SIAM J. Imaging Sciences
Vol. 11, No. 2, pp. 1560–1588

Parameterized Joint Reconstruction of the Initial Pressure and Sound Speed Distributions for Photoacoustic Computed Tomography*

Thomas P. Matthews†, Joemini Poudel†, Lei Li‡, Lihong V. Wang‡, and Mark A. Anastasio†

Abstract. Accurate estimation of the initial pressure distribution in photoacoustic computed tomography (PACT) depends on knowledge of the sound speed distribution. However, the sound speed distribution is typically unknown. Further, the initial pressure and sound speed distributions cannot both, in general, be stably recovered from PACT measurements alone. In this work, a joint reconstruction (JR) method for the initial pressure distribution and a low-dimensional parameterized model of the sound speed distribution is proposed. By employing a priori information about the structure of the sound speed distribution, both the initial pressure and sound speed can be accurately recovered. The JR problem is solved by use of a proximal optimization method that allows constraints and nonsmooth regularization functions for the initial pressure distribution. The gradients of the cost function with respect to the initial pressure and sound speed distributions are calculated by use of an adjoint state method that has the same per-iteration computational cost as calculating the gradient with respect to the initial pressure distribution alone. This approach is evaluated through two-dimensional computer-simulation studies for a small animal imaging model and by application to experimental in vivo measurements of a mouse.

Key words. image reconstruction, joint image reconstruction, photoacoustic tomography

AMS subject classifications. 92C55, 78A70, 35R30, 65K10

DOI. 10.1137/17M1153649

1. Introduction.

Photoacoustic computed tomography (PACT) is a hybrid imaging modality that combines optical absorption contrast with acoustic detection [58, 40, 57, 56]. PACT can provide both structural and functional information and has been employed for human and small animal imaging applications [30, 20, 63, 16, 31, 61, 17]. PACT has a number of advantages over other functional imaging modalities: compared with MRI, it offers faster image acquisition at lower costs; compared with PET and SPECT, it does not involve use of ionizing radiation; and compared with other optical imaging modalities, it has the ability to obtain high-resolution images at large penetration depths [59].

In PACT, typically, a short laser pulse is employed to illuminate the object. The laser light is absorbed by molecules within the object, giving rise to an initial pressure distribution via

*Received by the editors October 24, 2017; accepted for publication (in revised form) February 5, 2018; published electronically June 5, 2018.

http://www.siam.org/journals/siims/11-2/M115364.html

Funding: This work was funded in part by NSF award DMS1614305 and NIH award R01EB01696304. Computations were performed using the facilities of the Washington University Center for High Performance Computing, which were partially provided through grant NCRR 1S10RR022984-01A1.

†Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, MO 63130 (thomas.matthews@wustl.edu, jpoudel@wustl.edu, anastasio@wustl.edu).

‡Caltech Optical Imaging Laboratory, Andrew and Peggy Cherng Department of Medical Engineering, Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 (leili@caltech.edu, LVW@caltech.edu).

the photoacoustic effect. This pressure distribution then propagates according to the acoustic properties of the medium, and the resulting pressure waves are measured by a collection of ultrasonic transducers surrounding the object.

Traditional image reconstruction methods for PACT assume that the medium is acoustically homogeneous. However, this assumption is often violated for biomedical applications [64, 23, 43]. To mitigate artifacts induced by this assumption, half-time and partial-time image reconstruction methods, which seek to exploit redundant information in data and eliminate parts of the data that are disproportionately affected by inhomogeneities, have been proposed [3, 2, 43]. Additionally, some image reconstruction methods have the ability to compensate for sound speed variations if the sound speed distribution is known [22, 52, 24, 4, 44, 46, 31]. In practice, however, the sound speed distribution is not known. Adjunct imaging data, such as ultrasound tomography measurements, could be acquired to allow estimation of the sound speed distribution [27, 62, 34, 35]. While ultrasound tomography can be employed to accurately estimate the sound speed distribution [15, 60, 10, 19, 21], not all PACT imaging systems are capable of acquiring the necessary measurements. Thus, there is a need to estimate some approximation of the sound speed distribution from PACT measurements alone.

Various methods have been proposed to accomplish this. For image reconstruction methods assuming a constant sound speed, this sound speed can be tuned according to some image quality metric. For example, Treeby et al. describe a method by which this sound speed can be chosen automatically by maximizing a measure of image sharpness [51]. Such approaches can be difficult to apply when the assumed sound speed distribution is more complex and many sound speed values need to be estimated.

Several previous works have addressed joint estimation of the sound speed and initial pressure distributions [47, 25, 32, 29, 39, 65, 66] and the related problem of joint estimation of the sound speed and optical absorption coefficient distributions [13]. In particular, Stefanov and Uhlmann considered the joint estimation problem for the linearized wave equation and demonstrated that, in general, the sound speed and initial pressure distributions could not both be stably recovered [47]. Similar observations were made by Huang et al. via computer-simulation studies for the full acoustic wave equation [25].

In [66], Zhang et al. proposed use of a low-dimensional representation of the sound speed distribution in order to stabilize the joint reconstruction (JR) problem. The sound speed and initial pressure distributions were estimated by minimizing an objective function consisting of a data fidelity term, which describes the distance between the measured pressure and the estimated pressure given some model, and a pair of regularization terms. The objective function was minimized by use of a gradient-based alternating minimization approach. A generalized Radon transform model, which assumes that the acoustic heterogeneities are relatively weak, was employed for estimating the measured pressure. In addition, the gradient of the objective function with respect to the parameterized sound speed distribution was calculated using a finite-difference-based approach. Under this approach, the computational cost of calculating the gradient scaled linearly with the number of parameters in sound speed model.

In this work, the idea of employing a parameterized sound speed model to stabilize the JR problem is revisited and extended in several key ways. First, a full-wave acoustic model, which does not assume that the acoustic heterogeneities are weak, is employed. This allows more accurate simulation of wave propagation through bone and other materials whose

acoustic properties differ greatly from the background medium. Second, an efficient method for computing the gradients with respect to the initial pressure and sound speed distributions, whose computational cost does not depend on the number of parameters in the sound speed model, is provided. Third, the optimization method is updated to allow use of nonsmooth regularization terms. Fourth, the approach is applied to experimental data.

The remainder of the manuscript is organized as follows. In section 2, background information on the imaging model and existing image reconstruction methods is provided. In section 3, the proposed parameterized JR method is introduced. A description of the computer-simulation studies is given in section 4 and the results of these studies are shown in section 5. In section 6, a description of the experimental studies is provided. Results of these studies are shown in section 7. Finally, the results are summarized in section 8.

2. Background.

2.1. Imaging physics.

In PACT, an object is illuminated with a short optical pulse. Some of this optical energy is absorbed and converted into an initial pressure distribution via the photoacoustic effect [40, 56, 59]. This initial pressure distribution then propagates in a way that depends on the acoustic properties of the medium. The resulting pressure wavefield is recorded by a collection of ultrasonic transducers surrounding the object.

Assuming a lossless medium, the photoacoustic wave propagation is described by the following coupled differential equations [49, 50]:

$$\rho(\mathbf{r})\frac{\partial \mathbf{u}(\mathbf{r},t)}{\partial t} + \nabla p(\mathbf{r},t) = \mathbf{0}, \tag{1a}$$

$$\frac{1}{\rho(\mathbf{r})\,c(\mathbf{r})^2}\frac{\partial p(\mathbf{r},t)}{\partial t} + \nabla\cdot\mathbf{u}(\mathbf{r},t) = 0, \tag{1b}$$

$$p(\mathbf{r},0) = p_0(\mathbf{r}), \tag{1c}$$

$$\mathbf{u}(\mathbf{r},0) = \mathbf{0}, \tag{1d}$$

where $\mathbf{u}$ is the acoustic particle velocity, p is the acoustic pressure, ρ is the mass density, c is the sound speed, and p_0 is the initial pressure distribution.

The pressure recorded by the ith ultrasonic transducer can be denoted as [55]

$$g_i(t) = \mathcal{M}_i p(\mathbf{r},t), \tag{2}$$

where $\mathcal{M}_i$ is an operator that gives the pressure measured by the ith transducer given the pressure over the whole domain. Here, $\mathcal{M}_i$ is chosen to have the form

$$\mathcal{M}_i \equiv h_e(t) *_t \Lambda_i, \tag{3}$$

where $*_t$ denotes the temporal convolution operator, $h_e(t)$ is the electro-acoustic impulse response (EIR), which accounts for the frequency response of the transducers, and Λ_i is the restriction of the pressure to the location of the ith transducer. For simplicity, the EIR is assumed to be the same for all transducers. The spatial impulse response of the transducers is not considered in this work.

2.2. Forward model in discrete form. In discrete form, the forward model can be denoted as

$$\mathbf{g} = \mathbf{M}\mathbf{H}(\mathbf{c})\,\mathbf{p}_0, \tag{4}$$

where $\mathbf{c} \in \mathbb{R}^N$ is a discrete representation of the sound speed distribution $c(\mathbf{r})$, $\mathbf{p}_0 \in \mathbb{R}^N$ is a discrete representation of the initial pressure distribution $p_0(\mathbf{r})$, $\mathbf{H}(\mathbf{c}) \in \mathbb{R}^{NL\times NL}$ is an operator that gives a discrete approximation of the solution to the acoustic wave equation given in (1), $\mathbf{M} \in \mathbb{R}^{ML\times NL}$ is an operator that gives the pressure recorded by the transducers from the pressure over the whole domain (i.e., it is a concatenation of the discrete approximations of $\mathcal{M}_i$ for all transducers), and $\mathbf{g} \in \mathbb{R}^{ML}$ is the measured data at the transducers. Here, N is the total number of pixels in the simulation grid, M is the number of transducers acting as receivers, and L is the number of time points employed as part of the wavefield simulation.

2.3. Reconstruction of the initial pressure distribution. An estimate of the initial pressure distribution for some fixed sound speed distribution $\mathbf{c}$ can be obtained by solving the optimization problem [54]

$$\widehat{\mathbf{p}_0} = \underset{\mathbf{p}_0\geq 0}{\arg\min}\, F(\mathbf{p}_0, \mathbf{c}) + \beta R(\mathbf{p}_0), \tag{5}$$

where $F(\mathbf{p}_0, \mathbf{c})$ is the data fidelity term, which describes a distance between the measured data and the predicted data given the initial pressure and sound speed distributions, $R(\mathbf{p}_0)$ is a regularization term, which incorporates a priori information about the sought-after initial pressure distribution, and β is a regularization parameter that controls the relative weight of the data fidelity and regularization terms. Here, the data fidelity is given by

$$F(\mathbf{p}_0, \mathbf{c}) \equiv \frac{1}{2}\|\underline{\mathbf{g}} - \mathbf{M}\mathbf{H}(\mathbf{c})\,\mathbf{p}_0\|_2^2, \tag{6}$$

where $\underline{\mathbf{g}}$ is the recorded pressure data, which may not be consistent with the assumed forward model. Further, the regularization function is chosen to be the total-variation (TV) seminorm [11],

$$R(\mathbf{p}_0) \equiv \|\nabla \mathbf{p}_0\|_1. \tag{7}$$

The optimization problem given in (5) is solved by use of the FISTA optimization method with a backtracking line search and an adaptive restart procedure [24, 7, 6, 38]. As described in [6], the nonnegativity constraint is incorporated when solving the proximal problem corresponding to the TV regularization term.

3. Parameterized joint reconstruction.

3.1. Sound speed parameterization. To constrain the JR problem, a low-dimensional parameterized representation of the sound speed distribution is considered. Let $\mathbf{c}_p \in \mathbb{R}^Q$ denote the parameterized model, whose elements may be sound speed values or other parameters useful for calculating the sound speed distribution over the whole domain. Further, let the parameterized model be related to the sound speed distribution over the whole domain as

$$\mathbf{c} = \Phi(\mathbf{c}_p), \tag{8}$$

where $\Phi(\cdot)$ is continuously differentiable. In the case where the mapping from $\mathbf{c}_p$ to $\mathbf{c}$ is linear, the relationship can be described as

$$\mathbf{c} = \mathbf{\Phi}\mathbf{c}_p, \tag{9}$$

where $\mathbf{\Phi} \in \mathbb{R}^{N\times Q}$. For the computer-simulation and experimental studies described in sections 4–7, each pixel in the simulation grid was assumed to correspond to one of Q tissue types with each tissue type having a uniform sound speed. In this case, $\mathbf{c}_p$ is simply the vector of sound speed values for the different tissue types and $\mathbf{\Phi}$ is given by

$$[\mathbf{\Phi}]_{i,j} \equiv \begin{cases} 1 & i \in \mathcal{I}_j, \\ 0 & \text{otherwise}, \end{cases} \tag{10}$$

where $\mathcal{I}_j$ denotes the set of indices within the simulation grid that correspond to the jth parameterized sound speed value and $[\mathbf{A}]_{i,j}$ denotes the (i,j) element of the matrix $\mathbf{A}$.

3.2. Joint reconstruction. Given some choice for Φ, the parameterized JR problem is given by

$$\widehat{\mathbf{p}_0}, \widehat{\mathbf{c}_p} = \underset{\mathbf{p}_0\geq 0, \mathbf{c}_p}{\arg\min}\, F(\mathbf{p}_0, \Phi(\mathbf{c}_p)) + \beta R(\mathbf{p}_0). \tag{11}$$

The gradient of the data fidelity term with respect to the parameterized sound speed distribution can be related to the gradient with respect to the sound speed over the whole domain via the chain rule as

$$\nabla_{\mathbf{c}_p} F(\mathbf{p}_0, \Phi(\mathbf{c}_p)) = \left(\frac{\partial \mathbf{c}}{\partial \mathbf{c}_p}\right)^T \nabla_{\mathbf{c}} F(\mathbf{p}_0, \Phi(\mathbf{c}_p)), \tag{12}$$

where

$$\left[\frac{\partial \mathbf{c}}{\partial \mathbf{c}_p}\right]_{i,j} = \frac{\partial [\mathbf{c}]_i}{\partial [\mathbf{c}_p]_j}. \tag{13}$$

When Φ corresponds to a real-valued linear mapping, the gradient with respect to the parameterized sound speed distribution is given by

$$\nabla_{\mathbf{c}_p} F(\mathbf{p}_0, \mathbf{\Phi}\mathbf{c}_p) = \mathbf{\Phi}^T \nabla_{\mathbf{c}} F(\mathbf{p}_0, \mathbf{\Phi}\mathbf{c}_p). \tag{14}$$

For the choice of $\mathbf{\Phi}$ employed in this work, this expression reduces to

$$\left[\nabla_{\mathbf{c}_p} F(\mathbf{p}_0, \mathbf{\Phi}\mathbf{c}_p)\right]_j = \sum_{i\in\mathcal{I}_j} \left[\nabla_{\mathbf{c}} F(\mathbf{p}_0, \mathbf{\Phi}\mathbf{c}_p)\right]_i. \tag{15}$$

Based on (15), the magnitude of each component of $\nabla_{\mathbf{c}_p} F$ will depend on the number of pixels that correspond to each component of $\mathbf{c}_p$. This can result in a gradient that is poorly scaled, leading to slow convergence. To combat this, a diagonal, positive-definite scaling matrix is introduced,

$$[\mathbf{B}]_j = \frac{1}{|\mathcal{I}_j|}, \tag{16}$$

resulting in a new search direction $\mathbf{B}\nabla_{\mathbf{c}_p} F$ that represents the average gradient over the pixels for each component of the parameterized sound speed distribution. Since $\mathbf{B}$ is positive-definite by construction, this new search direction is always a descent direction. Other choices for $\mathbf{B}$ could result in even faster convergence; however, the proposed choice requires little additional computational cost and ensures that the magnitude of each gradient component does not depend strongly on the number of pixels corresponding to each parameter.

With this, a weighted proximal gradient descent method, summarized in Algorithm 1, is employed to solve (11). The gradients with respect to the initial pressure and sound speed distributions are calculated via the adjoint state method as described in Appendices A and B. Under this approach, both gradients can be computed with only two wave solver runs. This is far fewer than required by a finite-difference-based approach and fewer than if the gradients were calculated independently. For brevity, the following notation is employed:

$$F^{(k)} \equiv F\left(\mathbf{p}_0^{(k)},\ \Phi\left(\mathbf{c}_p^{(k)}\right)\right), \tag{17a}$$

$$\mathbf{G}_p^{(k)} \equiv \nabla_{\mathbf{p}_0} F\left(\mathbf{p}_0^{(k)}, \Phi\left(\mathbf{c}_p^{(k)}\right)\right), \tag{17b}$$

$$\mathbf{G}_c^{(k)} \equiv \nabla_{\mathbf{c}_p} F\left(\mathbf{p}_0^{(k)}, \Phi\left(\mathbf{c}_p^{(k)}\right)\right), \tag{17c}$$

where the superscript k refers to the kth iteration of the algorithm. In addition, the notation $\mathrm{prox}_R(\cdot)$ is employed for the proximity operator of R. A two-parameter line search method, described in Appendix C, is employed to choose separate scalar step sizes for updating the initial pressure and sound speed distributions.

Algorithm 1 Parameterized joint reconstruction method.

Input: $\mathbf{c}_p^{(0)}, \mathbf{p}_0^{(0)}, \beta$
Output: $\widehat{\mathbf{c}_p}, \widehat{\mathbf{p}_0}$

$k \leftarrow 0$ {k is the algorithm iteration number.}
while stopping criterion is not satisfied **do**
  Calculate $\mathbf{G}_p^{(k)}$ and $\mathbf{G}_c^{(k)}$
  Choose α_k^p and α_k^c via Algorithm 2
  $\mathbf{p}_0^{(k+1)} \leftarrow \mathrm{prox}_{\alpha_k^p \beta R}\left(\mathbf{p}_0^{(k)} - \alpha_k^p \mathbf{G}_p^{(k)}\right)$
  $\mathbf{c}_p^{(k+1)} \leftarrow \mathbf{c}_p^{(k)} - \alpha_k^c \mathbf{B}\mathbf{G}_c^{(k)}$
  $k \leftarrow k + 1$
end while
$\widehat{\mathbf{c}_p} \leftarrow \mathbf{c}_p^{(k)}$
$\widehat{\mathbf{p}_0} \leftarrow \mathbf{p}_0^{(k)}$

4. Description of computer-simulation studies. Two-dimensional computer-simulation studies were performed in order to validate the proposed approach. Most studies were based on a piecewise constant numerical mouse phantom. An additional study was also performed with a vessel-like phantom in order to investigate the impact of high frequency content in the measured pressure on the inverse problem.

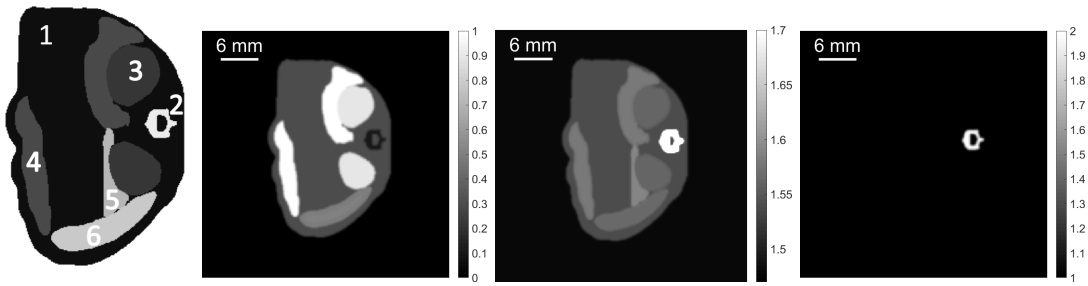

Figure 1. (a) *A schematic of the segmented tissue types within the mouse. The labels for each numbered tissue type are given in Table 1. Phantoms for* (b) *the normalized initial pressure distribution, given in arbitrary units,* (c) *the sound speed distribution, given in units of mm/μs, and* (d) *the mass density distribution, given in units of mg/mm*^{3}*. In order to better visualize the soft tissue variations for the sound speed distribution, the grayscale window for this phantom was set to* $[1.47, 1.7]$ *mm/μs resulting in saturation of the bone, which has a sound speed value of* 3.198 *mm/μs.*

Table 1

Parameter values for the initial pressure, sound speed, and mass density distributions for each tissue type. The normalized initial pressure values are roughly based on the relative concentration of blood for each tissue type [48, 33, 28, 18].

Index	Region	Initial pressure	Sound speed [mm/μs]	Mass density [mg/mm^3]
0	Water	0.0	1.480	1.00
1	Bulk tissue	0.3	1.540	1.00
2	Bone	0.1	3.198	1.99
3	Kidney	0.9	1.560	1.00
4	Liver	1.0	1.578	1.00
5	Pancreas	0.3	1.591	1.00
6	Spleen	0.5	1.567	1.00

4.1. Methods.

4.1.1. Numerical mouse phantom. The chosen numerical phantoms are based on segmented μCT images of a mouse produced by the DigiMouse project [14]. From one segmented slice, phantoms for the initial pressure and sound speed distributions were produced (see Figure 1). Each segmented region was assigned a constant initial pressure value and a constant sound speed value based on the tissue type. The initial pressure values were chosen based on the relative concentration of blood in each tissue type, as hemoglobin is one of the strongest endogenous optical absorbers in the body over the visible and near-infrared wavelength ranges [26]. The impact of optical fluence variations were not considered in this work. For most studies, the mass density was assumed to be homogeneous. These values are summarized in Table 1.

4.1.2. Imaging system. The imaging system was assumed to contain 512 ultrasonic transducers evenly distributed about a ring of radius 50 mm. The transducers were treated as

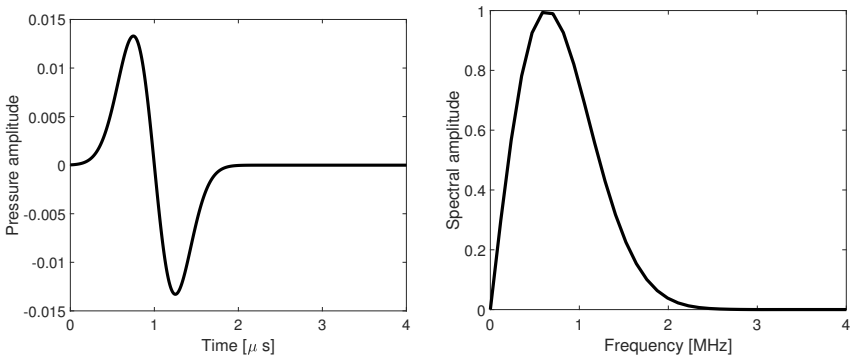

Figure 2. (a) *The pressure amplitude and* (b) *frequency spectrum of the EIR employed in the main computer-simulation studies.*

point-like detectors with the EIR given in Figure 2. The transducers were placed on the simulation grid by use of nearest neighbor interpolation.

4.1.3. Simulation of pressure data. In order to avoid inverse crime [12], the measured pressure data were simulated with different temporal and spatial sampling rates than employed for image reconstruction. When generating the measured data, the simulation grid consisted of a Cartesian grid with 1536×1536 pixels and a pixel size of 0.075 mm. The pressure data were recorded for 8000 time steps at a temporal sampling rate of 120 MHz. Additive Gaussian white noise with zero mean and a standard deviation of 2% of the maximum pressure amplitude was added to the measured data. The k-Wave toolbox was employed to simulate the measured pressure data [50].

4.1.4. Image reconstruction. During image reconstruction, the simulation grid consisted of 768×768 pixels with a pixel size of 0.15 mm. The pressure wavefield was simulated for 4000 time steps with a temporal sampling rate of 60 MHz. During the image reconstruction process, the pressure was simulated by use of a C-based numerical wave solver, implemented using NVIDIA's CUDA framework [1]. The gradients were estimated within a field-of-view of radius 20 mm.

4.2. Fixed constant sound speed. Initial pressure distributions were reconstructed by use of (5) for several fixed constant sound speed values and choices of β values. The stopping criterion was when the change in the ℓ_2-norm of the initial pressure distribution between successive iterations was less than 10^{-4}.

Most image reconstruction algorithms in PACT assume a constant speed. In practice, the value of this constant sound speed is tuned according to some criteria. The reconstructed initial pressure distribution images for several constant sound speeds were evaluated according to three such criteria: 1) the root-mean-square error (RMSE), 2) the value of the cost function given in (5), and 3) the Tenenbaum sharpness [51]. Use of the Tenenbaum sharpness for choosing a constant sound speed value was previously proposed by Treeby et al. [51]. Joint reconstruction with a 1-parameter model can be viewed as choosing the sound speed value

Table 2
Different sound speed parameterizations employed during image reconstruction.

Q	Regions
1	Constant sound speed
2	1) Background, 2) Mouse
3	1) Background, 2) Soft tissue, 3) Bone
7	1) Background, 2) Bulk tissue, 3) Bone, 4) Kidney, 5) Liver, 6) Pancreas, 7) Spleen
55869	All pixels within field-of-view

according to the value of the cost function. The agreement between these three measures is evaluated. The best criteria to employ to choose a constant sound speed value will depend on the task or tasks for which the image will be utilized [5]. Here, only a few simpler measures are considered.

4.3. Joint reconstruction. Several different sound speed parameterizations were considered in this work, each with a different number of parameters. These parameterizations are summarized in Table 2. For this initial investigation, the boundaries of these regions were based on the true segmented regions shown in Figure 1(a). When the assumed parameterization is too simple to describe the true sound speed variations within the object, model error will lead to errors in the reconstructed images. When the assumed parameterization is very complex, the inverse problem may be poorly conditioned and may have many local minima or saddle points. By considering a range of different parameterizations with different numbers of parameters, this trade-off between model error and the conditioning of the inverse problem is investigated.

The initial guess for initial pressure distribution was the vector of all zeros. The initial guess for the parameterized sound speed distribution depended on the number of parameters employed and the corresponding tissue types. The initial guesses for the different tissues were 1.48 mm/μs for the background, 1.50 mm/μs for all soft tissues, and 3.00 mm/μs for the bone. For fairness of comparison, the initial guess for the 55869-parameter model was chosen such that $\mathbf{\Phi c}_p^{(0)}$ was the same as for the 3-parameter and 7-parameter models, even though this model does not suggest any prior knowledge of the different tissue types present in the object.

The numerical wave solver can become inaccurate when high spatial frequencies are present in the sound speed distribution [49]. To avoid sharp discontinuities between the different parameterized regions, the sound speed distribution is smoothed with a Gaussian filter with a standard deviation of 1 pixel prior to running the numerical wave solver.

The stopping criterion was when the change in the ℓ_2-norm of the object between successive iterations was less than 10^{-4}. In this case, the object refers to the concatenation of the initial pressure and parameterized sound speed distributions.

4.4. Impact of mass density variations. In this work, mass density variations are ignored during the JR process. While the derived methods can naturally be extended to include mass density variations, for simplicity, the medium is assumed to have homogeneous mass density. However, the assumption of homogeneous mass density may be poor when bone or

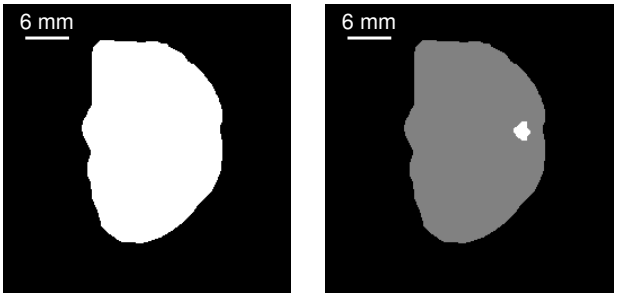

(a) (b)

Figure 3. *Estimated segmented regions for* (a) *a two-parameter sound speed model and* (b) *a three-parameter sound speed model.*

air voids are present in the object. To better understand the impact of ignoring mass density variations, the measured data were simulated with the heterogeneous mass density distribution shown in Figure 1(d). The mass density values for each of the tissue types are summarized in Table 1. During the image reconstruction process, a homogeneous mass density distribution was assumed.

4.5. Imperfect parameterization. To better understand the impact of errors in the assumed tissue type regions, approximate sound speed parameterizations were determined by segmenting a pair of reconstructed initial pressure distribution images. Two approximate parameterizations, shown in Figure 3, were employed. For the first, a two-parameter sound speed model was obtained by segmenting the outer boundary of the mouse from the reconstructed initial pressure image obtained by JR with a one-parameter sound speed model. The segmented region was determined by thresholding the initial pressure image with a cutoff of 0.1. For the second, a three-parameter sound speed model was employed. The outer boundary of the mouse was the same as for the first approximate parameterization. The outer boundary of the bone was estimated by manually segmenting a reconstructed initial pressure distribution image obtained by assuming a fixed constant sound speed value of 1.51 mm/μs. It was observed that the boundary of the bone could be more readily determined for this sound speed value, which was higher than that obtained by JR.

4.6. High-frequency pressure data. Due to the nature of the chosen initial pressure phantom, shown in Figure 1(b), the measured data consist of predominately low-frequency content. Accurate recovery of the sound speed distribution, and in turn accurate JR, can be challenging from high-frequency data due to the phenomenon of cycle skipping [53, 8]. Cycle skipping occurs when the cumulative phase error for a given frequency component exceeds π (or, equivalently, the traveltime error exceeds half of the period). In this case, the optimization method may align a given peak (or trough) in the simulated data with a neighboring peak in the measured data rather than with the true corresponding peak.

Several methods have been employed to address cycle skipping in ultrasound tomography and geophysics. When a frequency-domain wave solver is utilized, a common strategy is to use a frequency hopping method in which progressively higher frequencies are utilized to

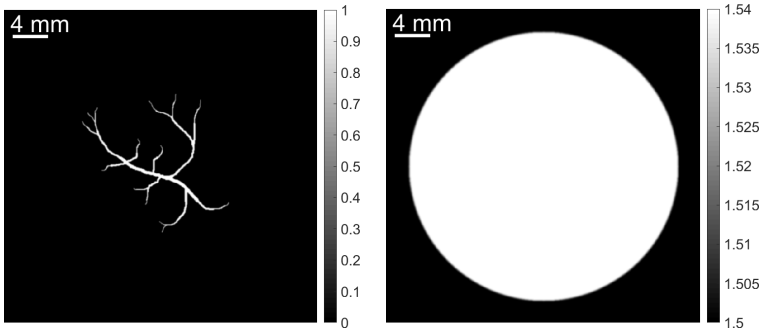

(a) (b)

Figure 4. *Phantoms for* (a) *the initial pressure distribution, given in arbitrary units, and* (b) *the sound speed distribution, given in units of mm/μs, for the high-frequency computer-simulation studies.*

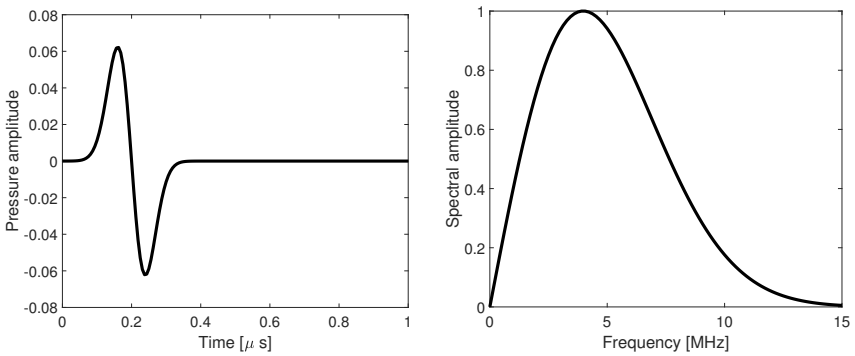

(a) (b)

Figure 5. (a) *The pressure amplitude and* (b) *frequency spectrum of the EIR employed in the high-frequency computer-simulation studies.*

estimate the sound speed distribution [45]. For time-domain wave solvers, a similar result can be achieved by low-pass filtering the measured data or through the use of multiscale methods [9, 8].

To investigate the utility of these approaches for parameterized JR, the topology of the data fidelity term was investigated for measured data that were filtered with several low-pass filters with different cutoff frequencies.

To obtain high-frequency measured data, a new vessel-like phantom, which contains much finer structures, is considered. This phantom was provided by the k-Wave toolbox [50]. The sound speed phantom consists of a circle with a radius of 15 mm and sound speeds of 1.50 mm/μs for the background and 1.54 mm/μs for the tissue-mimicking circle. Both phantoms are shown in Figure 4. In order to better capture the high-frequency content generated by the new phantom, a new EIR, shown in Figure 5, with a higher central frequency is employed. The measured pressure was simulated on a grid with 2048×2048 pixels and a pixel size of 0.05 mm. The pressure data were recorded for 8000 time points at a sampling rate of 160 MHz. Additive Gaussian white noise with zero mean and a standard deviation of 0.5% of

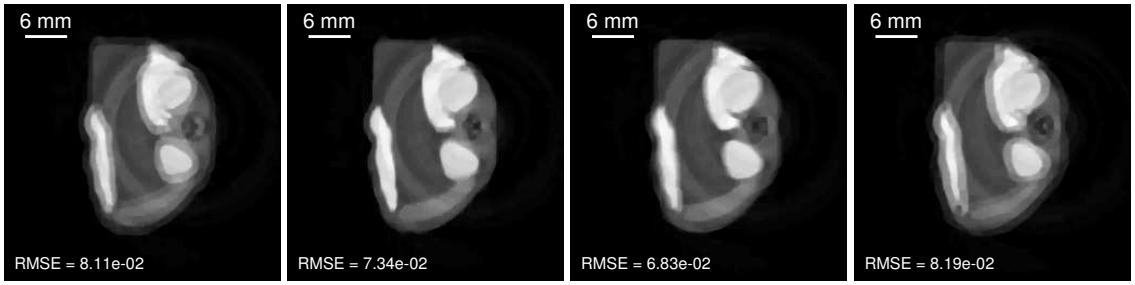

(a) (b) (c) (d)

Figure 6. *Reconstructed initial pressure distributions for* $\beta = 10^{-3}$ *and fixed constant sound speed values of* (a) 1.48 *mm/μs,* (b) 1.49 *mm/μs,* (c) 1.50 *mm/μs, and* (d) 1.51 *mm/μs. The images are shown in a grayscale window of* $[0.0, 1.1]$.

Table 3

Values of different metrics that could be employed for selecting a constant sound speed value. The optimal value according to each metric is shown in bold.

Metric	1.485	1.49	1.495	1.50	1.505	1.51
Sharpness $\times 10^3$	2.55	3.65	**4.00**	3.04	2.65	2.60
Cost	58.9	32.7	**22.2**	33.0	55.7	78.3
RMSE $\times 10^{-2}$	7.42	7.34	7.17	**6.83**	7.15	8.19

the maximum of the measured pressure was added to the data. The data were then filtered with several Hann low-pass filters with cutoff frequencies of 2 MHz, 4 MHz, and 6 MHz and downsampled to a sampling rate of 80 MHz. The value of data fidelity term was evaluated by employing a simulation grid of 1024×1024, a pixel size of 0.1 mm, 4000 time points, and a sampling rate of 80 MHz. The data fidelity term was evaluated at a collection for tissue sound speed values and separately for a collection of background sound speed values. The true initial pressure distribution and the remaining sound speed value were kept fixed at their true values.

5. Results of computer-simulation studies.

5.1. Fixed sound speed. The reconstructed initial pressure distributions for several fixed constant sound speed values and $\beta = 10^{-3}$ are shown in Figure 6. The RMSEs of each image are shown in the lower left-hand corner. The value of β was tuned in order to minimize the RMSEs for each sound speed.

Three criteria for selecting the optimal constant sound speed are compared: 1) the Tenenbaum sharpness [51], 2) the value of the cost function given in (11), and 3) the RMSEs. The values of the metrics for several constant sound speed values can be found in Table 3. For consistency, all results are reported for the $\beta = 10^{-3}$ case, even though the RMSE is lower for a lower value of β for the 1.495 mm/μs sound speed case.

The sharpness and cost metrics both suggest a constant sound speed value of 1.495 mm/μs, while the RMSE suggests a similar, though slightly higher, value of 1.500 mm/μs. In practice, it is not possible to tune the sound speed according to the RMSE as the true initial pressure

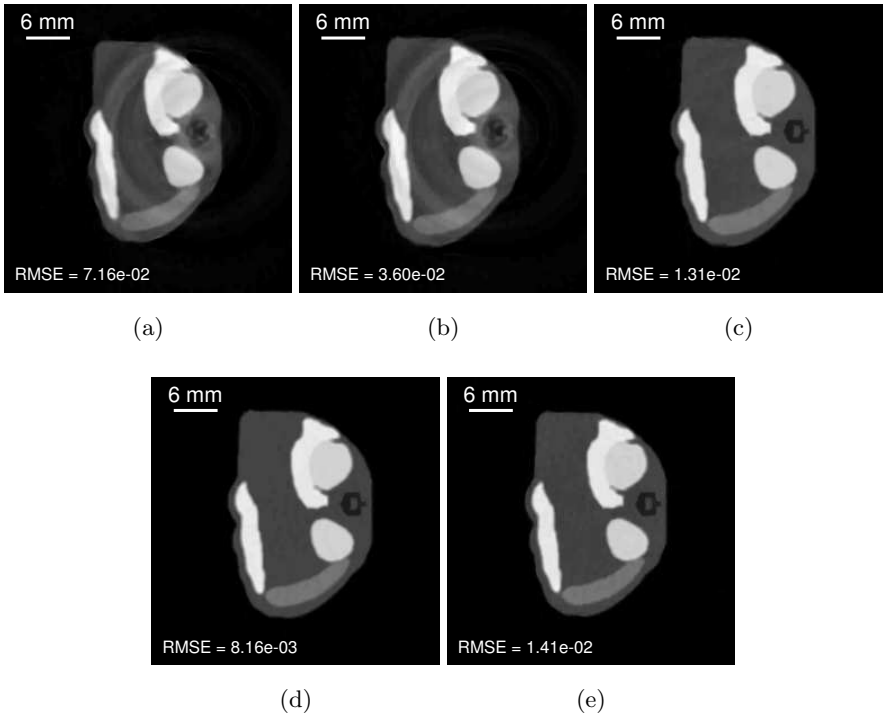

Figure 7. *Reconstructed initial pressure distributions for parameterized JR for* $\beta = 10^{-4}$ *with* (a) 1 *parameter,* (b) 2 *parameters,* (c) 3 *parameters,* (d) 7 *parameters, and* (e) 55869 *parameters. The images are shown in a grayscale window of* $[0.0, 1.1]$.

distribution is unknown. Both the sharpness and cost metrics may serve as adequate proxies. In some cases, performing JR, in which a single image is reconstructed, may be quicker than reconstructing a series of images for different sound speed values and evaluating them according to some metric. It may also be possible to design the cost function such that the cost metric more closely matches the relevant image quality measure.

5.2. Joint reconstruction. Initial pressure distributions and parameterized sound speed distributions were jointly reconstructed for the five parameterizations given in Table 2. The reconstructed initial pressure distributions are shown in Figure 7. For the purposes of evaluation and comparison with the fixed constant sound speed results, the value of β was tuned to minimize the RMSE of the initial pressure distribution. The accuracy of the reconstructed initial pressure distributions is greatly improved when the chosen sound speed parameterization allows compensation for the bone ($Q \geq 3$), which is the strongest source of acoustic heterogeneity in the phantom. This is reflected in both the apparent visual quality of the images and the RMSEs. Additionally accounting for sound speed variations within the soft tissue ($Q = 7$) leads to a minor improvement in the RMSE. The RMSE of reconstructed initial pressure distribution when the sound speed parameterization allows the sound speed values for each pixel to be independently estimated ($Q = 55869$) is worse than either the $Q = 3$ or $Q = 7$ cases. However, there is little obvious visual difference for the three cases where $Q \geq 3$.

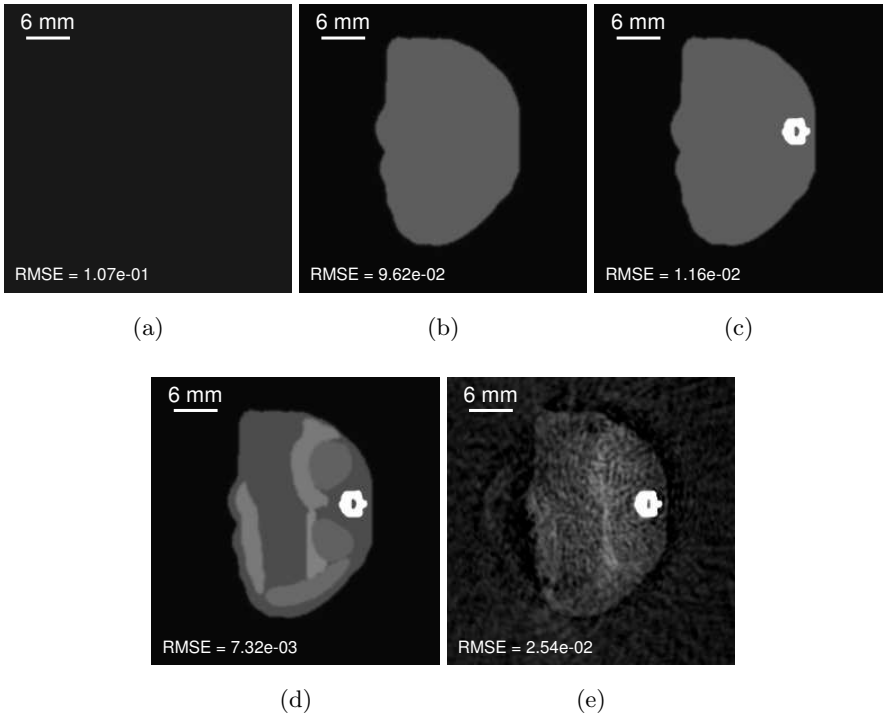

Figure 8. *Reconstructed sound speed distributions for parameterized JR for* $\beta = 10^{-4}$ *with* (a) 1 *parameter,* (b) 2 *parameters,* (c) 3 *parameters, and* (d) 7 *parameters, and* (e) 55869 *parameters. The images are shown in a grayscale window of* $[1.47, 1.70]$ *mm/μs.*

Table 4

Estimated sound speed values for different parameterized sound speed models. For models with $Q < 7$*, the estimated sound speed values represent effective sound speeds across several different tissue types.*

	Water	Bulk tissue	Bone	Kidney	Liver	Pancreas	Spleen
True	1.480	1.540	3.198	1.560	1.578	1.591	1.567
$Q = 1$	1.495	–	–	–	–	–	–
$Q = 2$	1.481	1.555	–	–	–	–	–
$Q = 3$	1.480	1.555	3.171	–	–	–	–
$Q = 7$	1.480	1.540	3.258	1.558	1.578	1.588	1.566

The corresponding reconstructed sound speed distributions are shown in Figure 8. For ease of visualization and comparison with the true sound speed phantom, the equivalent pixelwise representations of the reconstructed parameterized sound speed distributions are presented. The estimated sound speed values for $Q \leq 7$ are also summarized in Table 4. For the $Q = 7$ case, the estimated sound speed distribution closely matches the true sound speed distribution. For $Q < 7$, the true sound speed distribution cannot be recovered due to the choice of the parameterization. In these cases, effective sound speeds are estimated for certain regions. For example, for the $Q = 2$ and $Q = 3$ cases, a single sound speed value is estimated for the soft tissue of the mouse. In these cases, the estimated sound speed falls between the

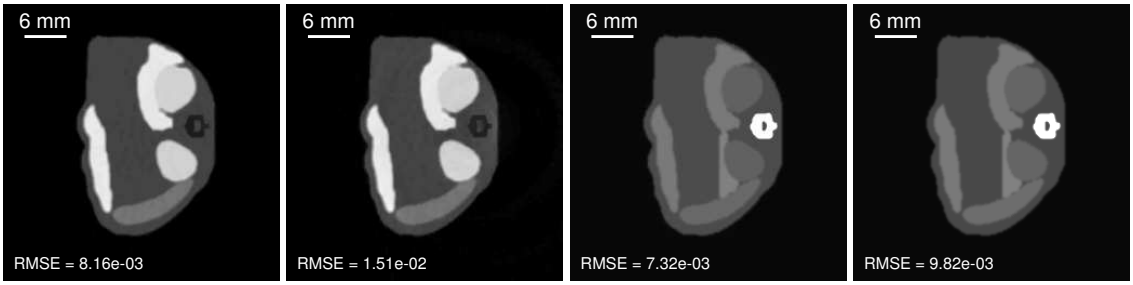

(a) (b) (c) (d)

Figure 9. *Initial pressure distributions reconstructed by parameterized JR for* $\beta = 10^{-4}$ *with* 7 *parameters from measured data that* (a) *ignored mass density variations and* (b) *included mass density variations. Sound speed distributions reconstructed by parameterized JR with* 7 *parameters from measured data that* (c) *ignored mass density variations and* (d) *included mass density variations. Mass density variations were not included in the reconstruction process. The initial pressure images are shown in a grayscale window of* $[0.0, 1.1]$. *The sound speed images are shown in a grayscale window of* $[1.47, 1.70]$ *mm/*μ*s.*

lower sound speed value of the bulk tissue and the higher sound speed values of the organs (1.555 mm/μs for $Q = 2$ and $Q = 3$). Similarly, for the $Q = 1$ case, a single effective sound speed is estimated, which is between the background sound speed and the bulk soft tissue sound speed of the mouse (1.495 mm/μs). When each pixel is allowed to independently vary, the estimated sound speed distribution is very noisy.

5.3. Mass density variations. The previous results were reconstructed from measured data that were generated assuming a homogeneous mass density distribution. However, the mass density of bone differs substantially from that of soft tissue. In principle, the proposed approach could be employed to jointly estimate the mass density distribution along with the initial pressure and sound speed distributions. This could result in an inverse problem that is more ill-conditioned. In order to evaluate the impact of ignoring mass density variations, the measured data were generated for the heterogeneous mass density distribution shown in Figure 1(d). Then, parameterized JR was performed assuming a homogeneous mass density distribution. The reconstructed initial pressure and sound speed distributions are shown in Figure 9. As references, the initial pressure and sound speed distributions reconstructed from measured data that did not include mass density variations are also shown. Ignoring mass density variations did have a minor impact on the accuracy of the reconstructed initial pressure and sound speed distributions. In particular, the estimated sound speed value of the bone was overestimated (3.340 vs. 3.198 mm/μs). Modeling the mass density variations during the simulation to generate the measured data increased the percentage of acoustic energy reflected at the bone soft tissue interface. Increasing the sound speed value of bone during the reconstruction, in the absence of mass density variations, also results in a larger percentage of acoustic energy being reflected. This could explain the overestimation of the bone sound speed.

5.4. Approximate segmentations. Figure 10 shows the reconstructed initial pressure and sound speed distributions obtained by parameterized JR for the approximate parameteriza-

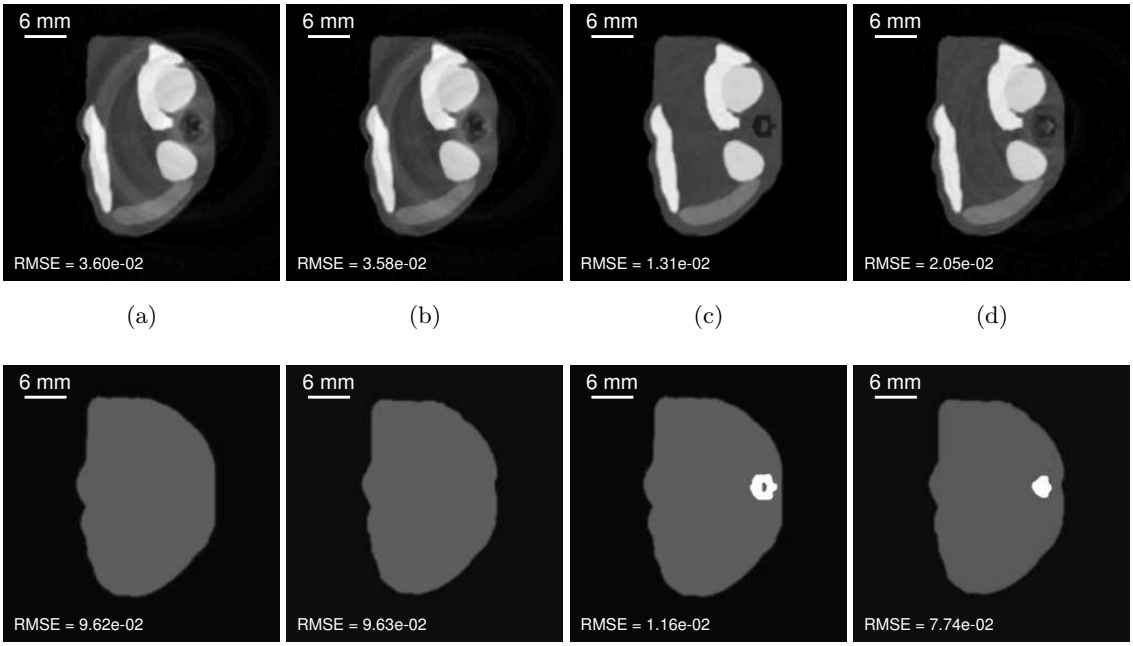

Figure 10. *Reconstructed initial pressure distributions obtained by parameterized JR for* $\beta = 10^{-4}$ *with* (a) 2 *parameters with perfect segmentation,* (b) 2 *parameters with imperfect segmentation,* (c) 3 *parameters with perfect segmentation, and* (d) 3 *parameters with imperfect segmentation. Reconstructed sound speed distributions obtained by parameterized JR with* (e) 2 *parameters with perfect segmentation,* (f) 2 *parameters with imperfect segmentation,* (g) 3 *parameters with perfect segmentation, and* (h) 3 *parameters with imperfect segmentation. The initial pressure images are shown in a grayscale window of* $[0.0, 1.1]$. *The sound speed images are shown in a grayscale window of* $[1.47, 1.70]$ *mm/μs.*

tions shown in Figure 3. As references, the reconstructed initial pressure and sound speed distributions obtained by use of the true parameterizations are also shown. For the $Q = 2$ case, the approximate parameterization is very close to the true parameterization, as the outer boundary of the mouse can be estimated relatively accurately. As a result, there are only minor differences between the reconstructed images for the exact and the approximate segmentations. For the $Q = 3$ case, the outer boundary of bone was not estimated particularly accurately. As a result, there are larger differences between the reconstructed images for the exact and approximate segmentations. Still, the reconstructed initial pressure and sound speed images for the $Q = 3$ approximate parameterizations both have lower RMSEs than the corresponding images for $Q = 2$ with exact parameterizations. This suggests that the method is at least somewhat robust to errors in the assumed parameterization.

5.5. Impact of frequency content. The values of the data fidelity term, evaluated at the true initial pressure distribution and the true background sound speed, for various values of the tissue sound speed are shown in Figure 11(a). The values are normalized in order to allow comparison for measured data subjected to low-pass filters with different cutoff frequencies. It can be seen that the width of the basin of attraction about the true tissue sound speed of

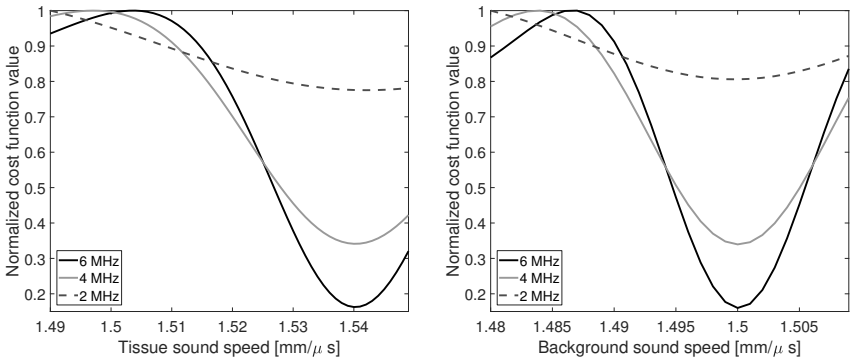

(a) (b)

Figure 11. *Normalized cost function values for* (a) *different tissue sound speeds and* (b) *different background sound speeds after applying a Hann low-pass filter to the measured data. Results are shown for several cutoff frequencies.*

1.54 mm/μs grows wider as the cutoff frequency for the low-pass filter is reduced. This provides some evidence of how low-pass filtering the measured data can help avoid local minima. For example, for the data with a cutoff frequency of 6 MHz, if an initial guess for the tissue sound speed of 1.5 mm/μs, the background value, was employed, the estimated value would become smaller instead of approaching its true value. This is not true for the cases with lower cutoff frequencies. However, there is a trade-off in filtering the data too aggressively. The basin of attraction also becomes shallower as the cutoff frequency is decreased. This could result in slower convergence. This is consistent with prior work that recommends maximizing bandwidth while avoiding cycle skipping [53].

Similar observations can be made for the case where the background sound speed is swept and the initial pressure distribution and the tissue sound speed are held fixed at their true values. In this case, the width of the basin of attraction is even narrower (see Figure 11(b)). This is because the propagation distance through the background is longer than that through the tissue for the considered phantom. As a result, phase errors can accumulate over longer distances, leading to a greater sensitivity to absolute errors in the estimated sound speed.

6. Description of experimental studies.

The value of the proposed approach was also evaluated through experimental studies.

6.1. Methods.

In vivo measurements were previously acquired from the trunk of an anesthetized adult nude mouse (Hsd:Athymic Nude-FoxlNU, Harlan) [31]. All experimental procedures were carried out in conformity with laboratory animal protocols approved by the Animal Studies Committee at Washington University in St. Louis.

The imaging system consisted of a circular ring array of radius 50 mm with 512 evenly distributed ultrasonic transducers. The transducers were elevationally focused and had a nominal central frequency of 5 MHz. The pressure data were recorded at a sampling rate of 40 MHz for 2000 time steps. By manual inspection, 12 transducers were identified as having low sensitivity and were not employed during the image reconstruction process. Illumination was performed by a 1064 nm pulsed laser with a 5-9 ns pulse width and pulse repetition

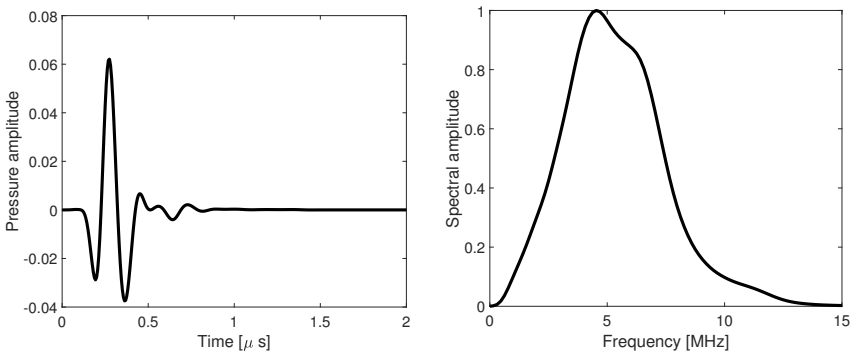

(a) (b)

Figure 12. (a) *The pressure amplitude and* (b) *frequency spectrum of the estimated EIR for the experimental system.*

rate of 50 Hz (DLS9050, Continuum). The fluence on the tissue surface was approximately 18 mJ/cm^2. More detailed information on this imaging system can be found in [31].

The EIR of the transducers was estimated by measuring the pressure from a tiny point-like absorber placed in the center of the ring array as described in the supplementary material of [31]. The measured pressure signals were shifted to align the peaks in order to account for small differences in the propagation distances. Then, the frequency response of a point source was deconvolved from the measured pressure by use of Wiener deconvolution. Finally, the responses were averaged across all transducers. The resulting estimated EIR is shown in Figure 12.

The measured data were upsampled by linear interpolation to an effective sampling frequency of 80 MHz in order to increase the numerical stability of the wave solver [49]. The resulting data were then filtered with a Butterworth low-pass filter with a cutoff of 12 MHz. In addition, the measured pressure data were shifted to account for any unwanted temporal shift induced by the EIR. To estimate the size of this shift, a narrow pulse was convolved with EIR and the shift in the envelope of the pulse was calculated.

6.2. Fixed sound speed. The initial pressure distribution was reconstructed for several fixed constant sound speed values by solving (5). The simulation grid consisted of 2048×2048 pixels with a pixel size of 0.05 mm. A temporal sampling rate of 80 MHz was employed. The radius of the field-of-view was 40 mm. The stopping criterion was when the change in the ℓ_2-norm of the initial pressure distribution between successive iterations was less than 10^{-4}.

6.3. Joint reconstruction. Joint reconstruction was performed for a two-parameter sound speed model. To obtain the parameterized model, the outer boundary of the mouse was manually segmented from the reconstructed initial pressure distribution image for a fixed constant sound speed of 1.50 mm/μs. The estimated segmentation is shown in Figure 13.

To avoid the phenomenon of cycle skipping discussed in section 4.6, the JR images were estimated by a multistage process. First, the measured data were filtered with a Hann low-pass filter with a cutoff frequency of 1 MHz. Due to the lower cutoff frequency, the temporal sampling frequency was reduced to 40 MHz from the 80 MHz employed when the cutoff was

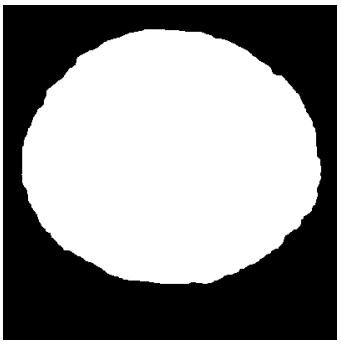

Figure 13. *Estimated segmented regions of the mouse for a 2-parameter sound speed model.*

12 MHz. The number of pixels in the simulation grid and the pixel size were also reduced to 1024×1024 and 0.1 mm, respectively. The initial guess for the initial pressure distribution was the vector of all zeros. The initial guess for the parameterized sound speed distribution was 1.48 mm/μs for the background and 1.54 mm/μs for the mouse. This stage was run for 600 iterations. It was observed that later iterations resulted in small changes to initial pressure distribution and minimal changes to the estimated parameterized sound speed distribution. Second, the estimated initial pressure and sound speed distributions were refined by JR by use of the measured data with the 12 MHz cutoff frequency. By increasing the cutoff frequency, finer structures in the object can be recovered. The initial guesses were the estimated initial pressure and sound speed distributions obtained from the earlier stage. This stage was run for 500 iterations. Finally, the initial pressure distribution was reconstructed for the fixed sound speed distribution obtained by the previous stage by use of FISTA [6]. In this case, the stopping criterion was when the change in the ℓ_2-norm of the initial pressure distribution between successive iterations was less than 10^{-4}. Since the final stage employed an accelerated first-order optimization method (FISTA), faster convergence could be achieved compared with the unaccelerated proximal gradient method employed as part of the JR procedure. This also avoids the line search procedure employed to choose a step size for updating the parameterized sound speed distribution.

7. Results of experimental studies. The reconstructed initial pressure distribution images for several fixed constant sound speed values are shown in Figure 14. In the reconstructed images, strong surface and interior vessel structures are observed. The interior vessel structures appear most in focus for a constant sound speed value of 1.500 mm/μs. However, for this sound speed value, some surface vessels, particularly along the lower right side of the mouse, appear as arcs rather than points, suggesting that they are out-of-focus. Thus, there is no single constant sound speed value for which all of the features of the image are in focus.

A multistage JR process was also employed to estimate the initial pressure and sound speed distributions. The estimated sound speed values following the first stage were 1.489 mm/μs for the background and 1.565 mm/μs for the body of the mouse. These were refined to 1.492 mm/μs and 1.561 mm/μs for the background and mouse body, respectively, during the second stage. In Figure 15, the initial pressure distribution images reconstructed by assum-

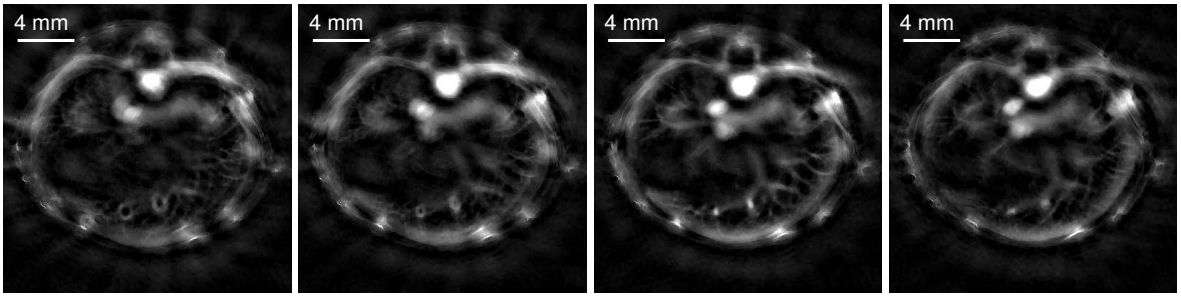

(a) (b) (c) (d)

Figure 14. *Reconstructed initial pressure distributions for several fixed constant sound speed values of* (a) 1.490 *mm/μs,* (b) 1.495 *mm/μs,* (c) 1.500 *mm/μs, and* (d) 1.505 *mm/μs. Results are shown for* $\beta = 10^{-1}$ *in a grayscale window of* $[0, 8000]$.

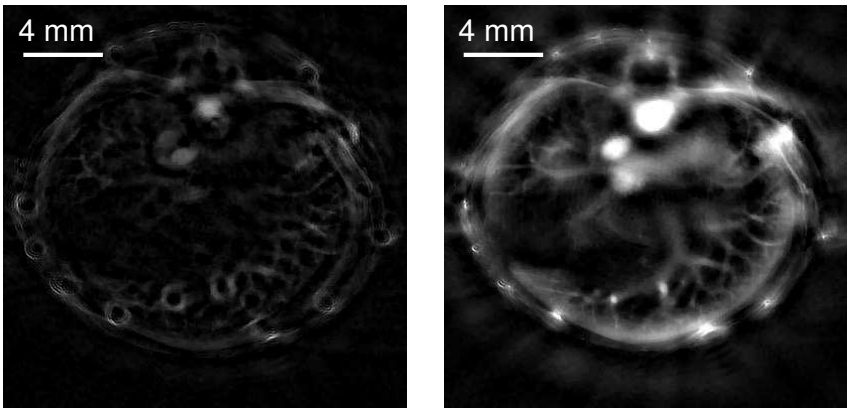

(a) (b)

Figure 15. *Reconstructed initial pressure distributions for* $\beta = 10^{-1}$ *assuming* (a) *the sound speed distribution employed as the initial guess for the first stage of JR and* (b) *the sound speed distribution obtained by parameterized JR for* $Q = 2$. *Results are shown in a grayscale window of* $[0, 8000]$.

ing the initial guess for the sound speed distribution and the final sound speed distribution obtained by JR are shown. The vessels can be much more clearly visualized in the image obtained with the final sound speed distribution compared with that corresponding to the initial guess. By employing a multistage reconstruction process, a focused initial pressure distribution image can be obtained even if there are moderate errors in the initial guess for the parameterized sound speed distribution.

The JR image also demonstrates improvement over the image obtained with a tuned constant sound speed. This is seen most clearly in the zoomed-in regions of the reconstructed initial pressure distributions shown in Figure 16. The largest difference can be seen in the rightmost surface vessel, which appears as an arc in the constant sound speed image and a point in the JR image. In addition, several interior vessels are better focused in the JR image.

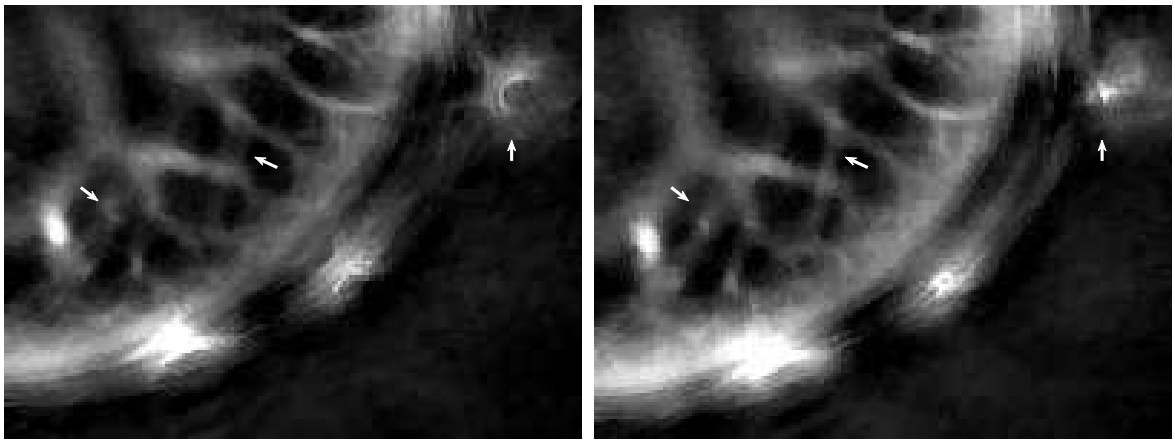

(a) (b)

Figure 16. *Zoomed-in region of the reconstructed initial pressure distributions for* $\beta = 10^{-1}$ *assuming* (a) *a tuned constant sound speed of* 1.500 *mm/μs and* (b) *the sound speed distribution obtained by parameterized JR for* $Q = 2$. *The arrows point to structures that are in focus in the JR image, but not in the tuned constant sound speed image. Results are shown in a grayscale window of* $[0, 6000]$.

8. Conclusions. Parameterized JR can allow estimation of the initial pressure distribution and a low-dimensional representation of the sound speed distribution from PACT measurements alone. This can result in more accurate reconstructed initial pressure distributions than assuming a constant sound speed. By considering a low-dimensional parameterized form for the sound speed distribution, the JR problem can be made more stable. In addition, parameterized JR may offer advantages compared with manually tuning a parameterized sound speed model, particularly as the number of parameters in the model is increased.

More generally, this approach can be thought of as a specific example of employing different discretizations of the object for the forward and inverse problems. By choosing the discretizations independently, representations of the object can be selected such that they are most suitable for their respective tasks. For example, the spatial sampling requirements of the k-space numerical wave solver employed in this work necessitated a small pixel size and a correspondingly high-dimensional representation of the sound speed distribution [49]. On the other hand, the inverse problem was seen to benefit from having fewer unknowns.

Still, parameterized JR does not eliminate many of the issues associated with JR of the initial pressure and sound speed distributions in a general setting. Local minima still represent a challenge, particularly for the case of high-frequency measured data. Practical methods for determining useful parameterizations for the sound speed distribution in an experimental setting are also needed. More careful design of cost function could help ensure that minimizing the cost function produces images that maximize the most relevant image quality measures. The assumed forward model also has a number of limitations, including 1) a two-dimensional imaging model is assumed while experimentally the wave propagation is a three-dimensional phenomenon, 2) the spatial impulse response of the transducers is ignored, and 3) acoustic attenuation and mass density variations are neglected. In addition, use of a forward model based on the elastic wave equation would permit consideration of shear waves, which can

occur in bone [36]. More efficient optimization methods for solving the JR problem, such as accelerated first-order or second-order methods, would reduce the computational cost and thus increase the attractiveness of this approach.

Appendix A. Overview of adjoint state method. The adjoint state method can be employed to calculate the gradient of a functional F with respect to some model parameters m. Excellent reviews of this method and its application to inverse scattering and geophysics can be found in [37] and [42], respectively. Consider the case where the functional depends on the model parameters through some state variables Ψ,

$$F(m) = f(\Psi(m), m). \tag{18}$$

For example, m could be the sound speed or initial pressure distributions, Ψ could be the acoustic pressure, and F could be the least squares error between the measured acoustic pressure and the estimated acoustic pressure given some model. The relationship between m and Ψ can be viewed as being enforced by some constraints,

$$h_k(\Psi, m) = 0 \tag{19}$$

for $k \in [0,\ K-1]$, where K is the total number of constraints. These constraints could correspond to the acoustic wave equation, given by (1).

The derivative of F with respect to m is then

$$\frac{dF}{dm} = \frac{\partial f}{\partial \Psi}\frac{\partial \Psi}{\partial m} + \frac{\partial f}{\partial m}. \tag{20}$$

This expression seems to indicate that calculation of $\frac{dF}{dm}$ requires calculation of $\frac{\partial \Psi}{\partial m}$, which is often high-dimensional and expensive to compute.

The adjoint state method provides a prescription for calculating $\frac{dF}{dm}$ without explicitly computing $\frac{\partial \Psi}{\partial m}$. The adjoint state method can be derived by use of perturbation theory or through the use of Lagrange multipliers. Here, the Lagrangian interpretation is given. Consider the augmented functional $\mathcal{L}$, given by

$$\mathcal{L}(\Psi, m, \lambda) = f(\Psi, m) - \sum_k \langle \lambda_k, h_k(\Psi, m) \rangle, \tag{21}$$

where λ_k are the adjoint state variables and λ is shorthand for the collection of all K adjoint state variables. This corresponds to the Lagrangian for the optimization problem of minimizing f with respect to Ψ subject to the constraints $h_k(\Psi, m) = 0$. Note that when $h_k(\Psi, m) = 0$, i.e., Ψ and m are related by the chosen constraints,

$$\frac{dF}{dm} = \frac{\partial \mathcal{L}}{\partial \Psi}\frac{\partial \Psi}{\partial m} + \frac{\partial \mathcal{L}}{\partial m}. \tag{22}$$

This holds independent of the choice of λ. As a result, λ can be chosen such that (22) can be evaluated efficiently. In particular, λ is chosen such that

$$\frac{\partial \mathcal{L}}{\partial \Psi} = \frac{\partial f}{\partial \Psi} - \sum_k \frac{\partial}{\partial \Psi}\langle \lambda_k, h_k \rangle = 0. \tag{23}$$

This is known as the adjoint state equation. Let the solution to (23) be denoted $\widetilde{\lambda}$ and let $\widetilde{\Psi}$ satisfy the constraints $h_k(\Psi, m) = 0$. Then, $(\widetilde{\Psi}, \widetilde{\lambda})$ represents a saddle point of $\mathcal{L}$ as $\frac{\partial \mathcal{L}}{\partial \Psi} = 0$ and $\frac{\partial \mathcal{L}}{\partial \lambda_k} = -h_k(\Psi, m) = 0$ at that point. With this, the derivative of F with respect to m can finally be calculated as

$$\frac{dF}{dm} = \frac{\partial \mathcal{L}\left(\widetilde{\Psi}, m, \widetilde{\lambda}\right)}{\partial m} = \frac{\partial f}{\partial m} - \sum_k \left\langle \widetilde{\lambda_k}, \frac{\partial h_k}{\partial m} \right\rangle. \tag{24}$$

The utility and applicability of this method depends on a number of factors. First, F, f, h_k, and Ψ must be continuously differentiable. Second, each m must correspond to a unique Ψ. Third, there must exist an efficient method for solving (23).

Appendix B. Application of adjoint state method to PACT. For simplicity, in this section, a continuous formulation of the joint optimization problem is considered, in keeping with the presentation of the previous section. The constraints correspond to the acoustic wave equation given by (1),

$$h_0(p, \mathbf{u}, p_0, c) = \rho(\mathbf{r}) \frac{\partial \mathbf{u}(\mathbf{r}, t)}{\partial t} + \nabla p(\mathbf{r}, t), \tag{25a}$$

$$h_1(p, \mathbf{u}, p_0, c) = \frac{1}{\rho(\mathbf{r}) c(\mathbf{r})^2} \frac{\partial p(\mathbf{r}, t)}{\partial t} + \nabla \cdot \mathbf{u}(\mathbf{r}, t), \tag{25b}$$

$$h_2(p, \mathbf{u}, p_0, c) = p(\mathbf{r}, 0) - p_0(\mathbf{r}), \tag{25c}$$

$$h_3(p, \mathbf{u}, p_0, c) = \mathbf{u}(\mathbf{r}, 0), \tag{25d}$$

where p and $\mathbf{u}$ are the state variables and c and p_0 are the sought-after model parameters.

In this work, the functional to-be-minimized is

$$f(p) = \frac{1}{2} \sum_{i=0}^{M-1} \int_0^T dt \left(\underline{g_i}(t) - \mathcal{M}_i p(\mathbf{r}, t)\right)^2, \tag{26}$$

where $\underline{g_i}(t)$ is the pressure recorded by the ith transducer and $\mathcal{M}_i$ is the restriction of the pressure over the whole domain to the location of the ith transducer. The augmented functional is given by

$$\begin{aligned} \mathcal{L} = {} & \frac{1}{2} \sum_{i=0}^{M-1} \int_0^T dt \left(g_i(t) - \mathcal{M}_i p(\mathbf{r}, t)\right)^2 - \int_0^T dt \left\langle \boldsymbol{\lambda}_1, \rho \frac{\partial \mathbf{u}}{\partial t} + \nabla p \right\rangle_U \\ & - \int_0^T dt \left\langle \lambda_2, \frac{1}{\rho c^2} \frac{\partial p}{\partial t} + \nabla \cdot \mathbf{u} \right\rangle_V - \langle \boldsymbol{\lambda}_3, \mathbf{u}(\mathbf{r}, 0) \rangle_U - \langle \lambda_4, p(\mathbf{r}, 0) - p_0(\mathbf{r}) \rangle_V, \end{aligned} \tag{27}$$

where

$$\langle \mathbf{x}, \mathbf{y} \rangle_U = \sum_{i=1}^n \int d\mathbf{r}\ x_i(\mathbf{r}) y_i(\mathbf{r}),$$

$$\langle a, b \rangle_V = \int d\mathbf{r}\ a(\mathbf{r}) b(\mathbf{r})$$

are the inner products for vector-valued and scalar-valued quantities, respectively. Here, n is

the number of components in the vector-valued quantity and x_i represents the ith component of $\mathbf{x}$.

To compute the derivatives needed to solve (23), it is useful to rearrange the constraint terms by use of integration-by-parts and the fact that the adjoint of the gradient operator is the negative divergence operator. The term in the Lagrangian corresponding to (25a) can be rewritten as

$$\begin{aligned}\int_0^T dt \left\langle \boldsymbol{\lambda}_1, \rho\frac{\partial \mathbf{u}}{\partial t} + \nabla p \right\rangle_U &= \int_0^T dt \left[\left\langle \boldsymbol{\lambda}_1, \rho\frac{\partial \mathbf{u}}{\partial t} \right\rangle_U + \langle \boldsymbol{\lambda}_1, \nabla p \rangle_U \right] \\ &= \int_0^T dt \left[\left\langle \rho\boldsymbol{\lambda}_1, \frac{\partial \mathbf{u}}{\partial t} \right\rangle_U - \langle \nabla \cdot \boldsymbol{\lambda}_1, p \rangle_V \right] \\ &= \langle \rho\boldsymbol{\lambda}_1, \mathbf{u} \rangle_U \,|_0^T - \int_0^T dt \left[\left\langle \rho\frac{\partial \boldsymbol{\lambda}_1}{\partial t}, \mathbf{u} \right\rangle_U + \langle \nabla \cdot \boldsymbol{\lambda}_1, p \rangle_V \right].\end{aligned}$$

Similarly, the term in the Lagrangian corresponding to (25b) can be rewritten as

$$\begin{aligned}\int_0^T dt \left\langle \lambda_2, \frac{1}{\rho c^2}\frac{\partial p}{\partial t} + \nabla \cdot \mathbf{u} \right\rangle_V &= \int_0^T dt \left[\left\langle \lambda_2, \frac{1}{\rho c^2}\frac{\partial p}{\partial t} \right\rangle + \langle \lambda_2, \nabla \cdot \mathbf{u} \rangle_V \right] \\ &= \int_0^T dt \left[\left\langle \frac{1}{\rho c^2}\lambda_2, \frac{\partial p}{\partial t} \right\rangle_V - \langle \nabla\lambda_2, \mathbf{u} \rangle_U \right] \\ &= \left\langle \frac{1}{\rho c^2}\lambda_2, p \right\rangle_V \Big|_0^T - \int_0^T dt \left[\left\langle \frac{1}{\rho c^2}\frac{\partial \lambda_2}{\partial t}, p \right\rangle_V + \langle \nabla\lambda_2, \mathbf{u} \rangle_U \right].\end{aligned}$$

With this, the Lagrangian can be rewritten as

$$\begin{aligned}\mathcal{L} = &\frac{1}{2}\sum_{i=0}^{M-1} \int_0^T dt \left(g_i(t) - \mathcal{M}_i p(\mathbf{r}, t)\right)^2 - \langle \rho\boldsymbol{\lambda}_1, \mathbf{u} \rangle_U \,|_0^T \\ &+ \int_0^T dt \left[\left\langle \rho\frac{\partial \boldsymbol{\lambda}_1}{\partial t}, \mathbf{u} \right\rangle_U + \langle \nabla \cdot \boldsymbol{\lambda}_1, p \rangle_V \right] - \left\langle \frac{1}{\rho c^2}\lambda_2, p \right\rangle_V \Big|_0^T \\ &+ \int_0^T dt \left[\left\langle \frac{1}{\rho c^2}\frac{\partial \lambda_2}{\partial t}, p \right\rangle_V + \langle \nabla\lambda_2, \mathbf{u} \rangle_U \right] - \langle \boldsymbol{\lambda}_3, \mathbf{u}(\mathbf{r}, 0) \rangle_U - \langle \lambda_4, p(\mathbf{r}, 0) - p_0(\mathbf{r}) \rangle_V .\end{aligned} \tag{28}$$

To obtain an expression for the adjoint state equation, compute $\frac{\partial \mathcal{L}}{\partial p}$ and $\frac{\partial \mathcal{L}}{\partial \mathbf{u}}$, set the expressions equal to zero, and group the terms by their time dependence. This yields

$$\rho(\mathbf{r})\frac{\partial \boldsymbol{\lambda}_1(\mathbf{r}, t)}{\partial t} + \nabla\lambda_2(\mathbf{r}, t) = \mathbf{0}, \tag{29a}$$

$$\nabla \cdot \boldsymbol{\lambda}_1(\mathbf{r}, t) + \frac{1}{\rho(\mathbf{r})\, c(\mathbf{r})^2}\frac{\partial \lambda_2(\mathbf{r}, t)}{\partial t} = -\tilde{S}(\mathbf{r}, t), \tag{29b}$$

$$\boldsymbol{\lambda}_1(\mathbf{r}, T) = 0, \tag{29c}$$

$$\lambda_2(\mathbf{r}, T) = 0, \tag{29d}$$

$$\boldsymbol{\lambda}_3(\mathbf{r}) = \rho(\mathbf{r})\,\boldsymbol{\lambda}_1(\mathbf{r}, 0), \tag{29e}$$

$$\lambda_4(\mathbf{r}) = \frac{1}{\rho(\mathbf{r})\, c(\mathbf{r})^2}\lambda_2(\mathbf{r}, 0), \tag{29f}$$

where

$$\tilde{S}(\mathbf{r},t) = \sum_{i=0}^{M-1} \mathcal{M}_i^{\dagger}\left(\mathcal{M}_i p(\mathbf{r},t) - \underline{g_i}(t)\right). \tag{30}$$

The superscript † is employed to denote the adjoint. The action of $\mathcal{M}_i^{\dagger}$ can be computed as

$$\mathcal{M}_i^{\dagger} g(t) = \Lambda_i^{\dagger}\left(h_e^{\dagger}(t) *_t g(t)\right), \tag{31}$$

where $\Lambda_i^{\dagger}$ amounts to placing the corresponding value at the location of the ith transducer and

$$h_e^{\dagger}(t) *_t g(t) = h_e(-t) *_t g(t) = \mathcal{F}^{-1}\left\{\mathcal{F}\left\{h_e\right\}^* \mathcal{F}\left\{g\right\}\right\},$$

where $\mathcal{F}$ is the Fourier transform and * denotes the complex conjugate.

A change of variables can be employed to change the final conditions to initial conditions

$$q_i(\mathbf{r},t) = \lambda_i(\mathbf{r},T-t). \tag{32}$$

This gives a set of differential equations that can be solved by the same method as employed to solve (1):

$$\rho(\mathbf{r})\frac{\partial \mathbf{q}_1(\mathbf{r},t)}{\partial t} + \nabla q_2(\mathbf{r},t) = \mathbf{0}, \tag{33a}$$

$$\nabla\cdot\mathbf{q}_1(\mathbf{r},t) + \frac{1}{\rho(\mathbf{r})\,c(\mathbf{r})^2}\frac{\partial q_2(\mathbf{r},t)}{\partial t} = -\tilde{S}(\mathbf{r},T-t), \tag{33b}$$

$$\mathbf{q}_1(\mathbf{r},0) = 0, \tag{33c}$$

$$q_2(\mathbf{r},0) = 0. \tag{33d}$$

Finally, the gradients with respect to the model parameters are given by

$$\frac{dF}{dp_0}(\mathbf{r}) = \lambda_4(\mathbf{r}) = \frac{1}{\rho(\mathbf{r})\,c(\mathbf{r})^2} q_2(\mathbf{r},T), \tag{34}$$

$$\frac{dF}{dc}(\mathbf{r}) = -\int_0^T dt \frac{2}{\rho(\mathbf{r})\,c(\mathbf{r})^3}\frac{\partial p(\mathbf{r},t)}{\partial t} q_2(\mathbf{r},T-t). \tag{35}$$

Appendix C. Two-parameter backtracking line search. A two-part backtracking line search procedure, detailed in Algorithm 2, is employed to select a scalar step size for the initial pressure distribution and a scalar step size for the parameterized sound speed distribution. First, a step size is chosen for updating the initial pressure distribution, following the method outlined in [41, 7]. During this step, the sound speed is kept fixed at its current estimate. Then, a step size for updating the parameterized sound speed distribution is chosen. During this step, the cost function is evaluated at the updated initial pressure value. In some cases,

if the gradient with respect to the parameterized sound speed distribution changes rapidly for small variations in the initial pressure distribution, the previously calculated gradient may not represent a descent direction at the updated initial pressure distribution value. To address this case, if a suitable step size for the sound speed cannot be found in l_{max} line search steps, the update to the parameterized sound speed distribution is skipped. This ensures that the chosen step sizes lead to a decrease in the value of the cost function. This situation could be avoided by recalculating the gradient with respect to the sound speed at the updated initial pressure value or by performing a grid search over the two step sizes. Whether this additional computational cost is justified will depend on how sensitive the gradient with respect to the sound speed is to changes in the initial pressure distribution, which may be problem-dependent.

Algorithm 2 Two-parameter backtracking line search.

Input: $\alpha_0^p,\ \alpha_0^c,\ F^{(k)}, \mathbf{G}_p^{(k)},\ \mathbf{G}_c^{(k)}, l_{max}$

Output: $\widehat{\alpha_k^p}, \widehat{\alpha_k^c}$

$\widehat{\alpha_k^p} \leftarrow \alpha_0^p$
found $\leftarrow$ **false**
while not found **do**
  $\widetilde{\mathbf{p}_0} \leftarrow \text{prox}_{\widehat{\alpha_k^p}\beta R}\left(\mathbf{p}_0^{(k)} - \widehat{\alpha_k^p}\mathbf{G}_p^{(k)}\right)$
  if $F\left(\widetilde{\mathbf{p}_0}, \Phi\left(\mathbf{c}_p^{(k)}\right)\right) > F^{(k)} + \left\langle \mathbf{G}_p^{(k)}, \widetilde{\mathbf{p}_0} - \mathbf{p}_0^{(k)} \right\rangle + \frac{1}{2\widehat{\alpha_k^p}} \|\widetilde{\mathbf{p}_0} - \mathbf{p}_0^{(k)}\|_2^2$ **then**
    $\widehat{\alpha_k^p} \leftarrow \widehat{\alpha_k^p}/2$
  else
    found $\leftarrow$ **true**
  end if
end while
$l \leftarrow 0$ {l is the line search step number.}
$\widehat{\alpha_k^c} \leftarrow \alpha_0^c$
found $\leftarrow$ **false**
while not found **and** $l < l_{max}$ **do**
  $\widetilde{\mathbf{c}_p} \leftarrow \mathbf{c}_p^{(k)} - \widehat{\alpha_k^c}\mathbf{B}\mathbf{G}_c^{(k)}$
  if $F\left(\widetilde{\mathbf{p}_0}, \Phi\left(\mathbf{c}_p^{(k)}\right)\right) < F\left(\widetilde{\mathbf{p}_0}, \Phi\left(\widetilde{\mathbf{c}_p}\right)\right)$ **then**
    $\widehat{\alpha_k^c} \leftarrow \widehat{\alpha_k^c}/2$
  else
    found $\leftarrow$ **true**
  end if
  $l \leftarrow l + 1$
end while
if l is l_{max} **then**
  $\widehat{\alpha_k^c} \leftarrow 0$
end if

REFERENCES

[1] *NVIDIA: CUDA C Programming Guide*, Technical report PG-02829-001_v7.5, 2015.

[2] M. Anastasio, J. Zhang, E. Sidky, Y. Zou, D. Xia, and X. Pan, *Feasibility of half-data image reconstruction in 3-D reflectivity tomography with a spherical aperture*, IEEE Trans. Med. Imaging, 24 (2005), pp. 1100–1112.

[3] M. Anastasio, J. Zhang, X. Pan, Y. Zou, G. Ku, and L. Wang, *Half-time image reconstruction in thermoacoustic tomography*, IEEE Trans. Med. Imaging, 24 (2005), pp. 199–210.

[4] S. R. Arridge, M. M. Betcke, B. T. Cox, F. Lucka, and B. E. Treeby, *On the adjoint operator in photoacoustic tomography*, Inverse Problems, 32 (2016), 115012.

[5] H. H. Barrett and K. J. Myers, *Foundations of Image Science*, Wiley Ser. Pure Appl. Opt., Wiley, New York, 2004.

[6] A. Beck and M. Teboulle, *Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems*, IEEE Trans. Image Process., 18 (2009), pp. 2419–2434.

[7] A. Beck and M. Teboulle, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM J. Imaging Sci., 2 (2009), pp. 183–202.

[8] S. Bernard, V. Monteiller, D. Komatitsch, and P. Lasaygues, *Ultrasonic computed tomography based on full-waveform inversion for bone quantitative imaging*, Phys. Med. Biol., 62 (2017), 7011.

[9] C. Bunks, F. Saleck, S. Zaleski, and G. Chavent, *Multiscale seismic waveform inversion*, Geophysics, 60 (1995), pp. 1457–1473.

[10] P. L. Carson, C. R. Meyer, A. L. Scherzinger, and T. V. Oughton, *Breast imaging in coronal planes with simultaneous pulse echo and transmission ultrasound*, Science, 214 (1981), pp. 1141–1143.

[11] A. Chambolle, V. Caselles, M. Novaga, D. Cremers, and T. Pock, *An Introduction to Total Variation for Image Analysis*, in Theoretical Foundations and Numerical Methods for Sparse Recovery, M. Fornasier, ed., Radon Ser. Comput. Appl. Math. 9, de Gruyter, Berlin, 2010, pp. 263–340.

[12] D. L. Colton, *Inverse Acoustic and Electromagnetic Scattering Theory*, 3rd ed., Springer, New York, 2013.

[13] T. Ding, K. Ren, and S. Valllian, *A one-step reconstruction algorithm for quantitative photoacoustic imaging*, Inverse Problems, 31 (2015), 095005.

[14] B. Dogdas, D. Stout, A. F. Chatziioannou, and R. M. Leahy, *Digimouse: A 3d whole body mouse atlas from CT and cryosection data*, Phys. Med. Biol., 52 (2007), 577.

[15] N. Duric, P. Littrup, O. Roy, S. Schmidt, C. Li, L. Bey-Knight, and X. Chen, *Breast imaging with ultrasound tomography: Initial results with SoftVue*, in Proceedings of the IEEE International Ultrasonics Symposium, 2013, pp. 382–385.

[16] S. A. Ermilov, T. Khamapirad, A. Conjusteau, M. H. Leonard, R. Lacewell, K. Mehta, T. Miller, and A. A. Oraevsky, *Laser optoacoustic imaging system for detection of breast cancer*, J. Biomed. Opt., 14 (2009), 024007.

[17] S. A. Ermilov, R. Su, A. Conjusteau, V. Ivanov, V. Nadvoretskiy, T. Oruganti, P. Talole, F. Anis, M. A. Anastasio, and A. A. Oraevsky, *3d laser optoacoustic ultrasonic imaging system for research in mice (LOUIS-3dm)*, in Proceedings of SPIE 8943 Photons Plus Ultrasound Imaging and Sensing, 2014, 89430J.

[18] N. B. Everett, B. Simmons, and E. P. Lasher, *Distribution of blood (Fe50) and plasma (I131) volumes of rats determined by liquid nitrogen freezing*, Circ. Res., 4 (1956), pp. 419–424.

[19] J. Greenleaf, S. Johnson, R. C. Bahn, and B. Rajagopalan, *Quantitative cross-sectional imaging of ultrasound parameters*, in Proceedings of the Ultrasonics Symposium, 1977, pp. 989–995.

[20] M. Heijblom, D. Piras, F. M. van den Engh, M. van der Schaaf, J. M. Klaase, W. Steenbergen, and S. Manohar, *The state of the art in breast imaging using the Twente photoacoustic mammoscope: Results from 31 measurements on malignancies*, Eur. Radiol., 26 (2016), pp. 3874–3887.

[21] T. Hopp, M. Zapf, E. Kretzek, J. Henrich, A. Tukalo, H. Gemmeke, C. Kaiser, J. Knaudt, and N. V. Ruiter, *3d Ultrasound Computer Tomography: Update from a Clinical Study*, 9790, 2016, 97900A.

[22] Y. Hristova, P. Kuchment, and L. Nguyen, *Reconstruction and time reversal in thermoacoustic tomography in acoustically homogeneous and inhomogeneous media*, Inverse Problems, 24 (2008), 055006.

[23] C. Huang, L. Nie, R. W. Schoonover, Z. Guo, C. O. Schirra, M. A. Anastasio, and L. V. Wang, *Aberration correction for transcranial photoacoustic tomography of primates employing adjunct image data*, J. Biomed. Opt., 17 (2012), 066016.

[24] C. Huang, K. Wang, L. Nie, L. Wang, and M. Anastasio, *Full-wave iterative image reconstruction in photoacoustic tomography with acoustically inhomogeneous media*, IEEE Trans. Med. Imaging, 32 (2013), pp. 1097–1110.

[25] C. Huang, K. Wang, R. W. Schoonover, L. V. Wang, and M. A. Anastasio, *Joint reconstruction of absorbed optical energy density and sound speed distributions in photoacoustic computed tomography: A numerical investigation*, IEEE Trans. Comput. Imaging, 2 (2016), pp. 136–149.

[26] S. L. Jacques, *Optical properties of biological tissues: A review*, Phys. Med. Biol., 58 (2013), pp. R37–R61.

[27] J. Jose, R. G. H. Willemink, W. Steenbergen, C. H. Slump, T. G. v. Leeuwen, and S. Manohar, *Speed-of-sound compensated photoacoustic tomography for accurate imaging*, Med. Phys., 39 (2012), pp. 7262–7271.

[28] N. Kaliss and D. Pressman, *Plasma and blood volumes of mouse organs, as determined with radioactive iodoproteins*, Exp. Biol. Med., 75 (1950), pp. 16–20.

[29] A. Kirsch and O. Scherzer, *Simultaneous reconstructions of absorption density and wave speed with photoacoustic measurements*, SIAM J. Appl. Math., 72 (2012), pp. 1508–1523.

[30] R. A. Kruger, C. M. Kuzmiak, R. B. Lam, D. R. Reinecke, S. P. Del Rio, and D. Steed, *Dedicated 3d photoacoustic breast imaging*, Med. Phys., 40 (2013), 113301.

[31] L. Li, L. Zhu, C. Ma, L. Lin, J. Yao, L. Wang, K. Maslov, R. Zhang, W. Chen, J. Shi, and L. V. Wang, *Single-impulse panoramic photoacoustic computed tomography of small-animal whole-body dynamics at high spatiotemporal resolution*, Nat. Biomed. Eng., 1 (2017), 0071.

[32] H. Liu and G. Uhlmann, *Determining both sound speed and internal source in thermo- and photo-acoustic tomography*, Inverse Problems, 31 (2015), 105005.

[33] G. D. Ludwig, *The velocity of sound through tissues and the acoustic impedance of tissues*, J. Acoust. Soc. Amer., 22 (1950), pp. 862–866.

[34] S. Manohar, R. G. H. Willemink, F. van der Heijden, C. H. Slump, and T. G. van Leeuwen, *Concomitant speed-of-sound tomography in photoacoustic imaging*, Appl. Phys. Lett., 91 (2007), 131911.

[35] T. P. Matthews and M. A. Anastasio, *Joint reconstruction of the initial pressure and speed of sound distributions from combined photoacoustic and ultrasound tomography measurements*, Inverse Problems, 33 (2017), 124002.

[36] K. Mitsuhashi, J. Poudel, T. Matthews, A. Garcia-Uribe, L. Wang, and M. Anastasio, *A forward-adjoint operator pair based on the elastic wave equation for use in transcranial photoacoustic computed tomography*, SIAM J. Imaging Sci., (2017), pp. 2022–2048.

[37] S. J. Norton, *Iterative inverse scattering algorithms: Methods of computing Frechet derivatives*, J. Acoust. Soc, Amer., 106 (1999), pp. 2653–2660.

[38] B. O'Donoghue and E. Candes, *Adaptive restart for accelerated gradient schemes*, Found. Comput. Math., 15 (2015), pp. 715–732.

[39] L. Oksanen and G. Uhlmann, *Photoacoustic and Thermoacoustic Tomography with an Uncertain Wave Speed*, preprint, arXiv:1307.1618[math], 2013.

[40] A. A. Oraevsky and A. A. Karabutov, *Optoacoustic Tomography*, in Biomedical Photonics Handbook, T. Vo-Dinh, ed., CRC Press, Boca Raton, FL, 2003.

[41] N. Parikh and S. Boyd, *Proximal Algorithms*, Found. Trends Optim. 3, Now Publishers, Delft, the Netherlands, 2013.

[42] R.-E. Plessix, *A review of the adjoint-state method for computing the gradient of a functional with geophysical applications*, Geophys. J. Int., 167 (2006), pp. 495–503.

[43] J. Poudel, T. P. Matthews, L. Li, M. A. Anastasio, and L. V. Wang, *Mitigation of artifacts due to isolated acoustic heterogeneities in photoacoustic computed tomography using a variable data truncation-based reconstruction method*, J. Biomed. Opt., 22 (2017), 041018.

[44] J. Qian, P. Stefanov, G. Uhlmann, and H. Zhao, *An efficient Neumann series-based algorithm for thermoacoustic and photoacoustic tomography with variable sound speed*, SIAM J. Imaging Sci., 4 (2011), pp. 850–883.

[45] L. Sirgue and R. G. Pratt, *Efficient waveform inversion and imaging: A strategy for selecting temporal frequencies*, Geophysics, 69 (2004), pp. 231–248.
[46] P. Stefanov and G. Uhlmann, *Thermoacoustic tomography with variable sound speed*, Inverse Problems, 25 (2009), 075011.
[47] P. Stefanov and G. Uhlmann, *Instability of the linearized problem in multiwave tomography of recovery both the source and the speed*, Inverse Probl. Imaging, 7 (2013), pp. 1367–1377.
[48] T. L. Szabo, *Diagnostic Ultrasound Imaging: Inside Out*, Academic Press, New York, 2013.
[49] M. Tabei, T. D. Mast, and R. C. Waag, *A k-space method for coupled first-order acoustic propagation equations*, J. Acoust. Soc. Amer., 111 (2002), pp. 53–63.
[50] B. Treeby and B. Cox, *k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave-fields*, J. Biomed. Opt., 15 (2010), 021314.
[51] B. E. Treeby, T. K. Varslot, E. Z. Zhang, J. G. Laufer, and P. C. Beard, *Automatic sound speed selection in photoacoustic image reconstruction using an autofocus approach*, J. Biomed. Opt., 16 (2011), 090501.
[52] B. E. Treeby, E. Z. Zhang, and B. T. Cox, *Photoacoustic tomography in absorbing acoustic media using time reversal*, Inverse Problems, 26 (2010), 115003.
[53] J. Virieux and S. Operto, *An overview of full-waveform inversion in exploration geophysics*, Geophysics, 74 (2009), pp. WCC1–WCC26.
[54] K. Wang and M. A. Anastasio, *Photoacoustic and Thermoacoustic Tomography: Image Formation Principles*, in Handbook of Mathematical Methods in Imaging, O. Scherzer, ed., Springer, New York, 2011, pp. 781–815.
[55] K. Wang, S. A. Ermilov, R. Su, H.-P. Brecht, A. A. Oraevsky, and M. A. Anastasio, *An imaging model incorporating ultrasonic transducer properties for three-dimensional optoacoustic tomography*, IEEE Trans. Med. Imaging, 30 (2011), pp. 203–214.
[56] L. V. Wang, ed., *Photoacoustic Imaging and Spectroscopy*, CRC Press, Boca Raton, FL.
[57] L. V. Wang, *Tutorial on photoacoustic microscopy and computed tomography*, IEEE J. Sel. Top. Quantum Electron., 14 (2008), pp. 171–179.
[58] L. V. Wang, *Multiscale photoacoustic microscopy and computed tomography*, Nat. Photonics, 3 (2009), pp. 503–509.
[59] L. V. Wang and H.-i. Wu, *Biomedical Optics: Principles and Imaging*, Wiley, New York, 2007.
[60] J. Wiskin, D. Borup, S. Johnson, and M. Berggren, *Non-linear inverse scattering: High resolution quantitative breast tissue tomography*, J. Acoust. Soc. Amer., 131 (2012), pp. 3802–3813.
[61] J. Xia, M. R. Chatni, K. Maslov, Z. Guo, K. Wang, M. Anastasio, and L. V. Wang, *Whole-body ring-shaped confocal photoacoustic computed tomography of small animals in vivo*, J. Biomed. Opt., 17 (2012), 050506.
[62] J. Xia, C. Huang, K. Maslov, M. A. Anastasio, and L. V. Wang, *Enhancement of photoacoustic tomography by ultrasonic computed tomography based on optical excitation of elements of a full-ring transducer array*, Opt. Lett., 38 (2013), pp. 3140–3143.
[63] Z. Xie, X. Wang, R. F. Morris, F. R. Padilla, G. L. Lecarpentier, and P. L. Carson, *Photoacoustic imaging for deep targets in the breast using a multichannel* $2d$ *array transducer*, in Proceedings of SPIE 7899 Photons Plus Ultrasound Imaging and Sensing, 2011, 789907.
[64] Y. Xu and L. V. Wang, *Effects of acoustic heterogeneity in breast thermoacoustic tomography*, IEEE Trans. Ultrason. Ferroelectr. Freq. Control, 50 (2003), pp. 1134–1146.
[65] Z. Yuan, Q. Zhang, and H. Jiang, *Simultaneous reconstruction of acoustic and optical properties of heterogeneous media by quantitative photoacoustic tomography*, Opt. Express, 14 (2006), pp. 6749–6754.
[66] J. Zhang, K. Wang, Y. Yang, and M. A. Anastasio, *Simultaneous Reconstruction of Speed-Of-Sound and Optical Absorption Properties in Photoacoustic Tomography via a Time-Domain Iterative Algorithm*, in Procedings of SPIE 6856 Photons Plus Ultrasound Imaging and Sensing, 2008, 68561F.