2-D Slicewise Waveform Inversion of Sound Speed and Acoustic Attenuation for Ring Array Ultrasound Tomography Based on a Block LU Solver

Rehman Ali, *Member, IEEE*, Trevor M. Mitcham, *Member, IEEE*, Thurston Brevett, *Student Member, IEEE*, Òscar Calderón Agudo, Cristina Durán Martinez, Cuiping Li, Marvin M. Doyley, *Fellow, IEEE*, and Nebojsa Duric

***Abstract*— Ultrasound tomography is an emerging imaging modality that uses the transmission of ultrasound through tissue to reconstruct images of its mechanical properties. Initially, ray-based methods were used to reconstruct these images, but their inability to account for diffraction often resulted in poor resolution. Waveform inversion overcame this limitation, providing high-resolution images of the tissue. Most clinical implementations, often directed at breast cancer imaging, currently rely on a frequency-domain waveform inversion to reduce computation time. For ring arrays, ray tomography was long considered a necessary step prior to waveform inversion in order to avoid cycle skipping. However, in this paper, we demonstrate that frequency-domain waveform inversion can reliably reconstruct high-resolution images of sound speed and attenuation without relying on ray tomography to provide an initial model. We provide a detailed description of our frequency-domain waveform inversion algorithm with open-source code and data that we make publicly available.**

***Index Terms*—Frequency domain, waveform inversion, tomography, ultrasound, ring array.**

Manuscript received 3 January 2024; revised 27 February 2024; accepted 27 March 2024. Date of publication 2 April 2024; date of current version 1 August 2024. This work was supported in part by the University of Rochester Medical Center and in part by the National Institute of Biomedical Imaging and Bioengineering (NIBIB) under Grant F32-EB034589 and Grant R01-EB028652. *(Corresponding author: Rehman Ali.)*

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Institutional Review Board (IRB) under Approval IRB No. 040912M1F and the Health Insurance Portability and Accountability Act (HIPAA).

Rehman Ali and Trevor M. Mitcham are with the Department of Imaging Sciences, University of Rochester Medical Center, Rochester, NY 14642 USA (e-mail: Rehman_Ali@URMC.Rochester.edu; Trevor_Mitcham@URMC.Rochester.edu).

Thurston Brevett is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: tbrevett@stanford.edu).

Òscar Calderón Agudo is with FrontWave Imaging, 08025 Barcelona, Spain, and also with the Department of Earth Science and Engineering, Imperial College London, SW7 2AZ London, U.K. (e-mail: ocalderon@frontwave.io).

Cristina Durán Martinez is with FrontWave Imaging, 08025 Barcelona, Spain (e-mail: cduran@frontwave.io).

Cuiping Li was with the Karmanos Cancer Institute, Wayne State University, Detroit, MI 48201 USA. She is now with Delphinus Medical Technologies, Novi, MI 48374 USA (e-mail: cli0713@gmail.com).

Marvin M. Doyley is with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627 USA (e-mail: m.doyley@rochester.edu).

Nebojsa Duric was with the Karmanos Cancer Institute, Wayne State University, Detroit, MI 48201 USA. He is now with the Department of Imaging Sciences, University of Rochester Medical Center, Rochester, NY 14642 USA, and also with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627 USA (e-mail: Nebojsa_Duric@URMC.Rochester.edu).

Digital Object Identifier 10.1109/TMI.2024.3383816

I. Introduction

ULTRASOUND tomography (UST) is an emerging imaging modality that uses the transmission of ultrasound through tissue to accurately reconstruct high-resolution images of tissue properties such as sound speed and acoustic attenuation. The primary application of UST is breast imaging [1], [2], [3], [4], where sound speed and acoustic attenuation are key biomarkers for identifying cancer in the breast. Unlike pulse-echo ultrasound, which solely relies on reflected signals to create a band-limited and often aberrated [5] estimate of the differential impedance in the tissue (i.e., the B-mode image), UST provides absolute estimates of material properties such as sound speed [1], density [6], [7], and acoustic attenuation [7], [8], [9]. Although significant progress has been made to characterize sound speed and acoustic attenuation with pulse-echo imaging [10], [11], [12], it is often difficult to estimate absolute values [11], [13] without first applying prior knowledge about the medium [14], [15], [16].

The first algorithms used in UST to reconstruct sound speed were ray methods that modeled the time-of-flight of the ultrasound wave as the line integral over slowness (i.e., the reciprocal of sound speed). Early work [17] assumed straight line paths for ease of reconstruction; however, "bent-ray" methods [1], [18] later incorporated refraction. The main difficulties with ray tomography were time-of-flight picking [19] and inherent resolution limits [20] that stemmed from the inability of time-of-flight to capture complex diffractive effects that can only be observed in the ultrasound waveform.

As a result, the field increasingly adopted waveform inversion, an approach that directly penalizes the mismatch between simulated and measured waveforms. The main advantage of the ray tracing methods is their fast reconstruction times relative to waveform inversion, especially for 3D imaging [21]. Ongoing improvements rely on various approximations to either further increase the speed of ray tracing methods [22] or improve their accuracy [23], [24]. Because the simulated waveforms are parameterized by estimates of the material properties, waveform inversion recovers images of these material properties. A significant portion of the literature is devoted to time-domain waveform inversion [9], [25], [26], [27], [28] because of its versatility; however, most clinical implementations of UST instead rely on frequency-domain waveform inversion to minimize computation time [4], [29].

Frequency domain modeling currently falls under three main categories: 1) direct solutions to the Helmholtz equation [7], [29], 2) iterative Born series solutions such as contrast source inversion [30], and 3) paraxial approximations [4], [31]. Although an open-source implementation of the paraxial approximation has been made available [31], the paraxial approximation often has limited utility outside of linear array UST. Instead, the current state-of-the-art in ring array UST [29] relies on a direct solution to the Helmholtz equation. Although the discretization of the Helmholtz equation is well-described in [29], the efficient solution of this discretized system is not fully explained. This is partly because the sparsity pattern of the discretized system often requires highly specialized solutions for fast performance, many of which are locked behind commercial mathematical softwares.

In this work, we describe an efficient approach that takes advantage of the block structure of the 2D discretized Helmholtz equation to leverage parallel computations on the GPU. We make the entire algorithm open source so that others can improve upon our work. We also make *in-vitro* and *in-vivo* experimental data publicly available with this work. We present many practical details regarding the application of waveform inversion to experimental data. As compared to prior work [29], this work does not rely on ray tomography to provide a starting model for waveform inversion and is the first work to demonstrate the frequency-domain waveform inversion of both sound speed and acoustic attenuation with *in-vitro* and *in-vivo* experimental data after an initial *in-silico* validation. The 2D-slice reconstructions of sound speed and acoustic attenuation are performed using a ring array transducer geometry.

II. Theory

A. Waveform Inversion Using the Helmholtz Equation

Using finite difference methods, an optimal 9-point stencil [32], [33] is used to discretize the two-dimensional Helmholtz equation with perfectly matched layers (PML) [34] to obtain a pressure field $\mathbf{u}$ for a given source δ, frequency $\omega = 2\pi f$, and slowness $\mathbf{s}$ (i.e., the reciprocal of sound speed):

$$(\nabla^2 + \omega^2 \mathbf{s}^2)\mathbf{u}(\omega, \mathbf{s}) = \delta(\omega), \tag{1}$$

After discretization, the $(\nabla^2 + \omega^2 \mathbf{s}^2)$ operator becomes a sparse matrix A that encodes the finite difference scheme so that equation (1) can be described using the following matrix-vector formulation (where $\mathbf{u}$ and δ are now vectors over spatial locations on the discretized grid):

$$A(\omega, \mathbf{s})\mathbf{u}(\omega, \mathbf{s}) = \delta(\omega), \tag{2}$$

Solving for $\mathbf{u}$ gives us that $\mathbf{u} = A^{-1}\delta$; however, this requires an efficient inversion of the discretized system A. This point regarding the inversion of A will be explained later in section II-E. In this section, the formulation in (2) is used to derive the overall waveform inversion algorithm. Differentiating both sides of equation (2) with respect to $\mathbf{s}$ gives us:

$$\frac{\partial \mathbf{u}}{\partial \mathbf{s}} = -A^{-1}\frac{\partial A}{\partial \mathbf{s}}\mathbf{u}. \tag{3}$$

Rather than observing the pressure field $\mathbf{u}$ over all space, we only observe the pressure field $\mathbf{p}$ at the ring array. The relationship between $\mathbf{p}$ and $\mathbf{u}$ is: $\mathbf{p} = K\mathbf{u}$, where K samples the pressure field at each element of the ring array. The derivative of $\mathbf{p}$ with respect to $\mathbf{s}$ is

$$\frac{\partial \mathbf{p}}{\partial \mathbf{s}} = -KA^{-1}\frac{\partial A}{\partial \mathbf{s}}\mathbf{u}. \tag{4}$$

Waveform inversion is posed as a least-squares problem:

$$E(\omega, \mathbf{s}) = \frac{1}{2}\sum_{i=1}^{N}\left\|\mathbf{p_i}(\omega, \mathbf{s}) - \mathbf{p_{obs,i}}(\omega)\right\|_2^2, \tag{5}$$

where $\mathbf{p_i}$ and $\mathbf{p_{obs,i}}$ are the simulated and observed pressure measurements for the i^{th} single-element transmission $\delta_\mathbf{i}$ on the ring array (i ranging from 1 to N, the number of ring array elements). The gradient of $E(\omega, \mathbf{s})$ with respect to $\mathbf{s}$ is:

$$\begin{aligned}\nabla_\mathbf{s} E &= \sum_{i=1}^{N} Re\left\{\left(\frac{\partial \mathbf{p_i}}{\partial \mathbf{s}}\right)^H \left(\mathbf{p_i} - \mathbf{p_{obs,i}}\right)\right\}\\ &= -\sum_{i=1}^{N} Re\left\{\left(\frac{\partial A}{\partial \mathbf{s}}\mathbf{u_i}\right)^H (A^H)^{-1}K^T\left(\mathbf{p_i} - \mathbf{p_{obs,i}}\right)\right\}.\end{aligned} \tag{6}$$

The residuals $(\mathbf{p_i} - \mathbf{p_{obs,i}})$ between the simulated and measured pressure fields are injected into the grid at the measurement locations via the adjoint of the sampling operator K^T. The residuals placed at the measurement locations then acts as sources for an adjoint Helmholtz equation $(A^H)^{-1}$, which yields a backpropagated field of the error between the simulation and the measurement. This backpropagated field is then demodulated by a virtual source $(\frac{\partial A}{\partial \mathbf{s}}\mathbf{u_i}) \approx (2\omega^2 \mathbf{s} \odot \mathbf{u_i})$, where $\odot$ refers to the pointwise multiplication between the slowness vector $\mathbf{s}$ and the full wavefield $\mathbf{u_i}$. Because the derivative of $A(\mathbf{s}) \approx (\nabla^2 + \omega^2\mathbf{s}^2)$ with respect to $\mathbf{s}$ is $2\omega^2\mathbf{s}$, $\frac{\partial A}{\partial \mathbf{s}}$ is a pointwise multiplication with $2\omega^2\mathbf{s}$

To implement the conjugate gradient (CG) algorithm for this objective function (5), we also need to calculate the linearized perturbation in the pressure field $\Delta \mathbf{p_i}$ for a given perturbation $\mathbf{d_k}$ to the slowness:

$$\Delta \mathbf{p_i} = \left(\frac{\partial \mathbf{p_i}}{\partial \mathbf{s}}\right) \mathbf{d_k} = -KA^{-1}\left(\frac{\partial A}{\partial \mathbf{s}}\mathbf{u_i}\right) \mathbf{d_k}. \tag{7}$$

This perturbation $\Delta \mathbf{p_i}$ allows us to properly scale the search direction $\mathbf{d_k}$ and update the slowness model $\mathbf{s_k}$:

$$\alpha_k = -\frac{\mathbf{d_k}^T \nabla_{\mathbf{s}} E(\mathbf{s}_k)}{\sum_{i=1}^{N} \|\Delta \mathbf{p_i}\|_2^2} \tag{8}$$

$$\mathbf{s_{k+1}} = \mathbf{s_k} + \alpha_k \mathbf{d_k} \tag{9}$$

Note that first iteration of CG is simply gradient descent (i.e., $\mathbf{d_1} = -\nabla_{\mathbf{s}} E(\mathbf{s}_1)$). Successive iterations of CG update the search direction according to the following formulas:

$$\beta_{FR,k} = \frac{\nabla_{\mathbf{s}} E(\mathbf{s}_k)^T \nabla_{\mathbf{s}} E(\mathbf{s}_k)}{\nabla_{\mathbf{s}} E(\mathbf{s}_{k-1})^T \nabla_{\mathbf{s}} E(\mathbf{s}_{k-1})}, \tag{10}$$

$$\beta_{PR,k} = \frac{\nabla_{\mathbf{s}} E(\mathbf{s}_k)^T \left(\nabla_{\mathbf{s}} E(\mathbf{s}_k) - \nabla_{\mathbf{s}} E(\mathbf{s}_{k-1})\right)}{\nabla_{\mathbf{s}} E(\mathbf{s}_{k-1})^T \nabla_{\mathbf{s}} E(\mathbf{s}_{k-1})}, \tag{11}$$

$$\beta_k = \min\left(\max\left(\beta_{PR,k}, 0\right), \beta_{FR,k}\right), \tag{12}$$

$$\mathbf{d}_k = -\nabla_{\mathbf{s}} E(\mathbf{s}_k) + \beta_k \mathbf{d}_{k-1}. \tag{13}$$

CG can be interpreted as gradient descent with a momentum β_k that depends on the current and previous gradient. We balance Fletcher-Reeves (FR) and Polak-Ribière (PR) to avoid the pitfalls associated with each formula [35].

B. Sound Speed and Attenuation

The slowness $\mathbf{s}$ can be complex-valued:

$$\mathbf{s} = \mathbf{s_R} + j\mathbf{s_I} \tag{14}$$

where the real part $\mathbf{s_R}$ corresponds to the reciprocal of sound speed and the imaginary part $\mathbf{s_I}$ corresponds to a frequency-dependent attenuation. Therefore, waveform inversion can be used to reconstruct sound speed and attenuation by replacing $\mathbf{s}$ with either $\mathbf{s_R}$ for sound speed or $\mathbf{s_I}$ for attenuation in all of the equations described in the previous section II-A. The only difference is in the virtual source term $(\frac{\partial A}{\partial \mathbf{s}}\mathbf{u_i})$:

$$\left(\frac{\partial A}{\partial \mathbf{s_R}}\mathbf{u_i}\right) = \frac{\partial \mathbf{s}}{\partial \mathbf{s_R}}\frac{\partial A}{\partial \mathbf{s}}\mathbf{u_i} = \frac{\partial A}{\partial \mathbf{s}}\mathbf{u_i} \approx (2\omega^2 \mathbf{s} \odot \mathbf{u_i}), \tag{15}$$

$$\left(\frac{\partial A}{\partial \mathbf{s_I}}\mathbf{u_i}\right) = \frac{\partial \mathbf{s}}{\partial \mathbf{s_I}}\frac{\partial A}{\partial \mathbf{s}}\mathbf{u_i} = j\frac{\partial A}{\partial \mathbf{s}}\mathbf{u_i} \approx (j2\omega^2 \mathbf{s} \odot \mathbf{u_i}), \tag{16}$$

where $(\frac{\partial A}{\partial \mathbf{s_R}}\mathbf{u_i})$ and $(\frac{\partial A}{\partial \mathbf{s_I}}\mathbf{u_i})$ differ by a factor of $j = \sqrt{-1}$.

C. Source Estimation

The waveform inversion described in section II-A assumes that the correct phase and amplitude is used for each single-element source $\delta_\mathbf{i}$; however, in practice, the phase and amplitude of each source must be estimated from measured waveform. The following projection formula provides a complex frequency-dependent scaling $\gamma_i(\omega; \mathbf{s})$ that accounts for both the amplitude and phase from each source i:

$$\gamma_i(\omega; \mathbf{s}) = \frac{\mathbf{p_i}(\omega, \mathbf{s})^H \mathbf{p_{obs,i}}(\omega)}{\mathbf{p_i}(\omega, \mathbf{s})^H \mathbf{p_i}(\omega, \mathbf{s})} \tag{17}$$

The source vector and pressure fields are scaled to $\{\delta_\mathbf{i} \to \gamma_i \delta_\mathbf{i}\}$, $\{\mathbf{u_i} \to \gamma_i \mathbf{u_i}\}$, and $\{\mathbf{p_i} \to \gamma_i \mathbf{p_i}\}$, respectively. The source scaling γ_i also depends on the current estimate of the complex slowness $\mathbf{s}$ and will be updated with each new estimate. Rather than estimating the source based on a calibration prior to waveform inversion, the source scaling is dynamically computed and applied during waveform inversion.

D. Ring Element Gridding Corrections

When discretizing the Helmholtz equation over a spatial grid, many ring elements will not fall exactly on the grid. In these cases, the ring elements are placed at the nearest grid point; however, this results in spatial discretization errors, which we correct by applying the following phase adjustment to data $p_{obs,i,v}(\omega)$ collected at receiver v for transmitter i:

$$p_{obs,i,v}(\omega) \to p_{obs,i,v}(\omega)\exp\left(j\omega\left(d_{iv}^{disc} - d_{iv}\right)/c\right) \tag{18}$$

where d_{iv} is the true distance between the transmitter and receiver, d_{iv}^{disc} is the distance between the discretized positions for the transmitter and receiver, and c is the sound speed near the transducer. In practice, c does not need to be known precisely and can be the same as the sound speed used in the homogeneous model to begin waveform inversion.

However, the phase adjustment described by equation (18) only corrects the impact of gridding errors on the signal transmitted directly from transmitter i to receiver v. It cannot correct for the impact of gridding errors on the backscattered received signals. Because this work only considers the through transmission, the issue of the backscattered received signals is beyond the scope of this work. Although not implemented in this work, carefully chosen interpolation schemes [36], [37] can be used to place point sources at fractional grid locations, solving the problem for both the through transmission and backscattered received signals.

E. Block LU Solver for the Helmholtz Equation

This section examines the structure of A to efficiently invert A for waveform inversion. With a 9-point stencil on a $M \times J$ grid, A has the following block-tridiagonal structure [38]:

$$A = \begin{bmatrix} D_1 & U_1 & & & \\ L_1 & D_2 & U_2 & & \\ & \ddots & \ddots & \ddots & \\ & & L_{J-2} & D_{J-1} & U_{J-1} \\ & & & L_{J-1} & D_J \end{bmatrix}, \tag{19}$$

where L_n, D_n, and U_n, are $M \times M$ matrix blocks (n ranging from 1 to J) each having their own tridiagonal structure. The algorithms in the following subsections describe the details of the block LU solver for the Helmholtz equation based on this block tridiagonal matrix structure. These algorithms were specifically developed in CUDA [39] to leverage parallel computations on the GPU.

1) Block LU Factorization of the Helmholtz Equation: The block LU factorization [38] $A = LU$ yields lower and upper block triangular matrices L and U:

$$L = \begin{bmatrix} T_1 & & & & \\ L_1 & T_2 & & & \\ & \ddots & \ddots & & \\ & & L_{J-2} & T_{J-1} & \\ & & & L_{J-1} & T_J \end{bmatrix}, \tag{20}$$

$$U = \begin{bmatrix} I & T_1^{-1}U_1 & & & \\ & I & T_2^{-1}U_2 & & \\ & & \ddots & \ddots & \\ & & & I & T_{J-1}^{-1}U_{J-1} \\ & & & & I \end{bmatrix}, \tag{21}$$

where the Schur complements T_n satisfy the recurrence

$$\begin{cases} T_1 = D_1 \\ T_n = D_n - L_{n-1}T_{n-1}^{-1}U_{n-1} \quad \text{for } n = 2, \ldots, J. \end{cases} \tag{22}$$

2) Applying the Block LU Factorization: The block LU factorization allows the discretized Helmholtz equation [$A\mathbf{u} = (LU)\mathbf{u} = L(U\mathbf{u}) = \delta$] (2) to be expressed as

$$\begin{cases} L\mathbf{v} &= \delta \\ U\mathbf{u} &= \mathbf{v}, \end{cases} \tag{23}$$

where

$$\mathbf{v} = \begin{bmatrix} \vec{v}_1 \\ \vec{v}_2 \\ \vdots \\ \vec{v}_{J-1} \\ \vec{v}_J \end{bmatrix} \quad \delta = \begin{bmatrix} \vec{\delta}_1 \\ \vec{\delta}_2 \\ \vdots \\ \vec{\delta}_{J-1} \\ \vec{\delta}_J \end{bmatrix} \quad \mathbf{u} = \begin{bmatrix} \vec{u}_1 \\ \vec{u}_2 \\ \vdots \\ \vec{u}_{J-1} \\ \vec{u}_J \end{bmatrix}. \tag{24}$$

The L-system in (23) can be solved using a forward sweep:

$$\begin{cases} \vec{v}_1 = T_1^{-1}\vec{\delta}_1 \\ \vec{v}_n = T_n^{-1}(\vec{\delta}_n - L_{n-1}\vec{v}_{n-1}) \quad \text{for } n = 2, \ldots, J. \end{cases} \tag{25}$$

The U-system in (23) can be solved using a backward sweep:

$$\begin{cases} \vec{u}_J = \vec{v}_J \\ \vec{u}_n = \vec{v}_n - T_n^{-1}U_n\vec{u}_{n+1} \quad \text{for } n = J-1, \ldots, 1. \end{cases} \tag{26}$$

Alternatively, if the adjoint system of equations [$A^H\mathbf{u} = (LU)^H\mathbf{u} = U^H(L^H\mathbf{u}) = \delta$] needs to be solved (e.g., the gradient calculation (6)), the block LU factorization yields:

$$\begin{cases} U^H\mathbf{v} &= \delta \\ L^H\mathbf{u} &= \mathbf{v}, \end{cases} \tag{27}$$

The U^H-system in (27) can be solved using a forward sweep:

$$\begin{cases} \vec{v}_1 = \vec{\delta}_1 \\ \vec{v}_n = \vec{\delta}_n - U_{n-1}^H(T_{n-1}^{-1})^H\vec{v}_{n-1} \quad \text{for } n = 2, \ldots, J. \end{cases} \tag{28}$$

The L^H-system in (27) can be solved using a backward sweep:

$$\begin{cases} \vec{u}_J = (T_J^{-1})^H\vec{v}_J \\ \vec{u}_n = (T_n^{-1})^H(\vec{v}_n - L_n^H\vec{u}_{n+1}) \quad \text{for } n = J-1, \ldots, 1. \end{cases} \tag{29}$$

Because each sweep relies on T_n^{-1}'s, we store T_n^{-1}'s rather than the T_n's when applying the recurrence relation (22). However, special care is needed when inverting the T_n's.

3) Inverting the Schur Complements: In each sweep, the T_n^{-1} matrices transform the pressure fields ($\vec{v}_n$ or $\vec{u}_n$) from one plane to the next one. However, the invertibility of each T_n is related to the boundary conditions applied to the Helmholtz equation. We apply the Dirichlet boundary condition that the pressure field is zero at the edges of the domain (i.e., first and last entries of each $\vec{v}_n$ or $\vec{u}_n$ vector are zero). This Dirichlet boundary condition is enforced by making the first and last entries of $\vec{\delta}_n$ zero. Additionally, the first and last rows of each L_n and U_n must be zero because there is no 9-point stencil for these entries corresponding to the boundary condition. As a result, the first and last entries of $\vec{\delta}_n - L_{j-1}\vec{v}_{n-1}$ and $U_n\vec{u}_{n+1}$ must also be zero. According to the forward (25) and backward (26) sweeps, this means that T_n^{-1} is only ever used to transform a vector whose first and last entries are zero to another vector whose first and last entries are zero.

Thus, it should be no surprise that T_n has a null space corresponding to the first and last entries of the input vector, which means that T_n is not invertible. However, there is no need to invert for the first and last entries because they should be zero; T_n only needs to be inverted over the interior entries of the vector, so we identify T_n^{int} the interior part of T_n that excludes the first and last row and column:

$$T_n = \begin{bmatrix} * & \ldots & * \\ \vdots & T_n^{int} & \vdots \\ * & \ldots & * \end{bmatrix}. \tag{30}$$

For all equations involving T_n^{-1}, we define T_n^{-1} as the inverse only over the interior entries in T_n^{int}, surrounded by zeros:

$$T_n^{-1} := \begin{bmatrix} 0 & \ldots & 0 \\ \vdots & (T_n^{int})^{-1} & \vdots \\ 0 & \ldots & 0 \end{bmatrix}. \tag{31}$$

III. Methods

A. Numerical Simulations

Figure 1 shows the ground-truth sound speed and attenuation maps for two numerical breast phantoms (a) and (b) used in k-Wave [40] to simulate single-element transmissions (1.0 MHz center frequency; 75% fractional bandwidth) from a ring array (22 cm diameter; 512 elements). Received channel data (time traces for the received signal from each individual element for each single-element transmission) from these k-Wave simulations were used to reconstruct the sound speed and attenuation in the medium using waveform inversion. No noise was added to the simulated channel data.

Phantom (a) was adapted and modified from a contrast-enhanced cone-beam breast CT image in work from the Diagnostic Breast Center Göttingen (https://doi.org/10.1016/j.tranon.2017.08.010) [41] under the Creative Commons Attribution-NonCommercial-No Derivatives License (CC BY NC ND) (https://creativecommons.org/licenses/by-nc-nd/4.0/). Phantom (b) was adapted from a 3D breast MRI dataset [42]. Root-mean-square (RMS) errors of the reconstructed sound

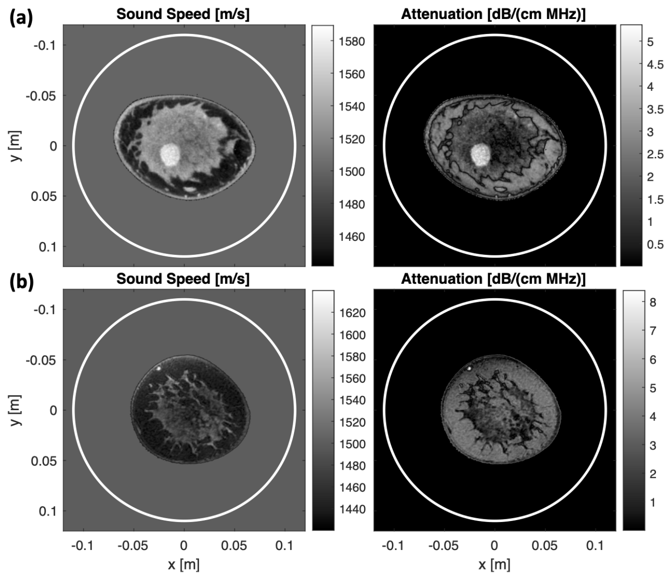

Fig. 1. Sound speed and attenuation maps [m/s] for two numerical breast phantoms (a) and (b). A 22 cm diameter ring array with 512 elements (white dots) was simulated in k-Wave for each phantom.

speed and attenuation are reported inside an 18 cm diameter circle at the center of the ring for each of the two phantoms.

B. Experiments

1) In-Vivo Acquisitions: Raw ultrasound channel data from an early UST prototype [43] was used to reconstruct sound speed and attenuation in the breast (data collected under IRB No. 040912M1F). This breast UST prototype consisted of a 22 cm diameter ring transducer with 1024 elements and a pulse center frequency of 2.5 MHz. Receive channel data from all 1024 elements were collected for each single-element transmit on the ring, resulting in a full 1024 × 1024 matrix capture of time series channel data (sampled at 12 MHz). The channel data was acquired with sufficient signal-to-noise ratio (SNR) for waveform inversion from 0.3 to 5.5 MHz.

2) In-Vitro Phantom Experiments: The complete experimental setup for the phantom experiments is shown in Figure 2. A 22 cm diameter ring transducer with 1024 elements (Transducer Works, Centre Hall, PA, USA) attached to a Verasonics Vantage 256 scanner via a UTA1024-MUX connector (Verasonics Inc., Kirkland, WA, USA) was used to collect the same full 1024 × 1024 matrix capture (described in III-B.1) from a Yezitronix Multi-modality Breast Phantom model B-MM-1.2 (Yezitronix Group Inc., Saint-Laurent, QC, Canada). Using the UTA1024-MUX, the acquisition sequence consisted of single-element transmissions, repeated four times for each 256-element subaperture of receivers (elements 1-256, 257-512, 513-768, and 769-1024), resulting in a total of 4096 transmit events to collect the full 1024 × 1024 matrix capture. The time-series channel data was sampled at 12.5 MHz.

However, as compared to the breast UST prototype, the Transducer Works ring only had content with sufficient SNR down to 0.7 MHz when excited at 2.5 MHz. In order to avoid cycle skipping [44], the Transducer Works ring was excited

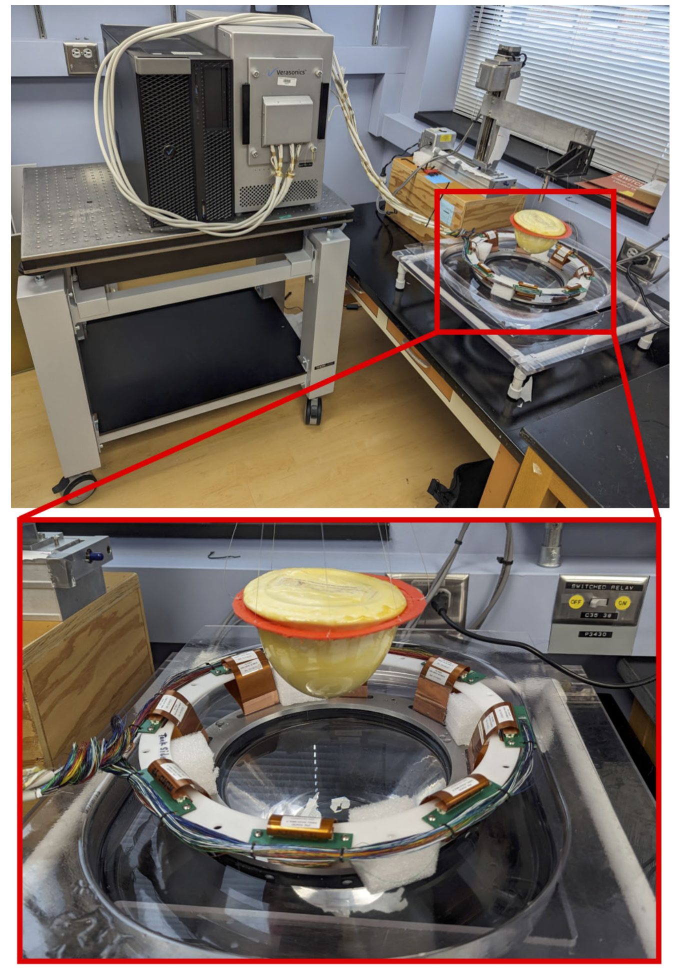

Fig. 2. Phantom experiment setup. The Verasonics data acquisition system is connected to the ring transducer (22 cm diameter, 1024 elements). The Yezitronix Multi-modality Breast Phantom was lowered into the water bath with the ring transducer using the vertically placed translation stage shown.

at 0.7 MHz to extract sufficient SNR down to 0.35 MHz. However, since 0.7 MHz is well outside the natural excitable range for the ring, nonlinearities cause the phase response to become significantly nonuniform across the ring, rendering content above 0.7 MHz unusable for waveform inversion.

C. Bootstrapping in Frequency

Waveform inversion is performed from low frequency to high frequency where the reconstruction resulting from the lowest frequency is used as the initial model for the reconstruction at the next frequency, and so forth. This bootstrapping approach begins at low frequencies to avoid cycle skipping [31], [44]. Higher frequencies then increase the resolution of the reconstructed images. In all imaging cases, our initial model prior to waveform inversion is a homogeneous 1480 m/s sound speed and zero attenuation in the medium.

Because sound speed is more closely related to the phase of the waveform, an imperfect attenuation model does not significantly impact sound speed reconstruction. On the other hand, attenuation tomography is far more sensitive to phase and amplitude errors caused by imperfections in the sound speed model. We therefore reconstruct images in two passes, each going from low to high frequency: on the first pass, we only update sound speed; on the second pass, we alternate between updates to sound speed and attenuation, first updating sound speed then attenuation at each frequency. We use 3 iterations of CG to update each material parameter at each frequency. The purpose of the first pass is to obtain an accurate sound speed estimate, which is used on the second pass to more accurately reconstruct attenuation.

We now describe the exact schedule of frequencies used during waveform inversion on our (1) *in-silico* simulation data, (2) *in-vivo* breast acquisitions, and (3) *in-vitro* phantom data. In each case, the frequencies used on the second pass are slightly shifted relative to the frequencies used on the first pass to avoid overfitting to data at the same frequencies.

1) In-Silico Simulation Data: For our simulated data, waveform inversion uses frequencies ranging from 0.3 to 1.25 MHz in 50 kHz steps on the first pass, and 0.325 to 1.275 MHz in 50 kHz steps on the second pass.

2) In-Vivo Breast Data: For our breast data, we applied the same schedule of frequencies used on the simulated data.

3) In-Vitro Phantom Data: Because the usable bandwidth of the transducer used for these experiments was much narrower (see section III-B.2), waveform inversion was done from 0.35 to 0.65 MHz in 50 kHz steps on the first pass and 0.375 to 0.675 MHz in 50 kHz steps on the second pass.

D. Isolating the Through-Transmission

For transmission tomography, we isolate the direct through-transmission by only using receive channels from the 75% of the ring opposite from the transmitter for waveform inversion. For the 512 × 512 full matrix capture, this is equivalent to excluding receive channel data from the transmitter and 63 elements on each side of the transmitting element (excluding a total of 127 receive channels). For the 1024 × 1024 full matrix capture, we exclude receive channels from the transmitter and 127 elements on either side (excluding a total of 255 receive channels). The main reason for excluding receive channels closest to the transmitter is that element directivity (on both transmission and reception) is not modeled in our simulations. Therefore, the simulation of the through transmission becomes less accurate as the receiver becomes closer to the transmitter. We empirically determined that only 25% of the receive channels closest to the transmitter need to be excluded in order to avoid this inaccuracy. Additionally, we apply a single-sided Gaussian window to each time trace in the received channel data to exclude events that appear in the data prior to the transmitted wave (e.g., multiplexer switching, precursor signals, parasitic waveguiding around ring, etc.). Figure 3 shows the entire process with channel data from an *in-vitro* phantom. After excluding the 25% of receive channels (receivers 1-128 and 897-1024) closest to the transmitter (transmit 1), Figure 3(b) eliminates the multiplexer artifact seen in the first 30 μs of Figure 3(a) by applying the single-sided Gaussian window.

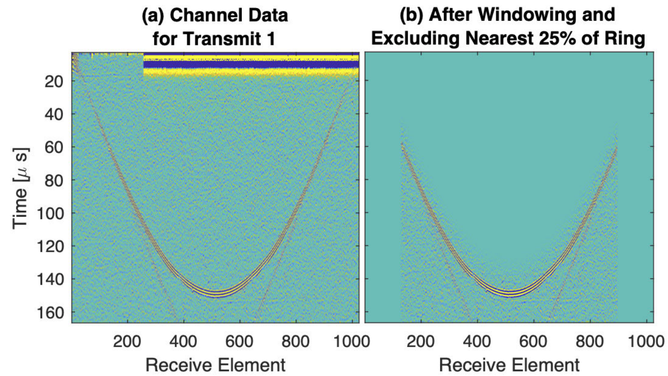

Fig. 3. Time-domain pre-processing to isolate through transmission. (a) Original *in-vitro* phantom channel data for pulse emitted from element 1. (b) Channel data after excluding of the 25% of receive elements closest to the transmitter and applying single-sided Gaussian windowing to remove artifacts and noise prior to the main arrival of the pulse.

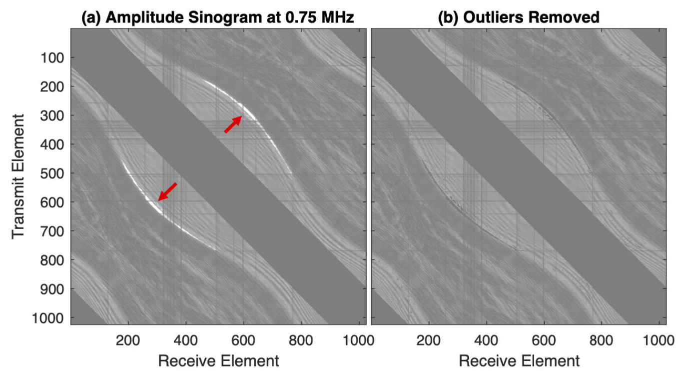

Fig. 4. Example of outlier removal with *In-Vivo* breast data. (a) 1024 × 1024 sinogram of phasor amplitudes, excluding the 25% of receivers closes to the transmitter (e.g., the blank region along the diagonal of the sinogram). Red arrows indicate the high-amplitude outliers caused by the glancing incidence of the ultrasound wave along the skin line surrounding the breast. (b) Sinogram after removing the top 1% of phasor measurements with the largest amplitudes.

E. Additional Pre-Processing Steps

After windowing the received channel data for each single-element transmission (see section III-D), we take a discrete-time Fourier transform (DTFT) to isolate the specific frequencies used in the waveform inversion. This gives us either a 512 × 512 (k-Wave simulation) or a 1024 × 1024 (experimental data) matrix of phasors at each frequency, and after removing everything except the opposing 75% of the ring for each transmitter, these become 512 × 385 and 1024 × 769 matrices, respectively. Afterwards, we exclude certain outliers in our remaining measurements that would cause problems for waveform inversion. We specifically remove the top 1% of phasor measurements having the largest magnitude at each frequency. These large values can be the most problematic because they dominate the gradient (6) image created when backprojecting the error $(\mathbf{p_i} - \mathbf{p_{obs,i}})$. Figure 4 shows an example of this outlier removal process with *in-vivo* breast data. The 1024 × 1024 sinogram in Figure 4(a) shows the amplitude of each phasor, with abnormally large values corresponding to a strip of skin along the surface of the breast. For certain glancing incidence angles, the skin around the breast

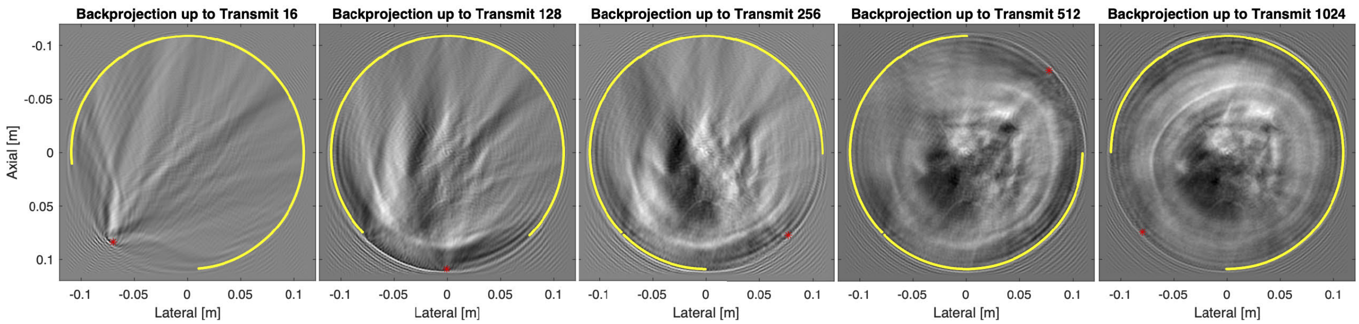

Fig. 5. Visual demonstration of backprojection using breast data at 0.3 MHz. The backprojection of the error in the received signal is accumulated over the transmitters to produce the final gradient image. The transmitter is indicated by a red asterisk, and active receivers are indicated in yellow.

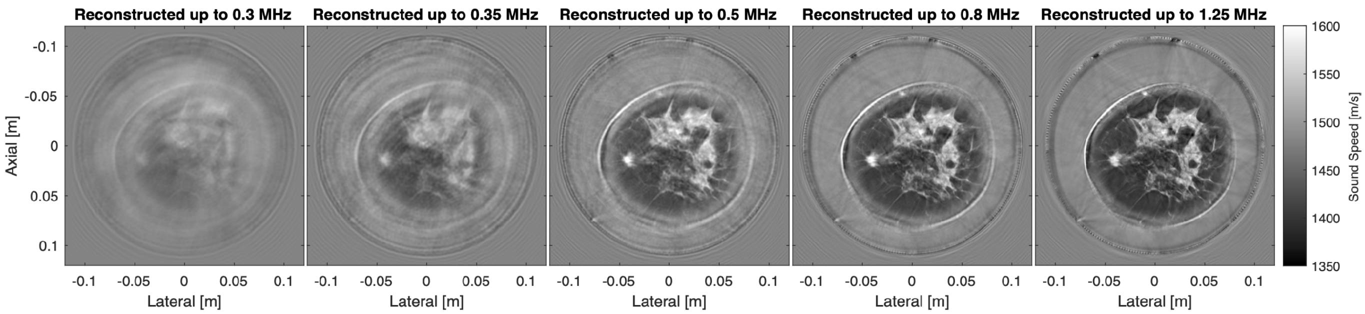

Fig. 6. Bootstrapping in frequency. Sound speed reconstructions are shown using data at frequencies up to 0.3, 0.35, 0.5, 0.8, and 1.25 MHz, respectively. The low frequencies are used to overcome cycle skipping, and as the frequency increases, image resolution improves. Attenuation is assumed to be zero throughout the medium in these reconstructions.

can act as waveguide due to its high sound speed. The wave, amplified by this waveguide, caused large signals at 0.75 MHz. Figure 4(b) shows that the 1% outlier removal eliminates most of these large signals. Currently, these outliers are simply set to zero; however, the waveform inversion algorithm may be modified such that the data residuals corresponding to the outliers are excluded from the backprojection in the gradient calculation.

IV. Results

A. Demonstration of Image Reconstruction Process

Using phasor measurements at 0.3 MHz from a breast dataset, Figure 5 visually demonstrates the backprojection calculation (6) when starting from a homogeneous 1480 m/s starting model. The error in the receive signal is backprojected into the estimate and accumulated over the transmitters, resulting in the final gradient image. This gradient image is then used to update the estimate of the sound speed in the medium. The image reconstruction begins at the lowest possible frequency. However, by iteratively applying this backprojection approach at progressively higher frequencies, image resolution improves as shown in Figure 6. Note that Figure 6 only shows the end result of the first pass through all the frequencies. The second pass further refines the estimated sound speed by incorporating attenuation. This will be shown later in section IV-D.

B. In-Silico Validation

The waveform inversion algorithm is first validated using simulated datasets where the true sound speed and attenuation profiles are known precisely. Figure 7 shows the reconstructed sound speed and attenuation profiles in each of the two simulation cases (a) and (b). When compared to the ground truth shown in Figure 1, the RMS error in sound speed is 11.0 and 8.3 m/s for cases (a) and (b), respectively. Similarly, the RMS error in the attenuation estimate is 0.36 and 0.46 dB/(cm MHz), respectively. Both sound speed and attenuation are reconstructed accurately in simulated data. Sections IV-C and IV-D show the simultaneous reconstruction of sound speed and attenuation in real experimental conditions.

C. In-Vitro Experiments

Figure 8 demonstrates the waveform inversion reconstructions of sound speed and attenuation on datasets collected from the *in-vitro* phantom. Figure 8(a)-(b) correspond to two different slices of the phantom. The inclusions in the images are numbered 1-8. Table I reports the sound speed and attenuation averaged over each inclusion, the background of the breast phantom, and the water surrounding the breast phantom. Lastly, Table I also provides manufacturer-reported sound speed and attenuation ranges for each type of material. Figure 8(c)-(d) plots the reconstructed sound speed and attenuation, respectively, as star-shaped points against their range of expected values, represented by the boxes.

The sound speed measured in inclusions 1, 3, 4, and 7 falls at the lower end expected in the high-sound speed lesions. However, the sound speed in inclusions 2 and the background material of the phantom falls the higher end of their expected ranges. According to the manufacturer, the chemical agent responsible for the high sound speed in the lesions is expected

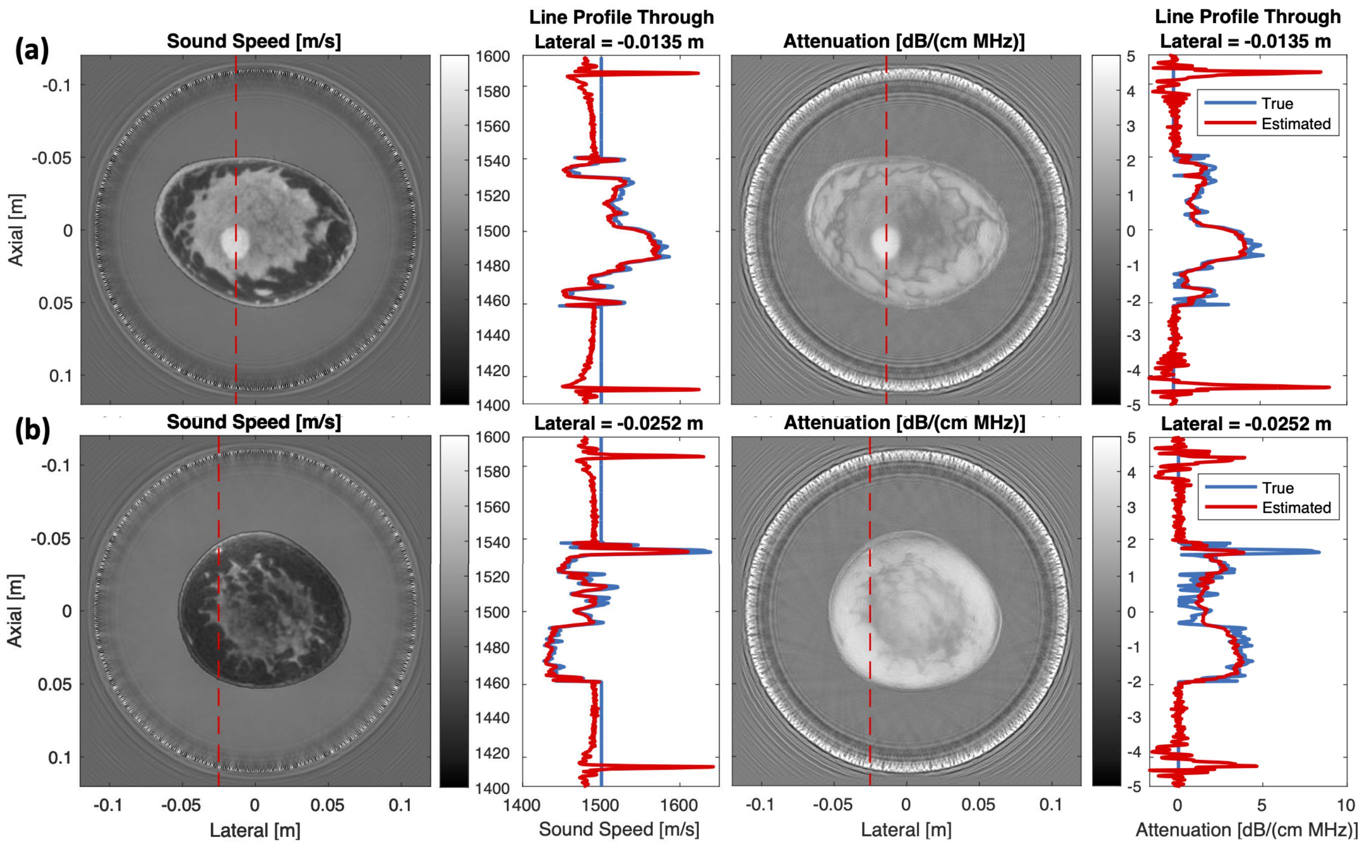

Fig. 7. Validation of waveform inversion algorithm using k-wave simulated datasets. The RMS error in the reconstructed sound speed and attenuation profiles are: (a) 11.0 m/s and 0.36 dB/(cm MHz), respectively; and (b) 8.3 m/s and 0.46 dB/(cm MHz), respectively. Line profiles at (a) Lateral = −0.0135 m and (b) Lateral = −0.0252 m comparing the true (blue) and estimated (red) sound speed and attenuation are shown.

to leak into the surrounding background as water diffuses into the lesions. Thus, the sound speed in the phantom approaches a homogeneous mean value as chemical agents within the phantom diffuse. An additional confounding factor is the unknown difference in temperature between manufacturer-reported values and our experiments. In the Discussion section, we address partial volume effects that impact both the sound speed and attenuation reconstructions in the phantom.

D. In-Vivo Demonstration

Figure 9 shows high-resolution waveform inversion reconstructions of sound speed and attenuation in two breast imaging examples: (a) an example of a benign cyst; and (b) an example of a spiculated malignancy. Much like the malignancy (1540-1600 m/s) in Figure 9(b), the fluid-filled cyst (Figure 9(a)) is characterized by a high sound speed (1520-1550 m/s) relative to the fatty background (1380-1420 m/s). However, the key differentiating factor between the benign cyst and the malignancy is the attenuation. The malignancy is characterized a high attenuation relative to the background tissue whereas the attenuation inside the cyst is close to the background breast fat. Similar to the *in-vitro* experiments in Figure 8(b), the large negative attenuation values seen in Figure 9(b) are most likely due to focusing effects that occur out of the imaging plane. The main distinction between the final result of Figure 6 and the sound speed image shown in Figure 8(b) is the simultaneous reconstruction of an attenuation map that corrects for deviations in the amplitude of the waveform.

V. Discussion

This is the first work to show high-resolution reconstructions of both sound speed and attenuation in the breast using an open-source frequency-domain waveform inversion algorithm with a ring-array transducer. We first present the mathematical theory behind frequency-domain waveform inversion, the practical aspects of the application of waveform inversion to real data (e.g., source estimation and gridding corrections), and the algorithmic details behind the Helmholtz equation solver that lies at the heart of frequency-domain waveform inversion. Some of these details have important implications for the results presented thus far. For example, the source estimation step estimates the phase and amplitude for each source by finding the complex scaling factor that maximizes the agreement between the simulated and measured channel data. However, this introduces two different non-uniqueness problems related to the phase and amplitude respectively.

The nonuniqueness in the phase affects the sound speed reconstruction. If a phase is applied to the simulated data at each frequency to minimize the discrepancy with the measured data, it is unclear whether this phase corresponds to the source or an error in the sound speed model. Fortunately, *in-silico* results shows that waveform inversion can robustly overcome

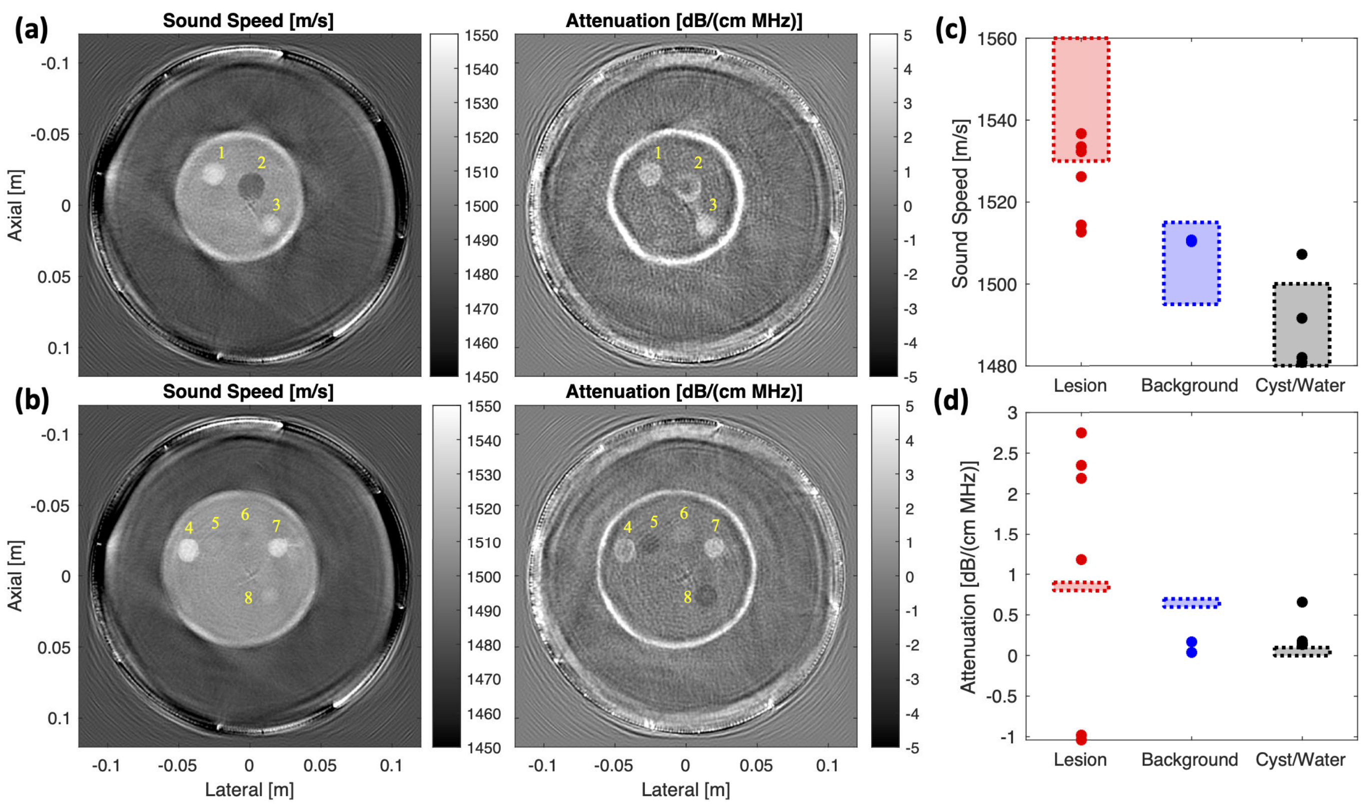

Fig. 8. Waveform inversion reconstruction of sound speed and attenuation *In-Vitro*. Two separate slices (a) and (b) from the same phantom were imaged. The reconstructed sound speed and attenuation values for each inclusion (numbered 1-8), the background, and the water is reported in Table I based on values averaged over the corresponding regions of each image. The expected sound speed and attenuation for each material type, as per the manufacturer, is also provided in Table I. Reconstructed sound speed (c) and attenuation (d) values were grouped according to whether they corresponded to a lesion inclusion, water (or water-filled cyst), or the background material of the phantom. In (c) and (d), these values were plotted as points with bands corresponding to expected ranges for each type of material.

TABLE I
RECONSTRUCTED SOUND SPEED AND ATTENUATION IN PHANTOM SLICES SHOWN IN FIGURE 8

Phantom Slice (a)	Material Type	Reconstructed Sound Speed [m/s]	Expected Sound Speed [m/s]	Reconstructed Attenuation [dB/(cm MHz)]	Expected Attenuation [dB/(cm MHz)]
Inclusion 1	Lesion	1532.3	1530-1560	2.19	0.8-0.9
Inclusion 3		1526.2		2.75	
Inclusion 2	Cyst/Water	1491.6	1480-1500	0.14	0.0-0.1
Outside Phantom		1482.0		0.14	
Phantom Background	Background	1510.3	1495-1515	0.17	0.6-0.7
Phantom Slice (b)	**Material Type**	**Reconstructed Sound Speed [m/s]**	**Expected Sound Speed [m/s]**	**Reconstructed Attenuation [dB/(cm MHz)]**	**Expected Attenuation [dB/(cm MHz)]**
Inclusion 4	Lesion	1536.7	1530-1560	1.18	0.8-0.9
Inclusion 5		1512.7		-1.04	
Inclusion 7		1533.5		2.35	
Inclusion 8		1514.4		-0.98	
Inclusion 6	Cyst/Water	1507.2	1480-1500	0.66	0.0-0.1
Outside Phantom		1480.8		0.18	
Phantom Background	Background	1510.7	1495-1515	0.04	0.6-0.7

the ambiguity between the phasing of each source and the sound speed estimate. Additionally, the phase rotation that minimizes the discrepancy between simulated and measured data acts to counteract cycle skipping by bringing the simulated waveforms as close as possible to the measure waveform.

The nonuniqueness in amplitude primarily affects the attenuation reconstruction. If the exact amplitude of each source is not known in advance, it becomes unclear whether the error in signal amplitudes is due to an incorrectly estimated source or an unknown attenuation profile in the medium. For this reason, both positive and negative values are seen in the attenuation images. The source scaling ambiguity means that absolute attenuation estimates cannot be guaranteed. The attenuation reconstruction should be interpreted as a relative measurement

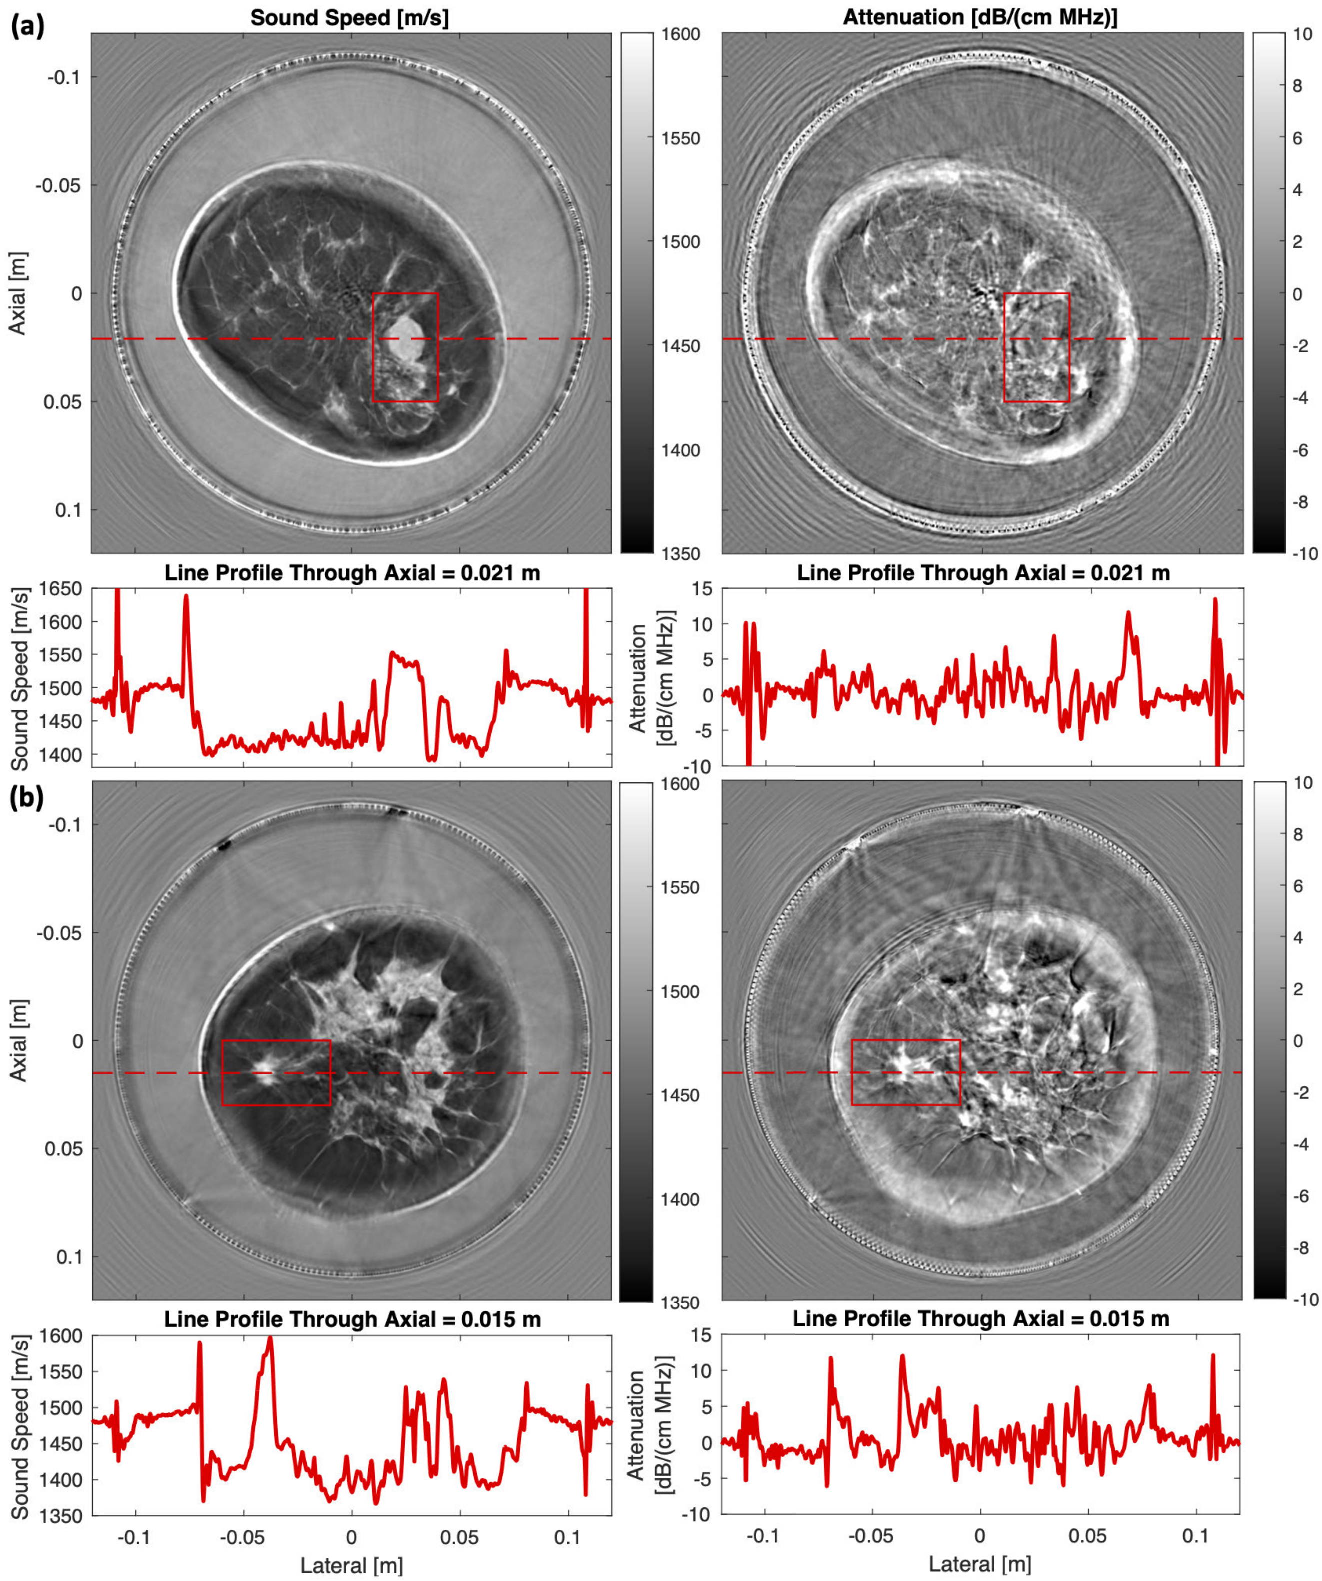

Fig. 9. Waveform inversion reconstructions of sound speed and attenuation in the breast. (a) Example of a breast with a benign cyst (surrounded by red box). Based on the line profile over the dotted line (at Axial = 0.021 m), the sound speed in the cyst (1520-1550 m/s) is higher than the surrounding breast fat (1400-1420 m/s); however, the attenuation in this cyst is similar to the attenuation in the surrounding breast fat. (b) Example of a breast with a cancerous spiculated mass (surrounded by red box). Based on the line profile over the dotted line (at Axial = 0.015 m), the sound speed in the spiculated mass (1540-1600 m/s) is much higher than the surrounding breast fat (1380-1420 m/s); similarly, the attenuation in the spiculated mass is much higher than the attenuation in the surrounding breast tissue.

based on a mean signal level that is not known a priori. Additionally, *in-vitro* experiments show that significant 3D partial-volume effects impact the attenuation reconstruction when applying 2D waveform inversion to experimental data.

These partial volume effects are caused by the spherical shape of the inclusions in the *in-vitro* phantom. The sound speed estimates of inclusions 5, 6, and 8 fall well outside their expected ranges and are much closer to the sound speed in the background material. This is most likely because only the glancing edges of inclusions 5, 6, and 8 fall within the imaging plane. Despite the low sound speed contrast of inclusions 5, 6, and 8 relative to the background material, they are well visualized in the attenuation image. It is interesting to note that the attenuation estimates are negative for inclusions 5 and 8. This implies the amplification of the ultrasound wave for paths going through these inclusions. We suspect that the wave focuses out of plane due to a waveguiding effect for the glancing angle of incidence with these high speed inclusions. Conversely, the wave diverges out of the plane when passing through inclusion 6, which is a glancing incidence with a low speed cyst, resulting in a positive attenuation estimate.

Across the remaining high sound speed lesions in the phantom (inclusions 1, 3, 4, and 7), attenuation is greatly overestimated; manufacturers report attenuation values of 0.8-0.9 dB/(cm MHz), while waveform inversion results in values ranging from 2-3 dB/(cm MHz) in most of the lesions. Thus, there is a tendency for the current 2D waveform inversion algorithm to overestimate attenuation values. This overestimation of attenuation may be the result of the geometric spreading of the wave outside the imaging plane, an effect that the current 2D algorithm cannot model. Lastly, inclusion 2 is a cyst with an average attenuation close to that expected in water. The attenuation is biased upward around the cyst due to the high impedance mismatch at the interface. Our current form of attenuation tomography is unable to distinguish between an impedance mismatch and viscous losses in the medium.

Another challenge for attenuation imaging is the choice of model [7]. Currently, we model attenuation as the imaginary part of slowness, resulting in a non-causal attenuation that is proportional to frequency. There are also frequency-domain power laws for attenuation that use Kramers-Kronig relationships to enforce causality. Because several attenuation models exist, it is unclear which model best generalizes to tissue.

Fortunately, the simultaneous imaging of both sound speed and attenuation provides radiologists with two different tissue properties to corroborate the identification of suspected malignancies. For example, the high sound speed cyst in Figure 9(a) could be mistaken for a cancerous lesion; however, the absence of spiculations and the low attenuation through the region would indicate that this is a fluid-filled cyst rather than a cancerous mass. On the other hand, the high sound speed and attenuation of the spiculated mass in Figure 9(b) are strongly indicative of cancer. The resolution of these images allows radiologists to visualize and compare structures that appear in both sets of images. For example, the spiculations that appear in both sound speed and attenuation images in Figure 9(b) are the strongest indicators that the mass is malignant [45]. One interesting avenue for future work is incorporating reflected signals into the reconstruction process. The current algorithm solely uses the through transmission; however, reflected signals could provide a secondary approach to image sound speed [10], [11], [12] or enable the imaging of mass density [6], [7]. Incorporating reflected signals could enable the imaging of 3 tissue properties: sound speed, mass density, and attenuation.

In addition to non-uniqueness issues, the dynamic nature of the source estimation used in this work also implicitly changes the gradient of the objective function with respect to slowness. Taking the source scaling (17) into account, the objective function (5) should actually be

$$E(\omega, \mathbf{s}) = \frac{1}{2}\sum_{i=1}^{N} \left\| \gamma_i(\omega; \mathbf{s})\mathbf{p_i}(\omega, \mathbf{s}) - \mathbf{p_{obs,i}}(\omega) \right\|_2^2 . \quad (32)$$

If the dependence of γ_i on slowness $\mathbf{s}$ is neglected, the gradient of this objective function would be

$$\nabla_{\mathbf{s}} E = \sum_{i=1}^{N} Re\left\{ \left(\gamma_i \frac{\partial \mathbf{p_i}}{\partial \mathbf{s}} \right)^H \left(\gamma_i \mathbf{p_i} - \mathbf{p_{obs,i}} \right) \right\}, \quad (33)$$

which is identical to our current gradient calculation with the source scalings. However, if the dependence of γ_i on slowness $\mathbf{s}$ were not neglected, the gradient formula would instead be

$$\nabla_{\mathbf{s}} E = \sum_{i=1}^{N} Re\left\{ \left(\gamma_i \frac{\partial \mathbf{p_i}}{\partial \mathbf{s}} + \frac{\partial \gamma_i}{\partial \mathbf{s}} \mathbf{p_i} \right)^H \left(\gamma_i \mathbf{p_i} - \mathbf{p_{obs,i}} \right) \right\}. \quad (34)$$

The $\frac{\partial \gamma_i}{\partial \mathbf{s}} \mathbf{p_i}$ term introduced by the source scaling complicates the backprojection in the gradient calculation. Future work on a similar source estimation technique should investigate the implementation and significance of this additional term.

The current forward-backward sweep approach to the Helmholtz equation also has interesting implications for future work on fast approximations to the Helmholtz equation, and hence, a faster waveform inversion algorithm. On an Intel i7-10700 system with an NVIDIA GeForce RTX 3060 GPU, the current per-slice reconstruction time is approximately 30 hours on CPU vs 80 minutes on the GPU. The reconstruction time could be reduced by limiting the highest frequency used in the reconstruction. This would not only reduce the number of iterations used to reconstruct the image but also allow for courser grids that reduce the computational cost at each frequency. Unfortunately, limiting the highest frequency used in the reconstruction also reduces image resolution. Previous work [38] notes the connection between the block LU decomposition and the analytic decomposition of the Helmholtz equation into forward and backward one-way wave equations. These one-way wave equations are the foundation of the paraxial approximation used in previous works [4], [31]. These paraxial solvers benefit from a Fourier split-step formulation that allows for rapid computation on the GPU. If an accurate Helmholtz equation solver could be designed based on forward and backward passes of a Fourier split-step method, there would be a major reduction in computation time and the paraxial approximation would find new applications outside of linear array UST. The most important application would be the extension of ring array UST to 3D imaging. The current block LU approach does not extend well to 3D because matrix storage requirements. However, the decomposition into one-way wave equations would result in a fast and much less memory-intensive approach to solve the Helmholtz equation. The paraxial approximation in [4], and [46] allows for an efficient 3D image reconstruction algorithm on the GPU; therefore, the full 3D Helmholtz equation could be solved in similar fashion using forward-backward sweeps of the one-way wave equation. A 3D reconstruction algorithm based on data from a multi-slice data acquisition or a multi-row ring array could overcome many of the partial volume and 3D out-of-plane effects encountered in our current 2D reconstruction algorithm.

Noise can significantly impact our reconstruction algorithm; however, the impact of noise on the image quality is often dependent on several circumstances surrounding it. For example, we typically assume an ideal Gaussian pulse excitation, especially in simulations. However, pulses produced by and recorded on physical systems may deviate significantly from this Gaussian shape. Channel data may have exceptionally lower signal at certain frequencies than one might expect based on a fractional bandwidth measurement that assumes an ideal Gaussian profile. Noise ultimately has a frequency-dependent effect on our waveform inversion algorithm because of the unique spectrum of the pulse. The windowing used to isolate the through transmission in the time domain signal can mitigate the impact of noise. Additionally, the large number of measurements (1024×1024) in our full matrix captures can further mitigate the impact of noise. Similar to coherent compounding in B-mode ultrasound imaging, accumulating the backprojections from several data residuals can overcome low signal-to-noise ratio (SNR) on individual channels. These are just a few of the ways noise can impact waveform inversion. For example, the impact of noise may be further complicated by the presence of cycle skipping.

Clinical implementations of UST [4], [29] have often preferred frequency-domain waveform inversion because time-domain waveform inversion often has greater storage requirements: whereas as frequency-domain waveform inversion only requires maintaining the periodic steady-state pressure field $\mathbf{u_i}$ for each source $\delta_\mathbf{i}$, time-domain waveform inversion requires the pressure field for each source as a function of time. In the case of 2D waveform inversion, frequency domain modeling often outperforms time-domain modeling because of its reduced storage requirements and the ease of solving the discretized system of equations produced by the 2D Helmholtz equation. However, the computations needed to solve the 3D discretized system often consumes much more memory than the storage requirement of time-domain model. As an alternative to the paraxial approximation [46], the WaveHoltz method [47] can solve the 3D Helmholtz equation by applying fixed-point iteration to the time-domain model. Without 3D modeling, it may still be possible to reconstruct 3D volumes based on our 2D-slice reconstruction approach by using regularization to enforce continuity across adjacent slices. The main drawback of this approach would be the inaccuracy of the 2D model at each slice.

VI. Conclusion

This is the first work to show high-resolution reconstructions of sound speed and attenuation in breast tissue using an open-source frequency-domain waveform inversion algorithm. This work explains many practical details related to the application of frequency-domain waveform inversion to ultrasound tomography with a ring-array transducer. The most computationally intensive aspect of the waveform inversion algorithm is the Helmholtz equation solver whose efficient implementation comes in the form of a forward-backward sweep through the computation grid based on the block LU decomposition. By using sufficiently low frequency content, the waveform inversions performed in this work do not require ray tomography to provide a starting model. All code and datasets used in this work have been made publicly available at https://github.com/rehmanali1994/WaveformInversionUST (DOI: https://zenodo.org/badge/latestdoi/684631232).

References

[1] C. Li, N. Duric, P. Littrup, and L. Huang, "In vivo breast sound-speed imaging with ultrasound tomography," *Ultrasound Med. Biol.*, vol. 35, no. 10, pp. 1615–1628, Oct. 2009.

[2] N. Duric et al., "Detection of breast cancer with ultrasound tomography: First results with the computed ultrasound risk evaluation (CURE) prototype," *Med. Phys.*, vol. 34, no. 2, pp. 773–785, 2007.

[3] J. S. Schreiman, J. J. Gisvold, J. F. Greenleaf, and R. C. Bahn, "Ultrasound transmission computed tomography of the breast," *Radiology*, vol. 150, no. 2, pp. 523–530, Feb. 1984.

[4] J. Wiskin, D. T. Borup, S. A. Johnson, and M. Berggren, "Nonlinear inverse scattering: High resolution quantitative breast tissue tomography," *J. Acoust. Soc. Amer.*, vol. 131, no. 5, pp. 3802–3813, May 2012.

[5] G. F. Pinton, G. E. Trahey, and J. J. Dahl, "Sources of image degradation in fundamental and harmonic ultrasound imaging using nonlinear, full-wave simulations," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 58, no. 4, pp. 754–765, Apr. 2011.

[6] P. Mojabi and J. LoVetri, "Ultrasound tomography for simultaneous reconstruction of acoustic density, attenuation, and compressibility profiles," *J. Acoust. Soc. Amer.*, vol. 137, no. 4, pp. 1813–1825, Apr. 2015.

[7] F. Faucher and O. Scherzer, "Quantitative inverse problem in visco-acoustic media under attenuation model uncertainty," *J. Comput. Phys.*, vol. 472, Jan. 2023, Art. no. 111685.

[8] C. Li, N. Duric, and L. Huang, "Comparison of ultrasound attenuation tomography methods for breast imaging," in *Medical Imaging: Ultrasonic Imaging and Signal Processing*, vol. 6920. Bellingham, WA, USA: SPIE, 2008, pp. 338–346.

[9] M. Pérez-Liva, J. L. Herraiz, J. M. Udías, E. Miller, B. T. Cox, and B. E. Treeby, "Time domain reconstruction of sound speed and attenuation in ultrasound computed tomography using full wave inversion," *J. Acoust. Soc. Amer.*, vol. 141, no. 3, pp. 1595–1604, Mar. 2017.

[10] R. Rau, D. Schweizer, V. Vishnevskiy, and O. Goksel, "Speed-of-sound imaging using diverging waves," *Int. J. Comput. Assist. Radiol. Surg.*, vol. 16, no. 7, pp. 1201–1211, Jul. 2021.

[11] R. Ali et al., "Sound speed estimation for distributed aberration correction in laterally varying media," *IEEE Trans. Comput. Imag.*, vol. 9, pp. 367–382, 2023. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10083328

[12] N. K. Martiartu, P. S. Yolgunlu, P. Stähli, M. Frenz, and M. Jaeger, "Towards attenuation imaging with computed ultrasound tomography in echo mode (CUTE)," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, Oct. 2022, pp. 1–4.

[13] A. S. Podkowa, "Spatiotemporal excitation sequence manipulation and processing for Doppler ultrasound and pulse-echo sound speed imaging applications," Ph.D. dissertation, Dept. Elect. Comput. Eng., Univ. Illinois Urbana-Champaign, Urbana-Champaign, IL, USA, 2020.

[14] R. Ali et al., "Local sound speed estimation for pulse-echo ultrasound in layered media," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 69, no. 2, pp. 500–511, Feb. 2022.

[15] A. V. Telichko et al., "Noninvasive estimation of local speed of sound by pulse-echo ultrasound in a rat model of nonalcoholic fatty liver," *Phys. Med. Biol.*, vol. 67, no. 1, Jan. 2022, Art. no. 015007.

[16] P. Stähli, M. Frenz, and M. Jaeger, "Bayesian approach for a robust speed-of-sound reconstruction using pulse-echo ultrasound," *IEEE Trans. Med. Imag.*, vol. 40, no. 2, pp. 457–467, Feb. 2021.

[17] J. F. Greenleaf, S. K. Kenue, B. Rajagopalan, R. C. Bahn, and S. A. Johnson, "Breast imaging by ultrasonic computer-assisted tomography," in *Acoustical Imaging: Ultrasonic Visualization and Characterization*. Cham, Switzerland: Springer, 1980, pp. 599–614.

[18] A. Hormati, I. Jovanović, O. Roy, and M. Vetterli, "Robust ultrasound travel-time tomography using the bent ray model," in *Medical Imaging: Ultrasonic Imaging, Tomography, and Therapy*, vol. 7629. Bellingham, WA, USA: SPIE, 2010, pp. 165–176.

[19] C. Li, L. Huang, N. Duric, H. Zhang, and C. Rowe, "An improved automatic time-of-flight picker for medical ultrasound tomography," *Ultrasonics*, vol. 49, no. 1, pp. 61–72, Jan. 2009.

[20] C. Li and N. Duric, "Resolution limitation of travel time tomography: Beyond the first Fresnel zone," in *Medical Imaging: Ultrasonic Imaging, Tomography, and Therapy*, vol. 8675. Bellingham, WA, USA: SPIE, 2013, pp. 359–368.

[21] A. Javaherian, F. Lucka, and B. T. Cox, "Refraction-corrected ray-based inversion for three-dimensional ultrasound tomography of the breast," *Inverse Problems*, vol. 36, no. 12, Dec. 2020, Art. no. 125010.

[22] M. Perez-Liva et al., "Speed of sound ultrasound transmission tomography image reconstruction based on Bézier curves," *Ultrasonics*, vol. 103, Apr. 2020, Art. no. 106097.

[23] T. Hopp, F. Zuch, P. Comby, and N. Ruiter, "Fat ray ultrasound transmission tomography: Preliminary experimental results with simulated data," in *Medical Imaging: Ultrasonic Imaging and Tomography*, vol. 11319. Bellingham, WA, USA: SPIE, 2020, pp. 125–133.

[24] X. Fang et al., "Study on acceleration schemes in Fresnel volume tomography for sound speed reconstruction," in *Medical Imaging: Ultrasonic Imaging and Tomography*, vol. 10955. Bellingham, WA, USA: SPIE, 2019, pp. 146–151.

[25] F. Lucka, M. Pérez-Liva, B. E. Treeby, and B. T. Cox, "High resolution 3D ultrasonic breast imaging by time-domain full waveform inversion," *Inverse Problems*, vol. 38, no. 2, Feb. 2022, Art. no. 025008.

[26] K. Wang, T. Matthews, F. Anis, C. Li, N. Duric, and M. A. Anastasio, "Waveform inversion with source encoding for breast sound speed reconstruction in ultrasound computed tomography," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 62, no. 3, pp. 475–493, Mar. 2015.

[27] L. Guasch, O. C. Agudo, M.-X. Tang, P. Nachev, and M. Warner, "Full-waveform inversion imaging of the human brain," *npj Digit. Med.*, vol. 3, no. 1, p. 28, Mar. 2020.

[28] C. Boehm, L. Krischer, I. Ulrich, P. Marty, M. Afanasiev, and A. Fichtner, "Using optimal transport to mitigate cycle-skipping in ultrasound computed tomography," in *Medical Imaging: Ultrasonic Imaging and Tomography*, vol. 12038. Bellingham, WA, USA: SPIE, 2022, pp. 48–57.

[29] G. Y. Sandhu, C. Li, O. Roy, S. Schmidt, and N. Duric, "Frequency domain ultrasound waveform tomography: Breast imaging using a ring transducer," *Phys. Med. Biol.*, vol. 60, no. 14, pp. 5381–5398, Jul. 2015.

[30] U. Taskin, A. Suzuki, T. Terada, Y. Tsubota, and K. W. Van Dongen, "Contrast source inversion on experimental data: Initial results," in *Medical Imaging: Ultrasonic Imaging and Tomography*, vol. 12038. Bellingham, WA, USA: SPIE, 2022, pp. 215–221.

[31] R. Ali, "Open-source full-waveform ultrasound computed tomography based on the angular spectrum method using linear arrays," in *Medical Imaging: Ultrasonic Imaging and Tomography*, vol. 12038. Bellingham, WA, USA: SPIE, 2022, pp. 187–205.

[32] Z. Chen, T. Wu, and H. Yang, "An optimal 25-point finite difference scheme for the Helmholtz equation with PML," *J. Comput. Appl. Math.*, vol. 236, no. 6, pp. 1240–1258, Oct. 2011.

[33] M. Patra and M. Karttunen, "Stencils with isotropic discretization error for differential operators," *Numer. Methods Partial Differ. Equ.*, vol. 22, no. 4, pp. 936–953, 2006.

[34] S. G. Johnson, "Notes on perfectly matched layers (PMLs)," 2021, *arXiv:2108.05348*.

[35] J. Nocedal and S. J. Wright, *Numerical Optimization*. Cham, Switzerland: Springer, 1999.

[36] G. J. Hicks, "Arbitrary source and receiver positioning in finite-difference schemes using Kaiser windowed sinc functions," *Geophysics*, vol. 67, no. 1, pp. 156–166, 2002.

[37] E. F. M. Koene, J. O. A. Robertsson, and F. Andersson, "A consistent implementation of point sources on finite-difference grids," *Geophys. J. Int.*, vol. 223, no. 1, pp. 1144–1161, May 2020.

[38] M. J. Gander and H. Zhang, "A class of iterative solvers for the Helmholtz equation: Factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized schwarz methods," *SIAM Rev.*, vol. 61, no. 1, pp. 3–76, Jan. 2019.

[39] B. Tuomanen, *Hands-On GPU Programming With Python and CUDA: Explore High-performance Parallel Computing With CUDA*. Birmingham, U.K.: Packt Publishing Ltd, 2018.

[40] B. E. Treeby and B. T. Cox, "K-wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields," *J. Biomed. Opt.*, vol. 15, no. 2, 2010, Art. no. 021314.

[41] J. Uhlig et al., "Contrast enhancement on cone-beam breast-CT for discrimination of breast cancer immunohistochemical subtypes," *Translational Oncol.*, vol. 10, no. 6, pp. 904–910, Dec. 2017.

[42] A. Saha et al., "A machine learning approach to radiogenomics of breast cancer: A study of 922 subjects and 529 DCE-MRI features," *Brit. J. Cancer*, vol. 119, no. 4, pp. 508–516, 2018.

[43] N. Duric, C. Li, O. Roy, and S. Schmidt, "Acoustic tomography: Promise versus reality," in *Proc. IEEE Int. Ultrason. Symp.*, Oct. 2011, pp. 2033–2041.

[44] R. Ali, T. Mitcham, and N. Duric, "Impact of starting model on waveform inversion in ultrasound tomography," in *Medical Imaging: Ultrasonic Imaging and Tomography*, vol. 12470. Bellingham, WA, USA: SPIE, 2023, pp. 89–98.

[45] C. J. Vyborny, T. Doi, K. F. O'Shaughnessy, H. M. Romsdahl, A. C. Schneider, and A. A. Stein, "Breast cancer: Importance of spiculation in computer-aided detection," *Radiology*, vol. 215, no. 3, pp. 703–707, Jun. 2000.

[46] J. W. Wiskin, D. T. Borup, E. Iuanow, J. Klock, and M. W. Lenox, "3-D nonlinear acoustic inverse scattering: Algorithm and quantitative results," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 64, no. 8, pp. 1161–1174, Aug. 2017.

[47] D. Appelö, F. Garcia, and O. Runborg, "WaveHoltz: Iterative solution of the Helmholtz equation via the wave equation," *SIAM J. Sci. Comput.*, vol. 42, no. 4, pp. A1950–A1983, Jan. 2020.