

BAYESIAN MESH ADAPTATION FOR ESTIMATING DISTRIBUTED PARAMETERS*

DANIELA CALVETTI[†], ANNA COSMO[‡], SIMONA PEROTTO[§], AND
ERKKI SOMERSALO[†]

Abstract. The problem of estimating numerically a distributed parameter from indirect measurements arises in many applications, and in that context the choice of the discretization plays an important role. In fact, guaranteeing a certain level of accuracy of the forward model that maps the unknown to the observations may require a fine discretization, adding to the complexity of the problem and to the computational cost. On the other hand, reducing the complexity of the problem by adopting a coarser discretization may increase the modeling error and can be very detrimental for ill-posed inverse problems. To balance accuracy and complexity, we propose an adaptive algorithm for adjusting the discretization level automatically and dynamically while estimating the unknown distributed parameter by an iterative scheme. In the Bayesian paradigm, all unknowns, including the metric that defines the discretization, are modeled as random variables. Our approach couples the discretization with a Bayesian hierarchical hyperparameter that is estimated simultaneously with the unknown parameter of primary interest. The viability of the proposed algorithm, the Bayesian mesh adaptation (BMA) is assessed on two test cases: a fan-beam X-ray tomography problem and an inverse source problem for a Darcy flow model.

Key words. a posteriori estimate, inverse problems, X-ray tomography, Darcy flow

AMS subject classifications. 65N20, 65N30

DOI. 10.1137/20M1326222

1. Introduction. A wide range of applications in, e.g., medical imaging, non-destructive material testing, geophysical sounding or remote sensing, require the estimation of a spatially distributed parameter from indirect and noisy measurements. Approximating the distributed parameter in terms of a mesh-based basis provides a natural discretization scheme, particularly when the connection between the observed quantity and the unknown parameter is governed by partial differential or integral equations. An example of the former is electrical impedance tomography, where the unknown conductivity distribution inside a body is estimated from boundary values of the solution of the second order elliptic equation for the voltage potential [14], while an example of the latter is the X-ray tomography problem, where the unknown density distribution inside the body is related to the data through the Radon transform [33].

The choice of the density of the computational mesh used for approximating the unknown is the result of a trade-off between stability and resolution on the one hand, and between accuracy and complexity of the inverse problem on the other

*Submitted to the journal's Methods and Algorithms for Scientific Computing section March 19, 2020; accepted for publication (in revised form) September 18, 2020; published electronically December 14, 2020.

<https://doi.org/10.1137/20M1326222>

Funding: The work of the first author was partially supported by NSF grant DMS-1522334. The work of the third author was supported by INdAM-GNCS Project 2020. The work of the fourth author was partially supported by NSF grant DMS-1714617.

[†]Mathematics, Applied Mathematics and Statistics, Case Western Reserve University, Cleveland, OH 44106 USA (dxc57@case.edu, ejs49@case.edu).

[‡]Microsoft Azure, Milan, Italy (cosmo.anna@libero.it).

[§]Dipartimento di Matematica, MOX - Modeling and Scientific Computing, Politecnico di Milano, Milano, I-20133, Italy (simona.perotto@polimi.it).

hand. In fact, while the reduction of the degrees of freedom attained with a coarse mesh representation may improve the stability of the discretized inverse problem, this comes at the potential loss of spatial resolution. An even bigger concern, however, is the fact that a coarse mesh discretization may introduce a significant discrepancy between the presumably accurate continuous model and the discretized one. The corresponding discretization error tends to behave like correlated noise and, if not accounted for properly, may have very detrimental effects on the solution of the inverse problem which, by its ill-posed nature, is sensitive to all noise. In contrast, a fine discretization yields higher accuracy, and the associated small discretization error can be safely lumped with the measurement noise at a higher computational cost. These considerations motivate the search for adaptive meshing strategies that refine the mesh dynamically, adding nodes, and hence increasing the accuracy, where a higher resolution is needed, while coarsening the mesh in uniform regions, where a too fine discretization would only increase the computational burden and risk fitting the model to noise. For example, if the distributed parameter is a piecewise constant function, ideally a fine mesh is required only near the discontinuities. The problem with this idea is that we usually do not know a priori where the refinement is required, as it is determined by the properties of the unknown function to be estimated. Thus, the discretization of the problem is part of the inverse problem. Addressing this challenge is the topic of this article.

Mesh adaptation is a standard procedure in numerical approximation of partial differential equations (PDEs), particularly when using the finite element method (FEM) [22, 40]. In that context, there is a good qualitative understanding of where refinement is needed: the solution of a PDE is expected to change rapidly near coefficient discontinuities corresponding to material interfaces, or near singular points of the boundary of the computational domain, all of which constitute the input of the problem. To guarantee also a quantitative control on the discretization error, rigorous mathematical tools, such as a priori and a posteriori error estimators, can be derived to control a suitable norm or a functional of the discretization error [1]. Here, we take a different approach based on the Bayesian paradigm. In that framework, conditionally Gaussian hierarchical prior models have been previously proposed as a viable way of dealing with inverse problems whose solution is expected to contain singularities, such as discontinuities [10, 11, 12], with the dynamically updated prior variance acting as an indicator of the presence of these singularities. In this work, we propose a novel way of combining some classical adaptive mesh strategies in FEM analysis with a dynamical updating of hierarchical models: The prior variance of the gradient, indicating the level of credibility of the gradient estimation in the current meshing, is given the additional role of signaling how the mesh should be adapted. The proposed Bayesian mesh adaptation (BMA) algorithm is illustrated with computed examples involving inverse problems in X-ray tomography and a PDE-based inverse source problem.

2. Statement of the problem. In this work, we consider the inverse problem of estimating a distributed parameter u over a bounded domain $\Omega \subset \mathbb{R}^d$ from a finite number of noisy observations of a quantity depending on u , that is,

$$(1) \quad b = F^*(u) + \varepsilon,$$

where $b \in \mathbb{R}^m$ comprises the data, F^* is the function mapping the parameter u to the data, and $\varepsilon \in \mathbb{R}^m$ represents additive exogenous noise independent of u . While the basic idea may be extended to any dimension, we restrict the discussion here to $d = 1$

and, in particular, $d = 2$.

We consider a discretized version of (1) starting from a tessellation $\mathcal{T}_h = \{K_j\}$ of the domain by nonoverlapping intervals in $d = 1$, or triangles or quadrilaterals in $d = 2$, covering a polygon Ω_h approximating the domain Ω , so that $\bar{\Omega}_h = \cup_j \bar{K}_j$, where the subindex h refers to the mesh density, which will be discussed in detail later. For the time being, the subindex h is included as a reminder of the approximation. Furthermore, we assume that the tessellation is conforming, that is, $\bar{K}_j \cap \bar{K}_i$ is empty or coincides with a vertex (for $d = 1, 2$) or an edge (for $d = 2$) [16]. This condition excludes meshes with hanging nodes.

Given a tessellation of the domain Ω with vertices $\{x_j\}_{j=1}^{n_v}$, we assume that a distributed parameter u defined over Ω can be approximated in terms of its nodal values, i.e.,

$$(2) \quad u(x) \approx \sum_{j=1}^{n_v} u_j \psi_j(x), \quad u_j = u(x_j),$$

where $\{\psi_j\}_{j=1}^{n_v}$ is any Lagrange basis, such that $\psi_j(x_k) = \delta_{jk}$ [16].

We approximate the continuous problem (1) by a computationally feasible finite dimensional approximation and write the model

$$(3) \quad b = F(u_h, \mathcal{T}_h) + \varepsilon_h,$$

where $u_h \in \mathbb{R}^d$ is the unknown grid function in (2) with components u_j , and F is a mesh-dependent computed approximation of the continuous forward model F^* . Typically, in PDE-based inverse problems, the computation of F requires solving a finite element problem. Observe that since the discretization introduces a discretization error, the additive noise ε_h is representing not only the exogenous noise ε in the continuous model (1) but also the discretization error. The estimation of the error due to discretization has been discussed extensively in the literature (see, e.g., [24, 23, 2, 4]) and will be briefly revisited below.

In this article, we consider the selection of the discretization mesh to be part of the inverse problem, and we propose a dynamical updating scheme to iteratively estimate the pair $(u_h, \mathcal{T}_h)$ from the data and the a priori information about the unknown. Our particular focus in the computed examples is on problems where the unknown distributed parameter is piecewise constant, with the discretized approximation thus allowing a sparse representation in a properly chosen basis. The methodology, however, does not require this particular assumption to be satisfied.

The method that we advocate will be illustrated in the context of two inverse problems that we introduce in the next subsections, and the results of numerical tests will be discussed in detail in section 5.

2.1. Example 1: X-ray tomography. Consider the X-ray tomography problem of estimating a density function $u(x) \geq 0$ in a domain $\Omega \subset \mathbb{R}^2$ from the attenuation of X-rays traversing the domain [33]. The setting assumes a fan-beam geometry with several views, with each view consisting of a finite number of rays traversing the domain Ω , as illustrated schematically in Figure 1. If the k th ray R_k is parametrized by the arc length t ,

$$R_k : \quad x = c_k(t), \quad 0 \leq t \leq L_k = \text{length of the ray segment},$$

the noiseless forward model consists of integrals of the form

$$b_k = F_k^*(u) = \int_0^{L_k} u(c_k(t)) dt, \quad 1 \leq k \leq m,$$




FIG. 1. The geometry of the tomography problem. The circular target Ω is illuminated from a given number of directions (here 3) by a fan-beam source. Each view consists of a finite number of rays R_j along which the density in Ω is integrated to obtain the attenuation by the target. The middle and right images show the targets used in the numerical simulations.

and if we approximate $u(x)$ according to (2), the forward model becomes

$$b_k \approx \sum_{j=1}^{n_v} \left(\int_0^{L_k} \psi_j(c_k(t)) dt \right) u_j = \sum_{j=1}^{n_v} A(\mathcal{T}_h)_{kj} u_j = F_k(u_h, \mathcal{T}_h),$$

where $A(\mathcal{T}_h) \in \mathbb{R}^{m \times n_v}$ is a discretization-dependent matrix.

2.2. Example 2: Inverse source problem. In this problem, we consider a linear PDE model over $\Omega = (0, 1) \times (0, 1)$,

$$\Delta f = -u,$$

with mixed boundary conditions

$$\frac{\partial f}{\partial n} \Big|_{\Gamma_1} = \frac{\partial f}{\partial n} \Big|_{\Gamma_3} = 0, \quad f|_{\Gamma_2} = f|_{\Gamma_4} = 0;$$

see Figure 2 for a schematic illustration of the problem setting. The data are assumed to consist of noisy observations of the solution f at discrete observation points $x_k \in \Omega$; i.e., the noiseless observation model is

$$b_k = f(x_k) = F_k(u), \quad 1 \leq k \leq m.$$

To discretize the problem, the domain Ω is divided into triangular elements, and the source term is approximated in terms of first order (piecewise linear) basis functions $\{\psi_j\}_{j=1}^{n_v}$, while in order to control the approximation error, the solution f is represented in terms of second order (piecewise quadratic) basis functions $\{\varphi_j\}_{j=1}^{N_v}$ corresponding to free second order nodes, excluding the Dirichlet nodes on $\Gamma_2 \cup \Gamma_4$. Observe that the boundary conditions for u are not yet specified. The standard FEM discretization yields the system of equations

$$\sum_{j=1}^{N_v} \underbrace{\left(\int \nabla \varphi_k \cdot \nabla \varphi_j \right)}_{(K_h)_{kj}} f_j = \sum_{j=1}^{n_v} \underbrace{\left(\int \varphi_k \psi_j \right)}_{(M_h)_{kj}} u_j, \quad 1 \leq k \leq N_v,$$




FIG. 2. *Left: The domain $\Omega = (0, 1) \times (0, 1)$ with the boundary conditions for f , and the true underlying source u used in the numerical simulations. Right: The computed solution f corresponding to the source u . The observation points are indicated by dots.*

whose solution,

$$f_h = K_h^{-1} M_h u_h,$$

is a stable approximation of f at the second order nodal points, while the data are obtained through an interpolation matrix $P_h \in \mathbb{R}^{m \times N_v}$,

$$b = P_h K_h^{-1} M_h u_h = F(u_h, \mathcal{T}_h).$$

Observe that in both examples, the distributed parameter can be characterized by specifying the discontinuities, which allow for a sparse representation of the discretized approximation; this will be specified in detail later. In the following section, we revisit the discretized inverse problems within the framework of Bayesian hierarchical models, formulate certain sparsity-promoting hierarchical priors, and outline how these connect naturally with the mesh determination problem.

3. Hierarchical Bayesian prior models. We begin by reviewing some results concerning Bayesian hypermodels in inverse problems [11, 6, 10, 12, 8, 7, 9]. We assume that the discrete forward model is linear. An extension to nonlinear models is possible but beyond the scope of the present paper.

Consider the linear discrete inverse problem of estimating $z \in \mathbb{R}^n$ from noisy indirect observations of the form

$$(4) \quad b = Az + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \Sigma),$$

where $A \in \mathbb{R}^{m \times n}$, and the noise covariance matrix $\Sigma \in \mathbb{R}^{m \times m}$ is symmetric positive definite. We assume a priori that z is sparse, that is, $\|z\|_0 = \text{card}\{\text{supp}(z)\} \ll n$, or, more generally, compressible, that is,

$$\text{card}\{z_j \mid |z_j| > \delta\} \ll n,$$

where the threshold δ corresponds to insignificant background values of the components. If Σ is known, the additive noise can be whitened by multiplying both sides of (4) by a matrix $S \in \mathbb{R}^{m \times m}$ satisfying $S\Sigma S^T = I_m$, where I_m is the identity matrix of size $m \times m$. The matrix S can be chosen, for instance, to be the Cholesky factor

of the precision matrix Σ^{-1} . Hence, without loss of generality, we may assume that $\Sigma = \mathbf{I}_m$, leading to a likelihood model

$$\pi_{b|z}(b | z) \propto \exp\left(-\frac{1}{2}\|b - Az\|^2\right),$$

where $\propto$ means proportionality up to a normalizing factor. Define the conditionally Gaussian prior model for z ,

$$\pi_{z|\theta}(z | \theta) \propto \frac{1}{(\theta_1\theta_2\cdots\theta_n)^{1/2}} \exp\left(-\frac{1}{2}\sum_{j=1}^n \frac{z_j^2}{\theta_j}\right),$$

where the quantities $\theta_j > 0$ are the a priori variances of the mutually independent components of z . Since we believe that the unknown z is either sparse or compressible, in order for the prior to favor such solutions, most of the variances should be close to zero, with only a few being significantly large. However, since we do not know a priori which components are significantly different from zero, we model θ as a random vector. To promote the desired properties for θ , we introduce a hyperprior that favors small values for θ_j while allowing for significantly large but rare outliers. Among the distributions with such properties, we concentrate on the three-parameter family of the generalized gamma distributions, with probability density function of the form

$$\pi(\theta_j) \propto \left(\frac{\theta_j}{\vartheta_j}\right)^{r\beta-1} \exp\left(-\left(\frac{\theta_j}{\vartheta_j}\right)^r\right).$$

In our setting, the parameter $r \neq 0$ identifies a member of the generalized gamma family, and the values of the scale parameters $\vartheta_j > 0$ are chosen separately for each j , while the same shape parameter $\beta > 0$ is used for all θ_j . Setting $r = 1$ gives the gamma distribution, while $r = -1$ is the inverse gamma distribution.

Combining the likelihood, conditional prior, and hyperprior according to Bayes' formula, we obtain the posterior distribution for the pair (z, θ) ,

$$\pi_{z,\theta|b}(z, \theta | b) \propto \exp\left(-\frac{1}{2}\|b - Az\|^2 - \frac{1}{2}\sum_{j=1}^n \frac{z_j^2}{\theta_j} + \left(r\beta - \frac{3}{2}\right)\sum_{j=1}^n \log\left(\frac{\theta_j}{\vartheta_j}\right) - \sum_{j=1}^n \left(\frac{\theta_j}{\vartheta_j}\right)^r\right).$$

In this work, we consider the maximum a posteriori (MAP) estimator, defined as the maximizer of the above expression or, equivalently, minimizer of the energy functional,

$$(5) \quad (z, \theta)_{\text{MAP}} = \operatorname{argmin} \left\{ \frac{1}{2}\|b - Az\|^2 + \frac{1}{2}\sum_{j=1}^n \frac{z_j^2}{\theta_j} - \eta \sum_{j=1}^n \log\left(\frac{\theta_j}{\vartheta_j}\right) + \sum_{j=1}^n \left(\frac{\theta_j}{\vartheta_j}\right)^r \right\},$$

where $\eta = r\beta - 3/2$. An iterative algorithm for effectively computing the MAP estimate, the iterative alternating sequential algorithm (IAS), has been suggested and extensively studied in the literature; see [6, 9, 12, 7, 8]. The structure of the algorithm can be summarized as follows:

Given the parameters $r \neq 0$, β , $\vartheta \in \mathbb{R}^n$:

Initialize: Set $\theta^0 = \vartheta$, $t = 0$.

Iterate until the stopping criterion is met:

(a) Update z ,

$$z^{t+1} = \operatorname{argmin} \left\{ \frac{1}{2} \|b - Az\|^2 + \frac{1}{2} \sum_{j=1}^n \frac{z_j^2}{\theta_j^t} \right\};$$

(b) Update θ ,

$$\theta^{t+1} = \operatorname{argmin} \left\{ \frac{1}{2} \sum_{j=1}^n \frac{(z_j^{t+1})^2}{\theta_j} - \eta \sum_{j=1}^n \log \left(\frac{\theta_j}{\vartheta_j} \right) + \sum_{j=1}^n \left(\frac{\theta_j}{\vartheta_j} \right)^r \right\};$$

(c) Advance the counter by one, $t \leftarrow t + 1$.

The stopping criterion requires the relative change in θ to fall below a given threshold value. As pointed out in the cited articles, the algorithm is particularly simple to implement: The update step (a) with fixed θ requires only the solution of a linear least squares problem, while in step (b), each component θ_j can be updated separately. The first order optimality condition yields a uniquely solvable condition for each θ_j (see [6, 9]), and in [7] it was shown that for all values $r \neq 0$, the updating can be done efficiently by solving a simple ordinary differential initial value problem. For particular values of r , including $r = \pm 1$, a closed form updating formula for θ_j can be found.

It was proved in [9] that for $r = 1$, the IAS algorithm converges to a unique minimizer of the energy functional, and the convergence rate was analyzed in [12]. For other values of r , the convergence properties were further studied in [7]. The articles [12, 7] established the connection between the proposed algorithm and a number of classical sparsity-promoting regularization schemes, including the ℓ^p penalty regularization, $p \leq 1$. Furthermore, in the cited articles, it was demonstrated that if an estimate of the signal-to-noise ratio (SNR) of the problem is available, under certain natural conditions the selection of the hyperparameter ϑ can be made automatic, and ϑ is directly related to the sensitivity of the inverse problem to the different components of z [9, 12, 7]. We summarize the results relevant to the present work in the following theorems.

THEOREM 3.1. *For $r = 1$ and $\eta > 0$, the objective function (5) is strictly convex, and the IAS algorithm converges to the global minimizer $(\hat{z}, \hat{\theta})$, where $\hat{\theta}$ satisfies*

$$\hat{\theta}_j = \frac{1}{2} \vartheta_j \left(\eta + \sqrt{\eta^2 + \frac{\hat{z}_j^2}{2\vartheta_j}} \right).$$

Moreover, when $\eta \rightarrow 0+$, the minimizer $\hat{z}$ converges to the minimizer of the functional

$$(6) \quad \frac{1}{2} \|b - Az\|^2 + \sqrt{2} \sum_{j=1}^n \frac{|z_j|}{\sqrt{\vartheta_j}}.$$

For the proof, see Theorem 3.7 in [12]. From this theorem it follows that in the limit, ϑ_j represents a weight of the ℓ^1 penalty. In the geophysical and biomedical literature, the weight in the penalty is usually introduced to account for the sensitivity of the data to different components in z ; see, e.g., [29, 27, 28]. Such weighting is often introduced based on heuristic arguments, and interpreting it as a prior in the Bayesian framework is problematic, since the prior would depend on A and hence on the measurement configuration. However, the following theorem establishes a proper Bayesian interpretation.

THEOREM 3.2.

- (a) Assuming that a support set $S \subset \{1, 2, \dots, n\}$ is given, the SNR conditional to the unknown z being supported on S , denoted by SNR_S , is given by

$$\text{SNR}_S = \frac{\sum_{j \in S} \nu(r, \beta) \vartheta_j}{\text{trace}(\Sigma)} + 1, \quad \nu(r, \beta) = \frac{\Gamma(\beta + 1/r)}{\Gamma(\beta)},$$

provided that $\beta > -1/r$.

- (b) Let $p_k = \mathbb{P}\{\|z\|_0 = k\}$ denote the probability that the support of the signal has cardinality k for $k = 1, 2, \dots, n$. Then the exchangeability condition ($\mathcal{E}$),

$$(\mathcal{E}) : \quad \text{SNR}_S = \text{SNR}_{S'} \text{ whenever } S \text{ and } S' \text{ are of the same cardinality,}$$

is satisfied if and only if ϑ_j is chosen as

$$\vartheta_j = \frac{C}{\|a^{(j)}\|^2}, \quad C = \frac{(\overline{\text{SNR}} - 1)\text{trace}(\Sigma)}{\nu(r, \beta)} \sum_{k=1}^n \frac{p_k}{k},$$

where $a^{(j)} = \mathbf{A}e_j$ is the j th column of the matrix $\mathbf{A}$.

For the proof, see Lemma 2.1 and Theorem 3.3 in [12]. We point out that since $\|\mathbf{A}e_j\|$ is the classical sensitivity of the data to the j th component of z , by choosing the weights ϑ_j as in the previous theorem, it turns out that the components z_j in the penalty are weighted as suggested in the literature, and the weighting is justified by a Bayesian argument rather than heuristics. The explanation for how the matrix $\mathbf{A}$ can enter the prior is that the concept of SNR implicitly carries information about the measuring configuration.

Now we address how to solve the least squares problem (a) in the IAS algorithm. The solution of the linear least squares problem can be restated as

$$z = \operatorname{argmin} \left\{ \left\| \begin{bmatrix} \mathbf{A} \\ \mathbf{D}_\theta^{-1/2} \end{bmatrix} z - \begin{bmatrix} b \\ 0 \end{bmatrix} \right\| \right\} = \operatorname{argmin} \left\{ \left\| \begin{bmatrix} \mathbf{A}\mathbf{D}_\theta^{1/2} \\ \mathbf{I}_n \end{bmatrix} (\mathbf{D}_\theta^{-1/2}z) - \begin{bmatrix} b \\ 0 \end{bmatrix} \right\| \right\},$$

where

$$\mathbf{D}_\theta = \operatorname{diag}(\theta) \in \mathbb{R}^{n \times n}, \quad \theta = \theta^t,$$

yielding a least squares problem that is formally equivalent to the standard Tikhonov regularized problem for the system

$$\mathbf{A}\mathbf{D}_\theta^{1/2}w = b, \quad \mathbf{D}_\theta^{-1/2}z = w,$$

with regularization parameter equal to one. Inspired by this observation, authors have suggested replacing the exact solution of the quadratic exact minimization problem (a) by an approximate one based on the regularization properties of the Krylov subspace iterations [10, 7]. In the approximate IAS algorithm, the update z^{t+1} is the *reduced Krylov subspace* (RKS) solution, defined as follows. Introduce the notation $\mathbf{A}_\theta = \mathbf{A}\mathbf{D}_\theta^{1/2}$, where $\theta = \theta^t$, and consider the sequence of approximate solutions by the conjugate gradient method for least squares (CGLS), defined as

$$(7) \quad w^{(k)} = \operatorname{argmin} \{ \|b - \mathbf{A}_\theta w\| \mid w \in \mathcal{K}_k(\mathbf{A}_\theta^\top b, \mathbf{A}_\theta^\top \mathbf{A}_\theta) \}, \quad z^{(k)} = \mathbf{D}_\theta^{1/2}w^{(k)},$$

where

$$\mathcal{K}_k(\mathbf{A}_\theta^\top b, \mathbf{A}_\theta^\top \mathbf{A}_\theta) = \operatorname{span} \{ (\mathbf{A}_\theta^\top \mathbf{A}_\theta)^\ell \mathbf{A}_\theta^\top b \mid 0 \leq \ell \leq k-1 \}$$

is the k th Krylov subspace associated with the vector $A_\theta^\top b$ and the matrix $A_\theta^\top A_\theta$. The quantity $b - A_\theta w^{(k)}$ whose norm is minimized is the discrepancy vector corresponding to $w^{(k)}$. The RKS solution is the k^* th iterate,

$$z^{t+1} = D_\theta^{1/2} w^{t+1}, \quad w^{t+1} = w^{(k^*)},$$

with k^* chosen to be the first index k satisfying the criterion

$$(\mathcal{E}) : \|b - A_\theta w^{(k+1)}\| \leq \sqrt{m} \quad \text{or} \quad G(w^{(k+1)}) > \tau G(w^{(k)}),$$

where $\tau - 1 = \epsilon > 0$ is a small safeguard parameter, and G is the norm of the discrepancy of the original linear system,

$$(8) \quad G(w) = \left\| \begin{bmatrix} AD_\theta^{1/2} \\ I_n \end{bmatrix} w - \begin{bmatrix} b \\ 0 \end{bmatrix} \right\|^2 = \|b - AD_\theta^{1/2} w\|^2 + \|w\|^2.$$

Extensive computed examples reported in the cited articles have shown that the approximate IAS solutions are consistently similar to those obtained with the exact IAS, at a much lower computational cost.

3.1. Sparse increments. In this article we focus on inverse problems of estimating piecewise constant distributed parameters with sparse gradients. In order to take advantage of the sparsity promotion of the IAS algorithm, the minimization problem must be expressed in terms of a discrete approximation of the function increments rather than the function itself. We denote by $u = u_h$ the discretized grid function, suppressing the subindex h to keep notation simpler. As a motivation, start by considering a one-dimensional problem over a unit interval. Denoting by $u(x)$ the unknown distributed parameter, $0 \leq x \leq 1$, with discretization $u_j = u(jh)$, $h = 1/n$, being the discretization parameter, we may write u_j in terms of the increments, i.e.,

$$u_j = u_0 + \sum_{k=1}^j (u_k - u_{k-1}), \quad 1 \leq j \leq n.$$

Assuming for simplicity that $u_0 = 0$, and defining $z_j = u_j - u_{j-1}$, we have

$$Lu = z, \quad L = \begin{bmatrix} 1 & & & \\ -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \end{bmatrix}.$$

Since the matrix L is invertible, we may reformulate the problem as estimating z from the observation model

$$b = AL^{-1}u + \varepsilon,$$

with the a priori belief that the solution is sparse. Accordingly, we modify the IAS update of the unknown in terms of the increment vector, namely,

$$z^{t+1} = \operatorname{argmin} \left\{ \frac{1}{2} \|b - AL^{-1}z\|^2 + \frac{1}{2} \sum_{j=1}^n \frac{z_j^2}{\theta_j^t} \right\}, \quad u^{t+1} = L^{-1}z^{t+1},$$




FIG. 3. Schematic illustration of the circulation condition $Mz = 0$. In this figure, the black dots indicate the free nodes, and the green squares are bound nodes, or Dirichlet nodes, where the grid function is assumed to vanish. Free edges, drawn as solid lines, have at least one free node as an endpoint, while bound edges, drawn as dashed lines, connect bound nodes. Bound edges give a zero contribution to the circulation. If z is the vector of increments corresponding to a grid function u , the sums of increments around edges of each element with a given orientation (here clockwise) must vanish.

followed by the update step of θ ,

$$\theta^{t+1} = \operatorname{argmin} \left\{ \frac{1}{2} \sum_{j=1}^n \frac{(z_j^{t+1})^2}{\theta_j} - \eta \sum_{j=1}^n \log \left(\frac{\theta_j}{\vartheta_j} \right) + \sum_{j=1}^n \left(\frac{\theta_j}{\vartheta_j} \right)^r \right\}.$$

A different approach is required in dimension $d \geq 2$ because the finite difference matrix is no longer square, and there is no one-to-one correspondence between x and z .

Consider a triangular mesh as shown in Figure 3, and assume that the values of the unknown function are known to vanish at some of the nodes, referred to as bound nodes (or Dirichlet nodes), whereas the remaining ones are referred to as free nodes. In order to keep the notation simple, let n_v denote from now on the number of free nodes, and n_e the number of free edges $e_\ell = \{v_{\ell_1}, v_{\ell_2}\} \in E_{\text{free}}$, for which at least one endpoint is a free node. Let $L \in \mathbb{R}^{n_e \times n_v}$ denote the matrix mapping the nodal values u at free nodes to increments z along free edges,

$$(9) \quad z = Lu.$$

Note that the increments along edges between bound nodes vanish. Since the nodal values at bound nodes vanish, it is straightforward to show that L has the null space $\mathcal{N}(L) = \{0\}$. Let n_t denote the number of all triangular elements in the mesh, and let $c \in \mathbb{R}^{n_t}$ be a vector containing the sums of the increments in a clockwise direction over the edges of each element. We denote by $M \in \mathbb{R}^{n_e \times n_t}$ the matrix of entries 0 or ± 1 , such that

$$(10) \quad c = Mz.$$

If z corresponds to a grid function u through (9), the circulations must vanish, that is, $z \in \mathcal{N}(M)$. Hence, the matrices L and M define a short exact chain,

$$\{0\} \longrightarrow \mathbb{R}^{n_v} \xrightarrow{L} \mathbb{R}^{n_e} \xrightarrow{M} \mathbb{R}^{n_t} \longrightarrow \{0\}.$$

To define a prior for the increments analogously to the one-dimensional case, we must secure that the prior is concentrated on the subspace $\mathcal{N}(\mathbf{M})$. To do this, let

$$(11) \quad \mathbf{L}_\theta = \mathbf{D}_\theta^{-1/2} \mathbf{L}, \quad \text{where} \quad \mathbf{D}_\theta = \text{diag}(\theta_1, \dots, \theta_{n_e}).$$

Given a grid function u , we introduce the auxiliary variable

$$(12) \quad \beta = \mathbf{L}_\theta u.$$

Observe that

$$\sum_{j=1}^{n_e} \frac{(\mathbf{L}u)_j^2}{\theta_j} = \|\beta\|^2.$$

In order that β corresponds to a grid function, we must require that $\beta \in \mathcal{H} = \mathcal{R}(\mathbf{L}_\theta)$. If this is the case, u can be solved through the pseudoinverse, i.e., $u = \mathbf{L}_\theta^\dagger \beta$. We therefore define the conditional density for β , setting

$$(13) \quad \pi_{\beta|\theta,b}(\beta | \theta, b) \propto \exp \left(-\frac{1}{2} \|b - \mathbf{A} \mathbf{L}_\theta^\dagger \beta\|^2 - \frac{1}{2} \|\beta\|^2 \right) \delta_{\mathcal{H}}(\beta),$$

where $\delta_{\mathcal{H}}(\beta)$ is the singular measure concentrated on the subspace $\mathcal{H}$. It was shown in [8] that for a fixed θ , the minimizer of the energy functional,

$$\beta = \operatorname{argmin} \left\{ \|b - \mathbf{A} \mathbf{L}_\theta^\dagger \beta\|^2 + \|\beta\|^2 \right\},$$

is automatically in the subspace $\mathcal{H}$, i.e., the compatibility condition need not be enforced. This observation leads to the following version of the IAS algorithm for the increments:

Given the parameters $r \neq 0$, $\beta, \vartheta \in \mathbb{R}^n$:

Initialize: Set $\theta^0 = \vartheta$, $t = 0$.

Iterate until convergence:

(a) Update u ,

$$u^{t+1} = \mathbf{L}_{\theta^t}^\dagger \beta^{t+1}, \quad \text{where} \quad \beta^{t+1} = \operatorname{argmin} \left\{ \|b - \mathbf{A} \mathbf{L}_{\theta^t}^\dagger \beta\|^2 + \|\beta\|^2 \right\}.$$

(b) Update θ ,

$$\theta^{t+1} = \operatorname{argmin} \left\{ \frac{1}{2} \sum_{j=1}^{n_e} \frac{(\mathbf{L}u^{t+1})_j^2}{\theta_j} - \eta \sum_{j=1}^{n_e} \log \left(\frac{\theta_j}{\vartheta_j} \right) + \sum_{j=1}^{n_e} \left(\frac{\theta_j}{\vartheta_j} \right)^r \right\};$$

(c) Advance the counter by one, $t \leftarrow t + 1$.

In order to update u , we need to solve in the least squares sense the linear system

$$(14) \quad \begin{bmatrix} \mathbf{A} \mathbf{L}_\theta^\dagger \\ \mathbf{I}_{n_e} \end{bmatrix} \beta = \begin{bmatrix} b \\ 0 \end{bmatrix}, \quad \theta = \theta^t.$$

Alternatively, as proposed in [7], we can consider the reduced problem

$$\mathbf{A} \mathbf{L}_\theta^\dagger \beta = b,$$

and as in the previous section, compute its approximate RKS solution by the CGLS algorithm with an appropriate stopping criterion. More precisely, if

$$\mathcal{K}_k = \mathcal{K}_k \left((\mathbf{A}\mathbf{L}_\theta^\dagger)^\top b, (\mathbf{A}\mathbf{L}_\theta^\dagger)^\top \mathbf{A}\mathbf{L}_\theta^\dagger \right) = \text{span} \left\{ \left[(\mathbf{A}\mathbf{L}_\theta^\dagger)^\top \mathbf{A}\mathbf{L}_\theta^\dagger \right]^j (\mathbf{A}\mathbf{L}_\theta^\dagger)^\top b \mid j = 0, 1, \dots, k-1 \right\}$$

defines the Krylov subspace of dimension k associated with the vector $(\mathbf{A}\mathbf{L}_\theta^\dagger)^\top b$ and the matrix $(\mathbf{A}\mathbf{L}_\theta^\dagger)^\top \mathbf{A}\mathbf{L}_\theta^\dagger$, we approximate the solution of (14) by the k th iterate of the CGLS algorithm,

$$\beta^{t+1} = \beta^{(k^*)} = \text{argmin} \left\{ \|b - \mathbf{A}\mathbf{L}_\theta^\dagger \beta\|, \beta \in \mathcal{K}_{k^*} \right\}, \quad \theta = \theta^t,$$

where k^* is the smallest index that satisfies the discrepancy criterion

$$(15) \quad (\mathcal{E}) : \|b - \mathbf{A}\mathbf{L}_\theta^\dagger \beta^{(k+1)}\| \leq \sqrt{m} \quad \text{or} \quad G(\beta^{(k+1)}) > \tau G(\beta^{(k)}),$$

with

$$G(\beta) = \|b - \mathbf{A}\mathbf{L}_\theta^\dagger \beta\|^2 + \|\beta\|^2.$$

Furthermore, it was shown in [7] that it is possible to add a positivity constraint to the IAS algorithm by a proper projection on the positive cone.

In the CGLS algorithm, the multiplication of a vector z with $\mathbf{A}\mathbf{L}_\theta^\dagger$ can be evaluated by first solving the linear system

$$\mathbf{L}_\theta \alpha = z$$

in the least squares sense, and then multiplying the solution α by $\mathbf{A}$. To implement the multiplication of b by $(\mathbf{A}\mathbf{L}_\theta^\dagger)^\top = (\mathbf{L}_\theta^\dagger)^\top \mathbf{A}^\top$, we recall the definitions of the pseudoinverse $\mathbf{L}_\theta^\dagger$ and its transpose,

$$\mathbf{L}_\theta^\dagger = (\mathbf{L}_\theta^\top \mathbf{L}_\theta)^{-1} \mathbf{L}_\theta^\top, \quad (\mathbf{L}_\theta^\dagger)^\top = \mathbf{L}_\theta (\mathbf{L}_\theta^\top \mathbf{L}_\theta)^{-1},$$

where the matrix $\mathbf{L}_\theta^\top \mathbf{L}_\theta$ in our application is very sparse. Therefore, for the multiplication with the transpose, we solve first

$$(\mathbf{L}_\theta^\top \mathbf{L}_\theta) w = \mathbf{A}^\top b,$$

for w , then multiply w by $\mathbf{L}_\theta$.

3.2. Modeling error. As part of the proposed algorithm that will be introduced in the next section, we need to solve iteratively the discretized inverse problem with a given discretization. As pointed out earlier, the discretization of a continuous problem introduces a discretization error that cannot be ignored when solving inverse problems unless the discretization is so fine that the induced error is dominated by the exogenous error in the data. The ways in which approximation errors are addressed in classical FEM applications and in inverse problems have some fundamental differences. In the former, the approximation error is the discrepancy between the true solution of a PDE and the computed approximation, and sophisticated methods for controlling the accuracy through mesh adaptation have been developed. In inverse problems, the intrinsic ill-posedness adds an extra layer of complexity, as small inaccuracies in the forward model, if not properly addressed, may propagate as huge errors in the computed solution. We briefly review here how discretization errors or, more generally, any modeling errors induced by replacing the exact model by an approximation, can be addressed in the Bayesian framework.

Given the exact model (1) and the corresponding discrete approximation (3), we write

$$\begin{aligned} b &= F^*(u) + \varepsilon \\ &= F(u_h, \mathcal{T}_h) + [F^*(u) - F(u_h, \mathcal{T}_h)] + \varepsilon, \end{aligned}$$

expressing the modeling error due to the discretization as

$$(16) \quad m_h = m(u, \mathcal{T}_h) = F^*(u) - F(u_h, \mathcal{T}_h) = \mathcal{F}_h(u).$$

The problem with addressing the modeling error is that it depends on the solution of the inverse problem, which is unknown. However, in the Bayesian framework, the prior distribution π_u of the unknown can be used to write an approximate model for the probability distribution of the modeling error. From formula (16), it follows that the a priori probability distribution of the modeling error can be expressed as the push-forward of the prior distribution π_u for u ,

$$(17) \quad \pi_{m_h} = (\mathcal{F}_h)_\# \pi_u.$$

In practice, the above formula leads to a useful algorithm through sampling: By generating an ensemble of independent realizations $\{u^{(1)}, \dots, u^{(K)}\}$ of the unknown drawn from the prior density π_u , a sample of realizations of the modeling error distributed according to π_{m_h} given by (17) can be generated by setting

$$m_h^{(j)} = \mathcal{F}_h(u^{(j)}), \quad 1 \leq j \leq K,$$

thus allowing an approximation of the distribution itself by using, e.g., a Gaussian approximation or an approximation by Gaussian mixtures. This approach requires, however, that the exact forward model F^* can be implemented, which usually is a challenge. In practice, a numerically accurate high-precision surrogate, such as a FEM model with a dense mesh, is used, and the time-consuming generation of the modeling error sample is performed off-line prior to the solution of the inverse problem. This idea, originally proposed in [24, 23], was further developed and applied to a variety of inverse problems (see, e.g., [2, 34, 35]) and generalized in [5, 4] through an iterative updating process taking into account the increasing information on u as the inverse problem is solved. In this paper, to contain the computational burden, we adopt a simplified approach and approximate only the order of magnitude of the modeling error in a straightforward manner, which will be described in section 5.

We point out that the modeling error related to reduced models is a recurrent topic in Bayesian inverse problems and sampling algorithms. In [38, 39], the authors consider approximate surrogate models in Markov chain Monte Carlo sampling, introducing a multifidelity correction to compensate for the approximation. In [13, 15], the error induced by an approximate model is controlled by a Bayes factor estimate, ensuring that the difference between the posterior density and its approximation remains insignificant. In [26, 17], the model reduction and the posterior estimation are coupled so that the reduced model does not significantly perturb the posterior estimation.

4. Mesh adaptation strategy. In classical finite element methods, mesh adaptation is a well-established tool for improving the accuracy and efficiency of a discretization scheme. The goal pursued by a mesh adaptation procedure is to minimize the number of mesh elements for a fixed accuracy of the discrete solution or, vice

versa, to maximize such accuracy for a given number of tiles in the mesh. The alternative to an adapted mesh is a uniform tessellation of the domain Ω_h , with a sufficiently small discretization size so that the complex features of the solution can be correctly identified. This last choice is clearly not optimal from a computational viewpoint, particularly in the presence of problems which are highly heterogeneous in space and/or in time.

The actual challenge when generating an adapted mesh is identifying the areas of the domain which deserve to be refined or coarsened. This goal can be pursued by resorting to a preliminary knowledge of the phenomenon at hand (if any), to a heuristic criterion (e.g., by tracking the gradient or the Hessian of the discrete solution), or to theoretically sound mathematical tools which are known in the literature as *a priori* or *a posteriori* error estimators [1]. The former are expressed in terms of the exact solution (e.g., via an interpolation error control); the latter involve the discrete solution and thus provide practical (i.e., directly computable) information. Among the several *a posteriori* error estimators available in the literature, we focus on the recovery-based error estimators first proposed by Zienkiewicz and Zhu [41]. This methodology inspired us to propose a new mesh adaptation for inverse problems.

In the following, u represents a solution of a PDE over a domain Ω , and u_h is the FEM approximation. According to a recovery-based approach [41, 43], the mesh adaptation is driven by a control of the H^1 -seminorm of the discretization error,

$$(18) \quad |e_h|_{H^1(\Omega)}^2 = \int_{\Omega} |\nabla u(x) - \nabla u_h(x)|^2 dx = \sum_{K_j \in \mathcal{T}_h} \int_{K_j} |\nabla u(x) - \nabla u_h(x)|^2 dx,$$

where for simplicity we have assumed $\Omega \equiv \Omega_h$. Since the gradient of the exact solution is not known, we need to replace ∇u in (18) by a computable quantity, $\nabla^* u$, known as *recovered gradient* [42]. This yields an *a posteriori* error control for e_h , given by

$$(19) \quad |e_h|_{H^1(\Omega)}^2 \approx \sum_{K_j \in \mathcal{T}_h} \int_{K_j} |\nabla^* u(x) - \nabla u_h(x)|^2 dx = \sum_{K_j \in \mathcal{T}_h} \eta_{K_j}^2.$$

Several procedures for recovering the gradient have been proposed in the literature, each yielding a different estimator in (19). In the pioneering work [41], the authors proposed an estimator based on either a least squares process or an averaging step. Independently of the adopted procedure, at first one computes the values $\nabla^* u(x_i)$ at certain points, $\{x_i\}_{i=1}^{m_r}$, known as *recovery points*, starting from the values assumed by ∇u_h on a set of *sampling points*, $\{s_j\}_{j=1}^{n_s}$, located in a patch associated with x_i . Subsequently, $\nabla^* u$ is defined over elements by a piecewise polynomial interpolation of the recovered values $\nabla^* u(x_i)$. If the recovery points x_i are chosen to be vertices of the mesh and the sampling points s_j are the triangle barycenters, the averaging method seeks a recovered gradient such that

$$(20) \quad \nabla^* u(x_i) = \frac{1}{W^i} \sum_{j=1}^{n_s} w_j^i \nabla u_h(s_j) \quad \text{with } i = 1, \dots, m_r,$$

where w_j^i are suitable weights, and $W^i = \sum_{j=1}^{n_s} w_j^i$. The values of the weights w_j^i can be set equal to the area, $|K_j|$, of the elements K_j in $\mathcal{T}_h$ containing x_i as a vertex, with j running over the triangles constituting the patch, $\Delta_{x_i} = \{K \in \mathcal{T}_h : x_i \in K\}$, associated with the vertex x_i , so that $W^i = |\Delta_{x_i}|$ is the area of the patch; for common choices in practical applications, see, e.g., [42, 25, 36, 3, 32, 37]. We remark that some

attention has to be paid in the definition of the recovered gradient when recovery points belong to the domain boundary (for details, see, e.g., [42]).

Once the recovered gradient has been computed, we can obtain the contribution of each element to the total a posteriori error estimate. According to the *equidistribution principle*, an ideal mesh is such that the error contribution is equally distributed among the elements. To attain this goal, define the scaled local error estimator as

$$\hat{\eta}_{K_j}^2 = \frac{1}{|K_j|} \eta_{K_j}^2,$$

which is thus defined to even out the effect of different sizes of the elements. If N_h is the cardinality of $\mathcal{T}_h$, the equidistribution principle requires that

$$\eta_{K_j}^2 = |K_j| \hat{\eta}_{K_j}^2 = \frac{\tau^2}{N_h},$$

where $\tau > 0$ is the accuracy required for the discrete solution. The value τ thus controls the overall accuracy of the FEM approximation and is not of importance for the goals of this paper. Assuming for simplicity that the mesh is isotropic, i.e., all triangles are approximately equilateral, and each element is a shifted, rotated, and scaled version of the reference element $\hat{K}$, an equilateral triangle with unit diameter, we have $|K_j| = h_j^2 |\hat{K}|$, where the diameter h_j of the element K_j defines the mesh density parameter $h = \min\{h_j = \text{diam}(K_j) \mid K_j \in \mathcal{T}_h\}$. Then the equidistribution principle can be translated into a condition for the diameter, i.e.,

$$(21) \quad h_j^2 = \frac{\tau^2}{N_h |\hat{K}| \hat{\eta}_{K_j}^2}.$$

These target values can be used to drive the mesh refinement and coarsening. In particular, as detailed in the next section, the new adapted mesh can be described in terms of a metric defined in Ω , providing us with a useful starting point for inverse problems.

4.1. Meshing and metric. Motivated by the ideas behind the mesh adaptation strategies in the classical FEM literature, we establish a connection between Bayesian hypermodels and the metric guiding the mesh adaptation algorithms. Let $\hat{K}$ denote the reference triangle defined in the previous section, and let K be a generic triangle in the tessellation. We define an affine mapping

$$F_K : \hat{K} \rightarrow K, \quad x \mapsto s_K + F_K x,$$

where $F_K \in \mathbb{R}^{2 \times 2}$ is the Jacobian of the map, and $s_K \in \mathbb{R}^2$ is the shift parameter. Let

$$F_K = P_K U_K$$

be the polar decomposition of F_K , where U_K is an orthogonal matrix, and P_K is a symmetric positive definite matrix whose eigenvalues, $0 < \lambda_{K,1} \leq \lambda_{K,2}$, are the singular values of F_K .

We associate a piecewise constant metric in Ω_h with a tessellation $\{K_j\}$ in terms of a piecewise constant metric tensor $G = [g_{ij}]$, such that

$$(22) \quad G|_{K_j} = P_{K_j}^{-2}.$$

To compute the diameter of the elements with respect to this metric, consider a line segment across the reference triangle, parametrized with respect to the arc length, $\xi = \xi(t) \in \hat{K}$, and the corresponding parametrized curve in K_j ,

$$\gamma(t) = s_{K_j} + \mathbf{F}_{K_j} \xi(t).$$

In the metric $\mathbf{G}$ defined by (22), the arc length element is given by

$$ds = (\gamma'(t)^T \mathbf{G}|_{K_j} \gamma'(t))^{1/2} dt = (\xi'(t)^T \mathbf{U}_{K_j}^T \mathbf{P}_{K_j} \mathbf{G}|_{K_j} \mathbf{P}_{K_j} \mathbf{U}_{K_j} \xi'(t))^{1/2} dt = \|\xi'(t)\| dt,$$

because of the orthogonality of $\mathbf{U}_{K_j}$; hence in this metric the elements K_j have the same diameter.

Conversely, given a metric tensor $\mathbf{G}$ on Ω , we can look for a tessellation so that the elements have approximately the same size with respect to this metric. The meshing problem can be formally written as follows.

PROBLEM 4.1. *Given a metric $\mathbf{G}$ in Ω and $0 < h_{\min} < h_{\max}$, find a tessellation $\{K_j\}$ of Ω such that $h_{\min} \leq \text{diam}(K_j) \leq h_{\max}$ by minimizing the discrepancy*

$$\max_j \|\mathbf{G}|_{K_j} - \mathbf{P}_{K_j}^{-2}\|_{\infty}$$

over the tessellation.

Observe that if for an element K , $\mathbf{P}_K = \lambda \mathbf{I}_2$, then the element K is, up to a translation and an orthogonal transformation, a scaled copy of $\hat{K}$, and its diameter is equal to λ . Consequently, if $h = h(x)$ is a given function defining the target distance between the vertices in the mesh and $x \in \Omega$, to find an isotropic nonhomogeneous mesh corresponding to $h(x)$, we define the metric as

$$\mathbf{G}(x) = \frac{1}{h^2(x)} \mathbf{I}_2.$$

In the next section, we propose an automatic mesh adaptation algorithm for inverse problems, where the current approximation is used, in a predictive way, to generate the function $h(x)$, and hence we have the metric yielding the new discretization.

4.2. Bayesian mesh adaptation (BMA). In this subsection we customize the mesh adaptation procedure to inverse problems, using the hierarchical IAS algorithm as a tool to convey the information for the mesh generation. The driving idea of this approach is the interpretation of large estimated variance of the increment of the solution along an edge as an indication of a potential jump in the solution that may require a finer discretization, while a small variance allows larger elements and thus reduces the complexity of the problem.

Consider the discretized inverse problem

$$(23) \quad b = \mathbf{A}_h u_h + \varepsilon_h,$$

where $\mathbf{A}_h$ is the discretized forward matrix based on the current tessellation $\mathcal{T}_h$, identified with the triplet $(\mathbf{V}_h, \mathbf{E}_h, \mathbf{T}_h)$ of matrices, where $\mathbf{V}_h \in \mathbb{R}^{2 \times \bar{n}_v}$ contains all $\bar{n}_v$ vertex coordinates, including the Dirichlet nodes, $\mathbf{E}_h \in \mathbb{N}^{\bar{n}_e \times 2}$ is the edge list of all $\bar{n}_e$ edges with pointers to the endpoint vertices, and $\mathbf{T}_h \in \mathbb{N}^{n_t \times 3}$ is the element topology matrix with pointers to the vertex nodes. Here u_h is the vector of the nodal values of the unknown function at the free nodes of the mesh, and ε_h is the noise vector with

zero mean and covariance Σ_h estimated from the approximation error model. Let L_h denote the full-rank matrix mapping the nodal values to the increments along the free edges.

Given the current model (23), we run the IAS algorithm for sparse increments, compute the current approximation $(x_h, \theta_h) \in \mathbb{R}^{n_v} \times \mathbb{R}^{n_e}$ for the free nodes and edges, and extend θ_h by zero to the boundary edges, i.e., the edges between the Dirichlet nodes. Observe that if $(\theta_h)_j$ is small, we expect the increment of the solution x_h along the j th edge to be small. In other words, a small $(\theta_h)_j$ is taken as an indication of unlikely discontinuity of the unknown along that edge, thus suggesting that there is no need to refine the mesh, or even that a coarser mesh would suffice. Conversely, a large value of $(\theta_h)_j$ indicates the potential for a large increment, and to better localize the point of discontinuity along the edge, a refinement is desirable. This heuristics suggests that we can use θ_h as a proxy for the metric, after interpolating the edge-bound values of θ_h to Ω , giving rise to the following mesh adaptation algorithm.

Algorithm 1: Variance-based meshing (VBM).

Input: Triplet (V_h, E_h, T_h) , a variance vector θ_h with cardinality equal to the number of edges.

- (a) Add an auxiliary vertex at the midpoint of each edge,

$$\widehat{v}_\ell = \frac{1}{2} [V_h(:, E(\ell, 1)) + V_h(:, E(\ell, 2))], \quad 1 \leq \ell \leq \overline{n}_e.$$

Denote by $\widehat{V}_h \in \mathbb{R}^{2 \times n_e}$ the matrix of all auxiliary nodes.

- (b) Generate an auxiliary triplet $\widehat{\mathcal{T}}_h = (\widehat{V}_h, \widehat{E}_h, \widehat{T}_h)$ by using Delaunay triangularization [20].
 (c) Define the metric

$$G(x) = \max\{\widehat{\theta}_h(x), \theta_*\} I_2,$$

where $\widehat{\theta}_h(x)$ is the piecewise linear extension of the nodal values θ_h in the auxiliary mesh, and $\theta_* > 0$ is a threshold value introduced to avoid excessively large triangles.

- (d) Given the metric, find a new triplet $(V_{h'}, E_{h'}, T_{h'})$.
-

We point out that in the above algorithm, slight adjustments of the nodes $\widehat{V}_h$ at the boundary of the domain may be necessary to ensure that the resulting polygonal domain properly approximates the domain Ω . In step (d), we generate the new mesh using the meshing tool included in the **FreeFem++** package [21]. The algorithm allows a nodal input of the metric in the current mesh, which in this case is the midpoint mesh constructed above, and requires the specification of a range $[h_{\min}, h_{\max}]$ for the mesh size, providing a way to control the complexity and the accuracy of the forward map.

We may now proceed iteratively by repeatedly recalculating the forward model, estimating the modeling error, and applying the IAS algorithm to the new discretized problem. We summarize how to organize the calculations in the following algorithm.

Algorithm 2: Bayesian mesh adaptation with IAS (BMA-IAS).

Input: The data $b \in \mathbb{R}^m$, the noise covariance matrix $\Sigma \in \mathbb{R}^{m \times m}$, the initial mesh parameter h .

Initialize: Generate the initial triplet (V_h, E_h, T_h) .

Iterate until stopping criterion is satisfied:

- (a) Compute the forward matrices A_h and L_h on the current mesh.
- (b) Estimate the modeling error, compute the covariance matrix Σ_h . Whiten the noise ε_h .
- (c) Compute the sensitivity vector corresponding to $A_h L_h^\dagger$; compute ϑ .
- (d) Initialize the IAS iteration, setting $\theta = \vartheta$.
- (e) **Run the IAS** algorithm with sparse increments by repeating until convergence:
 - (i) Update $\beta_h \rightarrow u_h$ by CGLS with early stopping criterion;
 - (ii) Update θ_h .
- (d) **Update the mesh** using Algorithm 1.

end

The IAS iteration can be augmented with the projection of the solution u_h to the positive cone if the parameter u is a priori known to be positive; see [7] for details. Above, the stopping criterion for the IAS algorithm is given by requiring that the relative change of the variance parameter θ is below a threshold value,

$$\frac{\|\theta^{t+1} - \theta^t\|_\infty}{\|\theta^t\|_\infty} < \delta,$$

where we use the value $\delta = 0.1$ for the tolerance. The stopping criterion for the mesh update, as well as the choice of parameters defining the mesh density, is discussed in light of computed examples, where preliminary versions of a systematic criterion are considered.

5. Computed examples. In this section we demonstrate the viability of Algorithm 2 by applying it to the linear inverse problems described in sections 2.1 and 2.2.

5.1. Tomography problem. In this example, we consider the X-ray tomography problem sketched in Figure 1. In the first test, we assume that the target is illuminated from several directions, choosing 15 equally spaced view directions, each view consisting of 300 rays, and thus the total dimensionality of the data is $m = 15 \times 300 = 4500$. The generative model used to compute the noiseless data, shown in Figure 1, consists of a unit disc Ω of zero background density with a circular inclusion with center at $(0, 0.5)$, radius 0.3, and density $u = 1$. We compute the noiseless data by solving analytically for each ray the length of the line segment inside the inclusion, thus avoiding any error due to discretization. In this experiment, the exogenous additive noise is white noise scaled by σ_{ex} equal to 0.5% of the maximum of the entries of the noiseless data.

To find a computationally operational model for the discretization error, let b_0 denote the exact noiseless data, and assume that the observation model is of the form

$$b = Au + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma_{\text{ex}}^2 \mathbf{I}_m),$$

where $\mathbf{A}$ denotes an integral operator. Given a discretization level $h > 0$, we write the scaled white noise likelihood model

$$\pi_{b|u}(b | u_h) \sim \mathcal{N}(\mathbf{A}_h u_h, \sigma_h^2 \mathbf{I}_m).$$

Choosing $\sigma_h^2 = \sigma_{\text{ex}}^2$, thus ignoring the modeling error, may not be a good choice if the discretization is coarse and the exogenous noise level is low, as the IAS algorithm with the discretized model may not be able to approximate the data within the noise level σ_{ex}^2 . Therefore, we increase artificially the noise level σ_h so that the modeling error, although possibly not zero mean Gaussian, is within the noise envelope. To find a suitable value for σ_h , we begin by writing

$$b = \mathbf{A}u + \varepsilon = \mathbf{A}_h u_h + (\mathbf{A}u - \mathbf{A}_h u_h) + \varepsilon = \mathbf{A}_h u_h + \varepsilon_h + \varepsilon,$$

where $\varepsilon_h = \mathbf{A}u - \mathbf{A}_h u_h$ is the modeling error due to discretization.

To estimate the level of the modeling error, we start by generating a sample of realizations u , assuming circular inclusions with random radius and at random positions, denoted by $\mathcal{S} = \{u^{(1)}, \dots, u^{(K)}\}$. Since the amplitude of the modeling error scales linearly, we assume for simplicity that the amplitude of the circular inclusion is one. For a given mesh size $h > 0$, we discretize the domain Ω by generating a uniform mesh, and approximate the densities $u^{(k)}$ by a piecewise linear function with nodal values given by the true density $u^{(k)}$. We compute the corresponding modeling error realizations

$$\varepsilon_h^{(k)} = \mathbf{A}u^{(k)} - \mathbf{A}_h u_h^{(k)}, \quad 1 \leq k \leq K.$$

We approximate the underlying noise distribution by scaled Gaussian white noise with variance $\tilde{\sigma}_h^2$, obtained by

$$\tilde{\sigma}_h^2 = \frac{1}{m} \text{trace} \left(\frac{1}{K} \sum_{k=1}^K \varepsilon_h^{(k)} (\varepsilon_h^{(k)})^\top \right) = \frac{1}{mK} \sum_{k=1}^K \|\varepsilon_h^{(k)}\|^2.$$

Observe that the above equation is exact in the limit $K \rightarrow \infty$ if the underlying distribution is indeed a scaled white noise. To account for the modeling error within a scaled white noise framework, we set

$$\sigma_h^2 = \tilde{\sigma}_h^2 + \sigma_{\text{ex}}^2.$$

Figure 4 shows the value of $\tilde{\sigma}_h$, computed by drawing 200 disc inclusions for the 15-view tomography data, for different uniform mesh sizes h . Since the dependency of $\tilde{\sigma}_h$ on h is almost linear, we can fit a linear model to the data and use $\sigma_h = 0.3h$ for the estimate of the noise model.

Since the modeling error scales linearly with respect to the inclusion, in order to have a reasonable scaling of the discretization error level, we need to have some a priori information about the density. To test how critical it is to know the amplitude, we run the IAS algorithm with fixed mesh for different choices of the approximation error noise variance. The hyperparameters in the IAS algorithm in this test, as well as in the numerical examples to ensue, are set to $r = 1$, corresponding to the gamma hyperprior, and $\eta = 0.001$. The vector ϑ is computed by using the sensitivity scaling. In this preliminary test, we use a fine, fixed mesh of size $h = 0.02$, with 17,813 elements and 9062 nodes. In this test, the exogenous noise level is set at zero; that is, we assume that the only error source is the discretization error in the forward model. To test




FIG. 4. Estimated level of the modeling error as a function of the grid parameter given by the solid curve. The dashed line is the least squares fit, given by $\sigma_h = \alpha h + \beta$ with $\alpha = 0.3035$, $\beta = -0.0005$.

how sensitive the solution is to the estimation of the noise level, we consider four different cases: $\sigma_h = 0.03h$ (underestimated), $\sigma_h = 0.15h$ (slightly underestimated), $\sigma_h = 0.3h$ (correctly estimated), and $\sigma_h = 0.6h$ (overestimated).

Figure 5 shows the results with the IAS algorithm; the hyperparameter controlling the focality of the solution was set at $\eta = 0.001$. Two observations are in order. First, when the modeling error is strongly underestimated, the edges of the target are well identified, but noise clutter is clearly visible in the solution, indicating that the IAS algorithm is fitting to the discretization noise not properly accounted for in the likelihood model. On the other hand, when the modeling error level is overestimated, the reconstruction of the inclusion is not sharp, and the structure of the underlying discretization mesh becomes visible by making the target hexagonal instead of round. These results indicate that the IAS algorithm stops the CGLS iterations too soon when the desired discrepancy is reached. Numerical tests suggest that a rough estimate of the modeling error level is sufficient for a reasonable quality of reconstruction, with a slight preference for underestimating the error level.

Next we ran the BMA algorithm using the modeling error formula $\sigma_h = 0.3 \min(h_j)$, with the minimum taken over all edges. In the **FreeFem++** code, the minimum and maximum mesh sizes were set at $h_{\min} = 0.006$ and $h_{\max} = 0.1$, respectively. In the metric update, we set

$$G_j = \max\{\theta_j, 0.5 \times \theta_{\max}\}, \quad \theta_{\max} = \max\{\theta_j\}.$$

Changes in the threshold value slightly affect the width of the region around the boundary of the inclusion discretized with the finer mesh. Figure 6 shows the results at each of the first four iterations: the left panel displays the underlying mesh, and the right shows the reconstruction based on that mesh. The IAS algorithm, terminated when the relative change in the θ update falls below 10% of the ℓ_∞ -norm of θ , required between four and eight outer iterations. The stopping criterion for the BMA algorithm is based on the three indicators shown in Figure 7, namely, from left to right, the norm of the discrepancy $\|b - A_h u_h\|$, the complexity measured in terms of the number of nodes of the mesh, and the product of these two numbers. Figure 7




FIG. 5. IAS reconstructions after two outer iteration rounds with fixed discretization mesh and with four different estimates of the discretization error, corresponding to, in lexicographical order, $\sigma_h = 0.03h$, $\sigma_h = 0.15h$, $\sigma_h = 0.3h$, and $\sigma_h = 0.6h$. The focality parameter of the IAS algorithm controlling the sharpness of the reconstruction is $\eta = 0.001$.

shows that the discrepancy using the initial coarse mesh is high, as one would expect, and subsequently it drops to a level where it stabilizes. While the discrepancy based on the second mesh is slightly smaller than that based on the third, the second mesh is significantly more complex, indicating an initial overrefinement. An indicator taking into account both the discrepancy and the complexity, such as the product of the two, seems to identify the third mesh as a good compromise. A visual inspection of the meshes confirms that after the third iteration, no significant changes take place, and therefore continuing the iterations would not improve the solution.

The X-ray tomography data with 15 views represent a rather complete set of projections, and it is of interest to see how the proposed methodology works with more limited data, as it is often collected, e.g., in some industrial and medical applications, where the number of views is limited due to geometric constraints, health considerations, or constraints in data acquisition speed. Figure 8 shows the results after three mesh adaptation iterations when the number of views is limited to three and seven, respectively; all other parameters remain the same. The reconstructions, very much in line with what could be expected, exhibit reduced sharpness along the boundary of the inclusion in directions not coinciding with the wave front sets.

Finally, we test the algorithm with a nonconvex inclusion by considering the configuration shown in Figure 1 (right). In the reconstructions, the same value for the modeling error as in the case of the convex disc-like object is used. The noise




FIG. 6. BMA-IAS iterations for the fan-beam tomography estimation problem, with 15 views and noise level 0.5%. Each row shows the discretization mesh (left) and the corresponding IAS estimate (right). The mesh in the next row is based on the prior variance of the estimate in the previous row.




FIG. 7. *Example 1. The norm of the discrepancy, $\|b - A_h u_h\|$ (left), the complexity measured in terms of the number, $n_{v,free}$, of degrees of freedom in the problem (center), and the product $n_{v,free} \times \|b - A_h u_h\|$ of the two (right), as a function of the iteration.*




FIG. 8. *The third iteration of the BMA-IAS iteration mesh and corresponding reconstruction of the density, when the number of views is limited to three (top row) and seven (bottom row). The noise level and the model parameters are the same as in Figures 6 and 7.*

level is 0.1% of the maximum of the noiseless signal. The reconstructions shown in Figure 9 and the adapted meshes suggest that the algorithm is robust with respect to the geometry.

5.2. Inverse source problem for Darcy flow. In the second example, we consider the inverse source problem for the Darcy flow introduced in section 2.2. The generative model consists of the piecewise constant source term shown in Figure 2, and the simulated data, collected in a vector of size $m = 400$, are generated by a second order FEM solver with a uniform mesh consisting of 4974 triangular elements, with




FIG. 9. BMA-IAS iterations for the fan-beam tomography estimation problem, with 15 views and noise level 0.1%, with the data corresponding to the nonconvex inclusion, shown in Figure 1 (right). As in Figure 6, each row shows the discretization mesh (left) and the corresponding IAS estimate (right). The mesh in the next row is based on the prior variance of the estimate in the previous row.

9897 degrees of freedom with exogenous additive scaled white noise with standard deviation σ_{ex} equal to 0.1% of the maximum of the noiseless signal.

We initialize the BMA algorithm by generating a coarse structured mesh with $n_t = 780$ triangular elements as shown in Figure 10 (first row). In the IAS iterations, we use the same parameter values as in the tomography problem, with $r = 1$ corresponding to the gamma hyperprior, and with

$$\eta = 0.001, \quad \vartheta_j = \frac{1 \times 10^{-3}}{\|A_h e_j\|},$$

and we stop the outer iterations when the relative change in the prior parameter θ is less than 10% of the previous value. As in the previous example, this stopping condition is satisfied after only 4–7 outer iterations. The modeling error is estimated by setting $\sigma_h = 0.3 \times h_{\min}$, in analogy with the previous example, where $h_{\min}$ is the minimum edge length in the current discretization. In the update of the metric, we set

$$G_j = \max\{\theta_j, 0.3 \times \theta_{\max}\}, \quad \theta_{\max} = \max\{\theta_j\}.$$

When computing the new mesh using **FreeFem++**, we specify the minimum and maximum allowed element sizes to be $h_{\min} = 0.015$ and $h_{\max} = 0.04$ to avoid overly dense meshes that make the FEM solver slow and badly scaled, and too coarse discretization that would make the FEM approximation error too large.

Figure 10 shows the progression of the algorithm by displaying the first four meshes and the corresponding reconstructions of the source term. As in the tomography problem, the first refinement significantly increases the complexity of the system, while the second update reduces the degrees of freedom by identifying the areas where the source term suggests a jump. After the second refinement, the iterations do not significantly change the discrepancy, and the solution of the inverse problem remains visually unaltered. Figure 11 shows the iteration history of the norm of the discrepancy $\|A_h u_h - b\|$, and the number of nodes of the mesh are plotted, revealing behavior similar to that in the tomography problem. Again, the product of the mesh size and the norm of the discrepancy provides a useful and informative indicator, suggesting that after few mesh adaptations, there is no reason to continue iterating.

6. Discussion. In this paper we propose a novel Bayesian method for simultaneously solving an inverse problem for a distributed parameter and updating the discretization mesh used for approximation of the unknown. We illustrate the performance of the proposed method on two test cases, both arising from linear inverse problems. The examples demonstrate the viability of the suggested method; however, the dependency on the choice of the various model parameters needs further investigation before the algorithm can be fully automatic. In particular, the approximation error model and the mesh size parameters $h_{\min}$ and $h_{\max}$ are addressed here only preliminarily, and further investigations are necessary to gain a comprehensive picture of their role. Extending the approach to nonlinear problems is possible but not immediate, as the underlying IAS algorithm requires a further linearization and iteration steps. In the second example, the mesh determined by the BMA-IAS algorithm is used also for solving the forward model with a higher order finite element method. The order of the FEM approximation is an important feature, as an underlying coarse mesh may be sufficient to accurately represent the distributed parameter of interest, but may fail to produce a sufficiently accurate approximation of the solution of the PDE unless higher order methods are used. To fully control the mesh adaptation




FIG. 10. *BMA-IAS iterations for the source estimation problem in a Darcy flow model. Each row shows the discretization mesh (left) and the corresponding IAS estimate (right). The mesh in the next row is based on the prior variance of the preceding estimate.*




FIG. 11. The norm of the discrepancy $\|b - A_h u_h\|$ as a function of the iteration (left), the complexity measured in terms of the number $n_{v,free}$ of degrees of freedom in the problem (center), and the product $n_{v,free} \times \|b - A_h u_h\|$ of the two (right).

process, in addition to the variance parameter of the inverse problem, the accuracy of the FEM solution via the a posteriori error estimates should be monitored. The accuracy of the forward model approximation constitutes an important part of a careful modeling error analysis that is not addressed in detail in this paper. A further extension of the modeling error analysis entails the correlation structure of the modeling error. It is well known that the modeling error due to discretization is typically strongly correlated and normally has a nonvanishing mean. A more careful analysis would require new ideas for dynamically assessing the modeling error without needing to resort to very dense discretization meshes that we aim to avoid by proposing the dynamic mesh updating. The algorithm was tested by using piecewise constant generative models that after discretization allow a sparse representation. From the point of view of the algorithm, the sparsity is not a necessary assumption, as the IAS algorithm can be adjusted for non-sparse problems by a proper selection of the model parameters [7]. However, the discretization error analysis requires more attention and is worth a separate discussion.

In our discussion, it was assumed that the target mesh at each adaptation should be close to isotropic. This assumption was encoded in the algorithm by assuming that the metric guiding the mesh calculation is such that the eigenvalues of the matrix defining the metric at each point are equal. In principle, this is not necessary, and, in fact, an anisotropic metric gives rise to meshes characterized by stretched elements that are used successfully to represent strongly directional features; see, e.g., [18, 19, 30, 31]. Extension of the present work to anisotropic meshes is a possible future direction but beyond of the scope of the present paper.

REFERENCES

- [1] M. AINSWORTH AND J. T. ODEN, *A Posteriori Error Estimation in Finite Element Analysis*, Wiley, New York, 2000.
- [2] S. R. ARRIDGE, J. P. KAIPPIO, V. KOLEHMAINEN, M. SCHWEIGER, E. SOMERSALO, T. TÄRVÄINEN, AND M. VAUHKONEN, *Approximation errors and model reduction with an application in optical diffusion tomography*, *Inverse Problems*, 22 (2006), pp. 175–195.
- [3] C. L. BOTTASSO, G. MAISANO, S. MICHELETTI, AND S. PEROTTO, *On some new recovery based a posteriori error estimators*, *Comput. Meth. Appl. Mech. Engrg.*, 195 (2006), pp. 4794–4815.
- [4] D. CALVETTI, M. DUNLOP, E. SOMERSALO, AND A. STUART, *Iterative updating of model error in Bayesian inversion*, *Inverse Problems*, 34 (2018), 025008.
- [5] D. CALVETTI, O. ERNST, AND E. SOMERSALO, *Dynamic updating of numerical model discrepancy using sequential sampling*, *Inverse Problems*, 30 (2014), 114019.
- [6] D. CALVETTI, H. HAKULA, S. PURSIANEN, AND E. SOMERSALO, *Conditionally Gaussian hypermodels for cerebral source localization*, *SIAM J. Imaging Sci.*, 2 (2009), pp. 879–909, <https://doi.org/10.1137/080723995>.

- [7] D. CALVETTI, M. PRAGIOLA, E. SOMERSALO, AND A. STRANG, *Sparse reconstructions from few noisy data: Analysis of hierarchical Bayesian models with generalized gamma hyperpriors*, *Inverse Problems*, 36 (2020), 025010.
- [8] D. CALVETTI, M. PRAGIOLA, AND E. SOMERSALO, *Hybrid Solver for Hierarchical Bayesian Inverse Problems*, preprint, <https://arxiv.org/abs/2003.06532v1>, 2020.
- [9] D. CALVETTI, A. PASCARELLA, F. PITOLLI, E. SOMERSALO, AND B. VANTAGGI, *A hierarchical Krylov-Bayes iterative inverse solver for MEG with physiological preconditioning*, *Inverse Problems*, 31 (2015), 125005.
- [10] D. CALVETTI AND E. SOMERSALO, *Priorconditioners for linear systems*, *Inverse Problems*, 21 (2005), pp. 1397–1418.
- [11] D. CALVETTI AND E. SOMERSALO, *Subjective knowledge or objective belief? An oblique look to Bayesian methods*, in *Large-Scale Inverse Problems and Quantification of Uncertainty*, L. Biegler et al., eds., John Wiley & Sons Ltd., 2010, pp. 33–70.
- [12] D. CALVETTI, E. SOMERSALO, AND A. STRANG (2019), *Hierarchical Bayesian models and sparsity: ℓ_2 -magic*, *Inverse Problems*, 35 (2019), 035003.
- [13] M. A. CAPISTRÁN, J. A. CHRISTEN, AND S. DONNET, *Bayesian Analysis of ODEs: Solver optimal accuracy and Bayes factors*, *SIAM/ASA J. Uncertain. Quantif.*, 4 (2016), pp. 829–849, <https://doi.org/10.1137/140976777>.
- [14] M. CHENEY, D. ISAACSON, AND J. C. NEWELL, *Electrical impedance tomography*, *SIAM Rev.*, 41 (1999), pp. 85–101, <https://doi.org/10.1137/S0036144598333613>.
- [15] J. A. CHRISTEN, M. A. CAPISTRAN, M. L. DAZA-TORRES, H. FLORES-ARGÜEDAS, AND J. CRICELIO MONTESINOS-LÓPEZ, *Posterior Distribution Existence and Error Control in Banach Spaces in the Bayesian Approach to UQ in Inverse Problems*, preprint, <https://arxiv.org/abs/1712.03299>, 2018.
- [16] PH. G. CIARLET, *The Finite Element Method for Elliptic Problems*, North-Holland, Amsterdam, 1978; reprinted as *Classics Appl. Math.* 40, SIAM, 2002, <https://doi.org/10.1137/1.9780898719208>.
- [17] T. CUI, Y. M. MARZOUK, AND K. WILLCOX, *Data-driven model reduction for the Bayesian solution of inverse problems*, *Internat. J. Numer. Methods Engrg.*, 102 (2015), pp. 966–990.
- [18] P. E. FARRELL, S. MICHELETTI, AND S. PEROTTO, *An anisotropic Zienkiewicz-Zhu type error estimator for 3D applications*, *Internat. J. Numer. Methods Engrg.*, 85 (2011), pp. 671–692.
- [19] L. FORMAGGIA, S. MICHELETTI, AND S. PEROTTO, *Anisotropic mesh adaption with application to CFD problems*, in *Proceedings of WCCM V, Fifth World Congress on Computational Mechanics*, Barcelona, Spain, H. Mang, F. Rammerstorfer, and J. Eberhardsteiner, eds., 2002, pp. 1481–1493.
- [20] P. L. GEORGE AND H. BOROCHAKI, *Delaunay Triangulation and Meshing: Application to Finite Elements*, Editions Hermès, 1993.
- [21] F. HECHT, *New development in FreeFem++*, *J. Numer. Math.*, 20 (2012), pp. 251–265.
- [22] C. JOHNSON, *Numerical Solutions of Partial Differential Equations by the Finite Element Method*, reprint of the 1987 edition, Dover Publications, Inc., 2009.
- [23] J. P. KAIPIO AND E. SOMERSALO, *Statistical inverse problems: Discretization, model reduction and inverse crimes*, *J. Comput. Appl. Math.*, 198 (2007), pp. 493–504.
- [24] J. KAIPIO AND E. SOMERSALO, *Statistical and Computational Inverse Problems*, Springer-Verlag, 2004.
- [25] M. KRÍZEK AND P. NEITTAANMÄKI, *Superconvergence phenomenon in the finite element method arising from averaging gradients*, *Numer. Math.*, 45 (1984), pp. 105–116.
- [26] J. LI AND Y. M. MARZOUK, *Adaptive construction of surrogates for the Bayesian solution of inverse problems*, *SIAM J. Sci. Comput.*, 36 (2014), pp. A1163–A1186, <https://doi.org/10.1137/130938189>.
- [27] Y. LI AND D. W. OLDENBURG, *3-D inversion of magnetic data*, *Geophysics*, 61 (1996), pp. 394–408.
- [28] Y. LI AND D. W. OLDENBURG, *3-D inversion of gravity data*, *Geophys.*, 63 (1998), pp. 109–119.
- [29] F. H. LIN, T. WITZEL, S. P. AHLFORS, S. M. STUFFLEBEAM, J. W. BELLIVEAU, AND M. S. HÄMÄLÄINEN, *Assessing and improving the spatial accuracy in MEG source localization by depth-weighted minimum-norm estimates*, *Neuroimage*, 31 (2006), pp. 160–171.
- [30] S. MICHELETTI, S. PEROTTO, AND P. E. FARRELL, *A recovery-based error estimator for anisotropic mesh adaptation in CFD*, *Bol. Soc. Esp. Mat. Apl. SeMA*, 50 (2010), pp. 115–137.
- [31] S. MICHELETTI AND S. PEROTTO, *Anisotropic adaptation via a Zienkiewicz-Zhu error estimator for 2D elliptic problems*, in *Numerical Mathematics and Advanced Applications*, G. Kreiss, P. Lötstedt, A. Målqvist, and M. Neytcheva, eds., Springer-Verlag, 2010, pp. 645–653.

- [32] A. NAGA AND Z. ZHANG, *A posteriori error estimates based on the polynomial preserving recovery*, SIAM J. Numer. Anal., 42 (2004), pp. 1780–1800, <https://doi.org/10.1137/S0036142903413002>.
- [33] F. NATTERER, *The Mathematics of Computerized Tomography*, Classics Appl. Math. 32, SIAM, 2001, <https://doi.org/10.1137/1.9780898719284>.
- [34] A. NISSINEN, L. M. HEIKKINEN, AND J. P. KAIPIO, *The Bayesian approximation error approach for electrical impedance tomography—experimental results*, Meas. Sci. Technol., 19 (2007), 015501.
- [35] A. NISSINEN, V. KOLEHMAINEN, AND J. P. KAIPIO, *Reconstruction of domain boundary and conductivity in electrical impedance tomography using the approximation error approach*, Int. J. Uncertain. Quantif., 1 (2011), pp. 203–222.
- [36] R. RODRÍGUEZ, *Some remarks on Zienkiewicz-Zhu estimator*, Numer. Methods Partial Differential Equations, 10 (1994), pp. 625–635.
- [37] N. YAN AND A. ZHOU, *Gradient recovery type a posteriori error estimation for finite element approximations on irregular meshes*, Comput. Methods Appl. Mech. Engrg., 190 (2001), pp. 4289–4299.
- [38] L. YAN AND T. ZHOU, *Adaptive multi-fidelity polynomial chaos approach to Bayesian inference in inverse problems*, J. Comput. Phys., 381 (2019), pp. 110–128.
- [39] L. YAN AND T. ZHOU, *An adaptive multi-fidelity PC-based ensemble Kalman inversion for inverse problems*, Int. J. Uncertain. Quantif., 9 (2019), pp. 205–220.
- [40] O. C. ZIENKIEWICZ AND K. MORGAN, *Finite Elements and Approximation*, reprint of the 1983 original, Dover Publications, Inc., 2006.
- [41] O. C. ZIENKIEWICZ AND J. Z. ZHU, *A simple error estimator and adaptive procedure for practical engineering analysis*, Int. J. Numer. Methods Engrg., 24 (1987), pp. 337–357.
- [42] O. C. ZIENKIEWICZ AND J. Z. ZHU, *The superconvergent patch recovery and a posteriori error estimates I: The recovery technique*, Int. J. Numer. Methods Engrg., 33 (1992), pp. 1331–1364.
- [43] O. C. ZIENKIEWICZ AND J. Z. ZHU, *The superconvergent patch recovery and a posteriori error estimates II: Error estimates and adaptivity*, Int. J. Numer. Methods Engrg., 33 (1992), pp. 1365–1382.