

# Regularized Dual Averaging Image Reconstruction for Full-Wave Ultrasound Computed Tomography

Thomas P. Matthews, Kun Wang, Cuiping Li, Neb Duric, and Mark A. Anastasio

**Abstract**—Ultrasound computed tomography (USCT) holds great promise for breast cancer screening. Waveform inversion-based image reconstruction methods account for higher order diffraction effects and can produce high-resolution USCT images, but are computationally demanding. Recently, a source encoding technique has been combined with stochastic gradient descent (SGD) to greatly reduce image reconstruction times. However, this method bundles the stochastic data fidelity term with the deterministic regularization term. This limitation can be overcome by replacing SGD with a structured optimization method, such as the regularized dual averaging method, that exploits knowledge of the composition of the cost function. In this paper, the dual averaging method is combined with source encoding techniques to improve the effectiveness of regularization while maintaining the reduced reconstruction times afforded by source encoding. It is demonstrated that each iteration can be decomposed into a gradient descent step based on the data fidelity term and a proximal update step corresponding to the regularization term. Furthermore, the regularization term is never explicitly differentiated, allowing nonsmooth regularization penalties to be naturally incorporated. The wave equation is solved by the use of a time-domain method. The effectiveness of this approach is demonstrated through computer simulation and experimental studies. The results suggest that the dual averaging method can produce images with less noise and comparable resolution to those obtained by the use of SGD.

**Index Terms**—Image reconstruction, sound speed imaging, ultrasound computed tomography (USCT), waveform inversion.

## I. INTRODUCTION

ULTRASOUND computed tomography (USCT) shows promise for a number of applications including breast cancer screening [1]–[5]. USCT is ideally suited to breast imaging as it offers novel tissue contrasts that can help differentiate benign masses from tumors [5]. It has several potential advantages over conventional imaging methods, as it is radiation free, breast compression free, and relatively inexpensive. In addition, ultrasound imaging may offer some

advantages over mammography for the detection of breast cancer in women with dense breasts [6], [7]. A variety of studies have been reported demonstrating the application of USCT to breast imaging [6], [8]–[16], with clinical measurements of breast cancer patients having already been performed [3], [17]. While USCT has several potential contrast mechanisms, in this study we focus on the estimation of the sound speed distribution.

Most USCT image reconstruction methods are based on linearized solutions to the acoustic wave equation [3], [6], [11], [18]–[21]. While such methods can possess computational efficient implementations, the spatial resolution of the resulting images can be severely limited by neglecting acoustic diffraction effects in the imaging model. This can hinder breast cancer screening where the ability to identify small tumors and fine features to distinguish cancerous and benign lesions is of great importance. To circumvent the limitations of linearized methods, waveform inversion methods seek to directly invert the acoustic wave equation without relying on linearizations [4], [18], [22]–[27]. Because they can accurately account for the acoustic wave physics, waveform inversion methods can produce high-resolution images; however, these nonlinear methods are computationally burdensome and generally correspond to nonconvex optimization problems. Waveform inversion methods can be classified by whether they solve the wave equation by the use of a time-domain method or a frequency-domain method. While frequency-domain methods have been successfully applied to USCT image reconstruction [14], here we focus on time-domain methods [28], [29].

Recently, an approach that combines waveform inversion with source encoding (WISE), which alleviates much of the computational burden, has been proposed [23], [30], [31]. In [30], the sound speed distribution was estimated by solving an optimization problem, where the cost function consisted of two terms. The first term is a data fidelity term. For this term, the pressure at the transducer locations is calculated based on the current estimate of the sound speed and an acoustic model described by the acoustic wave equation. This term quantifies how closely this estimated pressure matches the measured pressure. As described below, when the source encoding technique is employed, the data fidelity term corresponds to the expectation of a random quantity. The second term is a deterministic regularization term, which is used to incorporate *a priori* information about the image. This optimization problem was solved by the use of stochastic gradient descent (SGD). Under this approach, the stochastic data fidelity term and the deterministic regularization term

Manuscript received September 19, 2016; accepted March 8, 2017. Date of publication March 14, 2017; date of current version May 1, 2017. Computations were performed using the facilities of the Washington University Center for High Performance Computing, which were supported in part by NIH Grant 1S10RR022984-01A1 and Grant 1S10OD018091-01. This work was supported in part by NIH awards CA1744601 and EB01696301 and in part by NSF award DMS1614305.

T. P. Matthews, K. Wang, and M. A. Anastasio are with the Department of Biomedical Engineering, School of Engineering and Applied Science, Washington University in St. Louis, St. Louis, MO 63130 USA (e-mail: anastasio@wustl.edu).

C. Li is with Delphinus Medical Technologies, Plymouth, MI 48170 USA. N. Duric is with Delphinus Medical Technologies, Plymouth, MI 48170 USA, and also with Karmanos Cancer Institute, Wayne State University, Detroit, MI 48201 USA.

Digital Object Identifier 10.1109/TUFFC.2017.2682061

0885-3010 © 2017 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See [http://www.ieee.org/publications\\_standards/publications/rights/index.html](http://www.ieee.org/publications_standards/publications/rights/index.html) for more information.

are treated jointly as part of a single cost function. This approach ignores information about the structure of the cost function and requires the use of a differentiable regularization function [32].

Here, we propose the use of a structured optimization method, known as the regularized dual averaging (RDA) method, that considers the two terms in the cost function separately [32], [33]. This approach can mitigate the impact of the stochastic data fidelity on the deterministic regularization term and result in more effective regularization that offers superior tradeoffs between image resolution and noise variance by exploiting the structure of the cost function. It also provides the opportunity, for the first time, to employ nonsmooth penalties in the waveform inversion cost function, which can be designed to exploit certain sparseness properties of the object [34]–[36].

The remainder of this paper is organized as follows. In Section II, a discrete-to-discrete (D-D) USCT imaging model and the WISE method are reviewed. SGD is discussed briefly. In Section III, the RDA method and its application to USCT image reconstruction are described. The computer simulation studies and experimental results are presented in Sections IV and V, respectively. Finally, this paper concludes with a summary in Section VI.

## II. BACKGROUND

### A. Discrete-to-Discrete USCT Imaging Model

While digital imaging systems are naturally described by a continuous-to-discrete imaging model [37], it is typically necessary to approximate this model as a D-D mapping in order to facilitate the use of iterative image reconstruction algorithms. For simplicity, the D-D model is presented directly.

A canonical 2-D USCT imaging system that employs a circular transducer array [38] that surrounds the object is considered. Ultrasound pulses are transmitted through the object and measured by the transducers. Often, only one transducer will emit a pulse at a given time, with the pressure being recorded by all other transducers. A subset of the transducers will each serve as the emitter in turn, leading to a collection of measurements corresponding to different views of the object. The propagation of the ultrasound waves is governed by the acoustic wave equation, which can be solved by a numerical wave equation solver. This solver can be formulated as a D-D mapping as described below. In this study, the wave equation was solved by the k-space pseudospectral method [39]–[41].

Let  $\mathbf{c} \in \mathbb{R}^N$  denote the finite-dimensional representation, in a pixel basis, of the sought-after sound speed distribution. Here,  $N$  is the number of pixels in the simulation grid employed by the numerical wave solver. The propagation of the pressure wave through the object when the  $m$ th transducer is the emitter can be denoted by

$$\mathbf{g}_m = \mathbf{M}\mathbf{H}(\mathbf{c})\mathbf{s}_m \quad (1)$$

where  $\mathbf{s}_m \in \mathbb{R}^{NL}$  is the emitted pulse,  $\mathbf{g}_m \in \mathbb{R}^{N^{\text{rec}}L}$  is the pressure at each transducer,  $\mathbf{H}(\mathbf{c}) \in \mathbb{R}^{NL \times NL}$  is

the operator that denotes the action of the wave equation,  $\mathbf{M} \in \mathbb{R}^{N^{\text{rec}}L \times NL}$  is a sampling matrix that computes the pressure at the transducer locations from the pressure over the entire simulation grid,  $L$  is the number of time points employed by the wave solver, and  $N^{\text{rec}}$  is the number of transducers acting as receivers. The notation  $\mathbf{H}(\mathbf{c})$  is used to emphasize the dependence of  $\mathbf{H}$  on the sound speed  $\mathbf{c}$ .

An estimate of the sound speed can be obtained by solving the penalized least-squares optimization problem

$$\hat{\mathbf{c}} = \arg \min_{\mathbf{c}} \frac{1}{2} \sum_{m=0}^{M-1} \|\underline{\mathbf{g}}_m - \mathbf{M}\mathbf{H}(\mathbf{c})\mathbf{s}_m\|_2^2 + \lambda \mathcal{R}(\mathbf{c}) \quad (2)$$

where  $M$  is the total number of views,  $\underline{\mathbf{g}}_m$  is the measured pressure at each transducer,  $\mathcal{R}(\mathbf{c})$  is a regularization function, and  $\lambda$  is a regularization parameter, which controls the relative weight of the regularization term. The first term in (2), known as the data fidelity term, is a nonconvex function of  $\mathbf{c}$ , while the regularization function is assumed to be a convex function.

This approach can produce high-resolution images, but it is computationally very expensive. Each evaluation of the cost function requires the wave equation to be solved  $M$  times. This high computational cost has limited the wide spread use of time-domain-based waveform inversion methods.

### B. Waveform Inversion With Source Encoding

Recently, a source encoding technique has been employed to efficiently find the solution of (2) [23], [30]. In the WISE method [30], (2) is reformulated as the stochastic optimization problem

$$\hat{\mathbf{c}} = \arg \min_{\mathbf{c}} \mathbf{E}_{\mathbf{w}} \left\{ \frac{1}{2} \|\underline{\mathbf{g}}_{\mathbf{w}} - \mathbf{M}\mathbf{H}(\mathbf{c})\mathbf{s}_{\mathbf{w}}\|_2^2 \right\} + \lambda \mathcal{R}(\mathbf{c}) \quad (3)$$

where  $\mathbf{w}$  is a random encoding vector,  $\mathbf{E}_{\mathbf{w}}$  denotes the expectation with respect to  $\mathbf{w}$ , and

$$\underline{\mathbf{g}}_{\mathbf{w}} = \sum_{m=0}^{M-1} [\mathbf{w}]_m \underline{\mathbf{g}}_m \quad (4)$$

$$\mathbf{s}_{\mathbf{w}} = \sum_{m=0}^{M-1} [\mathbf{w}]_m \mathbf{s}_m \quad (5)$$

are the encoded measured pressure data and the encoded source term, respectively. Here,  $\mathbf{w}$  is chosen according to a Rademacher distribution as suggested in [42]. Under this formulation, evaluating the cost function for a particular choice of  $\mathbf{w}$  requires the wave equation to be solved only once. When the number of views is large, this can substantially reduce the computational time needed to reconstruct an image. The gradient of the data fidelity term is calculated using an adjoint state method as described in [30]. This approach allows the gradient to be estimated by solving the acoustic wave equation only one additional time (on top of what is already needed to evaluate the cost function). Knowledge of the gradient allows the use of a variety of optimization algorithms.

In [30], (3) was solved by the use of the SGD method, as described in Algorithm 1. In that approach, at each iteration, the gradient of the cost function is evaluated for a single

**Algorithm 1** Stochastic Gradient Descent**Input:**  $\mathbf{c}_0, \lambda$ **Output:**  $\hat{\mathbf{c}}$ 

- 1:  $k \leftarrow 0$  { $k$  is the algorithm iteration number.}
- 2: **while** stopping criterion is not satisfied **do**
- 3: Draw  $\mathbf{w}_k$  according to chosen distribution.
- 4: Calculate  $\mathbf{G}_k \leftarrow \nabla_{\mathbf{c}} f(\mathbf{c}_k, \mathbf{w}_k) + \lambda \nabla_{\mathbf{c}} \mathcal{R}(\mathbf{c}_k)$
- 5: Choose step size  $\alpha_k$
- 6:  $\mathbf{c}_{k+1} \leftarrow \mathbf{c}_k - \alpha_k \mathbf{G}_k$
- 7:  $k \leftarrow k + 1$
- 8: **end while**
- 9:  $\hat{\mathbf{c}} \leftarrow \mathbf{c}_k$

realization of the encoding vector. The update step for the  $(k + 1)$ th iteration for SGD is given by [43]

$$\mathbf{c}_{k+1} = \arg \min_{\mathbf{c}} \left\{ \langle \nabla_{\mathbf{c}} f(\mathbf{c}_k, \mathbf{w}_k), \mathbf{c} \rangle + \frac{1}{2\alpha_k} \|\mathbf{c} - \mathbf{c}_k\|_2^2 + \lambda \mathcal{R}(\mathbf{c}) \right\} \quad (6)$$

or equivalently

$$\mathbf{c}_{k+1} = \mathbf{c}_k - \alpha_k (\nabla_{\mathbf{c}} f(\mathbf{c}_k, \mathbf{w}_k) + \lambda \nabla_{\mathbf{c}} \mathcal{R}(\mathbf{c}_k)) \quad (7)$$

where  $\langle \cdot, \cdot \rangle$  denotes the standard Euclidean inner product,  $\alpha_k$  is the step size,  $\nabla_{\mathbf{c}}$  is the gradient with respect to  $\mathbf{c}$ , and

$$f(\mathbf{c}, \mathbf{w}) \equiv \frac{1}{2} \|\underline{\mathbf{g}}_{\mathbf{w}} - \mathbf{M}\mathbf{H}(\mathbf{c})\mathbf{s}_{\mathbf{w}}\|_2^2. \quad (8)$$

This approach has several limitations. First, it fails to exploit the structure of the objective function. In other words, SGD treats the cost function as a black box, ignoring potentially useful information about the nature of the cost function. For example, in the problem above, the cost function consists of two terms: a stochastic but differentiable data fidelity term and a deterministic regularization term. In SGD, this knowledge is ignored, and the gradients of the stochastic and deterministic terms are lumped together. Second, it assumes that all terms in the cost function are differentiable. This is not true of many sparsity-promoting regularization functions, such as the  $\ell_1$ -norm and the total variation (TV) seminorm. In [30], the TV seminorm was approximated by a smoothed differentiable version through the introduction of a small smoothing parameter. While this approach can be effective, modifications to other nonsmooth regularization functions could be more challenging. Third, it fails to exploit information from previous iterations. For SGD, at each iteration, only the gradient for a single realization of the encoding vector is considered when determining the search direction. When combined with a line search for choosing the step size, this can lead to overfitting [44]. In this case, the line search method will choose a large step that effectively minimizes the cost function evaluated for a single realization of the encoding vector, but which increases, or less effectively minimizes, the cost function evaluated for a large number of realizations. This problem can be overcome by the use of a fixed step size, at the expense of slowing the convergence rate.

**Algorithm 2** Regularized Dual Averaging Method**Input:**  $\mathbf{c}_0, \lambda$ **Output:**  $\hat{\mathbf{c}}$ 

- 1:  $k \leftarrow 0$  { $k$  is the algorithm iteration number.}
- 2:  $A_{-1} \leftarrow 0$
- 3: **while** stopping criterion is not satisfied **do**
- 4: Draw  $\mathbf{w}_k$  according to chosen distribution.
- 5: Calculate  $\mathbf{G}_k \leftarrow \nabla_{\mathbf{c}} f(\mathbf{c}_k, \mathbf{w}_k)$
- 6: Choose weight  $\alpha_k > 0$  {Unweighted case:  $\alpha_k = 1$ }
- 7:  $A_k \leftarrow A_{k-1} + \alpha_k$
- 8:  $\bar{\mathbf{G}}_k \leftarrow \left(1 - \frac{\alpha_k}{A_k}\right) \bar{\mathbf{G}}_{k-1} + \frac{\alpha_k}{A_k} \mathbf{G}_k$  {Compute weighted average of gradient.}
- 9: Choose  $\mu_k$  {For example,  $\mu_k = \gamma A_k$ , where  $\gamma > 0$  is a constant.}
- 10:  $\mathbf{c}_{k+1} \leftarrow \mathbf{c}_0 - \mu_k \bar{\mathbf{G}}_k$
- 11:  $\mathbf{c}_{k+1} \leftarrow \text{prox}_{\lambda, \mu_k \mathcal{R}}(\mathbf{c}_{k+1})$
- 12:  $k \leftarrow k + 1$
- 13: **end while**
- 14:  $\hat{\mathbf{c}} \leftarrow \mathbf{c}_k$

## III. REGULARIZED DUAL AVERAGING METHOD

The dual averaging method is a primal-dual optimization method originally developed by Nesterov [33]. Xiao [32] later extended this approach to include regularization. It can be employed to solve optimization problems of the same form as given in (3). Here, we review the RDA method and detail its application to waveform inversion. Our presentation is similar to that of Xiao [32] and Nesterov [33], but differs in several respects due to differences in the target application. In particular, the data fidelity term of our cost function is nonconvex. This affects how the step size or weights for each gradient term must be chosen. Further, for clarity, we do not attempt to describe the most general form of the RDA method, but merely one that has proved effective for waveform inversion. For the dual averaging method, as described in Algorithm 2, the update step for the  $(k + 1)$ th iteration is given by

$$\mathbf{c}_{k+1} = \arg \min_{\mathbf{c}} \left\{ \langle \bar{\mathbf{G}}_k, \mathbf{c} \rangle + \frac{1}{2\mu_k} \|\mathbf{c} - \mathbf{c}_0\|_2^2 + \lambda \mathcal{R}(\mathbf{c}) \right\} \quad (9)$$

where  $\bar{\mathbf{G}}_k$  is the average gradient of the data fidelity term over all past iterations and  $\mu_k > 0$  is a scalar. This differs from the update step for SGD in two key ways. First, instead of considering the gradient at a single point, the average gradient is employed. Second, the proximal term,  $(1/2)\|\mathbf{c} - \mathbf{c}_0\|_2^2$ , does not depend on the iteration number. In these ways, the RDA method is able to incorporate nonlocal information when determining the estimate of the object for the next iteration.

In the case of simple averaging, the average gradient is given by

$$\bar{\mathbf{G}}_k = \frac{1}{k+1} \sum_{i=0}^k \nabla_{\mathbf{c}} f(\mathbf{c}_i, \mathbf{w}_i). \quad (10)$$

A weighted average of gradient estimates can also be considered, as suggested in [33]. In this case

$$\bar{\mathbf{G}}_k = \frac{1}{\sum_{i=0}^k \alpha_i} \sum_{i=0}^k \alpha_i \nabla_{\mathbf{c}} f(\mathbf{c}_i, \mathbf{w}_i) \quad (11)$$

where  $\{\alpha_i\}_{i=0}^k$  are weights for each of the gradient estimates. Here, the weights are chosen using a line search. Since the search direction is given by the average gradient, adjusting this weight will determine the contribution of the most recent gradient estimate to the search direction. As a result, the direction of the line search changes as this parameter is adjusted. This is in contrast to most line search methods, where the direction is fixed and only the magnitude of the update is affected. As the search direction incorporates information from multiple iterations corresponding to different realizations of the encoding vector, this approach is less prone to overfitting than SGD with a line search. A detailed description of this approach is provided in the Appendix.

If  $\mathcal{R}$  is convex, the dual averaging update step can be written in terms of the proximity operator of  $\mathcal{R}$  as

$$\mathbf{c}_{k+1} = \text{prox}_{\lambda, \mu_k \mathcal{R}}(\mathbf{c}_0 - \mu_k \bar{\mathbf{G}}_k) \quad (12)$$

where the proximity operator is defined as [45]

$$\text{prox}_{\lambda, R}(\mathbf{x}) \equiv \min_{\mathbf{y}} \left\{ R(\mathbf{y}) + \frac{1}{2\lambda} \|\mathbf{x} - \mathbf{y}\|^2 \right\}. \quad (13)$$

From this expression, it becomes clear that the update step for the dual averaging method can be divided into two parts. First, a reference value is updated based on a weighted sum of all past gradient estimates. Second, regularization is incorporated by the use of the associated proximity operator. In this way, the stochastic estimates of the gradient of the data fidelity term are treated separately from the deterministic regularization term. Averaging the gradient estimates obtained over several iterations may help minimize the impact of the variance of the gradient estimates. Further, since the regularization term is not explicitly differentiated, nonsmooth penalties can be easily incorporated through the use of the corresponding proximity operator.

Unless otherwise noted, the regularization function is chosen to be the TV seminorm of the sound speed. The TV seminorm has been shown to be effective at mitigating noise while preserving sharp edges [46]. The proximity operator of the TV seminorm is computed using the fast gradient projection method described in [47] and [48]. Using this approach, the computational cost of applying the proximity operator is much less than that of computing the gradient, so that the computational cost of the RDA method is approximately the same as that of SGD on a per-iteration basis.

The sequence  $\{\mu_k\}$  determines the amount by which the algorithm steps in the search direction. Here, we choose  $\mu_k = \gamma A_k$ , where  $A_k = \sum_{i=0}^k \alpha_i$  and  $\gamma > 0$  is a constant. In this case, line 10 in Algorithm 2 becomes

$$\mathbf{c}_{k+1} \leftarrow \mathbf{c}_0 - \gamma \sum_{i=0}^k \alpha_i \mathbf{G}_i. \quad (14)$$




Fig. 1. Schematic of the measurement geometry. The measurement system consists of a circular ring array of ultrasonic transducers. These transducers are located in a larger rectangular simulation grid, over which the acoustic wave equation is solved. Within the ring array is a smaller rectangular region representing the reconstructed image. The estimated sound speed distribution is calculated within the gray circular field of view within that region.

The constant  $\gamma$  should be chosen to be sufficiently small to insure convergence. In the unweighted case,  $\gamma$  could be chosen to be the inverse of the Lipschitz constant of the gradient of the data fidelity term. It could be similarly chosen in the weighted case as the inverse of the product of the Lipschitz constant and the maximum allowable weight of the gradient,  $\alpha_{\max}$ .

## IV. COMPUTER SIMULATION STUDIES

### A. Methods

Two-dimensional computer simulation studies were performed to compare USCT image reconstruction methods based on SGD and RDA. Studies were performed for two numerical phantoms: 1) a numerical breast phantom [shown in Fig. 2(a)] and 2) a low-contrast phantom with two homogeneous bars [shown in Fig. 2(b)]. The first was employed to establish the potential utility of the proposed approaches for USCT breast imaging and the second was employed to perform a bias-variance analysis comparing SGD and RDA. For both phantoms, the same measurement geometry, excitation pulse, and numerical simulation methods were employed.

1) *Measurement Geometry*: The measurement system consisted of a circular transducer array with a radius of 110 mm and 256 evenly distributed elements. This geometry was chosen to match an existing USCT imaging system [17], [49], [50]. The wavefield data were simulated for 256 views using the first-order k-space method, as described in Section IV-A4 [30], [39], [41]. For each view, one transducer served as the emitter and the pressure was recorded by all 256 transducers. All transducers were modeled as point emitters and receivers. A schematic of this measurement geometry is shown in Fig. 1.

2) *Excitation Pulse*: The excitation pulse was given by

$$s(t) = \exp\left(-\frac{(t-t_c)^2}{2\sigma^2}\right) \sin(2\pi f_c t) \quad (15)$$

where  $f_c = 0.8$  MHz is the central frequency, and  $t_c = 3.2 \mu\text{s}$  and  $\sigma = 0.75 \mu\text{s}$  are the center and width of a Gaussian window, respectively. This corresponds to roughly three cycles. Since the transducers are treated as point emitters, when




Fig. 2. Sound speed distribution of (a) the numerical breast phantom and (b) the low-contrast two bar phantom employed in the bias-variance analysis, given in  $\text{mm}/\mu\text{s}$ .

nearest neighbor interpolation is employed, the source term for the  $m$ th view is given simply by

$$s_m(\mathbf{r}, t) = s(t)\delta(\mathbf{r} - \mathbf{r}_m) \quad (16)$$

where  $\mathbf{r}_m$  is the location of the pixel nearest to the emitter for the  $m$ th view.

3) *Numerical Phantoms*: The numerical breast phantom had a radius of 49 mm and was composed of eight structures representing adipose tissues, parenchymal breast tissues, cysts, benign tumors, and malignant tumors [see Fig. 2(a)]. A detailed description of the numerical breast phantom can be found in [30]. A phantom consisting of two low-contrast bars was created for the bias-variance analysis [see Fig. 2(b)]. The bars were placed far apart to minimize their influence on one another in the reconstructed images.

4) *Simulation of Pressure Data*: In order to avoid inverse crime [51], two related methods were employed to simulate the measured pressure. When generating the pressure data recorded by each transducer, the wave equation was solved by the use of a first-order pseudospectral method [39]. In this method, when attenuation and dispersion are neglected, the acoustic wave propagation is modeled by two coupled first-order differential equations

$$\frac{\partial \mathbf{u}(\mathbf{r}, t)}{\partial t} = -\nabla p(\mathbf{r}, t) \quad (17)$$

$$\frac{1}{c(\mathbf{r})^2} \frac{\partial p(\mathbf{r}, t)}{\partial t} = -\nabla \cdot \mathbf{u}(\mathbf{r}, t) + 4\pi \int_0^t dt' s(\mathbf{r}, t') \quad (18)$$

where  $\mathbf{u}(\mathbf{r}, t)$  is the acoustic particle velocity and  $p(\mathbf{r}, t)$  is the acoustic pressure. The calculation domain was of size  $512 \times 512 \text{ mm}^2$ , sampled on a  $2048 \times 2048$  uniform Cartesian grid with a spacing of 0.25 mm. Nearest neighbor interpolation was employed to place the transducers on the grid points. The pressure was simulated for 3600 time points at a sampling rate of 20 MHz. Additive Gaussian white noise was added to the pressure data. The noise had zero mean and a standard deviation of 5% of the maximum pressure amplitude received by the transducer opposite the emitter for a homogeneous medium.

When reconstructing the sound speed images, the operator  $\mathbf{H}(\mathbf{c})$  was computed by the use of the second-order pseudospectral k-space method [40]. This method solves a single second-order differential equation

$$\nabla^2 p(\mathbf{r}, t) - \frac{1}{c(\mathbf{r})^2} \frac{\partial^2 p(\mathbf{r}, t)}{\partial t^2} = -4\pi s(\mathbf{r}, t). \quad (19)$$

TABLE I  
SUMMARY OF IMAGE RECONSTRUCTION PARAMETERS

Parameter	Simulation	Experimental
Number of pixels	$1024 \times 1024$	$2560 \times 2560$
Grid spacing [mm]	0.5	0.2
Number of time points	1800	3500
Sampling frequency [MHz]	10	20
Number of transducers	256	976

Here, the calculation domain was of size  $512 \times 512 \text{ mm}^2$ , sampled on a  $1024 \times 1024$  uniform Cartesian grid with a spacing of 0.5 mm. The number of time points and the sampling rate were reduced to 1800 and 10 MHz, respectively. These reconstruction parameters are summarized in Table I. Both wave solvers were implemented using NVIDIA's CUDA platform [52]. These pseudospectral k-space methods were chosen for their high numerical accuracy for coarse spatial sampling rates [39], [40].

5) *Bias-Variance Analysis*: The statistical properties of the images produced by the two methods were compared by the use of a bias-variance analysis. The measured pressure was simulated as described above. Five different noise realizations were generated, each with 5% noise. Images were reconstructed for each noise realization for six different regularization parameter values by the use of both SGD with a constant step size and the unweighted RDA method. Each pixel in the reconstructed images can be treated as a random variable  $\hat{c}_i$  (for the  $i$ th pixel), whose true value in the original phantom is  $c_i$ . Due to the long reconstruction times (approximately 1 h for 250 iterations), it was not feasible to reconstruct images for a large number of noise realizations. Instead, each reconstructed image was divided into several regions, which were treated as independent samples for the purposes of this analysis. Specifically, each bar in the reconstructed image was divided into ten identical regions. Corresponding pixels in these regions were treated as having arisen from additional noise realizations. This yielded a total of 100 samples per regularization parameter value. In other words, if the set  $\hat{C}_i$  contains the values of the  $i$ th pixel for the five noise realizations, an augmented set  $\tilde{C}_i$  was created such that

$$\tilde{C}_i = \bigcup_{j=1}^{N_c} \hat{C}_{\mathcal{I}_i(j)} \quad (20)$$

where  $N_c$  is the total number of regions (20) and  $\mathcal{I}_i$  is an iterator that gives the indices of all pixels (across regions) that correspond to the  $i$ th pixel. The bias for a pixel was calculated by averaging these 100 samples and computing the difference between the average value and the corresponding value in the true phantom

$$\text{Bias}_i = \frac{1}{N_s} \sum_{\hat{c} \in \tilde{C}_i} \hat{c} - c_i \quad (21)$$

where  $N_s$  is the total number of samples. A summary measure of the bias was calculated by computing the  $\ell_2$ -norm of the bias values for each pixel. The sample variance of each pixel




Fig. 3. Plot of RMSE versus the number of iterations for images reconstructed by the use of SGD with a constant step size of 0.1 for several regularization parameter values.

across all samples was computed as

$$\text{Var}_i = \frac{1}{N_s - 1} \sum_{\hat{c} \in \hat{C}_i} \left( \hat{c} - \frac{1}{N_s} \sum_{\hat{c} \in \hat{C}_i} \hat{c} \right)^2. \quad (22)$$

The average variance for the pixels was computed as a summary measure. It should be noted that corresponding pixels in different regions may not have the same expected values and variances. In spite of this, the above bias and variance measures still provide insight into the ability of the two reconstruction algorithms to mitigate noise.

### B. Images Reconstructed by SGD

In order to provide a clear and fair point of comparison of the RDA and SGD methods, USCT image reconstruction from noisy data by the use of SGD was first considered and optimized. As seen in Fig. 3, the above choice of  $5 \times 10^{-4}$  for the regularization parameter value results in the most accurate reconstructed image for SGD, as quantified by the root-mean-square error (RMSE). As such, this value will be taken as the optimal value for SGD-based USCT image reconstruction and will be employed in all future comparisons with the results obtained by the use of the RDA method.

A similar methodology was employed to choose the optimal step size for the SGD method. Several constant step sizes were compared with the use of a line search method. As seen in Fig. 4, when a constant step size is too large, the optimization algorithm will diverge. However, when the step size is small, the convergence of the optimization algorithm will be slow. Use of a line search method can provide fast convergence but, as mentioned above, can result in reduced image quality. Since the use of a line search method introduces an additional computational cost, the convergence of these approaches are given both in terms of iteration number and the number of times the wave equation must be solved, referred to here as wave solver runs. Every step size considered as part of the line search will add one additional wave solver run. However, even when this additional computational effort is accounted for, the use of a line search can still produce a more accurate




Fig. 4. Plot of RMSE versus (a) number of iterations and (b) number of wave equation solver runs for images reconstructed by the use of SGD with a line search and with several constant step size values for a regularization parameter value of  $5 \times 10^{-4}$ .

reconstructed image for a given level of computational effort than use of a constant step size (at least, up to some threshold level of total computational effort). In addition, it removes the need to wisely choose the step size, a task that is often accomplished through trial and error. From Fig. 4, it can be seen from the constant step size results that a step size of 0.1 produces the fastest convergence rate while still resulting in an accurate reconstructed image.

In Section II, it was suggested that the use of a line search method may have a negative impact on the obtained solution for SGD. This is demonstrated in Fig. 4. Here, it is seen that the line search method results in oscillations in the RMSE of the reconstructed image, while the use of a constant step size produces a smoother convergence curve with fewer jumps. Also, note that the final RMSE is lower for the constant step size method ( $\text{RMSE} = 1.42 \times 10^{-3}$ ) than for the line search method ( $\text{RMSE} = 1.73 \times 10^{-3}$ ).

### C. Images Reconstructed by RDA

The optimal step size (or, equivalently, the value of  $\gamma$  in line 9 of Algorithm 2) and regularization parameter value for the RDA method will be determined in the same manner as employed for the SGD method. First, the regularization parameter value that resulted in the most accurate reconstructed image was determined. Example images reconstructed by RDA for several regularization parameter values are shown in Fig. 5. From Fig. 6(a), it can be seen that a regularization parameter value of  $1 \times 10^{-4}$  results in the most accurate reconstructed image. This is smaller than the value obtained for SGD. From Fig. 6(b), the optimal step size value is 0.1, the same value obtained for SGD.

The weighted RDA method can be used to accelerate the convergence of the RDA method. As was done for the unweighted implementation, images were reconstructed for several regularization parameter values (see Fig. 7). The impact of regularization appears unchanged by the weighting strategy. Once again, a regularization parameter value of  $1 \times 10^{-4}$  results in the smallest RMSE. While the ultimate image obtained after many iterations is largely unchanged by the weighting strategy, reconstructed images obtained at early iterations can be greatly improved. As seen in Fig. 8, the accuracy of the reconstructed images after 20, 50, and even 100 iterations is improved by the use of the weighted




Fig. 5. Images reconstructed by the unweighted RDA method with a fixed step size of 0.1 for the regularization parameter values of (a)  $1 \times 10^{-5}$ , (b)  $5 \times 10^{-5}$ , (c)  $1 \times 10^{-4}$ , and (d)  $5 \times 10^{-4}$ , shown after 300 iterations. All images are shown in a grayscale window of  $[1.47, 1.58] \text{ mm}/\mu\text{s}$ .




Fig. 6. Plot of RMSE versus the number of iterations for (a) several regularization parameter values and a fixed step size of 0.1 and (b) several constant step size values and a fixed regularization parameter value of  $1 \times 10^{-4}$  for images reconstructed by the use of the unweighted RDA method.

RDA method. This is seen in both the RMSE of the reconstructed images and in the apparent visual quality of the images. This improvement is reflected in the profiles through the reconstructed images shown in Fig. 9. This improvement is maintained even when the convergence of the reconstruction methods is viewed in terms of the number of wave solver runs as opposed to the number of iterations (see Fig. 10). After approximately 250 wave solver runs (or 250 iterations for the unweighted method), the weighted and unweighted approaches produce images of similar accuracy.

*D. Comparison of Images Reconstructed by SGD and RDA*

The images produced by the use of the SGD and RDA methods were compared directly. Images reconstructed by all four implementations are shown in Fig. 11: 1) SGD with a constant step size; 2) unweighted RDA; 3) SGD with a line search; and 4) weighted RDA. As indicated by the RMSEs noted in the bottom-left of each image, the initial convergence rates of SGD with a line search and the weighted RDA method are much faster than that of either SGD with a constant step size or the unweighted RDA method. However, the accuracies of the reconstructed images at later iterations are superior for the two RDA methods compared with those obtained for the SGD-based methods. In fact, the accuracy of the image reconstructed by the weighted RDA method is better than that obtained by SGD with a constant step size. This demonstrates that the weighted RDA method can provide both

fast convergence and more accurate images than was possible using the SGD method.

The improved accuracy of the weighted RDA method compared with that of SGD with a line search is reflected in the profiles through the reconstructed images (see Fig. 12). The profile obtained by the use of SGD is notably noisier than that obtained by the use of RDA. This suggests that the RDA method may be more effective in mitigating noise than SGD. This potential benefit will be considered more closely through the use of a bias–variance analysis, as detailed in Section IV-E.

The plots of the convergence rates, as shown in Fig. 13, further confirm the benefits provided by the RDA method. SGD with a line search has a fast initial convergence, but results in a less accurate final image. From this plot, it is also clear that the estimates of the object provided by SGD with a line search also exhibit a high variance, even at later iterations. This is likely due to the fact that the line search only evaluates the cost function for a single realization of the encoding vector. As a result, the line search will tend to choose a larger step size that effectively minimizes the cost function evaluated for that encoding vector, but which increases the cost function when all encoding vectors or a large number of encoding vectors are considered. This behavior is not seen for the weighted RDA method. Since for the RDA method the search direction is given by a weighted average of the gradient estimates for all past encoding vector realizations, it does not overfit the cost function evaluated for a single realization of the encoding vector. This is true even though the weight at a given iteration is chosen only by evaluating the cost function for a single realization. Thus, the high variance of the object estimates is eliminated, while the computational cost of selecting a weight for the RDA method is the same as that of performing the line search for SGD.

*E. Bias–Variance Analysis*

The investigations with a numerical breast phantom, described above, suggested that the RDA method could provide more effective regularization than SGD. However, care must be exhibited when evaluating this claim. Stronger regularization does not mean better image quality. It is not enough to compare two different reconstruction methods with the




Fig. 7. Images reconstructed by the weighted RDA method for the regularization parameter values of (a)  $1 \times 10^{-5}$ , (b)  $5 \times 10^{-5}$ , (c)  $1 \times 10^{-4}$ , and (d)  $5 \times 10^{-4}$ , shown after 300 iterations. All images are shown in a grayscale window of  $[1.47, 1.58] \text{ mm}/\mu\text{s}$ .




Fig. 8. Images reconstructed by the use of the unweighted dual averaging method with a fixed step size of 0.1 after (a) 20, (b) 50, (c) 100, and (d) 250 iterations. Images reconstructed by the use of the weighted dual averaging method after (e) 20, (f) 50, (g) 100, and (h) 250 iterations. All results are shown for a regularization parameter value of  $1 \times 10^{-4}$  and in a grayscale window of  $[1.47, 1.58] \text{ mm}/\mu\text{s}$ . The RMSE for each reconstructed image is displayed in the bottom-left of each subfigure.

same regularization parameter value. While one may appear to produce a superior image, the other may produce a comparable image when another regularization value is employed. Thus, it is necessary to consider a range of regularization parameter values when comparing any two methods. Furthermore, image quality is most properly evaluated through task-based measures of image quality [37]. However, such studies are substantial undertaking and are outside the scope of this paper. Instead, here, we use bias–variance curves as a proxy for this more complete assessment.

Bias–variance curves depict the inherent tradeoff between noise mitigation and close agreement with the measured data. As described above, an estimate of the sound speed is obtained

by solving a minimization problem consisting of two terms, the data fidelity term and the regularization term. The relative weight of these terms is controlled by varying a scalar regularization parameter. Noise can be more severely suppressed by increasing the value of the regularization parameter, but this can result in reduced resolution or other forms of bias.

The bias–variance curves for SGD with a constant step size and the unweighted RDA method are shown in Fig. 14. The curves are generated by reconstructing a collection of images across a range of regularization parameter values. As seen in Fig. 14, the RDA method consistently produces lower variance images (less noisy) for a given level of bias. This difference



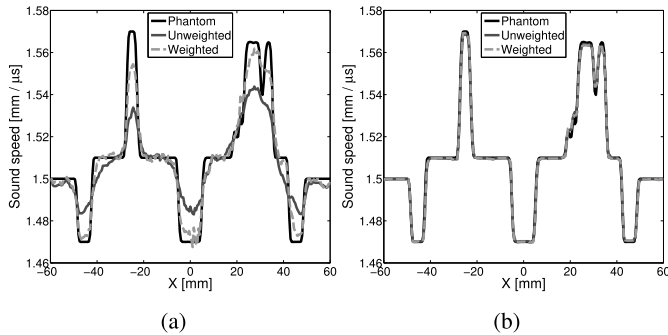


Fig. 9. (a) Profiles through  $y = -6.5$  mm for reconstructed images obtained by the use of the weighted RDA method and the unweighted RDA method with a fixed step size of 0.1, shown after 20 iterations. (b) Profiles through  $y = -6.5$  mm for reconstructed images obtained by the use of the weighted RDA method and the unweighted RDA method with a fixed step size of 0.1, shown after 250 iterations.




Fig. 10. Plot of RMSE versus (a) number of iterations and (b) number of wave solver runs for the weighted and unweighted RDA methods.

is seen in the reconstructed images. In Fig. 15, reconstructed images corresponding to the same bias level are shown. The image reconstructed by the use of SGD with a constant step size is notably noisier than the image obtained by the use of the unweighted RDA method.

## V. EXPERIMENTAL VALIDATION

### A. Methods

Clinical data were acquired previously by the use of the SoftVue USCT scanner [49]. The system consisted of a ring-shaped array with a radius of 110 mm, containing 2048 transducers. The transducers had a central frequency of 2.75 MHz with a pitch of 0.34 mm. Each element was elevationally focused to isolate a 3-mm-thick slice of the object. See [17], [49] for additional information regarding the measurement system and clinical studies.

Every other transducer element served as an emitter. The resulting pressure wave was then measured by the same set of 1024 transducers. The pressure was recorded with a sampling rate of 12 MHz for 2112 time points, corresponding to approximately 176  $\mu$ s. This measurement process was repeated with and without the object. Forty-eight transducers were identified as bad channels following manual inspection. The data from these channels were discarded, resulting in measurements from 976 transducers. The pressure data were upsampled to a sampling rate of 20 MHz by the use of linear interpolation in order to avoid the introduction of numerical

errors by the numerical wave solver [40]. The number of samples in each time trace was 3500. A Butterworth bandpass filter with the cutoff frequencies of 0.5 and 1.0 MHz was applied to each signal. The shape of the excitation pulse was estimated from the measured data without the object using the method described in [30].

An initial estimate of the object was reconstructed by the use of an adjoint state method (see Fig. 16) [53]. This estimate was used to generate a set of synthetic data. As detailed in [30], measurements near the emitter may not contribute positively to the reconstructed image due to mechanical cross talk, model mismatch, and measurement noise. The impact of these effects can be mitigated by replacing the measurements near the emitter with synthetic data. Unlike [30], here we substitute pressure data corresponding to an estimate of the object, provided by an adjoint state method [53], rather than a homogeneous medium. The 512 measurements from transducers opposite the emitter were kept. The others were replaced with the synthetic data.

The images were reconstructed by solving (3), where the operator  $\mathbf{H}(\mathbf{c})$  was calculated by the use of the second-order k-space pseudospectral wave equation solver, as described in Section IV-A4 [40]. The calculation domain was  $512 \times 512$  mm<sup>2</sup>, divided into a  $2560 \times 2560$  Cartesian grid with a spacing of 0.2 mm. The sound speed was updated within a circle of radius 105 mm. Reconstruction was performed on a platform consisting of dual hexa-core CPUs, 128 GB of RAM, and an NVIDIA Tesla K40 GPU. These reconstruction parameters are summarized in Table I.

While image quality is most objectively assessed using task-based methods of image quality [37], here, for reasons of expediency, the contrast-to-noise (CNR) ratio was employed as a proxy for the detectability of the tumor. The CNR of the reconstructed images was calculated by identifying three regions. The tumor was segmented manually. Regions of similar size corresponding to the parenchymal tissue and the water bath were also identified. The contrast was calculated based on the tumor and parenchymal tissue regions. The noise, however, was calculated based on the water bath to avoid misattributing any real variations within the parenchymal tissue to noise. The CNR was calculated as

$$\text{CNR} = \frac{\bar{c}_t - \bar{c}_p}{\sigma_n} \quad (23)$$

where  $\bar{c}_t$  is the average sound speed of the tumor,  $\bar{c}_p$  is the average sound speed over a comparably sized region of the parenchymal tissue, and  $\sigma_n$  is the standard deviation over a comparably sized region of the water bath.

### B. Clinical Results

As seen in Fig. 17, the weighted RDA method consistently produces reconstructed images with higher CNRs than SGD with a constant step size, as indicated by the CNR values that label each image. This is shown across a range of regularization parameter values. Further, the maximum CNR obtained by SGD is lower even when the regularization parameter value is optimized. This improvement in the CNR is likely due to the favorable noise mitigation properties of the RDA method observed in the computer simulation studies. While not shown,




Fig. 11. Images reconstructed by the use of SGD with a constant step size of 0.1 and a regularization parameter value of  $5 \times 10^{-4}$  after (a) 20, (b) 50, (c) 100, and (d) 250 iterations. Images reconstructed by the use of unweighted RDA with a fixed step size of 0.1 and a regularization parameter value of  $1 \times 10^{-4}$  after (e) 20, (f) 50, (g) 100, and (h) 250 iterations. Images reconstructed by the use of SGD with a line search and a regularization parameter value of  $5 \times 10^{-4}$  after (i) 20, (j) 50, (k) 100, and (l) 250 iterations. Images reconstructed by the use of weighted RDA with a regularization parameter value of  $1 \times 10^{-4}$  after (m) 20, (n) 50, (o) 100, and (p) 250 iterations. All images are shown in a grayscale window of  $[1.47, 1.58]$  mm/ $\mu$ s.

the CNRs for both methods do not continue to increase beyond 100 iterations. The CNR can serve as a proxy of detectability in cases where task-based measures of image quality cannot be performed [37]. While the CNRs of all the images shown

in Fig. 17 are quite high, the improvement in CNR could be more impactful for small or low-contrast tumors.

The reconstructed images as a function of iteration number are shown in Fig. 18. Since a nonconstant initial guess




Fig. 12. Profiles through  $y = -6.5$  mm for images reconstructed by the use of SGD with a line search and weighted RDA.

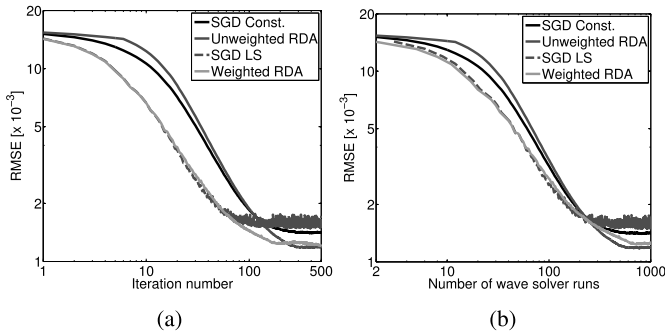


Fig. 13. Plot of RMSE versus (a) number of iterations and (b) number of wave solver runs for SGD with a line search, SGD with a constant step size of 0.1, unweighted RDA with a step size of 0.1, and weighted RDA.




Fig. 14. Bias-variance curve for SGD with constant step size and the unweighted RDA method. The corresponding regularization parameter values are given by each point.

was provided, the differences in the convergence rates of SGD with a constant step size and the weighted RDA method are less pronounced. However, a good initial guess is needed to avoid local minima since the data fidelity term is nonconvex. Still, the weighted RDA method produces a higher CNR at each iteration. The difference between the CNRs of the two methods continues to grow over the first 50 iterations. This gap is eventually decreased at later iterations. This suggests that the weighted RDA method is




Fig. 15. Example reconstructed images from the bias-variance analysis. (a) Image reconstructed by SGD with a regularization parameter value of  $5 \times 10^{-5}$ . (b) Image reconstructed by RDA with a regularization parameter value of  $1 \times 10^{-4}$ . The two images have approximately the same bias. Both images are shown in their full dynamic ranges. The sound speed values are given in mm/ $\mu$ s.




Fig. 16. Initial estimate of the object reconstructed by the use of an adjoint state method described in [53].

able to provide some initial improvement in the convergence rate. This is consistent with the computer simulation studies.

As discussed previously, the RDA method allows natural incorporation of nonsmooth penalties. This may allow the optimization problem be designed more optimally for a given image reconstruction task. While the determination of an optimal choice of regularization function (let alone the design of the entire optimization problem) is outside the scope of this paper, in Fig. 19, we show the results corresponding to an alternative nonsmooth penalty in order to emphasize the flexibility of this approach. The regularization function was chosen to be

$$\mathcal{R}(\mathbf{c}) = \|\Phi \mathbf{c}\|_1 \quad (24)$$

where  $\Phi$  is the 2-D wavelet transform of the object and the mother wavelet was the 12-tap Daubechies wavelet [54]. The wavelet transform was computed by the use of the GNU Scientific Library [55]. Images reconstructed with several regularization parameter values are shown.

## VI. CONCLUSION

WISE can produce high-resolution sound speed images without the computational burden of other time-domain waveform inversion approaches. Estimates of the sound speed




Fig. 17. Images reconstructed by the use of SGD with a constant step size of  $2.5 \times 10^5$  and regularization parameter values of (a)  $1 \times 10^{-10}$ , (b)  $3 \times 10^{-10}$ , (c)  $1 \times 10^{-9}$ , and (d)  $3 \times 10^{-9}$ . Images reconstructed by the use of the weighted RDA method with regularization parameter values of (e)  $1 \times 10^{-10}$ , (f)  $3 \times 10^{-10}$ , (g)  $1 \times 10^{-9}$ , and (h)  $3 \times 10^{-9}$ . Images are shown after 100 iterations and in a grayscale window of  $[1.38, 1.60]$  mm/ $\mu$ s.




Fig. 18. Images reconstructed by the use of SGD with a constant step size of  $2.5 \times 10^5$  after (a) 5, (b) 20, (c) 50, and (d) 100 iterations with a regularization parameter value of  $1 \times 10^{-9}$ . Images reconstructed by the use of weighted RDA after (e) 5, (f) 20, (g) 50, and (h) 100 iterations with a regularization parameter value of  $1 \times 10^{-9}$ . All images are shown in a grayscale window of  $[1.38, 1.60]$  mm/ $\mu$ s.

distribution can be obtained using this method by minimizing an objective function consisting of a data fidelity term and a regularization term. While this optimization problem can be solved using SGD, use of a structured optimization method, such as the RDA method, provides several advantages. First, it exploits knowledge of the structure of the cost function to

separate the stochastic data fidelity term from the deterministic regularization term. This appears to result in more effective regularization. In the case of the TV seminorm, noise is more effectively reduced while preserving the accuracy and contrast of the reconstructed images. Second, it does not assume that all terms in the objective function are differentiable, allowing




Fig. 19. Images reconstructed by the use of the weighted RDA method with a wavelet-based penalty and regularization parameter values of (a)  $3 \times 10^{-10}$ , (b)  $1 \times 10^{-9}$ , (c)  $3 \times 10^{-9}$ , and (d)  $1 \times 10^{-8}$ . Images are shown after 100 iterations and in a grayscale window of [1.38, 1.60] mm/ $\mu$ s.

natural incorporation of nonsmooth penalties, such as the total variation seminorm. Third, it exploits information from past iterations when determining the search direction. This allows the method to employ a line search while avoiding overfitting a particular realization of the encoding vector. This allows a fast initial convergence rate without sacrificing image quality. This was demonstrated through computer simulation studies involving a numerical breast phantom, generation of a bias-variance curve, and experimental studies involving clinical data.

Some reconstruction parameters were not strictly optimized, particularly for the clinical results. Similar results to those presented could potentially be obtained with coarser temporal or spatial sampling rates. In addition, the number of measurements kept as part of the data filling strategy may not be optimal. The optimal number of measurements will depend on the object and the degree of model mismatch and measurement noise. Further tuning of these parameters could lead to improved performance.

Opportunities for further improvement exist. The acoustic model employed in the calculation of the data fidelity term ignores a number of important factors that could lead to artifacts in the reconstructed images. In particular, the model ignores acoustic attenuation and dispersion and out-of-plane scattering. Since the assumed imaging model is 2-D, scattering out of the plane defined by the transducer ring array is not modeled. It also treats the transducers as ideal point detectors and emitters. Additional investigation of the numerical properties of this approach remains a topic for future study. As noted previously, the frequency content of the excitation pulse and the strength of the acoustic heterogeneities have a sizable impact on the reconstructed images [30]. A comparison with other image reconstruction methods is also needed [4], [14], [56], [57].

#### APPENDIX

##### LINE SEARCH FOR WEIGHTED RDA METHOD

The weights for the weighted RDA method were chosen via the line search method described in Algorithm 3. Other line search methods may produce similar, or even superior,

---

##### Algorithm 3 Line Search for RDA Method

---

**Input:**  $\mathbf{c}_0, A_{k-1}, \mathbf{w}_k, \mathbf{G}_k, \bar{\mathbf{G}}_{k-1}, f(\mathbf{c}_k, \mathbf{w}_k), \lambda, \alpha_{max}$

**Output:**  $\alpha_k$  {Weight for  $k$ -th iteration.}

```

1:  $\tilde{\alpha} \leftarrow \alpha_{max}$  { $\alpha_{max}$  is the initial guess for the weight.}
2:  $found \leftarrow \text{false}$ 
3: while not found do
4:    $\tilde{A} \leftarrow A_{k-1} + \tilde{\alpha}$ 
5:    $\tilde{\mathbf{G}} \leftarrow \left(1 - \frac{\tilde{\alpha}}{A}\right) \bar{\mathbf{G}}_{k-1} + \frac{\tilde{\alpha}}{A} \mathbf{G}_k$ 
6:    $\tilde{\mu} \leftarrow \gamma \tilde{A}$  {Should be consistent with Alg. 2.}
7:    $\tilde{\mathbf{c}} \leftarrow \text{prox}_{\lambda \tilde{\mu} \tilde{\mathcal{R}}}(\mathbf{c}_0 - \tilde{\mu} \tilde{\mathbf{G}})$ 
8:   if  $f(\tilde{\mathbf{c}}, \mathbf{w}_k) + \lambda \mathcal{R}(\tilde{\mathbf{c}}) < f(\mathbf{c}_k, \mathbf{w}_k) + \lambda \mathcal{R}(\mathbf{c}_k)$  then
9:      $found \leftarrow \text{true}$ 
10:  else
11:     $\tilde{\alpha} \leftarrow \tilde{\alpha}/2$ 
12:  end if
13: end while
14:  $\alpha_k \leftarrow \tilde{\alpha}$ 

```

---

results. Each weight value considered for a given iteration requires  $f(\mathbf{c}, \mathbf{w})$  to be evaluated one additional time. Since  $f(\mathbf{c}, \mathbf{w})$  is evaluated for only one realization of the encoding vector, this requires only one additional wave solver run. This is the same computational cost as for the line search procedure employed for SGD. The goal of the line search procedure is to find weights that improve the convergence rate of the algorithm while minimizing the computational cost needed to select those weights. Thus, it is neither practical nor advisable to choose weights that most minimize the cost function at each iteration. Here, we decrease the weight by a factor of two if the stopping criterion for the line search is not satisfied. This factor can be adjusted to perform the line search more coarsely (larger factor) or more finely (smaller factor).

##### ACKNOWLEDGMENT

The authors would like to thank F. Anis for her assistance in reconstructing an image from the clinical breast data by the use of the adjoint state method described in [53].

## 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. V. Rüter, G. Göbel, L. Berger, M. Zapf, and H. Gemmeke, "Realization of an optimized 3D USCT," *Proc. SPIE*, vol. 7968, p. 796805, Mar. 2011.
- [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, "Non-linear inverse scattering: High resolution quantitative breast tissue tomography," *J. Acoust. Soc. Amer.*, vol. 131, no. 5, pp. 3802–3813, 2012.
- [5] J. F. Greenleaf, S. A. Johnson, and R. C. Bahn, "Quantitative cross-sectional imaging of ultrasound parameters," in *Proc. Ultrason. Symp.*, Oct. 1977, pp. 989–995.
- [6] 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, Feb. 2007.
- [7] T. M. Kolb, J. Lichy, and J. H. Newhouse, "Comparison of the performance of screening mammography, physical examination, and breast us and evaluation of factors that influence them: An analysis of 27,825 patient evaluations," *Radiology*, vol. 225, no. 1, pp. 165–175, Oct. 2002.
- [8] 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*, vol. 214, no. 4525, pp. 1141–1143, Dec. 1981.
- [9] J. Zalev *et al.*, "Clinical feasibility study of combined optoacoustic and ultrasonic imaging modality providing coregistered functional and anatomical maps of breast tumors," *Proc. SPIE*, vol. 8223, p. 82230A, Feb. 2012.
- [10] 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.*, vol. 38, no. 16, pp. 3140–3143, 2013.
- [11] M. P. André, H. S. Janée, P. J. Martin, G. P. Otto, B. A. Spivey, and D. A. Palmer, "High-speed data acquisition in a diffraction tomography system employing large-scale toroidal arrays," *Int. J. Imag. Syst. Technol.*, vol. 8, no. 1, pp. 137–147, 1997.
- [12] S. A. Johnson *et al.*, "Apparatus and method for imaging with wavefields using inverse scattering techniques," U.S. Patent 6005 916, Dec. 21, 1999.
- [13] S. Manohar, R. G. H. Willeminck, F. van der Heijden, C. H. Slump, and T. G. van Leeuwen, "Concomitant speed-of-sound tomography in photoacoustic imaging," *Appl. Phys. Lett.*, vol. 91, no. 13, p. 131911, 2007.
- [14] 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.
- [15] N. Duric *et al.*, "Breast imaging with SoftVue: Initial clinical evaluation," *Proc. SPIE*, vol. 9040, p. 90400V, Mar. 2014.
- [16] G. Y. S. Sandhu, C. Li, O. Roy, S. Schmidt, and N. Duric, "High-resolution quantitative whole-breast ultrasound: In vivo application using frequency-domain waveform tomography," *Proc. SPIE*, vol. 9419, p. 94190D, Mar. 2015.
- [17] N. Duric *et al.*, "Clinical breast imaging with ultrasound tomography: A description of the SoftVue system," *J. Acoust. Soc. Amer.*, vol. 135, no. 4, p. 2155, Apr. 2014.
- [18] A. Hormati, I. Jovanović, O. Roy, and M. Vetterli, "Robust ultrasound travel-time tomography using the bent ray model," *Proc. SPIE*, vol. 7629, p. 76290I, Mar. 2010.
- [19] Y. Quan and L. Huang, "Sound-speed tomography using first-arrival transmission ultrasound for a ring array," *Proc. SPIE*, vol. 6513, p. 651306, Mar. 2007.
- [20] P. Huthwaite and F. Simonetti, "High-resolution imaging without iteration: A fast and robust method for breast ultrasound tomography," *J. Acoust. Soc. Amer.*, vol. 130, no. 3, pp. 1721–1734, 2011.
- [21] S. J. Norton and M. Linzer, "Correcting for ray refraction in velocity and attenuation tomography: A perturbation approach," *Ultrason. Imag.*, vol. 4, no. 3, pp. 201–233, Jul. 1982.
- [22] R. Lavarello and M. Oelze, "Density imaging using a multiple-frequency DBIM approach," *IEEE Trans. Ultrason., Ferroelect., Freq. Control*, vol. 57, no. 11, pp. 2471–2479, Nov. 2010.
- [23] Z. Zhang, L. Huang, and Y. Lin, "Efficient implementation of ultrasound waveform tomography using source encoding," *Proc. SPIE*, vol. 8320, p. 832003, Feb. 2012.
- [24] R. G. Pratt, L. Huang, N. Duric, and P. Littrup, "Sound-speed and attenuation imaging of breast tissue using waveform tomography of transmission ultrasound data," *Proc. SPIE*, vol. 6510, p. 65104S, Mar. 2007.
- [25] R. J. Lavarello and M. L. Oelze, "Density imaging using inverse scattering," *J. Acoust. Soc. Amer.*, vol. 125, no. 2, p. 793, 2009.
- [26] A. J. Hesford and W. C. Chew, "Fast inverse scattering solutions using the distorted Born iterative method and the multilevel fast multipole algorithm," *J. Acoust. Soc. Amer.*, vol. 128, no. 2, p. 679, 2010.
- [27] M. Pérez-Liva, J. L. Herraiz, J. M. Udías, B. T. Cox, and B. E. Treeby, "Full-wave attenuation reconstruction in the time domain for ultrasound computed tomography," in *Proc. IEEE 13th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2016, pp. 710–713.
- [28] W. A. Mulder and R.-E. Plessix, "Time-versus frequency-domain modelling of seismic wave propagation," in *Proc. 64th EAGE Conf. Exhibit.*, May 2002, pp. 1–4.
- [29] D. L. Colton and R. Kress, *Inverse Acoustic and Electromagnetic Scattering Theory* (Applied Mathematical Sciences), vol. 93, 2nd ed. New York, NY, USA: Springer, 1998.
- [30] K. Wang, T. P. 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., Ferroelect., Freq. Control*, vol. 62, no. 3, pp. 475–494, Mar. 2015.
- [31] J. R. Krebs *et al.*, "Fast full-wavefield seismic inversion using encoded sources," *Geophysics*, vol. 74, no. 6, pp. WCC177–WCC188, Nov. 2009.
- [32] L. Xiao, "Dual averaging methods for regularized stochastic learning and online optimization," *J. Mach. Learn. Res.*, vol. 11, pp. 2543–2596, Oct. 2010.
- [33] Y. Nesterov, "Primal-dual subgradient methods for convex problems," *Math. Program.*, vol. 120, no. 1, pp. 221–259, Aug. 2009.
- [34] J.-L. Starck, F. Murtagh, and J. M. Fadili, *Sparse Image Signal Processing: Wavelets, Curvelets, Morphological Diversity*. Cambridge, U.K.: Cambridge Univ. Press, 2010.
- [35] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Process. Mag.*, vol. 24, no. 4, pp. 118–121, Jul. 2007.
- [36] J. Bian *et al.*, "Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam CT," *Phys. Med. Biol.*, vol. 55, no. 22, pp. 6575–6599, Nov. 2010.
- [37] H. H. Barrett and K. J. Myers, *Foundations of Image Science* (Wiley Series in Pure and Applied Optics). Hoboken, NJ, USA: Wiley, 2004.
- [38] N. Duric *et al.*, "Ultrasound tomography systems for medical imaging," in *Emerging Imaging Technologies in Medicine*. Boca Raton, FL, USA: CRC Press, 2010, pp. 167–182.
- [39] M. Tabei, T. D. Mast, and R. C. Waag, "A  $k$ -space method for coupled first-order acoustic propagation equations," *J. Acoust. Soc. Amer.*, vol. 111, no. 1, pp. 53–63, 2002.
- [40] T. D. Mast, L. P. Souriau, D. L. D. Liu, M. Tabei, A. I. Nachman, and R. C. Waag, "A  $k$ -space method for large-scale models of wave propagation in tissue," *IEEE Trans. Ultrason., Ferroelect., Freq. Control*, vol. 48, no. 2, pp. 341–354, Mar. 2001.
- [41] 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, p. 021314, Mar. 2010.
- [42] T. van Leeuwen, A. Y. Aravkin, and F. J. Herrmann, "Seismic waveform inversion by stochastic optimization," *Int. J. Geophys.*, vol. 2011, pp. 1–18, 2011.
- [43] S. Sra, S. Nowozin, and S. J. Wright, Eds., *Optimization for Machine Learning*. Cambridge, MA, USA: MIT Press, 2012.
- [44] N. N. Schraudolph, J. Yu, and S. Günter, "A stochastic quasi-Newton method for online convex optimization," in *Proc. 11th Int. Conf. Artif. Intell. Statist.*, vol. 2, 2007, pp. 436–443.
- [45] N. Parikh and S. Boyd, *Proximal Algorithms* (Foundations and Trends in Optimization). Norwell, MA, USA: Now Publishers, 2013.
- [46] 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* (Radon Series on Computational and Applied Mathematics), M. Fornasier, Ed. de Gruyter, 2010, pp. 263–340.
- [47] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2419–2434, Nov. 2009.
- [48] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imag. Vis.*, vol. 20, no. 1, pp. 89–97, 2004.

- [49] N. Duric *et al.*, "Breast imaging with ultrasound tomography: Initial results with SoftVue," in *Proc. IEEE Int. Ultrason. Symp.*, Jul. 2013, pp. 382–385.
- [50] C. Li, N. Duric, and L. Huang, "Clinical breast imaging using sound-speed reconstructions of ultrasound tomography data," *Proc. SPIE*, vol. 6920, p. 692009, Mar. 2008.
- [51] D. L. Colton, *Inverse Acoustic and Electromagnetic Scattering Theory*, 3rd ed. New York, NY, USA: Springer, 2013.
- [52] "CUDA C programming guide," NVIDIA Corporation, Santa Clara, CA, USA, Tech. Rep. PG-02829-001 v7.5, Sep. 2015.
- [53] F. Anis *et al.*, "Investigation of the adjoint-state method for ultrasound computed tomography: A numerical and experimental study," *Proc. SPIE*, vol. 8943, p. 894337, Mar. 2014.
- [54] S. G. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, 3rd ed. Boston, MA, USA: Elsevier, 2009.
- [55] M. Galassi *et al.* *GNU Scientific Library Reference Manual*, accessed on Mar. 22, 2016. [Online]. Available: <http://www.gnu.org/software/gsl/>
- [56] H. Abdullah and A. Louis, "The approximate inverse for solving an inverse scattering problem for acoustic waves in an inhomogeneous medium," *Inverse Problems*, vol. 15, no. 5, pp. 1213–1229, 1999.
- [57] H. Gemmeke *et al.*, "Wave equation based transmission tomography," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, Sep. 2016, pp. 1–4.
- Thomas P. Matthews**, photograph and biography not available at the time of publication.
- Kun Wang**, photograph and biography not available at the time of publication.
- Cuiping Li**, photograph and biography not available at the time of publication.
- Neb Duric**, photograph and biography not available at the time of publication.
- Mark A. Anastasio**, photograph and biography not available at the time of publication.