Computational ultrasound tissue characterisation for brain tumour resection

Rémi Delaunay

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Medical Physics and Biomedical Engineering

University College London

June 2, 2022

I, Rémi Delaunay, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

In brain tumour resection, it is vital to know where critical neurovascular structures and tumours are located to minimise surgical injuries and cancer recurrence. The aim of this thesis was to improve intraoperative guidance during brain tumour resection by integrating both ultrasound standard imaging and elastography in the surgical workflow. Brain tumour resection requires surgeons to identify the tumour boundaries to preserve healthy brain tissue and prevent cancer recurrence. This thesis proposes to use ultrasound elastography in combination with conventional ultrasound B-mode imaging to better characterise tumour tissue during surgery. Ultrasound elastography comprises a set of techniques that measure tissue stiffness, which is a known biomarker of brain tumours. The objectives of the research reported in this thesis are to implement novel learning-based methods for ultrasound elastography and to integrate them in an image-guided intervention framework. Accurate and real-time intraoperative estimation of tissue elasticity can guide towards better delineation of brain tumours and improve the outcome of neurosurgery. We first investigated current challenges in quasi-static elastography, which evaluates tissue deformation (strain) by estimating the displacement between successive ultrasound frames, acquired before and after applying manual compression. Recent approaches in ultrasound elastography have demonstrated that convolutional neural networks can capture ultrasound high-frequency content and produce accurate strain estimates. We proposed a new unsupervised deep learning method for strain prediction, where the training of the network is driven by a regularised cost function, composed of a similarity metric and a regularisation term that preserves displacement continuity by directly optimising the strain smoothness. We further improved the accuracy of

our method by proposing a recurrent network architecture with convolutional long-short-term memory decoder blocks to improve displacement estimation and spatio-temporal continuity between time series ultrasound frames. We then demonstrate initial results towards extending our ultrasound displacement estimation method to shear wave elastography, which provides a quantitative estimation of tissue stiffness. Furthermore, this thesis describes the development of an open-source image-guided intervention platform, specifically designed to combine intra-operative ultrasound imaging with a neuronavigation system and perform real-time ultrasound tissue characterisation. The integration was conducted using commercial hardware and validated on an anatomical phantom. Finally, preliminary results on the feasibility and safety of the use of a novel intraoperative ultrasound probe designed for pituitary surgery are presented. Prior to the clinical assessment of our image-guided platform, the ability of the ultrasound probe to be used alongside standard surgical equipment was demonstrated in 5 pituitary cases.

Impact Statement

This thesis describes the development of novel ultrasound tissue characterisation methods to enable the intraoperative characterisation of brain tumours. Moreover, this work focused on the clinical translation of the presented research by developing open-source tools for intraoperative guidance and assessing their clinical feasibility.

From an academic perspective, the developed learning-based methodologies have helped improve the state-of-the-art in speckle tracking for ultrasound elastography. We proposed the first unsupervised training strategy applied to quasi-static ultrasound imaging (chapter 3). We also introduced a new regularisation term that preserves the displacement continuity by directly optimising the strain smoothness, which was later used in other studies [1, 2]. Furthermore, we made publicly available an *in vivo* database for quasi-static strain imaging, which consisted of 17271 RF data of blood vessels from the arm of a human volunteer[1]. We also released a numerical ultrasound simulation dataset for shear wave velocity estimation, consisting of 20,000 images[2]. These open-access datasets could help the development and validation of new learning-based strategies in ultrasound elastography for both academia and the industry.

This thesis also presents open-source software packages for the translation of the developed concepts into clinics. This comprises an ultrasound-aided neuronavigation module to integrate ultrasound imaging with a commercial neuronavigation system[3]. We also proposed the first real-time AI-driven ultrasound elastography module

[1]Open access *in vivo* database available on `https://www.synapse.org/InVivoDataForUSE`
[2]Open access simulation database available on `https://www.synapse.org/SimForSWEI`
[3]Open-source platform available on `https://github.com/UCL/SkullBaseNavigation`

designed for intraoperative use[4]. This stand-alone module can be connected to a live stream of data and allows the use of a pre-trained convolutional neural network to perform real-time inference. Those software programs could potentially be used for the clinical translation of new advances in ultrasound-based methodologies.

This work also studied the clinical feasibility of the use of intraoperative ultrasound for brain tumour resection. This comprises the development of a novel ultrasound transducer designed for transsphenoidal surgery in collaboration with Vermon (Tours, France), a company specialised in the manufacture of advanced transducer technologies. A clinical trial was performed on an initial cohort of 5 patients who underwent pituitary resection surgery. This study brought additional clinical evidence that ultrasound imaging could be a safe and effective technological adjunct to transsphenoidal surgery. Since then, the first intraoperative ultrasound transducer for transsphenoidal surgery was commercialised (BK Medical, UK) [3], suggesting that our project was in line with the current clinical needs. This work could lead to further clinical investigations regarding the intraoperative use of ultrasound imaging for the characterisation of brain tumours.

[4]Open-source platform & fine-tuning code available on `https://github.com/RemiDelaunay/DeepUSE`

Acknowledgements

My research was supported by the EPSRC-funded CDT for Intelligent, Integrated Imaging In Healthcare (i4health), the Department of Health's NIHR-funded Biomedical Research Centre at University College London Hospitals (UCLH) and the School of Biomedical Engineering & Imaging Sciences at King's College London.

First, I would like to express my deepest gratitude to my PhD supervisors, Prof. Tom Vercauteren and Dr. Yipeng Hu for giving me the opportunity to work on this project and for their valuable guidance, help and support throughout my PhD.

Moreover, I wish to acknowledge with much appreciation Dr. Neil Dorward, Dr. Hani J. Marcus and Dr. Ivan Cabrilo for their clinical expertise and participation in the pituitary tumour clinical study. My sincere appreciation also goes to Dr. Jonathan Shapey for his help and involvement in the ultrasound-guided navigation project.

Furthermore, I would like to acknowledge my colleagues from the CAI4CAI team at King's College London and people from the UCL Wellcome / EPSRC Centre for Interventional and Surgical Sciences. A special thank to Nora for her continuous help and encouragements during this academic journey.

Finally, my appreciation goes out to my family and friends for their encouragement and support all through my studies.

Contents

Research contribution **19**

1 Introduction **21**

1.1 Ultrasound-guided neurosurgery 21

1.1.1 Brain tumour resection 21

1.1.2 Image-guided neurosurgery 22

1.1.3 Intraoperative ultrasound imaging 23

1.1.4 Ultrasound elastography 23

1.2 General aims and motivation 24

1.3 Thesis structure 25

2 Tissue modelling and imaging for ultrasound elastography **27**

2.1 Medical ultrasound 27

2.1.1 Ultrasonic physics 28

2.1.2 Pulse-echo imaging 30

2.1.3 Coherent plane wave imaging 33

2.2 Biological tissues and elasticity 34

2.2.1 Background 34

2.2.2 Theory of elasticity 36

2.3 Ultrasound elastography 40

2.3.1 Quasi-static elastography imaging 41

2.3.2 Transient impulse shear wave elastography 43

2.4 Conclusion 51

3 An unsupervised approach to quasi-static elastography with end-to-end strain regularisation 53

3.1 Introduction . . . 53

3.1.1 Previous work in displacement estimation . . . 53

3.1.2 Contributions . . . 55

3.2 Methods . . . 56

3.2.1 Problem statement . . . 56

3.2.2 Displacement Estimation . . . 57

3.2.3 Implementation . . . 59

3.3 Experiments . . . 62

3.3.1 Simulation dataset . . . 62

3.3.2 *In vivo* dataset . . . 63

3.3.3 Performance comparison . . . 64

3.3.4 Results . . . 65

3.4 Discussion . . . 66

3.5 Conclusion . . . 68

4 Improving spatio-temporal consistency of strain estimates using recurrent neural networks 70

4.1 Introduction . . . 70

4.1.1 Background . . . 70

4.1.2 Contributions . . . 71

4.2 Methods . . . 72

4.2.1 Network architecture . . . 72

4.2.2 Training . . . 75

4.3 Experiments . . . 77

4.3.1 Experiments on numerical phantoms . . . 77

4.3.2 Experiments on *in vivo* human data . . . 79

4.4 Discussion . . . 82

4.5 Conclusion . . . 86

5 Learning-based methods for displacement estimation in shear wave elastography 87

5.1 Introduction . . . 87
5.1.1 Previous work in ARFI displacement estimation . . . 88
5.1.2 Contributions . . . 89
5.2 Methods . . . 89
5.2.1 Generating the numerical simulation dataset . . . 89
5.2.2 Training strategy . . . 91
5.2.3 Time-of-flight approach . . . 92
5.3 Experiments . . . 93
5.4 Discussion . . . 94
5.5 Conclusion . . . 97

6 Integration of intraoperative ultrasound with neuronavigation 98

6.1 Introduction . . . 99
6.1.1 Navigation in neurosurgery . . . 99
6.1.2 Ultrasound-guided surgery and related work on image-guided intervention software . . . 101
6.1.3 Contributions . . . 103
6.2 System design . . . 103
6.2.1 Hardware . . . 103
6.2.2 Software . . . 104
6.2.3 Neuronavigation Slicelet . . . 105
6.2.4 Stand-alone ultrasound elastography module . . . 107
6.2.5 Calibration . . . 108
6.2.6 3D ultrasound volume reconstruction . . . 109
6.3 Accuracy and workflow testing . . . 110
6.3.1 Laboratory testing of system accuracy . . . 110
6.3.2 Clinical simulation to test workflow integration . . . 110
6.4 Discussion . . . 113
6.5 Conclusion . . . 115

7 Evaluation of an intraoperative ultrasound transducer prototype for pituitary tumour resection **117**

7.1 Introduction . . . 118

7.1.1 Pituitary adenomas . . . 118

7.1.2 Transsphenoidal surgery . . . 118

7.1.3 Previous work on ultrasound-aided pituitary surgery . . . 120

7.1.4 Contributions . . . 122

7.2 Ultrasound transducer assessment . . . 122

7.2.1 Ultrasound probe specifications . . . 122

7.2.2 Image resolution assessment . . . 123

7.2.3 Ergonomics study . . . 126

7.3 First-in-human clinical study . . . 127

7.3.1 Study design . . . 127

7.3.2 Intervention . . . 128

7.3.3 Measures of outcomes . . . 128

7.3.4 Results . . . 129

7.4 Discussion . . . 130

7.5 Conclusion . . . 133

8 General conclusion and future perspectives **135**

8.1 General conclusion . . . 135

8.2 Limitations and future perspectives . . . 137

8.2.1 Learning-based methods for brain tissue characterisation . . . 137

8.2.2 Integration of intraoperative ultrasound in the surgical workflow . . . 138

List of Figures

2.1 Ultrasound wave propagation 29
2.2 Transmit and receive modes in pulse-echo imaging 31
2.3 Ultrasound display modes 32
2.4 B-mode image formation 33
2.5 Comparison between pulse-echo and plane wave imaging 35
2.6 Different deformation models and elastic moduli 40
2.7 Summary and classification of ultrasound elastography techniques . 41
2.8 Quasi-static elastography example. B-mode image (a) and strain elastogram (b) of a convexity meningioma acquired during brain tumour resection 43
2.9 Transient elastography with FibroScan^{TM} (Echosens, Paris, France) 47
2.10 Example of point shear wave elastography 49
2.11 Supersonic shear imaging 50

3.1 Overview of our unsupervised method 57
3.2 Network architecture of USENet 60
3.3 Finite element analysis pipeline 63
3.4 *In vivo* image acquisition pipeline. 64
3.5 Comparison of three axial strain fields computed from ultrasound simulations 66
3.6 Comparison of axial strains estimated from the *in vivo* testing dataset 67

4.1 Overview of the proposed recurrent model for ultrasound strain elastography (ReUSENet) 73

4.2 Network architecture of ReUSENet 74
4.3 Strain images of two numerical phantoms 79
4.4 Similarity, consistency, SNRe and NRMSE scores 80
4.5 Strain images of a numerical phantom temporal sequence with increasing deformation 81
4.6 Temporal strain image sequence from the testing *in vivo* dataset. . . 82
4.7 Similarity, consistency and SNRe scores for the *in vivo* testing dataset 84
4.8 Temporal strain image sequence from the testing *in vivo* dataset. . . 85

5.1 Shear wave displacement estimated in a heterogeneous phantom . . 95
5.2 Young's modulus elasticity maps estimated from ground truth displacement, NCC, USENet and ReUSENet 96

6.1 Picture of the Leksell stereotactic system. 100
6.2 Summary of the neurosurgical workflow for brain tumour resection . 101
6.3 Hardware and Software communication using IP/TCP network. . . . 106
6.4 Neuronavigation module user interface. 107
6.5 DeepUSE : Slicer ultrasound elastography module interface. 108
6.6 Neuronavigation module user interface. 109
6.7 Surgical workflow 111
6.8 Ultrasound navigation and reconstruction system accuracy 112
6.9 Intraoperative simulation of 3D Ultrasound reconstruction using a patient data and a patient-specific phantom model. 114

7.1 Anatomy of the pituitary gland. 119
7.2 Evolution of the transsphenoidal probe prototypes. 124
7.3 Photograph showing the novel forward-viewing ultrasound probe. . 125
7.4 Spatial resolution evaluation method. 125
7.5 Ergonomics study. 126
7.6 Photograph of the surgical field (1) 131
7.7 Photographs of the surgical field (2) 132
7.8 Intraoperative images of the clinical study (1) 133

7.9 Intraoperative images of the clinical study (2) 134

8.1 BK N20P6 ultrasound transducer 137

List of Tables

3.1 Mean and standard deviation of SNRe and CNR for the strain images of the simulation testing dataset obtained with finite element methods, OVERWIND, the unsupervised and supervised network. . 65

3.2 SNRe and CNR of strain estimates from our unsupervised method and OVERWIND, and local NCC scores for the three image pair results from the *in vivo* dataset. 66

4.1 Registration accuracy of USENet, GLUE and ReUSENet 83

5.1 Mean and standard deviation of SNR, CNR and Young's modulus MAE for the elasticity maps estimated from the simulation testing dataset 94

List of Abbreviations

ARFI Acoustic Radiation Force Imaging

BPTT Backpropagation Through Time

CNN Convolutional Neural Network

CNR Contrast-to-Noise Ratio

CT Computed Tomography

DDF Dense Displacement Field

FEM Finite Element Method

fps frames per second

I/Q In-phase/Quadrature signal

LSQSE Least-Squares Strain Estimator

MAE Mean Absolute Error

MRE Magnetic Resonance Elastography

MRI Magnetic Resonance Imaging

pSWE point Shear Wave Elastography

ReUSENet Recurrent Ultrasound Elastography Network

RMSE Root-Mean Square Error

SNR Signal-to-Noise Ratio

SWEI Shear Wave Elastography Imaging

TBTT Truncated Backpropagation Through Time

TE Transient Elastography

TTP Time-To-Peak

US Ultrasound Imaging

USE Ultrasound Elastography

USENet Ultrasound Elastography Network

List of Symbols

Z	acoustic impedance
ρ	tissue density
c	speed of sound
α	absorption coefficient
σ	stress
ε	strain
τ	time delay
G	shear modulus
K	bulk modulus
E	Young's modulus
ν	Poisson's ratio
λ	Lamé's first parameter
μ	Lamé's second parameter
C_L	longitudinal wave speed
C_s	shear wave speed
∇	nabla operator
u	displacement
F	radiation force
t	time

Research contribution

Journal publications

- **Rémi Delaunay**, Yipeng Hu, and Tom Vercauteren. An unsupervised learning approach to ultrasound strain elastography with spatio-temporal consistency. *Physics in Medicine and Biology*, 66(17), 2021.

- Ivan Cabrilo*, **Rémi Delaunay***, Clare L Heaysman, Sebastien Ourselin, Valentina Vitiello, Tom Vercauteren, Hani J Marcus, and Neil L Dorward. A novel intraoperative ultrasound probe for transsphenoidal surgery: First-in-human study. *Surgical Innovation*, page 15533506211031091, 2021.

Conference publications

- **Rémi Delaunay**, Yipeng Hu, and Tom Vercauteren. An unsupervised approach to ultrasound elastography with end-to-end strain regularisation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 573–582. Springer, 2020.

- **Rémi Delaunay**, Yipeng Hu, and Tom Vercauteren. An unsupervised learning-based shear wave tracking method for ultrasound elastography. In *SPIE in Medical Imaging 2022*. International Society for Optics and Photonics, 2022.

Co-authored publications

- Senay Mihcin, Hani J Marcus, **Rémi Delaunay**, Efthymios Maneas, Wenfeng Xia, Sebastien Ourselin, Adrien Desjardin, Neil Dorward, and Tom Vercauteren. Preclinical assessment of a novel intraoperative ultrasound probe

for transsphenoidal surgery. In *Computer Assisted Radiology and Surgery Proceedings of the 32nd International Congress and Exhibition*, pages 57–59. IJCARS, 2018.

- Yunguan Fu, Nina Montaña Brown, Shaheer U. Saeed, Adrià Casamitjana, Zachary M. C. Baum, **Rémi Delaunay**, Qianye Yang, Alexander Grimwood, Zhe Min, Stefano B. Blumberg, Juan Eugenio Iglesias, Dean C. Barratt, Ester Bonmati, Daniel C. Alexander, Matthew J. Clarkson, Tom Vercauteren, and Yipeng Hu. Deepreg: a deep learning toolkit for medical image registration. *Journal of Open Source Software*, 5(55):2705, 2020.

- N Brown, Yunguan Fu, Shaheer U Saeed, Adrià Casamitjana, Zachary MC Baum, **Rémi Delaunay**, Qianye Yang, Alexander Grimwood, Zhe Min, Ester Bonmati, Tom Vercauteren, Matthew Clarkson, and Yipeng Hu. Introduction to medical image registration with deepreg, between old and new. *arXiv preprint arXiv:2009.01924*, 2020.

- Jonathan Shapey, Thomas Dowrick, **Rémi Delaunay**, Eleanor C Mackle, Stephen Thompson, Mirek Janatka, Roland Guichard, Anastasis Georgoulas, David Pérez-Suárez, Robert Bradford, Shakeel Saeed, Sébastien Ourselin, Matthew Clarckson, and Tom Vercauteren. Integrated multi-modality image-guided navigation for neurosurgery: open-source software platform using state-of-the-art clinical hardware. *International Journal of Computer Assisted Radiology and Surgery*, pages 1–10, 2021.

* Equal contribution

Chapter 1

Introduction

1.1 Ultrasound-guided neurosurgery

1.1.1 Brain tumour resection

A brain tumour corresponds to an abnormal growth of cells located in the brain. The tumour can be malignant (as opposed to benign), which means it can metastasise and spread to other parts of the body. It can also be classified using a grading system from 1 to 4 depending on their growth rate and malignancy. Low-grade tumours (grade 1 or 2) tend to grow slowly and are less likely to spread to other regions. High-grade cancers (grade 3 or 4) are highly differentiated (i.e., different from normal cells) and usually spread more rapidly than low-grade tumours. Brain tumours can arise from the brain (i.e., primary tumours) or can spread from elsewhere in the body (i.e., secondary tumours). The National Health Service (NHS) estimates that about 11,000 people are diagnosed with a primary brain tumour every year in the United Kingdom (UK) [4]. The symptoms can vary and depend on the tumour location. They include headaches, nausea, hormonal imbalance, vision loss, speech disorder, or partial paralysis. The treatment options depend on the size, location, and type of the tumour and usually involves medical therapies, radiotherapy, or neurosurgery. Surgical resection is one of the main treatment options for brain tumours. The aim of tumour resection is to maximise tumour tissue removal, while preserving surrounding structures. The resection extent is an important prognostic indicator [5]. Incomplete resection can result in tumour recurrence and require further treatment.

Complete removal of the tumour can be challenging because margins of the tumour are not always clearly delineated. High-grade tumours can spread to nearby normal parts of the brain and cause high recurrence rates, while low-grade tumour cells can be difficult to distinguish from healthy ones. In addition, the tumour can be located close to critical neurovascular structures, which can further decrease the level of resection to avoid further neurological damage [6].

1.1.2 Image-guided neurosurgery

The introduction of computed tomography (CT) in the 1970s, and magnetic resonance imaging (MRI) in the 1990s, has led to considerable progress in the visualisation and localisation of brain tumours [7]. The development of frameless stereotaxy, also known as neuronavigation, arose in the late 1980s and enabled to further improve surgical guidance by mapping imaging data to the patient's anatomy [8]. Neuronavigation is now extensively used for surgery planning and intraoperative guidance [9]. In addition, the use of neuronavigation allows for smaller, more precisely positioned incisions and the accurate localisation of tumours and surrounding structural and functional regions which may be at risk during surgery [10].

However, intraoperative guidance based on preoperative imaging has significant limitations. Intraoperative tissue shift usually occurs, resulting in the displacement of anatomical structures, which may not be in the neurosurgeon's field of view [11]. Such brain shift typically arises from a change of intracranial pressure induced by craniotomy [12]. Navigation accuracy will also decrease during the course of the operation because of the brain shift induced by the removal of tumour tissue. These displacements can be inconsistent and hard to predict for neurosurgeons who can no longer rely on the preoperative images for guidance [13].

Intraoperative brain scanning has been used to provide a real-time localisation of brain structures during surgery. CT scanning provides a three-dimensional display of the brain, and can be easily integrated with neuronavigation systems [14, 15]. However, the ionising radiation doses affecting the patient and surgical team, and its low contrast in soft tissue, have limited its use in the operating room. Intraoperative MRI generally offers better contrast and resolution, allowing accurate visualisation

of brain structures, as well as the level of tumour resection. However, the significant cost attributed to a dedicated intraoperative MRI suite still makes it difficult to be routinely used in neurosurgery [16].

1.1.3 Intraoperative ultrasound imaging

The interest in ultrasound imaging for neurosurgical applications arose in the 1980s [17, 18] and has been constantly evolving since then [19, 20]. It has been used in many interventional settings and applications, mostly for surgical guidance and tissue characterisation [21, 22].

Ultrasound imaging has also shown potential for surgical guidance in tumour resection, especially for assessing the extent of resection [23–25]. The portability, cost-efficiency, and real-time imaging acquisition make ultrasound imaging a good alternative for the intraoperative assessment of brain tumours. Furthermore, several studies have emphasised that standard B-mode imaging allows for a tumour delineation similar or even more accurate than MRI and CT [26–29].

However, intraoperative ultrasound remains an under-utilised tool in neurosurgery. Intraoperative ultrasound remains challenging to use routinely for inexperienced users, because of its limited field-of-view and non-intuitive image orientation. The development of ultrasound-based navigation improved the readability of ultrasound data by providing information on the position and orientation of the imaging plane [19]. In addition, advances in navigation technology have enabled to reconstruct 3D ultrasound volumes and perform multimodal registration with preoperative scans such as MRI and CT [30]. The real-time overlay of ultrasound on preoperative images has proven to be useful during brain tumour resection by compensating for tissue shift [31].

1.1.4 Ultrasound elastography

Another important consideration for the delineation of brain tumours with ultrasound imaging is that tumour boundaries can sometimes not appear clearly on standard pulse-echo imaging. Surgeons do not only use image guidance to assess the tumour margin, but can also rely on the tissue mechanical properties, where tumours tend

to be stiffer than surrounding structures. Magnetic resonance elastography (MRE) is a technology that combines MRI with low-frequency vibrations to infer the mechanical properties of soft tissues. This technique has been used successfully for the preoperative assessment of tissue stiffness in brain tumour surgery [32]. However, the inherent limitations of MRI for intraoperative use (i.e., cost and prolonged operating time) also apply to MRE.

Ultrasound elastography is a real-time alternative that is derived from conventional ultrasound imaging and motion estimation techniques. Ultrasound elastography constitutes a set of techniques that measure and image tissue mechanical properties by applying a force on the tissue and monitor its following deformation [33]. It has been used clinically for the characterisation of lesions in different organs, such as the liver [34] or prostate [35], but also for the differentiation between benign and malignant tumours, such as those found in the thyroid [36] and breast [37]. Ultrasound elastography has also shown promising results in image-guided interventions, such as liver resection [38, 39]. Preliminary results suggest that the use of ultrasound elastography in combination with conventional B-mode ultrasound could be a useful adjunct to better delineate brain tumours [40, 41].

1.2 General aims and motivation

In brain tumour resection, it is vital to know where the tumour and critical neurovascular structures are located to minimise surgical injuries and cancer recurrence. The aim of this thesis is to improve intraoperative guidance during brain tumour resection by integrating both ultrasound standard imaging and elastography in the surgical workflow.

Recent approaches in ultrasound elastography have demonstrated that convolutional neural networks can capture ultrasound high-frequency content and produce an accurate measure of tissue stiffness, which is a known biomarker of brain tumours and could also potentially differentiate vital structures around the tumour site. Intraoperative characterisation of tumour tissue and surrounding structures could potentially help surgeons to identify tumour remnants and, therefore, improve the

extent of resection and minimise neurovascular injuries.

This thesis also proposes open-source tools to combine neuronavigation with intraoperative ultrasound imaging. We hypothesised that real-time ultrasound-based tissue characterisation could also benefit from preoperatively-acquired MRI and CT images, and that their combination in a fully integrated navigation system would improve neurosurgical guidance. To this end, we developed a research image-guided intervention platform, based on existing open-source software, designed for ultrasound-guided neuronavigation.

Finally, this project also studied the usefulness of intraoperative ultrasound in transsphenoidal surgery, a minimally invasive approach for the resection of pituitary tumours. This technique consists of removing the tumour by passing through the nasal cavity and the sphenoid sinus. At the time this study was initiated, there was no commercial transducer suitable for accessing the tumour site. To this end, we developed a novel ultrasound transducer which has been specifically designed for the transsphenoidal approach. Pre-clinical experiments and preliminary results of a clinical study on the feasibility and risk assessment of the use of this novel transducer are presented.

1.3 Thesis structure

This work is organised in 7 different chapters focusing on different aspects of this project. Chapter 2 gives an introduction to medical ultrasound and ultrasound elasticity imaging, including a review of the different ultrasound elastography techniques considered for this work. Chapter 3 focuses on the development of a novel learning-based method in ultrasound elastography. This technique is based on quasi-static elastography, a well-known ultrasound imaging modality that determines elastic tissue parameters and could potentially help differentiate brain tumour tissue. Chapter 4 is an extension of the work carried out in chapter 3 and focus on the development of a recurrent neural network architecture to improve spatio-temporal consistency of quasi-static strain estimation. Chapter 5 further investigates the use of convolutional neural networks for microscale displacement estimation in shear wave elastography,

by making use of the previous work accomplished in chapter 3 and 4. Chapter 6 contains the work carried out for the development of our prototype neuronavigation framework. Chapter 7 describes the technical challenges of pituitary surgery and the technical evaluation of a novel ultrasound transducer, specifically designed for pituitary tumour resection. We also present the preliminary results of a clinical study that aimed to assess the risk and feasibility of the intraoperative ultrasound transducer prototype. Finally, chapter 8 contains a general conclusion and provides future perspectives on the work carried out throughout this thesis.

Chapter 2

Tissue modelling and imaging for ultrasound elastography

2.1 Medical ultrasound

The first technological development of ultrasound imaging occurred during World War 1, with the use of the first sonar systems for under-water navigation. The research interest for the use of sonography in diagnostic imaging arose in the late 1940s [42]. However, it was not until the 1970s that diagnostic ultrasound imaging was used in medical practice. Nowadays, ultrasound is one of the most widely used medical imaging modalities. In 2011, it was estimated that nearly 25% of all medical imaging procedures involved ultrasound imaging [43].

The British Medical Ultrasound Society reported in 2010 that they considered diagnostic ultrasound to be safe, with no evidence of harmful effect to humans (including the fetus) [44]. Yet, the effect of ultrasound on biological tissue is still a topic of research [45]. Wood and Loomis were the first to report the detrimental effects of ultrasonic energy upon biological material, through experiments on small animals and fish [46]. In addition, recent studies have investigated the potential harmful effect of cranial ultrasound on the neural integrity and development of neonates [47, 48]. The two observed bio-effects of ultrasound imaging are heating and cavitation. Tissue heating results from the mechanical friction induced by the ultrasonic waves. Cavitation is the formation of small bubbles in response to the

change of pressure induced by the acoustic field. Presently, state-of-the-art ultrasound scanners display the thermal and mechanical index to provide the operator with an indication of the potential ultrasound bio-effects. The thermal index corresponds to the ratio of acoustic power to the power required to raise the tissue temperature by 1 °C. The mechanical index provides an estimation of the maximum amplitude of the pressure pulse in tissue, which is an indicator of the likelihood of cavitation [49].

Ultrasound imaging is inexpensive relative to MR and CT, portable and generally provides real-time visualisation. In addition, it does not produce any ionising radiation, unlike fluoroscopic and nuclear medicine procedures. It is used in a wide variety of diagnostic and therapeutic applications, including prenatal diagnosis, internal organ examination, blood flow measurements and tumour ablation. Last but not least, it is also used as a guiding tool for interventional procedures and often is an adjunct to other imaging modalities, such as X-rays and MRI.

2.1.1 Ultrasonic physics

The term 'ultrasound' refers, in physics, to all acoustic energy that oscillates at a frequency above the human auditory band (ranged between 20 Hz to 20 kHz). Ultrasound scanners typically use a frequency range of about 2 to 20 MHz. Ultrasound transducers are generally composed of an array of piezoelectric crystals, which convert the electrical voltage into mechanical vibrations (transmission mode) and vice versa (receiving mode). The ultrasound waves propagate through tissue along the particle motion (longitudinal waves) and oscillate successively between zones of compression and rarefaction [50]. An example of longitudinal wave propagation from an ultrasound transducer is shown in figure 2.1.

The way the ultrasound waves are transmitted will depend on certain tissue physical properties. The acoustic impedance Z describes how much resistance a longitudinal wave will encounter as it goes through a tissue, and is defined as the product of the tissue density ρ and speed of sound c:

$$Z = \rho c \tag{2.1}$$

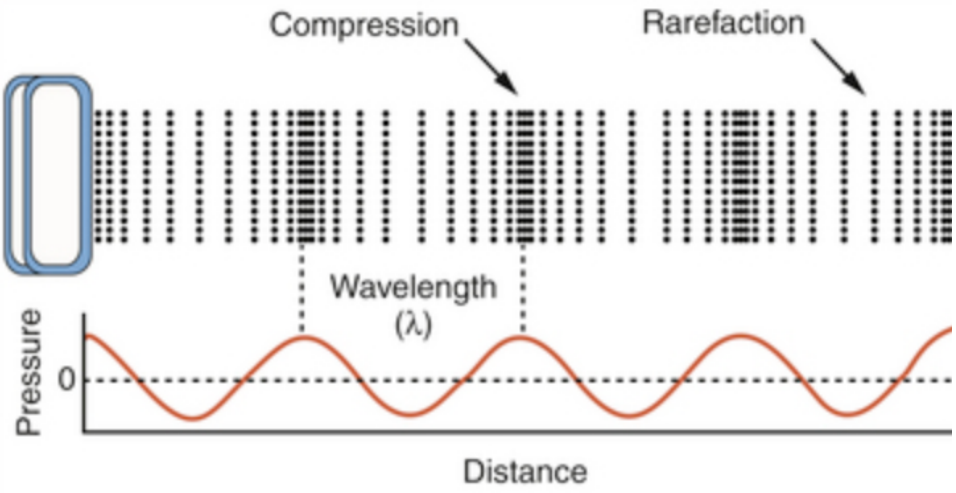

Figure 2.1: Ultrasound wave propagation. The ultrasound waves emitted from the transducer propagates in the longitudinal direction with successive zones of compression and rarefaction characterised by the wavelength λ. The fluctuations of the propagating wave occur with local negative/positive pressure changes. Courtesy of Zagzebski et al. [50]

At the interface of two tissue media with different acoustic impedance, part of the ultrasound wave will be reflected back to its source, while the remaining part will be transmitted through the second medium [51]. The difference of propagation speed between two media will result in a change of direction of the ultrasound waves, which corresponds to the refraction phenomenon. Ultrasonographers often use a water-like gel to to reduce the impedance difference at the air-tissue boundary and improve the transfer of acoustic energy from the transducer face to the body.

Scattering is another important aspect of ultrasound imaging and consists of the reflection of ultrasonic waves caused by localised heterogeneities in the tissue through which they pass (small molecules, cells, etc.). The scatterers that reflect back a portion of the waves are known as 'backscatters' and constitute part of the received ultrasonic signal. The insonification of the tissue also causes constructive and destructive interference between the scatterers, which creates the speckle pattern that gives the characteristic granular texture of ultrasound images [52].

As an ultrasound beam passes through tissue, its amplitude will be progressively attenuated. This is partly due to scattering, but most of the vibrational energy is lost in the form of heat, also known as the absorption phenomenon. Acoustic absorption

in tissue is mainly due to viscous loss (i.e., tissue resistance) and the conversion of the wave energy into internal molecular energy states. The attenuation of ultrasound in soft biological tissues is about 0.2 to 0.5 dB/cm^2/MHz [53]. The absorption coefficient α depends on the frequency of the transmitted ultrasonic waves, which is given by the following power law:

$$\alpha \approx \alpha_0 f^y \tag{2.2}$$

where y denotes the frequency dependence. It is unitless and varies according to the tissue viscosity and is usually comprised between 1 (e.g., muscle, heart and fat) and 2 (e.g., water) [54, 55]. The attenuation increases with frequency, shortening the depth of penetration of ultrasonic waves. Therefore, there is a trade-off between the image resolution, which is proportional to the transducer's central frequency, and the maximum image depth. Time-gain compensation is a way to partially compensate for signal attenuation in tissue by increasing the amplification gain according to the reflected echoes' depth.

2.1.2 Pulse-echo imaging

Pulse-echo imaging constitutes the standard medical ultrasound acquisition method for image diagnosis [56]. An ultrasound transducer (generally composed of multiple piezoelectric crystals) is used to generate a directional beam of short ultrasound pulses that are used to insonify the targeted tissue. Each transducer element can be excited individually at different times, which allows steering or focus the ultrasound beam via electronic manipulation (transmit mode). Those pulses are echoed back from the scatterers to the transducer that also acts as the receiver. The received signal is then used to form the image, which is also known as receive beamforming. Similar to the transmit mode, the signal at each pixel can be calculated by applying the time delays corresponding to this point in the image space [57]. The time delay τ corresponds to the time taken for the pulse to reach a specific target in depth and to be reflected back to the transducer. The target depth d can be easily retrieved by assuming a constant speed of sound throughout the medium, which

is assumed to be 1540 m/s in soft biological tissues. Therefore, the depth can be calculated as $d = c\tau/2$. Figure 2.2 illustrates the process of how the ultrasound signal is traditionally transmitted and received in pulse-echo imaging, also known as delay-and-sum beamforming [58].

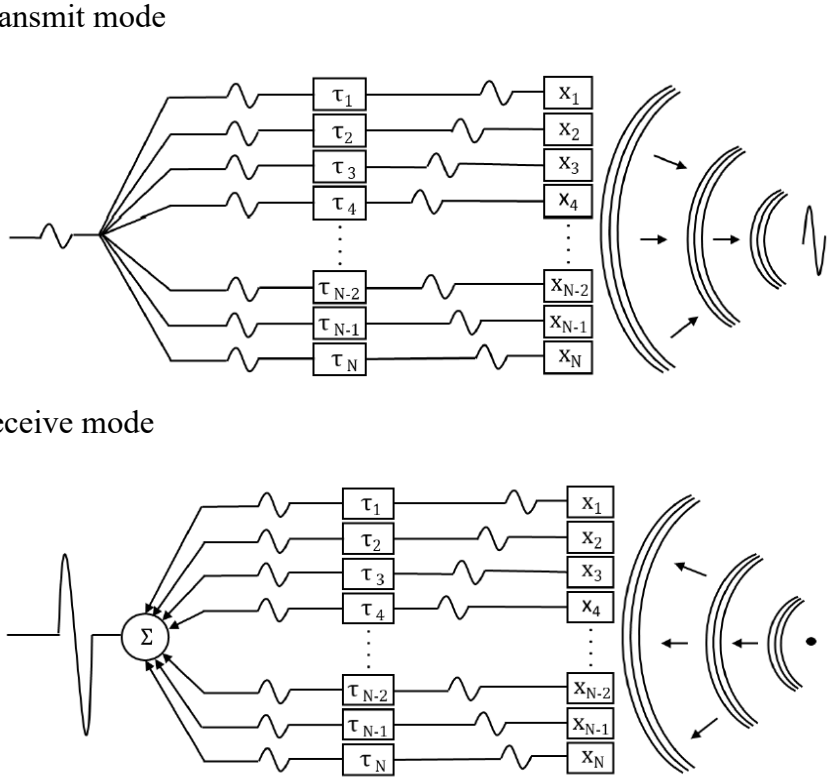

Figure 2.2: Transmit and receive modes in pulse-echo imaging. A typical transducer is composed of multiples elements X_N for which time-delays τ_N are applied to transmit and focus an ultrasound beam at a specific target depth. During receive beamforming, the acoustic reflections are received at each channel and are time aligned then summed to reconstruct the ultrasonic signal. Adapted from Engel et al. [58]

Different display modes have been developed for pulse-echo imaging. The three most common display modes are the Amplitude (A), Motion (M) and Brightness (B) modes [59]. An example of each display mode is showed in figure 2.3. The A-mode displays the echo amplitude of a single transducer scan line as a function of depth. The M-mode corresponds also to a single scan line, except that the transducer emits pulses multiple times, which allows tracking of anatomical structure motion over

time, such as cardiac pulsations. The B-mode is the most commonly used modality in clinical practice and consists in the combination of multiple A-lines to form a 2D grey-scale image of the imaged anatomy, where the brightness intensity corresponds to the echo amplitude.

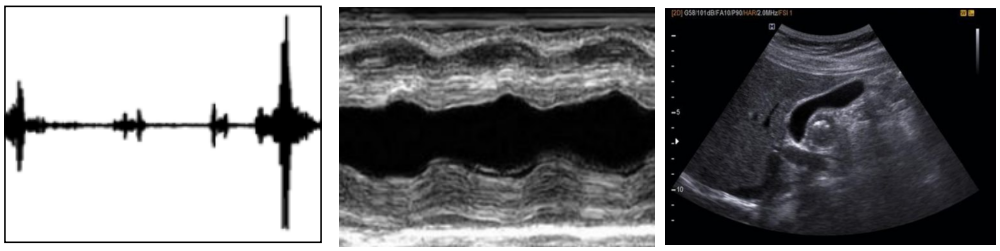

A-mode M-mode B-mode

Figure 2.3: Typical ultrasound display modes: amplitude (A) mode, motion (M) mode and brightness (B) mode.

The most common type of transducer used in clinical practice are multi-elements 1-D array transducers [56]. Each element can be fired independently, allowing the electronic focusing or steering of the transmitted ultrasound beam. Similarly, the focus of the received signal can be dynamically adjusted by applying the appropriate time delays. The reconstruction of the 2D image consists of summing and adjusting in time individual scan lines generated by sequentially firing subgroups of elements over the lateral plane. To improve image resolution, multiple transmit pulses can be used, each focused at a different imaging depth, and the resulting scan lines later combined. However, this also reduces the imaging frame rate, as a larger number of transmit-receive events are needed to form the image. The ultrasound transducers commonly used are linear and phased array transducers, which present different aperture shape and size. Linear arrays generally have a broad footprint and are usually used for diagnostic applications that have a large acoustic window. Phased array transducers feature a smaller aperture size, commonly with a convex shape, which creates a trapezoidal field-of-view [60].

After acquiring the raw ultrasound signal, different post-processing steps are necessary to obtain B-mode images. The first step aims to find the envelope of the raw

beamformed RF signal (Envelope detection). It allows extracting of the magnitude of the RF signal and filter the high frequency content. This is classically done using the Hilbert transformation. Finally, the envelope-detected signal is log-compressed to improve the image dynamic range (Log compression), and interpolated from polar to Cartesian space to improve the image display (Scan conversion) [61]. Figure 2.4 displays the result of the aforementioned processing steps.

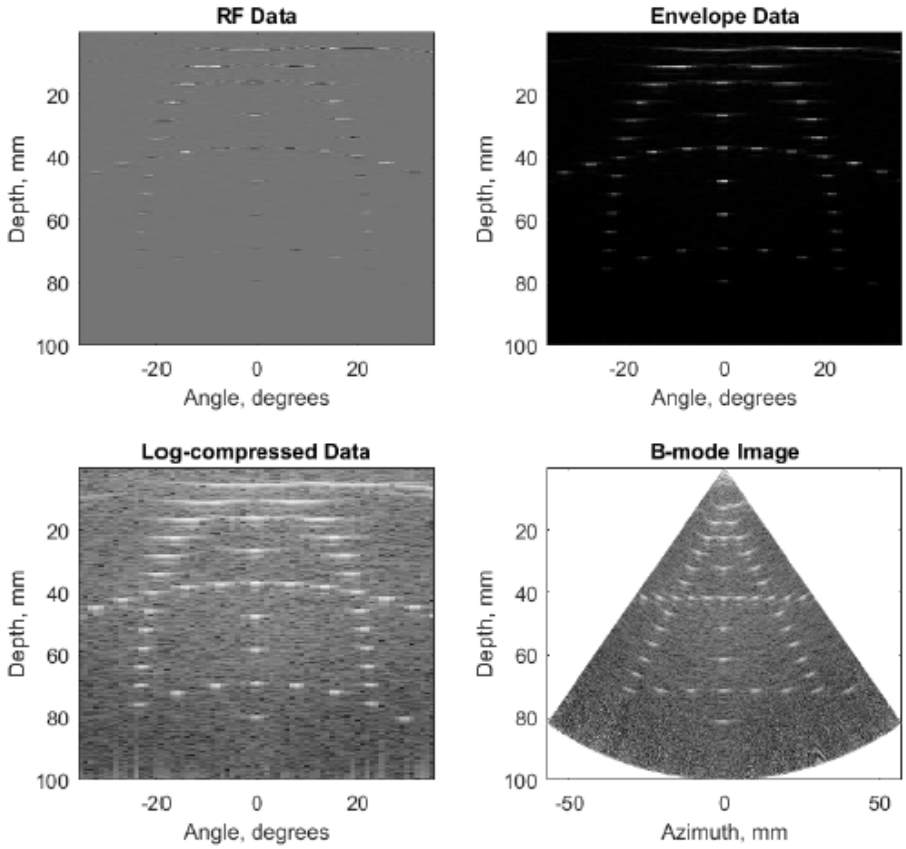

Figure 2.4: B-mode image formation from the beamformed radio-frequency signal. Image courtesy of Natalia Ilyina [59].

2.1.3 Coherent plane wave imaging

Plane-wave imaging is a different acquisition method that allows to image a 2D region of interest at a much faster rate, up to 100-fold faster than traditional B-mode imaging. Instead of sequentially scanning the image line-by-line, this technique fires the elements from the array all at once, creating an unfocused transmitting plane

wave over the entire field of view. This technique, also called 'Ultrafast imaging' was developed originally for shear wave elasticity imaging [62]. Although the plane wave imaging concept is not recent, its application was only made possible recently with the increase of computational power of current systems via the technological developments in graphical processing unit (GPU) [63].

However, the increased frame-rate due to the generation of unfocused plane waves results in a substantial loss in image resolution and contrast, as compared with traditional B-mode imaging that uses dynamic focusing. To improve the image quality, Montaldo et al. developed the coherent plane wave compounding method, which compensates for focused emission by taking several plane wave images at different angles [64]. The image acquisition process consists in acquiring plane wave frames at different angles (usually comprised between -8° to 8°, with a step of 1°) and compound the resulting signals to construct a high-quality image. Figure 2.5 illustrates the different modes of acquisition of plane wave and pulse-echo imaging.

The coherent plane wave compounding method still largely increases the imaging rate as compared to B-mode imaging, with frame rates up to 10000 frames per second. The increased frame rate enables to perform measurements of tissue and blood motion while improving the image quality [63].

2.2 Biological tissues and elasticity

2.2.1 Background

The change of tissue stiffness has been a known biomarker of several disease processes and is still used nowadays as a diagnostic tool. Palpation is the physical examination of tissue stiffness, which consists in applying pressure of the hand or fingers to the surface of the body to assess the size or consistency of an underlying organ. Palpation has been used for the diagnosis of a large variety of disease processes, from the presence of inflammation to the detection of fibrosis or tumours.

The study of tissue mechanical properties as a diagnostic tool can be dated to ancient Egyptian time with the Ebers papyrus from 1550 BC, where the palpation of pulses was already described. Palpation of the radial arteries for pulse examination

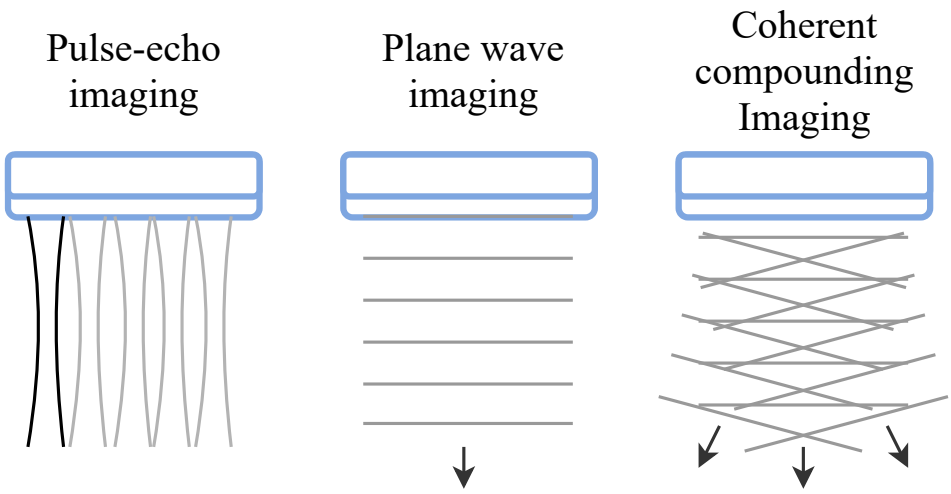

Figure 2.5: Comparison between pulse-echo and plane wave imaging. Pulse-echo imaging reconstructs the image line-by-line by sequentially firing groups of elements to create successive focused ultrasound beams. Plane-wave imaging generates unfocused transmitting plane wave over the entire image region. The coherent compounding method generates several plane wave images at different angles to compensate for unfocused emission, and reconstructs images with the same quality as the conventional line-by-line pulse-echo imaging method.

was also described in traditional Chinese medicine around 500 BC. In 440 BC, Hippocrates, the father of western medicine, also identified the relationship between tissue stiffness and different pathologies [65].

The alteration of tissue elasticity is known to appear in diseased tissue (e.g., brain tumours) by the exudation of fluid from the blood vessels to the extra- and intra-cellular space, which results in an increased interstitial pressure at the disease site, inducing a change of tissue stiffness. In the case of fibrosis, the diseased tissue is replaced with a stiffer collagen-based matrix.

Although palpation is very useful to assess the stiffness of superficial anatomical structures, such as the radial carotid arteries and the thyroid, it is limited by tissue depth. In addition, palpation is highly-user dependent and its subjective nature makes it prone to human errors. In addition, the difficulty to detect small alterations makes it challenging to provide an early diagnosis of disease formation.

This context has motivated the research of alternative techniques to quantitatively measure tissue mechanical properties. The emergence of noninvasive diag-

nostic imaging tools has inspired new tissue characterisation techniques, such as MR- and ultrasound elastography. In contrast with palpation, these tools present the advantage of providing a quantitative measurement of tissue stiffness and are less dependent on the size and location in depth of the targeted tissue.

2.2.2 Theory of elasticity

The theory of elasticity treats the relationship between forces (i.e., stress) applied to a material and the resulting deformation (i.e., strain). For linear materials, the relationship between stress and strain can be expressed by the generalised Hooke's law:

$$\sigma = C_{ijkl} \cdot \varepsilon_{kl} \tag{2.3}$$

where σ and ε are the stress and strain, respectively, and the indexes j,l indicate the direction of force (or strain) and i,k denote the plane on which it is acting. C is the fourth-order stiffness tensor, and is composed of 81 elastic constants, which account for three directions of deformation in three different planes.

Different types of elastic moduli are defined according to which direction stress and strain are to be measured (see figure 2.6). The three main types of elastic moduli are: the Young's modulus, shear modulus, and bulk modulus. The Young's modulus E describes how an object will be deformed along an axis where the stress, either a tensile or compressive force, is applied longitudinally. It corresponds to the coefficient of proportionality between the applied longitudinal stress σ_L and the longitudinal strain $\varepsilon_L = \Delta L/L$:

$$\sigma_L = E \cdot \varepsilon_L \tag{2.4}$$

The shear modulus G is defined for shear deformations, where an object will deform when acted upon by opposing forces. It is defined as the proportionality coefficient between the shear stress σ_s and strain ε_s:

$$\sigma_s = G \cdot \varepsilon_s \tag{2.5}$$

The bulk modulus K is the extension of the Young's modulus in three dimensions, and is used to describe the volume changes in all directions under a uniformly loaded pressure. It is defined as the coefficient between the volumetric stress σ_v and strain ε_v:

$$\sigma_v = K \cdot \varepsilon_v \tag{2.6}$$

For incompressible material, no volume change is observed. The high content of water in soft tissue makes it nearly incompressible and isotropic. For example, a cylindrical object undergoing a tensile load will be stretched in length, and vice-versa. The ratio between the longitudinal strain ε_L and radial strain $\varepsilon_r = \frac{\Delta r}{r}$, which corresponds to the deformation in the radial direction, is known as the Poisson's ratio ν:

$$\nu = \frac{\varepsilon_r}{\varepsilon_L} \tag{2.7}$$

The Poisson's ratio is an important parameter for ultrasound elastography and indicates the extent of volume change caused by the deformation. The Poisson's ratio ranges from 0.2 to 0.5 for most solids, where 0.5 corresponds to complete incompressibility. Soft tissues are usually assumed to be nearly incompressible, with a Poisson's ratio ranging from 0.490 to 0.495 [66].

The relationship between stress and strain according to Hooke's law (see Eq. 2.3) can be described solely by E and ν. The matrix C_{ijkl} of elastic constants can be reduced to 36 components by taking into account symmetries of the stress and strain tensors [66, 67]. For isotropic materials, the Young's modulus E can be factored out

of C_{ijkl}, which gives in matrix form:

$$\begin{bmatrix} \varepsilon_{11} \\ \varepsilon_{22} \\ \varepsilon_{33} \\ 2\varepsilon_{23} \\ 2\varepsilon_{13} \\ 2\varepsilon_{12} \end{bmatrix} = \frac{1}{E} \begin{bmatrix} 1 & -\nu & -\nu & 0 & 0 & 0 \\ -\nu & 1 & -\nu & 0 & 0 & 0 \\ -\nu & -\nu & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2(1+\nu) & 0 & 0 \\ 0 & 0 & 0 & 0 & 2(1+\nu) & 0 \\ 0 & 0 & 0 & 0 & 0 & 2(1+\nu) \end{bmatrix} \begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{bmatrix} \tag{2.8}$$

The inverse relation may be written as:

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{bmatrix} = \frac{E}{(1+\nu)(1-2\nu)} \begin{bmatrix} 1-\nu & \nu & \nu & 0 & 0 & 0 \\ \nu & 1-\nu & -\nu & 0 & 0 & 0 \\ \nu & \nu & 1-\nu & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1-2\nu}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1-2\nu}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1-2\nu}{2} \end{bmatrix} \begin{bmatrix} \varepsilon_{11} \\ \varepsilon_{22} \\ \varepsilon_{33} \\ 2\varepsilon_{23} \\ 2\varepsilon_{13} \\ 2\varepsilon_{12} \end{bmatrix} \tag{2.9}$$

This relation can be further simplified with the Lamé parameters [68], usually denoted as λ and μ:

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{bmatrix} = \begin{bmatrix} \lambda+2\mu & \lambda & \lambda & 0 & 0 & 0 \\ \lambda & \lambda+2\mu & \lambda & 0 & 0 & 0 \\ \lambda & \lambda & \lambda+2\mu & 0 & 0 & 0 \\ 0 & 0 & 0 & \mu & 0 & 0 \\ 0 & 0 & 0 & 0 & \mu & 0 \\ 0 & 0 & 0 & 0 & 0 & \mu \end{bmatrix} \begin{bmatrix} \varepsilon_{11} \\ \varepsilon_{22} \\ \varepsilon_{33} \\ 2\varepsilon_{23} \\ 2\varepsilon_{13} \\ 2\varepsilon_{12} \end{bmatrix} \tag{2.10}$$

with the Lamé parameters λ and μ corresponding to:

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)} \qquad \mu = \frac{E}{2(1+\nu)} \tag{2.11}$$

The second Lamé parameter μ corresponds to the shear modulus G. In the case

of an incompressible medium, the Poisson's ratio is usually approximated to 0.5. Therefore, the relation between the Young's and shear modulus can be further simplified to [69]:

$$E = 3G \tag{2.12}$$

In shear wave imaging, the elastic moduli also determine the propagation velocity of waves. Ultrasound imaging generally involves two types of wave propagation: longitudinal and transverse waves. The longitudinal waves are typically used to form the ultrasound image. The speed C_L of the longitudinal waves, which is typically about 1540 m/s in soft tissue, can be expressed in terms of the Bulk modulus and density of tissue ρ [70]:

$$c_L = \sqrt{\frac{K}{\rho}} \tag{2.13}$$

The speed of transverse waves c_s (1-10 m/s in soft tissue), also known as shear waves, can be determined by taking the square root of the ratio of the shear modulus G over the tissue density [70]:

$$c_s = \sqrt{\frac{G}{\rho}} \tag{2.14}$$

By combining Eq. 2.12 with Eq. 2.14, we can then define Young's modulus according to the shear wave speed and tissue density:

$$E = 3\rho c_s^2 \tag{2.15}$$

Compressive acoustic waves used for conventional ultrasound imaging travel at high speed in tissue (1540 m/s). On the contrary, mechanical shear waves travel relatively slowly (1-10 m/s), depending on tissue stiffness. According to Eq. 2.1, the acoustic impedance of soft biological tissues depends on their density and speed of sound C_L, which is related to the bulk modulus K with Eq. 2.13. Therefore, the contrast in pulse-echo imaging does not bear any direct information on the elastic

modulus E, which is related to shear wave speed c_s through Eq. 2.15.

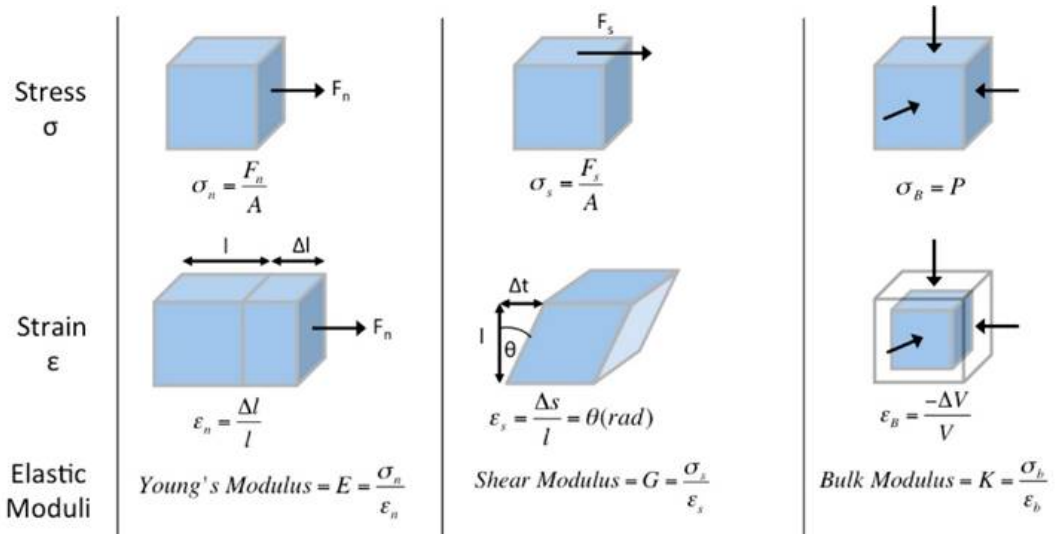

Figure 2.6: Different deformation models and elastic moduli. The elastic moduli can be defined as the ratio of the stress (force per unit area, noted σ) and the strain (expansion/compression per unit length, noted ε). The elastic moduli can be described by the direction of the force: normal to the surface (Young's modulus), tangential to the surface (shear modulus) and in all combined directions (bulk modulus). Image courtesy of Sigrist et al. [33].

2.3 Ultrasound elastography

Ultrasound elastography is an emerging ultrasound technology that was first described by Ophir et al. in 1990 [71]. Nowadays, it refers to a set of imaging techniques by which tissue mechanical properties are being estimated. The basic mechanism of ultrasound elastography consists of applying a force to a tissue and analysing its mechanical response to measure its elastic properties. The different techniques can be classified according to the nature of the force (quasi-static or acoustic radiation force) and the physical parameter being estimated (strain or elasticity modulus). In quasi-static elastography, the force is applied either by an external force, for example by manually pressing on the anatomy with the transducer, or by using physiological internal motions, such as blood vessel pulsation. In shear wave imaging, internal shear waves can be induced by a high-intensity pulsing push emitted by the ultrasound probe. The following sections provide an in-depth explanation

of the physics theory on which elastography is based, as well as a description of the different quasi-static and transient ultrasound elastography techniques (see figure 2.7).

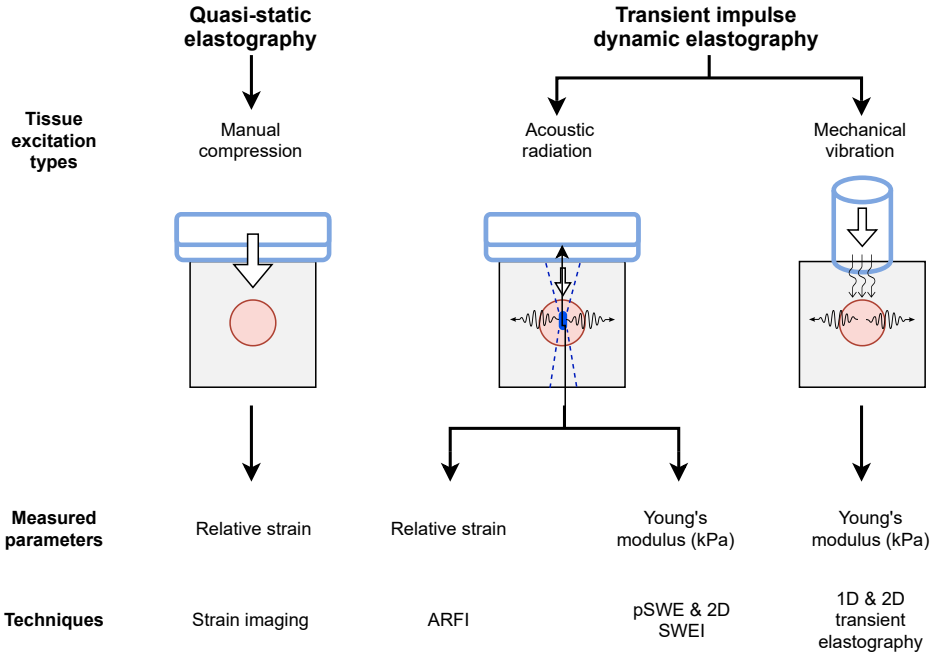

Figure 2.7: Summary and classification of the ultrasound elastography techniques discussed in this chapter.

2.3.1 Quasi-static elastography imaging

Quasi-static elastography was the first technique to be introduced by Ophir et al., in 1990 [71]. It consists in applying a slight force to the targeted tissue, and determining the tissue strain by analysing the displacement between two echo signals, acquired before and after applying compression. The strain typically investigated in quasi-static elastography is the longitudinal strain, which is linked to Young's modulus via Eq. 2.4. The longitudinal deformation relative to the applied stress is expressed as the ratio of the total deformation of the tissue ΔL to its initial length L.

In elastography, the strain is calculated by taking the derivative of the displacement u in the direction of the applied force:

$$\varepsilon_L = \frac{\Delta L}{L} = \frac{\delta u}{\delta x} \tag{2.16}$$

The longitudinal strain is also known as the *engineering* strain or the *Cauchy* strain. The displacement u is retrieved by acquiring an ultrasound image before and after applying the quasi-static force. This step is known as displacement estimation and is discussed in detail in chapter 3. It is important to note that, unlike Young's modulus, the strain is not an intrinsic mechanical property and does not provide a quantitative measure of tissue stiffness. This is because the strain is relative to the applied force, which is not being quantified in quasi-static elastography.

Different tissue excitation methods have been developed in quasi-static elastography. The most common technique is to apply force with a handheld probe on the surface of the tissue, which is similar to palpation [37]. This method is particularly useful for quantifying the stiffness of superficial anatomical structures. However, the palpation force is not efficiently transmitted to deeper structures due to the attenuation observed in soft tissue. In addition, the application of the stress is usually applied manually and is inherently highly user-dependent. A non-steady image acquisition can introduce some out-of-plane displacements and deteriorate the quality of the resulting strain images.

To induce uniform displacement, Hiltawsky et al. [72] developed a freehand applicator that can be used to place a hand-held probe perpendicular to the tissue surface and to apply a purely vertical displacement on the tissue. Another technique consists of using an additional low-frequency vibrator system (1-10 Hz) that generates quasi-static deformations [73, 74]. More recently, Patlan-Rosales et al. [75] developed a robotic system with haptic feedback to guide the clinician when performing palpation.

Internal physiological motion can also generate quasi-static displacements. These stimuli include respiratory and cardiovascular deformations [76]. This type of excitation can be useful to image tissue stiffness of deeper structures, where palpation is more challenging. However, physiological motions are hard to predict and introduce non-uniform deformation, which limits reproducibility.

Quasi-static elastography was the first ultrasound tissue characterisation technique to be translated into clinical use. This includes the detection of lesions in

liver disease [77, 78] and prostate cancer [79]. Quasi-static elastography has also proven useful to characterise tumour tissue by differentiating benign from malignant tumours in the thyroid [80], kidney [81], breast [37, 72, 82] and lymph nodes [83]. More recently, it has been used in image-guided interventions such as for liver [38] and brain tumour surgery [22, 40]. An example of a brain tumour elastogram (i.e., a tissue elasticity map) is showed in figure 2.8.

Most manufacturers of ultrasound systems have commercialised their own quasi-static elastography method, which makes it the most widely available ultrasound elastography technique. The information displayed is usually the strain map, also known as elastogram, which can be overlaid on the B-mode image. The elastogram is displayed as a pixel map with a colour scale that differs between different vendors. Other metrics can also be provided, such as the strain ratio which corresponds to the ratio between the strain measured in the target tissue region (i.e., a tumour) and a reference region (i.e., background). Although the strain ratio is not a quantitative measure, it has been extensively used for breast cancer diagnosis [84].

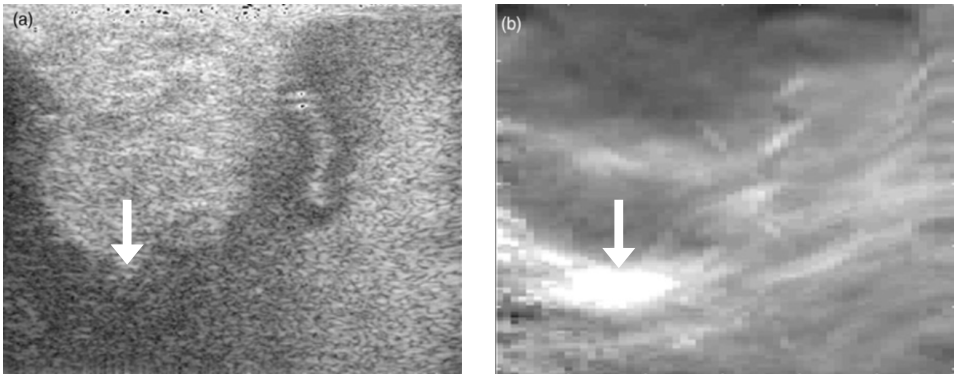

Figure 2.8: Quasi-static elastography example. B-mode image (a) and strain elastogram (b) of a convexity meningioma acquired during brain tumour resection. We can visualise a high strain at the inferior brain-tumour interface (white arrow). Adapted from Chakraborty et al. [40]

2.3.2 Transient impulse shear wave elastography

The quantitative evaluation of tissue stiffness in ultrasound elastography is provided by the Young's modulus. Shear wave imaging techniques estimate the Young's mod-

ulus by measuring the velocity of the induced shear wave (see Eq. 2.15). Transient impulse methods rely on the ability to track shear wave propagation by monitoring the tissue motion induced by the applied force. Similar to quasi-static strain estimation, the tissue motion is typically estimated by calculating the displacement between frames acquired before and after the excitation using correlation algorithms either in the phase or the time domain. A literature review on displacement estimation is available in chapter 5.

Once we track the shear wave propagation in tissue over time using ultrasonic displacement estimation, it is possible to use this information to recover quantitative elastic information, i.e, the shear and Young's moduli. The initial method proposed to recover elasticity mapping was based on the inversion of the second-order wave equation. The governing equation for wave propagation in an isotropic, elastic and locally homogeneous medium can be expressed as [85]:

$$G\nabla^2\vec{u} + (K+G)\nabla(\nabla\cdot\vec{u}) = \rho\frac{\partial^2\vec{u}}{\partial^2 t} \tag{2.17}$$

where G and K are the shear and bulk modulus, $\nabla^2\vec{u}$ is the spatial Laplacian of the shear wave displacement $\vec{u}$, ∇ is the del operator, ρ the tissue density (approximately 1000 kg/m^3 for soft biological tissue) and t is time (in seconds) [86]. Assuming that the propagated wave is pure shear wave and thus, divergence-free, $\nabla\cdot\vec{u} = 0$, then Eq. 2.17 can be simplified to the time-independent form of the wave equation, also known as the Helmholtz equation:

$$G\nabla^2\vec{u} = \rho\frac{\partial^2\vec{u}}{\partial^2 t} \tag{2.18}$$

The inversion of Eq. 2.18 can be implemented by computing the second-order spatial and temporal derivatives of the shear wave displacement. However, the 3D shear wave displacements cannot be estimated by imaging the medium with a 1D ultrasound probe. Bercoff et al. calculated the shear wave displacement only in the axial direction, hypothesising that the three components of the displacement are independent solutions of the wave equation [86], which gives:

$$G\nabla^2 u_z = \rho \frac{\partial^2 u_z}{\partial^2 t} \tag{2.19}$$

Since the imaging field of view corresponds to the (x,z) plane, only two components of the Laplacian can be calculated and, therefore, they assumed that:

$$\frac{\partial^2 u_z}{\partial^2 y} \ll \frac{\partial^2 u_z}{\partial^2 x} + \frac{\partial^2 u_z}{\partial^2 z} \tag{2.20}$$

Thus, the shear modulus can be calculated as :

$$G(x,z) = \rho \frac{\frac{\partial^2 u_z(x,z)}{\partial^2 t}}{\frac{\partial^2 u_z(x,z)}{\partial^2 x} + \frac{\partial^2 u_z(x,z)}{\partial^2 z}} \tag{2.21}$$

Although all the components of Eq. 2.21 can be obtained using ultrasound tracking techniques, the estimated displacements are in practice corrupted with signal noise. Therefore, it is challenging to obtain an accurate estimation of the spatial second derivatives from the ultrasonic displacement data. This requires different filtering operations to remove noise from displacement data, which can be computationally intensive, and thus challenging, for real-time processing.

The common approach to estimate the propagation speed of the shear wave consists in performing time-of-flight measurements, where the shear wave position is represented as a function of time. This is typically done by estimating the time of arrival of the shear wave at different lateral positions in the axial displacement data. The elastic modulus can then be estimated by using Eq. 2.15. The time of arrival can be estimated using either cross-correlation or time-to-peak methods.

The cross-correlation method estimates the time of arrival of the propagating shear wave front by determining the time lag between adjacent lateral positions [87]. The Time-To-Peak (TTP) method estimates the arrival time of the wave by taking the intensity peak of the displacement for each lateral position. The shear wave speed is then estimated by finding the slope of a fitted line estimated with linear regression on the time-to-peak displacement versus the lateral position [88–90].

2.3.2.1 Transient Elastography

Transient elastography (TE) uses low-frequency (50 Hz) mechanical vibration to generate shear waves in the targeted tissue. The external push is induced by a mechanical vibrating device, and the shear wave propagation in the tissue is tracked with a separate ultrasound transducer using motion-mode (M-mode) imaging. In contrast with B-mode imaging, which reconstructs a 2D image with multiple ultrasound scan lines across the lateral direction, M-mode imaging allows tracking of tissue displacement on one scan line with a high temporal resolution. This one-dimensional method allows one to follow the propagation of the shear wave and to measure its velocity along a scan line, which corresponds to a small tissue volume of about 1 cm wide and 4 cm long. The main drawback of this technique is that M-mode imaging does not allow imaging of the region of interest to guide tissue stiffness assessment. Ultrafast ultrasound imaging has later been implemented for transient elastography, allowing real-time 2D imaging of a larger portion of tissue [91].

This technique was the first shear wave ultrasound method to be commercially available (FibroScan^{TM}, Echosens, Paris, France) [92, 93]. It has been used mainly for liver diseases, with performance similar to serum biomarkers for the diagnosis of liver fibrosis [94]. This method has gained increasing attention to monitor the progression of chronic liver disease and could replace the more invasive liver biopsy.

2.3.2.2 Acoustic Radiation force imaging (ARFI)

Acoustic radiation force-based elasticity imaging techniques use a short duration (0.1-0.6 ms), high-intensity pulse (peak pulse average = 1400 W/cm^2) to induce an internal force to displace tissue. This force results from the transfer of momentum from the ultrasonic wave to the medium through which it is propagating due to scattering and absorption mechanisms [96]. This 'pushing' pulse can be generated by focusing several acoustic beams of a standard ultrasound transducer. According to the derivation by Nyborg et al [97, 98], the radiation force F for an acoustic plane wave can be related to the absorption α and speed of sound c in the tissue, as well as the local average intensity of the focused acoustic beam I so that:

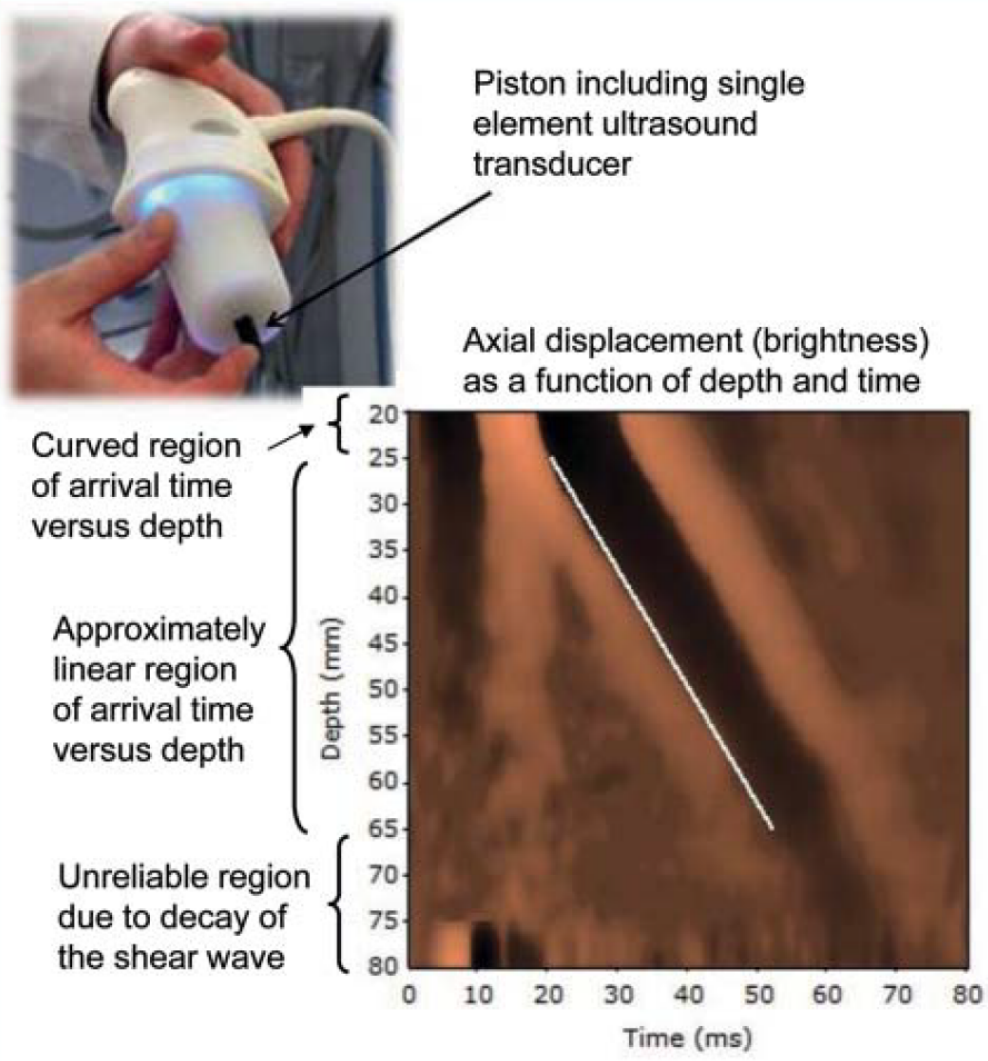

Figure 2.9: Transient elastography with FibroScan^{TM} (Echosens, Paris, France). The upper image shows the FibroScan ^{TM} applicator that contains the piston to generate the shear waves. The lower image represents axial displacement as a function of time and depth, obtained from an M-mode image sequence. The shear wave velocity by calculating the slope of the white line is drawn on the image. It corresponds to the best fit of shear wave time-of-arrival according to depth. Image courtesy of Bamber et al. [95].

$$F = \frac{2\alpha I}{c} \tag{2.22}$$

It is important to note that this derivation assumes that ultrasound waves in soft tissues are mainly attenuated by absorption and that the contribution of scattering can be neglected. The absorption of the radiation force by the tissue generates a displacement (about 10-30 µm) and normal shear waves that propagate perpendicular

to the imaging/excitation plane. Therefore, this excitation method allows the measurement the relative tissue deformation, i.e., strain, and the Young's modulus, which is proportional to the shear wave's propagation velocity. Excitation and imaging are performed with the same transducer and allow B-mode images to be displayed to guide the procedure.

For strain measurements, ARFI consists of tracking the tissue displacement in the region excited at a single lateral location by the pushing pulse. This is done by computing the displacement between a reference pulse, before the push, and a tracking pulse which is acquired directly after applying the acoustic radiation force. This procedure can be repeated across the aperture of the probe to recover a 2D strain elastogram [98]. In addition, this excitation method lowers signal decorrelation due to out-of-plane and large displacements, and is less user-dependent, compared to quasi-static elastography.

The shear wave ARFI method, also called point Shear Wave Elastography (pSWE), follows the same acquisition protocol but instead tracks the shear waves generated by the pushing pulse. Since the shear waves propagate away from the region of excitation, the reference/tracking pulses are acquired laterally to the pushing pulse. The shear wave speed in tissue has traditionally been calculated with time-of-flight methods [99]. These methods can provide a local average shear wave velocity measurement or the procedure can be repeated at different lateral locations to reconstruct a 2D quantitative map of elasticity (i.e., Young's modulus map). An example of velocity estimation with pSWE is shown in figure 2.10.

Both methods were developed by Nightingale et al. [100]. They are commercially available on the Siemens Acuson 2000 & 3000 system (Siemens Medical Solution, USA) as Virtual Touch^{TM} tissue imaging and Virtual Touch^{TM} tissue quantification, respectively [101]. ARFI has been used for many clinical applications. It has been considered for cardiovascular imaging, for example to guide thermal ablation procedures [102], or to identify arterial plaques [103]. It has also been investigated for tumour characterisation, such as breast [104] and prostate cancer [105].

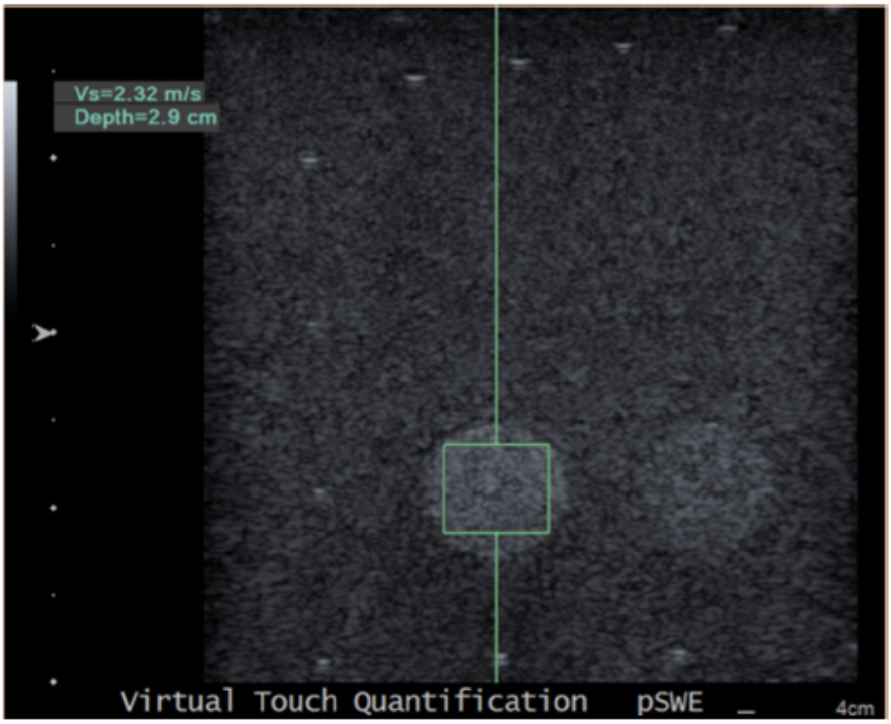

Figure 2.10: Example of point shear wave elastography in a tissue-mimicking phantom with the Virtual Touch^{TM} system. The shear wave speed is calculated in the green box, with the value indicated in the left top corner. Image courtesy of Ozturk et al. [106]

.

2.3.2.3 2D Shear wave elasticity imaging

The emergence of real-time 2D shear wave imaging has been made possible by the development of plane wave imaging, leading to ultrafast frame rates up to 10,000 frames per second (fps) [64]. This technique uses the same tissue excitation as ARFI, but the resultant shear wave front is readily detectable and tracked with ultrafast imaging in the tissue.

The first technique to make use of multiple acoustic radiation force excitations was supersonic shear imaging (SSI), developed by Bercoff et al. [86]. This method consists of inducing a conical quasi-plane shear wave front by applying consecutive acoustic radiation force excitations at different focusing depths. The resulting shear waves interfere constructively along the so-called 'Mach cone', which results in the

formation of two intense quasi-plane shear waves propagating in opposite directions. The name of this method refers to the 'sonic boom' created by supersonic aircraft. The formation of the Mach cone is made possible because the excitation foci are moved through the tissue above the supersonic speed, which is considerably higher than shear wave propagation speed (10-30 m/s). The shape and intensity of the resulting shear waves allow monitoring of simultaneous tissue response across a larger field of view than ARFI. An example of SSI acquisition is illustrated in figure 2.11.

This method was implemented on the Aixplorer ultrasound scanner (SuperSonic Imagine, France). The feasibility of SSI has been assessed for the characterisation of breast tumours [107] and liver viscoelasticity measurements [108]. Similar to ARFI, it has also been used to monitor thermal ablation procedures [86].

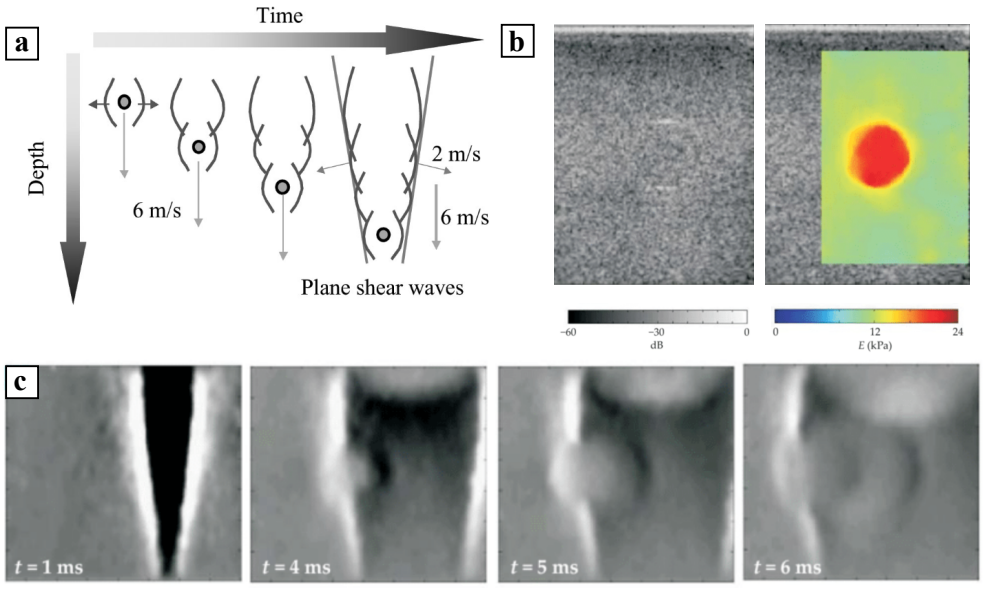

Figure 2.11: Supersonic shear imaging (SSI). (a) Sequential generation of the supersonic shear source. (b) B-mode image and its corresponding SSI 2-D elastogram acquired from a tissue-mimicking phantom. (c) Propagation of two quasi-plane shear waves generated from a supersonic acoustic radiation force. Adapted from Bercoff, Fink and Tanter [86, 109].

Another multi-stimuli acoustic radiation force imaging technique is Comb-push ultrasound shear elastography (CUSE), developed by Song et al. [110]. CUSE

generates shear waves by exciting the tissue with multiple, simultaneous pushing pulses in the lateral direction. This allows to obtain a 2D elasticity map of the whole field of view, and compensate for inaccurate shear wave speed in the focal beam region. The ARFI excitations can be focused or unfocused.

This method was further refined to be used on a standard ultrasound scanner with the time-aligned sequential tracking (TAST) method [111]. CUSE has been used for the *in vivo* assessment of breast tumour mass [112] and thyroid nodules [113].

2.4 Conclusion

This chapter presented an overview of ultrasound imaging, followed by fundamentals of linear soft tissue elasticity and their relationship with ultrasound elastography. This chapter also described different tissue stiffness evaluation techniques (e.g., strain imaging, ARFI, pSWE, SWEI) which can be classified into two ultrasound elastography modalities which are; quasi-static elastography and transient impulse shear wave elastography. Quasi-static elastography techniques can be used with most state-of-the-art ultrasound scanners and provide a qualitative estimation of tissue stiffness (i.e., strain). Transient impulse elastography techniques are based on the generation and study of shear waves in tissue to estimate tissue mechanical properties such as the strain, shear and Young's moduli.

Both quasi-static and shear wave elastography have been investigated in neurosurgery, showing promising results on their ability to delineate brain tumours [40, 41]. However, the lack of image analysis tools and integration with the surgical navigation system limited the usefulness of elastography in neurosurgery and its clinical translation. This thesis hypothesised that the combination of ultrasound elastography with neuronavigation could improve the location of brain tumours. The following chapters of this thesis will focus on the development of novel learning-based techniques to improve tissue stiffness evaluation (chapters 3, 4 and 5), and their integration into a neuronavigation module (chapter 6). We first considered the improvement of quasi-static elastography strain estimation (chapters 3 and 4), be-

cause its low hardware requirement makes it a good candidate for clinical translation. We then extended our method to shear wave elastography imaging (chapter 5), which offers operator-independent and quantitative estimation of tissue stiffness, and could better characterise brain tumours.

Chapter 3

An unsupervised approach to quasi-static elastography with end-to-end strain regularisation

This chapter focuses on the development of a new learning-based method for displacement estimation in quasi-static elastography. We first introduce the challenges and previous work related to speckle tracking for tissue strain estimation in section 3.1. We then introduce our unsupervised training strategy and network architecture in section 3.2. The sections 3.3 and 3.4 of this chapter respectively present and discuss our results on numerical simulation and *in vivo* data. The following work was presented at the 23rd international conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) in October 2020 [114].

3.1 Introduction

3.1.1 Previous work in displacement estimation

Various methods of displacement estimation for quasi-static ultrasound imaging have been proposed over the years. The most common approach consists of searching for the time-delay between local windows by maximising a function that quantifies the similarity between a pair of inputs. This maximisation has been operated both in the time [115–117] and phase domain [118, 119] of the RF data. Different similarity metrics have been proposed to match consecutive RF frames, such as the cross-

correlation function [120–122], the sum of square difference [123, 124] or the sum of absolute difference [125, 126]. Although windows-based methods have shown good performance in displacement estimation, working with local windows prevents the accurate prediction of large deformations and is sensitive to global decorrelation, i.e., the change of speckle appearance due to out-of-plane motion.

A different strategy involves minimising a cost function that combines image similarity and displacement regularity [127, 128]. These methods assume the displacement throughout the tissue to be smooth and, therefore, justify the use of a regularisation parameter that penalises the correlation function to prevent displacement discontinuity. However, this type of approach can be computationally expensive and is not suitable for real-time application. More recently, new approaches have been proposed to speed up the computation, such as by using dynamic programming [129] or by minimising a nonlinear cost function on the entire image [130]. Although these methods can be used in real-time, they are still sensitive to noise [131], which is of critical importance for real-time elastography.

The recent progress of learning-based methods for optical flow estimation has inspired novel strategies for ultrasound strain estimation [132, 133]. Wu et al. proposed a convolutional neural network (CNN) that takes as input patches sampled from a RF ultrasound data pair, and iteratively reconstruct the displacement and strain estimates [134]. Gao et al. further improved this method by proposing to use the Learning-Using- Privileged-Information (LUPI) paradigm [135]. LUPI makes use of intermediate displacement estimate obtained from numerical finite element method (FEM) simulation to refine the strain prediction during training. The authors showed that using privileged information during training improved the accuracy and generalisation of the network, as compared with their previous method. Kibria et al. used FlowNet 2.0, which was the first CNN architecture to outperform classical approaches in optical flow estimation [133]. The network was directly used on ultrasound elastography data to produce initial displacement estimates, which were further refined with an optimisation-based method [131]. Other methods directly used pre-trained optical flow networks and designed transfer learning strategies to

apply them on ultrasound RF data. Tehrani et al. presented a training strategy to fine-tune the Pyramid, Warping and Cost volume Network (PWCNet), another well-known optical flow network [136]. Their training strategy consists of minimising a loss function between the network's displacement estimates and their respective ground truth labels, generated from numerical ultrasound phantoms via finite element methods [137]. Peng et al. compared the performance of three different optical flow networks trained on ultrasound numerical simulation: FlowNet, PWC-Net, and Lite-FlowNet [138]. They assessed the performance of each network on tissue-mimicking phantoms and *in vivo* breast ultrasound data. The investigated networks appeared to provide similar performance as compared with state-of-the-art optimisation-based methods.

The aforementioned deep learning methods for quasi-static elastography all used FEM ground truth labels to train networks in a supervised way. Supervised learning refers to the task of training a network by using labeled datasets. This approach prevents the model from training on real-world ultrasound data because ground-truth displacement fields are not possible to obtain because the deformation applied to the tissue is usually difficult to quantify. Moreover, learning from real-world ultrasound data can improve the model's generalisation ability because this type of data often exhibits complex speckle patterns and echogenic features, which can be quite challenging to replicate in ultrasound simulation. The lack of ground truth data in medical imaging has driven the development of unsupervised training strategies in image registration [139, 140]. These methods usually train their network by optimising a similarity metric between input image pairs, by using the predicted displacement field to resample the moving image into the reference image space.

3.1.2 Contributions

In this work, we proposed an unsupervised method for speckle tracking that allows a neural network to be trained directly on clinical data and predict tissue displacement. The goal of unsupervised learning is to train a deep learning model without using labeled data. Therefore, our convolutional neural network (CNN) training procedure is performed without ground truth labels. Instead, the network weights are optimised

by minimising a dissimilarity function between the pre-compression image and the compressed image warped with the output displacement field. We also introduce a new regularisation term that preserves displacement continuity by directly optimising the smoothness of the strain prediction. We validated our method on *in vivo* ultrasound data acquired on the forearm of three human volunteers and on numerical ultrasound simulations. The performance of our method was evaluated by comparing our displacement field and strain prediction with the ground truth labels and a gold standard optimisation-based method [141]. To the best of our knowledge, this is the first unsupervised training strategy applied to quasi-static ultrasound imaging.

3.2 Methods

3.2.1 Problem statement

Our method follows an approach similar to a standard image registration framework, which aims to find a spatial transformation that maps a moving image into the space of a reference image. In the case of a non-rigid image registration solution, this transformation can be represented as a dense displacement field (DDF). This transformation is generally optimised in an iterative manner, by maximising an objective function that measures the similarity between the warped moving image and the reference image, and a regularisation parameter to ensure displacement continuity. From a learning-based approach perspective, the image mapping is predicted by a neural network instead of being directly optimised. An overview of the method is presented in figure 3.1.

If we consider a pair of RF images acquired from a tissue before and after undergoing compression, the aim of displacement estimation is to find the displacement field between the two images, which allows us to determine the tissue strain properties. In the case of a linear transducer, the acquired images are 2D and therefore the resulting DDF comprises the displacements in the axial and lateral direction. This allows the computation of the axial and lateral strain by differentiating the displacement field. However, since the tissue undergoes an axial compression, the elastic response will be more important and clinically useful in this direction. Therefore,

only axial displacement and strain maps are shown in this study.

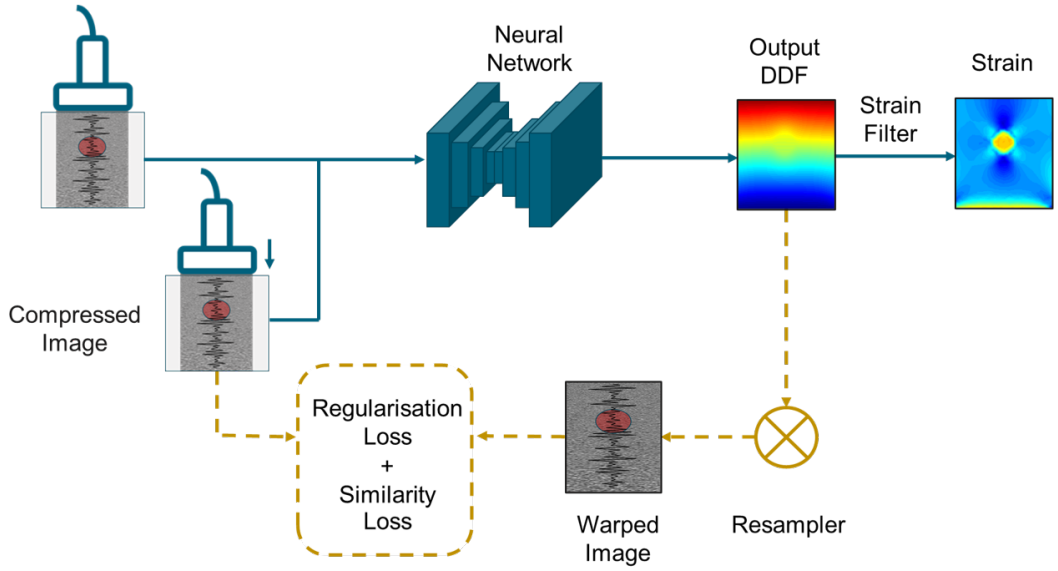

Figure 3.1: Overview of the method. This diagram illustrates the proposed unsupervised training strategy of our method, where the yellow dashed lines indicates data flows only required in training. During inference, the convolutional neural network (CNN) takes as input a pair of ultrasound radio-frequency data to output a dense displacement field (DDF). The strain is then derived from the displacement field in post-processing by using the least-squares strain estimator. The training stage only requires the input image pairs, where the output DDF is used to warp the compressed tissue image to align with the reference image.

3.2.2 Displacement Estimation

Our network takes as input a pair of pre- and post-compression 2D RF frames, here named *Pre* and *Post*, and predicts a DDF u. We named our network 'USENet' for 'Ultrasound Elastography Network'. The network parameters are learned by minimising a dissimilarity metric L_{sim} between pairs of 2D RF ultrasound data. Our training loss function also includes a regularisation term, L_{reg}, acting on the predicted displacement field u and associated with a weighting hyperparameter α, to ensure balance between the likelihood and smoothness of the predicted transformation. The optimisation problem can be written as:

$$\hat{\theta} = \arg\min_{\theta} \left[L_{sim}(Pre, Post; u) \; + \; \alpha \; \cdot \; L_{reg}(u) \right] \tag{3.1}$$

where θ represents the network parameters that are optimised through stochastic gradient descent [142]. The goal of stochastic gradient descent is to minimise the loss function by updating each parameter θ by a small amount at each iteration. The direction of fastest decrease of the objective function at each iteration is given by its negative gradient, which is then used to update the network's parameters.

The dissimilarity metric corresponds to a negative local normalised cross-correlation (NCC) which averages the NCC score between sliding windows sampled from the pre-compression image and the transformed post-compression image, re-sampled with the predicted displacement field. The NCC between two local image windows, W_{pre} and W_{post}, with i, j pixel components can be written as:

$$NCC = \frac{1}{N} \sum_{i,j} \frac{\left[W_{pre}(i,j) - \mu_{W_{pre}}\right] \cdot \left[W_{post}(i,j) - \mu_{W_{post}}\right]}{\sigma_{W_{pre}} \cdot \sigma_{W_{post}}} \tag{3.2}$$

where N is the number of pixels (i, j) and μ and σ correspond to the mean and standard deviation of the images, respectively. The local windows of size 160×16 pixels were sampled from the pre-compression and resampled post-compression images.

The regularisation term consists of the L1-norm of the second spatial derivatives of the predicted displacement field u. According to Mirzai et al., the L1-norm leads to substantial improvement in the sharpness of displacement estimates for quasi-static elastography [141]. Given that the strain modulus corresponds to the displacement gradient, minimising its second derivative allows to enforce the strain map smoothness. The regularisation term can be written as:

$$L_{reg} = \sum_{i,j} \left(\mid \partial_x^2 u_{i,j} \mid + \mid \partial_y^2 u_{i,j} \mid \right) \tag{3.3}$$

where ∂_x^2 and ∂_y^2 are the second partial derivatives in axial and lateral directions, respectively.

3.2.3 Implementation

3.2.3.1 Network architecture

The architecture of USENet is based on the U-Net, a fully convolutional network widely adopted in the medical imaging community [143, 144]. Furthermore, the use of this type of architecture has been demonstrated successfully for many medical image registration tasks [140, 145].

The network architecture is illustrated in figure 3.2, and consists of a contracting path (i.e., encoder) and an expansive path (i.e., decoder). The encoder part is composed of four down-sampling ResNet blocks [146], which capture the hierarchical features necessary to establish a correspondence between the pair of images. Each block corresponds to a residual unit composed of two sequential convolutional layers with a batch normalisation layer and a leaky rectified linear unit.

The contracting path starts with 8 feature map channels, which are then successively doubled after each ResNet unit. The contraction is performed by halving the feature map size after each down-sampling block. This is performed by using max pooling operation, which consists of using the maximum value of patches sampled from the feature map to reduce its spatial dimension. USENet uses max pooling with a filter size of 3×3 and a stride of 2.

Symmetrically, the decoder part is composed of four up-sampling ResNet blocks that consist of an additive up-sampling layer summed over a transpose convolutional layer. Skip connections are used to feed the output of each down-sampling block to its corresponding decoding block, and were introduced to prevent the gradient vanishing problem [146]. Finally, each up-sampling block outputs a displacement field that is convolved and resized to the input size, then summed to output the predicted displacement field.

This method was implemented in TensorFlow™ using NiftyNet [147]. USENet was trained with a 12 GB NVIDIA Pascal™ TITAN Xp general-purpose graphic processing unit (GPU). Image resampling was performed with the bilinear resampler module available in NiftyNet. Stochastic gradient descent was performed using the TensorFlow™ implementation of the Adam optimiser [148]. The regularisation

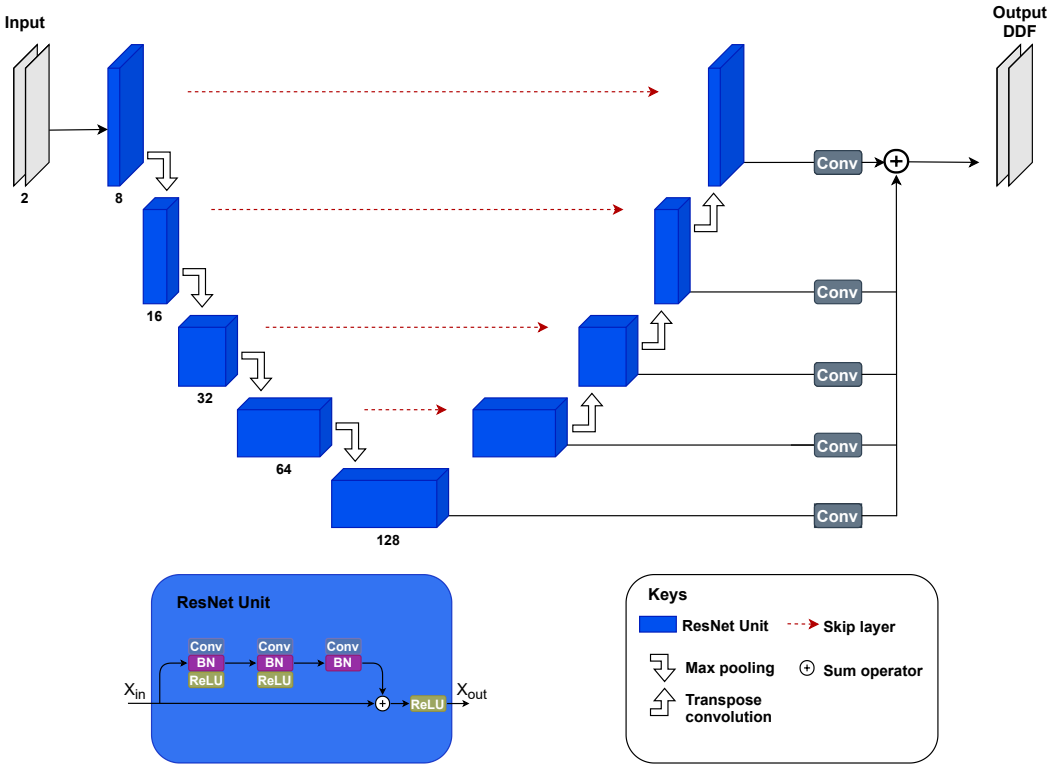

Figure 3.2: Network architecture of USENet. Each blue rectangle corresponds to a ResNet unit (illustrated below) with different channel size printed underneath. The output of each upsampling ResNet unit is then convolved and resized to the input size before being summed to output the displacement field.

weight was set to $\alpha = 1$ to ensure displacement continuity.

3.2.3.2 Strain Estimation

In quasi-static elastography, the strain estimates are obtained by computing the displacement field gradient. However, direct differentiation of the displacement field is rarely used because gradient operations generate a significant amount of noise in the resulting strain map. We used the least-squares strain estimator (LSQSE) to improve the elastogram elastographic Signal-to-Noise ratio [149]. The LSQSE method reduces the noise amplification due to gradient operation by achieving a linear curve fit to the displacement field estimates. This method samples 1D displacement kernels with scalable size n along depth to reconstruct the axial strain map. For soft tissue, we assume linear elasticity within the sampled kernel. Therefore, we can write the following relationship:

$$u(i) = a \cdot z(i) + b \tag{3.4}$$

where i corresponds the index of the kernel centre pixel, u the sampled 1D displacement field, z the tissue depth, and a and b are the constants to be estimated. This equation can be written in the following matrix notation:

$$\begin{bmatrix} u(i) \\ u(i+1) \\ u(i+2) \\ \vdots \\ u(i+n) \end{bmatrix} = \begin{bmatrix} z(i) & 1 \\ z(i+1) & 1 \\ z(i+2) & 1 \\ \vdots & \\ z(i+n) & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} \tag{3.5}$$

which is equivalent to this compact notation:

$$u = \mathbf{A} \begin{bmatrix} a \\ b \end{bmatrix} \tag{3.6}$$

where $\mathbf{A}$ is a nx2 matrix including the depth position $z(i)$ in the first row as well as a second row of ones. The slope of the fitted curve a corresponds to the estimated axial strain and is estimated by minimising the sum of squared error between the linear model and the displacement estimates, which is given by the well-known least-squares solution:

$$\begin{bmatrix} a \\ b \end{bmatrix} = [\mathbf{A}^T \mathbf{A}]^{-1} \mathbf{A}^T u \tag{3.7}$$

Since a represents the tissue strain in the middle of the kernel, the strain information contained at the top and bottom of the field (with length $n/2$) is not estimated. The parameter n also affects the smoothness of the resulting strain map and needs to be carefully adjusted. Since the LSQSE method assumes linear elasticity, a large kernel size will result in an over smooth strain field, which can prevent to clearly visualise a difference of stiffness at different tissue interfaces. In this study, we adjusted the kernel size n manually and find the best results for $n = 137$.

3.3 Experiments

3.3.1 Simulation dataset

The Field-II software [150, 151] was used to generate the ultrasound images. It is an ultrasound simulation package based on linear acoustics that calculates the emitted ultrasound field at a given spatial location, as a function of time using a convolution of the spatial impulse response with the excitation function. For pulse-echo imaging, the received signal is determined by convolving the transducer excitation function with the spatial impulse responses of the emitting and receiving apertures. The calculation of the spatial impulse response is based on the Tupholme-Stepanishen method [152, 153]. Both numerical phantom and ultrasound simulations were performed using the cyst phantom example available in Field-II [154].

The different processing steps required to create the ultrasound simulation dataset is shown in figure3.3. Each simulation consisted of a 3D rectangle of size $38 \times 40 \times 10$ mm, containing a cylindrical inclusion with a randomly assigned diameter (from 8 to 12 mm) and position. The speckle pattern typically observed in ultrasound imaging was obtained by randomly assigning a total of 400,000 scatterers across each digital phantom. The axial compression was assigned randomly to each phantom and represented between 0.5% and 4% of the phantom total length. The Young's modulus of the inclusion was set to different values (i.e., 8, 15, 45, and 75 kPa) while the background was fixed to 25 kPa. The difference of stiffness values between the phantom's inclusion and background were motivated by tissue elastic properties usually found in soft tissue such as in liver fibrosis [155], breast [156] and brain tumours [41]. The numerical phantoms used in this study did not specifically exhibit stiffness values typically found in brain tumours to artificially increase the amount of training data and increase the generalisation of the model. The digital phantoms were meshed by using the Iso2Mesh software [157]. The Poisson ratio of both background and inclusion were set to 0.495. We applied a uniform displacement to the top of the phantom, with a fixed boundary constraint applied to the bottom face, mimicking compression with a hand-held ultrasound probe.

Tissue displacements were estimated by finite element method using the Partial

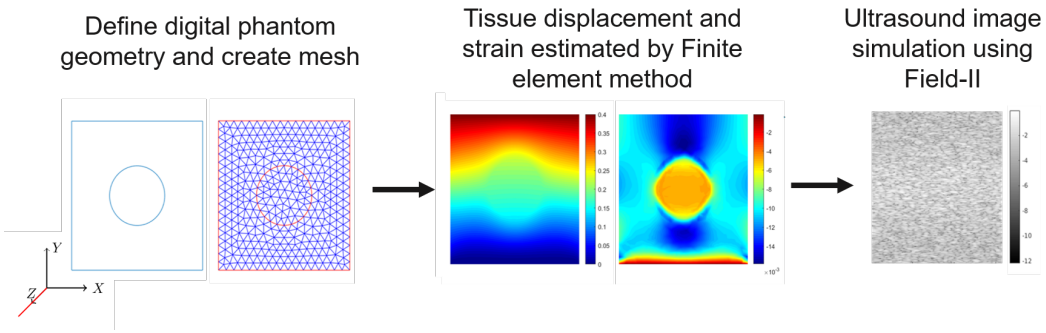

Figure 3.3: Finite element analysis pipeline.

Differential Equation Toolbox from MATLAB, and were used to interpolate the scatterers' position. A total of 192 RF lines were simulated for each image, with a probe central and sampling frequency of 7 and 40 MHz, respectively. Background scatterers were associated with a random intensity value from a Gaussian distribution to mimic homogeneous tissue. To increase the network's robustness to noise and image intensities, the inclusion scatterer intensities were either assigned to 0 or similar to the background. Moreover, white Gaussian noise with random signal power values from 5 to 20 dBW was added to each image to increase robustness to noise. Finally, 160 ultrasound image pairs were simulated, where 100 were used for training and 30 for validation and testing, respectively.

3.3.2 *In vivo* dataset

We created our own *in vivo* dataset by collecting images of the arms of three human volunteers. The RF data was acquired from a Cicada 128PX system equipped with a 10 MHz linear probe from Cephasonics (Cephasonics Inc., USA). The images were reconstructed using the delay-and-sum beamformer from SUPRA [158]. The *in vivo* dataset acquisition pipeline is illustrated in figure 3.4. The experimental protocol consisted of acquiring pairs of images while slowly applying an axial compression on the volunteer's arm with the handheld ultrasound probe. The *in vivo* dataset included 1300 image pairs for training and 300 pairs for validation and testing. The three image pairs presented in the results' section were taken from the testing partition and exhibit one or several blood vessels located in the arm.

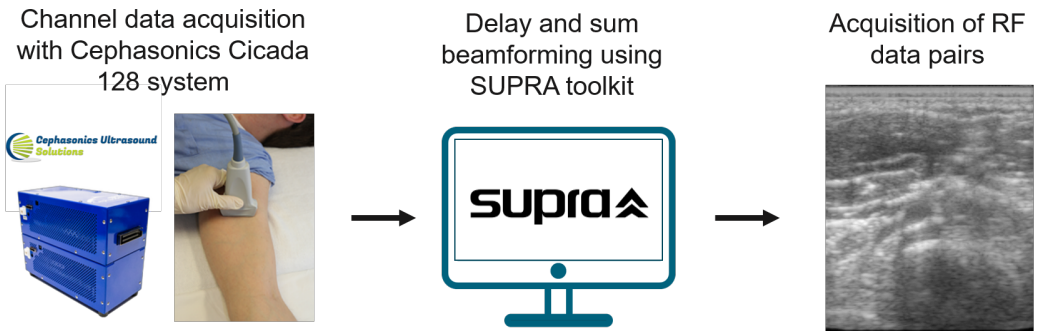

Figure 3.4: *In vivo* image acquisition pipeline.

3.3.3 Performance comparison

We compared our results on ultrasound simulations and *in vivo* data with a state-of-the-art strain elastography method called "tOtal Variation Regularization and WINDow-based time delay estimation" (OVERWIND) [141]. This method performs displacement estimation by using an iterative optimisation-based strategy applied to image windows. The OVERWIND results were obtained by using a MATLAB implementation made publicly available by the authors. The results on simulation were also compared with the strain estimates obtained by training our network with a supervised loss function. The supervised loss function was used previously in a learning-based method for quasi-static elastography [134]. It corresponds to the mean absolute error (MAE) between the network prediction and the ground truth labels. The MAE can be written as follows:

$$MAE = \frac{1}{N}\sum_{i}^{N}|Predicted_i - Label_i| \tag{3.8}$$

Where $Predicted$ and $Label$ correspond to the estimated and ground truth strain values, respectively.

We validated the performance of our method on our testing dataset. We used the root-mean-square error (RMSE) to measure the average magnitude of the errors between our displacement predictions and the ground truth labels. We also used the Contrast-to-Noise Ratio (CNR) and elastographic Signal-to-Noise Ratio (SNRe) to measure the quality of our strain estimation with respect to the ground truth. We

defined the CNR and SNRe as follows:

$$SNRe = \frac{\mu}{\sigma} \qquad CNR = \sqrt{\frac{2(\mu_b - \mu_i)^2}{\sigma_b^2 + \sigma_i^2}} \tag{3.9}$$

where μ and σ represent the mean and standard deviation of the elasticity map. Background and inclusion windows were used to compute the CNR for heterogeneous phantoms, where μ_b, σ_b and μ_i, σ_i are respectively the Young's modulus mean and standard deviation in the background and inclusion regions. The SNRe and CNR ratio are metrics commonly used in USE [141, 159]. We compared the SNRe and CNR between our prediction and ground truth labels by sampling both the variances and average in the stiffer inclusion and its background. The background and inclusion window were carefully selected as described in [141] to ensure meaningful statistical significance.

3.3.4 Results

Figure 3.5 shows axial strain images estimated from three simulated image pairs and obtained with FEM, OVERWIND and both supervised and unsupervised models. The averaged CNR and SNRe values of the entire testing dataset (i.e., 30 image pairs) are displayed in table 3.1.

Table 3.1: Mean and standard deviation of SNRe and CNR for the strain images of the simulation testing dataset obtained with finite element methods, OVERWIND, the unsupervised and supervised network.

	Ground truth (mean std)	Unsupervised (mean std)	Supervised (mean std)	OVERWIND (mean std)
SNRe	7.51 (2.61)	6.95 (2.54)	7.79 (2.01)	**9.31 (3.51)**
CNR	9.15 (2.73)	**7.70 (3.8)**	4.22 (2.08)	6.33 (3.6)

Figure 3.6 shows the estimated axial strains of three different *in vivo* image pairs. CNR and SNRe values for each case can be found in table 3.2. The local NCC values between the post-compression and resampled pre-compression images are also displayed, to indicate the quality of the predicted displacement fields. The mean of the three CNR values are 7 and 8.43 for our unsupervised model and OVERWIND, respectively. The average SNRe values are 0.31 and 0.7 for our

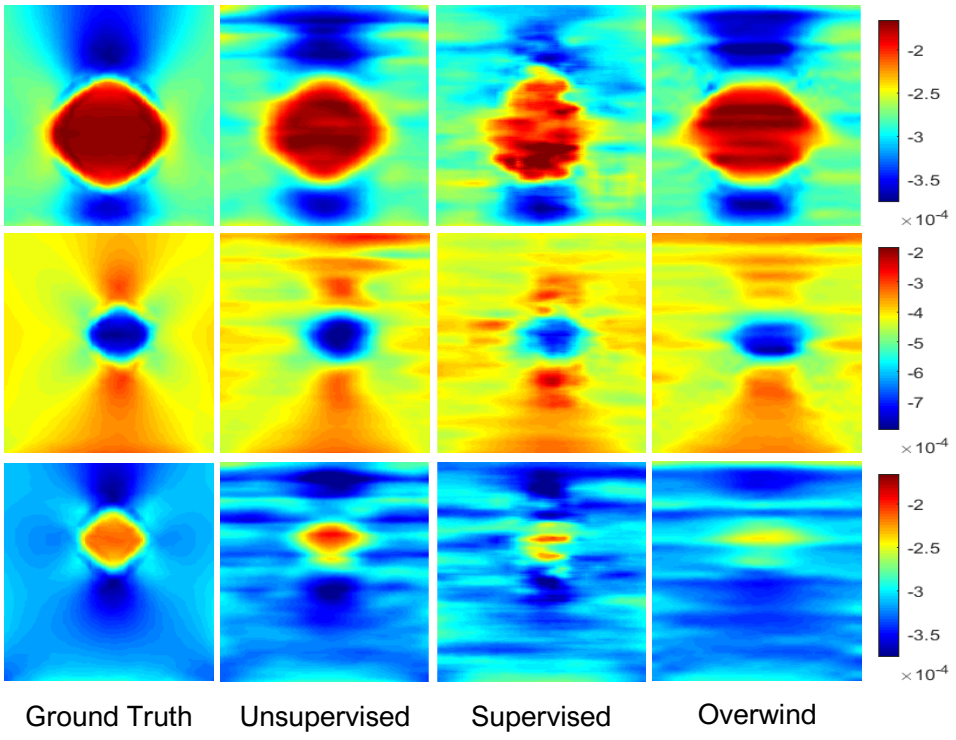

Figure 3.5: Comparison of three axial strain fields computed from ultrasound simulations by finite element method (i.e., ground truth), OVERWIND, the unsupervised and supervised network. Each row corresponds to a different simulation. OVERWIND's regularisation parameters : $\alpha 1 = \beta 1 = 20$ and $\alpha 2 = \beta 2 = 8$.

method and OVERWIND, respectively.

Table 3.2: SNRe and CNR of strain estimates from our unsupervised method and OVERWIND, and local NCC scores for the three image pair results from the *in vivo* dataset.

	Unsupervised			OVERWIND		
	Case 1	Case 2	Case 3	Case 1	Case 2	Case 3
SNRe	0.14	0.07	0.1	0.12	0.35	0.23
CNR	6.22	5.2	4.5	4.9	3.15	2.23
LNCC	0.87	0.81	0.94	0.90	0.81	0.93

3.4 Discussion

Our results on ultrasound simulation indicate that our method predicts strain estimates with a better CNR compared to the supervised network, with 7.70 and 4.22, respectively. This suggests that our training loss function, which includes a term that penalises the strain smoothness, improves the strain contrast. Our exper-

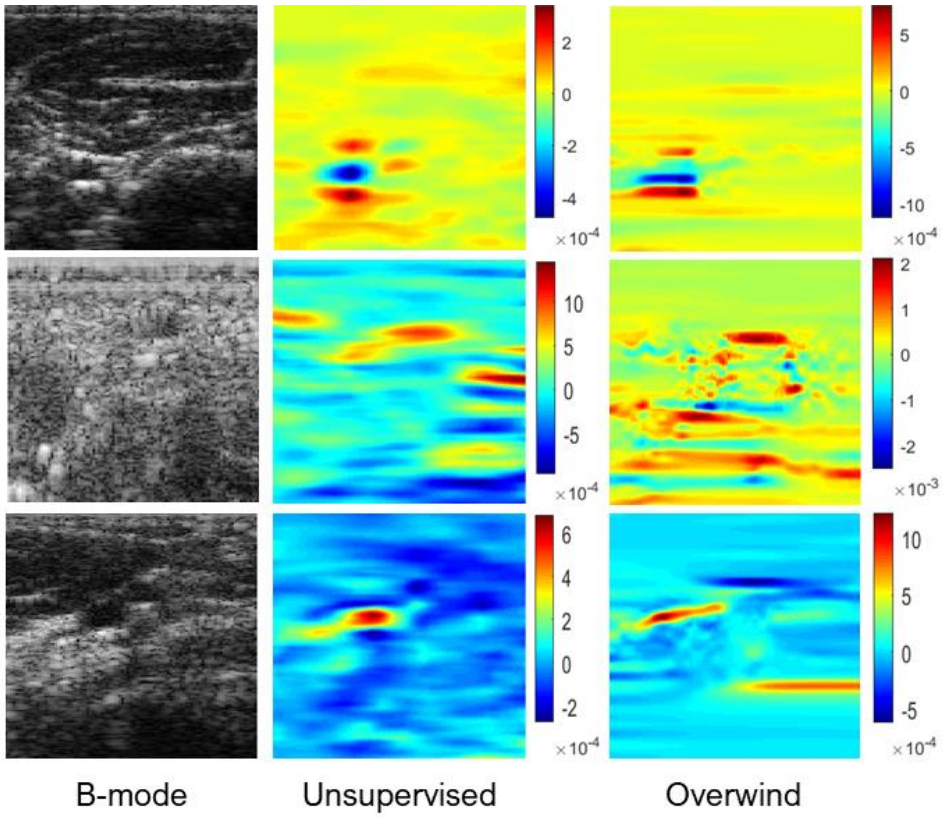

Figure 3.6: Comparison of axial strains estimated by our method and OVERWIND on three image pairs from the *in vivo* testing dataset. OVERWIND's regularisation parameters : $\alpha 1 = \beta 1 = 0.2$ and $\alpha 2 = \beta 2 = 0.05$.

iments showed comparable results to OVERWIND, a state-of-the-art method that has, in terms of CNR, already outperformed a previous classical approach [130] by 27.26%, 144.05%, and 49.90% on average in simulation, phantom, and *in vivo* data, respectively, as reported in [141].

The OVERWIND method predicted strain maps with a higher average SNRe score than USENet for the testing dataset. The high SNRe score associated with a lower average CNR, as compared with our method, suggests that the strain maps produced by OVERWIND were over-smoothed (see Fig.3.6). This can be explained by the choice of OVERWIND's regularisation parameters, i.e., α and β, which affect the smoothness of the predicted displacement field in the axial and lateral directions, respectively. Optimising OVERWIND's regularisation parameter for each case may be possible, but can be too time-consuming for real-time applications. Therefore, those parameters were set to $\alpha = 2$ and $\beta = 0.1$ for the numerical simulation dataset,

and $\alpha = 2$ and $\beta = 0.1$ for the *in vivo* data, by visually inspecting the resulting displacement fields. However, these results could have been further improved by automatically searching for the best regularisation parameter for each image pair, by optimising one or more metric scores such as CNR and SNRe.

Furthermore, the OVERWIND real-time performance had not been quantitatively reported, while our approach reached a strain prediction rate of about 13 frames per second on a 12 GB NVIDIA GTX-1080Ti GPU. The real-time inference of our network and its ability to be trained without ground truth labels represent great potential for the use of learning-based methods in ultrasound strain elastography.

3.5 Conclusion

In this chapter, we presented USENet, a new deep-learning approach for the estimation of the displacement and strain maps between a pair of ultrasound RF data undergoing axial compression. We validated our method on both ultrasound simulation and *in vivo* data acquired from the forearm of 3 human volunteers. Our method is completely unsupervised, and ground truth images collected from finite element analysis were only used to assess the performance of our method. Our results on simulation and *in vivo* data suggest that USENet is capable of estimating an accurate displacement and strain field between a pair of RF ultrasound data. Moreover, the CNR and SNR scores of our strain estimates suggest that penalising the norm of the displacement second derivative results in smoother elastograms.

Finally, our method only considered a pair of RF ultrasound data to estimate the tissue strain. However, diagnostic ultrasound systems typically feature real-time imaging at a high sampling rate (up to 100 frames per second). Although ultrasound strain elastography can estimate tissue stiffness with only two images, the resulting displacement is not always sufficient to obtain a relevant strain estimate. A non-uniform or small axial compression occurring between an image pair can greatly affect the signal-to-noise ratio and result in a strain map that does not effectively characterise the tissue stiffness. The following chapter will focus on extending our unsupervised training strategy to incorporate spatio-temporal consistency. We

proposed to use a recurrent neural network and showed that using a temporal stack of ultrasound images improves the strain prediction of the targeted tissue.

Chapter 4

Improving spatio-temporal consistency of strain estimates using recurrent neural networks

This chapter is built upon our unsupervised learning method for quasi-static elastography presented in chapter 3. Here, we propose to improve displacement estimation and spatio-temporal continuity between time series ultrasound frames in elastography by making use of recurrent neural networks. This work further explores existing challenges in quasi-static elastography in section 4.1 such as frame pair selection and temporal sampling. A novel recurrent network architecture and training strategy is presented in section 4.2. The performance of this new method was assessed on numerical simulation and *in vivo* time series in section 4.3. Finally, a discussion of this work is available in section 4.4. This work is adapted from our *Physics and Medicine in Biology* journal paper named 'An unsupervised learning approach to ultrasound strain elastography with spatio-temporal consistency' and published in July 2021 [160].

4.1 Introduction

4.1.1 Background

Typical ultrasound systems feature a high sampling rate (up to a 100 frames per second), which gives access to a considerable amount of data. The increase of

computing power, which has become available through the use of GPUs, has enabled the real-time processing of ultrasound data frames for strain estimation. Quasi-static elastography only requires two image frames to estimate the strain modulus, but the resulting information is not always relevant. A non-uniform or small axial compression occurring between an image pair can greatly affect the signal-to-noise ratio and result in a strain map that does not effectively characterise the tissue stiffness.

Different studies focused on finding the most suitable pair of images to be used for strain estimation. A common solution to this problem is to compute the strain between all image pairs and associate each resulting strain with a confidence score based on image similarity [161, 162] and/or tracking information [163]. In [164], the frame selection is performed before displacement estimation, by using a classifier that gives a binary decision on the suitability of the image pair for strain computation.

Finding the best image pair also means searching for the optimal interframe interval, i.e., the time interval between successive ultrasound frames, which greatly impacts the displacement estimation. A high interframe interval exacerbates decorrelation noise due to physiological motion, such as blood flow and muscle movement, which can greatly affect the performance of displacement estimation methods and reduce the quality of the resulting strain map. Therefore, this limits the range of possible image pairs in the temporal dimension for frame pairing methods [165]. Furthermore, commercial scanners can acquire images at a high frame rate and frame-pairing strategies discard a large proportion of the available data. The strain image quality can also be improved by accumulating successive displacement fields [166, 167] or by normalising strain images with an estimate of the applied stress [168].

4.1.2 Contributions

This chapter proposes a novel network architecture based on convolutional Long-Short-Term Memory units (convLSTM) [169] to improve displacement estimation accuracy for image pairs that are temporally distant, by making use of the intermediate frames. The use of intermediate ultrasound frames improved the displacement

estimation of our recurrent network for large range deformations, as well as the consistency between consecutive strain predictions. We called our method ReUSENet, which stands for 'Recurrent Ultrasound Strain Elastography Network'. An overview of ReUSENet is presented in figure 4.1. We also present a publicly available *in vivo* database which consists in 17271 RF data of blood vessels acquired on the arm of a human volunteer[1].

The training of ReUSENet is based on the unsupervised scheme used in chapter 3 but also incorporates a spatio-temporal consistency term to help the network to learn from previous predictions. At inference, ReUSENet takes a temporal sequence of RF ultrasound data as input and predicts the displacement and strain maps of consecutive image pairs by making use of the memory state of the convLSTM units captured from previous predictions. We compare the performance of ReUSENet with the standard feed-forward neural network architecture we developed in chapter 3, i.e., USENet. We validated our two models on ultrasound time series using numerical simulation and *in vivo* data, and compared our results with state-of-the-art deep learning-based and optimisation-based algorithms [130, 137]. Both networks can be run in real-time at a speed of about 20 frames per second with a standard 12 GB GPU.

4.2 Methods

4.2.1 Network architecture

4.2.1.1 ReUSENet

The network architecture of ReUSENet is presented in figure 4.2. The encoder part of the recurrent network is the same as USENet, described in chapter 3. In the decoder part, the up-sampling blocks from USENet are replaced by convLSTM units [169]. LSTMs are a type of neural network that have been designed to learn long-term dependencies and process temporal sequences of data [170]. A standard LSTM unit is composed of a memory cell c_t, also known as the internal state, and three 'gates' regulating the flow of information, i.e., the input gate i_t, output gate

[1] Open access database available on `https://www.synapse.org/InVivoDataForUSE`

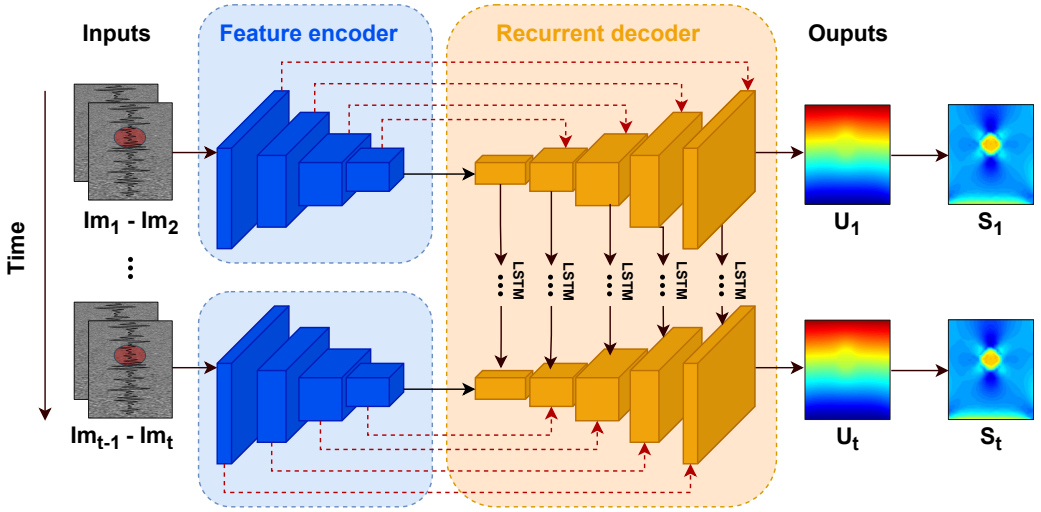

Figure 4.1: Overview of the proposed recurrent model for ultrasound strain elastography (ReUSENet). At each time step t, the network takes as inputs a pair of radio-frequency data frames (Im_{t-1}, Im_t) and outputs the dense displacement field U_t between the inputs images. The displacement field is then spatially derived from the strain field, by using the least-squares strain estimator (LSQSE).

o_t and forget gate f_t. Intuitively, the memory cell keeps track of the dependencies between the inputs of the temporal sequence, the input gate controls the incoming input flow, the forget gate controls the amount of information to keep in the cell and the output gate controls the amount of information to use for the output. The output of an LSTM is called the hidden state and is noted h_t.

A convLSTM cell differs from a standard LSTM unit by taking multidimensional data as input, such as videos. This is done by replacing the fully connected layer of each gate by a convolution operation to capture the image spatial features. The updated equations can be written as follows:

$$\begin{cases} i_t = \sigma_g(x_t * W_{xi} + h_{t-1} * W_{hi} + b_i) \\ f_t = \sigma_g(x_t * W_{xf} + h_{t-1} * W_{hf} + b_f) \\ o_t = \sigma_g(x_t * W_{xo} + h_{t-1} * W_{ho} + b_o) \\ \tilde{c}_t = \sigma_c(x_t * W_{xc} + h_{t-1} * W_{hc} + b_c) \\ c_t = c_{t-1} \odot f_t + i_t \odot \tilde{c}_t \\ h_t = o_t \odot \sigma_c(c_t) \end{cases} \tag{4.1}$$

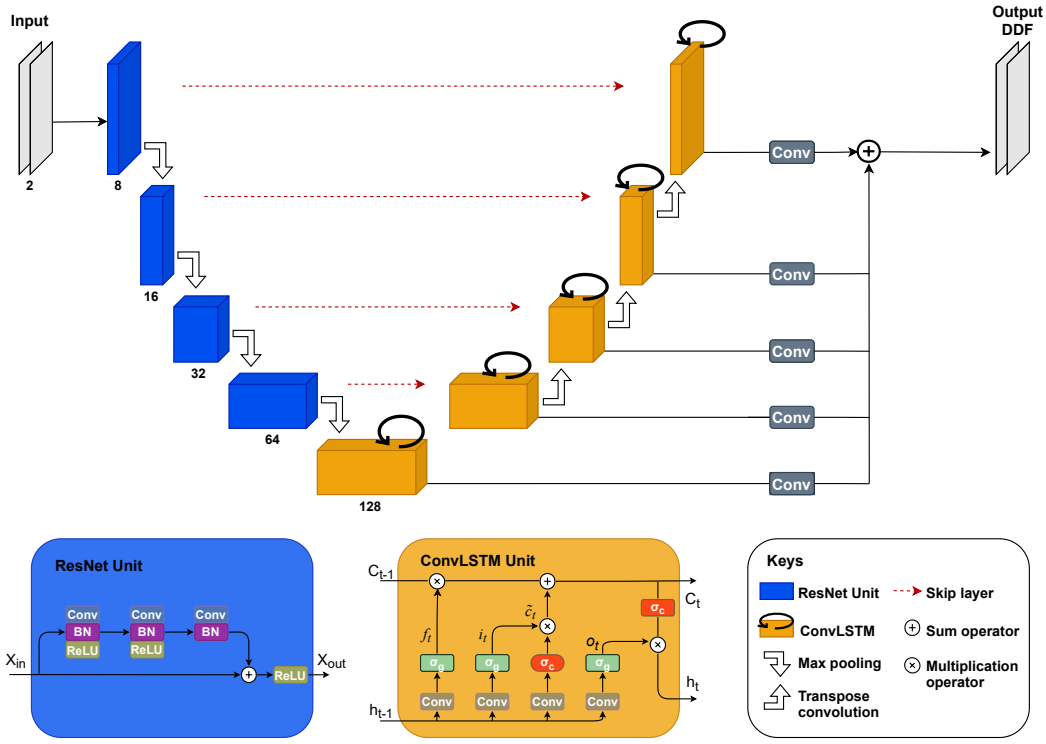

Figure 4.2: Network architecture of ReUSENet. Each blue rectangle constitutes the encoder part of the network and corresponds to a ResNet unit (illustrated below) with different channel size printed underneath. The decoder is made up of successive convLSTM units, which are represented in orange and are described in detail at the bottom of the figure. The output of each convLSTM unit is then convolved and resized to the input size before being summed to output the displacement field.

where $*$ and $\odot$ correspond to a convolution operation and an element-wise product, respectively. σ_g and σ_c are the logistic sigmoid and hyperbolic tangent functions. $\tilde{c}_t$ denotes the cell input activation vector. W_{**} and b_* correspond to the weight matrices and bias vector parameters, which are learned during training [169].

Inspired by Salvador et al., the encoded features along with the previous hidden state are fed to a convLSTM layer, which is then followed by four upsampling convLSTM blocks [171]. For a time step t, a convLSTM block i takes as input its temporal hidden state $h_{i,t-1}$ as well as the previous spatial hidden state $h_{i-1,t}$, which is up-sampled by a bilinear additive layer [172] and a transpose-convolution layer that are then added to the output of the symmetric encoding block output via a skip layer. Finally, each convLSTM block outputs a displacement field that is convolved and resized to the input size, then summed to output the predicted displacement field.

4.2.2 Training

The encoder of both ReUSENet and USENet takes a pair of pre- and post-compression 2D RF frames as input, here named *Pre* and *Post*, and the decoder predicts a dense displacement field (DDF). The parameters of our network are estimated by minimising a weighted loss function over the training set. The loss function L_{total} is composed of an image similarity term, a displacement regularisation term and a temporal consistency term which can be written as follows:

$$L_{total} = L_{sim} + \alpha L_{reg} + \beta L_{cons} \tag{4.2}$$

where L_{sim}, L_{reg} and L_{cons} are the similarity, regularisation and consistency terms, respectively. The terms α and β are the hyper-parameter used to specify the weight of the regularisation and consistency terms, respectively.

For any given training pair, the L_{sim} term is chosen as a negative local normalised cross-correlation (LNCC) function that averages the NCC score between sliding windows sampled from the pre-compression image and the post-compression image resampled with the predicted displacement field u. The NCC between two local image windows, W_1 and W_2, with i, j pixel components can be written as:

$$LNCC = \frac{1}{N} \sum_{i,j} \frac{[W_1(i,j) - \mu_{W_1}] \cdot [W_2(i,j) - \mu_{W_2}]}{\sigma_{W_1} \cdot \sigma_{W_2}} \tag{4.3}$$

where N is the number of pixels indexed by location (i, j) and μ and σ correspond to the mean and standard deviation of the images, respectively.

Given the LNCC, the similarity loss L_{sim} can be expressed as:

$$L_{sim} = LNCC(Pre, Post \circ T) \tag{4.4}$$

where $Post \circ T$ corresponds to the spatial transformation predicted by the network and applied to the post-compression image to map it in the pre-compression image space.

The regularisation term corresponds to the L1-norm of the strain spatial gradient.

Given that the strain modulus is defined as the displacement gradient, the strain field gradient corresponds to the second-order derivative of the predicted displacement and can be written as follows:

$$L_{reg} = \sum_{i,j} (\mid \partial_x^2 u_{i,j} \mid + \mid \partial_x \partial_y u_{i,j} \mid + \mid \partial_y^2 u_{i,j} \mid + \mid \partial_y \partial_x u_{i,j} \mid) \tag{4.5}$$

where u is the predicted axial displacement field and $\partial_x^2 u$, $\partial_x \partial_y u$, $\partial_y^2 u$ and $\partial_y \partial_x u$ are the second-order partial derivatives of u.

After displacement estimation, the axial strain map is computed directly during training. In quasi-static elastography, the strain estimates are obtained by computing the displacement field gradient. However, direct differentiation of the displacement field is rarely used because gradient operations generate a significant amount of noise in the resulting strain map. We used the least-squares strain estimator (LSQSE) to improve the elastogram signal-to-noise ratio [149].

Similar to our similarity loss, the strain consistency term computes the negative LNCC score between successive strain fields computed from a temporal sequence. It compares the current strain field S_t with the previously computed strain field S_{t-1}, where $S_t \circ T$ represents S_t warped by the predicted transformation T. Since the strain image is formed on the physical grid of the post-compression image, the same spatial transformation is used to perform the strain image mapping.

$$L_{cons} = LNCC(S_{t-1}, S_t \circ T) \tag{4.6}$$

Our consistency term is inspired by previous work, where it has been used as a metric to estimate the consistency between consecutive strain frames [161]. Jiang et al. motivated the use of this consistency metric by assuming that the noise in the strain image is uncorrelated with its underlying signal. Therefore, they suggest that a high correlation score between consecutive motion-compensated strain images indicates a relatively low noise level and, consequently, an improved image quality. This term is used only for the recurrent network, which deals with consecutive image pairs. Therefore, β is set to zero when training the USENet.

The presented method was implemented in PyTorch[2] and the following experiments were performed using a 12 GB NVIDIA GTX-1080Ti GPU. The network weights for both USENet and ReUSENet were fine-tuned independently for the numerical and *in vivo* databases. During training, the learning rate was initialised to 10^{-3} and was reduced by a factor of 0.8 when the validation loss stagnated for 10 epochs. The training was stopped when the difference between the new and previous learning rate was smaller than 10^{-8}. The regularisation loss weight was empirically set to $\alpha = 5$, while the consistency weight was set to $\beta = 0.2$ for ReUSENet. In inference, the strain map prediction rate reached a total of 20 images per second.

4.3 Experiments

4.3.1 Experiments on numerical phantoms

We first performed a quantitative comparison on numerical simulations of both USENet and ReUSENet together with two state-of-the-art elastography methods, namely RF Modified Pyramid, Warping and Cost volume Network (RFMPWC-Net) [137] and GLobal Ultrasound Elastography (GLUE) [130]. GLUE is an optimisation-based approach that relies on a regularised cost function to perform displacement estimation. We used the public Matlab implementation of GLUE to compute our results. RFMPWC-Net corresponds to a modified version of the well-known optical flow network PWC-Net [136]. We used the publicly available demo code and trained weights of the RFMPWC-Net for comparison. The network weights have been fine-tuned in a supervised way using an ultrasound simulation database the authors made publicly available, 'Ultrasound simulation database for deep learning" [137][3].

For reproducibility, we used the same ultrasound simulation database to train both USENet and ReUSENet. The database consists of 24 different phantoms with 10 different average strain values (from 0.5 to 4.5%) and 10 different simulations

[2]Fine-tuning code and pre-trained models are available at `https://github.com/RemiDelaunay/DeepUSE`

[3]The ultrasound simulation database, GLUE and RFMPWC-Net are available at `https://users.encs.concordia.ca/~impact/`

with different scatterer positions, which results in a total of 2400 simulated images. The displacements were obtained by Finite Element Methods using the ABAQUS software. The ultrasound images were simulated with a centre frequency of 5 MHz by using the publicly available field-II Matlab toolbox. Each digital phantom contains one or two inclusions with random positions and Young's modulus (from 40 to 60 kPa). The first 20 numerical phantoms were used for training, whereas the last four were used for testing. To mimic a temporal stack of ultrasound imaging data, the testing dataset consisted of sequences composed of 10 ultrasound images with increasing axial compression, i.e, with an average strain value ranging from 0.5 to 4.5%. All the numerical phantoms contained 10 different ultrasound simulations with different scatterer positions, except for the last one, which only had 6. Therefore, the testing dataset consisted of 36 sequences of 10 images.

During training, the entire sequence was fed to ReUSENet at each iteration, while the interframe interval was randomly assigned for USENet. Strain image quality was assessed in terms of normalised root-mean-squared error (NRMSE) and SNR (see Eq. 3.9). Both similarity and consistency scores are also displayed in figure 4.4. NRMSE corresponds to:

$$NRMSE = \sqrt{\frac{\sum_i^N (Predicted_i - Label_i)^2}{N}} \cdot \frac{100 \cdot N}{\sum_i^N (Label_i)} \tag{4.7}$$

where *Predicted* and *Label* are the axial displacement of the evaluated method and the ground truth label, respectively.

An example strain image of a simulated phantom with 1.5% of the average strain computed by the compared methods is shown in figure 4.3. Figure 4.4 shows the different metric score values plotted against the relative deformation (in % of strain) for the different methods. The shaded plots correspond to the average scores (in bold) with 25^{th} percentiles (shaded). Both feed-forward neural networks (RFMPWCNet and USENet) failed to compute an accurate displacement field for large compression. The NRMSE, SNRe, similarity and consistency scores dropped significantly after an average axial strain of 3.5% for USENet

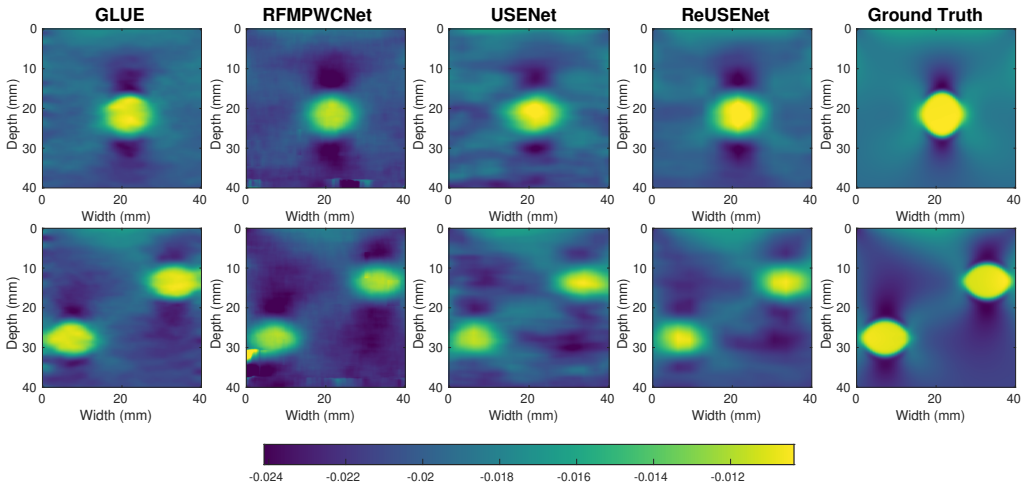

Figure 4.3: Strain images of two numerical phantoms with one inclusion (top row) and two inclusions (bottom row).

and 1.5% for RFMPWCNet. On the contrary, both ReUSENet and GLUE provide consistent and accurate results for all compression levels. Although GLUE has the lowest variance, average scores for ReUSENet (similarity=0.98 ± 0.001, consistency=0.98 ± 0.003, NRMSE=1.02 ± 0.05, SNRe=9.19 ± 1.10) are similar or slightly better than GLUE (similarity=0.98 ± 0.001, consistency=0.98 ± 0.001, NRMSE=1.10 ± 0.035, SNRe=8.87 ± 1.046). Finally, an example of a temporal axial strain estimation sequence (from 0.5% to 4.5% strain) computed with USENet and ReUSENet is compared to ground truth simulations in figure 4.5 to illustrate the degradation of performance of USENet with increasing compression.

4.3.2 Experiments on *in vivo* human data

The dataset we used in the following experiment consisted of image sequences acquired from the arm of a human volunteer. Data collection was approved by the King's College London Research Ethics Management Application System, ref HR-18/19-8881. The data were acquired by imaging the volunteer arm while slowly applying axial compression with the handheld ultrasound probe. We acquired raw channel ultrasound data from a Cicada 128PX system equipped with a 7.5 MHz linear probe from Cephasonics (Cephasonics Inc., USA). The images were generated using the delay-and-sum beamformer from SUPRA [173].

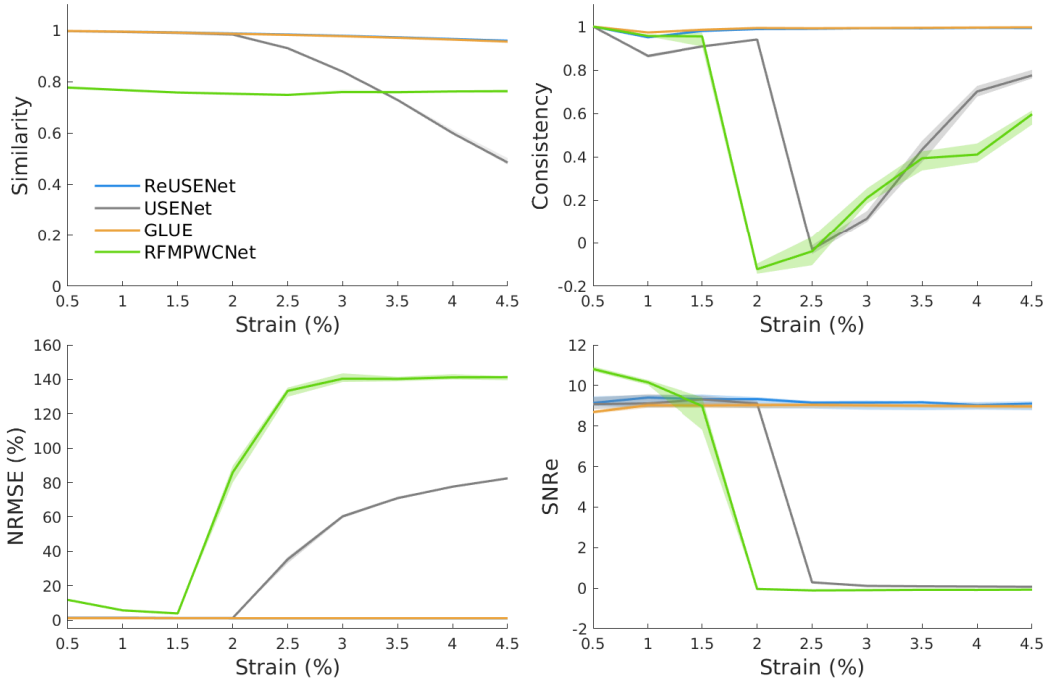

Figure 4.4: Similarity, consistency, SNRe and NRMSE scores with 25th percentiles of the compared methods for the testing simulation dataset (N=36) according to strain (in %).

The *in vivo* dataset included 310 sequences of variable length, i.e., from 19 to 127 images, for a total of 17271 images. The large variance in the image sequence size can be explained by the frame rate, which varied from 10 to 20 fps, and the time it took to perform the axial compression. Not all sequences exhibit a specific target region with a notable difference in stiffness. For example, a sequence can only show longitudinal muscle fibres of the forearm being compressed. We also decided to keep sequences with a large amount of lateral displacement or decorrelation noise for training our networks. However, each sequence selected for testing targeted at least one blood vessel. The *in vivo* dataset generally exhibits higher displacement and decorrelation noise between each frame compared to the simulated dataset. Therefore, the temporal sequences used as input for ReUSENet during training and inference corresponded to 6 successive frames. A total of 20 sequences of 6 images were used for testing, sampled from 13 different acquisition sequences.

The quality of the strain estimates were assessed in terms of consistency, similar-

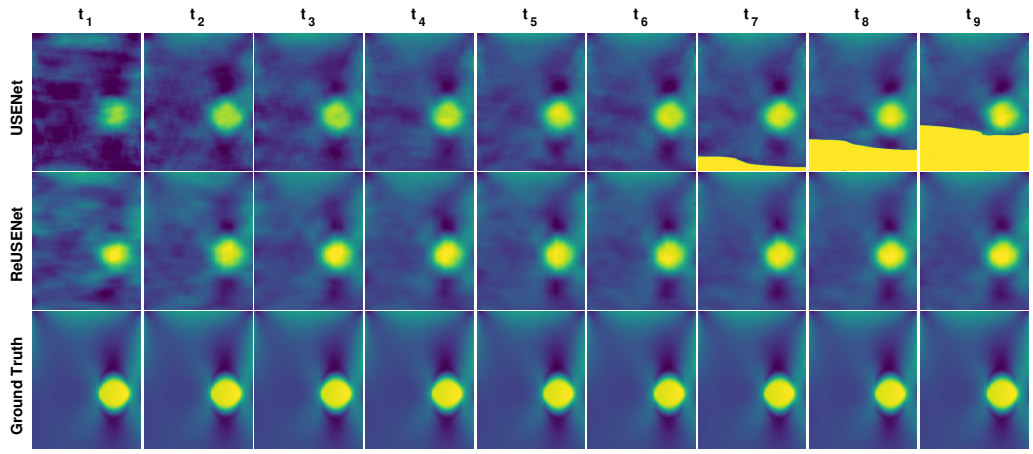

Figure 4.5: Strain images of a numerical phantom temporal sequence with increasing deformation (from 0.5% to 4.5% of strain).

ity, and SNRe. Since ground truth labels were not available, we further investigated the registration accuracy by computing a target registration error (TRE) in each case from the *in vivo* dataset. We manually identified 8 pairs of different corresponding landmarks between the first and last ultrasound frames in each temporal sequence. The TRE was then computed before and after resampling the last ultrasound frames of the sequence into the first one, by using the output displacement field predicted by USENet, GLUE and ReUSENet. The mean TRE for each case, measured in pixels, is summarised in table 4.1. In order to avoid unfair comparison, RFMPWCNet was not included in this performance comparison because the network's weights were not fine-tuned on the *in vivo* dataset, unlike USENet and ReUSENet.

Figure 4.6 and figure 4.8 both show an example of strain estimates of a temporal sequence computed by USENet, GLUE, and ReUSENet. The white arrows indicate a blood vessel that is also visible in the B-mode image. The similarity, consistency, and SNRe scores are represented in figure 4.7. Similar to the numerical phantom experiment, the consistency, similarity, and SNRe scores for the USENet gradually decrease as the interframe interval increments. Average scores for ReUSENet (Similarity=0.92 ± 0.03, Consistency=0.96 ± 0.04, SNRe=0.96 ± 0.2) are better than GLUE (Similarity=0.77 ± 0.07, Consistency=0.92 ± 0.10, SNRe=0.88 ± 0.25). The mean TREs for individual cases, measured in pixels, are summarised in table 4.1.

The average TRE and standard deviation for the entire testing dataset are the lowest for ReUSENet (2.87 ± 1.31), as compared with GLUE (3.80 ± 1.44) and USENet (4.37 ± 2.03).

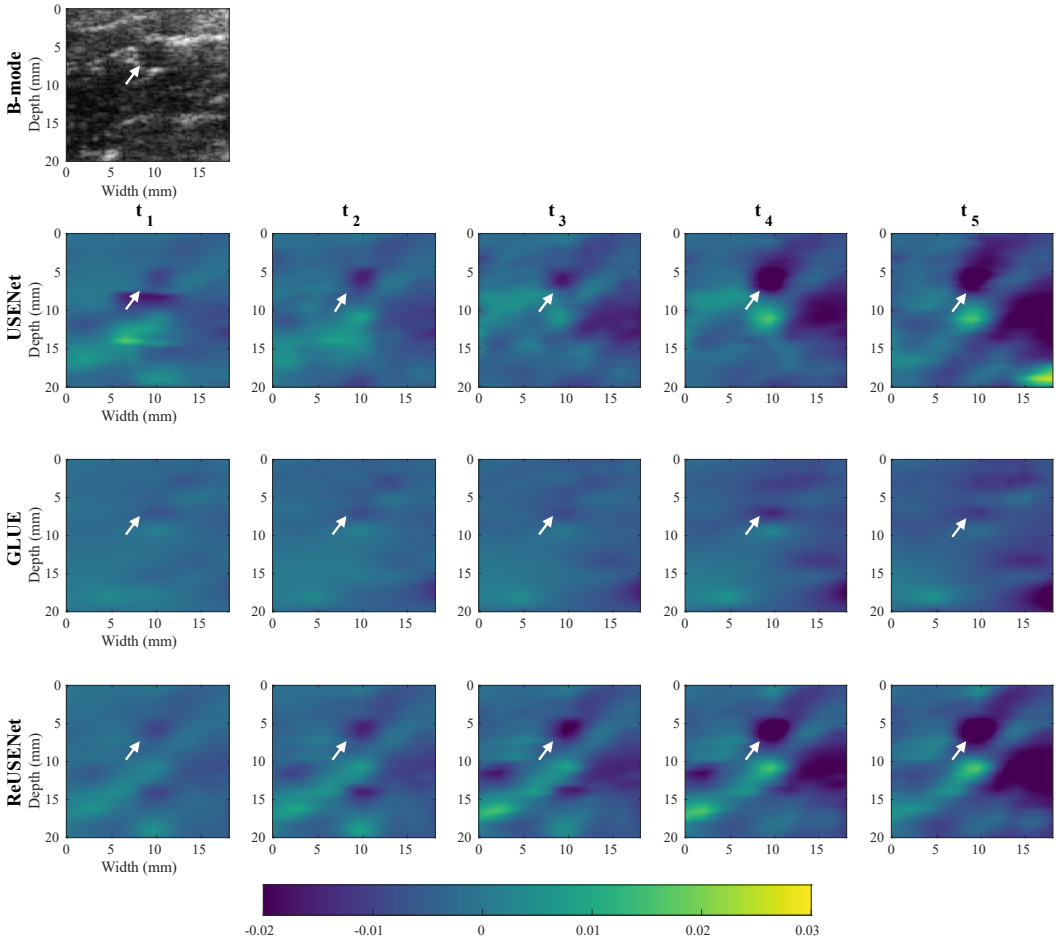

Figure 4.6: Temporal strain image sequence from the testing *in vivo* dataset.

4.4 Discussion

In chapter 3, we validated our method on relatively small deformations (up to 4% of strain). Here, our results suggest that incorporating temporal continuity by using convLSTM units improves displacement accuracy, especially for larger deformations. Experiments on numerical phantoms have highlighted the poor performance of standard feed-forward networks, such as USENet and RFMPWCNet, to estimate large range deformations. Indeed, they failed to estimate displacement fields for strain levels higher than 1.5% for RFPWCNet and 3.5% for USENet. On the other hand,

Table 4.1: Registration accuracy of USENet, GLUE and ReUSENet. The mean TRE is calculated in pixels for all cases (8 landmark pairs per cases) from the *in vivo* testing dataset.

Case	Initial	USENet	GLUE	ReUSENet
1	6.98	5.02	4.05	3.29
2	6.48	4.08	3.12	2.57
3	10.00	4.21	3.05	1.05
4	8.62	4.15	4.42	2.48
5	3.22	4.82	2.91	3.44
6	7.65	3.09	3.07	2.18
7	3.92	3.50	3.33	4.03
8	9.91	5.05	4.91	2.61
9	9.17	4.62	3.64	3.92
10	7.50	3.39	3.24	2.80
11	8.80	3.89	2.83	1.77
12	9.90	4.88	3.61	3.55
13	8.18	2.58	2.32	2.35
14	6.83	4.64	4.68	3.62
15	8.00	2.62	4.24	1.20
16	10.45	9.16	9.00	7.45
17	4.05	2.59	2.28	2.32
18	5.58	2.15	2.66	1.85
19	5.92	2.49	3.57	2.37
20	18.06	10.40	5.12	2.94
mean	7.96	4.37	3.80	**2.87**
stddev	3.09	2.03	1.44	**1.31**

ReUSENet utilises previous predictions to accurately estimate larger deformation (up to 4.5%). To the best of our knowledge, this is the first learning-based method to quantitatively reach the reported performance on such a large displacement search range.

Our results from the *in vivo* dataset showed that ReUSENet exhibited higher scores than USENet and GLUE in terms of SNRe, similarity, and consistency. Most interestingly, the performance gap increased with the interframe interval, which may suggest that the recurrent network did make use of previous memory states to predict the current displacement. The TRE results also suggest that ReUSENet performs better than GLUE and USENet in terms of registration accuracy. Most interestingly, the performance gap increases with the interframe interval, which suggests that the sequential information enabled by the recurrent network improved the displacement

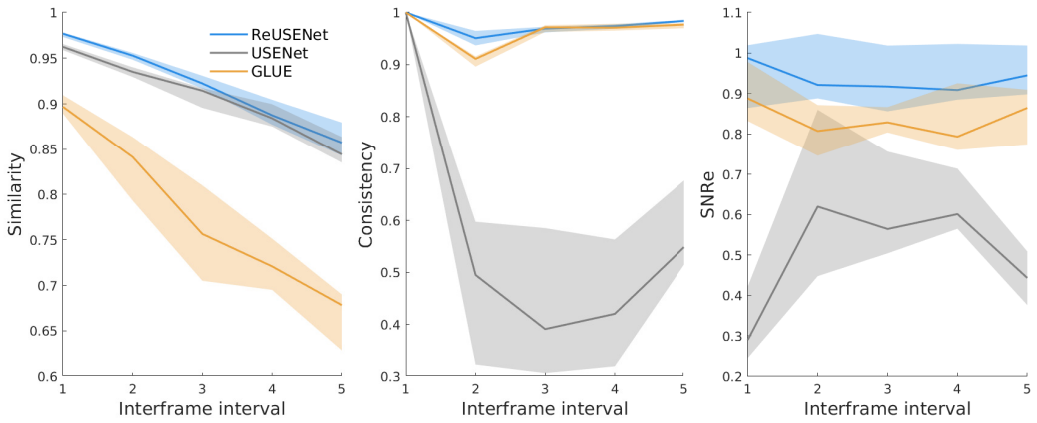

Figure 4.7: Similarity, consistency and SNRe scores with 25th percentiles for the *in vivo* testing dataset (N=20) according to the interframe interval.

estimation. It is important to note that the results from GLUE, RFPWCNet and USENet could have also been improved by applying the intermediate displacement fields at each time step. There is however no published best practice on how to exploit temporal context with these methods. We thus considered such possible extensions as out of scope for this work, and only compared with the publicly available implementations of GLUE and RFPWCNet.

We found that GLUE was sensitive to its regularisation parameter, α and β, which control the displacement field smoothness in axial and lateral directions, respectively. We used the default parameters suggested by the authors for the simulation dataset, i.e., $\alpha = 5$ and $\beta = 1$ [130]. However, those parameters tended to over-smooth the strain field for the *in vivo* dataset, which improved significantly the SNRe but also concealed the blood vessels in our experiments. Finding the optimum parameters for each case may be possible, but can be too time-consuming for real-time applications. Therefore, we selected the same parameters ($\alpha = 2$ and $\beta = 0.1$) for the entire testing dataset by visually inspecting the collection of output rather than automatically selecting the parameters that gave the best metric scores. Automating this process for optimisation-based elastography methods would be an interesting future research direction.

In conventional scanners, the strain elastogram is usually displayed next to, or directly overlayed onto the B-mode images. Therefore, processing time and real-time

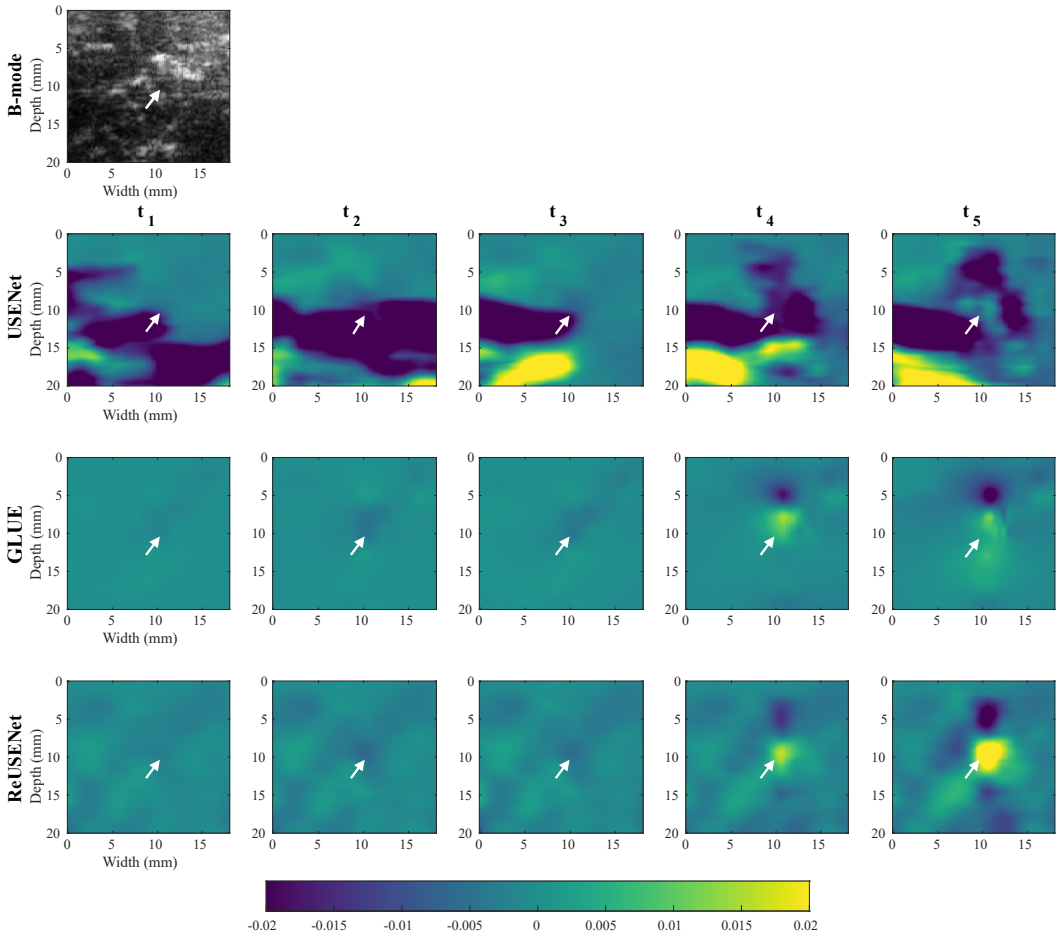

Figure 4.8: Temporal strain image sequence from the testing *in vivo* dataset.

visualisation are of high importance in quasi-static elastography. Both ReUSENet and USENet were able to achieve an inference speed of up to 20 frame-per-second (fps) on the DeepUSE Slicer extension, with a 12 GB NVIDIA GTX-1080Ti GPU. For comparison, RFMPWCNet achieved a framerate of 6 fps on the same GPU. The inference speed difference between ReUSENet and RFMPWC-Net can be partly explained by the input size, i.e., RFMPWCNet takes as input 3-channel images, but also the number of parameters. ReUSENet consists of 1.5 million parameters (0.8 million for USENet), whereas RFMPWCNet has approximately 9 million parameters. The MATLAB implementation of GLUE computed the strain field between one image pair in about 2 seconds on an Intel Core i7-7700HQ CPU. A GPU implementation of GLUE would significantly decrease the reported computation

time.

In addition, we have shown that the ability to incorporate temporal information in a neural network for quasi-static elastography can increase the robustness to decorrelation noise and improve displacement estimation between a pair of images that are temporally distant. We have also shown that including intermediate frames allows the recurrent network to measure larger deformation. Quasi-static elastography is highly user-dependent, and the displacement between each frame can not only be significant, but also variable, especially when the images are acquired at a high-frame rate. The use of a recurrent network that encodes the spatio-temporal information coupled with a frame-selecting method could also improve real-time visualisation.

4.5 Conclusion

In this chapter, two different neural networks were presented - USENet and ReUSENet. The former is presented in chapter 3 and is a feed-forward encoder-decoder, which takes a pair of images as input, while the latter has a recurrent architecture with decoding convLSTM units that allows a temporal RF data sequence to be used as input. Both networks were trained in an unsupervised way, which allows fine-tuning on *in vivo* data. The loss function used during training is the same as the one used in the previous chapter, except for the temporal continuity term used to train ReUSENet. The two networks were compared with a supervised network (RFMPWCNet) and a state-of-the-art optimisation-based method (GLUE).

Addressing temporal continuity in ultrasound elastography could also be of interest when using ultrafast ultrasound imaging technologies, i.e., plane wave imaging, to model fast tissue deformation [174]. This is particularly the case for shear wave elastography, which aims at tracking in time and space the propagation of shear waves generated by an acoustic radiation force. This technique has shown great promise in quantitative evaluation of soft tissue stiffness (see chapter 2), and more particularly in the characterisation of brain tumours [41]. The next chapter will focus on the evaluation of our unsupervised method for shear wave elasticity imaging.

Chapter 5

Learning-based methods for displacement estimation in shear wave elastography

This chapter investigates the use of the unsupervised learning-based approach, presented in chapter 3 and 4, for 2D shear wave elasticity imaging. Section 5.1 describes state-of-the-art methods in ARF-induced displacement estimation. Section 5.2 presents the approach we used to generate a new numerical ultrasound dataset that simulates the propagation of ARF-induced shear waves in homogeneous and heterogeneous tissue-mimicking media. Sections 3.3 and 3.4 of this chapter respectively present and discuss our results obtained with our unsupervised learning strategy and a state-of-the-art method on our numerical simulation dataset. This work was presented at the SPIE in Medical Imaging conference in February 2022 [175].

5.1 Introduction

This work focuses on shear wave elastography imaging (SWEI), where a focused ARF is emitted by the ultrasound probe and applied to the tissue (see chapter 2). The force induces small tissue displacements (about 10-40 µm), and generates quasi-planar shear waves that propagate through the tissue. The shear wave speed is relatively low in soft tissue (1-10 m/s), and is directly related to the Young's modulus (see Eq. 2.14). Ultrafast plane-wave ultrasound imaging has enabled tracking of

shear wave propagation with high temporal resolution (up to ten thousand frames per second) and reconstruction of 2D maps of tissue elasticity.

5.1.1 Previous work in ARFI displacement estimation

Shear wave estimation involves monitoring tissue motion spatially and temporally, which requires the acquisition of a large amount of ultrasound data at a high sampling rate (up to 200 frames per acquisition). Tissue motion can be monitored by calculating the displacement between a reference image, acquired when the tissue is at rest, and a set of images acquired after the induced ARF. Thereafter, the shear wave velocity is generally estimated with time-of-flight (TOF) methods, which determine the time of arrival of the shear wave at different lateral positions in the axial displacement fields [90]. Therefore, the accuracy of displacement measurement is crucial for shear wave estimation, which is challenging because of various sources of noise in the ultrasound data and the very small displacement magnitude.

Similarly to quasi-static elastography, displacement estimation in SWEI is usually performed by using a correlation-based approach on the beamformed radio-frequency (RF) ultrasound data [176]. Alternatively, optimisation-based approaches initially developed for quasi-static elastography have also been applied to transient dynamic elastography [177, 178]. Those methods make use of a regularisation parameter that penalises the correlation function by promoting the smoothness of the displacement estimates.

Another well-known method consists of calculating the shear wave propagation velocity by using a phase-shift estimator on the demodulated In-phase/Quadrature (I/Q) ultrasound data. In 1985, Loupas et al. developed a 2D auto-correlation method for the velocity estimation of ultrasound blood flow imaging [179]. The phase shift calculated between the acquired I/Q data is directly related to the displacement because of the relation between the ultrasound wavelength and the propagation speed of the ultrasound in tissue. The Loupas method is still employed in state-of-the-art shear wave elastography methods [180].

The recent advances of learning-based strategies applied to quasi-static elastography, including our work described in chapter 3 and 4, have raised new research

interest in ARF displacement estimation. Chan et al. proposed a fully convolutional neural network (CNN) trained on synthetic 3-D displacement data for displacement estimation in acoustic radiation force imaging (ARFI) [181]. Unlike SWEI, ARFI only tracks the tissue displacement across depth at a single lateral position and, therefore, does not allow tracking of the propagation of resulting shear waves (see chapter 2). Their network was trained in a supervised way, and allowed to estimate the displacement between two depth samples acquired before, and after applying the ARF.

Furthermore, Ahmed et al. developed a learning-based method for elasticity modulus reconstruction in SWEI [182]. Their method focused on the reconstruction of Young's modulus elasticity map by using FEM displacement data, and therefore, did not include a displacement estimation component. The network was also trained by using ground truth elasticity map labels.

5.1.2 Contributions

In this chapter, we investigated the use of CNNs to improve the displacement estimation accuracy for tracking shear wave propagation, directly from RF ultrasound data. We considered the two different network architectures proposed for strain elastography in chapters 3 and 4. We compared them with a conventional normalised cross-correlation method. The results presented were obtained from simulated ultrasound RF data. Our dataset consists of 150 shear wave propagation simulations in both homogeneous and heterogeneous media, which represents a total of 20,000 ultrasound images. To the best of our knowledge, this is the first study to explore the use of learning-based methods for ultrasonic shear wave tracking in SWEI.

5.2 Methods

5.2.1 Generating the numerical simulation dataset

Numerical simulation was used to train our models and assess their ability to track shear wave propagation. A publicly available finite element method package [183] was used to build our ultrasound simulation database. Soft tissue was modelled in a mesh based on hexahedral elements, with node spacing of 0.34 cm^3 and size

$35 \times 25 \times 10$ mm in axial, lateral, and out-of-plane directions, respectively. Each phantom presented a uniform background or included a spherical inclusion with variable location and radius (between 1.5-5 mm). The attenuation coefficient and Poisson's ratio were kept fixed for each phantom (0.45 Np/m and 0.495, respectively), while the background's Young's modulus varied between 15 kPa and 30 kPa, with a 1.5 to 4-fold stiffer inclusion. The distribution of acoustic pressure resulting from the ARF was simulated with Field-II [150]. The ARF was focused at a depth of 19 mm and emitted for 71 µs with a simulated linear transducer, with an excitation frequency of 7 MHz. The f-number, which corresponds to the ratio between the focal depth and the transducer's active aperture, was set to 2. The field-II pressure values were then calibrated to intensity values and converted to body-force values using the following expression:

$$F = \frac{2\alpha I}{c} \tag{5.1}$$

where α is the ultrasonic tissue attenuation (set to 0.45 Np/m), I the pulse-average intensity of the acoustic beam (in W/m^2) and c the speed of sound (set to 1540 m/s). The force field was then used as input to model the shear waves propagation. LS-Dyna (Ansys, Canonsburg, PA, USA), a commercially available finite element program, was used to solve the tissue displacement resulting from the simulated ARF. Perfect matching layers were positioned at the phantom boundaries to attenuate shear wave reflections in the medium.

The ultrasound images were generated using Field-II, by simulating an ultrasound linear transducer with 7 MHz central and 40 MHz sampling frequency. The speckle pattern typically observed in ultrasound imaging was obtained by randomly assigning a total of 30,000 scatterers per cm^3 across each digital phantom. The displaced scatterer phantoms were obtained by performing linear interpolation on the FEM displacement data. Each imaging sequence lasted 5 ms and consisted of 50 images acquired with a 10 kHz pulse repetition frequency. A total of 150 FEM simulations were performed, with only 110 numerical phantoms containing a spherical inclusion. The testing dataset is composed of 10 homogeneous and 15 inclusion phantoms. The other simulations were used for training, where three

different speckle organisations were generated for each phantom, resulting in 375 different imaging sequences for training, i.e., a total of 18750 images.

5.2.2 Training strategy

The aim of shear wave ultrasound tracking is to find the microscale displacements generated by shear wave propagation. This is done for each shear wave image acquisition by finding the set of dense displacement fields (DDF) between a fixed image I_F and a time series of moving images $I_{M(t)}$. We used the unsupervised training scheme presented in chapter 3.

The networks were trained by minimising a weighted loss function composed of an image similarity and displacement regularisation term. The image similarity term corresponds to a negative local normalised cross-correlation (LNCC) function, which averages the NCC score between sliding windows sampled from the fixed image, and the moving image resampled with the predicted displacement field. The regularisation term corresponds to the L1-norm of the second-order derivative of the predicted axial displacement field. The training loss can be written as follows:

$$Loss = LNCC(I_F, I_M \circ T) + \alpha \sum_{i,j} (|\nabla^2 u^{ax}_{i,j}|) \tag{5.2}$$

where T corresponds to the spatial transformation predicted by the network and applied to the moving image I_M to map it in the fixed image space. The second-order derivatives of the predicted axial displacement field are noted $\nabla^2 u^{ax}_{i,j}$.

The architectures of both neural networks used in this study, named USENet and ReUSENet, have been described in detail in chapter 3 and 4, respectively. They both consist of an encoder-decoder neural network with skip connections. The encoder part is the same for both networks and is composed of four downsampling ResNet blocks. The decoder part of ReUSENet is made of upsampling convolutional Long-short-term-memory (ConvLSTM) units, which take as input the encoded features as well as their temporal hidden state predicted from the previous displacement estimation. The convLSTM units are replaced by ResNet upsampling blocks for USENet. Finally, the DDF is predicted by first applying a convolution layer to the

output of each upsampling blocks, and then by summing their resized outputs.

Both USENet and ReUSENet followed the same training procedure and were implemented in PyTorch. At each time step t, a pair of fixed and moving images was fed to the network to predict the corresponding DDF. In chapter 4, we used backpropagation through time (BPTT), which consists of accumulating the network's prediction error at each time step for the entire time-series input rather than updating the network's weight at each iteration. The input size of the shear wave acquisition sequence (i.e., $1552 \times 128 \times 50$ pixels) was larger than quasi-static elastography time-series (i.e., $1552 \times 128 \times 6$ pixels). This difference in input size affected the computational load and caused exploding gradient issues. To overcome this issue, we used truncated backpropagation through time (TBTT) to update the network's weights, to reduce the computational load, and prevent from exploding/vanishing gradient problems. TBTT is a modified version of BPTT which consists of periodically updating the network's weights within a time-series for a fixed number of time steps. This allowed the network to capture the temporal structure of the time-series input while preventing the exploding/vanishing gradient problem. The regularisation loss weight and backpropagation step size were empirically set to $\alpha = 0.02$ and $\beta = 10$, respectively.

5.2.3 Time-of-flight approach

The shear wave speed was estimated from the displacement data using a time-of-flight method that computes the time lag between the adjacent lateral position via cross-correlation [107]. The main assumption for this method is that the overall shape of the shear waves is preserved as it propagates through the medium [184]. Therefore, the propagating shear wave front appearing in the displacement data should be highly correlated with the initial pulse. The correlation score C calculated between adjacent lateral displacement profile f_l and f_{l+1} is calculated as follows:

$$C_l(j) = \frac{\sum_{i=1}^{N} f_l(t_i) f_{l+1}(t_{i+j})}{[\sum_{i=1}^{N} f_l(t_i)^2 \sum_{i=1}^{N} f_l(t_{i+j})^2]^{1/2}} \tag{5.3}$$

where t_i represents the time step of each displacement prediction. The time lag is estimated for each lateral position and corresponds to the maximum correlation function.

Since the distance between adjacent lateral positions can be obtained from the image width and resolution, the shear wave arrival time for each lateral position (i.e., time lag) allows the estimation of the shear wave propagation speed (SWS) by using the following equation:

$$SWS = \frac{d}{\Delta_t} \tag{5.4}$$

where d is the distance between the adjacent lateral position and Δ_t is the time lag calculated by cross-correlation. Thereafter, the elasticity map was reconstructed by converting the shear wave velocities into Young's modulus values, and by applying a 9×9 median filter to mitigate the effects of outliers, as described in a previous study [90].

5.3 Experiments

The performance of our method was evaluated on our numerical simulation testing dataset. We compared our learning-based approach with ground truth FEM displacement fields and the displacement estimates obtained from a classical normalised cross-correlation (NCC) approach [176]. The performance of the compared techniques was assessed in terms of CNR and SNR of the elasticity map (see Eq. 3.9). The Young's modulus MAE (See Eq. 3.8 between the prediction and ground truth estimates for the background and inclusion values was also calculated.

Figure 5.1 shows an example of shear wave propagation at different time steps (0.5, 1.0, 1.5 and 2.0 ms) estimated from the testing dataset with ground-truth displacements, NCC, USENet, and ReUSENet. The left shear wave is propagating faster than the right one in the medium, which indicates that a stiffer spherical inclusion is located on the left side of the numerical phantom.

Figure 5.2 shows three examples of elasticity map estimated from the testing dataset with ground truth displacements, NCC, USENet and ReUSENet. The first

row corresponds to a numerical phantom with homogeneous background, while the two other rows show numerical phantoms that contain a stiffer inclusion. It is important to note that SWEI does not allow to characterise tissue stiffness at the excitation location. Therefore, the Young's modulus values displayed in the focal region are not to be taken into account for stiffness evaluation.

The CNR, SNR, and Young's modulus MAE values averaged over the testing dataset for the compared methods are shown in table 5.1. USENet, ReUSENet and NCC methods have similar Young's modulus prediction accuracy, with ReUSENet exhibiting the lowest mean Young's modulus MAE inclusion value, i.e., 9.61 kPa as compared with the ground truth values. On the other hand, the NCC method had the lowest MAE for the Young's modulus of the inclusion, i.e, MAE of 2.18 kPA. Our results also suggest that ReUSENet presents a higher SNR and CNR (0.85, 3.90), as compared with NCC (0.82, 3.32) and USENet (0.66, 3.59).

Table 5.1: Mean and standard deviation of SNR, CNR and Young's modulus MAE for the elasticity maps estimated from the simulation testing dataset (10 homogeneous background and 15 with stiffer inclusion). Results obtained from ground truth (GT), NCC, ReUSENET and USENet displacement estimations are showed.

	SNR (N=25)	CNR (N=15)	YM_t (kPa) (N=15)	YM_b (kPa) (N = 25)
GT	1.25 ($\pm$0.39)	4.59 ($\pm$1.80)	-	-
NCC	0.82 ($\pm$0.21)	3.32 ($\pm$1.91)	12.56 ($\pm$15.22)	**2.18** ($\pm$1.6)
USENet	0.66 ($\pm$0.12)	3.59 ($\pm$1.81)	11.22 ($\pm$13.32)	4.36 ($\pm$3.18)
ReUSENet	**0.85** ($\pm$0.23)	**3.90** ($\pm$1.85)	**9.61** ($\pm$12.2)	3.21 ($\pm$2.25)

5.4 Discussion

Our results suggest that learning-based methods could efficiently be used to track microscale deformations in SWEI. Our approach gives performance similar to the NCC method in terms of CNR, SNR, and YM MAE values. This suggests that the ability of ReUSENet to incorporate long-term dependencies with its convLSTM units improves the displacement estimation of time-series inputs.

Our results also indicate that ReUSENet showed higher performance than USENet in terms of SNR, CNR and YM MAE scores. However, the performance

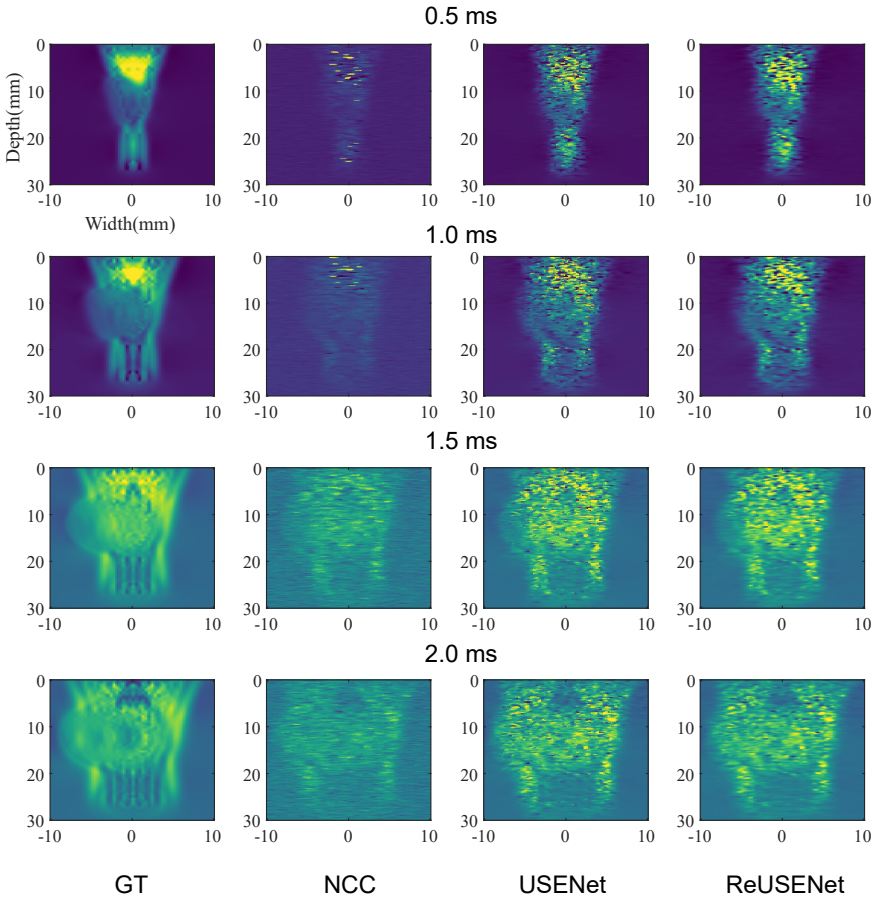

Figure 5.1: Shear wave displacement estimated in a heterogeneous phantom. Each column corresponds to a different displacement estimation method: FEM ground truth (GT), NCC, USENet and ReUSENet. Each row corresponds to the z-score normalised displacements at different time steps, from top to bottom: 0.5, 1.0, 1.5, 2.0 ms. Notice that the tracked shear waves are propagating faster on the left side, indicating the presence of a stiffer inclusion.

improvement is not as noteworthy as the one observed in chapter 4, where we showed that feed-forward networks failed to estimate large range quasi-static deformations (see Figure 4.5). In shear wave imaging, the displacement generated is of the order of a few micrometres, and the ultrasound time series are acquired with a high temporal sampling rate (about 5-10 ms). Therefore, our results suggest that both feed-forward and recurrent neural networks can efficiently track shear wave displacements, and that the use of convLSTM units is better suited for ultrasound time series that exhibit larger deformation.

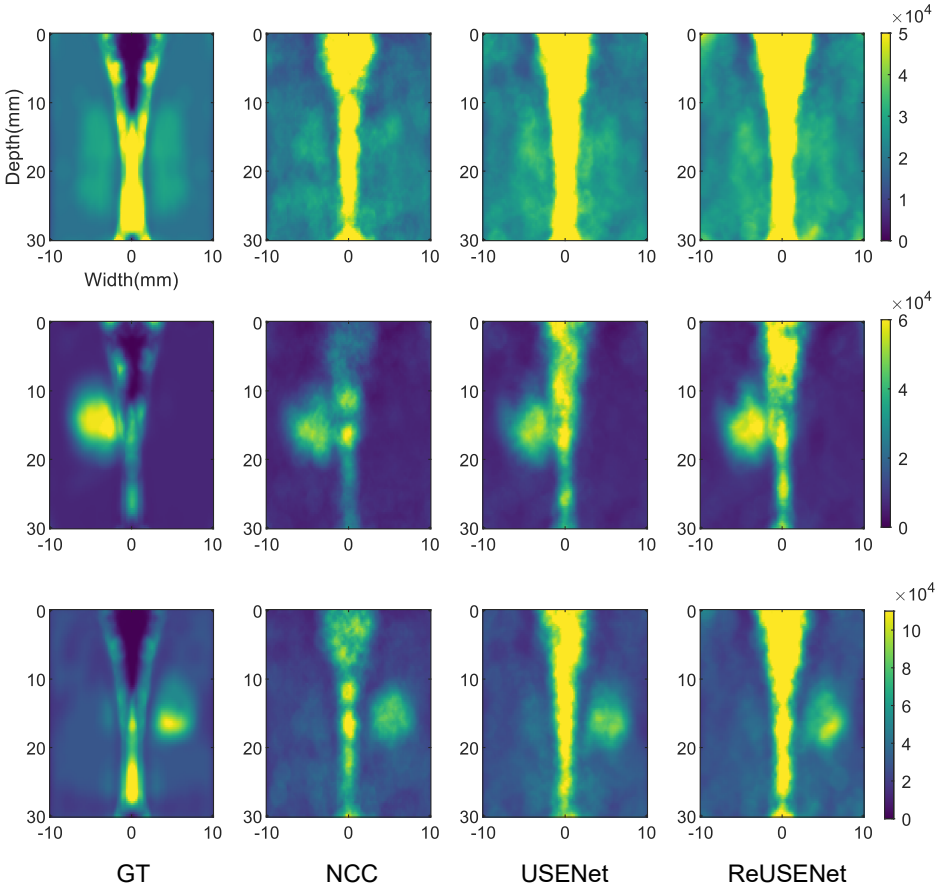

Figure 5.2: Young's modulus elasticity maps estimated from ground truth displacement, NCC, USENet and ReUSENet. Each row displays the Young's modulus map (colorbars in Pa) estimated from a different numerical phantom.

The processing times of ReUSENet and USENet have been reported in chapter 4, with a displacement prediction rate of 20 frames per second on a 12 GB NVIDIA GTX-1080ti GPU. On the other hand, the MATLAB implementation of the NCC approach processed a SWEI sequence in 7.3 s on an Intel Core i7-7700HQ CPU, which corresponds roughly to a prediction rate of 8 frames per second. However, it is important to note that a GPU implementation of the NCC approach would significantly decrease the reported computation time.

5.5 Conclusion

This chapter proposes the use of feed-forward and recurrent neural networks (i.e., USENet and ReUSENet) to track shear wave propagation in ultrasound elastography. We validated our method on a numerical data set that simulated the propagation of shear waves induced by a single-push acoustic radiation force. The training methodology is completely unsupervised, and ground-truth displacements estimated via FEM were only used for the method comparison. We also compared our method with a classical NCC approach in terms of displacement and elasticity estimation accuracy.

Last but not least, this work presents a publicly available ultrasound simulation dataset of 20,000 images[1]. Numerical ultrasound simulations are computationally expensive and, to the best of our knowledge, this is the first numerical SWEI dataset made publicly available. This dataset could be used to explore and validate further learning-based strategies for shear wave ultrasonic tracking.

[1] Open access simulation database available on `https://www.synapse.org/SimForSWEI`

Chapter 6

Integration of intraoperative ultrasound with neuronavigation

This chapter focuses on the integration of intraoperative ultrasound imaging and elastography into a neurosurgical workflow. It first introduces the concept of surgical navigation and its technical evolution in minimally invasive surgery. It also gives a brief introduction of the progress made in image-guided surgery and investigates the current software available (e.g., commercial and research open-source software) for ultrasound-guided surgery (section 6.1). The aim of this work was the design of an ultrasound-aided neuronavigation and elastography platform for clinical use. The system requirements and choice of open-source platform are presented in section 6.2. This chapter also presents the procedural assessment of the neuronavigation platform in a typical surgical workflow (section 6.3). Finally, the results of our feasibility study and its potential use in clinical settings are discussed in section 6.4.

This work is adapted from an *International Journal of Computer Assisted Radiology and Surgery* journal paper named 'Integrated multi-modality image-guided navigation for neurosurgery: open-source software platform using state-of-the-art clinical hardware', and published in May 2021 [185]. For this work, I performed the literature review on existing image-guided therapy software to determine the best suited open-source platform for our project, i.e., Slicer-IGT and PLUS toolkit. The design and development of the neuronavigation platform was a collaborative effort between our team and the UCL software development team.

My contribution was especially focused on the integration of the ultrasound system with the neuronavigation platform, and included establishing the streaming of data, the different calibration procedures, the reconstruction of volumetric ultrasound data and visualisation with preoperative scan images. The stand-alone module for elastography imaging is my own contribution and was presented separately in our *Physics and Medicine in Biology* journal paper named 'An unsupervised learning approach to ultrasound strain elastography with spatio-temporal consistency' [160]. The tissue-mimicking phantoms were fabricated by Eleanor Mackle and Dr. Jonathan Shapey. Finally, I participated in all the pre-clinical experiments presented in this chapter.

6.1 Introduction

6.1.1 Navigation in neurosurgery

Stereotactic surgery refers to surgical procedures that use a three-dimensional coordinate system to locate specific anatomical targets within the body. This term is used for a broad range of minimally invasive interventions such as needle biopsy, ablation, implantation, or stimulation. The use of stereotactic technologies has been firstly used in neurosurgery with the aim of performing safer and less invasive interventions while improving target localisation.

The first stereotactic system was developed by Leksell in 1949 [186] and consisted of an arc-centered frame to be fixed to the patient head (see figure 6.1). The target localisation was performed using a three polar coordinates system, expressed in terms of angle, depth and anterior-posterior location. However, frame-based stereotaxy presented several disadvantages. The frame attached to the patient was a source of discomfort, but also obstructed the view of the surgical field. Furthermore, once the frame was fixed, the target localisation accuracy could not be assessed intraoperatively. The advances in imaging technologies, especially with the invention of CT and MRI, and the progress in computer science allowed to overcome the limitations associated with frame-based stereotaxy.

Nowadays, typical neuronavigation systems have the ability to identify the

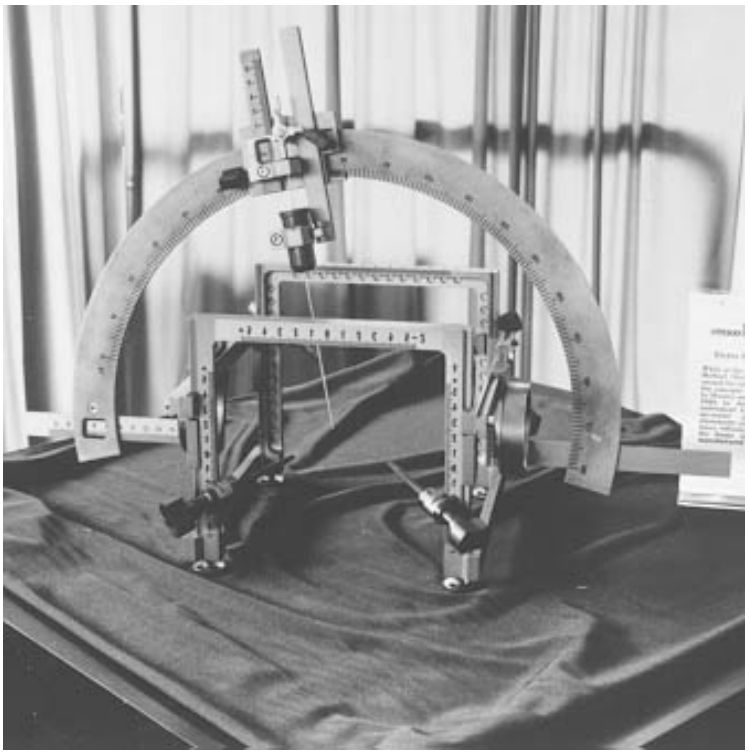

Figure 6.1: Picture of the Leksell stereotactic system. Courtesy of Lunsford et al. [186].

location of surgical tools relative to the patient position in real-time (via positional tracking), using a frameless localisation system (i.e., tracking device). The two main categories of localisation systems used in minimally invasive surgery are optical and electromagnetic tracking. Optical tracking typically consists of dual infra-red cameras that track the position of retro-reflective markers attached to an object (e.g., surgical instruments). The main limitation of optical tracking is the necessity of maintaining a direct line of sight between the infra-red camera and the tracked instruments and the reference frame. Electromagnetic tracking relies on the creation of a magnetic field by a field generator, and to measure its inhomogeneous intensities with sensors. These electromagnetic sensors are only a few millimetres in size (down to 0.3 mm). In addition, there is no need for direct line-of-sight visualisation. However, the field generator has to be very close to the operating site to ensure good localisation accuracy (e.g., working volume of $40 \times 40 \times 40$ cm). Furthermore, the use of ferro-magnetic elements can disrupt the localisation precision. The choice of localisation technology will depend on the type of intervention, but also the operator

preferences.

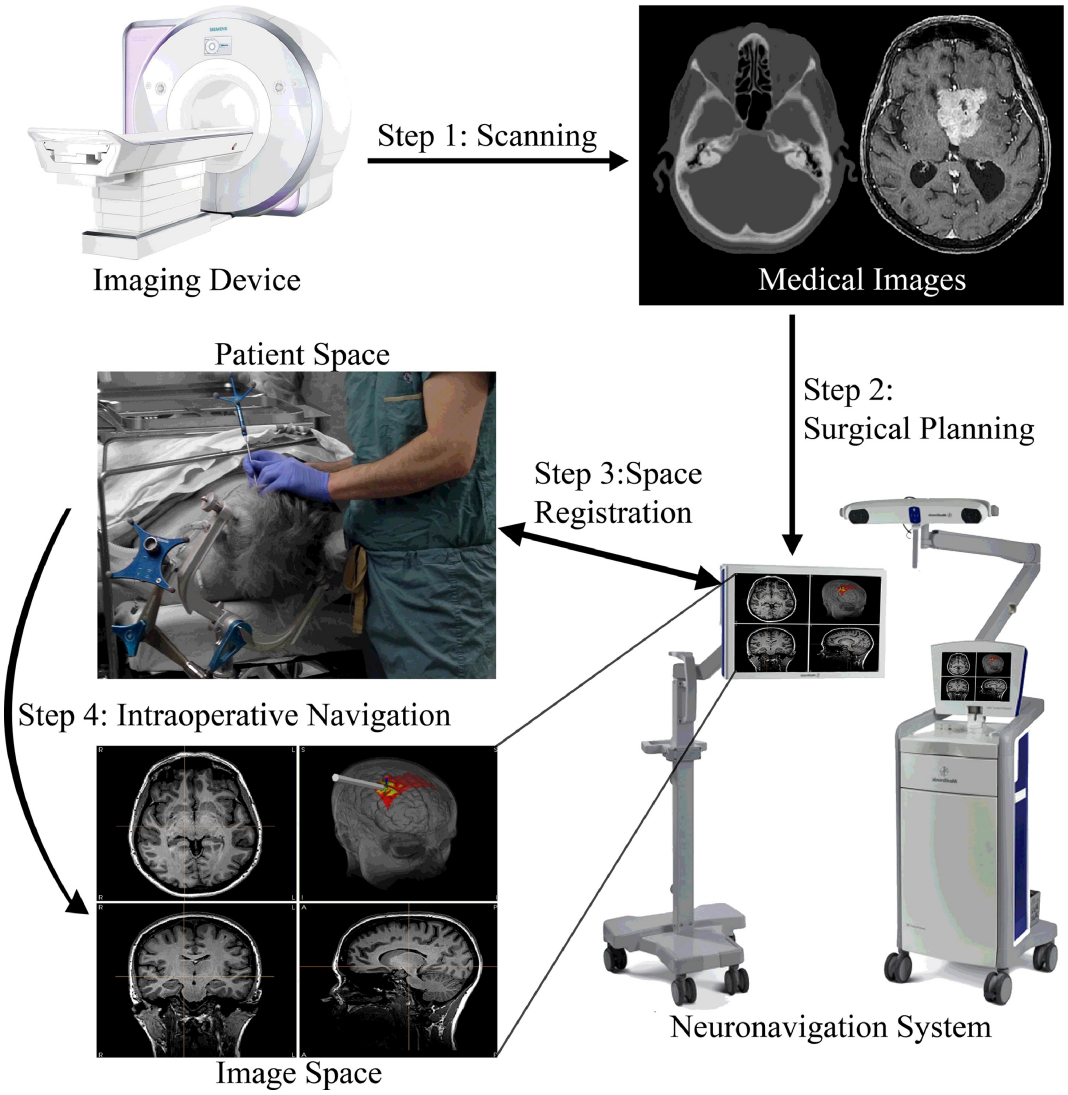

Figure 6.2: Summary of the neurosurgical workflow for brain tumour resection. Figure adapted from Gerard et al. [187]. Step 1: A scan of the patient is acquired. Step 2: The scan images are uploaded to the neuronavigation system. Step 3: Patient-to-image registration is performed on the day of surgery to map the patient's head with the scan images. Step 4: the neuronavigation system is used to guide the neurosurgeons during surgery.

6.1.2 Ultrasound-guided surgery and related work on image-guided intervention software

In neurosurgery, intraoperative ultrasound has shown promising results for real-time 3D visualisation, brain-shift adjustments and tissue characterisation [23–25]. However, despite this previous work, intraoperative ultrasound remains an under-

utilised tool in neurosurgery. Firstly, because neurosurgeons are not very familiar with US as an imaging modality and, secondly, because US is typically acquired and visualised in unfamiliar planes. Furthermore, ultrasound-based guidance for neurosurgery has not been widely used in clinics.

Intraoperative ultrasound-aided neurosurgery is not new, but the offer of commercial navigation systems is limited in this field. Only Brainlab (Brainlab, Germany) and Medtronic (SonoNav, Medtronic, USA) include an ultrasound imaging module in their neuronavigation system. However, their use is limited to a small range of commercial ultrasound devices. Sonowand [188] was the only company fully dedicated to the use of intraoperative ultrasound and navigation for minimally invasive surgery, but it has now gone out of business. On the other hand, the past 20 years have seen the emergence of different research image-guided software. These software programs aim to provide an open-source and customisable environment for a seamless integration of different medical equipment, including ultrasound imaging, and encourage the development of new image-guided applications. Those image-guided therapy software can be classified according to their flexibility and level of integration.

Different libraries have been focusing on one or different aspects of image-guided interventions and, therefore, can not be qualified as fully integrated platforms. For instance, we can note the Image-guided Surgery Toolkit (IGSTK) [189] provides different packages such as surgical navigation interface and image visualisation. One of the most popular libraries for open-source image-guided therapy software is openIGTLink [190], which consists of a network communication protocol for diverse medical equipment. The PLUS toolkit [191] is a library originally designed for ultrasound imaging research, which now also offers hardware compatibility for a multitude of medical imaging and tracking devices. Other open-source platforms aim to provide a highly-flexible environment for the development of application-specific modules such as NifTK [192], MITK [193] and 3D-Slicer [194]. Those applications offer core functionalities such as, a graphical user interface, different image visualisation and processing options, and hardware communication. In addition,

they also contain extensions specifically designed for image-guided therapy such as Slicer-IGT, NiftyIGI and MITK-IGT.

Finally, we can also find fully integrated image-guided therapy platforms such as IBIS [195] and CustusX [196]. These software programs have been designed to be translated into clinics with minimal effort. CustusX is focused on ultrasound-guided intervention, whereas IBIS has been specifically developed for intraoperative neurosurgery guidance, with an emphasis on augmented reality.

6.1.3 Contributions

This chapter describes the development of an open-source ultrasound-guided neuronavigation platform, designed to integrate with commercially-available neuronavigation and ultrasound systems. The system was assessed in terms of workflow integration and tracking accuracy using multi-modal tissue-mimicking phantoms. This chapter also presents a stand-alone 3D-Slicer ultrasound elastography module, which was build upon the methods presented in chapter 3, 4 and 5, and that could potentially be used alongside the neuronavigation platform.

6.2 System design

6.2.1 Hardware

6.2.1.1 Medtronic StealthStation

The Medtronic StealthStation is a neuronavigation platform that allows surgical navigation in real-time. The system incorporates different software, including the 'Cranial software', which is dedicated to neurosurgical interventions. The StealthStation includes both electromagnetic and optical tracking. For the purpose of this project, we decided to use optical tracking because it enables a bigger working field and a better accuracy than electromagnetic tracking. The spatial location of the ultrasound probe was recorded in real-time on the StealthStation using the SureTrack ^{TM} technology from Medtronic. It consists of a passive marker compatible with optical tracking which can be attached to an ultrasound transducer for the surgical procedure. A tracked navigation stylus was also used to perform patient-to-scan

registration.

Medtronic provides its own library, StealthLink, to extract tracking data from the StealthStation via a TCP/IP connection. It is an Application Program Interface (API) designed for research purposes only. The use of StealthLink is controlled via a licence research agreement that needs to be signed by the users to enable the server connection on the StealthStation system. For safety reasons, data can only be queried from the StealthStation, meaning that no modification can be executed on the system. The navigation information provided via StealthLink includes the coordinate transformations (i.e., the position) of the different tools in use, i.e., the tracker (SureTrack) attached to our ultrasound probe and any stylus or pointer used for registration, as well as the patient registration preoperative scan data.

6.2.1.2 Ultrasound systems

Two different ultrasound systems were used during the development of the open-source platform. The BK5000 (BK Ultrasound, UK) is a commercially available ultrasound system which provides transducers specifically designed for brain tumour surgery. The Burr-hole transducer (N11C5s, Bk Ultrasound, UK) is an ultrasound probe designed for intraoperative neurosurgery and was used for this study. The other ultrasound system available was the Ultrasonix Sonictouch Q+ (BK ultrasound, UK). This system is compatible with the custom-made ultrasound transducer that we specifically designed for pituitary tumour surgery (see chapter 7).

6.2.2 Software

Among the open-source image-guided software presented in section 6.1, Slicer-IGT appeared to be the most suitable platform for this project. This choice was driven by several considerations. One of the primary requirements was that the platform was compatible with the medical equipment used by our clinical collaborator at the National Hospital for Neurology and Neurosurgery (London, UK), for an easier translation to clinics. Slicer-IGT is fully integrated with the PLUS toolkit, which offers the largest hardware compatibility. In addition, PLUS provides a wide range of ultrasound image processing tools, including the calibration procedures used for

this study. Furthermore, Slicer-IGT is part of the 3D Slicer large community, which represented an important support to carry out this project (50+ regular contributors, 500+ papers published, 1000+ downloads/month according to Drouin et al. [195]). 3D-slicer also incorporates a wide range of modules dedicated to image processing and visualisation that is compatible with Slicer-IGT. For instance, 3D-Slicer includes a module called Slicelet, which allows creation of a simplified user interface that is optimised for a particular workflow.

The integration of all hardware and software components is illustrated in figure 6.3. The PLUS toolkit communicates via TCP/IP (using the openIGTlink protocol) with the Medtronic StealthStation and ultrasound system to stream the navigation and imaging data in real-time. PLUS provides a graphical user interface (GUI) denominated 'PLUS Server' to manage the hardware communication. It uses a 'device set configuration file' written in Extensible Markup Language (XML) to set up all the configurable parameters and enable the connection. The configurable parameters describe the device (imaging or tracking parameters), the connection, and any algorithms or applications PLUS needs to operate on the device. Images from both ultrasound systems are streamed through PLUS. PLUS supports both the Medtronic StealthStation and Ultrasonix Sonictouch systems, and is built on their respective research interfaces, respectively the Ulterius and StealthLink libraries. However, the BK5000 and its OEM research data streaming interface were not supported. Therefore, the BK surgery kit Python library [197] was used to interface the BK5000 with PLUS. Once the connection is established, the data from the ultrasound and neuronavigation systems are streamed in real-time to Slicer-IGT. The image calibration and 3D volume reconstruction are then performed in Slicer-IGT.

6.2.3 Neuronavigation Slicelet

A custom end-user interface was created using 3D Slicer in order to facilitate the use of Slicer-IGT during surgery. Slicer-IGT, which is part of 3D-slicer, can be tedious to navigate through because each step is performed on a different module. Therefore, we designed a light version of the original 3D slicer version by using the Slicelets application. A Slicelet allows to load only pre-defined 3D-slicer modules in a

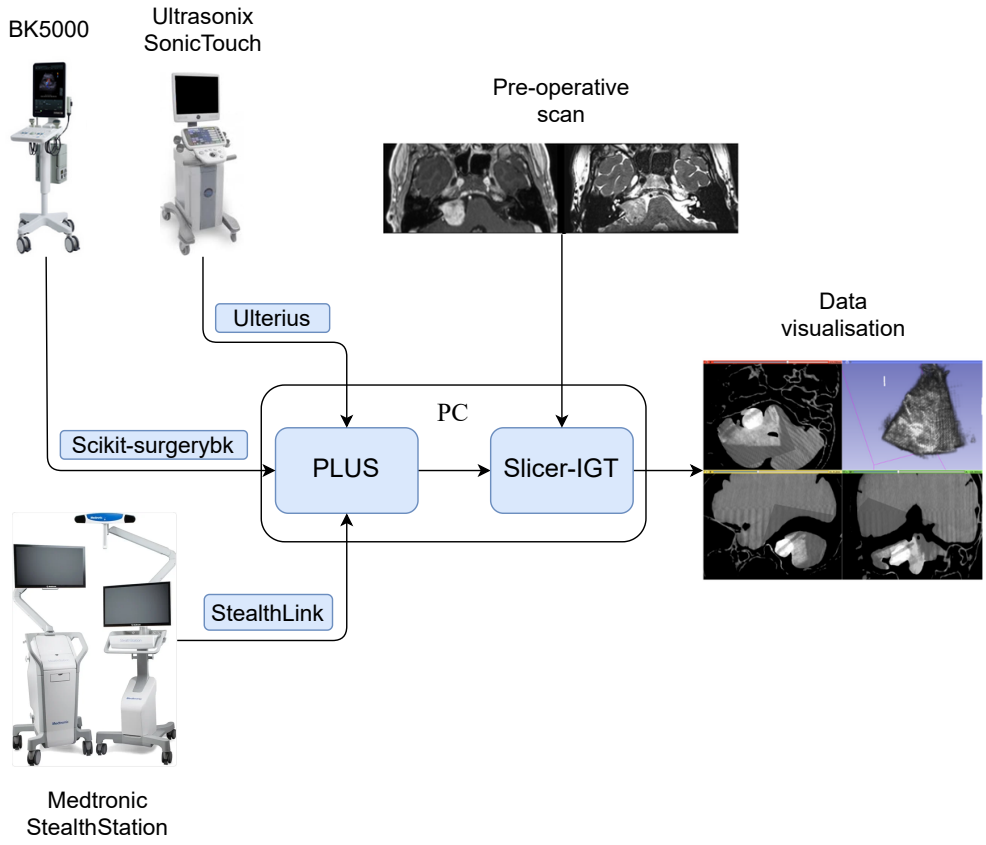

Figure 6.3: Hardware and Software communication using IP/TCP network. Data from the BK5000/Ultrasonix Touch and the Medtronic StealthStation is streamed to PLUS which enables the server connection between the different hardware and the PC. The Preoperative scan, ultrasound images and tracking data are uploaded to Slicer-IGT, which performed all the necessary processing steps: image calibration, patient-to-scan registration, ultrasound 3D reconstruction and real-time visualisation.

simplified interface and automate some processing steps (e.g., connection to different devices). We created a Slicelet module using the Python *ScriptedLoadableModule* base class provided by 3D Slicer. Features of the neuronavigation slicelet included display of preoperative MR/CT scans, real-time overlay of ultrasound data, tool locations and calibration process and volume reconstruction of ultrasound data. The Slicelet was made publicly available to allow others in the research community to test and build on this work[1].

[1]Open-source platform available at `https://github.com/UCL/SkullBaseNavigation`

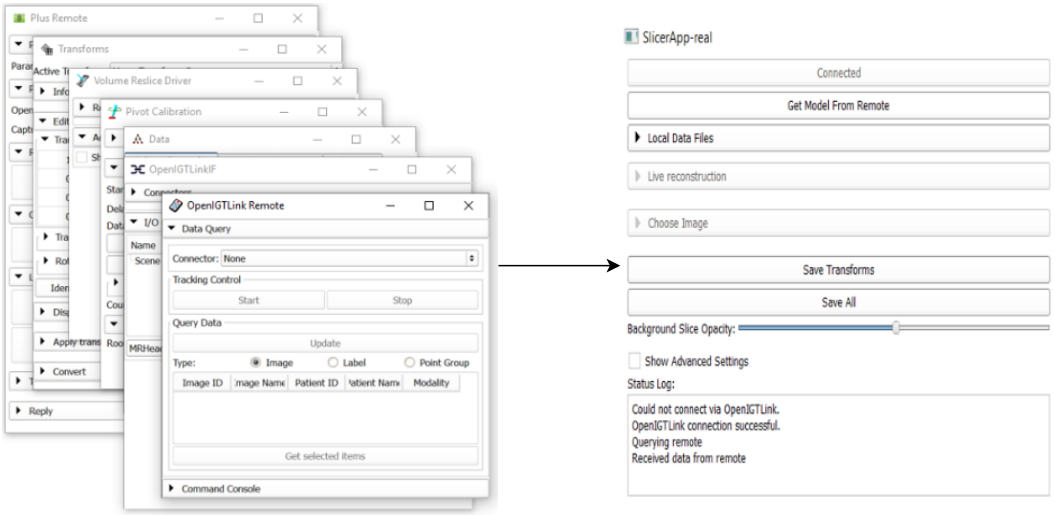

Figure 6.4: Neuronavigation module user interface. The 3D Slicer Slicelets module allows to combine different application into a simplified interface to improve the surgical workflow.

6.2.4 Stand-alone ultrasound elastography module

A 3D-slicer module was also implemented to perform intraoperative tissue characterisation with the learning-based elastography methods presented in chapters 3, 4 and 5. This stand-alone module has been named 'DeepUSE' and is also publicly available[2]. DeepUSE is written in Python using the *ScriptedLoadableModule* base class provided by 3D Slicer. The module is fully integrated with the PyTorch implementation of both ReUSENet and USENet, the two network architectures we developed for ultrasound elastography. DeepUSE's features include the loading of a trained model via a configuration file, the offline inference of a loaded RF data sequence and the real-time inference of a stream of RF data sent using the OpenIGTLink protocol [190]. In terms of visualisation, the strain image nodes are automatically displayed alongside the RF data converted into B-mode to facilitate data interpretation. Figure 6.5 shows a screenshot of the DeepUSE module, which shows a B-mode image and its respective elastogram side by side. The image displayed is from the publicly available dataset presented in chapter 4.

[2]Slicer DeepUSE module available at `https://github.com/RemiDelaunay/SlicerDeepUSE`

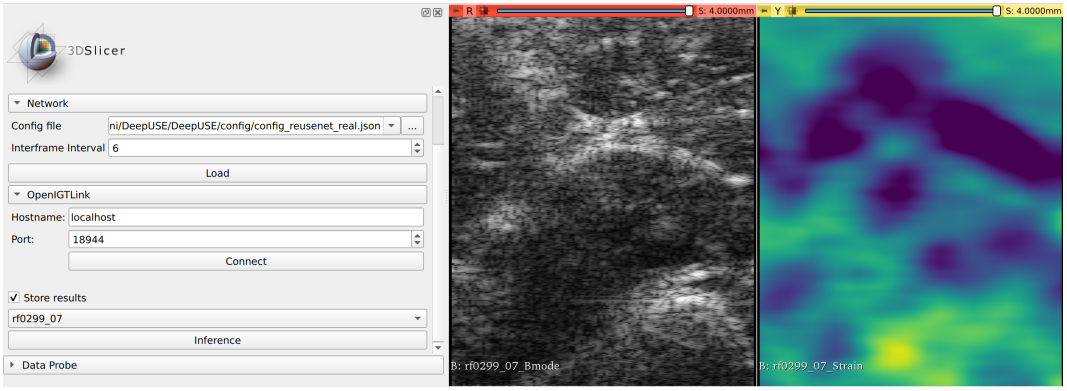

Figure 6.5: DeepUSE : Slicer ultrasound elastography module interface.

6.2.5 Calibration

Although the cranial software from the Medtronic StealthStation allows patient-to-MR registration and tool calibration, it does not provide ultrasound calibration algorithms. Therefore, all the ultrasound calibration steps were performed in Slicer-IGT.

The first calibration procedure consists of calculating the position of the tip of the ultrasound probe. Position tracking of the ultrasound transducer was calibrated at the start of each case using the pivot calibration algorithm provided as part of Slicer-IGT. Pivot calibration involved the use of the tracked stylus tool to determine the tip of each individual instrument relative to the SureTrack^{TM} marker. The calibration procedure was a two-step process that involved pivoting and then spinning the tracked instrument around a fixed point for 15 s each in turn. After pivot calibration, the RMSE between the predicted transform and the acquired points is displayed to assess the quality of the calibration. Typically, the RMSE for pivot and spin calibration should be less than 1 mm. As the tracker is securely attached to the tool, once calibration has been performed, the transformation matrix between the stylus tip and the tracker remains constant [198].

Since the tracking and ultrasound data are acquired by different devices, there can be a significant time delay between the different data streaming nodes. The streaming offset between the different devices can be estimated by performing a temporal calibration. The calibration procedure synchronises the data streams by scanning up and down a fixed object, such as the bottom of a water tank [198].

Finally, the last calibration procedure consists of computing the position of the image relative to the probe. This procedure is similar to the image-to-patient registration, and consists in finding the transformation between points on the ultrasound image and a tracked physical object, such as the position of a tracked stylus. Image calibration was performed using the stylus-aided calibration toolbox and involved imaging and registering the stylus tip in multiple locations within the ultrasound image [198]. Ultrasound image-to-probe calibration can be time-consuming and highly user-dependent. The position of the tracking stylus tip in the ultrasound image can be difficult to locate and can result in calibration errors. Therefore, image calibration was performed beforehand, and the transformation matrix was subsequently saved for use during surgery. Figure 6.6 shows the setup we used for the image calibration procedure.

Figure 6.6: Ultrasound image calibration method. A: Ultrasound transducer (attached to a clamp for illustration purposes only). B: Medtronic SureTrack ™ optical tracking marker. C: Medtronic Stylus. D: Medtronic reference frame. E: BK5000 Ultrasound machine. F: computer. G: water tank.

6.2.6 3D ultrasound volume reconstruction

We used the ultrasound reconstruction algorithm implemented in Slicer-IGT to reconstruct 3D ultrasound volumes [199]. There are two scanning steps to perform the 3D reconstruction. First, the users have to operate a 'scout scan', which consists of a low resolution reconstruction that defines the spatial boundaries of the region of

interest to image. Then, the high-resolution scan can be performed by repeatedly scanning the area of interest in real-time until the volume has a satisfactory resolution. An example of 3D reconstruction can be seen in figure 6.9.

6.3 Accuracy and workflow testing

6.3.1 Laboratory testing of system accuracy

An abstract phantom model was used to test the accuracy of the system's ultrasound navigation and 3D reconstruction. The phantom was made of polyvinyl alcohol cryogel (PVA-c), a tissue-mimicking material that is known to have similar acoustic and mechanical properties than soft tissue. The phantom consisted of two 15 mm diameter tumour-mimicking spheres embedded within parenchyma-mimicking tissue and was manufactured according to previously published methodology [200–202].

The phantom was imaged with a Medtronic O-arm^{TM} to provide a 3D volumetric x-ray contrast image for registration. Scan images were uploaded to the Medtronic StealthStation and registered with the phantom using the surface registration technique from the Medtronic Cranial software. The BK burr hole ultrasound transducer was calibrated using the method described above, and volumetric ultrasound data was acquired. Tumour spheres were segmented on the registered ultrasound and x-ray images using an intensity threshold. The system accuracy was measured in terms of target registration error (TRE) calculated between the two centres of the spheres.

6.3.2 Clinical simulation to test workflow integration

A patient-specific phantom model comprising skull, brain, and tumour created with tissue-mimicking ultrasound and X-ray properties was used to simulate the use of the navigation system in a clinical operating room. The skull was 3D printed in Polylactic acid (PLA), from an anonymised CT scan of the patient's brain. The brain tissues were made with Poly(vinyl) alcohol cryogel (PVA-c). The difference of acoustic contrast between the tumour and brain tissues was obtained by mixing different concentrations of glass spheres with the PVA-c [203]. This phantom was compatible with both CT and ultrasound imaging. Before phantom scanning, barium

sulphate, which is a clinical imaging contrast, was added to the tumour to further increase the image contrast. The time taken to set up the system and to perform probe calibration was recorded, and clinician feedback was obtained regarding the clinical utility and accuracy of the system.

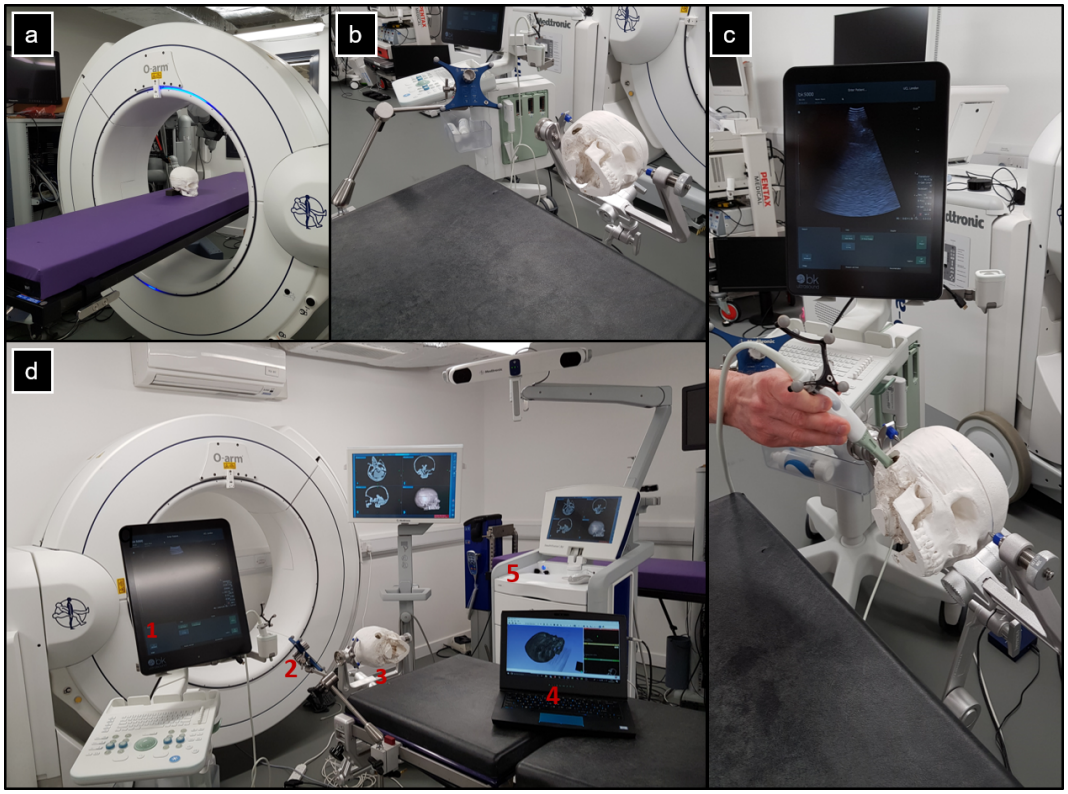

Figure 6.7: Software workflow and feasibility study images. a) Preoperative scanning with CT scan. b) Image-to-patient registration using the reference frame and patient-specific phantom. c) Intraoperative ultrasound acquisition using the BK5000 ultrasound system. d) Overview of the set-up with in 1. the BK5000, 2. the reference frame, 3. patient-specific phantom, 4. the operating room laptop, and 5. the Medtronic StealthStation.

6.3.2.1 Results

The phantom that was used to assess the system accuracy as well as the volumetric CT and ultrasound images are shown in figure 6.8. The distances between the centroids of the tumour spheres were 3.82 mm and 4.41 mm for tumour spheres 1 and 2, respectively.

The workflow integration assessment was performed in a mock operating room. The full experimental setup can be seen in figure 6.7. Prior to the clinical simulation,

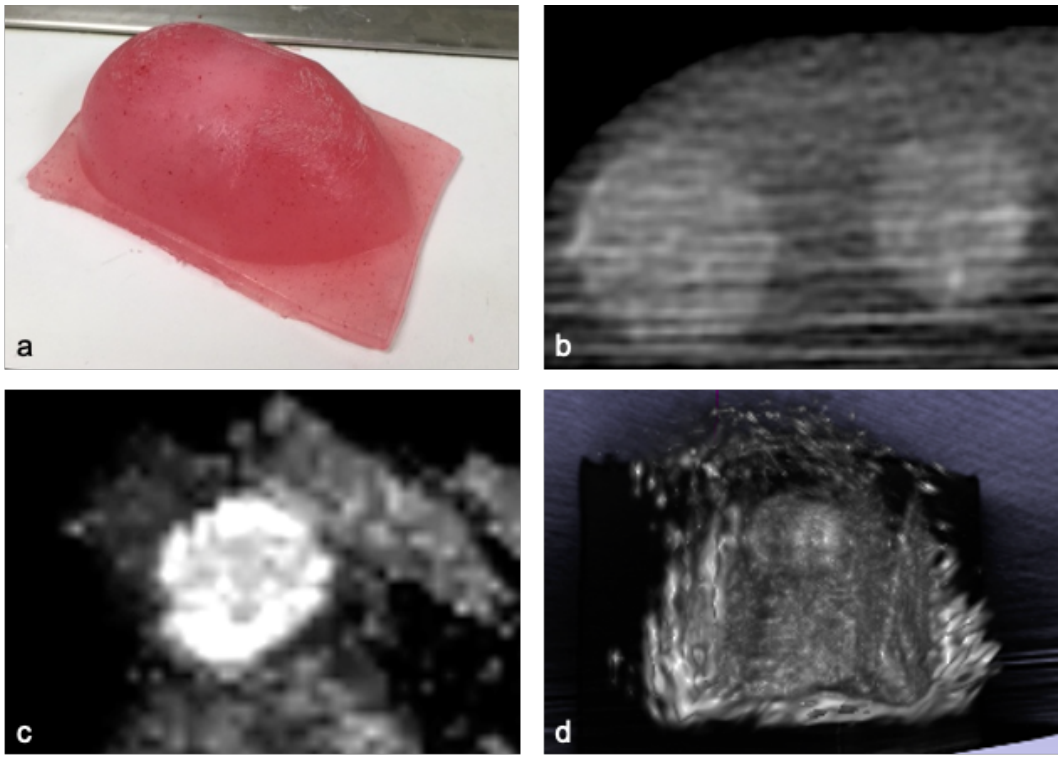

Figure 6.8: Ultrasound navigation and reconstruction system accuracy using a multi-modal polyvinyl alcohol cryogel (PVA-c) phantom. (a) Polyvinyl alcohol phantom. (b) Volumetric 3D x-ray image of phantom obtained with the Medtronic O-arm. (c) Reconstructed ultrasound image obtained with an intraoperative burr hole ultrasound transducer (N11C5s). (d) Volumetric reconstruction of ultrasound data.

the preoperative CT scan was uploaded for surgical planning and image calibration transform was estimated. Trained clinical staff correctly assembled the system hardware and completed initial setup in 10 min 19 s. Assembly of the system's sterile components was completed correctly in 1 min 22 s and intraoperative probe pivot and temporal calibration was completed in 1 min 43 s. Clinical evaluation of the system was undertaken independently by 2 consultant neurosurgeons using a patient-specific PVA-c phantom model of a patient with a vestibular schwannoma undergoing surgical resection via a simulated retrosigmoid craniotomy.

Figure 6.9 compares the patient-specific MRI data used to create the brain tumour phantom with CT and ultrasound images acquired during the workflow integration assessment. Ultrasound image acquisition included freehand real-time imaging and 3D volume reconstruction. Real-time ultrasound imaging was available

in our software, as a CT overlay or on a separate window. The neurosurgeons navigated with the ultrasound probe via the simulated retrosigmoid craniotomy and were able to visualise the simulated brain tumour. Both neurosurgeons considered the system to be very useful, with an intuitive display and clinically acceptable accuracy.

6.4 Discussion

The primary outcome of this work is the efficient integration of the neuronavigation Slicelet into the surgical workflow. The clinical simulation demonstrated the seamless integration of our software into the surgical workflow. The setup of the system was easy and quick to assemble for the clinical team. By only using existing hardware within its intended use, the platform complied with all relevant surgical safety requirements, including sterility and electrical standards. The additional pieces of hardware needed for this prototype system included a network switch and interconnecting cables. A laptop placed on a small surgical trolley could be positioned conveniently within the operating room as directed by the surgical team. The required intraoperative probe calibration procedure took only a couple of minutes. Both neurosurgeons considered the Slicelet user interface to be intuitive and easy to use.

The clinical simulation also involved assessing the neuronavigation platform's visualisation and processing features. The calibration of the tracked ultrasound probe enabled to display real-time ultrasound imaging into the patient' space and to reconstruct volumetric ultrasound data. Although the image streaming rate from the BK5000 system was acceptable (25 fps), it was still lower than the frame rate displayed on the ultrasound system (60 fps). The 3D volume reconstruction method is based on the number of frames acquired within the scout scan region. A low frame rate will result in a sparser volume and will likely increase the spacing between each frame and, therefore, will require a longer acquisition time to obtain an adequate image resolution. The streaming rate was considered satisfactory for real-time visualisation according to the neurosurgeons, but improving the image frame rate

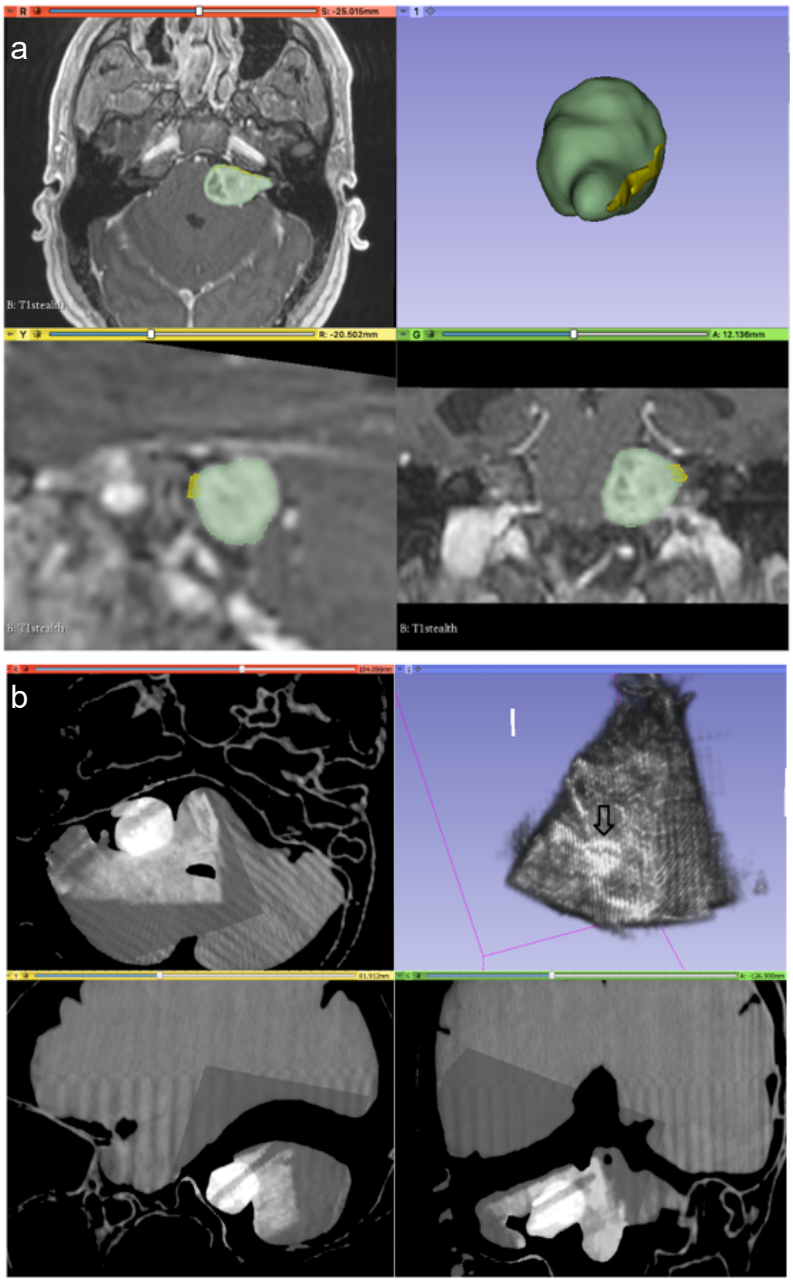

Figure 6.9: Intraoperative simulation of 3D Ultrasound reconstruction using patient data and a patient-specific phantom model. (a) Illustration of patient MRI data using the neuronavigation slicelet system, including an overlay of the tumour and nerve on the MRI data and a separate 3D model of those structures. Green: Tumour, Yellow: Nerve. (b) Neuronavigation slicelet display of the phantom's CT scan with overlaid 3D reconstructed ultrasound volume. A volumetric representation of the US data is displayed in the top right panel. Black arrow: Tumour.

would decrease the 3D reconstruction time. The reconstructed 3D ultrasound volume was also overlaid on the preoperative CT scan to facilitate navigation through the skull phantom and improve the visualisation of the simulated brain tumour. Both neurosurgeons found that the ultrasound and CT overlay was clinically useful and accurate.

Laboratory accuracy assessment demonstrated system accuracy levels comparable to previously reported research systems [196, 204], with an average TRE of 4.11 mm for the sphere phantom experiment. Nevertheless, the TRE achieved in our phantom experiment could be further improved. Future work will investigate the error sources that contribute to the current TRE and look at ways to reduce the most significant error sources. We currently believe that the calibration of the ultrasound probe is a significant source of error, so future work will investigate using alternative calibration methods, for example, phantomless auto-calibration [205]. This feature would be useful in the context of neurosurgery due to intraoperative brain tissue shift and tracking errors. It could considerably improve surgical guidance when there is a high mismatch between preoperative and ultrasound scans.

The stand-alone ultrasound elastography module was not evaluated during the accuracy and workflow assessment. This is due to the fact that the clinically-available ultrasound scanners used in this study did not allow us to stream the raw beamformed ultrasound signal that is used in ultrasound elastography. In chapter 4, we reported an inference speed of 20 fps on a 12 GB NVIDIA GTX-1080ti GPU for quasi-static strain prediction. This suggests that our elastography module could be used alongside our neuronavigation platform to perform real-time tissue elasticity characterisation. Further experiments will be required on tissue-mimicking phantoms to assess the integration of the ultrasound elastography module into the surgical workflow.

6.5 Conclusion

We presented in this chapter an integrated intraoperative navigation system for brain tumour surgery with the ability to incorporate neuronavigation with ultrasound-based guidance. The system was built around commercially-available CE-marked

hardware to facilitate clinical translation, although additional proprietary software for streaming the data out of the commercial devices was required. All other components of the system software, including the Slicer-IGT and PLUS libraries, are open-source.

Future work will focus on the integration of the intraoperative navigation platform in the surgical workflow. Chapter 7 presents a clinical study aiming to evaluate the use of ultrasound imaging during the resection of pituitary brain tumours. The first phase of this clinical trial focused on evaluating the feasibility and safety of a novel intraoperative ultrasound transducer, specifically designed to image pituitary tumours. The next phase of this clinical trial will aim at assessing the integration of the ultrasound-aided neuronavigation platform alongside this new ultrasound transducer.

Chapter 7

Evaluation of an intraoperative ultrasound transducer prototype for pituitary tumour resection

This chapter deals with the assessment of a new ultrasound probe designed specifically for transsphenoidal surgery. To better explain the different reasons lying behind the design and specifications of this prototype, previous work in intraoperative ultrasound imaging for pituitary surgery is presented (section 7.1). The first part of this work involves the pre-clinical assessment of our intraoperative ultrasound probe prototype, which aimed to evaluate the device's image resolution and ergonomics (section 7.2). The second part of this study presents the results of a clinical study performed on an initial cohort of 5 patients who underwent pituitary resection surgery. This clinical study aimed at assessing the feasibility and safety of our new intraoperative ultrasound transducer (section 7.3).

This chapter is adapted from different conference and journal publications and is a collaborative work between our team and our clinical collaborators. I actively participated in the design and implementation of the pre-clinical experiments. I also participated in all the clinical cases by helping set up the ultrasound system and acquiring the data shown in this chapter. I presented our pre-clinical feasibility assessment results at the Computer Assisted Radiology & Surgery (CARS) International Congress, which took place in Berlin, Germany in June 2018 [206]. In

addition, the clinical study preliminary results are adapted from our paper named 'Integrated multi-modality image-guided navigation for neurosurgery: open-source software platform using state-of-the-art clinical hardware', and published in July 2021 [207].

7.1 Introduction

7.1.1 Pituitary adenomas

Pituitary adenoma is a tumour that develops in the pituitary gland. It is located in the brain, behind the nose and eyes, in a depression of the sphenoid bone called the *sella turcica* (see figure 7.1). The pituitary is an endocrine (i.e., hormone-secreting) gland that also controls the production of hormones in the body, regulating diverse physiological functions such as reproduction, blood pressure and growth. Mostly benign, pituitary adenomas can cause different symptoms, depending on the type of tumour. Hormone-secreting (i.e., functioning) tumours can be responsible for a multitude of endocrine disorders such as, Cushing's disease, prolactinoma, or acromegaly [208]. As they grow, pituitary tumours can also compress the functioning gland or the optical chiasma, causing hormone imbalance or visual loss, respectively. Pituitary tumours are common, with a prevalence of 15-20% of the population according to a metaanalysis conducted on English-language articles from MEDLINE [209].

The treatment options depend on the size and type of the tumour and include chemotherapy, radiotherapy, and neurosurgery. Surgical resection is a common approach for treating pituitary adenomas [208, 210], and around 500 cases are performed each year in the United Kingdom [211]. This is due to the fact that the pituitary gland is located right behind the sinuses, and can be accessed by using a minimally invasive, transnasal approach. Among all resected brain tumours, 10% of them are pituitary adenomas [212].

7.1.2 Transsphenoidal surgery

The gold-standard method for pituitary tumour resection is called transsphenoidal surgery [213]. It consists of accessing the *sella turcica* by passing through the nasal

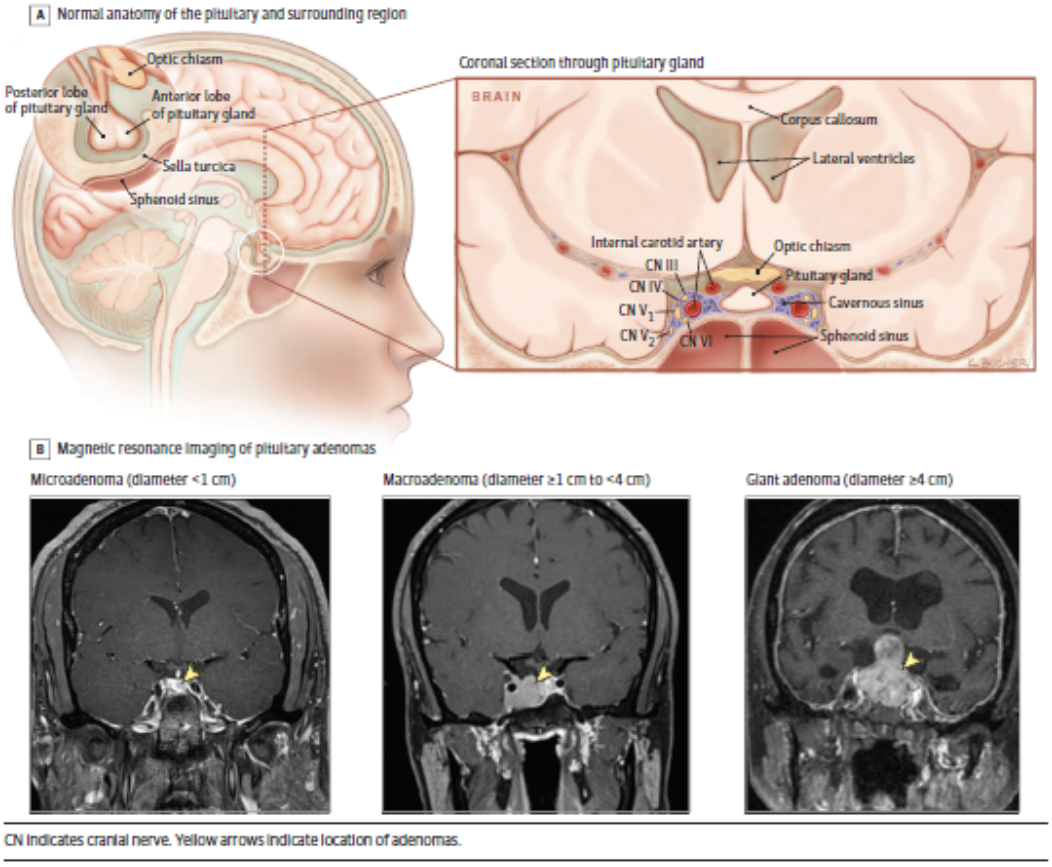

Figure 7.1: Anatomy of the pituitary gland. A: Labelled anatomy of the pituitary gland. B: Brain MRI showing adenomas of different size (from left to right: micro-, macro- and giant adenoma). Figure courtesy of Molitch et al. [208]).

cavity and sphenoid sinus. This method was developed in the early 1900s, with contributions from brain surgery pioneers such as Schloffer, Hirsch, and Cushing [214–216]. The transsphenoidal approach for pituitary surgery prevailed over traditional intra-cranial approaches in the 1960s, with the development of fluoroscopy and microscopy [217, 218]. In the past 20 years, the procedure became less and less invasive with the development of video endoscopy. In contrast with conventional microscopy, the endoscope can be inserted into the sphenoidal sinus without requiring any incisions (e.g., sublabial or tranfixion incision). In addition, endoscopy can provide a more detailed view of the operating site and images of better quality [219]. Furthermore, the use of preoperative scans and frameless stereotaxy further refined the method [220].

Although there have been consecutive improvements through the years, pitu-

itary tumour resection can still be challenging. Indeed, it still suffers from a high incomplete resection rate, which oscillates between 30 and 40%, depending on the study [221, 222]. This increases the risk of cancer recurrence and consequently the need for further treatment such as re-operation, hormonal therapy, and radiotherapy [223–225].

Incomplete resection arises from different factors, including the surgeon's experience, tumour characteristics and visualisation of the surgical field [226]. Although endoscopy improved intraoperative tumour visualisation, it still suffers from a narrow field-of-view and is limited to surface visualisation, which prevents from locating underlying structures such as, blood vessels, nerves and tumour remnants. There is also a risk that the tumour capsule collapses during resection, which can obstruct the tumour site view. Intraoperative blood can also reduce visualisation. Furthermore, pituitary tumour boundaries can be difficult to assess intraoperatively, especially if they are not in the surgeon's direct line of sight.

As a consequence, different intraoperative imaging modalities have been explored to address the high incomplete resection rate, including magnetic resonance imaging (MRI) and computed tomography (CT) [227, 228]. Although both intraoperative MRI (iMRI) and CT (iCT) can provide high-resolution images, their bulkiness, cost, and impact on the duration of surgery are obstacles to their regular and widespread use, not only in transsphenoidal surgeries but in neurosurgical procedures in general.

7.1.3 Previous work on ultrasound-aided pituitary surgery

Intraoperative ultrasound presents several advantages for intraoperative guidance and has been investigated during intracranial procedures for real-time intra-procedural orientation and for monitoring tumour resection procedures [20, 229, 230]. Different clinical studies have assessed the use of intraoperative ultrasound imaging for transsphenoidal surgery. These studies were performed with an ultrasound transducer designed for other clinical applications such as vascular [231], transesophageal [232] and transbronchial [233] imaging. This is due to the fact that the transsphenoidal approach creates a very narrow aperture, and only a few commercial transducers can

both fit into the nasal cavity and image the tumour site. At the time of this study and to the best of our knowledge, only one commercial ultrasound transducer was developed for transsphenoidal surgery [234]. However, this device has not been approved for the European market yet.

The need for ultrasound probes suited for the transsphenoidal approach led to the development of different transducer prototypes. Two non-CE marked prototypes were developed and tested in clinics [235, 236]. The first prototype consists of a 2D side-looking transducer with a tubular shape. The probe is composed of a 64 elements linear array with a 10 MHz central frequency. The tip's diameter and footprint were 4 mm and 12.8 mm, respectively. The prototype was tested on 9 patients with pituitary macroadenoma. The authors noted that the probe provided high-resolution images (0.19×0.22 mm), which allowed to identify tumour tissue and carotid arteries. However, the detection of major nervous structures such as the optical chiasma was not mentioned. In addition, the neurosurgeons indicated that the ultrasound image orientation was confusing. Indeed, the identification of brain structures and tumours are often done using MRI, which is decomposed into a standard image orientation, i.e., axial, coronal and sagittal planes. Therefore, the authors noted that the ultrasound image orientation affected image interpretation. This can be explained by the design of the probe, which features a side-viewing array to reduce the diameter of the tip.

The second prototype consisted of a forward-looking array with a bayonet shape [236]. The footprint of the probe is 12×8 mm, which is considerably larger than the previous side-viewing prototype. It also features a higher central frequency than the previous one, i.e., 12 MHz, to increase the image resolution in the near field. Another clinical trial assessed the probe on 24 pituitary resection cases. The authors argued that this probe provided a better image quality, as well as a more familiar image plane. Neurosurgeons managed to identify all major anatomical structures such as the optical chiasma, carotid arteries, and pituitary gland. In addition, the ultrasound images revealed tumour remnants not visible with endoscopy [236]. However, the bulky tip of the probe was too big to enter the sphenoid sinus in 2 cases, making the

imaging of the tumour site impossible. The authors suggested that both prototypes could be used interchangeably depending on the tumour type. Micro-adenomas are usually located around the gland, whereas the macro-adenomas are larger and can be widespread. Therefore, the second prototype could be more useful to image tumours located at the surface of the gland (i.e., micro-adenomas) and vice versa for the side-looking probe.

7.1.4 Contributions

In this work, we have developed a slender, forward-viewing ultrasound device for transsphenoidal access and plan to create novel software tools to integrate this device into a surgical navigation system. It is hoped that the system will be helpful in locating residual pituitary tumour tissue as well as neurovascular structures. Pre-clinical experiments including image resolution and ergonomics assessment of this new intraoperative ultrasound device are presented. Thereafter, we initiated a single-centre clinical study in the National Hospital for Neurology and Neurosurgery, Queen square, London. The study aimed at assessing the feasibility and safety of the new intraoperative ultrasound transducer in an initial patient cohort. The preliminary results, based on a cohort of five patients, are presented.

7.2 Ultrasound transducer assessment

7.2.1 Ultrasound probe specifications

Our intraoperative ultrasound transducer is a clinical-grade prototype manufactured by Vermon (Vermon S.A., Tours, France), a company specialising in advanced transducer technologies. The probe's design corresponds to a novel version of the two previous prototypes developed by Solheim et al., which were also manufactured by Vermon [235, 236]. This new transducer is designed to overcome the mentioned limitations of the two previous prototypes, which are the unfamiliar imaging plane and the bulky size. The difference of design of the three prototypes is displayed in figure 7.2. Our prototype is less bulky than the forward-viewing probe developed by Solheim et al., which was found to be too large to safely enter the sphenoid sinus in 2 patients. Moreover, the other side-viewing probe prototype needs to be inserted

further into the *sella turcica* to be able to scan the sellar structures, thereby limiting its use to larger tumours.

The development of thin ultrasound probes, dedicated to the transsphenoidal approach, has been based on a necessary trade-off between image quality (which depends on the number of ultrasound elements, greater in the side-viewing prototype) and image intuitiveness (best in the forward-viewing probe). It is with these surgically relevant considerations that we decided to centre further development around a forward-viewing prototype, and to work on optimising image quality for the purpose of transsphenoidal surgery while at the same time keeping the probe's diameter in proportion with the surgical corridor.

Accordingly, our novel transducer consists of an endoscopic-shaped probe with a shaft that is slenderer than in previous versions. The tip is 200 mm long and has a diameter of 7 mm. At the end of the probe, an 80-element convex array with a curvature radius of 8.48 mm is located along the tip, imaging both in the forward and side directions. The centre frequency is higher than previous prototypes, 18 MHz, to improve the image resolution. The transducer's connector is compatible with the CE-marked SonixTouch ultrasound console (BK Ultrasound, UK). A photograph of our quasi-forward-looking ultrasound probe is shown in figure 7.3.

7.2.2 Image resolution assessment

This section describes the method used to determine the image resolution of our high-frequency ultrasound probe. One of the improvements of the probe is its high-central frequency, which should allow the identification of neurovascular structures surrounding the tumour. In ultrasound imaging, the spatial resolution corresponds to the axial (vertical) and lateral (horizontal) resolution of the image. The axial resolution depends on the size of the sound wave pulse. Therefore, higher frequencies will generate shorter pulses and thus increase the axial resolution. However, shorter pulses are more attenuated in soft tissue and, therefore, will affect the imaging depth (see chapter 2).

Our method to assess the image resolution was inspired by an approach used in photo-acoustic imaging [237]. A 25 µm tungsten fibre (Quorum Technologies

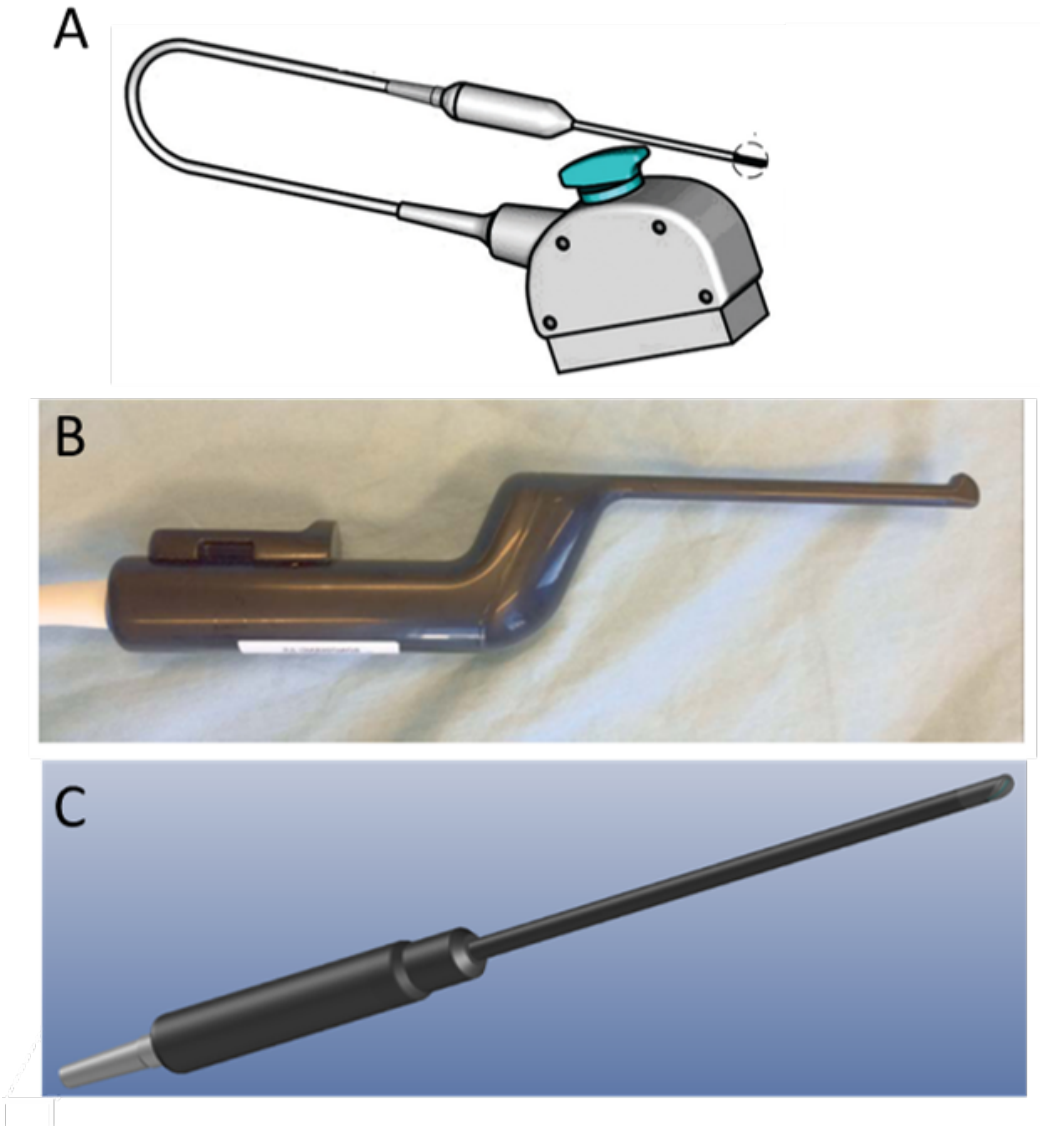

Figure 7.2: Evolution of the transsphenoidal probe prototypes. a) Design of the side-looking probe [235]. b) Picture of the forward-viewing probe with a bayonet shape [236]. c) Design of our prototype.

Ltd., East Sussex, United Kingdom) was anchored in the middle of a water tank (see figure 7.4). The fibre in the water tank was then scanned at different frequencies (8, 15 and 20 MHz), and depth (from 0.5 cm to 2.5 cm with intervals of 0.5 cm). The depth was accurately set by using a XY table. The maximum depth was set to 4.5 cm, because it corresponds to the maximum imaging depth required during pituitary surgery. Indeed, pituitary tumours rarely exceed 1 cm in diameter [238]. To ensure optimal lateral resolution, the ultrasound beam focus was set to the depth of

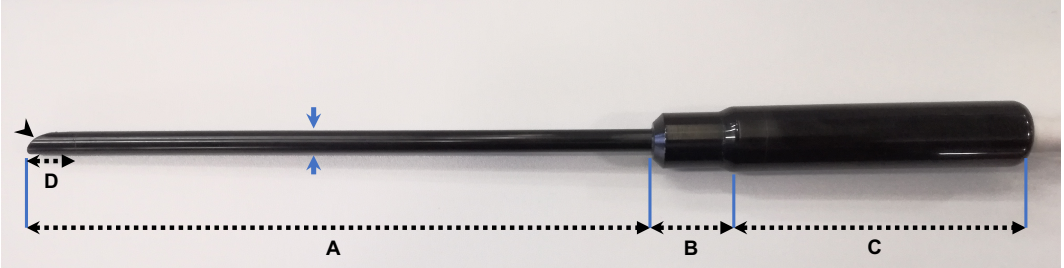

Figure 7.3: Photograph showing the novel forward-viewing ultrasound probe. (A) 200 mm. (B) 15 mm. (C) 85 mm. The probe tip's diameter (the distance between the two blue arrows) is 7 mm. (D) The probe's distal 16 mm contains an 80-element convex array. The radius of the transducer's curvature (indicated by the black arrowhead) is 8.48 mm.

the thread by using dynamic receive beamforming [239].

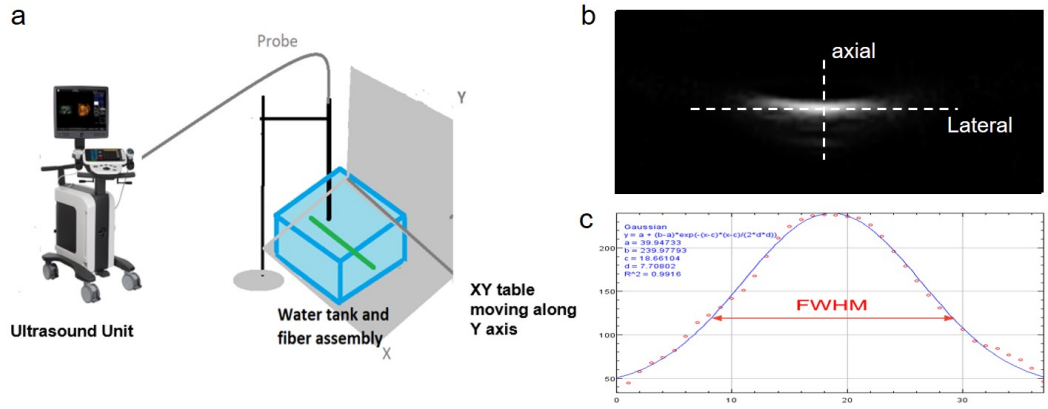

Figure 7.4: Spatial resolution evaluation method. (a) Illustration of the resolution experiment. The green line in the blue cube represents the tungsten wire immersed in the water tank. The water tank was moving along the Y-axis using an XY table, to scan the wire at different depth. (b) Example of the reconstructed wire image. (c) Associated intensity profile (red circles) of the reconstructed wire image with a Gaussian curve fitting (blue curve)

Ultrasound images of the water tank were acquired using a standard CE-marked ultrasound image display system (SonixTouch, BK Ultrasound, UK). A B-mode image was captured for each frequency and depth. The Tungsten fibre was visible on the images by displaying a point spread function shape. The spatial resolution was estimated by taking the full width at half maximum values (FWHM) of its intensity profile (see figure7.4). The size of the fibre in both axial and lateral direction was then converted from pixel to millimetre size. The average resolution was about 0.25 mm ($\pm$ 0.02 mm) in axial and about 1.3 mm ($\pm$ 0.3 mm) in the lateral direction.

7.2.3 Ergonomics study

Prior to the clinical study, a procedural test was performed to assess the device ergonomics. To this end, we mimicked a transsphenoidal surgery procedure. The mock surgery was performed on an existing Kezlex anatomical model (Ono and Co., Tokyo, Japan) (see Figure 7.5), which is a phantom designed for pituitary surgical training. To make the procedural test as realistic as possible, we also included all the equipment and tools commonly used in pituitary surgery. This includes a 0-degree endoscopic camera (Karl Storz SE and Co., Germany), surgical instruments (e.g., a curette), and a neuronavigation platform (Stealth Station S7, Medtronic, USA). In addition, a CT scan of the anatomical phantom was acquired beforehand in order to use it as a preoperative scan for neuronavigation. An optical passive marker (SureTrack^{TM}, Medtronic, USA) was attached to the ultrasound probe in order to track its position during the mock intervention. Finally, the entire procedure was performed and evaluated with the help of two experienced neurosurgeons specialised in transsphenoidal surgery, with over 50 cases performed.

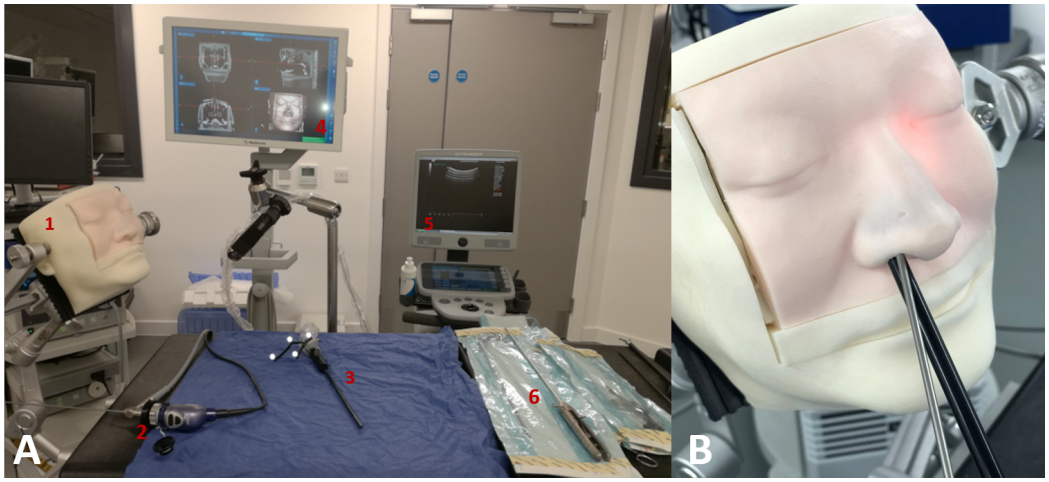

Figure 7.5: Ergonomics study. (a) Experimental set-up. 1: Anatomical phantom positioned in a Mayfield clamp. 2: Endoscope. 3: transsphenoidal probe. 4: Medtronic StealthStation with the phantom model displayed. 5: Ultrasonix system. 6: Surgical tools. (b) Picture of the ultrasound probe and endoscope fitting into the nose of the pituitary phantom.

The ergonomics study allowed us to assess the transducer's technical feasibility. The probe size and configuration enabled the surgeons to insert and manipulate it

easily into the anatomical phantom's nasal cavity. Both the endoscope and ultrasound probe could fit together into the sphenoid sinus (see figure 7.5). Although the probe was covered by a surgical drape, the probe appeared to be robust enough to not break when entering the sinus. The probe was also successfully tracked by the neuronavigation platform during the mock procedure (see figure 7.5). Overall, both neurosurgeons agreed that the probe was suitable for transsphenoidal surgery and were willing to use it in clinics.

7.3 First-in-human clinical study

7.3.1 Study design

The study began in October 2019 and involved patients undergoing transsphenoidal surgery for pituitary adenoma. Ethical approval was obtained from our local ethics committee (19/LO/0882), and the trial was registered on ClinicalTrials.gov (NCT03284775). The minimum inclusion criteria were 1) adult patients, 2) undergoing transsphenoidal surgery, 3) diagnosis of pituitary adenoma on preoperative MRI, and 4) able to provide consent. All patients gave their written informed consent to test our ultrasound probe after the risks, benefits, and alternatives were discussed. Patients were specifically informed that this was the first clinical application of this novel ultrasound probe.

Patients were recruited from October 2019 to March 2020, when the UK entered a lockdown in response to the COVID-19 pandemic. Eleven patients were approached to participate in this study and all eleven were recruited by the clinical team, but only five participated (2 female, 3 male; age range: 42–73 years), due to the study's discontinuation in the general response to the COVID-19 pandemic. Diagnosis included pituitary macroadenoma in three cases, acromegaly in one case, and re-operation due to incomplete resection in the remaining case.

All operations were performed by neurosurgeons who have a subspecialty interest in pituitary surgery, and were involved with the design and pre-clinical assessment of the ultrasound probe. All operations took place at the National Hospital for Neurology and Neurosurgery, which is a large regional neuroscience

centre, that carries out approximately 150-200 pituitary operations each year.

7.3.2 Intervention

The surgical technique for the endoscopic transsphenoidal approach is unchanged to that used without the ultrasound probe. General anaesthesia is administered, and nasal mucosal decongestants applied (Moffat's solution). Patients are positioned supine on the operating table, and the operative field prepared and draped. Access to the *sella turcica* is performed through a mononostril endoscopic approach. If needed, the floor and anterior wall of the *sella* are drilled off. A cruciate incision is made to the dura mater and the tumour removed using curettes and rongeurs.

The ultrasound probe is applied to the *sella* before and after the dural incision for visualisation of the tumour and surrounding neurovascular structures. In case of an intraoperative cerebrospinal fluid leak, a fat graft is harvested from the patient's abdomen and placed over the dural defect. Nasal packing is not placed in the absence of cerebrospinal fluid leak.

7.3.3 Measures of outcomes

The primary outcomes of the study were to demonstrate that the use of the ultrasound device in pituitary surgery is both technically feasible and safe. Technical feasibility was assessed by the ability of the probe to be used during transsphenoidal surgery alongside standard surgical equipment. Device safety was assessed by the number, type, and severity of adverse events related to use of the device, e.g., device malfunction. In addition, a short written assessment of the probe usage was performed by the surgeon carrying out the procedure after each operation.

The secondary outcomes were to demonstrate that the ultrasound device does not adversely affect the occurrence of postoperative complications, the operating time (minutes), and the length of stay (days). For this purpose, the operating times and length of stay were also collected for a series of 30 consecutive patients with pituitary adenoma, operated before and contemporaneously to this study, by the same surgical team and with the same surgical technique, but in whom intraoperative ultrasound had not been used. Postoperative complications to be recorded included death,

cerebrospinal fluid leak, meningitis, vascular complications, visual complications, diabetes insipidus, hypopituitarism, and cranial nerve injury. Vascular complications to be recorded included carotid or other vessel injury, or symptomatic haematoma. Venous bleeding from the cavernous sinus was to be considered a vascular complication only if it prevented completion of the surgical procedure. Epistaxis was only to be considered a vascular complication if it warranted return to the operating room. Cerebrospinal fluid leaks to be recorded included all postoperative leaks requiring lumbar drainage or repair. Patients with Cushing's disease receiving postoperative cortisol, or in whom a hypophysectomy was carried out, were also not included as surgical complications.

7.3.4 Results

7.3.4.1 Primary outcomes

A photograph of the surgical field is showed in figure 7.6. The probe's ability to be used alongside the standard surgical equipment for transsphenoidal procedures was demonstrated in all cases. The ultrasound probe was available at all times during the operation, and imaging the tumour site only took a few minutes each time. The probe was slender enough in all patients to access the *sella turcica*, even while the endoscope was being used (see figure 7.7). According to the neurosurgeons, the probe was easy to manipulate, but the surgical drape that enclosed it was too large and not steady enough. A smaller, thicker drape was also used, but it altered the image quality. Furthermore, no adverse events related to the device were encountered.

We did however experience some difficulty with flattening the superfluous folds of the probe's protective sterile sheath, actually intended for bulkier probes, and so adding unwanted volume to our prototype (see figure 7.7). This led to the aggregation of air bubbles near the tip of the probe, causing the formation of reflection artefacts in the image (see figure 7.9). We rapidly learned to use a sterile elastic band to maintain the sheath taut around the probe's tip. Nevertheless, a dedicated slimmer sheath or a transducer design compatible with standard sterilisation processes would be a preferable solution for future use. Finally, the probe did not enter into conflict with the endoscope, and so could be introduced under direct endoscopic vision (see

figure 7.7).

Three written assessments positively underlined that tumour and vessels were readily identifiable, and in two cases that an adenoma residual was evident. In one of these cases, ultrasound allowed to identify a residual that would not have otherwise been noticed. Additional remarks were that ultrasound clarity was best before opening the dura mater. One assessment observed that the probe would be best used alongside neuronavigation, in view of the difficulty to convincingly identify the anatomy on ultrasound in that particular case.

7.3.4.2 Secondary outcomes

After surgery, one patient represented with a delayed cerebrospinal fluid leak and subsequently developed meningitis. None of the patients had new pituitary hormonal deficits postoperatively, aside one patient who remained on hormonal replacement therapy on follow-up.

Operating time ranged from 68 to 110 minutes (mean: 93 min) and length of stay was 4 days in one case, 5 days in two cases, 6 days in one case, and 18 days in the case with a cerebrospinal fluid leak (median: 5 days). This represents an extra 20 minutes of operating time, when compared to the average duration of the thirty contemporaneous, consecutive cases of pituitary adenoma resection performed by the clinical team without intraoperative ultrasound (range: 45–168 min, mean: 73 min). The length of stay for this same group, on the other hand, was comparable (range: 2 – 20 days, median: 4 days).

7.4 Discussion

The ultrasound probe image resolution was considered satisfactory for the identification of brain neurovascular structure. The maximum average axial resolution was calculated as 0.26 mm ($\pm$0.02 mm) and the maximum average lateral resolution as 1.3 mm ($\pm$0.3 mm). Therefore, it is within the probe's design specifications to be able to image relevant anatomical structures such as the optic tracts and major intra-cranial arteries, which both exhibit a size of a few millimetres [240, 241].

The ergonomics study has demonstrated that the size and configuration of our

Figure 7.6: Photograph of the surgical field. 1: Endoscope display. 2: Ultrasound system. 3: Patient's head covered with sterile drapes.

prototype were appropriate and suitable for the transsphenoidal approach. It also demonstrated that the probe could also be used alongside neuronavigation platforms, which plays an important role for surgical guidance and is another important aspect of this project.

Those considerations on image quality and device ergonomics correlate with our results from the clinical study. The use of our novel ultrasound device appeared to be both technically feasible for the 5 cases. The use of the ultrasound probe did not significantly disrupt the surgical workflow and the additional average surgical time compared favourably with scanning times reported for other intraoperative imaging modalities such as, MRI [242, 243] and CT [244].

We did not encounter any intraoperative adverse events in this preliminary cohort. After surgery, one patient subsequently developed meningitis due to a

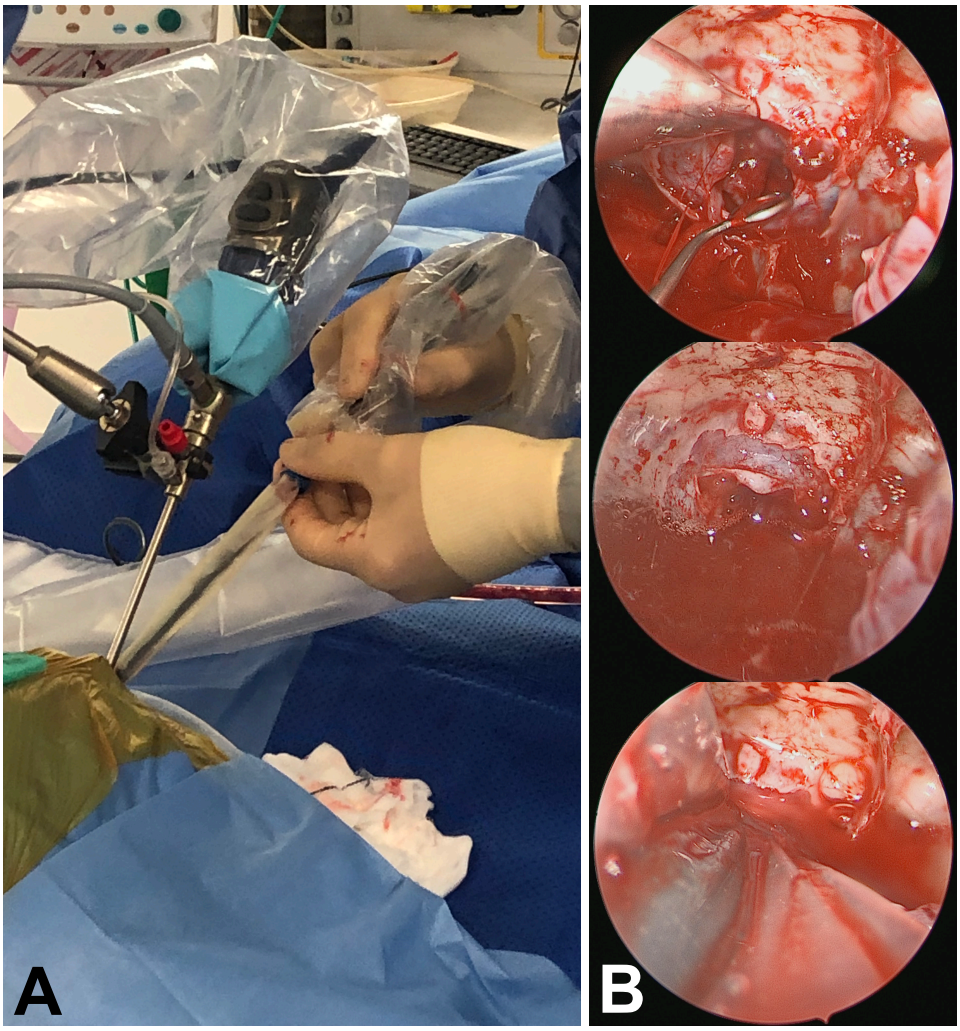

Figure 7.7: (A) Photograph of the surgical field, showing the ultrasound probe in use alongside the endoscope. (B) Top view: Transsphenoidal endoscopic view through left nostril, showing the opening made in an enlarged *sella turcica*. The pituitary tumour has been partially resected. Middle view: The sphenoid sinus is filled with water to reduce acoustic impedance. Lower view: The ultrasound probe's tip is inserted through the sellar opening under endoscopic vision.

delayed cerebrospinal fluid leak. The ultrasound probe was fully covered with a protective sterile sheath during surgery and no puncture of the drape occurred throughout the 5 cases. Cerebrospinal fluid leakage leading to intracranial infection represents one of the most common postoperative complication in transsphenoidal surgery, with an incidence rate ranging from 0.8% to 15% [245]. A larger cohort of patients would be required to determine if the use of intraoperative ultrasound induces a higher occurrence of postoperative infection.

Despite the difficulties encountered with the probe's protective sterile sheath, the

images of the tumour site acquired during surgery were clinically relevant. Operator feedback collected following surgery conveyed that vessels and tumour contours were readily identified on ultrasound (see figure 7.8 and 7.9); that clarity was felt to be best before opening the dura mater; and that the use of flowable haemostatic matrix obscured the ultrasound views. Although our project is not yet at a phase to assess our probe's intraoperative impact on surgical decision-making, anecdotally it did allow identifying an adenoma residual in one patient that would otherwise have been missed.

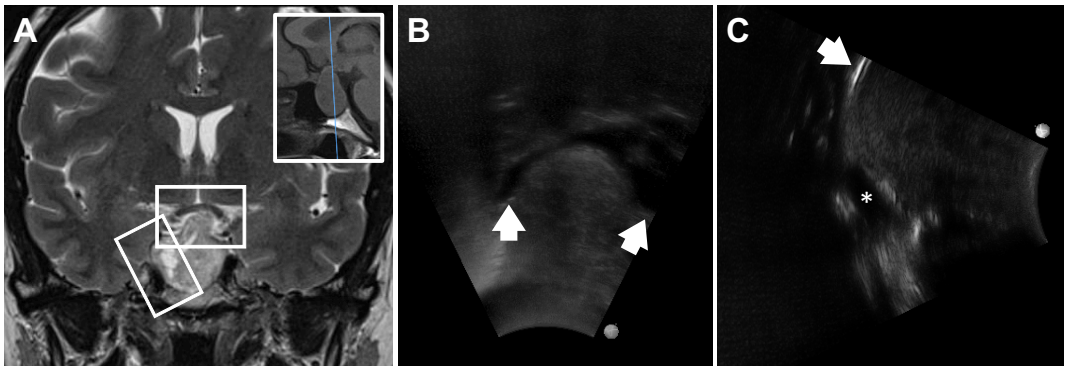

Figure 7.8: A: T2-weighted coronal magnetic imaging resonance slice of a patient with a pituitary macroadenoma. The inset in the top right corner of panel A is a T1-weighted sagittal slice in the same patient, indicating the level of the coronal slice (thin blue line). The white rectangular frame centred on the tumour's apex provides the approximative bearings for panel B, whilst the oblique frame centred on the cavernous segment of the right internal carotid artery provides the bearings for panel C. B: Intraoperative ultrasonographic caption showing the suprasellar portion of the tumour, protruding through the diaphragm sellae (i.e., a sheet of dura mater that forms the roof of the pituitary fossa). The diaphragm sellae is indicated on either side of the tumour by arrowheads. C: Intraoperative ultrasonographic caption of the right infrasellar portion of the tumour. The right diaphragm sellae is indicated by the white arrowhead. The asterisk indicates the cavernous segment of the right internal carotid artery.

7.5 Conclusion

This chapter demonstrates the safety and feasibility of a new ultrasound probe designed specifically for transsphenoidal surgery. Future work will aim at integrating this device into the ultrasound-aided neuronavigation presented in chapter 7, as a means of further enhancing the intraoperative understanding of the surgical field.

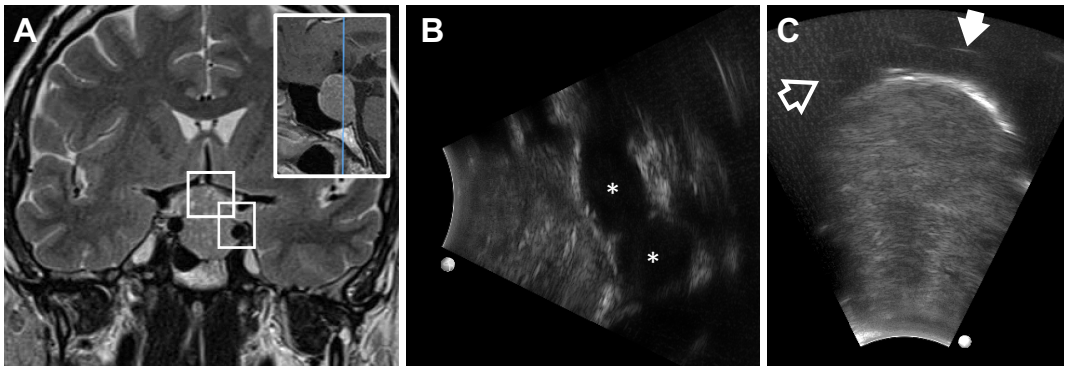

Figure 7.9: A: T2-weighted coronal magnetic imaging resonance slice of a patient with a pituitary macroadenoma. The inset in the top right corner of panel A is a T1-weighted sagittal slice in the same patient, indicated the level of the coronal slice (thin blue line). The white rectangular frame centred on the left internal carotid artery provides the approximative bearings for panel B, whilst the left frame centred on the tumour's apex provides the bearings for panel C. B: Intraoperative ultrasonographic caption of the left lateral aspect of the tumour and of the cavernous segment of the left internal carotid artery (asterisks). C: Intraoperative ultrasonographic caption showing the suprasellar portion of the tumour. A segment of the left anterior cerebral artery is characterised by the faint hyperechogenic line shown by the full white arrowhead. Another segment sliced transversally can be made out faintly as an ovoid circle (hollow arrowhead). The tumour is mostly of a homogenous appearance both on panels B and C. Panel C is also illustrative of reverberation artefact, due to the presence of air between a segment of the convex ultrasound array and the medium.

Chapter 8

General conclusion and future perspectives

8.1 General conclusion

This project focused on the integration of intraoperative ultrasound imaging in the surgical workflow and the development of new elasticity-based tissue characterisation techniques for brain tumour resection. It involved different aspects of ultrasound-based guidance, such as, the development of displacement estimation techniques for ultrasound elastography, the integration of ultrasound imaging into an image-guided platform and the assessment of a new ultrasound transducer prototype for transsphenoidal surgery.

Our work on ultrasound tissue characterisation focused on speckle tracking, which is one of the key challenge of elastographic tissue stiffness evaluation. We explored the use of CNNs for displacement estimation in the case of quasi-static deformations for tissue stiffness evaluation. To the best of our knowledge, we presented the first fully unsupervised learning-based method for displacement estimation in ultrasound elastography. We also extended our unsupervised training scheme to estimate the strain elastogram between ultrasound RF time series, and demonstrated that the use of recurrent units improves displacement estimation and temporal continuity for strain field prediction. Both approaches were validated on numerical ultrasound simulations and an *in vivo* datasets of RF data of a human forearm that we made

publicly available[1]. Our network architectures, USENet and ReUSENet, achieved state-of-the-art performance as compared with traditional approaches, and were able to achieve a prediction rate of 20 images per second with a 12 GB NVIDIA GTX-1080ti GPU. Therefore, it is highly suitable for real-time imaging and represents great potential for the use of learning-based methods in quasi-static elastography.

We also investigated the use of an unsupervised learning-based method for displacement estimation in shear wave elastography imaging. To this end, we generated an ultrasound simulation dataset of 20,000 images that simulates ARF-induced shear wave propagation in 150 different numerical tissue-mimicking phantoms[2]. ReUSENet and USENet showed similar performance to a traditional NCC approach, showing great promise in the use of learning-based methods for the estimation of microscale deformations in transient dynamic elastography.

This thesis also presented the development of an image-guided neurosurgery platform that combines intraoperative ultrasound with neuronavigation. Our system was built around commercially-available CE-marked hardware to facilitate clinical translation and was developed with open-source libraries such as, Slicer-IGT and the PLUS toolkit. The feasibility assessment of our neuronavigation module demonstrated the seamless integration of our software and intraoperative ultrasound imaging into the surgical workflow. The module included features that could potentially improve intraoperative surgical guidance, such as volumetric ultrasound reconstruction and the overlay of real-time ultrasound acquisition with preoperative images. Furthermore, we also developed a stand-alone elastography module that can perform real-time tissue characterisation with USENet and ReUSENet. This module has also been implemented in 3D-Slicer and could be used in combination with our ultrasound-guided neurosurgery platform.

Finally, we performed the clinical feasibility and risk assessment of an intraoperative probe prototype developed for pituitary tumour resection. The ultrasound images acquired during the clinical evaluation had a satisfying contrast and resolution, which allowed neurosurgeons to identify the tumour and surrounding neurovascular

[1]Open access *in vivo* database available on `https://www.synapse.org/InVivoDataForUSE`
[2]Open access simulation database available on `https://www.synapse.org/SimForSWEI`

structures. Our preliminary results suggest that the probe could be safe to use in transsphenoidal surgery, but more cases will be required to confirm its safety and usefulness. At the time this study was carried out, the first intraoperative ultrasound transducer for transsphenoidal surgery was commercialised (BK N20P6, BK Medical, UK) [3]. Although the specifications of the BK probe are not publicly available, it exhibits a similar shape and design to our prototype. This suggests that our project was in line with current clinical needs.

Figure 8.1: BK N20P6 ultrasound transducer for transsphenoidal surgery (BK Medical, UK)

8.2 Limitations and future perspectives

Future directions of this project should focus on the clinical validation and assessment of the learning-based ultrasound elastography techniques developed in this work, and the integration of intraoperative ultrasound imaging in brain tumour surgery.

8.2.1 Learning-based methods for brain tissue characterisation

The displacement estimation methods developed in this thesis for quasi-static elastography were validated on numerical simulations and *in vivo* dataset of images acquired from the arm of human volunteers. The application of our network's architectures to shear wave elastography was only validated using ultrasound simulation. However, we were not able to assess the performance of our methods on intraoperative brain tumour imaging.

Although we were able to record B-mode snapshots of pituitary tumours during the assessment of the intraoperative transducer prototype, we could not retrieve raw RF data from the ultrasound system as part of this ethically approved study. Accurate displacement estimation relies on the high-frequency content of the ultrasound signal, which is contained in the raw beamformed ultrasound data but not in the reconstructed B-mode images used for display. Therefore, future investigations

will aim at validating the use of our different methods on ultrasound data obtained from brain tumour resection. Our unsupervised training scheme allows our neural networks to be trained directly on *in vivo* ultrasound data, without requiring any ground truth labels. The acquisition of an ultrasound neurosurgery dataset for quasi-static and shear wave elastography will enable us to assess the performance of our approach on brain tumour characterisation.

Although the usefulness of quasi-static and shear wave elastography have been investigated previously in neurosurgery [40, 41], it would be interesting to compare their respective performance in a further clinical study. Both techniques determine different elasticity parameters and have different hardware requirements. Quasi-static elastography can be used with most conventional ultrasound systems (when raw data can be retrieved) and determines the strain, which is a qualitative measure of stiffness (see chapter 2). Shear wave imaging can measure the Young's modulus, which is an absolute measure of tissue stiffness, but requires ultrasound systems capable of generating sufficiently high-intensity acoustic radiation force. The comparison of both methods in neurosurgery will allow to better determine the clinical needs for brain tumour resection.

8.2.2 Integration of intraoperative ultrasound in the surgical workflow

The feasibility assessment of our neuronavigation module showed promising results, and future efforts will be made towards its clinical translation. In chapter 6, we highlighted the accuracy of the system, with an average TRE of 4.11 mm for the sphere phantom experiment, as the main limitation of our system. The most significant source of error experienced during our experiments originated from the image-to-probe calibration procedure. Further developments will aim at improving the accuracy of the ultrasound probe calibration by investigating, for example, phantomless auto-calibration techniques [205].

Further experiments would also be required to assess the integration of the stand-alone ultrasound elastography module in the surgical workflow. This will require the use of an ultrasound system that can allow us to retrieve RF data, such as

the Cicada 128PX ultrasound system (Cephasonics Inc., USA) we used to generate our *in vivo* dataset (see chapter 4). The PVA-c material used to make the tissue-mimicking phantoms in the preclinical assessment of the neuronavigation module have also been used for tissue stiffness evaluation [246]. This would allow us to reproduce the simulated brain tumour experiment by including realistic brain tissue and tumour stiffness properties.

Finally, the second part of the clinical study on the assessment of the intra-operative probe prototype will aim to evaluate the combination of the ultrasound probe prototype with the ultrasound-guided neurosurgery platform and investigate the usefulness of our intraoperative elastography module. This will allow us to clinically assess the feasibility and safety of our research platform in the surgical workflow.

References

[1] Md Ashikuzzaman, Ali Sadeghi-Naini, Abbas Samani, and Hassan Rivaz. Combining first and second order continuity constraints in ultrasound elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2021.

[2] Ali KZ Tehrani, Morteza Mirzaei, and Hassan Rivaz. Semi-supervised training of optical flow convolutional neural networks in ultrasound elastography. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 504–513. Springer, 2020.

[3] BK Ultrasound. Minimally Invasive Transducer. `https://www.bkmedical.com/transducers/minimally-invasive-n20p6/`. [Online; accessed 16-September-2021].

[4] NHS. U11: Pivot Calibration. `https://www.nhs.uk/conditions/brain-tumours/`. [Online; accessed 13-September-2021].

[5] Joao Paulo Almeida, Kaisorn L Chaichana, Jordina Rincon-Torroella, and Alfredo Quinones-Hinojosa. The value of extent of resection of glioblastomas: clinical evidence and current approach. *Current neurology and neuroscience reports*, 15(2):1–13, 2015.

[6] Hani J Marcus, Tom Vercauteren, Sebastien Ourselin, and Neil L Dorward. Intraoperative ultrasound in patients undergoing transsphenoidal surgery for pituitary adenoma: Systematic review. *World neurosurgery*, 106:680–685, 2017.

[7] B. Albayrak, A. F. Samdani, Peter M. Black, Yasuhiro Yonekawa, R. Fahlbusch, and Ch Nimsky. Intra-operative magnetic resonance imaging in neurosurgery. *Acta Neurochirurgica*, 146(6):543–557, 2004.

[8] David W. Roberts, John W. Strohbehn, John F. Hatch, William Murray, and Hans Kettenberger. A frameless stereotaxic integration of computerized tomographic imaging and the operating microscope. *Journal of Neurosurgery*, 65(4):545–549, 1986.

[9] Pankaj Daga, Gavin Winston, Marc Modat, M Jorge Cardoso, Mark White, Laura Mancini, Andrew W McEvoy, John Thornton, Caroline Micallef, Tarek Yousry, et al. Improved neuronavigation through integration of intraoperative anatomical and diffusion images in an interventional mri suite. In *International Conference on Information Processing in Computer-Assisted Interventions*, pages 168–178. Springer, 2011.

[10] N Kitchen and J Shapey. The operating theatre environment. In Ramez Kirollos, Adel Helmy, Simon Thomson, and Peter Hutchinson, editors, *Oxford Textbook of Neurological Surgery*, pages 45–56. Oxford University Press, 2019.

[11] David W. Roberts, Alexander Hartov, Francis E. Kennedy, Michael I. Miga, and Keith D. Paulsen. Intraoperative brain shift and deformation: A quantitative analysis of cortical displacement in 28 cases. *Neurosurgery*, 43(4):749–758, 1998.

[12] Derek L.G. Hill, Calvin R. Maurer, Robert J. Maciunas, John A. Barwise, J. Michael Fitzpatrick, and Matthew Y. Wang. Measurement of intraoperative brain surface deformation under a craniotomy. *Neurosurgery*, 43(3):514–526, 1998.

[13] Christopher Nimsky, Oliver Ganslandt, Peter Hastreiter, and Rudolf Fahlbusch. Intraoperative compensation for brain shift. *Surgical Neurology*, 56(6):357–364, 2001.

[14] L. Dade Lunsford and A. Julio Martinez. Stereotactic exploration of the brain in the era of computed tomography. *Surgical Neurology*, 22(3):222–230, 1984.

[15] P. Grunert, W. Müller-Forell, K. Darabi, R. Reisch, C. Busert, N. Hopf, and A. Perneczky. Basic principles and clinical applications of neuronavigation and intraoperative computed tomography. *Computer Aided Surgery*, 3(4):166–173, 1998.

[16] Jonathan H Sherman, Kathryn Hoes, Joshua Marcus, Ricardo J Komotar, Cameron W Brennan, and Philip H Gutin. Neurosurgery for brain tumors: update on recent technical advances. *Current neurology and neuroscience reports*, 11(3):313–319, 2011.

[17] George J. Dohrmann and Jonathan M. Rubin. Use of ultrasound in neurosurgical operations: a preliminary report. *Surgical Neurology*, 16(5):362–366, 1981.

[18] William F. Chandler, James E. Knake, John E. McGillicuddy, Kevin O. Lillehei, and Terry M. Silver. Intraoperative use of real-time ultrasonography in neurosurgery. *Journal of Neurosurgery*, 57(2):157–163, 1982.

[19] G. Unsgaard, O. M. Rygh, T. Selbekk, T. B. Müller, F. Kolstad, F. Lindseth, and T. A Nagelhus Hernes. Intra-operative 3D ultrasound in neurosurgery. *Acta Neurochirurgica*, 148(3):235–253, 2006.

[20] Rahul Sastry, Wenya Linda Bi, Steve Pieper, Sarah Frisken, Tina Kapur, William Wells, and Alexandra J. Golby. Applications of Ultrasound in the Resection of Brain Tumors. *Journal of Neuroimaging*, 27(1):5–15, 2017.

[21] 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. *International Journal of Computer Assisted Radiology and Surgery*, 12(9):1607–1619, 2017.

[22] Aabir Chakraborty, Gearóid Berry, Jeff Bamber, and Neil Dorward. Intraoperative ultrasound elastography and registered magnetic resonance imaging of brain tumours: a feasibility study. *Ultrasound*, 14(1):43–49, 2006.

[23] Maarouf A Hammoud, B Lee Ligon, Rabih Elsouki, Wei Ming Shi, Donald F Schomer, and Raymond Sawaya. Use of intraoperative ultrasound for localizing tumors and determining the extent of resection: a comparative study with magnetic resonance imaging. *Journal of neurosurgery*, 84(5):737–741, 1996.

[24] J. Regelsberger, F. Lohmann, K. Helmke, and M. Westphal. Ultrasound-guided surgery of deep seated brain lesions. *European Journal of Ultrasound*, 12(2):115–121, 2000.

[25] Geirmund Unsgaard, Steinar Ommedal, Tomm Muller, Aage Gronningsaeter, Toril A.Nagelhus Hernes, Patrick J. Kelly, and William F. Chandler. Neuronavigation by intraoperative three-dimensional ultrasound: Initial experience during brain tumor resection. *Neurosurgery*, 50(4):804–812, 2002.

[26] Jonathan M Rubin and GJ Dohrmann. Intraoperative neurosurgical ultrasound in the localization and characterization of intracranial masses. *Radiology*, 148(2):519–524, 1983.

[27] PD Le Roux, MS Berger, K Wang, LA Mack, and GA Ojemann. Low grade gliomas: comparison of intraoperative ultrasound characteristics with preoperative imaging studies. *Journal of Neuro-Oncology*, 13(2):189–98, 1992.

[28] V. Van Velthoven. Intraoperative ultrasound imaging: Comparison of pathomorphological findings in US versus CT, MRI and intraoperative findings. *Acta Neurochirurgica, Supplement*, 85(85):95–99, 2003.

[29] Geirmund Unsgaard, Aage Gronningsaeter, Steinar Ommedal, and Toril A Nagelhus Hernes. Brain operations guided by real-time two-dimensional ultrasound: new possibilities as a result of improved image quality. *Neurosurgery*, 51(2):402–412, 2002.

[30] Francesco Prada, Massimiliano Del Bene, Luca Mattei, Cecilia Casali, Assunta Filippini, Federico Legnani, Antonella Mangraviti, Andrea Saladino, Alessandro Perin, Carla Richetta, et al. Fusion imaging for intraoperative ultrasound-based navigation in neurosurgery. *Journal of ultrasound*, 17(3):243–251, 2014.

[31] Shiro Ohue, Yoshiaki Kumon, Shigeyuki Nagato, Shohei Kohno, Hironobu Harada, Kou Nakagawa, Keiichi Kikuchi, Hitoshi Miki, and Takanori Ohnishi. Evaluation of intraoperative brain shift using an ultrasound-linked navigation system for brain tumor surgery. *Neurologia medico-chirurgica*, 50(4):291–300, 2010.

[32] Adomas Bunevicius, Katharina Schregel, Ralph Sinkus, Alexandra Golby, and Samuel Patz. Mr elastography of brain tumors. *NeuroImage: Clinical*, 25:102109, 2020.

[33] Rosa MS Sigrist, Joy Liau, Ahmed El Kaffas, Maria Cristina Chammas, and Juergen K Willmann. Ultrasound elastography: review of techniques and clinical applications. *Theranostics*, 7(5):1303, 2017.

[34] Giovanna Ferraioli, Carlo Filice, Laurent Castera, Byung Ihn Choi, Ioan Sporea, Stephanie R Wilson, David Cosgrove, Christoph F Dietrich, Dominique Amy, Jeffrey C Bamber, et al. Wfumb guidelines and recommendations for clinical use of ultrasound elastography: Part 3: liver. *Ultrasound in medicine & biology*, 41(5):1161–1179, 2015.

[35] Mehdi Moradi, Parvin Mousavi, and Purang Abolmaesumi. Computer-aided diagnosis of prostate cancer with emphasis on ultrasound-based approaches: A review. *Ultrasound in medicine & biology*, 33(7):1010–1028, 2007.

[36] Yurong Hong, Xueming Liu, Zhiyu Li, Xiufang Zhang, Meifeng Chen, and Zhiyan Luo. Real-time ultrasound elastography in the differential diagnosis of benign and malignant thyroid nodules. *Journal of Ultrasound in Medicine*, 28(7):861–867, 2009.

[37] Timothy J Hall, Yanning Zhu, and Candace S Spalding. In vivo real-time freehand palpation imaging. *Ultrasound in medicine & biology*, 29(3):427–435, 2003.

[38] Koichi Kato, Hiroyuki Sugimoto, Naohito Kanazumi, Shuji Nomoto, Shin Takeda, and Akimasa Nakao. Intra-operative application of real-time tissue elastography for the diagnosis of liver tumours. *Liver International*, 28(9):1264–1271, 2008.

[39] Corin F Otesteanu, Valery Vishnevsky, and Orcun Goksel. Fem-based elasticity reconstruction using ultrasound for imaging tissue ablation. *International journal of computer assisted radiology and surgery*, 13(6):885–894, 2018.

[40] Aabir Chakraborty, Jeffrey C. Bamber, and Neil L. Dorward. Preliminary investigation into the use of ultrasound elastography during brain tumour resection. *Ultrasound*, 20(1):33–40, 2012.

[41] D Chauvet, M Imbault, L Capelle, C Demene, M Mossad, C Karachi, A-L Boch, J-L Gennisson, and M Tanter. In vivo measurement of brain tumor elasticity using intraoperative shear wave elastography. *Ultraschall in der Medizin-European Journal of Ultrasound*, 37(06):584–590, 2016.

[42] K Kirk Shung. Diagnostic ultrasound: Past, present, and future. *J Med Biol Eng*, 31(6):371–4, 2011.

[43] Peter NT Wells and Hai-Dong Liang. Medical ultrasound: imaging of soft tissue strain and elasticity. *Journal of the Royal Society Interface*, 8(64):1521–1549, 2011.

[44] Prepared by the Safety Group of the British Medical Ultrasound Society. Guidelines for the safe use of diagnostic ultrasound equipment. *Ultrasound*, 18(2):52–59, 2010.

[45] Zahra Izadifar, Paul Babyn, and Dean Chapman. Mechanical and biological

effects of ultrasound: A review of present knowledge. *Ultrasound in medicine & biology*, 43(6):1085–1104, 2017.

[46] William D O'Brien Jr and William D O'Brien Jr. Assessing the risks for modern diagnostic ultrasound imaging. *Japanese journal of applied physics*, 37(5S):2781, 1998.

[47] Charlie Demene, Jérome Baranger, Miguel Bernal, Catherine Delanoe, Stéphane Auvin, Valérie Biran, Marianne Alison, Jérome Mairesse, Elisabeth Harribaud, and Mathieu Pernot. Functional ultrasound imaging of brain activity in human newborns. *Science translational medicine*, 9(411):eaah6756, 2017.

[48] Jeroen Dudink, Sylke Jeanne Steggerda, and Sandra Horsch. State-of-the-art neonatal cerebral ultrasound: technique and reporting. *Pediatric Research*, 87(1):3–12, 2020.

[49] Francis A Duck. The meaning of thermal index (ti) and mechanical index (mi) values. *BMUS Bulletin*, 5(4):36–40, 1997.

[50] James A. Zagzebski. Physics and Instrumentation in Doppler and B-mode Ultrasonography. `https://radiologykey.com/physics-and-instrumentation-in-doppler-and-b-mode-ultrasonography/`. [Online; accessed 18-September-2021].

[51] Fikri M Abu-Zidan, Ashraf F Hefny, and Peter Corr. Clinical ultrasound physics. *Journal of Emergencies, Trauma and Shock*, 4(4):501, 2011.

[52] Chandra M Sehgal and James F Greenleaf. Scattering of ultrasound by tissues. *Ultrasonic imaging*, 6(1):60–80, 1984.

[53] Avinash C Kak and Kris A Dines. Signal processing of broadband pulsed ultrasound: measurement of attenuation of soft biological tissues. *IEEE transactions on biomedical engineering*, 1(4):321–344, 1978.

[54] PNT Wells. Absorption and dispersion of ultrasound in biological tissue. *Ultrasound in medicine & biology*, 1(4):369–376, 1975.

[55] Thomas L Szabo. Causal theories and data for acoustic attenuation obeying a frequency power law. *The Journal of the Acoustical Society of America*, 97(1):14–24, 1995.

[56] Peter NT Wells. Ultrasound imaging. *Physics in Medicine & Biology*, 51(13):R83, 2006.

[57] Vincent Chan and Anahi Perlas. Basics of ultrasound imaging. In *Atlas of ultrasound-guided procedures in interventional pain management*, pages 13–19. Springer, 2011.

[58] Aaron J Engel. *Development of cardiac atrial kick shear wave elastography to assess myocardial stiffness*. PhD thesis, The University of Nebraska-Lincoln, 2017.

[59] Natalia Ilyina. Reconstruction of acoustic tissue properties for ultrasound tissue characterization. 2017.

[60] Douglas G Wildes and L Scott Smith. Advanced ultrasound probes for medical imaging. In *AIP Conference Proceedings*, volume 1430, pages 801–808. American Institute of Physics, 2012.

[61] Vinayak Dutt and James F Greenleaf. Statistics of the log-compressed echo envelope. *The Journal of the Acoustical Society of America*, 99(6):3817–3825, 1996.

[62] Laurent Sandrin, Stefan Catheline, Mickaël Tanter, X Hennequin, and Mathias Fink. Time-resolved pulsed elastography with ultrafast ultrasonic imaging. *Ultrasonic imaging*, 21(4):259–272, 1999.

[63] Mickael Tanter and Mathias Fink. Ultrafast imaging in biomedical ultrasound. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 61(1):102–119, 2014.

[64] Gabriel Montaldo, Mickaël Tanter, Jérémy Bercoff, Nicolas Benech, and Mathias Fink. Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 56(3):489–506, 2009.

[65] James Adlard et al. *The genuine works of Hippocrates*, volume 1. Sydenham society, 1849.

[66] APC Choi and Yong Ping Zheng. Estimation of young's modulus and poisson's ratio of soft tissue from indentation using two different-sized indentors: finite element analysis of the finite deformation effect. *Medical and Biological Engineering and Computing*, 43(2):258–264, 2005.

[67] Charles Kittel and Paul McEuen. *Kittel's Introduction to Solid State Physics*. John Wiley & Sons, 2018.

[68] Kathryn Nightingale, Stephen McAleavey, and Gregg Trahey. Shear-wave generation using acoustic radiation force: in vivo and ex vivo results. *Ultrasound in medicine & biology*, 29(12):1715–1723, 2003.

[69] Tsuyoshi Shiina, Kathryn R Nightingale, Mark L Palmeri, Timothy J Hall, Jeffrey C Bamber, Richard G Barr, Laurent Castera, Byung Ihn Choi, Yi-Hong Chou, David Cosgrove, et al. Wfumb guidelines and recommendations for clinical use of ultrasound elastography: Part 1: basic principles and terminology. *Ultrasound in medicine & biology*, 41(5):1126–1147, 2015.

[70] Guo-Yang Li and Yanping Cao. Mechanics of ultrasound elastography. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2199):20160841, 2017.

[71] J. Ophir, I. Céspedes, H. Ponnekanti, Y. Yazdi, and X. Li. Elastography: A quantitative method for imaging the elasticity of biological tissues. *Ultrasonic Imaging*, 13(2):111–134, 1991.

[72] Karsten Mark Hiltawsky, Martin Krüger, Claudia Starke, Lothar Heuser, Helmut Ermert, and Arne Jensen. Freehand ultrasound elastography of breast lesions: clinical results. *Ultrasound in medicine & biology*, 27(11):1461–1469, 2001.

[73] Emre Turgay, Septimiu Salcudean, and Robert Rohling. Identifying the mechanical properties of tissue by ultrasound strain imaging. *Ultrasound in medicine & biology*, 32(2):221–235, 2006.

[74] A Pesavento, A Lorenz, S Siebers, and H Ermert. New real-time strain imaging concepts using diagnostic ultrasound. *Physics in Medicine & Biology*, 45(6):1423, 2000.

[75] Pedro A Patlan-Rosales and Alexandre Krupa. Robotic assistance for ultrasound elastography providing autonomous palpation with teleoperation and haptic feedback capabilities. In *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, pages 1018–1023. IEEE, 2020.

[76] Tomy Varghese. Quasi-static ultrasound elastography. *Ultrasound clinics*, 4(3):323, 2009.

[77] N Frulio and Herve Trillaud. Ultrasound elastography in liver. *Diagnostic and interventional imaging*, 94(5):515–534, 2013.

[78] Woo Kyoung Jeong, Hyo K. Lim, Hyoung Ki Lee, Jae Moon Jo, and Yongsoo Kim. Principles and clinical application of ultrasound elastography for diffuse liver disease. *Ultrasonography*, 33(3):149–160, 2014.

[79] J. M. Correas, A. M. Tissier, A. Khairoune, G. Khoury, D. Eiss, and O. Hélénon. Ultrasound elastography of the prostate: State of the art. *Diagnostic and Interventional Imaging*, 94(5):551–560, 2013.

[80] Jin Young Kwak and Eun Kyung Kim. Ultrasound elastography for thyroid nodules: Recent advances. *Ultrasonography*, 33(2):75–82, 2014.

[81] N Grenier, J-L Gennisson, F Cornelis, Y Le Bras, and L Couzi. Renal ultrasound elastography. *Diagnostic and interventional imaging*, 94(5):545–550, 2013.

[82] Hui Zhi, Bing Ou, Bao-Ming Luo, Xia Feng, Yan-Ling Wen, and Hai-Yun Yang. Comparison of ultrasound elastography, mammography, and sonography in the diagnosis of solid breast lesions. *Journal of ultrasound in medicine*, 26(6):807–815, 2007.

[83] Adrian Săftoiu, P. Vilmann, H. Hassan, and F. Gorunescu. Analysis of endoscopic ultrasound elastography used for characterisation and differentiation of benign and malignant lymph nodes. *Ultraschall in der Medizin*, 27(6):535–542, 2006.

[84] Ako Itoh, Ei Ueno, Eriko Tohno, Hiroshi Kamma, Hideto Takahashi, Tsuyoshi Shiina, Makoto Yamakawa, and Takeshi Matsumura. Breast disease: clinical application of us elastography for diagnosis. *Radiology*, 239(2):341–350, 2006.

[85] S Catheline, J-L Gennisson, G Delon, M Fink, R Sinkus, S Abouelkaram, and J Culioli. Measurement of viscoelastic properties of homogeneous soft solid using transient elastography: An inverse problem approach. *The Journal of the Acoustical Society of America*, 116(6):3734–3741, 2004.

[86] Jérémy Bercoff, Mickaël Tanter, and Mathias Fink. Supersonic shear imaging: A new technique for soft tissue elasticity mapping. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51(4):396–409, 2004.

[87] Joyce McLaughlin and Daniel Renzi. Shear wave speed recovery in transient elastography and supersonic imaging using propagating fronts. *Inverse Problems*, 22(2):681, 2006.

[88] Mark L Palmeri, Michael H Wang, Jeremy J Dahl, Kristin D Frinkley, and Kathryn R Nightingale. Quantifying hepatic shear modulus in vivo using

acoustic radiation force. *Ultrasound in medicine & biology*, 34(4):546–558, 2008.

[89] Michael H Wang, Mark L Palmeri, Veronica M Rotemberg, Ned C Rouze, and Kathryn R Nightingale. Improving the robustness of time-of-flight based shear wave speed reconstruction methods using ransac in human liver in vivo. *Ultrasound in medicine & biology*, 36(5):802–813, 2010.

[90] Ned C Rouze, Michael H Wang, Mark L Palmeri, and Kathryn R Nightingale. Parameters affecting the resolution and accuracy of 2-d quantitative shear wave images. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 59(8):1729–1740, 2012.

[91] Laurent Sandrin, Mickaël Tanter, Stefan Catheline, and Mathias Fink. Shear modulus imaging with 2-d transient elastography. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 49(4):426–435, 2002.

[92] Laurent Castéra, Julien Vergniol, Juliette Foucher, Brigitte Le Bail, Elise Chanteloup, Maud Haaser, Monique Darriet, Patrice Couzigou, and Victor De Lédinghen. Prospective comparison of transient elastography, fibrotest, apri, and liver biopsy for the assessment of fibrosis in chronic hepatitis c. *Gastroenterology*, 128(2):343–350, 2005.

[93] Laurent Sandrin, Bertrand Fourquet, Jean Michel Hasquenoph, Sylvain Yon, Céline Fournier, Frédéric Mal, Christos Christidis, Marianne Ziol, Bruno Poulet, Farad Kazemi, Michel Beaugrand, and Robert Palau. Transient elastography: A new noninvasive method for assessment of hepatic fibrosis. *Ultrasound in Medicine and Biology*, 29(12):1705–1713, 2003.

[94] Laurent Castera, Xavier Forns, and Alfredo Alberti. Non-invasive evaluation of liver fibrosis using transient elastography. *Journal of hepatology*, 48(5):835–847, 2008.

[95] Jeffrey Bamber, David Cosgrove, Christoph F Dietrich, Jérémie Fromageau, Jörg Bojunga, Fabrizio Calliada, V Cantisani, J-M Correas, M D'onofrio,

EE Drakonaki, et al. Efsumb guidelines and recommendations on the clinical use of ultrasound elastography. part 1: Basic principles and technology. *Ultraschall in der Medizin-European Journal of Ultrasound*, 34(02):169–184, 2013.

[96] GR Torr. The acoustic radiation force. *American Journal of Physics*, 52(5):402–408, 1984.

[97] Wesley Le Mars Nyborg. Acoustic streaming. In *Physical acoustics*, volume 2, pages 265–331. Elsevier, 1965.

[98] Joshua R Doherty, Gregg E Trahey, Kathryn R Nightingale, and Mark L Palmeri. Acoustic radiation force elasticity imaging in diagnostic ultrasound. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 60(4):685–701, 2013.

[99] Kathy Nightingale. Acoustic radiation force impulse (arfi) imaging: a review. *Current medical imaging*, 7(4):328–339, 2011.

[100] Kathryn Nightingale, Mary Scott Soo, Roger Nightingale, and Gregg Trahey. Acoustic radiation force impulse imaging: In vivo demonstration of clinical feasibility. *Ultrasound in Medicine and Biology*, 28(2):227–235, 2002.

[101] Yu-Ping Yang, Xiao-Hong Xu, Xiao-Wan Bo, Bo-Ji Liu, Le-Hang Guo, Jun-Mei Xu, Li-Ping Sun, and Hui-Xiong Xu. Comparison of virtual touch tissue imaging & quantification (vtiq) and virtual touch tissue quantification (vtq) for diagnosis of thyroid nodules. *Clinical hemorheology and microcirculation*, 65(2):137–149, 2017.

[102] Stephanie A Eyerly, Stephen J Hsu, Shruti H Agashe, Gregg E Trahey, Yang Li, and Patrick D Wolf. An in vitro assessment of acoustic radiation force impulse imaging for visualizing cardiac radiofrequency ablation lesions. *Journal of cardiovascular electrophysiology*, 21(5):557–563, 2010.

[103] Jeremy J Dahl, Douglas M Dumont, Jason D Allen, Elizabeth M Miller, and Gregg E Trahey. Acoustic radiation force impulse imaging for noninvasive characterization of carotid artery atherosclerotic plaques: a feasibility study. *Ultrasound in medicine & biology*, 35(5):707–716, 2009.

[104] Amy C Sharma, Mary Scott Soo, Gregg E Trahey, and Kathryn R Nightingale. Acoustic radiation force impulse imaging of in vivo breast masses. In *IEEE Ultrasonics Symposium, 2004*, volume 1, pages 728–731. IEEE, 2004.

[105] Liang Zhai, John Madden, Wen-Chi Foo, Mark L Palmeri, Vladimir Mouraviev, Thomas J Polascik, and Kathryn R Nightingale. Acoustic radiation force impulse imaging of human prostates ex vivo. *Ultrasound in medicine & biology*, 36(4):576–588, 2010.

[106] Arinc Ozturk, Joseph R Grajo, Manish Dhyani, Brian W Anthony, and Anthony E Samir. Principles of ultrasound elastography. *Abdominal Radiology*, 43(4):773–785, 2018.

[107] Mickael Tanter, Jeremy Bercoff, Alexandra Athanasiou, Thomas Deffieux, Jean-Luc Gennisson, Gabriel Montaldo, Marie Muller, Anne Tardivon, and Mathias Fink. Quantitative assessment of breast lesion viscoelasticity: initial clinical results using supersonic shear imaging. *Ultrasound in medicine & biology*, 34(9):1373–1386, 2008.

[108] Marie Muller, Jean-Luc Gennisson, Thomas Deffieux, Mickaël Tanter, and Mathias Fink. Quantitative viscoelasticity mapping of human liver using supersonic shear imaging: preliminary in vivo feasability study. *Ultrasound in medicine & biology*, 35(2):219–229, 2009.

[109] Mathias Fink and Mickaël Tanter. Multiwave imaging and super resolution. *Physics Today*, 63:28–33, 02 2010.

[110] Pengfei Song, Heng Zhao, Armando Manduca, Matthew W Urban, James F Greenleaf, and Shigao Chen. Comb-push ultrasound shear elastography (cuse):

a novel method for two-dimensional shear elasticity imaging of soft tissues. *IEEE transactions on medical imaging*, 31(9):1821–1832, 2012.

[111] Pengfei Song, Michael C Macdonald, Russell H Behler, Justin D Lanning, Michael H Wang, Matthew W Urban, Armando Manduca, Heng Zhao, Matthew R Callstrom, Azra Alizad, et al. Two-dimensional shear-wave elastography on conventional ultrasound scanners with time-aligned sequential tracking (tast) and comb-push ultrasound shear elastography (cuse). *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 62(2):290–302, 2015.

[112] Max Denis, Mohammad Mehrmohammadi, Pengfei Song, Duane D Meixner, Robert T Fazzio, Sandhya Pruthi, Dana H Whaley, Shigao Chen, Mostafa Fatemi, and Azra Alizad. Comb-push ultrasound shear elastography of breast masses: Initial results show promise. *PloS one*, 10(3):e0119398, 2015.

[113] Mohammad Mehrmohammadi, Pengfei Song, Duane D Meixner, Robert T Fazzio, Shigao Chen, James F Greenleaf, Mostafa Fatemi, and Azra Alizad. Comb-push ultrasound shear elastography (cuse) for evaluation of thyroid nodules: preliminary in vivo results. *IEEE transactions on medical imaging*, 34(1):97–106, 2014.

[114] **Rémi Delaunay**, Yipeng Hu, and Tom Vercauteren. An unsupervised approach to ultrasound elastography with end-to-end strain regularisation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 573–582. Springer, 2020.

[115] J. Ophir, I. Cespedes, B. Garra, H. Ponnekanti, Y. Huang, and N. Maklad. Elastography: Ultrasonic imaging of tissue strain and elastic modulus in vivo. *European Journal of Ultrasound*, 3(1):49–70, 1996.

[116] Hairong Shi and Tomy Varghese. Two-dimensional multi-level strain estimation for discontinuous tissue. *Physics in Medicine and Biology*, 52(2):389–401, 2007.

[117] Hao Chen, Hairong Shi, and Tomy Varghese. Improvement of elastographic displacement estimation using a two-step cross-correlation method. *Ultrasound in medicine & biology*, 33(1):48–56, 2007.

[118] T. Varghese, E. E. Konofagou, J. Ophir, S. K. Alam, and M. Bilgen. Direct strain estimation in elastography using spectral cross-correlation. *Ultrasound in Medicine and Biology*, 26(9):1525–1537, 2000.

[119] Sharmin R. Ara, Faisal Mohsin, Farzana Alam, Sharmin Akhtar Rupa, Rayhana Awwal, Soo Yeol Lee, and Md Kamrul Hasan. Phase-based direct average strain estimation for elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 60(11):2266–2283, 2013.

[120] Reza Zahiri-Azar and Septimiu E Salcudean. Motion estimation in ultrasound images using time domain cross correlation with prior estimates. *IEEE Transactions on Biomedical Engineering*, 53(10):1990–2000, 2006.

[121] Jianwen Luo and Elisa Konofagou. A fast normalized cross-correlation calculation method for motion estimation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(6):1347–1357, 2010.

[122] Jiaqi Wang, Qinghua Huang, and Xin Zhang. Ultrasound elastography based on the normalized cross-correlation and the PSO algorithm. *2017 4th International Conference on Systems and Informatics, ICSAI 2017*, 2018:1131–1135, 2018.

[123] Fai Yeung, Stephen F. Levinson, and Kevin J. Parker. Multilevel and motion model-based ultrasonic speckle tracking algorithms. *Ultrasound in Medicine and Biology*, 24(3):427–441, 1998.

[124] Francesco Viola and William F. Walker. A comparison of the performance of time-delay estimators in medical ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 50(4):392–401, 2003.

[125] Barry H. Friemel, Laurence N. Bohs, and Gregg E. Trahey. Relative performance of two-dimensional speckle-tracking techniques: normalized correlation, non-normalized correlation and sum-absolute-difference. *Proceedings of the IEEE Ultrasonics Symposium*, 2:1481–1484, 1995.

[126] Adrian Basarab, Pierre Gueth, Hervé Liebgott, and Philippe Delachartre. Two-dimensional least-squares estimation for motion tracking in ultrasound elastography. *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings*, pages 2155–2158, 2007.

[127] Claire Pellot-Barakat, Frédérique Frouin, Michael F Insana, and Alain Herment. Ultrasound elastography based on multiscale estimations of regularized displacement fields. *IEEE transactions on medical imaging*, 23(2):153–163, 2004.

[128] Andrey Kuzmin, Aaron M. Zakrzewski, Brian W. Anthony, and Victor Lempitsky. Multi-frame elastography using a handheld force-controlled ultrasound probe. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 62(8):1486–1500, 2015.

[129] Hassan Rivaz, Emad Boctor, Pezhman Foroughi, Richard Zellars, Gabor Fichtinger, and Gregory Hager. Ultrasound elastography: a dynamic programming approach. *IEEE transactions on medical imaging*, 27(10):1373–1377, 2008.

[130] Hoda Sadat Hashemi and Hassan Rivaz. Global Time-Delay Estimation in Ultrasound Elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 64(10):1625–1636, 2017.

[131] Md Golam Kibria and Hassan Rivaz. Gluenet: Ultrasound elastography using convolutional neural network. In *Simulation, Image Processing, and Ultrasound Systems for Assisted Diagnosis and Navigation*, pages 21–28. Springer, 2018.

[132] Alexey Dosovitskiy, Philipp Fischery, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. *Proceedings of the IEEE International Conference on Computer Vision*, 2015:2758–2766, 2015.

[133] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017.

[134] Sitong Wu, Zhifan Gao, Zhi Liu, Jianwen Luo, Heye Zhang, and Shuo Li. Direct reconstruction of ultrasound elastography using an end-to-end deep neural network. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11070:374–382, 2018.

[135] Zhifan Gao, Sitong Wu, Zhi Liu, Jianwen Luo, Heye Zhang, Mingming Gong, and Shuo Li. Learning the implicit strain reconstruction in ultrasound elastography using privileged information. *Medical image analysis*, 58:101534, 2019.

[136] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018.

[137] Ali KZ Tehrani and Hassan Rivaz. Displacement estimation in ultrasound elastography using pyramidal convolutional neural network. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2020.

[138] Bo Peng, Yuhong Xian, Quan Zhang, and Jingfeng Jiang. Neural-network-based motion tracking for breast ultrasound strain elastography: An initial

assessment of performance and feasibility. *Ultrasonic imaging*, 42(2):74–91, 2020.

[139] Shengyu Zhao, Tingfung Lau, Ji Luo, I Eric, Chao Chang, and Yan Xu. Unsupervised 3d end-to-end medical image registration with volume tweening network. *IEEE journal of biomedical and health informatics*, 24(5):1394–1404, 2019.

[140] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019.

[141] M. Mirzaei, A. Asif, and H. Rivaz. Combining Total Variation Regularization with Window-based Time Delay Estimation in Ultrasound Elastography. *IEEE Transactions on Medical Imaging*, 38(12):1–1, 2019.

[142] Jack Kiefer and Jacob Wolfowitz. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, pages 462–466, 1952.

[143] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[144] Nahian Siddique, Sidike Paheding, Colin P Elkin, and Vijay Devabhaktuni. U-net and its variants for medical image segmentation: A review of theory and applications. *IEEE Access*, 2021.

[145] Yipeng Hu, Marc Modat, Eli Gibson, Nooshin Ghavami, Ester Bonmati, Caroline M. Moore, Mark Emberton, J. Alison Noble, Dean C. Barratt, and Tom Vercauteren. Label-driven weakly-supervised learning for multimodal deformarle image registration. *Proceedings - International Symposium on Biomedical Imaging*, 2018-April(November):1070–1074, 2018.

[146] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[147] E Gibson, W Li, C Sudre, L Fidon, DI Shakir ... programs in Biomedicine, and undefined 2018. NiftyNet: a deep-learning platform for medical imaging. *arXiv:1709.03485 [cs]*, 158:113–122, 2017.

[148] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[149] Faouzi Kallel and Jonathan Ophir. A least-squares strain estimator for elastography. *Ultrasonic imaging*, 19(3):195–208, 1997.

[150] JΦrgen Arendt Jensen and Niels Bruun Svendsen. Calculation of Pressure Fields from Arbitrarily Shaped, Apodized, and Excited Ultrasound Transducers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 39(2):262–267, 1992.

[151] Jørgen Arendt Jensen. Field: A program for simulating ultrasound systems. In *10TH NORDICBALTIC CONFERENCE ON BIOMEDICAL IMAGING, VOL. 4, SUPPLEMENT 1, PART 1: 351–353*. Citeseer, 1996.

[152] Geoff E Tupholme. Generation of acoustic pulses by baffled plane pistons. *Mathematika*, 16(2):209–224, 1969.

[153] Peter R Stepanishen. Transient radiation from pistons in an infinite planar baffle. *The Journal of the Acoustical Society of America*, 49(5B):1629–1638, 1971.

[154] Jørgen Arendt Jensen. Calculation of B-mode image of cyst phantom. `https://field-ii.dk/?examples/cyst_phantom/cyst_phantom.html`. [Online; accessed 16-March-2022].

[155] N Frulio and Herve Trillaud. Ultrasound elastography in liver. *Diagnostic and interventional imaging*, 94(5):515–534, 2013.

[156] Abbas Samani, Judit Zubovits, and Donald Plewes. Elastic moduli of normal and pathological human breast tissues: an inversion-technique-based investigation of 169 samples. *Physics in medicine & biology*, 52(6):1565, 2007.

[157] Qianqian Fang and David A Boas. Tetrahedral mesh generation from volumetric binary and grayscale images. In *2009 IEEE international symposium on biomedical imaging: from nano to macro*, pages 1142–1145. Ieee, 2009.

[158] Rüdiger Göbl, Nassir Navab, and Christoph Hennersperger. SUPRA: open-source software-defined ultrasound processing for real-time applications: A 2D and 3D pipeline from beamforming to B-mode. *International Journal of Computer Assisted Radiology and Surgery*, 13(6):759–767, 2018.

[159] Yu Wang, Emily Helminen, and Jingfeng Jiang. Building a virtual simulation platform for quasistatic breast ultrasound elastography using open source software: A preliminary investigation. *Medical Physics*, 42(9):5453–5466, 2015.

[160] **Rémi Delaunay**, Yipeng Hu, and Tom Vercauteren. An unsupervised learning approach to ultrasound strain elastography with spatio-temporal consistency. *Physics in Medicine and Biology*, 66(17), 2021.

[161] Jingfeng Jiang, Timothy J Hall, and Amy M Sommer. A novel performance descriptor for ultrasonic strain imaging: A preliminary study. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 53(6):1088–1102, 2006.

[162] Graham Treece, Joel Lindop, Lujie Chen, James Housden, Richard Prager, and Andrew Gee. Real-time quasi-static ultrasound elastography. *Interface Focus*, 1(4):540–552, 2011.

[163] Pezhman Foroughi, Hyun-Jae Kang, Daniel A Carnegie, Mark G van Vledder, Michael A Choti, Gregory D Hager, and Emad M Boctor. A freehand ultra-

sound elastography system with tracking for in vivo applications. *Ultrasound in medicine & biology*, 39(2):211–225, 2013.

[164] Abdelrahman Zayed and Hassan Rivaz. Fast strain estimation and frame selection in ultrasound elastography using machine learning. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2020.

[165] R Chandrasekhar, J Ophir, T Krouskop, and K Ophir. Elastographic image quality vs. tissue motion in vivo. *Ultrasound in medicine & biology*, 32(6):847–855, 2006.

[166] Tomy Varghese and Jonathan Ophir. Performance optimization in elastography: Multicompression with temporal stretching. *Ultrasonic imaging*, 18(3):193–214, 1996.

[167] Mark A Lubinski, Stanislav Y Emelianov, and Matthew O'Donnell. Adaptive strain estimation using retrospective processing [medical us elasticity imaging]. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 46(1):97–107, 1999.

[168] Joel E Lindop, Graham M Treece, Andrew H Gee, and Richard W Prager. An intelligent interface for freehand strain imaging. *Ultrasound in medicine & biology*, 34(7):1117–1128, 2008.

[169] SHI Xingjian, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015.

[170] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[171] Amaia Salvador, Miriam Bellver, Victor Campos, Manel Baradad, Ferran Marques, Jordi Torres, and Xavier Giro-i Nieto. Recurrent neural networks for semantic instance segmentation. *arXiv preprint arXiv:1712.00617*, 2017.

[172] Zbigniew Wojna, Vittorio Ferrari, Sergio Guadarrama, Nathan Silberman, Liang-Chieh Chen, Alireza Fathi, and Jasper Uijlings. The devil is in the decoder: Classification, regression and gans. *International Journal of Computer Vision*, 127(11):1694–1706, 2019.

[173] Rüdiger Göbl, Nassir Navab, and Christoph Hennersperger. Supra: open-source software-defined ultrasound processing for real-time applications. *International Journal of Computer Assisted Radiology and Surgery*, 13(6):759–767, 2018.

[174] Jonathan Porée, Damien Garcia, Boris Chayer, Jacques Ohayon, and Guy Cloutier. Noninvasive vascular elastography with plane strain incompressibility assumption using ultrafast coherent compound plane wave imaging. *IEEE transactions on medical imaging*, 34(12):2618–2631, 2015.

[175] **Rémi Delaunay**, Yipeng Hu, and Tom Vercauteren. An unsupervised learning-based shear wave tracking method for ultrasound elastography. In *SPIE in Medical Imaging 2022*. International Society for Optics and Photonics, 2022.

[176] Mark L Palmeri, Stephen A McAleavey, Gregg E Trahey, and Kathryn R Nightingale. Ultrasonic tracking of acoustic radiation force-induced displacements in homogeneous media. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 53(7):1300–1313, 2006.

[177] Mahmoud Derakhshan Horeh, Amir Asif, and Hassan Rivaz. Regularized tracking of shear-wave in ultrasound elastography. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6264–6268. IEEE, 2017.

[178] Hassan Rivaz, Emad Boctor, Pezhman Foroughi, Richard Zellars, Gabor Fichtinger, and Gregory Hager. Ultrasound elastography: a dynamic programming approach. *IEEE transactions on medical imaging*, 27(10):1373–1377, 2008.

[179] Thanasis Loupas, JT Powers, and Robert W Gill. An axial velocity estimator for ultrasound blood flow imaging, based on a full evaluation of the doppler equation by means of a two-dimensional autocorrelation approach. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 42(4):672–688, 1995.

[180] Tingting He, Bo Peng, Pengcheng Chen, and Jingfeng Jiang. Performance assessment of motion tracking methods in ultrasound-based shear wave elastography. In *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 3643–3648. IEEE, 2020.

[181] Derek Y Chan, D Cody Morris, Thomas J Polascik, Mark L Palmeri, and Kathryn R Nightingale. Deep convolutional neural networks for displacement estimation in ARFI imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 68(7):2472–2481, 2021.

[182] Shahed Ahmed, Uday Kamal, and Md Kamrul Hasan. DSWE-Net: A deep learning approach for shear wave elastography and lesion segmentation using single push acoustic radiation force. *Ultrasonics*, 110:106283, 2021.

[183] Mark L Palmeri, Amy C Sharma, Richard R Bouchard, Roger W Nightingale, and Kathryn R Nightingale. A finite-element method model of soft tissue response to impulsive acoustic radiation force. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 52(10):1699–1712, 2005.

[184] Joyce McLaughlin and Daniel Renzi. Using level set based inversion of arrival times to recover shear wave speed in transient elastography and supersonic imaging. *Inverse Problems*, 22(2):707, 2006.

[185] Jonathan Shapey, Thomas Dowrick, **Rémi Delaunay**, Eleanor C Mackle, Stephen Thompson, Mirek Janatka, Roland Guichard, Anastasis Georgoulas, David Pérez-Suárez, Robert Bradford, Shakeel Saeed, Sébastien Ourselin, Matthew Clarckson, and Tom Vercauteren. Integrated multi-modality image-guided navigation for neurosurgery: open-source software platform using

state-of-the-art clinical hardware. *International Journal of Computer Assisted Radiology and Surgery*, pages 1–10, 2021.

[186] L Dade Lunsford and Dan Leksell. The leksell system. In *Modern stereotactic neurosurgery*, pages 27–46. Springer, 1988.

[187] Ian J. Gerard, Marta Kersten-Oertel, Kevin Petrecca, Denis Sirhan, Jeffery A. Hall, and D. Louis Collins. Brain shift in neuronavigation of brain tumors: A review. *Medical Image Analysis*, 35:403–420, 2017.

[188] Florian Kral, Elisabeth J. Puschban, Herbert Riechelmann, and Wolfgang Freysinger. Comparison of optical and electromagnetic tracking for navigated lateral skull base surgery. *International Journal of Medical Robotics and Computer Assisted Surgery*, 9(2):247–252, 2013.

[189] Andinet Enquobahrie, Patrick Cheng, Kevin Gary, Luis Ibanez, David Gobbi, Frank Lindseth, Ziv Yaniv, Stephen Aylward, Julien Jomier, and Kevin Cleary. The image-guided surgery toolkit igstk: an open source c++ software toolkit. *Journal of digital imaging*, 20(1):21–33, 2007.

[190] Junichi Tokuda, Gregory S. Fischer, Xenophon Papademetris, Ziv Yaniv, Luis Ibanez, Patrick Cheng, Haiying Liu, Jack Blevins, Jumpei Arata, Alexandra J. Golby, Tina Kapur, Steve Pieper, Everette C. Burdette, Gabor Fichtinger, Clare M. Tempany, and Nobuhiko Hata. OpenIGTLink: An open network protocol for image-guided therapy environment. *International Journal of Medical Robotics and Computer Assisted Surgery*, 5(4):423–434, 2009.

[191] Andras Lasso, Tamas Heffter, Adam Rankin, Csaba Pinter, Tamas Ungi, and Gabor Fichtinger. PLUS: Open-source toolkit for ultrasound-guided intervention systems. *IEEE Transactions on Biomedical Engineering*, 61(10):2527–2537, 2014.

[192] Matthew J. Clarkson, Gergely Zombori, Steve Thompson, Johannes Totz, Yi Song, Miklos Espak, Stian Johnsen, David Hawkes, and Sébastien Ourselin.

The NifTK software platform for image-guided interventions: platform overview and NiftyLink messaging. *International Journal of Computer Assisted Radiology and Surgery*, 10(3):301–316, 2015.

[193] Ivo Wolf, Marcus Vetter, Ingmar Wegner, Thomas Böttger, Marco Nolden, Max Schöbinger, Mark Hastenteufel, Tobias Kunert, and Hans Peter Meinzer. The medical imaging interaction toolkit. *Medical Image Analysis*, 9(6):594–604, 2005.

[194] Andriy Fedorov, Reinhard Beichel, Jayashree Kalpathy-Cramer, Julien Finet, Jean Christophe Fillion-Robin, Sonia Pujol, Christian Bauer, Dominique Jennings, Fiona Fennessy, Milan Sonka, John Buatti, Stephen Aylward, James V. Miller, Steve Pieper, and Ron Kikinis. 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging*, 30(9):1323–1341, 2012.

[195] Simon Drouin, Anna Kochanowska, Marta Kersten-Oertel, Ian J. Gerard, Rina Zelmann, Dante De Nigris, Silvain Bériault, Tal Arbel, Denis Sirhan, Abbas F. Sadikot, Jeffery A. Hall, David S. Sinclair, Kevin Petrecca, Rolando F. DelMaestro, and D. Louis Collins. IBIS: an OR ready open-source platform for image-guided neurosurgery. *International Journal of Computer Assisted Radiology and Surgery*, 12(3):363–378, 2017.

[196] Christian Askeland, Ole Vegard Solberg, Janne Beate Lervik Bakeng, Ingerid Reinertsen, Geir Arne Tangen, Erlend Fagertun Hofstad, Daniel Høyer Iversen, Cecilie Våpenstad, Tormod Selbekk, Thomas Langø, et al. Custusx: an open-source research platform for image-guided therapy. *International journal of computer assisted radiology and surgery*, 11(4):505–519, 2016.

[197] Tom Dowrick. scikit-surgerybk : Communicate with BK ultrasound scanners via TCP., 2019.

[198] Slicer-IGT developer community. U11: Pivot Calibration. `http://www.`

slicerigt.org/wp/user-tutorial/. [Online; accessed 13-September-2021].

[199] Slicer-IGT developer community. U34: Ultrasound volume reconstruction. http://www.slicerigt.org/wp/user-tutorial/. [Online; accessed 13-September-2021].

[200] Rivo Öpik, Andres Hunt, Asko Ristolainen, Patrick M Aubin, and Maarja Kruusmaa. Development of high fidelity liver and kidney phantom organs for use with robotic surgical systems. In *2012 4th IEEE RAS & EMBS international conference on biomedical robotics and biomechatronics (BioRob)*, pages 425–430. IEEE, 2012.

[201] M Janatka. *Structured Light Assisted Surface Registration for Augmented Reality in Robotic Minimally Invasive Surgery using the da Vinci Research Kit*. PhD thesis, Imperial College London, 2015.

[202] Patrick Brandao, Dimitris Psychogyios, Evangelos Mazomenos, Danail Stoyanov, and Mirek Janatka. Hapnet: hierarchically aggregated pyramid network for real-time stereo matching. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, 9(3):219–224, 2021.

[203] Eleanor C Mackle, Jonathan Shapey, Efthymios Maneas, Shakeel R Saeed, Robert Bradford, Sebastien Ourselin, Tom Vercauteren, and Adrien E Desjardins. Patient-specific polyvinyl alcohol phantom fabrication with ultrasound and x-ray contrast for brain tumor surgery planning. *JoVE (Journal of Visualized Experiments)*, page e61344, 2020.

[204] Laurence Mercier, Rolando F Del Maestro, Kevin Petrecca, Anna Kochanowska, Simon Drouin, Charles XB Yan, Andrew L Janke, Sean Jy-Shyang Chen, and D Louis Collins. New prototype neuronavigation system based on preoperative imaging and intraoperative freehand ultrasound: system description and validation. *International journal of computer assisted radiology and surgery*, 6(4):507–522, 2011.

[205] Matthew Toews and William M Wells. Phantomless auto-calibration and online calibration assessment for a tracked freehand 2-d ultrasound probe. *IEEE transactions on medical imaging*, 37(1):262–272, 2017.

[206] Senay Mihcin, Hani J Marcus, **Rémi Delaunay**, Efthymios Maneas, Wenfeng Xia, Sebastien Ourselin, Adrien Desjardin, Neil Dorward, and Tom Vercauteren. Preclinical assessment of a novel intraoperative ultrasound probe for transsphenoidal surgery. In *Computer Assisted Radiology and Surgery Proceedings of the 32nd International Congress and Exhibition*, pages 57–59. IJCARS, 2018.

[207] Ivan Cabrilo*, **Rémi Delaunay***, Clare L Heaysman, Sebastien Ourselin, Valentina Vitiello, Tom Vercauteren, Hani J Marcus, and Neil L Dorward. A novel intraoperative ultrasound probe for transsphenoidal surgery: First-in-human study. *Surgical Innovation*, page 15533506211031091, 2021.

[208] Mark E Molitch. Diagnosis and treatment of pituitary adenomas: a review. *Jama*, 317(5):516–524, 2017.

[209] Shereen Ezzat, Sylvia L Asa, William T Couldwell, Charles E Barr, William E Dodge, Mary Lee Vance, and Ian E McCutcheon. The prevalence of pituitary adenomas: a systematic review. *Cancer: Interdisciplinary International Journal of the American Cancer Society*, 101(3):613–619, 2004.

[210] OM Dekkers, AM Pereira, Fj Roelfsema, JHC Voormolen, KJ Neelis, MA Schroijen, JWA Smit, and JA Romijn. Observation alone after transsphenoidal surgery for nonfunctioning pituitary macroadenoma. *The Journal of Clinical Endocrinology & Metabolism*, 91(5):1796–1801, 2006.

[211] Hani J. Marcus, Tom Vercauteren, Sebastien Ourselin, and Neil L. Dorward. Intraoperative Ultrasound in Patients Undergoing Transsphoidal Surgery for Pituitary Adenoma: Systematic Review. *World Neurosurgery*, 106(December 2016):680–685, 2017.

[212] Sylvia L. Asa and Shereen Ezzat. The Cytogenesis and Pathogenesis of Pituitary Adenomas 1. *Endocrine Reviews*, 19(6):798–827, 1998.

[213] Ferdinand Roelfsema, Nienke R Biermasz, and Alberto M Pereira. Clinical factors involved in the recurrence of pituitary adenomas after surgical remission: a structured review and meta-analysis. *Pituitary*, 15(1):71–83, 2012.

[214] J K Liu, K Das, M H Weiss, E R Laws, and W T Couldwell. The history and evolution of transsphenoidal surgery. *Journal of neurosurgery*, 95(6):1083–96, 2001.

[215] Richard F Schmidt, Osamah J Choudhry, Ramya Takkellapati, Jean Anderson Eloy, William T Couldwell, and James K Liu. Hermann schloffer and the origin of transsphenoidal pituitary surgery. *Neurosurgical focus*, 33(2):E5, 2012.

[216] W. R. Henderson. The pituitary adenomata. A follow-up study of the surgical results in 338 cases. (DR. HARVEY CUSHING'S SERIES). *British Journal of Surgery*, 26(104):811–921, 1939.

[217] J Hardy and S M Wigser. Trans-sphenoidal surgery of pituitary fossa tumors with televised radiofluoroscopic control. *Journal Neurosurg*, 23(6):612–9, 1965.

[218] Chirag D Gandhi, Lana D Christiano, Jean Anderson Eloy, Charles J Prestigiacomo, and Kalmon D Post. The historical evolution of transsphenoidal surgery: facilitation by technological advances. *Neurosurgical focus*, 27(3):E8, 2009.

[219] Hae-Dong Jho. Endoscopic pituitary surgery. *Pituitary*, 2(2):139–154, 1999.

[220] W. Jeffrey Elias, James B. Chadduck, Tord D. Alden, and Laws Edward R. Frameless stereotaxy for transsphenoidal surgery. *Neurosurgery*, 45(2):271–277, 1999.

[221] Pietro Mortini, Marco Losa, Raffaella Barzaghi, Nicola Boari, and Massimo Giovanelli. Results of transsphenoidal surgery in a large series of patients with pituitary adenoma. *Neurosurgery*, 56(6):1222–1233, 2005.

[222] P Cappabianca, LM Cavallo, A Colao, M Del Basso De Caro, F Esposito, S Cirillo, G Lombardi, and E De Divitiis. Endoscopic endonasal transsphenoidal approach: outcome analysis of 100 consecutive procedures. *Min-Minimally Invasive Neurosurgery*, 45(04):193–200, 2002.

[223] Marco Locatelli, Mary Lee Vance, and Edward R. Laws. Clinical review: The strategy of immediate reoperation for transsphenoidal surgery for Cushing's disease. *Journal of Clinical Endocrinology and Metabolism*, 90(9):5478–5482, 2005.

[224] Annie S. Dubuisson, Albert Beckers, and Achille Stevenaert. Classical pituitary tumour apoplexy: Clinical features, management and outcomes in a series of 24 patients. *Clinical Neurology and Neurosurgery*, 109(1):63–70, 2007.

[225] Javier Estrada, Mauro Boronat, Mercedes Mielgo, Rosa Magallón, Isabel Millán, Santiágo Díez, Tomás Lucas, and Balbino Barceló. The long-term outcome of pituitary irradiation after unsuccessful transsphenoidal surgery in cushing's disease. *New England Journal of Medicine*, 336(3):172–177, 1997.

[226] Robert F. Dallapiazza, Yuval Grober, Robert M. Starke, Edward R. Laws, and John A. Jane. Long-term results of Endonasal Endoscopic Transsphenoidal resection of nonfunctioning pituitary Macroadenomas. *Neurosurgery*, 76(1):42–52, 2015.

[227] Prashant Chittiboina. imri during transsphenoidal surgery. *Neurosurgery Clinics*, 28(4):499–512, 2017.

[228] Peter T Sylvester and Michael R Chicoine. Intraoperative imaging for pituitary surgery. In *Transsphenoidal Surgery*, pages 259–275. Springer, 2017.

[229] Jared F Sweeney, Heather Smith, AmiLyn Taplin, Eric Perloff, and Matthew A Adamo. Efficacy of intraoperative ultrasonography in neurosurgical tumor resection. *Journal of Neurosurgery: Pediatrics*, 21(5):504–510, 2018.

[230] Richard Mair, James Heald, Ion Poeata, and Marcel Ivanov. A practical grading system of ultrasonographic visibility for intracerebral lesions. *Acta neurochirurgica*, 155(12):2293–2298, 2013.

[231] Yasushi Ota and Ishikawa Mami. Ultrasonography imaging during nasal endoscopic transsphenoidal surgery. *Orl*, 75(1):27–31, 2013.

[232] Kazunori Arita, Kaoru Kurisu, Atsushi Tominaga, Hitoshi Kawamoto, Koji Iida, Tatsuya Mizoue, Basant Pant, and Tohru Uozumi. Trans-sellar color doppler ultrasonography during transsphenoidal surgery. *Neurosurgery*, 42(1):81–85, 1998.

[233] Mami Ishikawa, Yasushi Ota, Naohiro Yoshida, Yukiko Iino, Yuichi Tanaka, and Eiju Watanabe. Endonasal ultrasonography-assisted neuroendoscopic transsphenoidal surgery. *Acta Neurochirurgica*, 157(5):863–868, 2015.

[234] UJ Knappe, M Engelbach, K Konz, H-J Lakomek, W Saeger, R Schönmayr, and WA Mann. Ultrasound-assisted microsurgery for cushing's disease. *Experimental and clinical endocrinology & diabetes*, 119(04):191–200, 2011.

[235] Ole Solheim, Tormod Selbekk, Lasse Løvstakken, Geir A. Tangen, Ole V. Solberg, Tonni F. Johansen, Johan Cappelen, and Geirmund Unsgård. Intrasellar ultrasound in transsphenoidal surgery: A novel technique. *Neurosurgery*, 66(1):173–185, 2010.

[236] Ole Solheim, Tonni Franke Johansen, Johan Cappelen, Geirmund Unsgård, and Tormod Selbekk. Transsellar ultrasound in pituitary surgery with a designated probe: early experiences. *Operative Neurosurgery*, 12(2):128–134, 2016.

[237] Wenfeng Xia, Daniil I. Nikitichev, Jean Martial Mari, Simeon J. West, Rosalind Pratt, Anna L. David, Sebastien Ourselin, Paul C. Beard, and Adrien E. Desjardins. Performance characteristics of an interventional multispectral photoacoustic imaging system for guiding minimally invasive procedures. *Journal of Biomedical Optics*, 20(8):086005, 2015.

[238] Ren Wen Ho, Hsiu Mei Huang, and Jih Tsun Ho. The influence of pituitary adenoma size on vision and visual outcomes after trans-sphenoidal adenectomy : A report of 78 cases. *Journal of Korean Neurosurgical Society*, 57(1):23–31, 2015.

[239] Charles J. Pavlin. *Diagnostic Ultrasound*. CRC Press, 2006.

[240] Jose Gutierrez, Ken Cheung, Ahmet Bagci, Tatjana Rundek, Noam Alperin, Ralph L. Sacco, Clinton B. Wright, and Mitchell S.V. Elkind. Brain Arterial Diameters as a Risk Factor for Vascular Events. *Journal of the American Heart Association*, 4(8):e002289, 2015.

[241] Joseph G Parravano, Abel Toledo, and Walter Kucharczyk. Dimensions of the optic nerves, chiasm, and tracts: Mr quantitative comparison between patients with optic atrophy and normals. *Journal of computer assisted tomography*, 17(5):688–690, 1993.

[242] Jon Ramm-Pettersen, Jon Berg-Johnsen, Per Kristian Hol, Sumit Roy, Jens Bollerslev, Thomas Schreiner, and Eirik Helseth. Intra-operative mri facilitates tumour resection during trans-sphenoidal surgery for pituitary adenomas. *Acta neurochirurgica*, 153(7):1367–1373, 2011.

[243] Hasan A Zaidi, Kenneth De Los Reyes, Garni Barkhoudarian, Zachary N Litvack, Wenya Linda Bi, Jordina Rincon-Torroella, Srinivasan Mukundan, Ian F Dunn, and Edward R Laws. The utility of high-resolution intraoperative mri in endoscopic transsphenoidal surgery for pituitary macroadenomas: early experience in the advanced multimodality image guided operating suite. *Neurosurgical focus*, 40(3):E18, 2016.

[244] Ryosuke Mori, Tatsuhiro Joki, Yoshinori Matsuwaki, Kostadin Karagiozov, Yuichi Murayama, and Toshiaki Abe. Initial experience of real-time intraoperative c-arm computed–tomography-guided navigation surgery for pituitary tumors. *World neurosurgery*, 79(2):319–326, 2013.

[245] Ben A Strickland, Joshua Lucas, Brianna Harris, Edwin Kulubya, Joshua Bakhsheshian, Charles Liu, Bozena Wrobel, John D Carmichael, Martin Weiss, and Gabriel Zada. Identification and repair of intraoperative cerebrospinal fluid leaks in endonasal transsphenoidal pituitary surgery: surgical experience in a series of 1002 patients. *Journal of Neurosurgery*, 129(2):425–429, 2017.

[246] Jérémie Fromageau, Jean-Luc Gennisson, Cédric Schmitt, Roch L Maurice, Rosaire Mongrain, and Guy Cloutier. Estimation of polyvinyl alcohol cryogel mechanical properties with four ultrasound elastography methods and comparison with gold standard testings. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 54(3):498–509, 2007.