti. “A Systematic Review and Best Evidence Synthesis of Effectiveness of Therapeutic Facet Joint Interventions in Managing Chronic Spinal Pain”. In: *Pain Physician* 18 (2015), E535–E582 (cit. on p. 66).

[34] Y.-C. Chang, Y.-S. Huang, C.-S. Huang, J.-H. Chen, and R.-F. Chang. “Intrinsic subtypes and tumor grades in breast cancer are associated with distinct 3-D power Doppler sonographic vascular features”. In: *European Journal of Radiology* 83.8 (Aug. 2014), pp. 1368–1374. DOI: 10.1016/j.ejrad.2014.05.005 (cit. on p. 103).

[35] P. Chatelain. “Quality-driven control of a robotized ultrasound probe”. Dissertation. Technische Universität München, Université de Rennes 1, 2016 (cit. on p. 74).

[36] P. Chatelain, A. Krupa, and N. Navab. “3D ultrasound-guided robotic steering of a flexible needle via visual servoing”. In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. IEEE, May 2015, pp. 2250–2255. DOI: 10.1109/ICRA.2015.7139497 (cit. on p. 74).

[37] P. Chatelain, A. Krupa, and N. Navab. “Confidence-Driven Control of an Ultrasound Probe: Target-Specific Acoustic Window Optimization”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. 2016 (cit. on pp. 74, 77, 87).

[38] P. Chatelain, A. Krupa, and N. Navab. “Optimization of ultrasound image quality via visual servoing”. In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2015, pp. 5997–6002. DOI: 10.1109/ICRA.2015.7140040 (cit. on p. 74).

[39] F. Chaumette and S. Hutchinson. “Visual servo control. I. Basic approaches”. In: *IEEE Robotics & Automation Magazine* 13.4 (Dec. 2006), pp. 82–90. DOI: 10.1109/MRA.2006.250573 (cit. on pp. 68, 74).

[40] F. Chaumette and S. Hutchinson. “Visual servo control. II. Advanced approaches”. In: *IEEE Robotics & Automation Magazine* 14.1 (2007), pp. 109–118. DOI: 10.1109/MRA.2007.339609 (cit. on p. 68).

[41] R. S. C. Cobbold. *Foundations of Biomedical Ultrasound*. Oxford University Press New York, 2007 (cit. on pp. 7, 10, 12, 14).

[42] D. W. Cool, J. Bax, C. Romagnoli, A. D. Ward, L. Gardi, V. Karnik, J. Izawa, J. Chin, and A. Fenster. “Fusion of MRI to 3D TRUS for Mechanically-Assisted Targeted Prostate Biopsy: System Design and Initial Clinical Experience”. In: *Prostate Cancer Imaging. Image Analysis and Image-Guided Interventions*. Springer, 2011, pp. 121–133. DOI: 10.1007/978-3-642-23944-1_13 (cit. on p. 37).

[43] M. Correia, J. Provost, M. Tanter, and M. Pernot. “4D ultrafast ultrasound flow imaging: in vivo quantification of arterial volumetric flow rate in a single heartbeat”. In: *Physics in Medicine and Biology* 61.23 (Dec. 2016), pp. L48–L61. DOI: 10.1088/0031-9155/61/23/L48 (cit. on p. 105).

[44] A. Criminisi and J. Shotton. *Decision forests for computer vision and medical image analysis*. Springer, 2013 (cit. on p. 49).

[45] J. Dang, B. Frisch, P. Lasaygues, D. Zhang, S. Tavernier, N. Felix, P. Lecoq, E. Auffray, J. Varela, S. Mensah, and M. Wan. “Development of an Anthropomorphic Breast Phantom for Combined PET, B-Mode Ultrasound and Elastographic Imaging”. In: *IEEE Transactions on Nuclear Science* 58.3 (June 2011), pp. 660–667. DOI: 10.1109/TNS.2011.2105279 (cit. on p. 50).

[46] G. R. de Oliveira Filho. “The Construction of Learning Curves for Basic Skills in Anesthetic Procedures: An Application for the Cumulative Sum Method”. In: *Anesthesia & Analgesia* 95.2 (Aug. 2002), pp. 411–416. DOI: 10.1213/00000539-200208000-00033 (cit. on p. 66).

[47] P. De Rango. “Aneurysm Diameter Measurement: A Challenging and Frustrating Task”. In: *European Journal of Vascular and Endovascular Surgery* 43.1 (Jan. 2012), p. 34. DOI: 10.1016/j.ejvs.2011.10.004 (cit. on p. 67).

[48] J. De Schutter and H. Van Brussel. “Compliant robot motion II. A control approach based on external control loops”. In: *The International Journal of Robotics Research* 7.4 (1988), pp. 18–33 (cit. on p. 77).

[49] N. B. Delongchamps, M. Peyromaure, A. Schull, F. Beuvon, N. Bouazza, T. Flam, M. Zerbib, N. Muradyan, P. Legman, and F. Cornud. “Prebiopsy Magnetic Resonance Imaging and Prostate Cancer Detection: Comparison of Random and Targeted Biopsies”. In: *The Journal of Urology* 189.2 (Feb. 2013), pp. 493–499. DOI: 10.1016/j.juro.2012.08.195 (cit. on p. 36).

[50] F. A. Duck. *Physical Properties of Tissues: A Comprehensive Reference Book*. Academic press, 2013 (cit. on pp. 11, 14).

[51] B. Dunmire, K. Beach, K.-H. Labs, M. Plett, and D. Strandness. “Cross-beam vector Doppler ultrasound for angle-independent velocity measurements”. In: *Ultrasound in Medicine & Biology* 26.8 (Oct. 2000), pp. 1213–1235. DOI: 10.1016/S0301-5629(00)00287-8 (cit. on p. 104).

[52] N. Duric, P. Littrup, O. Roy, C. Li, S. Schmidt, X. Cheng, and R. Janer. “Clinical breast imaging with ultrasound tomography: A description of the SoftVue system”. In: *The Journal of the Acoustical Society of America* 135.4 (Apr. 2014), pp. 2155–2155. DOI: 10.1121/1.4876990 (cit. on p. 73).

[53] M. Eiber, S. G. Nekolla, T. Maurer, G. Weirich, H.-J. Wester, and M. Schwaiger. “68Ga-PSMA PET/MR with multimodality image analysis for primary prostate cancer”. In: *Abdominal Imaging* 40.6 (Aug. 2015), pp. 1769–1771. DOI: 10.1007/s00261-014-0301-z (cit. on p. 36).

[54] I. K. Ekroll, T. Dahl, H. Torp, and L. Lovstakken. “In vivo vector flow imaging with retrospective pulsed wave doppler”. In: *2012 IEEE International Ultrasonics Symposium*. IEEE, Oct. 2012, pp. 1886–1889. DOI: 10.1109/ULTSYM.2012.0473 (cit. on p. 105).

[55] J. I. Epstein, L. Egevad, M. B. Amin, B. Delahunt, J. R. Srigley, and P. A. Humphrey. “The 2014 International Society of Urological Pathology (ISUP) Consensus Conference on Gleason Grading of Prostatic Carcinoma”. In: *The American Journal of Surgical Pathology* 40.2 (Oct. 2015), p. 1. DOI: 10.1097/PAS.0000000000000530 (cit. on p. 35).

[56] A. P. Erikson and K. Åström. “On the Bijectivity of Thin-Plate Splines”. In: *Analysis for Science, Engineering and Beyond*. Springer, 2012, pp. 93–141. DOI: 10.1007/978-3-642-20236-0_5 (cit. on p. 45).

[57] D. H. Evans, J. A. Jensen, and M. B. Nielsen. “Ultrasonic colour Doppler imaging”. In: *Interface Focus* 1.4 (Aug. 2011), pp. 490–502. DOI: 10.1098/rsfs.2011.0017 (cit. on p. 102).

[58] D. H. Evans and W. N. McDicken. *Doppler Ultrasound: Physics, Instrumentation and Signal Processing*. 2nd ed. John Wiley & Sons Ltd, 1999, p. 456 (cit. on p. 18).

[59] K. Evans, S. Roll, and J. Baker. “Work-Related Musculoskeletal Disorders (WRMSD) Among Registered Diagnostic Medical Sonographers and Vascular Technologists: A Representative Sample”. In: *Journal of Diagnostic Medical Sonography* 25.6 (Nov. 2009), pp. 287–299. DOI: 10.1177/8756479309351748 (cit. on p. 66).

[60] A. Fedorov, S. Khallaghi, C. A. Sánchez, A. Lasso, S. Fels, K. Tuncali, E. N. Sugar, T. Kapur, C. Zhang, W. Wells, P. L. Nguyen, P. Abolmaesumi, and C. Tempany. “Open-source image registration for MRI–TRUS fusion-guided prostate interventions”. In: *International Journal of Computer Assisted Radiology and Surgery* 10.6 (June 2015), pp. 925–934. DOI: 10.1007/s11548-015-1180-7 (cit. on pp. 38, 57, 61).

[61] A. Fenster, D. B. Downey, and H. N. Cardinal. “Three-dimensional ultrasound imaging”. In: *Physics in Medicine and Biology* 46.5 (May 2001), R67–R99. DOI: 10.1088/0031-9155/46/5/201 (cit. on p. 20).

[62] G. Fichtinger, J. P. Fiene, C. W. Kennedy, G. Kronreif, I. Iordachita, D. Y. Song, E. C. Burdette, and P. Kazanzides. “Robotic assistance for ultrasound-guided prostate brachytherapy”. In: *Medical Image Analysis* 12.5 (Oct. 2008), pp. 535–545. DOI: 10.1016/j.media.2008.06.002 (cit. on p. 72).

[63] S. W. Fine and V. E. Reuter. “Anatomy of the prostate revisited: implications for prostate biopsy and zonal origins of prostate cancer”. In: *Histopathology* 60.1 (Jan. 2012), pp. 142–152. DOI: 10.1111/j.1365-2559.2011.04004.x (cit. on p. 35).

[64] M. Fox. “Multiple crossed-beam ultrasound Doppler velocimetry”. In: *IEEE Transactions on Sonics and Ultrasonics* 25.5 (Sept. 1978), pp. 281–286. DOI: 10.1109/T-SU.1978.31028 (cit. on p. 104).

[65] V. Freire, D. Grabs, M. Lepage-Saucier, and T. P. Moser. “Ultrasound-Guided Cervical Facet Joint Injections: A Viable Substitution for Fluoroscopy-Guided Injections?” In: *Journal of Ultrasound in Medicine* 35.6 (June 2016), pp. 1253–1258. DOI: 10.7863/ultra.15.07062 (cit. on pp. 64, 66).

[66] C. Freschi, E. Troia, V. Ferrari, G. Megali, A. Pietrabissa, and F. Mosca. “Ultrasound guided robotic biopsy using augmented reality and human-robot cooperative control”. In: *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, Sept. 2009, pp. 5110–5113. DOI: 10.1109/IEMBS.2009.5332720 (cit. on p. 72).

[67] B. Frisch, E. Storz, O. Zettinig, A. Shah, H. Kübler, N. Navab, H.-J. Wester, M. Schwaiger, M. Eiber, and T. Maurer. “PET/MRI/TRUS image fusion guided prostate biopsy: Development of a research platform and initial clinical results”. In: *Journal of Nuclear Medicine* 56 (2015). URL: http://jnm.snmjournals.org/content/56/supplement_3/510 (cit. on p. 134).

[68] B. Frisch, O. Zettinig, B. Fuerst, S. Virga, C. Hennersperger, and N. Navab. “Collaborative Robotic Ultrasound: Towards Clinical Application”. In: *Radiological Society of North America (RSNA) Annual Meeting*. 2016, SSM12–04. URL: http://archive.rsna.org/2016/16005047.html (cit. on p. 134).

[69] F. A. Fröhlich, G. Passig, A. Vazquez, and G. Hirzinger. “Robot assisted internal mammary artery detection for coronary revascularisation surgery”. In: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. IEEE, Oct. 2010, pp. 1849–1855. DOI: 10.1109/IROS.2010.5653269 (cit. on pp. 74, 77).

[70] B. Fuerst, W. Wein, M. Müller, and N. Navab. “Automatic ultrasound–MRI registration for neurosurgery using the 2D and 3D LC2 Metric”. In: *Medical Image Analysis* 18.8 (Dec. 2014), pp. 1312–1319. DOI: 10.1016/j.media.2014.04.008 (cit. on pp. 26, 27, 29, 33, 39, 46, 81, 98).

[71] J. J. Fütterer, A. Briganti, P. De Visschere, M. Emberton, G. Giannarini, A. Kirkham, S. S. Taneja, H. Thoeny, G. Villeirs, and A. Villers. “Can Clinically Significant Prostate Cancer Be Detected with Multiparametric Magnetic Resonance Imaging? A Systematic Review of the Literature”. In: *European Urology* 68.6 (Dec. 2015), pp. 1045–1053. DOI: 10.1016/j.eururo.2015.01.013 (cit. on p. 35).

[72] K. Galiano, A. A. Obwegeser, G. Bodner, M. Freund, H. Maurer, F. S. Kamelger, R. Schatzer, and F. Ploner. “Ultrasound Guidance for Facet Joint Injections in the Lumbar Spine: A Computed Tomography-Controlled Feasibility Study”. In: *Anesthesia & Analgesia* 101.2 (Aug. 2005), pp. 579–583. DOI: 10.1213/01.ANE.0000158609.64417.93 (cit. on pp. 99, 100, 129).

[73] D. Garcia, J. C. del Álamo, D. Tanné, R. Yotti, C. Cortina, É. Bertrand, J. C. Antoranz, E. Pérez-David, R. Rieu, F. Fernández-Avilés, and J. Bermejo. “Two-Dimensional Intraventricular Flow Mapping by Digital Processing Conventional Color-Doppler Echocardiography Images”. In: *IEEE Transactions on Medical Imaging* 29.10 (Oct. 2010), pp. 1701–1713. DOI: 10.1109/TMI.2010.2049656 (cit. on p. 105).

[74] A. Gee, R. Prager, G. Treece, and L. Berman. “Engineering a freehand 3D ultrasound system”. In: *Pattern Recognition Letters* 24.4-5 (Feb. 2003), pp. 757–777. DOI: 10.1016/S0167-8655(02)00180-0 (cit. on p. 20).

[75] E. A. Geiser, L. G. Christie, D. A. Conetta, C. R. Conti, and G. S. Gossman. “A mechanical arm for spatial registration of two-dimensional echocardiographic sections”. In: *Catheterization and Cardiovascular Diagnosis* 8.1 (1982), pp. 89–101. DOI: 10.1002/ccd.1810080114 (cit. on p. 19).

[76] B. Georgescu, S. Rapaka, T. Mansi, O. Zettinig, A. Kamen, and D. Comaniciu. “Towards Real-Time Cardiac Electrophysiology Computations Using GP-GPU Lattice-Boltzmann Method”. In: *MICCAI Workshop on High Performance Computing for Biomedical Image Analysis–HPC-MICCAI*. 2013 (cit. on pp. 135, 141).

[77] S. Ghose, A. Oliver, R. Martí, X. Lladó, J. C. Vilanova, J. Freixenet, J. Mitra, D. Sidibé, and F. Meriaudeau. “A survey of prostate segmentation methodologies in ultrasound, magnetic resonance and computed tomography images”. In: *Computer Methods and Programs in Biomedicine* 108.1 (Oct. 2012), pp. 262–287. DOI: 10.1016/j.cmpb.2012.04.006 (cit. on p. 38).

[78] S. Ghose, A. Oliver, J. Mitra, R. Martí, X. Lladó, J. Freixenet, D. Sidibé, J. C. Vilanova, J. Comet, and F. Meriaudeau. “A supervised learning framework of statistical shape and probability priors for automatic prostate segmentation in ultrasound images”. In: *Medical Image Analysis* 17.6 (Aug. 2013), pp. 587–600. DOI: 10.1016/j.media.2013.04.001 (cit. on p. 38).

[79] M. Giarrè, B. Dousse, and J.-J. Meister. “Velocity vector reconstruction for color flow Doppler: Experimental evaluation of a new geometrical method”. In: *Ultrasound in Medicine & Biology* 22.1 (Jan. 1996), pp. 75–88. DOI: 10.1016/0301-5629(95)02020-9 (cit. on p. 104).

[80] B. Glocker, N. Komodakis, N. Navab, G. Tziritas, and N. Paragios. “Dense Registration with Deformation Priors”. In: *Information Processing in Medical Imaging*. 2009, pp. 540–551. DOI: 10.1007/978-3-642-02498-6_45 (cit. on p. 38).

[81] D. G. Gobbi and T. M. Peters. “Interactive Intra-operative 3D Ultrasound Reconstruction and Visualization”. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. 2002, pp. 156–163. DOI: 10.1007/3-540-45787-9_20 (cit. on p. 41).

[82] S. Gold, A. Rangarajan, C.-P. Lu, S. Pappu, and E. Mjolsness. “New algorithms for 2D and 3D point matching”. In: *Pattern Recognition* 31.8 (Aug. 1998), pp. 1019–1031. DOI: 10.1016/S0031-3203(98)80010-1 (cit. on p. 24).

[83] A. Gomez, K. Pushparajah, J. M. Simpson, D. Giese, T. Schaeffter, and G. Penney. “A sensitivity analysis on 3D velocity reconstruction from multiple registered echo Doppler views”. In: *Medical Image Analysis* 17.6 (Aug. 2013), pp. 616–631. DOI: 10.1016/j.media.2013.04.002 (cit. on pp. 105, 107, 108, 110).

[84] A. Gomez, A. de Vecchi, M. Jantsch, W. Shi, K. Pushparajah, J. M. Simpson, N. P. Smith, D. Rueckert, T. Schaeffter, and G. P. Penney. “4D Blood Flow Reconstruction Over the Entire Ventricle From Wall Motion and Blood Velocity Derived From Ultrasound Data”. In: *IEEE Transactions on Medical Imaging* 34.11 (Nov. 2015), pp. 2298–2308. DOI: 10.1109/TMI.2015.2428932 (cit. on p. 105).

[85] L. Grady. “Random Walks for Image Segmentation”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28.11 (Nov. 2006), pp. 1768–1783. DOI: 10.1109/TPAMI.2006.233 (cit. on p. 53).

[86] E. G. Grant, C. B. Benson, G. L. Moneta, A. V. Alexandrov, J. D. Baker, E. I. Bluth, B. A. Carroll, M. Eliasziw, J. Gocke, B. S. Hertzberg, S. Katanick, L. Needleman, J. Pellerito, J. F. Polak, K. S. Rholl, D. L. Wooster, and R. E. Zierler. “Carotid artery stenosis: gray-scale and Doppler US diagnosis–Society of Radiologists in Ultrasound Consensus Conference.” In: *Radiology* 229.2 (Nov. 2003), pp. 340–6. DOI: 10.1148/radiol.2292030516 (cit. on p. 102).

[87] C. Graumann, B. Fuerst, C. Hennersperger, F. Bork, and N. Navab. “Robotic ultrasound trajectory planning for volume of interest coverage”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 2016-June. IEEE, May 2016, pp. 736–741. DOI: 10.1109/ICRA.2016.7487201 (cit. on p. 73).

[88] Y. Guo, Y. Gao, and D. Shen. “Deformable MR Prostate Segmentation via Deep Feature Learning and Sparse Patch Matching”. In: *IEEE Transactions on Medical Imaging* 35.4 (Apr. 2016), pp. 1077–1089. DOI: 10.1109/TMI.2015.2508280 (cit. on p. 128).

[89] G. Guthart and J. Salisbury. “The Intuitive (TM) telesurgery system: overview and application”. In: *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation.* Vol. 1. IEEE, 2000, pp. 618–621. DOI: 10.1109/ROBOT.2000.844121 (cit. on p. 71).

[90] R. Hara, A. Nagai, T. Fujii, K. Fukumoto, S. Ohira, Y. Jo, T. Yokoyama, and Y. Miyaji. “Practical application of color Doppler ultrasonography in patients with ejaculatory dysfunction”. In: *International Journal of Urology* 22.6 (June 2015), pp. 609–611. DOI: 10.1111/iju.12754 (cit. on p. 103).

[91] A. Harloff, F. Albrecht, J. Spreer, A. Stalder, J. Bock, A. Frydrychowicz, J. Schöllhorn, A. Hetzel, M. Schumacher, J. Hennig, and M. Markl. “3D blood flow characteristics in the carotid artery bifurcation assessed by flow-sensitive 4D MRI at 3T”. In: *Magnetic Resonance in Medicine* 61.1 (Jan. 2009), pp. 65–74. DOI: 10.1002/mrm.21774 (cit. on p. 104).

[92] M. Hastenteufel. “Neue Methoden des 3D Ultraschalls zur Geschwindigkeitsrekonstruktion und intraoperativen Navigation”. PhD thesis. Universität Fridericiana, Karlsruhe, Germany, 2005 (cit. on p. 112).

[93] I. Hein and W. O’Brien. “Current time-domain methods for assessing tissue motion by analysis from reflected ultrasound echoes-a review”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 40.2 (Mar. 1993), pp. 84–102. DOI: 10.1109/58.212556 (cit. on p. 104).

[94] M. P. Heinrich, M. Jenkinson, M. Bhushan, T. Matin, F. V. Gleeson, S. M. Brady, and J. A. Schnabel. “MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration”. In: *Medical Image Analysis* 16.7 (Oct. 2012), pp. 1423–1435. DOI: 10.1016/j.media.2012.05.008 (cit. on pp. 28, 29).

[95] C. L. Hennersperger. “Domain-Specific Modeling for Vascular Freehand Ultrasound”. Dissertation. Technische Universität München, 2015. URL: http://mediatum.ub.tum.de?id=1241152 (cit. on pp. 9, 10, 15).

[96] C. Hennersperger, M. Baust, D. Mateus, and N. Navab. “Computational Sonography”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Springer International Publishing, 2015, pp. 459–466. DOI: 10.1007/978-3-319-24571-3_55 (cit. on pp. 23, 79).

[97] C. Hennersperger*, B. Fuerst*, S. Virga*, O. Zettinig, B. Frisch, T. Neff, and N. Navab. "Towards MRI-Based Autonomous Robotic US Acquisitions: A First Feasibility Study". In: *IEEE Transactions on Medical Imaging* 36.2 (2017), pp. 538–548. DOI: 10.1109/TMI.2016.2620723 (cit. on pp. 71, 73, 75, 86, 87, 95, 133).

[98] C. Hennersperger, A. Karamalis, and N. Navab. "Vascular 3D+T Freehand Ultrasound Using Correlation of Doppler and Pulse-Oximetry Data". In: *Information Processing in Computer-Assisted Interventions*. Ed. by D. Stoyanov, D. L. Collins, I. Sakuma, P. Abolmaesumi, and P. Jannin. Vol. 8498. Springer International Publishing, 2014, pp. 68–77. DOI: 10.1007/978-3-319-07521-1_8 (cit. on pp. 20, 114).

[99] C. Hennersperger, D. Mateus, M. Baust, and N. Navab. "A Quadratic Energy Minimization Framework for Signal Loss Estimation from Arbitrarily Sampled Ultrasound Data". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Springer, 2014, pp. 373–380. DOI: 10.1007/978-3-319-10470-6_47 (cit. on pp. 79, 80).

[100] M. Heras. "Atherosclerosis: a systemic condition that requires a global approach". In: *European Heart Journal* 21.11 (June 2000), pp. 872–873. DOI: 10.1053/euhj.1999.2071 (cit. on p. 102).

[101] H. Ho, P. Mohan, E. Lim, D. Li, J. Yuen, W. Ng, W. Lau, and C. Cheng. "Robotic ultrasound-guided prostate intervention device: system description and results from phantom studies". In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 5.1 (Mar. 2009), pp. 51–58. DOI: 10.1002/rcs.232 (cit. on p. 72).

[102] P.-W. Hsu, R. W. Prager, A. H. Gee, and G. M. Treece. "Freehand 3D Ultrasound Calibration: A Review". In: *Advanced Imaging in Biology and Medicine*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 47–84. DOI: 10.1007/978-3-540-68993-5_3 (cit. on pp. 21, 22).

[103] Y. Hu, H. U. Ahmed, Z. Taylor, C. Allen, M. Emberton, D. Hawkes, and D. Barratt. "MR to ultrasound registration for image-guided prostate interventions". In: *Medical Image Analysis* 16.3 (Apr. 2012), pp. 687–703. DOI: 10.1016/j.media.2010.11.003 (cit. on pp. 37, 38).

[104] Y. Hu, E. Gibson, H. U. Ahmed, C. M. Moore, M. Emberton, and D. C. Barratt. "Population-based prediction of subject-specific prostate deformation for MR-to-ultrasound image registration". In: *Medical Image Analysis* 26.1 (Dec. 2015), pp. 332–344. DOI: 10.1016/j.media.2015.10.006 (cit. on p. 38).

[105] B. Hussain, E. Y. Wong, and T. L. Poepping. "Numerical design and evaluation of Vector Doppler Ultrasound systems for accurate 3-D velocity measurements". In: *2010 IEEE International Ultrasonics Symposium*. IEEE, Oct. 2010, pp. 1311–1314. DOI: 10.1109/ULTSYM.2010.5936012 (cit. on p. 104).

[106] B. Hussain, B. Y. S. Yiu, A. C. H. Yu, J. C. Lacefield, and T. L. Poepping. "Investigation of Crossbeam Multi-receiver Configurations for Accurate 3-D Vector Doppler Velocity Estimation". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 63.11 (Nov. 2016), pp. 1786–1798. DOI: 10.1109/TUFFC.2016.2597135 (cit. on p. 104).

[107] T. Iwao, A. Toyonaga, K. Oho, C. Tayama, H. Masumoto, T. Sakai, M. Sato, and K. Tanikawa. "Value of Doppler ultrasound parameters of portal vein and hepatic artery in the diagnosis of cirrhosis and portal hypertension." In: *The American journal of gastroenterology* 92.6 (June 1997), pp. 1012–7. URL: http://www.ncbi.nlm.nih.gov/pubmed/9177521 (cit. on p. 102).

[108] B.-H. Jung, B.-H. Kim, and S.-M. Hong. "Respiratory motion prediction with extended Kalman filters based on local circular motion model". In: *International Journal of Bio-Science and Bio-Technology* 5.1 (2013), pp. 51–58 (cit. on p. 99).

[109] R. E. Kalman. "A new approach to linear filtering and prediction problems". In: *Journal of basic Engineering* 82.1 (1960), pp. 35–45 (cit. on p. 99).

[110] H. Kalmus, A. Hedrich, and D. Pardue. "The Acoustic Flowmeter Using Electroric Switching". In: *Ultrasonic Engineering, Transactions of the IRE Professional Group on* 1.1 (1954), pp. 49–63. DOI: 10.1109/TPGUE.1954.1538197 (cit. on p. 7).

[111] I. Kaplan, N. E. Oldenburg, P. Meskell, M. Blake, P. Church, and E. J. Holupka. “Real time MRI-ultrasound image guided stereotactic prostate biopsy”. In: *Magnetic Resonance Imaging* 20.3 (Apr. 2002), pp. 295–299. DOI: 10.1016/S0730-725X(02)00490-3 (cit. on p. 37).

[112] A. Karamalis, W. Wein, T. Klein, and N. Navab. “Ultrasound confidence maps using random walks”. In: *Medical Image Analysis* 16.6 (Aug. 2012), pp. 1101–1112. DOI: 10.1016/j.media.2012.07.005 (cit. on pp. 74, 76, 79).

[113] A. Karamalis, W. Wein, O. Kutter, and N. Navab. “Fast hybrid freehand ultrasound volume reconstruction”. In: *SPIE Medical Imaging 2009: Visualization, Image-Guided Procedures, and Modeling*. Ed. by M. I. Miga and K. H. Wong. International Society for Optics and Photonics. Feb. 2009, p. 726114. DOI: 10.1117/12.813835 (cit. on pp. 23, 26, 78, 79).

[114] J. Kettenbach, G. Kronreif, M. Figl, M. Fürst, W. Birkfellner, R. Hanel, and H. Bergmann. “Robot-assisted biopsy using ultrasound guidance: initial results from in vitro tests”. In: *European Radiology* 15.4 (Apr. 2005), pp. 765–771. DOI: 10.1007/s00330-004-2487-x (cit. on p. 72).

[115] A. Kikuchi. “Medical ultrasound diagnosis in the near future as we move toward the era of the singularity”. In: *Journal of Medical Ultrasonics* 43.3 (July 2016), pp. 315–316. DOI: 10.1007/s10396-016-0712-2 (cit. on p. 8).

[116] R. Kojcev, B. Fuerst, O. Zettinig, J. Fotouhi, S. C. Lee, B. Frisch, R. Taylor, E. Sinibaldi, and N. Navab. “Dual-robot ultrasound-guided needle placement: closing the planning-imaging-action loop”. In: *International Journal of Computer Assisted Radiology and Surgery* 11.6 (June 2016), pp. 1173–1181. DOI: 10.1007/s11548-016-1408-1 (cit. on p. 134).

[117] R. Kojcev, A. Khakzar, B. Fuerst, O. Zettinig, C. Fahkry, R. DeJong, J. Richmon, R. Taylor, E. Sinibaldi, and N. Navab. “On the reproducibility of expert-operated and robotic ultrasound acquisitions”. In: *International Journal of Computer Assisted Radiology and Surgery* 12.6 (June 2017), pp. 1003–1011. DOI: 10.1007/s11548-017-1561-1 (cit. on p. 133).

[118] E. Kokkalis, P. Hoskins, G. Corner, P. Stonebridge, A. Doull, A. Melzer, and G. Houston. “Vector doppler imaging and secondary flow patterns in vascular prostheses”. In: *2012 IEEE International Ultrasonics Symposium*. IEEE, Oct. 2012, pp. 1–4. DOI: 10.1109/ULTSYM.2012.0475 (cit. on p. 104).

[119] G. Kossoff. “Display techniques in ultrasound pulse echo investigations: A review”. In: *Journal of Clinical Ultrasound* 2.1 (Mar. 1974), pp. 61–72. DOI: 10.1002/jcu.1870020111 (cit. on p. 6).

[120] S. Koundinyan, R. Toth, A. Madabhushi, and T. Maguire. “A statistical deformation model based regularizer for registration of histology and MRI”. In: *2014 40th Annual Northeast Bioengineering Conference (NEBEC)*. IEEE, Apr. 2014, pp. 1–2. DOI: 10.1109/NEBEC.2014.6972845 (cit. on p. 39).

[121] A. Krupa. “Automatic calibration of a robotized 3D ultrasound imaging system by visual servoing”. In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2006, pp. 4136–4141. DOI: 10.1109/ROBOT.2006.1642338 (cit. on p. 22).

[122] A. Krupa. “3D Steering of a Flexible Needle by Visual Servoing”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Springer, 2014, pp. 480–487. DOI: 10.1007/978-3-319-10404-1_60 (cit. on p. 74).

[123] A. Krupa, G. Fichtinger, and G. D. Hager. “Real-time Motion Stabilization with B-mode Ultrasound Using Image Speckle Information and Visual Servoing”. In: *The International Journal of Robotics Research* 28.10 (Oct. 2009), pp. 1334–1354. DOI: 10.1177/0278364909104066 (cit. on p. 74).

[124] A. Lasso, T. Heffter, A. Rankin, C. Pinter, T. Ungi, and G. Fichtinger. “PLUS: Open-Source Toolkit for Ultrasound-Guided Intervention Systems”. In: *IEEE Transactions on Biomedical Engineering* 61.10 (Oct. 2014), pp. 2527–2537. DOI: 10.1109/TBME.2014.2322864 (cit. on pp. 49, 88).

[125] D. Liepsch and S. Moravec. “Pulsatile flow of non-Newtonian fluid in distensible models of human arteries.” In: *Biorheology* 21.4 (1984), pp. 571–86. URL: http://www.ncbi.nlm.nih.gov/pubmed/6487768 (cit. on p. 104).

[126] Y. Liu, C. Zeng, M. Fan, L. Hu, C. Ma, and W. Tian. “Assessment of respiration-induced vertebral motion in prone-positioned patients during general anaesthesia”. In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 12.2 (June 2016), pp. 214–218. DOI: 10.1002/rcs.1676 (cit. on pp. 94, 98).

[127] N. Lonjon, E. Chan-Seng, V. Costalat, B. Bonnafoux, M. Vassal, and J. Boetto. “Robot-assisted spine surgery: feasibility study through a prospective case-matched analysis”. In: *European Spine Journal* 25.3 (Mar. 2016), pp. 947–955. DOI: 10.1007/s00586-015-3758-8 (cit. on p. 66).

[128] V. Mallapragada, N. Sarkar, and T. Podder. “Robot-Assisted Real-Time Tumor Manipulation for Breast Biopsy”. In: *IEEE Transactions on Robotics* 25.2 (Apr. 2009), pp. 316–324. DOI: 10.1109/TRO.2008.2011418 (cit. on p. 72).

[129] L. Manchikanti, V. Pampati, B. Fellows, and C. E. Bakhit. “Prevalence of lumbar facet joint pain in chronic low back pain.” In: *Pain physician* 2.3 (Oct. 1999), pp. 59–64. URL: http://www.ncbi.nlm.nih.gov/pubmed/16906217 (cit. on p. 65).

[130] L. Manchikanti, V. Singh, F. J. E. Falco, R. M. Benyamin, and J. A. Hirsch. “Epidemiology of Low Back Pain in Adults”. In: *Neuromodulation: Technology at the Neural Interface* 17 (Oct. 2014), pp. 3–10. DOI: 10.1111/ner.12018 (cit. on p. 65).

[131] T. Mansi, R. Beinart, O. Zettinig, S. Rapaka, B. Georgescu, A. Kamen, Y. Dori, M. M. Zviman, D. A. Herzka, H. R. Halperin, and D. Comaniciu. “A Framework for the Pre-clinical Validation of LBM-EP for the Planning and Guidance of Ventricular Tachycardia Ablation”. In: *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges*. Vol. 8330 LNCS. Springer Verlag, 2014, pp. 253–261. DOI: 10.1007/978-3-642-54268-8_30 (cit. on pp. 135, 138).

[132] L. Marks, S. Young, and S. Natarajan. “MRI–ultrasound fusion for guidance of targeted prostate biopsy”. In: *Current Opinion in Urology* 23.1 (Jan. 2013), pp. 43–50. DOI: 10.1097/MOU.0b013e32835ad3ee (cit. on p. 35).

[133] K. Masabni, S. Raza, E. H. Blackstone, H. L. Gornik, and J. F. Sabik. “Does preoperative carotid stenosis screening reduce perioperative stroke in patients undergoing coronary artery bypass grafting?” In: *The Journal of Thoracic and Cardiovascular Surgery* 149.5 (May 2015), pp. 1253–1260. DOI: 10.1016/j.jtcvs.2015.02.003 (cit. on p. 102).

[134] T. Maurer, A. J. Beer, H.-J. Wester, H. Kübler, M. Schwaiger, and M. Eiber. “Positron emission tomography/magnetic resonance imaging with 68 Gallium-labeled ligand of prostate-specific membrane antigen: Promising novel option in prostate cancer imaging?” In: *International Journal of Urology* 21.12 (Dec. 2014), pp. 1286–1288. DOI: 10.1111/iju.12577 (cit. on p. 36).

[135] T. Maurer, M. Eiber, M. Schwaiger, and J. E. Gschwend. “Current use of PSMA–PET in prostate cancer management”. In: *Nature Reviews Urology* 13.4 (Feb. 2016), pp. 226–235. DOI: 10.1038/nrurol.2016.26 (cit. on pp. 36, 128).

[136] G. Megali, O. Tonet, C. Stefanini, M. Boccadoro, V. Papaspyropoulos, L. Angelini, and P. Dario. “A Computer-Assisted Robotic Ultrasound-Guided Biopsy System for Video-Assisted Surgery”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*. 2001, pp. 343–350. DOI: 10.1007/3-540-45468-3_41 (cit. on p. 72).

[137] S. Merouche, L. Allard, E. Montagnon, G. Soulez, P. Bigras, and G. Cloutier. “A Robotic Ultrasound Scanner for Automatic Vessel Tracking and Three-Dimensional Reconstruction of B-Mode Images”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 63.1 (Jan. 2016), pp. 35–46. DOI: 10.1109/TUFFC.2015.2499084 (cit. on p. 73).

[138] F. Milletari, N. Navab, and S.-A. Ahmadi. “V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation”. In: *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, Oct. 2016, pp. 565–571. DOI: 10.1109/3DV.2016.79 (cit. on p. 128).

[139] F. Milletari, M. Yigitsoy, N. Navab, and S.-A. Ahmadi. “Left Ventricle Segmentation in Cardiac Ultrasound Using Hough-Forests With Implicit Shape and Appearance Priors”. In: *MIDAS Journal* (2014) (cit. on p. 42).

[140] J. Mitra, Z. Kato, R. Martí, A. Oliver, X. Lladó, D. Sidibé, S. Ghose, J. C. Vilanova, J. Comet, and F. Meriaudeau. “A spline-based non-linear diffeomorphism for multimodal prostate registration”. In: *Medical Image Analysis* 16.6 (Aug. 2012), pp. 1259–1279. DOI: 10.1016/j.media.2012.04.006 (cit. on p. 37).

[141] M. Modat, G. R. Ridgway, Z. A. Taylor, M. Lehmann, J. Barnes, D. J. Hawkes, N. C. Fox, and S. Ourselin. “Fast free-form deformation using graphics processing units”. In: *Computer Methods and Programs in Biomedicine* 98.3 (June 2010), pp. 278–284. DOI: 10.1016/j.cmpb.2009.09.002 (cit. on p. 81).

[142] J. Moore, C. Clarke, D. Bainbridge, C. Wedlake, A. Wiles, D. Pace, and T. Peters. “Image Guidance for Spinal Facet Injections Using Tracked Ultrasound”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009*. Springer, 2009, pp. 516–523. DOI: 10.1007/978-3-642-04268-3_64 (cit. on p. 70).

[143] N. Mottet, J. Bellmunt, E. Briers, R. C. N. van den Bergh, M. Bolla, N. J. van Casteren, P. Cornford, S. Culine, S. Joniau, T. Lam, M. D. Mason, V. Matveev, H. van der Poel, T. H. van der Kwast, O. Rouvière, and T. Wiegel. *Guidelines on Prostate Cancer*. European Association of Urology, 2015 (cit. on pp. 34, 35).

[144] E. Moult, T. Ungi, M. Welch, J. Lu, R. C. McGraw, and G. Fichtinger. “Ultrasound-guided facet joint injection training using Perk Tutor”. In: *International Journal of Computer Assisted Radiology and Surgery* 8.5 (Sept. 2013), pp. 831–836. DOI: 10.1007/s11548-012-0811-5 (cit. on pp. 66, 94).

[145] D. Mozaffarian, E. J. Benjamin, A. S. Go, et al. “Executive Summary: Heart Disease and Stroke Statistics—2016 Update:” in: *Circulation* 133.4 (Jan. 2016), pp. 447–454. DOI: 10.1161/CIR.0000000000000366 (cit. on p. 103).

[146] D. M. Muratore and R. L. Galloway. “Beam calibration without a phantom for creating a 3-D freehand ultrasound system”. In: *Ultrasound in Medicine & Biology* 27.11 (Nov. 2001), pp. 1557–1566. DOI: 10.1016/S0301-5629(01)00469-0 (cit. on p. 22).

[147] A. S. B. Mustafa, T. Ishii, Y. Matsunaga, R. Nakadate, H. Ishii, K. Ogawa, A. Saito, M. Sugawara, K. Niki, and A. Takanishi. “Development of robotic system for autonomous liver screening using ultrasound scanning device”. In: *2013 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE. IEEE, Dec. 2013, pp. 804–809. DOI: 10.1109/ROBIO.2013.6739561 (cit. on p. 73).

[148] A. Myronenko and Xubo Song. “Point Set Registration: Coherent Point Drift”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32.12 (Dec. 2010), pp. 2262–2275. DOI: 10.1109/TPAMI.2010.46 (cit. on pp. 25, 33, 39, 44, 45, 60, 86).

[149] C. Nadeau and A. Krupa. “Intensity-Based Ultrasound Visual Servoing: Modeling and Validation With 2-D and 3-D Probes”. In: *IEEE Transactions on Robotics* 29.4 (Aug. 2013), pp. 1003–1015. DOI: 10.1109/TRO.2013.2256690 (cit. on p. 74).

[150] C. Nadeau, H. Ren, A. Krupa, and P. Dupont. “Intensity-Based Visual Servoing for Instrument and Tissue Tracking in 3D Ultrasound Volumes”. In: *IEEE Transactions on Automation Science and Engineering* 12.1 (Jan. 2015), pp. 367–371. DOI: 10.1109/TASE.2014.2343652 (cit. on p. 74).

[151] L. Nahlawi, F. Imani, M. Gaed, J. A. Gomez, M. Moussa, E. Gibson, A. Fenster, A. D. Ward, P. Abolmaesumi, H. Shatkay, and P. Mousavi. “Prostate Cancer: Improved Tissue Characterization by Temporal Modeling of Radio-Frequency Ultrasound Echo Data”. In: *19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part I*. Springer International Publishing, 2016, pp. 644–652. DOI: 10.1007/978-3-319-46720-7_75 (cit. on p. 36).

[152] R. Nakadate, J. Solis, A. Takanishi, E. Minagawa, M. Sugawara, and K. Niki. “Out-of-plane visual servoing method for tracking the carotid artery with a robot-assisted ultrasound diagnostic system”. In: *2011 IEEE International Conference on Robotics and Automation*. IEEE. IEEE, May 2011, pp. 5267–5272. DOI: 10.1109/ICRA.2011.5979594 (cit. on p. 74).

[153] R. Narayanan, J. Kurhanewicz, K. Shinohara, E. D. Crawford, A. Simoneau, and J. S. Suri. "MRI-ultrasound registration for targeted prostate biopsy". In: *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, June 2009, pp. 991–994. DOI: 10.1109/ISBI.2009.5193221 (cit. on p. 37).

[154] S. Natarajan, L. S. Marks, D. J. Margolis, J. Huang, M. L. Macairan, P. Lieu, and A. Fenster. "Clinical application of a 3D ultrasound-guided prostate biopsy system". In: *Urologic Oncology: Seminars and Original Investigations* 29.3 (May 2011), pp. 334–342. DOI: 10.1016/j.urolonc.2011.02.014 (cit. on p. 36).

[155] National Health Service England. *Diagnostic Imaging Dataset 2015-16*. 2016. URL: https://www.england.nhs.uk/statistics/statistical-work-areas/diagnostic-imaging-dataset/ (visited on Jan. 27, 2016) (cit. on p. 8).

[156] A. Nchimi, J. Biquet, D. Brisbois, P. Reginster, K. Bouali, C. Saive, and P. Magotteaux. "Duplex ultrasound as first-line screening test for patients suspected of renal artery stenosis: prospective evaluation in high-risk group". In: *European Radiology* 13.6 (2003), pp. 1413–1419. DOI: 10.1007/s00330-002-1685-7 (cit. on p. 64).

[157] T. R. Nelson, A. Tran, H. Fakourfar, and J. Nebeker. "Positional calibration of an ultrasound image-guided robotic breast biopsy system." In: *Journal of Ultrasound in Medicine* 31.3 (Mar. 2012), pp. 351–359. URL: http://www.ncbi.nlm.nih.gov/pubmed/22368125 (cit. on p. 72).

[158] W. Ng, V. Chung, S. Vasan, and P. Lim. "Robotic radiation seed implantation for prostatic cancer". In: *Proceedings of 18th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. Vol. 1. IEEE, 1996, pp. 231–233. DOI: 10.1109/IEMBS.1996.656930 (cit. on p. 72).

[159] S. Nouranian, M. Ramezani, I. Spadinger, W. J. Morris, S. E. Salcudean, and P. Abolmaesumi. "Learning-Based Multi-Label Segmentation of Transrectal Ultrasound Images for Prostate Brachytherapy". In: *IEEE Transactions on Medical Imaging* 35.3 (Mar. 2016), pp. 921–932. DOI: 10.1109/TMI.2015.2502540 (cit. on p. 128).

[160] J. A. Onofrey, L. H. Staib, S. Sarkar, R. Venkataraman, and X. Papademetris. "Learning nonrigid deformations for constrained point-based registration for image-guided MR-TRUS prostate intervention". In: *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*. Vol. 2015-July. IEEE. IEEE, Apr. 2015, pp. 1592–1595. DOI: 10.1109/ISBI.2015.7164184 (cit. on p. 38).

[161] S. Onogi, T. Yoshida, Y. Sugano, T. Mochizuki, and K. Masuda. "Robotic Ultrasound Guidance by B-scan Plane Positioning Control". In: *Procedia CIRP* 5 (2013), pp. 100–103. DOI: 10.1016/j.procir.2013.01.020 (cit. on p. 73).

[162] C. M. Otto. *Textbook of Clinical Echocardiography*. 5th ed. Elsevier Health Sciences, 2013. URL: https://books.google.de/books?id=cwIwSSAksssC (cit. on p. 102).

[163] S. G. Patel. "Outcome, observer reliability, and patient preferences if CTA, MRA, or Doppler ultrasound were used, individually or together, instead of digital subtraction angiography before carotid endarterectomy". In: *Journal of Neurology, Neurosurgery & Psychiatry* 73.1 (July 2002), pp. 21–28. DOI: 10.1136/jnnp.73.1.21 (cit. on p. 104).

[164] G. Pedrizzetti, G. La Canna, O. Alfieri, and G. Tonti. "The vortex—an early predictor of cardiovascular outcome?" In: *Nature Reviews Cardiology* 11.9 (June 2014), pp. 545–553. DOI: 10.1038/nrcardio.2014.75 (cit. on p. 102).

[165] K. Pfister, W. Schierling, E. M. Jung, H. Apfelbeck, C. Hennersperger, and P. M. Kasprzak. "Standardized 2D ultrasound versus 3D/4D ultrasound and image fusion for measurement of aortic aneurysm diameter in follow-up after EVAR". In: *Clinical Hemorheology and Microcirculation* 62.3 (Apr. 2016), pp. 249–260. DOI: 10.3233/CH-152012 (cit. on pp. 67, 87, 92).

[166] M. J. Pihl, M. B. Stuart, B. G. Tomov, J. M. Hansen, M. F. Rasmussen, and J. A. Jensen. "Preliminary examples of 3D vector flow imaging". In: ed. by J. G. Bosch and M. M. Doyley. Mar. 2013, 86750H. DOI: 10.1117/12.2006845 (cit. on p. 104).

[167] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat. “Comparing ICP variants on real-world data sets”. In: *Autonomous Robots* 34.3 (Apr. 2013), pp. 133–148. DOI: 10.1007/s10514-013-9327-2 (cit. on pp. 24, 25).

[168] M. Postema. *Fundamentals of Medical Ultrasonics*. CRC Press, 2011 (cit. on p. 12).

[169] M. J. D. Powell. “The BOBYQA algorithm for bound constrained optimization without derivatives”. In: *Cambridge NA Report NA2009/06, University of Cambridge, Cambridge* (2009) (cit. on pp. 26, 46, 81).

[170] A. Priester, S. Natarajan, and M. Culjat. “Robotic ultrasound systems in medicine”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 60.3 (Mar. 2013), pp. 507–523. DOI: 10.1109/TUFFC.2013.2593 (cit. on pp. 71–73).

[171] J. Provost, C. Papadacci, J. E. Arango, M. Imbault, M. Fink, J.-L. Gennisson, M. Tanter, and M. Pernot. “3D ultrafast ultrasound imaging in vivo”. In: *Physics in Medicine and Biology* 59.19 (Oct. 2014), pp. L1–L13. DOI: 10.1088/0031-9155/59/19/L1 (cit. on pp. 100, 105, 129).

[172] R. N. Rankin, A. Fenster, D. B. Downey, P. L. Munk, M. F. Levin, and A. D. Vellet. “Three-dimensional sonographic reconstruction: techniques and diagnostic applications.” In: *American Journal of Roentgenology* 161.4 (Oct. 1993), pp. 695–702. DOI: 10.2214/ajr.161.4.8372741 (cit. on p. 19).

[173] A. Rasoulian, J. Osborn, S. Sojoudi, S. Nouranian, V. A. Lessoway, R. N. Rohling, and P. Abolmaesumi. “A System for Ultrasound-Guided Spinal Injections: A Feasibility Study”. In: *Information Processing in Computer-Assisted Interventions*. Springer, 2014, pp. 90–99. DOI: 10.1007/978-3-319-07521-1_10 (cit. on p. 70).

[174] J. W. S. Rayleigh. *The Theory of Sound I, II*. MacMillan and Co, London, 1877 (cit. on p. 5).

[175] D. M. Reif, A. A. Motsinger, B. A. McKinney, J. E. Crowe, and J. H. Moore. “Feature Selection using a Random Forests Classifier for the Integrated Analysis of Multiple Data Types”. In: *2006 IEEE Symposium on Computational Intelligence and Bioinformatics and Computational Biology*. IEEE. IEEE, Sept. 2006, pp. 1–8. DOI: 10.1109/CIBCB.2006.330987 (cit. on p. 42).

[176] K. Rematas and B. Leibe. “Efficient object detection and segmentation with a cascaded Hough Forest ISM”. In: *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*. IEEE, Nov. 2011, pp. 966–973. DOI: 10.1109/ICCVW.2011.6130356 (cit. on p. 42).

[177] M. Reyes, M. A. Gonzalez Ballester, Zhixi Li, N. Kozic, See Chin, R. M. Summers, and M. G. Linguraru. “Anatomical variability of organs via principal factor analysis from the construction of an abdominal probabilistic atlas”. In: *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE. IEEE, June 2009, pp. 682–685. DOI: 10.1109/ISBI.2009.5193139 (cit. on p. 86).

[178] C. Reynier, J. Troccaz, P. Fourneret, A. Dusserre, C. Gay-Jeune, J.-L. Descotes, M. Bolla, and J.-Y. Giraud. “MRI/TRUS data fusion for prostate brachytherapy. Preliminary results”. In: *Medical Physics* 31.6 (2004), p. 1568. DOI: 10.1118/1.1739003 (cit. on pp. 37, 53).

[179] F. Ringel, J. Villard, Y.-M. Ryang, and B. Meyer. “Navigation, Robotics, and Intraoperative Imaging in Spinal Surgery”. In: *Advances and Technical Standards in Neurosurgery*. 2014, pp. 3–22. DOI: 10.1007/978-3-319-01830-0_1 (cit. on p. 72).

[180] L. M. Rios and N. V. Sahinidis. “Derivative-free optimization: a review of algorithms and comparison of software implementations”. In: *Journal of Global Optimization* 56.3 (July 2013), pp. 1247–1293. DOI: 10.1007/s10898-012-9951-y (cit. on p. 26).

[181] F. Rousseau, P. Hellier, and C. Barillot. “A fully automatic calibration procedure for freehand 3D ultrasound”. In: *Proceedings IEEE International Symposium on Biomedical Imaging*. IEEE, 2002, pp. 985–988. DOI: 10.1109/ISBI.2002.1029428 (cit. on p. 22).

[182] D. Rueckert, A. Frangi, and J. Schnabel. “Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration”. In: *IEEE Transactions on Medical Imaging* 22.8 (Aug. 2003), pp. 1014–1025. DOI: 10.1109/TMI.2003.815865 (cit. on pp. 38, 46).

[183] D. Rueckert and J. A. Schnabel. “Medical Image Registration”. In: *Biomedical Image Processing*. Ed. by Thomas Martin Deserno. Springer Berlin Heidelberg, 2010. Chap. 5, pp. 131–154. DOI: 10.1007/978-3-642-15816-2_5 (cit. on pp. 24, 26).

[184] S. E. Salcudean, G. Bell, S. Bachmann, W. H. Zhu, P. Abolmaesumi, and P. D. Lawrence. “Robot-Assisted Diagnostic Ultrasound – Design and Feasibility Experiments”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI’99*. 1999, pp. 1062–1071. DOI: 10.1007/10704282_115 (cit. on pp. 69, 73).

[185] C. Samson, B. Espiau, and M. L. Borgne. *Robot Control: The Task Function Approach*. Oxford University Press, 1991 (cit. on p. 69).

[186] A. P. Sarvazyan, O. V. Rudenko, S. D. Swanson, J. Fowlkes, and S. Y. Emelianov. “Shear wave elasticity imaging: a new ultrasonic technology of medical diagnostics”. In: *Ultrasound in Medicine & Biology* 24.9 (Dec. 1998), pp. 1419–1435. DOI: 10.1016/S0301-5629(98)00110-0 (cit. on p. 8).

[187] W. Schäberle, L. Leyerer, W. Schierling, and K. Pfister. “Ultrasound diagnostics of the abdominal aorta”. In: *Gefässchirurgie* 20.S1 (Jan. 2015), pp. 22–27. DOI: 10.1007/s00772-014-1411-1 (cit. on pp. 67, 92).

[188] T. Schmidt, N. Mühlberger, I. Chemelli-Steingruber, A. Strasak, B. Kofler, A. Chemelli, and U. Siebert. “Benefit, Risks and Cost-Effectiveness of Screening for Abdominal Aortic Aneurysm”. In: *RöFo - Fortschritte auf dem Gebiet der Röntgenstrahlen und der bildgebenden Verfahren* 182.07 (July 2010), pp. 573–580. DOI: 10.1055/s-0029-1245140 (cit. on p. 68).

[189] C. Schneider, A. Okamura, and G. Fichtinger. “A robotic system for transrectal needle insertion into the prostate with integrated ultrasound”. In: *IEEE International Conference on Robotics and Automation (ICRA), 2004.* Vol. 1. IEEE, 2004, pp. 365–370. DOI: 10.1109/ROBOT.2004.1307177 (cit. on p. 72).

[190] D. Schnell, M. Reynaud, M. Venot, A. L. Le Maho, M. Dinic, M. Baulieu, G. Ducos, J. Terreaux, F. Zeni, E. Azoulay, F. Meziani, J. Duranteau, and M. Darmon. “Resistive Index or color-Doppler semi-quantitative evaluation of renal perfusion by inexperienced physicians: results of a pilot study.” In: *Minerva anestesiologica* 80.12 (Dec. 2014), pp. 1273–81. URL: http://www.ncbi.nlm.nih.gov/pubmed/24569358 (cit. on p. 102).

[191] C. Schulte zu Berge, D. Declara, C. Hennersperger, M. Baust, and N. Navab. “Real-time uncertainty visualization for B-mode ultrasound”. In: *2015 IEEE Scientific Visualization Conference (SciVis)*. IEEE, Oct. 2015, pp. 33–40. DOI: 10.1109/SciVis.2015.7429489 (cit. on p. 79).

[192] C. Schulte zu Berge, A. Grunau, H. Mahmud, and N. Navab. *CAMPVis - A Game Engine-inspired Research Framework for Medical Imaging and Visualization*. Tech. rep. Technische Universität München, 2014. URL: http://campar.in.tum.de/Main/CAMPVis (cit. on p. 49).

[193] C. Schulte zu Berge, A. Kapoor, and N. Navab. “Orientation-Driven Ultrasound Compounding Using Uncertainty Information”. In: *Information Processing in Computer-Assisted Interventions, 5th International Conference, IPCAI 2014*. 2014, pp. 236–245. DOI: 10.1007/978-3-319-07521-1_25 (cit. on p. 23).

[194] A. Shah, O. Zettinig, T. Maurer, C. Precup, C. Schulte zu Berge, J. Weiss, B. Frisch, and N. Navab. “An Open Source Multimodal Image-Guided Prostate Biopsy Framework”. In: *Clinical Image-Based Procedures. Translational Research in Medical Imaging*. Vol. 8680. Springer LNCS, 2014, pp. 1–8. DOI: 10.1007/978-3-319-13909-8_1 (cit. on pp. 39, 41, 49, 51, 135).

[195] A. Shah, O. Zettinig, E. Storz, T. Maurer, M. Eiber, N. Navab, and B. Frisch. “Challenges in Multimodal Image-guided Targeted Prostate Biopsy”. In: *Hamlyn Symposium on Medical Robotics, London, UK*. 2015, pp. 21–22. URL: http://hamlyn.doc.ic.ac.uk/hsmr/sites/default/files/HSMR-2015-Proceedings-FINAL.pdf (cit. on p. 134).

[196] P. Shankar. “Ultrasonic tissue characterization using a generalized Nakagami model”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 48.6 (2001), pp. 1716–1720. DOI: 10.1109/58.971725 (cit. on p. 13).

[197] K. K. Shung. “Ultrasound: Past, Present and Future”. In: *The Third International Conference on the Development of Biomedical Engineering in Vietnam. IFMBE Proceedings*. Ed. by V. Van Toi and T. Khoa. Vol. 27. Springer, Berlin, Heidelberg, 2010, pp. 10–13. DOI: 10.1007/978-3-642-12020-6_3 (cit. on p. 8).

[198] M. M. Siddiqui, S. Rais-Bahrami, B. Turkbey, A. K. George, J. Rothwax, N. Shakir, C. Okoro, D. Raskolnikov, H. L. Parnes, W. M. Linehan, M. J. Merino, R. M. Simon, P. L. Choyke, B. J. Wood, and P. A. Pinto. “Comparison of MR/Ultrasound Fusion–Guided Biopsy With Ultrasound-Guided Biopsy for the Diagnosis of Prostate Cancer”. In: *JAMA* 313.4 (Jan. 2015), p. 390. DOI: 10.1001/jama.2014.17942 (cit. on pp. 36, 38).

[199] R. L. Siegel, K. D. Miller, and A. Jemal. “Cancer statistics, 2016”. In: *CA: A Cancer Journal for Clinicians* 66.1 (Jan. 2016), pp. 7–30. DOI: 10.3322/caac.21332 (cit. on p. 34).

[200] S. Smith, H. Pavy, and O. von Ramm. “High-speed ultrasound volumetric imaging system. I. Transducer design and beam steering”. In: *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on* 38.2 (Mar. 1991), pp. 100–108. DOI: 10.1109/58.68466 (cit. on p. 7).

[201] M. Soehl, R. Walsh, A. Rankin, A. Lasso, and G. Fichtinger. “Tracked ultrasound calibration studies with a phantom made of LEGO bricks”. In: *SPIE Medical Imaging 2014: Image-Guided Procedures, Robotic Interventions, and Modeling*. Ed. by Z. R. Yaniv and D. R. Holmes. Mar. 2014, 90362R. DOI: 10.1117/12.2044121 (cit. on p. 22).

[202] U. Soergel. *Radar remote sensing of urban areas*. Vol. 15. Springer Science & Business Media, 2010 (cit. on p. 13).

[203] O. V. Solberg, F. Lindseth, H. Torp, R. E. Blake, and T. A. Nagelhus Hernes. “Freehand 3D Ultrasound Reconstruction Algorithms—A Review”. In: *Ultrasound in Medicine & Biology* 33.7 (July 2007), pp. 991–1009. DOI: 10.1016/j.ultrasmedbio.2007.02.015 (cit. on pp. 20, 22, 23).

[204] C. G. Solomon and K. C. Kent. “Abdominal Aortic Aneurysms”. In: *New England Journal of Medicine* 371.22 (Nov. 2014), pp. 2101–2108. DOI: 10.1056/NEJMcp1401430 (cit. on pp. 66, 67).

[205] N. J. Soni, R. Franco-Sadud, D. Schnobrich, R. Dancel, D. M. Tierney, G. Salame, M. I. Restrepo, and P. McHardy. “Ultrasound guidance for lumbar puncture”. In: *Neurology: Clinical Practice* 6.4 (2016), pp. 358–368. DOI: 10.1212/CPJ.0000000000000265 (cit. on pp. 64, 66).

[206] G. A. Sonn, D. J. Margolis, and L. S. Marks. “Target detection: Magnetic resonance imaging-ultrasound fusion–guided prostate biopsy”. In: *Urologic Oncology: Seminars and Original Investigations* 32.6 (Aug. 2014), pp. 903–911. DOI: 10.1016/j.urolonc.2013.08.006 (cit. on pp. 20, 36).

[207] A. Sotiras, C. Davatzikos, and N. Paragios. “Deformable Medical Image Registration: A Survey”. In: *IEEE Transactions on Medical Imaging* 32.7 (July 2013), pp. 1153–1190. DOI: 10.1109/TMI.2013.2265603 (cit. on pp. 26, 37, 38).

[208] R. Sparks, B. N. Bloch, E. Feleppa, D. Barratt, and A. Madabhushi. “Fully automated prostate magnetic resonance imaging and transrectal ultrasound fusion via a probabilistic registration metric”. In: *SPIE Medical Imaging*. Ed. by D. R. Holmes and Z. R. Yaniv. International Society for Optics and Photonics. Mar. 2013, 86710A. DOI: 10.1117/12.2007610 (cit. on p. 37).

[209] D. Sperling. “MRI-Ultrasound Fusion Imaging”. In: *Image Guided Prostate Cancer Treatments*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 115–123. DOI: 10.1007/978-3-642-40429-0_11 (cit. on p. 37).

[210] M. W. Spong. “Modeling and Control of Elastic Joint Robots”. In: *Journal of Dynamic Systems, Measurement, and Control* 109.4 (1987), p. 310. DOI: 10.1115/1.3143860 (cit. on p. 77).

[211] P. Stather, D. Sidloff, I. Rhema, E. Choke, M. Bown, and R. Sayers. “A Review of Current Reporting of Abdominal Aortic Aneurysm Mortality and Prevalence in the Literature”. In: *European Journal of Vascular and Endovascular Surgery* 47.3 (Mar. 2014), pp. 240–242. DOI: 10.1016/j.ejvs.2013.11.007 (cit. on pp. 66, 67).

[212] E. Storz, A. Shah, O. Zettinig, M. Eiber, H.-J. Wester, H. Kübler, J. E. Gschwend, M. Schwaiger, B. Frisch, and T. Maurer. “PSMA-PET/MRI-guided transrectal fusion biopsy for the detection of prostate cancer”. In: *European Urology Supplements* 14.2 (Apr. 2015), e217. DOI: 10.1016/S1569-9056(15)60217-2 (cit. on pp. 39, 134).

[213] Y. Sun, J. Yuan, M. Rajchl, W. Qiu, C. Romagnoli, and A. Fenster. “Efficient Convex Optimization Approach to 3D Non-rigid MR-TRUS Registration”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*. Springer, 2013, pp. 195–202. DOI: 10.1007/978-3-642-40811-3_25 (cit. on pp. 38, 58, 61).

[214] A. Swillens, P. Segers, H. Torp, and L. Lovstakken. “Two-dimensional blood velocity estimation with ultrasound: speckle tracking versus crossed-beam vector doppler based on flow simulations in a carotid bifurcation model”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 57.2 (Feb. 2010), pp. 327–339. DOI: 10.1109/TUFFC.2010.1413 (cit. on p. 104).

[215] T. Szabo. *Diagnostic Ultrasound Imaging: Inside Out*. 2nd ed. Academic Press, 2014 (cit. on pp. 3, 6–8, 10–13, 15, 17).

[216] A. M. Tahmasebi, R. Sharifi, H. K. Agarwal, B. Turkbey, M. Bernardo, P. Choyke, P. Pinto, B. Wood, and J. Kruecker. “A statistical model-based technique for accounting for prostate gland deformation in endorectal coil-based MR imaging”. In: *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, Aug. 2012, pp. 5412–5415. DOI: 10.1109/EMBC.2012.6347218 (cit. on p. 39).

[217] T. Taube. “Robotic Acquisition for Advanced 3D Reconstruction of Doppler Ultrasound”. Master’s thesis. Technische Universität München, 2017 (cit. on pp. 101, 111).

[218] R. Taylor, B. Mittelstadt, H. Paul, W. Hanson, P. Kazanzides, J. Zuhars, B. Williamson, B. Musits, E. Glassman, and W. Bargar. “An image-directed robotic system for precise orthopaedic surgery”. In: *IEEE Transactions on Robotics and Automation* 10.3 (June 1994), pp. 261–275. DOI: 10.1109/70.294202 (cit. on p. 71).

[219] R. Taylor and D. Stoianovici. “Medical robotics in computer-integrated surgery”. In: *IEEE Transactions on Robotics and Automation* 19.5 (Oct. 2003), pp. 765–781. DOI: 10.1109/TRA.2003.817058 (cit. on p. 71).

[220] J. Tokuda, G. S. Fischer, X. Papademetris, Z. Yaniv, L. Ibanez, P. Cheng, H. Liu, J. Blevins, J. Arata, A. J. Golby, T. Kapur, S. Pieper, E. C. Burdette, G. Fichtinger, C. M. Tempany, and N. Hata. “OpenIGTLink: an open network protocol for image-guided therapy environment”. In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 5.4 (Dec. 2009), pp. 423–434. DOI: 10.1002/rcs.274 (cit. on p. 88).

[221] G. E. Trahey, J. W. Allison, and O. T. von Ramm. “Angle Independent Ultrasonic Detection of Blood Flow”. In: *IEEE Transactions on Biomedical Engineering* BME-34.12 (Dec. 1987), pp. 965–967. DOI: 10.1109/TBME.1987.325938 (cit. on p. 104).

[222] D. Tran, A. A. Kamani, E. Al-Attas, V. A. Lessoway, S. Massey, and R. N. Rohling. “Single-operator real-time ultrasound-guidance to aim and insert a lumbar epidural needle”. In: *Canadian Journal of Anesthesia/Journal canadien d’anesthésie* 57.4 (Apr. 2010), pp. 313–321. DOI: 10.1007/s12630-009-9252-1 (cit. on p. 70).

[223] R. Tsai and R. Lenz. “A new technique for fully autonomous and efficient 3D robotics hand/eye calibration”. In: *IEEE Transactions on Robotics and Automation* 5.3 (June 1989), pp. 345–358. DOI: 10.1109/70.34770 (cit. on p. 91).

[224] B. Turkbey, P. A. Pinto, and P. L. Choyke. “Imaging techniques for prostate cancer: implications for focal therapy”. In: *Nature Reviews Urology* 6.4 (Apr. 2009), pp. 191–203. DOI: 10.1038/nrurol.2009.27 (cit. on pp. 35, 36).

[225] J. Tyndall. *Sound*. 2014th ed. Cambridge University Press, 1867 (cit. on p. 5).

[226] S. Umeyama. “Least-squares estimation of transformation parameters between two point patterns”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13.4 (Apr. 1991), pp. 376–380. DOI: 10.1109/34.88573 (cit. on pp. 24, 55).

[227] T. Ungi, P. Abolmaesumi, R. Jalal, M. Welch, I. Ayukawa, S. Nagpal, A. Lasso, M. Jaeger, D. P. Borschneck, G. Fichtinger, and P. Mousavi. “Spinal Needle Navigation by Tracked Ultrasound Snapshots”. In: *IEEE Transactions on Biomedical Engineering* 59.10 (Oct. 2012), pp. 2766–2772. DOI: 10.1109/TBME.2012.2209881 (cit. on p. 70).

[228] F. Viazzi, G. Leoncini, L. E. Derchi, and R. Pontremoli. “Ultrasound Doppler renal resistive index”. In: *Journal of Hypertension* 32.1 (Jan. 2014), pp. 149–153. DOI: 10.1097/HJH.0b013e328365b29c (cit. on p. 102).

[229] P. Viola and W. M. Wells III. “Alignment by Maximization of Mutual Information”. In: *International Journal of Computer Vision* 24.2 (1997), pp. 137–154. DOI: 10.1023/A:1007958904918 (cit. on p. 29).

[230] S. Virga*, O. Zettinig*, M. Esposito, K. Pfister, B. Frisch, T. Neff, N. Navab, and C. Hennersperger. “Automatic force-compliant robotic ultrasound screening of abdominal aortic aneurysms”. In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Oct. 2016, pp. 508–513. DOI: 10.1109/IROS.2016.7759101 (cit. on pp. 64, 80, 85, 88, 96–99, 133).

[231] C. Wachinger, T. Klein, and N. Navab. “Locally adaptive Nakagami-based ultrasound similarity measures”. In: *Ultrasonics* 52.4 (Apr. 2012), pp. 547–554. DOI: 10.1016/j.ultras.2011.11.009 (cit. on p. 13).

[232] I. Waechter, J. Bredno, R. Hermans, J. Weese, D. C. Barratt, and D. J. Hawkes. “Model-based blood flow quantification from rotational angiography”. In: *Medical Image Analysis* 12.5 (Oct. 2008), pp. 586–602. DOI: 10.1016/j.media.2008.06.003 (cit. on pp. 105, 106, 108).

[233] A. Walker, E. Olsson, B. Wranne, I. Ringqvist, and P. Ask. “Accuracy of spectral Doppler flow and tissue velocity measurements in ultrasound systems”. In: *Ultrasound in Medicine & Biology* 30.1 (Jan. 2004), pp. 127–132. DOI: 10.1016/j.ultrasmedbio.2003.08.020 (cit. on p. 122).

[234] Y. Wang, D. Ni, J. Qin, M. Lin, X. Xie, M. Xu, and P. A. Heng. “Towards Personalized Biomechanical Model and MIND-Weighted Point Matching for Robust Deformable MR-TRUS Registration”. In: *Computer-Assisted and Robotic Endoscopy, Lecture Notes in Computer Science*. Vol. 8899. Springer, 2014, pp. 121–130. DOI: 10.1007/978-3-319-13410-9_12 (cit. on p. 38).

[235] P. Weber, J. Schlegel, J. Meiche, L. Peter, and U. Harland. “A system for ultrasound based intraoperative navigation in spine surgery”. In: *2001 IEEE Ultrasonics Symposium. Proceedings.* Vol. 2. IEEE, 2001, pp. 1361–1364. DOI: 10.1109/ULTSYM.2001.991973 (cit. on p. 72).

[236] W. Wein, S. Brunke, A. Khamene, M. R. Callstrom, and N. Navab. “Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention”. In: *Medical Image Analysis* 12.5 (Oct. 2008), pp. 577–585. DOI: 10.1016/j.media.2008.06.006 (cit. on p. 83).

[237] W. Wein and A. Khamene. “Image-based method for in-vivo freehand ultrasound calibration”. In: *SPIE Medical Imaging 2008: Ultrasonic Imaging and Signal Processing*. Ed. by S. A. McAleavey and J. D’hooge. International Society for Optics and Photonics. Mar. 2008, 69200K. DOI: 10.1117/12.769948 (cit. on pp. 22, 78, 114).

[238] W. Wein, A. Ladikos, B. Fuerst, A. Shah, K. Sharma, and N. Navab. “Global Registration of Ultrasound to MRI Using the LC2 Metric for Enabling Neurosurgical Guidance”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2013*. 2013, pp. 34–41. DOI: 10.1007/978-3-642-40811-3_5 (cit. on p. 27).

[239] W. Wein, F. Pache, B. Röper, and N. Navab. “Backward-Warping Ultrasound Reconstruction for Improving Diagnostic Value and Registration”. In: *Med. Image Comput. Comput. Assist. Invterv. - MICCAI 2006*. Chair for Computer Aided Medical Procedures (CAMP), TU Munich Boltzmannstr. 3, 85748 Garching, Germany. wein@cs.tum.edu. 2006, pp. 750–757. DOI: 10.1007/11866763_92 (cit. on p. 23).

[240] J. J. Wild and J. M. Reid. “Application of Echo-Ranging Techniques to the Determination of Structure of Biological Tissues”. In: *Science* 115.2983 (Feb. 1952), pp. 226–230. DOI: 10.1126/science.115.2983.226 (cit. on p. 6).

[241] T. Wong, K. Tay, M. Sebastian, and S. Tan. “Duplex ultrasonography arteriography as first-line investigation for peripheral vascular disease”. In: *Singapore Medical Journal* 54.5 (May 2013), pp. 271–274. DOI: 10.11622/smedj.2013107 (cit. on p. 64).

[242] J. Wouters, E. D’Agostino, F. Maes, D. Vandermeulen, and P. Suetens. “Non-rigid brain image registration using a statistical deformation model”. In: *SPIE Medical Imaging 2006: Image Processing*. Ed. by J. M. Reinhardt and J. P. W. Pluim. International Society for Optics and Photonics. Mar. 2006, p. 614411. DOI: 10.1117/12.653081 (cit. on p. 38).

[243] Wu Qiu, Jing Yuan, E. Ukwatta, Yue Sun, M. Rajchl, and A. Fenster. “Prostate Segmentation: An Efficient Convex Optimization Approach With Axial Symmetry Using 3-D TRUS and MR Images”. In: *IEEE Transactions on Medical Imaging* 33.4 (Apr. 2014), pp. 947–960. DOI: 10.1109/TMI.2014.2300694 (cit. on pp. 38, 54, 60).

[244] S. Xu, J. Kruecker, P. Guion, N. Glossop, Z. Neeman, P. Choyke, A. K. Singh, and B. J. Wood. “Closed-Loop Control in Fused MR-TRUS Image-Guided Prostate Biopsy”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 128–135. DOI: 10.1007/978-3-540-75757-3_16 (cit. on p. 37).

[245] C. X. B. Yan, B. Goulet, J. Pelletier, S. J.-S. Chen, D. Tampieri, and D. L. Collins. “Towards accurate, robust and practical ultrasound-CT registration of vertebrae for image-guided spine surgery”. In: *International Journal of Computer Assisted Radiology and Surgery* 6.4 (July 2011), pp. 523–537. DOI: 10.1007/s11548-010-0536-2 (cit. on p. 70).

[246] K. Yan, C. Li, X. Wang, A. Li, Y. Yuan, D. Feng, M. Khadra, and J. Kim. “Automatic prostate segmentation on MR images with deep network and graph model”. In: *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. IEEE, Aug. 2016, pp. 635–638. DOI: 10.1109/EMBC.2016.7590782 (cit. on p. 38).

[247] C. Yang, Z. He, and W. Yu. “Comparison of public peak detection algorithms for MALDI mass spectrometry data analysis”. In: *BMC Bioinformatics* 10.1 (2009), p. 4. DOI: 10.1186/1471-2105-10-4 (cit. on p. 114).

[248] S. H. Yoon, S. L. O’Brien, and M. Tran. “Ultrasound Guided Spine Injections: Advancement Over Fluoroscopic Guidance?” In: *Current Physical Medicine and Rehabilitation Reports* 1.2 (June 2013), pp. 104–113. DOI: 10.1007/s40141-013-0016-z (cit. on pp. 66, 70).

[249] O. Zettinig*, B. Frisch*, S. Virga, M. Esposito, A. Rienmüller, B. Meyer, C. Hennersperger, Y.-M. Ryang*, and N. Navab*. “3D ultrasound registration-based visual servoing for neurosurgical navigation”. In: *International Journal of Computer Assisted Radiology and Surgery* 12.9 (Sept. 2017), pp. 1607–1619. DOI: 10.1007/s11548-017-1536-2 (cit. on pp. 64, 77, 81, 82, 85, 90, 91, 93–95, 133).

[250] O. Zettinig, B. Fuerst, R. Kojcev, M. Esposito, M. Salehi, W. Wein, J. Rackerseder, E. Sinibaldi, B. Frisch, and N. Navab. “Toward real-time 3D ultrasound registration-based visual servoing for interventional navigation”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2016, pp. 945–950. DOI: 10.1109/ICRA.2016.7487226 (cit. on pp. 64, 78, 83, 84, 133).

[251] O. Zettinig*, C. Hennersperger*, C. Schulte zu Berge, M. Baust, and N. Navab. “3D Velocity Field and Flow Profile Reconstruction from Arbitrarily Sampled Doppler Ultrasound Data”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Vol. 8674 LNCS. PART 2. Springer Verlag, 2014, pp. 611–618. DOI: 10.1007/978-3-319-10470-6_76 (cit. on pp. 101, 106, 110, 111, 116, 117, 119, 120, 135).

[252] O. Zettinig, T. Mansi, B. Georgescu, E. Kayvanpour, F. Sedaghat-Hamedani, A. Amr, J. Haas, H. Steen, B. Meder, H. Katus, N. Navab, A. Kamen, and D. Comaniciu. “Fast Data-Driven Calibration of a Cardiac Electrophysiology Model from Images and ECG”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*. Vol. 8149. Springer LNCS, 2013, pp. 1–8. DOI: 10.1007/978-3-642-40811-3_1 (cit. on pp. 135, 140).

[253] O. Zettinig, T. Mansi, B. Georgescu, S. Rapaka, A. Kamen, J. Haas, K. S. Frese, F. Sedaghat-Hamedani, E. Kayvanpour, A. Amr, S. Hardt, D. Mereles, H. Steen, A. Keller, H. A. Katus, B. Meder, N. Navab, and D. Comaniciu. “From Medical Images to Fast Computational Models of Heart Electromechanics: An Integrated Framework towards Clinical Use”. In: *Functional Imaging and Modeling of the Heart*. Vol. 7945. Springer LNCS, 2013, pp. 249–258. DOI: 10.1007/978-3-642-38899-6_30 (cit. on pp. 9, 135, 139).

[254] O. Zettinig, T. Mansi, D. Neumann, B. Georgescu, S. Rapaka, P. Seegerer, E. Kayvanpour, F. Sedaghat-Hamedani, A. Amr, J. Haas, H. Steen, H. Katus, B. Meder, N. Navab, A. Kamen, and D. Comaniciu. “Data-driven estimation of cardiac electrical diffusivity from 12-lead ECG signals”. In: *Medical Image Analysis* 18.8 (Dec. 2014), pp. 1361–1376. DOI: 10.1016/j.media.2014.04.011 (cit. on pp. 134, 137).

[255] O. Zettinig, J. Rackerseder, B. Lentes, T. Maurer, K. Westenfelder, M. Eiber, B. Frisch, and N. Navab. “Preconditioned intensity-based prostate registration using statistical deformation models”. In: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*. IEEE, Apr. 2017, pp. 853–857. DOI: 10.1109/ISBI.2017.7950651 (cit. on pp. 34, 39, 46, 56–59, 133).

[256] O. Zettinig, A. Shah, C. Hennersperger, M. Eiber, C. Kroll, H. Kübler, T. Maurer, F. Milletarì, J. Rackerseder, C. Schulte zu Berge, E. Storz, B. Frisch, and N. Navab. “Multimodal image-guided prostate fusion biopsy based on automatic deformable registration”. In: *International Journal of Computer Assisted Radiology and Surgery* 10.12 (Dec. 2015), pp. 1997–2007. DOI: 10.1007/s11548-015-1233-y (cit. on pp. 34, 36, 39, 41, 48–51, 53–56, 60, 134).

[257] Zhe Luo, Junfeng Cai, Su Wang, Qiang Zhao, T. M. Peters, and Lixu Gu. “Magnetic Navigation for Thoracic Aortic Stent-graft Deployment Using Ultrasound Image Guidance”. In: *IEEE Transactions on Biomedical Engineering* 60.3 (Mar. 2013), pp. 862–871. DOI: 10.1109/TBME.2012.2206388 (cit. on p. 20).

[258] D. Zikic, M. Baust, A. Kamen, and N. Navab. “A general preconditioning scheme for difference measures in deformable registration”. In: *2011 International Conference on Computer Vision*. IEEE. IEEE, Nov. 2011, pp. 49–56. DOI: 10.1109/ICCV.2011.6126224 (cit. on p. 47).

List of Figures

List of Tables