

EFFICIENT SOLUTION OF PARAMETER IDENTIFICATION PROBLEMS WITH H^1 REGULARIZATION*

JAN BLECHTA[†] AND OLIVER G. ERNST[‡]

Abstract. We consider the identification of spatially distributed parameters under H^1 regularization. Solving the associated minimization problem by Gauss–Newton iteration results in linearized problems to be solved in each step that can be cast as boundary value problems involving a low-rank modification of the Laplacian. Using an algebraic multigrid as a fast Laplace solver, the Sherman–Morrison–Woodbury formula can be employed to construct a preconditioner for these linear problems which exhibits excellent scaling w.r.t. the relevant problem parameters. We first develop this approach in the functional setting, thus obtaining a consistent methodology for selecting boundary conditions that arise from the H^1 regularization. We then construct a method for solving the discrete linear systems based on combining any fast Poisson solver with the Woodbury formula. The efficacy of this method is then demonstrated with scaling experiments. These are carried out for a common nonlinear parameter identification problem arising in electrical resistivity tomography.

Key words. inverse problem, parameter identification, H^1 regularization, preconditioning, electrical resistivity tomography

MSC codes. 65F08, 65F22, 65N21

DOI. 10.1137/22M1520591



See reproducibility of computational results at end of the article.

1. Introduction. The problem of reconstructing a distributed parameter by the standard output least squares approach leads, after discretization, to a system of (nonlinear) algebraic equations which is typically solved using Newton-type methods, requiring the solution of a linearized problem at each step. For Gauss–Newton iteration, this linearized problem involves the Jacobian of the parameter-to-observation map, resulting in a discrete least squares problem. In the common setting where a high-dimensional unknown parameter is to be reconstructed from a small number of observations, this least squares problem is typically rank-deficient. The underdetermined nature combined with the ill-posedness of the underlying continuous inverse problem make it necessary to regularize the least squares formulation by adding a penalty term, usually involving norms of spatial derivatives of the unknown. In the high-dimensional setting, the linearized problems are also solved iteratively, usually by some variant of Krylov subspace projection methods adapted to least squares problems. The ill-posed nature of the underlying inverse problem as well as the spectral distribution of the regularization operator combine to make the preconditioning of the least squares iteration highly challenging, and many strategies have been proposed. For parameter identification problems arising from PDEs, the Jacobian is typically a compact operator (cf. [56], [23, Theorem 4.21]), and it is known that Krylov subspace

*Submitted to the journal’s Numerical Algorithms for Scientific Computing section September 7, 2022; accepted for publication (in revised form) September 7, 2023; published electronically April 2, 2024.

<https://doi.org/10.1137/22M1520591>

Funding: This work was supported by the European Union (EU) European Social Fund (ESF) and the Free State of Saxony (project GEOSax, grant 100310486).

[†]Faculty of Mathematics and Physics, Charles University, 186 75 Prague, Czech Republic (blechta@karlin.mff.cuni.cz).

[‡]Department of Mathematics, Chemnitz University of Technology, Chemnitz 09126, Germany (oernst@math.tu-chemnitz.de).

methods such as LSQR converge very slowly for the discretized problem. Tikhonov regularization, by adding an L^2 -norm penalty term, changes the setting to a compact perturbation of the identity, which in turn is fundamentally favorable for fast convergence of Krylov subspace iterations [55, 39, 24, 9]. The spectral properties of the iteration matrix, however, become much more challenging when regularization involving smoothing terms is employed, leading to large Gauss–Newton inner iteration counts.

In this work, we address the efficient solution of the nonlinear least squares problem arising from distributed parameter estimation problems regularized by the H^1 norm, sometimes referred to as *smoothness regularization*. The function to be reconstructed from observations is represented as piecewise constant w.r.t. a given triangulation of the domain and, following an idea proposed by Schwarzbach and Haber [47], define its gradient in the regularization term by duality. We derive this formulation in the continuous setting by expressing the Gauss–Newton updates as the solution of boundary value problems, which are then discretized using a mixed finite element method. After discretization, the linear systems arising in each Gauss–Newton step have a saddle-point structure and are solved by preconditioned MINRES [44, 26] iteration. Our preconditioning strategy is based on a known spectrally equivalent preconditioner for the Laplacian, which enters the problem by way of the regularization term, combined with an application of the Sherman–Morrison–Woodbury formula to account for the low-rank perturbation arising from the data misfit term. As a result, we obtain a mesh-independent solver for the nonlinear least squares problem which is also robust w.r.t. a large range of regularization parameters.

Background references for PDE-based nonlinear parameter identification problems are [53] and, with a focus on geoelectromagnetic exploration problems, [31]. While less common for parameter identification, H^1 regularization is also used in optimal control problems involving control and state constraints [2, 19]. The extensive literature on Krylov projection methods for least squares problems is summarized in [6, 7]. Finite-precision effects are analyzed in [8], and more recent developments include extensions of these methods to the class of *symmetric quasi-definite* problems in [43], an analysis of LSQR for compact operators in Hilbert space in [17], and an extensive numerical comparison of the state of the art in Krylov methods and preconditioners for sparse linear least squares problems in [30]. A popular construction principle for preconditioning matrices of saddle-point structure is based on the observation that suitable block-triangular and block-diagonal preconditioners result in a system matrix with a minimal polynomial of degree 2 or 3 [36, 41, 34], for which Krylov subspace projection will return the exact solution in the same number of steps. A more comprehensive review of operator preconditioning techniques with special emphasis on mixed discretizations and saddle-point problems can be found in [38]. A large class of preconditioning techniques for general least squares problems are based on incomplete factorizations [4, 13, 1, 49, 48, 18] as well as inner-outer iteration [40]. Closer to the approach proposed in this work, the idea of using a suitable Laplace preconditioner for variational inverse problems involving a compact operator, when the Laplacian is used as a regularization for the normal equations, is explored in [32] (cf. also [54]). A refinement of this approach is described in [35] and a further variant proposed in [14]. Image restoration problems are also close to our setting in that the origin of the least squares problem is a continuous inverse problem and regularization is a necessity. However, given that there the unknown is an unblurred image, the basic problem is typically not one of least squares since there are typically as many measurements (pixel values) as unknowns. Once regularization by penalty terms is added,

however, the formulation is typically that of a minimization problem [5, 20], and a successful approach here is the class of *hybrid projection methods* [22, 21]. Particularly in connection with statistical inverse problems, using preconditioners derived from covariance matrices has recently drawn increased attention [16, 15]. Finally, methods employing the Sherman–Morrison–Woodbury formula for constructing preconditioners have been considered by Yin [57] and Benzi and Faccio [3]. In [57], a recursive factorization technique is employed to apply a preconditioner for Tikhonov-regularized least squares problems with a Euclidean penalty term. Benzi and Faccio [3] discuss preconditioning strategies for linear systems with a matrix of the form $\mathbf{A} + \gamma \mathbf{U} \mathbf{U}^\top$ with a tall-and-skinny matrix \mathbf{U} , a class of problems which includes that addressed in this paper. It is reported there that efforts to construct a preconditioner using the Woodbury formula with an approximation of the factor \mathbf{A}^{-1} occurring therein proved unsuccessful [3, page 4]. By contrast, we will demonstrate this approach for our problem, in which \mathbf{A} is a discrete Laplacian, to be quite effective. It is fair to remark that [3] considered a broader class of problems, including the case of singular \mathbf{A} .

The structure of the paper is as follows. Section 2 introduces the problem setting of H^1 -regularized parameter estimation, derives the operator equations to be solved in each Gauss–Newton step, and briefly presents its mixed discretization based on an inf-sup stable mixed discretization for the Poisson equation. Section 3 presents three variants of the solution algorithm: (i) a direct approach based on the use of the Woodbury formula and a factorization of the Laplacian; (ii) a MINRES iteration preconditioned by the Woodbury formula and a suitable Laplace preconditioner; (iii) a simplified variant of (ii) that omits the low-rank modification in the Woodbury formula. Section 4 contains an extensive numerical illustration, in which our solution approach is applied to an electrical resistivity tomography (ERT) problem from geophysical exploration in two and three space dimensions. Realistic measurement setups are considered involving up to thousands of observational data points. The algorithm is seen to perform efficiently and robustly across a variety of settings. Finally, in section 5, we summarize our findings and indicate further aspects to be investigated in subsequent research.

2. Problem formulation. We consider the output-least-squares formulation for estimating a distributed parameter $m \in L^2(\Omega)$ defined on a bounded domain $\Omega \subset \mathbb{R}^d$ and a (typically nonlinear) parameter-to-observation map $g: L^2(\Omega) \rightarrow \mathbb{R}^M$ assigning to each parameter m a set of M observations, from which m is to be reconstructed by minimizing the misfit $\sum_{i=1}^M |g^i(m) - g_{\text{obs}}^i|^2$ w.r.t. a vector $\mathbf{g}_{\text{obs}} = \{g_{\text{obs}}^i\}_{i=1}^M$ of observations. For example, when $\exp m > 0$ is the diffusion coefficient of an elliptic forward problem, this minimization is an ill-posed and severely underdetermined problem which can be addressed by adding a regularizing penalty term to the data misfit functional. In this work, we develop efficient computational methods for determining m when the regularization term is the H^1 norm, a common device for promoting smoothness of the reconstructed function. This leads to the task of minimizing the objective function

$$(2.1) \quad \sum_{i=1}^M |g^i(m) - g_{\text{obs}}^i|^2 + \beta \int_{\Omega} |\nabla(m - m_{\text{ref}})|^2,$$

where $\beta > 0$ is a regularization parameter and m_{ref} denotes a reference or background value for the unknown m . The regularization thus penalizes the gradient of the deviation from the known background value m_{ref} , a common setting in, e.g., geophysical inverse problems. Consequently, in order for the regularization term to make sense,

this formulation, which we shall weaken in the following, would require $m - m_{\text{ref}}$ to lie in the smaller space $H^1(\Omega) \subset L^2(\Omega)$.

To develop a Gauss–Newton iteration for the minimization of (2.1), we will reformulate the first-order optimality condition as a set of normal equations in the function space setting. The gradient acting on m then becomes a Laplacian, for which spectrally equivalent preconditioners are available, allowing efficient iterative solution of the linearized equation in each Gauss–Newton step. In addition, we recast the optimality equations in a mixed formulation, which is well-defined also for $m \in L^2(\Omega)$, where the gradient in the regularization term is defined by duality.

2.1. Assumptions and notation. We assume that $\Omega \subset \mathbb{R}^d$ is a bounded Lipschitz domain with boundary partitioned into $\partial\Omega = \overline{\Gamma_D} \cup \overline{\Gamma_N}$, with Γ_N, Γ_D open and disjoint. For simplicity, we assume $|\Gamma_D| > 0$ to exclude the pure Neumann problem. We denote by $L^2(\Omega)$ the space of measurable functions $f: \Omega \rightarrow \mathbb{R}$ with finite norm

$$\|f\|_2 := \left(\int_{\Omega} |f|^2 \right)^{\frac{1}{2}}$$

and by $H^1(\Omega)$ the Sobolev space of such functions $f \in L^2(\Omega)$ with finite norm

$$\|f\|_{1,2} := \left(\int_{\Omega} |\nabla f|^2 + |f|^2 \right)^{\frac{1}{2}}.$$

The subspace of functions vanishing on Γ_D is denoted by $H^1_{\Gamma_D}(\Omega) \subset H^1(\Omega)$, and

$$\|\nabla f\|_2 := \left(\int_{\Omega} |\nabla f|^2 \right)^{\frac{1}{2}}$$

is a norm on $H^1_{\Gamma_D}(\Omega)$ that is equivalent to $\|\cdot\|_{1,2}$. The space $H(\text{div}; \Omega)$ consists of all vector fields $\vec{f}: \Omega \rightarrow \mathbb{R}^d$ such that $\vec{f} \in L^2(\Omega)$ and $\text{div} \vec{f} \in L^2(\Omega)$, and it is equipped with the norm $\|\vec{f}\|_{\text{div}} := \|\vec{f}\|_2 + \|\text{div} \vec{f}\|_2$. The subspace of $H(\text{div}; \Omega)$ consisting of vector fields with vanishing normal trace on Γ_N is denoted by $H_{\Gamma_N}(\text{div}; \Omega)$.

Next we assume that the parameter-to-observation map g is given and Gâteaux-differentiable with the derivative denoted by $J: L^2(\Omega) \rightarrow [L^2(\Omega)]^M$ so that

$$\langle J(m), \delta m \rangle = \left[\frac{d}{dt} g(m + t\delta m) \right]_{t=0}, \quad m, \delta m \in L^2(\Omega).$$

Therefore, the mapping $\delta m \mapsto \langle J(m), \delta m \rangle$ is assumed to be linear and bounded. The individual components of g and J are denoted by $g^i: L^2(\Omega) \rightarrow \mathbb{R}$ and $J^i: L^2(\Omega) \rightarrow L^2(\Omega)'$, respectively, so that

$$\langle J^i(m), \delta m \rangle = \left[\frac{d}{dt} g^i(m + t\delta m) \right]_{t=0}, \quad m, \delta m \in L^2(\Omega), \quad i = 1, 2, \dots, M.$$

2.2. Primal and mixed regularized least squares formulation. To simplify the following expressions, we rescale the regularized least squares functional (2.1) by $1/\beta > 0$ and obtain the objective function

$$(2.2) \quad \Phi_{\beta}(m) = \frac{1}{\beta} \sum_{i=1}^M |g^i(m) - g^i_{\text{obs}}|^2 + \int_{\Omega} |\nabla(m - m_{\text{ref}})|^2, \quad m - m_{\text{ref}} \in H^1_{\Gamma_D}(\Omega).$$

Besides requiring the deviation $m - m_{\text{ref}}$ to lie in the smoother space $H^1(\Omega)$, we impose an essential boundary condition on the portion Γ_D of the boundary of the domain Ω . This is a modeling decision which depends on the type of assumptions or a priori information available on the unknown parameter m ; in this case, m is assumed to coincide with the background value m_{ref} on Γ_D . As we will see below, this choice also implicitly imposes a natural boundary condition on Γ_N .

Taking the first variation (Gâteaux derivative) of (2.2) in a direction $\phi \in H_{\Gamma_D}^1(\Omega)$ and setting it to zero, we arrive at the first-order necessary optimality condition for minimizing (2.2):

$$(2.3) \quad \begin{aligned} &\text{Find } m \in H^1(\Omega) \text{ such that } m - m_{\text{ref}} \in H_{\Gamma_D}^1(\Omega) \text{ and} \\ &\frac{1}{\beta} \sum_{i=1}^M (g^i(m) - g_{\text{obs}}^i) \langle J^i(m), \phi \rangle + \int_{\Omega} \nabla(m - m_{\text{ref}}) \cdot \nabla \phi = 0 \\ &\hspace{15em} \text{for all } \phi \in H_{\Gamma_D}^1(\Omega). \end{aligned}$$

Assuming sufficient regularity, (2.3) can be interpreted as a weak formulation of the boundary value problem

$$(2.4a) \quad \frac{1}{\beta} \sum_{i=1}^M (g^i(m) - g_{\text{obs}}^i) (J^i(m))' - \Delta(m - m_{\text{ref}}) = 0 \quad \text{in } \Omega,$$

$$(2.4b) \quad m - m_{\text{ref}} = 0 \quad \text{on } \Gamma_D,$$

$$(2.4c) \quad \frac{\partial}{\partial n}(m - m_{\text{ref}}) = 0 \quad \text{on } \Gamma_N.$$

Here $(J^i(m))' \in L^2(\Omega)$ denotes the Riesz representer of $J^i(m) \in L^2(\Omega)'$, i.e.,

$$(2.5) \quad \langle J^i(m), \phi \rangle = \int_{\Omega} (J^i(m))' \phi \quad \text{for all } \phi \in L^2(\Omega).$$

Recall that $g^i(m) - g_{\text{obs}}^i$ is a number for any fixed m . Hence, for a fixed m , the first term in (2.4a) is an $L^2(\Omega)$ -function in the present setting.

Gauss–Newton linearization of (2.3) is obtained by applying Newton’s method to (2.3) and neglecting the Hessian of g , which is given by

$$\langle H(m)\phi, \delta m \rangle = \left[\frac{d}{dt} \langle J(m + t\delta m), \phi \rangle \right]_{t=0}, \quad m \in H^1(\Omega), \quad \phi, \delta m \in H_{\Gamma_D}^1(\Omega).$$

Given an initial value m , associated model-generated responses $g_m^i = g^i(m)$, and the derivatives $J_m^i = J^i(m)$, $i = 1, 2, \dots, M$, one step of Gauss–Newton iteration determines an update $m + \delta m$ by solving the following problem:

$$(2.6) \quad \begin{aligned} &\text{Find } \delta m \in H^1(\Omega) \text{ such that } \delta m + m - m_{\text{ref}} \in H_{\Gamma_D}^1(\Omega) \text{ and} \\ &\frac{1}{\beta} \sum_{i=1}^M \langle J_m^i, \delta m \rangle \langle J_m^i, \phi \rangle + \int_{\Omega} \nabla \delta m \cdot \nabla \phi \\ &= -\frac{1}{\beta} \sum_{i=1}^M (g_m^i - g_{\text{obs}}^i) \langle J_m^i, \phi \rangle - \int_{\Omega} \nabla(m - m_{\text{ref}}) \cdot \nabla \phi \\ &\hspace{15em} \text{for all } \phi \in H_{\Gamma_D}^1(\Omega). \end{aligned}$$

The variational equation (2.6) in turn is a weak formulation of the boundary value problem for the Gauss–Newton correction δm :

$$\begin{aligned}
 (2.7a) \quad & \left[\frac{1}{\beta} \sum_{i=1}^M J_m^{i\prime} \langle J_m^i, \cdot \rangle - \Delta \right] \delta m = -\frac{1}{\beta} \sum_{i=1}^M J_m^{i\prime} (g_m^i - g_{\text{obs}}^i) + \Delta(m - m_{\text{ref}}) \quad \text{in } \Omega, \\
 (2.7b) \quad & \delta m = -(m - m_{\text{ref}}) \quad \text{on } \Gamma_D, \\
 (2.7c) \quad & \frac{\partial}{\partial n} \delta m = -\frac{\partial}{\partial n} (m - m_{\text{ref}}) \quad \text{on } \Gamma_N.
 \end{aligned}$$

In view of (2.5), the first operator in (2.7a) acting on δm can be expressed as

$$\left[\sum_{i=1}^M J_m^{i\prime} \langle J_m^i, \delta m \rangle \right] (x) = \int_{\Omega} \left[\sum_{i=1}^M J_m^{i\prime}(x) J_m^{i\prime}(y) \right] \delta m(y) \, dy,$$

i.e., as a finite-rank integral operator with kernel $\sum_{i=1}^M J_m^{i\prime}(x) J_m^{i\prime}(y) \in L^2(\Omega \times \Omega)$.

To weaken the regularity requirements on m , we next recast problem (2.6) in a mixed formulation by introducing the flux variable $\vec{\zeta} := \nabla(\delta m + m - m_{\text{ref}})$:

$$\begin{aligned}
 & \text{Find } (\vec{\zeta}, \delta m) \in H_{\Gamma_N}(\text{div}; \Omega) \times L^2(\Omega) \text{ such that} \\
 (2.8) \quad & \int_{\Omega} \vec{\zeta} \cdot \vec{\psi} + \int_{\Omega} \delta m \, \text{div } \vec{\psi} = - \int_{\Omega} (m - m_{\text{ref}}) \, \text{div } \vec{\psi}, \\
 & \frac{1}{\beta} \sum_{i=1}^M \langle J_m^i, \delta m \rangle \langle J_m^i, \phi \rangle - \int_{\Omega} \phi \, \text{div } \vec{\zeta} = -\frac{1}{\beta} \sum_{i=1}^M (g_m^i - g_{\text{obs}}^i) \langle J_m^i, \phi \rangle \\
 & \quad \text{for all } \vec{\psi} \in H_{\Gamma_N}(\text{div}; \Omega) \text{ and } \phi \in L^2(\Omega).
 \end{aligned}$$

Introducing the operators

$$(2.9a) \quad \langle Q\vec{\zeta}, \vec{\psi} \rangle := \int_{\Omega} \vec{\zeta} \cdot \vec{\psi}, \quad \vec{\zeta}, \vec{\psi} \in H_{\Gamma_N}(\text{div}; \Omega),$$

$$(2.9b) \quad \langle D\vec{\zeta}, \phi \rangle := \int_{\Omega} \phi \, \text{div } \vec{\zeta}, \quad \vec{\zeta} \in H_{\Gamma_N}(\text{div}; \Omega), \phi \in L^2(\Omega),$$

we can rewrite (2.8) in the block operator form

$$(2.10) \quad \begin{bmatrix} Q & D' \\ D & -\frac{1}{\beta} J_m' J_m \end{bmatrix} \begin{bmatrix} \vec{\zeta} \\ \delta m \end{bmatrix} = \begin{bmatrix} -D'(m - m_{\text{ref}}) \\ \frac{1}{\beta} J_m'(\mathbf{g}_m - \mathbf{g}_{\text{obs}}) \end{bmatrix},$$

where the occurrences of J_m and J_m' are expressed using duality as

$$\begin{aligned}
 \langle J_m' J_m \delta m, \phi \rangle &= \sum_{i=1}^M \langle J_m^i, \delta m \rangle \langle J_m^i, \phi \rangle, & \delta m, \phi \in L^2(\Omega), \\
 \langle J_m'(\mathbf{g}_m - \mathbf{g}_{\text{obs}}), \phi \rangle &= \sum_{i=1}^M (g_m^i - g_{\text{obs}}^i) \langle J_m^i, \phi \rangle, & \phi \in L^2(\Omega).
 \end{aligned}$$

In an analogous way, by defining the operator

$$\langle L\delta m, \phi \rangle := \int_{\Omega} \nabla \delta m \cdot \nabla \phi, \quad \delta m, \phi \in H_{\Gamma_D}^1(\Omega),$$

we can rewrite the primal formulation (2.6) as the operator equation

$$(2.11) \quad \left(L + \frac{1}{\beta} J'_m J_m \right) \delta m = -L(m - m_{\text{ref}}) - \frac{1}{\beta} J'_m (\mathbf{g}_m - \mathbf{g}_{\text{obs}}).$$

The primal formulation (2.11) can be seen as Schur complement reduction of the mixed formulation (2.10), in terms of which the Laplacian is represented as $L = DQ^{-1}D'$. Indeed, block elimination of $\vec{\zeta}$ in (2.10) gives

$$(2.12) \quad \left(DQ^{-1}D' + \frac{1}{\beta} J'_m J_m \right) \delta m = -DQ^{-1}D'(m - m_{\text{ref}}) - \frac{1}{\beta} J'_m (\mathbf{g}_m - \mathbf{g}_{\text{obs}}).$$

We note that Schwarzbach and Haber [47, section 3.2.1] also formulated the H^1 regularization using the mixed formulation (2.10). Their approach consisted of discretizing by lowest-order Raviart–Thomas elements and approximating Q^{-1} by a diagonal matrix in the Schur complement formulation (2.12). We will instead proceed by considering the mixed formulation (2.10) and design a solution strategy for this system.

2.3. Finite element discretization. The $H^1(\Omega)$ formulation (2.11) suggests an $H^1(\Omega)$ -conforming discretization for the parameter m , using, e.g., continuous Lagrange elements. Instead, to allow for parameters $m \in L^2(\Omega)$, we will employ a standard discretization of the mixed formulation (2.8). Let us assume in the following that Ω is polyhedral so that we can consider its simplicial partitions \mathcal{T}_h . Further, let finite element spaces $V_h \times Q_h \subset H_{\Gamma_N}(\text{div}; \Omega) \times L^2(\Omega)$ be chosen as

$$V_h \times Q_h := RT_k(\mathcal{T}_h, \Gamma_N) \times dP_k(\mathcal{T}_h)$$

for some order $k \in \mathbb{N}_0$, where RT_k and dP_k denote the finite element spaces of Raviart–Thomas and discontinuous Lagrange of order k counted such that $k = 0$ corresponds to the lowest-order case. This is an inf-sup stable discretization for the Poisson equation in mixed formulation, i.e., the operator given by (2.10) (or, equivalently, (2.8)) without the $J'_m J_m$ term; see [11].

Let $\{\vec{\psi}_i^h\}_{i=1}^K$ and $\{\phi_i^h\}_{i=1}^N$ denote bases of V_h and Q_h , respectively, so that

$$\text{span}\{\vec{\psi}_i^h\}_{i=1}^K = V_h, \quad K = \dim V_h,$$

$$\text{span}\{\phi_i^h\}_{i=1}^N = Q_h, \quad N = \dim Q_h.$$

Inserting the basis elements into (2.9) yields the matrices

$$(2.13a) \quad \mathbf{Q} \in \mathbb{R}^{K \times K}, \quad (\mathbf{Q})_{ij} := \langle Q \vec{\psi}_j^h, \vec{\psi}_i^h \rangle,$$

$$(2.13b) \quad \mathbf{D} \in \mathbb{R}^{N \times K}, \quad (\mathbf{D})_{ij} := \langle D \vec{\psi}_j^h, \phi_i^h \rangle.$$

Assuming that $m, m_{\text{ref}} \in Q_h$, these can then be expressed as

$$(2.14a) \quad \mathbf{m} \in \mathbb{R}^N, \quad m(x) = \sum_{j=1}^N (\mathbf{m})_j \phi_j^h(x), \quad x \in \Omega,$$

$$(2.14b) \quad \mathbf{m}_{\text{ref}} \in \mathbb{R}^N, \quad m_{\text{ref}}(x) = \sum_{j=1}^N (\mathbf{m}_{\text{ref}})_j \phi_j^h(x), \quad x \in \Omega.$$

We seek to determine $(\vec{\zeta}, \delta m) \in V_h \times Q_h$ so that

$$(2.14c) \quad \zeta \in \mathbb{R}^K, \quad \vec{\zeta}(x) = \sum_{j=1}^K (\zeta)_j \vec{\psi}_j^h(x), \quad x \in \Omega,$$

$$(2.14d) \quad \delta m \in \mathbb{R}^N, \quad \delta m(x) = \sum_{j=1}^N (\delta m)_j \phi_j^h(x), \quad x \in \Omega.$$

Naturally, g and J are restricted to Q_h , which gives rise to the vector and the matrix

$$(2.14e) \quad \mathbf{g}_m \in \mathbb{R}^M, \quad (\mathbf{g}_m)_i := g^i(m),$$

$$(2.14f) \quad \mathbf{J}_m \in \mathbb{R}^{M \times N}, \quad (\mathbf{J}_m)_{ij} := \langle J^i(m), \phi_j^h \rangle,$$

where $m \in Q_h$ on the right-hand sides is given by (2.14a). We thus arrive at the discrete counterpart of (2.10), the block linear system

$$(2.15) \quad \begin{bmatrix} \mathbf{Q} & \mathbf{D}^\top \\ \mathbf{D} & -\frac{1}{\beta} \mathbf{J}_m^\top \mathbf{J}_m \end{bmatrix} \begin{bmatrix} \zeta \\ \delta m \end{bmatrix} = \begin{bmatrix} -\mathbf{D}^\top (m - m_{\text{ref}}) \\ \frac{1}{\beta} \mathbf{J}_m^\top (\mathbf{g}_m - \mathbf{g}_{\text{obs}}) \end{bmatrix}.$$

3. Solution of the linear systems using the Woodbury formula. The linear system (2.15) to be solved for the Gauss–Newton updates is a low-rank perturbation of a Poisson problem in the mixed formulation. In this section, we employ the Woodbury matrix identity (see, e.g., [29, section 2.1.4]) to construct algorithms for efficiently solving this system. We first consider a direct solution approach which can benefit from reusing the factorization for the unperturbed problem. As a second approach, we propose two preconditioners for an iterative solution which can take advantage of any available efficient solution method for the unperturbed problem.

To this end, we note that any formulation and discretization for the Laplacian which has an efficient solution method can be used to proceed along the following lines. Recall that the operators in the infinite-dimensional linear systems (2.10), (2.11), and (2.12) are seen to be finite-rank perturbations of the Laplacian. But in the following, we will focus entirely on the discrete mixed formulation (2.15).

3.1. Direct solution. We introduce the following matrices (cf. (2.15)), which will be useful in constructing the solution schemes:

$$(3.1) \quad \mathbf{A} := \begin{bmatrix} \mathbf{Q} & \mathbf{D}^\top \\ \mathbf{D} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{(K+N) \times (K+N)}, \quad \mathbf{S} := \mathbf{D} \mathbf{Q}^{-1} \mathbf{D}^\top \in \mathbb{R}^{N \times N}, \\ \mathbf{P}_2 := \begin{bmatrix} \mathbf{0} & \mathbf{I}_N \end{bmatrix} \in \mathbb{R}^{N \times (K+N)},$$

where \mathbf{I}_N denotes the $N \times N$ identity. The solution of a linear system with coefficient matrix \mathbf{A} and right-hand side blocks \mathbf{y}_1 and \mathbf{y}_2 are related via the Schur complement \mathbf{S} as

$$(3.2) \quad \mathbf{P}_2 \mathbf{A}^{-1} \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix} = \mathbf{S}^{-1} (\mathbf{D} \mathbf{Q}^{-1} \mathbf{y}_1 - \mathbf{y}_2), \quad \mathbf{y}_1 \in \mathbb{R}^K, \mathbf{y}_2 \in \mathbb{R}^N.$$

The Schur complement matrix \mathbf{S} is, in general, dense, and hence linear systems with matrix \mathbf{S} are impractical to assemble and solve. On the other hand, (3.2) implies that

$$(3.3) \quad \mathbf{S}^{-1} \mathbf{y}_2 = -\mathbf{P}_2 \mathbf{A}^{-1} \begin{bmatrix} \mathbf{0} \\ \mathbf{y}_2 \end{bmatrix}, \quad \mathbf{y}_2 \in \mathbb{R}^N;$$

i.e., the solution of the dense system $\mathbf{S} \mathbf{x}_2 = -\mathbf{y}_2$ can be expressed as the solution of the sparse saddle-point system

$$\mathbf{A} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{y}_2 \end{bmatrix}.$$

This is the setting we encounter in the Gauss–Newton update step, where the second block of the solution of (2.15) is needed and the saddle-point matrix $\mathbf{A}_{\beta, \mathbf{m}}$ is a low-rank modification of \mathbf{A} in the second block:

$$(3.4) \quad \delta \mathbf{m} = \mathbf{P}_2 \mathbf{A}_{\beta, \mathbf{m}}^{-1} \begin{bmatrix} -\mathbf{D}^\top (\mathbf{m} - \mathbf{m}_{\text{ref}}) \\ \frac{1}{\beta} \mathbf{J}_m^\top (\mathbf{g}_m - \mathbf{g}_{\text{obs}}) \end{bmatrix}, \quad \mathbf{A}_{\beta, \mathbf{m}} := \mathbf{A} - \frac{1}{\beta} \begin{bmatrix} \mathbf{0} \\ \mathbf{J}_m^\top \end{bmatrix} \begin{bmatrix} \mathbf{0} & \mathbf{J}_m \end{bmatrix}.$$

Using the Woodbury formula, we may express $\mathbf{A}_{\beta, \mathbf{m}}^{-1}$ in terms of \mathbf{A}^{-1} as

$$\mathbf{A}_{\beta, \mathbf{m}}^{-1} = \mathbf{A}^{-1} + \frac{1}{\beta} \mathbf{A}^{-1} \begin{bmatrix} \mathbf{0} \\ \mathbf{J}_m^\top \end{bmatrix} \left(\mathbf{I}_M - \frac{1}{\beta} \begin{bmatrix} \mathbf{0} & \mathbf{J}_m \end{bmatrix} \mathbf{A}^{-1} \begin{bmatrix} \mathbf{0} \\ \mathbf{J}_m^\top \end{bmatrix} \right)^{-1} \begin{bmatrix} \mathbf{0} & \mathbf{J}_m \end{bmatrix} \mathbf{A}^{-1},$$

where \mathbf{I}_M denotes the $M \times M$ identity. Defining the matrix

$$(3.5) \quad \mathbf{H}_m := \mathbf{S}^{-1} \mathbf{J}_m^\top = -\mathbf{P}_2 \mathbf{A}^{-1} \begin{bmatrix} \mathbf{0} \\ \mathbf{J}_m^\top \end{bmatrix} \in \mathbb{R}^{N \times M}$$

and observing $\begin{bmatrix} \mathbf{0} & \mathbf{J}_m \end{bmatrix} = \mathbf{J}_m \mathbf{P}_2$, we arrive at the expression for the matrix whose action is required in the update (3.4):

$$(3.6) \quad \mathbf{P}_2 \mathbf{A}_{\beta, \mathbf{m}}^{-1} = \left(\mathbf{I}_N - \frac{1}{\beta} \mathbf{H}_m \left(\mathbf{I}_M + \frac{1}{\beta} \mathbf{J}_m \mathbf{H}_m \right)^{-1} \mathbf{J}_m \right) \mathbf{P}_2 \mathbf{A}^{-1}.$$

Combining (3.2) and (3.5), we obtain for the unperturbed problem

$$(3.7) \quad \mathbf{P}_2 \mathbf{A}^{-1} \begin{bmatrix} -\mathbf{D}^\top (\mathbf{m} - \mathbf{m}_{\text{ref}}) \\ \frac{1}{\beta} \mathbf{J}_m^\top (\mathbf{g}_m - \mathbf{g}_{\text{obs}}) \end{bmatrix} = -(\mathbf{m} - \mathbf{m}_{\text{ref}}) - \frac{1}{\beta} \mathbf{H}_m (\mathbf{g}_m - \mathbf{g}_{\text{obs}}).$$

Equations (3.4), (3.6), and (3.7) now yield an expression for the update vector as

$$\delta \mathbf{m} = -(\mathbf{m} - \mathbf{m}_{\text{ref}}) + \frac{1}{\beta} \mathbf{H}_m \left(\mathbf{I}_M + \frac{1}{\beta} \mathbf{J}_m \mathbf{H}_m \right)^{-1} \left(\mathbf{J}_m (\mathbf{m} - \mathbf{m}_{\text{ref}}) - (\mathbf{g}_m - \mathbf{g}_{\text{obs}}) \right).$$

The computations for constructing this vector within a complete Gauss–Newton minimization are summarized in Algorithm 1. It requires a single \mathbf{LDL}^\top factorization of the large sparse matrix \mathbf{A} (line 2). This is done once, prior to the nonlinear iteration; hence, its computational cost is amortized over the nonlinear solution process. On the other hand, the fill-in resulting in the factors of \mathbf{A} , especially in three dimensions, makes application of \mathbf{A}^{-1} expensive with complexity considerably larger than $O(N)$. This occurs M times on line 5 and thus potentially becomes a bottleneck of the algorithm if N and/or M are large. Once \mathbf{H}_m is computed, the construction of the capacitance matrix $\mathbf{C}_{\beta, \mathbf{m}}$ on line 5 can proceed very efficiently in $O(M^2 N)$ operations as a Level 3 BLAS operation. The dense solve on line 6 costs $O(M^3)$ and can be efficiently performed by LAPACK.

The evaluation of the model response and its derivative on line 4 is assumed to be available as a given function $\mathbf{m} \mapsto (\mathbf{g}_m, \mathbf{J}_m)$. In many contexts where the mapping is based on a PDE model, the evaluation of \mathbf{g}_m requires the solution of a

Algorithm 1 (Gauss–Newton with direct solver).

Input: Parameter-to-observation map $\mathbf{g}: \mathbb{R}^N \rightarrow \mathbb{R}^M$, observational data $\mathbf{g}_{\text{obs}} \in \mathbb{R}^M$, reference parameter $\mathbf{m}_{\text{ref}} \in \mathbb{R}^N$, simplicial partition \mathcal{T}_h of Ω , Neumann boundary $\Gamma_N \subset \partial\Omega$, initial guess $\mathbf{m} \in \mathbb{R}^N$, regularization parameter $\beta > 0$

Output: Final $\mathbf{m} \in \mathbb{R}^N$

- 1 Assemble matrices $\mathbf{Q} \in \mathbb{R}^{K \times K}$ and $\mathbf{D} \in \mathbb{R}^{N \times K}$ according to (2.13)
- 2 Use a sparse direct solver (e.g., sparse \mathbf{LDL}^\top) to factorize

$$\mathbf{A} = \begin{bmatrix} \mathbf{Q} & \mathbf{D}^\top \\ \mathbf{D} & \mathbf{0} \end{bmatrix}$$

3 repeat

- 4 Compute model response $\mathbf{g}_m \in \mathbb{R}^M$ and its derivative $\mathbf{J}_m \in \mathbb{R}^{M \times N}$ according to (2.14e, f, and a)
- 5 Use the factorization of \mathbf{A} from above to construct $\mathbf{H}_m \in \mathbb{R}^{N \times M}$ and $\mathbf{C}_{\beta, m} \in \mathbb{R}^{M \times M}$ such that

$$\mathbf{H}_m := -\mathbf{P}_2 \mathbf{A}^{-1} \begin{bmatrix} \mathbf{0} \\ \mathbf{J}_m^\top \end{bmatrix}, \quad \mathbf{C}_{\beta, m} := \mathbf{I}_M + \frac{1}{\beta} \mathbf{J}_m \mathbf{H}_m$$

- 6 Solve the capacitance system directly, i.e., find $\mathbf{y} \in \mathbb{R}^M$ such that

$$\mathbf{C}_{\beta, m} \mathbf{y} = \mathbf{J}_m (\mathbf{m} - \mathbf{m}_{\text{ref}}) - (\mathbf{g}_m - \mathbf{g}_{\text{obs}})$$

- 7 Compute $\delta \mathbf{m} \in \mathbb{R}^N$ as

$$\delta \mathbf{m} := -(\mathbf{m} - \mathbf{m}_{\text{ref}}) + \frac{1}{\beta} \mathbf{H}_m \mathbf{y}$$

8 $\mathbf{m} := \mathbf{m} + \delta \mathbf{m}$

9 until happy

forward PDE problem, and the computation of the associated derivative \mathbf{J}_m can be performed efficiently using adjoint techniques. This will be the case in the numerical examples presented in section 4.

3.2. Iterative solution. We consider the block-diagonal preconditioners

$$(3.8) \quad \mathbf{P} := \begin{bmatrix} \mathbf{Q} & \\ & \mathbf{S} \end{bmatrix} \quad \text{and} \quad \mathbf{P}_{\beta, m} := \begin{bmatrix} \mathbf{Q} & \\ & \mathbf{S} + \frac{1}{\beta} \mathbf{J}_m^\top \mathbf{J}_m \end{bmatrix},$$

with the Laplace Schur complement \mathbf{S} as in (3.1). These are “ideal” preconditioners for \mathbf{A} and $\mathbf{A}_{\beta, m}$, respectively. Indeed, the minimal polynomial of $\mathbf{A} \mathbf{P}^{-1}$ has degree at most 3 [41, Proposition 1], and, as a consequence, minimum residual Krylov subspace iteration applied to $\mathbf{A} \mathbf{P}^{-1}$ converges in at most three iterations, as shown by Murphy, Golub, and Wathen [41] (see also [37, Theorem 2.2.3]). This does not hold for $\mathbf{A}_{\beta, m} \mathbf{P}_{\beta, m}^{-1}$, but it is known that the spectrum of $\mathbf{A}_{\beta, m} \mathbf{P}_{\beta, m}^{-1}$ is contained in $[-1, -\frac{1}{\phi}] \cup [1, \phi]$, where $\phi = \frac{1+\sqrt{5}}{2}$; see [45, Theorem 4]. This inclusion guarantees two-step linear convergence of MINRES for $\mathbf{A}_{\beta, m} \mathbf{P}_{\beta, m}^{-1}$ independently of M , N , β , and the right-hand side; see, e.g., [33, section 3.1] or [25, section 4.2.4].

The action of \mathbf{P}^{-1} and $\mathbf{P}_{\beta, m}^{-1}$ is essentially as expensive as that of \mathbf{A}^{-1} and $\mathbf{A}_{\beta, m}^{-1}$, respectively; hence, we seek a good and inexpensive approximation of \mathbf{P}^{-1} and $\mathbf{P}_{\beta, m}^{-1}$. Consider

$$\hat{\mathbf{P}}^{-1} := \begin{bmatrix} \hat{\mathbf{Q}}^{-1} & \\ & \hat{\mathbf{S}}^{-1} \end{bmatrix} \quad \text{and} \quad \hat{\mathbf{P}}_{\beta, m}^{-1} := \begin{bmatrix} \hat{\mathbf{Q}}^{-1} & \\ & \hat{\mathbf{S}}_{\beta, m}^{-1} \end{bmatrix},$$

Algorithm 2 (Gauss–Newton with iterative solver).

Input: Parameter-to-observation map $\mathbf{g}: \mathbb{R}^N \rightarrow \mathbb{R}^M$, observational data $\mathbf{g}_{\text{obs}} \in \mathbb{R}^M$, reference parameter $\mathbf{m}_{\text{ref}} \in \mathbb{R}^N$, simplicial partition \mathcal{T}_h of Ω , Neumann boundary $\Gamma_N \subset \partial\Omega$, initial guess $\mathbf{m} \in \mathbb{R}^N$, regularization parameter $\beta > 0$

Output: Final $\mathbf{m} \in \mathbb{R}^N$

- 1 Assemble matrices $\mathbf{Q} \in \mathbb{R}^{K \times K}$ and $\mathbf{D} \in \mathbb{R}^{N \times K}$ according to (2.13)
- 2 Prepare a mass term preconditioner

$$\hat{\mathbf{Q}}^{-1} := (\text{diag } \mathbf{Q})^{-1}$$

- 3 Prepare a Schur complement preconditioner using an algebraic blackbox, e.g., algebraic multigrid,

$$\hat{\mathbf{S}}^{-1} := \text{AMG}(\mathbf{D}(\text{diag } \mathbf{Q})^{-1} \mathbf{D}^\top)$$

4 **repeat**

- 5 Compute model response $\mathbf{g}_m \in \mathbb{R}^M$ and its derivative $\mathbf{J}_m \in \mathbb{R}^{M \times N}$ according to (2.14e, f, and a)
 - 6 Compute $\delta \mathbf{m} \in \mathbb{R}^N$ using Algorithm 3 or Algorithm 4
 - 7 $\mathbf{m} := \mathbf{m} + \delta \mathbf{m}$
 - 8 **until** *happy*
-

where

$$\begin{aligned} \hat{\mathbf{Q}}^{-1} &:= (\text{diag } \mathbf{Q})^{-1}, \\ \hat{\mathbf{S}}^{-1} &:= \text{AMG}(\mathbf{D} \hat{\mathbf{Q}}^{-1} \mathbf{D}^\top), \\ \hat{\mathbf{S}}_{\beta, \mathbf{m}}^{-1} &:= \hat{\mathbf{S}}^{-1} - \frac{1}{\beta} \hat{\mathbf{S}}^{-1} \mathbf{J}_m^\top (\mathbf{I}_M + \frac{1}{\beta} \mathbf{J}_m \hat{\mathbf{S}}^{-1} \mathbf{J}_m^\top)^{-1} \mathbf{J}_m \hat{\mathbf{S}}^{-1}. \end{aligned}$$

The preconditioner $\hat{\mathbf{P}}^{-1}$ was introduced by Powell and Silvester [46] for preconditioning the mixed Laplacian \mathbf{A} . We can employ this preconditioner also for $\mathbf{A}_{\beta, \mathbf{m}}$, which is, in view of (3.4), a perturbation of \mathbf{A} by at most rank M . The expressions for $\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}$ follow easily by requiring, in analogy to (3.8), that $\hat{\mathbf{S}}_{\beta, \mathbf{m}} = \hat{\mathbf{S}} + \frac{1}{\beta} \mathbf{J}_m^\top \mathbf{J}_m$ and by using the Woodbury matrix identity.

Algorithm 2 summarizes the Gauss–Newton procedure based on iterative solution of the linearized problems. It invokes either Algorithm 3, which employs $\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}$ as a preconditioner, or Algorithm 4, which uses $\hat{\mathbf{P}}^{-1}$. The latter omits the correction due to the Woodbury formula, hence bypassing the computations involving the capacitance matrix $\mathbf{C}_{\beta, \mathbf{m}} = \mathbf{I}_M + \frac{1}{\beta} \mathbf{J}_m \hat{\mathbf{S}}^{-1} \mathbf{J}_m^\top$, and thus results in a less expensive preconditioner. As it fails to account for the low-rank modification due to the data misfit term, it is expected to deteriorate with increasing M . We will confirm this experimentally in section 4. Additionally, we will see that $\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}$, in contrast to $\hat{\mathbf{P}}^{-1}$, provides robustness w.r.t. β ; see Figure 6.

The preconditioned MINRES procedures in Algorithms 3 and 4 correspond to different minimization problems,

$$(3.9) \quad \begin{aligned} \text{MINRES}(\mathbf{A}_{\beta, \mathbf{m}}, \hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}, \mathbf{b}, \mathbf{x}_0) &: \|\mathbf{r}_k\|_{\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}} = \min_{p \in \mathcal{P}_k^0} \|p(\mathbf{A}_{\beta, \mathbf{m}} \hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}) \mathbf{r}_0\|_{\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}}, \\ \text{MINRES}(\mathbf{A}_{\beta, \mathbf{m}}, \hat{\mathbf{P}}^{-1}, \mathbf{b}, \mathbf{x}_0) &: \|\mathbf{r}_k\|_{\hat{\mathbf{P}}^{-1}} = \min_{p \in \mathcal{P}_k^0} \|p(\mathbf{A}_{\beta, \mathbf{m}} \hat{\mathbf{P}}^{-1}) \mathbf{r}_0\|_{\hat{\mathbf{P}}^{-1}}, \end{aligned}$$

Algorithm 3 (MINRES with the Laplace–Woodbury preconditioner $\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}$).

Input: $\mathbf{m} \in \mathbb{R}^N$, $\mathbf{g}_m \in \mathbb{R}^M$, $\mathbf{J}_m \in \mathbb{R}^{M \times N}$, $\beta > 0$

Output: $\delta \mathbf{m} \in \mathbb{R}^N$

- 1 $\hat{\mathbf{H}}_m := \hat{\mathbf{S}}^{-1} \mathbf{J}_m^\top$
- 2 $\mathbf{C}_{\beta, \mathbf{m}} := \mathbf{I}_M + \frac{1}{\beta} \mathbf{J}_m \hat{\mathbf{H}}_m$
- 3 Compute Cholesky factor $\mathbf{L}_{\beta, \mathbf{m}} \in \mathbb{R}^{M \times M}$ such that $\mathbf{L}_{\beta, \mathbf{m}} \mathbf{L}_{\beta, \mathbf{m}}^\top = \mathbf{C}_{\beta, \mathbf{m}}$
- 4 Run MINRES:

$$\begin{bmatrix} \zeta \\ \delta \mathbf{m} \end{bmatrix} := \text{MINRES}(\mathbf{A}_{\beta, \mathbf{m}}, \hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}, \mathbf{b}, \mathbf{x}_0),$$

where the system operator $\mathbf{A}_{\beta, \mathbf{m}}: \mathbb{R}^{K+N} \rightarrow \mathbb{R}^{K+N}$ and the preconditioner $\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}: \mathbb{R}^{K+N} \rightarrow \mathbb{R}^{K+N}$ are represented matrix-free by formulas

$$\mathbf{A}_{\beta, \mathbf{m}} \begin{bmatrix} \zeta \\ \delta \mathbf{m} \end{bmatrix} = \begin{bmatrix} \mathbf{Q}\zeta + \mathbf{D}^\top \delta \mathbf{m} \\ \mathbf{D}\zeta - \frac{1}{\beta} (\mathbf{J}_m^\top (\mathbf{J}_m \delta \mathbf{m})) \end{bmatrix},$$

$$\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1} \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{Q}}^{-1} \mathbf{y}_1 \\ \hat{\mathbf{S}}^{-1} \mathbf{y}_2 - \frac{1}{\beta} (\hat{\mathbf{H}}_m (\mathbf{L}_{\beta, \mathbf{m}}^{-\top} (\mathbf{L}_{\beta, \mathbf{m}}^{-1} (\hat{\mathbf{H}}_m^\top \mathbf{y}_2)))) \end{bmatrix}$$

and the right-hand side $\mathbf{b} \in \mathbb{R}^{K+N}$ and the initial guess $\mathbf{x}_0 \in \mathbb{R}^{K+N}$ are given by

$$\mathbf{b} := \begin{bmatrix} -\mathbf{D}^\top (\mathbf{m} - \mathbf{m}_{\text{ref}}) \\ \frac{1}{\beta} \mathbf{J}_m^\top (\mathbf{g}_m - \mathbf{g}_{\text{obs}}) \end{bmatrix}, \quad \mathbf{x}_0 := \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}.$$

Algorithm 4 (MINRES with the Laplace preconditioner $\hat{\mathbf{P}}^{-1}$).

Input: $\mathbf{m} \in \mathbb{R}^N$, $\mathbf{g}_m \in \mathbb{R}^M$, $\mathbf{J}_m \in \mathbb{R}^{M \times N}$, $\beta > 0$

Output: $\delta \mathbf{m} \in \mathbb{R}^N$

- 1 Run MINRES:

$$\begin{bmatrix} \zeta \\ \delta \mathbf{m} \end{bmatrix} := \text{MINRES}(\mathbf{A}_{\beta, \mathbf{m}}, \hat{\mathbf{P}}^{-1}, \mathbf{b}, \mathbf{x}_0),$$

where the system operator $\mathbf{A}_{\beta, \mathbf{m}}: \mathbb{R}^{K+N} \rightarrow \mathbb{R}^{K+N}$ and the preconditioner $\hat{\mathbf{P}}^{-1}: \mathbb{R}^{K+N} \rightarrow \mathbb{R}^{K+N}$ are represented matrix-free by formulas

$$\mathbf{A}_{\beta, \mathbf{m}} \begin{bmatrix} \zeta \\ \delta \mathbf{m} \end{bmatrix} = \begin{bmatrix} \mathbf{Q}\zeta + \mathbf{D}^\top \delta \mathbf{m} \\ \mathbf{D}\zeta - \frac{1}{\beta} (\mathbf{J}_m^\top (\mathbf{J}_m \delta \mathbf{m})) \end{bmatrix},$$

$$\hat{\mathbf{P}}^{-1} \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{Q}}^{-1} \mathbf{y}_1 \\ \hat{\mathbf{S}}^{-1} \mathbf{y}_2 \end{bmatrix}$$

and the right-hand side $\mathbf{b} \in \mathbb{R}^{K+N}$ and the initial guess $\mathbf{x}_0 \in \mathbb{R}^{K+N}$ are given by

$$\mathbf{b} := \begin{bmatrix} -\mathbf{D}^\top (\mathbf{m} - \mathbf{m}_{\text{ref}}) \\ \frac{1}{\beta} \mathbf{J}_m^\top (\mathbf{g}_m - \mathbf{g}_{\text{obs}}) \end{bmatrix}, \quad \mathbf{x}_0 := \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}.$$

where \mathcal{P}_k^0 denotes the set of polynomials of degree at most k normalized to $p(0) = 1$, $\mathbf{r}_k = \mathbf{b} - \mathbf{A}_{\beta, \mathbf{m}} \mathbf{x}_k$ are the true residuals corresponding to the k th iterates $\mathbf{x}_k = [\zeta_k^\top, \delta \mathbf{m}_k^\top]^\top$, and the norm $\|\mathbf{x}\|_M = (\mathbf{x}^\top \mathbf{M} \mathbf{x})^{1/2}$ for a symmetric positive definite \mathbf{M} . In particular, one can see that different residual norms are used.

To assess the complexity of Algorithm 3, we assume that the black-box preconditioners $\hat{\mathbf{Q}}^{-1}$ and $\hat{\mathbf{S}}^{-1}$ are optimal; i.e., the actions $\hat{\mathbf{Q}}^{-1} \mathbf{y}_1$, $\hat{\mathbf{S}}^{-1} \mathbf{y}_2$ on vectors $\mathbf{y}_1 \in \mathbb{R}^K$, $\mathbf{y}_2 \in \mathbb{R}^N$ are performed in $O(K)$ and $O(N)$ floating-point operations, respectively. In the settings under consideration, we have $M \leq N$ (typically $M \ll N$) and $O(K) = O(N)$. Moreover, we do not distinguish between complexity for number of floating point operations and execution times. A breakdown of the complexity of the steps in Algorithm 3 is as follows:

- line 1. M applications of $\hat{\mathbf{S}}^{-1}$, i.e., $O(MN)$;
- line 2. dense matrix-matrix multiply; $O(M^2N)$;
- line 3. dense Cholesky factorization; $O(M^3)$;
- line 4. cost per one MINRES step is $O(MN)$ because $\mathbf{Q}\zeta$, $\mathbf{D}^\top \delta \mathbf{m}$, $\mathbf{D}\zeta$, $\hat{\mathbf{Q}}^{-1} \mathbf{y}_1$, and $\hat{\mathbf{S}}^{-1} \mathbf{y}_2$ are $O(N)$; $\mathbf{J}_m^\top (\mathbf{J}_m \delta \mathbf{m})$, $\hat{\mathbf{H}}_m^\top \mathbf{y}_2$, and $\hat{\mathbf{H}}_m \cdot$ are $O(MN)$; and $\mathbf{L}_{\beta, \mathbf{m}}^{-\top} \cdot$ and $\mathbf{L}_{\beta, \mathbf{m}}^{-1} \cdot$ are $O(M^2)$.

If the number of MINRES iterations remains constant independent of M and N , one observes that the overall complexity of Algorithm 3 is dominated by $O(M^2N)$ due to the assembly of the capacitance matrix on line 2. On the other hand, this operation would typically be carried out by the Level 3 BLAS routine `gemm`, thus very efficiently (in terms of utilizing the theoretical floating point capability of the CPU). Note that one must not assemble $(\mathbf{J}_m^\top \mathbf{J}_m)$, which would be a dense $\mathbb{R}^{N \times N}$ matrix and thus would degrade the complexity to $O(N^2)$. We will demonstrate via the numerical experiments in section 4 that the number of MINRES iterations in Algorithm 3 tends to be constant.

On the other hand, the simplified Algorithm 4 has, by the same reasoning, complexity of only $O(MN)$ per MINRES iteration, but the number of MINRES iterations tends to increase as M and N grow, which we will see confirmed in section 4. Moreover, Algorithm 4 is not robust w.r.t. β ; see Figure 6.

4. An application: ERT. Consider a conducting medium occupying a domain $\Omega \subset \mathbb{R}^d$ characterized by an unknown spatially varying electrical conductivity $\sigma_{\text{true}}: \Omega \rightarrow (0, \infty)$. ERT (also known as *the direct current (DC) resistivity method* in the geophysical exploration literature) reconstructs the unknown σ_{true} from voltage measurements of stationary electric fields excited by known synthetic DC sources. We model the excitation current by a source-sink pair of point sources of known DC current strength. This corresponds physically to a current source connected to the medium at two distinct points by way of cables (conductors), while the cables themselves are not part of the conductivity model σ_{true} but are rather represented as a point source and point sink, respectively. The response of the medium to this excitation can be measured as a voltage (potential difference) at two other points in the medium. By varying the placement of current source-sink and/or the voltage electrode positions, one can perform multiple measurements. Ultimately, one wishes to reconstruct a conductivity distribution $\tilde{\sigma}$ which is consistent with these measurements. A finite set of such measurements is likely to be explained equally well by multiple different values of $\tilde{\sigma}$, indicating that the problem is underdetermined. Moreover, the (inverse) problem of reconstructing conductivity from potential measurements is well known to be ill-posed. As a selection criterion, one can ask for extra smoothness of $\tilde{\sigma}$ and thus regularize the inverse problem. In any case, it is clear that, except for special cases, it cannot be expected that $\tilde{\sigma} = \sigma_{\text{true}}$.

Consider a bounded Lipschitz domain $\Omega \subset \mathbb{R}^d$, $d = 2, 3$, and electrical conductivity $\sigma \in L^\infty(\Omega)$, $\sigma \geq \sigma_0 > 0$. Assume that $\partial\Omega = \overline{\gamma_D} \cup \overline{\gamma_N}$ with open and disjoint γ_D, γ_N and such that $|\gamma_D| > 0$. Note that γ_D and γ_N are, in general, different from Γ_D and Γ_N from (2.2). We consider the diffusion equation for the stationary electric potential u ,

$$(4.1a) \quad -\operatorname{div} \sigma \nabla u = f \quad \text{in } \Omega,$$

$$(4.1b) \quad u = 0 \quad \text{on } \gamma_D,$$

$$(4.1c) \quad \frac{\partial}{\partial n} u = 0 \quad \text{on } \gamma_N,$$

where we employ the homogeneous boundary conditions (4.1b, c) for simplicity. The electric potential $u_{x_A x_B}$ for a unit current source-sink pair $x_A \neq x_B$ in $\Omega \cup \gamma_N$ is then defined as the distributional solution of (4.1) with $f := \delta_{x_A} - \delta_{x_B}$. Note that it makes sense to place x_A and/or x_B on γ_N . The distributional solution $u_{x_A x_B}$ does not belong to the Sobolev space $H^1(\Omega)$. Nevertheless, $u_{x_A x_B}$ is continuous in $\Omega \setminus (x_A \cup x_B)$; see [42, equation (3)]. One can therefore define the voltage difference

$$(4.2) \quad u_{x_A x_B}(x_M) - u_{x_A x_B}(x_N) = \langle \delta_{x_M} - \delta_{x_N}, u_{x_A x_B} \rangle$$

between any two points $x_M, x_N \in \Omega \setminus (x_A \cup x_B)$. We define the solution operator for (4.1):

$$A_\sigma^{-1}: f \mapsto u \quad \text{such that } \sigma, f, \text{ and } u \text{ satisfy (4.1) in the sense of distributions.}$$

With this definition, we may express the quantity in (4.2) as

$$(4.3) \quad \begin{aligned} \langle \delta_{x_M} - \delta_{x_N}, A_\sigma^{-1}(\delta_{x_A} - \delta_{x_B}) \rangle &= \langle \delta_{x_A} - \delta_{x_B}, A_\sigma^{-1}(\delta_{x_M} - \delta_{x_N}) \rangle \\ &= \int_\Omega \sigma \nabla A_\sigma^{-1}(\delta_{x_A} - \delta_{x_B}) \cdot \nabla A_\sigma^{-1}(\delta_{x_M} - \delta_{x_N}), \end{aligned}$$

and the Gâteaux derivative of this quantity is readily expressed as¹

$$(4.4) \quad \delta\sigma \mapsto - \int_\Omega \delta\sigma \nabla A_\sigma^{-1}(\delta_{x_A} - \delta_{x_B}) \cdot \nabla A_\sigma^{-1}(\delta_{x_M} - \delta_{x_N}),$$

which is a linear functional.

It is convenient to introduce the change of variables $m = \log \sigma$ for the conductivity so that for $m \in L^\infty(\Omega)$, one has $0 < \exp(\operatorname{ess\,inf}_\Omega m) \leq \sigma \leq \exp(\operatorname{ess\,sup}_\Omega m)$. The solution map $A_{\exp(m)}^{-1}$ is then well-defined for all $m \in L^\infty(\Omega)$, as the boundedness condition $0 < \underline{\sigma} \leq \sigma \leq \bar{\sigma} < \infty$ is equivalent to $m \in L^\infty(\Omega)$.

A practical ERT survey consists of multiple measurements using different combinations of points x_A^i, x_B^i, x_M^i , and x_N^i for $i = 1, 2, \dots, M$. Following (4.3) and (4.4), we express the quantity of interest and its derivative as

$$(4.5) \quad \begin{aligned} g^i(m) &:= k_i \int_\Omega \exp(m) \nabla A_{\exp(m)}^{-1}(\delta_{x_A^i} - \delta_{x_B^i}) \cdot \nabla A_{\exp(m)}^{-1}(\delta_{x_M^i} - \delta_{x_N^i}), \\ J^i(m) \delta m &:= -k_i \int_\Omega \delta m \exp(m) \nabla A_{\exp(m)}^{-1}(\delta_{x_A^i} - \delta_{x_B^i}) \cdot \nabla A_{\exp(m)}^{-1}(\delta_{x_M^i} - \delta_{x_N^i}) \\ &\quad \text{for } i = 1, 2, \dots, M. \end{aligned}$$

¹This follows along the lines of the formula $d(A^{-1}) = -A^{-1} dA A^{-1}$, which is valid for any invertible matrix A . Concerning the Gâteaux derivative of the singular integral (4.3), additional rigor and care in the choice of the function spaces are needed, but this is beyond the scope of this work; hence, we proceed just formally.

Here we have introduced additional scaling factors k_i given by

$$(4.6) \quad k_i := \begin{cases} \frac{\pi}{-\log|x_A^i - x_M^i| + \log|x_B^i - x_M^i| + \log|x_A^i - x_N^i| - \log|x_B^i - x_N^i|}, & d=2, \\ \frac{2\pi}{|x_A^i - x_M^i|^{-1} - |x_B^i - x_M^i|^{-1} - |x_A^i - x_N^i|^{-1} + |x_B^i - x_N^i|^{-1}}, & d=3, \end{cases} \quad \text{for } i=1, 2, \dots, M.$$

These *geometric factors* only depend on the coordinates of the electrodes. Their purpose is that the original voltage measurement (4.3) is transformed into a quantity known as *apparent resistivity*.² This is a commonly applied method of scaling the data $g^i(m) - g_{\text{obs}}^i$, $i=1, \dots, M$.

In practice, the placement of electrodes x_A^i , x_B^i , x_M^i , and x_N^i is critical for the goal of approximating the original conductivity distribution, i.e., $\tilde{\sigma} \approx \sigma_{\text{true}}$. The geophysics literature contains a number of established electrode placement designs; see, e.g., [51, section 8.5] or [52] and the references therein. In the examples below, we consider what is known as a *pole-dipole configuration*. Figure 1 shows a sequence of one-dimensional electrode configurations, which exhibit increasing measurement resolution but with sensitivity only in regions increasingly closer to the surface as the configuration is refined. This sequence was chosen to obtain a series of problems which are each meaningful for the underlying inverse problem and at the same time illustrate the performance of the preconditioners across a wide range of values for the finite element mesh size, number of measurements, and regularization parameter. In the framework of (4.5) and (4.6), the x_B^i -electrode is modeled as an electrode placed at ∞ (resp., at γ_D) in the context of the boundary datum (4.1b). Hence, $\delta_{x_B^i}$ does not contribute to (4.5), and the factors k_i are obtained by taking limit $|x_B^i| \rightarrow \infty$ in (4.6). These one-dimensional configurations are typically used in ERT surveys along the upper boundary of a two-dimensional vertical cross section. For surveys over a three-dimensional region, it is common to construct a two-dimensional surface electrode configuration as the Cartesian product of the one-dimensional pattern.

In the following, we illustrate the performance of the aforementioned algorithms with a sequence of parameter identification experiments in an idealized ERT setting. We aim to reconstruct an a priori known conductivity anomaly against a homogeneous background. We consider a sequence of problems involving a checkerboard anomaly structure of increasing complexity with decreasing depth in accordance with the sensitivity and resolution capability of the chosen electrode configuration designs. Sequences of problems in two (the left column in Figure 2) and three (the top row in Figure 3) spatial dimensions are carefully chosen to work well with the aforementioned electrode configuration. In particular, because the spacing of electrodes decreases with finer configurations, the survey is only sensitive in an increasingly shallow region below the surface. This rather artificial scenario allows us to reconstruct an increasingly finer pattern with only $M = O(N)$ measurements (see Figure 1), thus allowing us to increase the parameters M and N many times before exhausting computer resources (see Table 1).

²A measurement $g^i(m)$ from (4.5) gives apparent resistivity of a homogeneous half space. Precisely, it holds true that $g^i(\log \sigma_0) = 1/\sigma_0$, for a constant $\sigma_0 > 0$, half-space domain $\Omega = \{x \in \mathbb{R}^d, x_d > 0\}$, and $x_A^i, x_B^i, x_M^i, x_N^i \in \{x \in \mathbb{R}^d, x_d = 0\}$. This is derived using Green's functions for the Laplace Dirichlet problem in a half space.




FIG. 1. The first five pole-dipole electrode configurations used in the numerical examples, starting from $N^{\text{ele}} = 17$ distinct electrode positions (left) and reaching as many as $N^{\text{ele}} = 257$ positions (right). The vertical axes enumerate measurement number $i = 1, 2, \dots, M$, and the horizontal axes represent electrode x -position with N^{ele} distinct equidistant positions in the interval $x \in [-50, 50]$. Transmitter electrodes x_A^i (∇), receiver electrodes x_M^i (\times), second transmitter electrode x_B^i placed at ∞ (not shown), $i = 1, 2, \dots, M$. The measurements $i = 1, 2, \dots, 2N^{\text{ele}} - 8$ use the spacing 2; e.g., for $i = 1$, the electrodes are at positions 1, 3, 5; for $i = 2$, the electrodes are at positions 2, 4, 6; etc. The measurements $i = 2N^{\text{ele}} - 7, \dots, 4N^{\text{ele}} - 24$ use the spacing 4, and the measurements $i = 4N^{\text{ele}} - 23, \dots, 6N^{\text{ele}} - 56$ use the spacing 8. Total number of measurements is $M = 6N^{\text{ele}} - 56$.

Two-dimensional test case. We consider the half-disk domain $\Omega := \{(x, z) \in \mathbb{R}^2, z > 0, \sqrt{x^2 + z^2} < 80\}$. The line $\{z = 0\}$ represents the ground surface where measurements are taken using electrodes placed as described in Figure 1. Following geophysical convention, the half-space $\{z > 0\}$ represents the subsurface consisting of a medium with (here a priori known) conductivity distribution σ_{true} as in Figure 2 (left), displaying a series of increasingly finer anomalous conductivity patterns imposed on a background medium of constant conductivity. The opposite side $\{z < 0\}$ represents the air half-space of negligible conductivity, which is thus excluded from the domain and modeled by a vanishing normal component of the electric field (4.1c) on $\gamma_N := \{z = 0\}$. For simplicity, we consider (4.1b) on $\gamma_D := \partial\Omega \setminus \gamma_N$. This description fully specifies the functions $m \mapsto g^i(m)$ and $m \mapsto J^i(m)$, $i = 1, 2, \dots, M$.

For the configurations, the first five of which are indicated in Figure 2, we compute the finite element approximations of the quantities $g_{\text{obs}}^i := g^i(\log \sigma_{\text{true}})$, $i = 1, 2, \dots, M$,




FIG. 2. True resistivities $\frac{1}{\sigma_{\text{true}}}$ (on the left) and the result of inversion using Algorithm 3 (on the right) for the two-dimensional example. The series of configurations (from top to bottom) corresponds to the first five configurations of electrodes; see Figure 1. The majority of the medium has background resistivity $\frac{1}{\sigma_{\text{ref}}} = 3500$ (blue, on the left) with presence of anomaly resistivity 7000 (red, on the left). The positions of the electrodes at the surface are indicated by the black vertical bar (\cdot , on the left).

which serve as the (synthetic) observational data for inversion. Note that these data are noisy due to the discretization error (although the meshes used to generate the values g_{obs}^i are finer compared to the meshes for the inversion). The reference value is taken to be $m_{\text{ref}} := \log \sigma_{\text{ref}}$ as in Figure 2, and Γ_{D} in (2.2) is taken as $\Gamma_{\text{D}} := \partial\Omega$. Two

TABLE 1

Performance characteristics of the numerical experiments. Timings t_1 , t_2 , and t_3 for substeps of Algorithm 3 for each Gauss–Newton step i . Number of MINRES iterations n_{iter} and overall runtime t_{norm} for Algorithms 3 and 4 to solve the normal equations $\mathbf{A}_{\beta, \mathbf{m}} \mathbf{x} = \mathbf{b}$ within tolerance $\|\mathbf{r}_k\|_2 / \|\mathbf{b}\|_2 \leq 10^{-7}$ in the Euclidean norm. Cases marked † did not converge to the prescribed tolerance in $2(K+N)$ iterations (recall that $\mathbf{A}_{\beta, \mathbf{m}} \in \mathbb{R}^{(K+N) \times (K+N)}$); tolerance 10^{-6} was reached in all these cases.

	N^{ele}	N	M	i	Algorithm 3					Algorithm 4	
					t_1 [s]	t_2 [s]	t_3 [s]	n_{iter}	t_{norm} [s]	n_{iter}	t_{norm} [s]
two dimensions ($\beta = 0.1$)	17	840	46	1	0.01	0.0019	0.0001	4	0.02	79	0.02
				2	0.01	0.0007	0.0001	11	0.02	286	0.07
	33	1584	142	1	0.03	0.0029	0.0050	4	0.07	737	0.38
				2	0.03	0.0019	0.0024	12	0.05	3526	1.58
	65	3140	334	1	0.16	0.0116	0.0010	4	0.19	1921	6.46
				2	0.15	0.0113	0.0007	14	0.22	15878†	39.69
	129	6012	718	1	0.61	0.0930	0.0026	4	0.80	3441	33.71
				2	0.63	0.0885	0.0118	14	1.09	29883†	331.63
	257	11644	1486	1	2.49	0.4523	0.0133	4	3.24	5700	180.17
				2	2.64	0.4656	0.0139	17	3.88	58798†	2016.60
	513	22884	3022	1	11.01	1.8966	0.0668	4	13.98		
				2	10.78	1.9269	0.0567	14	15.40		
	1025	44848	6094	1	44.22	14.6608	0.3928	4	63.39		
				2	44.93	14.6799	0.3932	12	69.61		
	2049	89608	12238	1	190.69	117.3031	3.0011	6	330.03		
				2	199.42	117.3992	2.9831	14	357.53		
4097	178232	24526	1	848.10	932.6321	23.4441	7	1943.20			
			2	847.55	932.3276	22.3924	14	1963.20			
three dimensions ($\beta = 10^5$)	81	120192	216	1	5.58	0.1238	0.0005	130	24.19	156	15.77
				2	5.61	0.1209	0.0004	132	20.98	160	15.43
	169	262464	728	1	54.31	2.6834	0.0036	123	139.58	210	94.94
				2	53.76	2.6758	0.0028	125	169.42	215	72.15
	289	452736	1564	1	171.92	10.1975	0.0155	124	394.06	291	277.57
				2	176.60	10.2716	0.0153	126	392.27	297	314.99
	441	679296	2940	1	536.86	53.4599	0.1874	136	1208.20	425	1363.20
				2	516.93	53.4413	0.0981	139	1259.60	438	1204.00
	625	937408	4700	1	1146.50	184.4371	0.2116	131	2840.50	526	2989.20
				2	1270.10	183.5881	0.2154	136	2988.60	542	3100.70

Gauss–Newton steps with a fixed value of the regularization parameter $\beta := 0.1$ are performed, and the resistivity distributions in Figure 2 (on the right) are obtained. The meshes for the inversion (see Figure 2) are a priori refined around the electrode positions, which are at the surface $\{z = 0\}$, so that the meshes scale as $N = O(N^{\text{ele}})$; see Table 1.

Three-dimensional test case. Here we consider the semispherical domain $\Omega := \{(x, y, z) \in \mathbb{R}^3, z > 0, \sqrt{x^2 + y^2 + z^2} < 80\}$. The measurements are again taken on the surface $\{z = 0\}$ using the grid of electrodes shown in Figure 3. One uses the pole-dipole scheme (as described in Figure 1) along the x -direction for all possible $y = \text{const}$ profiles and then the same in the y -direction for all possible $x = \text{const}$ profiles. By analogy, the true resistivity model is also constructed in a Cartesian product fashion; see Figure 3.

The remaining details are analogous to the two-dimensional test case above with the exception that different values of the regularization parameter β were necessary to obtain good reconstructions. The question of choosing the best value of the regulariza-




FIG. 3. Series of three-dimensional computational examples of increasing difficulty (left to right). Domain indicated by slices $x = 0$, $y = 0$, and $z = 0$ (top row); electrode positions indicated by black dots (top row). True resistivities $\frac{1}{\sigma_{\text{true}}}$ in the majority of the medium is $\frac{1}{\sigma_{\text{ref}}} = 3500$, with presence of anomaly resistivity 7000 (top row, red blocks). Result of inversion using Algorithm 3 with $\beta = 10^5$ (remaining rows; sections through $z = \text{const}$ planes); the slices framed in the red frame correspond to the top and the bottom of the anomalous resistivity (red blocks).




FIG. 4. A fixed three-dimensional example (third column from Figure 3, $N^{\text{ele}} = 289$, $N = 452736$, $M = 1564$) computed for series of regularization parameters β using Algorithm 3.

tion parameter is beyond the scope of this paper. Nevertheless, we experimented with a number of choices and noticed how this affects the performance of the algorithms. For the sake of illustration, we indicate in Figure 4 the effect of the regularization parameter on the reconstructed conductivity.




FIG. 5. Time for solving normal equations t_{norm} for the two-dimensional (top) and three-dimensional (bottom) example. Each timing corresponds to the runtime of Algorithm 3, Algorithm 4, or lines 5–7 of Algorithm 1 in each Gauss–Newton step i .

The linear systems resulting from two-dimensional discretizations of (4.1) were solved using a sparse direct method and in three dimensions using conjugate gradient iteration preconditioned by an algebraic multigrid cycle. These choices make the approximation of $A_{\text{exp}(m)}^{-1}$ in (4.5) and in turn the computation of \mathbf{g}_m and \mathbf{J}_m sufficiently inexpensive and scalable, leaving the main effort in the solution of (2.15), which is the primary concern of this work.

The numerical experiments were implemented using MATLAB, HSL_MI20 [12], and Gmsh [28]. The plots were produced using matlab2tikz [27] and PyVista [50]. The complete code for reproducing the experiments is available in [10].

Table 1 and Figure 5 show that the computational cost of the examples agrees with the expected complexity as predicted in subsection 3.2. In particular, we can see that the dominating cost of Algorithm 3 is $O(MN)$, but we can see the $O(M^2N)$ term becoming effective for larger values of M . The Cholesky factorization for $O(M^3)$ (value t_3 in Table 1) and the matrix-matrix product for $O(M^2N)$ (value t_2 in Table 1) have a small multiplicative constant, as these would typically run very efficiently in LAPACK and BLAS, respectively. Nevertheless, it is clear that $O(M^2N)$ will dominate for larger problems.

Figure 6 shows the performance of a fixed three-dimensional test case depending on the value of the regularization parameter β . In particular, Algorithm 3 is seen to exhibit robust performance independent of β . For this, one has to pay the price of computing and factoring the capacitance matrix. Algorithm 4, on the other hand, shows strong dependence of the required number of MINRES steps on the value of β and MN . Although Algorithm 4 may sometimes be a less expensive alternative,




FIG. 6. Time for solving normal equations t_{norm} (top) and number of MINRES iterations (bottom) for a fixed three-dimensional example (third column from Figure 3) with series of regularization parameters β .

TABLE 2

	Algorithm 3	Algorithm 4
employed preconditioner	$\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}$	$\hat{\mathbf{P}}^{-1}$
$\mathbf{J}_m^\top \mathbf{J}_m$ handled by	Woodbury	Krylov
cost per MINRES iteration	$O(M^2 N)$	$O(MN)$
number of MINRES iterations	$O(1)$	$O(M^{\gamma_1} \beta^{-\gamma_2}), \gamma_1, \gamma_2 > 0$
overall robustness	✓	×

Algorithm 3 should generally be preferred for its robustness. To this end, we also note that with a better implementation of the solver for $\hat{\mathbf{S}}^{-1}$, one might achieve more favorable timings for the computation of the capacitance matrix. We have used HSL_MI20 [12], which is fully sequential in contrast to the threaded BLAS used in other parts of the code; eight threads were used where applicable. Moreover, HSL_MI20 only implements $\hat{\mathbf{S}}^{-1} \mathbf{z}$ for a single-column vector \mathbf{z} , but we need, on line 1 in Algorithm 3, to apply $\hat{\mathbf{S}}^{-1}$ to all the M columns of \mathbf{J}_m^\top . This operation therefore runs sequentially column by column, which is certainly not optimal in utilizing the theoretical floating-point performance and memory bandwidth of the machine. This implementation drawback penalizes Algorithm 3 in this experimental performance assessment, and it should be kept in mind that Algorithm 3 can be more favorable than Algorithm 4 whenever a suitable AMG implementation is available.

Furthermore, we have observed that for lower values of the regularization parameter β (for example, the two-dimensional case with $\beta = 0.001$, which is not shown in the paper), the solutions produced with preconditioners $\hat{\mathbf{P}}_{\beta, \mathbf{m}}^{-1}$ and $\hat{\mathbf{P}}^{-1}$ may differ

significantly, although they were solved to the same residual accuracy in the Euclidean norm: $\|\mathbf{r}_k\|_2/\|\mathbf{b}\|_2 \leq 10^{-7}$. Note that this has always been used as the stopping criterion in MINRES, although the minimization intrinsic to the preconditioned MINRES process minimizes a different quantity;³ see (3.9).

5. Conclusion and outlook. We have formulated a nonlinear parameter identification problem subject to H^1 regularization and its Gauss–Newton linearization as a second-order boundary value problem including a consistent interpretation of possible choices of boundary conditions as they result from the nature of the regularization procedure. For a standard inf-sup stable mixed discretization, we have proposed a number of efficient and robust solution strategies of the linear systems arising from the Gauss–Newton linearization. The proposed methods included a direct method, a preconditioned iterative scheme based on the Woodbury formula, and a preconditioned iterative scheme in which the low-rank perturbation is not accounted for by the preconditioner and must be compensated by the Krylov iteration. In a series of extensive numerical experiments, we have performed scaling tests w.r.t. the relevant problem parameters for a challenging parameter identification problem arising in ERT.

In Table 2, we summarize our findings concerning the interplay between the efficiency and quality of the two considered preconditioners. Note that regarding the indicated scaling of MINRES iterations required to solve the linear system to prescribed tolerance, we do not have a rigorous theoretical argument but merely empirical findings specific to the class of problems we solved; see Table 1 for the observed dependence on M and Figure 6 for the dependence on β . Specifically, the MINRES convergence behavior for $\hat{\mathbf{P}}^{-1}$ is sure to be problem dependent and likely depends on the distribution of singular values of \mathbf{J}_m . The simple parametrization M^{γ_1} observed here may only apply when \mathbf{J}_m 's are selected from a narrow class. We have observed that the variant with the full preconditioner $\hat{\mathbf{P}}_{\beta,m}^{-1}$ exhibits robustness of convergence for a range of parameter values M , N , and β , while the cheaper preconditioner $\hat{\mathbf{P}}^{-1}$ can suffer from slow convergence (Figure 6) or even stagnation (Table 1).

In future work, we would like to investigate data sparse approximation and fast solution of the capacitance matrix equation in applying the Woodbury formula, e.g., using \mathcal{H} -matrix methods and/or randomized low-rank approximations. This would allow applying the preconditioner $\hat{\mathbf{P}}_{\beta,m}^{-1}$ with a lower complexity than $O(M^2N)$.

Reproducibility of computational results. This paper has been awarded the “SIAM Reproducibility Badge: code and data available”, as a recognition that the authors have followed reproducibility principles valued by SISC and the scientific computing community. Code and data that allow readers to reproduce the results in this paper are available at <https://doi.org/10.5281/zenodo.6855783>.

Acknowledgment. The authors are grateful to Mathias Scheunert (Technische Universität Bergakademie Freiberg) for programming mesh generation for the computational examples.

³This is MATLAB's actual behavior: `MINRES(A, b, tol, maxit, M1, M2, x0)` mathematically means, for \mathbf{A} symmetric and \mathbf{P} symmetric positive definite, $\|\mathbf{r}_k\|_{\mathbf{P}^{-1}} = \min_{\mathbf{p} \in \mathcal{P}_k^0} \|\mathbf{p}(\mathbf{A}\mathbf{P}^{-1})\mathbf{r}_0\|_{\mathbf{P}^{-1}}$, $\|\mathbf{r}_k^{\mathbf{P}}\|_{\mathbf{P}} = \min_{\mathbf{p} \in \mathcal{P}_k^0} \|\mathbf{p}(\mathbf{P}^{-1}\mathbf{A})\mathbf{r}_0^{\mathbf{P}}\|_{\mathbf{P}}$, or $\|\mathbf{r}_k^{\mathbf{L}}\|_2 = \min_{\mathbf{p} \in \mathcal{P}_k^0} \|\mathbf{p}(\mathbf{L}^{-1}\mathbf{A}\mathbf{L}^{-\top})\mathbf{r}_0^{\mathbf{L}}\|_2$, where all of these are equivalent formulations through $\mathbf{r}_k = \mathbf{b} - \mathbf{A}\mathbf{x}_k$, $\mathbf{r}_k^{\mathbf{P}} = \mathbf{P}^{-1}\mathbf{r}_k$, $\mathbf{r}_k^{\mathbf{L}} = \mathbf{L}^{-1}\mathbf{r}_k$, $\mathbf{L}\mathbf{L}^{\top} = \mathbf{P}$ and preconditioner \mathbf{P} is given by $\mathbf{M1}$ and $\mathbf{M2}$ as per the function's docstring. On the other hand, the function uses the Euclidean stopping criterion $\|\mathbf{r}_k\|_2 \leq \text{tol}\|\mathbf{b}\|_2$, regardless of the preconditioner and the initial guess.

REFERENCES

- [1] M. ARIOLI AND I. S. DUFF, *Preconditioning linear least-squares problems by identifying a basis matrix*, SIAM J. Sci. Comput., 37 (2015), pp. S544–S561, <https://doi.org/10.1137/140975358>.
- [2] A. T. BARKER, T. REES, AND M. STOLL, *A fast solver for an \mathcal{H}_1 regularized PDE-constrained optimization problem*, Commun. Comput. Phys., 19 (2016), pp. 143–167, <https://doi.org/10.4208/cicp.190914.080415a>.
- [3] M. BENZI AND C. FACCIO, *Solving Linear Systems of the Form $(A + \gamma UU^T)\mathbf{x} = \mathbf{b}$ by Preconditioned Iterative Methods*, preprint, arXiv:2206.10444v1, 2022.
- [4] M. BENZI AND M. TÜMA, *A robust preconditioner with low memory requirements for large sparse least squares problems*, SIAM J. Sci. Comput., 25 (2003), pp. 499–512, <https://doi.org/10.1137/S106482750240649X>.
- [5] S. BERISHA AND J. G. NAGY, *Iterative methods for image restoration*, in Academic Press Library in Signal Processing, Vol. 4, A. S. Joel Trussell, A. K. Roy-Chowdhury, A. Srivastava, P. A. Naylor, R. Chellappa, and S. Theodoridis, eds., Academic Press, New York, 2014, pp. 193–247, <https://doi.org/10.1016/B978-0-12-396501-1.00007-8>.
- [6] Å. BJÖRCK, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, 1996, <https://doi.org/10.1137/1.9781611971484>.
- [7] Å. BJÖRCK, *Numerical Methods in Matrix Computations*, Texts in Applied Mathematics 59, Springer, Cham, Switzerland, 2015, <https://doi.org/10.1007/978-3-319-05089-8>.
- [8] Å. BJÖRCK, T. ELFVING, AND Z. STRAKOŠ, *Stability of conjugate gradient and Lanczos methods for linear least squares problems*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 720–736, <https://doi.org/10.1137/S089547989631202X>.
- [9] J. BLECHTA, *Stability of linear GMRES convergence with respect to compact perturbations*, SIAM J. Matrix Anal. Appl., 42 (2021), pp. 436–447, <https://doi.org/10.1137/20M1340848>.
- [10] J. BLECHTA AND M. SCHEUNERT, *Supporting Code for Paper Efficient Solution of Parameter Identification Problems with H^1 Regularization*, 2022, <https://doi.org/10.5281/zenodo.6855783>.
- [11] D. BOFFI, F. BREZZI, AND M. FORTIN, *Mixed Finite Element Methods and Applications*, Springer Series in Computational Mathematics 44, Springer, Cham, Switzerland, 2013, <https://doi.org/10.1007/978-3-642-36519-5>.
- [12] J. BOYLE, M. MIHAJLOVIĆ, AND J. SCOTT, *HSL_MI20: An efficient AMG preconditioner for finite element problems in 3D*, Internat. J. Numer. Methods Engr., 82 (2010), pp. 64–98, <https://doi.org/10.1002/nme.2758>.
- [13] R. BRU, J. MARÍN, J. MAS, AND M. TÜMA, *Preconditioned iterative methods for solving linear least squares problems*, SIAM J. Sci. Comput., 36 (2014), pp. 2002–A2022, <https://doi.org/10.1137/130931588>.
- [14] A. BUNSE-GERSTNER, V. GUERRA-ONES, AND H. M. DE LA VEGA, *An improved preconditioned LSQR for discrete ill-posed problems*, Math. Comput. Simul., 73 (2006), pp. 65–75, <https://doi.org/10.1016/j.matcom.2006.06.023>.
- [15] D. CALVETTI, F. PITOLLI, E. SOMERSALO, AND B. VANTAGGI, *Bayes meets Krylov: Statistically inspired preconditioners for CGLS*, SIAM Rev., 60 (2018), pp. 429–461, <https://doi.org/10.1137/15M1055061>.
- [16] D. CALVETTI AND E. SOMERSALO, *Priorconditioners for linear systems*, Inverse Problems, (2005), pp. 1397–1418, <https://doi.org/10.1088/0266-5611/21/4/014>.
- [17] N. A. CARUSO AND P. NOVATI, *Convergence analysis of LSQR for compact operator equations*, Linear Algebra Appl., 583 (2019), pp. 146–164, <https://doi.org/10.1016/j.laa.2019.08.024>.
- [18] J. CERDÁN, D. GUERRERO, J. MARÍN, AND J. MAS, *Preconditioners for rank deficient least squares problems*, J. Comput. Appl. Math., 372 (2020), 112621, <https://doi.org/10.1016/j.cam.2019.112621>.
- [19] Y. CHEN AND F. HUANG, *Spectral method approximation of flow optimal control problems with H^1 -norm state constraint*, Numer. Math. Theory Methods Appl., 10 (2017), pp. 614–648, <https://doi.org/10.4208/nmtma.2017.m1419>.
- [20] J. CHUNG AND S. GAZZOLA, *Computational Methods for Large-Scale Inverse Problems: A Survey on Hybrid Projection Methods*, preprint, arXiv:2105.07221, 2021.
- [21] J. CHUNG, S. KNEPPER, AND J. G. NAGY, *Large-scale inverse problems in imaging*, in Handbook of Mathematical Methods in Imaging, 2nd ed., O. Scherzer, ed., Springer, New York, 2015, pp. 47–90, https://doi.org/10.1007/978-1-4939-0790-8_2.
- [22] J. CHUNG AND K. PALMER, *A hybrid LSMR algorithm for large-scale Tikhonov regularization*, SIAM J. Sci. Comput., 37 (2015), pp. S562–S580, <https://doi.org/10.1137/140975024>.

- [23] D. COLTON AND R. KRESS, *Inverse Acoustic and Electromagnetic Scattering Theory*, 3rd ed., Springer, New York, 2013, <https://doi.org/10.1007/978-1-4614-4942-3>.
- [24] M. EIERMANN AND O. G. ERNST, *Geometric aspects of the theory of Krylov subspace methods*, *Acta Numer.*, 10 (2001), pp. 251–312, <https://doi.org/10.1017/S0962492901000046>.
- [25] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, 2nd ed., Numerical Mathematics and Scientific Computation, Oxford University Press, Oxford, 2014, <https://doi.org/10.1093/acprof:oso/9780199678792.001.0001>.
- [26] B. FISCHER, *Polynomial Based Iteration Methods for Symmetric Linear Systems*, Wiley-Teubner Series Advances in Numerical Mathematics, Vieweg+Teubner Verlag Wiesbaden, Wiesbaden, Germany, 1996, <https://doi.org/10.1007/978-3-663-11108-5>.
- [27] E. GEERARDYN, N. SCHLÖMER, P. PABLO, O. KOMAROV, M. S. COSTA, B. LINGNER, L. JEUB, K. BROELEMANN, D. HORSLEY, D. MITREWSKI, R. PESCHKE, AND J. T. SVEJDA, *matlab2tikz*, 2020, <https://doi.org/10.5281/zenodo.593413>.
- [28] C. GEUZAINÉ AND J.-F. REMACLE, *Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities*, *Internat. J. Numer. Methods Engrg.*, 79 (2009), pp. 1309–1331, <https://doi.org/10.1002/nme.2579>.
- [29] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 4th ed., Johns Hopkins University Press, Baltimore, 2013.
- [30] N. GOULD AND J. SCOTT, *The state-of-the-art of preconditioners for sparse linear least-squares problems*, *ACM Trans. Math. Software*, 43 (2017), 36, <https://doi.org/10.1145/3014057>.
- [31] E. HABER, *Computational Methods in Geophysical Electromagnetics*, Mathematics in Industry, SIAM, Philadelphia, 2014, <https://doi.org/10.1137/1.9781611973808>.
- [32] M. HANKE AND C. R. VOGEL, *Two-level preconditioners for regularized inverse problems I: Theory*, *Numer. Math.*, 83 (1999), pp. 385–402, <https://doi.org/10.1007/s002110050455>.
- [33] 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>.
- [34] I. C. F. IPSEN, *A note on preconditioning nonsymmetric matrices*, *SIAM J. Sci. Comput.*, 23 (2001), pp. 1050–1051, <https://doi.org/10.1137/S1064827500377435>.
- [35] M. JACOBSEN, P. C. HANSEN, AND M. A. SAUNDERS, *Subspace preconditioned LSQR for discrete ill-posed problems*, *BIT*, 43 (2003), pp. 975–989, <https://doi.org/10.1023/B:BITN.0000014547.88978.05>.
- [36] A. KLAWONN AND G. STARKE, *Block triangular preconditioners for nonsymmetric saddle point problems: Field-of-values analysis*, *Numer. Math.*, 81 (1999), pp. 577–594, <https://doi.org/10.1007/s002110050405>.
- [37] J. LIESEN AND Z. STRAKOŠ, *Krylov Subspace Methods: Principles and Analysis*, Numerical Mathematics and Scientific Computation, Oxford University Press, Oxford, 2013, <https://doi.org/10.1093/acprof:oso/9780199655410.001.0001>.
- [38] K.-A. MARDAL AND R. WINTHER, *Preconditioning discretizations of systems of partial differential equations*, *Numer. Linear Algebra Appl.*, 18 (2011), pp. 1–40, <https://doi.org/10.1002/nla.716>.
- [39] I. MORET, *A note on the superlinear convergence of GMRES*, *SIAM J. Numer. Anal.*, 34 (1997), pp. 513–516, <https://doi.org/10.1137/S0036142993259792>.
- [40] K. MORIKUNI AND K. HAYAMI, *Inner-iteration Krylov subspace methods for least squares problems*, *SIAM J. Matrix Anal. Appl.*, 34 (2013), pp. 1–22, <https://doi.org/10.1137/110828472>.
- [41] M. F. MURPHY, G. H. GOLUB, AND A. J. WATHEN, *A note on preconditioning for indefinite linear systems*, *SIAM J. Sci. Comput.*, 21 (2000), pp. 1969–1972, <https://doi.org/10.1137/S1064827599355153>.
- [42] J. NASH, *Continuity of solutions of parabolic and elliptic equations*, *Amer. J. Math.*, 80 (1958), pp. 931–954, <https://doi.org/10.2307/2372841>.
- [43] D. ORBAN AND M. ARIOLI, *Iterative Solution of Symmetric Quasi-Definite Linear Systems*, *SIAM Spotlights 3*, SIAM, Philadelphia, 2017, <https://doi.org/10.1137/1.9781611974737>.
- [44] 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>.
- [45] J. W. PEARSON, *Fast Iterative Solvers for PDE-constrained Optimization Problems*, Ph.D. thesis, University of Oxford, 2013, <https://kar.kent.ac.uk/48153/>.
- [46] C. E. POWELL AND D. SILVESTER, *Optimal preconditioning for Raviart-Thomas mixed formulation of second-order elliptic problems*, *SIAM J. Matrix Anal. Appl.*, 25 (2003), pp. 718–738, <https://doi.org/10.1137/S0895479802404428>.

- [47] C. SCHWARZBACH AND E. HABER, *Finite element based inversion for time-harmonic electromagnetic problems*, Geophys. J. Int., 193 (2013), pp. 615–634, <https://doi.org/10.1093/gji/ggt006>.
- [48] J. SCOTT, *On using Cholesky-based factorizations and regularization for solving rank-deficient sparse linear least-squares problems*, SIAM J. Sci. Comput., 39 (2017), pp. C319–C339, <https://doi.org/10.1137/16M1065380>.
- [49] J. SCOTT AND M. TUMA, *Preconditioning of linear least squares by robust incomplete factorization for implicitly held normal equations*, SIAM J. Sci. Comput., 38 (2016), pp. C603–C623, <https://doi.org/10.1137/16M105890X>.
- [50] C. B. SULLIVAN AND A. KASZYNSKI, *PyVista: 3d plotting and mesh analysis through a streamlined interface for the visualization toolkit (VTK)*, J. Open Source Softw., 4 (2019), 1450, <https://doi.org/10.21105/joss.01450>.
- [51] W. M. TELFORD, L. P. GELDART, AND R. E. SHERIFF, *Applied Geophysics*, 2nd ed., Cambridge University Press, Cambridge, 1990, <https://doi.org/10.1017/CBO9781139167932>.
- [52] S. UHLEMANN, P. B. WILKINSON, H. MAURER, F. M. WAGNER, T. C. JOHNSON, AND J. E. CHAMBERS, *Optimized survey design for electrical resistivity tomography: Combined optimization of measurement configuration and electrode placement*, Geophys. J. Int., 214 (2018), pp. 108–121, <https://doi.org/10.1093/gji/ggy128>.
- [53] C. R. VOGEL, *Computational Methods for Inverse Problems*, Frontiers in Applied Mathematics, SIAM, Philadelphia, 2002, <https://doi.org/10.1137/1.9780898717570>.
- [54] C. R. VOGEL AND M. HANKE, *Two-Level Preconditioners for Regularized Inverse Problems II: Implementation and Numerical Results*, manuscript, 1998, <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=bd1e92928fb81ca10a796d275ead399b566235db>.
- [55] R. WINTHER, *Some superlinear convergence results for the conjugate gradient method*, SIAM J. Numer. Anal., 17 (1980), pp. 14–17, <https://doi.org/10.1137/0717002>.
- [56] J. A. WORTHEN, *Inverse Problems in Mantle Convection: Models, Algorithms, and Applications*, Ph.D. thesis, University of Texas at Austin, 2012, <http://hdl.handle.net/2152/19458>.
- [57] J.-F. YIN, *Preconditioner based on the Sherman–Morrison formula for regularized least squares problems*, Appl. Math. Comput., 215 (2009), pp. 3007–3016, <https://doi.org/10.1016/j.amc.2009.09.048>.