Accurate Simulation of Low-Intensity Transcranial Ultrasound Propagation for Neurostimulation

James Robertson

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

October 18, 2017

I, James Robertson, 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.

Essentially, all models are wrong, but some are useful

–George Box

Abstract

Neural stimulation with low-intensity ultrasound is a potentially transformative technology with applications in therapy and research. To develop, it will require ultrasound to be tightly focused on brain structures with accurate spatial targeting and fine control over the ultrasound amplitude at the target. However, the skull is an impediment to the effective focusing of ultrasound. Simulations of ultrasound propagation through acoustic property maps derived from medical images can be used to derive focusing drive signals for multi-element arrays. Focusing effectiveness is dependent on the fidelity of the numerical simulations. In combination with MRI based treatment verification, model based focusing has been used to focus high-intensity ultrasound onto the brain for ablation. This thesis presents a thorough and systematic study of the simulation parameters required to achieve effective transcranial focusing. The literature on ultrasonic neurostimulation, transcranial ultrasonic focusing, and the derivation of property maps from medical images is reviewed. The sampling criteria required to ensure numerical accuracy for the k-space pseudospectral time domain simulation scheme is established through testing of individual sources of numerical error, and convergence testing of a simulated time-reversal protocol. With numerical accuracy assured, the importance of acoustic property maps is examined through simulations to determine the sensitivity of intracranial fields to the properties of the skull layer. These results are corroborated by matching experimental measurements of ultrasound propagation through skull bone phantoms with spatially registered simulations. Finally, the impact of image related homogenisation and loss of internal bone structure is determined using simulations through co-registered clinical CT and micro-CT data of the skull.

Impact Statement

This thesis comprises an exhaustive study of the potential sources of error and inaccuracy that can arise when using numerical simulations of ultrasound propagation through the skull to focus ultrasound on the deep brain. Model based focusing, as well as simulation of transcranial ultrasound propagation in general, are important tools in the development of clinical treatments and research methods. For example, models are currently employed in the transcranial focusing of high intensity ultrasound for ablation to treat Essential Tremor.

The thesis begins with a literature review of ultrasonic neurostimulation, the unique challenges involved in applying ultrasound to the brain, and simulation-based methods for transcranial focusing. This review may be used as a comprehensive introduction to the state-of-the-art of this developing technology, and to identify future areas of inquiry.

The first stage of the experimental process focuses on numerical error and establishes the sampling criteria required to reduce error below defined bounds. The initial examination of different sources of numerical error in isolation also serves as a mechanistic introduction to the simulation scheme used, and can be used to determine the specific source of simulation inaccuracy. For those seeking to optimise transcranial propagation simulations, this can help ensure the numerical accuracy of simulations while minimising the computational burden. Likewise, the final numerical convergence criteria computed will be of interest to anyone attempting to use numerical simulations to focus ultrasound onto the brain, and will allow them to carry out simulation of ultrasound transmission through the skull with numerical accuracy assured.

The remainder of the thesis focuses on the simulation inaccuracies that occur due to incorrect mapping of acoustic medium properties inside the simulated domain. Experiments establish the relevant impact of errors in assigned acoustic properties, gross geometry of the skull layer, and the loss of internal structure due to limited image resolution. The results suggest development priorities for future exploration into methods of deriving acoustic property maps from medical image data. This stage also involved the creation and validation of ultrasonic bone phantoms using computer aided design and resin casting techniques. This demonstrated that matching ultrasound propagation to numerical simulation is possible with high accuracy when computer aided stereotaxis is employed. This work also examined the relative merits of two bone phantom materials, including their acoustic property profiles and ease of use for phantom creation. Phantoms are ubiquitous tools in the testing and development of imaging and therapeutic technologies, and the methods and materials explored in this thesis will be of general utility in the design and production of models with a well described geometry.

Acknowledgements

I will begin by thanking my family, through whom all things are possible, and my copy editor and sometime girlfriend, Annie. I couldn't have done it without you.

Since its inception, it has been a great privilege to be a part of the growing Biomedical Ultrasound Group at UCL, and all the members have supported me at one time or another, even if they weren't aware of it. The work contained in this thesis wouldn't have been feasible without many useful technical discussions with Elly Martin and Ben Cox, who helped teach me to enjoy the work I was doing. I should also acknowledge the moral support of Michael Brown, who was almost as bad an influence on me as I was on him.

Within the wider Medical Physics department I have to thank Daniil Nikitichev and Eve Hatten, for their assistance with 3D printing design and production, and Terence Leung and Clare Elwell for first encouraging me to pursue a PhD. It goes without saying that I would not have been able to do so if not for the support of an ESPRC sponsored studentship, which I was very grateful to receive. Further afield, I have to thank Nishant Ravikumar and Zeike Taylor from Imperial College London for the provision of skull and brain meshes, and Jillian Urban and Joel Stitzel from Wake Forest Baptist Health, who were kind enough to walk me through their fantastic CT dataset.

Finally, I have to acknowledge the unwavering support of Dr. Bradley Treeby and take the opportunity to thank him. His combination of patience, humour and guidance has kept me going over the last few years, and I am lucky to have had such an excellent supervisor.

Contents

List of Figures

List of Tables

List of Abbreviations

Acronym	Term
CNS	Central nervous system
HIFU	High intensity focused ultrasound
DBS	Deep brain stimulation
AP/DV/ML	Antero-posterior/Dorso-ventral/Medio-lateral
MCP	Mid-commisural point
TDCS	Transcranial direct current stimulation
TMS	Transcranial magnetic stimulation
UNMS	Ultrasonic neuromodulation & stimulation
EEG	Electroencephalogram
(f)MRI	(Functional) magnetic resonance imaging
S1/S2	Primary/secondary somatosensory cortex
V1	Primary visual cortex
M1	Primary motor cortex
DC	Duty cycle
CW	Continuous wave
TB	Toneburst
LFP	Local field potential
BDNF	Brain derived neurotrophic factor
GABA	Gamma-aminobutyric acid
EMG	Electromyogram
SEP/VEP	Sensory/visual evoked potential
FEF	Frontal eye field
MC	Motor cortex
CSF	Cerebrospinal fluid
IT'IS	Foundation for Research on Information Technologies in Society
CT	X-ray computed tomography
HU	Hounsfield units
FDTD	Finite difference time domain
PSTD	Pseudospectral time domain
BLI	Bandlimited interpolant
PML	Perfectly matched layer
CFL	Courant-Friedreichs-Lewy number
PPW	Points per wavelength
PPP	Points per period
FWHM	Full width half maximum
CAD	Computer aided design

Chapter 1

Introduction

1.1 Brain Stimulation

1.1.1 Neurological Disorders

Neurological and neuropsychiatric disorders encompass a wide range of conditions affecting every segment of the population. Combined, they represent a significant burden in modern society, especially in the developed world. Gustavsson et al. [17] and Wittchen et al. [18] show that disorders of the brain are the largest single cause of morbidity in the EU in terms of disability adjusted life years, and represent a significant economic burden, a burden that is expected to grow over the first half of the 21st century, and compound the effects of an ageing population.

The scientific understanding of the aetiology and mechanisms of neurological disorders varies widely depending on the disease and is often incomplete, although recent advances in functional imaging and cellular neuroscience have driven clinical progress [19]. The brain is an exceedingly complicated organ, the normal function of which involves the interdependent function of multiple neural information circuits of varying scales and complexities alongside more mundane physiological processes such as respiration. Accordingly, the causes of pathology can vary from the gross destruction or loss of neural tissue in the case of stroke and serious degenerative disease, to subtle imbalances in the expression of neurotransmitters and abnormal neuronal firing patterns, to the metabolic and mass effects of tumours. Neurological disorders can affect any level of neural circuit function with a corre-

spondingly wide variation in disease manifestation and symptoms.

The large range of neurological and neuropsychiatric disorders precludes a description of the variety of treatments available, however, many will respond to pharmacological treatments. For disorders that are not specific to the nervous system, these can include more generic treatments such as chemotherapeutic agents for cancers, or immunosuppresants for autoimmune disorders. However, many treatments aim to directly influence neuronal function by enhancing or suppressing the expression or effects of neurotransmitters, or groups of neurotransmitters. This modulation of neurotransmitter function will then either directly or indirectly counteract the pathology and restore normal function or provide relief from symptoms. Examples include the treatment of Parkinson's disorder by attempting to enhance dopaminergic transmission, and the treatment of depression by increasing serotonin and monoamine transmission [20].

While pharmaceuticals remain a core pillar of treatment for neurological disorders, they are subject to a number of serious limitations. Most notable is the presence of the blood-brain barrier, the unique endothelial cell lining found in blood vessels that supply the central nervous system (CNS). The blood-brain barrier allows highly selective communication between the blood vessels and the cells of the CNS, and protects delicate CNS tissues from potentially dangerous agents in the bloodstream [20, 21]. However, an important consequence of its presence is a marked difficulty in using pharmaceuticals introduced intravenously to achieve therapeutic effect in the brain [22]. An important additional consideration is that pharmaceuticals that target neurotransmitters have the potential for serious side effects, due to the lack of specificity in intravenously infused drugs and the ubiquity of neurotransmitters across the central and peripheral nervous systems [20].

More invasive treatment options, such as surgery and radiotherapy, that aim to destroy or remove pathological tissue are also employed when appropriate. A recent development is the Federal Drug Administration's approval of biomedical ultrasound, previously used for transcranial imaging through the temporal bone [23], for the treatment of essential tremor via non-invasive thalamotomy with high in-

tensity focused ultrasound (HIFU) [24]. However, these options have their own associated risks and side effects, compounded by the delicacy and importance of the brain. Furthermore, while often effective medical treatments, these methods are not able to reproduce the more nuanced modulation of neuronal function produced by pharmaceutical agents.

1.1.2 Deep Brain Stimulation

Deep brain stimulation (DBS) of neurological structures such as the thalamus and basal ganglia with implanted electrodes is an established treatment for neurological disorders, most notably Parkinson's disease, with over 100,000 patients treated in the last twenty years [19]. The exact effect of DBS on neural tissue remains unclear despite extensive investigation. There is conflicting evidence as to whether DBS stimulates or suppresses activity at the cellular level [25, 19, 26]. The complexity of neuronal interactions complicates matters, as stimulation or suppression of an individual neurone or nucleus will cause a more complex effect at the circuit level and over varying timescales, for instance via changes in synaptic plasticity [19]. A more recent view is that neurological disorders able to be effectively treated by DBS in fact correspond to pathological firing patterns propagating within a larger neural circuit, such as the well described motor circuits in the case of Parkinson's disease. This pathological firing leads to dysfunction and disease, and targeted DBS is able to disrupt the production of these firing patterns, creating an informational lesion and preventing pathological circuit activity without disabling the function of the circuit itself [19, 27, 28].

Apart from its use as a therapeutic tool, the ability to stimulate deep brain nuclei or specific areas of neural tissue selectively has significant potential as a neuroscientific research technique. The stimulation carried out to determine electrode placement has already been used to advance mapping of deep brain structures and develop understanding of motor circuit function [29]. Likewise, the investigation of DBS as a treatment for other disorders such as obsessive-compulsive disorder and major depressive disorder has gone hand in hand with new discoveries about the role of neural circuits in these conditions [28, 19]. A list of proposed and existing

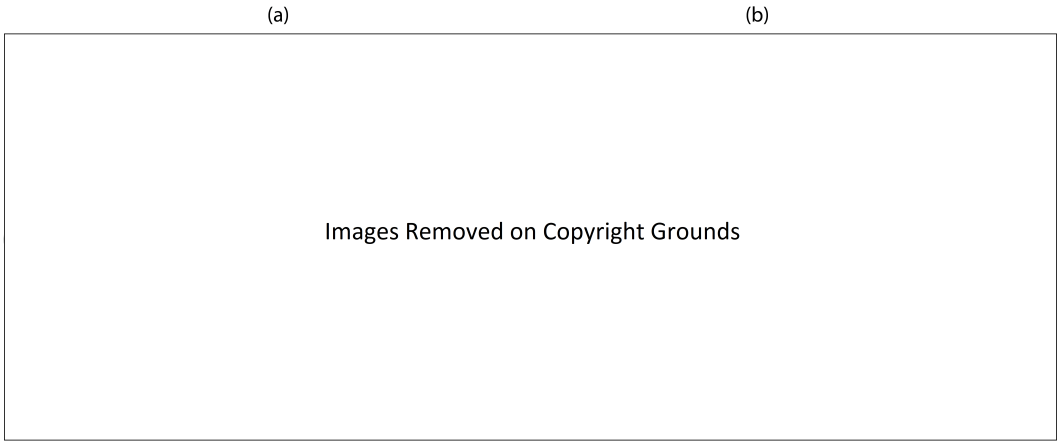

Figure 1.1: Brain anatomy & location of deep brain targets. (a) The gross position of the basal ganglia and the thalamus within the brain. (b) A coronal section showing the arrangement of the different components (reproduced from Leisman et al. [1]).

	Target Structure	AP × DV × ML [mm]	MCP Deviation [mm]
Thalamus	Ventral-Intermediate	10 × 15.8 × 11	17
	Ventral-Anterior	7 × 12.6 × 10	15
	Centro-Medial	8 × 4.5 × 4	14
BG	Globus Pallidus (Ext.)	21.5 × 10 × 3	23
	Globus Pallidus (Int.)	12.5 × 8 × 6	20
	Nucleus Accumbens	9.5 × 10 × 12	21

Table 1.1: A selection of neurological structures in the thalamus and basal ganglia (BG), their Antero-Posterior (AP) × Dorso-Ventral (DV) × Medio-Lateral (ML) dimensions and their distance from the mid-commisural point (MCP) [11].

target structures for DBS and their dimensions is shown in Table 1.1 & Fig. 1.1.

However, the highly invasive nature of DBS with electrodes means that it is typically only considered to supplement limited or failing drug regimens, and it is not employed as a research tool in healthy human subjects [26, 30]. A review of the clinical literature showed that the overall incidence of adverse effects from the surgery or hardware was 9%, with the risk of haemhorrage at 3%, and general potential for serious morbidity [31]. The application of the electrical stimulation can also lead to neurological side effects, including speech and movement problems, and sensory disturbances [25].

In the laboratory, optogenetic stimulation can be used to directly stimulate in-

dividual neurons *in-vivo*, by genetically encoding an optical sensitivity to neuronal ion channels [19]. While this allows unparalleled precision, due to the nature of these techniques they are unlikely to be extended to human subjects. As a result of the clear importance of neurostimulation as a clinical and scientific tool, research into non-invasive alternatives for the purposes of both research and therapy is ongoing.

Transcranial direct current stimulation (TDCS), and transcranial magnetic stimulation (TMS) are non-invasive neurostimulation methods that utilise electric current and strong magnetic fields, respectively, to affect neural tissue through the skull [32]. TMS takes advantage of the inductive properties of magnetic fields and, using static or changing fields, is able to modify neuronal excitability and stimulate neural tissue to generate action potentials [33]. It is a powerful research tool, and is increasingly used in clinical practice as a treatment for a number of neurological conditions, most notably depression, and for improving recovery following stroke [34]. TDCS is only capable of neuromodulation and has found more limited application for the treatment of disease and as a research tool [32].

Unfortunately, these methods are limited in their ability to penetrate deep into brain tissue and generate a tight stimulatory focus, compared to traditional DBS, and are limited to inducing effects on relatively large areas of the cerebral cortex on the surface of the brain [32]. While this has demonstrated potential for the treatment of disorders that are not strongly anatomically localised, such as major depressive disorder or diffuse tissue damage [33, 35], it precludes the tightly focused stimulation possible with electrodes.

1.1.3 Ultrasonic Neurostimulation

History of Ultrasonic Neurostimulation

Ultrasound was first shown to have a galvanistic effect on muscular tissue as early as the late 1920s [36], but the first direct research into the effects of ultrasound on neural tissue occurred in the 1950s. Fry et al. showed that ultrasound was capable of reversibly blocking action potential generation in crayfish ventral abdominal nerves [37], and of suppressing light evoked visual potentials in the lateral genicu-

late nucleus of a cat, a key component of the visual system [38]. Work continued over the latter half of the 20th century, notably by Gavrilov et al. [39] who examined the effects of ultrasound on the peripheral nervous system. Theoretically, ultrasound neuromodulation and stimulation (UNMS) has several advantages over other non-invasive methods such as TMS and TDCS. Firstly, the high transmission of ultrasonic energy through soft-tissue gives it superior depth penetration, even when transmitting through the skull. Secondly, ultrasound has better spatial focusing, with a focal spot size determined by the acoustic wavelength, which could theoretically allow focal stimulation on the scale of millimetres. Renewed contemporary interest in ultrasonic stimulation of the CNS dates from a series of papers published between 2008 and 2010, which demonstrated the feasibility of using ultrasound to directly stimulate neuronal action potentials in *ex-vivo* neural tissue, rather than simply modulate activity. Khraiche et al. [40] and Muratore et al. [41] applied ultrasound to *ex-vivo* hippocampal tissue and observed changes in aggregate excitability, and individual neuronal activity similar to that produced by direct stimulation with electrodes. Tyler et al. [42] examined this effect in a detailed microbiological study, and established that the effects of ultrasound on hippocampal tissue involved activation of voltage gated sodium and calcium channels, transient currents from calcium flow, and the triggering of synaptic vesicle release. Finally, Tufail et al. [43] extended experiments to targeting of the mouse brain *in-vivo*, eliciting observable electroencephalogram (EEG) signals, and a motor response when targeting the motor cortex. Post-sonication staining of the brain confirmed neuronal activity in the targeted locations, and their stimulation protocol was published separately to encourage new work [44]. These initial papers were followed by a range of research examining *in-vivo* stimulation of the CNS in the small animal model. Generally, these papers confirmed earlier observations that ultrasound applied to the brain could elicit motor responses and alterations in brain function confirmed by EEG. Novel effects observed included the use of ultrasound to suppress chemically induced epileptic activity in rats [45], and alteration of extracellular levels of neurotransmitters [46, 47]. Kim et al. [48] combined UNMS with positron emis-

sion tomography to observe increased glucose uptake in thalamic cells, while Yoo et al. [49] used functional magnetic resonance imaging (fMRI) imaging to confirm a neurostimulatory effect in the brains of rabbits. More recently, stimulation was extended to non-human primates by Deffieux et al. [50], who also demonstrated a neuromodulatory effect on high level cognitive behaviour in the form of visual anti-saccade. Neurostimulation was also repeatedly demonstrated to be possible without causing any tissue damage, although the presence of microhaemorrhages in one study at 250 kHz [49] demonstrates the importance of effectively controlling the ultrasound field. Since 2013, UNMS has been extended to human subjects and used to elicit a variety of responses, although studies have thus far been limited to stimulating the superficial, cortical areas of the brain using single element transducers. Mueller et al. [51] and Legon et al. [52] demonstrated that UNMS of the primary somatosensory cortex (S1) resulted in observable change in recorded EEG signals during median nerve stimulation. In a series of papers, Lee et al. [53, 54, 55] extended stimulation to multiple areas of the sensory cortex (S1 and S2) in an attempt to elicit specific sensory responses, and to the visual cortex (V1) to examine the effect on visual processing. S1 and S2 stimulation resulted in perceived sensations in the hands and other parts of the body, including warmth, pain, and parenthesia, although the sensation type could not be controlled as intended. V1 stimulation resulted in EEG responses and the perception of phosphenes. Human UNMS has also been successfully combined with fMRI to monitor brain tissue activation, although results have sometimes been inconsistent [56, 55]. A (non-exhaustive) selection of recent papers demonstrating ultrasonic neurostimulation *in-vivo* in animals and humans are shown in Table 1.2 and Table 1.3, respectively. These tables include the estimated ultrasonic intensity at the stimulation target, and the characteristics of the ultrasonic pulse shape in terms of the parameters detailed in Fig. 1.2.

Mechanisms of Ultrasonic Neurostimulation

Renewed interest has resulted in extensive speculation & investigation into the mechanism underlying ultrasonic neurostimulation. Broadly, the effects of ultrasound on biological tissue can be divided into mechanical and thermal [70]. The

Year	Author	Freq [MHz]	PRF [kHz]	DC [%]	t_{stim} [s]	I_{sppa} [Wcm^{-2}]	Target	Effect
2010	Tufail et al. [43]	0.25-0.5	1.2-3	19-86	0.026-0.333	0.08-0.23	Mouse motor cortex	Motor response, LFP & EMG recordings
		0.25 & 0.35	1.5 & 2	32 & 21	0.333	0.26 & 0.17	Mouse hippocampus	LFP recordings & BDNF release
		0.35	1.5	21		0.17	Mouse whole brain	Neural activation imaged via antibody staining
2011	Min et al. [45]	0.69	0.1	5	180	2.6	Rat epileptic focus	Suppression of epileptic EEG signals
2011	Min et al. [47]	0.65	0.1	5	1200	3.5	Rat thalamus	Alteration of extracellular serotonin & dopamine
2011	Yoo et al. [49]	0.69	0.01-1	5-50	0.5-2	3.3-12.6	Rabbit cortex	Cortical activation imaged via fMRI, suppression of VEP, motor response at higher intensities
2012	Yang et al. [46]	0.65	0.1	5	1200	3.5	Rat thalamus	Change in extracellular GABA levels
2013	King et al. [57]	0.25-0.6	CW	100	0.02-0.48	0.02-16.8	Mouse motor cortex	EMG spiking, motor response
		0.5	0.1-3	2-60	0.04-1.2	1-79		
2013	Younan et al. [58]	0.32	2	50	0.25	17.25	Rat Cortex	Tail and whisker twitches above intensity threshold
2013	Deffieux et al. [50]	0.32	CW	100	0.1	4	Macaque FEF	Change in visual antisaccade latency
2014	Kim et al. [48]	0.35	1	50	0.3	6	Rat thalamus	Motor response, glucose uptake in thalamus
2014	Kim et al. [59].	0.35	CW	100	0.15	7.73	Rat motor cortex	Motor response of tail. CW & high frequencies increase intensity threshold
		0.35 & 0.65	0.06-2.8	30-70	0.2-0.4	5-14		
2014	Mehic et al. [60]	0.5	1.5	15 & 30	0.012-0.059	6.7-26.7	Mouse brain	Variety of motor responses
2014	King et al. [61]	0.5	CW	100	0.3	3	Mouse motor cortex	Motor response, measured via EMG
2015	Kim et al. [62]	0.35	0.02-0.166	1-8.3	150	1-5	Rat visual cortex	Controllable VEP elevation & suppression
2015	Lee et al. [63]	0.25	0.5	50	0.3	3.4-11.8	Sheep sensiromotor cortex	Motor response
						1.7-14.3	Sheep visual cortex	EEG response in the absence of visual stimulation
2016	Ye et al. [64]	0.3-2.9	CW	100	0.08	0.03-127	Mouse brain	Motor response, easier to elicit at low frequencies
2016	Darvas et al. [65]	2	1.05	20	1	7	Rat brain	Distinctive EEG response
2016	Li et al. [66]	1 & 5	1	50	0.3	0.26-0.46	Mouse brain	Motor response, measured via EMG, scaling with intensity
2016	Kamimura et al. [67]	1.9	1	50	1	66-158	Mouse brain	Motor response, eye movement & pupil dilation
2016	Yuan et al. [68]	2.5	0.5	0.08	0.16	9.8	Rat hippocampus	Modulation of gamma oscillations
2017	Dallapiazza et al. [69]	0.22-1.14	0.01	43.7	40	25-30 (SA)	Swine thalamus	Selective suppression of SEP in different thalamic nuclei

Table 1.2: Review of selected recent ultrasonic neuromodulation and neurostimulation literature for animal models *in-vivo*. Some values taken from Naor et al. [12]. PRF — pulse repetition frequency, DC — duty cycle, I_{sppa} — spatial peak pulse average intensity, t_{stim} — stimulation duration, SA — spatial average intensity, LFP — local field potential, BDNF — brain derived neurotrophic factor, GABA — gamma-aminobutyric acid, EEG — electroencephalogram, EMG — Electromyogram, fMRI – functional magnetic resonance imaging, SEP/VEP — sensory/visual evoked potential, FEF — frontal eye field, S1 — primary somatosensory cortex, MC — motor cortex.

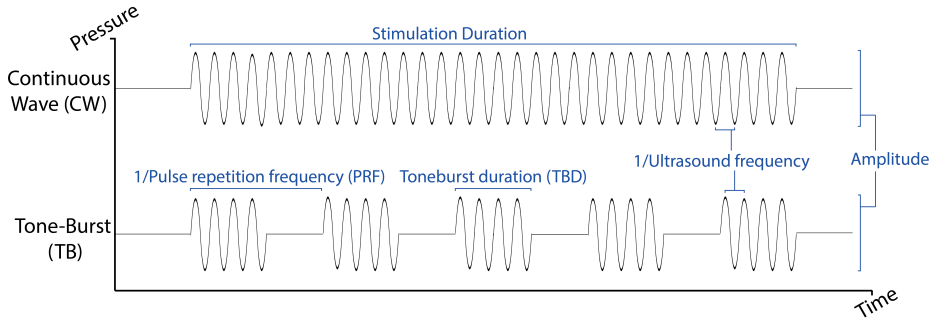

Figure 1.2: Schematic of continuous wave and toneburst waveforms, and the definition of various parameters that define ultrasonic waveforms used in the neurostimulation papers outlined in Table 1.3. Duty Cycle is calculated Duty Cycle = Pulse Repetition Frequency × Toneburst Duration.

Year	Author	Freq [MHz]	PRF [kHz]	DC [%]	t_{stim} [s]	I_{sppa} [Wcm^{-2}]	Target	Effect
2014	Mueller et al. [51]	0.5	1	36	0.5	5.9	S1	Change in EEG potentials after Median nerve stimulation
2015	Lee et al. [53]	0.25	0.5	50	0.3	0.3-2.5	S1 (hand)	Evoked sensations in hand, but also in other areas
2016	Ai et al. [56]	0.5 & 0.86	1	37 & 50	0.5	6	M1/S1	Activation imaged with fMRI
2016	Lee et al. [54]	0.21	0.1	50	0.5	7.0-8.8	S1/S2	Evoked sensations, including warmth & paresthesia
2016	Lee et al. [55]	0.27	0.5	50	0.3	11.6	V1	fMRI activation & perception of phosphenes

Table 1.3: Review of previous neurostimulation studies carried out on humans. PRF — pulse repetition frequency, DC — duty cycle, I_{sppa} — spatial peak pulse average intensity, t_{stim} — stimulation duration, EEG — electroencephalogram, fMRI – functional magnetic resonance imaging, M1/S1/V1 — Primary motor/sensory/visual cortex, S2 — somatosensory cortex.

observation of successful neurostimulation at low intensities and frequencies, in opposition to the established neuro-suppressive effects of heating [71, 72], or with no change in temperature measured [73, 49], strongly suggest a mechanical effect. It should be noted that ultrasonic heating can certainly affect neural tissue function, either through cell necrosis similar to that generated in HIFU treatments [74], or transient effects on neural conduction properties due to a smaller change in temperature [72, 71]. However, the effects of temperature appear to be limited primarily to suppression of neuronal activity, and are not always reversible.

While UNMS is generally taken to be a mechanical effect, the exact mechanism remains unclear, due to the large range of ultrasound parameters shown to be effective and variety of responses observed (see Tables 1.2 & 1.3). Experimentally, several studies have attempted to analyse quantitatively the effectiveness of ultra-

sound applied with different pulse characteristics and at different intensities. Early work by King et al. [57] concluded that intensity and stimulation duration above certain minimum thresholds are required for successful neurostimulation, and that success rates increase above those thresholds. These observations are supported across the literature, although an intensity ceiling for effective neurostimulation related to heating has been observed [71]. Kim et al. [59] extended this work to a wide range of ultrasound pulse parameters, and determined that there were certain "sweet spots" for neurostimulation in terms of stimulation duration, duty cycle and tone burst duration.

In terms of ultrasonic frequency, ultrasound below 1 MHz, usually around 500 kHz, has been used in the majority of UNMS research and is generally taken to be more effective. Several studies have confirmed that neurostimulation at higher frequencies ($>$ 1 MHz) is possible [66] with the extreme example of direct stimulation of retinal neurons with 43 MHz ultrasound by Menz et al. [75], although this seems to require higher ultrasonic intensities. However, it is not clear whether increased action at lower frequencies corresponds to a stronger neurostimulatory effect, an indirect increase in energy transmission (due to reduced absorption at low frequencies), reduction of conflicting effects (such as heating at higher intensities), or larger focal spot size exciting a greater number of neurons. Recently, Ye et al. [64], attempted to examine frequency dependence while controlling for some of these factors, and did conclude that lower frequencies more reliably elicited responses independent of focal spot size, attenuation or targeting.

Based on these observations, a mechanism based on acoustic cavitation has been suggested by multiple researchers. Notably, Plaskin et al. described a "bilayer sonophore" model for UNMS, wherein low frequency ultrasound causes stable cavitation of bubbles between the phospholipid membrane of neurons. However, this remains a theory. In two papers, Mueller & Tyler [76, 77] outline the wide range of potential interactions between mechanical forces and neural tissue which could be active just at the level of the individual neuron. These include mechanically sensitive ion channels, interations with the internal cytoskeleton, activation of membrane

proteins, and cell membrane piezoelectricity. Ultimately, the context of the nervous system and the complex system of interactions it entails makes determining or describing a single unified effect difficult and will presumably cause the systemic outcomes of UNMS to vary based on the target site, as well as the exact cellular mechanism at the target. For example, suppression of activity could be caused by either a suppressive effect of ultrasound, or an excitatory effect of ultrasound on a group of inhibitory neurons. Ultimately, while determining the exact underlying mechanism remains an important milestone, the presence of a stimulatory effect is sufficient motivation for the present interest.

1.1.4 Challenges for Ultrasonic Neurostimulation

UNMS benefits from a unique combination of advantages as a non-invasive neurostimulation method for research and clinical purposes that recommend it for further research and development. However, looking forward, there remain numerous practical difficulties to be overcome to allow tight transcranial focusing of ultrasound through the skull onto small deep brain structures, and in achieving the reliable stimulation of more invasive methods. Independent of the problem of transcranial focusing, understanding of the mechanisms and potential utility of ultrasonic neurostimulation is still developing. In order to fully exploit this nascent, transformative technology, the relationship between applied ultrasound characteristics, the targeted nervous tissue, and the stimulatory, suppressive or modulatory effects needs to be quantified.

The status of UNMS as a developing technology will remain a consideration when extending UNMS to deep brain targets. Given the influence of the skull, and the scale of neural structures of interest relative to the acoustic wavelengths of frequencies suitable for transcranial transmission, highly effective focusing through the skull will be necessary. Effective spatial targeting will ensure that the correct structure is stimulated, while a tight focus will avoid activation of secondary sites. Furthermore, tight control over the amplitude will be of crucial importance to allow safe and effective neurostimulation, and to allow investigations into the mechanisms and utility of UNMS in order to develop it as a clinical and research technique.

The low-intensities associated with effective and safe neural stimulation preclude the use of MRI thermometry or adaptive focusing methods described in Section 1.3.2. Therefore non-invasive focusing using simulations and image-derived acoustic property maps is likely to be required to achieve the necessary accuracy. The primary goal of this thesis is to establish the simulation parameters necessary to achieve the desired focusing quality, in terms of the spatial and temporal discretisation of the simulated domain, and the maps of acoustic properties used.

1.2 Structure and Acoustic Properties of the Head

1.2.1 Relevance for Transcranial Focusing

The historical usefulness of ultrasound as an imaging and therapeutic technique has derived from two essential assumptions about the tissues present in the ultrasonic field. The first is the consideration of biological tissues as fluid, which do not support shear waves. Although soft-tissues such as muscle and fat do in fact support limited shear mode propagation, the shear speed and absorption of soft-tissues are such that their impact on imaging & therapy is negligible [9]. The second is the low variation in sound speed between different soft-tissues, typically on the order of less than 200 $\mathrm{m/s}$. For imaging, this simplifies the relationship between the time-of-flight of reflections and the distance to interfaces. For therapy, it simplifies the focusing and targeting of applied ultrasound. However, in the case of propagation through the skull to target structures in the deep brain, these assumptions break down [78].

As might be expected, the properties of the skull lead to significant aberrations in propagating wavefronts when attempting to transmit ultrasound through tissues containing bone such as the head or thorax. This reduces the effectiveness of simple geometric focusing methods. Compounding this, bone is also a relatively potent absorber of ultrasound energy. As well as further degrading focusing through uneven attenuation of the wavefront, this can lead to reductions in the effective transmission of ultrasonic energy to transcranial targets, and can potentially cause hazardous heating of the skull and superficial tissues.

In the following section, the anatomical structure and acoustic properties of the skull are discussed with reference to the specific difficulties encountered when focusing ultrasound through the skull. Methods for deriving maps of the acoustic properties of the skull for use in numerical simulations for transcranial focusing are then reviewed.

1.2.2 Gross Anatomy and Skull Structure

When applying ultrasound from the surface of the head to the brain, the relevant anatomical areas are the brain and its coverings, the cranial cavity they sit inside, the skull bone, and the scalp. The scalp covers the top of the head, and is the most superficial layer of the head. It is a layered structure, comprised of multiple soft-tissues beneath the skin. Across most of its surface it is made up of connective tissue, aponeurotic fibrous tissue, and periosteum, which lies on the surface of the cranium. Some portions of the scalp do contain or cover muscle, most notably the temporalis muscle, which covers the temple and most of the side of the head, and the occipital and frontal portions of the occipitofrontalis muscle, which partly cover the frontal and occipital bones, and are connected by the aponeurotic tissue of the scalp [3].

The skull is made up of multiple distinct parts, the most relevant of which is the calvaria - the dome shaped portion that covers the brain. During foetal development the calvaria forms as multiple separate flat bones, with spaces between known as fontanelles [3]. They are the frontal bone, the occipital bone, the paired parietal bones and the paired temporal bones, shown in Fig. 1.3 (a). Shortly after birth the fontanelles close, and the various flat bones fuse, held tightly together by distinctive dense fibrous joints known as sutures, which gradually ossify throughout life [3]. The general structure of calvarial bone is of three layers. The internal and external “tables” are composed of dense cortical bone, which sandwich the diploe - a layer of cancellous trabecular bone and bone marrow. The thickness and structure of the skull, including the relative thicknesses of the different layers, varies across the surface [79]. A recent statistical analysis put the average skull thickness across the skull at 6.32 mm, with a standard deviation of 1.45 mm [80]. There are two

(a)

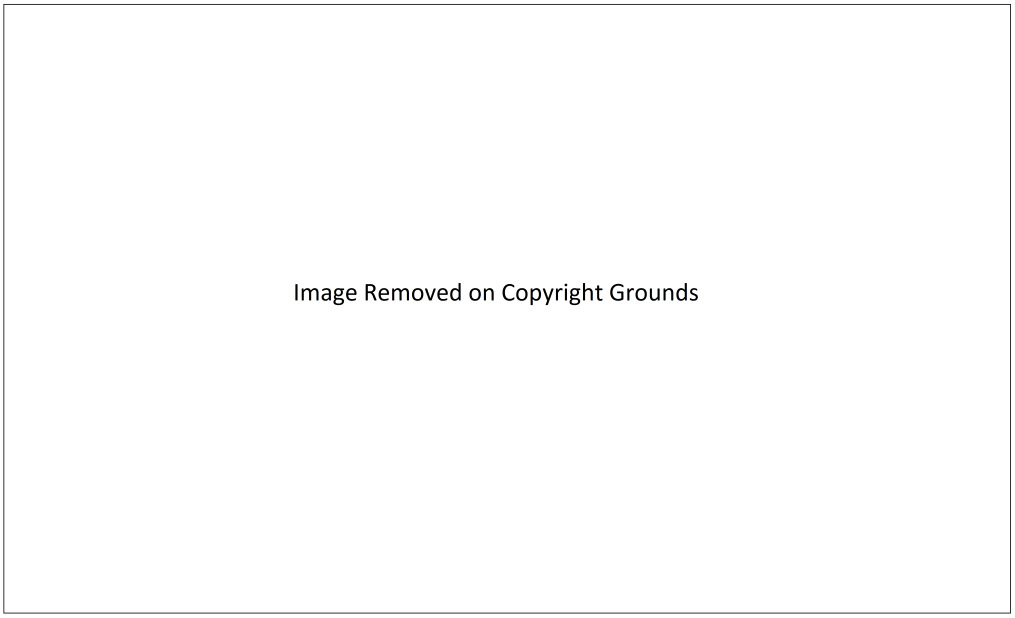

(b)

Figure 1.3: Anatomy of the skull. (a) Medial aspect of a sagittal section through a skull with the soft-tissues removed, with constituent bones labelled (reproduced from Gray's Anatomy [2]), (b) Structure of the skull & meninges (reproduced from Gray's Anatomy for Students [3]).

notable exceptions. The first is the thin, squamous portion of the temporal bone, which is only around 3 mm thick, and is mainly cortical bone [81, 82], leading to its historical use as an "acoustic window" for ultrasound imaging of the brain. The second is the paired frontal sinuses, air filled cavities in the frontal bone just above the eyes [83].

The brain sits within the cranial cavity covered in a series of membranes called the meninges - the dura, pia and arachnoid mater - which support the brain. The brain is also supported by cerebrospinal fluid (CSF), a fluid produced within the brain that sits between the arachnoid and pia mater. The brain itself is composed of three distinct units: the brainstem, the cerebellum and the cerebrum. The cranial cavity is dominated by the cerebrum, which is made up of the two cerebral hemispheres connected by a large white matter tract (the corpus callosum) and with a ventricular system that contains and produces CSF [21]. The general structure of the cerebral hemisphere is a grey matter surface - the cerebral cortex - which is multiply folded into different lobes which are, roughly, functional subdivisions. These communicate with a series of deeper grey matter nuclei, including the thalamus and basal ganglia by a layer of white matter, shown in Fig. 1.1 (b). Generally, the cerebral cortex acts as the centre of most neural processing, while the thalamus & basal ganglia act as relay stations between different areas of the cortex and cerebellum (which is involved in motor control and memory), and ascending and descending impulses which travel through the brainstem [21].

1.2.3 Acoustic Properties of Head Tissues

The full range of historical measured values of the acoustical properties of tissues in the head is too extensive to be enumerated here. Documentation of the literature can be found in Duck [84], and is available online as part of the Foundation for Research on Information Technologies in Society's (IT'IS) tissue properties database [10]. Selected ranges for the relevant tissues of the human head drawn from this database are shown in Table 1.4. Likewise, the range of techniques for measuring acoustical properties and the corresponding technical considerations is extensive, and comprehensively reviewed in Zequiri et al. [85] and Bamber [86]. A key technique is the transmission-substitution method, which is outlined in detail in Section 5.1.3, and allows for the measurement of sound speed and acoustic absorption. At this point an important distinction between acoustic absorption and acoustic attenuation should be described. Acoustic absorption refers specifically to losses occurring as a result of viscosity and other relaxation processes, and involves the conversion of

ultrasonic energy to heat. Attenuation refers to the sum of all amplitude losses in a propagating beam, and includes scattering and reflective losses. Density measurement is much simpler and can be carried out by displacement experiments, direct measurement with a picnometer, or derived from x-ray computed tomography (CT) images as discussed further in Section 1.2.4.

There are a few key experimental studies examining the properties of the skull that should be discussed. In an early experimental study, Fry & Barger [13] used the transmission-substitution method to characterise the compressional acoustic properties of skull bone extensively. Across a range of frequencies from 0.25-5 MHz, they measured the transmission loss, compressional sound speed and phase dispersion. Measurements were performed on whole bone, and also on the inner, trabecular and outer layers of the bone, sectioned separately. In addition to the direct measurements, a key observation was the majority of attenuation of ultrasound occurs in the trabecular layer. This indicates that scattering, rather than absorption, is the primary mechanism by which the skull attenuates ultrasound. This observation was confirmed by Pinton et al. [16], who carried out transcranial propagation experiments in tandem with thermographic measurements of heating due to ultrasound absorption. Full-wave acoustic simulations of propagation through high-resolution (40 μm) maps of the bone were then carried out. The absorption coefficient assigned to the bone of the cortical and trabecular layers was then adjusted until simulation output matched the heat and transmission measurements. The resulting value for longitudinal absorption of ultrasound in bone was 2.7 dB cm^{-1}, compared to an attenuation value of 13.3 dB cm^{-1} for the skull as a whole.

In a series of papers, White et al. [8, 9] examined the elastic properties of bone using transmission-substitution experiments at multiple angles of incidence. Combined with optimisation of a simplified simulation model treating the skull as a single layer with parallel surfaces, they calculated a shear speed in whole skull bone of 1500 m s^{-1}, and attenuation from 8.2-18.5 dB cm^{-1} across the frequency range tested. They also calculated values for the longitudinal sound speed and attenuation of the whole skull bone.

Tissue	Density [kg m^{-3}]	Sound Speed [m s^{-1}]	1 MHz Attenuation [dB cm^{-1}]
Skin	1100-1125	1537-1720	1.84
Muscle	1041-1178	1545-1631	0.62
Cortical Bone	1800-2100	2660-4180	3.5-5.2
Cancellous Bone	1080-1350	1854-2450	10.3
Brain	1039-1050	1506-1565	0.59
CSF	1007	1502-1507	$8.7x10^{-3}$

Table 1.4: Ranges of measured compressional sound speed and density values taken from the IT'IS database [10]. Attenuation measurements for skull bone are taken from Fry & Barger [13].

1.2.4 Medical Images and Acoustic Property Maps

Segmented and Simple Models

When simulating the propagation of ultrasound in the head, maps of the acoustic properties are required as inputs, with accurate maps resulting in more accurate simulations. When the simulations are to be used in treatment planning for patients, direct measurements of the acoustic properties are not feasible. Instead maps can be derived from medical image data using a variety of processing methods.

The simplest model of the skull that can be derived from 3D medical images of the head is homogeneous layer where the head is divided into bone and soft-tissue components. Acoustic parameters can then be applied based on empirical values for whole skull sound speed, density and attenuation, or based on properties derived from the CT image itself [87]. The simplicity of these models means that they can be more easily produced using a range of image processing methods and tools, and are often suitable to be used with more basic techniques for estimating phase shifts or simulating ultrasound propagation [88]. Another advantage of homogeneous skull maps is that they can be derived from MRI images. In their non-invasive focusing method discussed below, Sun & Hynynen [89] presented a single layer model based on the mapping of the inner and outer surfaces of the skull from MRI images. Many patients undergoing transcranial therapy will have an MRI scan as a matter of course, and MRI has a significant advantage over CT imaging in that it does not involve the application of ionising radiation, which researchers and

clinicians typically seek to avoid.

The analytical focusing method presented by Clement & Hynynen [87] experimented with treating the skull as either a single layer, or three layers based on the cortical-trabecular-cortical structure of the bone. Sound speed was assigned alternately to the whole skull layer and the three internal layers *en-bloc* based on the values measured by Fry & Barger [13]. They also experimented with calculating effective sound speed at different points over the skull surface based on CT derived density, and phase shift measurements used in a polynomial fit for both the single and three layer models. Modelling the skull as a three-layer structure and accounting for variations in sound speed across the surface of the skull both resulted in improved focusing.

Fully heterogeneous models

While the skull models described above can be used as simulation inputs, more accurate models of the acoustic properties of the skull aim to take account of the internal heterogeneities of the skull, and the corresponding spatial variations in sound speed, density and absorption. To that end, 3D medical imaging technologies can be used in voxel based acoustic property mapping. X-ray CT imaging has two clear advantages for this application. Firstly, it is known to give excellent discrimination and contrast when imaging bone in soft-tissue. Secondly, due to the relationship between photoelectric attenuation coefficient and density, CT Hounsfield Units (HU) can be used to infer mass density on a voxel by voxel basis accurately based on the CT calibration [6].

Historically, a wide variety of methods have been proposed to derive longitudinal sound speed from HU, or from density values derived from HU, based on analytical calculations or empirical measurements. The use of "effective" sound speed values based on HU for the creation of homogeneous or three-layer skull models has already been discussed. In an extensive, multi-frequency study, Pichardo et al. [7] measured sound speed and attenuation of whole skull bone, and then used image derived apparent density and a simulation based optimisation method to calculate relationships between density and other acoustic properties. Jones & Hynynen [90]

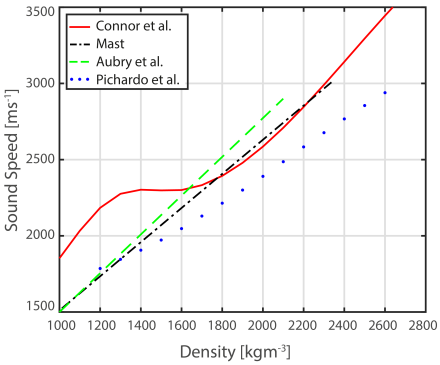

Figure 1.4: Relationship between the density of the skull bone, derived from CT, and the compressional sound speed. Values taken from Connor et al. [4], Mast [5], Aubry et al. [6], and Pichardo et al. [7]. Reproduced from Robertson et al. [116].

went on to use these values in simulation based aberration correction for focusing through the skull, and demonstrated improved performance relative to homogeneous models. In a similar study, Connor et al. [4] used a genetic algorithm in combination with high resolution CT images and experimental transmission measurements to derive a non-linear relationship between density and sound speed. More recently, Mast [5] reviewed acoustical property values of a range of soft-tissues to calculate linear relationships between density and different properties. These relationships between density and sound-speed are shown in Fig. 1.4, demonstrating the potential for variation in sound speed values calculated from density.

Arguably the most comprehensive method for the conversion of CT images was developed by Aubry et al. [6] for their initial work on transcranial focusing via simulated time-reversal. Based on the acoustic property measurements by Fry & Barger [13], their algorithm converts CT HU directly into fully heterogeneous maps of sound speed, density and acoustic absorption in the skull. Firstly, CT images normalised between water (0 HU) and cortical bone (1000 HU) are used to derive a map of porosity

$$\phi = 1 - (\mathrm{HU}/1000) \quad . \tag{1.1}$$

The porosity value, ϕ is then used to interpolate linearly between reference values for density ($\rho_{\min}, \rho_{\max}$) and sound speed ($c_{\min}, c_{\max}$), derived from the values for

water and cortical bone

$$\rho = \phi\rho_{\min} + (1-\phi)\rho_{\max} \qquad c = c_{\min} + (c_{\max} - c_{\min})(1-\phi) \ . \tag{1.2}$$

Porosity is also used to calculate the acoustic absorption based on maximum and minimum values. However, the assigned absorption value is inflated to counteract the loss of fine detail due to the limited resolution of clinical CT images. The loss of the internal detail of the skull trabecular layer, reduces the scattering and attenuation that would occur due to variation in acoustic impedance. Therefore the assigned absorption value is, strictly speaking, an attenuation value, although it is incorporated into numerical schemes in the same way as absorption [6, 13]. Accordingly, unlike sound speed and density, high porosity equates to a high absorption value, and a power law is used to interpolate between the values

$$\alpha = \alpha_{\min} + (\alpha_{\max} - \alpha_{\min})\phi^{0.5} \ . \tag{1.3}$$

This technique has been used successfully in multiple studies, including a simulated full wave time-reversal study by Marquet et al. [91]. Variations include alterations to the reference property values [91], and the use of a homogeneous absorption map based on a skull layer segmentation [15].

1.3 Methods for Focusing through the Skull

1.3.1 Background and Focusing Priorities

Due to its structure and acoustic properties, the skull poses a serious obstacle for those aiming to focus ultrasound in the brain. Fortunately, this issue has previously been addressed for the purpose of applying HIFU for tissue ablation or thrombolysis as far back as the 1950s. Initial approaches to the problem of focusing across the skull simply circumvented it. For example, Fry et al. [37] used craniotomy to avoid the aberrating bone layer entirely, simplifying focusing [92, 93, 94]. Although this method is still occasionally used [95], it renders the non-invasive nature of ultrasound somewhat redundant. Yin & Hynynen [96] demonstrated another method

for avoiding the effect of the skull, taking advantage of the reduced aberration and absorption of low frequency ultrasound to improve focusing. However, the longer wavelength prevents tight focusing, and the reduced absorption of ultrasound also had the effect of reducing heating at the target.

Interest in actually correcting for the effect of the skull for the purposes of focusing ultrasound rather than simply avoiding it dates back to the 90s [97]. The advent of large, multi-element phased arrays introduced the opportunity for electronic aberration correction through independent adjustment of the driving signals of individual elements. For high intensity applications they have the added advantage of spreading the applied ultrasound over a larger area, and reducing heating of the skull - a primary limitation of transcranial HIFU. Various focusing techniques can then be derived to determine the relative phases, and in some cases amplitudes, necessary to drive the transducer elements and correct for the aberration of the skull.

The acoustic frequency used does remain an important concern, due to the tradeoff between transmission and absorption in the skull, absorption at the target, the theoretical minimum focus spot size, and the acoustic cavitation threshold (which increases at lower frequencies). Deffieux et al. [15] identified frequencies in the range 500 kHz to 800 kHz as the optimal frequencies that allow focusing and effective transcranial transmission for low-intensity applications, although transcranial transmission with a wider range of frequencies is feasible, subject to the tradeoffs mentioned.

The requirements for any focusing system will depend on the specific application. For example, in HIFU, production of a tight focus in the desired position and the absence of standing waves is of paramount importance to avoid damaging healthy tissue or inducing side effects such as haemorrhage. However, fine control over the intensity at the target is less important, as the amplitude of the applied ultrasound can simply be increased to achieve the desired ablative effects. There is also the concern mentioned above, of reducing heating of the skull during treatment, which can be addressed with certain focusing methods. For low-intensity methods, heating of the skull is not a concern, while spatial targeting remains a priority to

ensure that the desired area is being sonicated. Fine control over the ultrasonic intensity at the target is likely to be of primary importance in the development and implementation of these techniques, due to the proximity between therapeutic and potentially damaging ultrasound intensities [63, 98], and to aid studies into the underlying mechanisms of novel therapeutic applications such as neurostimulation and opening the blood brain barrier with ultrasound. The importance of focal spot size will vary depending on the desired target, but ultrasound for neurostimulation should aim to focus on the scale of the neurological structures in Table 1.1.

1.3.2 Adaptive Focusing

An important set of focusing techniques that should be mentioned before discussing more advanced methods involve real-time correction of applied ultrasound based on the imaging of ultrasonic fields. Specifically, the development of MRI compatible, multi-element transcranial arrays means that real-time MRI thermometry can be used to image heating in the brain during treatment for feedback and focusing. Typically, an initial lower-intensity sonication is applied, and the heat distribution imaged. Any necessary corrections can then be applied by correcting the position of the transducer or adjusting the phase of the transducer elements until a tight focus is achieved at the target location, and then the ablative sonication is carried out [99, 100].

MRI based focusing has been used to carry out transcranial HIFU surgery on the thalamus as a treatment for essential tremor [24], and ablation of glioblastoma tumours [101]. More recently, Hertzberg et al. [102] demonstrated the effectiveness of alternate MRI imaging sequences that are able to measure tissue displacement as a result of acoustic radiation force for the purposes of adaptive focusing. During sonications from ExAblate HIFU transducers at 220 kHz, 710 kHz and 1 MHz, tissue displacements of up to 10 μm were imaged and used to correct the phase of the transducer elements. In order to improve the focusing of the initial sonication, these methods are also often combined with some of the non-invasive focusing methods discussed below.

The use of real time monitoring of ultrasonic heating in the brain remains a

valuable tool in transcranial HIFU, and adaptive focusing based on these measurements will likely continue to be used to correct targeting errors and tighten focusing. However, while they are effective for high-intensity treatments, the requirement for heating at the target or tissue displacement prevents their use for low-intensity applications where treatment cannot be monitored and focusing adjustments must be calculated ahead of time.

1.3.3 Phase Correction and Time Reversal

In addition to adaptive focusing based on feedback from MRI images, or as an alternative when applying low-intensity ultrasound, a variety of methods for deriving the transducer drive signals have been developed for measuring or estimating the aberration of the skull.

Several methods employ the principle of acoustic time-reversal, first described by Fink [103], which takes advantage of the fact that acoustic wave propagation is time-invariant in a lossless medium, and so propagating acoustic waves behave the same way in forward and reverse time through scattering media. Generally, time-reversal entails placing an ultrasonic source at a desired target location and allowing the ultrasound to propagate outwards through the aberrating media towards a multi-element transducer array acting in receive mode [104]. For broadband acoustic waves, the full time-varying signals are reversed in time and used as drive signals, while for single frequency CW waves the inverted signals are just used to derive the necessary adjustments in the relative phase and driving amplitude of each element [88]. A schematic illustrating the principle of time-reversal is shown in Fig. 1.5.

It should be noted that, while time-reversal can be used in absorbing media, absorption is not time-invariant. This means that amplitude changes due to scattering (which will reverse in time) should ideally be separated from those due to absorption (which won't). Historically, this has been complicated by a desire to avoid excessive ultrasound absorption in the skull and attendant heating. White et al. [105] experimented with different methods of amplitude correction with hydrophone based focusing. They observed that, while driving elements with low transmission through the skull with higher amplitude (correctly reversing absorp-

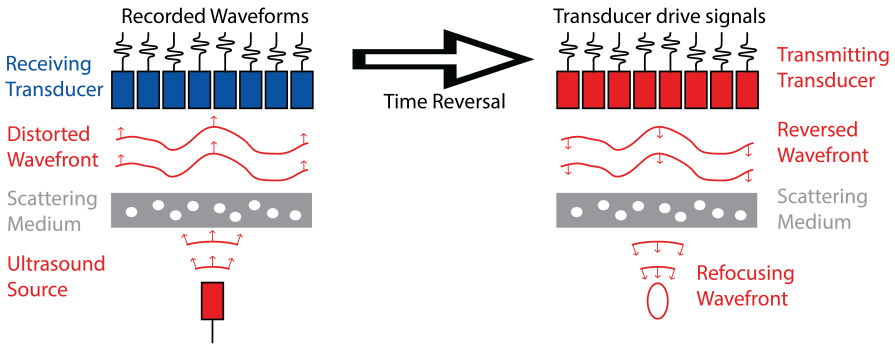

Figure 1.5: Schematic illustrating the time-reversal process. An ultrasonic source is placed at the target, and ultrasound propagated through the scattering medium to a receiving transducer. The recorded pressure signals are time reversed, and used as drive signals for the reversal propagation. The time-reversed wavefront retraces its path through the scattering medium and converges on the target point in space and time.

tion losses) leads to slightly better focusing, doing the opposite can still achieve effective focusing, with higher peak intensity at the target. Thomas & Fink [104] confirmed the merits of time-reversal focusing through both *ex-vivo* skull bone and an aberrating lossless phase screen using a curved transducer array operating at 3.3 MHz and 1.5 MHz. They employed amplitude correction to try and account for absorption, and obtained -6 dB beamwidths of 1-5 mm.

An important variation involves placing a measurement hydrophone at the target location rather than a source. Transducer elements are driven in turn, and the pressure recorded at the target is inverted to derive element drive signals. Hynynen & Jolesz [94] first used this technique to calculate phase corrections for transcranial propagation experiments, attaining a targeting error of less than 3 mm for frequencies below 1 MHz. Pernot et al. [106] used this technique in an *in-vivo* study where an implanted hydrophone was used to focus ultrasound onto the brains of sheep to produce thermal lesions. Hydrophone based focusing of this kind has been widely used in experimental tests of transcranial transducers, and has become the "gold-standard" against which novel focusing methods are tested [88].

1.3.4 Non-Invasive and Simulation Based Methods

While time-reversal and hydrophone based correction methods represent powerful techniques that allow focusing through the skull, their clinical and *in-vivo* experimental usefulness is severely limited by the need to introduce a hydrophone or acoustic source to the target. The desire to apply transcranial ultrasound clinically has therefore driven the development of a variety of methods for deriving the necessary phase and amplitude adjustments non-invasively, taking advantage of the various methods for deriving the acoustic properties of the skull discussed in Section 1.2.4. These range in complexity from analytical methods, through relatively simple numerical methods such as the angular spectrum method, through to full-wave time-domain simulation of ultrasound propagation. Each method is accompanied by individual limitations in terms of the accuracy with which ultrasound propagation is modelled, the complexity of the acoustic property maps that can be used, and the computational burden and simulation times.

The most notable analytical method was presented by Clement & Hynynen [87], who used CT images to derive the thickness of the skull across the skull surface. The average sound speed of the skull beneath each transducer element, calculated using their analytical method described in Section 1.2.4, was then used to derive the corresponding shift in phase based on a time-of-flight calculation. This method demonstrated strong prediction of the trend in shift in phase as a function of skull thickness. However, due to the simple nature of the scheme it could not capture all the variations in phase, even when attempting to account for average sound speed variation across the skull . A version of this method is used in some current clinical targeting software, with phase shifts for each element calculated based on the effective sound speed and resulting phase shift for a ray propagating from a transducer element to the target point [88].

While analytical methods can be used for rapid calculation of phase shifts, they are necessarily limited in terms of the range of acoustic propagation phenomena that they can account for. Numerical methods for the simulation of ultrasound propagation through the skull can be adapted to account for a range of different aspects

of transcranial propagation. Simulations imitating the arrangement of the invasive techniques described above can then be used to derive phase and amplitude shifts for transducer elements. In a series of papers, Sun & Hynynen [89, 107] presented non-invasive transcranial focusing using a method for the rapid computation of ultrasonic fields based on the Rayleigh integral. Ultrasound sources were treated as a collection of point sources, with the field computed based on the sum of the fields radiating from those points. The model was extended to layered heterogeneous media by computing the field across the interface surface, and then treating this as a source term for a new propagation. This model was used to simulate the propagation of ultrasound from transducer elements to the target location to calculate the necessary phase shifts. In a later study, Clement & Hynynen [78] used a more advanced angular spectrum method to simulate transcranial propagation. While this method continued to use homogeneous models of the skull, it could also take into account the angle of incidence of applied ultrasound to calculate refraction more accurately.

Simulation schemes such as these give increased fidelity in wave propagation at the cost of some computation time, but are still unable to account for some acoustic phenomena. Full-wave numerical simulation of ultrasound, which involves solving a system of equations that describe the propagation of acoustic waves, can give superior accuracy at the cost of further computation time [88]. Time-reversal focusing using full-wave numerical simulations of ultrasound was first adapted for transcranial focusing by Aubry et al. [6]. Finite difference time-domain (FDTD) simulations were used to simulate the propagation of the desired ultrasound signal from a virtual target point, through a model of the skull derived from 3D CT images, to a simulated transducer surface. The pressure signals recorded at the position of the transducer elements were then used to derive the corresponding drive signals. Chauvet et al. [108] confirmed the potential for model-driven TR-based focusing inside the human head to millimeter precision, verified by MRI thermometry. Marquet et al. [91] showed that model-driven TR is capable of restoring 90% of the peak pressure that can be obtained with gold-standard hydrophone based methods

when focusing through *ex-vivo* skull bone.

The primary limitation of full-wave simulations is the time that it takes to carry them out [88]. In order to ensure numerical accuracy, the simulated domain needs to be discretised at a certain spatial resolution and, for time-domain methods, the pressure must be updated iteratively at each time step. When simulating more complex acoustic phenomena such as absorption, non-linearity or shear mode propagation, the computational burden can increase [109]. One method to reduce total computation time, is to combine full-wave simulations through the skull layer with simpler methods. For example, Pinton et al. [110] combined FDTD simulations of propagation through a heterogeneous skull map with a ray propagation scheme that allowed the phase to be projected through soft-tissue, which was assumed to be homogeneous. Pulkkinen et al. [111] used a three part simulation scheme to model skull base heating during ultrasound therapy. They used Rayleigh integral methods to propagate ultrasound through superficial soft-tissues, an elastic full wave finite-element code to propagate ultrasound through the skull, and then an angular spectrum code to propagate ultrasound through intracranial soft-tissue. These hybrid methods greatly reduced computation time by only using full-wave numerical simulations when necessary. Reduced computational burden can also be achieved by using full-wave methods that can achieve acceptable accuracy with less restrictive grid sampling requirements [112].

1.4 Goals of the Thesis

Stimulation of neural tissue with ultrasound represents a nascent, potentially transformative technology with applications in clinical therapy and neuroscientific research. The ability to stimulate and modulate arbitrary deep brain structures with a spatial resolution on the scale of acoustic wavelengths would allow unprecedented mapping of brain anatomy and function in healthy subjects, and could be used to treat conditions such as Parkinson's disease and essential tremor.

In order to stimulate human brain tissue effectively, ultrasound must be focused on deep brain structures. However focusing is achieved, the focus will need to be spatially targeted on the correct area, and the focal spot will need to be sufficiently tight to ensure stimulatory specificity. Furthermore, the limited understanding of the mechanism(s) of UNMS and resulting ambiguity in the most effective ultrasound intensities and pulse shapes, mean that fine control or prediction of the time-varying ultrasound at the target location will likely be necessary. Precise amplitude control will ensure that stimulation occurs while avoiding any tissue damage, will aid research into the mechanisms underlying UNMS, and may even allow control over the specific type of stimulation, suppression or modulation.

Based on the dimensions of the structures in Table 1.1, this suggests a need for a focal spot size of approximately 3 mm in diameter, and a maximum spatial targeting error of approximately 1.5 mm. While the exact level of accuracy in ultrasound amplitude cannot be known due to the limited understanding of the mechanistic underpinning, the range of values shown in Table 1.3 suggest that an error of less than 10% in intensity (or approximately 5% in pressure) will ensure that stimulation occurs while preventing tissue damage, and could potentially be sufficient to control the neural modulation at the target.

Non-invasive methods for the transcranial focusing of ultrasound have previously been used to focus ultrasound for HIFU. However, despite the risks associated with ablative therapy, the ability to monitor treatment outcomes in real time via MR imaging and the relative unimportance of control over the ultrasound amplitude have limited the importance of highly accurate transcranial focusing. Transcranial

ultrasound focusing for neurostimulation will require refinement of the simulated time-reversal process. To that end, a taxonomy of potential sources of error and inaccuracy when using full-wave numerical simulations of ultrasound propagation for time-reversal focusing can be defined. They comprise:

1. **Physical**: The underlying equations used to describe the propagation of ultrasound in the head, and how well they represent the physical reality.

2. **Numerical**: The accuracy of the simulation scheme used to solve the physical equations, and the spatial and temporal sampling necessary to achieve required accuracy.

3. **Inputs**: The acoustic property maps used as inputs to the numerical scheme, including the values of the acoustic properties and the geometry of the bone.

4. **Outputs**: How accurate simulations of transcranial propagation are used within a wider time-reversal simulation scheme. For example, how the simulated source term relates to the desired pressure at the focus and how amplitude changes due to absorption are accounted for.

This thesis examines the second and third items on this list, making use of a k-space pseudospectral time-domain (PSTD) scheme for the simulation of ultrasound, and examining both homogeneous and heterogeneous models of the skull acoustic properties. The physical equations used to describe ultrasound propagation are reviewed, but examining them in detail or attempting to improve them is outside the scope of the present work. Likewise the usage of numerical simulations for time-reversal focusing is limited to driving experimental design and informing analysis of results.

The overall goal of this thesis is to determine the parameters necessary to simulate the transmission of ultrasound through skull bone with sufficient accuracy. This is examined in terms of the spatial and temporal sampling required to achieve numerical accuracy for the numerical method used, the accuracy required in the acoustic properties of the skull bone, and the impact of limited resolution of property maps derived from clinical CT images. Practical simulations are also carried

out to demonstrate the potential for accurate mapping between simulated and experimental transcranial sonications. While the intended application of the work is focusing for the purposes of neurostimulation, results are relevant to other applications requiring the effective transcranial focusing of ultrasound, as well as general simulation of ultrasound through skull bone.

In Chapter 2 the k-space PSTD numerical scheme is described, and the impact of various individual sources of numerical error are tested independently to determine the spatial or temporal sampling criteria required to reduce their influence to an acceptable level. The combined impact of numerical errors on simulations used for the specific application of time-reversal focusing is then examined through fully simulated time-reversal in Chapter 3. Together, these two chapters establish the relative impact of different sources of numerical error and determine the spatial sampling criteria necessary to ensure numerical accuracy. These criteria are used throughout the remainder of the thesis.

With numerical accuracy assured, the impact of variation in the properties of an idealised skull bone layer on the transcranial transmission of ultrasound in terms of the change in the intracranial field is examined through an extensive numerical sensitivity analysis in Chapter 4. This is used to determine the accuracy in acoustic medium properties and skull layer geometry required to achieve the desired accuracy in focused intracranial fields. The results of the tests of numerical error and the sensitivity of intracranial fields are validated through a set of practical experiments in Chapter 5. Sonication experiments are carried out examining the propagation of ultrasound through multiple skull bone phantoms created using computer aided design, 3D printing and resin casting techniques, including a phantom based on the idealised skull geometry used in the sensitivity analysis in Chapter 4. These experiments are compared with registered simulations to determine the accuracy with which the transcranial transmission of ultrasound can be simulated given an accurate map of acoustic properties. Finally, Chapter 6 examines the impact of the limited resolution of clinical CT images on simulations through acoustic property maps derived from those images. Acoustic property maps are derived from

micro-CT data of multiple skull bone samples, and the impact of progressive homogenisation of these maps established. Simulation through maps derived from micro-CT data is also compared with simulation through property maps derived from co-registered clinical CT data.

1.5 List of Publications

The work contained in this thesis has previously been presented in the following publications.

1. J. L. B. Robertson, B. T. Cox, and B. E. Treeby. Quantifying numerical errors in the simulation of transcranial ultrasound using pseudospectral methods. In *2014 IEEE Int. Ultrason. Symp.*, pages 2000–2003. IEEE, sep 2014

2. J. L. B. Robertson, E. Hill, A. A. Plumb, S. Choong, S. J. West, and D. I. Nikitichev. 3D printed ultrasound phantoms for clinical training. In *IS&T's 32nd Int. Conf. Print. Fabr.*, volume NIP32, pages 6–8, sep 2016

3. J. L. B. Robertson, B. T. Cox, J. Jaros, and B. E. Treeby. Accurate simulation of transcranial ultrasound propagation for ultrasonic neuromodulation and stimulation. *J. Acoust. Soc. Am.*, 141(3):1726–1738, mar 2017

4. J. L. B. Robertson, E. Martin, B. T. Cox, and B. E. Treeby. Sensitivity of simulated transcranial ultrasound fields to acoustic medium property maps. *Phys. Med. Biol.*, 62(7):2559–2580, apr 2017

5. J. L. B. Robertson, J. E. Urban, J. D. Stitzel, B. T. Cox, and B. E. Treeby. Influence of trabecular microstructure on the simulation of ultrasound through the skull (Under Preparation). In *7th Eur. Symp. Ultrason. Charact. Bone*, Bad Staffelstein, 2017. IEEE

Chapter 2

Numerical Accuracy of k-space Simulations

When using simulations of ultrasound propagation in simulated time-reversal, or to predict ultrasonic fields in the treatment region, the effectiveness of the simulation scheme used is of fundamental importance. Throughout this thesis, numerical simulations of ultrasound propagation are carried out using the k-Wave toolbox, which employs a k-space corrected pseudospectral time domain (PSTD) method for the time domain simulation of ultrasound fields [118].

The effectiveness of a numerical method can be evaluated in a number of ways. Primarily the objective is that a numerical method should accurately calculate the acoustic pressure and particle velocity at a given position in space and time given a set of initial conditions and a simulated ultrasonic source. The accuracy inherent to a numerical scheme can be thought of using three key concepts: consistency, stability and convergence. A consistent scheme is one that mathematically reduces to the continuous differential equations in the limit as the spatial and temporal grid steps approach zero [119]. This means that the discrete equations are a true representation of the mathematical model that described the physical system, but is not sufficient to ensure accuracy. Numerical stability means that any errors in the simulation do not grow exponentially with successive time steps, although stable schemes that are not consistent, or that do not meet certain other requirements will not necessarily be accurate [119, 120]. If a numerical scheme is both consistent and stable, such

as the k-space PSTD method [121, 118], it is said to be convergent via the Lax equivalence theorem. This means mean that the solution calculated by the scheme will approach the correct solution as the size of the grid spacing approaches zero [119, 120]. However, simply being convergent does not ensure that a numerical scheme will be sufficiently accurate for a given application. It will still require a certain grid spacing to reach the required accuracy, dependent on the sources of numerical error inherent in the discrete approximation. A numerical method should ideally also be computationally efficient, meaning that the numerical scheme does not require an excessive investment of time, or present an excessive computational burden in terms of memory or processing requirements.

In the following section, the equations that describe linear wave propagation in a fluid medium are described, along with the discrete equations used in the k-Wave toolbox. Individual sources of numerical error found in the k-space scheme are then outlined and examined in a series of numerical tests to study their impact on transcranial propagation.

2.1 Numerical Methods for Ultrasound Simulation

2.1.1 Governing Equations

The primary mode of acoustic wave propagation in biological tissue is of compressional, or longitudinal waves, which serves as the justification for generally modelling biological tissues as fluids. Compressional wave propagation through a fluid medium involves fluctuations in physical variables including pressure, density, and acoustic particle velocity and displacement. For ultrasound at low intensities, a description of propagation of ultrasound in heterogeneous fluid materials can be derived from linearised equations for the conservation of mass

$$\frac{\partial \rho}{\partial t} = -\rho_0 \nabla \cdot \boldsymbol{v} - \boldsymbol{v} \cdot \nabla \rho_0 \ , \tag{2.1}$$

the conservation of momentum

$$\frac{\partial \boldsymbol{v}}{\partial t} = -\frac{1}{\rho_0}\nabla p \ , \tag{2.2}$$

and a linearised pressure density relationship [122]

$$p = c_l^2(\rho + \mathbf{u}\cdot\nabla\rho_0) \ . \tag{2.3}$$

Here ρ_0 is the ambient density, c_l is longitudinal sound speed, p is the acoustic pressure, ρ is the acoustic density, and $\mathbf{u}$ and $\boldsymbol{v}$ are the acoustic particle displacement and velocity, respectively. The linearisation of Eqs.(2.1)-(2.3) is based on the assumption that the fluctuations in density caused by an acoustic wave are small compared to ambient density, and that the particle velocity is small compared to the sound speed of the medium [123].

When shear mode propagation in the skull is likely to affect the transcranial transmission of ultrasound, the alternate system of equations described in Appendix A.1.2 can be used. However, the requirement to model field tensors such as stress and strain leads to a significant increase in complexity compared with the fluid case, with a resulting increase in computational burden [124]. Furthermore, the importance of shear mode propagation depends on the angle of incidence of applied ultrasound [8], and is therefore highly case dependent. As a result, in the following chapters, elastic wave propagation is only modelled in situations where it is likely that it will affect transmission, or to test whether it will do so. Otherwise the fluid equations are used.

2.1.2 k-Wave and Pseudospectral Methods

The k-space PSTD scheme used within the k-Wave toolbox solves Eqs.(2.1)-(2.3) using the k-space corrected PSTD method described by Tabei et al. [121]. In this scheme, spatial gradients of field variables are calculated by a spectral collocation method using a Fourier basis, while the iterative time-stepping is performed by a finite difference calculation. k-Space correction refers to the introduction of a term that corrects for numerical dispersive error arising from the finite difference time

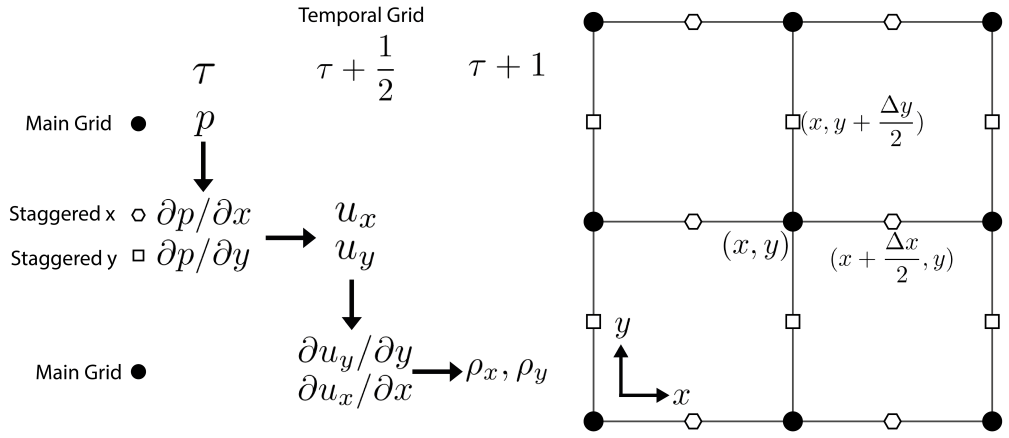

Figure 2.1: Schematic of a 2D spatial grid, showing the spatial and temporal stepping of the k-space PSTD numerical scheme and the different times and positions where acoustic variables are calculated, including staggered spatial and temporal grids.

step calculation. A full derivation of the k-space scheme as an extension of finite difference time domain (FDTD) and PSTD methods can be found in Appendix A, including details of how acoustic absorption is modelled, and the governing equations that describe elastic propagation in solid media.

In terms of practical implementation, the k-Wave toolbox uses the following set of first-order coupled equations based on Eqs.(2.1)-(2.3) to update pressure and particle velocity across the simulation grid in a time stepping manner. Here ξ is used to represent a single spatial dimension, with each step repeated for the dimensions x,y & z as necessary. The term Δt is the temporal discretisation, $\Delta\xi$ is the spatial discretisation (in the direction ξ), τ is the temporal index and S_{F_ξ} and S_{M_ξ} represent force and mass based source terms, respectively. The subscript ξ is used to indicate ambient acoustic properties defined at the staggered grid positions. A visual representation of the scheme and the different computational steps is shown in Fig. 2.1 for a 2D grid.

First, the spatial gradient of pressure is calculated for positions on the staggered grid:

$$\frac{\partial p^\tau}{\partial \xi} = F^{-1}\left[ik_\xi \kappa e^{ik_\xi \Delta\xi/2} F[p^\tau]\right] \quad . \tag{2.4}$$

Here, F and F^{-1} are the discrete Fourier transform and inverse Fourier transform,

respectively. k_ξ is the discrete, bounded set of wavenumbers in the ξ direction for which the discrete Fourier transform produces coefficients. This Fourier series corresponds to a continuous function that agrees with the discretely sampled pressure field at each grid point. This is the Band-Limited interpolant (BLI), explored further in Section 2.2.2. The term κ is the k-space correction factor $\mathrm{sinc}(c_{\mathrm{ref}} k \Delta t/2)$, where

$$k^2 = \sum_\xi k_\xi^2 \ . \tag{2.5}$$

This factor corrects the error generated by the finite difference approximation of the temporal gradient, although the correction is only exact for media with the sound speed c_{ref} [121]. When $c_{\mathrm{ref}} = 0$, the PSTD scheme is uncorrected. The term $e^{ik_\xi \Delta\xi/2}$ translates the results of the gradient calculation by $\Delta\xi/2$ to allow particle velocity to be calculated on a staggered grid. The use of a Fourier series allows the gradient of the interpolating function to be calculated anywhere inside the grid space, meaning that a staggered grid is not strictly necessary. However, a staggered grid helps suppress the appearance of oscillations due to the BLI, and is also thought to improve the representation of medium discontinuities [121].

Following calculation of the spatial gradient of pressure via spectral methods, the equation for conservation of momentum (Eq.(2.2)) is used to update each Cartesian component of the particle velocity defined on the staggered grid in time by using a finite difference approximation of the temporal gradient

$$v_\xi^{\tau+\frac{1}{2}} = v_\xi^{\tau-\frac{1}{2}} - \frac{\Delta t}{\rho_{0_\xi}} \frac{\partial p^\tau}{\partial x} + \Delta t S_{F_\xi}^\tau \ . \tag{2.6}$$

Note that this equation requires the density map to be separately defined on the staggered grid. A source term representing applied force can be introduced at this stage. The gradient of the newly evaluated particle velocity is then calculated spectrally. This calculation includes another exponential operator to shift the calculated

gradients back onto the main grid

$$\frac{\partial v_\xi^{\tau+\frac{1}{2}}}{\partial \xi} = F^{-1}\left[ik_\xi \kappa e^{-ik_\xi \Delta\xi/2} F[v_\xi^{\tau+\frac{1}{2}}]\right]. \tag{2.7}$$

The particle velocity gradient is then used in another finite difference update step based on the conservation of mass (Eq. (2.1)), calculating the new acoustic density values on the grid, at the next full time step:

$$\rho_\xi^{\tau+1} = \rho_\xi^{\tau} - \Delta t \rho_0 \frac{\partial}{\partial \xi} v_\xi^{\tau+\frac{1}{2}} + \Delta t S_{M_\xi}^{\tau+\frac{1}{2}} \quad . \tag{2.8}$$

Note that the density values calculated are separated into Cartesian components, based on their calculation from the Cartesian components of particle velocity. This is done to allow the use of a split field perfectly matched layer (PML) for the absorption of waves at the edge of the simulated domain [125], which is explored further in Section 2.2.3. A mass source term can be introduced into the numerical scheme at this point. The final step is the summing of the Cartesian components of the density, and the conversion of the updated density field into the new pressure field,

$$p^{\tau+1} = c_0^2\big(\sum_\xi \rho_\xi^{\tau+1} - L_d\big) \quad . \tag{2.9}$$

Here L_d is a loss term based on a fractional Laplacian term derived by Treeby et al. [126] and described further in Appendix A.1.1. For simulations in more than one dimension, gradient calculations and update steps are repeated in each dimension as necessary. The k-space PSTD scheme has been shown to be exact for homogeneous media independent of the temporal discretisation, subject to the conditions of the BLI [121].

The general stability condition for the k-space PSTD scheme is given by

$$|\sin(c_{\text{ref}} k \Delta t/2)| \leq \frac{c_{\text{ref}}}{c_0} \quad . \tag{2.10}$$

This demonstrates how the scheme can be made unconditionally stable for $c_{\text{ref}} = c_0$,

as the sin term will never be larger that 1. When $c_{\text{ref}} \leq c_0$, the stability condition can be reformulated to give an expression for the largest stable timestep,

$$\Delta t \leq \frac{2}{c_{\text{ref}} k_{\text{max}}} \sin^{-1}\left(\frac{c_{\text{ref}}}{c_{\text{max}}}\right) \ , \tag{2.11}$$

where k_{max} is the maximum wavenumber supported on the grid. A useful criterion when discussing simulation stability is the Courant-Friedreichs-Lewy (CFL) number [121]. The CFL in 1D is defined as

$$\text{CFL} = \frac{c_{\text{max}} \Delta t}{\Delta x} \ . \tag{2.12}$$

This definition is used in this thesis for time domain solutions of the acoustic wave equation for simulations in all dimensions.

2.1.3 Finite Difference Schemes

When appropriate, tests of numerical errors in the k-space scheme are compared with those arising from widely used FDTD numerical schemes for the solution of the same coupled system of equations. These schemes are briefly described here. A more detailed discussion of FDTD methods can be found in Appendix A.2. A second-order accurate in time, fourth-order accurate in space (2-4) scheme, described by Strikwerda et al. [127] has previously been used in multiple studies involving the simulation of ultrasound through skull bone, including the practical application of time-reversal focusing [108, 91]. A version of this numerical scheme is available as part of the k-Wave toolbox.

By combining the linearised pressure-density relationship with the conservation of mass (Eqs.(2.9) &(2.8) for the k-space PSTD scheme) the non-absorbing 2-4 FDTD scheme can be summarised as two coupled first-order equations. Here j is used to represent the spatial index of variable values on the grid. The first step updates the particle velocity on staggered spatial and temporal grids based on the

pressure distribution across the grid

$$v_{j+\frac{1}{2}}^{\tau+\frac{1}{2}} = \frac{\Delta t}{p_{0_\xi}} \frac{p_{j-1}^{\tau} - 27p_j^{\tau} + 27p_{j+1}^{\tau} - p_{j+2}^{\tau}}{24\Delta x} + v_{j+\frac{1}{2}}^{\tau-\frac{1}{2}} \quad . \tag{2.13}$$

The second step uses the particle velocity values calculated at the staggered grid positions to compute the updated pressure field across the grid

$$p_j^{\tau+1} = \Delta t \rho_0 c_0^2 \frac{v_{j-\frac{3}{2}}^{\tau+\frac{1}{2}} - 27v_{j-\frac{1}{2}}^{\tau+\frac{1}{2}} + 27v_{j+\frac{1}{2}}^{\tau+\frac{1}{2}} - v_{j+\frac{3}{2}}^{\tau+\frac{1}{2}}}{24\Delta x} + p_j^{\tau} \quad . \tag{2.14}$$

A second-order accurate in space, second-order accurate in time (2-2) scheme is used for a specific subsets of numerical tests in Section 2.2.3. This gives equivalant equations for the updating of acoustic particle velocity

$$v_{j+\frac{1}{2}}^{\tau+\frac{1}{2}} = \frac{\Delta t}{p_{0_\xi}} \frac{p_{j+1}^{\tau} - p_j^{\tau}}{\Delta x} + v_{j+\frac{1}{2}}^{\tau-\frac{1}{2}} \quad , \tag{2.15}$$

and for the calculation of the updated pressure field based on the spatial gradient of particle velocty

$$p_j^{\tau+1} = \Delta t \rho_0 c_0^2 \frac{v_{j+\frac{1}{2}}^{\tau+\frac{1}{2}} - v_{j-\frac{1}{2}}^{\tau+\frac{1}{2}}}{\Delta x} + p_j^{\tau} \quad . \tag{2.16}$$

The stability conditions for FDTD schemes vary depending on the number of dimensions being simulated, and the number of coefficients used to calculate spatial and temporal gradients. Using the definition of the CFL in Eq.(2.12), the stability condition for the 2-4 scheme is CFL $\leq (4\mathrm{D}/3)^{-1/2}$ while the stability condition for the 2-2 scheme is CFL $\leq \mathrm{D}^{-1/2}$ where D is the number of dimensions being simulated [128].

2.2 Testing of Individual Numerical Errors

2.2.1 Overview

In the following sections, the impact of various factors which affect the convergence of FDTD and PSTD models for the case of transcranial ultrasound simulation is presented. These comprise:

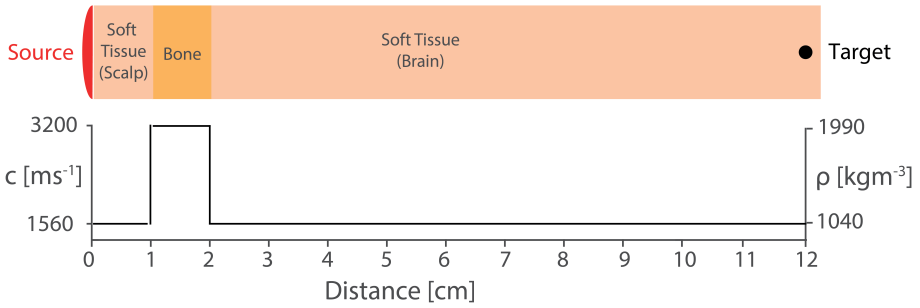

Figure 2.2: A scaled schematic of the simulation model used to evaluate the impact of numerical errors.

1. The impact of bandlimited interpolation of the field variables.

2. Errors due to the use of a PML to truncate the simulation domain.

3. Numerical dispersive phase error.

4. Errors in reflection and transmission from step-change discontinuities in acoustic properties.

5. Staircasing errors in the representation of continuous geometries on a discrete Cartesian grid.

The first two represent fundamental considerations in PSTD simulations, and were dealt with independently, evaluating the spatial and/or temporal sampling necessary to reduce erroneous pressure signals below -60 dB. For the subsequent tests of error, the specific inaccuracies occurring when simulating the propagation of ultrasound from a source in the deep brain to an external transducer were established. This was modelled as consisting of 10 cm propagation through cerebral soft-tissue, 1 cm propagation through bone, and 1 cm additional propagation through superficial soft-tissue, as shown in Fig. 2.2. When examined, reverberations within the skull were considered to consist of an additional 20 cm propagation through soft-tissue, and 2 cm propagation through bone. The acoustic properties used in simulations were taken from Duck [84]. When homogeneous simulation grids were used, they were assigned the acoustic properties of brain tissue, a density of 1040 kg m^{-3} and a sound speed of 1560 m s^{-1}, which were also used to represent soft-tissue in general.

For heterogeneous simulations, bone tissue was assigned a density of 1990 kg m^{-3} and a sound speed of 3200 m s^{-1}. Frequency filtered Kronecker delta functions, like that shown in Fig. 2.3 (b), were used to create broadband pressure sources. The detailed parameters of the different simulation grids are described in Appendix B.1. Accuracy was quantified in terms of the resulting error in the amplitude and time-of-flight of the temporal maximum intensity at the target position. Intensity is calculated from pressure via a plane wave assumption,

$$\mathrm{I} = \frac{p^2}{\rho_0 c_0} \quad . \tag{2.17}$$

The spatial or temporal sampling criteria necessary to constrain the error in intensity amplitude below 10% and the error peak position by 1.5 mm, are established. Positional error is calculated from the error in time-of-flight using the speed of sound in soft-tissue.

In this thesis, spatial and temporal sampling criteria are defined in terms of spatial points per wavelength (PPW) and temporal points per period (PPP), where

$$\mathrm{PPW} = \frac{\lambda}{\Delta x} \quad , \qquad \mathrm{PPP} = \frac{1}{f \Delta t} \quad . \tag{2.18}$$

A constant relationship between the local spatial PPW and the temporal PPP is defined by the CFL,

$$\mathrm{CFL} = \frac{c_0 \mathrm{PPW}}{c_{\max} \mathrm{PPP}} \quad , \tag{2.19}$$

where c_0 is the local sound speed.

Higher values of PPW & PPP will lead to more accurate numerical simulation, but also increase the computational burden and resulting time taken for simulations. When it was necessary to define a specific ultrasonic frequency of interest to calculate the required sampling criteria, 500 kHz was used. This frequency has seen extensive use in studies of UNMS (see Table 1.3), sits within the range of ultrasound frequencies demonstrating optimal transcranial transmission [15], and has a theoretical minimum focus size of approximately 1.5 mm diameter in soft-tissue. PPW values were calculated using the sound speed in soft-tissue.

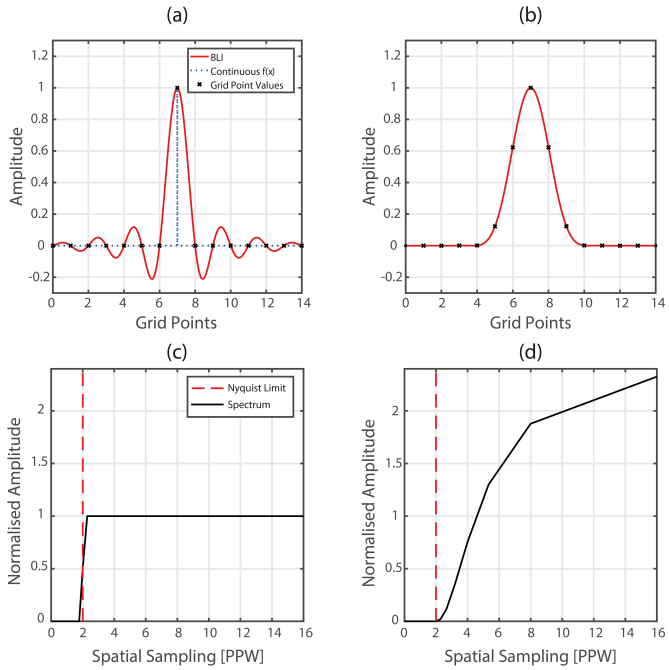

Figure 2.3: The effect of bandlimited interpolation on various functions defined on the discrete grid, and the impact of smoothing on the power spectra of those functions. (a) The BLI of an unsmoothed delta function, (b) the BLI of the same delta function following frequency filtered with a Blackman window, (c) the normalised frequency spectrum of the unsmoothed delta (d) the spectrum of the smoothed delta, showing the redistribution of energy away from the Nyquist limit.

2.2.2 The Bandlimited Interpolant

Introduction

The BLI is the continuous representation of field parameters when approximated by discrete Fourier transform. For an idealised 1D grid with N grid points, length 2π, spatial discretisation Δx, and grid indices j, the discrete Fourier transform is written

$$F_k = \Delta x \sum_{j=0}^{N-1} e^{-x_j i k_x} f_j \quad , \tag{2.20}$$

where

$$k_x = -\frac{\pi}{\Delta x}, -\frac{\pi}{\Delta x} + \frac{2\pi}{\Delta x N_x}, \ldots, \frac{\pi}{\Delta x} - \frac{2\pi}{\Delta x N_x} \quad . \tag{2.21}$$

Here f_j is the value of the discretely sampled function at position x_j, and F_k is the set of coefficients for the series of wavenumbers, k_x. To explicitly calculate the

continuous BLI of a discretely sampled function, the inverse of the discrete Fourier transform is evaluated continuously over $0 < x < 2\pi$, rather than just at the discrete points x_j. This is written

$$b(x) = \frac{1}{2\pi} \sum_{k_x=-\pi/\Delta x}^{\pi/\Delta x-1} e^{ik_x x} F_k \quad , \tag{2.22}$$

where $0 < x < 2\pi$. This will give the continuous BLI $b(x)$ that the discrete set of wavenumbers computed by Eq.(2.20) actually represent [109]. Equation (2.22) can be used to understand the effect that the BLI has on gradient calculation by examining the BLI of a delta function at $j = 0$;

$$b(x) = \frac{\Delta x}{2\pi} \left\{ \cot\left(\frac{x}{2}\right) \sin\left(\frac{\pi}{\Delta x}\right) \right\} \quad . \tag{2.23}$$

This is the periodic sinc function S_n, which constitutes the BLI of the discrete delta function. As any discretely sampled function may be considered a sum of delta functions, this allows the BLI to be computed for any discrete function f_j

$$b(x) = \sum_{j=1}^{N} f_j S_n(x - x_j) \quad . \tag{2.24}$$

This is a convolution of the discretely sampled function with the periodic sinc function. Figure 2.3 (a) & (b) demonstrate how this can lead to Gibbs oscillations in the BLI around discontinuities in the discretely sampled pressure. These exist because the Fourier coefficients of the sampled function do not decay to zero before the Nyquist limit of the grid. It is important to understand that this representation is not erroneous *per se*, but that there is a disparity between the intended input to the PSTD scheme, and what the scheme is capable of representing via a bandlimited Fourier series. To reduce the size of the disparity, smoothing of the intended field can be used to force the Fourier coefficients to decay [129]. This is shown in Fig. 2.3 (b) for a delta function frequency filtered with a Blackman window. Although this remains an inexact representation of the original delta function, the non-oscillating BLI more closely matches the intended underlying pressure distribution as defined

by the values at the discrete grid points.

Testing Method

Band-limited interpolation, as described above, can result in a discrepancy between the intended pressure field and the representation of that field within PSTD schemes when the Fourier coefficients of the intended field have not decayed sufficiently. This manifests as undesired, oscillating pressure values across the simulation grid (see Fig. 2.3 (a)). Therefore, to examine the impact of BLI effects, it is necessary to determine the amplitude of these undesired pressures relative to that of an intended input.

In practice, the error in the representation of a particular pressure distribution will depend on how well it can be represented by a discrete, bandlimited Fourier series at a specific spatial discretization [109]. Tonebursts have a well-defined frequency spectrum determined by their length and central frequency. Therefore, to approximate the BLI error likely to be generally observed, a series of time-varying 10, 30 and 50 cycle acoustic toneburst sources with central wavenumbers approaching the spatial Nyquist limit were used as input signals. These sources have 22.7%, 7.4% and 4.3% full width at half maximum (FWHM) bandwidth as a percentage of central frequency, respectively. The source was positioned a quarter of the way along a homogeneous 1D computational grid with no PML, to prevent any attenuation of the global BLI. The simulations were run for the time taken for waves to travel from the source to the centre of the grid, $t = \frac{\mathrm{N}\Delta\mathrm{x}}{4c_{\mathrm{st}}}$ where c_{st} is the sound speed of brain/soft-tissue, and N and Δx are the number of grid points and the spatial discretization, respectively.

The pressure was recorded at every grid point of the other half of the grid which, according to causality, should have remained quiescent if the BLI of the pressure field matched the intended input of compactly supported tonebursts. Error was quantified as the maximum pressure recorded across the second half of the grid relative to the peak pressure of the source toneburst.

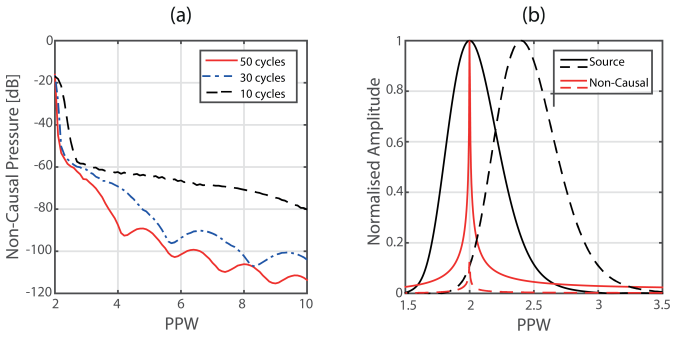

Figure 2.4: Results from tests of the impact of the BLI. (a) Non-causal pressure amplitude as a function of toneburst central PPW for different toneburst lengths and (b) Normalised amplitude spectra of source tonebursts with central frequencies of 2 PPW (solid) and 2.4 PPW (dashed) and the corresponding spectra of the non-causal pressure signals, demonstrating how non-casual pressures relate to the 2 PPW component of the source signal.

Results

Results are shown in Fig. 2.4 (a) as a function of the PPW of the central wavenumber of the toneburst. The amplitude of the non-causal pressure drops rapidly as the number of PPW increases from the Nyquist limit. Reducing the error requires a higher number of PPW for shorter tonebursts due to their wider frequency spectra, but for all three toneburst lengths the error drops to below -60 dB by 3 PPW. An additional observation was that wavenumbers corresponding to less than 2 PPW are not aliased or otherwise propagated on the grid. To determine what frequencies comprised the observed non-causal pressure, the results obtained from 10 cycle tonebursts were examined further. Time-varying pressure signals were recovered from the grid points closest to the wave front, which experienced the peak non-causal pressures. The normalised amplitude spectra of these signals resulting from source tonebursts with central wavenumbers sampled at 2 and 2.4 PPW are displayed in Fig. 2.4 (b), alongside the corresponding amplitude spectra of the source tonebursts. The recorded spectra demonstrate a sharp peak at 2 PPW regardless of the central frequency of the source toneburst, and the amplitude of the peak is proportional to the amplitude of the 2 PPW component of the source toneburst. Practically, these results demonstrate that this error reduces very rapidly as the spatial sampling of the pressure distribution increases, and at 3 PPW BLI errors are

reduced to below -60 dB.

2.2.3 The Perfectly Matched Layer

Introduction

The edges of the computational domain pose particular problems when simulating ultrasound propagation. The discrete Fourier transform used to calculate spatial gradients for PSTD and k-space corrected schemes makes an implicit assumption that the grid is periodic. This leads to waves passing out of the computational grid reappearing on the opposite side of the grid, a phenomena referred to as wrapping. In the FDTD scheme used, outgoing pressure waves are perfectly reflected from the edge of the grid, although in general the exact response depends on the boundary conditions enforced at the edges of the domain.

Fortunately, the use of the coupled first-order equations in k-Wave permits implementation of Berengers split-field PML which surrounds the computational domain and absorbs waves propagating towards the edge of the grid [125]. By splitting the density and particle velocity into their Cartesian components, absorption can be applied selectively to the normally incident components of the wavefield approaching the edge of the domain. The absorption profile of the PML must be selected carefully. If the acoustic absorption increases too rapidly within the PML, the sharp change in acoustic properties will lead to reflections, whereas a low absorption may not attenuate waves sufficiently to prevent reflections or wrapping at the edge of the domain. Within k-Wave, a variable absorption profile is used, which starts low to minimise reflection and increases towards the outer edge of the PML, described by the following formula [121]

$$a_\xi = \frac{c_j a_0}{\Delta\xi}\left(\frac{j - j_p}{d}\right)^m \quad . \tag{2.25}$$

Here d is the thickness of the PML in grid points, j is the spatial index of a point on the grid within the PML, j_p is the spatial index at the start of the PML and a_0 is a normalised absorption coefficient with units of Nepers per grid point. The m variable determines the absorption gradient - several examples are shown in Fig.

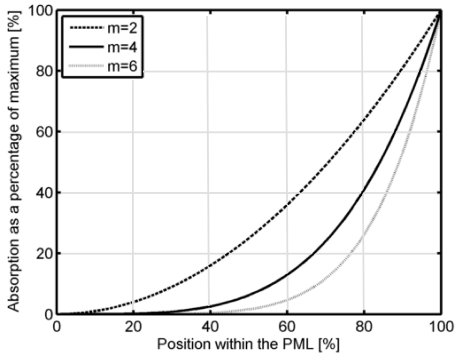

Figure 2.5: Profiles of the PML produced by Eq.(2.25) for different values of m. By default, k-Wave uses a value of $m = 4$

2.5.

Testing Method

When examining the accuracy of assumptions made when implementing a PML, the first consideration is how well the layer absorbs normally incident waves while preventing reflections from its surface. The splitting of the incident wavefield into Cartesian components means that the PML is intended to specifically and totally absorb these waves, as they contain no components parallel to the boundary layer.

The response of the PML to normally incident waves was examined by propagating broadband pressure sources towards the PML on a 1D simulation grid, and recording the incident, transmitted and reflected time-varying pressure signals. These signals were spectrally decomposed, and used to calculate reflection and transmission in decibels as a function of the spatial PPW. In order to prevent numerical dispersion affecting the results when testing the FDTD scheme, the 2-2 FDTD numerical scheme described in Section 2.1.3 was used, with the CFL set to 1. It is a peculiarity of this scheme (explored further in Appendix A.2) that for 1D simulations on homogeneous grids, a CFL of 1 prevents all numerical dispersion. However, due to stability constraints, a CFL of 1 cannot be used with this scheme in higher dimensions and it will not prevent dispersion in heterogeneous media.

The PML was also tested in 2D to determine its dependence on the angle of incidence of incoming waves. A broadband point source was placed close to the

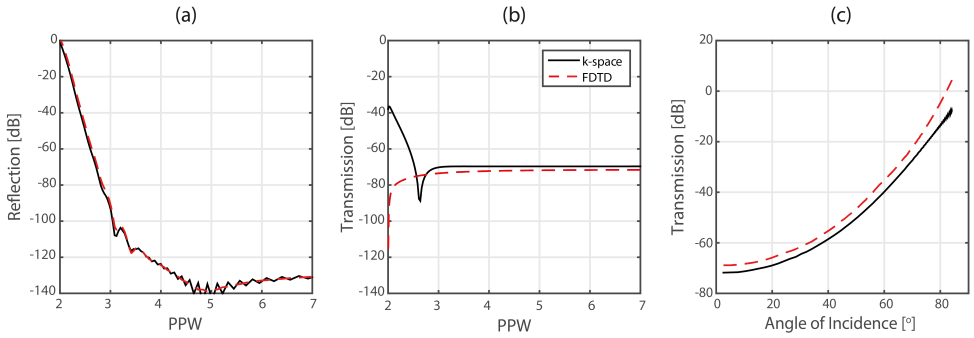

Figure 2.6: Results of tests of the effectiveness of the PML. (a) Reflection from and (b) transmission through the PML in 1D as a function of PPW. (c) Transmission through the PML as a function of angle of incidence.

edge of the PML on the 2D grid and propagated into, and across the surface of, the PML. The time-varying pressure was recorded across the nearby edge of the grid, with each recording position corresponding to a different angle of incidence. A reference simulation was carried out with PML absorption set to zero, and the peak temporal pressure calculated for each recording position. PML absorption was turned on for the test simulation, and the change in the peak temporal pressure at each recording position was used to calculate the transmission factor in dB as a function of angle of incidence.

Results

Results from simulations in both 1D and 2D are shown in Fig. 2.6 for k-space corrected PSTD and FDTD schemes. The results for normally incident waves, demonstrate a clear dependance on spatial sampling at frequencies close to the Nyquist limit of the grid. For both schemes, the pressure reflection coefficient rises steadily from below -120 dB for frequencies sampled at above 4 PPW to total reflection at 2 PPW. Transmission to the edge of the grid remains constant at below -70 dB for both schemes until spatial sampling drops beneath 3 PPW, below which the k-space scheme shows an increase in transmission and the FDTD scheme shows a reduction in transmission. These results indicate that the effectiveness of the PML is greatly reduced for wavenumbers sampled at below 3 PPW, and it cannot be relied on at these PPW values. However, erroneous reflection and transmission reduce rapidly as sampling increases. It should be noted that pressure reaching the edge of the grid

for both schemes is subject to further attenuation within the PML when reflected or wrapped back into the grid. Furthermore, the BLI will have influenced the behavior of these tests for frequencies sampled at close to the Nyquist limit, which may explain why the k-space scheme appears to show an increase in both reflection and transmission close to 2 PPW.

Results for the 2D simulations are shown in Fig. 2.6 (c) as a function of angle of incidence for both numerical schemes, demonstrating a clear dependence of transmission through the PML as a function of angle of incidence. Transmission is lowest for normally incident waves, rising with increasing angle of incidence crossing to above -60 dB at 40°. No clear relationship between angle of incidence and reflection from the PML was observed. These results should be considered when designing acoustic sources and considering the angles at which pressure waves will impinge on the PML.

2.2.4 Numerical Dispersion

2.2.4.1 Introduction

Phase error, or numerical dispersion error, refers to an unphysical dependence of sound speed on frequency within the numerical scheme (in contrast to physical dispersion, a distinct acoustic phenomenon related to acoustic absorption), and was identified by Pulkkinen et al. [130] as a potential source of error when using numerical simulations to derive the phase shifts necessary for focusing ultrasound through the skull.

Finite difference approximations of gradients can be considered a rearrangement of a truncated Taylor series to solve for derivatives, and introduce an inherent error, explored in Appendix A.2. When used to solve the wave equation, this manifests as a dependence of the sound speed on the wavenumber and causes a cumulative error in the phase of travelling acoustic waves. In FDTD schemes dispersion error arises as a result of both the spatial and temporal sampling, while for PSTD schemes dispersion arises dependent on the temporal sampling alone.

With regards to k-space corrected schemes, when the correct value for sound speed is used in the correction factor, which will be referred to as perfect k-space

correction, numerical dispersion is completely eliminated for homogeneous media. However, only one sound speed can be employed and perfectly corrected for. This means that in heterogeneous media, phase error can still accumulate dependent on the difference between c_{ref} and the local sound speed. When the local sound speed is above that used in the correction factor, dispersion error is bounded between zero and that obtained from an uncorrected PSTD scheme. When c_{ref} is larger than the local sound speed, the phase error can grow to be larger than in the uncorrected case [131]. However, as shown by Eq.(2.10), setting c_{ref} to the highest sound speed in the medium give the k-space scheme unconditional stability, potentially permitting the use of coarser temporal discretisation with a resulting decrease in computational burden. Therefore it is important to test the error resulting from both under-correction and over-correction, within the context of simulations of ultrasound propagation in a head model.

Testing Methods

To examine the impact of numerical dispersion, a broadband pressure source was defined on a homogeneous 1D grid and propagated a distance of 1 cm. The time-varying pressure was recorded, and the phase spectra of the recorded pressure signals were compared with a dispersion-free reference simulation obtained using perfect k-space correction. Grids with the properties of both bone and brain tissue were tested. As discussed above, when k-space correction is used for simulations in homogeneous media it prevents numerical dispersion, although only a single sound speed can be corrected for. Therefore, to examine the maximum discrepancy between the reference sound speed and the sound speed of the grid, c_{ref} was set to the sound speed of bone for simulations in soft-tissue, and c_{ref} was set to the sound speed in soft-tissue for simulations in bone.

For the 2-4 FDTD scheme used, the temporal and spatial dispersive errors oppose each other, with reduced dispersive error at higher CFL numbers [127]. Therefore the CFL was set to 0.5 for these simulations calculated with reference to the sound speed of bone, the highest value at which both schemes would be stable in 3D for a heterogeneous simulations of propagation through both bone and soft-

tissue [127, 121].

To compare the schemes used, the change in the phase spectra was used to calculate the phase error in radians per cm propagated, ϕ_{error}, as a function of spatial and temporal sampling. However, when transmitting ultrasound across the skull for the purposes of neurostimulation, the ultimate aim is to focus ultrasound on a small area of brain tissue. Therefore when quantifying the error resulting from numerical dispersion in terms of its practical significance, it will be how it affects the position of the simulated pressure wave in soft-tissue, regardless of the tissue in which the dispersion arises. The phase error calculated from the time-trace was then converted to a corresponding positional error in soft-tissue per cm propagated in both bone and soft-tissue by multiplying the phase error by the wavelength in soft-tissue:

$$\text{Positional Error}(f) = \frac{c_{\text{st}}\phi_{\text{error}}}{2\pi f} \quad . \tag{2.26}$$

This was then used to calculate the final positional error using the model for transcranial propagation of ultrasound to a deep brain target shown in Fig. 2.2, and for up to two reverberations within the skull. This was used to calculate the sampling criteria required to obtain $<$1.5 mm positional error.

Results

The phase error in the recorded pressure signal per cm propagated in both brain and soft-tissue is shown as a function of PPP in Fig 2.7 for the 2-4 FDTD scheme, uncorrected PSTD scheme, and k-space corrected PSTD schemes. It should be noted that, for the FDTD scheme, each PPP value is related to a PPW value by the CFL. Two main findings are apparent from these results. Firstly, the FDTD scheme consistently performs worse than both uncorrected and k-space corrected PSTD schemes. Secondly, k-space correction of the PSTD scheme leads to reduced error for propagating through bone when c_{ref} is set to the sound speed of soft-tissue, but increased error for propagation through soft-tissue when c_{ref} is set to the speed of sound in bone. Based on these results it can be concluded that, while setting c_{ref} to the sound speed of bone would result in unconditional stability, the dispersion

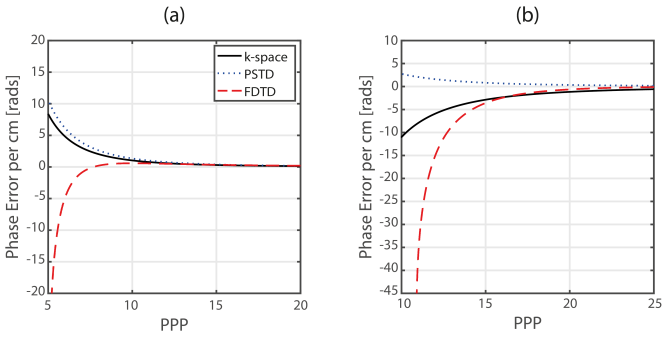

Figure 2.7: Phase error per centimetre propagated in (a) bone, and (b) soft-tissue for FDTD, uncorrected PSTD and k-space schemes.

Target	Direct Path [PPP]	First Reverb. [PPP]	Second Reverb. [PPP]
FDTD	17.9	21.9	23.9
PSTD	11.2	18.9	24.2
k-space	4.0	4.0	4.6

Table 2.1: Temporal sampling required to obtain less than 1.5 mm targeting error for direct propagation to a deep brain target, and for multiple reverberations within the skull.

error alone would necessitate finer temporal sampling than would be necessary to maintain stability in the opposite case. Moreover, in transcranial simulations the majority of the simulated domain will have the acoustic properties of soft-tissue, increasing the impact of dispersion errors arising in soft-tissue. Therefore it is apparent that simulations using the k-space method should employ k-space correction based on a c_{ref} value derived from the soft-tissue.

The values for phase error per cm propagated in each tissue can be used to calculate a positional error at the target. The temporal sampling values required to attain an error of less than 1.5 mm at a deep brain target based on the phase error for the different numerical schemes are shown in Table 2.1. The positional error for the k-space scheme is calculated based on the total error that would occur when c_{ref} is set to the speed of sound in soft-tissue.

2.2.5 Reflection and Transmission from Discontinuities

Introduction

When an acoustic wave is incident on an interface between two media with different acoustic properties, such as bone and soft-tissue, it will undergo varying degrees of reflection and transmission. Depending on the angle of incidence, and the ability of the different media to support shear waves, there may also be conversion between compressional and shear modes [9].

When compressional waves are normally incident on an interface, there is no mode conversion. Pressure reflection and transmission coefficients are defined with respect to incident, reflected and transmitted pressures, indicated in Fig. 2.8 (a). Based on the boundary conditions at a discontinuous interface, coefficients can be expressed in terms of the characteristic acoustic impedance of the two media, $Z_n = \rho_n c_n$ [122]:

$$R = \frac{p_{\text{refl}}}{p_{\text{init}}} = \frac{Z_2 - Z_1}{Z_2 + Z_1} \qquad T = \frac{p_{\text{trans}}}{p_{\text{init}}} = \frac{2Z_2}{Z_2 + Z_1} \ . \tag{2.27}$$

Normal incidence pressure reflection and transmission coefficients can be used to calculate intensity reflection and transmission coefficients, R_i & T_i respectively :

$$R_i = R^2 \qquad T_i = 1 - R_i \ . \tag{2.28}$$

These reflection coefficients are derived from, and assume, a discontinuous interface. When two media are connected by a smoothly varying interface, the behaviour is more complicated. Reflection and transmission becomes dependent both on the impedance profile and the frequency of incident ultrasound. The behaviour varies depending on the impedance profile, but the general rule is that high frequencies are reflected less from smooth interfaces [132]. An added complication arises when acoustic particle velocity is calculated on a staggered grid, as the grid density map must be defined separately on the staggered grid, which is offset by $\frac{\Delta x}{2}$ relative to the main grid.

Ideally, numerical simulations should reproduce the behaviour from Eqs.(2.27)

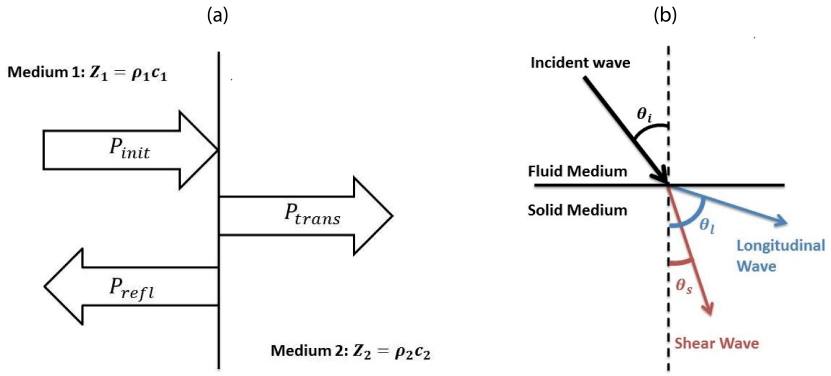

(c)

Image Removed on Copyright Grounds

Figure 2.8: Reflection and transmission of ultrasound. (a) Reflection and transmission at normal incidence. (b) Generation of shear and compressional waves for non-normal incidence on a solid. (c) Relative amplitudes of shear and compressional waves as a function of angle of incidence (reproduced from Clement et al. [8]).

when acoustic plane waves are normally incident on a discontinuity defined on a discrete grid. However, due to the difficulty in representing discontinuities using spectral methods, the representation of step-changes in acoustic media has previously been identified as a weakness of PSTD methods [121]. As discussed in Section 2.1.2, during the iterative time-steps of the k-space scheme the pressure field is multiplied by the maps of medium density and sound speed, before being represented as a truncated Fourier series for the purpose of calculation of gradients. Step changes in medium properties will therefore introduce Gibbs phenomenon into the pressure field.

For non-normal incidence of compressional waves onto an interface, the trans-

mission angle and the amplitude of the transmitted and reflected waves will change dependent on the angle of incidence and difference in compressional sound speed. If the second media can support shear waves, as in the case of propagation of ultrasound through the scalp into the skull, there may also be a degree of acoustic mode conversion, shown in Fig. 2.8. The angles of the transmitted waves are related to the angles of incident waves (shown in Fig. 2.8) by the different sound speed values via Snell's law:

$$\frac{\sin\theta_i}{c_i} = \frac{\sin\theta_l}{c_l} = \frac{\sin\theta_s}{c_s} \quad . \tag{2.29}$$

As a result of the dominance of compressional modes for normal or near-normal incidence, and the high absorption of shear waves in bone, shear mode propagation is often neglected in models of transcranial transmission [8].

Testing Method

To examine the error in reflection and transmission from medium discontinuities, broadband pressure sources were propagated from soft-tissue into bone (propagation direction across the interface was observed to make no difference) on a 1D simulation grid. The incident, reflected, and transmitted waves were recorded and the power spectra used to calculate intensity reflection and transmission coefficients for each wavenumber. Percentage error in these coefficients was calculated through comparison with analytical values for the intensity reflection and transmission coefficients. The error in the intensity transmission coefficient was then used to calculate the error in intensity when ultrasound is transmitted through the skull across two interfaces between bone and soft-tissue

$$\text{Transmission Error}(f) = \frac{\tilde{\text{Te}}(f)^2 - \text{Te}^2}{\text{Te}^2} \quad . \tag{2.30}$$

Here Te is the true intensity transmission coefficient calculated using Eqs.(2.27) & (2.28), and $\tilde{\text{Te}}(f)$ is the simulated intensity transmission coefficient as a function of frequency.

Two additional tests were carried out to further characterise the response of the k-space scheme to interfaces between acoustic media. Firstly, the interface between

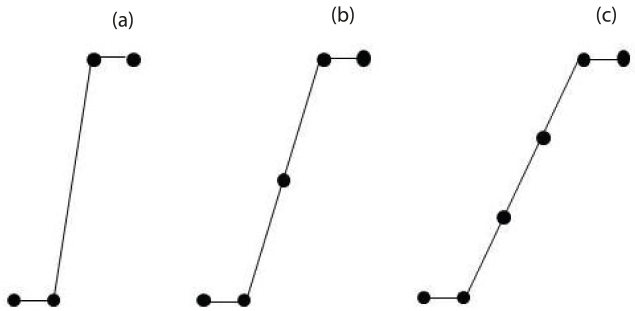

Figure 2.9: Smoothing of the interface between bone and soft-tissue. (a) Unsmoothed, (b) one point smoothing, (c) two point smoothing.

bone and soft-tissue interface was progressively smoothed via linear interpolation, to determine if smoothing of the interface leads to an improved representation of the discontinuous interface, similar to that shown in Fig. 2.3 for a delta function pressure distribution. The acoustic properties of the grid were smoothed via two stages of linear interpolation, shown in Fig. 2.9, and simulations were carried out using the k-space PSTD scheme only. Secondly, a test was carried out to evaluate the impact of the magnitude of the impedance change, to determine if the performance of the k-space corrected PSTD scheme relative to the FDTD scheme is dependent on the size of the medium discontinuity. The soft-tissue portion of the domain was assigned a sound speed of 1000 m s^{-1} and a density of 1000 kg m^{-3}. The impedance of the bone portion of the grid was then varied up to ten times that of the soft-tissue, with sound speed and density varied independently to determine their relative impact. Furthermore, to determine the impact of the linear interpolation and smoothing of acoustic density values onto the staggered grid, simulations across density interfaces were repeated with this interpolation disabled, so that density changes were represented as step changes on both the main and staggered grids.

Results

Results for the error in intensity transmission and reflection coefficients as a function of PPW are shown in Fig. 2.10. These results demonstrate the increasing error in both transmission and reflection at lower levels of spatial sampling. Frequencies sampled at close to the Nyquist limit of the grid are reflected more from the

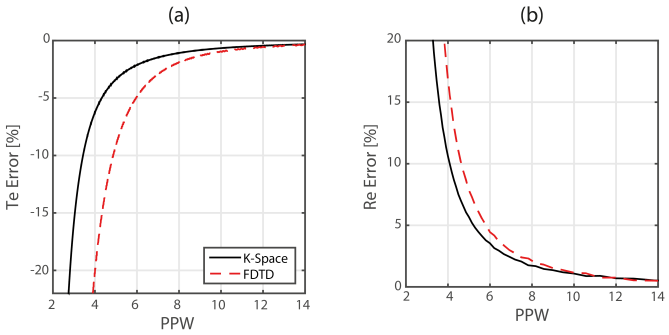

Figure 2.10: Error in simulated intensity (a) transmission (Te) and (b) reflection (Re) coefficients as a function of spatial PPW.

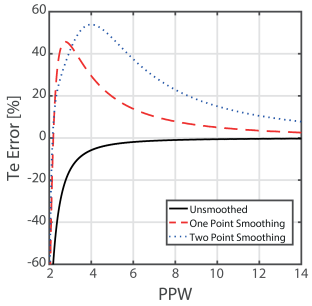

Figure 2.11: The error in intensity transmission coefficient at different levels of interface smoothing as a function of PPW.

interface, and transmitted less. The results also demonstrate that the FDTD scheme performs worse than the k-space PSTD scheme. This result is notable, given the supposed weakness of PSTD schemes at representing discontinuities. In order to obtain less that 10% error in intensity transmitted through the skull, calculated using Eq.(2.30), the FDTD scheme requires 5.9 PPW while the k-space and PSTD schemes require 4.3 PPW. The results of tests on the impact of smoothing the interface are shown in Fig. 2.11. Increasing the levels of smoothing leads to increased simulated transmission coefficients and a reduction in simulated reflection coefficient, and it is clear from these results that smoothing the interface does not result in increased accuracy insofar as approximating Eqs.(2.27) & (2.28). As mentioned above, smoothly varying interfaces between acoustic media will have a frequency dependent reflection and transmission, with increased transmission at high frequencies, so this is likely a manifestation of that change, rather than an effect related

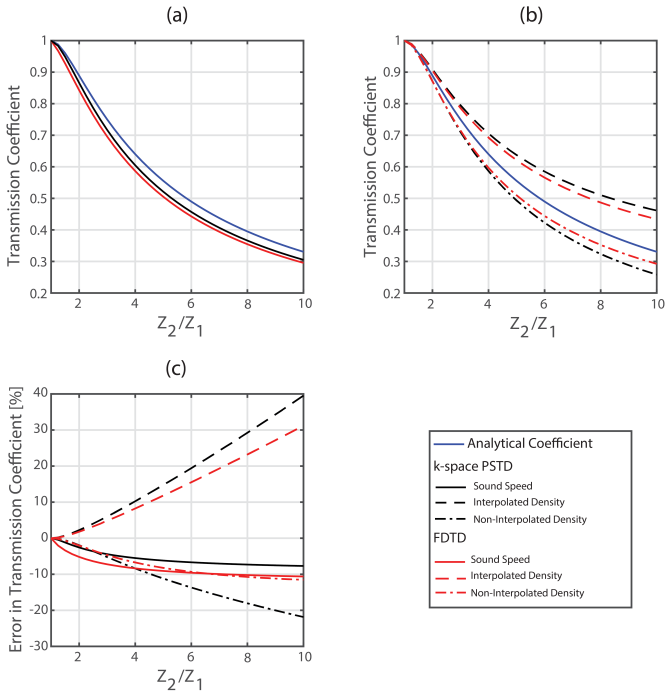

Figure 2.12: Effect of increasing impedance change on simulated transmission coefficients. Transmission coefficients are calculated spectrally for the ultrasound wavelength sampled at 6 PPW. (a) Simulated and analytical intensity transmission coefficients for changes in sound speed. (b) Simulated and analytical intensity transmission coefficients for changes in density, including simulations where the step change in density is not linearly interpolated onto the staggered grid. (c) Percentage error in simulated transmission coefficients for changes in sound speed and density.

to the numerical accuracy of the schemes. Results for the test of the impact of increasing impedance change are shown in Fig. 2.12, showing the simulated intensity transmission coefficients relative to the analytical coefficients and the percentage error in the simulated coefficient, both as a function of impedance change for both for FDTD and k-space corrected PSTD schemes. The simulated transmission coefficients shown were calculated spectrally for a PPW value of 6. When the impedance change is the result of a step change in sound speed, shown in Fig. 2.12 (a), the simulated transmission coefficient is lower than the analytical value, and the absolute error in the coefficient appears to approach a constant value with increasing impedance change. When the impedance interface is the result of a change in density, and the density has been linearly interpolated onto the staggered grid, the transmission coefficient is erroneously large, and the disparity increases at larger

impedance changes. However, when the staggered grid representation of density has not been linearly interpolated, the transmission coefficient is erroneously small. When compared with the impact of smoothing the interface shown in Fig. 2.11, this indicates that the interpolation of density onto the staggered grid results in the increase in simulated transmission coefficient observed. Generally FDTD and PSTD schemes show similar behaviour with increasing impedance change for both density and sound speed interfaces, with the k-space PSTD scheme performing better when simulating changes in sound speed, and the FDTD scheme performing better when simulating changes in density.

2.2.6 Staircasing

Introduction

Staircasing refers to the error in the representation of continuous geometries on a discrete, regular grid when performing simulations in more than one dimension. Curved surfaces and lines at an angle to the Cartesian axes of the grid will be mapped in a stair-stepped manner, but any geometry that cannot be perfectly represented on the grid will be subject to some approximation. Furthermore certain vertex and edge positions will not correspond exactly to points on the grid and will be offset.

The impact of staircasing was examined separately for acoustic sources and heterogeneous media with the acoustic properties of bone and soft-tissue. Tests involved recording the time-varying pressure at a number of positions across the field resulting from a staircased representation of a source or medium, and comparison of these signals with references obtained from a staircase free simulation. Error was then quantified as the percentage error in the amplitude of the temporal peak intensity (calculated using a plane wave assumption) and its positional error (derived from the change in the time-of-arrival of the intensity peak) as a percentage of wavelength. A positional error of 50% of wavelength corresponds to 1.5 mm for a source frequency of 500 kHz in brain tissue. These errors were calculated for each recording position, and then averaged across the field to give mean errors in peak intensity amplitude and position. No notable difference in error was observed

Figure 2.13: Simulation layouts used to test the impact of staircasing (not to scale). Sources are shown in red, and pressure recording positions in black. (a) Non-staircased line source used as reference, and (b) staircased line source used to examine error. (c) Non-staircased medium map used as a reference, and (d) staircased medium used to examine error. (e) High resolution map of a bone-tissue layer used as reference, and (f) downsampled medium. Reproduced from Robertson et al. [115].

between FDTD, PSTD, and k-space schemes across all tests.

Source Staircasing

The impact of staircasing on acoustic sources was examined using line-sources with a length of $65 \times \Delta x$, where Δx is the spatial discretization step, at a series of angles to the Cartesian grid. These included four angles that form Pythagorean triangles on the grid, specifically 14.3° (with Pythagorean triple 16, 63, 65), 22.6° (25, 60, 65), 30.5° (33, 56, 65), and 36.9° (39, 52, 65). For these angles, the line-source endpoints are coincident with specific grid point positions, and any error is only due to the staircased representation of the line, rather than endpoint misregistration (both are aspects of staircasing error). A source defined parallel to a Cartesian axis was used as a non-staircased reference. The sources were excited with 10 cycle acoustic tonebursts with central wavenumbers ranging from 3 PPW to 100 PPW. The amplitudes of the source signals were normalised based on any change in the

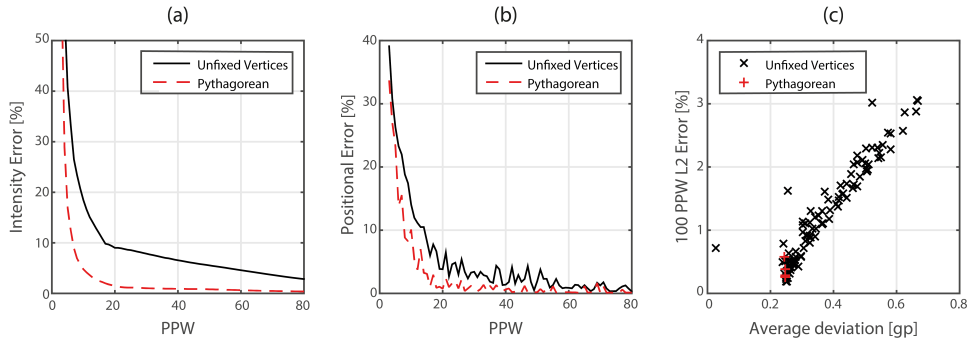

Figure 2.14: Error in staircased line sources as a function of spatial PPW. (a) Percentage error in peak intensity. (b) Positional error as a percentage of wavelength. (c) L2 error at 100 PPW sampling against average deviation from ideal line source. Pythagorean sources have both vertices of the parametric line source exactly defined on the grid.

number of distinct source points used when defining an angled line source when compared to the aligned case. The time-varying field was recorded at 100 points positioned in front of the line source, and the sensor map was rotated with the line source to maintain source-sensor geometry. The simulation layout is shown in Figs. 2.13 (a) & (b). Mean errors in the amplitude and position of the peak intensity across the sensor field were calculated independently for each angle tested relative to the aligned, nonstaircased reference case.

The maximum mean errors across the range of angles tested for each PPW value are shown in Fig. 2.14. These results demonstrate that staircasing errors worsen with lower spatial sampling and are less serious for Pythagorean angles, when endpoints are correctly registered. The error in the position of the intensity peak never rises above 50% of the wavelength for any source. Seventeen PPW are required to obtain less than 10% error in the amplitude of the intensity peak for all angles tested, while Pythagorean angles require only 7 PPW. Although the error examined here does not relate directly to the model for transcranial ultrasound propagation described in Section 2.2.1, these results do indicate that staircasing and spatial sampling must be considered when defining acoustic source distributions, and that error can be reduced by ensuring endpoint registration. The testing of multiple angles also allowed examination of how the exact mapping of the staircased line relates to the error observed in the resulting field. No clear relationship between

the angle of the line source and the level of staircasing error was observed. However, a staircasing metric was defined as the average distance between the staircased source points and their equivalent equispaced points on an ideal angled line. It was observed that the convergence rate of the error of the staircased source maps, quantified as the average L2 error in the recorded pressure signals at the maximum spatial sampling tested, showed a strong dependence on this staircasing metric (shown in Fig. 2.14 (c)). Although only a simple metric, this indicates that the severity of staircasing error can be predicted through comparison of an ideal or parametric map of the intended geometry with its staircased representation.

Medium Staircasing

The impact of staircasing of heterogeneous medium properties was examined in two separate tests. The first was conceptually similar to the examination of source staircasing. An acoustic point source excited by 10 cycle acoustic tonebursts with central frequencies ranging from 3 to 80 PPW was propagated across a planar medium boundary (soft-tissue to bone), defined at varying angles to the Cartesian axes. The time-varying pressure field was recorded at 100 sensor points following the interaction with the medium boundary. The sensor points were rotated with the medium boundary to maintain the simulation geometry. A nonstaircased boundary defined along a Cartesian axis was used as a reference map. This simulation layout is shown in Fig. 2.13 (c) & (d).

The second test was designed as a more accurate model of staircasing in transcranial transmission. A 10 cycle, 2520 PPW toneburst was propagated through a medium map comprising a quarter circle bone layer. The medium was then artificially staircased through spatial downsampling, before being remapped to the original grid. A 3780×3780 simulation grid was used due to the large number of integer factors of 3780, which allowed the medium to be successively downsampled while maintaining positioning. The timevarying pressure was recorded across a quarter-circle, and error metrics computed through comparison with the least staircased medium distribution. This simulation layout is shown in Figs. 2.13 (e) & (f). An effective PPW value for each level of downsampling was calculated through

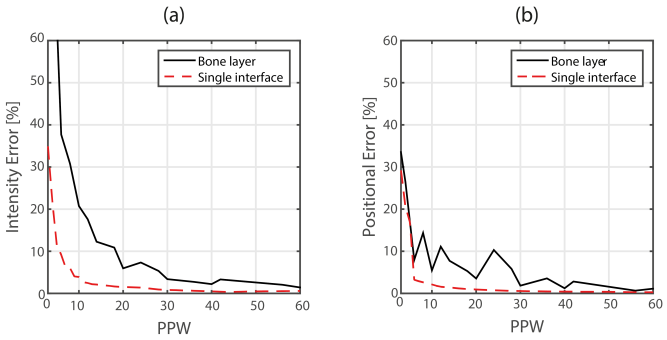

Figure 2.15: Error resulting from propagation through bone-layer and single interface staircased medium boundaries. (a) Percentage error in peak intensity magnitude. (b) Positional error of peak intensity as a percentage of wavelength.

comparison of the source PPW with the new effective spatial discretization.

Mean error measurements across the recorded fields as a function of PPW are shown in Fig. 2.15 for both medium staircasing tests. For the single interface model, the maximum mean errors across the range of angles tested are shown. In terms of the impact on the model for transcranial propagation, the results for the bone layer model shown in Fig. 2.15 suggest that 20 PPW or above are required to obtain less than 10% mean error in intensity transmitted to an external transducer surface. As might be expected, the errors for the single interface are lower, with 6 PPW required to obtain the same error in peak intensity amplitude. The error in peak intensity position is less serious, with mean positional error never rising above 50% of wavelength (1.5 mm for 500 kHz ultrasound in brain tissue) for both tests, as with source staircasing. This may be due to staircasing introducing a random error in acoustic pathlength, leading to a defocusing and change in amplitude rather than a shifting in the peak position. To place this in context, the voxel size of clinical CT images is on the order of 0.5 mm at best. This corresponds to 6 PPW at 500 kHz, and fewer at higher frequencies. This suggests that staircasing may have a significant impact on simulations using image derived medium property maps. When considered alongside the results for source staircasing, these results indicate that staircasing error is likely the most serious of the numerical errors tested.

2.3 Discussion

In this chapter, initial simulations examined the reduction in the effectiveness of the PML, and the impact of the BLI when using k-space and PSTD methods. Both the PML and BLI lead to erroneous pressures appearing on the grid (in the form of transmission/reflection and non-causal propagation, respectively) when simulating frequencies sampled at close to the spatial Nyquist limit. Although both of these effects have the potential to seriously reduce the accuracy of the simulations, they decrease in severity rapidly as the rate of spatial sampling increases. Above 3 PPW, erroneous pressures resulting from both BLI and PML effects were at least -60 dB below the amplitudes of the ultrasound sources being simulated.

Numerical dispersion has a serious effect on the accuracy of FDTD and PSTD schemes, resulting in high temporal sampling requirements to reduce positional error. However, this was not the case for the k-space scheme, where approximately 3 PPW will serve to limit dispersion sufficiently for transcranial transmission for any stable CFL value. Errors in reflection and transmission from discontinuous medium properties manifest in the magnitude of reflected and transmitted simulated intensities. Despite the representation of step changes in media previously being identified as a key limitation of PSTD schemes, the error was shown to be more severe for the 2-4 FDTD scheme tested. To reduce error in the intensity below 10% following transcranial transmission, k-space and PSTD schemes require 4.3 PPW, while FDTD requires 5.9 PPW. Further testing of reflection and transmission indicated that the FDTD scheme does not give improved performance relative to the k-space PSTD method at higher impedance changes, and that sound speed and density changes result in different errors in reflection and transmission. Smoothing of the interface does not represent an improvement in reflection and transmission from the interface, and the interpolation of the density map results in the disparity between sound speed and density observed in the tests of increasing impedance changes.

Staircasing of source and medium geometries was shown to require the most stringent sampling criteria to obtain the required accuracy, and affects FDTD, PSTD

and k-space schemes equally. Both source and medium staircasing were shown to have a greater impact on the intensity amplitude of the toneburst signal being examined than the position of the intensity peak. The results shown in Fig. 2.14 and Fig. 2.15 indicate that 20 PPW are required to reduce the error in peak intensity following transcranial transmission below 10%. The preliminary examination of a potential staircasing metric also suggests that the error resulting from a particular staircased geometry is directly related to its deviation from the ideal geometry.

The work in the present Chapter is subject to some limitations, primarily the degree to which the examination of individual numerical errors can be generalised to different simulation set-ups, although trends and qualitative observations remain valid. Many of the tests only examine toneburst sources, and the error is evaluated over a small field, with pressure recorded at a limited number of sensor positions (See Fig. 2.13). Furthermore, the impact of shear wave propagation was not examined. This will not have affected 1D or homogenous simulations, but a more thorough examination of medium staircasing should include testing of elastic wave propagation. Similarly, no effort was made to examine the manifestation of numerical errors when modeling nonlinear propagation or acoustic absorption, which will become relevant for applications requiring the simulation of high-amplitude ultrasound.

Chapter 3

Convergence Testing of Simulated Time-Reversal Focusing

In the previous chapter, sources of numerical error that affect FDTD and k-space corrected PSTD schemes were examined in isolation to determine their relative impact, and to identify the spatial and temporal sampling necessary to constrain errors in time-of-flight and ultrasound intensity below acceptable levels. However, in practice, simulations of transcranial ultrasound propagation will be subject to all numerical errors acting in concert, and the interaction between different errors might serve to increase or decrease the overall sampling requirements. Furthermore the specific simulation type used in time reversal focusing might affect how the impact of numerical errors manifests in terms of their effect on practical focusing quality. Convergence testing is a method of establishing the sampling required to ensure accuracy for a numerical scheme that is known to be convergent. The spatial and temporal sampling of test simulations are progressively increased, and the error in the numerical simulations evaluated. As sampling increases, numerical simulations will approach the true solution and are said to have converged when the rate of reduction in the error with increasing sampling is reduced below what is considered an acceptable level [119, 120].

In the present chapter, convergence testing of simulated time-reversal in 2D and 3D is presented, to show the combined effects of numerical errors on the effectiveness of transcranial time-reversal focusing. A model of transcranial focusing of

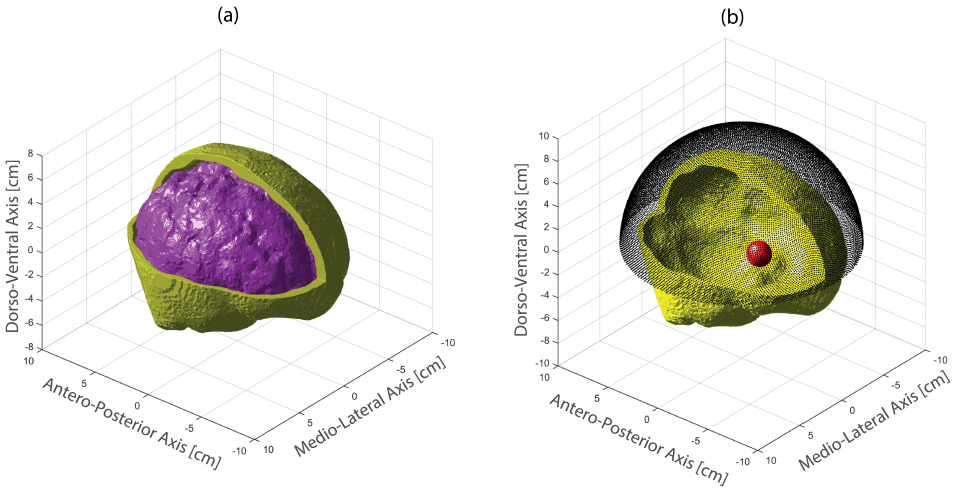

Figure 3.1: Skull and brain meshes used in convergence testing. (a) Rendering of skull mesh used to define bone layer in yellow, and of the brain mesh over which the pressure field was evaluated in magenta. (b) Skull mesh shown with representation of the hemispherical transducer distribution in black, and the position of the time-reversal focus in red

ultrasound from a hemispheric transducer onto a deep brain target is described based on parametric mesh models of the brain and skull and a geometrical description of the transducer. The convergence of the k-space PSTD scheme is then examined based on full 3D simulations of a time reversal protocol, to give a comprehensive picture of the realistic impact of numerical error in transcranial simulations in terms of the error in final focusing quality. The results of simulations carried out in 2D are also presented, which allow the examination of multiple frequencies and higher maximum spatial and temporal sampling without excessive use of computational resources, and serve to verify the results of the 3D simulations. Simulations in 2D also allow comparison between a 2-4 FDTD scheme and the k-space PSTD scheme, as the k-Wave toolbox does not have the facility to carry out large 3D FDTD simulations in a computationally efficient manner [133].

3.1 Skull and Simulation Design

3.1.1 Simulation Layout

The medium property map used in the convergence tests was derived from a T1-weighted MR image obtained from the Imperial College brain development dataset

[134]. Brain and skull volumes were extracted using the FSL MRI processing toolbox [135] and converted into a surface mesh using the iso2mesh toolbox [136]. Renderings of both skull and brain meshes are shown in Fig. 3.1, alongside the hemispheric transducer and location of the focusing target. For 3D simulations, maps of the acoustic properties for both forward and reversal simulations were generated by sampling the skull mesh onto grids at the required spatial discretisation. For reversal simulations, the brainmesh was also sampled onto the 3D grid, to allow the pressure distribution across the brain to be isolated from the rest of the simulation domain. For 2D simulations, acoustic property maps were created in the same way, based on coronal slices through the head like those shown in Fig. 3.2.

For both forward and reversal simulations, the skull was modelled as a single homogeneous layer, assigned the properties of cortical bone. Although heterogeneous models have been shown to give tighter focusing when used for time-reversal, homogeneous models are still effective [90], and are necessary for a convergence test of the numerical accuracy of a time-reversal protocol. A homogeneous map allowed the use of a mesh as a parametric, geometrical map of the skull that can be resampled onto spatial grids of arbitrary discretisation. Furthermore, a heterogeneous skull would have necessitated resampling of heterogeneous acoustic properties derived from a clinical CT image with a spatial resolution on the scale of the acoustic wavelength onto different spatial grids, which would have affected the convergence rate of simulations [112]. The use of a homogeneous map ensures that the convergence of the time-reversal tests is dependent solely on the numerical accuracy of the simulations. The acoustic properties of the skull were taken from Duck [84], with cortical bone assigned a sound speed of 3200 m s^{-1} and a density of 1990 kg m^{-3}. The remainder of the simulated domain was modelled as homogeneous soft-tissue, assigned the acoustic properties of brain tissue, with a density of 1040 kg m^{-3} and a sound speed of 1560 m s^{-1}.

The transcranial transducer was modelled as a hemisphere covering the superior surface of the skull, with a radius of 95 mm, focused on the geometric centre of the skull mesh. The 3D extent of the transducer is shown in Fig. 3.1 (b). In

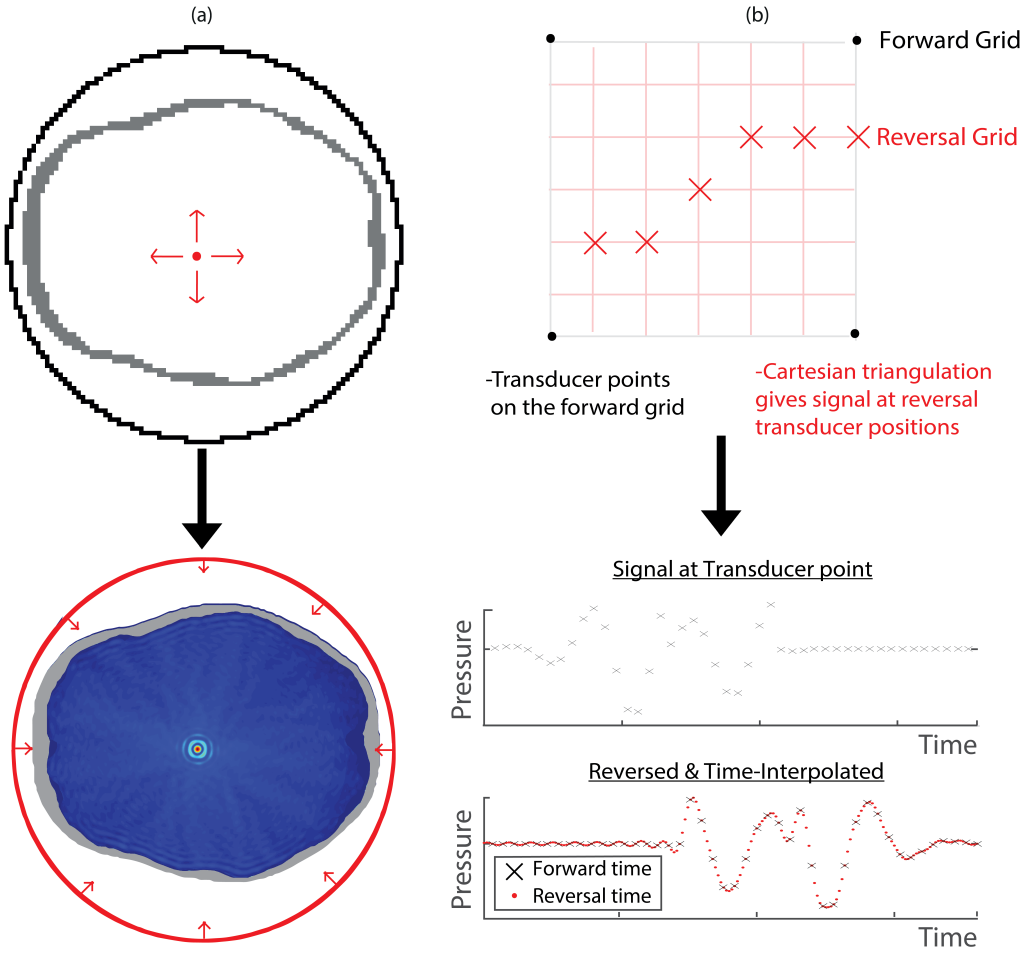

Figure 3.2: Method for convergence testing. (a) Coarsely sampled forward simulation of ultrasound from a target point to simulated transducer surface. Followed by reversal simulation back into the finely sampled head model, with an example of the recorded intracranial pressure field. (b) The pressure signal is spatially triangulated onto the reversal spatial grid, before being time reversed and interpolated onto the finer temporal grid to create drive signals.

2D, the transducer was modelled as a circle surrounding the coronal section of the head, shown in Fig. 3.2 (a) & (c). The radius of the 2D transducer was 91 mm. As these tests were to examine the specific impact of numerical errors on simulation effectiveness, the simulated transducer array surface was modelled as a continuous surface made up of point transducers at the resolution of the reversal spatial grid, and no attempt was made to replicate real transducer characteristics. As with the use of a homogeneous medium map, this ensures that convergence is dependent only on the numerical accuracy of the forward simulations and will apply to any

chosen source conditions.

The target of the time-reversal process, modelled as a source term in the forward simulations, was placed at a slight offset to the geometrical focus of the simulated transducer, to replicate the off-centre targeting that might be necessary to focus ultrasound on some of the neural structures detailed in Table. 1.1. For the 2D simulations, the target was offset by 2 mm laterally and 3 mm anteriorly. For the 3D simulations the target was offset by 5 mm anteriorly, 5 mm laterally and 5 mm superiorly. The approximate position of the source term in the skull, relative to the 3D transducer is shown in Fig. 3.1 (b). It should be noted that, due to the variety of spatial grids, the position of the desired target did not always correspond to a point on the spatial grid that could be modelled as a source point. In those instances the nearest point on the spatial grid was used instead.

3.1.2 Convergence Testing Protocol

The general time-reversal protocol used for convergence testing in both 2D and 3D is shown in Fig. 3.2. The first stage consists of forward propagation of a 10 cycle toneburst from a source point inside a virtual skull model to the circular (2D) or hemispherical (3D) virtual transducer-sensor array.

In order to evaluate the impact of the numerical accuracy of simulations of the time-reversal problem, these forward simulations were carried out with varied spatial and temporal discretisations corresponding to a range of PPW and PPP values for the frequencies simulated. Simulations in 3D were carried out with toneburst sources with central frequencies of 500 kHz, with the spatial sampling ranging from 2 to 16.7 PPW in soft-tissue in steps of 0.5 PPW. Simulations in 2D were carried out with both 250 kHz and 500 kHz source tonebursts with spatial discretisations ranging from 2 to 50 PPW for 500 kHz, and 2 to 101 PPW for 250 kHz, both in steps of 0.5 PPW. The forward simulations were run for the time taken for an acoustic wave to propagate across the grid three times, plus the duration of the source toneburst. The CFL was kept constant at 0.3 for all simulations, and the detailed parameters of the simulations are described in Appendix B.2.

The time-varying pressure signals at the transducer surface were recorded,

time-reversed, and used as source terms for a reversal simulation modelling the propagation of ultrasound from the transducer, through the skull and into the deep brain carried out at the highest spatial and temporal resolution feasible - the maximum values from the range used in the forward simulation. Due to the change in spatial and temporal discretization, it was necessary to interpolate the pressure signals recorded in the forward simulations onto the spatial and temporal grids used in the reversal simulations, shown in Fig. 3.2 (b). All reversal simulations were carried out using the k-space corrected PSTD scheme for increased numerical accuracy.

The reversal simulation was run for the duration of the forward simulations, plus an the additional time of five acoustic cycles in order to fully capture the reconstructed toneburst. The peak pressure occurring in a time window of twenty acoustic cycles centered on the expected time of refocusing was recorded across the brain volume. To determine the numerical convergence of the k-space scheme for the time-reversal, the focusing quality of the reversal simulations was calculated as a function of the spatial sampling of the forward simulations. The focusing metrics examined were the spatial peak temporal peak pressure across the brain volume, the distance of the spatial pressure peak from the location of the source point in the forward simulation, and the volume of the focus - quantified as the number of grid points with a peak temporal pressure value $\geq 50\%$ of the spatial peak temporal peak pressure across the brain volume.

In order to determine the convergence rate of the time-reversal protocol in the absence of an aberrating skull layer, the 3D convergence testing protocol was repeated using the same transducer and target geometry, but with a homogeneous grid with the properties of soft-tissue.

3.2 Results

3.2.1 Two-Dimensional Time-Reversal

The results of the 2D convergence testing are shown in Fig. 3.3. Spatial peak pressure and focus FWHM values are normalised relative to the values obtained from highest resolution forward simulation.

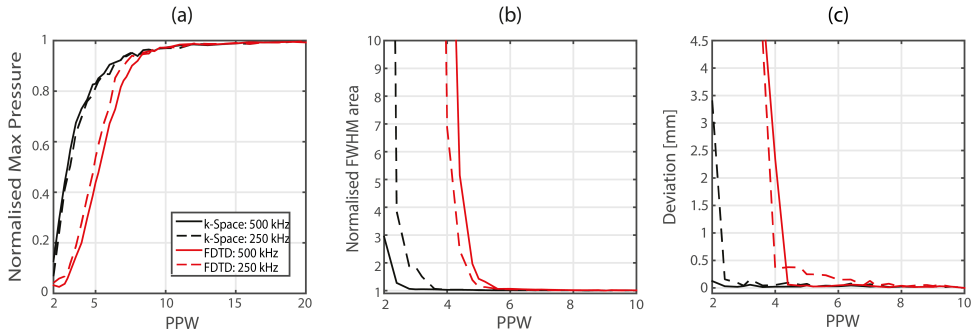

Figure 3.3: Results of 2D convergence testing. (a) Peak pressure recorded across the brain volume. (b) The deviation of the peak pressure from the location of the forward simulation source. (c) Normalised FWHM area of the focal spot size. Normalization is relative to results obtained with the most highly sampled forward simulation.

Refocusing quality in the reversal simulations increases with the spatial PPW of the simulated frequency in the forward simulations. Several key points can be derived from these results. Firstly, for all three metrics, the k-space scheme demonstrates convergence at approximately 2 PPW below the FDTD scheme. Secondly, although there is some difference in the position and size of the focus at very low sampling, both 250 kHz and 500 kHz demonstrate similar behavior as a function of spatial sampling. This indicates that these results can, to some extent, be generalised, and suggests that the reversal simulations have converged for both frequencies. Furthermore it can be used to validate the lower PPW values used out of necessity for 3D simulations, as all metrics are well converged by 16 PPW. The peak pressure at the focus has reached 99% of the maximum value, the error in the position of the focus is less than 0.03 mm, the error in the FWHM area is less than 1.5%, and the errors do not appear to reduce further with increased sampling. Finally, of the three refocusing metrics examined, normalised peak pressure across the brain volume (shown in Fig. 3.3 (a)) requires a higher spatial sampling to converge than either focal volume or the deviation of the focus from the target.

3.2.2 Three-Dimensional Time-Reversal

The results for 3D time-reversal convergence testing are shown in Fig. 3.4, while sections through the 3D simulated fields for different reversal simulations using

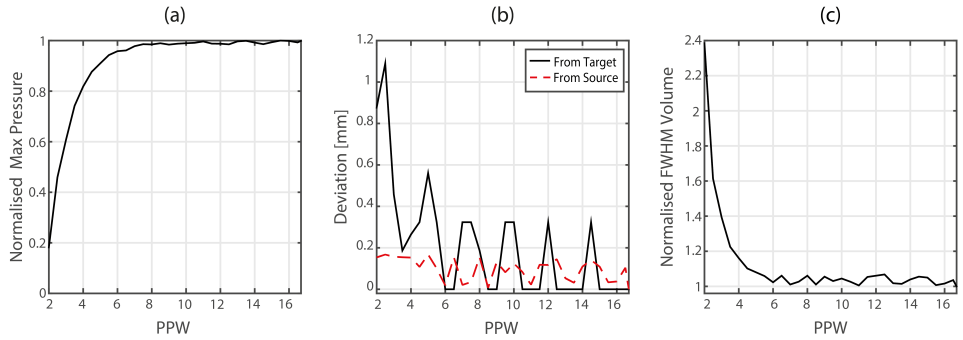

Figure 3.4: 3D convergence testing results. (a) Normalised peak pressure amplitude across the brain. (b) Deviation of pressure peak from both the parametrical defined target and the source used in the forward simulation. (c) Normalised half-maximum focal volume. Normalization is relative to results obtained with the most highly resolved simulation.

source terms derived from forward simulations with increasing spatial sampling are shown in Fig. 3.5. The results for normalised peak pressure amplitude in Fig. 3.4 (a) show similar trends to the 2D results. 6 PPW are required to obtain 95% reconstruction of pressure at the target and 10 PPW to attain convergence, with 99% peak pressure reconstruction and no improvement with increasing sampling. The convergence of the volume of the focal spot in Fig 3.4(c) shows similar behaviour, although it only requires 6 PPW to fully converge. Given the known impact of staircasing in 2D, the faster convergence here is likely due to a reduced staircasing error for the 3D geometry. Notably, the spatial targeting, calculated with reference to the position of the source in the forward simulation, does not improve with increasing spatial sampling, indicating that this error has fully converged by 2 PPW. For comparison, the distance of the reversal simulation from the parametrically defined target location is also shown. Again, this indicates that simulations with coarse spatial sampling may be sufficient when accurate spatial targeting is the only concern. However, difficulties can still arise in spatial targeting when the desired target coordinates do not correspond to a specific point on the spatial simulation grid.

Results for time-reversal using a homogeneous simulation domain are not shown, but demonstrated complete convergence of focus volume and position at 2 PPW, with no improvements observed with increasing spatial sampling, while the peak pressure amplitude was converged by 3.5 PPW. This supports the results

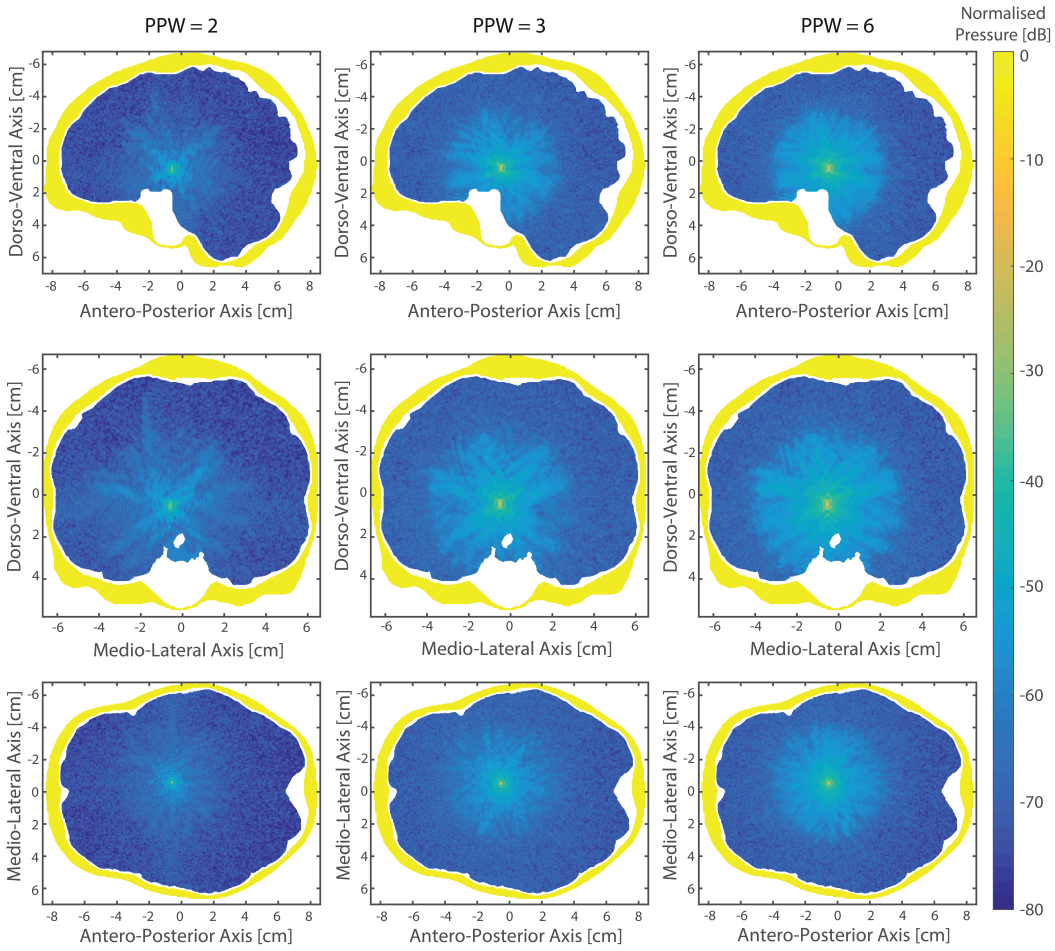

Figure 3.5: Sections through the pressure fields obtained from the convergence testing from reversal simulations, used source terms derived from forward simulations sampled at the indicated spatial PPW. Pressure is normalised to the peak spatial and temporal pressure recorded across the brain for all reversal simulations. The skull map is shown in yellow.

shown in Chapter 2 for the sampling required to prevent errors due to the PML and BLI, which are the main sources of error in homogeneous simulations.

3.3 Discussion

Convergence testing of a fully simulated time-reversal protocol using 2D and 3D head models was used to examine the impact of all numerical errors in concert. Testing in 2D was carried out to compare k-space corrected PSTD and FDTD model, to examine multiple frequencies, and to ensure that the numerical simulations have converged by the highest PPW value that is also feasible for simulations in 3D.

Testing in 2D for 250 kHz and 500 kHz ultrasound showed a faster rate of convergence for all focusing metrics for the k-space scheme when compared to FDTD, and generally close agreement in the convergence between the two frequencies. In addition, the error in the peak pressure amplitude at the focus showed slower convergence than both the positional error, and the volume of the focus. This is likely due to the most serious source of error, medium staircasing, which was shown to have a greater impact on the peak intensity amplitude of transcranially transmitted ultrasound, than the position of the peak.

Results in 3D showed similar trends to the 2D results for the convergence of the peak pressure amplitude. The focal spot size showed slightly slower convergence in the 3D case, while the positional error demonstrated no dependence on the sampling rate of the forward simulation once the variable positioning of the source point on the forward simulation grid was accounted for. This indicates that less stringent sampling may suffice for applications concerned only with the position of the focus, rather than the size of the focal spot and the exact amplitude at the target, such as high intensity focused ultrasound (HIFU) for ablation. When fine control over the pressure amplitude is required, stricter sampling may be necessary. Interestingly, despite the relatively severe error resulting from staircasing shown in Section 2.2.6, all three metrics of focusing quality were well converged at below 20 PPW. This discrepancy may be a due to the differences between the convergence testing protocol and the specific test used to examine staircasing across a bone layer, and suggests that the influence of staircasing is case specific.

Generally, these results confirm that when targeting accuracy is the prime concern, spatial sampling requirements are laxer than when a tight focal volume with a known peak pressure amplitude is required. This is in agreement with previous studies which have demonstrated that good spatial targeting of HIFU can be obtained via simulated TR using relatively coarse spatial sampling. Based on these spatial sampling required to achieve convergence in the 3D tests, all simulations of transcranial propagation using the k-space PSTD scheme should employ at least 10 spatial PPW to ensure numerical accuracy. While these simulations were carried

out using a CFL number of 0.3 and therefore a constant ratio of PPW to PPP, the 10 PPW value should be valid for alternate CFL values. This is because altering the temporal sampling should only affect the error due to numerical dispersion. Based on the results presented in Chapter 2, this should not lead to significant error for any stable CFL value while the simulated domain is sampled at 10 spatial PPW or more.

Chapter 4

Sensitivity of Intracranial Fields to Skull Acoustic Medium Properties

In the previous chapters, extensive testing of the numerical accuracy of the k-space corrected PSTD scheme determined the minimum spatial and temporal sampling requirements necessary to ensure numerical convergence for a frequency of interest simulated using the k-space PSTD scheme. As identified in the taxonomy of errors in Section 1.4, once numerical accuracy is assured, the next source of error in the use of numerical simulations is the representation of the acoustic medium within the simulation.

The level of fidelity required in acoustic property maps of the head is currently unclear [137] although, as discussed in Section 1.2, the acoustic properties of the skull are known to be of particular importance. Mueller et al. [138] recently examined the impact of modelling brain and meningeal geometry on simulated intracranial fields, concluding that it is not necessary to model detailed brain tissue geometry. They further concluded that modelling different soft tissues separately was of limited importance, although large changes in brain tissue sound speed and density (scaling by a factor of two or more) did have a significant effect. They also observed that changing skull geometry from flat to curved reduced peak RMS intensity by 40%, and generally affirmed that the skull was the major material influence on the intracranial field.

In the present chapter, the impact of the acoustic properties and geometry of

the bone is examined through a numerical sensitivity analysis. An idealised forward model of ultrasound transmission from a single element transducer through a skull bone layer to the brain is constructed. The sensitivity of the intracranial field to changes in the medium properties and geometric parameters of the bone layer is then established in terms of the change in the amplitude and location of the intracranial spatial pressure peak, and the size of the simulated focus.

4.1 Sensitivity Analysis Design

4.1.1 Simulation Setup

To carry out the numerical sensitivity analysis, a virtual model of transcranial ultrasound propagation was created. A single element transducer with a radius of curvature of 75 mm and a diameter of 60 mm was placed 27.3 mm from a curved layer of homogeneous bone tissue within the simulated environment. This was defined as a 6.5 mm thick layer taken from the surface of a sphere of 80 mm radius, in order to approximate the geometry of human skull bone [80]. The transducer geometry was chosen to match the HIFU transducer used in the experimental sonications in Chapter 5, and is similar to the low frequency single element transducers previously used in studies of neuromodulation in humans [53, 56]. The sensitivity analysis simulation layout is shown in Fig. 4.1.

To simplify the variation of acoustic properties and allow gross variation in the geometry of the layer in the sensitivity analysis, the bone map was modelled as a single homogeneous layer. Accordingly, the reference medium properties assigned to the bone layer were based on values obtained by White et al. [9] for whole skull compressional and shear wave transmission. Bone compressional sound speed was 2850 $\mathrm{m\,s^{-1}}$ and density was 1732 $\mathrm{kg\,m^{-3}}$. Absorption values were assigned using power law factors derived from this data, where $a_0 = 8.83\,\mathrm{dB\ MHz^{-y}\ cm^{-1}}$ and $y = 1.43$ [139]. When simulated, the reference sound speed of shear waves in the bone layer was set to 1500 $\mathrm{m\,s^{-1}}$, and the absorption power law prefactor was $a_0 = 19.5\,\mathrm{dB\ MHz^{-y}\ cm^{-1}}$ [9]. The remainder of the simulation domain was assigned the acoustic properties of water, with $c_\mathrm{w} = 1481\ \mathrm{m\,s^{-1}}$ and $\rho_\mathrm{w} = 1000$

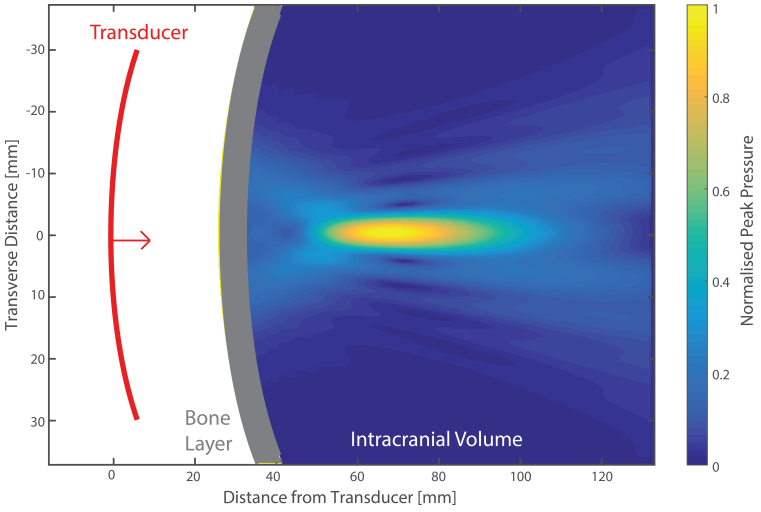

Figure 4.1: Section through the sensitivity analysis layout, demonstrating the relative positions of the transducer, idealised bone layer and the resulting intracranial focus. The peak pressure field shown is that obtained when the transducer is driven at 500 khz, with reference parameters assigned to the skull bone layer.

kg m^{-3}.

During simulations, the virtual transducer was excited with continuous wave (CW) signals at ultrasonic frequencies of 500 kHz, 750 kHz and 1 MHz. These correspond to a range of frequencies that are feasible for transcranial transmission, and have been demonstrated effective for UNMS [15]. Once the simulation had reached a steady state, the maximum pressure was recorded across the 3D field following transmission through the bone layer. This was used to evaluate the impact of changes to the bone layer on the intracranial field. Errors in the intracranial fields due to changes in the skull map were quantified relative to a reference simulation. Three error metrics were used: the change in the peak focal pressure, the change in the -6 dB volume of the pressure focus (calculated using the number of voxels with pressure amplitude above 50% of the peak pressure), and the spatial deviation of the centre of mass of the focal volume. As established, errors in peak pressure of less than 5% and error in focus position of less than 1.5 mm were used as criteria for accurate simulation. The importance of focal volume will vary according to target, frequency, and transducer characteristics, but an error of less than 20% was considered appropriate.

The simulations were performed using both the fluid and elastic codes within the k-Wave toolbox [133, 124] as laid out in Section 2.1 & Appendix A. The fluid code was used for the majority of sensitivity simulations, with the elastic code used to confirm that shear waves could be neglected, and to examine the impact of variation in shear speed and absorption. Only linear wave propagation was considered (even for high intensity applications, the wave field is generally linear at the skull surface [140]). Elastic wave simulations were conducted in 2D, while fluid simulations were conducted in 3D. Simulation parameters are described in Appendix B.3.

4.1.2 Testing of Variation in Layer Properties

4.1.2.1 Shear Mode Propagation

The simulation setup was first used to determine whether modelling shear wave propagation in the skull layer was required. Due to the increased computational burden of modelling elastic wave propagation, the 3D simulation setup was translated into 2D. Simulations were carried out using the k-Wave 2D elastic PSTD code for 500 kHz, 750 kHz and 1 MHz ultrasound [124]. Two reference simulations were carried out at each frequency: a fluid simulation with shear speed set to zero throughout the domain, and an elastic simulation with appropriate shear speed and absorption values assigned to the bone layer. In order to determine the sensitivity of the intracranial field to the elastic acoustic properties of the bone, additional simulations were carried out with the shear sound speed and absorption varied by $\pm 40\%$ in steps of 2%. Error was computed relative to the non-shear reference simulation.

4.1.2.2 Sensitivity to the Acoustic Properties of the Skull

The sensitivity of the intracranial field to changes in the acoustic properties of the skull was examined in three separate numerical experiments. First, the values of compressional sound speed, density, and absorption assigned to the skull layer were separately perturbed by a linear variation. The maximum variation assigned to sound speed and density was $\pm$ 60% based on the range of acoustic properties reported in the IT'IS material properties database (see Fig. 4.2). A wider range

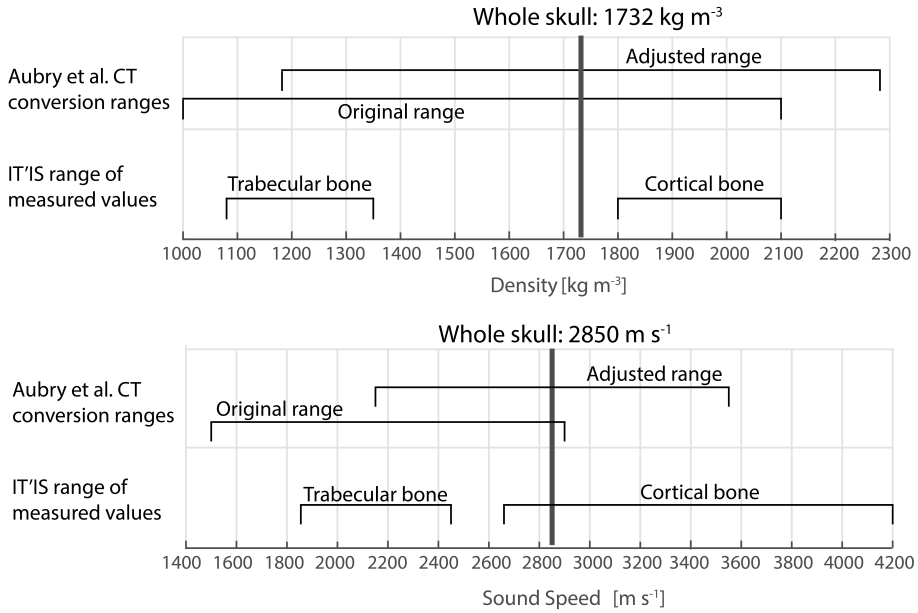

Figure 4.2: Whole skull values of sound speed and density used as reference values in the sensitivity analysis (taken from White et al. [9]). The range of measured values for bone taken from the IT'IS database [10] are shown for comparison. The original range of sound speed and density values used in the Aubry et al. CT conversion algorithm [6], and the adjusted values which yield the reference whole skull values for an HU value of 500, are also shown. Reproduced from Robertson et al. [116].

of $\pm$ 90% was tested for absorption, due to the difficulty in accurately measuring tissue absorption [86]. Acoustic properties were varied in steps of 1% for the range $\pm$10%, and in steps of 2% beyond that.

Second, the acoustic medium properties were individually perturbed by Gaussian noise with a coefficient of variation ranging from 1% to 20% in steps of 1%. The average property value assigned to the skull layer was kept constant. It should be noted that this noise was applied at the grid resolution (see Fig. 4.3(b)), and as such does not correspond to realistic heterogeneities in skull bone. Instead it represents an uncertainty in the assigned medium properties, such as the impact of noise in pixel intensity when properties are derived from medical images. Sound speed and density values were constrained between the minimum and maximum values for cortical and trabecular bone reported in the IT'IS tissue property database [10] (see Fig. 4.2). Absorption was constrained between zero and three times the refer-

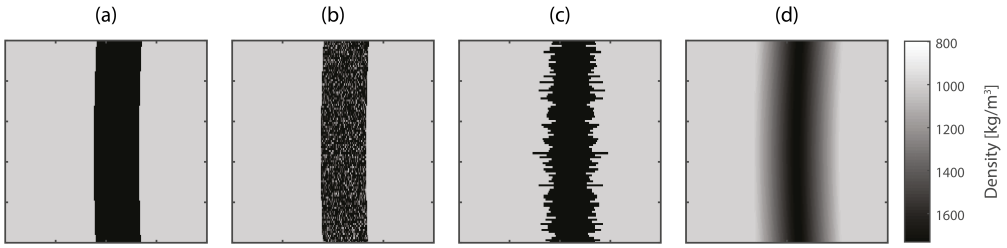

Figure 4.3: Selected density maps used in sensitivity analysis. (a) Reference bone layer. (b) Noisy variation in the medium properties. (c) Random variation in the skull thickness. (d) Smoothing of the bone layer with a cube smoothing kernel of 5.85 mm. Reproduced from Robertson et al. [116].

ence calculated using the power law parameters outlined in Section 4.1.1. Finally, a multi-dimensional sensitivity analysis was performed, with variations in HU used to generate coupled changes in density, sound speed, and absorption based on the Aubry et al. conversion algorithm [6] described in Section 1.2.4. For the present sensitivty analysis, the minimum and maximum values for each parameter were chosen such that Eqs. (1.1), (1.2) & (1.3) give the reference whole skull properties outlined in Sec. 4.1.1 for an HU value of 500. The adjusted interpolation ranges are shown in Fig. 4.2 alongside the original ranges used in Aubry et al. [6]. The value of HU was then perturbed to give coupled changes in density, sound speed, and absorption. Previous studies into variation in CT HU have described up to 25% HU error for densities of over 1100 kg m^{-3} [141], dependent on the calibration. Consequently, a range of $\pm$ 250 HU, in steps of 2 HU up to 42 HU and steps of 4 HU from there, was tested.

4.1.2.3 Sensitivity to the Geometry of the Skull

The sensitivity of the intracranial field to changes in the geometry of the bone layer was examined in three separate numerical experiments. First, the effect of systematic changes in the bone thickness, which can occur when segmenting CT or MRI images [142], was examined by increasing and decreasing the thickness of the bone layer up to $\pm$ 3.9 mm in steps of 0.15 mm. Second, the distribution of thickness across the length of the bone layer was varied randomly according to a Gaussian distribution with an increasing coefficient of variance, in steps of 1% up to a max-

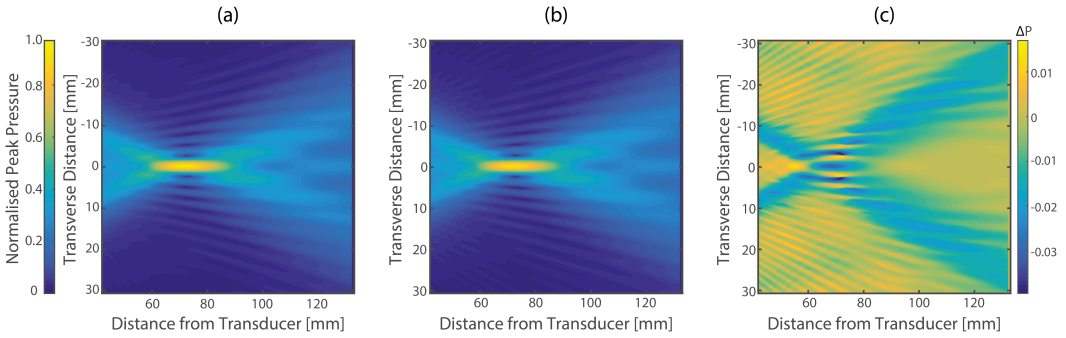

Figure 4.4: Results of elastic vs fluid comparison at 750 kHz showing the steady state maximum pressure. (a) Fluid simulation. (b) Elastic simulation. (c) Difference plot. Reproduced from Robertson et al. [116].

imum of 20%. This was done to examine the potential influence of uncertainties in the exact position of boundaries on transcranial propagation, as shown in Fig. 4.3 (c). The random thickness of the bone layer was constrained between 1-12 mm based on previously reported maximum and minimum values for skull thickness [80]. Third, the potential impact of CT partial volume effects was examined via spatial smoothing of the skull layer. Smoothing was achieved by applying a square smoothing kernel of increasing size, in steps of 0.15 mm up to a maximum of 5.85 mm, to maps of acoustic properties, as a simulacrum of CT pixel size. An example is shown in Fig. 4.3 (d).

4.2 Sensitivity Analysis Results

4.2.1 Shear Mode Propagation & Variation in Shear Properties

The change in the peak pressure field due to including or excluding shear waves in the bone layer for a source frequency of 750 kHz is shown in Fig. 4.4, which was the frequency where the greatest impact was observed. The maximum change in peak intracranial pressure was +0.09% for 500 kHz, -2.81% for 750 kHz, and -0.80% for 1 MHz, respectively. Changes in the position of the centre of mass of the -6 dB focus area all occurred along the beam axis, and were all less than 0.3 mm. The change in the -6 dB focal area was less than 1.5% for all frequencies. An L2 error norm was also calculated across the peak pressure field between elastic and fluid simulations for each frequency. The maximum L2 error was 4.94% for the

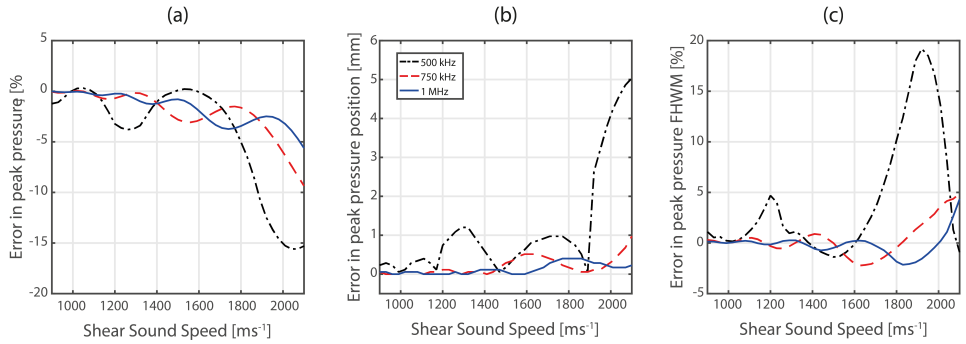

Figure 4.5: Errors in the intracranial field resulting from assigning different shear speed values to the bone layer, computed relative to a reference simulation without shear propagation modelled. (a) Percentage error in peak pressure. (b) Error in focus position. (c) Percentage error in focal volume.

750 kHz frequency, which also demonstrated the maximum error in peak pressure amplitude and position. L2 errors were 3.51% and 1.77% for 500 kHz and 1 MHz ultrasound, respectively. These results serve as a justification for neglecting shear in the 3D numerical sensitivity analysis. However, it should be noted that the limited impact of shear mode propagation may not hold for geometries involving more oblique angles of incidence. While these results indicate that shear mode propagation in the skull can be neglected for the reference shear speed and absorption values simulated, a sensitivity analysis was performed by varying the values of shear speed and absorption, to determine if different values for these parameters have a greater effect on the intracranial field. Results for variation in shear sound speed are shown in Fig. 4.5. The effect of accounting for shear absorption, and of variation in shear absorption, was negligible. Based on the results shown, it appears that changes in the field caused by the introduction of shear mode propagation begin to increase in magnitude at higher shear speeds, and lead to significant errors once shear speed in the layer is higher than 1800 m s^{-1}.

4.2.2 Linear Variation in Acoustic Properties

The errors in the intracranial field resulting from linear variation in acoustic properties are shown in Fig. 4.6. Several conclusions can be drawn from these results. Overestimation of the sound speed, density and absorption all lead to a correspond-

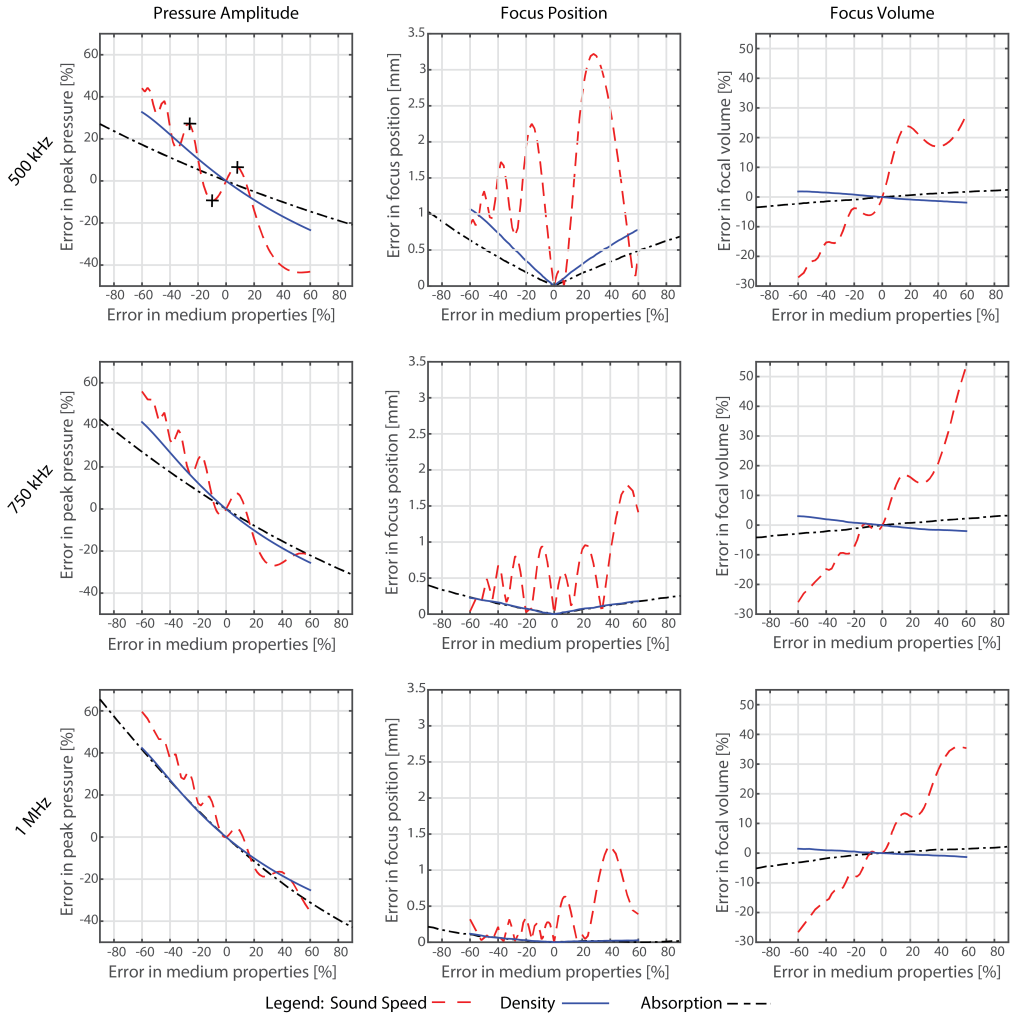

Figure 4.6: Errors in the intracranial field resulting from a linear variation of the acoustic medium properties. Rows correspond to frequencies, columns to the error metrics evaluated. Points marked on the 500 kHz pressure amplitude plot correspond to the field plots shown in Fig. 4.7. Reproduced from Robertson et al. [116].

ing decrease in the simulated transcranial pressure, and vice versa. This is due, at least in part, to the increased acoustic impedance (a product of sound speed and density) and absorption, resulting in increased attenuation. To obtain an error in peak pressure of less than 5% for all frequencies, sound speed and density must be assigned with a change of less than 4% and 9%, respectively. To obtain less than a 5% error in peak pressure, absorption must be assigned with a change of less than 9% at 1 MHz, increasing to 20% at 500 kHz. This is due to the increase in

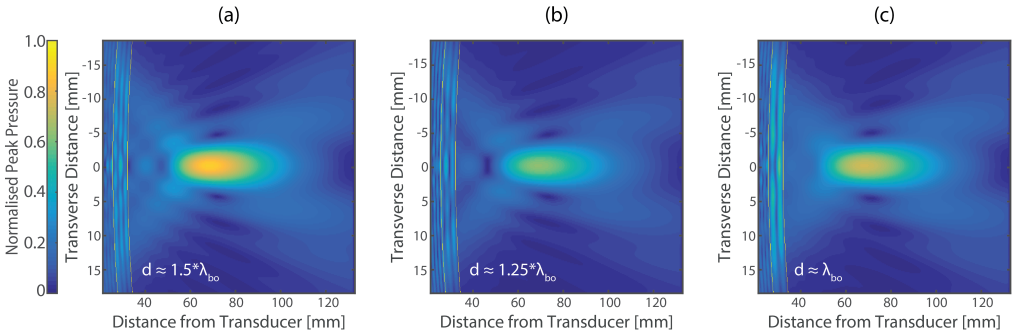

Figure 4.7: Slices through the 500 kHz peak pressure field at increasing sound speeds, where d is the thickness of the bone and λ_{bo} is the acoustic wavelength. (a) Field when sound speed in the bone layer is 2109 m s^{-1} (d $\approx 1.5\lambda_{bo}$). (b) Field when the sound speed is 2565 m s^{-1} (d $\approx 1.25\lambda_{bo}$). (c) Field when the sound speed is 3050 m s^{-1} (d $\approx \lambda_{bo}$). Reproduced from Robertson et al. [116].

absorption with frequency.

The position and volume of the focus also depend on all three acoustic properties. However, changes in density and absorption do not cause errors of larger than 1.5 mm in position or 5% in focal volume. To obtain less than 20% error in focal volume across all frequencies, sound speed must be defined with less than 11% variation. To obtain an error in focus position of less than 1.5 mm at 500 kHz, sound speed must be defined with a variation of less than 9%. As the sound speed changes further, the position of the focus can deviate up to 3 mm from the reference position, but does not vary steadily. The focus position oscillates around a fixed point for all frequencies, with the amplitude of the oscillation and the dependence on sound speed informed by the ultrasonic frequency. Error in focus position is highest at 500 kHz, due to the longer wavelength and the distance of the reference focus from the fixed point around which the focus oscillates. Errors in peak pressure and focus volume also oscillate with changing sound speed. An example of the oscillating peak pressure is shown in Fig. 4.7, with the bone layer highlighted. These plots correspond to maxima and minima in the oscillating pressure pattern (indicated by the plus symbols, +, in Fig. 4.6), and occur when the thickness of the bone layer corresponds to particular multiples of the wavelength of 500 kHz ultrasound in the layer. This is supported by a similar relationship between transcranial transmission and ultrasound frequency previously described by White et al.

[143], which would also lead to wavelength related interference in the bone layer. This relationship to wavelength does not hold for the oscillations in focus position in the same way, which is likely due to the lensing behaviour of the bone layer changing with sound speed instead. These results show that, for the skull geometry simulated, small systematic variation in the assigned medium properties can result in significant changes in the size and position of the focus. They also indicate that accurate assignation of sound speed is of primary importance when modelling transcranial propagation. Changes in the assigned absorption and density can lead to comparable errors to sound speed in the simulated peak intracranial pressure.

4.2.3 Noisy Variation in Acoustic Properties

The errors in the intracranial field resulting from noisy perturbation of the acoustic properties of the skull layer are shown in Fig. 4.8. Several conclusions can be drawn from these results. First, the errors in the field resulting from noisy variation in medium properties are smaller than those from linear changes. Errors in peak pressure and focus volume remain below 5%, while the error in focus position remains below 1 mm. Second, the introduction of noise into the map of acoustic absorption (while keeping the average absorption constant) has almost no effect on the intracranial field. This would appear to vindicate previous studies where a homogeneous map of absorption has been used alongside fully heterogeneous maps of sound speed and density [15]. Finally, errors in the intracranial field are generally less severe at higher frequencies. This is most apparent for the error in focus volume, but can also be observed across the other error metrics.

4.2.4 Variation in CT Hounsfield Units

The errors in the intracranial field resulting from variation in the CT HU used to derive the acoustic properties are shown in Fig. 4.9. Across the range of HU tested, the maximum error in peak pressure was an increase of 17% for 1 MHz ultrasound. To obtain less than 5% error in peak pressure, a change of less than 60 HU is required. For 500 kHz, a variation of no greater than 180 HU results in a positional error of less than 1.5 mm, while the error remains below 1 mm for higher frequencies. The

focal volume increases at higher HU values, with a variation of 200 HU causing an error in focal volume of 20%, while reducing HU appears to have a lesser effect on focal volume. Errors in the intracranial field oscillate, in line with the linear changes in sound speed shown in Fig. 4.6. Within the range of HU tested, the maximum errors across all metrics are lower than the errors shown in Fig. 4.6 for individual changes in assigned medium properties. This suggests that the primary concern should be the robustness of any method for the conversion of CT images to density, absorption and especially sound speed, with the accuracy of HU measurements be-

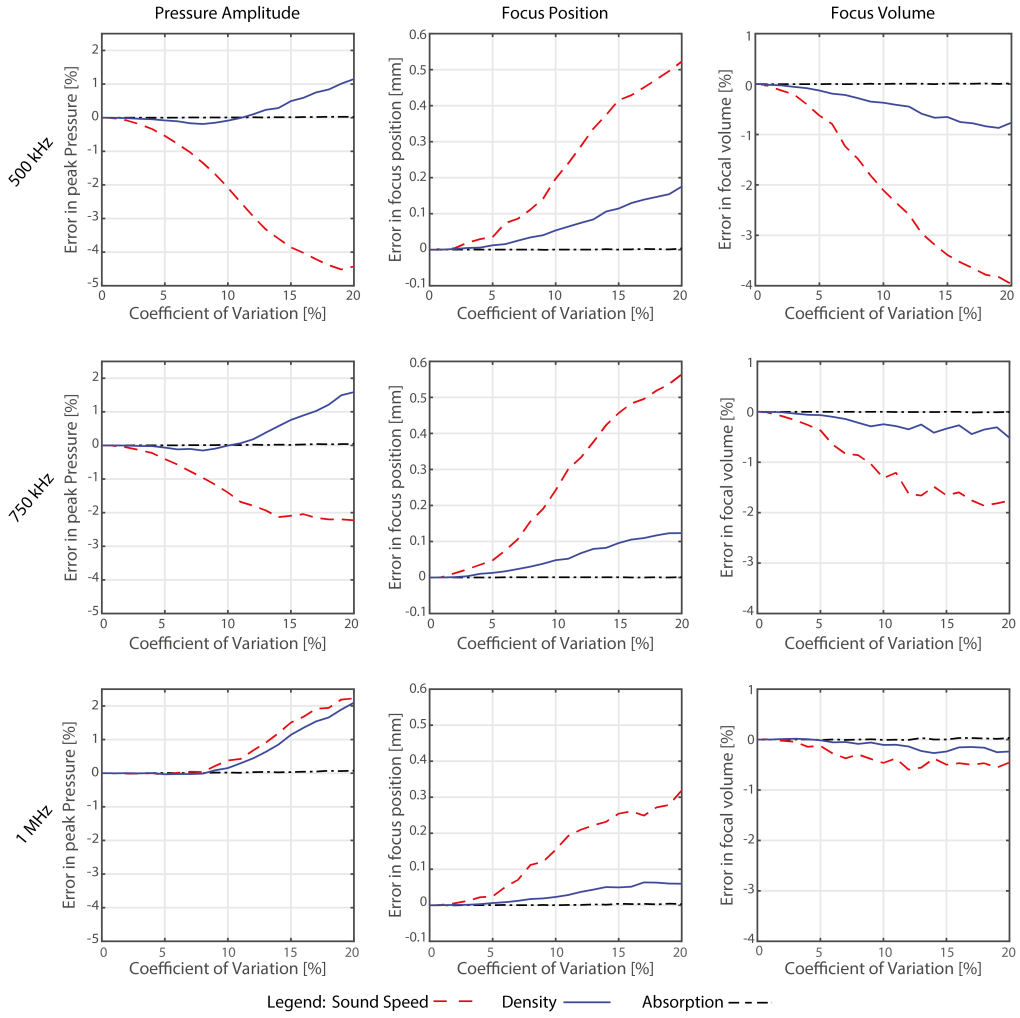

Figure 4.8: Errors in field resulting from noisy perturbation of acoustic properties with an increasing coefficient of variation. Rows correspond to frequency tested and columns correspond to error metric evaluated. Reproduced from Robertson et al. [116].

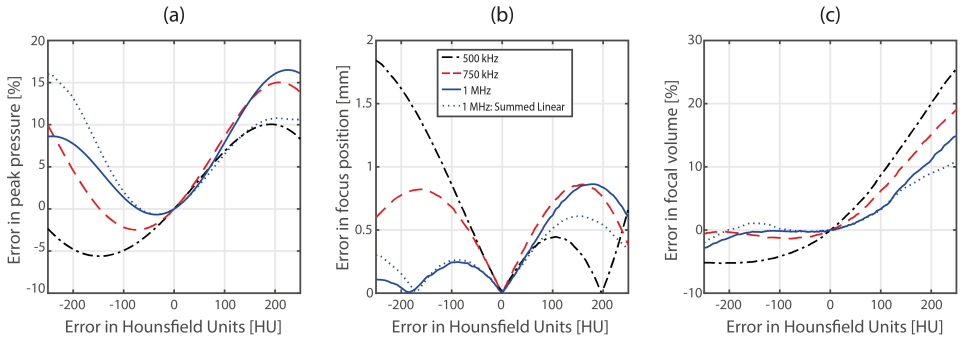

Figure 4.9: Errors in the intracranial field resulting from variation in Hounsfield Units used to derive the bone layer medium properties, and a comparison with the sum of linear changes shown in Fig. 4.6 for 1 MHz. (a) Percentage error in peak pressure. (b) Error in focus position. (c) Percentage error in focal volume. Reproduced from Robertson et al. [116].

ing a secondary concern. Figure 4.9 also compares the error resulting from the sum of equivalent individual linear errors (from Fig. 4.6) for 1 MHz ultrasound. While they match well for low HU variation, when reducing HU the sum of linear errors is less than the error resulting from the simultaneous alteration of all properties, and vice versa.

4.2.5 Linear Variation in Thickness

The errors in the intracranial field resulting from linear variations in bone layer thickness are shown in Fig. 4.10. A change in layer thickness of 0.1 mm is sufficient to cause a peak pressure error of 5% for all frequencies, while a change in skull thickness of $\pm$ 1 mm can lead to an error of up to $\pm$ 20%, depending on frequency. Beyond this level, underestimation of the layer thickness has a greater effect on the peak intracranial pressure than overestimation. The position of the focus changes with thickness up to maximum errors of 2.5 mm, 1 mm, and 0.5 mm for 500 kHz, 750 kHz, and 1 MHz, respectively. A change in thickness of less than 0.8 mm is required to obtain a positional error of less than 1.5 mm for 500 kHz. The error in focal volume shows a general trend of increasing with larger layer thickness and vice versa, but did not exceed $\pm$ 20% across the range of thicknesses tested. Oscillations in the error metrics can be observed, similar to the oscillations seen for changes in sound speed in Fig. 4.6. These are related to wavelength, as can be seen

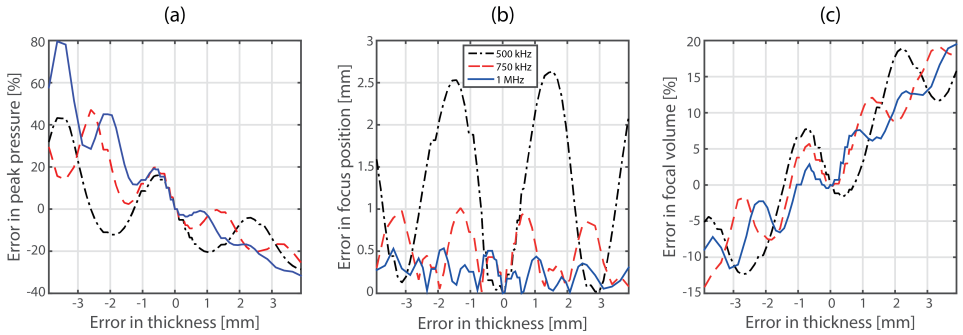

Figure 4.10: Errors in field resulting from systematic over/under estimation of the bone layer thickness, for a reference thickness of 6.5 mm. (a) Percentage error in peak intracranial pressure. (b) Error in the position of the focus. (c) Percentage error in focus volume. Reproduced from Robertson et al. [116].

by the relative periods of the oscillations in Fig. 4.10 (b). This supports previous conjecture regarding the influence of the wavelength of applied ultrasound relative to the thickness of the layer, and the lensing effects of the layer.

4.2.6 Noisy Variation in Thickness

The errors in the intracranial field resulting from random perturbations to the skull layer thickness are shown in Fig. 4.11. Introducing a noisy variation in the thickness can have a large effect on the peak intracranial pressure, with an overestimation of peak pressure with increasing perturbation of the thickness. Coefficients of variation in thickness of 1%, 2% and 3% cause an error in peak pressure of 5% for 500 kHz, 750 kHz and 1 MHz, respectively, rising to a maximum error of around 50%. It is not clear why this should be the case, as scattering could be expected to reduce focused transmission. It is possible that the random scattering reduces existing destructive interference effects, or that the random perturbation alters the scattering profile to increase the forward transmission of ultrasound. Regardless, these results do suggest that uncertainty in the profile of the bone layer, and modelling the layer as closer to or further from a smooth layer of even thickness, can alter the predicted peak intracranial pressure. The error in focus position was below 0.5 mm for 1 MHz and 750 kHz, but increased to over 1.4 mm for 500 kHz ultrasound. The increased impact at higher frequencies is likely due to the larger scale of the changes in thickness relative to the wavelength of the ultrasound. The error in focal volume

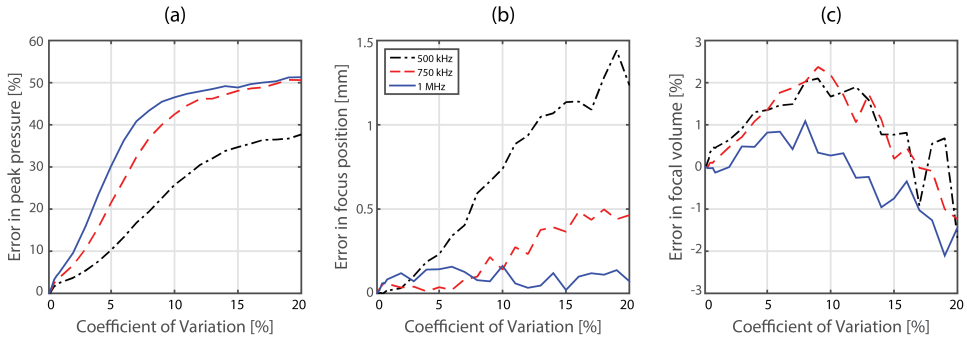

Figure 4.11: Errors in field resulting from noisy perturbation of the bone layer thickness at multiple frequencies. (a) Percentage error in peak pressure. (b) Error in focus position. (c) Percentage error in focal volume. Reproduced from Robertson et al. [116].

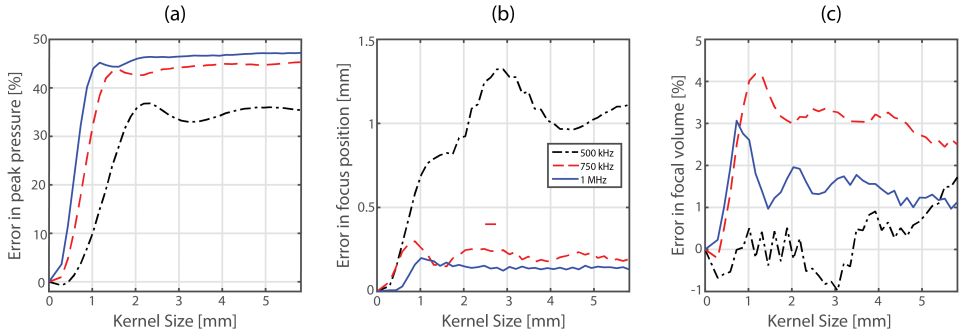

Figure 4.12: Errors in field resulting from kernel smoothing of the bone layer at multiple frequencies. a) Percentage error in peak pressure. (b) Error in focus position. (c) Percentage error in focal volume. Reproduced from Robertson et al. [116].

is low, not rising above 3% across the frequencies and range of errors tested.

4.2.7 Smoothing of the Bone Layer

The errors in the intracranial field resulting from smoothing the map of medium properties are shown in Fig. 4.12. There is a clear overestimation of transcranial transmission dependent on the level of smoothing and the frequency of the applied ultrasound, with higher frequencies more affected. Kernels of 0.35 mm, 0.45 mm, and 0.8 mm cause an error in peak pressure of greater than 5% for 500 kHz, 750 kHz, and 1 MHz ultrasound, respectively. For a kernel size of 1 mm the peak transcranial pressure is increased by over 50% for 1 MHz ultrasound. Given that the scale of modern CT imaging systems falls well within this range, this implies that partial volume smoothing of CT images could potentially have a dramatic effect

on simulated transcranial fields if the smoothing of the CT images is translated into the map of acoustic medium properties. There is some error in the position of the focus with increasing smoothing, although it only rises above 1 mm for 500 kHz ultrasound. The error in the focal volume is small, never rising above 5%.

4.3 Discussion

The impact of the accurate mapping of acoustic medium properties on the intracranial field was examined via an idealised model of transcranial transmission. A sensitivity analysis was carried out, examining the impact of systematic changes and random uncertainty in the assigned medium properties, changes in the geometry of the layer, and smoothing of the bone layer. Before the sensitivity analysis, the impact of modelling shear mode propagation was examined, and determined to have a negligible impact for shear speed values of $<$1800 m s^{-1}. The results of the numerical sensitivity analysis are summarised in Table 4.1.

In terms of the assignation of acoustic medium properties there are several relevant observations. The peak intracranial pressure is affected by the values of sound speed, density, and absorption. The impact is similar for each medium property, and the effect is larger at higher frequencies. The position and volume of the focus is most dramatically affected by changes in sound speed, although changes in density and absorption will have a minor effect. When altering sound speed, oscillations in each error metric were observed. This is due to the pattern of wavelength interference within the bone tissue layer, but is also influenced by changes in diffraction through the layer, and the resulting interference pattern in the intracranial field. To obtain less than 5% error in peak pressure, less than 1.5 mm error in focus position, and less than 20% error in focal volume, sound speed must be defined with less than 4% change, or 114 m s^{-1}, and density must be defined with less than 9% change, or 156 kg m^{-3}. Absorption has a significant effect on intracranial pressure only, and a change in the absorption coefficient of less than 9% is required to obtain less than 5% error in peak pressure for the highest frequency of 1 MHz. Random, noisy uncertainty in medium properties has a lesser effect on the intracranial field,

Frequency [kHz]	c [%]	ρ [%]	α [%]	HU	Thickness [mm]	Kernel [mm]
500	4/9/12	10/×/×	18/×/×	71/150/200	0.10/0.74/×	0.35/×/×
750	4/46/39	9/×/×	12/×/×	62/×/×	0.10/×/×	0.45/×/×
1000	4/×/33	8/×/×	7/×/×	71/×/×	0.12/×/×	0.80/x/×

Table 4.1: Maximum permissible errors in the acoustic medium properties and geometry of the skull needed ensure less than 5% error in peak pressure / 1.5 mm in focus position / 20% error in focal volume. An × symbol indicates that these error criterion were not exceeded for the range of variation tested.

and has almost no effect for uncertainty in absorption, indicating that the use of homogeneous absorption maps based on average absorption or transmission loss is justified. The impact of combined changes in medium assignment was also examined, with reference to a widely used algorithm for the conversion of HU to density, sound speed, and absorption. Across the range of typical uncertainty in HU tested, the overall errors in the intracranial field were less than observed for expected variations in individual medium properties. This strongly suggests that uncertainty in the HU of a CT image used to create a map of medium properties is less important than the robustness of the conversion algorithm. For variation in medium geometry, systematic changes in layer thickness and smoothing on the scale of clinical CT resolution (0.5-1 mm) result in significant changes in simulated peak intracranial pressure. A systematic change in thickness of 0.1 mm, or a smoothing kernel of 0.35-0.8 mm is sufficient to cause an error in peak pressure of 5%. Focal volume and position are also affected, especially by systematic change in layer thickness. Noisy uncertainty in the layer thickness also affected peak intracranial pressure, with a coefficient of variation of 1% enough to increase the peak pressure by 5% for 1 MHz ultrasound. Consistently, the properties of the skull layer have a more significant effect on the peak pressure amplitude than on focus position or volume. This indicates that accurate maps of acoustic properties are likely to be less important for ablative applications, where fine control over the pressure is secondary to accurate targeting.

A potential limitation of this work is the extent to which the results are generalisable. The sensitivity analysis was carried out with a single element transducer

with no aberration correction, and a specific skull geometry designed to approximate the varied shaped of the skull. A homogeneous model with aggregate skull properties was used to simplify the introduction of different errors into the layer. The impact of variation in nonlinearity parameter was not examined, which may be relevant for high intensity applications. Different skull geometries may have different specific sensitivities to the errors examined in the numerical sensitivity analysis. Different transducer configurations may also influence the results. However, it is expected that the general scale of the resulting errors in the field, and the relative importance of different medium properties and aspects of medium geometry will be maintained. Moreover, as the transcranial field used in the sensitivity analysis was not strongly aberrated, the results should also extend to fields in which aberration correction is applied.

Chapter 5

Registered Simulations Through Ultrasonic Bone Phantoms

In the previous chapters, the criteria required to ensure numerical accuracy in simulations of transcranial ultrasound propagation, and the fidelity necessary in maps of skull acoustic properties were examined in a wide range of numerical experiments. However, to confirm the potential for highly accurate transcranial simulations using the k-space PSTD method, validation of the accuracy of simulations through practical experiments is necessary.

In the present chapter, multiple acoustic skull phantoms are designed for use in practical experiments examining the transmission of ultrasound through the phantoms. The resulting experimental pressure fields are then compared with numerical simulations spatially registered to the experimental setup in order to determine simulation accuracy. Computer aided design (CAD) techniques are used in combination with optical mounting equipment to allow highly accurate spatial registration of the experimental setup with a simulated environment. Bone phantoms are created from two materials based on the previously reported proximity of their acoustic property values to those of skull bone, and their amenability to being produced from CAD generated models. The first is a 3D printing material, VeroBlack (Stratasys, Eden Prairie, MN), the acoustic properties of which were recently characterised by Nikitichev et al. [144]. The second is Araldite 1302 (Farnell, Leeds, United Kingdom) resin, first characterised by Truscott et al. [145], which has previously

been used to develop bone phantoms [146]. The acoustic properties of the Araldite 1302 material are measured for use as inputs to numerical simulations, including an estimation of uncertainty in the measured values.

A series of practical sonication experiments examining the propagation of 1 MHz ultrasound from a focused transducer through the skull bone phantoms are carried out. Rigid stereotactic positioning is used to allow the translation of the experimental domain into a simulated environment to carry out registered simulations with the k-space PSTD scheme. The pressure fields observed in the practical sonication experiments are then compared with those predicted by the registered simulations, to determine the level of simulation accuracy in the context of the estimated uncertainty in the maps of acoustic properties.

5.1 Design of an Ultrasonic Bone Phantom

5.1.1 Computer Aided Design

Computer aided design, in combination with 3D printing, laser cutting and resin casting techniques, allow 3D models with user defined geometry to be directly converted to physical components. This allows highly accurate matching between practical experiments and numerical simulation environments. Here, 3D models were created in FreeCAD (Juergen Riegel, Werner Mayer, and Yorik van Havre; Open-Source, www.freecad.com) software, a freely available CAD software. A range of the 3D designs produced are shown in Fig. 5.1, rendered in the FreeCAD environment alongside finished components. Two skull phantoms were designed for the experimental validation using rapid prototyping and casting techniques. The first phantom was designed to match the idealised skull geometry used in the numerical sensitivity analysis in Chapter 4, and will be referred to as the “parametric” phantom. As in the sensitivity analysis, this was designed as a 6.5 mm thick layer taken from the surface of a sphere of 80 mm radius. This curved piece had a diameter of 106 mm, and was appended with four square mounting blocks to allow it to be locked into a mounting system.

The second phantom was designed to match the anatomical geometry of the

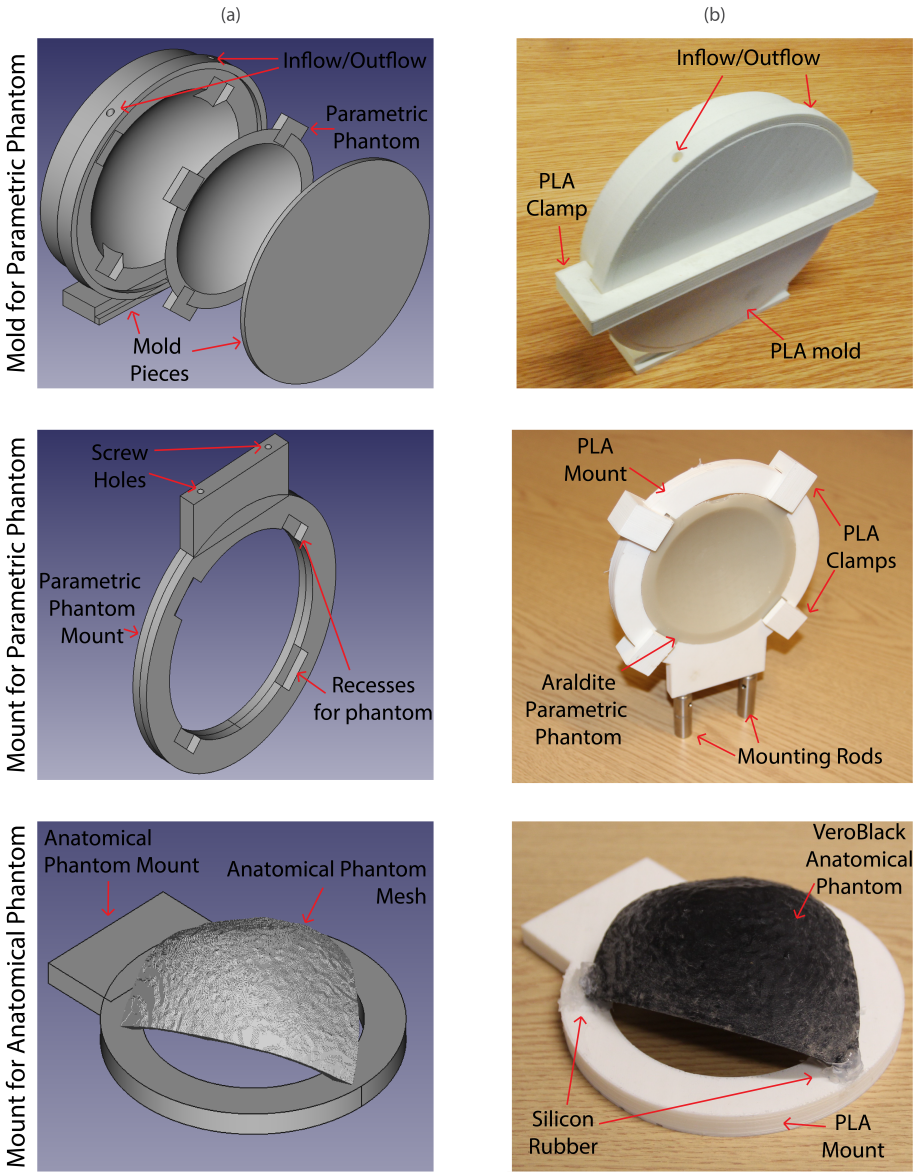

Figure 5.1: Computer aided design of phantoms and mounting components. (a) Rendering of molds and mounts designed around the intended phantom geometry in the FreeCAD environment. (b) Finished prints in PLA and VeroBlack.

skull, and was based on a portion of the 3D mesh of the human skull used in Chapter 3 for convergence testing, roughly corresponding to the posterior portion of the parietal bone. This was derived from a T1-weighted MR image obtained from the Imperial College brain development dataset. Brain and skull volumes were extracted using the FSL MRI processing toolbox and converted into a surface mesh using the iso2mesh toolbox [136].

Several additional 3D models were created using CAD techniques, also shown in Fig. 5.1. Firstly, in order to allow the casting of the parametric phantom in the Araldite 1302 material, a mold was designed in FreeCAD around the geometry of the parametric phantom. The mold was designed as two seperate pieces held together by a clamp, with two vent holes to allow resin to be pumped in and air to escape. Secondly, in order to allow these phantoms to be registered sterotactically during practical experiments so that their positions could be accurately translated into 3D simulations, holders were designed for use with an optical breadboard mounting system. Holders for both the parametric and anatomical phantoms were designed as circular pieces. The phantoms would attach to the edges of the circle, and the ultrasound beam would pass through the annular centre of the mount. The holder for the parametric phantom was designed entirely in FreeCAD, with recesses corresponding to the mounting blocks appended to the phantom, and were held in with clamps. A mounting system for the anatomical phantom was designed in the FreeCAD software, and then exported into the Blender 3D design software (Stichting Blender Foundation, Amsterdam, the Netherlands; OpenSource www.blender.org/) alongside the skull mesh. Blender allows operations on multiple large meshes, and a boolean subtraction operation was carried out between the corners of the bone mesh, and the circular mount. This allowed the phantom to be locked into the mounting system and secured with silicon gel. Finished models for 3D printing were exported from the FreeCAD environment as stereolithography files, and checked for errors using the NetFabb software (Autodesk, San Rafael, CA). Mounts to allow stereotactic positioning of the transducer were designed in Adobe Illustrator (Adobe, San Jose, CA), for production with a laser cutter.

5.1.2 3D Printing and Resin Casting

Phantoms and mounting systems were created from the CAD models using 3D printing, resin casting and laser cutting techniques. The mold for the parametric phantom and the mounting systems for both phantoms were printed on an Ultimaker 2.0 3D printer (Ultimaker, Chorley Lancashire, UK), in PolyLactic Acid (PLA)(PolyMax; Polymakr, Changshu, China). 3D printing material. The Ulti-

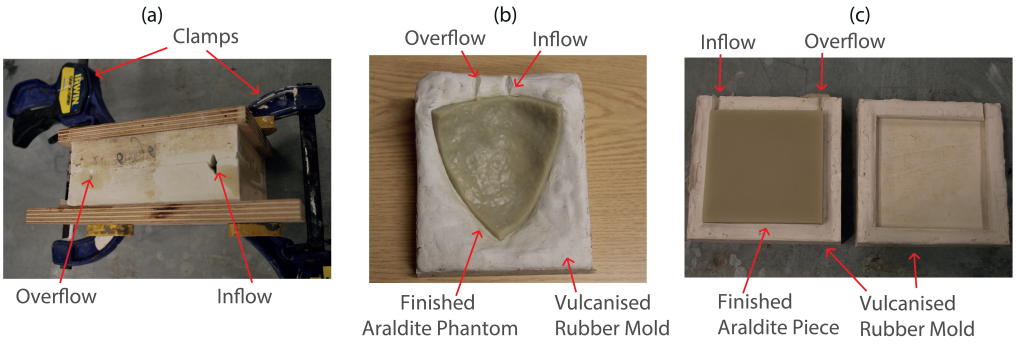

Figure 5.2: Casting of Araldite 1302 resin in vulcanised rubber molds. (a) Mold clamped durin casting process, with visible overflow. (b) Araldite phantom following casting. (c) Cuboid Araldite piece for medium property measurements.

maker is a filament deposition printer, that creates prints by iterative laying down of a filament of liquid material, which hardens into a finished print [147]. The resulting trapped air limits the usefulness of these prints as ultrasound phantoms, but the speed and low cost make them useful for molds and mounting systems.

Mounts for the phantoms were printed in PLA, while mounts for ultrasound transducers and the cuboid araldite samples were laser cut from 5 mm PMA perspex on a PLS6.75 laser cutter (Universal Laser Systems, Scottsdale, AZ). Mounts were designed with integrated holes, that were tapped to allow optical mounting equipment to be screwed into them. The parametric phantom model was produced in the Araldite 1302 material only, using the PLA mold shown in Fig. 5.1. The internal surface of the PLA mold was coated in PVA mold release agent (Easy-Composites, Stoke on Trent, United Kingdom) which was left to dry. The Araldite was prepared by mixing the resin and hardener in the required proportions as specified by the manufacturer, before being degassed in a vacuum chamber. The mold was sealed and clamped shut, and the Aralidite mixture pumped in. The mold was left for 48 hours to allow the resin to cure, before removing the finished phantom piece from the mold.

The anatomical phantom was produced in both VeroBlack, and the Araldite 1302 resin. The VeroBlack phantom was printed using an Objet30 polyjet printer (Stratasys, as above), and can be seen in Fig. 5.1. Polyjet printing works on a similar principle to 2D inkjet printing, wherein liquid 3D printing material is iteratively

sprayed onto a print, and then cured with UV light [147]. Following cleaning and removal of support material, this was used as a positive to create a vulcanised rubber mold, shown in Fig. 5.2 alongside the finished Araldite phantom.

VeroBlack and Araldite samples were also created for acoustic property measurements. A single cuboid VeroBlack sample with a thickness of 6 mm was printed, while two cuboid Araldite samples were cast in rubber molds. The casting process of these samples is shown in Fig. 5.2. The two Araldite samples were then milled to thicknesses of 9.40 mm and 14.05 mm, measured using a micrometer (Mitutoyo, Kawasaki, Japan). The milling process ensured that the faces of the samples were flat and parallel. Araldite samples and resin phantoms were all cast in separate batches due to limitations on the amount of resin mixture that could be degassed at one time.

5.1.3 Acoustic Property Measurements

While previous studies have given values for the acoustic properties of Araldite 1302, the sound speed and absorption at 1 MHz were measured to ensure the maximum possible accuracy in the registered simulations of propagation through phantoms. The uncertainty in the medium property measurements was estimated, to allow comparison between the results of the sensitivity analysis in Chapter 4 and the error in the registered simulations. The density of the mixture is known from the manufacturer information. As discussed above, the fluid acoustic properties of the VeroBlack 3D printing material were recently characterised by Nikitichev et al. [144] and are not examined here.

Measurements were made using the transmission-substitution method, described in detail by Zequiri et al. [85], which will be reviewed here. First, ultrasound transmission from a transducer in a water bath is measured, usually in the far field. A material sample of known thickness is then inserted into the field, perpendicular to the field axis. The change in the amplitude and time-of-flight of transmitted ultrasound is measured. Sound speed can then be calculated directly based on the

change in the ultrasound time-of-flight

$$\frac{1}{c_t} = \frac{1}{c_w} - \frac{\Delta t}{\Delta x} \quad . \tag{5.1}$$

Here c_t is the sound speed of the sample, c_w is the speed of sound in the background medium, Δx is the thickness of the sample and Δt is the shift in time-of-flight.

The absorption coefficient can be calculated from the change in amplitude when the sample is introduced. However, it requires an understanding of the distinction between attenuation and absorption of the ultrasound wave, described in Section 1.2.3. A beam propagating through a single sample can only be used to calculate acoustic attenuation in the form of transmission loss. To calculate the attenuation coefficient of the material independent of reflective losses, two samples of different thicknesses are required, and the transmission loss in dB calculated for each, where

$$\text{Transmission Loss} = -20\log_{10}\frac{p_w}{p_t} \quad . \tag{5.2}$$

Here p_t is the pressure following transmission through the sample, and p_w is the pressure following transmission through the background medium (i.e. in the free field). The change in transmission loss divided by the change in thickness between the two samples then gives the absorption coefficient, assuming no scattering within the sample. To characterise the acoustic attenuation and compressional sound speed of Araldite, transmission measurements were conducted in a scanning tank. The layout is shown in Fig. 5.3 (a). Acoustic waves were generated with a plane piston transducer with an active diameter of 40 mm (Olympus, Tokyo, Japan), excited with a single cycle toneburst at its central frequency of 1 MHz. This transducer was mounted on a five axis positioning arm to allow the field to be scanned. Time varying pressure signals were recorded with a calibrated 0.5 mm needle hydrophone (Precision Acoustics, Dorchester, United Kingdom). Measurements were repeated at 0.5 mm intervals along the transducer beam axis between 420 mm to 580 mm from the transducer face (to allow field shift compensation, discussed below). For each measurement, the acoustic time-of-flight was calculated from the central max-

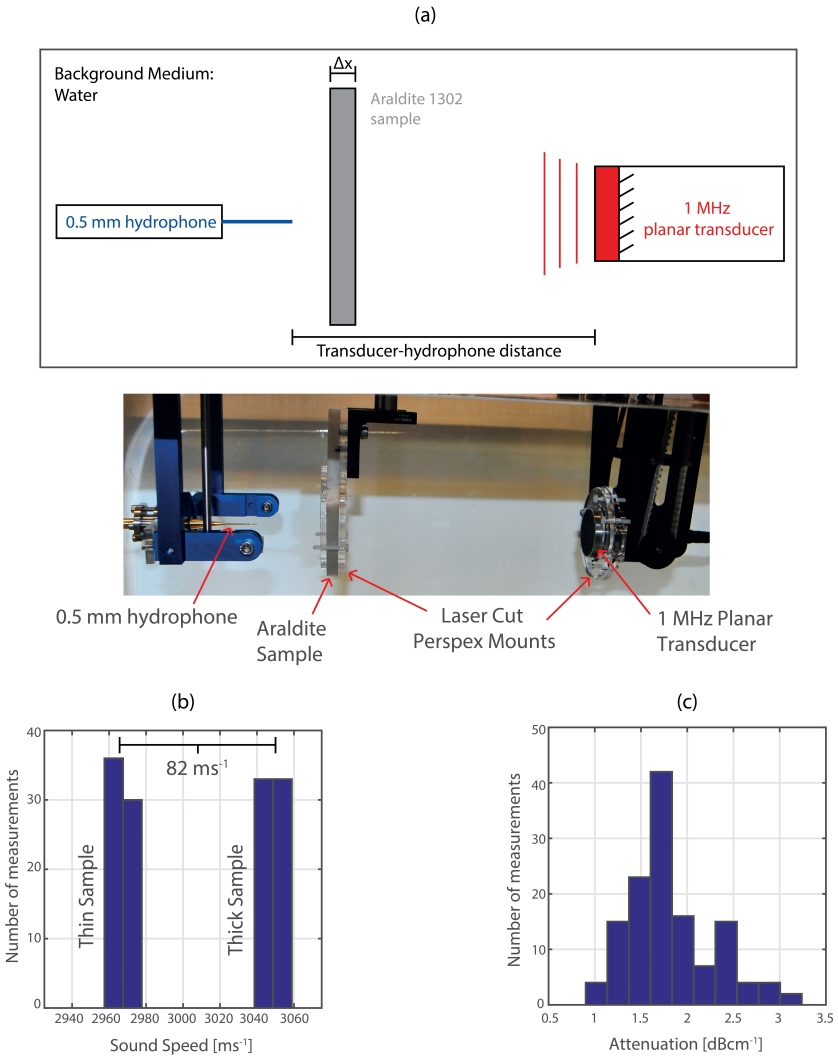

Figure 5.3: Measurements of the sound speed and attenuation of Araldite 1302 resin. (a) Schematic & photograph showing experimental setup for medium property measurements. The transducer-hydrophone distance was varied from 420-580 mm for free field measurements, and from 450-550 mm when the sample was in place. (b) Results for measured sound speed, showing the difference between samples. (c) Attenuation values calculated from the change in transmission loss between samples.

ima and the two adjacent minima of the pulse. Resin samples were then placed between the transducer and the hydrophone in the far field of the transducer. Laser cut PMMA Perspex holders were used to ensure that the samples were orthogonal to the beam axis. The measurements were repeated at 10 mm intervals along the beam axis between 450 mm to 550 mm from the transducer face. Measurements were also repeated with transmission through different positions on the sample, and with the

Acoustic Property	Araldite 1302	VeroBlack
Compressional Speed [m s^{-1}]	3008 $\pm$220	2495 $\pm$8
Shear Speed [m s^{-1}]	1562 $\pm$30	1081 $\pm$31
Absorption [dB cm^{-1}]	1.81 $\pm$0.6	3.70 $\pm$0.1
Density [kg m^{-3}]	1637	1180

Table 5.1: Compressional & shear speed, attenuation and density of Araldite 1302 resin and VeroBlack 3D printing material [14] including estimated/reported uncertainty.

samples at three distances from the hydrophone. Frequency dependent transmission loss was calculated by spectrally decomposing the signals and comparing them with free field measurements, selected following field shift compensation based on the measured sound speed.

The measured sound speed and attenuation values are shown in Fig. 5.3. Sound speed values were slightly different for the two samples (which were cast separately) due to minor variations in the Araldite mixture. The mean sound speed of the two samples is 3008 m s^{-1}. The corresponding attenuation values are shown in Fig. 5.3(b). As attenuation is calculated through a comparison of the transmission between the two samples, any inter-sample variation in the Araldite mixture could also lead to a systematic error in the calculated attenuation coefficient. The magnitude of this uncertainty is difficult to estimate. The other major source of uncertainty is the measurement of sample thickness. Based on the variation in measurements made at different points on the samples, this is estimated to be no greater than 200 μ m. This uncertainty could lead to a change in sound speed of up to $\pm$56 m s^{-1}, and in attenuation of $\pm$0.15 dB cm^{-1}. The sound speed of the deionised water is known accurately based on the temperature, while the absorption of ultrasound at frequencies below 5 MHz is negligible [85]. The impact of non-linearity was quantified by examining the spectra of free field measurements taken at the maximum distance from the transducer, and was found to be negligible. Based on the variation in material properties both within and between samples and the uncertainty of the thickness measurements, the estimated uncertainty in the sound speed is $\pm$220 m s^{-1} (or 7.3%) and in the attenuation of $\pm$0.6 dB cm^{-1} (or 33.1%).

The shear speeds of both Araldite 1302 and VeroBlack were also measured.

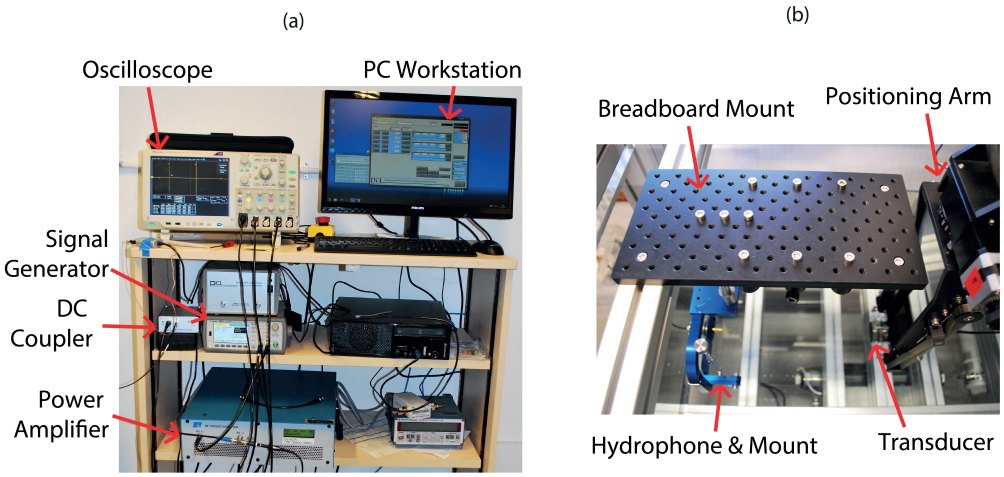

Figure 5.4: Practical experiment setup. (a) Ultrasound signal generation and monitoring equipment, including PC control system. (b) Stereotactic experimental setup, with optical breadboard for rigid registration and a computer controlled positioning arm.

Shear speed measurements were carried out with a 2.25 MHz shear wave transducer (SWT) and a pulser-receiver (Olympus, Tokyo, Japan). The transducer was excited with a single 2.25 MHz cycle. The time-of-flight of the first and second reflections of the pulse were used with the pulse input time to calculate the shear speed. Results for Araldite resin and VeroBlack are shown in Table 5.1. Both Araldite and VeroBlack have shear speeds lower than 1800 ms^{-1}, which was established as the sound speed relevant for transmission through the skull model used in Chapter 5, upon which the parametric phantom was based.

5.2 Practical Sonications and Registered Simulations

5.2.1 Practical Sonications

The pressure field generated by a 1 MHz focused transducer was measured following transmission through the skull bone phantoms. Comparison of the experimental and simulated fields was then used to evaluate the impact of uncertainties in the modelled sound speed and attenuation, which was modelled as absorption within the numerical scheme.

The ultrasound source used was a Precision Acoustics PZT HIFU transducer

with a diameter of 60 mm and a focal length of 75 mm. The transducer was driven at its nominal centre frequency of 1 MHz with a 3 cycle toneburst, generated by an Agilent 33522A Arbitrary Waveform Generator (Agilent, Berkshire, UK) and amplified by an E&I A075 RF power amplifier (Electronics & Innovation Ltd., Rochester, NY). The drive signal was monitored by a Tektronix DPO-5034B oscilloscope (Tektronix Inc., Beaverton, OR), also used to process hydrophone measurements. The ultrasound generation and monitoring setup is shown in Fig. 5.4, alongside the tank and positioning system. Measurements were carried out in a tank of degassed and deionised water with an integrated three-axis positioning system (Precision Acoustics, Dorchester, UK). Signals were acquired with a calibrated 0.2 mm needle hydrophone, connected to the oscilloscope via a submersible preamplifier and DC coupler (Precision Acoustics, as above). Water temperature was monitored continuously with a thermocouple (National Instruments Corporation, Austin, TX).

To examine the experimental peak pressure field, the time-varying pressure was recorded across a plane following transmission through the phantom. The experimental setup with a resin phantom in place is shown in Fig. 5.5 (a). The scan planes were 41×41 mm, scanned in steps of 0.2 mm, and were positioned at 58 mm and 85 mm from the transducer surface for the parametric and anatomical phantoms, respectively. The time-varying pressure plane was then bandpass filtered, and projected forward and (in the case of the mesh based phantom) backward using linear acoustic holography to establish the pressure field across a volume encompassing the transducer focus. The position of the scan planes and 3D projection volumes are shown in Fig. 5.5. The time-varying pressure was also recorded across a plane 45 mm from the centre of the transducer without the phantoms in place, to establish the source conditions for simulation.

5.2.2 Registered Simulations

Simulations designed to mimic the experimental sonications were performed based on the sterotactic alignment. The source conditions were calculated by bandpass filtering and time-reversing the measured free-field pressure signals and backpro-

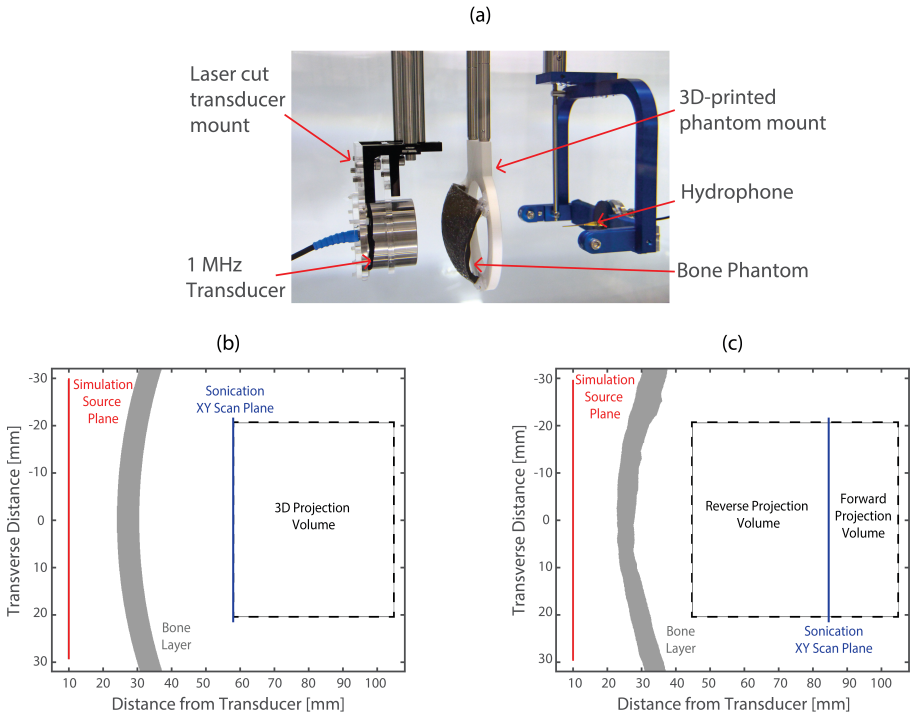

Figure 5.5: Sonication & registered simulation layout. (a) Sonication setup for scanning following transmission through phantom. (b) Registered simulation setup for the parametric skull phantom. (c) Registered simulation setup for the anatomical skull phantom.

jecting them to a plane positioned 10 mm from the transducer surface. To create the map of acoustic medium properties, the parametric phantom was incorporated into the simulation based on *a priori* knowledge of its 3D shape and position. For the anatomical phantom, the 3D mesh was directly converted into a binary map of medium properties. The background was assigned the properties of water based on the recorded temperature of the water in the tank [148]. Simulations used a spatial PPW of 15, to satisfy the sampling requirements suggested by Chapter 3. Detailed simulation parameters are described in Appendix B.4.

To examine the accuracy of the registered simulations with reference to the experimental sonications, the maximum pressure was recorded across a volume matching the 3D projection volumes shown in Fig. 5.5. For both experimental and simulated volumes, the time-varying pressure was also extracted at the position of the peak pressure in the practical sonications.

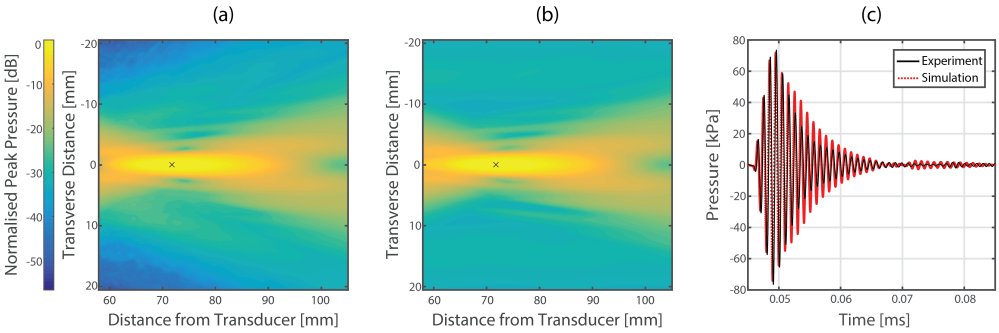

Figure 5.6: Results of experiments and simulation through parametric phantoms. (a) Experimental peak positive pressure measured on the axial-lateral plane. (b) Simulated peak positive pressure on the axial-lateral plane. (c) Comparison of time-varying pressures at the measured spatial peak pressure (marked by an $\times$ symbol). Reproduced from Robertson et al. [116].

5.3 Results

5.3.1 Parametric Phantom

A comparison of the experimental and simulated fields for the parametric skull phantom is shown in Fig. 5.6. The peak pressure in the simulation relative to the experiment was underestimated by 2.1%, the focus position moved by 0.137 mm, and the simulated focal volume was 2.03% larger. When evaluated with reference to the uncertainty in sound speed and attenuation of the parametric phantoms shown in Table 5.1, and the results of the numerical sensitivity analysis described in Chapter 4, all metrics fall well within the expected tolerances.

The experimental and simulated time-varying pressure signals at the position indicated in the peak pressure fields is shown in Fig. 5.6(c). The registered simulation predicts a slightly earlier time-of-flight for the waveform shown compared to the experiment. There is also a small difference in the pulse shape between experimental and simulated signals. The simulated pulse has a slightly lower peak amplitude, and a longer ring down. Based on this discrepancy, and the general trends shown in Chapter 4, it is probable that one or both of the attenuation and sound speed values assigned to the bone layer were slightly too high. However, overall there is good quantitative agreement between the two signals.

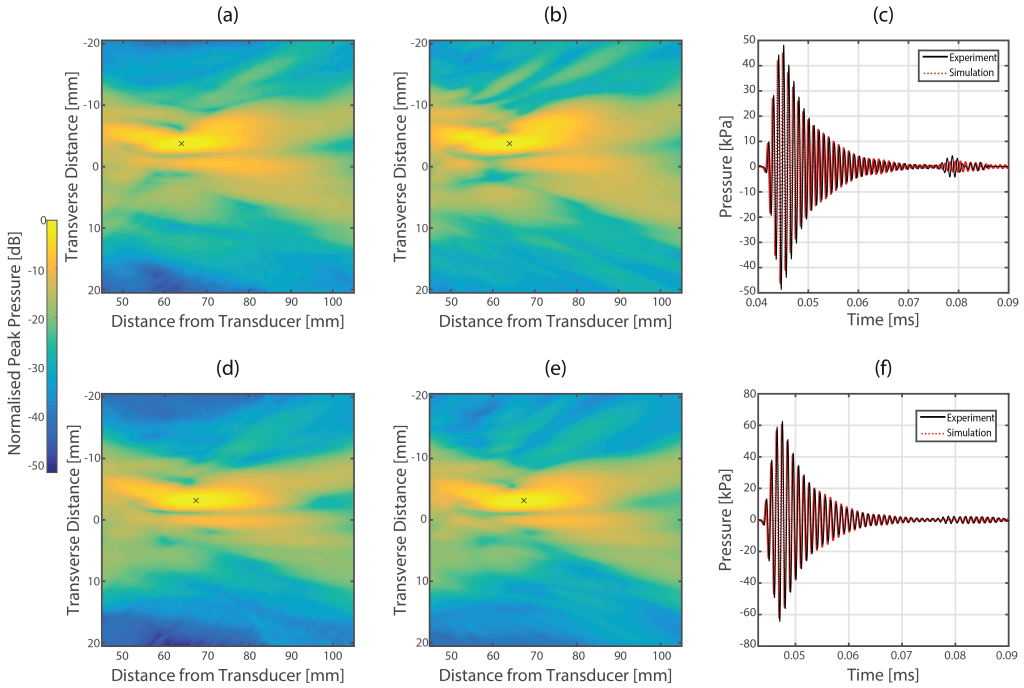

Figure 5.7: Results of experiments and simulations through anatomical phantoms. The top row shows the results for the Araldite phantom, and the bottom row shows the results for the VeroBlack phantom. (a) & (d) Experimental peak positive pressure measured across the axial-lateral plane. (b) & (e) Simulated peak positive pressure across the axial-lateral plane. (c) & (f) Comparison of the time-varying pressures at the measured peak pressure position (marked by an $\times$ symbol). Reproduced from Robertson et al. [116].

5.3.2 Anatomical Phantom

The results of the experimental sonications and registered simulations through the anatomical phantoms are shown in Fig. 5.7. Simulation of propagation through the Araldite anatomical phantom underestimated the peak pressure by 0.8% compared to the experimental measurement, the focus position was moved by 0.61 mm, and the focus volume was overestimated by 11.6%. While the errors in both peak pressure and focus position remain within expected variations, the error in focal volume is greater than would be expected for the maximum uncertainty of 220 m s^{-1}. Furthermore, based on the results of the numerical sensitivity analysis, this error in focal volume should be accompanied by a correspondingly large change in, at least, peak pressure.

The simulation of propagation through the VeroBlack mesh based phantom

underestimated peak pressure by 0.8%, the focus position was moved by 0.87 mm, and the focus volume was underestimated by 2.1%. The closer agreement is likely due to more accurate medium property measurements. Figure 5.7 also shows the experimental and simulated time-varying pressure at the peak pressure position in the experimental field. These results show good agreement in terms of pulse shape, amplitude and peak pressure, especially for the VeroBlack phantom.

The discrepancy in the results for the Araldite phantom could have a variety of causes. It is possible that the difference in medium geometry between the numerical sensitivity analysis and the anatomical phantom has altered the sensitivity of the intracranial field to the phantom acoustic properties. However, the VeroBlack phantom has comparable errors in peak pressure and position without the larger error in focal volume. It is more likely that there is inaccuracy in the medium maps arising from the casting process, or errors in phantom alignment. It is also possible that the influence of shear mode propagation, determined not to influence transmission for the parametric model in Chapter 4, has increased influence for the geometry of the anatomical phantom.

5.4 Discussion

In this chapter, multiple ultrasonic skull bone phantoms were created using computer aided design, 3D printing and resin casting techniques for use in practical experiments of transcranial ultrasound propagation. Phantoms were created in both VeroBlack and Araldite 1302 resin, and the acoustic properties of Araldite were measured.

The first "parametric" phantom was created in Araldite only, and was designed to mimic the idealised bone geometry used for the numerical sensitivity analysis in Chapter 4. The shape of simulated field agreed qualitatively with the measured field. Quantitatively, the errors in the peak pressure amplitude, size and position of the focus were all within the expected error based on the estimated uncertainty and the results of the sensitivity analysis. The second "anatomical" phantom was based on the skull mesh used for convergence testing in Chapter 3, and was cre-

ated in both VeroBlack and Araldite 1302. Generally, simulations showed good agreement with experimental measurements, including in the time-varying pressure at the experimental peak pressure position, with the exception of the error in the focus volume through the Araldite phantom. Notably, simulation through the VeroBlack achieved a very high degree of fidelity both across the simulation field and in the time-varying pressure at the focus, likely due to the accuracy of the measured acoustic property values. This highlights the importance of accurate medium property value measurements and consistent phantom materials and speaks to a primary limitation of the work described above, which is the uncertainty in the acoustic property measurements of Araldite 1302. Overall these results indicate that highly accurate simulation of ultrasound propagation through phantoms with geometry and acoustic properties close to that of skull bone is achievable when the numerical accuracy and spatial registration are well controlled, and the acoustic properties are known with a high degree of accuracy.

There are also a number of corollary observations regarding the creation process for bone phantoms and the stereotactic registration systems. The use of CAD models was integral to the simulation validation, and likely contributed greatly to the fidelity acheived in registered simulations. PLA and perspex mounts were easy to produce and served very well. While the 3D printed PLA mold for the parametric phantom led to an acceptable finished piece, the PLA binds strongly to the Araldite resin (despite the use of a release agent), and it was necessary to damage the mold during removal of the phantom, preventing it from being reused. Generally, the use of vulcanised rubber molds led to an easier casting and release process, and allows the molds to be reused for the creation of multiple phantoms. Both phantom materials demonstrated unique advantages for the production of skull bone phantoms. Due to its usage as a 3D printing substrate, the VeroBlack material can be directly printed from a 3D model with high spatial fidelity and a wide range of potential phantom geometries. Furthermore, the comparison between practical and simulated fields shown in Fig. 5.7 indicates that the acoustic properties of VeroBlack are described accurately by Nikitichev et al. [144], and are consistent between prints. However,

the acoustic properties themselves are not ideal for a bone phantom. The density of VeroBlack lies within the range of reported values of trabecular bone from the IT'IS database [10], shown in Fig. 4.2 and Table 1.4. The compressional sound speed, identified as the most important acoustic property by the results in Chapter 4, lies between the reported ranges for cortical and trabecular bone. Acoustic absorption at 1 MHz (3.7 dB cm^{-1}) is above that reported by Pinton et al. [16] for cortical bone, which may help counteract the increased transmission due to the relatively low impedance. While VeroBlack may be acceptable when aiming to approximate whole bone acoustic properties, the density and compressional sound speed, and shear sound speed are all below those reported by White et al. [9] for whole skull bone. Likewise VeroBlack is unsuitable for phantoms designed to mimic cortical bone, or trabecular bone microstructure.

The acoustic property profile of Araldite 1302 resin can be contrasted with that for VeroBlack. The compressional sound speed of 3008 m s^{-1} falls within the range of reported values for cortical bone, the density sits between the ranges for cortical and trabecular bone and the shear speed is close to that reported by White et al. [9]. The acoustic absorption measured in Section 5.1.3 is lower than that given for cortical bone by Pinton et al. [16]. However, in their original paper, Truscott et al. [145] reported an absorption value of 3.7 dB cm^{-1}, so it is possible that the value calculated in Section 5.1.3 has been influenced by the specific batch or mixing process. Based on these measurements, it is clear that the acoustic properties of Araldite are more suitable for whole bone and cortical bone phantoms, and even for mimicking cortical bone microstructure. This supports the previous usage of the resin to produce phantoms, notably Clarke et al. [146], who combined Araldite with flax seeds (which approximate the acoustic properties of human skull marrow) to produce trabecular bone phantoms. The primary limitation of Araldite is the necessity of mixing and casting final phantom pieces. This limits the range of geometries that can be produced, does not guarantee the spatial fidelity possible with 3D printing, and can evidently lead to inconsistency in the sound speed and possibly in the other acoustic properties.

Chapter 6

Homogenisation and Imaging

In the previous chapter, simulations spatially registered to practical experiments demonstrated that numerical models can effectively predict fields following transmission through skull bone phantoms, given a sufficiently accurate map of the acoustic properties. As described in Section 1.2.4, in practice these maps are usually derived from clinical CT images using a variety of methods. CT intensity can be converted to heterogeneous maps of density, sound speed and attenuation, or homogeneous maps can be produced based on a segmentation, although these necessarily neglect all internal structure of the bone.

The use of CT images to produce fully heterogeneous acoustic property maps is subject to some limitations. Notably, the spatial resolution of clinical CT images is limited to an approximate scale of 0.5 mm, which leads to a loss of the fine internal microstructure of the skull and smoothing of interfaces between bone and soft-tissue due to partial volume effects. In their method for deriving acoustic properties from CT images, Aubry et al. [6] attempt to account for this by artificially increasing the acoustic absorption value assigned to the skull bone to account for the scattering that occurs in the trabecular layer, based on the attenuation measurements made by Fry et al.[13]. Furthermore, in this thesis, spatial smoothing of acoustic property maps has already been demonstrated to have a strong influence on the transmission of ultrasound through skull bone. In Section 2.2.5 the impact of smoothing discontinuous interfaces on the accuracy of simulated reflection and transmission coefficients is shown to have a greater impact than the numerical er-

ror resulting from the use of a k-space PSTD scheme for frequencies sampled at close to 2 PPW. Likewise, in Chapter 4, kernel based smoothing of an idealised skull bone layer was shown to have a significant effect on the transmission of ultrasound. However, simulations through acoustic property maps derived from medical images have not been carried out, and the impact of limited imaging resolution of the transmission of ultrasound through skull bone not properly explored.

In this chapter, simulations of ultrasound propagation through the skull are performed to examine the impact of smoothing and loss of microstructure detail due to the limits of medical image resolution. Acoustic property maps are derived from a dataset comprising clinical CT images of skull bone and corresponding micro-CT data of sections of skull bone co-registered to the original clinical CT images. These simulations are used to examine the impact of progressive downsampling of high resolution acoustic property maps based on micro-CT data, which mimics the effects of image related homogenization. Simulations through maps based on micro-CT data are compared with simulations through heterogeneous acoustic property maps derived from co-registered clinical CT data. Different methods for segmenting CT data and resampling it onto simulation grids at the resolution of the micro-CT images are compared. The degree to which assigning whole skull attenuation to the skull layer *en-bloc* is able to correct for the loss of microstructure in clinical CT data is also examined. Simulations through homogeneous medium maps derived from the same clinical CT data are carried out to examine the difference from heterogeneous maps relative to the simulations through micro-CT data. The impact of elastic wave propagation through the microstructure of skull bone is examined using simulations in 2D.

6.1 Maps of Acoustic Properties

6.1.1 Clinical-CT and Micro-CT Datasets

Co-registered clinical and micro-CT data of human skull bones were provided by partners from Wake Forest University School of Medicine [79] obtained as part of their Whole Body Donation Program. The clinical CT data comprised two clinical

CT images of complete heads obtained post-mortem with isotropic resolution of 0.625 mm. Following acquisition of the clinical CT images, small samples of bone (four from skull A, three from skull B) were extracted for high-resolution imaging in a micro-CT scanner with an isotropic resolution of 50 μm. Micro-CT data were then manually aligned and co-registered with the clinical and micro-CT datasets using GEOMAGIC STUDIO (v 12.1.0 Geomagic, Research Triangle Park, NC, USA), described in Lille et al. [79]. This alignment was used to create an affine transform matrix, provided with each micro-CT image, that allowed them to be aligned and registered with their corresponding clinical CT datasets.

Following visual inspection, two of the datasets were rejected from the study due to their geometry being unsuitable for the planned transmission simulations. This left a total of two micro-CT images for skull A, and three for skull B. These comprised a section of frontal and parietal bone for each skull, in addition to a portion of temporal bone from skull B. 3D renderings of the individual bone samples are shown in Fig. 6.1, alongside sections through the raw micro-CT data, demonstrating the three-layered structure and internal heterogeneity of the skull. Here, it can be seen that the section of temporal bone has a different structure than the samples of frontal and parietal bone, comprising mainly dense cortical bone with very little trabecular microstructure. This sample serves as a comparison to the other bone fragments, and is used to determine the extent to which any effects observed are due to the loss of bone microstructure, rather than other aspects of the smoothing and homogenization process.

Micro-CT and clinical CT datasets were aligned inside the Matlab environment. Due to the need to interpolate the clinical CT data onto spatial grids of the same spatial sampling as the micro-CT data for simulations, the affine transform matrix was inverted and used to rotate the clinical CT data. Interpolation of the clinical CT onto the micro-CT resolution grids was carried out using both nearest neighbor and linear interpolation, to determine which method results in more accurate transcranial simulations.

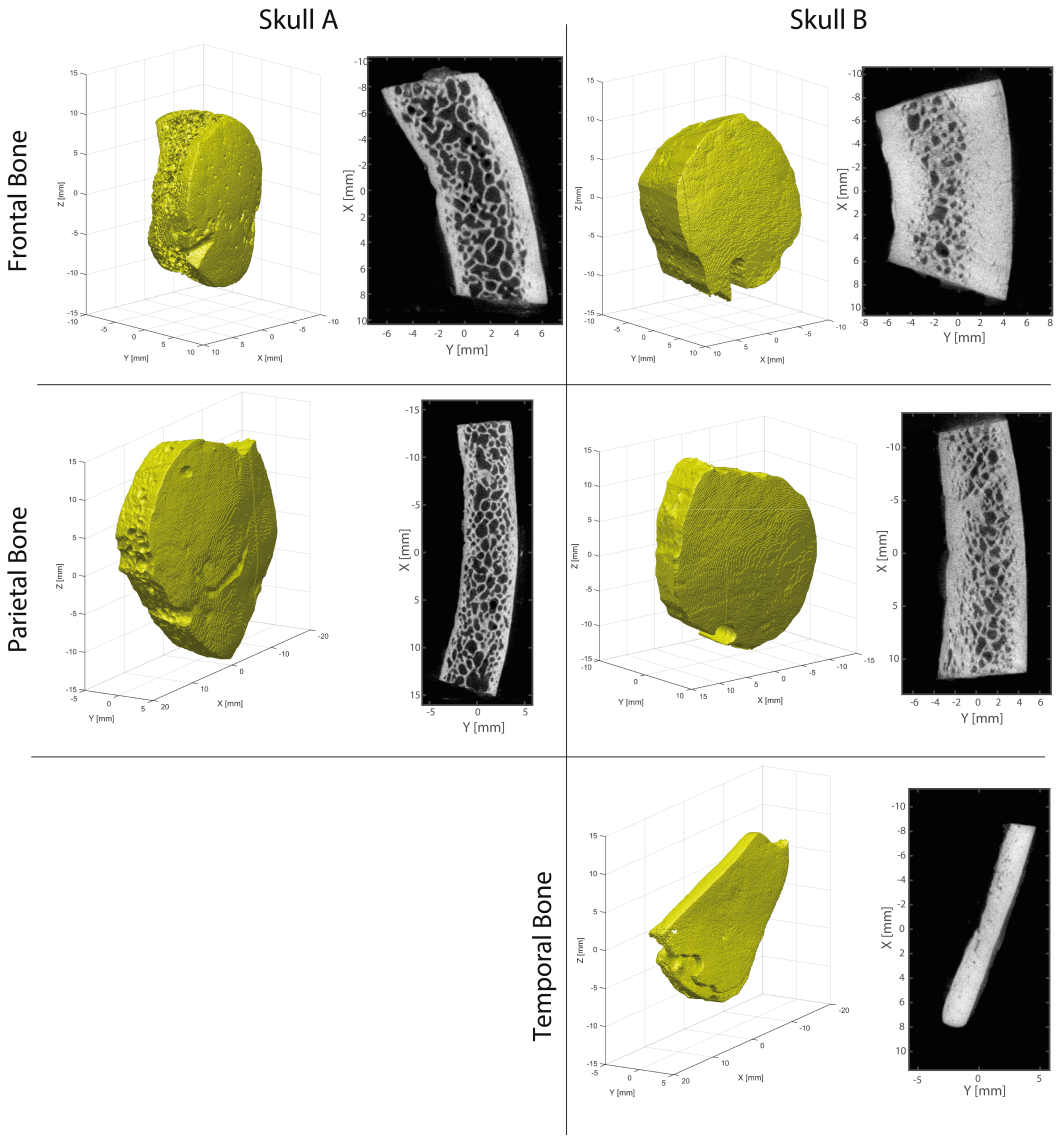

Figure 6.1: 3D renderings of the micro-CT skull sections produced in the Seg3D segmentation toolbox, alongside sections through the raw micro-CT data, demonstrating the internal structure of the skull bone.

6.1.2 Assignation of Acoustic Medium Properties

Acoustic property maps were derived from the micro-CT datasets by segmenting the 3D image volume into cortical bone and soft-tissue components based on the CT intensity. A 40% threshold intensity was used to created the segmentation, followed by the application of a connected component algorithm. Segmentations of all micro-CT datasets were carried out using the Seg3D segmentation toolbox [149] and exported directly into the Matlab environment. A separate, solid segmentation

Tissue	c [ms^{-1}]	ρ [kgm^{-3}]	1 MHz Absorption/Attenuation [dBcm^{-1}]
Soft-Tissue	1540	1000	0
Bone Marrow	1540	1000	1.1
Cortical Bone	3100	2190	2.7
Whole Bone	N/A	N/A	13.3

Table 6.1: Acoustic property values used in the derivation of property maps from CT images. Sound speed and density values for soft-tissue and cortical bone are taken from Deffieux et al. [15]. Absorption values for cortical bone, and attenuation values for whole bone are taken from Pinton et al. [16]. The absorption value for bone marrow is taken from the IT'IS database [10].

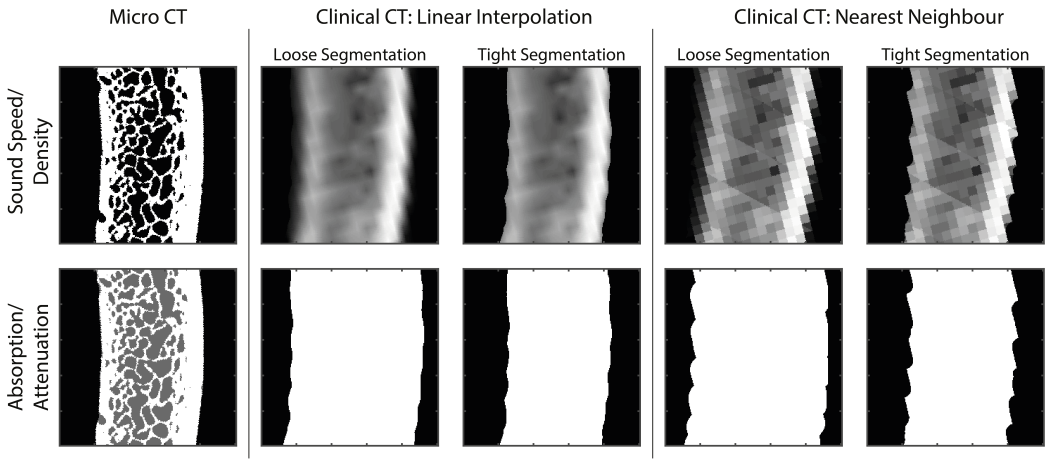

Figure 6.2: The different methods for homogenisation of the acoustic property maps based on segmentation of micro-CT data into cortical bone and soft-tissue segments. Homogenisation at two different spatial resolutions are shown, alongside clinicat CT data resampled onto a spatial grid with 50 μm.

delineating the bone segment as a whole was also created to allow soft-tissue within the bone layer (bone marrow) to be separated from the rest of the volume. The acoustic properties of cortical bone were assigned to the initial segmentation of the bone, soft-tissue within the bone volume was assigned the properties of bone marrow, and the rest of the domain was assigned the properties of water. To ensure that acoustic property maps derived from micro-CT data could be compared to those derived from clinical CT data using the algorithm described below, the absorption value assigned to cortical bone was based on the measurements by Pinton et al. [16], and the same sound speed and density were assigned to water and bone marrow. The acoustic property values used are shown in Table 6.1.

Acoustic properties were assigned to clinical CT data in accordance with the algorithm developed by Aubry et al. [6], which uses HU values to interpolate between acoustic sound speed and density values for cortical bone and soft-tissue as described in Section 1.2.4. In accordance with previous usage of this algorithm by Deffieux et al. [15], the first step is segmentation of clinical CT data into skull & non-skull segments using a threshold segmentation of 50% of the peak intensity in the image, referred to herein the "tight" segmentation. Segmentations of the clinical CT data were carried out in the Matlab environment separately for linear and nearest neighbour interpolations following the rotation and interpolation of the clinical CT images onto grids of micro-CT data. To determine the influence of accounting for smoothing around the edges of acoustic property maps, a separate "loose segmentation was created by dilating the "tight segmentation by 1 mm in all directions. Within the loose segmentation volume, intensity values of the clinical CT data were then rescaled to vary between 0 and 1, and used to interpolate between the sound speed and density values used for cortical bone and marrow/water in the micro-CT data. This ensures that the basis for assigning acoustic property values remained consistent between clinical and micro-CT datasets. The tight segmentation was then applied to the resulting acoustic property maps, to determine the impact of enforcing a discontinuous medium boundary at the edges of the skull. The range of different acoustic property maps produced from both clinical and micro-CT data are shown in Fig. 6.2.

Simulations through acoustic property maps derived from the clinical CT data were carried out both with and without attenuation, modelled within the simulations as absorption, to determine how effectively it accounted for the reduced scattering due to loss of microstructure detail at clinical CT resolutions. As in Deffieux et al. [15], attenuation was assigned homogeneously to the bone volume using the corresponding skull segmentation. The whole skull attenuation value assigned was that measured by Pinton et al. [16] during their measurement of the cortical bone absorption that was used in micro-CT simulations. In addition, a series of homogeneous acoustic property maps were created based on the average values of sound

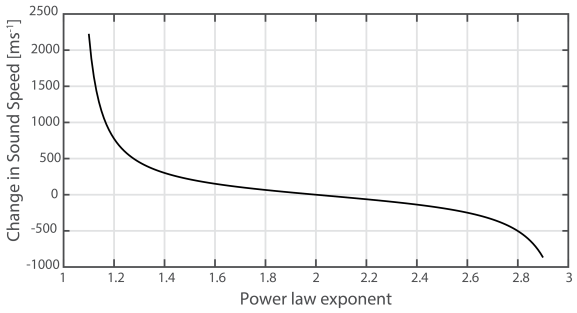

Figure 6.3: Change in the sound speed of cortical bone for 1 MHz ultrasound resulting from the introduction of absorption in the k-Wave toolbox, as a function of the power law exponent used.

speed and density across the heterogeneous maps derived from clinical CT data. The segmentations used to delineate the bone volume were used to create these homogeneous maps, and the same maps of attenuation were used.

For all simulations where absorption was modelled (including attenuation modelled as absorption), the power law exponent was set to $y = 2$. As all simulations used either CW sources, or tonebursts with a narrow frequency spectrum, this should not influence the absorption of ultrasound. However, when absorption is simulated in k-Wave, it is accompanied by physical dispersion, which causes a frequency dependent change in sound speed determined by the power law exponent, shown in Fig. 6.3. Using $y = 2$ results in no dispersion, which allows direct comparison between lossless and absorbing simulations through clinical CT data.

6.1.3 Homogenisation of Property Maps

The maps of acoustic properties derived from micro-CT data were progressively homogenised in order to mimic the effects of imaging resolution. Homogenisation was applied directly to the maps of acoustic properties rather than the original micro-CT images, and consisted of spatially downsampling the maps by averaging medium properties across neighboring groups of grid points in order to imitate the voxels of images obtained at a lower spatial resolution. These downsampled medium maps were then resampled back onto the simulation grid at the micro-CT resolution using either linear interpolation or nearest neighbor interpolation, resulting in both smooth and pixelated property maps. The initial acoustic property maps,

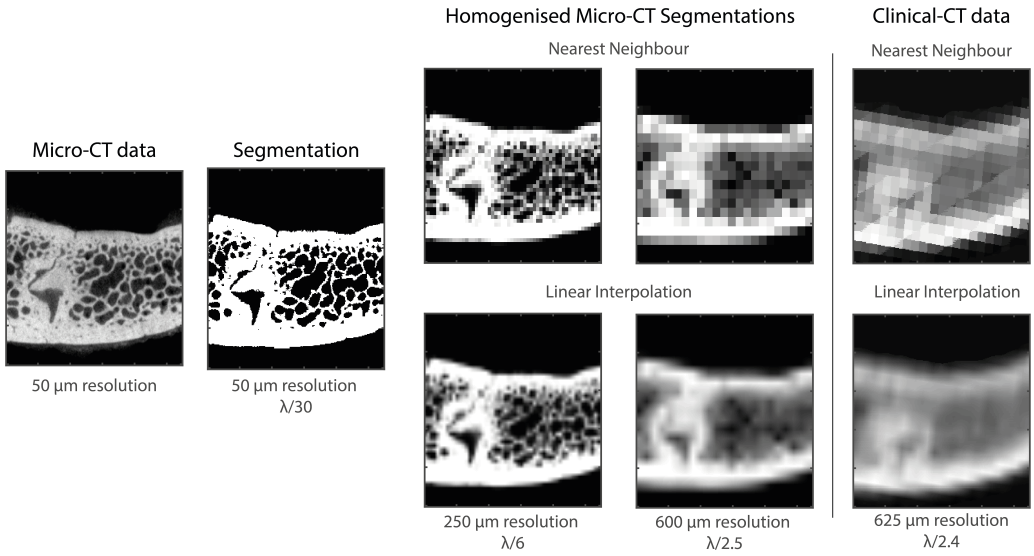

Figure 6.4: The different methods for homogenisation of the acoustic property maps based on segmentation of micro-CT data into cortical bone and soft-tissue segments. Homogenisation at two different spatial resolutions are shown, alongside clinicat CT data resampled onto a spatial grid with 50 μm. Also shows the scale of the virtual voxel dimensions relative to the wavelength of 1 MHz ultrasound in soft-tissue, λ.

with a spatial resolution of 50 μm, were homogenised to effective voxel sizes of up to 1 mm in 50 μm increments. Figure 6.4 shows examples of the resulting medium maps at different effective resolutions alongside those derived from clinical CT.

6.2 Simulation Setup and Testing Method

Simulations of ultrasound propagation were carried out to evaluate the impact of image related homogenization of acoustic property maps derived from micro-CT data. The results for micro-CT data homogenised at different spatial sampling levels were compared with simulations through both heterogeneous and homogeneous maps derived from clinical CT data. Simulations through the highest resolution, non-homogenised micro-CT data, where acoustic properties were assigned based on cortical bone segmentations, were considered reference simulations and errors were computed relative to them. Simulation through maps based on clinical CT data when attenuation was not simulated were compared with simulations through non-absorbing micro-CT data, while simulations through attenuating clinical CT

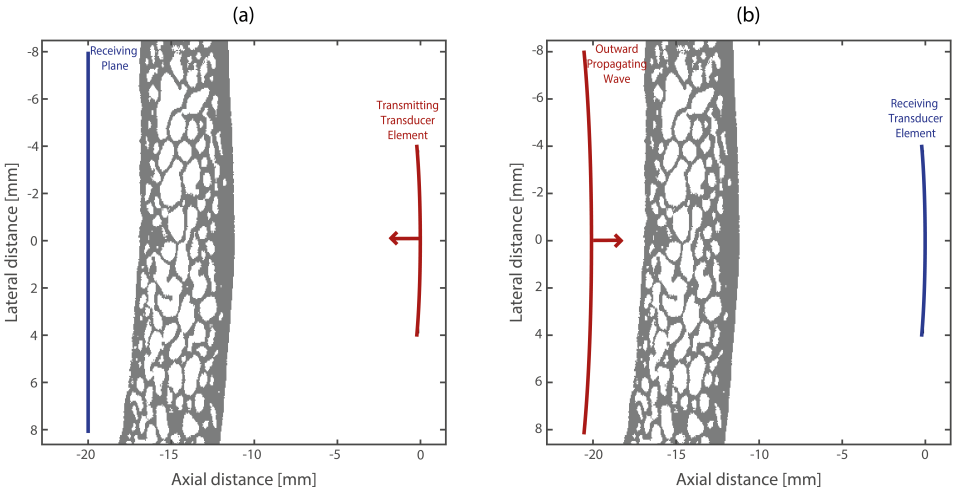

Figure 6.5: Layout of different testing regimes. (a) Forward simulation of propagation from extracranial transducer to recording plane inside the skull. (b) Reversal simulation comprising propagation from outwardly propagating spherical wave to external transducer.

maps were compared with simulations through absorbing micro-CT data.

In keeping with the main uses of numerical simulation of transcranial ultrasound propagation, two testing schemes were used. The first was based on simulations designed to examine the intracranial ultrasound field produced by applied ultrasound, designated the forward problem. The second was based on time-reversal simulations used to derive drive signals for transducer elements of large transcranial arrays, designated the reverse problem. The layout for both forward and reversal simulations is shown in Fig. 6.5, demonstrating the position of the transducers relative to the bone and the source and sensor distributions placed on the internal side of the bone map.

The same simulated transducer element was used in both forward and reversal simulations (in transmit and receive modes, respectively) based on individual elements of large hemispherical transcranial arrays like that presented by Clement et al. [92]. It was modelled as a circular single element focused transducer with a diameter of 8 mm and a radius of curvature of 10 cm, operating at a frequency of 1 MHz, which lies at the upper end of ultrasound frequencies feasible for transcranial transmission [15]. Forward simulations were carried out using 10 cycle toneburst sources, while reversal simulations used both continuous wave (CW) and ten-cycle

toneburst (TB) ultrasound sources. Detailed simulation parameters for both simulation setups can be found in Appendix B.5.

6.2.1 Forward Propagation Simulation

The forward simulation modelled the propagation of ultrasound into the skull from the single element transducer. The use of a 50 μm spatial discretization step prevented extending the simulation domain to include the theoretical focus of the transducer, as the grid size would become too large to execute the simulations. Therefore the simulated field was evaluated by recording the time-varying pressure field over a 2D transverse plane with a size of 11×11 mm, located 2 cm from the centre of the virtual transducer.

For each grid point across the recording plane, the TB data signal envelope was computed and used to calculate the amplitude and time-of-flight of the toneburst peak across the plane. An L2 error norm in the temporal peak amplitude at each grid point, as well as the average error in time-of-flight was computed across the recording field for each simulation relative to the corresponding reference simulation.

6.2.2 Time-Reversal Simulation

For the reversal simulation, the source term was designed to imitate an outwardly propagating spherical wave originating from the theoretical focus of the transducer element. The source was modelled as a curved surface spanning the entire simulated domain, with a radius of curvature of 8 cm with its center placed at 2 cm from the center of the transducer element on the internal side of the bone map, shown in Fig. 6.5 (b). After the simulation was completed, the time-varying pressure recorded across the virtual transducer was integrated into a single signal. In practical focusing simulations, this signal would normally be time-reversed, and used to derive drive signals for a transcranial sonication. This could involve using the raw time-reversed signal as an input, or analysing it spectrally to calculate phase and amplitude drive parameters. The integrated signal was analysed accordingly. For CW simulations, it was decomposed spectrally to give phase and amplitude values, while for TB

data the signal envelope was computed, and the error in the amplitude and time-of-flight of the envelope peak was calculated. An L2 error was also calculated for the entire TB signal, which is more reflective of the error that would be introduced in full time-reversal of the recorded signal. It should be noted that, to allow analysis of phase values for the CW simulations, the results for the progressive homogenization were unwrapped to determine the trend in phase. For plotting, phase values for the individual CT simulations were shifted by increments of 2π to move them onto the same plotting axis.

6.2.3 Elastic Simulations

As with the sensitivity analysis in Chapter 4, additional testing was carried out in 2D to examine the impact of shear mode propagation on transmission through acoustic property maps derived from the micro-CT data. The "reversal" simulation setup was converted into 2D for testing, and simulations were carried out using toneburst sources only.

Simulations were carried out using the elastic PSTD code in the k-Wave toolbox, and simulations were carried out both with and without elastic properties assigned to bone layer to determine the impact of shear mode propagation and the homogenisation of elastic properties. When simulated, a shear sound speed of 1500 m s^{-1} was assigned to the cortical bone segment based on the measurements made by White et al. [9], and an absorption value of 19.5 dB cm^{-1} for ultrasound absorption at 1 MHz, based on the power law derived by Treeby et al. [139]. Shear sound speed and absorption were set to zero throughout the rest of the domain. Homogenisation of the acoustic property maps was carried out using linear interpolation, as shown in Fig. 6.4. Elastic simulations were not carried out using the clinical CT data, as a reliable method for converting clinical CT images of the skull into maps of elastic acoustic properties is not available.

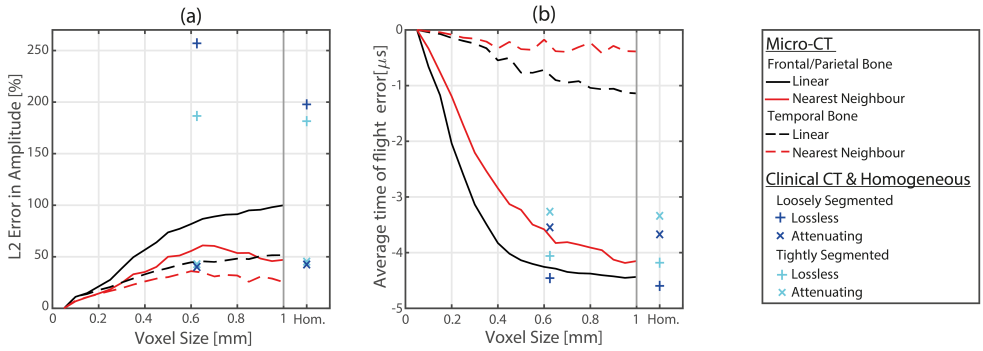

Figure 6.6: Errors in the ultrasound signals recorded across the intracranial plane for the forward simulation setup, computed from signal envelopes as a function of homogenisation voxel size. Results for simulations through both heterogeneous and homogeneous skull models derived from linearly interpolated clinical CT data are also shown, with those from homogeneous maps plotted separately. (a) L2 error in peak amplitude recorded across the plane. (b) Average error in ultrasound time-of-flight

6.3 Results

6.3.1 Forward Propagation

Progressive homogenisation of acoustic property maps derived from micro-CT data leads to an increase in the amplitude of transmitted ultrasound, a reduction in time-of-flight, and a loss of fine spatial detail in the intracranial field. Figure 6.6 summarises the error metrics computed for the forward simulations, averaging the results across frontal and parietal bone samples. These are compared with the results of simulations through temporal bone, which demonstrated reduced errors relative to the bone samples with more internal microstructure, especially in terms of the effect on the time-of-flight. The errors shown are calculated for non-absorbing simulations only, as absorbing data showed identical trends only with a slight reduction in error. When downsampling acoustic property maps derived from micro-CT data, the use of a nearest neighbor algorithm for resampling onto the high-resolution grid consistently results in reduced error in amplitude, including for simulations through temporal bone.

The used of a nearest neighbour algorithm for interpolation of CT data also resulted in reduced error when the property maps are loosely segmented, but not when

Interpolation Method	Segmentation	Bone Type	Heterogeneous Lossless	Attenuating	Homogeneous Lossless	Attenuating
Linear	Tight	Frontal/ Parietal	187%	43%	182%	45%
		Temporal	62 %	46%	50%	49%
	Loose	Frontal/ Parietal	257%	40%	198%	42%
		Temporal	92%	44%	114%	47%
N.N.	Tight	Frontal/ Parietal	195%	45%	187%	47%
		Temporal	64%	51%	57%	55%
	Loose	Frontal/ Parietal	246%	40%	209%	45%
		Temporal	89%	46%	154%	52%

Table 6.2: Average error in toneburst amplitude (percentage error) across the intracranial scan plane for forward simulations through acoustic property maps derived from clinical CT data. Average errors in frontal & parietal bone are shown alongside the corresponding error observed for simulation through temporal bone. N.N. — nearest neighbour

Interpolation Method	Segmentation	Bone Type	Heterogeneous Lossless	Attenuating	Homogeneous Lossless	Attenuating
Linear	Tight	Frontal/ Parietal	-4.06 μs	-3.26 μs	-4.18 μs	-3.34 μs
		Temporal	-1.28 μs	-1.29 μs	-1.13 μs	-1.27 μs
	Loose	Frontal/ Parietal	-4.46 μs	-3.55 μs	-4.60 μs	-3.67 μs
		Temporal	-1.57 μs	-1.50 μs	-1.61 μs	-1.55 μs
N.N.	Tight	Frontal/ Parietal	-3.88 μs	-3.07 μs	-4.04 μs	-3.18 μs
		Temporal	-1.11 μs	-1.23 μs	-0.86 μs	-1.14 μs
	Loose	Frontal/ Parietal	-4.31 μs	-3.48 μs	-4.57 μs	-3.65 μs
		Temporal	-1.50 μs	-1.44 μs	-1.55 μs	-1.52 μs

Table 6.3: Average error in toneburst time-of-flight across the intracranial scan plane for forward simulations through acoustic property maps derived from clinical CT data. Average errors in frontal & parietal bone are shown alongside the corresponding error observed for simulation through temporal bone. N.N. — nearest neighbour

tightly segmented. Errors in toneburst amplitude and time-of-flight resulting from simulations through clinical CT data are summarised in Tables 6.2 & 6.3, comparing average errors in simulations through frontal & parietal bone with simulations through temporal bone. Tight segmentation of property maps gives reduced error when attenuation is not simulated, and the reduced error for attenuating, loosely segmented property maps is likely due to the larger segmentation volume over which attenuation is introduced. Contrary to expectations, homogeneous models derived from clinical CT data tend to perform approximately the same as heterogeneous models, at least in terms of the error metrics computed here.

Results for clinical CT maps that result from linear interpolation on to the simulation grid are also shown alongside the homogenised micro-CT data in Fig. 6.6 for comparison with the micro-CT data. Errors calculated for simulations through heterogeneous property map derived from clinical CT data are plotted on the axis in a position corresponding to the resolution of the CT images, while results for ho-

mogeneous medium maps are plotted separately. Simulations through maps derived from clinical CT data showed similar errors in average time-of-flight to the micro-CT data homogenised at the same resolution, but the error in amplitude over the recording plane is much larger. Accounting for attenuation does lead to a large reduction in the error in transmitted amplitude, and a small reduction in time-of-flight error.

Examples of the peak pressure maps across the intracranial scan planes are shown in Fig. 6.7 for simulations through the parietal bone sample from skull B. They show fields resulting from propagation through acoustic property maps derived from the micro-CT, at various levels of downsampling, and through heterogeneous and homogeneous maps derived from clinical CT data (non-absorbing maps are shown). They clearly demonstrate the increase in transmitted amplitude with homogenisation of the micro-CT maps, and the improved performance obtained from using a nearest neighbour algorithm. In terms of the fields resulting from transmission through clinical CT data, while tightly segmenting the acoustic property maps and using a nearest neighbor algorithm did restore some finer spatial detail to the fields, it is not apparent that the new field shape more accurately represents the reference simulations. Individual results for each bone segment can be found in Appendix C, including maps of the intracranial peak pressure across the scan plane for each bone sample.

6.3.2 Time-Reversal

Generally, the results for the time-reversal simulations follow the same trends as those observed for the forward simulation tests. Homogenisation of micro-CT leads to an increase in the transmitted amplitude for both CW and TB simulations. It leads to a decrease in simulated time-of-flight for TB data, which manifests in the CW data as an increasing error in phase. Figure 6.8 summarises these results, showing the average of each error metric across frontal and parietal bone segments, with results for the temporal bone plotted separately to illustrate any difference. As with forward simulations, the error in amplitude is lower for temporal bone, where it shows a steady increase with homogenisation, rather than the rapid increase be-

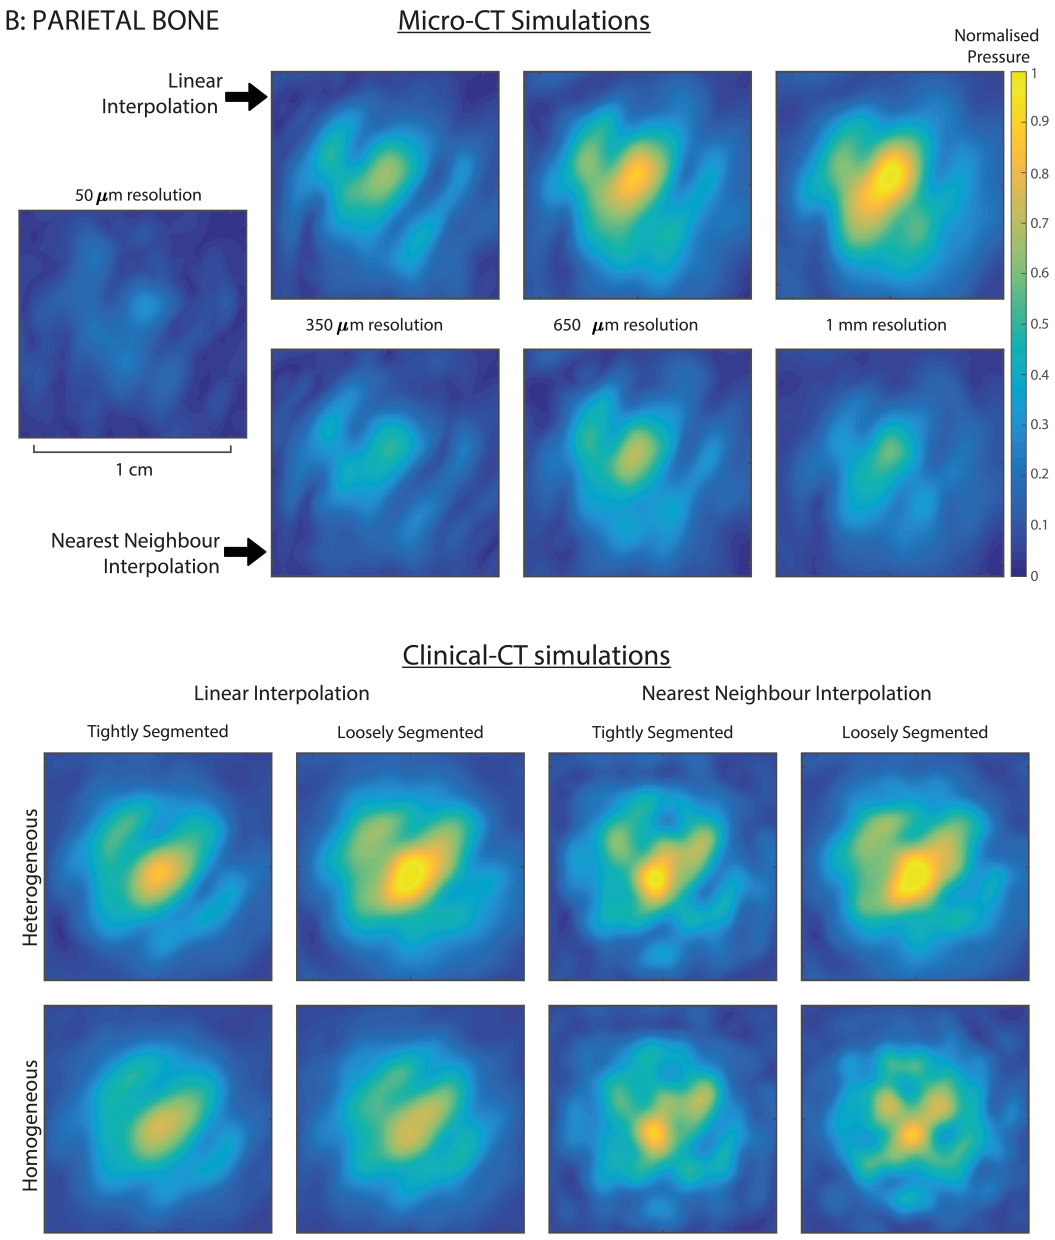

Figure 6.7: Peak pressure fields across the intracranial scan plane following forward simulation through various acoustic property maps of the section of parietal bone taken from skull B. Plots for each bone sample can be found in Appendix C.

tween 0.05 mm and 0.5 mm observed for frontal and parietal bone. The difference is more pronounced for errors in time-of-flight and phase, to the extent that the simulations through temporal bone show an increase in time-of-flight, rather than the reduction observed for frontal and parietal bone.

Average errors for TB simulations resulting from maps derived from clinical

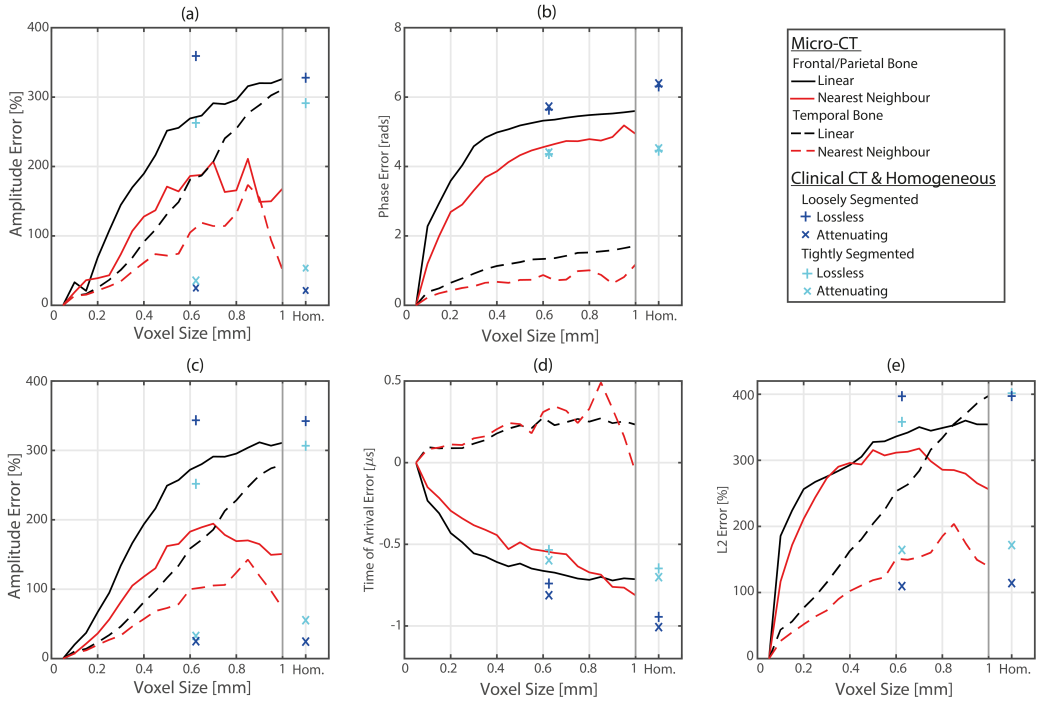

Figure 6.8: Errors in the signal across the surface of a simulated transducer following a time-reversal simulation through bone models, as a function of homogenisation voxel size. Results for simulations through both heterogeneous and homogeneous skull models derived from linearly interpolated clinical CT data are also shown, with those from homogeneous maps plotted separately. Error in continuous wave (a) amplitude and (b) phase of signal calculated at transducer. Error in toneburst (c) amplitude and (d) time-of-flight calculated at transducer. (e) L2 error in time-varying signal computed relative to the signal resulting from the reference simulation.

CT data for frontal and parietal samples are shown in Tables 6.4 & 6.5. For the heterogeneous maps without attenuation, the errors in time-of-flight, phase (not shown) and amplitude are similar to those obtained from micro-CT data homogenised at approximately the same spatial resolution. As with forward simulations, the use of nearest neighbour interpolation and tightly segmenting acoustic property maps leads to reduced error, while adding attenuation leads to a large reduction in amplitude error. For TB simulations, introducing attenuation leads to a slightly worse error in time-of-flight, but almost no change in phase for single frequency, CW simulations. This could be due to frequency dependent absorption leading to a change in the pulse shape for TB simulations, or due to greater attenuation of longer ultrasound

Interpolation Method	Segmentation	Bone Type	Heterogeneous Lossless	Heterogeneous Attenuating	Homogeneous Lossless	Homogeneous Attenuating
Linear	Tight	Frontal/ Parietal	252%	33%	307%	55%
		Temporal	270%	111%	235%	98%
	Loose	Frontal/ Parietal	343%	25%	342%	24%
		Temporal	285%	66%	135%	1%
N.N.	Tight	Frontal/ Parietal	260%	34%	302%	52%
		Temporal	218%	81%	177%	65%
	Loose	Frontal/ Parietal	328%	12%	346%	19%
		Temporal	276%	56%	187%	18%

Table 6.4: Error in toneburst amplitude (percentage error) for reversal simulations through acoustic property maps derived from clinical CT data. Average errors in Frontal & Parietal bone are shown alongside the corresponding error observed for simulation through temporal bone. N.N. — Nearest Neighbour

Interpolation Method	Segmentation	Bone Type	Heterogeneous Lossless	Heterogeneous Attenuating	Homogeneous Lossless	Homogeneous Attenuating
Linear	Tight	Frontal/ Parietal	-0.53 μs	-0.60 μs	-0.65 μs	-0.70 μs
		Temporal	0.20 μs	0.09 μs	0.10 μs	0.05 μs
	Loose	Frontal/ Parietal	-0.74 μs	-0.81 μs	-0.94 μs	-1.01 μs
		Temporal	-0.04 μs	-0.15 μs	-0.16 μs	-0.28 μs
N.N.	Tight	Frontal/ Parietal	-0.55 μs	-0.60 μs	-0.68 μs	-0.72 μs
		Temporal	0.13 μs	0.05 μs	0.04 μs	0.01 μs
	Loose	Frontal/ Parietal	-0.67 μs	-0.76 μs	-0.98 μs	-1.04 μs
		Temporal	0.03 μs	-0.10 μs	-0.17 μs	-0.29 μs

Table 6.5: Error in toneburst time-of-flight for reversal simulations through acoustic property maps derived from clinical CT data. Average errors in Frontal & Parietal bone are shown alongside the corresponding error observed for simulation through temporal bone. N.N. — Nearest Neighbour

paths through the bone leading to a shift in the signal recorded at the transducer surface. As with the forward simulation setup, propagation through homogenous models was at least as erroneous as simulations through heterogeneous models. Raw results for the homogenisation of each for each bone sample are presented in Appendix C.

For reversal simulations, it should be noted that the relationship between the error in the signal recorded at the virtual transducer and the resulting impact on the intracranial fields when using these simulations to focus is not straightforward. If it is necessary to reconstruct the exact time-varying pressure signal, then errors in the amplitude and time-of-flight assigned to the transducer element will lead to corresponding errors in the amplitude and timing of the pressure field at the target. However, if the error in the amplitude and phase (or time-of-flight) is constant across each element, the relative phase and amplitude will be conserved, and the intracranial field should be effectively spatially focused. In this respect, it is not the absolute error in amplitude and phase/time-of-flight that is important, but the

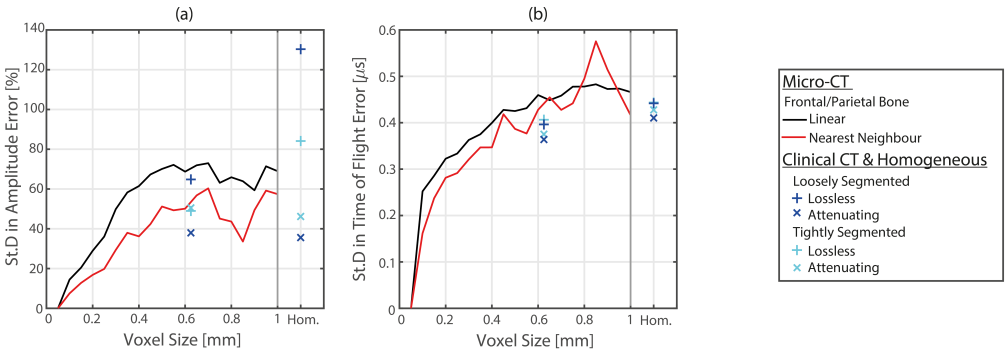

Figure 6.9: Standard deviation in the error metrics computed for reversal simulations using toneburst sources, calculated across all bone samples. (a) Standard deviation in amplitude error. (b) Standard deviation in time-of-flight error.

variation in this error between different transducer elements.

Figure 6.9 shows the standard deviation in the error in both amplitude and time-of-flight for reversal simulations using a toneburst source, for all bone samples (including temporal bone). It demonstrates how the standard deviation in both amplitude and time-of-flight error increases with increasing homogenization. Amplitude error shows a standard deviation of over 60% for homogenization on the scale of clinical CT images, while time-of flight varies by over 0.3 s. The variation in the error resulting from simulations through linearly interpolated clinical CT data is also shown. Heterogeneous property maps show similar variance in error to the homogenised micro-CT data, while homogeneous property maps show greater variance in the non-attenuating case.

6.3.3 Elastic Simulations

Results for simulations through 2D property maps of the skull, with and without shear mode propagation modelled, are shown in Fig. 6.10. Error in the signal across the transducer surface is computed relative to simulations through non-homogenised micro-CT data without elastic medium properties assigned. For the non-homogenised micro-CT data, the introduction of acoustic properties leads to an average reduction in the amplitude of 33% for frontal and parietal bone, and an average reduction in time-of-flight of 0.45 μs. However, the difference in amplitude error between elastic and non-elastic simulations does not increase with

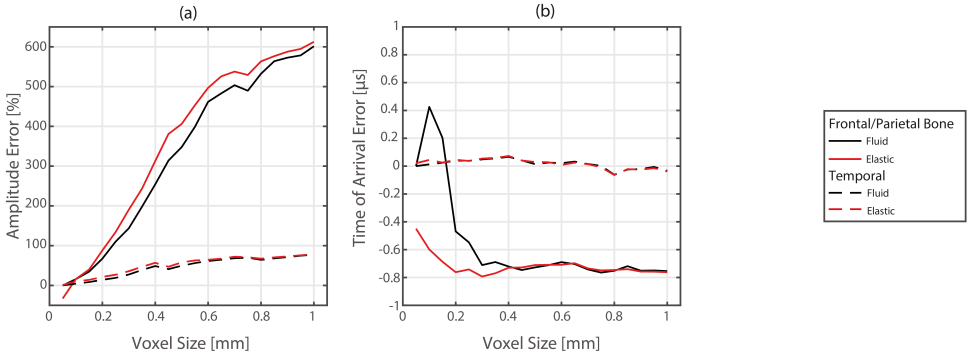

Figure 6.10: Comparison between errors resulting from reversal simulations through progressively homogenised micro-CT data for elastic and non-elastic simulations, showing error in the toneburst signal across the simulated transducer surface relative to the non-homogenised, non-elastic case. (a) Error in the amplitude of the toneeburst signal. (b) Error in the time-of-flight of the toneburst signal.

homogenisation, and the difference in time-of-flight error reduces with increasing homogenisation.

For temporal bone simulations, the difference between elastic and non-elastic simulations is much smaller, and simulations through the non-homogenised map result in an error in amplitude of 1.2%, and an error in time-of-flight of 0.02 μs when introducing shear mode propagation. This indicates that the influence of shear mode properties is mainly the result of wave propagation in the trabecular bone layer. Overall, these results indicate that, accounting for elastic wave propagation in the bone tissue can have a significant effect when modelling the propagation of ultrasound across the internal microstructure of the skull. However, the inclusion of elastic wave propagation does not appear to modify the change in transmission through the skull that occurs when acoustic property maps are progressively homogenised.

6.4 Discussion

Homogenisaton of medium maps derived from micro-CT data universally leads to an increase in simulated transmitted amplitude, and errors in time-of flight and phase on the scale of the acoustic period of 1 MHz ultrasound. Comparative simulations through temporal bone showed a large reduction in the error in simulated

time-of-flight & phase. Amplitude error increased more slowly with increasing homogenization, but was still significant. This suggests that the effects of homogenization on simulated time-of-flight are due to the loss of heterogeneous internal microstructure not found in temporal bone, possibly due to changes in the mean pathlength through the skull layer. Likewise, it suggests that the error in simulated amplitude is due to a combination of loss of internal structure and other effects of homogenization such as smoothing. Another trend observed across all simulations is that the use of a nearest neighbor algorithm for the down sampling of micro-CT property maps results in reduced error. This trend was observed for simulations through heterogeneous maps derived from clinical CT data when loosely segmented, but not when tightly segmented. At the same time, simulations through tightly segmented clinical CT data consistently give reduced error before attenuation is considered. Taken together, these results highlight the importance of defining the edges of the bone layer as a discontinuous, rather than smoothly varying, interface. This is likely because when acoustic properties vary over a finite distance (rather than forming sharp, discontinuous interfaces), the resulting reflection and trasmission coefficients demonstrate frequency dependent behavior determined by the spatial rate of change of acoustic properties [132]. Ultrasound with wavelengths close to the spatial extent of the smoothly varying interface (or shorter) experience reduced reflection and increased transmission, consistent with the results shown here.

For the time-reversal simulation setup, the errors resulting from simulation through heterogeneous medium property maps derived from clinical CT data are on the same scale as the errors resulting from micro-CT data homogenised at the same resolution. For forward simulations, the error across the intracranial field is consistently larger for simulations through clinical CT data. Modelling attenuation leads to a large reduction in amplitude error across all simulations, indicating that the attenuation value measured by Pinton et al. [16] is effectively able to account for a portion of the impact of image related homogenisation. For simulations using a toneburst source, the addition of attenuation actually leads to a change in the error

in time-of-flight, increasing it for reversal simulations, and reducing it for forward simulations. The reason for the change in time-of-flight and the difference between forward and reversal simulations demonstrate opposite trends is unclear.

Generally, the error resulting from propagation through homogeneous medium maps relative to the reference simulations is similar to propagation through heterogeneous medium map, although the variation in these errors seems to be higher. This is notable, because homogeneous acoustic property maps have previously been identified as giving reduced performance when using simulations of ultrasound proapgation in time-reversal focusing [90]. This suggests that the improved performance of heterogeneous maps is not due to their ability to account for small scale acoustic property variations inside the skull layer, but instead due to the fact that a heterogeneous model will account for gross variations in acoustic properties across the skull surface. It is likely that the homogeneous models used in this chapter demonstrate improved performance because they were based on average acoustic property values across a small volume of bone tissue. When an entire skull map is assigned a single sound speed and density value, it fails to account for the variation in average property values beneath each transducer element. The fact that homogeneous models ensure a discontinuous interface between bone and soft-tissue may also contribute to their performance.

Chapter 7

General Conclusions

The unique potential of truly non-invasive, controllable stimulation of neural structures with ultrasound serves as ample justification for research into the technical challenges associated with transcranial focusing. Given the complexity, delicacy and lack of understanding of the mechanisms of this phenomenon, highly effective control of intracranial ultrasound fields will be a crucial step in ensuring the development of ultrasonic neurostimulation. Former work on the transcranial focusing of ultrasound through the skull for the purposes of high-intensity ablation indicates that numerical simulations of ultrasound propagation through the skull will be an essential tool in achieving this control. The main contribution of this thesis is a comprehensive and rigorous analysis of sources of error, which has led to the establishment of criteria that are sufficient to ensure that the model accuracy is adequate for applications in neurostimulation.

The numerical accuracy of simulations is a fundamental concern, and the k-space PSTD scheme was selected as a candidate for testing based on its potential for rapid, accurate computation of time-varying ultrasound fields. Results for tests of fundamental numerical simulation phenomena such as the bandlimited interpolant and the use of a perfectly matched layer indicated that a spatial sampling of 3 PPW is necessary for all simulations using the k-space PSTD method. Tests of numerical dispersion indicated that it would also not be a concern for the k-space PSTD numerical scheme at these sampling levels, given appropriate assignation of the reference sound speed used in the k-space correction factor. The error resulting from

the transmission across an acoustic property interface between bone and soft-tissue was shown to require more stringent sampling, and testing of this error served as the first demonstration of the impact of spatial smoothing of acoustic properties on transmission across the interface between bone and soft-tissue. Finally, staircasing of source and medium geometries was shown to be the most serious source of error, requiring 20 PPW to reduce errors below the established criteria. All tests of individual numerical errors confirmed that the k-space PSTD method performed as well as or better than the widely used 2-4 FDTD scheme.

The combined effect of numerical errors on the practical application of time-reversal was examined in the convergence testing in Chapter 3, which demonstrated that numerical accuracy can be assured for the application of simulated time-reversal. Initial convergence testing in 2D confirmed the improved performance of the k-space PSTD method relative to the 2-4 FDTD numerical scheme. Based on the results of convergence testing in 3D, a PPW value of 10, with a CFL number of 0.3 is sufficient to ensure convergence for time-reversal focusing. This is less restrictive than suggested by the tests of staircasing, indicating that staircasing has a reduced impact on either 3D simulations, or for the specific application of time-reversal. Furthermore it is clear from the results in both 2D and 3D that certain focusing metrics converge faster than others. Reconstruction of the pressure amplitude at the focus requires the highest spatial sampling to acheieve convergence, followed by focus volume, and then the spatial position of the focus. In fact, in 3D, spatial targeting appears to have almost no dependence on spatial sampling, once the variable position of the source in the forward simulations is accounted for.

While numerical accuracy can be assured given sufficient sampling and computational resources, the accuracy of acoustic property maps is dependent on the capabilities of medical imaging methods, and the effectiveness of techniques for converting those images to acoustic property maps. The numerical sensitivity analysis conducted in Chapter 4 examined the influence of the different acoustic properties and geometrical parameters of an idealised bone layer on intracranial fields. The results demonstrated that even limited changes in acoustic properties and bone layer

geometry can lead to significant changes in intracranial fields. Sound speed was observed to be the most influential acoustic property, although changes in density and attenuation can significantly affect the intracranial field. Changes in the geometry of the layer, including systematic and random errors in thickness and kernel smoothing, can also have significant effects. It was also observed that the focusing metric most sensitive to changes in the skull layer was the focus amplitude, with focus volume and spatial targeting generally less affected by alterations to the skull layer, which mirrors the results of the numerical convergence testing.

The results of the sensitivity analyses clearly underline the importance of obtaining accurate maps of acoustic properties and robust methods for the conversion of CT intensity to acoustic property values. Fortunately, the results in Chapter 5 confirm that accurate simulation through media with the approximate acoustic properties of bone tissue can be achieved, when the map of acoustic properties is accurate, and the simulated domain is stereotactically registered. The work carried out in Chapter 5 also allowed an exploration of different bone phantom materials, and methods for the creation of phantoms with well-defined geometry. In that respect, the accuracy of the registered simulations in Chapter 5 also serves as a vindication of the use of CAD techniques for the creation and registration of bone phantoms.

During the testing of both the numerical accuracy of the k-space PSTD scheme and the sensitivity of intracranial fields to acoustic property maps, the influence of spatial smoothing of acoustic properties on transmission through the skull was demonstrated repeatedly. Chapter 6 represents a more detailed examination of the importance of acoustic property maps, with greater consideration of the practical implications of deriving property maps from medical images with limited spatial resolution, including the established influence of the internal microstructure. The results demonstrate that homogenisation of acoustic property values can have a profound effect on the transmission of ultrasound in terms of both transmission amplitude and time-of-flight. These errors manifest differently depending on whether the bone sample has a trabecular layer, indicating that effects of homogenisation on time-of-flight are due to the loss of internal microstructure, while changes in ampli-

tude are due to a mix of loss of microstructure and other effects of homogenisation. Simulations through clinical CT data show comparable errors to micro-CT data homogenised at the same resolution for reversal simulations. Artificial inflation of the acoustic attenuation modelled within the bone layer can counteract transmission errors to some extent. However, it is not clear if sufficiently accurate simulation of ultrasound propagation can be achieved using acoustic property maps derived from clinical CT image data. In particular, the results presented in both Chapter 4 and Chapter 6 seem to indicate that even minor errors in the skull bone map can lead to large changes in intracranial field.

The work of this thesis also suggests certain future areas of inquiry, either building on the results and conclusions presented or developing the practical methods used. For example, the dataset of co-registered clinical and micro-CT images of skull bone could potentially be used in a range of numerical experiments. It has already been used to establish the impact of image related homogenisation, and the degree to which simulations through clinical CT data can approximate propagation through detailed property maps derived from micro-CT data. It would be interesting to examine if the method through which acoustic property values are assigned based on clinical CT intensity could be optimised to give improved performance, including for homogeneous acoustic property maps. The production of skull bone phantoms in Chapter 5 could also be extended based on previous usage of the Araldite 1302 phantom material by Clarke et al. [146]. A more realistic skull bone phantom, designed using CAD techniques and with both cortical and trabecular bone layers would be of use as a practical prop in transmission and focusing experiments. These could then be compared with registered simulations through acoustic property maps derived from a micro-CT image of the Araldite phantom, where acoustic properties could be assigned based on the known properties of Araldite and a segmentation of the micro-CT image. The final natural extension of this work would be comparison of practical transcranial propagation through real skull bone samples with registered simulations through acoustic property maps derived from micro-CT data. This could be used to definitively validate the ability to accurately simulate propa-

gation through skull bone using the k-space PSTD method, and could be used to test different methods for deriving acoustic property maps, and the impact of different image processing methods, against a practical experimental reference.

It remains unclear if ultrasound can be reliably simulated or focused with the accuracy initially envisaged when limited to acoustic property maps that are derived from clinical CT images with limited resolution, using algorithms that do not necessarily assign acoustic properties with perfect accuracy. However, this thesis represents a clear advancement in the understanding of the acoustic phenomena and practical considerations encountered when attempting to simulate ultrasound propagation through the skull. The results and conclusions presented form a useful set of guidelines that can help ensure that such simulations are accurate, and give an understanding of the potential sources of discrepancy between simulations and practical sonications. As ultrasonic neuromodulation continues to be extended to human subjects with the aid of numerical simulations of ultrasound, this knowledge forms a key component of our shared capability for exploring the limitations and potential of this extraordinary technology.

Appendix A

Numerical Methods for Ultrasound Simulation

A.1 Governing Equations

A.1.1 Fluid Equations

The primary mode of acoustic wave propagation in soft biological tissue is of compressional, or longitudinal waves, which serves as the justification for generally modelling biological tissues as fluids. Compressional wave propagation through a fluid medium involves fluctuations in physical variables including pressure, density, and acoustic particle velocity and displacement. For ultrasound at low intensities, the linear second-order wave equation for lossless, heterogeneous media [123], is written

$$\nabla \cdot \left(\frac{1}{\rho_0} \nabla p \right) = \frac{1}{\rho_0 c_l^2} \frac{\partial^2 p}{\partial t^2} \ . \tag{A.1}$$

This is derived from linearised equations for the conservation of mass

$$\frac{\partial \rho}{\partial t} = -\rho_0 \nabla \cdot \boldsymbol{v} - \boldsymbol{v} \cdot \nabla \rho_0 \ , \tag{A.2}$$

the conservation of momentum

$$\frac{\partial \boldsymbol{v}}{\partial t} = -\frac{1}{\rho_0} \nabla p \ , \tag{A.3}$$

and a linearised pressure density relationship [122]

$$p = c_l^2(\rho + \mathbf{u} \cdot \nabla \rho_0) \quad . \tag{A.4}$$

Here ρ_0 is the ambient density, c_l is longitudinal sound speed, p is the acoustic pressure, ρ is the acoustic density, and $\mathbf{u}$ and $\boldsymbol{v}$ are the acoustic particle displacement and velocity, respectively. The linearisation of Eqs.(A.2) - (A.4) is based on the assumption that the fluctuations in density caused by an acoustic wave are small compared to the ambient density, and that the particle velocity is small compared to the sound speed of the medium (the Mach number is low) [123]. Non-linear effects occur at high acoustic pressures when this assumption breaks down and the sound speed begins to vary with acoustic density, leading to harmonics and shock-wave formation.

When compressional waves propagate they lose energy to the media they are propagating in by viscous absorption and other relaxation processes, which result in conversion of acoustic energy to heat. When taking together with other sources of energy loss, such as scattering, the combined effect is termed acoustic attenuation. In biological tissues, absorption has been observed to follow a power law of the form $\alpha = \alpha_0 \omega^y$ where α_0 is the power law prefactor in Np $(\mathrm{rad/s})^{-\mathrm{y}}\ \mathrm{m}^{-1}$, and y is the dimensionless power law exponent. These are both properties of the medium, which have been experimentally quantified for various tissues [84]. The above equations do not include terms to account for power law absorption, but it can be modelled by the introduction of a loss term based on a fractional Laplacian into Eq.(A.4), which becomes

$$p = c_l^2(\rho + \mathbf{u} \cdot \nabla \rho_0 - \mathrm{L}\rho) \quad , \tag{A.5}$$

where the operator L is defined [126]

$$\mathrm{L} = -2\alpha_0 c_l^{y-1} \frac{\partial}{\partial t}(-\nabla^2)^{\frac{y}{2}-1} + 2\alpha_0 c_l^y \tan(\pi y/2)(-\nabla^2)^{\frac{y+1}{2}-1} \quad . \tag{A.6}$$

The derivation of the fractional Laplacian loss term is discussed in detail in Treeby et al. [126].

A.1.2 Elastic Equations

Although the majority of the head is soft-tissue and can be modelled accurately as a fluid, the most comprehensive description of wave propagation in the head requires a viscoelastic model due to the ability of skull bone to support shear waves [8]. Such a model describes the propagation of both shear and compressional waves in terms of the relationship between stress, strain and particle displacement. A modified Kelvin-Voigt equation can be used to describe the relationship between stress and strain in an isotropic medium [124]. Written using Einstein summation notation, this is given by

$$\sigma_{ij} = \lambda \delta_{ij} \varepsilon_{kk} + 2\mu \varepsilon_{ij} + \chi \delta_{ij} \frac{\partial}{\partial t} \varepsilon_k k + 2\eta \frac{\partial}{\partial t} \varepsilon_{ij} \ , \tag{A.7}$$

where σ is the stress tensor, ε is the strain tensor and χ and η are the compressional and shear viscosity coefficients. The first and second Lamé parameters μ and λ are related to the shear speed of sound and longitudinal speed of sound

$$\mu = c_s^2 \rho_0 \ , \qquad \lambda + 2\mu = c_l^2 \rho_0 \ . \tag{A.8}$$

Equation (A.7) can be used to model acoustic wave propagation when combined with expressions for the relationship between strain and particle displacement and the conservation of momentum

$$\varepsilon_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \ , \qquad \frac{\partial v_i}{\partial t} = \frac{1}{\rho} \frac{\partial \sigma_{ij}}{\partial x_j} \ . \tag{A.9}$$

These equations account for acoustic absorption with quadratic frequency dependence for both shear and compressional waves. They can be further adapted to model power law absorption using a fractional Laplacian, as in the fluid case [139].

A.2 Finite Difference Methods

Finite difference methods are a class of techniques that can be used to calculate the gradients of a continuous function that is discretely sampled in space and/or time

for the purpose of simulating, for example, ultrasound propagation. They have their basis in the approximation of a function $f(x)$ by a Taylor series

$$f(x \pm \Delta x) = f(x) \pm \Delta x \frac{df(x)}{dx} + \frac{\Delta x^2}{2!}\frac{d^2 f(x)}{dx^2} \pm \frac{\Delta x^3}{3!}\frac{d^3 f(x)}{dx^3} + \ldots \ . \tag{A.10}$$

By discarding the higher derivative terms, and rearranging, this gives the first-order forward and backward difference approximations, with an error term resulting from the discounting of higher derivatives

$$\frac{df(x)}{dx} = \frac{f(x+\Delta x) - f(x)}{\Delta x} + \mathrm{O}(\Delta x) \ , \quad \frac{df(x)}{dx} = \frac{f(x) - f(x-\Delta x)}{\Delta x} + \mathrm{O}(\Delta x) \ . \tag{A.11}$$

The magnitude of the error term is dependent on the higher derivative components of the original function, and is therefore lower for more smoothly varying functions. Based on Eq.(A.10), the Taylor series for $f(x+\Delta x)$ and $f(x-\Delta x)$ can be combined to get the central difference approximation:

$$\frac{df(x)}{dx} \approx \frac{f(x+\Delta x) - f(x-\Delta x)}{2\Delta x} \ . \tag{A.12}$$

Here, the even order derivatives of the Taylor series cancel, giving a reduced error term. It is apparent that Eq.(A.12) for the first derivative of $f(x)$ is consistent, as when $\Delta x \to 0$ it forms the basic definition of the derivative of a function. Finite difference schemes can also be easily visualized as interpolation between the discretised points of a function to estimate gradients, as shown in Figure A.1. This also helps demonstrate the increased accuracy of central difference over forward and backward difference schemes. More complicated schemes for gradient calculations in higher dimensions, or to calculate gradients to a greater degree of accuracy also exist, but the present examples demonstrate the fundamental principle. Finite difference approximations can be used to calculate the spatial and temporal gradients and solve the equations describing acoustic wave propagation in a time-stepping

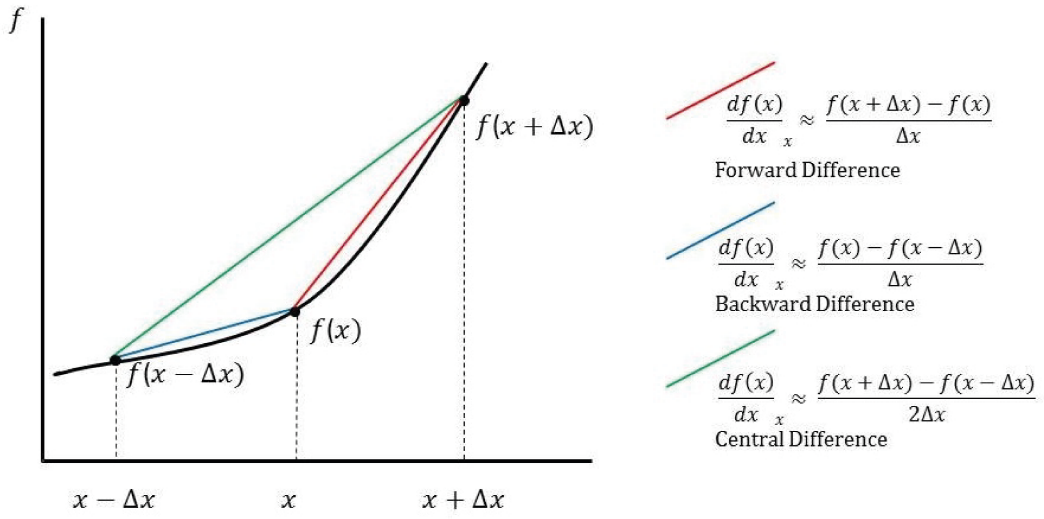

Figure A.1: Finite difference methods visualised as interpolation between points sampled from a function $f(x)$.

fashion. For example, the 1D wave equation for homogeneous media is written:

$$c_0^2 \frac{\partial^2 p}{\partial x^2} = \frac{\partial^2 p}{\partial t^2} \quad . \tag{A.13}$$

Using central-difference approximations for the second-order spatial and temporal derivatives, a discrete, time stepping formulation of Eq.(A.13) in one dimension can be written

$$c_0^2 \frac{p_{j+1}^{\tau} - 2p_j^{\tau} + p_{j-1}^{\tau}}{\Delta x^2} = \frac{p_j^{\tau+1} - 2p_j^{\tau} + p_j^{\tau-1}}{\Delta t^2} \quad . \tag{A.14}$$

Here τ is the temporal grid index, j is the spatial grid index, and Δx and Δt are the spatial and temporal discretisation steps. This equation can easily be rearranged so that, subject to knowledge of the pressure field at time τ and time $\tau - 1$, the pressure field at time $\tau + 1$ can be computed.

In terms of practical implementation, coupled first-order equations equivalent to Eq.(A.14) can be used, that calculate spatial gradients of pressure and particle velocity, and update them sequentially. On normal spatial and temporal grids, this would mean that each acoustic variable would only be calculated at every other spatial/temporal step in turn. To allow a value for each acoustic variable to be calculated for every point on the grid, the following scheme can be used. It consists

of an expression for the calculation of the updated particle velocity field based on the spatial gradient of pressure

$$v_{j+\frac{1}{2}}^{\tau+\frac{1}{2}} = -\frac{\Delta t}{\rho_0}\frac{p_{j+1}^{\tau} - p_j^{\tau}}{\Delta x} + v_{j+\frac{1}{2}}^{\tau-\frac{1}{2}} \tag{A.15}$$

followed by an expression for the calculation of the updated pressure field based on the spatial gradient of particle velocity

$$p_j^{\tau+1} = -\Delta t \rho_0 c_0^2 \frac{v_{j+\frac{1}{2}}^{\tau+\frac{1}{2}} - v_{j-\frac{1}{2}}^{\tau+\frac{1}{2}}}{\Delta x} + p_j^{\tau} \ . \tag{A.16}$$

Here, values for acoustic particle velocity are calculated at spatial points corresponding to a staggered grid with a spatial offset of Δx and a temporal offset of Δt compared to the main grid. This permits calculation of values of both variables for each point on the spatial and temporal grids, and increases the accuracy of the scheme by effectively halving the spatial and temporal discretisation [121]. This is the 2-2 FDTD scheme, described in Section 2.1.3.

As discussed above, finite difference methods have an inherent error due to the truncation of the Taylor series, and the resulting inexact calculation of both temporal and spatial gradients. An expression for this error may be obtained by comparing analytical solutions for the propagation of a wave with solutions obtained using finite difference approximations for either the second-order derivatives in Eq.(A.14), or the two first-order derivatives in Eqs.(A.15) & (A.16).

The result is an unphysical dependence of the sound speed on the spatial and temporal discretisation and on the wavenumber k, known as numerical dispersion. For the 2-2 FDTD scheme described in Eq.(A.14), the resulting expression for sound speed as a function of wavenumber on the grid is obtained

$$c(k) = \frac{\mathrm{sinc}(k\Delta x/2)}{\mathrm{sinc}(c_0 k \Delta t/2)} c_0 \ , \tag{A.17}$$

where the sound speed in the exact case should be c_0 for all wavenumbers. This error leads to a requirement for fine spatial and temporal discretisation relative to

the wavelength and period, respectively, of the acoustic frequencies used in the simulation.

The stability conditions for finite difference methods vary according to the scheme used, and the number of dimensions being simulated. The 2-2 FDTD scheme in 1D has a requirement that CFL ≤ 1, where CFL is the Courant-Friedreichs-Lewy number, defined

$$\mathrm{CFL} = \frac{c_{\max}\Delta t}{\Delta x} \quad . \tag{A.18}$$

It should be noted that for CFL = 1, the error in sound speed resulting from the spatial and temporal discretisation shown in Eq.(A.17) will cancel, rendering the second-order finite difference method exactly accurate for homogenous media. However, simulations in higher dimensions have more restrictive CFL requirements, rendering simulations where CFL $= 1$ unstable.

A.3 Pseudospectral Methods and k-Space Correction

Knowledge of the distribution of pressure and particle velocity across the spatial grid at each time step can be used to create a more accurate time-stepping numerical scheme without the need for finer discretisation, by implementing spectral methods for the calculation of spatial gradients. The following description of PSTD methods is based on Trefethen [109], which provides an excellent exploration of the application of spectral methods.

Spectral methods involve approximation of the distributions of pressure and velocity across the discrete grid through summations of global basis functions. A Fourier basis can be used, which are sets of complex exponentials of the form $n_k e^{ikx}$ where k is the wavenumber and n_k is the basis function weight. These are computed using a discrete version of the Fourier transform. The spatial derivative of each exponential may then be easily calculated as $n_k ike^{ikx}$ and an inverse Fourier transform performed to give the spatial gradient of the approximating sum of basis functions at each point.

By replacing the finite difference calculations for spatial gradients in Eq.(A.14)

with spectral calculation of gradients, the pseudospectral time-domain (PSTD) method is obtained [121]

$$c_0^2 F^{-1}\Big[-ik_x F[p^\tau]\Big] = \frac{p^{\tau+1} - 2p^\tau + p^{\tau-1}}{\Delta t^2} \tag{A.19}$$

Here, F and F^{-1} are the discrete Fourier transform and inverse Fourier transform, respectively. k_x refers to the discrete, bounded set of wavenumbers for which the discrete Fourier transform produces coefficients, which correspond to a continuous function that agrees with the discretely sampled pressure field at each grid point - the Band-Limited Interpolant (BLI).

The PSTD method removes the error from the finite difference approximation of spatial gradients, subject to the conditions of the BLI. However, error from the finite-difference approximation of the temporal step remains, giving a numerical dispersion dependent on the temporal discretisation

$$c(k_x) = \frac{c_0}{\mathrm{sinc}(c_0 k_x \Delta t/2)} \tag{A.20}$$

The k-space corrected PSTD method demonstrated by Tabei et al. [121] can be used to correct this for homogeneous media. The k-space method takes advantage of the fact that, in homogeneous media, the error in the PSTD scheme resulting from the discretisation of the time step can be described analytically (as per Eq.(A.20)). The known value of this error may then be used to introduce a correction factor into the spectral calculation of the spatial gradients. For simulations in one dimension, the discretized, second-order k-space model is written as

$$\frac{p^{\tau+1} - 2p^\tau + p^{\tau-1}}{\Delta t^2} = c_0^2 F^{-1}\Big[-k_x^2 \mathrm{sinc}^2(c_{\mathrm{ref}} k_x \Delta t/2) F[p^\tau]\Big] \tag{A.21}$$

The term $\mathrm{sinc}(c_{ref} k_x \Delta t/2) \equiv \kappa$ is known as the k-space correction factor, and corrects the error generated by the finite difference approximation of the temporal gradient. It has been shown to be exactly accurate for homogeneous media independent of the size of the time steps. This makes it preferable to both FDTD and PSTD

methods, which require finer temporal and (in the case of FDTD) spatial discretisation to attain the same level of error, at the expense of simulation efficiency. This is especially useful for long, large-scale simulations where numerical errors will accumulate and the decrease in computational efficiency that would result from finer spatial and temporal sampling will be compounded.

Appendix B

Simulation Parameters

This appendix contains the computational details for the simulations carried out in each chapter of this thesis. All simulations were carried out using the open source k-Wave toolbox [118]. The toolbox includes k-space, PSTD, and 2-4 FDTD codes for the time-domain simulation of acoustic fields. For every simulation, the size of the spatial and temporal grids used is quoted. When simulations are carried out with respect to some real world geometry, spatial and temporal sampling criteria are quoted. Otherwise a CFL value is used to describe the ration of temporal to spatial discretisation. For simulations in 2D & 3D, a simulation runtime is quoted to give an idea of the computational burden, however, this number is approximate, and will vary between simulation sessions, computing hardware and versions of k-Wave.

In terms of computing resources, 1D simulations were carried out in the MATLAB environment on a Dell Precision T1700 with an Intel Xeon E3-1240 3.40GHz CPU and 16 GB of RAM running Windows 10 64 bit. 2D simulations, including elastic simulations, were carried out in the MATLAB environment with CUDA hardware acceleration on a Dell PowerEdge R730 compute server with 2 $\times$ 6-core Xeon E5-2620 2.4GHz CPUs, 64 GB of 1866MHz memory, on an Nvidia Titan X GPU with 3072 CUDA cores and 12 GB of memory. Large 3D fluid simulations were run using the MPI version of k-Wave on the IT4I Salomon supercomputing cluster [133]. Each 3D simulation was carried out on a compute node with an 8-core Intel Xeon E5-4627v2 3.3 GHz CPU, and 256 GB of RAM. Computational resources were provided by the IT4Innovations Centre of Excellence project

(CZ.1.05/1.1.00/02.0070), funded by the European Regional Development Fund and the national budget of the Czech Republic via the Research and Development for Innovations Operational Programme, as well as Czech Ministry of Education, Youth and Sports via the project Large Research, Development and Innovations Infrastructures (LM2011033).

B.1 Numerical Accuracy of k-Space Simulations

Individual tests of the impact of specific sources of numerical accuracy consisted of a range of simulations in 1D and 2D. Apart from some exceptions detailed below, 1D tests were carried out on a spatial grid of 4096 grid points, 2D tests were carried out on a 1024×1024 grid and the CFL was set to 0.3, as defined in Eq. 2.12. The 2D simulations used in testing of the PML had a runtime of approximately 30 seconds. The tests of source staircasing had a runtime of 7 minutes, while the tests of medium staircasing across a single interface had a simulation runtime of 14 minutes.

The first exception to the rules described above is in the testing of the PML in 1D. As discussed in Section 2.2.3 the 2-2 FDTD numerical scheme was used, with a CFL number of 1 to prevent numerical dispersion. The other exception is the second test examining the impact of medium staircasing, described in Section 2.2.6. This used an oversized spatial grid to allow progressive resampling of the acoustic property map without altering the spatial or temporal sampling. These large 2D simulations used a spatial grid of 3780×3780, and ran for 258462 time steps, with a total runtime of 10.6 h.

B.2 Convergence Testing

The general convergence testing protocol is outline in Section 3.1.2. It involved multiple forward simulations at variegated spatial and temporal sampling including those sampling the frequency of interest at the Nyquist limit of 2 PPW, combined with reversal simulations at the highest sampling feasible. In 2D, seventy-six different forward grids were used, for simulations carried out at 250 kHz and 500 kHz, using both FDTD and k-space PSTD numerical schemes, each with a corresponding reversal simulation. In 3D, thirty-one forward simulations were carried out,

		Grid Size	Δx [μm]	Δt [ns]	Timesteps	Runtime [hr:min]
2D	250 kHz	108×108	3100	292	138	0:01
		3096×3096	62	6	73375	0:48
	500 kHz	162×162	1600	146	2658	0:01
		3096×3096	62	6	68144	0:44
3D	500 kHz	144×144×144	1600	146	2653	0:04
		1024×1024×1024	190	18	22718	112:18

Table B.1: Simulation parameters used for Convergence Testing in 2D and 3D in Chapter 3.

	Grid Size	Δx [μm]	Δt [ns]	Time steps	Runtime [hr:min]
Fluid 3D	512×512×1024	145.4	5.1	33502	19:53
Elastic 2D	1728×2592	57.2	1.0	169706	3:47

Table B.2: Simulation parameters used for the sensitivity analysis. Fluid simulations were conducted with a minimum spatial sampling of 10 points per wavelength (PPW), and elastic simulations with 25 PPW.

at 500 kHz only using the k-space PSTD scheme only. Simulation parameters for convergence testing are shown in Table B.1. Numbers quoted are for the forward simulations with the coarsest spatial and temporal discretisation, and the reversal simulations, which were the same parameters used as for the most finely sampled forward simulation.

B.3 Sensitivity Analysis

Simulations for the numerical sensitivity analyses were carried out using both the elastic and fluid simulation models in k-Wave. Elastic simulations were all 2D, while fluid simulations were 3D. The parameters used in the different simulations are shown in Table B.2. The spatial and temporal discretisations used conform to the previously established criteria to ensure numerical accuracy for transcranial simulation, with a minimum spatial sampling step of ten spatial points per wavelength (PPW) at 1 MHz for fluid simulations, and 25 PPW for elastic simulations. A reference simulation and 567 test simulations were carried out for each frequency tested.

	Simulation	Grid size	Time steps	Runtime [hr:min]
	Source projection	720×720×400	25465	7:31
Parametric	Scan projection	450×450×510	18010	6:17
	Registered simulation	720×720×1000	41046	46:12
Anatomical	Scan forward projection	450×450×240	16421	2:13
	Scan reverse projection	450×450×450	31452	6:16
	Registered simulation	720×720×1000	41095	46:30

Table B.3: Simulation parameters used for the registered simulations in Chapter 5.

B.4 Registered Simulations

A range of fluid 3D simulations were carried out to acheive the final registered simulations through Araldite and VeroBlack phantoms. The first simulation carried out was a projection of the free field scan back to a plane 10 mm from the surface of the transducer position, for use as a source term in later simulations.

Then a simulations were carried out to determine the 3D ultrasound fields during the practical experiments and compare them with fields predicted from propagation through simulated phantom models. First, there were projections of the scan planes recorded following transmission through the phantoms across the 3D volumes shown in Fig. 5.5. This consisted of a single forward projection for scans following transmission through the parametric phantom, and both forward and reversal projection for scans following transmission through the anatomical phantoms. Finally, full size registered simulations were carried out using source plane computed from projection of the free field scan through virtual phantom models. All simulations used a spatial discretisation of 0.1 mm, corresponding to 15 spatial PPW for the 1 MHz ultrasound being simulated, and a temporal discretisation of 5 ns. The simulation specific parameters, including the simulation runtime and the sizes of the spatial and temporal grids are shown in Table B.3.

Bone Sample	Grid Size	Lossless Runtime [hr:min]	Attenuating Runtime [hr:min]	Elastic 2D Runtime [hr:min]
A: Frontal	248×486×274	7:37	10:46	2:26
A: Parietal	342×486×368	14:12	19:53	2:47
B: Frontal	278×486×252	7:41	10:42	2:00
B: Parietal	304×486×340	9:21	13:11	2:18
B: Temporal	266×486×256	6:54	9:33	2:28

Table B.4: Simulation parameters used for examining propagation through acoustic property maps derived from micro-CT data, homogenised micro-CT data, and co-registered clinical-CT data in Chapter 6. The elastic 2D simulations were the size of the first two dimensions of the 3D grids.

B.5 Homogenisation and Imaging

For each simulation setup, forward and reverse, CW and TB sources, a total of 94 simulations were carried out for each bone sample. These comprised 47 lossless simulations and 47 corresponding attenuating/absorbing simulations. The 48 consisted of one reference simulation based on micro-CT data, thirty-eight simulations examining progressive downsampling using either nearest neighbour or linear interpolation, and eight simulations based on clinical CT data, four heterogeneous, four homogeneous. All simulations used a spatial step of 50 microns, which corresponds to 30 spatial PPW for the 1 MHz ultrasound being simulated. A temporal step of 0.65 ns was used, and the simulations were run for a total of 97832 timesteps. Elastic simulations were based on sections through the 3D simulation setup, and so are sized the same as the first two dimensions of the 3D fluid simulations, use the same spatial and temporal discretisations, and run for the same number of timesteps. The different spatial grids used for each bone sample, and the corresponding runtime, are given in Table B.4.

Appendix C

Raw Homogenisation Results

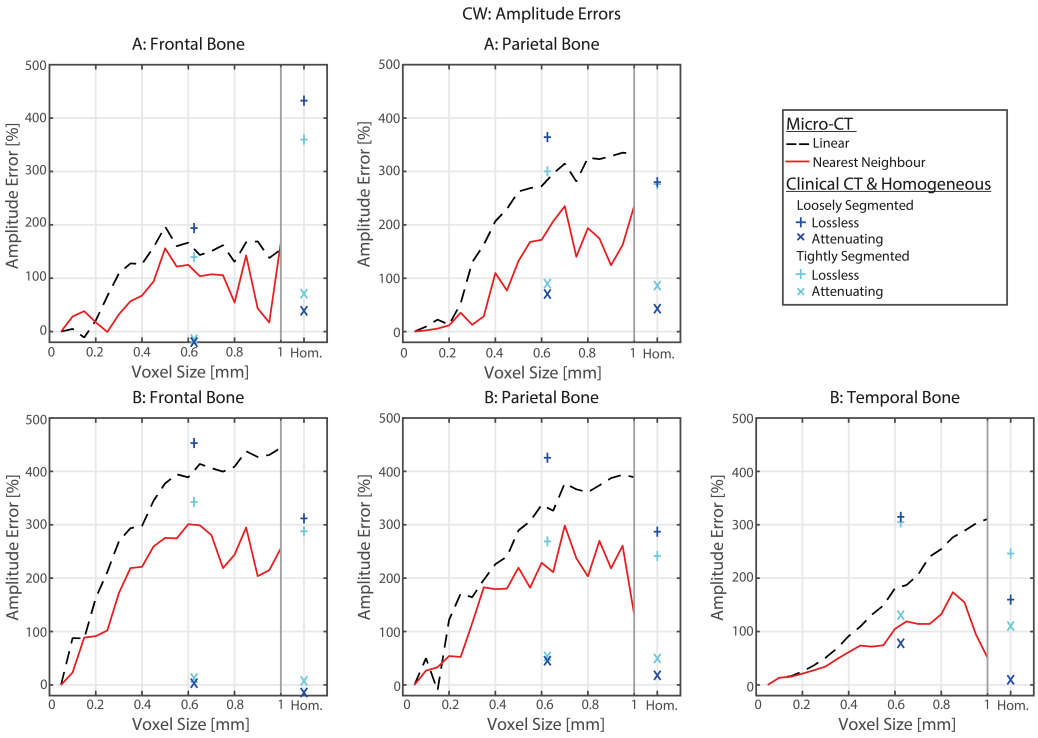

Figure C.1: Error in amplitude across the transducer surface for continuous wave time-reversal simulations as a function of homogenisation voxel size applied to the acoustic property maps for each different bone sample.

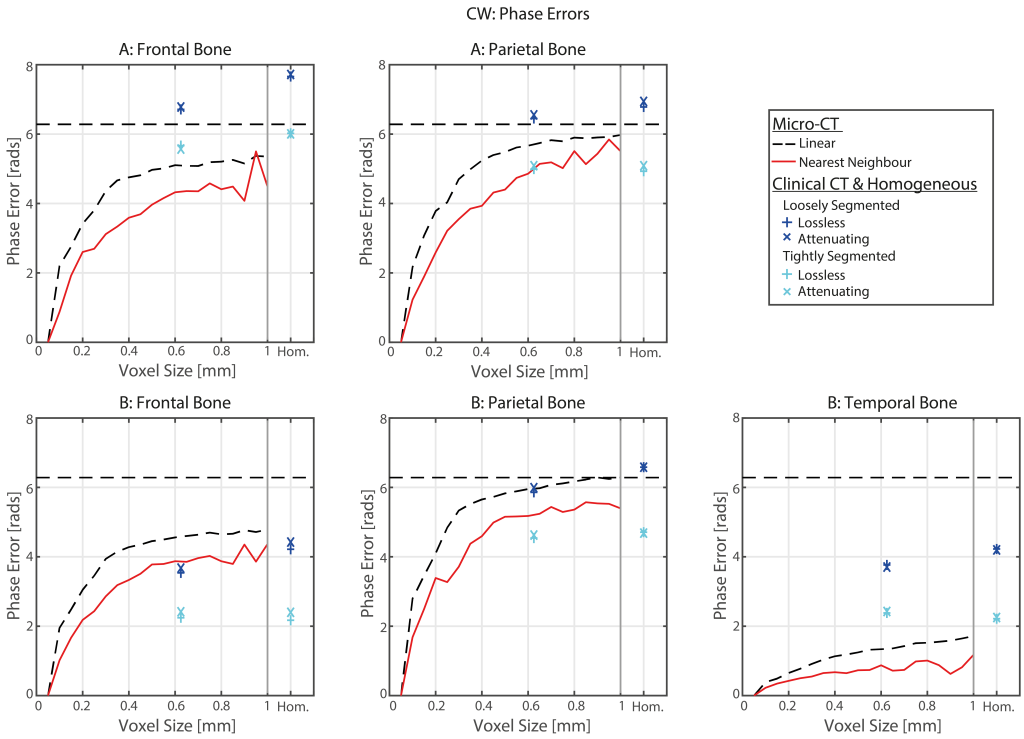

Figure C.2: Error in phase across the transducer surface for continuous wave time-reversal simulations as a function of homogenisation voxel size applied to the acoustic property maps for each different bone sample.

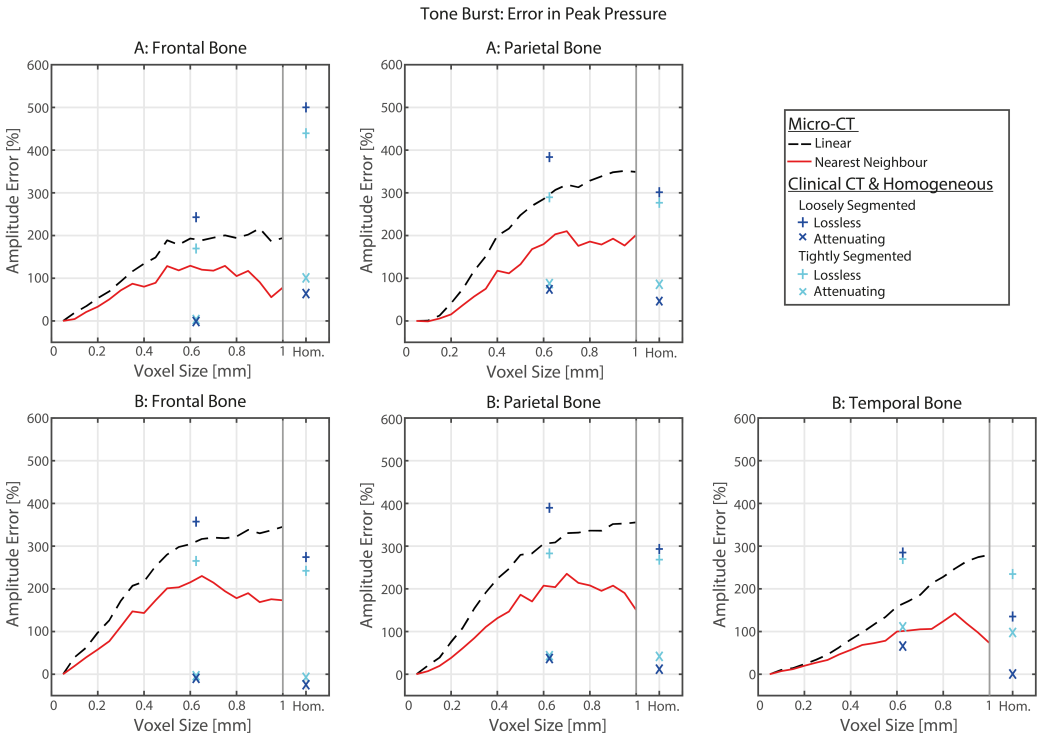

Figure C.3: Error in amplitude across the transducer surface for toneburst time-reversal simulations as a function of homogenisation voxel size applied to the acoustic property maps for each different bone sample.

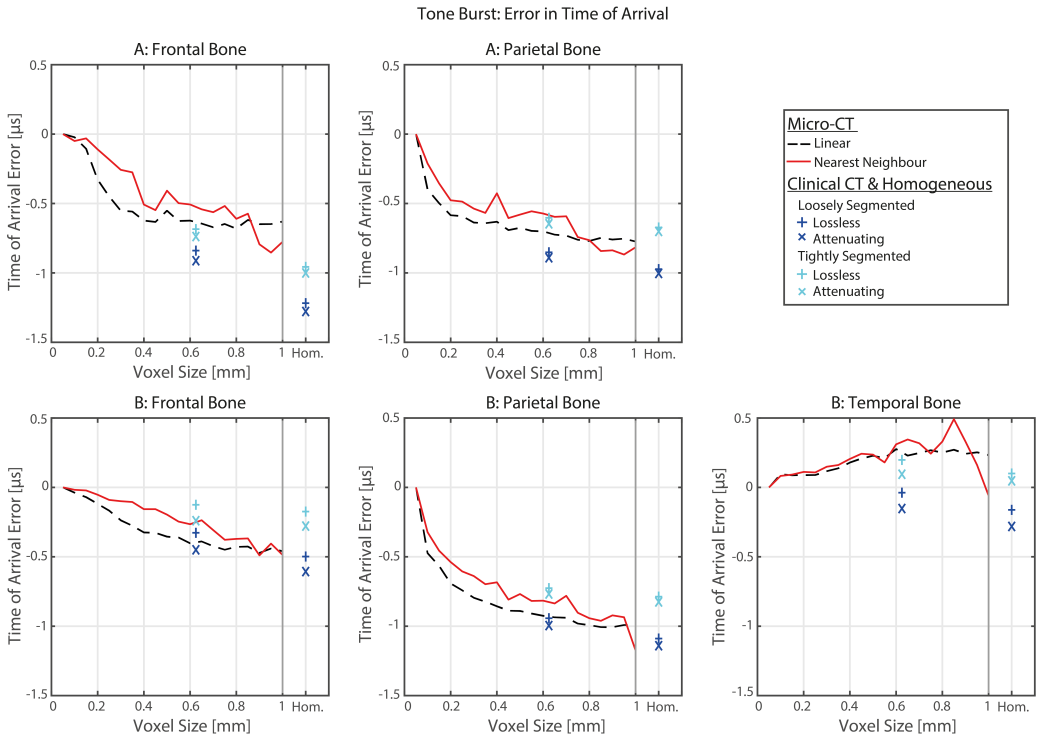

Figure C.4: Error in time of flight across the transducer surface for toneburst time-reversal simulations as a function of homogenisation voxel size applied to the acoustic property maps for each different bone sample.

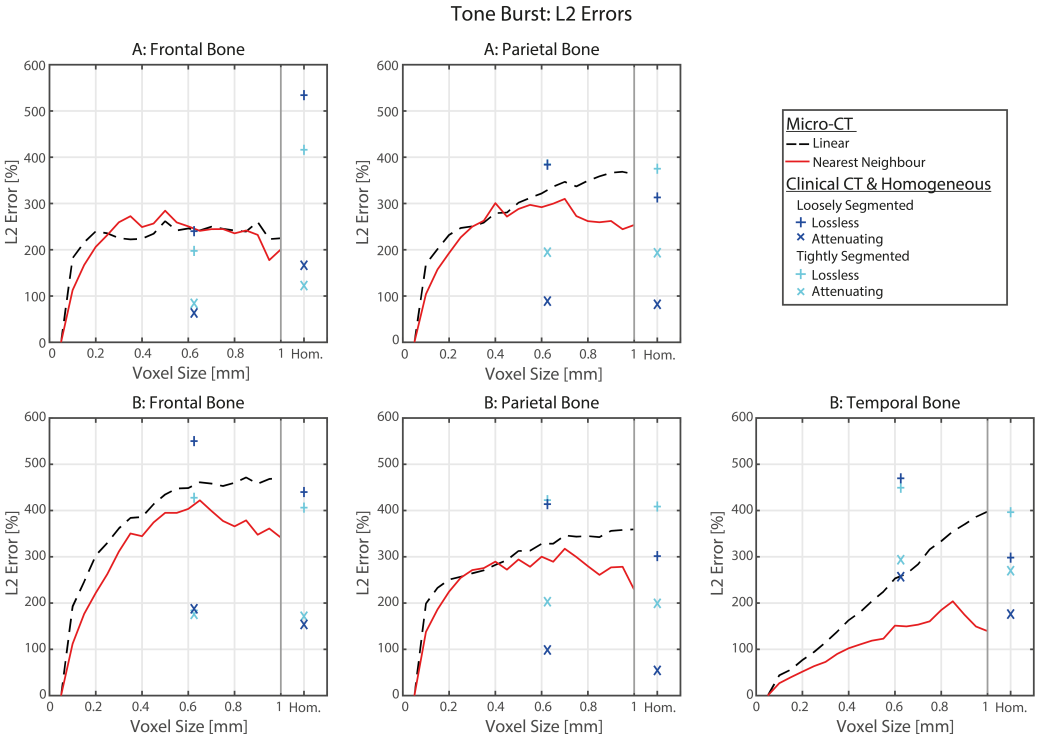

Figure C.5: L2 error in the time-varying signal across the transducer surface for toneburst time-reversal simulations as a function of homogenisation voxel size applied to the acoustic property maps for each different bone sample.

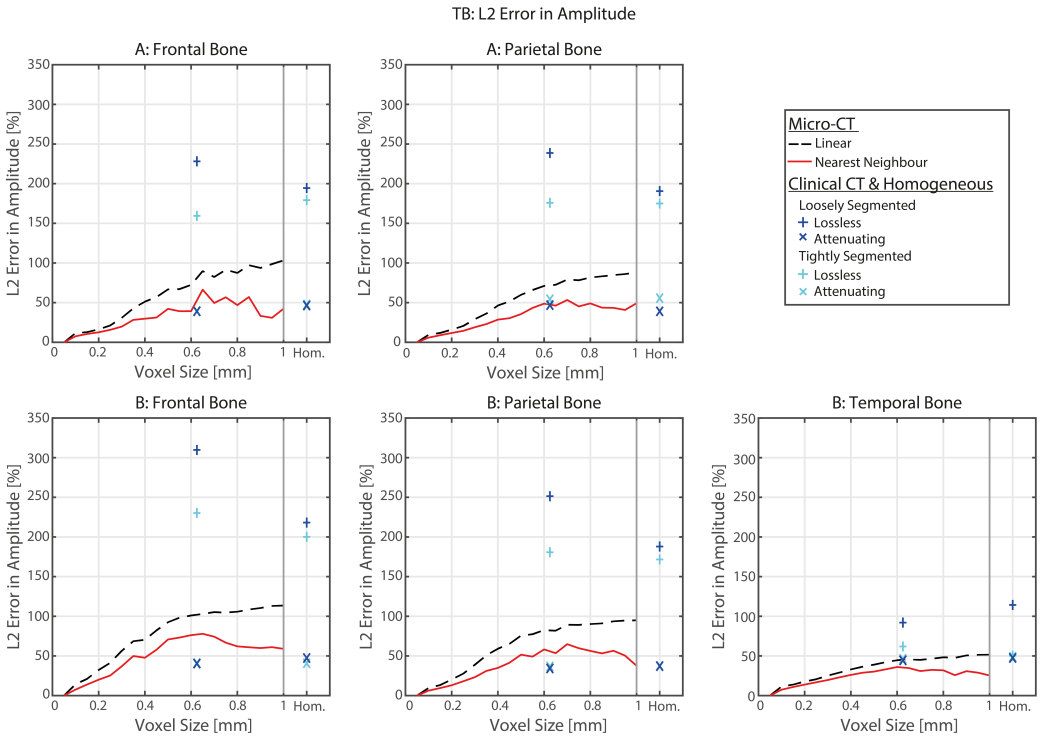

Figure C.6: L2 error in the peak pressure amplitude across the intracranial scan plane for forward propagation simulations as a function of homogenisation voxel size applied to the acoustic property maps for each different bone sample.

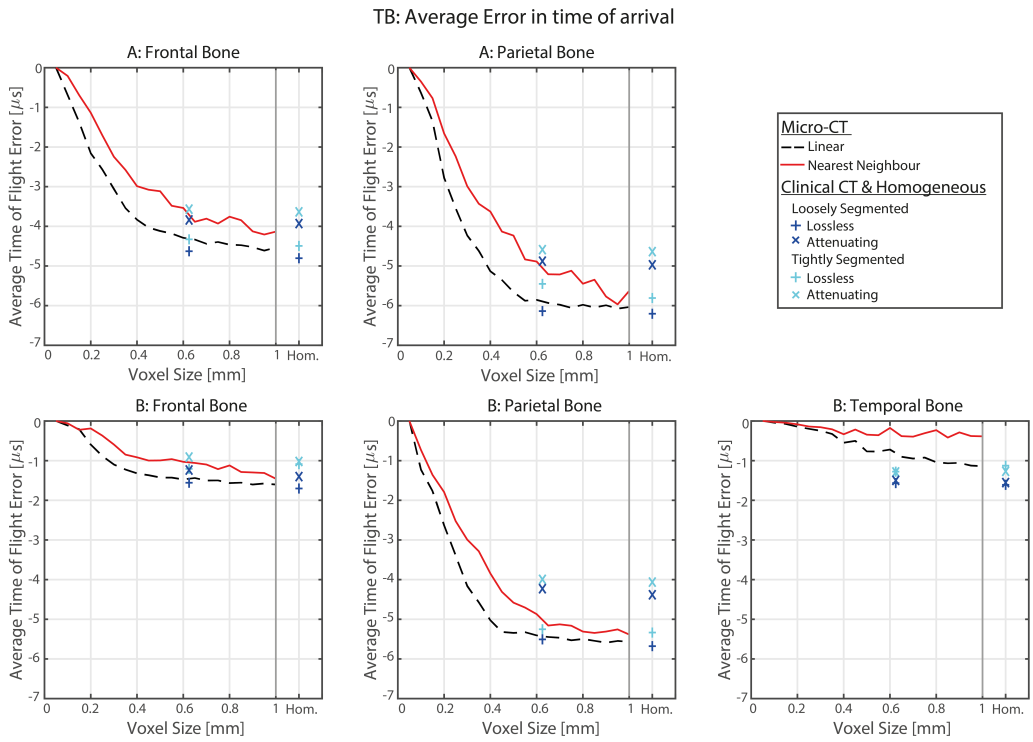

Figure C.7: Average error in time of flight across the intracranial scan plane for forward propagation simulations as a function of homogenisation voxel size applied to the acoustic property maps for each different bone sample.

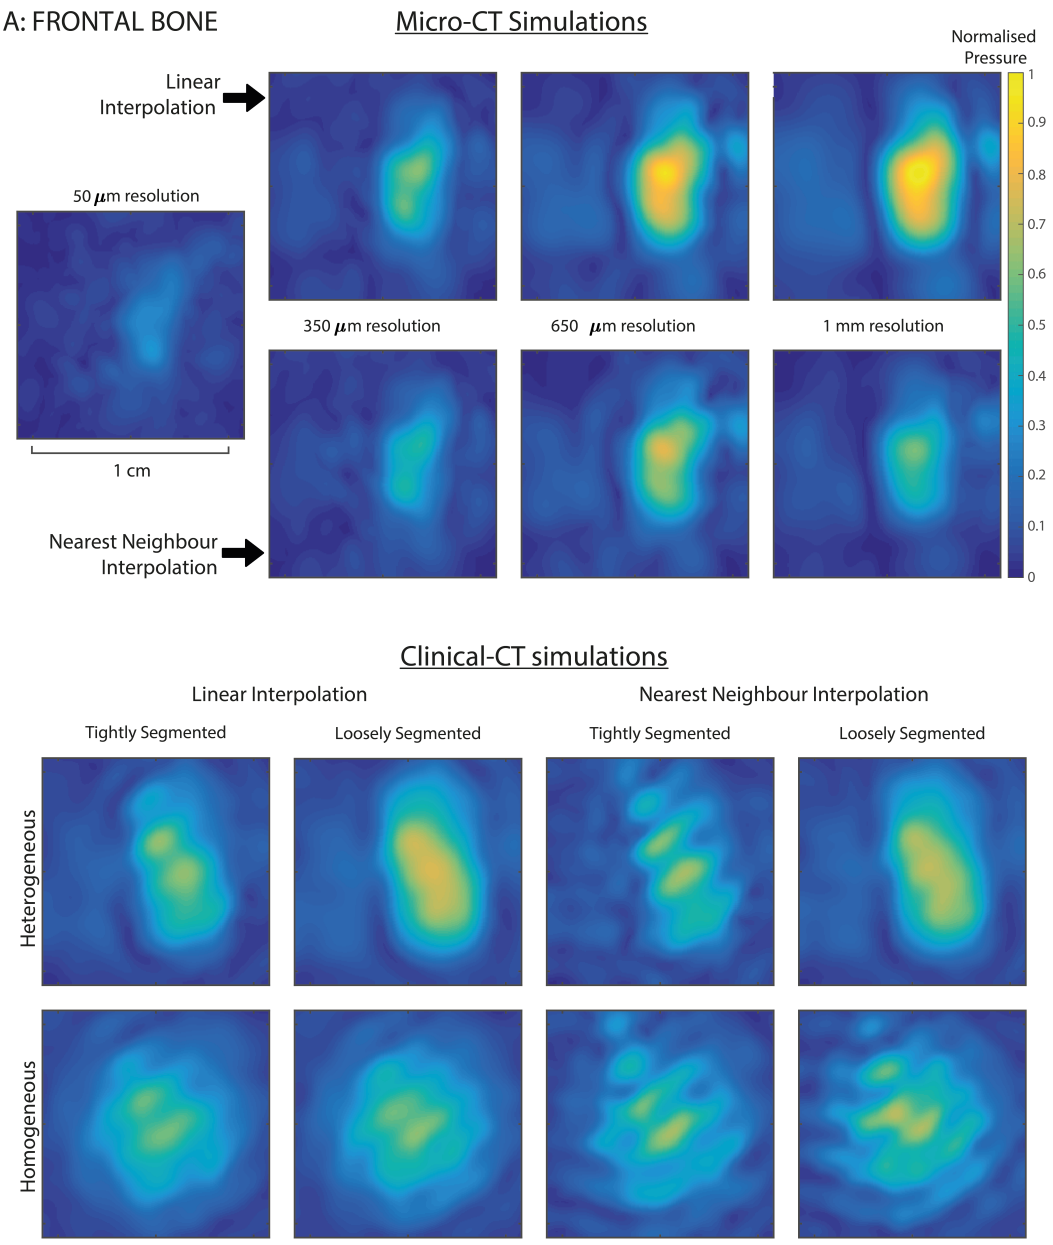

Figure C.8: Peak pressure field across the intracranial scan planes for non-attenuating and non-absorbing simulations through the frontal bone sample from skull A.

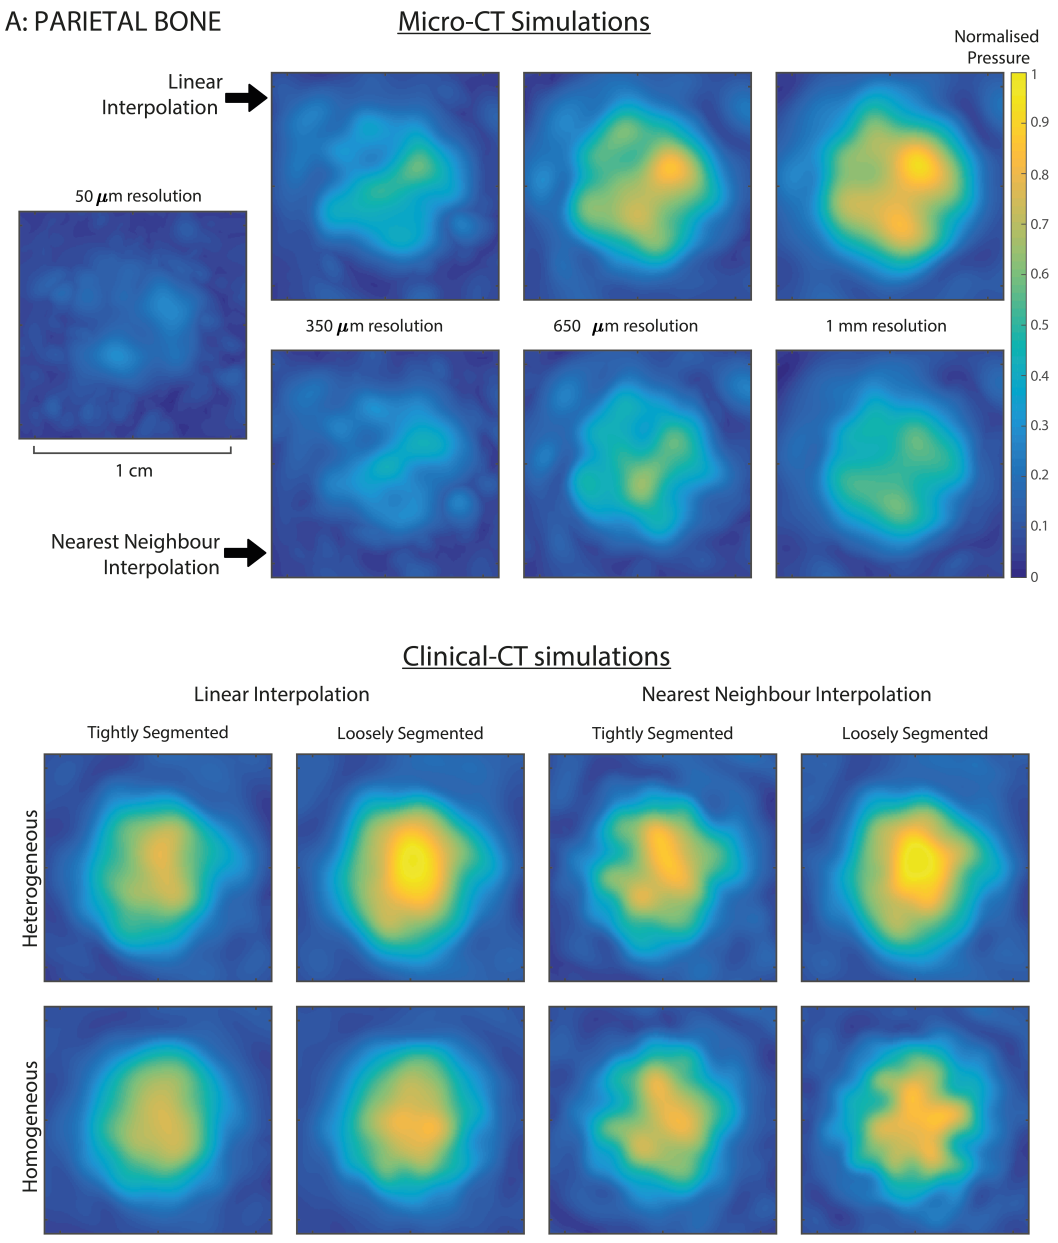

Figure C.9: Peak pressure field across the intracranial scan planes for non-attenuating and non-absorbing simulations through the parietal bone sample from skull A.

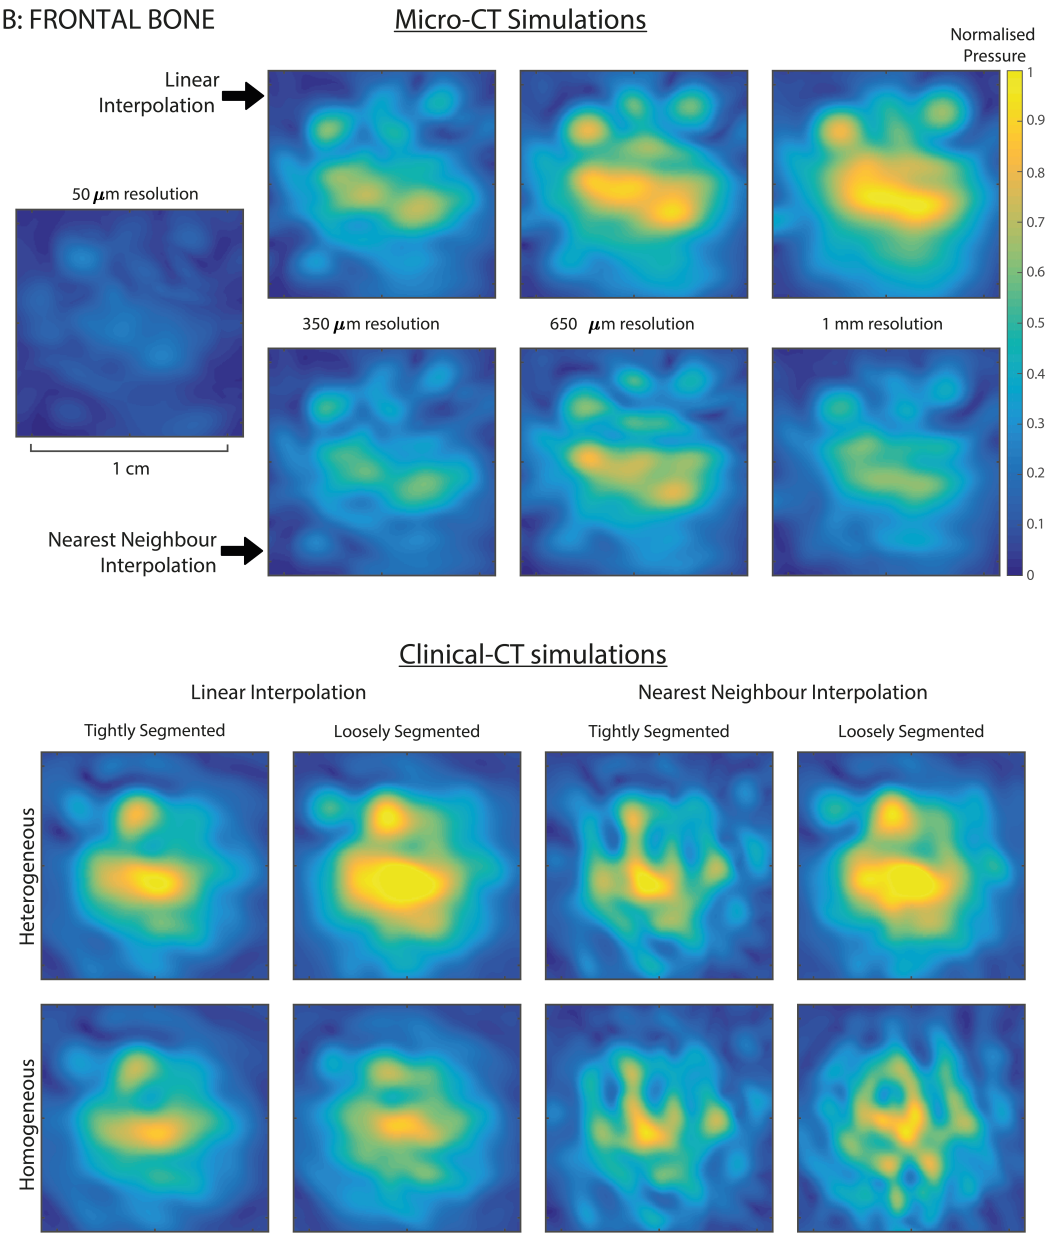

Figure C.10: Peak pressure field across the intracranial scan planes for non-attenuating and non-absorbing simulations through the frontal bone sample from skull B.

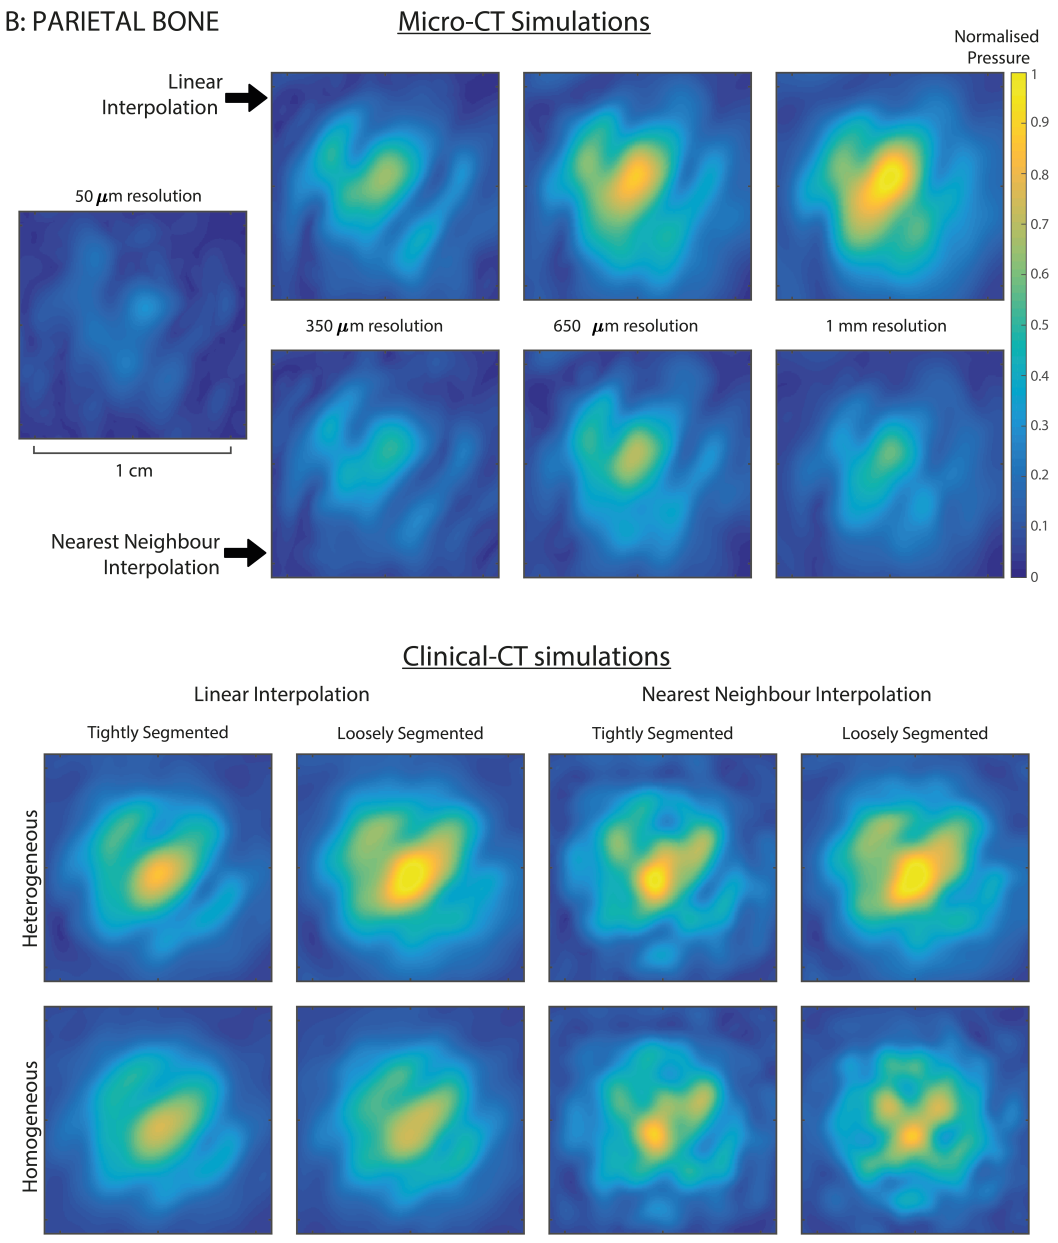

Figure C.11: Peak pressure field across the intracranial scan planes for non-attenuating and non-absorbing simulations through the parietal bone sample from skull B.

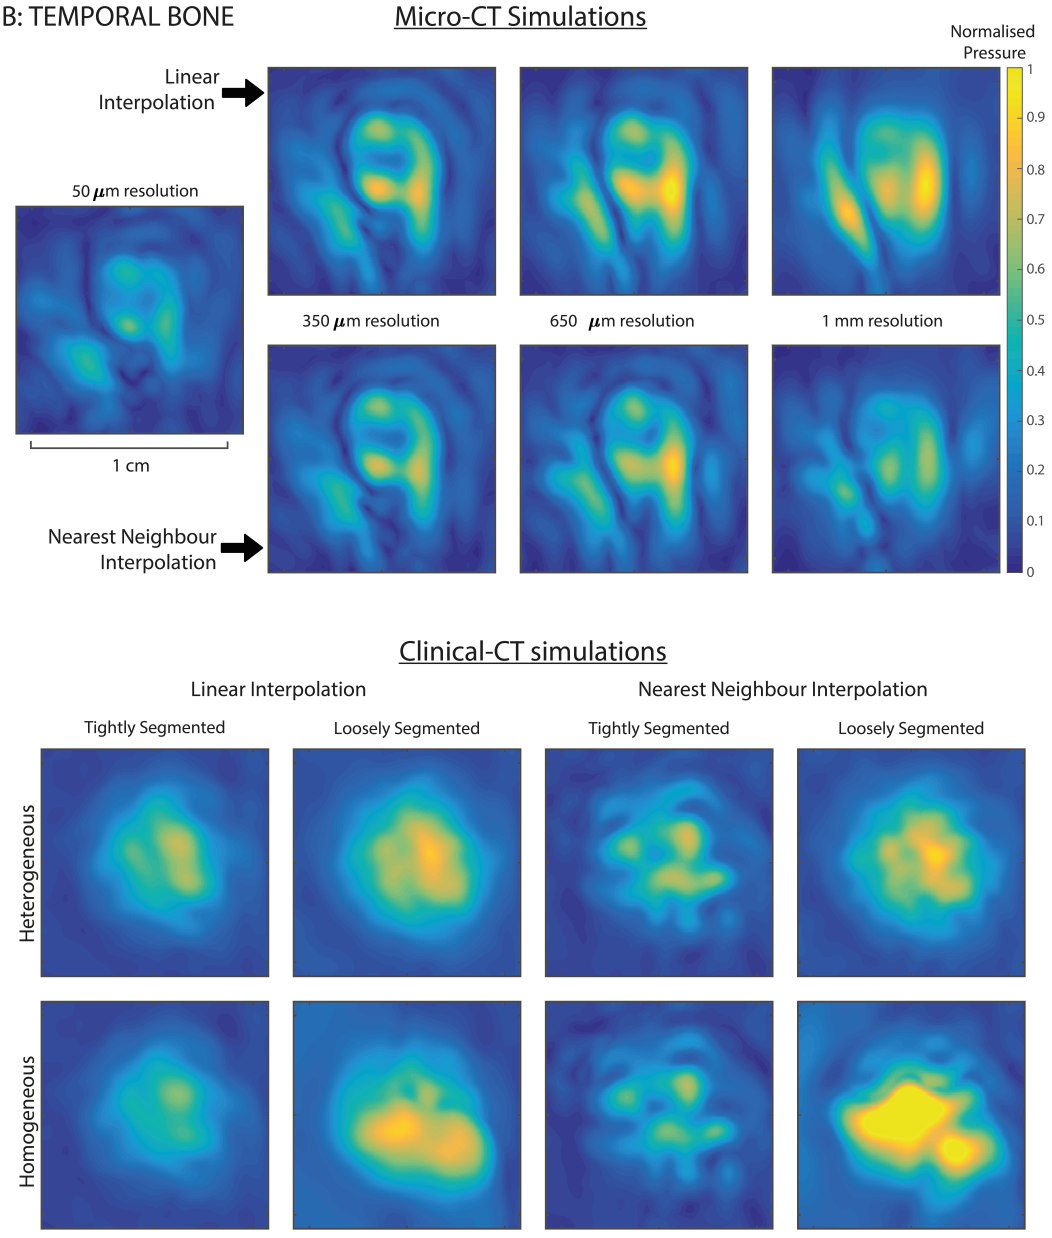

Figure C.12: Peak pressure field across the intracranial scan planes for non-attenuating and non-absorbing simulations through the temporal bone sample from skull B.

Bibliography

[1] G. Leisman and R. Melillo. The basal ganglia: Motor and cognitive relationships in a clinical neurobehavioral context. *Rev. Neurosci.*, 24(1):9–25, 2013.

[2] H. Gray. *Anatomy of the Human Body*. Bartleby, New York, 20th edition, 2000.

[3] R. Drake, A. W. Vogl, and A. W. M. Mitchell. *Gray's Anatomy for Students*. Churchill Livingstone, Philadelphia, 2nd edition, 2010.

[4] C. W. Connor, G. T. Clement, and K. Hynynen. A unified model for the speed of sound in cranial bone based on genetic algorithm optimization. *Phys. Med. Biol.*, 47(22):3925–44, nov 2002.

[5] T. D. Mast. Empirical relationships between acoustic parameters in human soft tissues. *Acoust. Res. Lett. Online*, 1(2):37–42, oct 2000.

[6] J-F. Aubry, M. Tanter, M. Pernot, J. L. Thomas, and M. Fink. Experimental demonstration of noninvasive transskull adaptive focusing based on prior computed tomography scans. *J. Acoust. Soc. Am.*, 113(1):84–93, jan 2003.

[7] S. Pichardo, V. W. Sin, and K. Hynynen. Multi-frequency characterization of the speed of sound and attenuation coefficient for longitudinal transmission of freshly excised human skulls. *Phys. Med. Biol.*, 56(1):219–250, jan 2011.

[8] G. T. Clement, P. J. White, and K. Hynynen. Enhanced ultrasound transmission through the human skull using shear mode conversion. *J. Acoust. Soc. Am.*, 115(3):1356–1364, mar 2004.

[9] P. J. White, G. T. Clement, and K. Hynynen. Longitudinal and shear mode ultrasound propagation in human skull bone. *Ultrasound Med. Biol.*, 32(7):1085–1096, jul 2006.

[10] P. A. Hasgall, F. Di Gennaro, C. Baumgartner, E. Neufeld, M. C. Gosselin, D. Payne, A. Klingenböck, and N. Kuster. IT'IS Database for thermal and electromagnetic parameters of biological tissues, 2015.

[11] A. Morel. *Stereotactic Atlas of the Brain and Basal Ganglia*. Informa Healthcare, New York, New York, USA, 1st edition, 2007.

[12] O. Naor, S. Krupa, and S. Shoham. Ultrasonic neuromodulation. *J. Neural Eng.*, 13(3):031003, jun 2016.

[13] F. J. Fry and J. E. Barger. Acoustical properties of the human skull. *J. Acoust. Soc. Am.*, 63(5):1576–1590, may 1978.

[14] D. Nikitichev, E. Alles, S. Noimark, W. Xia, and A. Desjardins. Acoustic properties of 3D printed materials and their application. *(Under Prep.*, 2016.

[15] T. Deffieux and E. E. Konofagou. Numerical study of a simple transcranial focused ultrasound system applied to blood-brain barrier opening. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 57(12):2637–2653, dec 2010.

[16] G. Pinton, J-F. Aubry, E. Bossy, M. Muller, M. Pernot, and M. Tanter. Attenuation, scattering, and absorption of ultrasound in the skull bone. *Med. Phys.*, 39(1):299–307, dec 2011.

[17] A. Gustavsson, M. Svensson, F. Jacobi, C. Allgulander, J. Alonso, E. Beghi, R. Dodel, M. Ekman, C. Faravelli, L. Fratiglioni, B. Gannon, D. H. Jones, P. Jennum, A. Jordanova, L. Jönsson, K. Karampampa, M. Knapp, G. Kobelt, T. Kurth, R. Lieb, M. Linde, C. Ljungcrantz, A. Maercker, B. Melin, M. Moscarelli, A. Musayev, F. Norwood, M. Preisig, M. Pugliatti, J. Rehm, L. Salvador-Carulla, B. Schlehofer, R. Simon, H-C. Steinhausen, L. J. Stovner, J-M. Vallat, P. V. den Bergh, Ji. van Os, P. Vos, W. Xu, H-U.

Wittchen, B. Jönsson, and J. Olesen. Cost of disorders of the brain in Europe 2010. *Eur. Neuropsychopharmacol.*, 21(10):718–779, oct 2011.

[18] H. U. Wittchen, F. Jacobi, J. Rehm, A. Gustavsson, M. Svensson, B. Jönsson, J. Olesen, C. Allgulander, J. Alonso, C. Faravelli, L. Fratiglioni, P. Jennum, R. Lieb, A. Maercker, J. van Os, M. Preisig, L. Salvador-Carulla, R. Simon, and H-C. Steinhausen. The size and burden of mental disorders and other disorders of the brain in Europe 2010. *Eur. Neuropsychopharmacol.*, 21(9):655–679, sep 2011.

[19] A. M. Lozano and N. Lipsman. Probing and Regulating Dysfunctional Circuits Using Deep Brain Stimulation. *Neuron*, 77(3):406–424, feb 2013.

[20] H. P. Rang, M. Dale, M. Ritter, R. Flower, and G. Henderson. *Pharmacology*. Elsevier/Churchill Livingstone, Edinburgh, 7th edition, 2012.

[21] A. R. Crossman and D. Neary. *Neuroanatomy: an Illustrated Colour Text*. Churchill Livingstone, London, 4th edition, 2010.

[22] C. Poon, D. McMahon, and K. Hynynen. Noninvasive and targeted delivery of therapeutics to the brain using focused ultrasound. *Neuropharmacology*, 4(4):519–526, feb 2016.

[23] A. Y. Ammi, T. D. Mast, I-H. Huang, T. A. Abruzzo, C-C. Coussios, G. J. Shaw, and C. K. Holland. Characterization of Ultrasound Propagation Through Ex-vivo Human Temporal Bone. *Ultrasound Med. Biol.*, 34(10):1578–1589, oct 2008.

[24] W. J. Elias, D. Huss, T. Voss, J. Loomba, M. Khaled, E. Zadicario, R. C. Frysinger, S. A. Sperling, S. Wylie, S. J. Monteith, J. Druzgal, B. B. Shah, M. Harrison, and M. Wintermark. A Pilot Study of Focused Ultrasound Thalamotomy for Essential Tremor. *N. Engl. J. Med.*, 369(7):640–648, aug 2013.

[25] M. D. Johnson, S. Miocinovic, C. C. McIntyre, and J. L. Vitek. Mechanisms and targets of deep brain stimulation in movement disorders. *Neurotherapeutics*, 5(2):294–308, apr 2008.

[26] P. J. Karas, C. B. Mikell, E. Christian, M. A. Liker, and S. A. Sheth. Deep brain stimulation: a mechanistic and clinical update. *Neurosurg. Focus*, 35(5):E1, nov 2013.

[27] W. M. Grill, A. N. Snyder, and S. Miocinovic. Deep brain stimulation creates an informational lesion of the stimulated nucleus. *Neuroreport*, 15(7):1137–1140, may 2004.

[28] C. C. McIntyre and P. J. Hahn. Network perspectives on the mechanisms of deep brain stimulation. *Neurobiol. Dis.*, 38(3):329–337, jun 2010.

[29] Andres M. Lozano, Brian J. Snyder, Clement Hamani, William D. Hutchison, and Jonathan O. Dostrovsky. Basal ganglia physiology and deep brain stimulation. *Mov. Disord.*, 25(S1):S71–S75, jan 2010.

[30] S. K. Kalia, T. Sankar, and A. M. Lozano. Deep brain stimulation for Parkinsons disease and other movement disorders. *Curr. Opin. Neurol.*, 26(4):374–380, aug 2013.

[31] C. Hamani, E. Richter, J. M. Schwalb, and A. M. Lozano. Bilateral Subthalamic Nucleus Stimulation for Parkinson's Disease: A Systematic Review of the Clinical Literature. *Neurosurgery*, 56(6):1313–1324, jun 2005.

[32] T. Wagner, A. Valero-Cabre, and A. Pascual-Leone. Noninvasive Human Brain Stimulation. *Annu. Rev. Biomed. Eng.*, 9(1):527–565, aug 2007.

[33] M. Hallett. Transcranial Magnetic Stimulation: A Primer. *Neuron*, 55(2):187–199, jul 2007.

[34] J-P. Lefaucheur. Stroke recovery can be enhanced by using repetitive transcranial magnetic stimulation (rTMS). *Neurophysiol. Clin. Neurophysiol.*, 36(3):105–115, may 2006.

[35] W. Paulus. Transcranial brain stimulation: potential and limitations. *e-Neuroforum*, 5(2):29–36, sep 2014.

[36] E. N. Harvey. The effect of high frequency sound waves on heart muscle and other irritable tissues. *Am. J. Physiol.*, 91(1):284–290, 1929.

[37] W. J. Fry, V. J. Wulff, D. Tucker, and F. J. Fry. Physical Factors Involved in Ultrasonically Induced Changes in Living Systems: I. Identification of NonTemperature Effects. *J. Acoust. Soc. Am.*, 22(6):867–876, nov 1950.

[38] F. J. Fry, H. W. Ades, and W. J. Fry. Production of Reversible Changes in the central nervous system by ultrasound. *Science (80-.).*, 127(3289):83–4, jan 1958.

[39] L. R. Gavrilov and E. M. Tsirulnikov. Focused ultrasound as a tool to input sensory information to humans (Review). *Acoust. Phys.*, 58(1):1–21, jan 2012.

[40] M. L. Khraiche, W. B. Phillips, N. Jackson, and J. Muthuswamy. Ultrasound induced increase in excitability of single neurons. In *2008 30th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, volume 2008, pages 4246–4249. IEEE, aug 2008.

[41] R. Muratore, J. LaManna, E. Szulman, A. Kalisz, M. Lamprecht, M. Simon, Z. Yu, N. Xu, and B. Morrison. Bioeffective Ultrasound at Very Low Doses: Reversible Manipulation of Neuronal Cell Morphology and Function in Vitro. In E. S. Ebbini, editor, *AIP Conf. Proc.*, volume 1113, pages 25–29, 2009.

[42] W. J. Tyler, Y. Tufail, M. Finsterwald, M. L. Tauchmann, E. J. Olson, and C. Majestic. Remote Excitation of Neuronal Circuits Using Low-Intensity, Low-Frequency Ultrasound. *PLoS One*, 3(10):e3511, oct 2008.

[43] Y. Tufail, A. Matyushov, N. Baldwin, M. L. Tauchmann, J. Georges, A. Yoshihiro, S. I. H. Tillery, and W. J. Tyler. Transcranial Pulsed Ultrasound Stimulates Intact Brain Circuits. *Neuron*, 66(5):681–694, jun 2010.

[44] Y. Tufail, A. Yoshihiro, S. Pati, M. M. Li, and W. J. Tyler. Ultrasonic neuromodulation by brain stimulation with transcranial ultrasound. *Nat. Protoc.*, 6(9):1453–1470, sep 2011.

[45] B-K. Min, A. Bystritsky, K-I. Jung, K. Fischer, Y. Zhang, L-S. Maeng, S. In Park, Y-A. Chung, F. A. Jolesz, and S-S. Yoo. Focused ultrasound-mediated suppression of chemically-induced acute epileptic EEG activity. *BMC Neurosci.*, 12(1):23, mar 2011.

[46] P. S. Yang, H. Kim, W. Lee, M. Bohlke, S. Park, T. J. Maher, and S-S. Yoo. Transcranial Focused Ultrasound to the Thalamus Is Associated with Reduced Extracellular GABA Levels in Rats. *Neuropsychobiology*, 65(3):153–160, 2012.

[47] B-K. Min, P. S. Yang, M. Bohlke, S. Park, D. R. Vago, T. J. Maher, and S-S. Yoo. Focused ultrasound modulates the level of cortical neurotransmitters: Potential as a new functional brain mapping technique. *Int. J. Imaging Syst. Technol.*, 21(2):232–240, jun 2011.

[48] H. Kim, S. D. Lee, A. Chiu, S-S. Yoo, and S. Park. Estimation of the spatial profile of neuromodulation and the temporal latency in motor responses induced by focused ultrasound brain stimulation. *Neuroreport*, 25(7):1, dec 2013.

[49] S-S. Yoo, A. Bystritsky, J-H. Lee, Y. Zhang, K. Fischer, B-K. Min, N. J. McDannold, A. Pascual-Leone, and F. A. Jolesz. Focused ultrasound modulates region-specific brain activity. *Neuroimage*, 56(3):1267–1275, jun 2011.

[50] T. Deffieux, Y. Younan, N. Wattiez, M. Tanter, P. Pouget, and J-F. Aubry. Low-Intensity Focused Ultrasound Modulates Monkey Visuomotor Behavior. *Curr. Biol.*, 23(23):2430–2433, dec 2013.

[51] J. Mueller, W. Legon, A. Opitz, T. F. Sato, and W. J. Tyler. Transcranial Focused Ultrasound Modulates Intrinsic and Evoked EEG Dynamics. *Brain Stimul.*, 7(6):900–908, nov 2014.

[52] W. Legon, T. F. Sato, A. Opitz, J. Mueller, A. Barbour, A. Williams, and W. J. Tyler. Transcranial focused ultrasound modulates the activity of primary somatosensory cortex in humans. *Nat. Neurosci.*, 17(2):322–329, jan 2014.

[53] W. Lee, H. Kim, Y. Jung, I-U. Song, Y. A. Chung, and S-S. Yoo. Image-Guided Transcranial Focused Ultrasound Stimulates Human Primary Somatosensory Cortex. *Sci. Rep.*, 5(1):8743, aug 2015.

[54] W. Lee, Y. A. Chung, Y. Jung, I-U. Song, and S-S. Yoo. Simultaneous acoustic stimulation of human primary and secondary somatosensory cortices using transcranial focused ultrasound. *BMC Neurosci.*, 17(1):68, dec 2016.

[55] W. Lee, H-C. Kim, Y. Jung, Y. A. Chung, I-U. Song, J-H. Lee, and S-S. Yoo. Transcranial focused ultrasound stimulation of human primary visual cortex. *Sci. Rep.*, 6(1):34026, dec 2016.

[56] L. Ai, J. K. Mueller, A. Grant, Y. Eryaman, and W. Legon. Transcranial focused ultrasound for BOLD fMRI signal modulation in humans. In *2016 38th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, number January, pages 1758–1761. IEEE, aug 2016.

[57] R. L. King, J. R. Brown, W. T. Newsome, and K. B. Pauly. Effective Parameters for Ultrasound-Induced In Vivo Neurostimulation. *Ultrasound Med. Biol.*, 39(2):312–331, feb 2013.

[58] Y. Younan, T. Deffieux, B. Larrat, M. Fink, M. Tanter, and J-F. Aubry. Influence of the pressure field distribution in transcranial ultrasonic neurostimulation. *Med. Phys.*, 40(8):082902, jul 2013.

[59] H. Kim, A. Chiu, S. D. Lee, K. Fischer, and S-S. Yoo. Focused Ultrasound-mediated Non-invasive Brain Stimulation: Examination of Sonication Parameters. *Brain Stimul.*, 7(5):748–756, sep 2014.

[60] E. Mehić, J.M. Xu, C. J. Caler, N. K. Coulson, C. T. Moritz, and P. D. Mourad. Increased Anatomical Specificity of Neuromodulation via Modulated Focused Ultrasound. *PLoS One*, 9(2):e86939, feb 2014.

[61] R. L. King, J. R. Brown, and K. B. Pauly. Localization of Ultrasound-Induced In Vivo Neurostimulation in the Mouse Model. *Ultrasound Med. Biol.*, 40(7):1512–1522, jul 2014.

[62] H. Kim, M. Y. Park, S. D. Lee, W. Lee, A. Chiu, and S-S. Yoo. Suppression of EEG visual-evoked potentials in rats through neuromodulatory focused ultrasound. *Neuroreport*, 26(4):211–215, mar 2015.

[63] W. Lee, S. D. Lee, M. Y. Park, L. Foley, E. Purcell-Estabrook, H. Kim, K. Fischer, L-S. Maeng, and S-S. Yoo. Image-Guided Focused Ultrasound-Mediated Regional Brain Stimulation in Sheep. *Ultrasound Med. Biol.*, 42(2):459–470, feb 2016.

[64] P. P. Ye, J. R. Brown, and K. B. Pauly. Frequency Dependence of Ultrasound Neurostimulation in the Mouse Brain. *Ultrasound Med. Biol.*, 42(7):1512–1530, jul 2016.

[65] F. Darvas, E. Mehić, C. J. Caler, J. G. Ojemann, and P. D. Mourad. Toward Deep Brain Monitoring with Superficial EEG Sensors Plus Neuromodulatory Focused Ultrasound. *Ultrasound Med. Biol.*, 42(8):1834–1847, aug 2016.

[66] G-F. Li, H-X. Zhao, H. Zhou, F. Yan, J-Y. Wang, C-X. Xu, C-Z. Wang, L-L. Niu, L. Meng, S. Wu, H-L. Zhang, W-B. Qiu, and H-R. Zheng. Improved Anatomical Specificity of Non-invasive Neuro-stimulation by High Frequency (5 MHz) Ultrasound. *Sci. Rep.*, 6(1):24738, jul 2016.

[67] H. A. S. Kamimura, S. Wang, H. Chen, Q. Wang, C. Aurup, C. Acosta, A. A. O. Carneiro, and E. E. Konofagou. Focused ultrasound neuromodulation of cortical and subcortical brain structures using 1.9 MHz. *Med. Phys.*, 43(10):5730–5735, sep 2016.

[68] Y. Yuan, J. Yan, Z. Ma, and X. Li. Effect of noninvasive focused ultrasound stimulation on gamma oscillations in rat hippocampus. *Neuroreport*, 27(7):508–515, may 2016.

[69] R. F. Dallapiazza, K. F. Timbie, S. Holmberg, J. Gatesman, M. B. Lopes, R. J. Price, G. W. Miller, and W. J. Elias. Noninvasive neuromodulation and thalamic mapping with low-intensity focused ultrasound. *J. Neurosurg.*, pages 1–10, apr 2017.

[70] D. Dalecki. Mechanical Bioeffects of Ultrasound. *Annu. Rev. Biomed. Eng.*, 6(1):229–248, aug 2004.

[71] P-H. Tsui, S-H. Wang, and C-C. Huang. In vitro effects of ultrasound with different energies on the conduction properties of neural tissue. *Ultrasonics*, 43(7):560–565, jun 2005.

[72] V. A. Velling and S. P. Shklyaruk. Modulation of the functional state of the brain with the aid of focused ultrasonic action. *Neurosci. Behav. Physiol.*, 18(5):369–375, sep 1988.

[73] P. C. Rinaldi, J. P. Jones, F. Reines, and L. R. Price. Modification by focused ultrasound pulses of electrically evoked responses from an in vitro hippocampal preparation. *Brain Res.*, 558(1):36–42, aug 1991.

[74] J. L. Foley, J. W. Little, and S. Vaezy. Image-Guided High-Intensity Focused Ultrasound for Conduction Block of Peripheral Nerves. *Ann. Biomed. Eng.*, 35(1):109–119, dec 2006.

[75] M. D. Menz, O. Oralkan, P. T. Khuri-Yakub, and S. A. Baccus. Precise Neural Stimulation in the Retina Using Focused Ultrasound. *J. Neurosci.*, 33(10):4550–4560, mar 2013.

[76] W. J. Tyler. Noninvasive Neuromodulation with Ultrasound? A Continuum Mechanics Hypothesis. *Neurosci.*, 17(1):25–36, feb 2011.

[77] J. K. Mueller and W. J. Tyler. A quantitative overview of biophysical forces impinging on neural function. *Phys. Biol.*, 11(5):051001, aug 2014.

[78] G. T. Clement and K. Hynynen. A non-invasive method for focusing ultrasound through the human skull. *Phys. Med. Biol.*, 47(8):1219–1236, apr 2002.

[79] E. M. Lillie, J. E. Urban, A. A. Weaver, A. K. Powers, and J. D. Stitzel. Estimation of skull table thickness with clinical CT and validation with microCT. *J. Anat.*, 226(1):73–80, jan 2015.

[80] A. Moreira-Gonzalez, F. E. Papay, and J. E. Zins. Calvarial Thickness and Its Relation to Cranial Bone Harvest. *Plast. Reconstr. Surg.*, 117(6):1964–1971, may 2006.

[81] A. D. Wijnhoud, M. Franckena, A. van der Lugt, P. J. Koudstaal, and D. W. J. Dippel. Inadequate Acoustical Temporal Bone Window in Patients with a Transient Ischemic Attack or Minor Stroke: Role of Skull Thickness and Bone Density. *Ultrasound Med. Biol.*, 34(6):923–929, jun 2008.

[82] M. Marinoni, A. Ginanneschi, P. Forleo, and L. Amaducci. Technical limits in transcranial Doppler recording: Inadquate acoustic windows. *Ultrasound Med. Biol.*, 23(8):1275–1277, jan 1997.

[83] C. Fatu, M. Puisoru, M. Rotaru, and A.M. Truta. Morphometric evaluation of the frontal sinus in relation to age. *Ann. Anat. - Anat. Anzeiger*, 188(3):275–280, may 2006.

[84] F. Duck. *Physical Properties of Tissue. A Comprehensive Reference Book.* Academic Press, London, jul 1991.

[85] B. Zeqiri, W. Scholl, and S. P. Robinson. Measurement and testing of the acoustic properties of materials: a review. *Metrologia*, 47(2):S156–S171, apr 2010.

[86] J. C. Bamber. Attenuation and Absorption. In *Phys. Princ. Med. Ultrason.*, chapter 4, pages 93–166. John Wiley & Sons, Ltd, Chichester, UK, 2 edition, oct 2005.

[87] G. T. Clement and K. Hynynen. Correlation of ultrasound phase with physical skull properties. *Ultrasound Med. Biol.*, 28(5):617–624, may 2002.

[88] A. Kyriakou, E. Neufeld, B. Werner, M. M. Paulides, G. Szekely, and N. Kuster. A review of numerical and experimental compensation techniques for skull-induced phase aberrations in transcranial focused ultrasound. *Int. J. Hyperth.*, 30(1):36–46, feb 2014.

[89] J. Sun and K. Hynynen. Focusing of therapeutic ultrasound through a human skull: A numerical study. *J. Acoust. Soc. Am.*, 104(3):1705–1715, sep 1998.

[90] R. M. Jones and K. Hynynen. Comparison of analytical and numerical approaches for CT-based aberration correction in transcranial passive acoustic imaging. *Phys. Med. Biol.*, 61(1):23–36, jan 2016.

[91] F. Marquet, M. Pernot, J-F. Aubry, G. Montaldo, L. Marsac, M. Tanter, and M. Fink. Non-invasive transcranial ultrasound therapy based on a 3D CT scan: protocol validation and in vitro results. *Phys. Med. Biol.*, 54(9):2597–2613, may 2009.

[92] G. T. Clement, J. White, and K. Hynynen. Investigation of a large-area phased array for focused ultrasound surgery through the skull. *Phys. Med. Biol.*, 45(4):1071–1083, apr 2000.

[93] F. Marquet, A-L. Boch, M. Pernot, G. Montaldo, D. Seilhean, M. Fink, M. Tanter, and J-F. Aubry. Non-invasive ultrasonic surgery of the brain in non-human primates. *J. Acoust. Soc. Am.*, 134(2):1632–1639, aug 2013.

[94] K. Hynynen and F. A. Jolesz. Demonstration of Potential Noninvasive Ultrasound Brain Therapy Through an Intact Skull. *Ultrasound Med. Biol.*, 24(2):275–283, feb 1998.

[95] S. Chauhan, M. T. Tan, G. Seet, and I. Ng. Minimally invasive robotic HIFU neurosurgical applications. In *2009 38th Annu. Symp. Ultrason. Ind. Assoc.*, pages 1–5. IEEE, mar 2009.

[96] X. Yin and K. Hynynen. A numerical study of transcranial focused ultrasound beam propagation at low frequency. *Phys. Med. Biol.*, 50(8):1821–1836, apr 2005.

[97] A. Kyriakou, E. Neufeld, B. Werner, G. Székely, and N. Kuster. Full-wave acoustic and thermal modeling of transcranial ultrasound propagation and investigation of skull-induced aberration correction techniques: a feasibility study. *J. Ther. Ultrasound*, 3(1):11, dec 2015.

[98] K. Hynynen, N. McDannold, N. A. Sheikov, F. A. Jolesz, and N. Vykhodtseva. Local and reversible bloodbrain barrier disruption by noninvasive focused ultrasound at frequencies suitable for trans-skull sonications. *Neuroimage*, 24(1):12–20, jan 2005.

[99] K. Hynynen, N. McDannold, G. Clement, F. A. Jolesz, E. Zadicario, R. Killiany, T. Moore, and D. Rosen. Pre-clinical testing of a phased array ultrasound system for MRI-guided noninvasive surgery of the brainA primate study. *Eur. J. Radiol.*, 59(2):149–156, aug 2006.

[100] K. Hynynen. MRIgHIFU: A tool for image-guided therapeutics. *J. Magn. Reson. Imaging*, 34(3):482–493, sep 2011.

[101] N. McDannold, G. T. Clement, P. Black, F. Jolesz, and K. Hynynen. Transcranial Magnetic Resonance Imaging Guided Focused Ultrasound Surgery of Brain Tumors: Initial findings in three patients. *Neurosurgery*, 66(2):323–332, feb 2010.

[102] Y. Hertzberg, A. Volovick, Y. Zur, Y. Medan, S. Vitek, and G. Navon. Ultrasound focusing using magnetic resonance acoustic radiation force imaging: Application to ultrasound transcranial therapy. *Med. Phys.*, 37(6):2934–2942, may 2010.

[103] M. Fink. Time reversal of ultrasonic fields. I. Basic principles. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 39(5):555–566, sep 1992.

[104] J-L. Thomas and M. Fink. Ultrasonic beam focusing through tissue inhomogeneities with a time reversal mirror: application to transskull therapy. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 43(6):1122–1129, nov 1996.

[105] J. White, G. T. Clement, and K. Hynynen. Transcranial ultrasound focus reconstruction with phase and amplitude correction. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 52(9):1518–1522, sep 2005.

[106] M. Pernot, J-F. Aubry, M. Tanter, J-L. Thomas, and M. Fink. High power transcranial beam steering for ultrasonic brain therapy. *Phys. Med. Biol.*, 48(16):2577–2589, aug 2003.

[107] J. Sun and K. Hynynen. The potential of transskull ultrasound therapy using the maximum available skull surface area. In *1998 IEEE Ultrason. Symp. Proc.*, volume 2, pages 1537–1540. IEEE, apr 1999.

[108] D. Chauvet, L. Marsac, M. Pernot, A-L. Boch, R. Guillevin, N. Salameh, L. Souris, L. Darrasse, M. Fink, M. Tanter, and Jean-F. Aubry. Targeting accuracy of transcranial magnetic resonanceguided high-intensity focused ultrasound brain therapy: a fresh cadaver model. *J. Neurosurg.*, 118(5):1046–1052, may 2013.

[109] L. N. Trefethen. *Spectral Methods in MATLAB*. Society for Industrial and Applied Mathematics, Philadelphia, 1st edition, jan 2000.

[110] G. Pinton, J-F. Aubry, and M. Tanter. Direct phase projection and transcranial focusing of ultrasound for brain therapy. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 59(6):1149–1159, jun 2012.

[111] A. Pulkkinen, Y. Huang, J. Song, and K. Hynynen. Simulations and measurements of transcranial low-frequency ultrasound therapy: skull-base heating and effective area of treatment. *Phys. Med. Biol.*, 56(15):4661–4683, aug 2011.

[112] Y. Jing, F. C. Meral, and G. T. Clement. Time-reversal transcranial ultrasound beam focusing using a k-space method. *Phys. Med. Biol.*, 57(4):901–917, feb 2012.

[113] J. L. B. Robertson, B. T. Cox, and B. E. Treeby. Quantifying numerical errors in the simulation of transcranial ultrasound using pseudospectral methods. In *2014 IEEE Int. Ultrason. Symp.*, pages 2000–2003. IEEE, sep 2014.

[114] J. L. B. Robertson, E. Hill, A. A. Plumb, S. Choong, S. J. West, and D. I. Nikitichev. 3D printed ultrasound phantoms for clinical training. In *IS&T's 32nd Int. Conf. Print. Fabr.*, volume NIP32, pages 6–8, sep 2016.

[115] J. L. B. Robertson, B. T. Cox, J. Jaros, and B. E. Treeby. Accurate simulation of transcranial ultrasound propagation for ultrasonic neuromodulation and stimulation. *J. Acoust. Soc. Am.*, 141(3):1726–1738, mar 2017.

[116] J. L. B. Robertson, E. Martin, B. T. Cox, and B. E. Treeby. Sensitivity of simulated transcranial ultrasound fields to acoustic medium property maps. *Phys. Med. Biol.*, 62(7):2559–2580, apr 2017.

[117] J. L. B. Robertson, J. E. Urban, J. D. Stitzel, B. T. Cox, and B. E. Treeby. Influence of trabecular microstructure on the simulation of ultrasound through

the skull (Under Preparation). In *7th Eur. Symp. Ultrason. Charact. Bone*, Bad Staffelstein, 2017. IEEE.

[118] B. E. Treeby and B. T. Cox. k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields. *J. Biomed. Opt.*, 15(2):021314:1–12, 2010.

[119] G. H. Golub and J. M. Ortega. *Scientific Computing and Differential Equations: An Introduction to Numerical Methods.* Academic Press, Orlando, Florida, 1st edition, apr 2014.

[120] J. W. Thomas. *Numerical Partial Differential Equations: Finite Difference Methods*, volume 22 of *Texts in Applied Mathematics*. Springer New York, New York, NY, 1st edition, 1995.

[121] M. Tabei, T. D. Mast, and R. C. Waag. A k -space method for coupled first-order acoustic propagation equations. *J. Acoust. Soc. Am.*, 111(1):53–63, jan 2002.

[122] A. D. Pierce. *Acoustics: an introduction to its physical principles and applications*. McGraw-Hill, New York, 2nd edition, 1989.

[123] D. T. Blackstock and A. A. Atchley. Fundamentals of Physical Acoustics. *J. Acoust. Soc. Am.*, 109(4):1274–1276, apr 2001.

[124] B. E. Treeby, J. Jaros, D. Rohrbach, and B. T. Cox. Modelling elastic wave propagation using the k-Wave MATLAB Toolbox. In *2014 IEEE Int. Ultrason. Symp.*, number 5, pages 146–149. IEEE, sep 2014.

[125] J-P. Berenger. A perfectly matched layer for the absorption of electromagnetic waves. *J. Comput. Phys.*, 114(2):185–200, oct 1994.

[126] B. E. Treeby and B. T. Cox. Modeling power law absorption and dispersion for acoustic propagation using the fractional Laplacian. *J. Acoust. Soc. Am.*, 127(5):2741–2748, may 2010.

[127] J. Strikwerda. *Finite difference schemes and partial differential equations*. SIAM: Society for Industrial and Applied Mathematics, Pacific Grove, CA, 2nd edition, 2004.

[128] L. R. Lines, R. Slawinski, and R. P. Bording. A recipe for stability of finite-difference waveequation computations. *GEOPHYSICS*, 64(3):967–969, may 1999.

[129] B. E. Treeby and B. T. Cox. A k -space Green's function solution for acoustic initial value problems in homogeneous media with power law absorption. *J. Acoust. Soc. Am.*, 129(6):3652–3660, jun 2011.

[130] A. Pulkkinen, B. Werner, E. Martin, and K. Hynynen. Numerical simulations of clinical focused ultrasound functional neurosurgery. *Phys. Med. Biol.*, 59(7):1679–1700, apr 2014.

[131] B. E. Treeby, J. Jaros, A. P. Rendell, and B. T. Cox. Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k -space pseudospectral method. *J. Acoust. Soc. Am.*, 131(6):4324–4336, jun 2012.

[132] M. Cieszko, R. Drelich, and M. Pakula. Acoustic wave propagation in equivalent fluid macroscopically inhomogeneous materials. *J. Acoust. Soc. Am.*, 132(5):2970–2977, nov 2012.

[133] J. Jaros, A. P. Rendell, and B. E. Treeby. Full-wave nonlinear ultrasound simulation on distributed clusters with applications in high-intensity focused ultrasound. *Int. J. High Perform. Comput. Appl.*, 30(2):137–155, may 2016.

[134] R. A. Heckemann, S. Keihaninejad, P. Aljabar, D. Rueckert, J. V. Hajnal, and A. Hammers. Improving intersubject image registration using tissue-class information benefits robustness and accuracy of multi-atlas based anatomical segmentation. *Neuroimage*, 51(1):221–227, may 2010.

[135] M. Jenkinson, C. F. Beckmann, T. E. J. Behrens, M. W. Woolrich, and S. M. Smith. FSL. *Neuroimage*, 62(2):782–790, aug 2012.

[136] Q. Fang and D. A. Boas. Tetrahedral mesh generation from volumetric binary and grayscale images. In *2009 IEEE Int. Symp. Biomed. Imaging From Nano to Macro*, pages 1142–1145, Boston, jun 2009. IEEE.

[137] Q. Grimal, D. Rohrbach, J. Grondin, R. Barkmann, C-C. Glüer, K. Raum, and P. Laugier. Modeling of Femoral Neck Cortical Bone for the Numerical Simulation of Ultrasound Propagation. *Ultrasound Med. Biol.*, 40(5):1015–1026, may 2014.

[138] J. K. Mueller, L. Ai, P. Bansal, and W. Legon. Computational exploration of wave propagation and heating from transcranial focused ultrasound for neuromodulation. *J. Neural Eng.*, 13(5):056002:1–14, oct 2016.

[139] B. E. Treeby and B. T. Cox. Modeling power law absorption and dispersion for acoustic propagation using the fractional Laplacian. *J. Acoust. Soc. Am.*, 127(5):2741–2748, may 2010.

[140] G. Pinton, J-F. Aubry, M. Fink, and M. Tanter. Effects of nonlinear ultrasound propagation on high intensity brain therapy. *Med. Phys.*, 38(3):1207–1216, feb 2011.

[141] B. Zurl, R. Tiefling, P. Winkler, P. Kindl, and K. S. Kapp. Hounsfield units variations. *Strahlentherapie und Onkol.*, 190(1):88–93, jan 2014.

[142] M. Wintermark, N. J. Tustison, W. J. Elias, J. T. Patrie, W. Xin, N. Demartini, M. Eames, S. Sumer, B. Lau, A. Cupino, J. Snell, A. Hananel, N. Kassell, and J-F. Aubry. T1-weighted MRI as a substitute to CT for refocusing planning in MR-guided focused ultrasound. *Phys. Med. Biol.*, 59(13):3599–3614, jul 2014.

[143] P. J. White, G. T. Clement, and K. Hynynen. Local frequency dependence in transcranial ultrasound transmission. *Phys. Med. Biol.*, 51(9):2293–2305, may 2006.

[144] D. Nikitichev, A. Barburas, K. McPherson, J-M. Mari, S. J. West, and A. E. Desjardins. Construction of 3-Dimensional Printed Ultrasound Phantoms With Wall-less Vessels. *J. Ultrasound Med.*, 35(6):1333–1339, jun 2016.

[145] J. G. Truscott, R. Milner, A. J. Clarke, and J. A. Evans. Towards an Ultrasonic Bone Phantom. In IPSM, editor, *Abstr. IPSM Golden Jubil. Congr.*, page 151, York, 1993.

[146] A. J. Clarke, J. A. Evans, J. G. Truscott, R. Milner, and M. A. Smith. A phantom for quantitative ultrasound of trabecular bone. *Phys. Med. Biol.*, 39(10):1677–1687, oct 1994.

[147] D. B. Saint John. *3D printing primer : a survey of technologies and principles*. O'Reilly Media, Sebastopol, CA, 2015.

[148] W. Marczak. Water as a standard in the measurements of speed of sound in liquids. *J. Acoust. Soc. Am.*, 102(5):2776–2779, nov 1997.

[149] CIBC. Seg3D: Volumetric Image Segmentation and Visualization. Scientific Computing and Imaging Institute (SCI), Download from: http://www.seg3d.org, 2016.