

RANDOMIZED APPROACH TO NONLINEAR INVERSION COMBINING RANDOM AND OPTIMIZED SIMULTANEOUS SOURCES AND DETECTORS*

SELIN S. ASLAN[†], ERIC DE STURLER[†], AND MISHA E. KILMER[‡]

Abstract. In partial differential equation-based (PDE-based) inverse problems with many measurements, many large-scale discretized PDEs must be solved for each evaluation of the misfit or objective function. In the nonlinear case, evaluating the Jacobian requires solving an additional set of systems. This leads to a tremendous computational cost, and this is by far the dominant cost for these problems. Several authors have proposed randomization and stochastic programming techniques to drastically reduce the number of system solves by estimating the objective function using only a few appropriately chosen random linear combinations of the sources. While some have reported good solution quality at a greatly reduced cost, for our problem of interest, diffuse optical tomography, the approach often does not lead to sufficiently accurate solutions. We propose two improvements. First, to efficiently exploit Newton-type methods, we modify the stochastic estimates to include random linear combinations of detectors, drastically reducing the number of adjoint solves. Second, after solving to a modest tolerance, we compute a few simultaneous sources and detectors that maximize the Frobenius norm of the sampled Jacobian to improve the rate of convergence and obtain more accurate solutions. We complement these optimized simultaneous sources and detectors by random simultaneous sources and detectors constrained to a complementary subspace. Our approach leads to solutions of the same quality as obtained using all sources and detectors but at a greatly reduced computational cost, as the number of large-scale linear systems to be solved is significantly reduced.

Key words. DOT, PaLS, stochastic programming, randomization, inverse problems, optimization

AMS subject classifications. 65F22, 65N21, 65N22, 65M32, 62L20, 90C15

DOI. 10.1137/17M1143319

1. Introduction. The solution of nonlinear inverse problems requires solving many large-scale discretized PDEs in the evaluation of the forward problem. In parameterized inverse problems, we can compute the response of the system for a particular input by numerically solving the PDE. The forward model used in this paper (see section 2) is already regularized using the parametric level set (PaLS) approach [1], and we focus on efficiently solving the nonlinear least squares problem

$$(1.1) \quad \min_{\mathbf{p}} f(\mathbf{p}) := \min_{\mathbf{p}} \frac{1}{2} \|\mathbf{M}(\mathbf{p}) - \mathbf{d}\|_2^2,$$

where $\mathbf{M}(\mathbf{p})$ is the vector of computed measurements given by the forward model for the parameter vector \mathbf{p} , and \mathbf{d} is the vector of measured data at the detectors.

Each evaluation of $f(\mathbf{p})$ requires the solution of the PDE for all inputs and each frequency. Moreover, to efficiently compute derivative information using the costate approach [29], we also need to solve linear systems with the adjoint for each detector

*Submitted to the journal's Computational Methods in Science and Engineering section August 14, 2017; accepted for publication (in revised form) November 12, 2018; published electronically March 26, 2019.

<http://www.siam.org/journals/sisc/41-2/M114331.html>

Funding: This material is based upon work supported by the National Science Foundation under grants NSF DMS 1217156 and 1217161, and NIH R01-CA154774.

[†]Department of Mathematics, Virginia Tech, Blacksburg, VA 24061-0123 (selin@vt.edu, sturler@vt.edu).

[‡]Department of Mathematics, Tufts University, Medford, MA 02155 (misha.kilmer@tufts.edu).

and each frequency. This leads to an enormous computational bottleneck, as rapid advances in technology allow for large numbers of sources and detectors. Multiply this by the number of frequencies, and the number of linear systems to solve in the solution of (1.1) is very large indeed. For the main application discussed in this paper, diffuse optical tomography (DOT), the number of sources and the number of detectors may each be a thousand or more; the number of frequencies used is typically modest (less than ten) [9].

To solve the minimization problem (1.1), we use the trust region algorithm with Gauss–Newton regularized model solution (TREGS) [10] that has proven very effective for parameterized problems of the type we consider in this paper. In [9], we use reduced order models to approximate both the function evaluation as well as its derivatives to compute regularized Gauss–Newton steps in TREGS. Here, we explore an alternative approach, following the work by Haber, Chung, and Herrmann [12]. The main idea in their paper was to drastically reduce the number of systems to be solved by exploiting randomization [12], posing the problem as a stochastic optimization problem [26]. In their approach, the misfit or objective function is estimated using only a few random linear combinations of the sources, referred to as *random simultaneous sources*, that are kept fixed over many optimization steps. In [26], this approach is referred to as the sample average approximation (SAA) method.

The use of random simultaneous sources has been well studied in several papers; see [3, 22, 18, 25, 28] and the references therein. While replacing the original objective function by the stochastic optimization problem seems to work well for direct current resistivity and seismic tomography [12], we find that the approach does not lead to accurate recovery of the parameters for the DOT problem. Therefore, we propose *two innovations* to the use of random simultaneous sources.

First, we extend the idea of random simultaneous sources to the randomized treatment of the detector solves for efficiently computing the Jacobian in Newton-type methods. Second, we propose to combine random simultaneous sources and detectors with optimized simultaneous sources and detectors to best capture the sensitivity and hence obtain more accurate estimates of the dominant singular components of the Jacobian and the corresponding components of the gradient.

The first innovation drastically reduces the cost of Newton-type methods. In particular, we derive a stochastic optimization problem, analogous to randomized simultaneous sources, that allows us to reduce the number of adjoint solves for the detectors.

The second innovation avoids stagnation in the residual norm decrease of the stochastic optimization approach due to poor or less effective estimates of derivative information. This is typically more important closer to the solution than early in the optimization, and several authors have addressed this problem by dynamically varying the sample size in the stochastic algorithm or increasing it slowly; see, for example, [6, 25]. We propose an alternative method to improve the estimates of derivative information that keeps the sample size fixed (and small) for efficiency. Comparing the two approaches in detail is future work. For the DOT problem, using random simultaneous sources and detectors does provide moderately accurate parameter solution estimates at a drastically reduced number of linear system solves. Thus, in our new approach, we first solve with a fixed set of random simultaneous sources and detectors to an intermediate tolerance. After reaching this intermediate tolerance, we compute a few simultaneous sources and detectors that maximize the Frobenius norm of the sampled Jacobian (see section 3); we refer to these as *optimized simultaneous sources and detectors*. We complement these optimized sources and detectors by ran-

dom simultaneous sources and detectors constrained to a complementary subspace (see section 3). After this update, the optimization converges rapidly to a solution of the same quality as obtained using all sources and detectors. Our use of optimized simultaneous sources and detectors is based on two motivations. First, the regularized model problem solves in TREGS [10] focus on the directions corresponding to the large singular values of the Jacobian. Second, the directions corresponding to the large singular values are best informed by the data. More details follow at the end of section 2.

This paper is organized as follows. In section 2, we briefly review DOT, PaLS, and TREGS. In section 3, we introduce an alternative stochastic problem that includes random simultaneous detectors, to reduce the number of adjoint solves. In section 3.2, we introduce optimized simultaneous sources and detectors combined with random simultaneous sources and detectors constrained to a complementary subspace. In section 3.3, we provide an overall algorithm for our approach. We also give an outline of our implementation strategies. In section 4, we demonstrate the effectiveness of combining random and optimized simultaneous sources and detectors using a 2-dimensional (2D) and a 3-dimensional (3D) experiment. Finally, we draw some conclusions and discuss future work in section 5.

2. Background on DOT, PaLS, and TREGS. We assume that the region to be imaged is a rectangular prism with sources and detectors on the top and or the bottom. We consider the diffusion model for the photon flux $\eta(\mathbf{x})$ obtained by an input source $g(\mathbf{x})$ as in [2]. Let the diffusion (or the scattering) and the absorption coefficients be given by $D(\mathbf{x})$ and $\mu(\mathbf{x})$, respectively. Then, the mathematical model of the problem in the frequency domain is given by

(2.1)

$$\begin{aligned} -\nabla \cdot (D(\mathbf{x})\nabla\eta(\mathbf{x})) + \mu(\mathbf{x})\eta(\mathbf{x}) + \frac{i\omega}{\nu}\eta(\mathbf{x}) &= g(\mathbf{x}) \\ \text{for } \mathbf{x} = (x_1, x_2, x_3)^T \text{ and } -a < x_1 < a, \quad -b < x_2 < b, \quad 0 < x_3 < c, \\ \eta(\mathbf{x}) &= 0 \text{ if } 0 \leq x_3 \leq c \text{ and either } x_1 = \pm a \text{ or } x_2 = \pm b, \\ 0.25\eta(\mathbf{x}) + \frac{D(\mathbf{x})}{2} \frac{\partial\eta(\mathbf{x})}{\partial\xi} &= 0 \text{ for } x_3 = 0 \text{ or } x_3 = c, \end{aligned}$$

where ξ is the outward unit normal, ω is the frequency modulation of light, and ν is the speed of light in the medium.

Assuming that the diffusion coefficient is known (a common assumption for breast imaging), we use measurements and the forward model to recover the absorption coefficient of the medium, which can be used to distinguish healthy tissue from tumors [4]. Typical inversion methods would optimize for the desired physical quantity over a collection of grid points/voxels resulting in a parameter vector with at least $O(10^6)$ unknowns. Instead, we assume that the absorption field, $\mu(\mathbf{x})$, is expressible as $\mu(\mathbf{x}; \mathbf{p})$ with a modest number of (unknown) parameters, $\mathbf{p} = [p_1, p_2, \dots, p_{n_p}]^T$, where n_p is the number of parameters. We use the PaLS approach [1, 9] and parameterize the absorption $\mu(\mathbf{x}; \mathbf{p})$ as follows.

Let $\varphi: \mathbb{R}^+ \rightarrow \mathbb{R}$ be a smooth, compactly supported radial basis function (CSRBF),¹ γ be a positive, small, real number, and $\|\mathbf{x}\|^\dagger := \sqrt{\|\mathbf{x}\|_2^2 + \gamma^2}$ denote the (regularized) Euclidean norm. Then the PaLS function ϕ with a vector of unknown parameters \mathbf{p}

¹The CSRBF used here, $\varphi(r)$, is the Wendland function $\psi_{2,1}(r) = (1-r)^4(4r+1)$ [1, Table 1].




FIG. 2.1. (a) Surface and contour plot of a test anomaly on 100×100 mesh with 25 basis functions where the cutoff is at $c = 0.15$. (b) The PaLS function of the test anomaly on the left. If $\phi(\mathbf{x}, \mathbf{p}) \geq 0.15$, then \mathbf{x} is inside the anomaly (light) and if $\phi(\mathbf{x}, \mathbf{p}) < 0.15$, then \mathbf{x} is outside the anomaly (dark).

consisting of expansion coefficients α_j , dilation coefficients β_j , and center locations χ_j is defined as

$$(2.2) \quad \phi(\mathbf{x}, \mathbf{p}) := \sum_{j=1}^{m_0} \alpha_j \varphi(\|\beta_j(\mathbf{x} - \chi_j)\|^\dagger).$$

The PaLS approach uses an approximate Heaviside function $H_\epsilon(r)$, where r is a scalar, to create a differentiable, but sharp transition from anomaly to background. The absorption $\mu(\mathbf{x}, \mathbf{p})$ takes the value $\mu_{in}(\mathbf{x})$ if \mathbf{x} is inside the region and $\mu_{out}(\mathbf{x})$ if \mathbf{x} is outside the region,

$$(2.3) \quad \mu(\mathbf{x}, \mathbf{p}) = \mu_{in}(\mathbf{x})H_\epsilon(\phi(\mathbf{x}, \mathbf{p}) - c) + \mu_{out}(\mathbf{x})(1 - H_\epsilon(\phi(\mathbf{x}, \mathbf{p}) - c)),$$

where $c \in \mathbb{R}$ is a chosen cutoff parameter for the level set.

Figure 2.1 illustrates how PaLS represents the absorption field. Using PaLS, edges and complex boundaries can be captured with relatively few basis functions. Moreover, the PaLS representation with a modest number of basis functions regularizes the problem, hence, no further regularization is needed. Since there is no point in reducing the misfit below the (known or estimated) norm of the noise in the data, we stop the optimization when the objective function reaches this noise level. This is called the discrepancy principle [29]. For further discussion of the PaLS parameters for DOT, we refer the reader to [1, 9].

Let n_d , n_s , and n_ω denote the number of detectors, sources, and frequencies, respectively. The discretization of (2.1) leads to computed measurements, $\mathbf{m}_i(\omega_j, \mathbf{p}) \in \mathbb{C}^{n_d}$, for each source term, \mathbf{b}_i ,

$$(2.4) \quad \mathbf{m}_i(\omega_j, \mathbf{p}) = \mathbf{C}^T \left(\frac{i\omega_j}{\nu} \mathbf{E} + \mathbf{A}(\mathbf{p}) \right)^{-1} \mathbf{b}_i,$$

where the rows of \mathbf{C}^T correspond to the detectors.² $\mathbf{A}(\mathbf{p})$ derives from a finite difference discretization of the diffusion and absorption terms in (2.1), and \mathbf{E} derives from

²In practice, we also split \mathbf{m}_i into its real and imaginary parts.

the frequency term in (2.1). \mathbf{E} is almost the identity except that it has zero rows for points on the boundary, $x_3 = 0, x_3 = c$, in (2.1); so, \mathbf{E} is singular.

For simplicity, we consider the nonlinear residual for a single frequency, $\omega_j = 0$. In vector form, the residual is defined as follows:

$$(2.5) \quad \mathbf{r}(\mathbf{p}) = \begin{bmatrix} \mathbf{r}_1(\mathbf{p}) \\ \vdots \\ \mathbf{r}_{n_s}(\mathbf{p}) \end{bmatrix} = \begin{bmatrix} \mathbf{m}_1(\mathbf{p}) - \mathbf{d}_1 \\ \vdots \\ \mathbf{m}_{n_s}(\mathbf{p}) - \mathbf{d}_{n_s} \end{bmatrix} = \begin{bmatrix} \mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \mathbf{b}_1 - \mathbf{d}_1 \\ \vdots \\ \mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \mathbf{b}_{n_s} - \mathbf{d}_{n_s} \end{bmatrix},$$

where $\mathbf{r}_i \in \mathbb{R}^{n_d}$, \mathbf{d}_i is the data vector with the measurements from the detectors corresponding to source \mathbf{b}_i , and the nonlinear least squares problem (1.1) becomes

$$(2.6) \quad \min_{\mathbf{p}} \frac{1}{2} \|\mathbf{r}(\mathbf{p})\|_2^2.$$

Let \mathbf{J} be the Jacobian of $\mathbf{r}(\mathbf{p})$,

$$(2.7) \quad \mathbf{J} = \frac{\partial \mathbf{r}(\mathbf{p})}{\partial \mathbf{p}} = \begin{bmatrix} \frac{\partial \mathbf{r}(\mathbf{p})}{\partial p_1} & \dots & \frac{\partial \mathbf{r}(\mathbf{p})}{\partial p_{n_p}} \end{bmatrix},$$

where the components of \mathbf{J} are given by the small vectors

$$(2.8) \quad \mathbf{J}_{jk}(\mathbf{p}) = \frac{\partial}{\partial p_k} (\mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \mathbf{b}_j) = -\mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \frac{\partial \mathbf{A}(\mathbf{p})}{\partial p_k} \mathbf{A}^{-1}(\mathbf{p}) \mathbf{b}_j \in \mathbb{R}^{n_d}.$$

Evaluating the objective function at a given \mathbf{p} requires solving $n_s \cdot n_\omega$ large linear systems. Once $\mathbf{r}(\mathbf{p})$ and $\mathbf{A}^{-1}(\mathbf{p}) \mathbf{b}_j$ are available, evaluating \mathbf{J} using the costate approach [29] requires solving an additional $n_d \cdot n_\omega$ adjoint systems for the detectors. As a result, standard optimization approaches require $O(10^3 - 10^4)$ large linear system solves at each optimization step. The size of a realistic linear system is at least $O(10^6)$. This leads to an enormous computational bottleneck, and new computational techniques are needed.

We use TREGS [10] to solve the nonlinear least squares problem (2.6). The TREGS algorithm combines a trust region method with a regularized minimization of the Gauss–Newton model [11]. The local Gauss–Newton model at the current parameter vector, \mathbf{p}_c , is given by

$$(2.9) \quad f(\mathbf{p}_c + \boldsymbol{\delta}) \approx m_{GN}(\mathbf{p}_c + \boldsymbol{\delta}) = \frac{1}{2} \mathbf{r}_c^T \mathbf{r}_c + \mathbf{r}_c^T \mathbf{J}_c \boldsymbol{\delta} + \frac{1}{2} \boldsymbol{\delta}^T \mathbf{J}_c^T \mathbf{J}_c \boldsymbol{\delta},$$

and its minimization is equivalent to the least squares problem

$$(2.10) \quad \min_{\boldsymbol{\delta}} \|\mathbf{J}_c \boldsymbol{\delta} + \mathbf{r}(\mathbf{p}_c)\|_2^2.$$

The TREGS algorithm favors updates corresponding to (1) the large singular values and (2) the left singular vectors with large components in \mathbf{r} as determined by a generalized cross validation-like criterion. Since the Jacobian tends to be ill-conditioned, the emphasis on large singular values leads to relatively small steps that provide relatively large reductions in the Gauss–Newton model (2.9). We refer the reader to [10] for more details of TREGS.

3. A randomized approach. We recast the nonlinear least squares problem as a stochastic optimization problem using randomization to drastically reduce the number of large linear system solves in (2.5) and (2.8). The columns of $\mathbf{B} = [\mathbf{b}_1 \cdots \mathbf{b}_{n_s}]$ are source terms, and we refer to any linear combination of these sources as a simultaneous source. Simultaneous random sources, $\mathbf{B}\mathbf{w}$, with $\mathbf{w} \in \mathbb{R}^{n_s}$ a random vector, have been used in several areas [3, 21, 22, 12]. In this section, we introduce the concept of *optimized simultaneous sources and detectors* to improve the rate of convergence of the optimization and the quality of the inverse solution.

3.1. A stochastic optimization approach. To recast (2.5)–(2.6) as a stochastic optimization problem, we first write the residual in matrix form. For a single frequency, we get

$$(3.1) \quad \mathbf{R}(\mathbf{p}) = [\mathbf{r}_1(\mathbf{p}) \quad \mathbf{r}_2(\mathbf{p}) \quad \cdots \quad \mathbf{r}_{n_s}(\mathbf{p})] = \mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \mathbf{B} - \mathbf{D},$$

where the vectors $\mathbf{r}_i \in \mathbb{R}^{n_d}$ are defined in (2.5) and, consequently, $\mathbf{r}(\mathbf{p}) = \text{vec}(\mathbf{R}(\mathbf{p}))$.³ The columns of $\mathbf{D} = [\mathbf{d}_1 \cdots \mathbf{d}_{n_s}]$ are the measurements corresponding to source \mathbf{b}_i . We have

$$(3.2) \quad \min_{\mathbf{p}} \|\mathbf{r}(\mathbf{p})\|_2^2 = \min_{\mathbf{p}} \sum_{j=1}^{n_s} \|\mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \mathbf{b}_j - \mathbf{d}_j\|_2^2 = \min_{\mathbf{p}} \|\mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \mathbf{B} - \mathbf{D}\|_F^2.$$

Each evaluation of the objective function requires solving $n_s \cdot n_\omega$ linear systems. Haber, Chung, and Herrmann [12] reduce this cost using simultaneous random sources combined with (stochastic) trace estimators, following Hutchinson [15].

Let \mathbf{w} be a random vector with mean $\mathbf{0}$ and identity covariance matrix, and let \mathbb{E} denote the expected value with respect to the random vector \mathbf{w} . Then

$$\mathbb{E} [\mathbf{w}^T \mathbf{R}(\mathbf{p})^T \mathbf{R}(\mathbf{p}) \mathbf{w}] = \text{trace} (\mathbf{R}(\mathbf{p})^T \mathbf{R}(\mathbf{p})) = \|\mathbf{R}(\mathbf{p})\|_F^2.$$

As a particular choice, we choose \mathbf{w} to be a realization from the Rademacher distribution, where each component of \mathbf{w} is independent and identically distributed (i.i.d.) taking values from $\{-1, +1\}$, each with probability $\frac{1}{2}$. Then, as shown in [15], $\mathbf{w}^T \mathbf{R}(\mathbf{p})^T \mathbf{R}(\mathbf{p}) \mathbf{w}$ is a minimum variance and unbiased estimator of the trace of $\mathbf{R}(\mathbf{p})^T \mathbf{R}(\mathbf{p})$. Thus, the nonlinear least squares problem can be written as a stochastic minimization problem

$$(3.3) \quad \min_{\mathbf{p}} \|\mathbf{R}(\mathbf{p})\|_F^2 = \min_{\mathbf{p}} \text{trace} \mathbf{R}(\mathbf{p})^T \mathbf{R}(\mathbf{p}) = \min_{\mathbf{p}} \mathbb{E} (\mathbf{w}^T \mathbf{R}(\mathbf{p})^T \mathbf{R}(\mathbf{p}) \mathbf{w}).$$

For a random vector \mathbf{w} and simultaneous random source $\mathbf{B}\mathbf{w}$, we have

$$(3.4) \quad \mathbf{R}(\mathbf{p}) \mathbf{w} = (\mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \mathbf{B} - \mathbf{D}) \mathbf{w} = \mathbf{C}^T \mathbf{A}^{-1}(\mathbf{p}) \mathbf{B} \mathbf{w} - \mathbf{D} \mathbf{w}.$$

So, computing $\|\mathbf{R}(\mathbf{p}) \mathbf{w}\|_2^2$ requires a single PDE solve rather than n_s solves, which drastically reduces the cost of a function evaluation.

In contrast to the approach in [12], we use a Newton-type method, so we also need to reduce the cost of Jacobian evaluations. Therefore, we propose a variation

³For multiple frequencies, we need to compute the residual for each frequency, $[\mathbf{R}(\omega_1, \mathbf{p}) \quad \mathbf{R}(\omega_2, \mathbf{p}) \quad \cdots] = [\mathbf{C}^T \mathbf{A}^{-1}(\omega_1, \mathbf{p}) \mathbf{B} - \mathbf{D}_1 \quad \mathbf{C}^T \mathbf{A}^{-1}(\omega_2, \mathbf{p}) \mathbf{B} - \mathbf{D}_2 \quad \cdots]$.

that also drastically reduces the cost of computing $\mathbf{A}^{-T}(\mathbf{p})\mathbf{C}$ for the Jacobian. Let $\mathbf{v} \in \mathbb{R}^{n_d}$ and $\mathbf{w} \in \mathbb{R}^{n_s}$ with all components i.i.d. uniformly from $\{-1, +1\}$. Using the well-known cyclic property of the trace [20, p. 110], we get

$$\begin{aligned} \mathbb{E} \left[(\mathbf{v}^T \mathbf{R} \mathbf{w})^2 \right] &= \mathbb{E} \left[(\mathbf{v}^T \mathbf{R} \mathbf{w}) (\mathbf{w}^T \mathbf{R}^T \mathbf{v}) \right] = \mathbb{E} \left[\text{trace} (\mathbf{v} \mathbf{v}^T \mathbf{R} \mathbf{w} \mathbf{w}^T \mathbf{R}^T) \right] \\ (3.5) \qquad \qquad &= \text{trace} (\mathbb{E} (\mathbf{v} \mathbf{v}^T \mathbf{R} \mathbf{w} \mathbf{w}^T \mathbf{R}^T)) = \text{trace} (\mathbf{R} \mathbf{R}^T) = \|\mathbf{R}\|_F^2, \end{aligned}$$

which requires a single additional adjoint solve rather than an additional n_d solves for the Jacobian.

Typically, we need multiple random samples \mathbf{w}_j and \mathbf{v}_j to make the variance in our stochastic estimates sufficiently small. Hence, we set

$$(3.6) \qquad \mathbf{W} = \frac{1}{\sqrt{\ell_s}} (\mathbf{w}_1 \mathbf{w}_2 \cdots \mathbf{w}_{\ell_s}) \in \mathbb{R}^{n_s \times \ell_s},$$

where each column vector \mathbf{w}_j is i.i.d. with zero expectation and covariance equal to the identity and $\ell_s \ll n_s$. Similarly, we set

$$(3.7) \qquad \mathbf{V} = \frac{1}{\sqrt{\ell_d}} (\mathbf{v}_1 \mathbf{v}_2 \cdots \mathbf{v}_{\ell_d}) \in \mathbb{R}^{n_d \times \ell_d},$$

where each column vector \mathbf{v}_j is i.i.d. with zero expectation and covariance equal to the identity and $\ell_d \ll n_d$. It is easily verified that these choices give

$$(3.8) \qquad \mathbb{E}[\mathbf{W} \mathbf{W}^T] = \mathbf{I}_{n_s} \text{ and } \mathbb{E}[\mathbf{V} \mathbf{V}^T] = \mathbf{I}_{n_d}.$$

Next, we replace the sources \mathbf{B} by simultaneous random sources $\mathbf{B} \mathbf{W}$ and the detectors \mathbf{C} by simultaneous random detectors $\mathbf{C} \mathbf{V}$. Assume that \mathbf{W} and \mathbf{V} are independent and we compute unbiased estimates for $\|\mathbf{R}(\mathbf{p})\|^2$.

THEOREM 3.1. *Let $\mathbf{W} \in \mathbb{R}^{n_s \times \ell_s}$ and $\mathbf{V} \in \mathbb{R}^{n_d \times \ell_d}$ be as given above. Let $\mathbf{R} \in \mathbb{R}^{n_d \times n_s}$. Then*

$$(3.9) \qquad \mathbb{E} \left[\|\mathbf{V}^T \mathbf{R} \mathbf{W}\|_F^2 \right] = \|\mathbf{R}\|_F^2.$$

Proof.

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{V}^T \mathbf{R} \mathbf{W}\|_F^2 \right] &= \mathbb{E} \left[\text{trace} (\mathbf{W}^T \mathbf{R}^T \mathbf{V} \mathbf{V}^T \mathbf{R} \mathbf{W}) \right] = \mathbb{E} \left[\text{trace} (\mathbf{W} \mathbf{W}^T \mathbf{R}^T \mathbf{V} \mathbf{V}^T \mathbf{R}) \right] \\ (3.10) \qquad \qquad &= \text{trace} (\mathbb{E} [\mathbf{W} \mathbf{W}^T] \mathbf{R}^T \mathbb{E} [\mathbf{V} \mathbf{V}^T] \mathbf{R}) = \text{trace} (\mathbf{R}^T \mathbf{R}) = \|\mathbf{R}\|_F^2. \quad \square \end{aligned}$$

Since TREGS has proven very effective for the nonlinear least squares problem in DOT with PaLS, we continue to use the TREGS algorithm in the stochastic minimization problem

$$(3.11) \qquad \min_{\mathbf{p}} \mathbb{E} \left[\|\mathbf{V}^T \mathbf{R}(\mathbf{p}) \mathbf{W}\|^2 \right] = \min_{\mathbf{p}} \|\mathbf{R}(\mathbf{p})\|^2.$$

We derive the least squares problem used in TREGS to compute a regularized Gauss–Newton update for the stochastic problem as follows. For any \mathbf{p} ,

$$(3.12) \qquad \text{vec} (\mathbf{V}^T \mathbf{R}(\mathbf{p}) \mathbf{W}) = (\mathbf{W}^T \otimes \mathbf{V}^T) \text{vec}(\mathbf{R}(\mathbf{p})) = (\mathbf{W}^T \otimes \mathbf{V}^T) \mathbf{r}(\mathbf{p});$$

see [14, Lemma 4.3.1]. Using a first order approximation to $\mathbf{r}(\mathbf{p} + \boldsymbol{\delta})$ gives

$$(\mathbf{W}^T \otimes \mathbf{V}^T) \mathbf{r}(\mathbf{p} + \boldsymbol{\delta}) \approx (\mathbf{W}^T \otimes \mathbf{V}^T) (\mathbf{r}(\mathbf{p}) + \mathbf{J} \boldsymbol{\delta}),$$

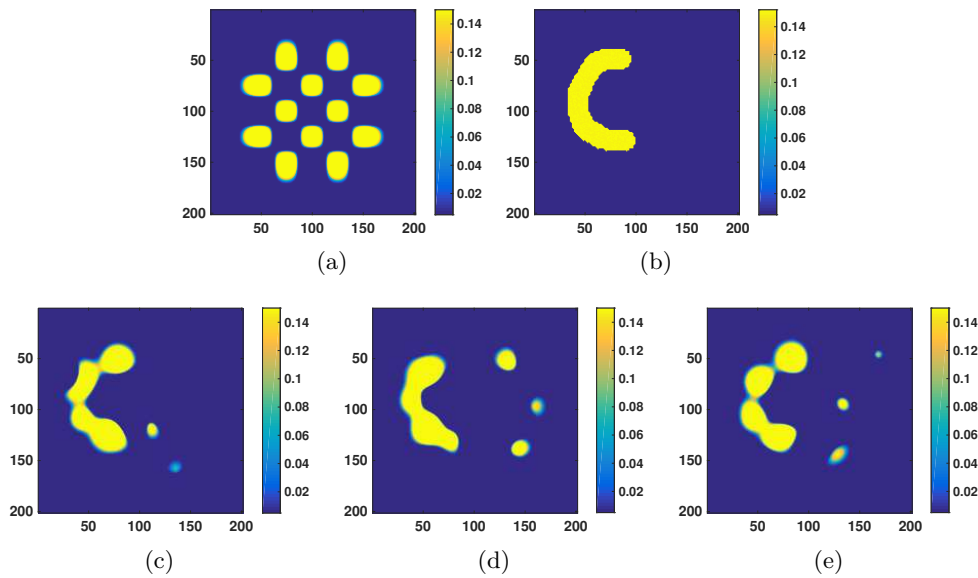


FIG. 3.1. Reconstruction of a test anomaly on a 201×201 mesh with 32 sources, 32 detectors, using only the zero frequency. (a) Initial configuration with 25 basis functions arranged in a 5×5 grid where 12 basis functions have positive expansion factors (visible as high absorption regions) and 13 basis functions have negative expansion factors (invisible). (b) True shape of the anomaly. (c) Reconstruction using all sources and detectors. (d) and (e) Two reconstruction results using simultaneous random sources and detectors with $\ell_s = \ell_d = 10$.

which leads to the (sampled) least squares problem

$$(3.13) \quad \min_{\delta} \|(\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{J}\delta + (\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{r}(\mathbf{p})\|_2^2,$$

replacing (2.10). Note that setting up the least squares problem (3.13) does not require any computations beyond $\mathbf{A}^{-1}(\mathbf{p})(\mathbf{B}\mathbf{W})$ and $\mathbf{A}^{-T}(\mathbf{p})(\mathbf{C}\mathbf{V})$. In addition, (3.13) has the following desirable properties for the sampled Jacobian and residual, which follow directly from (3.8) and well-known properties of the Kronecker product:

$$(3.14) \quad \mathbb{E} \left[((\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{J})^T (\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{r} \right] = \mathbf{J}^T \mathbb{E}[\mathbf{W}\mathbf{W}^T \otimes \mathbf{V}\mathbf{V}^T]\mathbf{r} = \mathbf{J}^T \mathbf{r},$$

$$(3.15) \quad \mathbb{E} \left[((\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{J})^T (\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{J} \right] = \mathbf{J}^T \mathbb{E}[\mathbf{W}\mathbf{W}^T \otimes \mathbf{V}\mathbf{V}^T]\mathbf{J} = \mathbf{J}^T \mathbf{J}.$$

So, the proposed randomization provides unbiased estimates for the gradient and the Gauss–Newton Hessian.

Two approaches to stochastic optimization are commonly used [26]. One approach, stochastic approximation (SA), uses a new random vector (or small batch of random vectors) in each optimization step. The other approach, SAA, uses a fixed set of random vectors over multiple (or many) optimization steps. In this paper, we focus on the SAA approach [26] to solve the stochastic problem (3.11). The SAA approach approximates (3.11) by the sample average problem. At each iteration, this approach requires solving only $\ell_s + \ell_d$ linear systems for each frequency to estimate the objective function and the Jacobian rather than $n_s + n_d$.

We give two representative solutions for our problem using the SAA approach in Figure 3.1. For DOT, the use of simultaneous random sources and detectors initially

leads to good progress. However, later in the iteration the convergence slows down, and in many cases, for our problem, it does not lead to sufficiently accurate solutions. In fact, with the SAA approach, (typically) the residual norm does not reach the noise level, the stopping criterion used (see the discussion of PaLS in section 2), while the standard optimization using all sources and all detectors does converge to the noise level. We will demonstrate this in section 4. In the next section, we provide a solution to this problem.

3.2. Improving the randomized approach. In the standard SAA approach, when convergence slows down or a minimum is found for the chosen sample (but not for the true problem), a new sample is chosen to improve the approximate solution. However, for our problem this approach leads to slow convergence and stagnation, unless we obtain more accurate estimates of the dominant singular components of the Jacobian and the corresponding components of the gradient (see also [10]). Hence, after exploiting the relatively fast initial convergence for our problem, we want to avoid stagnation of convergence in the next phase. One approach is to add additional random simultaneous sources and detectors, that is, increase the sample size, as proposed in [25, 6, 5, 23, 24], with good results. However, this requires progressively more, expensive, solves. Therefore, for efficiency, we choose to keep the number of simultaneous sources and detectors fixed. To improve convergence, we exploit the (often fairly good) approximate solution at a chosen modest intermediate tolerance to obtain a small number of simultaneous sources and detectors that are optimized to provide more accurate estimates of the desired derivative information. Since this optimization is local for the current \mathbf{p} , we complement these optimized vectors by random simultaneous sources and detectors. We make this precise below.

The nonlinear least squares algorithm TREGS focuses on the dominant singular values of the Jacobian to compute good updates to the parameter vector [10]. The corresponding right singular vectors capture the directions in parameter space of largest sensitivity in the objective function. Hence, we want to update \mathbf{W} and \mathbf{V} so as to best approximate the dominant right singular subspace of \mathbf{J} while respecting the Kronecker product structure in (3.13). This is important for two reasons. First, for the same (fixed) small number of simultaneous sources and detectors, this gives us locally (at the current \mathbf{p}) the best approximation to what TREGS would do using all sources and detectors. Second, the directions corresponding to the dominant right singular vectors are best informed by the data.

So, when a chosen intermediate tolerance is reached, our method computes the full Jacobian \mathbf{J} once, which requires a total of $n_s + n_d$ solves. Then, we compute a small number, q_s respectively q_d , of orthonormal, optimized simultaneous sources ($\widehat{\mathbf{W}}$) and detectors ($\widehat{\mathbf{V}}$). In practice, small q_s and q_d , 2 to 4, seem to be sufficient. We provide some experimental results regarding the number of optimized directions in section 4. Since \mathbf{J} is typically of rank substantially lower than n_p [1, 10], we expect that computing optimized directions can be done with an accurate approximation to \mathbf{J} that is computed efficiently using one of the sampling approaches from Halko, Martinsson, and Tropp [13] combined with knowledge of the structure of \mathbf{J} [1, 10]. The effort will be substantially less than solving for all sources and detectors (for each frequency) [13]. However, the details are beyond the scope of the current paper.

We would like to maximize

$$(3.16) \quad \|(\widehat{\mathbf{W}}^T \otimes \widehat{\mathbf{V}}^T)\mathbf{J}\|_F^2.$$

However, the Kronecker product structure combined with the constraints that $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{V}}$ be isometric matrices leads to a nonlinear constrained optimization problem. For efficiency, we exploit the tensor structure of this problem, that is, we consider the Jacobian as a third order array, \mathcal{J} , with components

$$(3.17) \quad \mathcal{J}_{ijk} = \mathbf{c}_i^T \mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial p_k} \mathbf{A}^{-1} \mathbf{b}_j.$$

This allows us to use an alternating least squares algorithm, a variant of the higher order orthogonal iteration (HOOI) [8], to find $\widehat{\mathbf{V}} \in \mathbb{R}^{n_d \times q_d}$ and $\widehat{\mathbf{W}} \in \mathbb{R}^{n_s \times q_s}$ that approximately maximize (3.16); details follow in the next section. The algorithm is only guaranteed to find a local maximum [19, 8]. However, in a number of numerical tests carried out, the tensor algorithm discussed below seems to always converge to the global maximum. A similar observation is reported in [8]. In our experiments, the algorithm also attains the same solutions as the MATLAB `fmincon` routine, which optimizes for $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{V}}$ simultaneously.

In the remainder of this section, we first discuss computing the optimized simultaneous detectors ($\widehat{\mathbf{V}}$) and sources ($\widehat{\mathbf{W}}$) and then complementing these with random simultaneous sources constrained to $\text{Range}(\widehat{\mathbf{W}})^\perp$ and random simultaneous detectors constrained to $\text{Range}(\widehat{\mathbf{V}})^\perp$.

3.2.1. Computing optimized simultaneous sources and detectors. This problem is closely related to the truncated higher-order SVD [7, 17] or, more precisely, to a truncated Tucker2 decomposition [27, 17], as we do not need truncation in the parameter-derivative direction (the columns of \mathbf{J}). As $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{V}}$ are both isometric matrices, so is their Kronecker product, $\mathbf{X} = \widehat{\mathbf{W}} \otimes \widehat{\mathbf{V}}$. Let $[\mathbf{X} \ \mathbf{X}_c]$ be a (real) orthogonal matrix, and let $S^{k \times \ell} = \{\Theta \in \mathbb{R}^{k \times \ell} \mid \Theta^T \Theta = \mathbf{I}_\ell\}$ (the set of all $k \times \ell$ isometric matrices). Then, it follows from standard properties of the Frobenius norm that $\|\mathbf{J}\|_F = \|[\mathbf{X} \ \mathbf{X}_c]^T \mathbf{J}\|_F$ and

$$\|\mathbf{J}\|_F^2 = \|\mathbf{X}^T \mathbf{J}\|_F^2 + \|\mathbf{X}_c^T \mathbf{J}\|_F^2.$$

Since $\|(\mathbf{I} - \mathbf{X}\mathbf{X}^T)\mathbf{J}\|_F = \|\mathbf{X}_c^T \mathbf{J}\|_F$, we have that the maximization problem

$$(3.18) \quad \arg \max_{\substack{\widetilde{\mathbf{W}} \in \mathbf{S}^{n_s \times \ell_s} \\ \widetilde{\mathbf{V}} \in \mathbf{S}^{n_d \times \ell_d}}} \|(\widetilde{\mathbf{W}}^T \otimes \widetilde{\mathbf{V}}^T)\mathbf{J}\|_F^2$$

is equivalent to the minimization problem

$$\arg \min_{\substack{\widetilde{\mathbf{W}} \in \mathbf{S}^{n_s \times \ell_s} \\ \widetilde{\mathbf{V}} \in \mathbf{S}^{n_d \times \ell_d}}} \|\mathbf{J} - (\widetilde{\mathbf{W}} \otimes \widetilde{\mathbf{V}})(\widetilde{\mathbf{W}} \otimes \widetilde{\mathbf{V}})^T \mathbf{J}\|_F^2;$$

see also [8, Theorem 4.1] and [17, p. 477]. We solve the maximization problem by a slight variation of the HOOI algorithm [8] to compute a truncated Tucker2 decomposition [27, 17]. We do not need any compression of the column dimension of \mathbf{J} (the direction corresponding to the parameter derivatives).

We briefly outline the algorithm in tensor form; a MATLAB pseudocode is given in subsection 3.3. Let \mathcal{J} be the tensor representation of the Jacobian \mathbf{J} , where $\mathcal{J}_{ijk} = \mathbf{c}_i^T \mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial p_k} \mathbf{A}^{-1} \mathbf{b}_j$; the dependence on \mathbf{p} in \mathbf{A} , \mathbf{J} , and \mathcal{J} is suppressed for brevity.

(1) We compute the SVD of the matrix obtained by unfolding the tensor \mathcal{J} along its lateral slices $\mathcal{J}_{:,j,:}$ and computing the SVD of the resulting matrix $\tilde{\mathbf{J}}$:

$$(3.19) \quad \mathbf{U}\Sigma\mathbf{Y}^T = \tilde{\mathbf{J}} = [\mathcal{J}_{:,1,:} \quad \mathcal{J}_{:,2,:} \quad \cdots \quad \mathcal{J}_{:,n_s,:}] .$$

Next, we set the initial $\tilde{\mathbf{V}} = [\mathbf{u}_1 \ \mathbf{u}_2 \ \cdots \ \mathbf{u}_{q_d}]$ (the leading q_d left singular vectors of $\tilde{\mathbf{J}}$).

Subsequently, we iterate the following two steps until the change in the approximate solution, the tuple $(\tilde{\mathbf{V}}, \tilde{\mathbf{W}})$, is sufficiently small.

(2) Let the tensor \mathcal{K} be defined as

$$\mathcal{K}_{jkl} = \sum_{m=1}^{n_d} (\tilde{\mathbf{v}}_\ell)_m \mathcal{J}_{m,j,k} \quad \text{for } \ell = 1, \dots, q_d,$$

where the vectors $\tilde{\mathbf{v}}_\ell$ are the columns of $\tilde{\mathbf{V}}$ ($\ell = 1, \dots, q_d$). We compute the SVD of the matrix $\tilde{\mathbf{K}}$ obtained by unfolding \mathcal{K} along its frontal slices $\mathcal{K}_{:,:,\ell}$,

$$(3.20) \quad \Phi\Omega\Psi^T = \tilde{\mathbf{K}} = [\mathcal{K}_{:,:,1} \quad \mathcal{K}_{:,:,2} \quad \cdots \quad \mathcal{K}_{:,:,q_d}] ,$$

and we set $\tilde{\mathbf{W}} = [\varphi_1 \ \varphi_2 \ \cdots \ \varphi_{q_s}]$ (the leading q_s left singular vectors of $\tilde{\mathbf{K}}$).

(3) Define the tensor \mathcal{L} by

$$\mathcal{L}_{ikl} = \sum_{m=1}^{n_s} (\tilde{\mathbf{w}}_\ell)_m \mathcal{J}_{i,m,k} \quad \text{for } \ell = 1, \dots, q_s,$$

where the vectors $\tilde{\mathbf{w}}_\ell$ are the columns of $\tilde{\mathbf{W}}$ ($\ell = 1, \dots, q_s$). We compute the SVD of the matrix $\tilde{\mathbf{L}}$ obtained by unfolding \mathcal{L} along its frontal slices $\mathcal{L}_{:,:,\ell}$,

$$(3.21) \quad \mathbf{U}\Sigma\mathbf{Y}^T = \tilde{\mathbf{L}} = [\mathcal{L}_{:,:,1} \quad \mathcal{L}_{:,:,2} \quad \cdots \quad \mathcal{L}_{:,:,q_s}] ,$$

and we set (the new) $\tilde{\mathbf{V}} = [\mathbf{u}_1 \ \mathbf{u}_2 \ \cdots \ \mathbf{u}_{q_d}]$.

Finally, after convergence, we set $\hat{\mathbf{V}} = \tilde{\mathbf{V}}$ and $\hat{\mathbf{W}} = \tilde{\mathbf{W}}$.

While this procedure gives good approximate solutions, in general only a local maximum of (3.16) is guaranteed [8]. However, in a number of numerical tests carried out, the tensor algorithm seems to always converge to the globally optimal solution. This experience was also reported in [8]. Moreover, we have the following useful result for the case that $\|\mathbf{J}\|_F$ can be preserved exactly.

THEOREM 3.2. *If isometric matrices $\tilde{\mathbf{W}} \in \mathbb{R}^{n_s \times q_s}$ and $\tilde{\mathbf{V}} \in \mathbb{R}^{n_d \times q_d}$ exist such that*

$$(3.22) \quad \|(\tilde{\mathbf{W}}^T \otimes \tilde{\mathbf{V}}^T)\mathbf{J}\|_F = \|\mathbf{J}\|_F,$$

then, in one iteration, steps (1)–(3) above compute isometric matrices $\hat{\mathbf{W}} \in \mathbb{R}^{n_s \times q_s}$ and $\hat{\mathbf{V}} \in \mathbb{R}^{n_d \times q_d}$ such that

$$\|(\hat{\mathbf{W}}^T \otimes \hat{\mathbf{V}}^T)\mathbf{J}\|_F = \|\mathbf{J}\|_F.$$

Proof. This result follows from the error bound (Property 10) in [7], which in the (matrix) notation of this paper is given by

$$(3.23) \quad \|\mathbf{J} - (\widetilde{\mathbf{W}} \otimes \widetilde{\mathbf{V}})(\widetilde{\mathbf{W}} \otimes \widetilde{\mathbf{V}})^T \mathbf{J}\|_F^2 \leq \sum_{i=q_d+1}^{n_d} \sigma_i^2 + \sum_{j=q_s+1}^{n_s} \omega_j^2.$$

The assumption (3.22) implies that $\text{Range}(\widetilde{\mathbf{J}}) \subseteq \text{Range}(\widetilde{\mathbf{V}})$, which in turn implies that $\text{Rank}(\widetilde{\mathbf{J}}) \leq q_d$ and hence that $\sigma_{q_d+1} = 0, \sigma_{q_d+2} = 0, \dots, \sigma_{n_d} = 0$. In a similar fashion, from assumption (3.22) combined with the choice for $\widetilde{\mathbf{V}}$ from step (1) (see (3.19)), it follows that $\text{Range}(\widetilde{\mathbf{K}}) \subseteq \text{Range}(\widetilde{\mathbf{W}})$, which implies that $\text{Rank}(\widetilde{\mathbf{K}}) \leq q_s$ and hence $\omega_{q_s+1} = 0, \omega_{q_s+2} = 0, \dots, \omega_{n_s} = 0$.

Substitution of $\sigma_{q_d+1}, \dots, \sigma_{n_d}, \omega_{q_s+1}, \dots, \omega_{n_s}$ into (3.23) shows that the algorithm will have converged at this point. The algorithm then sets $\widehat{\mathbf{W}} = \widetilde{\mathbf{W}}$ and $\widehat{\mathbf{V}} = \widetilde{\mathbf{V}}$. \square

Since $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{V}}$ are only optimal at the current \mathbf{p} , we complement these optimized simultaneous sources and detectors with a new set of random simultaneous sources and detectors constrained to the orthogonal complement of the span of the optimized directions, keeping the total number of columns in \mathbf{W} and \mathbf{V} the same as before. This procedure can be carried out periodically or for a sequence of prescribed tolerances, but in our experiments it never needs to be done more than once.

3.2.2. Computing complementary random simultaneous sources and detectors. We extend the optimized sources and detectors with random simultaneous sources and detectors. Let $\mathbf{W}_f = [\widehat{\mathbf{W}} \ \mathbf{W}_c] \in \mathbb{R}^{n_s \times n_s}$ and $\mathbf{V}_f = [\widehat{\mathbf{V}} \ \mathbf{V}_c] \in \mathbb{R}^{n_d \times n_d}$ be orthogonal matrices. We have

$$(3.24) \quad \begin{aligned} \|\mathbf{R}(\mathbf{p})\|_F &= \|\mathbf{V}_f^T \mathbf{R}(\mathbf{p}) \mathbf{W}_f\|_F \\ &= \left\| \begin{bmatrix} \widehat{\mathbf{V}}^T \mathbf{R}(\mathbf{p}) \widehat{\mathbf{W}} & \widehat{\mathbf{V}}^T \mathbf{R}(\mathbf{p}) \mathbf{W}_c \\ \mathbf{V}_c^T \mathbf{R}(\mathbf{p}) \widehat{\mathbf{W}} & \mathbf{V}_c^T \mathbf{R}(\mathbf{p}) \mathbf{W}_c \end{bmatrix} \right\|_F. \end{aligned}$$

The (1, 1)-block of this matrix can be computed using the known optimized sources and detectors. We estimate the remaining blocks, proceeding more or less as before. We pick random matrices $\mathbf{Y} = (\ell_s - q_s)^{-1/2} [\mathbf{y}_1 \ \mathbf{y}_2 \ \dots \ \mathbf{y}_{\ell_s - q_s}]$, where each column vector $\mathbf{y}_j \in \mathbb{R}^{\ell_s - q_s}$ is i.i.d. with zero mean and identity covariance and $\mathbf{Z} = (\ell_d - q_d)^{-1/2} [\mathbf{z}_1 \ \mathbf{z}_2 \ \dots \ \mathbf{z}_{\ell_d - q_d}]$, where each column vector $\mathbf{z}_j \in \mathbb{R}^{n_d - q_d}$ is i.i.d. with zero mean and identity covariance. In our numerical experiments, we again use the Rademacher distribution. Next, we set the new matrices \mathbf{W} and \mathbf{V} to

$$(3.25) \quad \mathbf{W} = [\widehat{\mathbf{W}} \ (\mathbf{W}_c \mathbf{Y})],$$

$$(3.26) \quad \mathbf{V} = [\widehat{\mathbf{V}} \ (\mathbf{V}_c \mathbf{Z})].$$

We have the following result.

THEOREM 3.3. *Let $\mathbf{W} \in \mathbb{R}^{n_s \times \ell_s}$ and $\mathbf{V} \in \mathbb{R}^{n_d \times \ell_d}$ be given in (3.25) and (3.26), respectively. Let $\mathbf{R} \in \mathbb{R}^{n_d \times n_s}$. Then,*

$$\begin{aligned}
 (3.27) \quad & \mathbb{E}[\mathbf{W}\mathbf{W}^T] = \mathbf{I}_{n_s}, \\
 (3.28) \quad & \mathbb{E}[\mathbf{V}\mathbf{V}^T] = \mathbf{I}_{n_d}, \\
 (3.29) \quad & \mathbb{E}[\|\mathbf{V}^T\mathbf{R}(\mathbf{p})\mathbf{W}\|_F^2] = \|\mathbf{R}(\mathbf{p})\|_F^2.
 \end{aligned}$$

Proof. For (3.27), we have

$$\begin{aligned}
 \mathbb{E} \left[[\widehat{\mathbf{W}}(\mathbf{W}_c\mathbf{Y})][\widehat{\mathbf{W}}(\mathbf{W}_c\mathbf{Y})]^T \right] &= \widehat{\mathbf{W}}\widehat{\mathbf{W}}^T + \mathbb{E}[\mathbf{W}_c\mathbf{Y}\mathbf{Y}^T\mathbf{W}_c^T] \\
 &= \widehat{\mathbf{W}}\widehat{\mathbf{W}}^T + \mathbf{W}_c\mathbb{E}[\mathbf{Y}\mathbf{Y}^T]\mathbf{W}_c^T = \mathbf{I}_{n_s}.
 \end{aligned}$$

An analogous derivation holds for (3.28). The proof for the last equation follows the proof for Theorem 3.1, using the results above. □

As a result of Theorem 3.3, the new \mathbf{W} and \mathbf{V} again give for the expectation of the sampled gradient and the expectation of the sampled Gauss–Newton Hessian the true gradient and Gauss–Newton Hessian,

$$\begin{aligned}
 \mathbb{E}[(\mathbf{W}^T \otimes \mathbf{V})\mathbf{J}]^T(\mathbf{W}^T \otimes \mathbf{V})\mathbf{r} &= \mathbf{J}^T\mathbf{r}, \\
 \mathbb{E}[(\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{J}]^T((\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{J}) &= \mathbf{J}^T\mathbf{J}.
 \end{aligned}$$

3.3. Implementation. In this section, we first outline the efficient computation of the sampled residual and sampled Jacobian. Next, we give the overall algorithm for solving the inverse problem in Algorithm 4.1. We give the details of computing optimized sources and detectors in Algorithm 4.2, which is based on the matrix representation of the Jacobian as defined in (2.8). We estimate the norm of the residual using

$$(3.30) \quad (\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{r}(\mathbf{p}) = \begin{bmatrix} \mathbf{V}^T\mathbf{C}^T\mathbf{A}^{-1}(\mathbf{p})\mathbf{B}\mathbf{w}_1 - \mathbf{V}^T\mathbf{D}\mathbf{w}_1 \\ \vdots \\ \mathbf{V}^T\mathbf{C}^T\mathbf{A}^{-1}(\mathbf{p})\mathbf{B}\mathbf{w}_{\ell_s} - \mathbf{V}^T\mathbf{D}\mathbf{w}_{\ell_s} \end{bmatrix} = \begin{bmatrix} \mathbf{V}^T\mathbf{C}^T\mathbf{z}_1 - \mathbf{V}^T\mathbf{D}\mathbf{w}_1 \\ \vdots \\ \mathbf{V}^T\mathbf{C}^T\mathbf{z}_{\ell_s} - \mathbf{V}^T\mathbf{D}\mathbf{w}_{\ell_s} \end{bmatrix},$$

where we solve $\mathbf{A}(\mathbf{p})\mathbf{z}_i = \mathbf{B}\mathbf{w}_i$ for $\mathbf{z}_i, i = 1, \dots, \ell_s$. This reduces the number of large solves from n_s to ℓ_s per frequency. To compute the Jacobian, we solve the systems, $\mathbf{A}^T(\mathbf{p})\mathbf{y}_j = \mathbf{C}\mathbf{v}_j$ for $\mathbf{y}_j, j = 1, \dots, \ell_d$. This reduces the additional number of large solves from n_d to ℓ_d per frequency. We can use iterative solvers or sparse direct solvers depending on the size of the system [16]. To obtain the k th column of $(\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{J}$, we compute

$$(3.31) \quad \left[\mathbf{y}_1^T \frac{\partial \mathbf{A}(\mathbf{p})}{\partial \mathbf{p}_k} \mathbf{z}_1 \cdots \mathbf{y}_{\ell_d}^T \frac{\partial \mathbf{A}(\mathbf{p})}{\partial \mathbf{p}_k} \mathbf{z}_1 \quad \mathbf{y}_1^T \frac{\partial \mathbf{A}(\mathbf{p})}{\partial \mathbf{p}_k} \mathbf{z}_2 \cdots \mathbf{y}_{\ell_d}^T \frac{\partial \mathbf{A}(\mathbf{p})}{\partial \mathbf{p}_k} \mathbf{z}_{\ell_s} \right]^T,$$

where $\partial \mathbf{A}(\mathbf{p})/\partial \mathbf{p}_k$ is a diagonal matrix if we only invert for absorption. If we also invert for diffusion, this matrix has 5 (in 2 dimensions) or 7 (in 3 dimensions) diagonals. Moreover, after a few iterations, the changes in $\mathbf{A}(\mathbf{p})$ are highly localized, and $\partial \mathbf{A}(\mathbf{p})/\partial \mathbf{p}_k$ contains mostly zero coefficients; see [9]. In that case, first we find the few nonzero components of $\partial \mathbf{A}(\mathbf{p})/\partial \mathbf{p}_k$ for each k and the corresponding nonzeros in \mathbf{z}_j and \mathbf{y}_i , and next we efficiently compute $\mathbf{y}_i^T \frac{\partial \mathbf{A}(\mathbf{p})}{\partial \mathbf{p}_k} \mathbf{z}_j$ referencing only the few nonzero components in $(\partial \mathbf{A}(\mathbf{p})/\partial \mathbf{p}_k)\mathbf{z}_j$.

Algorithm 4.1 Inversion using random and optimized simultaneous sources and detectors.

Pick standard initial guess \mathbf{p} ▷ see Fig. 3.1
 $\tau = \delta$ ▷ set intermediate tolerance τ based on noise level δ (section 4)

$\mathbf{W} = \ell_s^{-1/2}(\mathbf{w}_1 \cdots \mathbf{w}_{\ell_s})$ ▷ pick ℓ_s random simultaneous sources, cf. (3.6)
 $\mathbf{V} = \ell_d^{-1/2}(\mathbf{v}_1 \cdots \mathbf{v}_{\ell_d})$ ▷ pick ℓ_d random simultaneous detectors, cf. (3.7)
 $\widehat{\mathbf{D}} = \mathbf{V}^T \mathbf{D} \mathbf{W}$ ▷ with columns $\widehat{\mathbf{d}}_j$ (used below)

Solve $\mathbf{A}(\mathbf{p})\mathbf{z}_j = \mathbf{B}\mathbf{w}_j$, for $j = 1, \dots, \ell_s$
 $\widehat{\mathbf{r}}(\mathbf{p}) = (\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{r}(\mathbf{p}) = \text{vec}(\widehat{\mathbf{r}}_1(\mathbf{p}) \cdots \widehat{\mathbf{r}}_{\ell_s}(\mathbf{p}))$ with $\widehat{\mathbf{r}}_j(\mathbf{p}) = (\mathbf{C}\mathbf{V})^T \mathbf{z}_j - \widehat{\mathbf{d}}_j$

while $\|\widehat{\mathbf{r}}(\mathbf{p})\|_2^2 > \tau$ **do**
 { *compute optimization step* }
 Solve $\mathbf{A}^T(\mathbf{p})\mathbf{y}_i = \mathbf{C}\mathbf{v}_i$, for $i = 1, \dots, \ell_d$
 Compute $\widehat{\mathbf{J}} = (\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{J}$ using $\mathbf{y}_i^T \frac{\partial \mathbf{A}(\mathbf{p})}{\partial \mathbf{p}_k} \mathbf{z}_j$ ▷ see discussion starting below (3.30)
 Compute update $\delta\mathbf{p}$ by solving (3.13) using $\widehat{\mathbf{J}}$ and $\widehat{\mathbf{r}}$ ▷ with TREGS or alternative
 $\mathbf{p} = \mathbf{p} + \delta\mathbf{p}$; Update $\mathbf{A}(\mathbf{p})$ (if explicitly stored)
 Solve $\mathbf{A}(\mathbf{p})\mathbf{z}_j = \mathbf{B}\mathbf{w}_j$, for $j = 1, \dots, \ell_s$
 $\widehat{\mathbf{r}}(\mathbf{p}) = (\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{r}(\mathbf{p})$ ▷ implemented as above

end while
{ *compute optimized simultaneous sources and detectors* }
Compute $\mathbf{J} = \frac{\partial \mathbf{r}(\mathbf{p})}{\partial \mathbf{p}}$ following (2.8) or an accurate approximation by sampling
Compute $\widehat{\mathbf{V}}$ and $\widehat{\mathbf{W}}$ using Algorithm 4.2 ▷ see discussion in section 3.2.1
{ *compute the complementary random simultaneous sources and detectors* }
 $\mathbf{W} = \begin{bmatrix} \widehat{\mathbf{W}} & (\mathbf{W}_c \mathbf{Y}) \end{bmatrix}$; $\mathbf{V} = \begin{bmatrix} \widehat{\mathbf{V}} & (\mathbf{V}_c \mathbf{Z}) \end{bmatrix}$ ▷ see discussion in section 3.2.2
 $\widehat{\mathbf{D}} = \mathbf{V}^T \mathbf{D} \mathbf{W}$
Solve $\mathbf{A}(\mathbf{p})\mathbf{z}_j = \mathbf{B}\mathbf{w}_j$, for $j = 1, \dots, \ell_s$
 $\widehat{\mathbf{r}}(\mathbf{p}) = (\mathbf{W}^T \otimes \mathbf{V}^T)\mathbf{r}(\mathbf{p})$ ▷ implemented as above
while $\|\widehat{\mathbf{r}}(\mathbf{p})\|_2^2 > \tau^2$ **do**
 { *compute optimization step* } ▷ same steps as above
end while

4. Numerical experiments. In this section, we provide two numerical experiments, a 2D and a 3D test case, to demonstrate the effectiveness of combining random simultaneous sources and detectors with optimized simultaneous sources and detectors. The results show that our approach not only produces reconstruction results that are close to those obtained using all sources and all detectors, but it also substantially reduces the computational cost.

Our experimental set up is the same as that in [9], where model reduction was used to reduce the cost of inversion in DOT. The absorption images for the initial sets of parameters for the 2D and 3D experiments are given in Figure 4.1. For each test case, we construct anomalies in the pixel basis, and we add a small normally distributed random heterogeneity to both the background and the anomaly to make the medium inhomogeneous. We use this pixel-based absorption image to compute the (true) measured data and add 0.1% white noise to the measured data. This is the same noise level as used in [9]. Then, we reconstruct the absorption images using the PaLS [1] representation and TREGS [10] for optimization.

2D experiment. We use a 201×201 grid, which yields 40,401 unknowns in the discretized PDE (2.1). The model has 32 sources, 32 detectors, and we use only the zero frequency. Our model has 25 CSRBFs, which leads to 100 parameters (four per

Algorithm 4.2 Compute optimized sources and detectors.

```

{ compute  $\tilde{\mathbf{J}}$  and its  $q_d$  leading left singular vectors }
 $\tilde{\mathbf{J}} = []$ ; ▷ initialize  $\tilde{\mathbf{J}}$  as empty matrix
for  $j = 1:n_s$  do
     $\tilde{\mathbf{J}} = [\tilde{\mathbf{J}} \mathbf{J}_{(j-1)n_d+1:jn_d,1:k}]$  ▷ add next block from tensor  $\mathcal{J}$ 
end for
 $[\mathbf{U}, \Sigma, \mathbf{Y}] = \text{SVD}(\tilde{\mathbf{J}})$ ;  $\tilde{\mathbf{V}} = [\mathbf{u}_1 \mathbf{u}_2 \cdots \mathbf{u}_{q_d}]$  ▷ keep  $q_d$  leading left singular vectors of  $\tilde{\mathbf{J}}$ 
while not converged do
    { compute  $\tilde{\mathbf{K}}$  and its  $q_s$  leading left singular vectors }
     $\tilde{\mathbf{K}} = []$  ▷ initialize  $\tilde{\mathbf{K}}$  as empty matrix
    for  $\ell = 1:q_d$  do
         $\tilde{\mathbf{K}}_\ell = []$ 
        for  $j = 1:n_s$  do
             $(\tilde{\mathbf{K}}_\ell)_{j,1:k} = \mathbf{v}_\ell^T \mathbf{J}_{(j-1)n_d+1:jn_d,1:k}$  ▷ compute row  $j$  of  $\tilde{\mathbf{K}}_\ell$ 
        end for
         $\tilde{\mathbf{K}} = [\tilde{\mathbf{K}} \tilde{\mathbf{K}}_\ell]$  ▷ add next block of  $\tilde{\mathbf{K}}$ 
    end for
     $[\Phi, \Omega, \Psi] = \text{SVD}(\tilde{\mathbf{K}})$ ;  $\tilde{\mathbf{W}} = [\varphi_1 \varphi_2 \cdots \varphi_{q_s}]$  ▷ keep  $q_s$  leading left singular vectors of  $\tilde{\mathbf{K}}$ 
    { compute  $\tilde{\mathbf{L}}$  and its  $q_s$  leading left singular vectors }
     $\tilde{\mathbf{L}} = []$ 
    for  $\ell = 1:q_s$  do
         $\tilde{\mathbf{L}}_\ell = \sum_{j=1}^{n_s} (\tilde{\mathbf{w}}_\ell)_j \mathbf{J}_{(j-1)n_d+1:jn_d,1:k}$ 
         $\tilde{\mathbf{L}} = [\tilde{\mathbf{L}} \tilde{\mathbf{L}}_\ell]$  ▷ add next block of  $\tilde{\mathbf{L}}$ 
    end for
     $[\mathbf{U}, \Sigma, \mathbf{Y}] = \text{SVD}(\tilde{\mathbf{L}})$ ;  $\tilde{\mathbf{V}} = [\mathbf{u}_1 \mathbf{u}_2 \cdots \mathbf{u}_{q_d}]$  ▷ keep  $q_d$  leading left singular vectors of  $\tilde{\mathbf{L}}$ 
end while
 $\hat{\mathbf{V}} = \tilde{\mathbf{V}}$ ;  $\hat{\mathbf{W}} = \tilde{\mathbf{W}}$ 

```

2D basis function) for the nonlinear optimization. We use the same starting guess for each trial, (see Figure 4.1(a), with the 25 basis functions arranged in a 5×5 grid, where 12 basis functions have a positive expansion coefficient (visible as high absorption regions) and 13 basis functions have a negative expansion coefficient (invisible).

We use 10 random simultaneous sources and detectors. We update the simultaneous sources and detectors as discussed in section 3 after a chosen intermediate tolerance has been reached. We find that, in general, the noise level, δ , is a good choice as the intermediate tolerance: $\|\mathbf{r}(\mathbf{p})\|_2^2 = \delta$. Since the PaLS representation regularizes the problem, we consider the problem converged when $\|\mathbf{r}(\mathbf{p})\|_2^2 \leq \delta^2$. This is called the discrepancy principle (the factor $\frac{1}{2}$ in (1.1) is dropped for convenience). We run the 2D experiment for 50 trials. In each trial, the random simultaneous sources and detectors are chosen independently to get representative reconstruction results.

Example 1. The true absorption image for Example 1 is given in Figure 4.2(a). We also include the reconstruction results using all sources and detectors for comparison (see Figure 4.2(b)). As can be seen in Figure 4.2(c) at the intermediate tolerance, SAA gives a good localization of the anomaly; however, there is little further improvement using SAA (see Figures 3.1(d)–(e) and Figure 4.3(a)). Figures 4.2(d)–(f) show that using optimized simultaneous sources and detectors leads to solutions of the same quality as obtained using all sources and detectors. We report the total number of PDE solves required for each approach in Table 4.1 for a representative result from 50 trials.

While initially the SAA estimate is unbiased, a systematic underestimation of the residual/misfit (bias) [26, section 5.1.2] arises, since we optimize for a specific




FIG. 4.1. (a) Initial configuration for the 2D experiment with 25 basis functions arranged in a 5×5 grid, where 12 basis functions have positive expansion factors (visible) and 13 basis functions have negative expansion factors (invisible). (b) Initial configuration for the 3D experiment with 27 basis functions arranged in a $3 \times 3 \times 3$ grid, where 13 basis functions have positive expansion factors (visible) and 14 basis functions have negative expansion factors (invisible).

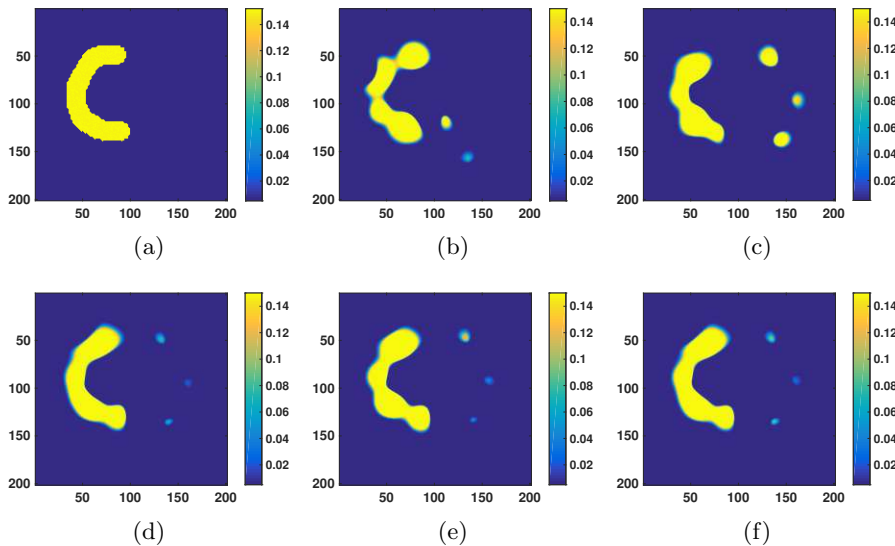


FIG. 4.2. Results for Example 1. Reconstruction of a test anomaly on a 201×201 mesh with 32 sources and detectors, 25 basis functions, and using only the zero frequency. The SAA approach uses 10 random simultaneous sources and detectors. (a) True shape of the anomaly. (b) Reconstruction using all sources and all detectors. (c) Reconstruction using the SAA approach at a chosen intermediate tolerance. (d) Reconstruction with SAA and 1 optimized simultaneous source and detector. (e) Reconstruction with SAA and 2 optimized simultaneous sources and detectors. (f) Reconstruction with SAA and 3 optimized simultaneous sources and detectors.

small set of random simultaneous sources and detectors. As a result, the algorithm generally stops prematurely. This can make a big difference, since often substantial improvement in the shape of the anomaly occurs towards the end of the optimization. Figure 4.3 demonstrates how poor the reconstructions using only the SAA approach can be at the convergence tolerance when underestimation of the residual norm is severe. To make a fair comparison in terms of the number of large systems solved, we check the true function evaluation of the SAA approach on the side. Table 4.2 shows that in terms of the true function evaluation, the SAA approach does not reach the




FIG. 4.3. Two examples of poor SAA reconstructions for the 2D problem after the maximum number of iterations. Reconstruction of a test anomaly on a 201×201 mesh with 32 sources and detectors, 25 basis functions, and using only the zero frequency.

TABLE 4.1

Example 1 results. The total number of iterations, function evaluations, Jacobian evaluations, and PDE solves required on average for 50 trials to reach the stopping criterion, $\|\mathbf{r}(\mathbf{p})\|_2^2 = \delta^2$. Parentheses indicate that the SAA approach does not reach the tolerance.

	Iteration number	Function evaluations	Jacobian evaluations	Total PDE solves	Tol
SAA* (intermediate tol)	10	11	6	170	δ
1 opt simult src/det	18	19	10	524	δ^2
2 opt simult srcs/dets	18	19	10	524	δ^2
3 opt simult srcs/dets	16	17	8	484	δ^2
All srcs/All dets	71	72	47	3808	δ^2
SAA**	(32)	(33)	(19)	(520)	δ^2
SAA ***	(92)	(93)	(67)	(1700)	δ^2

* The first row gives the cost to reach the intermediate tolerance for the SAA approach, $\|\mathbf{r}(\mathbf{p})\|_2^2 = \delta$.

** Since the SAA estimate becomes biased and underestimates the objective function, the algorithm stops prematurely.

***The SAA approach measuring the convergence with the true objective function.

convergence tolerance. Once we use a few optimized sources and detectors, this is no longer an issue (see Table 4.2).

The main purpose of the SAA approach and our modification is to reduce the large number of discretized PDE solves required for the inversion. In Table 4.1, we give a comparison of the total number of PDE solves for Example 1. Our approach drastically reduces the number of large-scale linear systems that need to be solved. Additionally, it substantially improves the reconstruction results of the SAA approach.

3D experiment. We use a $32 \times 32 \times 32$ grid, which gives 32,768 unknowns in the discretized PDE (2.1). The model has 225 sources at the top and 225 detectors on the bottom, and we use only the zero frequency. In the PaLS approach, we use 27 CSRBFs, which leads to 135 parameters (five per 3D basis function) for the nonlinear optimization. The absorption image using the initial set of parameters is given in Figure 4.1(b) where 13 basis functions have a positive expansion coefficient (visible as high absorption regions) and 14 basis functions have a negative expansion coefficient (invisible). In our approach, we use only 12 random simultaneous sources and detectors.

Example 2. The true absorption image for Example 2 is given in Figure 4.4(a). The reconstruction using all sources and detectors is given in Figure 4.4(b). Figure 4.4(c) shows that the SAA approach gives a good localization of the anomaly at the intermediate tolerance. However, little further improvement appears using the SAA

TABLE 4.2

Subset of results for Example 1. The comparison of the true objective function $\|\mathbf{r}\|_2^2$ and its SAA estimate relative to the stopping criterion (δ^2) for selected iterations. For the SAA approach, the estimated residual is obtained with 10 random simultaneous sources and detectors. Parentheses indicate that the SAA approach does not reach the tolerance. The estimated residual norms using random and optimized simultaneous sources and detectors are obtained using 3 optimized simultaneous sources and detectors.

SAA approach			Rand and optimized simult src/det		
Iter	True $\ \mathbf{r}\ _2^2$ (δ^2)	Estimated $\ \mathbf{r}\ _2^2$ (δ^2)	Iter	True $\ \mathbf{r}\ _2^2$ (δ^2)	Estimated $\ \mathbf{r}\ _2^2$ (δ^2)
1	118940	38820	1-5	(SAA)*	(SAA)*
6	1192.5	391.15	6	1194.9	1197.3
11	56.550	22.575	13	118.73	118.56
15	3.748	0.8650	16	19.894	19.929
(99)	—	—	22	0.8403	0.8389

* indicates that we initially use the SAA approach up to the intermediate tolerance.




FIG. 4.4. Results for Example 2. Reconstruction of a test anomaly on a $32 \times 32 \times 32$ mesh with 225 sources and detectors, 27 basis functions, and using only the zero frequency. The SAA approach uses 15 random simultaneous sources and detectors. (a) True shape of the anomaly. (b) Reconstruction using all sources and all detectors. (c) Reconstruction using the SAA approach at the intermediate tolerance. (d) Reconstruction using the SAA approach after the maximum iterations. (e) Reconstruction with SAA and 2 optimized simultaneous sources and detectors. (f) Reconstruction with SAA and 4 optimized simultaneous sources and detectors.

approach, even after the maximum iterations; see Figure 4.4(d). In Figures 4.4(e)–(g), we show the reconstruction results when combining random and optimized simultaneous sources and detectors.

The straightforward inversion using all sources and detectors requires 9,225 large linear solves. Table 4.3 shows that our approach reduces the number of large linear solves by *about a factor 12* compared with using all sources and detectors, while approximating the original shape well. Clearly, there is a large improvement to be gained by using a small number of optimized simultaneous sources and detectors. For larger problems with many sources and detectors and using multiple frequencies, we expect much larger gains.

TABLE 4.3

Example 2 Results. The total number of iterations, function evaluations, Jacobian evaluations, and PDE solves required to reach the stopping criterion, $\|\mathbf{r}(\mathbf{p})\|_2^2 = \delta^2$. Parentheses indicate that the SAA approach does not reach the tolerance.

	Iteration number	Function evaluations	Jacobian evaluations	Total PDE solves	Tol
SAA* (intermediate tol)	4	5	6	132	δ
2 opt simult srcs/dets	12	13	6	726	δ^2
4 opt simult srcs/dets	9	10	3	762	δ^2
All srcs/all dets	25	26	15	9225	δ^2
SAA **	(99)	(100)	(67)	(2505)	δ^2

*The first row gives the costs to reach the intermediate tolerance for the SAA approach, $\|\mathbf{r}(\mathbf{p})\|_2^2 = \delta$.

**The SAA approach measuring the convergence with the true objective function.

Overall, our approach improves the rate of convergence of the optimization and reduces the number of large-scale linear system solves. Moreover, combining random and optimized simultaneous sources and detectors improves the quality of the inverse solution.

5. Conclusions and future work. We use the SAA approach to estimate the objective function, the Jacobian, and the gradient using only a few simultaneous random sources and detectors in DOT problems. While this approach is reasonably effective for the application in [12], it does not work quite that well for DOT. Since convergence slows down in later iterations before the noise level is reached, and the standard SAA approach regularly does not converge to the noise level, we propose using optimized simultaneous sources and detectors. With the addition of optimized directions, we observe faster convergence, good quality reconstructions, and robustness. This technique could be quite useful in other applications as well.

Several further improvements should be considered in the future. In particular, approximating the, typically low rank, Jacobian at low cost but sufficiently accurately to compute effective optimized simultaneous sources and detectors would lead to a further substantial reduction in the number of large linear solves. Potentially, computing such a low rank approximation can be combined more efficiently with the tensor form of the Jacobian.

Although our approach has proved successful experimentally, we aim to understand the underlying theory better. In the future, we plan to analyze, more fundamentally, what are the most effective simultaneous sources and detectors for fast convergence of the inverse problem: randomized, optimized (and in what sense), and their combination. An alternative approach to improve convergence, studied in multiple papers [25, 6, 5, 23, 24], is to slowly increase the sample size as the optimization progresses or dynamically choose the sample size. In future work, we plan to compare these approaches with the approach proposed in this paper. We also plan to test and evaluate SA approaches for inversion in DOT.

We intend to update the TREGS algorithm and study how small we can make the number of simultaneous sources and detectors (random and optimized) and still obtain good solutions and fast convergence. Moreover, finding more appropriate stopping criteria for the randomized approach may also improve our results.

As shown in [9], using parameterized interpolatory model reduction can also reduce the cost of inversion for DOT. We plan to combine model reduction with the randomized approach.

Acknowledgment. We sincerely thank the reviewers for many helpful suggestions that substantially improved this paper.

REFERENCES

- [1] A. AGHASI, M. KILMER, AND E. L. MILLER, *Parametric level set methods for inverse problems*, SIAM J. Imaging Sci., 4 (2011), pp. 618–650.
- [2] S. R. ARRIDGE, *Optical tomography in medical imaging*, Inverse Problems, 16 (1999), pp. R41–R93.
- [3] C. J. BEASLEY, *A new look at marine simultaneous sources*, Leading Edge, 27 (2008), pp. 914–917.
- [4] D. BOAS, D. BROOKS, E. MILLER, C. DIMARZIO, M. KILMER, R. GAUDETTE, AND Q. ZHANG, *Imaging the body with diffuse optical tomography*, IEEE Signal Process. Mag., 18 (2001), pp. 57–75.
- [5] R. BOLLAPRAGADA, R. H. BYRD, AND J. NOCEDAL, *Exact and inexact subsampled Newton methods for optimization*, IMA J. Numer. Anal., to appear.
- [6] R. H. BYRD, G. M. CHIN, J. NOCEDAL, AND Y. WU, *Sample size selection in optimization methods for machine learning*, Math. Program., 134 (2012), pp. 127–155.
- [7] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *A multilinear singular value decomposition*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1253–1278.
- [8] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *On the best rank-1 and rank- (R_1, R_2, \dots, R_N) approximation of higher-order tensors*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1324–1342.
- [9] E. DE STURLER, S. GUGERCIN, M. E. KILMER, S. CHATURANTABUT, C. BEATTIE, AND M. O’CONNELL, *Nonlinear parametric inversion using interpolatory model reduction*, SIAM J. Sci. Comput., 37 (2015), pp. B495–B517.
- [10] E. DE STURLER AND M. E. KILMER, *A regularized Gauss-Newton trust region approach to imaging in diffuse optical tomography*, SIAM J. Sci. Comput., 33 (2011), pp. 3057–3086.
- [11] J. DENNIS AND R. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Classics Appl. Math. 16, SIAM, Philadelphia, 1996.
- [12] E. HABER, M. CHUNG, AND F. HERRMANN, *An effective method for parameter estimation with PDE constraints with multiple right-hand sides*, SIAM J. Optim., 22 (2012), pp. 739–757.
- [13] N. HALKO, P. G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288.
- [14] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, 1991.
- [15] M. F. HUTCHINSON, *A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines*, Comm. Statist. Simulation Comput., 18 (1990), pp. 1059–1076.
- [16] M. E. KILMER AND E. DE STURLER, *Recycling subspace information for diffuse optical tomography*, SIAM J. Sci. Comput., 27 (2006), pp. 2140–2166.
- [17] T. G. KOLDA AND B. W. BADER, *Tensor decompositions and applications*, SIAM Rev., 51 (2009), pp. 455–500.
- [18] J. R. KREBS, J. E. ANDERSON, D. HINKLEY, R. NEELAMANI, S. LEE, A. BAUMSTEIN, AND M.-D. LACASSE, *Fast full-wavefield seismic inversion using encoded sources*, Geophysics, 74 (2009), pp. WCC177–WCC188.
- [19] P. M. KROONENBERG AND J. DE LEEUW, *Principal component analysis of three-mode data by means of alternating least squares algorithms*, Psychometrika, 1 (1980), pp. 69–97.
- [20] C. D. MEYER, *Matrix Analysis and Applied Linear Algebra*, SIAM, Philadelphia, 2000.
- [21] S. A. MORTON AND C. C. OBER, *Faster shot-record depth migrations using phase encoding*, Society of Exploration Geophysicists 68th Annual International Meeting, Expanded Abstracts, 37 (1998), pp. 1131–1134.
- [22] R. NEELAMANI, C. E. KROHN, J. R. KREBS, J. K. ROMBERG, M. DEFFENBAUGH, AND J. E. ANDERSON, *Efficient seismic forward modeling using simultaneous random sources and sparsity*, Geophysics, 75 (2010), pp. WB15–WB27.
- [23] F. ROOSTA-KHORASANI AND M. W. MAHONEY, *Sub-sampled Newton Methods I: Globally Convergent Algorithms*, preprint, <https://arxiv.org/abs/1601.04737> (2016).
- [24] F. ROOSTA-KHORASANI AND M. W. MAHONEY, *Sub-sampled Newton Methods II: Local Convergence Rates*, preprint, <https://arxiv.org/abs/1601.04738> (2016).
- [25] F. ROOSTA-KHORASANI, K. VAN DEN DOEL, AND U. ASCHER, *Stochastic algorithms for inverse problems involving PDEs and many measurements*, SIAM J. Sci. Comput., 36 (2014), pp. S3–S22.

- [26] A. SHAPIRO, D. DENTCHEVA, AND A. RUSZCZYŃSKI, *Lectures on Stochastic Programming: Modeling and Theory*, MPS-SIAM Ser. Optim. 9, SIAM, Philadelphia, 2009.
- [27] L. R. TUCKER, *Some mathematical notes on three-mode factor analysis*, *Psychometrika*, 31 (1966), pp. 279–311.
- [28] T. VAN LEEUWEN, A. Y. ARAVKIN, AND F. J. HERRMANN, *Seismic waveform inversion by stochastic optimization*, *Internat. J. Geophys.*, (2011), 689041.
- [29] C. R. VOGEL, *Computational Methods for Inverse Problems*, *Front. Appl. Math.* 23, SIAM, Philadelphia, 2002.