SIAM J. SCI. COMPUT.
Vol. 39, No. 5, pp. A2365–A2393

A DATA SCALABLE AUGMENTED LAGRANGIAN KKT PRECONDITIONER FOR LARGE-SCALE INVERSE PROBLEMS*

NICK ALGER†, UMBERTO VILLA†, TAN BUI-THANH‡, AND OMAR GHATTAS†

Abstract. Current state-of-the-art preconditioners for the reduced Hessian and the Karush–Kuhn–Tucker (KKT) operator for large-scale inverse problems are typically based on approximating the reduced Hessian with the regularization operator. However, the quality of this approximation degrades with increasingly informative observations or data. Thus the best case scenario from a scientific standpoint (fully informative data) is the worse case scenario from a computational perspective. In this paper we present an augmented Lagrangian-type preconditioner based on a block diagonal approximation of the augmented upper left block of the KKT operator. The preconditioner requires solvers for two linear subproblems that arise in the augmented KKT operator, which we expect to be much easier to precondition than the reduced Hessian. Analysis of the spectrum of the preconditioned KKT operator indicates that the preconditioner is effective when the regularization is chosen appropriately. In particular, it is effective when the regularization does not overpenalize highly informed parameter modes and does not underpenalize uninformed modes. Finally, we present a numerical study for a large data/low noise Poisson source inversion problem, demonstrating the effectiveness of the preconditioner. In this example, three MINRES iterations on the KKT system with our preconditioner results in a reconstruction with better accuracy than 50 iterations of CG on the reduced Hessian system with regularization preconditioning.

Key words. PDE constrained inverse problems, data scalability, augmented Lagrangian, preconditioning, KKT matrix, Krylov subspace methods

AMS subject classifications. 65J22, 49K20, 65F08, 65N21, 65F22, 65K10

DOI. 10.1137/16M1084365

1. Introduction.

Here we focus on data scalable Karush–Kuhn–Tucker (KKT) preconditioners for large-scale linear[1] inverse problems in which one seeks to reconstruct a parameter field from observations of an associated state variable. Specifically, suppose we have observations y of a state variable u that have been corrupted by some noise ζ,

$$y = Bu + \zeta, \tag{1}$$

where B is a linear operator encoding the action of the observation process (i.e., the observation operator). Further, let the state variable u depend on a parameter q through a linear state equation,

$$Tq + Au = f, \tag{2}$$

*Submitted to the journal's Methods and Algorithms for Scientific Computing section July 12, 2016; accepted for publication (in revised form) August 1, 2017; published electronically October 24, 2017.

http://www.siam.org/journals/sisc/39-5/M108436.html

Funding: This work was funded by DOE grants DE-SC0010518 and DE-SC0009286, AFOSR grant FA9550-12-1-0484, and NSF grant CBET-1508713.

†Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78712 (nalger@ices.utexas.edu, uvilla@ices.utexas.edu, omar@ices.utexas.edu).

‡Department of Aerospace Engineering and Engineering Mechanics, and Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78712 (tanbui@ices.utexas.edu).

[1]The preconditioner presented here is also applicable to nonlinear inverse problems, which give rise to linear systems of the form we address here at each iteration of a (Gauss) Newton or sequential quadratic programming method.

where A is the forward operator, T maps the parameter to the residual of the state equation, and f is the known right-hand-side source. We seek to reconstruct the parameter q based on the observations y. Under an independent and identically distributed (i.i.d.) Gaussian noise model,[2] this inverse problem naturally leads to a least squares optimization problem of the following form:

$$\begin{aligned} \min_{q,u} \quad & \frac{1}{2}\|Bu - y\|^2 + \frac{\alpha}{2}\|Rq\|^2 \\ \text{such that} \quad & Tq + Au = f, \end{aligned} \tag{3}$$

where R is a suitably chosen regularization operator and α is a regularization parameter. We focus on the computationally difficult case in which the observations y are highly informative about the parameter q, and the regularization parameter α is correspondingly small. The KKT system expressing the first order necessary condition for an optimal solution of (3) is

$$\underbrace{\begin{bmatrix} \alpha R^*R & & T^* \\ & B^*B & A^* \\ T & A & \end{bmatrix}}_{K} \begin{bmatrix} q \\ u \\ \eta \end{bmatrix} = \begin{bmatrix} 0 \\ B^*y \\ f \end{bmatrix}. \tag{4}$$

Here all operators are linear, possibly unbounded, maps between suitable Hilbert spaces. The symbols A^*, B^*, R^*, and T^* denote the adjoints (in the Hilbert space sense) of A, B, R, and T, respectively, and η denotes the adjoint variable (or Lagrange multiplier) for the state equation. More details regarding our setting and notation are provided in section 1.5. We denote the KKT operator by K. For large-scale problems, direct factorization of K is not computationally feasible due to both memory and CPU time considerations, and iterative methods must be used. Krylov methods, MINRES [51] in particular, are the gold standard for iteratively solving this kind of large-scale symmetric indefinite system. The performance of the MINRES method strongly depends on the clustering of the spectrum of the preconditioned operator, the more clustered the better [65]. In this paper we propose clustering the spectrum of the KKT operator by using the following block diagonal preconditioner:

$$P := \begin{bmatrix} \alpha R^*R + \rho T^*T & & \\ & B^*B + \rho A^*A & \\ & & \frac{1}{\rho} I \end{bmatrix}, \tag{5}$$

where I denotes the identity map associated with the appropriate inner product (in the computations, a mass matrix). We further propose choosing $\rho = \sqrt{\alpha}$ based on theoretical results and numerical evidence. In our theory and numerical experiments we assume that A and R are invertible maps. Although the application of preconditioner (5) and the abstract theory we present in section 3 do not depend on invertibility of T, much of the intuition behind the assumptions of the theory is lacking in the case where T is noninvertible. Remedies for this case are the subject of ongoing research. While existing data scalable KKT preconditioners usually require regularization operators R that are spectrally equivalent to the identity,[3] our preconditioner (5) performs well even if R is a discretization of an unbounded operator (e.g., Laplacian regularization).

[2] The least squares formulation in (3) also applies to general (non-i.i.d.) Gaussian noise models after reweighing the data misfit equation (1) by the inverse square root of the noise covariance.

[3] A review of existing work is presented in section 1.3. In particular, see sections 1.3.2 and 1.3.3.

1.1. Overview of results. In section 3.2 we prove that, using our preconditioner (5), the symmetrically preconditioned KKT operator satisfies the condition number bound

$$\operatorname{cond}\left(P^{-1/2}KP^{-1/2}\right) \leq \frac{3}{(1-\beta)\delta} ,$$

where δ and β are bounds on the eigenvalues of the arithmetic and geometric means of certain damped projectors.[4] Based on the nature of the damped projectors, we expect these eigenvalue bounds to be satisfied with good constants δ and β for inverse problems that are *appropriately regularized*. By "appropriately regularized," we mean that the regularization is chosen so that components of the parameter that are highly informed by the data are not overpenalized, and components of the parameter that are poorly informed by the data are not underpenalized. In section 5 we derive quantitative bounds on δ and β for the special case of source inversion problems with spectral filtering regularization. When the regularization is chosen appropriately, these bounds are independent of the mesh size and of the information content in the data.

In section 6 we numerically demonstrate the effectiveness of the preconditioner on a Poisson source inversion problem with highly informative data and Laplacian regularization. Preconditioning the KKT system with our preconditioner results in greater accuracy in three MINRES iterations than the widely used regularization preconditioning on the reduced Hessian system achieves in 50 conjugate gradient iterations. Even though the regularization is not a spectral filter, our preconditioner still exhibits mesh independence and good scalability with respect to a decrease in the regularization parameter by 10 orders of magnitude. As suggested by our theory, we see that the performance of the preconditioner in the small regularization regime actually improves as more data are included in the inversion.

1.2. Desirable properties of KKT preconditioners for inverse problems. To evaluate the quality of a KKT preconditioner, it is useful to consider its performance with respect to the following desired properties:

(a) *Problem generality*: A KKT preconditioner exhibits problem generality if it applies to a wide variety of inverse problems.

(b) *Efficient solvers for preconditioner subproblems*: If applying the inverse of the preconditioner to a vector involves solving subproblems, efficient solvers for those subproblems are required.

(c) *Mesh scalability*: Finite dimensional inverse problems often arise from discretizations of infinite dimensional inverse problems. Preconditioners for such problems are *mesh scalable* if the effectiveness of the preconditioner (as measured in terms of either the condition number of the preconditioned KKT operator, the clustering of the spectrum of the preconditioned KKT operator, or the number of Krylov iterations required to converge to a fixed tolerance) does not degrade substantially as the meshes used to discretize the problem are refined.

(d) *Regularization robustness*: KKT preconditioners are *regularization robust* if their effectiveness does not degrade substantially as the regularization parameter α is made smaller.

[4] The condition number provides an upper bound on the required number of MINRES iterations. An even sharper bound could be obtained by characterizing all four extreme eigenvalues (endpoints of the positive and negative intervals in which the eigenvalues reside) of the preconditioned system [28].

(e) *Data scalability*: KKT preconditioners are *data scalable* if their effectiveness does not degrade substantially as more data—or rather, more informative data—are included in the inverse problem.

Currently, there is no known preconditioner that exhibits uniformly good performance with respect to all of these properties. In this paper, we provide a combination of theoretical results and numerical evidence demonstrating that our preconditioner provides substantial improvements over existing preconditioners, especially with respect to problem generality and data scalability.

Within the scope of inverse problems, we view the goal of robustness to arbitrarily chosen values of the regularization parameter, (d), to be unwarranted and unnecessarily restrictive. In particular, for properly regularized inverse problems the regularization operator and regularization parameter are not arbitrary. Rather, they are chosen in response to the data available in the problem: that is, to constrain parameter modes that are not informed by the data, while minimally modifying components of the parameter that are informed by the data. Thus it is important that the preconditioner performs well as the informative content of the data increases while the regularization parameter decreases correspondingly. However, it is not important for the preconditioner to perform well in the under-regularized regime in which the regularization parameter is small but the data are uninformative. In this under-regularized regime, a good preconditioner would simply accelerate convergence to noise, i.e., more rapid solution of the wrong problem. Instead, we advocate designing preconditioners that perform well with increasingly informative data, (e), for which the regularization parameter is considered a dependent parameter chosen so that the inverse problem is neither substantially over- nor under-regularized. This extra flexibility permits design of the preconditioner to better address the entire set of desired properties (a)–(e).

Among data (e) and mesh scalable (c) preconditioners, ours is the most general (a). The subproblems that must be solved (b) while applying the preconditioner are of similar difficulty to those encountered by existing (less general) data scalable preconditioners. What remains for our preconditioner to fully satisfy all of the remaining desirable properties, (a), (b), (c), and (e), is to generalize it to noninvertible T. As mentioned above, this is ongoing research; nevertheless, there are many inverse problems characterized by invertible T operators. In addition to source inversion problems (addressed in sections 5 and 6), coefficient inverse problems in which the state and parameter share the same discretization often give rise to invertible T.

1.3. Review of existing work. A wide variety of preconditioners for KKT operators similar to (4) have been developed in a number of different contexts including parameter estimation, optimal control, PDE constrained optimization, optimal design, and saddle point systems arising in mixed discretizations of forward problems [11, 23, 45]. In the following subsections we discuss existing preconditioners based on the reduced Hessian (section 1.3.1), the adjoint Schur complement (section 1.3.2), block scaling (section 1.3.3), and multigrid (section 1.3.4).

We will see that existing preconditioners either scale poorly with increasing data and decreasing regularization, or they only apply to specific problems, or they make restrictive assumptions about the B, R, and T operators. In particular, in the literature it is common to assume that the parameter and/or observation spaces are L^2 spaces, and one or more of the operators B, R, and T are spectrally equivalent to either identity maps (I), or restriction maps (Γ) that restrict functions to a subdomain. These assumptions on B, R, and T may be inappropriate for the inverse problem at hand. For example, they prevent one from using observations of derived quan-

tities such as flux, using smoothing Laplacian-like regularization, and inverting for material coefficients. We will regularly note such assumptions by following references with a parenthetical expression; e.g., "[61] (L^2, $R \approx I$, $B \approx \Gamma$, $T \approx -I$)" means that the preconditioner in [61] assumes that the parameter and observation spaces are L^2 spaces, R is spectrally equivalent to an identity map (L^2 regularization), B is spectrally equivalent to a restriction map (direct observations of the state on a subdomain), and T is spectrally equivalent to a negative identity map (the parameter enters the state equation on the right-hand side as a source term).

1.3.1. The reduced Hessian. The reduced Hessian is the Hessian of the unconstrained reformulation of optimization problem (3), in which the constraint is eliminated by viewing the state u as an implicit function of the parameter q via solution of the state equation. We discuss this reduced space problem in more detail in section 4.1. For linear inverse problems (as considered in this paper), the reduced Hessian is equivalent to the Schur complement of the KKT operator with respect to the parameter. In other words, it is the operator remaining when the state and adjoint variables (and corresponding equations) are solved for and eliminated from the KKT system. Likewise, the KKT operator can be derived by starting with the reduced Hessian, defining auxiliary variables, and performing simple algebraic manipulations. Thus performing solves with the reduced Hessian and performing KKT solves are equivalent: If one can efficiently solve the former, then one can efficiently solve the latter and vice versa. For this reason, a popular class of methods for solving (4) relies on approximations or preconditioners for the reduced Hessian [13, 14, 37].

The most popular class of general purpose preconditioners for the reduced Hessian is based on approximating this operator with just the regularization operator, and either neglecting the data misfit term or dealing with it through some form of low rank approximation. The regularization is typically an elliptic operator and can be inverted using multigrid or other standard techniques. Furthermore, for ill-posed inverse problems the data misfit portion of the reduced Hessian at the optimal solution is usually a compact operator in the infinite-dimensional limit [20, 21, 22, 64]. Thus Krylov methods preconditioned by the regularization operator usually yield mesh independent, superlinear convergence rates[5] [7, 33, 40]. However, the importance of the regularization term in the reduced Hessian decreases as the regularization parameter is made smaller, and the importance of the data misfit term increases as the informativeness of the data increases. Indeed, the numerical rank of the data misfit portion of the reduced Hessian is roughly the number of parameter modes that are "informed" by the data. In addition, the eigenvalues of the regularization preconditioned Hessian are typically well-separated, which means that this approach will still require large numbers of Krylov iterations on problems with highly informative data. Thus, the best case scenario from a scientific standpoint (highly informative data) is the worst case scenario from a computational standpoint (large numbers of Krylov iterations required).

Other problem-specific reduced Hessian solvers and preconditioners have been developed using a diverse set of techniques including analysis of the symbol of the reduced Hessian [5], matrix probing [24], approximate sparsity in curvelet frames [38], and analytic expressions derived for model problems [2, 4, 31, 32, 34].

[5]Here, by *superlinear*, we mean that the norm of the error decays superlinearly with respect to the number of Krylov iterations.

1.3.2. Schur complement for the adjoint variable. In contrast to the approaches based on the reduced Hessian described above (where the state and adjoint are eliminated), another class of preconditioners of increasing interest in recent years is based on block factorizations that eliminate the parameter and state, resulting in a Schur complement operator for the adjoint variable. This approach requires one to design preconditioners for the objective block (the 2×2 block corresponding to q and u in (4)) and for the Schur complement associated with the adjoint variable. In the case of limited observations, the objective block is singular and requires special handling; a common approach is to add a small positive diagonal shift to the block.

Mesh independent block diagonal preconditioners based on approximating the objective block with mass matrices and the adjoint Schur complement with AA^* have been proposed for L^2 regularized optimal control problems with the Poisson equation as a constraint and a control objective targeting the state variable directly [55, 56] (L^2, $R \approx I, B \approx I, T \approx -I$), and extended to problems with parabolic PDE constraints and limited observations [61] (L^2, $R \approx I, B \approx \Gamma, T \approx -I$). More nuanced approximations of the Schur complement have been shown to yield robustness with respect to the regularization parameter for problems in the elliptic case in [54] (L^2, $R \approx I, B \approx I, T \approx -I$) and the parabolic case in [52] (L^2, $R \approx \Gamma, B \approx \Gamma, T \approx -\Gamma$). Regularization robust adjoint Schur complement based KKT preconditioners have also been developed for optimal control problems in cases where there are additional box constraints on the control and state variables [53] (L^2, $R \approx I$, $T \approx -I$). A general framework for using the Schur complement for the adjoint variable to precondition optimal control problems with box constraints is analyzed in an abstract function space setting in [58], with only minimal assumptions on the operators B, R, and T. However, the specific Schur complement preconditioners presented in [58] are not regularization robust.

Certain KKT preconditioners that are not block diagonal [8] have been used to precondition elliptic PDE constrained optimal control problems with L^2 regularization and observations [39, 59] (L^2, $B \approx I$, $R \approx I$, $T \approx -I$). Preconditioners of this type have also been shown to be Hermitian positive definite in certain nonstandard inner products, allowing the use of conjugate gradient as a Krylov solver [18, 60].

Inner-outer methods where the Schur complement solve is performed (exactly or approximately) with an additional inner stationary iteration have also been proposed for several problems. These include optimal Stokes control with L^2 regularization and observations [57] (L^2, $B \approx I$, $R \approx I$, $T \approx -I$), and optimal transport with a problem-specific diagonal regularization operator [12]. Recently, a method of this type was proposed for optimal control problems with elliptic and parabolic PDE constraints and smoothing regularization (L^2, $R^*R \approx \Delta + I$) [9]. Regularization robustness was demonstrated for the case $B = I$.

1.3.3. Block scaling. An abstract framework for constructing parameter independent (e.g., regularization robust) block diagonal preconditioners for saddle point systems is studied in [67] and applied to optimal control problems with elliptic and Stokes PDE constraints, with $B \approx I$, $R \approx I$, $T \approx -I$. In [47], a certain class of block diagonal KKT preconditioners for inverse problems (satisfying many assumptions) was shown to be mesh independent and only weakly dependent on the regularization parameter.[6] One of the central assumptions of the theory for this block diagonal preconditioner is that the spectrum of the observation operator decays exponentially.

[6] Note that in several of the papers cited in this subsection, the meanings of B and T are switched relative to their use here.

In a subsequent paper this assumption was replaced with the similar assumption that the spectrum of the unregularized KKT system decays exponentially [48]. Since the decay rates of these spectra depend on the informativeness of the data, these assumptions are not applicable (with good constants) in the context of inverse problems with highly informative data. To overcome this limitation, recently the block diagonal preconditioner

$$\begin{bmatrix} \alpha I & & \\ & B^*B + \alpha\widehat{A^*A} & \\ & & \frac{1}{\alpha}I \end{bmatrix} \tag{6}$$

was proposed in [44], where $\widehat{A^*A}$ is a fourth order elliptic operator that is spectrally equivalent to A^*A. This preconditioner was proven to be mesh and regularization robust for a specific source inversion problem with L^2 regularization (L^2, $R \approx I$, $T \approx -I$). Despite substantial differences in motivation and analysis, our proposed preconditioner (5) could be considered as a generalization of this work to more general operators R and T. Specifically, setting $\rho = \alpha$ (instead of our suggestion $\rho = \sqrt{\alpha}$), our preconditioner has the same second and third diagonal blocks as the preconditioner (6), but contains a more elaborate operator depending on R and T in the first block.

1.3.4. Multigrid. Another family of KKT preconditioners for parameter estimation problems are based on multigrid (see the review paper [17] and references therein). These techniques are classically categorized into three main categories: (1) speeding up or preconditioning forward and adjoint solves, (2) using multigrid to precondition the reduced Hessian, and (3) collective smoothing.

Methods in the first category do not use multigrid to address the fundamental difficulties stemming from highly data informed inverse problems: Speeding up the forward (and adjoint) solves does not address the challenge of creating a preconditioner that is data scalable, because the number of forward/adjoint solves that must be done scales with the informativeness of the data.

The primary difficulty with category (2) is that when the regularization is chosen appropriately, the regularization and data misfit terms of the reduced Hessian tend to "fight" each other (more on this in section 4). Thus smoothers for the regularization term tend to be roughers for the data misfit term, and vice versa. As a result, multigrid methods belonging to the second category tend to be restricted to the case $R \approx I$. We note in particular the papers [1, 2, 3, 26, 27] on elliptic, parabolic, and Stokes source inversion problems with this restriction.

In collective smoothing (3), one designs multigrid smoothers for the entire KKT system (parameter, forward, and adjoint) at once [15, 16]. Collective smoothers also tend to either require $R \approx I$, e.g., [62], or substantially degrade in performance as the regularization parameter decreases, e.g., [6].

1.4. Commentary on solving the preconditioner subsystems. Applying our preconditioner (5) requires the solution of two subsystems with coefficient operators

$$\alpha R^*R + \rho T^*T \tag{7}$$

and

$$B^*B + \rho A^*A, \tag{8}$$

respectively. This can be a challenge. However, reduced Hessian preconditioning and KKT preconditioning for large-scale inverse problems with highly informative data are fundamentally difficult endeavors, and the operators (7) and (8) have many advantages over the alternatives.

To begin with, we typically have easy access to the entries of the concrete matrix representations of these operators.[7] Thus we have at our disposal the entire arsenal of symmetric positive definite sparse preconditioning techniques that deal with matrix entries, e.g., incomplete factorizations, factorized sparse approximate inverses [30], and modern multilevel techniques including algebraic multigrid and hierarchical interpolative factorizations [41]. This stands in direct contrast to the reduced Hessian, which is dense owing to the inverses of the forward and adjoint operators within it, and as such may be accessed only via matrix-vector multiplies.

Additionally, the data misfit Hessian (which often acts as a compact operator) and the regularization operator (which often acts as a differential operator) tend to act in opposition to each other by construction.[8] Since the reduced Hessian is the sum of these operators, it is difficult to design preconditioners that are effective for both terms in the reduced Hessian at the same time. In contrast, the different terms in our subsystems tend not to act in opposition to each other.

In typical applications R^*R is chosen to be an elliptic differential operator, and T is either identity-like, or acts like a differential operator. Thus there is good reason to believe that multilevel techniques will be effective on the system $\alpha R^*R + \rho T^*T$ in situations of practical interest. A similar argument applies to $B^*B + \rho A^*A$ whenever the forward operator A is amenable to multilevel techniques. In the numerical results section (section 6), we see that for a source inversion problem with an elliptic PDE constraint, replacing the two subsystem solves with a few algebraic multigrid V-cycles results in nearly the same convergence rate as performing the solves exactly.

Of course, the operators in our subsystems are squared, and such squaring should always be done with caution. However, subsystems involving squared operators are also present in state-of-the-art preconditioners that have been proposed in the literature (see sections 1.3.2 and 1.3.3). In particular, a matrix spectrally equivalent to $B^*B + \rho A^*A$ shows up in the preconditioner proposed in [44].

1.5. Setting and notation. For the purposes of this paper we consider the case for which all spaces are finite dimensional Hilbert spaces, as might arise in stable discretize-then-optimize methods [36] for infinite dimensional problems. To fix ideas, consider the case of an infinite dimensional function space $\mathcal{U}_\infty$ approximated by a finite dimensional function space $\mathcal{U}$, the elements of which are in turn represented on a computer by lists of degrees of freedom in $\mathbb{R}^n$ corresponding to a potentially nonorthogonal basis $\Theta : \mathbb{R}^n \to \mathcal{U}$. Schematically,

$$\underset{\substack{\infty\text{-dimensional}\\ \text{function space}}}{\mathcal{U}_\infty} \approx \underset{\substack{n\text{-dimensional}\\ \text{function space}}}{\mathcal{U}} \overset{\Theta^{-1}}{\underset{\Theta}{\rightleftharpoons}} \underset{\substack{\text{representation}\\ \text{space}}}{\mathbb{R}^n}.$$

Here we work in intermediate finite dimensional function spaces like $\mathcal{U}$. In a representation space associated with a particular nonorthogonal basis, all formulas from this

[7] Although (dense) inverses of mass matrices can arise in concrete representations of these subsystems due to the adjoint operation, these inverse mass matrices can typically be replaced with spectrally equivalent sparse lumped mass approximations.

[8] By "act in opposition," we mean that modes that are amplified by one operator tend to be diminished by the other operator, and vice versa. This is discussed more in section 4.

paper remain essentially the same, except linear operators are replaced with matrix representations (arrays of numbers), abstract vectors are replaced with their concrete representations (lists of numbers), and Gram matrices (mass matrices) and their inverses appear in various locations to account for the Riesz representation theorem for adjoints in a nonorthogonal basis.

The parameter q, state u, adjoint η, and observations y are assumed to reside in finite dimensional Hilbert spaces $\mathcal{Q}$, $\mathcal{U}$, $\mathcal{V}$, and $\mathcal{Y}$ with dimensions n_q, n_u, n_u, and n_{obs}, respectively. Linear operators, e.g., $A : \mathcal{U} \to \mathcal{V}$, are viewed as abstract mappings between vector spaces, without reference to any particular basis, except in the case where the domain and/or range are of the form $\mathbb{R}^n$. Although we work with operators, we make routine use of standard results for matrices that are easily extended to the finite dimensional linear operator case, such as the existence and properties of eigenvalues of certain classes of operators, and the existence of the singular value decomposition. Transferring these results from the matrix setting to the finite dimensional linear operator setting is a straightforward process that involves working with the matrix representations of the operators in bases that are orthonormal with respect to the inner products on their domains and ranges.[9] Concatenation of linear operators such as BA denotes composition of linear maps, and concatenation of a linear operator with a vector, as in Au, denotes the action of the operator on the vector. Adjoints of operators are denoted by superscript stars, as in A^*. Superscript stars on a vector denote the linear functional that takes inner products with that vector. Namely, $u^* : v \mapsto (u, v)$, where $(\cdot, \cdot)$ is the inner product for the space u resides in. Functions of a linear operator such as inverses and square roots (where defined) are denoted in the standard way, i.e., $A^{-1}, A^{1/2}$. Unless otherwise noted, the norm of a vector, e.g., $\|u\|$, is the norm associated with the Hilbert space the vector resides in, and the norm of an operator, e.g., $\|A\|$, is the induced norm associated with the norms on the domain and range spaces of the operator. Block operators, such as

$$\begin{bmatrix} X & Y \\ Z & W \end{bmatrix} : \text{domain}(X) \oplus \text{domain}(Y) \to \text{range}(X) \oplus \text{range}(Z),$$

are defined by the blockwise action of their constituent operators, in the usual way, and with the expected consistency restrictions on the domains and ranges of the various blocks. Empty blocks are assumed to contain the zero operator with the appropriate domain and range. We use the notation $\Lambda = \text{diag}(\lambda_k)_{n,m}$ to denote the linear map $\Lambda : \mathbb{R}^m \to \mathbb{R}^n$ whose matrix representation in the standard basis is diagonal, with kth diagonal entry λ_k. Likewise, when we write $\Phi = \begin{bmatrix} \phi_1 & \phi_2 & \dots & \phi_m \end{bmatrix}$ for an operator $\Phi : \mathbb{R}^m \to \mathcal{X}$ and vectors $\phi_k \in \mathcal{X}$, we mean that ϕ_k is the result of applying Φ to the kth standard basis vector in $\mathbb{R}^k$ (ϕ_k is the "kth column" of Φ). An operator is said to be square if the dimensions of the domain and range are equal, and rectangular if the dimensions of the domain and range might differ.

The maximum and minimum singular values of an operator Y are denoted $\sigma_{\max}(Y)$ and $\sigma_{\min}(Y)$, respectively. Similarly, the maximum and minimum eigenvalues of an operator X with strictly real eigenvalues are denoted $\lambda_{\max}(X)$ and $\lambda_{\min}(X)$, respectively. The condition number of an operator X is denoted $\text{cond}(X)$.

2. Derivation of the preconditioner. The preconditioner in (5) is derived from a block diagonal approximation to the KKT operator associated with an aug-

[9] Note that such matrix representations with respect to orthonormal bases are generally not the same as the matrix representations that arise computationally within, say, a finite element method.

mented Lagrangian formulation of optimization problem (3). In the following derivation, it will be convenient to group the parameter and state variables into a single vector $x := \begin{bmatrix} q \\ u \end{bmatrix}$. With this grouping, optimization problem (3) takes the following standard quadratic programming form:

$$\begin{aligned} \min_x \quad & \frac{1}{2} x^* M x - g^* x \\ \text{such that} \quad & Cx = f, \end{aligned} \tag{9}$$

where $g := \begin{bmatrix} 0 \\ B^* y \end{bmatrix}$, $C := \begin{bmatrix} T & A \end{bmatrix}$, and M is the (generally singular) operator,

$$M := \begin{bmatrix} \alpha R^* R & \\ & B^* B \end{bmatrix}.$$

The KKT operator from (4) then becomes

$$K := \begin{bmatrix} \alpha R^* R & & T^* \\ & B^* B & A^* \\ T & A & \end{bmatrix} = \begin{bmatrix} M & C^* \\ C & \end{bmatrix}. \tag{10}$$

For nonsingular M, it is well established [46] that the positive definite block diagonal preconditioner

$$\begin{bmatrix} M & \\ & C M^{-1} C^* \end{bmatrix} \tag{11}$$

clusters the eigenvalues of the preconditioned operator onto at most three distinct values. Note that the positive operator $CM^{-1}C^*$ is the negative Schur complement for the adjoint variable. Since the objective block M is singular whenever B is not full rank (i.e., in the case of limited observations), we cannot directly use this result. However, (9) has the same solution as the following augmented optimization problem:

$$\begin{aligned} \min_x \quad & \frac{1}{2} x^* M x - g^* x + \frac{\rho}{2} \left\| Cx - f \right\|^2 \\ \text{such that} \quad & Cx = f, \end{aligned}$$

where the constraint is enforced strictly, but an additional quadratic penalty term is added to the objective function to further penalize constraint violations when an iterate is away from the optimal point. The KKT operator for this augmented optimization problem is

$$\begin{bmatrix} M + \rho C^* C & C^* \\ C & \end{bmatrix}. \tag{12}$$

With this augmentation, the objective block is now nonsingular, provided that M is coercive on the null space of C (i.e., the optimization problem is well-posed).

The positive definite block diagonal preconditioner analogous to (11) but based on the augmented KKT operator (12) is

$$\begin{bmatrix} M + \rho C^* C & \\ & C(M + \rho C^* C)^{-1} C^* \end{bmatrix}. \tag{13}$$

This preconditioner clusters the spectrum of the original (nonaugmented) KKT operator onto the union of two well-conditioned intervals [35]. However, this preconditioner is not practical since it is computationally difficult to perform solves $(M + \rho C^*C)^{-1}$, as well as apply the Schur complement $C(M + \rho C^*C)^{-1}C^*$ and its inverse. Thus we construct the preconditioner in (5) by replacing these blocks with cheaper approximations.

Intuitively, when ρ is large, constraint violations are more strongly penalized by the objective, so the adjoint variable does not need to "work as hard" to enforce the constraint. This manifests in better conditioning of the Schur complement for the adjoint, $C(M + \rho C^*C)^{-1}C^*$. Indeed, it is easy to see that $C(M + \rho C^*C)^{-1}C^* \to \frac{1}{\rho} I$ as $\rho \to \infty$. To this end, we expect the approximate preconditioner,

$$\begin{bmatrix} M + \rho C^*C & \\ & \frac{1}{\rho} I \end{bmatrix}, \tag{14}$$

to perform well when ρ is large. The preconditioner (14) is, essentially, a mechanism for using an unconstrained penalty method to precondition a constrained optimization problem.

The augmented objective block, $M + \rho C^*C$, takes the form

$$M + \rho C^*C = \begin{bmatrix} \alpha R^*R + \rho T^*T & \rho T^*A \\ \rho A^*T & B^*B + \rho A^*A \end{bmatrix}.$$

Since this 2×2 block operator is difficult to solve, we cannot use preconditioner (14) directly, and must make further approximations. In particular, the off-diagonal blocks are scaled by ρ, so when ρ is small we expect the relative importance of these blocks to be reduced. Dropping the off-diagonal blocks in $M + \rho C^*C$ and then substituting the result into (14) yields our overall 3×3 block diagonal preconditioner (5),

$$P := \begin{bmatrix} \alpha R^*R + \rho T^*T & & \\ & B^*B + \rho A^*A & \\ & & \frac{1}{\rho} I \end{bmatrix}.$$

One hopes that it is possible to choose ρ large enough that the Schur complement is well-approximated by $\frac{1}{\rho} I$, but at the same time small enough that the objective block is well-preconditioned by the block diagonal approximation. Our theory and numerical results in subsequent sections suggest that these competing interests can be balanced by choosing $\rho = \sqrt{\alpha}$, provided that the inverse problem is appropriately regularized. In the next section we provide an abstract theoretical analysis of the preconditioner without making any assumptions about the value of ρ. A more specific analysis for source inversion problems with spectral filtering regularization, which motivates our choice of ρ, is performed in section 5.

3. Abstract analysis of the preconditioner. In this section we analyze the preconditioned KKT operator, showing that it is well-conditioned if bounds on the arithmetic and geometric means of certain damped projectors are satisfied. First, we highlight the structure of the preconditioned KKT operator, state the necessary arithmetic and geometric mean bounds, and recall a prerequisite result from Brezzi theory. Then we prove bounds on the condition number of the preconditioned KKT operator based on the arithmetic and geometric mean bounds.

3.1. Prerequisites.

3.1.1. Preconditioned KKT operator. Let E denote the symmetrically preconditioned KKT operator,

$$E := P^{-1/2} K P^{-1/2},$$

with P and K defined in (5) and (10), respectively. Direct calculation shows that the symmetrically preconditioned KKT operator has the following block structure:

$$E = \begin{bmatrix} I - F^*F & & F^* \\ & I - G^*G & G^* \\ F & G & \end{bmatrix}, \tag{15}$$

where the operators F and G are defined as

$$F := T \left(\frac{\alpha}{\rho} R^* R + T^* T \right)^{-1/2}, \quad G := A \left(\frac{1}{\rho} B^* B + A^* A \right)^{-1/2}.$$

For convenience, we further denote the objective and constraint blocks of the preconditioned system by X and Y, respectively, where

$$X := \begin{bmatrix} I - F^*F & \\ & I - G^*G \end{bmatrix}, \quad Y := \begin{bmatrix} F & G \end{bmatrix}, \tag{16}$$

so that the preconditioned KKT operator takes the form

$$E = \begin{bmatrix} X & Y^* \\ Y & \end{bmatrix}. \tag{17}$$

3.1.2. Arithmetic and geometric mean assumptions. The quality of the preconditioner depends on the arithmetic and geometric means of the following two *damped projectors*:[10]

$$Q_R := FF^* = T \left(\frac{\alpha}{\rho} R^* R + T^* T \right)^{-1} T^* \tag{18}$$

and

$$Q_J := GG^* = A \left(\frac{1}{\rho} B^* B + A^* A \right)^{-1} A^*.$$

Note that if T is invertible, we have

$$Q_J = T \left(\frac{1}{\rho} J^* J + T^* T \right)^{-1} T^*, \tag{19}$$

where

$$J := -BA^{-1}T \tag{20}$$

is the *parameter-to-observable map* that transforms candidate parameter fields into predicted observations.

[10] Recall that $X(\gamma I + X^*X)^{-1}X^*$ approximates the orthogonal projector onto the column space of X for small γ. With this in mind, one can view an operator of the form $X(Y^*Y + X^*X)^{-1}X^*$ as an approximate projector onto the column space of X, damped by the operator Y. We call such operators *damped projectors*.

As damped projectors, it is easy to show that the eigenvalues of Q_R and Q_J are bounded between 0 and 1. The degree to which the eigenvalues of Q_R are damped below 1 is controlled by the strength of the damping term $\frac{\alpha}{\rho}R^*R$ and its interaction with the eigenstructure of T. Similarly, the degree of damping of the eigenvalues of Q_J is controlled by the strength of the damping term $\frac{1}{\rho}J^*J$ and its interaction with the eigenstructure of T (or the interaction of the damping term $\frac{1}{\rho}B^*B$ with the eigenstructure of A, when T is not invertible).

ASSUMPTION 1 (damped projector AM-GM bounds). *We assume there exist constants β, δ such that the following bounds on the spectrum of the arithmetic and geometric means of the damped projectors hold:*

(a) $$0 < \delta \leq \frac{1}{2}\lambda_{min}\left(Q_R + Q_J\right),$$

(b) $$\lambda_{max}\left(Q_R Q_J\right)^{1/2} \leq \beta < 1.$$

Theorem 7 will establish that the larger δ is and the smaller β is, the more effective preconditioner (5) is.

Qualitatively, if T is invertible and the regularization is chosen to act in opposition to the data misfit, as desired for the problem to be properly regularized based on the analysis that will be performed in section 4, then R will act strongly on vectors that J acts weakly on, and vice versa. Thus we expect the damping in Q_R to be strong where the damping in Q_J is weak, and vice versa. Consequently, it is reasonable to hypothesize that Assumption 1 will be satisfied with good constants for inverse problems that are properly regularized. Making this intuition precise requires careful analysis of the interaction between the eigenstructures of R, J, and T, which must be done on a case-by-case basis. We perform this analysis for the special case of source inversion problems with spectral filtering regularization in section 5, and expect similar behavior to hold in more general situations.

3.1.3. Brezzi theory for well-posedness of saddle point systems. The proof of the coercivity bound for our preconditioned KKT operator invokes Brezzi theory for saddle point systems [19, 25, 66]. In particular, we use a recently discovered bound in [42], which is slightly sharper than bounds derived from the classical theory. Here we state the prerequisite theorem (without proof) and refer the reader to [42] for more details. This theory can be stated in much greater generality than what we present here.

THEOREM 2 (see Krendl, Simoncini, and Zulehner [42]). *Let E be the saddle point system*

$$E = \begin{bmatrix} X & Y^* \\ Y & \end{bmatrix},$$

where $X : \mathcal{X} \to \mathcal{X}$ is self-adjoint and positive semidefinite, and $Y : \mathcal{X} \to \mathcal{Y}$. Further, suppose that

- *X is coercive on the kernel of Y, i.e.,*

$$0 < a \leq \inf_{\substack{z \in Ker(Y) \\ z \neq 0}} \frac{z^* X z}{\|z\|^2};$$

- *X is bounded, i.e., $\|X\| < b$;*

- *the singular values of Y are bounded from below, i.e.,*

$$0 < c \leq \sigma_{min}(Y).$$

Then the minimum singular value of E is bounded from below, with the bound

$$\frac{a}{1 + \left(\frac{b}{c}\right)^2} \leq \sigma_{min}(E). \tag{21}$$

3.2. Bound on the condition number of the preconditioned KKT operator. To apply Brezzi theory (Theorem 2) to our problem, we need a coercivity bound for X on the kernel of Y, a continuity bound for X on the whole space, and a coercivity bound on Y, where the constants for these bounds are denoted a, b, and c, respectively. We use the particular structure of the KKT operator (10), along with Assumption 1, to derive these bounds in section 3.2.1. In Proposition 3 we derive bounds for a and b, and then in Proposition 4 we derive a bound for c.

In section 3.2.2 we derive well-posedness and continuity bounds on the preconditioned KKT operator, E, and then combine these bounds to provide an upper bound on the condition number of E. Well-posedness of E is proven in Proposition 5, using Brezzi theory in the form of Theorem 2. Continuity of E is proven directly in Proposition 6. Finally, the overall condition number bound for E is given in Theorem 7.

3.2.1. Bounds on X and Y.

PROPOSITION 3 (bounds a, b for X). *The eigenvalues of X restricted to the kernel of Y are bounded below by $1-\beta$, where β is defined in Assumption 1. That is,*

$$0 < 1 - \beta \leq \inf_{\substack{z \in Ker(Y) \\ z \neq 0}} \frac{z^* X z}{\|z\|^2}.$$

Additionally,

$$\|X\| \leq 1.$$

Proof. For vectors $z \in \mathrm{Ker}(Y)$, we have

$$z^* X z = z^*(X + Y^* Y)z \geq \lambda_{\min}(X + Y^* Y)||z||^2. \tag{22}$$

This augmented operator has the following block structure:

$$X + Y^* Y = \begin{bmatrix} I - F^* F & \\ & I - G^* G \end{bmatrix} + \begin{bmatrix} F^* \\ G^* \end{bmatrix} \begin{bmatrix} F & G \end{bmatrix} = \begin{bmatrix} I & F^* G \\ G^* F & I \end{bmatrix}.$$

Thus the eigenvalues λ of $X + Y^* Y$ satisfy

$$\begin{bmatrix} I & F^* G \\ G^* F & I \end{bmatrix} \begin{bmatrix} v \\ \xi \end{bmatrix} = \lambda \begin{bmatrix} v \\ \xi \end{bmatrix}$$

or

$$\begin{bmatrix} & F^* G \\ G^* F & \end{bmatrix} \begin{bmatrix} v \\ \xi \end{bmatrix} = (\lambda - 1) \begin{bmatrix} v \\ \xi \end{bmatrix}. \tag{23}$$

Solving for u from the block equation associated with the first row block of (23) and substituting into the second yields

$$G^* F F^* G \xi = (\lambda - 1)^2 \xi.$$

Thus, the magnitudes of the shifted eigenvalues, $|\lambda - 1|$, are the square roots of the eigenvalues of G^*FF^*G. By a similarity transform, the eigenvalues of G^*FF^*G are the same as the eigenvalues of the operator FF^*GG^*, and by the second part of Assumption 1, we know that these eigenvalues are bounded above by β. Thus,

$$|\lambda - 1| \leq \lambda_{\max}(FF^*GG^*)^{1/2} \leq \beta,$$

which implies

$$1 - \beta \leq \lambda,$$

so that

$$z^*Xz \geq (1 - \beta) \|z\|^2,$$

from which the inf-sup bound directly follows.

Since FF^* and GG^* are damped projectors, their eigenvalues reside in the interval $[0, 1]$, as do the eigenvalues of F^*F and G^*G. Using the definition of X in (16), this implies that the singular values of X reside in the interval $[0, 1]$, and so we have the upper bound $||X|| \leq 1$. ∎

PROPOSITION 4 (bound c for Y). *The singular values of the preconditioned constraint are bounded below, with bound*

$$0 < \sqrt{2\delta} \leq \sigma_{min}(Y).$$

Proof. Since G is invertible, $Y = \begin{bmatrix} F & G \end{bmatrix}$ has full row rank. Thus the singular values of Y are the square roots of the eigenvalues of

$$YY^* = FF^* + GG^*.$$

Recalling the arithmetic mean assumption (Assumption 1(a)), we have

$$0 < \delta \leq \frac{1}{2}\lambda_{\min}(FF^* + GG^*) = \frac{1}{2}\lambda_{\min}(YY^*)$$

or

$$0 < \sqrt{2\delta} \leq \sigma_{\min}(Y).$$ ∎

3.2.2. Well-posedness, continuity, and conditioning of the preconditioned KKT operator, E.

PROPOSITION 5 (well-posedness of E). *The singular values of E have the following lower bound:*

$$0 < \frac{2}{3}(1 - \beta)\delta \leq \sigma_{min}(E).$$

Proof. Based on the results of Propositions 3 and 4, and the block structure of E from (17), we can apply bound (21) from Theorem 2 to E with $a = 1 - \beta$, $b = 1$, and $c^2 = 2\delta$. Doing this and then using the fact that $0 < \delta \leq 1$, we get the desired lower bound on the minimum singular value:

$$\sigma_{\min}(E) \geq \frac{1 - \beta}{1 + \frac{1}{2\delta}} = \frac{2(1 - \beta)\delta}{1 + 2\delta} \geq \frac{2}{3}(1 - \beta)\delta.$$ ∎

PROPOSITION 6 (continuity of E). *The singular values of E are bounded above by 2; i.e.,*

$$\sigma_{max}(E) \leq 2.$$

Proof. To prove the upper bound, we directly estimate the quantity $|w_1^* E w_2|$ for arbitrary w_1, w_2. Denote the blocks of w_1 and w_2 by

$$w_1 = \begin{bmatrix} p_1 \\ v_1 \\ \xi_1 \end{bmatrix}, \quad w_2 = \begin{bmatrix} p_2 \\ v_2 \\ \xi_2 \end{bmatrix}.$$

Recalling the blockwise definition of E from (15) and using the triangle inequality, we have

$$\begin{aligned} |w_1^* E w_2| &= \left| \begin{bmatrix} p_1^* & v_1^* & \xi_1^* \end{bmatrix} \begin{bmatrix} I - F^*F & & F^* \\ & I - G^*G & G^* \\ F & G & \end{bmatrix} \begin{bmatrix} p_2 \\ v_2 \\ \xi_2 \end{bmatrix} \right| \\ &= |p_1^*(I - F^*F)p_2 + p_1^* F^* \xi_2 + v_1^*(I - G^*G)v_2 + v_1^* G^* \xi_2 + \xi_1^* F p_2 + \xi_1^* G v_2| \\ &\leq |p_1^*(I - F^*F)p_2| + |p_1^* F^* \xi_2| + |v_1^*(I - G^*G)v_2| + |v_1^* G^* \xi_2| + |\xi_1^* F p_2| \\ &\quad + |\xi_1^* G v_2|. \end{aligned} \tag{24}$$

Since the operators F and G have singular values between zero and one, we can eliminate all of the intermediate operators in (24), yielding

$$|w_1^* E w_2| \leq \|p_1\| \|p_2\| + \|p_1\| \|\xi_2\| + \|v_1\| \|v_2\| + \|v_1\| \|\xi_2\| + \|\xi_1\| \|p_2\| + \|\xi_1\| \|v_2\|. \tag{25}$$

By Cauchy–Schwarz, three of the terms on the right-hand side of (25) can be estimated as follows:

$$\begin{aligned} &\|p_1\| \|p_2\| + \|v_1\| \|\xi_2\| + \|\xi_1\| \|v_2\| \\ &\qquad \leq \left(\|p_1\|^2 + \|v_1\|^2 + \|\xi_1\|^2 \right)^{1/2} \left(\|p_2\|^2 + \|v_2\|^2 + \|\xi_2\|^2 \right)^{1/2} \\ &\qquad = \|w_1\| \|w_2\|. \end{aligned}$$

The other three terms can be estimated similarly:

$$\|p_1\| \|\xi_2\| + \|v_1\| \|v_2\| + \|\xi_1\| \|p_2\| \leq \|w_1\| \|w_2\|.$$

Thus we have the overall estimate

$$|w_1^* E w_2| \leq 2 \|w_1\| \|w_2\|,$$

which implies $\sigma_{\max}(E) \leq 2$, as required. ∎

THEOREM 7 (conditioning of E). *We present*

$$\operatorname{cond}(E) \leq \frac{3}{(1-\beta)\delta}.$$

Proof. Divide the upper bound from Proposition 6 by the lower bound from Proposition 5. ∎

4. Spectral filtering and appropriate regularization assumptions. To better characterize the constants δ and β in the condition number bound in Theorem 7, in this section we propose *appropriate regularization assumptions* (Assumption 9) that limit the degree to which the inverse problem can be over- or under-regularized. These assumptions are motivated by an analysis of the error in the reconstruction

of the parameter (sections 4.1 and 4.2) and apply to spectral filtering regularization operators (Definition 8). Since one part of Assumption 9 (specifically, Assumption 9(b)) is novel, we discuss that part in greater detail. Much of the development we present leading up to (but not including) Assumption 9 mirrors the classical treatment presented in [29].

Since construction of spectral filtering regularization operators is too expensive for large-scale inverse problems with highly informative data, Assumption 9 is used for theoretical analysis only. In section 5 we will prove that satisfying Assumption 9 implies the existence of good constants δ and β for source inversion problems, thereby guaranteeing that our preconditioner will perform well on these problems.

4.1. The reduced problem and decomposition of error. Although we take a full space approach for solving optimization problem (3), for the purpose of analysis it is useful to consider the reduced version of the problem in which the constraint is eliminated by viewing the state u as an implicit function of the parameter q via solution of the state equation. This yields the following unconstrained optimization problem in q only:

$$\min_q \quad \frac{1}{2}\left\|Jq - y\right\|^2 + \frac{\alpha}{2}\left\|Rq\right\|^2, \tag{26}$$

where we recall from (20) that the *parameter-to-observable map* J is defined as $J := -BA^{-1}T$. The solution q to this reduced problem is the solution to the normal equations,

$$Hq = J^*y, \tag{27}$$

where

$$H := J^*J + \alpha R^*R \tag{28}$$

is the Hessian of the reduced optimization problem (26), which we call the *reduced Hessian*. The reduced Hessian has been the target of much of the previous work on preconditioners for inverse problems (see section 1.3.1), including the method to which we numerically compare our preconditioner in section 6.

From an optimization perspective, the purpose of the regularization is to make optimization problem (26) well-posed by introducing curvature in the objective function in directions that are in the (numerical) null space of J. However, in the context of inverse problems the regularization is primarily seen as a means of stabilizing the inversion with respect to noise in the observations.

Recall from (1) that the observations we use for the inversion are corrupted by additive noise ζ via the formula

$$y = y_{\text{true}} + \zeta = Jq_{\text{true}} + \zeta, \tag{29}$$

where q_{true} is the unknown true parameter and $y_{\text{true}} = Jq_{\text{true}}$ are the observations that would have been obtained if there were no noise. Substituting (29) into (27) and then subtracting the result from q_{true}, we see that the error takes the form

$$q_{\text{true}} - q = e_\zeta + e_q,$$

consisting of a term

$$e_\zeta := -\left(J^*J + \alpha R^*R\right)^{-1} J^*\zeta \tag{30}$$

that depends on the noise, and a term

$$e_q := \left(I - \left(J^*J + \alpha R^*R\right)^{-1} J^*J\right) q_{\text{true}} \tag{31}$$

that does not. From the form of (30) and (31), a trade-off is evident: Strengthening the regularization tends to reduce e_ζ at the expense of increasing e_q, and weakening the regularization tends to reduce e_q at the expense of increasing e_ζ. To achieve a good reconstruction of the parameter, it is desirable for both of these terms to be as small in magnitude as possible. To investigate this trade-off in more detail, we restrict our subsequent analysis to the special case of spectral filtering regularization, which we define and discuss in the following section. This will provide convenient bases to diagonalize the operators $-\left(J^*J + \alpha R^*R\right)^{-1} J^*$ and $\left(I - J^*J + \alpha R^*R\right)^{-1} J^*J$, and hence allow us to understand the errors e_ζ and e_q in a per-component manner.

4.2. Spectral filtering regularization.

DEFINITION 8. *An operator R is a spectral filtering regularization operator for a linear inverse problem with parameter-to-observable map J if R and J share a common basis of right singular vectors. That is, there exist*

- *unitary operators $U : \mathbb{R}^{n_{obs}} \to \mathcal{Y}$, $V : \mathbb{R}^{n_q} \to \mathcal{Q}$, and $\Phi : \mathbb{R}^{n_q} \to \mathcal{Q}$, and*
- *nonnegative diagonal operators $\Sigma_J = \text{diag}(d_k)_{n_{obs},n_q}$ and $\Sigma_R = \text{diag}(r_k)_{n_q,n_q}$*

such that

$$\begin{cases} J = U\Sigma_J\Phi^*, \\ R = V\Sigma_R\Phi^*. \end{cases} \tag{32}$$

By convention we order the singular values d_k of J in descending order ($d_k \geq d_{k+1}$). In the case where $n_{obs} < n_q$, for convenience we define $d_k := 0$ for $k = n_{obs}+1, \ldots, n_q$. Note that the descending order for d_k forces an order (possibly nonmonotone) for the singular values r_k of R. We use ϕ_k to denote the kth right singular vector shared by J and R. That is, $\Phi = \begin{bmatrix}\phi_1 & \phi_2 & \ldots & \phi_{n_q}\end{bmatrix}$.

Spectral filtering regularization is ideally suited for inverse problems—by manipulating the regularization singular values r_k, one can selectively filter out undesirable components of the parameter from the reconstruction without affecting the reconstruction of the desirable components. The larger r_k is, the more component ϕ_k is penalized, and vice versa. Limiting cases of spectral filtering regularization include the following:

- identity regularization ($R = I$), where all singular vectors are penalized equally and
- truncated SVD, where singular vectors ϕ_k are not penalized at all if d_k is above a given threshold, but are penalized infinitely[11] otherwise.

Spectral filtering regularization is routinely used for small to moderate sized inverse problems, and for large inverse problems that admit low-rank approximations to the parameter-to-observable map. However, aside from identity regularization, spectral filtering regularization is generally computationally infeasible for large-scale inverse problems with highly informative data. In fact, spectral filtering regularization requires computing the dominant singular vectors and singular values of J in order to construct R, and the number of dominant singular vectors of J scales with the informativeness of the data. Thus we view spectral filtering as an idealized form of

[11] That is, the reconstruction of the component of q in the direction ϕ_k is set to zero.

regularization that practical regularization operators attempt to approximate. For a more comprehensive discussion of spectral filtering and its relation to other regularizations, we refer the reader to the classic monograph [29].

For spectral filtering regularization, we can formulate expressions for the errors in the reconstruction on a per-component manner. Substituting the singular value decomposition factors from (32) into the error expressions from (30) and (31), and then performing some algebraic manipulations yields

$$e_\zeta = -\Phi \operatorname{diag}\left(\frac{d_k}{d_k^2 + \alpha r_k^2}\right) U^* \zeta, \tag{33}$$

$$e_q = \Phi \operatorname{diag}\left(\frac{\alpha r_k^2}{d_k^2 + \alpha r_k^2}\right) \Phi^* q_{\text{true}}. \tag{34}$$

From (33), we see that the regularization should not be weak (small αr_k^2) in directions ϕ_k to which the observations are insensitive (small d_k^2). Otherwise the noise associated with observations of those directions will be highly amplified, leading to large errors. In such a scenario we say that the problem is *under-regularized*.

On the other hand, (34) shows that strong regularization can also lead to large errors. In directions ϕ_k for which observation data is lacking (or dominated by noise), there is no hope to reconstruct the component of the parameter in that direction, so some degree of error in e_q is to be expected. However, if d_k is large, then the observations are highly sensitive to changes to the parameter in direction ϕ_k, so it is likely that the observations associated with direction ϕ_k contain more signal than noise. That is, when d_k is large, it is likely that the component of the parameter q_{true} in direction ϕ_k can, in principle, be inferred from the data. Hence, if the regularization is strong (large αr_k^2) in directions for which the parameter-to-observable map is also strong (large d_k^2), the reconstruction will contain substantial *unnecessary* error due to the regularization. In this scenario we say that the problem is *over-regularized*. To simultaneously avoid under- and over-regularization, the regularization should be strong in directions where the parameter-to-observable map is weak, and weak in directions where the parameter-to-observable map is strong.

4.3. Appropriate regularization assumptions. In light of the preceding discussion of over- and under-regularization error and spectral filtering, we propose the following *appropriate regularization assumptions*.

ASSUMPTION 9 (appropriate regularization). *There exist constants μ and ν such that*

(a) $0 < \mu \leq d_k^2 + \alpha r_k^2$,

(b) $d_k r_k \leq \nu < \infty$

for all $k = 1, 2, \ldots, n_q$.

Assumption 9(a) is already required for linear optimization problem (3) to be well-posed. It says that the regularization cannot be arbitrarily small in basis directions ϕ_k to which the observations are insensitive, but allows the regularization to be arbitrarily small in directions ϕ_k to which the observations are sensitive. In contrast, Assumption 9(b) prevents the regularization from being large in basis directions ϕ_k to which the observations are sensitive, but still allows the regularization singular values to diverge ($r_k \to \infty$ as $k \to \infty$), as long as the sensitivity of the observations to changes to

the parameter, d_k, goes to zero in the inverse manner. Informally, Assumption 9(a) says that the problem is not under-regularized, and Assumption 9(b) says that the problem is not over-regularized.

Since Assumption 9(a) is standard, we do not discuss it further. The motivation for Assumption 9(b) is less obvious, so we provide a more in-depth discussion of it. To begin with, the multiplicative nature of Assumption 9(b) makes it a relatively weak assumption compared to other possible candidates for preventing over-regularization. In particular, observe that the eigenvalues of the regularization preconditioned reduced Hessian, $R^{-*}HR^{-1}$, are $d_k^2/r_k^2 + \alpha$. Thus situations in which the strength of the regularization operator on a mode is inversely proportional to how informed that mode is (i.e., $r_k \approx \frac{1}{d_k}$) can lead to arbitrarily poor conditioning of the regularization preconditioned reduced Hessian while still satisfying Assumption 9(b) with a constant of order one.

An instructive model problem that illustrates Assumption 9(b) is the Poisson source inversion problem on a rectangular domain, with Laplacian regularization, zero Dirichlet boundary conditions for both A and R, and distributed observations of the first n_{obs} Fourier modes of the state variable in the domain. That is,

- $T = I$,
- $A = R = \Delta_D$, where Δ_D is the Laplacian operator with zero Dirichlet boundary conditions, and
- $B : \mathcal{U} \to \mathbb{R}^{n}_{\text{obs}}$ is a wide rectangular operator with Fourier modes as right singular vectors (the same as A and R), but with singular values $\sigma_k = 1$, $k = 1, \dots, n_{\text{obs}}$.

Recalling that $J = -BA^{-1}T = -B\Delta_D^{-1}$, we see that

$$d_k = \begin{cases} \frac{1}{\lambda_k}, & k = 1, \dots, n_{\text{obs}}, \\ 0, & k > n_{\text{obs}}, \end{cases}$$

where λ_k is the kth eigenvalue of the Laplacian Δ_D. At the same time, the singular values of R are $r_k = \lambda_k$. Thus $d_k r_k = 1$ for $k = 1, \dots, n_{\text{obs}}$ and $d_k r_k = 0$ for $k > n_{\text{obs}}$, so Assumption 9(b) holds with constant $\nu = 1$, regardless of the number of observations, n_{obs}.

5. Analysis of the source inversion problem with spectral filtering regularization. In section 3.1.2 we hypothesized that the damped projector arithmetic and geometric mean assumptions (Assumption 1) are satisfied with good constants δ and β whenever an inverse problem is properly regularized. Then in section 4 we formulated another assumption (Assumption 9) that quantifies the concept of proper regularization for spectral filtering regularization operators. Here we show that Assumption 9 implies Assumption 1 for the source inversion problem. Specifically, in Theorem 11 and Corollary 12 we prove quantitative bounds on the constants δ and β for source inversion problems that are neither over- nor under-regularized in the manner made precise by Assumption 9. The more appropriate to the problem the regularization is, the better the bounds.

DEFINITION 10. *An inverse problem is a source inversion problem if the parameter q being inverted for is the right-hand side of the state equation. That is, $T = -I$, and state equation* (2) *takes the form*

$$Au = q.$$

THEOREM 11. *Let R be a spectral filtering regularization operator for a source inversion problem (see Definitions 8 and 10). If R satisfies appropriate regularization Assumption 9 with constants μ and ν, then Assumption 1 is also satisfied with constants*

$$\delta = \frac{1}{2}\left(1 + \frac{\alpha}{\rho^2}\nu^2\right)^{-1} \quad \text{and} \quad \beta = \left(1 + \frac{1}{\rho}\mu\right)^{-1/2}.$$

Proof. For δ, we seek a lower bound on the eigenvalues of the arithmetic mean of the damped projectors Q_R and Q_J (as defined in (18) and (19), respectively), while for β we seek an upper bound on their geometric mean. For source inversion problems these damped projectors take the forms

$$Q_R = \left(\frac{\alpha}{\rho}R^*R + I\right)^{-1} \quad \text{and} \quad Q_J = \left(\frac{1}{\rho}J^*J + I\right)^{-1}.$$

Furthermore, for spectral filtering regularization, R^*R and J^*J share the same eigenvectors and have eigenvalues r_k^2 and d_k^2, respectively. Thus the eigenvalues δ_k of the arithmetic mean $\frac{1}{2}(Q_R + Q_J)$ can be estimated as

$$\delta_k = \frac{1}{2}\left(\frac{1}{\frac{\alpha}{\rho}r_k^2 + 1} + \frac{1}{\frac{1}{\rho}d_k^2 + 1}\right) \geq \frac{1}{2}\left(1 + \frac{\alpha}{\rho^2}d_k^2 r_k^2\right)^{-1} \geq \frac{1}{2}\left(1 + \frac{\alpha}{\rho^2}\nu^2\right)^{-1}.$$

In the first inequality we have combined fractions, and used the nonnegativity of r_k^2, d_k^2 and monotonicity of the function $f(x) = x/(a+x)$. In the second inequality we have used Assumption 9(b).

Similarly, we use the Assumption 9(a) to bound the eigenvalues β_k of the geometric mean $(Q_R Q_J)^{1/2}$ as

$$\beta_k = \left(\frac{1}{\frac{\alpha}{\rho}r_k^2 + 1} \cdot \frac{1}{\frac{1}{\rho}d_k^2 + 1}\right)^{1/2} \leq \left(1 + \frac{\alpha}{\rho}r_k^2 + \frac{1}{\rho}d_k^2\right)^{-1/2} \leq \left(1 + \frac{1}{\rho}\mu\right)^{-1/2}.$$

∎

The following corollary of Theorem 11 shows that the preconditioner will be effective in the low-to-moderate regularization regime ($\alpha \leq 1$) if we choose $\rho = \sqrt{\alpha}$.

COROLLARY 12. *If the conditions of Theorem 11 are satisfied, and $\alpha \leq 1$, and the regularization parameter is chosen as $\rho = \sqrt{\alpha}$, then Assumption 1 is satisfied, with constants*

$$\delta = \frac{1}{2}\left(1 + \nu^2\right)^{-1} \quad \text{and} \quad \beta = (1 + \mu)^{-1/2}.$$

Proof. Substituting $\rho = \sqrt{\alpha}$ into the results of Theorem 11, we immediately have the desired lower bound on the arithmetic mean of damped projectors with constant $\delta = \frac{1}{2}\left(1 + \nu^2\right)^{-1}$. For the geometric mean, Theorem 11 implies

$$\lambda_{\max}\left(Q_R Q_J\right)^{1/2} \leq \left(1 + \alpha^{-1/2}\mu\right)^{-1/2}.$$

But note that for $\alpha \leq 1$ we have

$$\left(1 + \alpha^{-1/2}\mu\right)^{-1/2} \leq (1 + \mu)^{-1/2}, \tag{35}$$

and so we get the desired upper bound with $\beta = (1 + \mu)^{-1/2}$. ∎

6. Numerical results. We apply our method to a Poisson source inversion problem with pointwise observations randomly distributed throughout a rectangular domain $\Omega = [0, 1.45] \times [0, 1]$, using Laplacian regularization. Specifically, we take q, u, and v to reside in the space of continuous piecewise linear functions on a uniform triangular mesh with mesh size parameter h, with the L^2 inner product. The state equation,

$$Au := \Delta_D u = q,$$

is the Poisson equation discretized by the finite element method, with homogeneous Dirichlet boundary conditions enforced by the symmetric Nitsche method [49]. Pointwise observations of the form

$$y_k = u(x_k)$$

are taken for a collection of points $\{x_k \in \Omega\}_{k=1}^{n_{\text{obs}}}$, shown in Figure 1. Noise is not included in the inverse problem since we are interested in preconditioners for the low noise, high data, small regularization limit. The regularization operator is defined by

$$R^*R := \Delta_N + tI,$$

where Δ_N is the Laplacian operator with Neumann boundary conditions discretized by the finite element method, and $t = 1/10$. The combined operator R^*R is used directly; in fact, the solution algorithm does not require R explicitly.[12]

FIG. 1. *Left: True source field q_{true} used for all inversions. Center: Reconstruction q for the case of $n_{obs} = 2000$ observations with regularization parameter $\alpha = 10^{-8}$ and mesh size $h = \sqrt{2} \cdot 10^{-2}$. Right: Observation locations x_k, denoted by dots.*

The true source field, q_{true}, used to generate the observations, y_k, is a grayscale image of the Peter O'Donnell Jr. Building at the University of Texas at Austin, scaled to contain values in $[0, 1]$, shown in Figure 1. The combination of sharp edges and smooth features in this image make this an ideal test case for highly informative data and small regularization.

Abstract vectors q, u, η are represented concretely by lists of nodal degrees of freedom $\mathbf{q}, \mathbf{u}, \boldsymbol{\eta}$, respectively. The norm of a concrete vector, e.g., $\|\mathbf{q}\|$, is the Euclidean norm (square root of the sum of the squares of the entries). Since we use uniform meshes and present only relative errors, this is spectrally equivalent to using the function space L^2 norm on the underlying function being represented by the concrete vector. We use the FEniCS [43] package to assemble concrete matrix representations of A, R^*R, T, and I, which are denoted $\mathbf{A}$, $\mathbf{R^*R}$, $\mathbf{T}$, and $\mathbf{W}$, respectively. The diagonal lumped mass matrix is denoted $\mathbf{W}_L$, with diagonal entries given by row sums of the mass matrix: $(\mathbf{W}_L)_{ii} = \sum_j \mathbf{W}_{ij}$. The concrete sparse matrix representation of the

[12] Both A and R^*R should be viewed as finite dimensional discretizations of densely defined unbounded operators acting $L^2(\Omega)$.

observation operator is denoted $\mathbf{B}$. Its (i,j) entry, $\mathbf{B}_{ij}$, equals the evaluation of the jth basis function at the ith observation location.

In a concrete basis, the KKT operator (4) becomes

$$\begin{bmatrix} \alpha\mathbf{R}^*\mathbf{R} & & -\mathbf{W} \\ & \mathbf{B}^T\mathbf{B} & \mathbf{A}^T \\ -\mathbf{W} & \mathbf{A} & \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \mathbf{u} \\ \boldsymbol{\eta} \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{B}^T y \\ 0 \end{bmatrix}. \tag{36}$$

The reconstructed function q based on the exact[13] solution of this KKT system with regularization parameter $\alpha = 10^{-8}$ is shown in Figure 1.

In a concrete basis the preconditioner (5) becomes

$$\mathbf{P} = \begin{bmatrix} \alpha\mathbf{R}^*\mathbf{R} + \rho\mathbf{W} & & \\ & \mathbf{B}^T\mathbf{B} + \rho\mathbf{A}^T\mathbf{W}^{-1}\mathbf{A} & \\ & & \frac{1}{\rho}\mathbf{W} \end{bmatrix}. \tag{37}$$

In our numerical experiments, we consider three variants of this preconditioner.

- *BDAL, exact:* All solves in preconditioner (37) are performed exactly.
- *BDAL, lumped mass, exact:* The mass matrix $\mathbf{W}$ is replaced with the lumped mass matrix $\mathbf{W}_L$, but preconditioner solves are performed exactly with this replacement.
- *BDAL, lumped mass, multigrid:* The mass matrix is replaced by the lumped mass matrix, and the solves for $\alpha\mathbf{R}^*\mathbf{R} + \rho\mathbf{W}_L$ and $\mathbf{B}^T\mathbf{B} + \rho\mathbf{A}^T\mathbf{W}_L^{-1}\mathbf{A}$ are replaced by a small number of algebraic multigrid V-cycles.

For algebraic multigrid we use the root-node smoothed aggregation [50, 63] method implemented in PyAMG [10], with the default settings. One V-cycle is used for $\alpha\mathbf{R}^*\mathbf{R} + \rho\mathbf{W}_L$, and three V-cycles are used for $\mathbf{B}^T\mathbf{B} + \rho\mathbf{A}^T\mathbf{W}_L^{-1}\mathbf{A}$.

6.1. Convergence comparison. In Figure 2, we show a convergence comparison between between MINRES on the KKT system preconditioned by our block diagonal augmented Lagrangian preconditioner, and conjugate gradient on the reduced Hessian preconditioned by the regularization term (CG-HESS). For our block diagonal augmented Lagrangian preconditioner, we also show results for lumped mass and algebraic multigrid approximations to the subsystems being solved. The regularization, forward, and adjoint solves used for the reduced Hessian solve are all performed exactly. The mesh size is $h = \sqrt{2} \cdot 10^{-2}$, the number of observations is 2000, and the regularization parameter is $\alpha = 10^{-8}$. Error is measured with respect to the converged solution to the linear system (36), i.e., $\|\mathbf{q} - \mathbf{q}_k\| / \|\mathbf{q}\|$. This allows us to make a fair comparison between the reduced and full space methods.

In terms of Krylov iteration count, our preconditioner far outperforms regularization preconditioning on the reduced Hessian. The error in our method after three iterations is much less than the error after 50 iterations of regularization preconditioning on the reduced Hessian. Performance with the lumped mass approximation is almost identical to performance with exact solves. In the case with the multigrid approximation, we see roughly the same asymptotic convergence rate as the exact solve, but with a lag of 10 to 20 iterations. In our numerical experiments we also observed that MINRES with our "BDAL, lumped mass, multigrid" preconditioner takes considerably less time per iteration than CG on the reduced Hessian, which is expected since applying the reduced Hessian requires solving the forward and adjoint equations to a high tolerance within each CG iteration.

[13] By "exact," we mean that the result of a computation is accurate to tolerance 10^{-12} or smaller.

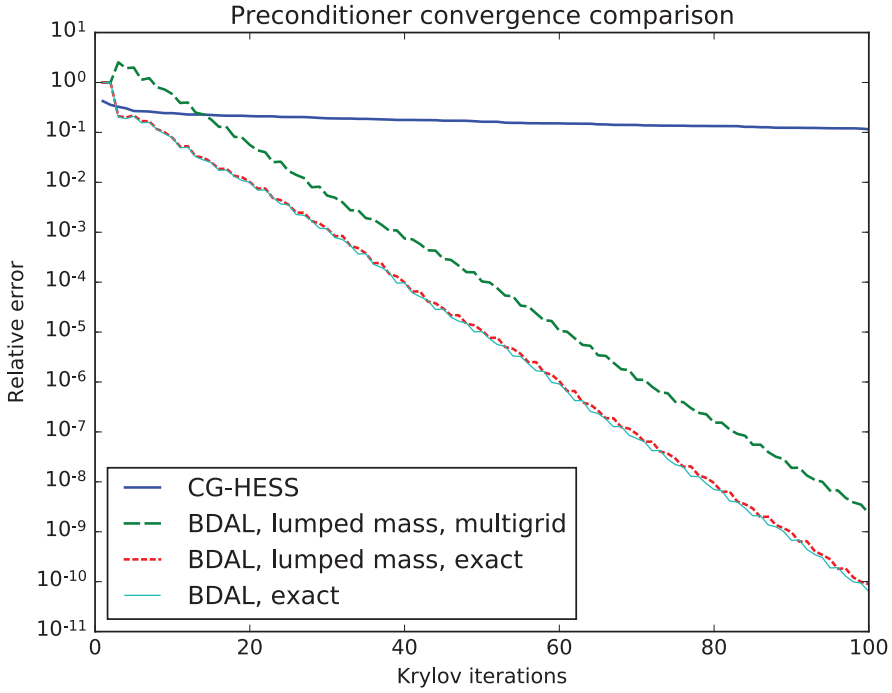

FIG. 2. *Relative error in the parameter,* $\|\mathbf{q} - \mathbf{q}_k\| / \|\mathbf{q}\|$*, for the high data Poisson source inversion problem, as a function of the number of Krylov iterations. The observation locations, regularization parameter, and mesh size are the same as in Figure* 1 *(*$n_{obs} = 2000$*,* $\alpha = 10^{-8}$*,* $h = \sqrt{2} \cdot 10^{-2}$*).*

In Figure 3, we see that the reconstruction using the reduced Hessian starts off smooth, then slowly includes information from successively higher frequency parameter modes as the CG iterations progress. In contrast, our preconditioner applied to the KKT system reconstructs low and high frequency information simultaneously.

6.2. Mesh scalability. To test mesh scalability, we solve the Poisson source inversion problem on a sequence of progressively finer meshes using MINRES with our block diagonal augmented Lagrangian preconditioner. The same regularization parameter, $\alpha = 10^{-8}$, and observation locations, $\{x_k\}_{k=1}^{2000}$, are used for all meshes. The numbers of iterations k required to achieve a relative error of $\|\mathbf{q} - \mathbf{q}_k\| / \|\mathbf{q}\| < 10^{-5}$ are shown in Table 1. All meshes are uniform triangular meshes. The coarsest mesh has size $h = 5.7 \cdot 10^{-2}$ with $1,800$ triangles, and the finest mesh has $h = 5.7 \cdot 10^{-3}$ with $181,000$ triangles. To quantify the error, the exact solution $\mathbf{q}$ was computed for each mesh using a sparse factorization of the KKT matrix. All results are based on the lumped mass approximation for mass matrices within the preconditioner.

The results clearly demonstrate mesh independence. The number of MINRES iterations required remains essentially constant over a two orders of magnitude increase in problem size, differing by at most one iteration across all mesh sizes.

6.3. Regularization and data scalability. A data and regularization robustness study is shown in Figure 4. The number of MINRES iterations k required for the method to converge to an error $\|\mathbf{q} - \mathbf{q}_k\| / \|\mathbf{q}\| < 10^{-5}$ is plotted for values of the regularization parameter in the range $\alpha \in [10^{-10}, 1.0]$, and number of observations

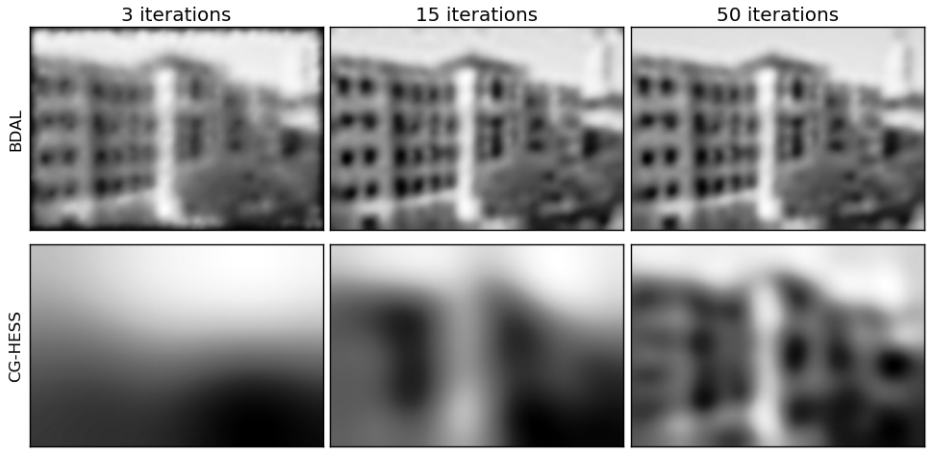

FIG. 3. *Visual comparison of the 3rd, 15th, and 50th Krylov iterates ($n_{obs} = 2000$, $\alpha = 10^{-8}$, $h = \sqrt{2} \cdot 10^{-2}$). Top row: Reconstruction using MINRES on the KKT system with our "BDAL, lumped mass, exact" preconditioner. Bottom row: Reconstruction using CG on the reduced Hessian with regularization preconditioning.*

TABLE 1

Mesh scalability study for our "BDAL, lumped mass, exact" preconditioner over a range of meshes. The table shows the number of MINRES iterations required to achieve parameter convergence to relative error 10^{-5}. The number of observations is $n_{obs} = 2000$, and the regularization parameter is $\alpha = 10^{-8}$. The observation locations x_k are the same for all mesh sizes.

h	# triangles	MINRES iterations
5.68e-02	1800	51
2.84e-02	7200	50
1.89e-02	16200	51
1.41e-02	29000	51
1.13e-02	45250	51
9.44e-03	65100	51
8.09e-03	88550	51
7.07e-03	116000	51
6.29e-03	146700	51
5.66e-03	181000	51

$n_{\text{obs}} \in \{150, 600, 2400, 9600\}$. The mesh size is fixed at $h = \sqrt{2} \cdot 10^{-2}$, and for each value of n_{obs}, the observation locations, x_k, are fixed as the regularization parameter varies.

The overall performance of the preconditioner is relatively steady over a broad range of values of α and n_{obs}. The performance of the method does decrease as the regularization parameter goes to zero for a fixed number of observations (upper left, Figure 4). However, the combination of small regularization parameter and small number of observations corresponds to the under-regularized regime, which we would not find ourselves in for an appropriately regularized problem. As the number of observations increases, the performance of the method improves in the small regularization regime while slightly worsening in the large regularization (over-regularized) regime, as suggested by our theory. This behavior is consistent with a data scalable method: one can take small values for the regularization parameter if that choice is supported by the data available in the problem.

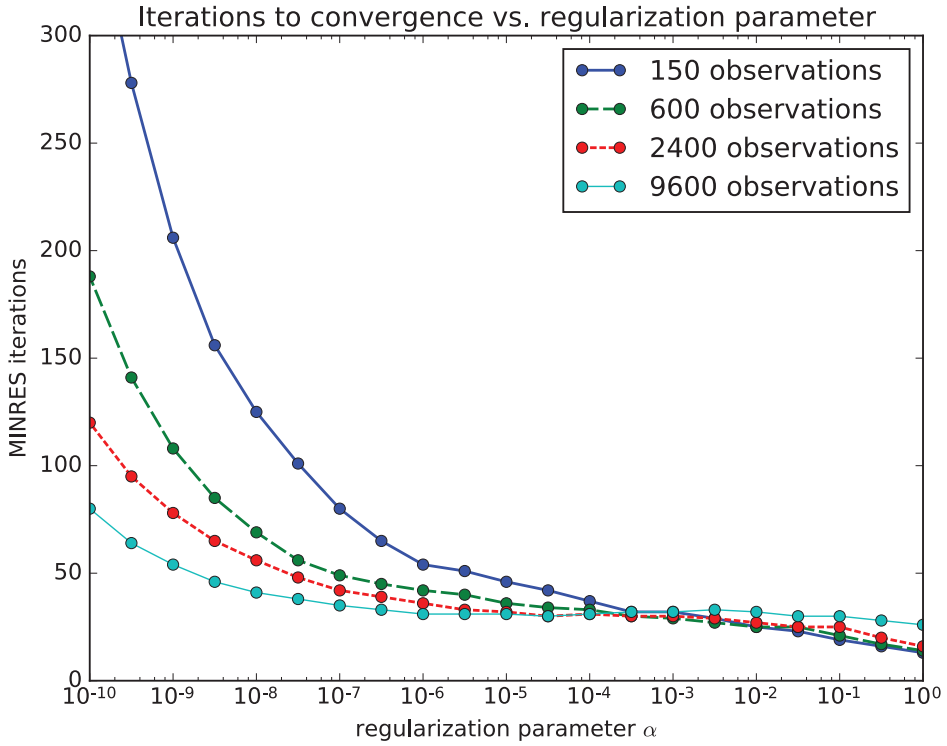

FIG. 4. *Regularization and data scalability study for our "BDAL, lumped mass, exact" preconditioner. Plot shows the number of MINRES iterations k required to achieve relative error* $\|\mathbf{q} - \mathbf{q}_k\| / \|\mathbf{q}\| < 1e-5$.

7. Conclusion. Traditional methods for solving linear inverse problems either scale poorly with increasing data and decreasing regularization, or are restricted to specific forms of regularization that may not be appropriate for the inverse problem at hand, or apply only to very specific problems. To overcome these limitations, we proposed a preconditioner based on a block diagonal approximation to the augmented Lagrangian KKT operator. We proved bounds on the condition number of the preconditioned system in an abstract setting, specialized the analysis to the case of source inversion problems with spectral filtering regularization, and tested the preconditioner numerically on a Poisson source inversion problem with highly informative data and small regularization parameter. Our analysis and numerical results indicate that the preconditioner is mesh and data scalable when the regularization does not over-penalize highly informed parameter modes and does not under-penalize uninformed modes.

Acknowledgments. We thank James Martin and Toby Isaac for helpful discussions, and James Martin for editing suggestions on an early draft of this paper. We thank the anonymous reviewers for their helpful comments. We would in particular like to thank one of the reviewers for bringing the improved Brezzi theory bound in [42] to our attention, as this allowed the constant in Theorem 7 to be decreased from $2 + 2\sqrt{2}$ to 3.

REFERENCES

[1] S. S. Adavani and G. Biros, *Multigrid algorithms for inverse problems with linear parabolic PDE constraints*, SIAM J. Sci. Comput., 31 (2008), pp. 369–397, https://doi.org/10.1137/070687426.

[2] S. S. Adavani and G. Biros, *Fast algorithms for source identification problems with elliptic PDE constraints*, SIAM J. Imaging Sci., 3 (2010), pp. 791–808, https://doi.org/10.1137/080738064.

[3] V. Akçelik, G. Biros, A. Drăgănescu, O. Ghattas, J. Hill, and B. van Bloeman Waanders, *Dynamic data-driven inversion for terascale simulations: Real-time identification of airborne contaminants*, in Proceedings of the 2005 ACM/IEEE Conference on Supercomputing, Seattle, IEEE, Washington, DC, 2005.

[4] E. Arian and A. Iollo, *Analytic Hessian derivation for the quasi-one-dimensional Euler equations*, J. Comput. Phys., 228 (2009), pp. 476–490.

[5] E. Arian and S. Ta'asan, *Analysis of the Hessian for aerodynamic optimization: Inviscid flow*, Comput. & Fluids, 28 (1999), pp. 853–877.

[6] U. M. Ascher and E. Haber, *A multigrid method for distributed parameter estimation problems*, Electron. Trans. Numer. Anal., 15 (2003), pp. 1–17, Tenth Copper Mountain Conference on Multigrid Methods (Copper Mountain, CO, 2001).

[7] O. Axelsson and J. Karatson, *Mesh independent superlinear PCG rates via compact-equivalent operators*, SIAM J. Numer. Anal., 45 (2007), pp. 1495–1516, https://doi.org/10.1137/06066391X.

[8] R. E. Bank, B. Welfert, and H. Yserentant, *A class of iterative methods for solving saddle point problems*, Numer. Math., 56 (1990), pp. 645–666.

[9] A. Barker, T. Rees, and M. Stoll, *A fast solver for an H_1 regularized PDE-constrained optimization problem*, Commun. Comput. Phys., 19 (2016), pp. 143–167.

[10] W. N. Bell, L. N. Olson, and J. Schroder, *PyAMG: Algebraic multigrid solvers in Python*, Version 2.1, 2013, http://www.pyamg.org.

[11] M. Benzi, G. H. Golub, and J. Liesen, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137, https://doi.org/10.1017/S0962492904000212.

[12] M. Benzi, E. Haber, and L. Taralli, *A preconditioning technique for a class of PDE-constrained optimization problems*, Adv. Comput. Math., 35 (2011), pp. 149–173.

[13] G. Biros and O. Ghattas, *Parallel Lagrange–Newton–Krylov–Schur methods for PDE-constrained optimization. Part* I: *The Krylov–Schur solver*, SIAM J. Sci. Comput., 27 (2005), pp. 687–713, https://doi.org/10.1137/S106482750241565X.

[14] G. Biros and O. Ghattas, *Parallel Lagrange–Newton–Krylov–Schur methods for PDE-constrained optimization. Part* II: *The Lagrange–Newton solver and its application to optimal control of steady viscous flows*, SIAM J. Sci. Comput., 27 (2005), pp. 714–739, https://doi.org/10.1137/S1064827502415661.

[15] A. Borzi, *Multigrid methods for parabolic distributed optimal control problems*, J. Comput. Appl. Math., 157 (2003), pp. 365–382.

[16] A. Borzi and R. Griesse, *Experiences with a space-time multigrid method for the optimal control of a chemical turbulence model*, Internat. J. Numer. Methods Fluids, 47 (2005), pp. 879–885.

[17] A. Borzi and V. Schulz, *Multigrid methods for PDE optimization*, SIAM Rev., 51 (2009), pp. 361–395, https://doi.org/10.1137/060671590.

[18] J. H. Bramble and J. E. Pasciak, *A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems*, Math. Comp., 50 (1988), pp. 1–17, https://doi.org/10.1090/S0025-5718-1988-0917816-8.

[19] F. Brezzi and M. Fortin, *Mixed and Hybrid Finite Element Methods*, Springer Ser. Comput. Math. 15, Springer-Verlag, New York, 1991.

[20] T. Bui-Thanh and O. Ghattas, *Analysis of the Hessian for inverse scattering problems. Part* I: *Inverse shape scattering of acoustic waves*, Inverse Problems, 28 (2012), 055001, https://doi.org/10.1088/0266-5611/28/5/055001.

[21] T. Bui-Thanh and O. Ghattas, *Analysis of the Hessian for inverse scattering problems. Part* II: *Inverse medium scattering of acoustic waves*, Inverse Problems, 28 (2012), 055002, https://doi.org/10.1088/0266-5611/28/5/055002.

[22] T. Bui-Thanh and O. Ghattas, *Analysis of the Hessian for inverse scattering problems. Part* III: *Inverse medium scattering of electromagnetic waves*, Inverse Probl. Imaging, 7 (2013), pp. 1139–1155.

[23] Y. Choi, *Simultaneous Analysis and Design in PDE-Constrained Optimization*, Ph.D. thesis, Stanford University, Stanford, CA, 2012.

[24] L. Demanet, P.-D. Létourneau, N. Boumal, H. Calandra, J. Chiu, and S. Snelson, *Matrix probing: A randomized preconditioner for the wave-equation Hessian*, Appl. Comput. Harmon. Anal., 32 (2012), pp. 155–168.

[25] L. Demkowicz, *Babuška ⇔ Brezzi?*, Tech. Report 06-08, Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX, 2006.

[26] A. Drăgănescu and T. Dupont, *Optimal order multilevel preconditioners for regularized ill-posed problems*, Math. Comp., 77 (2008), pp. 2001–2038.

[27] A. Drăgănescu and A. M. Soane, *Multigrid solution of a distributed optimal control problem constrained by the Stokes equations*, Appl. Math. Comput., 219 (2013), pp. 5622–5634.

[28] H. C. Elman, D. J. Silvester, and A. J. Wathen, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, Oxford University Press, Oxford, UK, 2014.

[29] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*, Math. Appl. 375, Kluwer Academic Publishers Group, Dordrecht, 1996.

[30] M. Ferronato, *Preconditioning for sparse linear systems at the dawn of the 21st century: History, current developments, and future perspectives*, IRSN Appl. Math., 2012 (2012), 127647.

[31] P. H. Flath, *Hessian-Based Response Surface Approximations for Uncertainty Quantification in Large-Scale Statistical Inverse Problems, with Applications to Groundwater Flow*, Ph.D. thesis, The University of Texas at Austin, Austin, TX, 2013.

[32] P. H. Flath, L. C. Wilcox, V. Akçelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas, *Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations*, SIAM J. Sci. Comput., 33 (2011), pp. 407–432, https://doi.org/10.1137/090780717.

[33] Z. Fortuna, *Some convergence properties of the conjugate gradient method in Hilbert space*, SIAM J. Numer. Anal., 16 (1979), pp. 380–394, https://doi.org/10.1137/0716031.

[34] A. Gholami, A. Mang, and G. Biros, *An inverse problem formulation for parameter estimation of a reaction–diffusion model of low grade gliomas*, J. Math. Biol., 72 (2016), pp. 409–433.

[35] G. H. Golub, C. Greif, and J. M. Varah, *An algebraic analysis of a block diagonal preconditioner for saddle point systems*, SIAM J. Matrix Anal. Appl., 27 (2006), pp. 779–792, https://doi.org/10.1137/04060679X.

[36] M. D. Gunzburger, *Perspectives in Flow Control and Optimization*, SIAM, Philadelphia, 2003, https://doi.org/10.1137/1.9780898718720.

[37] E. Haber and U. Ascher, *Preconditioned all-at-once methods for large, sparse parameter estimation problems*, Inverse Problems, 17 (2001), pp. 1847–1864.

[38] F. J. Herrmann, P. Moghaddam, and C. C. Stolk, *Sparsity- and continuity-promoting seismic image recovery with curvelet frames*, Appl. Comput. Harmon. Anal., 24 (2008), pp. 150–173.

[39] R. Herzog and E. Sachs, *Preconditioned conjugate gradient method for optimal control problems with control and state constraints*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 2291–2317, https://doi.org/10.1137/090779127.

[40] R. Herzog and E. Sachs, *Superlinear convergence of Krylov subspace methods for self-adjoint problems in Hilbert space*, SIAM J. Numer. Anal., 53 (2015), pp. 1304–1324, https://doi.org/10.1137/140973050.

[41] K. L. Ho and L. Ying, *Hierarchical interpolative factorization for elliptic operators: Differential equations*, Comm. Pure Appl. Math., 69 (2015), pp. 1415–1451.

[42] W. Krendl, V. Simoncini, and W. Zulehner, *Stability estimates and structural spectral properties of saddle point problems*, Numer. Math., 124 (2013), pp. 183–213.

[43] A. Logg, K.-A. Mardal, and G. Wells, *Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book*, Lecture Notes in Comput. Sci. Engrg. 84, Springer Science & Business Media, Springer-Verlag, Berlin, Heidelberg, 2012.

[44] K.-A. Mardal, B. Nielsen, and M. Nordaas, *Robust preconditioners for PDE-constrained optimization with limited observations*, BIT, 57 (2017), pp. 405–431.

[45] K.-A. Mardal and R. Winther, *Preconditioning discretizations of systems of partial differential equations*, Numer. Linear Algebra Appl., 18 (2010), pp. 1–40.

[46] M. F. Murphy, G. H. Golub, and A. J. Wathen, *A note on preconditioning for indefinite linear systems*, SIAM J. Sci. Comput., 22 (2000), pp. 1969–1972, https://doi.org/10.1137/S1064827599355153.

[47] B. Nielsen and K.-A. Mardal, *Efficient preconditioners for optimality systems arising in connection with inverse problems*, SIAM J. Control Optim., 48 (2010), pp. 5143–5177, https://doi.org/10.1137/080719960.

[48] B. Nielsen and K.-A. Mardal, *Analysis of the minimum residual method applied to ill posed optimality systems*, SIAM J. Sci. Comput., 35 (2012), pp. A785–A814, https://doi.org/10.1137/120871547.

[49] J. Nitsche, *Über ein Variationsprinzip zur Lösung Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind*, Abh. Math. Sem. Univ. Hamburg, 36 (1971), pp. 9–15.

[50] L. Olson, J. Schroder, and R. Tuminaro, *A general interpolation strategy for algebraic multigrid using energy minimization*, SIAM J. Sci. Comput., 33 (2011), pp. 966–991, https://doi.org/10.1137/100803031.

[51] C. C. Paige and M. A. Saunders, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629, https://doi.org/10.1137/0712047.

[52] J. W. Pearson, M. Stoll, and A. Wathen, *Regularization-robust preconditioners for time-dependent PDE-constrained optimization problems*, SIAM J. Matrix Anal. Appl., 33 (2012), pp. 1126–1152, https://doi.org/10.1137/110847949.

[53] J. W. Pearson, M. Stoll, and A. J. Wathen, *Preconditioners for state constrained optimal control problems with Moreau-Yosida penalty function*, Numer. Linear Algebra Appl., 21 (2014), pp. 81–97.

[54] J. W. Pearson and A. Wathen, *A new approximation of the Schur complement in preconditioners for PDE-constrained optimization*, Numer. Linear Algebra Appl., 19 (2012), pp. 816–829.

[55] T. Rees, H. S. Dollar, and A. J. Wathen, *Optimal solvers for PDE-constrained optimization*, SIAM J. Sci. Comput., 32 (2010), pp. 271–298, https://doi.org/10.1137/080727154.

[56] T. Rees, M. Stoll, and A. Wathen, *All-at-once preconditioning in PDE-constrained optimization*, Kybernetika, 46 (2010), pp. 341–360.

[57] T. Rees and A. Wathen, *Preconditioning iterative methods for the optimal control of the Stokes equations*, SIAM J. Sci. Comput., 33 (2011), pp. 2903–2926, https://doi.org/10.1137/100798491.

[58] A. Schiela and S. Ulbrich, *Operator preconditioning for a class of inequality constrained optimal control problems*, SIAM J. Optim., 24 (2014), pp. 435–466, https://doi.org/10.1137/120877532.

[59] J. Schöberl and W. Zulehner, *Symmetric indefinite preconditioners for saddle point problems with applications to PDE-constrained optimization problems*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 752–773, https://doi.org/10.1137/060660977.

[60] M. Stoll and A. Wathen, *Combination preconditioning and the Bramble–Pasciak$^+$ preconditioner*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 582–608, https://doi.org/10.1137/070688961.

[61] M. Stoll and A. Wathen, *All-at-Once Solution of Time-Dependent PDE-Constrained Optimization Problems*, Tech. report, 2010.

[62] S. Takacs and W. Zulehner, *Convergence analysis of multigrid methods with collective point smoothers for optimal control problems*, Comput. Vis. Sci., 14 (2011), pp. 131–141.

[63] P. Vaněk, J. Mandel, and M. Brezina, *Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems*, Computing, 56 (1996), pp. 179–196.

[64] C. R. Vogel, *Computational Methods for Inverse Problems*, Frontiers Appl. Math. 23, SIAM, Philadelphia, 2002, https://doi.org/10.1137/1.9780898717570.

[65] A. J. Wathen, B. Fischer, and D. J. Silvester, *The convergence rate of the minimal residual method for the Stokes system*, Numer. Math., 71 (1995), pp. 121–134.

[66] J. Xu and L. Zikatanov, *Some Observations on Babuška and Brezzi Theories*, Tech. Report AM222, Penn State University, University Park, PA, 2000, http://www.math.psu.edu/ccma/reports.html.

[67] W. Zulehner, *Nonstandard norms and robust estimates for saddle point problems*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 536–560, https://doi.org/10.1137/100814767.