

COMPUTING REDUCED ORDER MODELS VIA INNER-OUTER KRYLOV RECYCLING IN DIFFUSE OPTICAL TOMOGRAPHY*

MEGHAN O'CONNELL[‡], MISHA E. KILMER[‡], ERIC DE STURLER[†], AND SERKAN GUGERCIN[†]

Abstract. In nonlinear imaging problems whose forward model is described by a partial differential equation (PDE), the main computational bottleneck in solving the inverse problem is the need to solve many large-scale discretized PDEs at each step of the optimization process. In the context of absorption imaging in diffuse optical tomography, one approach to addressing this bottleneck proposed recently (de Sturler, et al, 2015) reformulates the viewing of the forward problem as a differential algebraic system, and then employs model order reduction (MOR). However, the construction of the reduced model requires the solution of several full order problems (i.e. the full discretized PDE for multiple right-hand sides) to generate a candidate global basis. This step is then followed by a rank-revealing factorization of the matrix containing the candidate basis in order to compress the basis to a size suitable for constructing the reduced transfer function. The present paper addresses the costs associated with the global basis approximation in two ways. First, we use the structure of the matrix to rewrite the full order transfer function, and corresponding derivatives, such that the full order systems to be solved are symmetric (positive definite in the zero frequency case). Then we apply MOR to the new formulation of the problem. Second, we give an approach to computing the global basis approximation dynamically as the full order systems are solved. In this phase, only the incrementally new, relevant information is added to the existing global basis, and redundant information is not computed. This new approach is achieved by an inner-outer Krylov recycling approach which has potential use in other applications as well. We show the value of the new approach to approximate global basis computation on two DOT absorption image reconstruction problems.

Key words. Krylov recycling, diffuse optical tomography, model order reduction, nonlinear inverse problem

AMS subject classifications. 65F10, 65F22

1. Introduction. Nonlinear inverse problems are becoming ubiquitous but remain very expensive to solve, including medical image reconstruction, identification of anomalous regions such as unexploded ordinance, land mines, and contaminant plumes in the subsurface. In inverse problems we recover an image of an unknown quantity of interest inside a given medium, such as the light absorption coefficient in human tissue, using a mathematical model, the *forward model*, that relates the image of the unknown quantity to the measured data. In this paper, we consider the problem of imaging in diffuse optical tomography (DOT), where the forward model is a large-scale, discretized, partial differential equation describing the photon fluence/flux through tissue. Other imaging problems (e.g. electrical impedance/resistance tomography, hydraulic tomography) are similar in that the forward problems are described by PDEs.

The need to solve these large-scale, forward problems many times to recover two or three-dimensional images represents the largest computational impediment to effective, practical use of DOT. To this end, the use of reduced order modeling in the context of DOT imaging was considered in [13]. The key observation in that work is that function evaluations for the underlying optimization problem (which

*This material is based upon work supported by the National Science Foundation under Grants No. NSF-DMS 1025327, NSF DMS 1217156 and 1217161, NSF-DMS 0645347, and NIH R01-CA154774.

[†]Department of Mathematics, Virginia Tech, Blacksburg, VA 24061.

[‡]Department of Mathematics, Tufts University, Medford, MA 02115.

requires solving for the parameters defining the parametric level set image) may be viewed as transfer function evaluations along the imaginary axis, which motivates the use of system-theoretic model order reduction methods. Specifically, interpolatory parametric reduced models as surrogates for the full forward model were used. In the DOT setting, such surrogate models were found to approximate both the cost functional and the associated Jacobian with very little loss of accuracy while at the same time drastically reducing the cost of the overall inversion process.

Nevertheless, difficulties remain. In order to determine the global basis to be used for the interpolatory projection, several large-scale linear systems must still be solved. Because of the structure and size of the systems, the solves are typically done iteratively [9]. In recent work [2, 16, 17], the authors investigate the use of Krylov subspace recycling for these systems to generate the model-order reduction (MOR) basis. Recycling for shape based inversion for DOT was investigated in [21], but the goal was solving the sequence of systems in the inversion process, rather than for use in computing global basis matrices for use with MOR.

In the present work, we present several computational advances for applying MOR to this inverse problem, but note that *the results are applicable in a more general interpolatory MOR setting*. The paper is organized as follows. In Section 2, we give the system theoretic notation, introduction to DOT, and relevant information on computing reduced order models. In Section 3, we show that in a particular geometry and discretization, the transfer function may be reformulated as a transfer function for a slightly smaller symmetric problem. In Section 4, we consider innovations in leveraging Krylov recycling techniques to reduce the total amount of computation for the reduced global basis matrix by eliminating redundant information dynamically. An analysis is given in Section 5. Numerical results are given in Section 6 and conclusions and future work are outlined in Section 7.

2. Background. We begin by describing the forward model for DOT. Then, we express the forward problem in the systems theoretic notation that we will use throughout the paper. The first two sections borrow heavily from the presentation in our previous paper [13]. The third subsection gives background on the generation of the reduced transfer function, in preparation for the presentation of our new techniques for computing the reduced model global projection bases given in Sections 3 and 4.

2.1. The DOT Problem. Our image domain is a rectangular slab, $\Omega = [a_1, b_1] \times [a_2, b_2] \times [a_3, b_3]$, with the top ($x_3 = b_3$) and bottom ($x_3 = a_3$) surfaces denoted by $\partial\Omega_+$ and $\partial\Omega_-$, respectively. We use a diffusion model for the photon flux/fluence [6] $\eta(\mathbf{x}, t)$ driven by an input light source $g(\mathbf{x}, t)$. The input light source is one out of a set of n_{src} possible sources that are each physically stationary. This means there are functions, $b_j(\mathbf{x})$, $j = 1, \dots, n_{src}$ such that $g(\mathbf{x}, t) = b_j(\mathbf{x})u_j(t)$ for selected j . Additionally the observations are made with a limited number of n_{det} detectors observations. We use $m_i(t)$ to denote the presumed stationary observations located on the bottom surface. Using these definitions, the model for the diffusion and absorption

of light is described via

$$\frac{1}{\nu} \frac{\partial}{\partial t} \eta(\mathbf{x}, t) = \nabla \cdot (D(\mathbf{x}) \nabla \eta(\mathbf{x}, t)) - \mu(\mathbf{x}) \eta(\mathbf{x}, t) + b_j(\mathbf{x}) u_j(t), \quad \text{for } \mathbf{x} \in \Omega, \quad (2.1)$$

$$0 = \eta(\mathbf{x}, t) + 2 \mathcal{A} D(\mathbf{x}) \frac{\partial}{\partial \xi} \eta(\mathbf{x}, t), \quad \text{for } \mathbf{x} \in \partial\Omega_{\pm}, \quad (2.2)$$

$$0 = \eta(\mathbf{x}, t), \quad \text{for } x_1 = a_1 \text{ or } x_1 = b_1 \text{ or } x_2 = a_2 \text{ or } x_2 = b_2, \quad (2.3)$$

$$m_i(t) = \int_{\partial\Omega} c_i(\mathbf{x}) \eta(\mathbf{x}, t) d\mathbf{x} \quad \text{for } i = 1, \dots, n_{det} \quad (2.4)$$

(see [6, p. R56]). Here, the vector $\mathbf{x} = (x_1, x_2, x_3)^T$ refers to spatial location, \mathcal{A} is a constant defining the particular diffusive boundary reflection (see [6, p.R50]), and $D(\mathbf{x})$ and $\mu(\mathbf{x})$ denote diffusion and absorption coefficients, respectively. Also, ξ denotes the outward unit normal and ν is the speed of light in the medium.

The inverse problem consists of utilizing observations, $\mathbf{m}(t)$, made when the system is illuminated by a variety of source signals, $\mathbf{u}(t)$, to more accurately determine $D(\mathbf{x})$ and $\mu(\mathbf{x})$. For purposes of this work, we assume the diffusivity $D(\mathbf{x})$ is known (a common assumption in DOT breast tissue imaging) and that only the absorption field, $\mu(\mathbf{x})$, must be recovered. We also assume that the absorption field, $\mu(\cdot)$, although unknown, is expressible in terms of a finite set of parameters, $\mathbf{p} = [p_1, \dots, p_\ell]^T$. We will assume that parametric level sets (PaLS), developed in [1] and used in the context of DOT imaging in [13], have been used for $\mu(\cdot) = \mu(\cdot, \mathbf{p})$. The inverse problem is often solved by moving to the frequency domain and using frequency domain data (e.g. data using frequency modulated light). However, we adopt the approach in [13] of reformulating the forward problem in dynamical systems notation before moving to the frequency domain, so that we can see how to employ model reduction in the context of solving the inverse problem.

The discretization of (2.1)-(2.4) can be done by finite element or finite difference techniques. This gives the following differential-algebraic system,

$$\frac{1}{\nu} \mathbf{E} \dot{\mathbf{y}}(t; \mathbf{p}) = -\mathbf{A}(\mathbf{p}) \mathbf{y}(t; \mathbf{p}) + \mathbf{B} \mathbf{u}(t) \quad \text{with} \quad \mathbf{m}(t; \mathbf{p}) = \mathbf{C}^T \mathbf{y}(t; \mathbf{p}) \quad (2.5)$$

where \mathbf{y} denotes the discretized photon flux, $\mathbf{m} = [m_1, \dots, m_{n_{det}}]^T$ is the vector of detector outputs, $\mathbf{C}^T \mathbf{y}$ constitutes a set of quadrature rules with the approximate photon flux; the columns of \mathbf{B} are discretizations of the source ‘‘footprints’’ $b_j(\mathbf{x})$ for $j = 1, \dots, n_{src}$; $\mathbf{A}(\mathbf{p}) = \mathbf{A}_0 + \mathbf{A}_1(\mathbf{p})$ with \mathbf{A}_0 and $\mathbf{A}_1(\mathbf{p})$ discretizations of the diffusion and absorption terms, respectively ($\mathbf{A}_1(\mathbf{p})$ inherits the absorption field parametrization, $\mu(\cdot, \mathbf{p})$). \mathbf{E} is singular due to the inclusion of the discretized Robin condition (2.2) as an algebraic constraint. This fact will become important in Section 3.

Let $\hat{\mathbf{y}}(\omega; \mathbf{p})$, $\hat{\mathbf{u}}(\omega)$, and $\hat{\mathbf{m}}(\omega; \mathbf{p})$ denote the Fourier transforms of $\mathbf{y}(t; \mathbf{p})$, $\mathbf{u}(t)$, and $\mathbf{m}(t; \mathbf{p})$, respectively. Taking the Fourier transform of (2.5) and rearranging, we get

$$\hat{\mathbf{m}}(\omega; \mathbf{p}) = \Psi(\omega; \mathbf{p}) \hat{\mathbf{u}}(\omega) \quad \text{where} \quad \Psi(\omega; \mathbf{p}) = \mathbf{C}^T \left(\frac{i\omega}{\nu} \mathbf{E} + \mathbf{A}(\mathbf{p}) \right)^{-1} \mathbf{B}, \quad (2.6)$$

where $\omega \in \mathbb{R}$, and $\Psi(\omega; \mathbf{p})$ is known as the *frequency response* of the dynamical system defined in (2.5), though we will often refer to it as the transfer function¹.

¹In describing linear dynamical systems, usually the *transfer function* $\Psi(s; \mathbf{p}) = \mathbf{C}^T \left(\frac{s}{\nu} \mathbf{E} + \mathbf{A}(\mathbf{p}) \right)^{-1} \mathbf{B}$ is used where $s \in \mathbb{C}$ and is not restricted to the imaginary axis. Here, though, the measurements are made only on the imaginary axis and it is enough to take $s = i\omega$ with $\omega \in \mathbb{R}$.

For any absorption field, $\mu(\cdot, \mathbf{p})$, associated with \mathbf{p} , the vector of (estimated) observations for the i^{th} input source at frequency ω_j , as predicted by the forward model in the frequency domain, will be denoted as $\widehat{\mathbf{m}}_i(\omega_j; \mathbf{p}) \in \mathbb{C}^{n_{det}}$. Stacking the predicted observation vectors for all n_{src} sources and n_ω frequencies, we obtain

$$\mathcal{M}(\mathbf{p}) = [\widehat{\mathbf{m}}_1(\omega_1; \mathbf{p})^T, \dots, \widehat{\mathbf{m}}_1(\omega_{n_\omega}; \mathbf{p})^T, \widehat{\mathbf{m}}_2(\omega_1; \mathbf{p})^T, \dots, \widehat{\mathbf{m}}_{n_{src}}(\omega_{n_\omega}; \mathbf{p})^T]^T, \quad (2.7)$$

which is a (complex) vector of dimension $n_{det} \cdot n_{src} \cdot n_\omega$. We construct the corresponding empirical data vector, \mathbb{D} , from acquired data. The optimization problem that must be solved is

$$\min_{\mathbf{p} \in \mathbb{R}^\ell} \|\mathcal{M}(\mathbf{p}) - \mathbb{D}\|_2. \quad (2.8)$$

We note that mathematically, what was just described is equivalent to the usual approach for DOT imaging using frequency modulated data. What is different is using the dynamical systems interpretation, as was developed in [13], so that we can leverage efficiencies from a systems theoretic perspective on solving the parametric inversion problem.

2.2. A systems theoretic perspective on parametric inversion. From (2.6) and (2.7), it follows that a single evaluation of $\mathcal{M}(\mathbf{p}) - \mathbb{D}$ involves computing (for all i and j)

$$\widehat{\mathbf{m}}_i(\omega_j; \mathbf{p}) = \Psi(\omega_j) \widehat{\mathbf{u}}_i(\omega_j), \quad (2.9)$$

where $\Psi(\omega; \mathbf{p})$ is the frequency response defined in (2.6) and $\widehat{\mathbf{u}}_i(\omega_j) = e^{i\omega_j} \mathbf{e}_i$ where \mathbf{e}_i is the i^{th} column of the identity matrix; i.e., $\widehat{\mathbf{u}}_i(\omega_j)$ excites the i^{th} source location. The key observation is that as a result, an objective function evaluation at parameter vector \mathbf{p}_k requires solving the block linear systems

$$\left(\frac{i\omega_j}{\nu} \mathbf{E} + \mathbf{A}(\mathbf{p}_k) \right) \mathbf{Y}_{k,j} = \mathbf{B} \quad \mathbf{B} \in \mathbb{R}^{n \times n_{src}}, \quad j = 1, \dots, n_\omega. \quad (2.10)$$

Here,

$$\mathbf{Y}_{k,j} := [\widehat{\mathbf{y}}_1(\omega_j, \mathbf{p}_k), \dots, \widehat{\mathbf{y}}_{n_s}(\omega_j, \mathbf{p}_k)].$$

It is important to note that each column of the matrix \mathbf{B} is, in our application, a multiple of the i^{th} column of the identity matrix, corresponding to the i^{th} source location.

To solve the nonlinear inverse problem (2.8) for the parameters, the Jacobian is constructed using an adjoint-type (or co-state) approach that exploits the fact that the number of detectors is roughly equal to the number of sources, as discussed in [19] and [26, p. 88]. Using (2.6) and (2.9) and differentiating $\widehat{\mathbf{m}}_i(\omega_j; \mathbf{p})$ with respect to the k^{th} component of \mathbf{p} ,

$$\frac{\partial}{\partial p_k} \widehat{\mathbf{m}}_i(\omega_j; \mathbf{p}) = \frac{\partial}{\partial p_k} [\Psi(\omega_j; \mathbf{p})] \widehat{\mathbf{u}}_i(\omega_j) = -\mathbf{Z}(\omega_j; \mathbf{p})^T \frac{\partial}{\partial p_k} \mathbf{A}(\mathbf{p}) \widehat{\mathbf{y}}_i(\omega_j; \mathbf{p}), \quad (2.11)$$

where, for each ω_j and any \mathbf{p} we can compute $\mathbf{Z}(\omega_j; \mathbf{p})$ from

$$\left(\frac{i\omega_j}{\nu} \mathbf{E} + \mathbf{A}(\mathbf{p}) \right)^T \mathbf{Z}(\omega_j; \mathbf{p}) = \mathbf{C}, \quad \mathbf{C} \in \mathbb{R}^{n \times n_{src}}. \quad (2.12)$$

Here, the columns of \mathbf{C} also correspond to columns of the identity matrix, with a non-zero entry appearing at the j th detector location.

The matrices $\frac{\partial}{\partial p_k} \mathbf{A}(\mathbf{p})$ need to be computed only once for all parameters. Except for the cost of computing $\widehat{\mathbf{y}}_i(\omega_j; \mathbf{p})$ (which was completed already for the function evaluation), the computational cost of evaluating the Jacobian at a parameter vector, for all ω_j , consists mainly of the cost of computing the solution to the block systems (2.12).

In sum, the critical bottleneck in solving the inverse problem is associated with repeatedly solving (2.10) and (2.12).

2.3. Global Basis Computation. As done in [13] for the DOT problem, we therefore seek a surrogate function $\Psi_r(\omega; \mathbf{p})$ that is much cheaper to evaluate, yet provides a high-fidelity approximation to $\Psi(\omega; \mathbf{p})$ *over parameters and frequencies of interest*. The frequencies are dictated to us by the experimental set-up. The parameters of interest are those which would be chosen by the optimization routine if it were run using the full order model. Likewise, we require that $\nabla_{\mathbf{p}} \Psi_r(\omega; \mathbf{p})$ is much cheaper to evaluate and that $\nabla_{\mathbf{p}} \Psi(\omega; \mathbf{p}) \approx \nabla_{\mathbf{p}} \Psi_r(\omega; \mathbf{p})$ over the same range of arguments.

The surrogate parametric model is obtained using projection (see [10, 13]). Suppose $\mathbf{A} \in \mathbb{R}^{n \times n}$ and full rank matrices $\mathbf{V} \in \mathbb{C}^{n \times r}$ and $\mathbf{W} \in \mathbb{C}^{n \times r}$ are given. Assuming the full state $\mathbf{y}(t; \mathbf{p})$ evolves near the r -dimensional subspace $\text{Range}(\mathbf{V})$, one has $\mathbf{y}(t; \mathbf{p}) \approx \mathbf{V} \widehat{\mathbf{y}}(t; \mathbf{p})$. If we enforce a Petrov-Galerkin condition we obtain a reduced system with the reduced matrices given by

$$\mathbf{E}_r = \mathbf{W}^T \mathbf{E} \mathbf{V}, \quad \mathbf{A}_r(\mathbf{p}) = \mathbf{W}^T \mathbf{A}(\mathbf{p}) \mathbf{V}, \quad \mathbf{B}_r = \mathbf{W}^T \mathbf{B}, \quad \text{and} \quad \mathbf{C}_r = \mathbf{V}^T \mathbf{C}. \quad (2.13)$$

The reduced transfer function is then $\Psi_r = \mathbf{C}_r^T (\frac{i\omega}{\nu} \mathbf{E}_r + \mathbf{A}_r(\mathbf{p}))^{-1} \mathbf{B}_r$. The goal is to choose \mathbf{V}, \mathbf{W} so that the transfer function evaluated at \mathbf{p}, ω (likewise the Jacobian) will match (or be close to) the reduced transfer function evaluated at those parameters and frequencies of interest.

Referring to (2.10), we let

$$\mathbf{Y} := [\mathbf{Y}_{1,1}, \dots, \mathbf{Y}_{1,n_\omega}, \mathbf{Y}_{2,1}, \dots, \mathbf{Y}_{2,n_\omega}, \dots, \mathbf{Y}_{K,n_\omega}]. \quad (2.14)$$

Then, ideally, we define \mathbf{V} to correspond to the left singular vectors corresponding to the non-zero singular values of the matrix \mathbf{Y} . Similar steps are applied to construct \mathbf{W} from $\mathbf{Z}_{i,j}$ via (2.12) for $i = 1, \dots, K$ and $j = 1, \dots, n_\omega$. Taking the exact left singular vectors for non-zero singular values to define both \mathbf{V}, \mathbf{W} ensures [8] that the corresponding reduced transfer function will match the transfer function evaluation at *every* $(\omega, \mathbf{p}) = (\omega_j, \boldsymbol{\pi}_i)$ for $j = 1, \dots, n_\omega$ and $i = 1, \dots, K$. A similar result holds for the respective derivative computations.

The one-sided global basis approach, which is used in the context of DOT in [13], refers to using $\mathbf{V} \leftarrow [\mathbf{V}, \mathbf{W}]$ and $\mathbf{W} \leftarrow [\mathbf{V}, \mathbf{W}]$ in (2.13). Clearly, if $\mathbf{A}(\mathbf{p}), \mathbf{E}$ are symmetric (Hermitian), the reduced counterparts defined in (2.13) are also symmetric (Hermitian). For the use of model reduction in other optimization and inverse problem applications, we refer the reader to [5, 20, 22, 3, 4, 15, 7, 11, 28] and the references therein.

In the preceding description, we have assumed (as elsewhere in the literature), that several full order problems for the \mathbf{Y}_{ij} (\mathbf{Z}_{ij}) have been solved prior to selecting the global basis. In the context of solving our inverse problem, however, we do not a priori pick parameter values, but rather use the parameter sequence defined at the

start of the optimization to construct \mathbf{V} from the required full order model (FOM) solves. It was observed in [13] for our DOT application that then the singular values of \mathbf{Y} (\mathbf{Z}) fall off by several orders of magnitude. Thus, truncated SVDs are used to obtain the \mathbf{V} , \mathbf{W} which are subsequently concatenated to form one large projection matrix \mathbf{V} . This means that in first solving for all the \mathbf{Y}_{ij} , \mathbf{Z}_{ij} , we must have computed redundant information that we subsequently must squeeze out through via SVDs on the concatenated matrices. The purpose of this paper is to construct an approximate symmetric global basis matrix *without first solving each linear system in 2.10 and 2.12* so that we can also avoid the additional post processing step of computing truncated SVDs.

First, however, we show in §3 how we can capitalize on the structure of the system matrix in our case to design a ROM scheme on a slightly different system matrix/variables. Then in Section 4 we exploit this property to develop a particularly efficient algorithm and corresponding analysis for generating columns of the projection matrix \mathbf{V} .

3. Rewriting Transfer Function and Derivatives. In this section, we show that the structure of the linear system allows us to express the transfer function in terms of a symmetric positive definite matrix. This leads to a more efficient method for generating the global basis, and is of benefit in some theoretical arguments.

We follow the same discretization as in [21], in which we use second order centered differences away from the boundary, and first order discretization to implement the Robin boundary condition. The matrix $\frac{i\omega}{\nu}\mathbf{E} + \mathbf{A}(\mathbf{p})$ appearing in the definition of the full order transfer function (2.6) in the 2D case, has the following block structure:

$$\begin{bmatrix} \mathbf{G} & \mathbf{D}_1 \\ \mathbf{D}_2 & (\mathbf{F}(\mathbf{p}) + \frac{i\omega h^2}{\nu}\mathbf{I}) \end{bmatrix}, \quad (3.1)$$

where we have ordered the $N_x N_y$ unknowns such that the boundary unknowns (lexicographically ordered) appear first, followed by lexicographical ordering of internal points.

Furthermore, regarding the blocks in (3.1),

- \mathbf{G} is an invertible diagonal matrix,
- \mathbf{D}_1 has at most one nonzero per row, and these occur only in the first $N_x N_y$ and last $N_x N_y$ columns,
- \mathbf{D}_2 , although it has different entries, has the same sparsity pattern as \mathbf{D}_1^T .

We note that the matrix $\mathbf{A}(\mathbf{p})$ is not symmetric; however, as is shown in [21], the Schur complement (for the $\omega = 0$ case), which is given by $\mathbf{F}(\mathbf{p}) - \mathbf{D}_2 \mathbf{G}^{-1} \mathbf{D}_1$, is symmetric and positive definite. We will now investigate why this fact is particularly useful with regard to specification of the transfer function.

3.1. Transfer Function Revisited. The columns of \mathbf{C} and \mathbf{B} are scaled columns from an $N_x N_y \times N_x N_y$ identity matrix. Since sources and detectors appear on the boundary, this means that partitioning conformably with $\frac{i\omega}{\nu}\mathbf{E} + \mathbf{A}(\mathbf{p})$ as in (3.1) we obtain

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}_1 \\ \mathbf{0} \end{bmatrix}, \mathbf{B} = \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{0} \end{bmatrix}.$$

Since the sources and detectors are not co-located, it follows that $\mathbf{C}_1^T \mathbf{B}_1 = 0$.

Let us assume for ease of exposition that $\omega = 0$ in (3.1) above. In this case, it is the inverse of $\mathbf{A}(\mathbf{p})$ that appears in the definition of the transfer function. Let the

matrix $\mathbf{A}(\mathbf{p})^{-1}$ have the following block structure:

$$\begin{bmatrix} \mathbf{H} & \mathbf{S}_1 \\ \mathbf{S}_2 & \mathbf{N} \end{bmatrix}.$$

From $\mathbf{A}(\mathbf{p})\mathbf{A}(\mathbf{p})^{-1} = \mathbf{I}$, we have the following three expressions that will be helpful in rewriting the transfer function and in computing corresponding measurement derivatives:

$$\mathbf{H} = [\mathbf{G} - \mathbf{D}_1^T \mathbf{F}^{-1} \mathbf{D}_2]^{-1}, \quad (3.2)$$

$$\mathbf{S}_1 = -\mathbf{G}^{-1} \mathbf{D}_1 [\mathbf{F} - \mathbf{D}_2 \mathbf{G}^{-1} \mathbf{D}_1]^{-1}, \quad (3.3)$$

$$\mathbf{S}_2 = -[\mathbf{F} - \mathbf{D}_2 \mathbf{G}^{-1} \mathbf{D}_1]^{-1} \mathbf{D}_2 \mathbf{G}^{-1}. \quad (3.4)$$

It is straightforward to show that

$$\boldsymbol{\Psi}(0, \mathbf{p}) = \mathbf{C}^T (\mathbf{A}(\mathbf{p}))^{-1} \mathbf{B} = \mathbf{C}_1^T \mathbf{H} \mathbf{B}_1.$$

Now $\mathbf{F}(\mathbf{p}) = \mathbf{L} + \text{diag}(\mu(\mathbf{p}))$, where \mathbf{L} is the discretization of the Laplacian at the internal nodes multiplied by the (constant) diffusion coefficient, and $\text{diag}(\mu(\mathbf{p}))$ is a non-negative diagonal matrix. Thus $\mathbf{F}(\mathbf{p})$ is SPD, so we express it in terms of its eigendecomposition, $\mathbf{F} = \mathbf{Q} \boldsymbol{\Lambda} \mathbf{Q}^T$, so that

$$\begin{aligned} \mathbf{H} &= [\mathbf{G} - \mathbf{D}_1^T \mathbf{Q} \boldsymbol{\Lambda}^{-1/2} \boldsymbol{\Lambda}^{-1/2} \mathbf{Q}^T \mathbf{D}_2]^{-1} \\ &= \mathbf{G}^{-1} + \mathbf{G}^{-1} \mathbf{D}_1^T [\mathbf{F} - \mathbf{D}_2 \mathbf{G}^{-1} \mathbf{D}_1]^{-1} \mathbf{D}_2 \mathbf{G}^{-1}, \end{aligned}$$

by the Sherman-Morrison-Woodbury formula.

Since the sources and detectors are not co-located, and because \mathbf{G} is diagonal, it follows that

$$\boldsymbol{\Psi}(0, \mathbf{p}) = \underbrace{\mathbf{C}_1^T \mathbf{G}^{-1} \mathbf{D}_1^T}_{\tilde{\mathbf{C}}^T} \underbrace{[\mathbf{F} - \mathbf{D}_2 \mathbf{G}^{-1} \mathbf{D}_1]^{-1}}_{\tilde{\mathbf{A}}(\mathbf{p})^{-1}} \underbrace{\mathbf{D}_2 \mathbf{G}^{-1} \mathbf{B}_1}_{\tilde{\mathbf{B}}}.$$

Importantly, the structure of the matrices involved means that $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{B}}$ maintain the same structure as \mathbf{C}_1 and \mathbf{B}_1 : namely, they contain multiples of columns of the $N_x N_y$ identity matrix, so those matrices can be considered as ‘effective’ sources and receivers.

Using a similar argument, it is also straightforward to show that if $\omega \neq 0$,

$$\boldsymbol{\Psi}(\omega, \mathbf{p}) = \tilde{\mathbf{C}}^T \left(\frac{i\omega}{\nu} \mathbf{I} + \tilde{\mathbf{A}}(\mathbf{p}) \right)^{-1} \tilde{\mathbf{B}}^T.$$

This means that the transfer function (for the 0 frequency case) is in fact expressed using an SPD matrix to represent the system matrix (complex symmetric if ω is non-zero).

3.2. Systems Theoretic Interpretation. Even though the new expression in the last section was derived from matrix analysis, there is also a systems theoretic interpretation that will lead to the same expression.

Specifically, using the same reordering of unknowns as in the previous section, we observe that the singular matrix \mathbf{E} has the structure

$$\begin{bmatrix} 0 & 0 \\ 0 & \mathbf{I} \end{bmatrix},$$

so that the system is a differential-algebraic system and, due to the structure of \mathbf{E} , of index 1. If we partition the state vector conformably with the matrices $\mathbf{A}(\mathbf{p})$, \mathbf{E} , and \mathbf{B} , then $\mathbf{E}\dot{\mathbf{y}} = \mathbf{A}(\mathbf{p})\mathbf{y} + \mathbf{B}\mathbf{u}$ becomes

$$0 = \mathbf{G}\mathbf{y}_1 + \mathbf{D}_1\mathbf{y}_2 \quad (3.5)$$

$$\dot{\mathbf{y}}_2 = \mathbf{D}_2\mathbf{y}_1 + \mathbf{F}\mathbf{y}_2. \quad (3.6)$$

If we solve for \mathbf{y}_1 in (3.5) and plug the result into (3.6) and rearrange, we obtain

$$\dot{\mathbf{y}}_2 = \tilde{\mathbf{A}}(\mathbf{p})\mathbf{y}_2 - \tilde{\mathbf{B}}\mathbf{u}.$$

Now the same calculations in the previous subsection lead us to

$$\mathbf{m}(t) = \mathbf{C}^T \mathbf{y} = \mathbf{C}_1^T \mathbf{y}_1 = -\tilde{\mathbf{C}}^T \mathbf{y}_2.$$

Posing the two previous equations in the Fourier domain yields the transfer function

$$\Psi(s, \mathbf{p}) = \tilde{\mathbf{C}}^T (s\mathbf{I} - \tilde{\mathbf{A}}(\mathbf{p}))^{-1} \tilde{\mathbf{B}}.$$

This manipulation here corresponds to decomposing the transfer function of a DAE as the sum of the strictly proper part and polynomial part. A satisfactory reduced model should match the polynomial component exactly. In the interpolatory MOR setting, this means that the strictly proper part of the reduced transfer function interpolates that of the full-order model. For details on interpolatory MOR of DAEs, we refer the reader to [18]. For the DOT problem we consider here, this derivation illustrates that the transfer function of the index-1 DAE does not contain any polynomial part and is strictly proper.

3.3. Derivative Computation Revisited. Again using $\omega = 0$ for simplicity, define the $n_{src} \times n_{det}$ matrix

$$M(0, \mathbf{p}) = [\hat{\mathbf{m}}_1(0, \mathbf{p}), \dots, \hat{\mathbf{m}}_{n_{det}}(0, \mathbf{p})].$$

We use the “vec” command to map a matrix in $\mathbb{R}^{m_1 \times m_2}$ to a vector in $\mathbb{R}^{m_1 m_2}$ by unstacking the columns of the argument from left to right. Note that the vector $\text{vec}(\frac{\partial}{\partial p_k} M) \in \mathbb{R}^{n_{src} n_{det}}$ gives the k^{th} column of the Jacobian matrix. Using (2.10 - 2.12), we can write

$$\frac{\partial}{\partial p_k} M(\mathbf{p}) = -\mathbf{C}^T \mathbf{A}(\mathbf{p})^{-1} \frac{\partial}{\partial p_k} \mathbf{A}(\mathbf{p}) \mathbf{A}(\mathbf{p})^{-1} \mathbf{B}.$$

This calculation can be carried out with only reference to $\tilde{\mathbf{A}}(\mathbf{p})$, as follows.

Since the boundary terms have no absorption, then under the same finite difference discretization scheme and ordering of unknowns as before,

$$\frac{\partial}{\partial p_k} \mathbf{A}(\mathbf{p}) = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \Delta \end{bmatrix},$$

for a diagonal matrix $\Delta := \frac{\partial}{\partial p_k} \tilde{\mathbf{A}}(\mathbf{p})$. Thus,

$$\begin{aligned} \frac{\partial}{\partial p_k} M(0, \mathbf{p}) &= - [\mathbf{C}_1^T \quad \mathbf{0}] \begin{bmatrix} \mathbf{H} & \mathbf{S}_1 \\ \mathbf{S}_2 & \mathbf{N} \end{bmatrix} \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \Delta \end{bmatrix} \begin{bmatrix} \mathbf{H} & \mathbf{S}_1 \\ \mathbf{S}_2 & \mathbf{N} \end{bmatrix} \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{0} \end{bmatrix} \\ &= -\mathbf{C}_1^T \mathbf{S}_1 \Delta \mathbf{S}_2 \mathbf{B}_1. \end{aligned}$$

Now using (3.3) and (3.4), we have

$$\frac{\partial}{\partial p_k} M(0, \mathbf{p}) = - \underbrace{\mathbf{C}_1^T \mathbf{G}^{-1} \mathbf{D}_1}_{\tilde{\mathbf{C}}^T} \underbrace{[\mathbf{F} - \mathbf{D}_2 \mathbf{G}^{-1} \mathbf{D}_1]^{-1}}_{\tilde{\mathbf{A}}(\mathbf{p})^{-1}} \Delta \underbrace{[\mathbf{F} - \mathbf{D}_2 \mathbf{G}^{-1} \mathbf{D}_1]^{-1}}_{\tilde{\mathbf{A}}(\mathbf{p})^{-1}} \underbrace{\mathbf{D}_2 \mathbf{G}^{-1} \mathbf{B}_1}_{\tilde{\mathbf{B}}}. \quad (3.7)$$

This means that the necessary derivatives, for the 0 frequency case, can be computed from the same SPD matrix as the transfer function.

3.4. A New View on Generating the ROM. We have expressed the original transfer function and corresponding derivatives in terms of an SPD (for $\omega = 0$) matrix of size $(N_y - 2)N_x \times (N_y - 2)N_x$. So, we look for a ROM corresponding to this slightly smaller system matrix and its system variables.

Following the discussion in Section 2 of the one-sided global basis projection approach, we need $\mathbf{V} \in \mathbb{C}^{n \times r}$ and we define²

$$\tilde{\mathbf{E}}_r = \mathbf{V}^T \mathbf{V}, \quad \tilde{\mathbf{A}}_r(\mathbf{p}) = \mathbf{V}^T \tilde{\mathbf{A}}(\mathbf{p}) \mathbf{V}, \quad \tilde{\mathbf{B}}_r = \mathbf{V}^T \tilde{\mathbf{B}}, \quad \tilde{\mathbf{C}}_r = \mathbf{V}^T \tilde{\mathbf{C}}; \quad (3.8)$$

so, the reduced transfer function is $\tilde{\Psi}_r = \tilde{\mathbf{C}}_r^T (\frac{i\omega}{\nu} \tilde{\mathbf{E}}_r + \tilde{\mathbf{A}}_r(\mathbf{p}))^{-1} \tilde{\mathbf{B}}_r$.

Since $\tilde{\mathbf{A}}$ is SPD we do not need to solve the forward and adjoint problems separately to generate \mathbf{V} . Instead, we can solve

$$\left(\tilde{\mathbf{A}}(\mathbf{p}_k) + \frac{i\omega_j}{\nu} \mathbf{I} \right) \mathbf{x}_{k,j} = [\tilde{\mathbf{B}}, \tilde{\mathbf{C}}] \quad (3.9)$$

for appropriate choices of parameters \mathbf{p}_k , $k = 1, \dots, K$ and frequencies $\omega_j, j = 1, \dots, \nu_\omega$.

For the remainder of the paper, we will assume that only data for $\omega = 0$ has been provided. This helps with the remaining exposition of our new approach. Indeed, only results for the zero frequency case are provided in [13]. Moreover, treatment of non-zero ω is not a trivial extension of the algorithm provided, and in the interest of both paper focus and space, we relegate extensions for non-zero ω to forthcoming work.

The objective in reduced order modeling is to create a surrogate transfer function, $\tilde{\Psi}_r(\omega; \mathbf{p})$, that provides a high-fidelity approximation to $\tilde{\Psi}(\omega; \mathbf{p})$ as well as ensuring $\nabla_{\mathbf{p}} \tilde{\Psi}(\omega; \mathbf{p}) \approx \nabla_{\mathbf{p}} \tilde{\Psi}_r(\omega; \mathbf{p})$. The following theorem, which follows from [8], shows how to construct \mathbf{V} to guarantee this for the symmetric DOT-PaLs in the zero frequency case.

THEOREM 3.1. *Suppose $\tilde{\mathbf{A}}(\mathbf{p})$ is continuously differentiable in a neighborhood of $\hat{\mathbf{p}} \in \mathbb{R}^l$. Let both $\tilde{\mathbf{A}}(\hat{\mathbf{p}})$ and $\tilde{\mathbf{A}}_r(\hat{\mathbf{p}})$ be invertible. If $\tilde{\mathbf{A}}(\hat{\mathbf{p}})^{-1} \mathbf{B}$ and $(\mathbf{C} \tilde{\mathbf{A}}(\hat{\mathbf{p}})^{-1})^T$ are in $\text{range}(\mathbf{V})$, then the reduced parametric model satisfies $\tilde{\Psi}(0; \hat{\mathbf{p}}) = \tilde{\Psi}_r(0; \hat{\mathbf{p}})$ and $\nabla_{\mathbf{p}} \tilde{\Psi}(0; \hat{\mathbf{p}}) = \nabla_{\mathbf{p}} \tilde{\Psi}_r(0; \hat{\mathbf{p}})$.*

²The \mathbf{V} as we generate it will typically not have orthonormal columns, hence the need to specify $\tilde{\mathbf{E}}_r$.

In accordance with the discussion at the end of §2.3, the approach to obtaining \mathbf{V} for the solution to the DOT problem would consist of the following steps. Note that what follows is the approach of [13], but modified for our newly formulated transfer function representation.

Algorithm 1: Generate Symmetric Global Basis via Truncated SVD

1. Solve the systems (3.9) for $k = 1, \dots, K_*$ for the first \mathbf{p}_k parameter vectors produced by the optimization;
2. Concatenate the block solutions into a large block matrix

$$[\mathbf{X}_{1,0}, \mathbf{X}_{2,0}, \dots, \mathbf{X}_{K,0}].$$

3. Set \mathbf{V} to be the matrix of the first r left singular vectors of the above matrix. This gives a reduced order model, (3.8), of dimension r .
-

On the one hand, we need to compute each $\mathbf{X}_{k,0}$, because we need this to compute function and Jacobian evaluations at steps 1 through K_* of the optimization problem. Clearly, as these are large and sparse systems with SPD matrices, we should employ a Krylov subspace algorithm to solve the individual systems. On the other hand, for suitable K_* , some of the information that we generate and put into the concatenated matrix is redundant (or nearly so), as evidenced by the rapid decay of the singular values of the concatenated matrix. This means in terms of generating a global basis matrix \mathbf{V} , we have computed information that we don't really need (adding to the cost) and we thus incur the cost of the postprocessing via rank-revealing information.

Therefore, in the next section, we propose a method that is iterative in nature and which has the two-fold advantage of minimizing the work for computing only what we need, in terms of a) approximating the $\mathbf{X}_{k,0}(:, j)$ where it's needed for the optimization, b) generating an approximate global basis matrix by augmenting an initial estimate with only the most non-redundant information. The second part means we eliminate the need for step 3 above in the process for computing \mathbf{V} , because we build \mathbf{V} up to have r columns as we go, rather than overbuild and then compress to r terms.

4. Inner-Outer Krylov Recycling. Summarizing the previous section, we see that with $\tilde{\mathbf{A}}_k := \tilde{\mathbf{A}}(\mathbf{p}_k)$, $\mathbf{X}_k := \mathbf{X}_{k,0}$, $\mathbf{B} := [\tilde{\mathbf{B}}, \tilde{\mathbf{C}}]$, the relevant information for generating the basis comes from the systems

$$\tilde{\mathbf{A}}_k \mathbf{X}_k = \mathbf{B}, \tag{4.1}$$

for several values of k in order to determine the global basis. (Note that the matrix \mathbf{B} has a new definition from that of Section 2.) Moreover, since

$$\tilde{\mathbf{A}}(\mathbf{p}) = \underbrace{(\mathbf{L} - \mathbf{D}_2 \mathbf{G}^{-1} \mathbf{D}_1)}_{\tilde{\mathbf{A}}_*} + \text{diag}(\mu(\mathbf{p})), \tag{4.2}$$

the changes to the matrix as a function of parameter are restricted to the diagonal.

Recycling for a sequence of systems of the form (4.1) was shown to be efficient in [21] in the context of optimization for shape parameters in diffuse optical tomographic imaging. In that work, the authors solve (4.1) for all parameters selected during the course of the optimization, using different recycle spaces for each right-hand side.

However, the focus of this paper is different since we not only want to solve the (shorter) sequence of full order problems efficiently, we need to construct this global basis, and we want to do it in such a way that near redundant information is detected on the fly. In this section, we first cover the basic information about recycling for a single system, then introduce our new inner-outer recycling method which avoids adding nearly redundant information to \mathbf{V} and having to discard it (through a rank-revealing factorization) afterwards.

4.1. Recycling Basics. For this discussion, consider the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$, with symmetric $\mathbf{A} \in \mathbb{R}^{N \times N}$ and $\mathbf{b} \in \mathbb{R}^N$. Let $\mathbf{U} \in \mathbb{R}^{N \times n_c}$, be given such that $\mathbf{A}\mathbf{U} = \mathbf{K}$ and $\mathbf{K}^T\mathbf{K} = \mathbf{I}$. The approximate solution in $\text{Range}(\mathbf{U})$ that minimizes the 2-norm of the residual is

$$\mathbf{z} = \mathbf{U}\mathbf{K}^T\mathbf{b}, \quad (4.3)$$

which yields the residual $\mathbf{r} = \mathbf{b} - \mathbf{K}\mathbf{K}^T\mathbf{b}$ that is orthogonal to $\text{Range}(\mathbf{K})$. If this solution is not adequate, we expand the subspace as follows [12]. Let $\mathbf{v}_1 = (\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{b} / \|(\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{b}\|_2$ be the normalized initial residual. We use a Lanczos recurrence with $(\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{A}$ and \mathbf{v}_1 to generate the recurrence relation³

$$\begin{aligned} (\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{A}\mathbf{V}_m &= \mathbf{V}_{m+1}\underline{\mathbf{T}}_m \Leftrightarrow \\ \mathbf{A}\mathbf{V}_m &= \mathbf{K}\mathbf{K}^T\mathbf{A}\mathbf{V}_m + \mathbf{V}_{m+1}\underline{\mathbf{T}}_m, \end{aligned} \quad (4.4)$$

where $\underline{\mathbf{T}}_m$ is $(m+1) \times m$ tridiagonal, since \mathbf{A} is symmetric. Next, we compute the approximate solution in $\text{Range}([\mathbf{V}_m \ \mathbf{U}])$ that minimizes the 2-norm of the residual, $\|\mathbf{b} - \mathbf{A}(\mathbf{V}_m\mathbf{y} + \mathbf{U}\mathbf{z})\|_2$, as follows:

$$\begin{aligned} & \min_{\mathbf{y}, \mathbf{z}} \left\| \mathbf{b} - \mathbf{A}[\mathbf{U} \ \mathbf{V}_m] \begin{bmatrix} \mathbf{z} \\ \mathbf{y} \end{bmatrix} \right\|_2 \\ &= \min_{\mathbf{y}, \mathbf{z}} \left\| \mathbf{b} - [\mathbf{K} \ \mathbf{V}_{m+1}] \begin{bmatrix} \mathbf{I} & \mathbf{K}^T\mathbf{A}\mathbf{V}_m \\ 0 & \underline{\mathbf{T}}_m \end{bmatrix} \begin{bmatrix} \mathbf{z} \\ \mathbf{y} \end{bmatrix} \right\|_2 \\ &= \min_{\mathbf{y}, \mathbf{z}} \left\| \begin{bmatrix} \mathbf{K}^T\mathbf{b} \\ \xi\mathbf{e}_1 \end{bmatrix} - \begin{bmatrix} \mathbf{I} & \mathbf{K}^T\mathbf{A}\mathbf{V}_m \\ 0 & \underline{\mathbf{T}}_m \end{bmatrix} \begin{bmatrix} \mathbf{z} \\ \mathbf{y} \end{bmatrix} \right\|_2, \end{aligned} \quad (4.5)$$

where \mathbf{e}_1 denotes the first Cartesian basis vector in \mathbb{R}^{m+1} and $\xi = \|(\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{b}\|_2$. The minimization in (4.5) corresponds to a small least squares problem, whose solution requires only the QR decomposition of the submatrix $\underline{\mathbf{T}}_m$, which can be easily updated from the QR decomposition at step $m-1$. The block nature of the system means that the solution can be obtained by a three-step process: find \mathbf{y} that solves the projected minimization problem

$$\min_{\mathbf{y}} \|\underline{\mathbf{T}}_m\mathbf{y} - \xi\mathbf{e}_1\|_2$$

compute \mathbf{z} that satisfies (4.5), set $\mathbf{x} = \mathbf{V}_m\mathbf{y} + \mathbf{U}\mathbf{z}$. In reality, $\mathbf{y}_m := \mathbf{V}_m\mathbf{y}$ is computed via short term recurrences (MINRES), so the \mathbf{V}_m are not explicitly stored (see also [27, 23]).

³We note that the matrix in the recurrence is formally equivalent to $(\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{A}(\mathbf{I} - \mathbf{K}\mathbf{K}^T)$, but the right-most projector is suppressed for clarity.

4.2. Global Basis Construction. We start by trying to kill two birds with one stone: if we have an estimate of the global basis matrix \mathbf{V} , we can investigate its use as a candidate recycling space. If the basis can be improved, we have more full order systems to solve and we should consider solving the remaining systems by recycling. However, \mathbf{V} may already have too many columns for it to prove computationally feasible to use as a recycle space.

4.2.1. Recycling on the Updated Equations. To keep the notation simple and consistent with the previous subsection, let k be fixed and set $\mathbf{A} := \tilde{\mathbf{A}}(\mathbf{p}_k)$, $\mathbf{b}_j = \tilde{\mathbf{B}}(:, j)$, and $\mathbf{x}_j = \tilde{\mathbf{X}}_k(:, j)$. Next, assume $\mathbf{A}\mathbf{V} = \tilde{\mathbf{K}}$, find the QR factorization $\tilde{\mathbf{K}} = \mathbf{K}\mathbf{R}$, and then set⁴ $\mathbf{U} = \mathbf{V}\mathbf{R}^{-1}$, so that $\mathbf{A}\mathbf{U} = \mathbf{K}$.

According to the previous recycling discussion, the optimal solution in $\text{Range}(\mathbf{U}) = \text{Range}(\mathbf{V})$ is $\mathbf{z} = \mathbf{U}\mathbf{K}^T\mathbf{b}_j = \mathbf{V}(\mathbf{R}^{-1}\mathbf{K}^T\mathbf{b}_j)$, and the initial residual is $(\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{b}_j$. If this initial residual is small in a relative sense, then there is no need to go further, we need no iteration. If the initial residual is not small enough, then according to the previous section, we should expand the search space by running Lanczos to form a basis for the Krylov subspace generated by the projected matrix $(\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{A}$ and projected right-hand side $(\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{b}_j$. With this basis in the columns of \mathbf{V}_m , we find an approximate solution in $\text{Range}([\mathbf{U}, \mathbf{V}_m])$. However, unless the number of columns in \mathbf{V} , hence \mathbf{K} , is relatively small, we cannot afford to do this, because the re-orthogonalization is too expensive! Moreover, this presupposes that we actually need to find an approximation to \mathbf{x} directly. In fact, as far as updating our global basis approximation is concerned, we need only information *that is not already reconstructable from $\text{Range}(\mathbf{V}) = \text{Range}(\mathbf{U})$* .

An important fact comes to light if we decompose \mathbf{b}_j using the orthogonal projector $\mathbf{K}\mathbf{K}^T$:

$$\begin{aligned} \mathbf{A}\mathbf{x}_j &= (\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{b}_j + \mathbf{K}\mathbf{K}^T\mathbf{b}_j \\ \mathbf{A}\mathbf{x}_j - \mathbf{K}\mathbf{K}^T\mathbf{b}_j &= (\mathbf{I} - \mathbf{K}\mathbf{K}^T)\mathbf{b}_j \\ \mathbf{A}\mathbf{x}_j - \mathbf{A}\mathbf{U}\mathbf{K}^T\mathbf{b}_j &= \mathbf{r}_j \\ \mathbf{A} \underbrace{(\mathbf{x}_j - \mathbf{U}\mathbf{K}^T\mathbf{b}_j)}_{\mathbf{g}_j} &= \mathbf{r}_j. \end{aligned} \tag{4.6}$$

The vector \mathbf{g}_j is the correction to the initial guess $\mathbf{U}\mathbf{K}^T\mathbf{b}_j$. So, to obtain the incremental information to construct the global basis matrix, we should consider an iterative solution to (4.6). As we have one such system for right-hand sides $j = 1, \dots, n_{src} + n_{det}$, and a sequence of these as we iterate over the parameters, we can employ a recycling type of approach, if we choose our recycle space carefully.

Specifically, suppose \mathbf{r}_j is not already suitably small, so we want to solve (4.6), or alternatively, we want to find

$$\min_{\mathbf{g}_j \in \mathcal{S}} \|\mathbf{r}_j - \mathbf{A}\mathbf{g}_j\|_2,$$

for suitable subspace \mathcal{S} ⁵. We cannot afford to use all the columns of \mathbf{V} as a recycle space in generating a \mathbf{V}_m to use, because of the cost of the orthogonalization against $\text{Range}(\mathbf{K})$. Instead, we will use a subset of the columns for right-hand side j to be the recycle space, which we'll immediately need to expand.

⁴In practice, \mathbf{U} is formed without inverting \mathbf{R} explicitly.

⁵Note that \mathcal{S} should not be $\text{Range}(\mathbf{U})$, as then the solution is zero.

Specifically, this amounts to finding $\mathbf{U}_j \in \mathbb{R}^{N \times n_j}$, $\mathbf{U}_j \subseteq \text{Range}(\mathbf{U})$, and such that $\mathbf{A}\mathbf{U}_j = \mathbf{K}_j$, where $\mathbf{K}_j^T \mathbf{K}_j = \mathbf{I}$. Now we use $\mathcal{S} = \text{Range}([\mathbf{U}_j, \mathbf{V}_m^{(j)}])$ where the $\mathbf{V}_m^{(j)}$ are the Lanczos vectors for $K_m((\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T)\mathbf{A}, (\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T)\mathbf{r}_j)$ (compare to (4.4)). That is, we want to solve

$$\min_{\mathbf{z}, \mathbf{y}} \left\| \mathbf{r}_j - \mathbf{A}[\mathbf{V}_m^{(j)}, \mathbf{U}_j] \begin{bmatrix} \mathbf{y} \\ \mathbf{z} \end{bmatrix} \right\|.$$

Importantly, this choice for \mathbf{U}_j gives $(\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T)\mathbf{r}_j = \mathbf{r}_j$. Thus, we observe $\mathbf{v}_1^{(j)} = \mathbf{r}_j / \|\mathbf{r}_j\|$. Next we use the Lanczos recurrence with $(\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T)\mathbf{A}$ and $\mathbf{v}_1^{(j)}$ to generate the recurrence relation

$$\begin{aligned} (\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T)\mathbf{A}\mathbf{V}_m^{(j)} &= \mathbf{V}_{m+1}^{(j)}\underline{\mathbf{T}}_m^{(j)} \Leftrightarrow \\ \mathbf{A}\mathbf{V}_m^{(j)} &= \mathbf{K}_j \mathbf{K}_j^T \mathbf{A}\mathbf{V}_m^{(j)} + \mathbf{V}_{m+1}^{(j)}\underline{\mathbf{T}}_m^{(j)}. \end{aligned} \quad (4.7)$$

Now \mathbf{y}, \mathbf{z} are found by (compare to (4.5)) solving

$$\min_{\mathbf{y}, \mathbf{z}} \left\| \begin{bmatrix} 0 \\ \xi \mathbf{e}_1 \end{bmatrix} - \begin{bmatrix} \mathbf{I} & \mathbf{K}_j^T \mathbf{A}\mathbf{V}_m^{(j)} \\ 0 & \underline{\mathbf{T}}_m^{(j)} \end{bmatrix} \begin{bmatrix} \mathbf{z} \\ \mathbf{y} \end{bmatrix} \right\|_2.$$

Hence, $\mathbf{g}_j = \mathbf{V}_m^{(j)}\mathbf{y} + \mathbf{U}_j\mathbf{z}$, where $\mathbf{z} = -\mathbf{K}_j^T \mathbf{A}\mathbf{V}_m^{(j)}\mathbf{y}$ and $\mathbf{y}_m^{(j)} := \mathbf{V}_m^{(j)}\mathbf{y}$ is generated by a short term recurrence. So $\mathbf{x}_j = \mathbf{y}_m^{(j)} - \mathbf{U}_j \mathbf{K}_j^T \mathbf{A} \mathbf{y}_m^{(j)} + \mathbf{U} \mathbf{K}^T \mathbf{b}_j$. Since $\text{Range}(\mathbf{U}_j) \subseteq \text{Range}(\mathbf{U}) = \text{Range}(\mathbf{V})$, the relevant incremental information about \mathbf{x}_j that cannot already be expressed using the columns of \mathbf{V} is $\mathbf{V}_m^{(j)}\mathbf{y}$. Therefore, we extend the global basis with this vector. We repeat this process for any j for which the initial residual \mathbf{r}_j is not already small enough⁶. At most, for a fixed k , we have added $n_{src} + n_{det}$ columns to \mathbf{V} , but in theory, we may add substantially fewer.

Now we consider the remaining issues: a) what to use as the initial guess to the global basis \mathbf{V} and b) how to choose the individual columns to specify the \mathbf{U}_j when we do need to solve (4.6).

4.2.2. Identifying the Recycling Spaces. In [21] when working on a different parametric inverse model problem for DOT, the authors observed that a good recycle subspace for one right-hand side did not necessarily make a good recycle space for the next right-hand side. Instead, they employed a different recycle space for each right-hand side in the sequence, with the caveat that the recycle spaces did have a common subspace pertaining to a specific (approximate) invariant subspace. That invariant subspace corresponded to the smallest eigenvalues, because that subspace remained relatively unchanged through the optimization process.

Even though in the present paper we are using a different image parameterization than what was used in [21], we also observe that the invariant subspace due to the smallest several eigenvalues for the $\tilde{\mathbf{A}}(\mathbf{p}_k)$ remains unchanged. Thus, we adopt the same approach here. We use a different recycle space \mathbf{U}_j for each right-hand side, but seed each with the same invariant subspace, plus right-hand-side specific information.

First, we compute (approximate) eigenvectors of $\tilde{\mathbf{A}}_0$ that correspond to the smallest eigenvalues. The small eigenvalues of the $\tilde{\mathbf{A}}_k$ matrices remain close from one system to the next suggesting that the corresponding invariant subspaces also remain

⁶The integer m for which the solution estimate is good enough will vary depending on the system – that is, $m = m_j$ – but for ease in notation we have omitted the subscript on m .

close. We refer the reader to [21] for more details and theory. We have found experimentally that 10 eigenvectors is sufficient for the invariant subspace, while keeping the recycle space small. So, we set $\mathbf{U}_0 \in \mathbb{R}^{n \times 10}$ to contain (estimates of) those 10 vectors. Spending the computational resources to get an accurate invariant subspace so that it can be deflated from the right-hand side has been observed to be worthwhile in other large-scale applications, such as QCD, as well [25, 24].

Initially, we let $\mathbf{V} = ([\mathbf{U}_0, \mathbf{X}_0])$ and $\mathbf{U}_j = [\mathbf{U}_0, \mathbf{X}_0(:, j)]$. Note that \mathbf{V} has the initial solutions to all the right-hand sides, while \mathbf{U}_j has only the solution from the j^{th} right-hand side. If we find that \mathbf{r}_j is not suitably small, perform the recycling outlined above on (4.6), and we update both \mathbf{V} and \mathbf{U}_j with $\mathbf{y}_m^{(j)}$. The $\mathbf{y}_m^{(j)}$ is appended to \mathbf{V} every time we need to do recycling, while it is only appended to \mathbf{U}_j if we are working on the j^{th} right-hand side. This ensures that \mathbf{V} contains information pertinent to the entire system, while \mathbf{U}_j is kept small.

4.3. The Algorithm. Algorithm 2 describes our dynamic process. Details on efficient implementation of various steps in the algorithm will be addressed in the next section.

Algorithm 2: Recycling and Global Basis Construction

```

1  $\mathbf{U}_0 \leftarrow$  10 eigenvectors of  $\tilde{\mathbf{A}}_0$ ,  $\mathbf{X}_0$  solves  $\tilde{\mathbf{A}}_0 \mathbf{X}_0 = \mathbf{B}$ 
2  $\mathbf{V} \leftarrow$  basis for  $\text{Range}([\mathbf{U}_0, \mathbf{X}_0])$ 
3  $\mathbf{U}_j \leftarrow [\mathbf{U}_0, \mathbf{X}_0(:, j)]$ 
4 for  $i = 1 : K$  % for each interpolation point i do
5   for  $j = 1 : nrhs$  do
6     % Check if V is a good enough space
7      $\tilde{\mathbf{K}} = \tilde{\mathbf{A}}_i \mathbf{V}$ 
8      $[\mathbf{K}, \mathbf{R}] = qr(\tilde{\mathbf{K}}, 0)$ 
9      $\mathbf{V} = \mathbf{V} / \mathbf{R}$  %implicit only; now U, V same
10     $\mathbf{r}_j = \mathbf{B}(:, j) - \mathbf{K} \mathbf{K}^T \mathbf{B}(:, j)$ 
11    if  $\frac{\|\mathbf{r}_j\|}{\|\mathbf{B}(:, j)\|} > tol$  then
12      % MINRES recycling using U_j
13       $\tilde{\mathbf{K}}_j = \tilde{\mathbf{A}}_i \mathbf{U}_j$  % have already done this product
14       $[\mathbf{K}_j, \mathbf{R}] = qr(\tilde{\mathbf{K}}_j, 0)$  % need not be done from scratch
15       $\mathbf{U}_j = \mathbf{U}_j / \mathbf{R}$ 
16      Solve  $(\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T) \tilde{\mathbf{A}}_i \mathbf{y}_m^{(j)} = \mathbf{r}_j$  with MINRES
17       $\mathbf{V} \leftarrow [\mathbf{V}, \mathbf{y}_m^{(j)}]$ 
18       $\mathbf{U}_j \leftarrow [\mathbf{U}_j, \mathbf{y}_m^{(j)}]$ 
19    end
20  end
21 end

```

5. Algorithm Analysis. Algorithm 2 given in the previous section describes how we solve the sequence of systems while generating the approximate global basis. Our approach not only solves the sequence of systems efficiently but also builds the reduced global basis with only non-redundant information, therefore eliminating unnecessary solves as well as the need for an expensive SVD and a corresponding ad hoc

approach for truncation. Additionally, we can keep the cost down by exploiting the fact that we update \mathbf{V} and \mathbf{U}_j one column at a time.

Next, we explain how our approach differs from [13]. We also show that it is different from simply applying the recycling as in [21] to solve the systems and then doing an SVD to get the global basis.

5.1. System Solves. Our new approach is an improvement over Algorithm 1. To see this, we consider implementation of Step 1. To find the \mathbf{X}_k , one can of course use MINRES directly for each right-hand side. Alternatively, one can use recycling with the \mathbf{U}_j for the respective right-hand side, across all the systems⁷. The recycling would consist of generating $\hat{\mathbf{V}}_m^{(j)}$ as a basis for $\mathcal{K}_m((\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T) \mathbf{A}, (\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T) \mathbf{b}_j)$, with the intent of approximating \mathbf{x}_j over $\text{Range}([\mathbf{U}_j, \hat{\mathbf{V}}_m^{(j)}])$.

In contrast, in our new approach, we solve (4.6). Now $\mathbf{V}_m^{(j)}$ is generated as a basis for $\mathcal{K}_m((\mathbf{I} - \mathbf{K}_j \mathbf{K}_j^T) \mathbf{A}, (\mathbf{I} - \mathbf{K} \mathbf{K}^T) \mathbf{b}_j)$. The two Krylov spaces differ in the two approaches by the right-hand sides. Also, we want to approximate $\mathbf{x}_j - \mathbf{x}_{0,j}$, with $\mathbf{x}_{0,j} = \mathbf{U} \mathbf{K}^T \mathbf{b}_j$ rather than \mathbf{x}_j . A further numerical comparison is provided in Section (6.3)

It should be clear that recycling using \mathbf{U}_j must have some advantage over not recycling at all. In the first place, since \mathbf{U}_j contains an approximate invariant subspace, MINRES convergence on the projected systems would behave as if part of the spectrum has been deflated.

Additionally, the right-hand sides of (4.6) are residuals that have been already made small across spectral components *other than just those included in the invariant subspace*, since these residuals are \mathbf{b}_j 's orthogonalized against the entire \mathbf{K} , not just the (much smaller) \mathbf{K}_j . An argument for why \mathbf{r}_j is small in norm follows along the lines of subsection 5.2 in [21], and assumes that $\tilde{\mathbf{A}}_0 - \tilde{\mathbf{A}}_1$ is small over the invariant subspace of $\tilde{\mathbf{A}}_0$ corresponding to the smallest eigenvalues (smooth modes), and makes use of the fact that the columns of \mathbf{X}_0 are relatively smooth. If we assume that the values of absorption in our object and in the background are known and optimize only for the shape parameters, $\tilde{\mathbf{A}}_0 - \tilde{\mathbf{A}}_1$ is both diagonal, possibly low rank, with smooth modes made small by the operator.

We claim that using space \mathbf{U}_j instead of the (expensive) \mathbf{V} does not require many more iterations.

5.2. Global Basis. It is natural to ask why \mathbf{V} is a good global basis for the ROM problem. In addition to solution information, our \mathbf{V} contains an (approximate) invariant subspace corresponding to the smoothest modes for $\tilde{\mathbf{A}}_0$. By Theorem 4.1 in [21], we know that if the changes to $\tilde{\mathbf{A}}_k$ are concentrated over the high frequency modes, the invariant subspace consisting of eigenvectors corresponding to the smallest eigenvalues (low frequency modes) remain close. The \mathbf{X}_k are expected to be smooth. From a ROM standpoint, it would only be helpful to include this information in \mathbf{V} if we expect the solutions to be well represented in terms of the smooth modes. Figure 5.1 shows, in log scale, the absolute values of the coefficients of the solutions \mathbf{X}_5 in the directions of \mathbf{V} for Experiment 1 in Section 6. Note that these solutions were not used to build \mathbf{V} . You can see that the solutions have large components in the directions of the invariant subspace of $\tilde{\mathbf{A}}_0$, as well as the corresponding column of \mathbf{X}_0 .

⁷This is essentially the approach in [21], with recycle spaces having common invariant subspace information but tailored to the particular right-hand side. But they do further tuning of the recycle spaces to account for where one is in the optimization process.




Fig. 5.1: Logarithmically scaled image of the absolute values of the coefficients of solutions \mathbf{X}_5 in the directions of \mathbf{V} , where \mathbf{V} has 197 columns. Recall that \mathbf{X}_5 has 32 columns to source positions, and 32 columns corresponding to receiver positions. The first 10 columns of \mathbf{V} contain the invariant subspace of $\tilde{\mathbf{A}}_0$ corresponding to its smallest 10 eigenvalues, the next 64 columns correspond to \mathbf{X}_0 , and the remaining columns have been constructed using the update procedure in Algorithm 2 for $K_* = 3$.

5.3. Implementation Issues. In subsection 4.3, we gave an algorithm for constructing a global basis using the full order model solves using recycling. We will now discuss the cost of Algorithm 2 and how adding one column at a time to \mathbf{V} and \mathbf{U}_j helps to keep the cost down.

We will assume that \mathbf{U}_0 is a matrix of k_u columns containing a basis for the invariant subspace of $\tilde{\mathbf{A}}_0$ corresponding to the smallest eigenvalues and that \mathbf{X}_0 is known. Note that $\mathbf{X}_0, \mathbf{U}_0$ can be precomputed off-line and can be reused for other experiments. Again, we will use the fact (see (4.2)) that $\tilde{\mathbf{A}}_k = \tilde{\mathbf{A}}_* + \Delta_k$, where the first term in the sum is fixed and $\Delta_k = \text{diag}(\mu(\mathbf{p}_k))$.

Let's consider solving system 1 using Algorithm 2. In Algorithm 2 line 7, $\tilde{\mathbf{K}} := (\tilde{\mathbf{A}}_* + \Delta_1)[\mathbf{U}_0, \mathbf{X}_0]$. We can precompute and save the products $\tilde{\mathbf{A}}_*\mathbf{X}_0$ and $\tilde{\mathbf{A}}_*\mathbf{U}_0$, so only updating by $\Delta_1\mathbf{U}_0$ and $\Delta_1\mathbf{X}_0$ is required.

Next, in line 8 we need to compute the QR factorization of $\tilde{\mathbf{K}}$. We will compute this in a way that does not require full re-orthogonalization each time it is repeated. First, partition $\tilde{\mathbf{K}} = [\tilde{\mathbf{K}}_a, \tilde{\mathbf{K}}_b]$, where the first block corresponds to the number of columns of \mathbf{U}_0 . Next compute, $\mathbf{Q}_1\mathbf{R}_1 = \tilde{\mathbf{K}}_a$. Then, compute $(\mathbf{I} - \mathbf{Q}_1\mathbf{Q}_1^T)\tilde{\mathbf{K}}_b = \tilde{\mathbf{K}}_b - \mathbf{Q}_1(\mathbf{Q}_1^T\tilde{\mathbf{K}}_b) = \mathbf{Q}_2\mathbf{R}_2$. So, $\tilde{\mathbf{K}}_b = \mathbf{Q}_2\mathbf{R}_2 + \mathbf{Q}_1(\mathbf{Q}_1^T\tilde{\mathbf{K}}_b)$. It follows that a QR

factorization is

$$\begin{aligned} [\tilde{\mathbf{K}}_a, \tilde{\mathbf{K}}_b] &= [\mathbf{Q}_1 \mathbf{R}_1, \mathbf{Q}_1 (\mathbf{Q}_1^T \tilde{\mathbf{K}}_b) + \mathbf{Q}_2 \mathbf{R}_2] \\ &= [\mathbf{Q}_1, \mathbf{Q}_2] \begin{bmatrix} \mathbf{R}_1 & \mathbf{Q}_1^T \tilde{\mathbf{K}}_b \\ 0 & \mathbf{R}_2 \end{bmatrix}. \end{aligned}$$

We call $\mathbf{V} = [\mathbf{U}_0, \mathbf{X}_0] \mathbf{R}^{-1}$, noting that \mathbf{R}^{-1} need not be applied explicitly. Now we have $(\tilde{\mathbf{A}}_* + \Delta_1) \mathbf{V} = \mathbf{K}$, and \mathbf{K} has orthonormal columns.

Suppose we do recycling for system 1 and right-hand side 1. In Algorithm 2 line 13, we need $\tilde{\mathbf{K}}_1 := (\tilde{\mathbf{A}}_* + \Delta_1) [\mathbf{U}_0, \mathbf{X}_0(:, 1)]$. We have already formed this product above, so we just have to select the right columns of $\tilde{\mathbf{K}}$. Moreover, the QR factorization of $\tilde{\mathbf{K}}_1$ is computed from \mathbf{Q}_1 and \mathbf{R}_1 . All we need to compute is $\tilde{\mathbf{K}}_1(:, k_u + 1) - \mathbf{Q}_1 (\mathbf{Q}_1^T \tilde{\mathbf{K}}_1(:, k_u + 1))$ and then normalize it. The normalization constant becomes the lower right corner component of the upper triangular matrix.

Next, we solve the projected problem with MINRES and append the solution, $\mathbf{y}_m^{(1)}$, to \mathbf{V} and \mathbf{U}_j . We will check to see if the newly enlarged \mathbf{V} is sufficient to represent the solution for the second right-hand side. In order to do this, we compute

$$(\tilde{\mathbf{A}}_* + \Delta_1) [\mathbf{V}, \mathbf{y}_m] = [(\tilde{\mathbf{A}}_* + \Delta_1) \mathbf{V}, (\tilde{\mathbf{A}}_* + \Delta_1) \mathbf{y}_m] = [\mathbf{K}, \mathbf{z}].$$

Since \mathbf{K} already has orthogonal columns, we need only to compute

$$[\mathbf{K}, \mathbf{z}] = [\mathbf{K}, \mathbf{q}] \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{K}^T \mathbf{z} \\ 0 & \rho \end{bmatrix}}_{\tilde{\mathbf{R}}},$$

where $\mathbf{z} - \mathbf{K}(\mathbf{K}^T \mathbf{z}) = \rho \mathbf{q}$, so we have $\mathbf{K} \leftarrow [\mathbf{K}, \mathbf{q}]$, $\mathbf{V} \leftarrow [\mathbf{V}, \mathbf{y}_m] \tilde{\mathbf{R}}^{-1}$.

For a recycle solve for right-hand side 2, we follow the same procedure. For additional right-hand sides, only incremental new calculations are needed.

6. Numerical Results. We present two experiments on a 201×201 mesh (which gives us 40401 degrees of freedom) for the forward problem. We use 32 sources and 32 detectors in the model, meaning that \mathbf{B} in (4.1) will have 64 columns. The image space is parameterized using parametric level sets (PaLS) (see [1] for details). We use 25 compactly supported radial basis functions to define the PaLS image, which results in a total of 100 parameters for the optimization problem (2.8).

The ground truth images for Experiments 1 and 2 are given in Figure 6.2a and 6.3a, respectively. We note that these images cannot be exactly reconstructed via the image space parameterization we are using. Thus, we avoid the so-called inverse crime. To obtain the noisy data we added 1% noise to the simulated true measured data in each of our experiments.

We solve the optimization problems using the TREGS [14] algorithm. We stop the optimization when the residual norm falls below 1.1 times the noise level. We report and compare the results for two cases in each experiment. First, we report results assuming that the full order problem was used to compute the function and Jacobian evaluation at each step. Then we report results using the ROM to replace the function and Jacobian evaluation. Figure 6.1 gives the absorption image using the initial set of parameters. We used $K_* = 3$ systems for each experiment to create the reduced order model space. The tolerance in line 11 of Algorithm 1 was set to be 10^{-7} . All of the experiments were run using a laptop with a 3.20 GHz processor and 16.0 GB RAM using MATLAB R2014a.





Fig. 6.1: Initial Absorption Image

6.1. Experiment 1. In Experiment 1, we needed to solve 187 large, single right-hand side systems to generate what we needed to construct the global basis matrix (note that the 64 of these corresponding to \mathbf{X}_0 could have been pre-computed off-line). Including the additional 10 eigenvectors of $\tilde{\mathbf{A}}_0$ that were used as the first 10 columns of \mathbf{V} , \mathbf{V} has 197 columns and thus the reduced model has order 197. Therefore, the reduced models require solutions to linear systems of size 197×197 rather than 40401×40401 for the full order model.

The optimization using the full order model for Experiment 1 required 30 function evaluations and 15 Jacobian evaluations. In comparison, the optimization run using the reduced order model, once it's been generated, for function and Jacobian evaluations required 28 function evaluations and 14 Jacobian evaluations, indicating that using a ROM in place of FOM does not greatly impact convergence rate of the optimization. The bottom line is that solving the optimization using the full order model requires the solution of 1440 systems of size 40401×40401 . On the other hand, solving using our approach requires solution of 187 systems of size 40401×40401 , which are used to construct \mathbf{V} during the first few optimization steps. The remainder of the work is in solving systems of size 197×197 until the convergence tolerance for the optimization is achieved.

Figure 6.2 shows the reconstructions for Experiment 1. Figure 6.1 also includes the number of (unpreconditioned) MINRES iterations for each experiment with and without recycling. Although the tables only show a sample of results, it is clear that the iterations decrease from one right-hand side to the next, and system to system, using our approach. The jump in number of iterations for right-hand-side 33 comes from the fact that we concatenated $\tilde{\mathbf{B}}$ and $\tilde{\mathbf{C}}$ to form one right-hand-side for the symmetric transfer function, so the 33rd right-hand side corresponds to the first column in $\tilde{\mathbf{C}}$.

6.2. Experiment 2. In Experiment 2, a total 188, 40401×40401 single-right-hand side systems were solved to compute our global basis. The reduced order model has order 198. Therefore, the reduced models require solutions to linear systems of size 198×198 rather than 40401×40401 for the full order model. The optimization using the full order model required 126 function evaluations and 78 Jacobian evaluations. The optimization run using our reduced order model took 123 function evaluations and 76 Jacobian evaluations to converge to our stopping criterion, so again, there is no negative impact on convergence rate by replacing the FOM with the ROM. The difference in the total number of large (40401×40401) single-right-hand side systems that need to be solved, though, is even more pronounced in this example than in the last: 6,528 are needed for the FOM approach vs. only 188 for the ROM approach. Moreover, the work involved in solving for the latter systems is reduced,



(a) Original anomaly with normally distributed noise added. (b) Reconstruction using the full order model. (c) Reconstruction using the reduced order model.

Fig. 6.2: Results for Experiment 1. Reconstruction on a 201×201 mesh, resulting in 40401 degrees of freedom in the forward model and 197 degrees of freedom in the reduced model for the forward model. 32 sources, 32 detectors, and 25 basis functions were used.

System	RHS	Experiment 1		Experiment 2	
		MINRES Its	MINRES Its with Recycling	MINRES Its	MINRES Its with Recycling
1	1	463	140	470	127
	20	541	52	506	53
	32	487	0	493	0
	33	467	124	470	120
	53	528	57	514	48
	64	489	0	494	0
2	1	474	118	501	124
	20	513	35	545	38
	32	497	0	526	5
	33	474	105	500	132
	53	526	37	567	127
	64	497	0	532	0

Table 6.1: Number of MINRES iterations for Experiments 1 and 2.

since MINRES requires fewer iterations due to the recycling.

Figure 6.3 shows the reconstructions for Experiment 2. Again, Figure 6.1 shows the number of unpreconditioned MINRES iterations for each experiment with and without our inner-outer recycling approach.

6.3. Value of Inner-Outer Recycling. There is a significant benefit to using two levels of recycling information. To see this, consider Algorithm 1 to construct the global basis. We could solve the full order model systems in Step 1 (e.g. systems in 4.1) with the unpreconditioned MINRES recycling approach in [21]. It is important to note that the recycle spaces would be different than those used in our new method. Furthermore, in the new method we solve the correction equations (4.6) as opposed to solving (4.1). For $j > 1$, the recycle spaces for the [21] approach do not incorporate information from other systems corresponding to other right-hand sides. In contrast,

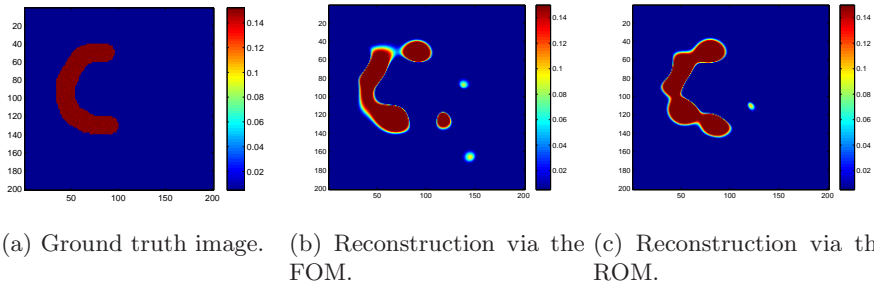


Fig. 6.3: Results for Experiment 2. Reconstruction on a 201×201 mesh, resulting in 40401 degrees of freedom in the forward model and 198 degrees of freedom in the reduced order model. 32 sources, 32 detectors, and 25 basis functions were used.

since we augment \mathbf{V} from information about right-hand side j , we update \mathbf{K} . The update in \mathbf{K} then causes updates to $\mathbf{r}_{j+1}, \dots, \mathbf{r}_{n_{src}+n_{det}}$, which are the right-hand sides in (4.6).

Table 6.2 compares the recycling of [21] with our new approach. The results show that with our approach, the number of iterations and the relative residuals decrease as you move from one right-hand-side to the next and also as you move from system to system. The jump at right-hand-side 33 is due to the fact that you are moving to the second half of the concatenated right-hand-sides, so these correspond to solving the adjoint problem. Using the approach in [21], however, does not speed up convergence across right-hand-sides. In our approach, the reduced global basis, \mathbf{V} , is already constructed when we are done with the full order model solves. We note there is a big difference in total number of MINRES iterations to squeeze all information from systems 1 and 2. It took our approach 5,006 iterations, while it took 22,659 iterations for the recycling method in [21].

7. Conclusions and Future Work. First, we established that our transfer function at zero frequency for DOT could be re-written in terms of a SPD matrix. Then, we developed an inner-outer Krylov recycling approach to update the global basis matrix relative to our new formulation of the transfer function. Two numerical experiments illustrate the success in using the ROM in place of the FOM during the optimization.

In this paper, we only considered the 0 frequency case. It is non-trivial to extend the algorithm to the case when ω is non-zero, and it is therefore the subject of a forthcoming paper.

Clearly, the performance of our method depends on the values of some parameters, such as the residual tolerance and the number of systems K_* . We found, for example, that if we dropped the tolerance slightly, the number of system solves, and therefore the reduced model order, was even further reduced, without too much degradation in the reconstruction. Likewise, using a larger value of K_* gave slightly larger reduced order models, but with no improvement in the quality of the reconstruction. The trade-offs in performance due to these selections are currently under investigation. Finally, preliminary results indicate that solving the systems corresponding to different right-hand sides in a different ordering may also have an impact on the model

System	RHS	Our Approach		Recycling from [21]	
		Its	Initial Relative Residual	Its	Initial Relative Residual
1	1	140	7.523115e-05	140	7.523329e-05
	20	52	1.077645e-06	185	6.584893e-04
	32	0	8.866398e-08	152	1.164802e-04
	33	124	4.776975e-05	127	5.213653e-05
	53	57	1.692149e-06	191	6.114295e-04
	64	0	9.960690e-08	151	1.450962e-04
2	1	118	4.673235e-05	131	5.091877e-05
	20	35	6.653493e-07	190	6.171251e-04
	32	0	8.754708e-08	153	1.051619e-04
	33	105	2.588292e-05	129	3.388392e-05
	53	37	8.454303e-07	188	6.405050e-04
	64	0	7.960101e-08	151	9.958796e-05

Table 6.2: Comparison of MINRES recycling using \mathbf{U}_j as described in [21] vs. the inner-outer approach using both \mathbf{V} and \mathbf{U}_j as described in Algorithm 2. Note that the two approaches lead to different choices for \mathbf{U}_j as well as different systems to solve.

order, and we will continue to investigate this phenomenon.

REFERENCES

- [1] A. Aghasi, E. Miller, and M. E. Kilmer. Parametric level set methods for inverse problems. *SIAM Journal on Imaging Science*, 4:618–650, 2011.
- [2] K. Ahuja, E. de Sturler, and P. Benner. Recycling bicgstab with an applicaiton to parametric model order reduction. *SIAM J. Sci. Comput.*, 37:S429–S446, 2015.
- [3] H. Antil, M. Heinkenschloss, and R. H. W. Hoppe. Domain decomposition and balanced truncation model reduction for shape optimization of the Stokes system. *Optimization Methods and Software*, 26(4–5):643–669, 2011.
- [4] H. Antil, M. Heinkenschloss, R. H. W. Hoppe, C. Linsenmann, and A. Wixforth. Reduced order modeling based shape optimization of surface acoustic wave driven microfluidic biochips. *Mathematics and Computers in Simulation*, 82(10):1986–2003, 2012.
- [5] E. Arian, M. Fahl, and E. Sachs. Trust-region proper orthogonal decomposition models by optimization methods. In *Proceedings of the 41st IEEE Conference on Decision and Control*, pages 3300–3305, Las Vegas, NV, 2002. IEEE.
- [6] S. R. Arridge. Optical tomography in medical imaging. *Inverse Problems*, Vol. 16:R41–R93, 1999.
- [7] O. Bashir, K. Willcox, O. Ghattas, B. van Bloemen Waanders, and J. Hill. Hessian-based model reduction for large-scale systems with initial condition inputs. *International Journal for Numerical Methods in Engineering*, 73(6):844–868, 2008.
- [8] U. Baur, C. Beattie, P. Benner, and S. Gugercin. Interpolatory projection methods for parameterized model reduction. *SIAM Journal on Scientific Computing*, 33:2489–2518, 2011.
- [9] C. Beattie, S. Gugercin, and S. Wyatt. Inexact solves in interpolatory model reduction. *Linear Algebra and its Applications*, 2011. Appeared on-line as doi:10.1016/j.laa.2011.07.015.
- [10] P. Benner, S. Gugercin, and K. Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Review*, 57(4):483–531, 2015.
- [11] L. Borcea, V. Druskin, A. V. Mamonov, and M. Zaslavsky. A model reduction approach to numerical inversion for a parabolic partial differential equation. *arXiv preprint arXiv:1210.1257*, 2012.
- [12] E. de Sturler. Nested Krylov methods based on GCR. *J. Comput. Appl. Math.*, 67(1):15–41, 1996.

- [13] E. de Sturler, S. Gugercin, M. E. Kilmer, S. Chaturantabut, C. Beattie, and M. O’Connell. Non-linear parametric inversion using interpolatory model reduction. *SIAM J. Sci. Comput.*, 37, 2015.
- [14] E. de Sturler and M. E. Kilmer. A regularized Gauss-Newton trust region approach to imaging in diffuse optical tomography. *SIAM Journal on Scientific Computing*, 33:3057 – 3086, 2011.
- [15] V. Druskin, V. Simoncini, and M. Zaslavsky. Solution of the time-domain inverse resistivity problem in the model reduction framework Part I. One-dimensional problem with SISO data. *SIAM Journal on Scientific Computing*, 35(3):A1621–A1640, 2013.
- [16] L. Feng, P. Benner, and J. G. Korvink. Parametric model order reduction accelerated by subspace recycling. In *Proceedings of the 48th IEEE Conference on Decision and Control*. IEEE, IEEE, 2009.
- [17] L. Feng, P. Benner, and J. G. Korvink. Subspace recycling accelerates the parametric macro-modeling of MEMS. *Int. J. Numer. Methods Eng.*, 94:84–110, 2013.
- [18] S. Gugercin, T. Stykel, and S. Wyatt. Model reduction of descriptor systems by interpolatory projection methods. *SIAM Journal on Scientific Computing*, 35(5):B1010–B1033, 2013.
- [19] E. Haber, U. M. Ascher, and D. Oldenburg. On optimization techniques for solving nonlinear inverse problems. *Inverse Problems*, 16:1263–1280, 2000.
- [20] M. Hinze and S. Volkwein. Proper orthogonal decomposition surrogate models for nonlinear dynamical systems: Error estimates and suboptimal control. In *Dimension Reduction of Large-Scale Systems*, pages 261–306. Springer, 2005.
- [21] M. Kilmer and E. de Sturler. Recycling subspace information for diffuse optical tomography. *SIAM Journal on Scientific Computing*, 27(6):2140–2166, 2006.
- [22] K. Kunisch and S. Volkwein. Proper orthogonal decomposition for optimality systems. *ESAIM: Mathematical Modelling and Numerical Analysis*, 42(1):1–23, 1 2008.
- [23] L. A. M. Mello, E. de Sturler, G. H. Paulino, and E. C. N. Silva. Recycling Krylov subspaces for efficient large-scale electrical impedance tomography. *Computer Methods in Applied Mechanics and Engineering*, 199:3101–3110, 2010.
- [24] A. Stathopoulos, A. M. Abdel-Rehim, and K. Orginos. Deflation for inversion with multiple right-hand sides in QCD. *Journal of Physics: Conference Series 180*, 2009.
- [25] A. Stathopoulos, A. M. Abdel-Rehim, and W. Wilcox. Deflated bicgstab for linear equations in QCD. pages 026/1–026/7, 2007. Proceedings of Science LAT2007.
- [26] C. R. Vogel. *Computational Methods for Inverse Problems*. SIAM, Philadelphia, 2002.
- [27] S. Wang, E. de Sturler, and G. Paulino. Large-scale topology optimization using preconditioned Krylov subspace methods with recycling. *International Journal for Numerical Methods in Engineering*, 69:2441–2461, 2007.
- [28] Y. Yue and K. Meerbergen. Accelerating optimization of parametric linear systems by model order reduction. *SIAM Journal on Optimization*, 23(2):1344–1370, 2013.