

Acoustoelastic Full-Waveform Inversion for Transcranial Ultrasound Computed Tomography

Conference Paper**Author(s):**

Marty, Patrick; Boehm, Christian; Fichtner, Andreas

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000499222>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

Proceedings of SPIE 11602, <https://doi.org/10.1117/12.2581029>

Acoustoelastic Full-Waveform Inversion for Transcranial Ultrasound Computed Tomography

Patrick Marty^a, Christian Boehm^a, and Andreas Fichtner^a

^aETH Zurich, Sonneggstrasse 5, 8092 Zurich, Switzerland

ABSTRACT

Full-waveform inversion applied to ultrasound computed tomography is a promising technique to provide high-resolution quantitative images of soft human tissues, which are otherwise difficult to illuminate by conventional ultrasound imaging. A particular challenge which arises within transcranial ultrasound is the imprint of the solid skull on the measured wavefield. We present an acoustoelastic approach to full-waveform inversion for transcranial ultrasound computed tomography that accurately accounts for the solid-fluid interactions along the skull-tissue interfaces. Using the spectral-element method on cubical meshes, we obtain a scalable and performant method to resolve such a coupled physical system. Moreover, since the volume of the skull is small compared to the entire simulation domain, solving a coupled system of the acoustoelastic wave equation increases the computational cost only by a small margin compared to the acoustic approximation. We perform an *in silico* forward and inverse modeling study that reveals significant coupling effects at the skull-tissue interfaces when considering the skull as an elastic medium as opposed to an acoustic medium. Applying full-waveform inversion to a set of synthetically generated acoustoelastic forward data allows for favorable reconstructions to be achieved when considering an acoustoelastic prior model of the skull.

Keywords: Full-waveform inversion, transcranial ultrasound, ultrasound computed tomography, acoustoelastic, spectral-element method

1. INTRODUCTION

The application of ultrasound computed tomography (USCT) for imaging the human brain poses considerable challenges due to the complex interactions that occur between the ultrasound wavefield and the skull.¹⁻⁴ The strong material contrast at the tissue-bone interface causes a significant reflection of the incoming ultrasound wave, which severely limits its transmission.

Many of the earlier studies which analyzed the acoustic properties of the skull primarily consider the propagation of compressional waves (P-waves) within the skull and largely neglected the effects of shear waves (S-waves).^{1,5} The influence of S-waves was typically ignored due to the assumption that the coupling effects between P- and S-waves at the tissue-bone interface caused an insignificant change to the resulting P-wave amplitudes for wavefronts with a low incidence angle relative to the skull.^{1,6} Other subsequent studies have shown that there may be considerable coupling effects between the skull and the surrounding tissue due to the elastic characteristics of bone (non-negligible S-wave velocity) when compared to the surrounding acoustic tissues (negligible S-wave velocity).^{4,7}

Imaging algorithms such as full-waveform inversion (FWI) which utilize the full wavefield within the reconstruction process have been developed in the area of seismology for several decades⁸⁻¹⁰ and have also been subsequently applied in the context of imaging soft tissues using ultrasound.¹¹⁻¹⁶ Mitsuhashi *et al.*¹⁷ have analyzed the propagation of ultrasound waves through the skull using the elastic wave equation in the context of photoacoustic computed tomography. FWI has shown to be an effective tool in geophysics for subsalt imaging where the significant material contrast between salt layers and adjacent sandstone formations render many other inversion methods ineffective.¹⁸ Given the qualitative similarities between subsalt imaging and transcranial USCT,

Further author information: (Send correspondence to P.M.)

E-Mail: patrick.marty@erdw.ethz.ch, christian.boehm@erdw.ethz.ch, andreas.fichtner@erdw.ethz.ch

studies such as those by Guasch *et al.*¹⁹ and Agudo²⁰ have proposed applying FWI to image the human brain within an acoustic regime.

This *in silico* study analyzes the use of acoustoelastic FWI in conjunction with the spectral-element method (SEM) for reconstructing a 2D model of the brain. The primary objectives of this study are as follows: (a) determine the differences in the forward ultrasound wavefields between cases where the skull is treated as an acoustic versus elastic medium, (b) compare different source time functions and FWI strategies, and (c) assess the efficacy of acoustoelastic FWI for reconstructing the brain tissue using both acoustic and acoustoelastic forward data.

2. BACKGROUND

The two primary modeling methods that are used throughout this study are the SEM and FWI. The SEM is used as the underlying framework for computing the forward and adjoint wavefields while FWI is the method used for determining appropriate model updates over a series of inversion iterations.

2.1 Spectral-Element Method

The spectral-element method (SEM) is a type of finite-element method which possesses certain computational benefits over other finite-element or finite-difference schemes when modeling the wave equation in a coupled medium.^{21–25} One of the primary advantages of the SEM is that the spatial domain is constructed such that the resulting global mass matrix is diagonal, which allows for the use of explicit time stepping schemes without having to solve a linear system. Furthermore, sharp discontinuities can be effectively captured within the geometry of the problem given how the domain is discretized in space.

Consider a domain Ω of spatial dimension $n = 2, 3$ with coordinates $\mathbf{x} \in \Omega \subset \mathbb{R}^n$ for some time interval $t \in T \subset \mathbb{R}$. The acoustic wavefield within an inviscid fluid can be parameterized in terms of the density ρ and P-wave velocity v_p described by the strong form of the acoustic wave equation as

$$\rho^{-1}(\mathbf{x})v_p^{-2}(\mathbf{x})\partial_t^2\phi(\mathbf{x}, t) - \nabla \cdot (\rho^{-1}(\mathbf{x})\nabla\phi(\mathbf{x}, t)) = f(\mathbf{x}, t) \quad (1)$$

for a scalar potential ϕ and source term f .²¹ A similar parameterization of the strong form of the elastic wave equation can be defined as

$$\rho(\mathbf{x})\partial_t^2\mathbf{u}(\mathbf{x}, t) - \nabla \cdot \boldsymbol{\sigma}(\mathbf{x}, t) = \mathbf{0} \quad (2)$$

for displacement field \mathbf{u} and stress tensor $\boldsymbol{\sigma}$.²⁶ Hooke's law describes the constitutive relation between stress and strain modeled by the action of the elastic tensor, which contains information on the material properties of the domain such as the P-wave velocity, S-wave velocity, and density. The elastic medium in equation (2) is assumed to be source-free since the ultrasound transducers are only placed in the acoustic coupling fluid surrounding the head of the patient. Furthermore, all numerical examples consider an isotropic and non-dissipative elastic tensor.

Equations (1) and (2) for acoustic and elastic regions of the domain are coupled through interface conditions along the interfaces $\partial\Omega_i$. Since this boundary $\partial\Omega_i$ is constructed within the mesh to precisely follow the fluid-solid (acoustic-elastic) interface, it is possible to accurately account for the coupling effects at these boundaries. In particular, the accompanying interface conditions^{21, 22} between the acoustic and the elastic portions of the domain can be described by the continuity of traction given by

$$\boldsymbol{\sigma}(\mathbf{x}, t) \cdot \hat{\mathbf{n}}(\mathbf{x}) = \rho(\mathbf{x})\partial_t^2\phi(\mathbf{x}, t)\hat{\mathbf{n}}(\mathbf{x}), \quad \text{on } \partial\Omega_i \quad (3)$$

as well as the kinematic boundary condition of continuity of the normal component of the displacement, as described by

$$\hat{\mathbf{n}}(\mathbf{x}) \cdot \nabla\phi(\mathbf{x}, t) = \hat{\mathbf{n}}(\mathbf{x}) \cdot \mathbf{u}(\mathbf{x}, t), \quad \text{on } \partial\Omega_i. \quad (4)$$

Note that a domain which consists of such a coupled system and thus possesses both independent acoustic and elastic constituents is referred to as an *acoustoelastic* domain throughout this work.

Furthermore, absorbing boundary conditions must be defined for the outermost boundary of the domain $\partial\Omega_{\text{abs}}$. To simplify the notation, the equations below only consider the Sommerfeld radiation condition,²⁷ which mitigates

most of the near-normal incidence waves to the boundary. The numerical examples include additional sponge layers²⁸ to dissipate any residual reflections off of the boundary. It is assumed that the absorbing boundaries within this setup will be employed exclusively within the acoustic coupling fluid, meaning that the absorbing boundary conditions only need to be applied to the acoustic wave equation. Thus, the first-order Sommerfeld radiation condition can be expressed in terms of the scalar potential $\phi(\mathbf{x}, t)$ as

$$\nabla\phi(\mathbf{x}, t) \cdot \hat{\mathbf{n}}(\mathbf{x}) = -v_p^{-1}\partial_t\phi(\mathbf{x}, t), \quad \text{on } \partial\Omega_{\text{abs}}. \quad (5)$$

As mentioned, sponge layers are also applied in addition to the Sommerfeld condition. However, the expression for this additional dampening criterion is omitted here for ease of notation.

Although equations (1) and (2) are instructive from a conceptual point of view, the spectral-element method requires the transformation into their so-called weak or variational forms. Solving the weak formulation allows for one to find a field that satisfies the wave equation for any choice of test function. The weak forms of both the acoustic and elastic wave equations are obtained by multiplying the strong form by some test function and then applying Green's first identity. The strong form of the acoustic wave equation from (1) can be expressed in its weak form as

$$\int_{\Omega} \rho^{-1}v_p^{-2}w\partial_t^2\phi \, d^n\mathbf{x} - \oint_{\partial\Omega} (\rho^{-1}w\nabla\phi) \cdot \hat{\mathbf{n}} \, d^{n-1}\mathbf{x} + \int_{\Omega} \rho^{-1}\nabla w \cdot \nabla\phi \, d^n\mathbf{x} = \int_{\Omega} wf \, d^n\mathbf{x} \quad (6)$$

with w being a scalar test function. The surface integral term in equation (6) describes the behavior at both the absorbing boundaries $\partial\Omega_{\text{abs}}$ as well as the interfaces $\partial\Omega_i$. Thus, these interface conditions from equations (4) and (5) can be explicitly included within the weak formulation as

$$\int_{\Omega} \rho^{-1}v_p^{-2}w\partial_t^2\phi \, d^n\mathbf{x} + \oint_{\partial\Omega_{\text{abs}}} \rho^{-1}v_p^{-1}w\partial_t\phi \, d^{n-1}\mathbf{x} - \oint_{\partial\Omega_i} \rho^{-1}w\mathbf{u} \cdot \hat{\mathbf{n}} \, d^{n-1}\mathbf{x} + \int_{\Omega} \rho^{-1}\nabla w \cdot \nabla\phi \, d^n\mathbf{x} = \int_{\Omega} wf \, d^n\mathbf{x} \quad (7)$$

Similar to the acoustic wave equation, the strong form of the elastic wave equation from equation (2) can be formulated in its weak form by multiplying the strong form by some vector-valued test function \mathbf{w} and applying Green's first identity. This yields the weak form as expressed by

$$\int_{\Omega} \rho\mathbf{w} \cdot \partial_t^2\mathbf{u} \, d^n\mathbf{x} - \oint_{\partial\Omega} \mathbf{w} \cdot (\boldsymbol{\sigma} \cdot \hat{\mathbf{n}}) \, d^{n-1}\mathbf{x} + \int_{\Omega} \nabla\mathbf{w} : \boldsymbol{\sigma} \, d^n\mathbf{x} = 0. \quad (8)$$

The interface conditions from equation (3) can be substituted into equation (8) to yield

$$\int_{\Omega} \rho\mathbf{w} \cdot \partial_t^2\mathbf{u} \, d^n\mathbf{x} - \oint_{\partial\Omega_i} \rho\mathbf{w} \cdot \hat{\mathbf{n}} \partial_t^2\phi \, d^{n-1}\mathbf{x} + \int_{\Omega} \nabla\mathbf{w} : \boldsymbol{\sigma} \, d^n\mathbf{x} = 0. \quad (9)$$

In contrast to the weak form of the acoustic wave equation as seen in equation (7), the weak form of the elastic wave equation in equation (9) does not contain an absorbing boundary term due to the absorbing boundaries only being applied to the outer-most boundary of the acoustic coupling fluid.

The reader is referred to works such as Komatitsch *et al.*²¹ and Afanasiev *et al.*²⁵ for further details on the discretization of the weak formulations.

The SEM is widely used in the seismology community for applications ranging from the exploration scale to the global scale due to the efficacy with which the interfaces along distinct boundaries can be represented. The accurate depiction of these interfaces allows for numerical discrepancies to be avoided, which are introduced by the *stair-stepping* observed at sharp discontinuities within rectilinear meshes.²⁹

Although the accurate representation of these fluid-solid interfaces is a considerable asset of the SEM, constructing such a computational mesh can be quite challenging. All meshes considered herein consist of hexahedral elements since their inherent tensorized structure leads to a reduction in the computational cost of

the subsequent wave simulation when compared to tetrahedral elements.²⁴ Furthermore, dispersion errors for tetrahedral meshes are typically higher when compared to hexahedral meshes unless the tetrahedral mesh uses a higher polynomial degree.

The size of the elements throughout the mesh also plays a critical role throughout this process as the elements must be sufficiently small to resolve the shorter wavelengths within the wavefield, while being large enough to avoid restrictions of the time step due to the Courant-Friedrichs-Lewy stability criterion. The global time step is limited by the ratio of element size and local velocity, meaning that a single unnecessarily small element within the mesh could significantly increase the number of time steps required for the simulation to be stable.

2.2 Full-Waveform Inversion

The following is intended to provide the reader with a high-level conceptual overview of full-waveform inversion. The reader is referred to works such as Fichtner³⁰ and Virieux & Operto¹⁰ for more details.

Full-waveform inversion (FWI) is an inversion technique which takes a set of observations $\mathbf{u}(\mathbf{x}^r, t)$ and computes an optimal set of model parameters $\tilde{\mathbf{m}}$ which can explain the set of observed data. This process involves defining some misfit functional $\chi(\mathbf{m})$, which acts as a measure of discrepancy between the observed data $\mathbf{u}(\mathbf{x}^r, t)$ and a set of synthetically generated forward data $\mathbf{u}(\mathbf{m}; \mathbf{x}^t, t)$ for some set of model parameters \mathbf{m} . These model parameters are iteratively adjusted such as to minimize the misfit functional. In the context of ultrasound imaging, \mathbf{m} represents the set of material parameters of the various tissues within the domain and $\mathbf{u}(\mathbf{x}^r, t)$ represents the observed ultrasound measurements at some receiver transducer positions \mathbf{x}^r .

The strategy that is typically employed for minimizing $\chi(\mathbf{m})$ is to define some starting model \mathbf{m}_0 and then iteratively update these model parameters until the observed measurements and the synthetically generated measurements are within some acceptable threshold. However, determining an optimal model $\tilde{\mathbf{m}}$ is quite challenging due to there being a multitude of local minima throughout $\chi(\mathbf{m})$. Although such a local minimum may decrease the misfits, the corresponding model may be physically meaningless. Thus, it is important that an initial set of model parameters \mathbf{m}_0 is chosen so that the model is initialized in the neighborhood of $\tilde{\mathbf{m}}$, thus increasing the likelihood that the optimization algorithm reaches a meaningful optimum.

The general optimization problem aims to iteratively update the model \mathbf{m}_i by

$$\mathbf{m}_{i+1} = \mathbf{m}_i + \mathbf{s}_i \quad \text{with} \quad \chi(\mathbf{m}_{i+1}) < \chi(\mathbf{m}_i). \quad (10)$$

This ensures that the model update \mathbf{s}_i is chosen such that the new model \mathbf{m}_{i+1} explains the observed data more effectively than the previous model. A trust-region L-BFGS algorithm is employed here in order to minimize this objective functional.^{16,31} The model update is determined as an approximate solution to the auxiliary trust-region subproblem

$$\mathbf{s}_i = \arg \min_{\mathbf{s}} \chi(\mathbf{m}_i) + \nabla_{\mathbf{m}} \chi(\mathbf{m}_i)^T \mathbf{s} + \frac{1}{2} \mathbf{s}^T \mathbf{H}_i \mathbf{s} \quad \text{s.t.} \quad \|\mathbf{s}\| \leq \Delta_i, \quad (11)$$

where \mathbf{H}_i is the approximate Hessian of the objective functional constructed from the previous models and gradients, and preconditioned by a smoothing operator, and Δ_i is the trust-region radius that limits the magnitude of the changes between two subsequent models.

Gradients of the misfit functional with respect to the model parameters can be computed efficiently using the *adjoint method*. The derivation of the adjoint method is described in detail in the literature and is thus omitted for the sake of brevity; the reader is referred to Fichtner³⁰ for more details. The adjoint method involves solving an auxiliary equation to obtain the adjoint wavefield \mathbf{u}^\dagger , which can then be used to express the gradient as

$$\nabla_{\mathbf{m}} \chi \delta \mathbf{m} = \int_{\Omega} \int_T \mathbf{u}^\dagger \cdot \nabla_{\mathbf{m}} \mathbf{L} \delta \mathbf{m} \quad dt \quad d^n \mathbf{x}. \quad (12)$$

Here, \mathbf{L} denotes the wave operator of the forward wavefield. Thus, the two underlying pieces of information that are required for computing the gradients of the misfit functional are the forward wavefield and the adjoint

wavefield. It is important to note that the gradient in this case is computed independently for each source position and is summed in a subsequent step to obtain a cumulative gradient across all source positions.

Implementing the adjoint equation requires the terminal condition of $\mathbf{u}^\dagger = 0$ for time $t = t_1$ given a time interval $[t_0, t_1] \in T$. Unlike the forward wavefield which is computed from t_0 to t_1 , this condition requires the adjoint wavefield to be computed backwards in time from t_1 to t_0 such that the terminal conditions of the adjoint run correspond with the effective initial condition at t_0 . A consequence of using the adjoint equation is that both the forward and the adjoint wavefields need to be known simultaneously in time. This can be particularly problematic for cases where a large number of time steps are required since this may introduce infeasible memory requirements. A checkpointing strategy³² is employed in order to significantly reduce the memory requirements of computing the gradients of the misfit functional at the consequence of increasing the computational cost slightly.

3. METHODS

The primary tool used throughout this study for meshing the domain as well as conducting the forward and inverse modeling is the SEM package *Salvus*.²⁵

3.1 Cranial Phantom

The multimodal imaging-based detailed anatomical (MIDA) model is a segmented cranial phantom which has been created from a combination of scans taken with magnetic resonance imaging, magnetic resonance angiography, and diffusion tensor imaging.³³ Within the literature, tissues which are volumetrically significant within the brain possess several reference estimates of their approximate acoustic properties.^{34,35} However, the acoustic properties of less volumetrically significant tissues can be quite problematic to measure in practice (especially *in vivo*), which leads to many of these tissues possessing little to no documented reference acoustical properties. In cases where no material properties of the relevant tissue are available, approximate acoustic properties are applied to the relevant tissue using some perturbation from the material parameters of water (typically within the range of $v_p \approx 1500 \pm 75$ m/s and $\rho \approx 1000 \pm 50$ kg/m³). The coupling fluid surrounding the head of the patient is assumed to be water. The P-wave velocity and density models used for generating synthetic *observed* data can be seen in figure 1.

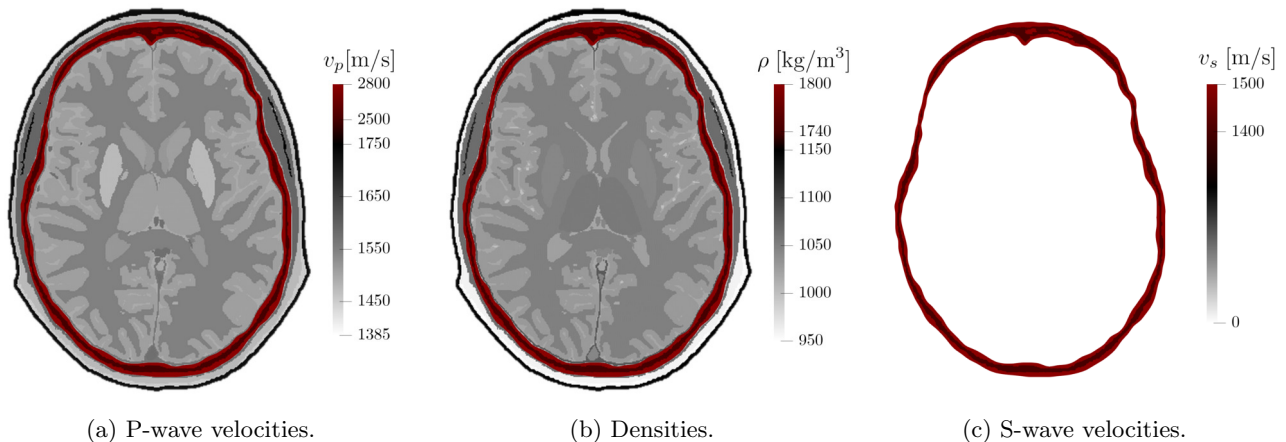


Figure 1: The models used for generating the *observed* forward data. Note that the v_s velocity model is only relevant for acoustoelastic simulations and is otherwise ignored in purely acoustic simulations. All regions denoted with red correspond to the skull whereas grayscale regions correspond to soft tissue.

3.2 Mesh Generation

Two possible meshing strategies are considered for constructing a model of the skull, namely (a) an equidistantly-spaced rectilinear mesh similar to grids of finite-difference methods and (b) a mesh that explicitly aligns elements

with the skull-tissue interface. Rectilinear meshes were initially investigated due to the ease with which they can be implemented. However, preliminary testing yielded similar results to those observed in Afanasiev *et al.*²⁹ where the *stair-stepping* structure along the strong material contrast at the skull-tissue interface introduced significant numerical artifacts in both the reflected and transmitted wavefields. Such numerical errors could be completely eliminated when using a mesh which has its element boundaries conform with the shape of the skull using a cubed sphere mesh. As a result, all models shown herein explicitly mesh the skull and are based on a cubed sphere meshing algorithm.

The general strategy used for constructing such a cubed sphere mesh which precisely follows the shape of the skull is to first create a perfectly circular mesh which possesses three distinct layers: (a) the inner-most acoustic (fluid) layer representing the brain tissue, (b) the middle elastic (solid) layer representing the skull, and (c) the outer-most acoustic (fluid) layer representing both the scalp as well as the coupling fluid surrounding the head of the patient. This perfectly spherical mesh can then be deformed such that the skull boundaries conform with the inner and outer contours of the skull. An overview of this meshing strategy can be seen in figure 2.

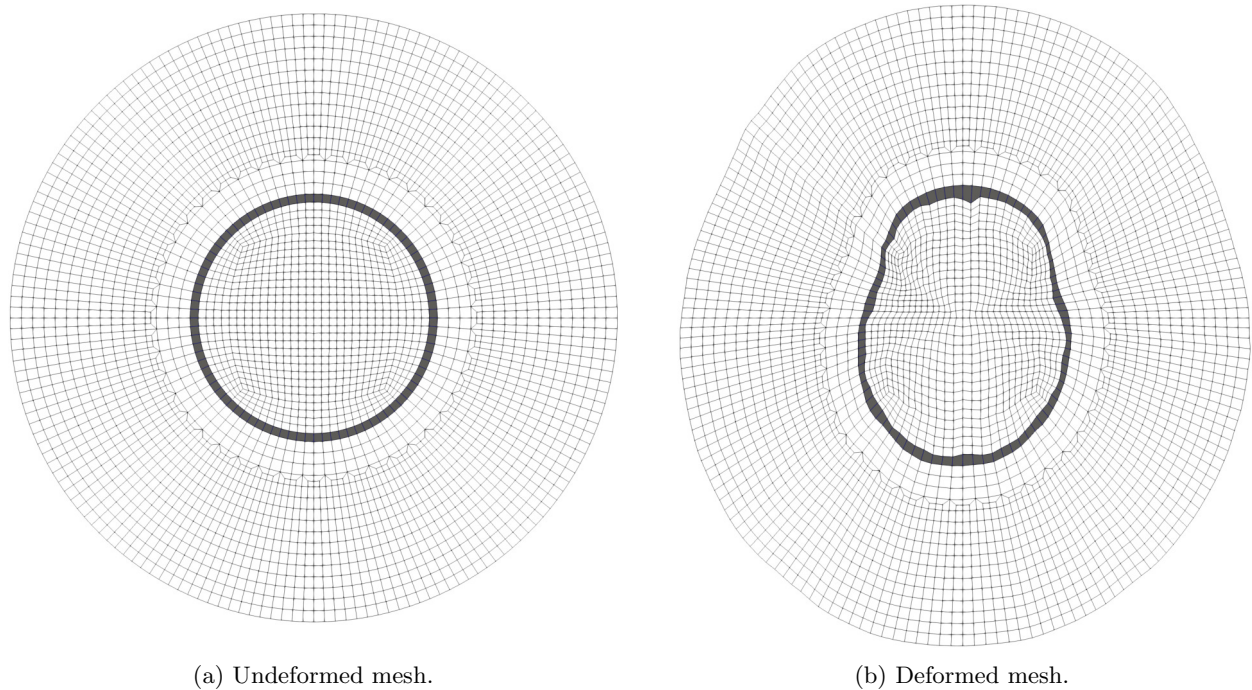
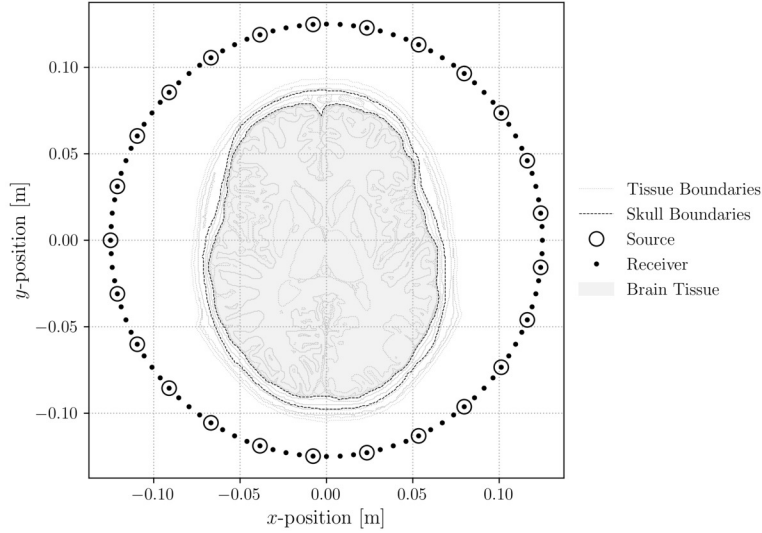


Figure 2: An example of how a cubed sphere can be constructed and subsequently deformed to ensure that the element boundaries conform with the shape of the skull.

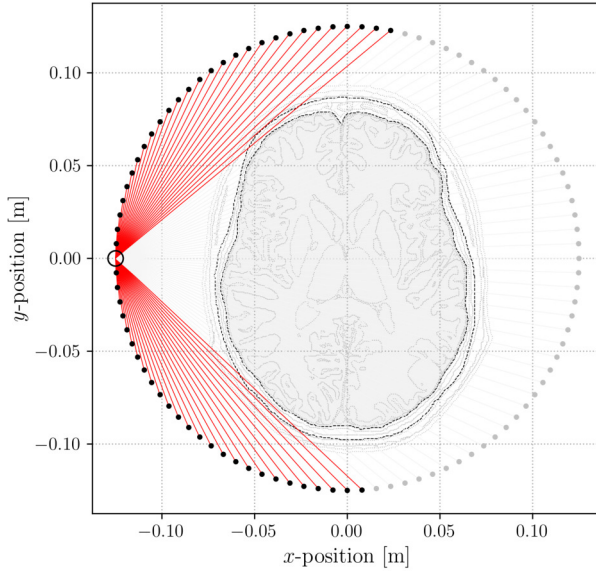
3.3 Source-Receiver Apertures

A circular ring-like source-receiver aperture consisting of 25 source and 100 possible receiver transducer positions is utilized throughout this study. Three possible source-receiver configurations are considered: (a) a full ring of source-receiver transducers, (b) select receivers proximal to the source position to primarily capture reflection information, and (c) select receivers opposite the source position to primarily capture transmission information. The source-receiver apertures of (b) and (c) are assigned using a simple ray-casting algorithm between each source-receiver pair. The rays that pass through the brain tissue denote transmission source-receiver pairs while those that do not intersect the brain tissue denote reflection source-receiver pairs. An example of such an aperture is shown in figure 3.


Previous studies in the literature suggest that the ideal frequency range for transmitting ultrasound waves through the skull is within the range of approximately 500 kHz to 1 MHz.^{1,37} Two source time functions (STFs) are considered in this study, both with center frequency 500 kHz. The first STF is that of a Ricker wavelet which



(a) All possible source-receiver positions.




(b) Possible aperture favoring reflection information.




(c) Possible aperture favoring transmission information.

Figure 3: An overview of the source-receiver apertures considered. All possible source-receiver transducer positions are shown in a) with a total of 25 source locations and 100 possible receiver locations. The examples shown in b) and c) demonstrate possible apertures for a source at location $[-0.125, 0.0]$ depending on whether reflection or transmission data is preferred. Figures adapted from Marty.³⁶

possesses a bell-shaped frequency spectrum with a bandwidth that is slightly less than twice the center frequency (approx. 1 MHz bandwidth). The second STF considered is that of a Ricker wavelet which has been filtered to discard the lower and higher frequencies within the wavelet, thus reducing the bandwidth of the source (approx. 500 kHz bandwidth). In both cases, the STFs are assumed to be emitted from omnidirectional point sources. Examples of the Ricker wavelet and the filtered Ricker wavelet can be seen in figure 4.



(a) Ricker wavelet.



(b) Bandpass filtered Ricker wavelet.

Figure 4: A comparison between the two STF’s used for this analysis. Both wavelets possess similar bell-shaped frequency spectra with the Ricker wavelet in a) containing a broader bandwidth when compared to the filtered Ricker wavelet in b). The Ricker wavelet in a) represents an *ideal* source from the point of view that the power spectrum contains lower frequency information, which significantly aids in the convergence of FWI. The bandpass filtered Ricker wavelet in b) possesses a somewhat more realistic STF which does not contain these lower frequencies.

3.4 Inversion Workflows

Two different inversion workflows are considered, namely (a) inverting for the brain tissue directly by using a full ring of sources and receivers (see figure 3a) forcing all model updates to be mapped to the interior of the brain and (b) inverting first for the scalp in one series of inversions using mainly reflection information (see figure 3b) and then subsequently inverting for the brain tissue using mainly transmission information (see figure 3c). The case described in (a) will likely allow for the reconstruction to be conducted with fewer inversion iterations when compared to (b), although the quality of the reconstruction in (a) will likely be decreased given that the presence of the scalp tissue is neglected. The setups described in (a) and (b) are referred to as *single-step* and *two-step* inversion schemes, respectively. Figure 5 provides an overview of these inversion workflows.

The underlying workflow between both the single-step and two-step inversions are relatively similar. First, a set of source and receiver positions (denoted by \mathbf{x}^s and \mathbf{x}^r , respectively) are defined. As discussed in section 3.3, the source-receiver apertures considered here consists of a ring of 25 sources and 100 possible receivers. In cases where recorded data is available, the transducer positions would be provided here instead.

Next, a mesh of the synthetic *true* model is generated which can be used for computing the forward wavefield $\mathbf{u}(\mathbf{x}, t)$. The wavefield is sampled at coordinates \mathbf{x}^r so that the set of recordings $\mathbf{u}(\mathbf{x}^r, t)$ can act as the *observed* data for testing the respective inversion algorithm. This process can be skipped in cases where recorded data is available instead.

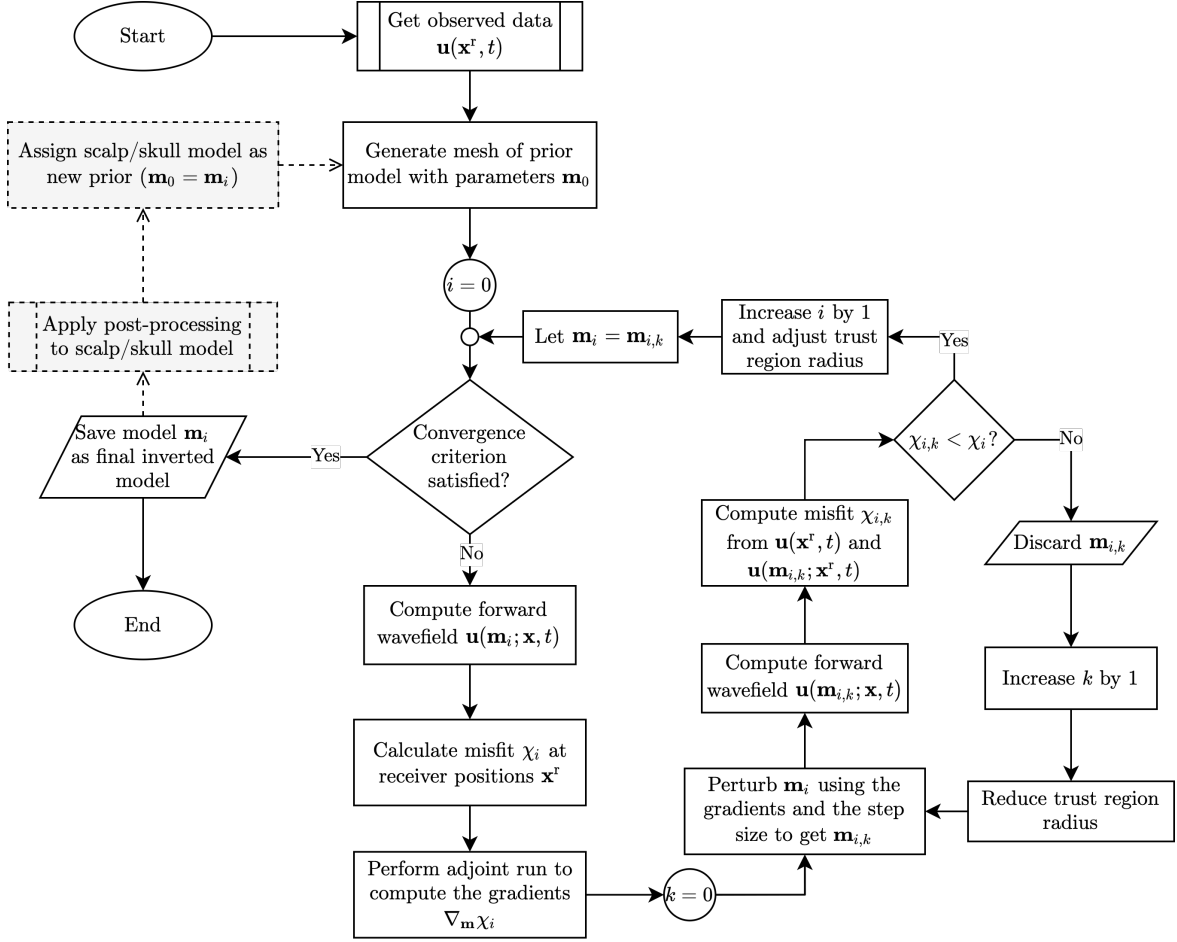


Figure 5: Boxes in white denote the general workflow for conducting a single-step inversion process. The additional hatched boxes marked in gray correspond to the additional steps required for conducting a two-step inversion process. Figure adapted from Marty.³⁶

Now that a set of observed data is available, the starting mesh (referred to as the *prior model*) is constructed with some material parameters \mathbf{m}_0 ; the model parameters subscript refers to the current inversion iteration $i = 0$. This mesh is used for computing the forward wavefield $\mathbf{u}(\mathbf{m}_i; \mathbf{x}, t)$, which is then sampled at the same receiver positions \mathbf{x}^r as the observed data. The misfit χ_i between the observed data $\mathbf{u}(\mathbf{x}^r, t)$ and the synthetically generated data $\mathbf{u}(\mathbf{m}_i; \mathbf{x}^r, t)$ can then be computed using this forward data from the prior model. Throughout this study, the L^2 norm is used exclusively as the misfit functional.

Next, the gradients of the misfit functional $\nabla_{\mathbf{m}} \chi_i$ relative to one of the inverted parameters \mathbf{m} (eg. v_p or ρ) can be computed during the adjoint run by computing the adjoint field $\mathbf{u}^\dagger(\mathbf{m}_i; \mathbf{x}, t)$. These gradients can then be used in conjunction with some step size to perturb the prior model which should, ideally, reduce the model misfits. As previously discussed, the model update step employs a trust region L-BFGS method for creating the updated model $\mathbf{m}_{i,k}$, where the index k denotes the current trial model. In order to assess whether the misfits for the updated model have in fact decreased, the forward wavefield for model $\mathbf{m}_{i,k}$ must be computed to assess whether the misfits of $\chi_{i,k}$ (based on the misfit between $\mathbf{u}(\mathbf{x}^r, t)$ and $\mathbf{u}(\mathbf{m}_{i,k}; \mathbf{x}^r, t)$) has reduced relative to χ_i . In the event that $\chi_{i,k}$ has not decreased relative to χ_i , the model $\mathbf{m}_{i,k}$ is discarded and a new trust region radius is chosen. This process of updating the model \mathbf{m}_i to $\mathbf{m}_{i,k}$ is repeated until the misfits $\chi_{i,k}$ have reduced sufficiently. Once a satisfactory model $\mathbf{m}_{i,k}$ has been found, $\mathbf{m}_{i+1} = \mathbf{m}_{i,k}$ such that the entire updating process can be repeated until some set number of inversion iterations are reached.

The two-step inversion process (see the additional gray hatched cells in figure 5) is quite similar to the single-step inversion process, with the exception being that the first set of inversion iterations attempts to invert for the scalp tissue and the second set of iterations then attempts to delineate the brain tissue. Different source-receiver apertures are also selected for these two different setups with the inversion of the scalp primarily utilizing the reflection information while the brain tissue mainly relies on the transmission data. Utilizing such a two-step inversion process allows for one to leverage an increase in sensitivity achieved through modifying the source-receiver aperture to improve the quality of model updates for specific regions of the domain. Utilizing an aperture which primarily captures reflection information should, in principle, allow for an increase in the sensitivities in the scalp layer whereas using the transmission information should allow for an increase in the sensitivities within the brain tissue itself.

An important consideration in this two-step inversion process is to minimize numerical discrepancies within the coupling fluid, which are introduced by the first set of inversions prior to commencing with the second set of inversions. In cases where the geometry of the inversion domain is known, it is relatively simple to force the model updates to be within some specific region of the domain by setting the gradients of the misfit functional to zero outside of this so-called *region of interest*. However, since the geometry of the scalp is not known *a priori*, such a region of interest cannot be effectively defined. Thus, erroneous model updates within the coupling fluid are to be expected. The proposed strategy for mitigating these numerical discrepancies involves extracting the collection of points that approximate the scalp and then apply a convex hull algorithm to delineate the shape of the epidermis. All material parameters outside of this region which have been erroneously updated during the inversion can then be reset back to the initial material parameters consistent with those of the coupling fluid.

In order to test the efficacy of the aforementioned inversion strategies, five inversion configurations are considered. The differences in inversion quality are compared between different choices in STF (Ricker or filtered Ricker wavelets), inversion algorithm (single-step or two-step), and the type of physics considered within the prior model (purely acoustic versus acoustoelastic). A summary of the inversion configurations considered throughout this section can be seen in table 1.

Table 1: An overview of the primary inversions conducted throughout this study. The choice in STF, inversion algorithm, and physics incorporated in the skull model are varied between different inversion configurations.

	STF		Algorithm		Skull Model	
Inversion No.	Ricker	Filtered	Single-Step	Two-Step	Purely Acoustic	Acoustoelastic
1	✓		✓			✓
2	✓			✓		✓
3		✓	✓			✓
4		✓		✓		✓
5	✓			✓	✓	

Several input parameters are kept constant between all inversions conducted. In particular, the source center frequency is fixed at approximately 500 kHz, a fixed 25 source and 100 receiver aperture distributed radially at 12.5 cm from the domain’s center is used (see figure 3), and only acoustoelastic observed data is inverted for. It is assumed that the skull is immersed in a coupling fluid with material properties equivalent to water ($v_p = 1500$ m/s, $\rho = 1000$ kg/m³, and $v_s = 0$ m/s).

The incorporation of prior knowledge within the inversion process is vital to the success of FWI. An effective method which can be used for incorporating such prior knowledge is the choice of starting model used for the inversion process. Preliminary testing explored the possibility of using a simple homogeneous starting model, which consisted only of the material parameters of the coupling fluid. However, it quickly became apparent that the strong influence of the skull did not allow for the model to converge to a meaningful reconstruction due to the significant differences between the observed and the prior data. A prior model which incorporates the geometry

and material parameters of the inner and outer tables of the skull are assumed to be known *a priori* such as to initialize the model in model space within the neighborhood of the true material parameters. This choice of starting model is certainly a strong prior which is likely infeasible for many clinical applications outside of long-term patient monitoring where a baseline of the patient’s skull is first obtained using an alternate imaging modality. Strategies which may be employed in future studies to relax this restrictive prior are discussed in section 5.

Although the geometry of the inner and outer tables are assumed to be known *a priori* for all inversion configurations, the diploë layer is only assumed to be known within the prior model for the case where the purely acoustic prior model is utilized; all other inversion configurations use a prior model of the skull with uniform material parameters. The motivation for this assumption is to reduce the amount of inverse crime within the reconstruction given that the material parameters of the skull are effectively assumed to be different from those within the true model. The diploë layer is, however, included in the skull model for the case where a purely acoustic prior is considered such as to improve the likelihood that this model converges to a meaningful reconstruction.

Additional strategies which are employed to lessen the effects of inverse crime within the reconstruction include (1) using a higher meshing frequency for the true model than the one used during the inversion, and (2) representing the tissue parameters in the true model with a higher polynomial degree. This ensures that the two meshes are not identical such that the nodal positions between the meshes differ.

Further sources of uncertainty such as those introduced by the effects of attenuation as well as instrument noise are neglected for this study. These effects are certainly non-trivial (especially attenuation in the diploë; see Barger³⁷), but incorporating these effects is outside the scope of this study.

4. RESULTS

Having an intuition of the physical phenomena observed within the forward problem is important for understanding the strong suits as well as the shortcomings of the resulting waveform inversions. In particular, the differences in considering the skull as being either an acoustic or elastic medium are compared. An analysis of the gradients is also conducted such as to show the impact that one’s choice in source-receiver aperture has on the resulting sensitivities. Finally, the inversion results for the configurations outlined in table 1 are compared.

4.1 Acoustic versus Acoustoelastic Forward Problem

First, the wavefields for a simple ellipse shaped skull model are considered in order to highlight the differences between considering the skull as being either an acoustic or elastic medium. A simple setup is considered where an omnidirectional point source is positioned 5.75 cm from the left-hand side of the 5 mm thick skull. The resulting wavefields can be seen in figure 6.

The primary transmitted and reflected wavefronts are relatively similar for both the acoustic and acoustoelastic cases for low incidence angles. However, the acoustoelastic wavefield (containing both the acoustic and elastic effects of the domain) observed in figure 6b possesses significantly more secondary reflections when compared to the purely acoustic case in figure 6a. These strong secondary arrivals can be attributed to the coupling effects at the skull-tissue interface in the case where the skull is considered as an elastic medium; these effects are completely neglected in the purely acoustic case.

Similar effects are observed when considering the more realistic MIDA phantom as seen in figure 7. As described previously within figure 6, the primary transmitted wavefronts are relatively similar between the acoustic (figure 7a) and the acoustoelastic (figure 7b) wavefields at low incidence angles with the primary discrepancies between the two wavefields being observed in the later arrivals.

The consequence of these additional lower-amplitude arrivals which propagate behind the primary transmitted wavefront are particularly relevant when comparing the shot gathers in figure 8. Within figure 8a, the reflections which result from the interfaces between the different portions of the brain tissue appear to be similar in amplitude to the additional arrivals which pollute the wavefields within figure 8b. This suggests that the reflections which occur within the brain tissue will be quite difficult to discern from the coupling effects introduced by the elastic

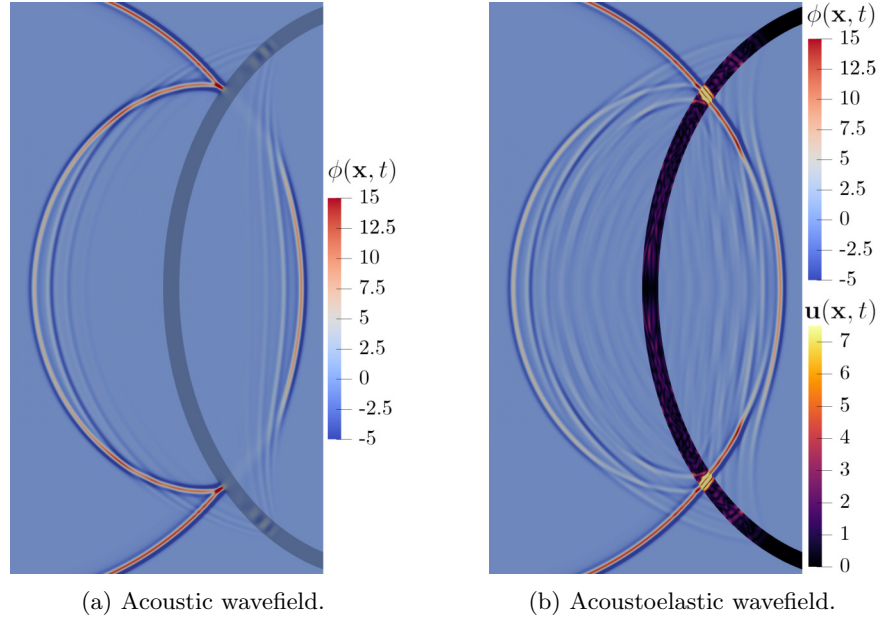


Figure 6: A comparison between the wavefields observed when considering a simple ellipse shaped skull with a source emanating from the left-hand side of the skull. The skull in this scenario possesses a uniform thickness of 5 mm, P-wave velocity of 2800 m/s, and density of 1800 kg/m³; the background material is assumed to be water. An S-wave velocity of 1500 m/s is assigned to the elastic skull in figure 6b.

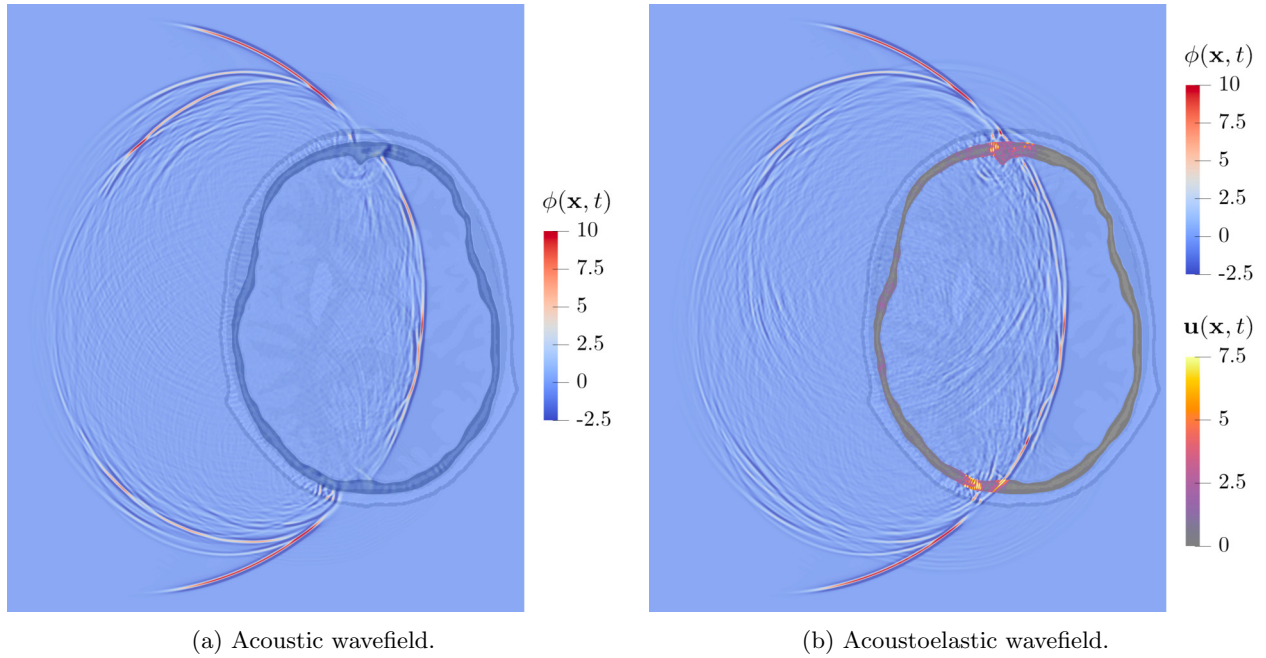


Figure 7: A comparison between the acoustic and the acoustoelastic wavefields observed for the MIDA phantom using a 500 kHz Ricker STF. Similar to the ellipse shaped skull in figure 6, the secondary arrivals within the acoustoelastic wavefield are of considerably higher amplitude when compared to those within the purely acoustic wavefield.

skull due to their similar amplitudes. It may be possible to discard these later arrivals in favor of preserving the primary reflected and transmitted wavefronts, which are similar between the acoustic and acoustoelastic cases for




Figure 8: A comparison between the shot gathers for a set of purely acoustic versus acoustoelastic data corresponding to the wavefields shown in figure 7. Note the absence of data at angle 0 due to this transducer position corresponding with the source position.

low incidence angles. However, discarding these subsequent arrivals may lead to an elimination of potentially useful information from the wavefield.

4.2 Sensitivity Analysis

Next, the impact on the gradients of the misfit functional are analyzed when considering the different source-receiver apertures discussed in section 3.3. Comparing the gradients $\nabla_{v_p} \chi$ for a) the aperture which favors reflection information versus b) the aperture that favors transmission information shows that considerably higher gradients are obtained within the scalp layer in a) whereas higher gradients are observed within the interior brain tissue for b). As one might expect, the magnitudes of the gradients for the reflection setup are considerably higher when compared to the gradients for the transmission setup. A comparison between these two gradients can be seen in figure 9.

Two important conclusions can be drawn from these gradients. First, adjusting the source-receiver apertures has a profound effect on the sensitivities within the regions inside and outside of the skull. Modifying the aperture to increase the sensitivities within the desired areas of the model allows for model updates to be preferentially assigned to specific regions of the domain. The second conclusion that can be deduced is that there is a fairly strong source imprint within the gradients, regardless of the aperture chosen. This source imprint is particularly strong closer to the source locations which demonstrates the need for applying a radial mask to the domain to suppress model updates that would otherwise be erroneously mapped to the source positions.

4.3 Inversion Results

The first set of inversions conducted utilize a Ricker STF within a single-step inversion setup. A region of interest mask is applied to force all model updates to be mapped to within the interior brain tissue; the scalp is not inverted for in this inversion configuration. Many of the larger features within the brain tissue are apparent (putamen or thalamus), although other features (smaller features at the interfaces between the gray and white matter) are not as clearly resolved. Areas proximal to local thickening in the scalp are particularly affected within the reconstruction. The reconstruction in the gray/white matter interfaces near the temporalis are particularly poorly resolved. This suggests that accounting for the presence of the scalp within the two-step inversion should lead to improvements in the reconstruction within these areas of the domain. The final inversion result for this setup can be seen in figure 11b.




Figure 9: The gradients $\nabla_{v_p}\chi$ for the MIDA phantom for two different source-receiver apertures. Considerably higher sensitivities are observed within the scalp tissue when an aperture which favors reflection information is used. Similarly, the sensitivities within the brain tissue are generally higher when using an aperture which favors transmission information. Notice the difference in color scales between both figures, thus demonstrating that the model is significantly more sensitive to the tissue outside of the skull.

The next inversion configuration utilizes a two-step inversion with a Ricker STF. First, the reflection source-receiver aperture is used to invert for the scalp tissue. Model updates within the interior brain tissue are still permitted during this first set of inversions, which allows for some structure to be recovered within the brain tissue before transitioning to the second set of inversions using the transmission aperture. As previously discussed, the high gradients near the source locations are removed using a simple radial mask, although some inversion artifacts are still erroneously introduced into the reconstruction within the coupling fluid. These are removed using the procedure described in section 3 to yield the intermediate results shown in figure 10.

The corrected intermediate results shown in figure 10b are then used as the new prior model within the second set of inversions where the transmission aperture is used to illuminate the interior brain tissue. Here the region of interest is modified to enforce all model updates to be within the brain tissue; the material parameters in the scalp are treated as fixed. The final reconstruction obtained for this inversion configuration shows well defined boundaries between the different tissue types, most notably between the white/gray matter interfaces which were poorly resolved within the previous single-step inversion. The final inversion result for this two-step inversion can be seen in figure 11c.

Next, the single-step inversion using a filtered Ricker STF is considered. The presence of lower frequencies within a regular Ricker wavelet is a key component which allows for the previous two inversions to converge to a meaningful solution without the use of any additional form of regularization. After eliminating the low-frequency content of the STF, the non-convexity of the optimization problem becomes more pronounced, thus increasing the likelihood that the optimization algorithm becomes trapped in a local minimum. This local minimum has the potential to correspond with a set of physically meaningless model parameters. To reduce the issues introduced by this increased non-convexity of the optimization problem, the gradients are smoothed between model updates using the diffusion equation to act as a Tikhonov-style regularization. It was determined empirically that a uniform smoothing length relative to the wavelength λ on the order of approximately λ to 2λ yielded favorable results for this inversion scheme. The inverted results for this single-step inversion are somewhat comparable to the case of the single-step inversion using a conventional Ricker wavelet, although the example using the Ricker wavelet possesses sharper material boundaries given that the gradients were not smoothed between inversion

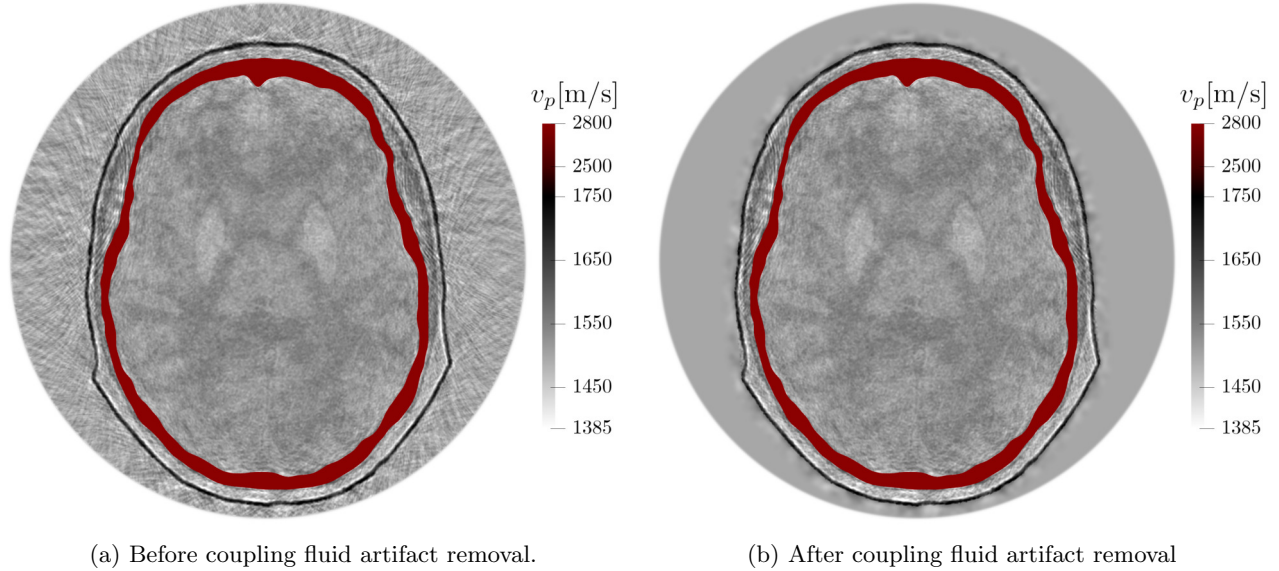


Figure 10: The removal of inversion artifacts within the coupling fluid within a set of intermediate results of the two-step inversion. This intermediate reconstruction is obtained from using the reflection aperture shown in figure 3b.

iterations. The final reconstruction for this inversion configuration can be seen in figure 11e.

The two-step inversion using the filtered Ricker STF was conducted using a similar smoothing technique as in the previous single-step case. Resolving the scalp to a level which is comparable in quality to that obtained using a conventional Ricker STF proved to be challenging; significantly more inversion artifacts became apparent within the coupling fluid when inverting for the scalp tissue. These increased inversion artifacts within the coupling fluid, in turn, caused an increased number of artifacts to appear within the scalp layer. In order to reduce the persistence of these inversion artifacts within the scalp tissue, a larger smoothing length of approximately 2λ was initially chosen for the set of inversions which utilize the reflection aperture. After removing the artifacts from the coupling fluid within the intermediate inversion result, a region of interest was defined which allowed for model updates to be mapped to both the scalp and the brain tissue. This strategy allowed for many of the remnant artifacts within the scalp layer to be eliminated during the second phase of inversions where the transmission information was utilized. The smoothing length is also reduced when considering the transmission data such as to allow for sharper boundaries to be resolved between tissues.

The majority of the features within the brain tissues that are apparent within the two-step inversion using a Ricker wavelet are also apparent within the reconstruction using a filtered Ricker wavelet. Similar to the comparison between the previous single-step inversions, however, shows that the features within the reconstruction using a filtered Ricker STF causes the boundaries between tissues to be less sharp given the need to smooth the model between inversion iterations. One of the primary areas of discrepancy between both of these two-step inversions is within the scalp tissue. The quality of this portion of the reconstruction is significantly blurred when using a filtered Ricker STF, although this appears to have a relatively limited influence on the quality of the subsequently recovered brain tissue. The final inversion result for this two-step inversion can be seen in figure 11f.

Finally, a set of inversions are conducted using a purely acoustic prior as opposed to an acoustoelastic prior model (as was the case in all other inversions conducted). The purpose of this inversion is to verify if it is still possible to converge to a meaningful reconstruction without applying any sort of additional processing step to mitigate or otherwise account for the elastic effects of the skull. The most numerically favorable aspects of the previous inversions are combined such as to increase the likelihood of the inversion yielding a physically realistic result. In particular, a skull prior which incorporates the structure of the inner and outer tables as well as the diploë layer is provided within the prior, an unfiltered Ricker wavelet is used, and a two-step inversion is executed.

As can be seen in figure 11d, the inversion is able to delineate the scalp layer with considerable accuracy, although the interior brain tissue remains unresolvable. There are certainly other strategies which may be employed here to mitigate these elastic effects within an acoustic-based inversion; this is discussed in more detail in section 5.

One of the considerable benefits of using the SEM for modeling media with distinct fluid-solid regions is that the elastic wave equation only needs to be solved within the elastic elements of the domain. Modeling the entire domain as a purely elastic medium may cause the computational cost to increase by orders of magnitude compared to the purely acoustic simulation. However, given that the proportion of elastic to acoustic elements within the domain is very low (typically around 2% of elements are elastic), the increase in computational cost is relatively small; much of the increase in computational cost can be attributed to the increased number of time steps. A comparison between acoustic versus elastic simulations can be seen in table 2.

Table 2: A comparison of the computation times required for forward simulations within an acoustic or acoustoelastic domain using 4th-order SEM. The ratio of the number of elements represents the percentage of elastic to acoustic elements within the mesh. Similarly, the ratio of computation time represents the increase in computation time going from a purely acoustic simulation to an acoustoelastic simulation. Each simulation was run on two Nvidia® Tesla P100 graphical processing units (GPUs).

Frequency [kHz]	Type of Physics	No. Elements			Time Steps		Comp. Time	
		Acoustic	Elastic	Ratio	No.	Ratio	Runtime [s]	Ratio
250	Acoustic	81196	N/A		14155	1.157	13	1.69
250	Acoustoelastic	80348	1696	2.07%	16376		22	
500	Acoustic	318780	N/A		28020	1.145	58	1.41
500	Acoustoelastic	315420	6720	2.09%	32095		82	
750	Acoustic	717952	N/A		41889	1.067	175	1.24
750	Acoustoelastic	710368	13904	1.92%	44710		217	
1000	Acoustic	1270080	N/A		55740	1.079	432	1.31
1000	Acoustoelastic	1256640	25200	1.97%	60148		568	

5. DISCUSSION

The results presented in section 4 allow for several important conclusions to be drawn with respect to the source type, inversion algorithm, and treating the domain as an acoustic versus acoustoelastic medium. In comparing the inversion cases using a Ricker STF (figures 11b and 11c) versus those using a filtered STF (figures 11e and 11f), it is clear that the narrower bandwidth has a profound effect on the convergence of the reconstruction. The absence of lower frequencies led to the requirement for additional regularization to ensure that a meaningful reconstruction could be obtained. Applying adequate regularization allows for the solution to converge to a more physically meaningful reconstruction at the cost of smoothing the final reconstruction obtained.

The inversion algorithm chosen (single-step or two-step inversion configuration; see figure 5) has a profound effect on the quality of the reconstruction obtained. The single-step inversion has the advantage of generally requiring fewer iterations to converge to a moderately accurate reconstruction when compared to the two-step inversion, but the single-step inversion is generally unable to resolve smaller-scale features within the reconstruction due to its tendency to converge to a local minimum. The quality of the inversion is particularly degraded in areas that are proximal to local thickening of the scalp tissue given that the presence of the scalp tissue is completely ignored within the single-step inversion. The two-step inversion produces comparatively superior results at the consequence of conducting two separate sets of inversions.

Considering the skull as either an acoustic or elastic medium appears to have a profound influence on the quality of the final inversion reconstruction when a set of acoustoelastic observed data is considered. Even




Figure 11: A comparison between the inverted results for the different inversion configurations outlined in table 1. The coupling fluid surrounding the skull has been removed here for the sake of visualization. The skull (which is included within the prior model) is indicated in red within these inversions.

under ideal conditions (Ricker STF and precise skull geometry known *a priori*), it is not possible to recover a meaningful reconstruction of the brain tissue without applying additional intermediate steps to account for the acoustoelastic coupling effects. A number of well-established processing methods exist within the geophysics community that allow for an acoustic-based approximation of the elastic wave equation, which may allow for these elastic effects to be sufficiently suppressed within an acoustic inversion scheme.^{38,39} The motivation for using such an acoustic-based approximation of the elastic wave equation within a geophysical context is to avoid incurring a significant increase in the computational cost of the simulation given that geophysical domains typically consist of predominantly elastic elements. However, the limited number of elastic elements within the inversion scheme presented within this study yields a comparatively small penalty in the computational cost when comparing the fully acoustic versus acoustoelastic wave simulations.

Quantitatively defining a stopping criterion for the inversion process remains a challenge for these *in silico* inversions. Most of the major structures within the domain are recovered within the first few iterations (usually <5 iterations), with the subsequent iterations allowing for the refinement of already delineated structures. Given the under-determined nature of this inverse problem, conducting the inversion process for too many iterations causes the model to introduce erroneous small-scale structures, which are not constrained by the data. This has the consequence of obscuring features which may have been apparent in earlier inversion iterations. Future work into assessing a suitable stopping criterion would certainly be an important consideration for automating such an FWI scheme.

Undoubtedly the most restrictive assumption within the current inversion scheme is the requirement for one to know the geometry of the skull *a priori*. As previously mentioned, this assumption would severely restrict the applicability of this inversion scheme in a clinical setting. A promising area of research which may allow for the geometry of the skull to be inverted for within a separate set of inversions is using some shape optimization routine,^{40,41} such as those used in salt body reconstructions within exploration seismology.⁴²

6. CONCLUSIONS

This *in silico* study assessed the ability to image the brain using transcranial USCT by employing an acoustoelastic FWI scheme. A SEM solver was used for effectively accounting for the interface conditions between the soft tissue (acoustic material) and the skull (elastic material) to accurately represent the acoustoelastic coupling effects within the domain. Several inversion configurations were considered using a 500 kHz STF consisting of either a Ricker wavelet or a bandpassed Ricker wavelet at $500 \text{ kHz} \pm 250 \text{ kHz}$. Different inversion strategies where the presence of the scalp is either considered or ignored were also compared within the so-called *single-step* or *two-step* inversion configurations. It was shown that using a Ricker STF yielded strong inversion results regardless of the inversion configuration chosen given the presence of lower frequencies within the unfiltered STF. Discarding the lower and higher frequencies within the STF degraded the overall quality of the reconstruction, although good results could still be achieved by smoothing the gradients of the misfit functional between model updates. Employing a purely acoustic prior by considering the skull as an acoustic medium introduced significant inversion artifacts within the reconstruction. Additional processing of the observed acoustoelastic data would be required to account for the elastic nature of the skull when utilizing a purely acoustic prior model.

Further research into expanding this method to 3D is required. Preliminary testing suggests that there are strong out-of-plane effects introduced by the skull, which cannot be accurately captured within a purely 2D model. The development of further pre-processing methods which can be used to construct a reasonable prior model of the skull are also required to relax the relatively restrictive assumption that the geometry of the skull needs to be known *a priori*. Finally, the incorporation of attenuation has been neglected for this study, although previous studies in the literature have demonstrated that the attenuating effects within such a domain (particularly within the skull) are significant.

ACKNOWLEDGMENTS

The authors would like to thank the support from the Swiss National Supercomputing Center (CSCS) in the form of the computing time grant s1040.

REFERENCES

- [1] Fry, F. J. and Barger, J. E., “Acoustical properties of the human skull,” *The Journal of the Acoustical Society of America* **63**, 1576–1590 (May 1978).
- [2] Dines, K. A., Fry, F. J., Patrick, J. T., and Gilmor, R. L., “Computerized ultrasound tomography of the human head: Experimental results,” *Ultrasonic Imaging* **3**, 342–351 (Oct. 1981).
- [3] Ylitalo, J., Koivukangas, J., and Oksman, J., “Ultrasonic reflection mode computed tomography through a skullbone,” *IEEE Transactions on Biomedical Engineering* **37**, 1059–1066 (Nov. 1990). Conference Name: IEEE Transactions on Biomedical Engineering.
- [4] White, P. J., Clement, G. T., and Hynynen, K., “Longitudinal and shear mode ultrasound propagation in human skull bone,” *Ultrasound in Medicine & Biology* **32**, 1085–1096 (July 2006).
- [5] Martin, B. and McElhaney, J. H., “The acoustic properties of human skull bone,” *Journal of Biomedical Materials Research* **5**(4), 325–333 (1971). eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jbm.820050405>.
- [6] Hayner, M. and Hynynen, K., “Numerical analysis of ultrasonic transmission and absorption of oblique plane waves through the human skull,” *The Journal of the Acoustical Society of America* **110**, 3319–3330 (Dec. 2001).
- [7] Salahshoor, H., Shapiro, M. G., and Ortiz, M., “Transcranial focused ultrasound generates skull-conducted shear waves: Computational model and implications for neuromodulation,” *Applied Physics Letters*, **6** (2020).
- [8] Tarantola, A., “Inversion of seismic reflection data in the acoustic approximation,” *Geophysics* **49**, 1259–1266 (Aug. 1984).
- [9] Mora, P., “Nonlinear two-dimensional elastic inversion of multioffset seismic data,” *Geophysics* **52**, 1211–1228 (Sept. 1987).
- [10] Virieux, J. and Operto, S., “An overview of full-waveform inversion in exploration geophysics,” *Geophysics* **74**, WCC1–WCC26 (Nov. 2009). Publisher: Society of Exploration Geophysicists.
- [11] Pratt, R. G., Huang, L., Duric, N., and Littrup, P., “Sound-speed and attenuation imaging of breast tissue using waveform tomography of transmission ultrasound data,” in [*Medical Imaging 2007: Physics of Medical Imaging*], Hsieh, J. and Flynn, M. J., eds., **6510**, 1523 – 1534, International Society for Optics and Photonics, SPIE (2007).
- [12] Wiskin, J., Borup, D., Johnson, q. S., Berggren, M., Abbott, T., and Hanover, R., “Full-wave, non-linear, inverse scattering,” in [*Acoustical Imaging*], 183–193, Springer (2007).
- [13] Wang, K., Matthews, T., and C. Li, F. A., Duric, N., and Anastasio, M., “Waveform inversion with source encoding for breast sound speed reconstruction in ultrasound computed tomography,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **62** (2014).
- [14] Goncharsky, A., Romanov, S. Y., and Seryozhnikov, S. Y., “A computer simulation study of soft tissue characterization using low-frequency ultrasonic tomography,” *Ultrasonics* **67**, 136–150 (2016).
- [15] Perez-Liva, M., Herraiz, J. L., Udías, J. M., Cox, B. T., and Treeby, B. E., “Full-wave attenuation reconstruction in the time domain for ultrasound computed tomography,” in [*2016 IEEE 13th International Symposium on Biomedical Imaging*], 710–713, IEEE (2016).
- [16] Boehm, C., Martiartu, N. K., Vinard, N., Balic, I. J., and Fichtner, A., “Time-domain spectral-element ultrasound waveform tomography using a stochastic quasi-Newton method,” in [*Medical Imaging 2018: Ultrasonic Imaging and Tomography*], Duric, N. and Byram, B. C., eds., **10580**, 92 – 100, International Society for Optics and Photonics, SPIE (2018).
- [17] Mitsuhashi, K., Poudel, J., Matthews, T. P., Garcia-Uribe, A., Wang, L. V., and Anastasio, M. A., “A Forward-Adjoint Operator Pair Based on the Elastic Wave Equation for Use in Transcranial Photoacoustic Computed Tomography,” *SIAM Journal on Imaging Sciences* **10**, 2022–2048 (Jan. 2017).
- [18] Shen, X., Ahmed, I., Brenders, A., Dellinger, J., Etgen, J., and Michell, S., “Full-waveform inversion: The next leap forward in subsalt imaging,” *The Leading Edge* **37**(1), 67b1–67b6 (2018). Publisher: Society of Exploration Geophysicists.
- [19] Guasch, L., Calderón Agudo, O., Tang, M.-X., Nachev, P., and Warner, M., “Full-waveform inversion imaging of the human brain,” *npj Digital Medicine* **3**, 28 (Feb. 2020).

- [20] Agudo, O. C., *Acoustic full-waveform inversion in geophysical and medical imaging*, PhD thesis, Imperial College London (Apr. 2018). <https://doi.org/10.25560/62620>.
- [21] Komatitsch, D., Barnes, C., and Tromp, J., “Wave propagation near a fluid-solid interface: A spectral-element approach,” *Geophysics* **65**, 623–631 (Mar. 2000).
- [22] Chaljub, E. and Valette, B., “Spectral element modelling of three-dimensional wave propagation in a self-gravitating earth with an arbitrarily stratified outer core,” *Geophysical Journal International* **158**, 131–141 (07 2004).
- [23] Nissen-Meyer, T., Fournier, A., and Dahlen, F. A., “A 2-D spectral-element method for computing spherical-earth seismograms-II. Waves in solid-fluid media,” *Geophysical Journal International* **174**, 873–888 (Sept. 2008).
- [24] Ferroni, A., Antonietti, P. F., Mazzieri, I., and Quarteroni, A., “Dispersion-dissipation analysis of 3-D continuous and discontinuous spectral element methods for the elastodynamics equation,” *Geophysical Journal International* **211**, 1554–1574 (Dec. 2017).
- [25] Afanasiev, M., Boehm, C., van Driel, M., Krischer, L., Rietmann, M., May, D. A., Knepley, M. G., and Fichtner, A., “Modular and flexible spectral-element waveform modelling in two and three dimensions,” *Geophysical Journal International* **216**, 1675–1692 (Mar. 2019).
- [26] Aki, K., Richards, P., and Books, U. S., [*Quantitative Seismology*], University Science Books (2002).
- [27] Sommerfeld, A., “Die Greensche Funktion der Schwingungsgleichung,” in [*Jahresbericht der Deutschen Mathematiker-Vereinigung*], **21**, 309–353 (1912).
- [28] Kosloff, R. and Kosloff, D., “Absorbing boundaries for wave propagation problems,” *Journal of Computational Physics* **63**, 363–376 (Apr. 1986).
- [29] Afanasiev, M., Boehm, C., van Driel, M., Krischer, L., and Fichtner, A., “Flexible high-performance multiphysics waveform modeling on unstructured spectral-element meshes,” in [*SEG Technical Program Expanded Abstracts 2018*], 4035–4039, Society of Exploration Geophysicists, Anaheim, California (Aug. 2018).
- [30] Fichtner, A., [*Full Seismic Waveform Modelling and Inversion*], Advances in Geophysical and Environmental Mechanics and Mathematics, Springer, Berlin Heidelberg (2011).
- [31] Nocedal, J. and Wright, S. J., [*Numerical optimization*], Springer series in operations research, Springer, New York, 2nd ed ed. (2006). OCLC: ocm68629100.
- [32] Griewank, A. and Walther, A., “Algorithm 799: Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation,” *ACM Trans. Math. Softw.* **26**, 19–45 (Mar. 2000).
- [33] Iacono, M. I., Neufeld, E., Akinagbe, E., Bower, K., Wolf, J., Vogiatzis Oikonomidis, I., Sharma, D., Lloyd, B., Wilm, B. J., Wyss, M., Pruessmann, K. P., Jakab, A., Makris, N., Cohen, E. D., Kuster, N., Kainz, W., and Angelone, L. M., “MIDA: A Multimodal Imaging-Based Detailed Anatomical Model of the Human Head and Neck,” *PLOS One* **10** (Apr. 2015).
- [34] Duck, F. A., “Chapter 4 - Acoustic Properties of Tissue at Ultrasonic Frequencies,” in [*Physical Properties of Tissues*], Duck, F. A., ed., 73–135, Academic Press, London (Jan. 1990).
- [35] IT’IS Foundation, “Tissue Properties Database V4.0,” (2018). DOI: 10.13099/VIP21000-04-0.
- [36] Marty, P., *Acoustoelastic Full-Waveform Inversion of the Brain with Transcranial Ultrasound*, Master’s thesis, TU Delft, ETH Zurich, RWTH Aachen (Aug. 2020).
- [37] Barger, J. E., “Attenuation and dispersion of ultrasound in cancellous bone,” in [*Ultrasonic Tissue Characterization II*], 197–201, National Bureau of Standards, Maryland (1977).
- [38] Chapman, C. H., Hobro, J. W. D., and Robertsson, J. O. A., “Correcting an acoustic wavefield for elastic effects,” *Geophysical Journal International* **197**, 1196–1214 (03 2014).
- [39] Calderón Agudo, O., da Silva, N. V., Warner, M., and Morgan, J., “Acoustic full-waveform inversion in an elastic world,” *GEOPHYSICS* **83**(3), R257–R271 (2018).
- [40] Cea, M., “Optimal design for 2d wave equations,” *Optimization Methods and Software* **32**(1), 86–108 (2017).
- [41] Cagnol, J. and Zolésio, J.-P., “Shape derivative in the wave equation with dirichlet boundary conditions,” *Journal of Differential Equations* **158**(2), 175 – 210 (1999).
- [42] Dahlke, T., Biondi, B., and Clapp, R., “Applied 3D salt body reconstruction using shape optimization with level sets,” *Geophysics* **85**(5), 437–446 (2020).