

Estimation of the Robin coefficient field in a Poisson problem with uncertain conductivity field*

Ruanui Nicholson^{1,4} , Noémi Petra² and Jari P Kaipio³

¹ Department of Engineering Science, University of Auckland, Private Bag 92019, Auckland Mail Center, Auckland 1142, New Zealand

² School of Natural Sciences, University of California, Merced, 5200 North Lake Road, Merced, CA 95343, United States of America

³ Department of Mathematics, University of Auckland, Private Bag 92019, Auckland Mail Center, Auckland 1142, New Zealand

E-mail: ruanui.nicholson@auckland.ac.nz

Received 21 December 2017, revised 3 July 2018

Accepted for publication 9 August 2018

Published 31 August 2018



CrossMark

Abstract

We consider the reconstruction of a heterogeneous coefficient field in a Robin boundary condition on an inaccessible part of the boundary in a Poisson problem with an uncertain (or unknown) inhomogeneous conductivity field in the interior of the domain. To account for model errors that stem from the uncertainty in the conductivity coefficient, we treat the unknown conductivity as a nuisance parameter and carry out approximative premarginalization over it, and invert for the Robin coefficient field only. We approximate the related modelling errors via the Bayesian approximation error (BAE) approach. The uncertainty analysis presented here relies on a local linearization of the parameter-to-observable map at the maximum *a posteriori* (MAP) estimates, which leads to a normal (Gaussian) approximation of the parameter posterior density. To compute the MAP point we apply an inexact Newton conjugate gradient approach based on the adjoint methodology. The construction of the covariance is made tractable by invoking a low-rank approximation of the data misfit component of the Hessian. Two numerical experiments are considered: one where the prior covariance on the conductivity is isotropic, and one where the prior covariance on the conductivity is anisotropic. Results are compared to those based on standard error models, with particular emphasis on the feasibility of the posterior uncertainty estimates. We show that the BAE approach is a feasible one in the sense that the predicted posterior uncertainty is consistent with the actual estimation errors, while neglecting the

* This work was partially supported by NSF grants ACI-1550547 and 1654311.

⁴ Author to whom any correspondence should be addressed.

related modelling error yields infeasible estimates for the Robin coefficient. In addition, we demonstrate that the BAE approach is approximately as computationally expensive (measured in the number of PDE solves) as the conventional error approach.

Keywords: estimation of Robin coefficient, Modelling errors, adjoint-based Hessian, low rank approximation, Bayesian approximation error approach, approximate marginalization, Bayesian framework

(Some figures may appear in colour only in the online journal)

1. Introduction

In this paper, we consider the problem of estimating the unknown Robin coefficient field in a Poisson problem with uncertain conductivity field from available measurements. This problem setup is inspired from physical applications, e.g. the detection of corrosion of an electrostatic conductor [1–3] and the estimation of thermal parameters [4, 5]. The Poisson problem with uncertain Robin boundary coefficient has received considerable attention, both from a theoretical standpoint as well as from a numerical point of view [1–10]. However, the standard assumption in these studies is that the internal conductivity (which is generally a distribute parameter) is known. In this paper, we consider both the conductivity and boundary condition to be uncertain. A common approach would be to invert for both fields simultaneously, however this results in a highly ill-posed and potentially untrackable problem. To avoid the need for a joint inversion, we *premarginalize* over the internal conductivity, and then invert for the Robin coefficient. Furthermore, we employ a discretisation invariant method for the inversions as in [11, 12]. Thus the methods developed should be immediately applicable to the case in which the Robin parameter is high-dimensional (which is often the case for real applications).

There is a rich body of literature on theoretical and computational aspects of the so-called inverse Robin problem, i.e. the problem of inferring the (distributed) Robin coefficient given measurements of the potential. In [2], the authors develop a direct reconstruction method based on a thin plate approximation. In [6], several results on stability, uniqueness and identifiability are established, while in [3] a more general stability estimate is proved. Numerical methods developed to solve the inverse Robin problem include a quasi-reversibility method [7] and an approach based on an L^1 -tracking functional [8]. More recently, in [9], a regularized least-squares approach is taken via a variational formulation and in [10], a regularised least-squares problem is solved using an adjoint based approach, similar to the methods considered in the present paper.

The previous studies consider the following inverse problem: given noisy (partial point) measurements of u on the boundary of a bounded domain Ω , with $\Omega \in \mathbb{R}^d$, $d \in \{2, 3\}$, determine the Robin coefficient field, $\beta(\mathbf{x})$. The field u satisfies the forward problem,

$$\begin{aligned} -\Delta u(\mathbf{x}) &= 0 && \text{in } \Omega, \\ \nabla u(\mathbf{x}) \cdot \mathbf{n}_t &= g(\mathbf{x}) && \text{on } \Gamma_t \\ \nabla u(\mathbf{x}) \cdot \mathbf{n}_b + \exp(\beta(\mathbf{x}))u(\mathbf{x}) &= 0 && \text{on } \Gamma_b \\ u(\mathbf{x}) &= 0 && \text{on } \Gamma_s, \end{aligned} \quad (1)$$

where (in the context of the present paper), Γ_t is referred to as the top of the domain, Γ_s the sides of the domain, and Γ_b the bottom of the domain. As such, we have

$\Gamma_t \cap \Gamma_s = \Gamma_s \cap \Gamma_b = \Gamma_t \cap \Gamma_b = \emptyset$, and (eventually noisy pointwise) measurements of u are available on Γ_t . In the literature, Dirichlet boundary conditions are often replaced by Neumann boundary conditions on Γ_s , see for example [8, 9].

However, in essentially all practical problems, the conductivity cannot be assumed to be a constant, leading to the spatially inhomogeneous problem

$$\begin{aligned} -\nabla \cdot (\exp(a(\mathbf{x})) \nabla u(\mathbf{x})) &= 0 && \text{in } \Omega, \\ \exp(a(\mathbf{x})) \nabla u(\mathbf{x}) \cdot \mathbf{n}_t &= g(\mathbf{x}) && \text{on } \Gamma_t \\ \exp(a(\mathbf{x})) \nabla u(\mathbf{x}) \cdot \mathbf{n}_b + \exp(\beta(\mathbf{x})) u(\mathbf{x}) &= 0 && \text{on } \Gamma_b \\ u(\mathbf{x}) &= 0 && \text{on } \Gamma_s. \end{aligned} \quad (2)$$

For example, in the case of detecting corrosion of an electrostatic volume conductor, $\exp(a)$ models the electrical conductivity and it is generally an unknown (distributed) parameter.

The estimation of the Robin coefficient $\beta(\mathbf{x})$ under the forward model (2) would typically necessitate the simultaneous estimation of the conductivity $a(\mathbf{x})$. Such a task carries with it several challenges: the ill-posedness of the problem would be increased significantly, and there is the potential for issues around the identifiability of $\beta(\mathbf{x})$. In this paper, however, our main concern is the drastically increased computational cost associated with estimating the parameter $a(\mathbf{x})$ which is distributed over the entire volume Ω rather than estimating $\beta(\mathbf{x})$ only which is distributed on Γ_b only.

The approach in this paper is based on the (initially) infinite-dimensional formulation of the inverse problem, the adjoint method for the computation of the related first and second order derivative information, and the subsequent Bayesian approximation error approach. A brief review of these topics is given below.

The infinite dimensional approach to Bayesian inverse problems that was developed in [13] is receiving considerable attention, and has been successfully applied to several inverse problems. The method is particularly suited to the case when the parameter of interest is high-dimensional (stemming from the discretization of the the unknown infinite-dimensional parameter field), and it ensures convergence under discretization. An efficient computational framework was developed in [11, 12], based on an adjoint approach [14, 15], to implement the theoretical work put forward in [13], and was applied, for example, to global seismic inversions in [11] and ice sheet flow inverse problems in [12]. The approach has also successfully been applied to inverse acoustic obstacle scattering problems in [16]. The infinite-dimensional Bayesian setup has also been applied to optimal experimental design (OED) for Bayesian non-linear inverse problems governed by partial differential equations [17]. The goal of the OED problem was to find an optimal placement of sensors (for measurements) so as to minimize the uncertainty in the inferred parameter field.

The Bayesian approximation error (BAE) approach [18, 19] was originally used as a means to take into account the modelling errors induced by the use of reduced order models. The approach is based on approximate premarginalization over modelling errors, which refers to a process similar to the marginalization over additive errors to obtain the likelihood. However, a particular advantage of this method is the ability to approximately premarginalize also over parameters which are not of primary interest. In the context of electrical impedance tomography (EIT), the BAE approach has been used to simultaneously premarginalize over the unknown domain shape and the contact impedances of the electrodes [20]. Furthermore, in [21], the approach was used to premarginalize over the distributed scattering coefficient in diffuse optical tomography (DOT) and, in [22], the method was used to premarginalize over both the scattering and absorption coefficients in the context of fluorescence diffuse optical

tomography (fDOT). The BAE method has also been applied to x-ray tomography to premarginalize over distributed parameters outside a region of interest [23].

In the context of premarginalizing over distributed parameters, the BAE approach has thus far only been used to marginalize over unknowns defined on spatial dimensions at most equal to that of the primary parameter of interest. In this paper, we show that the BAE approach is also feasible for the premarginalization over a distributed parameter in the entire domain when the parameter of primary interest is defined only on (a subset of) the boundary of the domain. We also show that the infinite-dimensional framework for inverse problems posed in the Bayesian setting is an effective method for solving the so-called Robin inverse problem under an unknown (distributed) conductivity.

The paper is organized as follows. In section 2, we review the Bayesian framework for inverse problems, including a limited discussion on the infinite dimensional set up, the computation of the maximum *a posteriori* estimate and the approximate posterior covariance. In section 3, we review the Bayesian approximation error approach and, in section 4, we formulate the problem of estimating the Robin coefficient in the case of an unknown conductivity. In section 5, we consider two numerical experiments: a conductivity with spatially isotropic smooth covariance and one with an anisotropic smooth covariance. The results are compared to those based on standard error models, with particular emphasis on the feasibility of the posterior uncertainty estimates.

2. Background on the Bayesian approach to inverse problems

In this section, we give a brief review of the formulation of Bayesian inverse problems following [11, 12] to an extent that is relevant to the present paper. For a more in-depth discussion on some of the technical difficulties in dealing with inverse problems posed in infinite dimensions, see for example [11–13]. We consider the problem of finding $\beta(\mathbf{x}) \in \mathcal{H} \subset L^2(\Omega_\beta)$, from observed measurements $\mathbf{d}^{\text{obs}} \in \mathbb{R}^q$, with β and $\mathbf{d}^{\text{obs}}$ linked by

$$\mathbf{d}^{\text{obs}} = \mathbf{f}_{a_*}(\beta) + \mathbf{e}, \quad (3)$$

where $\mathbf{f}_{a_*} : \mathcal{H} \rightarrow \mathbb{R}^q$ is the parameter-to-observable map, and $\mathbf{e}$ represents additive errors in the measurements. The slightly unconventional notation $\mathbf{f}_{a_*}$ used for the so-called *parameter-to-observable map* will be explained in section 3.

In this paper, we take the prior to be a Gaussian measure, $\mu_\beta = \mathcal{N}(\beta_*, \mathcal{C}_\beta)$ on $L^2(\Omega_\beta)$, where β_* is the prior mean, which lives in $\mathcal{H}$, and $\mathcal{C}_\beta$ is the prior covariance operator. As outlined in [13], the prior must be chosen to satisfy certain regularity assumptions to ensure the Bayesian inverse problem is well-defined. We employ a *weighted* squared inverse elliptic operator as our prior covariance operator [24], with the addition of homogeneous Robin (or Neumann) boundary conditions. This is a slight modification to that used in [11, 12], with the aim of mitigating any artefacts in estimates or prior samples due to the enforcement of boundary conditions. Specifically, for $s \in L^2(\Omega_\beta)$, the weak solution of $\mathcal{A}\beta = s$ satisfies

$$\alpha_\beta \int_{\Omega_\beta} \gamma_\beta \nabla \beta \cdot \nabla v + \beta v \, \mathbf{d}\mathbf{y} + \int_{\partial\Omega_\beta} \kappa_\beta \beta v \, \mathbf{d}\mathbf{t} = \int_{\Omega_\beta} s v \, \mathbf{d}\mathbf{y} \quad \text{for all } v \in H^1(\Omega_\beta), \quad (4)$$

where $\alpha_\beta > 0$ is inversely proportional to the prior variance, γ_β is a symmetric positive definite uniformly bounded matrix controlling the correlation [11], and $\kappa_\beta \geq 0$. Then we take the prior covariance operator to be

$$\mathcal{C}_\beta = \mathcal{W} \mathcal{A}^{-2} \mathcal{W} \quad (5)$$

where

$$\mathcal{W} := \frac{\sigma_\beta}{\sqrt{\mathcal{A}^{-2}(\mathbf{x}, \mathbf{x})}} \quad \text{and} \quad \sigma_\beta := \sqrt{\frac{\text{Ga}(\nu)}{(4\pi)^{d/2} \gamma_\beta^\nu \alpha_\beta^2}}. \quad (6)$$

Here Ga denotes the gamma function, d is the spatial dimension, and $\nu + d/2 = 2$ [24], and σ is the point-wise standard deviation, as in [24, 25]. The addition of the weights normalizes the variance across the domain as discussed below in section 5. For efficient methods to extract $\mathcal{A}^{-2}(\mathbf{x}, \mathbf{x})$, see for example, [26–30]. We also note that there are other methods aimed at mitigating the boundary effects, see for example, [24, 31, 32].

In this paper, we consider a normal noise model $\mathbf{e} \sim \mu_{\text{noise}} = \mathcal{N}(0, \mathbf{\Gamma}_e)$, which results in the likelihood

$$\pi_{\text{like}}(\mathbf{d}^{\text{obs}} | \beta) \propto \exp \left\{ -\frac{1}{2} \left(\mathbf{f}_{a_*}(\beta) - \mathbf{d}^{\text{obs}} \right)^T \mathbf{\Gamma}_e^{-1} \left(\mathbf{f}_{a_*}(\beta) - \mathbf{d}^{\text{obs}} \right) \right\}. \quad (7)$$

2.1. Discretization of Bayesian inverse problems

In this section, we briefly review some of the key results on the finite-dimensional approximation of Bayesian inverse problems. The discussion here follows [11, 12, 17], where it is shown that the mass-weighted inner product space is the *correct* space to work in.

Discretization of the inverse problem results in the parameter of interest being approximated as $\beta_h = \sum_{j=1}^n \beta_j \phi_j \in V_h$, where V_h denotes a finite-dimensional subspace of $L^2(\Omega_\beta)$ induced by a finite element discretization with continuous Lagrange basis functions. We then seek to invert for $\beta = [\beta_1, \beta_2, \dots, \beta_n]^T \in \mathbb{R}^n$. Furthermore any inner product between nodal coefficients will be weighted by a mass matrix $\mathbf{M}$ so as to correctly approximate the infinite-dimensional L^2 inner product. We denote the mass-weighted inner product by $\langle \cdot, \cdot \rangle_{\mathbf{M}}$, with $\langle \mathbf{y}, \mathbf{z} \rangle_{\mathbf{M}} = \mathbf{y}^T \mathbf{M} \mathbf{z}$ and the symmetric positive definite mass matrix given by

$$M_{ij} = \int_{\Omega_\beta} \phi_i(\mathbf{y}) \phi_j(\mathbf{y}) \, d\mathbf{y} \quad i, j \in \{1, 2, \dots, n\}. \quad (8)$$

To distinguish between the Euclidean space $\mathbb{R}^n$ and $\mathbb{R}^n$ endowed with the mass-weighted inner product, we introduce the notation $\mathbb{R}_{\mathbf{M}}^n$ to denote $\mathbb{R}^n$ equipped with the mass-weighted inner product.

As shown in for example [11], there are several differences which must be taken into account when working in $\mathbb{R}_{\mathbf{M}}^n$ rather than in $\mathbb{R}^n$, which revolve largely around the concepts of adjoints and transposes. For instance, for an operator $\mathbf{B} : \mathbb{R}_{\mathbf{M}}^n \rightarrow \mathbb{R}_{\mathbf{M}}^n$, we denote the matrix transpose by $\mathbf{B}^T$. The mass-weighted inner product adjoint, however, denoted here by $\mathbf{B}^*$, satisfies $\langle \mathbf{B}\mathbf{y}, \mathbf{z} \rangle_{\mathbf{M}} = \langle \mathbf{y}, \mathbf{B}^*\mathbf{z} \rangle_{\mathbf{M}}$, for $\mathbf{y}, \mathbf{z} \in \mathbb{R}_{\mathbf{M}}^n$, implying

$$\mathbf{B}^* = \mathbf{M}^{-1} \mathbf{B}^T \mathbf{M}. \quad (9)$$

We also require two further adjoint type operators below in section 2.2. For this, let $\mathbb{R}^q$ and $\mathbb{R}^r$ for some q, r , be endowed with the Euclidean inner product. We then define the adjoints $\mathbf{F}^\natural$ of $\mathbf{F} : \mathbb{R}_{\mathbf{M}}^n \rightarrow \mathbb{R}^q$ and $\mathbf{V}^\diamond$ of $\mathbf{V} : \mathbb{R}^r \rightarrow \mathbb{R}_{\mathbf{M}}^n$ as

$$\mathbf{F}^\natural = \mathbf{M}^{-1} \mathbf{F}^T, \quad (10)$$

$$\mathbf{V}^\diamond = \mathbf{V}^T \mathbf{M}. \quad (11)$$

With these definitions in hand, the finite-dimensional approximation of the operator $\mathcal{A}$ in (4) is $\mathbf{A} = \mathbf{M}^{-1} \mathbf{K} \mathbf{W}^{-1}$ where

$$K_{ij} = \alpha_\beta \int_{\Omega_\beta} (\gamma_\beta \nabla \phi_i \cdot \nabla \phi_j + \phi_i \phi_j) \, d\mathbf{y} + \int_{\partial\Omega_\beta} \kappa_\beta \phi_i \phi_j \, d\mathbf{t}, \quad (12)$$

$$W_{ij} = \frac{1}{\sigma_\beta} \sqrt{K_{ij}^{-1}} \delta_{ij}, \quad i, j \in \{1, 2, \dots, n\}, \quad (13)$$

where σ_β is as in (6) and δ_{ij} is the Kronecker delta.

We can now express the finite-dimensional approximation μ_β^h of the prior Gaussian measure, μ_β , as

$$\pi_{\text{prior}}(\beta) \propto \exp\left(-\frac{1}{2} \|\mathbf{A}(\beta - \beta_*)\|_{\mathbf{M}}^2\right), \quad (14)$$

where β_* is the discretization of the prior mean, β_* , and the prior covariance matrix is by definition $\Gamma_\beta = \mathbf{A}^{-2}$ (since $\mathbf{A}$ is symmetric). We can now state the finite-dimensional Bayes' formula

$$\pi_{\text{post}}(\beta) := \pi_{\text{post}}(\beta | \mathbf{d}^{\text{obs}}) \propto \pi_{\text{prior}}(\beta) \pi_{\text{like}}(\mathbf{d}^{\text{obs}} | \beta), \quad (15)$$

where $\pi_{\text{post}}(\beta)$ is the density of the finite-dimensional approximation μ_d^h of the posterior measure μ_d and π_{like} is the likelihood given by (7), see [11, 12]. Thus we can express the posterior density explicitly as

$$\pi_{\text{post}}(\beta) \propto \exp\left(-\frac{1}{2} \left\| \mathbf{f}_{a_*}(\beta) - \mathbf{d}^{\text{obs}} \right\|_{\Gamma_e^{-1}}^2 - \frac{1}{2} \|\mathbf{A}(\beta - \beta_*)\|_{\mathbf{M}}^2\right). \quad (16)$$

2.2. The MAP estimate and the approximate posterior covariance

In the Bayesian framework, one seeks to determine the posterior density of the parameter of interest. In principle, then, one would explore the posterior density typically with Markov chain Monte Carlo sampling methods. However, in large-scale problems with limited computational resources, one often has to be content with a single representative point estimate of the parameter along with an (approximate) posterior covariance and credibility intervals. Standard point estimates for the posterior include the conditional mean (CM) and the maximum *a priori* (MAP) estimates. In this paper, we aim at computing the finite-dimensional MAP estimate β_{MAP} and the Laplace (local normal) approximation for the covariance which also yields approximate marginal distribution for $\beta_{\text{MAP},k}$. For discussions on the extension of the MAP estimate to infinite dimensions, see, for example [33, 34].

Following on from (16), the MAP estimate is defined as the point in parameter space that maximizes the posterior probability density function [18, 35], that is,

$$\beta_{\text{MAP}} := \arg \min_{\beta \in \mathbb{R}^n} \frac{1}{2} \left\| \mathbf{f}_{a_*}(\beta) - \mathbf{d}^{\text{obs}} \right\|_{\Gamma_e^{-1}}^2 + \frac{1}{2} \|\mathbf{A}(\beta - \beta_*)\|_{\mathbf{M}}^2. \quad (17)$$

In line with [12], in (17) and the remainder of the paper, we denote by $\mathbf{f}_{a_*}(\beta)$ the parameter-to-observable map evaluated at the finite element function corresponding to the parameter vector β .

Assuming that the parameter-to-observable map is Fréchet differentiable, we can linearize (3) around β_{MAP} and discretize to obtain the affine approximation

$$\mathbf{d}^{\text{obs}} \approx \mathbf{f}_{a_*}(\beta_{\text{MAP}}) + \mathbf{F}_{a_*}(\beta_{\text{MAP}})(\beta - \beta_{\text{MAP}}) + \mathbf{e}, \quad (18)$$

where $\mathbf{F}_{a_*}(\beta_{\text{MAP}})$ is the Fréchet derivative of $\mathbf{f}_{a_*}(\beta)$ with respect to β evaluated at β_{MAP} . The resulting approximation for the discrete posterior measure μ_d^h of β conditional on $\mathbf{d}^{\text{obs}}$ is then necessarily a Gaussian measure, given by $\mathcal{N}(\beta_{\text{MAP}}, \mathbf{\Gamma}_{\text{post}})$, with covariance matrix

$$\mathbf{\Gamma}_{\text{post}} = \left(\mathbf{F}_{a_*}^\dagger \mathbf{\Gamma}_e^{-1} \mathbf{F}_{a_*} + \mathbf{\Gamma}_\beta^{-1} \right)^{-1}, \quad (19)$$

where $\mathbf{F}_{a_*}^\dagger$ is the adjoint of $\mathbf{F}_{a_*}$, see (10). We also note that the (approximate) posterior covariance matrix is the inverse of the Gauss–Newton Hessian of the negative log posterior (referred to simply as the Hessian from this point onwards), denoted by $\mathbf{H}$, that is,

$$\mathbf{\Gamma}_{\text{post}} = \mathbf{H}^{-1}. \quad (20)$$

2.2.1. Low rank approximation for the approximate posterior covariance matrix. For an efficient calculation of the MAP estimate, and efficient action of the approximate posterior covariance to vectors and action of the square root of the Hessian on vectors (the latter is needed for computing samples from the approximate posterior), we apply a low rank approximation of the Hessian as detailed in [11, 12, 36] and summarized below.

We note that the Hessian of the negative log posterior (16) can be split into the sum of a data misfit term, $\mathbf{H}_{\text{mis}}$, and the inverse of the prior covariance, $\mathbf{\Gamma}_\beta^{-1}$. By factorizing the prior covariance as $\mathbf{\Gamma}_\beta = \mathbf{L}\mathbf{L}^*$, we can rewrite the Hessian as

$$\mathbf{H} = \mathbf{H}_{\text{mis}} + \mathbf{\Gamma}_\beta^{-1} = \mathbf{H}_{\text{mis}} + \mathbf{L}^{-*} \mathbf{L}^{-1} = \mathbf{L}^{-*} (\mathbf{L}^* \mathbf{H}_{\text{mis}} \mathbf{L} + \mathbf{I}) \mathbf{L}^{-1}. \quad (21)$$

The final form of (21) allows for an efficient method of approximating $\mathbf{\Gamma}_{\text{post}}$. The procedure relies on exploiting the discretization invariant and often low rank nature of $\mathbf{H}_{\text{mis}}$ [11, 12, 36, 37]. Thus, in this paper, we consider a low rank approximation for the so-called *prior-preconditioned data misfit Hessian* [36] which takes the form

$$\tilde{\mathbf{H}}_{\text{mis}} = \mathbf{L}^* \mathbf{H}_{\text{mis}} \mathbf{L} \approx \mathbf{V}_r \mathbf{\Lambda}_r \mathbf{V}_r^\circ, \quad (22)$$

where $\mathbf{V}_r \in \mathbb{R}^{n \times r}$ contains the r eigenvectors of the prior-preconditioned data misfit Hessian corresponding to the r largest eigenvalues⁵ λ_i , $i = 1, 2, \dots, r$, and $\mathbf{\Lambda}_r = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_r)$. Then, by using the Sherman–Morrison–Woodbury formula [38], we have

$$\mathbf{H}^{-1} \approx \mathbf{L} (\mathbf{I} - \mathbf{V}_r \mathbf{D}_r \mathbf{V}_r^\circ) \mathbf{L}^*, \quad (23)$$

where $\mathbf{D}_r = \text{diag}(\lambda_1/(\lambda_1 + 1), \lambda_2/(\lambda_2 + 1), \dots, \lambda_r/(\lambda_r + 1)) \in \mathbb{R}^{r \times r}$. Furthermore, for drawing samples from the Gaussian approximation for the posterior, we have

$$\mathbf{H}^{-1/2} \approx \mathbf{L} (\mathbf{V}_r \mathbf{P}_r \mathbf{V}_r^\circ + \mathbf{I}) \mathbf{M}^{-1/2}, \quad (24)$$

where $\mathbf{P}_r = \text{diag}(\lambda_1/\sqrt{\lambda_1 + 1} - 1, \lambda_2/\sqrt{\lambda_2 + 1} - 1, \dots, \lambda_r/\sqrt{\lambda_r + 1} - 1) \in \mathbb{R}^{r \times r}$, see [11].

3. Background on model discrepancy and the Bayesian approximation error approach

Solving inverse problems in any framework, especially in the large-scale case, can be computationally prohibitively expensive. To overcome this challenge, several classes of reduced-order and surrogate models have been introduced. In [39], methods to reduce the computational cost of solving inverse problems in the statistical setting are divided into three broad methods:

⁵ The truncation value r is chosen such that $\lambda_i \ll 1$ for $i > r$ [11, 12, 36].

reducing the cost of a forward simulation, reducing the dimension of the input parameter space, and reducing the number of samples required. In line with our goal of keeping the computational cost trackable, in this paper we focus on approximations of the posterior rather than on sampling-based methods.

As an approach, the Bayesian approximation error approach (BAE), developed in [18, 19], can be seen to lie in the first and second category. However, along with lowering computational costs of forward simulations and reducing the dimension of the parameter space, the BAE method can simultaneously take into account a vast array of uncertainties in the forward model, see, for example, [21–23, 40–46]. Below, we summarize the concept and implementation of the BAE method.

In the BAE approach, any errors induced by the use of surrogate models, reduction of the parameter dimension, and/or model uncertainties are propagated to a single additive error term. Hence the form of the posterior will be as in (15) with a redefined likelihood density. In what follows, let $\beta(\mathbf{x})$ be our parameter of interest, and take $a(\mathbf{x})$ to denote a secondary (nuisance) parameter. By secondary, we mean that we do not wish to estimate the unknown $a(\mathbf{x})$ but attempt to take the related uncertainty into account and propagate the effects into the estimate for the parameter of interest and the posterior uncertainty. Except for jointly normal linear models, it is not possible to *exactly premarginalize* over $a(\mathbf{x})$ [18, 23]. In the following, we outline how one can approximately premarginalize over $a(\mathbf{x})$. To this end, let

$$(a, \beta) \mapsto \mathbf{f}(a, \beta) \quad (25)$$

denote an accurate forward model, and let $\mathbf{e}$ again denote noise which is additive and mutually independent with β and a such that $\mathbf{e} \sim \mu_{\text{noise}} = \mathcal{N}(0, \mathbf{\Gamma}_e)$. Then the accurate relationship between the parameters and measurements is

$$\mathbf{d}^{\text{obs}} = \mathbf{f}(a, \beta) + \mathbf{e}. \quad (26)$$

In the BAE approach, rather than using the accurate model $\mathbf{f}(a, \beta)$, we instead set $a = a_*$ and use the approximate forward model

$$\beta \mapsto \mathbf{f}_{a_*}(\beta). \quad (27)$$

We note that in many applications the dimension of the parameter of interest is also reduced by projecting onto some reduced basis, see, for example, [23] for more details. In general, replacing the accurate model with the approximate model introduces what has become known as *model discrepancy*, the difference between the predictions of the two models. To take into account this model discrepancy we rewrite (26) as

$$\mathbf{d}^{\text{obs}} = \mathbf{f}(a, \beta) + \mathbf{e} = \mathbf{f}_{a_*}(\beta) + \mathbf{e} + \underbrace{(\mathbf{f}(a, \beta) - \mathbf{f}_{a_*}(\beta))}_{=\boldsymbol{\varepsilon}(a, \beta)} = \mathbf{f}_{a_*}(\beta) + \mathbf{e} + \boldsymbol{\varepsilon} = \mathbf{f}_{a_*}(\beta) + \boldsymbol{\nu}, \quad (28)$$

where the discrepancy in the models, $\boldsymbol{\varepsilon}$, is a random variable with the same dimensions as the measurements, and is known as the *approximation error* [18, 19, 40]. The sum $\boldsymbol{\nu} = \mathbf{e} + \boldsymbol{\varepsilon}$ is called the *total error* here.

At this point in the BAE approach, the following approximation is made:

$$\boldsymbol{\varepsilon}|\beta \sim \mu_{\boldsymbol{\varepsilon}|\beta} \approx \mathcal{N}(\boldsymbol{\varepsilon}_{*|\beta}, \mathbf{\Gamma}_{\boldsymbol{\varepsilon}|\beta}), \quad (29)$$

that is, the conditional density of the approximation error $\boldsymbol{\nu}$ given the parameter of interest β is approximated as Gaussian⁶. The computation of $\boldsymbol{\varepsilon}_{*|\beta}$ and $\mathbf{\Gamma}_{\boldsymbol{\varepsilon}|\beta}$ is outlined in section 5.

⁶ There is some work on retaining the full conditional density, see, for example, [47, 48].

The fact that the approximation error depends on β implies that formally ε and β cannot be taken as mutually independent. However, in several cases, such a further approximation of independence leads to similar estimates for a significantly smaller cost as explained in [23]. With this further approximation, we have

$$\nu_* = e_* + \varepsilon_* \quad \text{and} \quad \Gamma_\nu = \Gamma_e + \Gamma_\varepsilon, \quad (30)$$

which was originally referred to as the *enhanced error model* in [18, 19]⁷.

The BAE approach results in both a revised functional, which the MAP estimate minimizes, and a reformulated approximate posterior covariance matrix. Specifically, we now have

$$\beta_{\text{MAP}} = \min_{\beta \in \mathbb{R}^n} \frac{1}{2} \left\| f_{a_*}(\beta) - d^{\text{obs}} + \nu_* \right\|_{\Gamma_\nu^{-1}}^2 + \frac{1}{2} \|\mathbf{A}(\beta - \beta_*)\|_{\mathbf{M}}^2, \quad (31)$$

$$\Gamma_{\text{post}} = \left(F_{a_*}^{\text{H}} \Gamma_\nu^{-1} F_{a_*} + \Gamma_\beta^{-1} \right)^{-1}. \quad (32)$$

We note that the infinite-dimensional counterparts to (31) and (32) can be formulated naturally.

As an indicator as to whether or not inclusion of the approximation errors is appropriate the following rule of thumb can be adopted [23]: if

$$\|e_*\|^2 + \text{trace}(\Gamma_e) < \|\varepsilon_*\|^2 + \text{trace}(\Gamma_\varepsilon) \quad (33)$$

holds, then the approximation errors *dominate* the noise and neglecting the approximation errors will generally result in meaningless reconstructions, as demonstrated in section 5.2. Moreover, if

$$e_*^2(k) + \Gamma_e^2(k, k) < \varepsilon_*^2(k) + \Gamma_\varepsilon^2(k, k) \quad (34)$$

for any k , then neglecting the approximation errors can still lead to meaningless results [23].

4. Recovery of the Robin coefficient

In this section, we formulate the inverse Robin problem with a spatially varying conductivity coefficient $a(\mathbf{x})$ which will later be interchanged for a fixed conductivity a_* . In the chosen geometry, we refer to the Robin coefficient as ‘basal’ since this condition is posed only on the bottom part of a slab. The measurements are taken to be pointwise (noisy) potential measurements on the top of the domain. We solve the inverse problem with Newton’s method. Therefore, in what follows, we formulate the forward problem and derive the corresponding first and second order adjoint problems for the gradient and the action of the Hessian to a vector needed by the optimization method.

4.1. The forward problem

As a model problem, let $\Omega = [0, L] \times [0, L] \times [0, H] \in \mathbb{R}^3$ with $0 < H \ll L < \infty$ denote the domain of the problem (a thin slab) with boundary $\partial\Omega$. In our regime, a flux is prescribed on $\Gamma_t := [0, L] \times [0, L] \times \{H\}$, while the potential is measured at q points on

⁷ The actual form of $\varepsilon_{*|\beta}$ is $\varepsilon_{*|\beta} = \varepsilon_* + \Gamma_{\varepsilon\beta} \Gamma_{\beta\beta}^{-1}(\beta - \beta_*)$ which incorporates the full covariance structure of the normal approximation for $\pi(\varepsilon, \beta)$. The prior covariance of β cannot, however, be used in place of $\Gamma_{\beta\beta}$ in this conditional expectation. Rather, it must be based on the same draws as those used to compute $\Gamma_{\varepsilon\varepsilon}$ as in section 5 below. In practise, this leads to a semidefinite estimate for $\Gamma_{\beta\beta}$, and the associated rank-deficient forms for the conditional expectations need to be employed [23].




Figure 1. Set up for the model problem. A prescribed flux is set through the top of the domain and measurements of the potential are taken at points on the top of the domain (blue circles). A Robin boundary condition is enforced at the bottom of the domain, while the sides of the domain are prescribed homogeneous Dirichlet boundary conditions.

Γ_t , see figure 1 for a schematic representation. A homogenous Robin boundary condition is prescribed on $\Omega_\beta = \Gamma_b := [0, L] \times [0, L] \times \{0\}$ while on the remainder of the boundary, $\Gamma_s := \partial\Omega \setminus (\Gamma_t \cup \Gamma_b)$, homogeneous Dirichlet boundary conditions are specified. The conductivity coefficient $\exp(a(\mathbf{x}))$ is taken to be a spatially distributed random field in Ω , while the Robin coefficient is taken to be a spatially varying random field on Γ_b . To summarize, the forward problem reads

$$\begin{aligned} -\nabla \cdot (\exp(a(\mathbf{x})) \nabla u(\mathbf{x})) &= 0 && \text{in } \Omega \\ \exp(a(\mathbf{x})) \nabla u(\mathbf{x}) \cdot \mathbf{n}_t &= g(\mathbf{x}) && \text{on } \Gamma_t \\ \exp(a(\mathbf{x})) \nabla u(\mathbf{x}) \cdot \mathbf{n}_b + \exp(\beta(\mathbf{x})) u(\mathbf{x}) &= 0 && \text{on } \Gamma_b, \\ u(\mathbf{x}) &= 0 && \text{on } \Gamma_s \end{aligned} \quad (35)$$

where $u(\mathbf{x})$ the potential, $g(\mathbf{x})$ is the flux through Γ_t with unit normal $\mathbf{n}_t$, and Γ_b has unit normal $\mathbf{n}_b$. We employ the finite element method (FEM) for the numerical approximation of the forward problem, with the standard Lagrange piecewise linear nodal basis functions.

4.2. Adjoint-based Gauss–Newton method for solving the inverse problem

We employ an inexact Newton–CG approach to solve the minimization problem (17) which requires both the gradient and the Hessian of the negative log prior and likelihood. To avoid calculations of forward sensitivities, which would require as many forward solves as the dimension of the parameter, we employ the adjoint approach [14, 15, 49, 50] to compute the derivatives, which we show next.

We denote the observation operator with $\mathcal{B}$ so that the parameter-to-observable mapping can be written as $\mathbf{f}_{a_s} = \mathcal{B}u$. Hence the infinite-dimensional counterpart of the functional to be minimized in (17) can be rewritten as

$$\mathcal{J}(\beta) = \frac{1}{2} \left\| \mathcal{B}u - \mathbf{d}^{\text{obs}} + \boldsymbol{\nu}_* \right\|_{\Gamma_t^{-1}}^2 + \frac{1}{2} \left\| \mathcal{A}\mathcal{W}^{-1}(\beta - \beta_*) \right\|_{L^2(\Gamma_b)}^2, \quad (36)$$

where $u(\mathbf{x})$ solves the forward problem (35). Furthermore, let us define the space,

$$\mathcal{V} := \{v \in H^1(\Omega) : v|_{\Gamma_s} = 0\}, \quad (37)$$

then we can define the Lagrangian functional $\mathcal{L} : \mathcal{V} \times \mathcal{V} \times \mathcal{E} \rightarrow \mathbb{R}$,

$$\mathcal{L}(u, p, \beta) := \mathcal{J}(\beta) + \int_{\Omega} \exp(a(\mathbf{x})) \nabla u \cdot \nabla p \, d\mathbf{x} - \int_{\Gamma_t} gp \, d\mathbf{s}_t + \int_{\Gamma_b} \exp(\beta) up \, d\mathbf{s}_b. \quad (38)$$

The space $\mathcal{E}$ is the Cameron–Martin space $\mathcal{E} = \text{range}(\mathcal{C}^{\frac{1}{2}}) = \text{dom}(\mathcal{A})$, induced by the prior measure, see [13] for full details, or for example [17] for a brief overview. We note that due to the regularity of the boundary of Γ_b in the current set up, $\mathcal{E} = H^2(\Gamma_b)$ [51, 52].

Determining the gradient of $\mathcal{J}$ is achieved by requiring that variations of the Lagrangian $\mathcal{L}$ with respect to the forward potential u and the so-called *adjoint potential* p vanish. This results in the following strong form of the gradient $\mathcal{G}$ for the variations with respect to β

$$\mathcal{G}(\beta) := \mathcal{W}^{-1} \mathcal{A}^2 \mathcal{W}^{-1} (\beta - \beta_*) + \exp(\beta) u p, \quad (39)$$

with u being the solution of the *forward Poisson problem* (35) for given β , while p satisfies the following *adjoint Poisson problem* for given $u(\mathbf{x})$ and $\beta(\mathbf{x})$

$$\begin{aligned} -\nabla \cdot (\exp(a(\mathbf{x})) \nabla p(\mathbf{x})) &= -\mathcal{B}^* \Gamma_{\nu}^{-1} (\mathcal{B} u(\mathbf{x}) - \mathbf{d}^{\text{obs}} + \nu_*) && \text{in } \Omega, \\ \exp(a(\mathbf{x})) \nabla p(\mathbf{x}) \cdot \mathbf{n}_t &= 0 && \text{on } \Gamma_t, \\ \exp(a(\mathbf{x})) \nabla p(\mathbf{x}) \cdot \mathbf{n}_b + \exp(\beta(\mathbf{x})) p(\mathbf{x}) &= 0 && \text{on } \Gamma_b, \\ p(\mathbf{x}) &= 0 && \text{on } \Gamma_s. \end{aligned} \quad (40)$$

The action of the Gauss–Newton approximation of the Hessian operator evaluated at β in the direction $\hat{\beta}$ is given by

$$\mathcal{H}(\beta)(\hat{\beta}) := \mathcal{W}^{-1} \mathcal{A}^2 \mathcal{W}^{-1} \hat{\beta} + \exp(\beta) \hat{\beta} u \hat{p}, \quad (41)$$

where the *incremental adjoint potential* $\hat{p}$ satisfies the so-called *incremental adjoint (or second order adjoint) Poisson problem*

$$\begin{aligned} -\nabla \cdot (\exp(a(\mathbf{x})) \nabla \hat{p}(\mathbf{x})) &= -\mathcal{B}^* \Gamma_{\nu}^{-1} \mathcal{B} \hat{u}(\mathbf{x}) && \text{in } \Omega \\ \exp(a(\mathbf{x})) \nabla \hat{p}(\mathbf{x}) \cdot \mathbf{n}_t &= 0 && \text{on } \Gamma_t \\ \exp(a(\mathbf{x})) \nabla \hat{p}(\mathbf{x}) \cdot \mathbf{n}_b + \exp(\beta(\mathbf{x})) \hat{p}(\mathbf{x}) &= 0 && \text{on } \Gamma_b, \\ \hat{p}(\mathbf{x}) &= 0 && \text{on } \Gamma_s, \end{aligned} \quad (42)$$

and the *incremental forward potential* $\hat{u}$ satisfies the *incremental forward Poisson problem*

$$\begin{aligned} -\nabla \cdot (\exp(a(\mathbf{x})) \nabla \hat{u}(\mathbf{x})) &= 0 && \text{in } \Omega \\ \exp(a(\mathbf{x})) \nabla \hat{u}(\mathbf{x}) \cdot \mathbf{n}_t &= 0 && \text{on } \Gamma_t \\ \exp(a(\mathbf{x})) \nabla \hat{u}(\mathbf{x}) \cdot \mathbf{n}_b + \exp(\beta(\mathbf{x})) \hat{u}(\mathbf{x}) &= -\hat{\beta} \exp(\beta(\mathbf{x})) u(\mathbf{x}) && \text{on } \Gamma_b, \\ \hat{u}(\mathbf{x}) &= 0 && \text{on } \Gamma_s. \end{aligned} \quad (43)$$

The resulting system to be solved for the Gauss–Newton search direction, $\hat{\beta}$, is

$$\mathcal{H}(\beta)(\hat{\beta}) = -\mathcal{G}(\beta). \quad (44)$$

In the case of a high-dimensional parameter, it may not be feasible to solve the system (44) directly, and iterative methods are usually employed. To this end, we employ the conjugate gradient (CG) method, which requires only the action of the (discrete) Hessian, that is, Hessian-vector products. By examining (41), we see that each Hessian-vector product amounts to solving both the adjoint Poisson problem (43) and the incremental adjoint Poisson problem (42). Moreover, the computation of the gradient (39) requires the solutions of the forward Poisson problem (35) and the adjoint Poisson problem (40). The dominant cost of solving (44) is in computing the solution to Poisson equations, and thus the computational cost of the method can be roughly measured in number of Poisson problems solved.

Carrying out the premarginalization over a avoids solving for a since we replace $a(\mathbf{x})$ with $a_*(\mathbf{x})$. Indeed, solving for a would require the computation of another search direction, which in the Gauss–Newton case would be computed by solving

$$\mathcal{H}_a(a)(\hat{a}) = -\mathcal{G}_a(a) \quad (45)$$

for $\hat{a}$, with

$$\mathcal{G}_a(a) := \mathcal{A}_a^2(a - a_*) + \exp(a)\nabla u \cdot \nabla p \quad \text{and} \quad \mathcal{H}_a(a)(\hat{a}) := \mathcal{A}_a^2\hat{a} + \exp(a)\hat{a}\nabla u \cdot \nabla \hat{p} \quad (46)$$

the gradient and Hessian respectively, and $\mathcal{A}_a$ is defined in (47). In addition, the system (45) would need to be solved over the entire domain Ω , necessitating an extra order of magnitude of computational complexity. Such an approach is outlined for the nonlinear Stokes flow problem in [53].

5. Numerical examples

In this section, we consider two numerical experiments, one with a conductivity with isotropic homogeneous covariance structure and one with an anisotropic structure. The latter structure is akin to horizontally layered (stochastic) strata in which the correlation length is smaller in the vertical direction than in the horizontal plane. We will pay particular attention to the feasibility of the posterior error estimates, that is, we will investigate whether the posterior models (essentially) support the actual Robin coefficient.

5.1. Problem setup

In both experiments, the domain $\Omega \in \mathbb{R}^3$ is a rectangular parallelepiped with thickness $H = 0.01$ and width $L = 1$, such that $L/H = 100$. The measurements consists of $q = 33$ point measurements on the top of the domain, as illustrated in figure 1. To avoid the so-called inverse crime, we use a finer FEM discretization to generate the synthetic data than the FEM discretization used in the inversions. Moreover, the mesh used in the second example to generate the data is finer than the corresponding mesh used in the first example to ensure refinement of the stratified conductivity in the volume. The details of the meshes are presented in table 1, in all cases Lagrange piece-wise linear basis functions are used.

In both numerical examples, zero mean white noise is added to the simulated measurements, with the noise covariance matrix given by $\mathbf{\Gamma}_e = \delta_e^2 \mathbf{I}$, with $\delta_e = (\max(\mathcal{B}\mathbf{u}) - \min(\mathcal{B}\mathbf{u})) \times 1/100$, that is, the noise level is 1% of the range of the noiseless measurements.

5.1.1. Prior models. The prior density (normal random field) imposed on β is the same in both experiments, as outlined in section 2. We assign the parameters that fix the mean β_* and the covariance operator $\mathcal{C}_\beta$ as follows: $\beta_* = 1$, $\alpha_\beta = 7$, $\gamma_\beta = 0.01\mathbf{I}$ and $\kappa_\beta = 0$. On the far right of figure 2, we show the resulting spatial variance structure of $\mathbf{\Gamma}_\beta$ with the weighting. For comparison we also show the typically implemented case of homogeneous Neumann boundary conditions without weighting (far left), the case of enforcing Dirichlet boundary conditions (centre left), and the case of applying a homogeneous Robin boundary condition following the method of [32] (centre right). We note that the weighted covariance approach nullifies all boundary effects. Figure 3 shows three samples drawn from the prior density, μ_β , along with the actual (distributed) Robin coefficient β_{true} used to generate the synthetic data in both experiments.




Figure 2. The diagonal of the prior covariance operator with $\alpha_\beta = 7$ and $\gamma_\beta = 0.01$. Far left: with homogeneous Neumann boundary conditions. Centre left: with homogeneous Dirichlet boundary conditions. Centre right: with homogeneous Robin boundary conditions and $\kappa_\beta = 1.42\sqrt{\gamma_\beta/\alpha_\beta}$ as in [32]. Far right: the weighted approach of the current paper as discussed in section 2, with $\kappa_\beta = 0$ (homogeneous Neumann).




Figure 3. Far left to centre right: samples from the prior distribution on β . Right: the true value β_{true} used to compute the data.

The BAE approach is (in part) based on (approximate) marginalization over the nuisance parameter. Technically, this involves drawing samples from the joint prior density of the nuisance and the primary parameters $\pi(a, \beta)$ to compute the second order statistics of (ε, β) which, in turn, involves the computation of the forward problem $f(a, \beta)$ for the draws. In this paper, we take (a, β) to be mutually independent. In addition, we take the prior on $a(\mathbf{x})$ to be a Gaussian measure, $\mu_a = \mathcal{N}(a_*, \mathcal{C}_a)$ on $L^2(\Omega)$ with $\mathcal{C}_a$ defined similarly as $\mathcal{C}_\beta$ in section 2. To be precise, we use a squared inverse elliptic operator as our prior covariance operator with homogeneous Neumann boundary conditions. We neglect to weight the covariance operator $\mathcal{C}_a$ by the diagonal of the associated solution operator and simply impose homogeneous Neumann boundary conditions in a bid to reduce the computational cost. Any boundary effects caused by this are of no consequence as we do not wish to reconstruct a . Formally, the prior covariance operator of a is defined as $\mathcal{C}_a = \mathcal{A}_a^{-2}$, with the operator $\mathcal{A}_a$ defined (similarly to $\mathcal{C}_\beta$ in section 2.1) through the variational problem: for $s \in L^2(\Omega)$, the solution of $\mathcal{A}_a a = s$ satisfies

$$\alpha_a \int_{\Omega} (\gamma_a \nabla a \cdot \nabla v + av) \, d\mathbf{x} = \int_{\Omega} sv \, d\mathbf{x} \quad \text{for all } v \in H^1(\Omega). \quad (47)$$

For the first numerical example with isotropic correlation structure, we use $\alpha_a = 100$ and $\gamma_a = 10^{-3}\mathbf{I}$, while for the second numerical example with anisotropic correlation structure, we take $\gamma_a = \text{diag}(10^{-2}, 10^{-2}, 10^{-8})$. In figure 4, three samples drawn from μ_a for the first




Figure 4. Far left to centre right: slice plots of samples from the prior distribution on $a(\mathbf{x})$ with isotropic covariance structure (first numerical example). Right: slice plot of the true value a_{true} used to compute the data.




Figure 5. Far left to centre right: slice plot of samples from the prior distribution on a with anisotropic covariance structure (second numerical example). Right: slice plot of the true value a_{true} used to compute the data.

numerical example are shown, along with the true value, a_{true} , used to generate the synthetic data. Similarly, in figure 5, three samples drawn from μ_a for the second numerical example are shown along with the true value a_{true} used to generate the synthetic data. A standard requirement when designing the prior is that the priors should not be ‘too narrow’, which here is reflected in the shown draws when compared to the actual $a(\mathbf{x})$ and $\beta(\mathbf{x})$.

5.1.2. Estimation of approximation error statistics. In the linear normal case, that is, $\mathbf{f}(a, \beta)$ and $\mathbf{f}_{a_*}(\beta)$ both linear and $\pi(a, \beta)$ is normal, $\boldsymbol{\varepsilon}_*$ and $\Gamma_{\boldsymbol{\varepsilon}}$ can be computed analytically. If this is not the case, both $\boldsymbol{\varepsilon}_*$ and $\Gamma_{\boldsymbol{\varepsilon}}$ must be estimated using sample statistics using samples drawn from (the not necessarily jointly Gaussian) joint prior model $\pi(a, \beta)$. With an ensemble of r samples, $(\beta^{(\ell)}, a^{(\ell)})$ from the associated prior densities, we compute

$$\boldsymbol{\varepsilon}^{(\ell)} = \mathbf{f}(a^{(\ell)}, \beta^{(\ell)}) - \mathbf{f}_{a_*}(\beta^{(\ell)}), \quad \ell = 1, 2, \dots, r, \quad (48)$$

and take the mean and covariance as

$$\boldsymbol{\varepsilon}_* = \frac{1}{r} \sum_{\ell=1}^r \boldsymbol{\varepsilon}^{(\ell)} \quad \text{and} \quad \Gamma_{\boldsymbol{\varepsilon}} = \frac{1}{r-1} \sum_{\ell=1}^r (\boldsymbol{\varepsilon}^{(\ell)} - \boldsymbol{\varepsilon}_*)(\boldsymbol{\varepsilon}^{(\ell)} - \boldsymbol{\varepsilon}_*)^T. \quad (49)$$

The number of samples r required depends on the models, the variance of the approximation error, and the joint prior model, see for example [42]. However, we remark that all samples




Figure 6. Second order statistics of the noise and approximation errors for the first (isotropic conductivity) numerical example. Far left: the density of the noise, μ_{noise} (mean shown in red) with higher probability density indicated by darker shading. Centre left: the covariance matrix of the noise Γ_e . Centre right: the density of the total errors, μ_ν (mean shown in red). Far right: the covariance matrix of the approximation errors Γ_ϵ .




Figure 7. Second order statistics of the noise and approximation errors for the second (anisotropic conductivity) numerical example. Far left: the density of the noise, μ_{noise} (mean shown in red) with higher probability density indicated by darker shading. Centre left: the covariance matrix of the noise Γ_e . Centre right: the density of the total errors, μ_ν (mean shown in red) with higher probability density indicated by darker shading. Far right: the covariance matrix of the approximation errors Γ_ϵ .

(and all accurate forward simulations) are carried out at the offline stage and that the accurate forward model is never used in the inversion, with only the approximate model evaluated at the online stage. For the current problem, 1000 samples were drawn for both numerical examples to compute the approximation error statistics. Figures 6 and 7 show the statistics of the noise and of the approximation errors for the first and second examples, respectively. It is clear that the approximation errors dominate the noise in both cases, entries on the diagonal of Γ_ϵ being almost two orders of magnitude larger than those on the diagonal of Γ_e . Thus, not including the approximation errors will most likely lead to meaningless results, as confirmed below in section 5.2. Note the difference in the mean and covariance of ϵ due to different priors on α .

5.1.3. The MAP and approximate posterior covariance estimates. To compare the solutions of the inverse problems formulated with the approximation error noise model and with the conventional error model, we compute the respective MAP estimates and posterior covariances matrices. For both the isotropic and anisotropic conductivity cases, we compute the following three MAP estimates:

- MAP-REF: the reference maximum *a posteriori* estimate for β with the conventional error model and use of the *actual* value of a in the model, that is, $\mathbf{d}^{\text{obs}} = \mathbf{f}(a_{\text{true}}, \beta) + \mathbf{e} = \mathbf{f}_{a_{\text{true}}}(\beta) + \mathbf{e}$. This reconstruction is taken as the benchmark one, as it is computed with no modelling errors present. The estimate is computed as

$$\beta_{\text{MAP}}^{\text{REF}} = \min_{\beta \in \mathbb{R}^n} \frac{1}{2} \left\| \mathbf{f}_{a_{\text{true}}}(\beta) - \mathbf{d}^{\text{obs}} \right\|_{\Gamma_e^{-1}}^2 + \frac{1}{2} \|\mathbf{A}(\beta - \beta_*)\|_{\mathbf{M}}^2. \quad (50)$$

- MAP-CEM: the maximum *a posteriori* estimate for β with the conventional error model, using the (incorrect) fixed $a = a_*$ in the forward model, that is, $\mathbf{d}^{\text{obs}} = \mathbf{f}_{a_*}(\beta) + \mathbf{e}$. This estimate is computed as

$$\beta_{\text{MAP}}^{\text{CEM}} = \min_{\beta \in \mathbb{R}^n} \frac{1}{2} \left\| \mathbf{f}_{a_*}(\beta) - \mathbf{d}^{\text{obs}} \right\|_{\Gamma_e^{-1}}^2 + \frac{1}{2} \|\mathbf{A}(\beta - \beta_*)\|_{\mathbf{M}}^2. \quad (51)$$

- MAP-BAE: the maximum *a posteriori* estimate for β with the approximation error model and using fixed $a = a_*$ in the model, that is, $\mathbf{d}^{\text{obs}} = \mathbf{f}_{a_*}(\beta) + \nu$. This estimate is computed as

$$\beta_{\text{MAP}}^{\text{BAE}} = \min_{\beta \in \mathbb{R}^n} \frac{1}{2} \left\| \mathbf{f}_{a_*}(\beta) - \mathbf{d}^{\text{obs}} + \nu_* \right\|_{\Gamma_\nu^{-1}}^2 + \frac{1}{2} \|\mathbf{A}(\beta - \beta_*)\|_{\mathbf{M}}^2. \quad (52)$$

The related approximate posterior covariance matrices are then as follows.

- $\Gamma_{\text{post}}^{\text{REF}}$: the reference posterior covariance matrix is computed using the conventional error model and using the actual value of a in the model, i.e. $\mathbf{d}^{\text{obs}} = \mathbf{f}(a_{\text{true}}, \beta) + \mathbf{e}$. Since, in this case, there are no modelling errors present, we expect the reference posterior covariance matrix to be smaller (in the sense of quadratic forms) than the posterior covariance matrix obtained using the approximation error model. This posterior covariance matrix is

$$\Gamma_{\text{post}}^{\text{REF}} = \left(\mathbf{F}_{a_{\text{true}}}^\dagger \Gamma_e^{-1} \mathbf{F}_{a_{\text{true}}} + \Gamma_\beta^{-1} \right)^{-1}. \quad (53)$$

- $\Gamma_{\text{post}}^{\text{CEM}}$: the posterior covariance matrix with the conventional error model is computed using the conventional error model with the fixed $a = a_*$ in the model, that is, $\mathbf{d}^{\text{obs}} = \mathbf{f}_{a_*}(\beta) + \mathbf{e}$. This posterior covariance matrix is

$$\Gamma_{\text{post}}^{\text{CEM}} = \left(\mathbf{F}_{a_*}^\dagger \Gamma_e^{-1} \mathbf{F}_{a_*} + \Gamma_\beta^{-1} \right)^{-1}. \quad (54)$$

- $\Gamma_{\text{post}}^{\text{BAE}}$: the approximation error model posterior covariance matrix is computed using the approximation error model and using fixed $a = a_*$ in the model, i.e. $\mathbf{d}^{\text{obs}} = \mathbf{f}_{a_*}(\beta) + \nu$. This posterior covariance matrix is

$$\Gamma_{\text{post}}^{\text{BAE}} = \left(\mathbf{F}_{a_*}^\dagger \Gamma_\nu^{-1} \mathbf{F}_{a_*} + \Gamma_\beta^{-1} \right)^{-1}. \quad (55)$$

5.2. Results

The computation of the reference MAP estimate, $\beta_{\text{MAP}}^{\text{REF}}$ given in (50), the conventional error model MAP estimate, $\beta_{\text{MAP}}^{\text{CEM}}$ given in (51), and the approximation error model MAP estimate, $\beta_{\text{MAP}}^{\text{BAE}}$ given in (52), is done by applying an inexact adjoint-based Gauss–Newton method outlined in section 4.2. We start each of the optimization procedures with the prior mean as the initial guess, that is $\beta_0 = \beta_* = 1$. A preconditioned conjugate gradient (CG) method is used




Figure 8. Isotropic conductivity. Far left: the true Robin coefficient β_{true} . Centre left to far right: the reference MAP estimate $\beta_{\text{MAP}}^{\text{REF}}$, the conventional error model MAP estimate $\beta_{\text{MAP}}^{\text{CEM}}$ and the approximation error model MAP estimate $\beta_{\text{MAP}}^{\text{BAE}}$, respectively. The cross sections are shown in figures 9 and 10.

with an Einsentat–Walker condition [54] which terminates the CG iterations early, when the norm of the gradient is sufficiently reduced. In line with [11], we use the prior operator as a preconditioner for the CG iterations.

5.2.1. The estimates with isotropic conductivity. On the far left of figure 8 we show the true basal Robin coefficient, β_{true} , which is used to generate the measurements. We also show the reconstructed reference MAP estimate (centre left), the reconstructed conventional error model MAP estimate (centre right), and the reconstructed approximation error model MAP estimate (far right). The images of the reconstructions in figure 8 also show the (dotted) lines p - p^* and q - q^* which are the locations of the cross sections shown in figures 9 and 10. We now discuss several observations that can be made from figures 9 and 10. Firstly, the reference estimate is clearly *feasible* in the sense that the posterior uncertainty supports the actual Robin coefficient. Conversely, the estimates with the conventional error model (severely) underestimate the true basal Robin coefficient. In particular, the estimate is clearly *infeasible*: the actual coefficient has almost vanishing posterior density at almost all points along the cross sections. On the other hand, the estimate with the approximation error model is clearly a feasible one with the posterior marginals supporting the actual Robin coefficient. Finally, we see that marginalization over the conductivity results in the widening of the posterior density which is evident when comparing the marginal densities of the reference and approximation error estimates.

5.2.2. The estimates with anisotropic conductivity. The corresponding results in the case of the anisotropic conductivity are shown in figures 11–13. The results are qualitatively similar to the isotropic case. In this case, the conventional error severely overestimates the actual Robin coefficient, the only difference between the two cases being the spatial covariance structure of the conductivity a . With the approximation error model, the estimates are still feasible but slightly worse when comparing to the isotropic conductivity case.

5.2.3. Computational costs. In this section, we compare the computational cost of the inverse solution method applied for the three methods: the reference case; the conventional error procedure; and the BAE approach. We measure this cost in terms of number of Poisson solves needed for the optimization algorithm to converge. We note that to compute the MAP estimates, the number of Poisson solves needed per Gauss–Newton iteration can be calculated




Figure 9. Isotropic conductivity. Far left: the prior mean and marginal distributions of the β along the line p - p^* . Centre left to far right: the cross sections of figure 8 along the line p - p^* . The reference MAP estimate $\beta_{\text{MAP}}^{\text{REF}}$, the conventional error model MAP estimate $\beta_{\text{MAP}}^{\text{CEM}}$ and the approximation error model MAP estimate $\beta_{\text{MAP}}^{\text{BAE}}$. The true β and the MAP estimates are shown in blue and red, respectively, along with the approximate posterior marginal distributions of β .




Figure 10. Isotropic conductivity. The cross sections of figure 8 along the line q - q^* . The description otherwise as in figure 9.

as $\# \text{Poisson} = 2 + 2\# \text{CG} + \# \text{back}$, where $\# \text{CG}$ is the number of CG iterations and $\# \text{back}$ is the number of back-tracks needed to get a sufficient decrease in the objective function. In all three cases, the convergence of the Gauss–Newton iterations is established when the norm of the gradient (relative to the initial norm of the gradient) is decreased by a factor of 10^7 . The results shown in table 2 reveal that, at the inversion stage, the BAE approach is approximately as expensive as the conventional error approach. For completeness, we also show the costs in the reference reconstructions.

5.2.4. Interpretation of the posterior covariance matrices. Here we discuss and compare the three posterior covariance matrices corresponding to the reference case, the conventional error with reduced model case, and the approximation error model case, which are defined in (53)–(55). We begin the discussion by analyzing the spectrum of the respective prior-preconditioned data misfit Hessian components of the posterior covariance matrices for both numerical examples.

On the left of figure 14 is shown the dominant spectrum of the prior-preconditioned data misfit Hessian for the three cases evaluated at the respective MAP estimates for the first example, while the same results for the second example are shown on the right (reference case with blue circles, conventional error with reduced model with red diamonds, and approximation error model with yellow crosses). In all three cases we are only required to retain relatively few eigenvalues to compute a reasonable low rank approximation of the Hessian. Specifically, in the first example, for the reference case we need about 30 eigenvalues, while for both the

Table 1. Comparison of the computational costs. The cost of solving for the MAP estimates in the reference case (REF), the case of using the conventional error model (CEM) with fixed $a = a_*$ and the case of using the approximation error model (BAE) with fixed $a = a_*$, measured in number of Poisson solves. The first column (MAP) refers to which MAP estimate we are solving for, MAP-REF (50), MAP-CEM (51), or MAP-BAE (52); the second column (#GN) reports the number of Gauss–Newton iterations; the third (#CG) and fourth (avg.CG) columns show total and the average (per Gauss–Newton iteration) number of CG iterations; the fifth column (#back) reports the total number of backtracks needed throughout the Gauss–Newton iterations; and the last column (#Poisson) reports the total number of Poisson solves (from forward, adjoint, and incremental forward and adjoint problems). The Gauss–Newton iterations are terminated when the norm of the gradient is decreased by a factor of 10^7 , while the CG iterations are terminated inline with the Einstat–Walker condition [54]. These results illustrate that the use of the approximation error approach can be carried out at no additional cost compared to the conventional error approach and reference case.

MAP	#GN	#CG	avg.CG	#back	#Poisson
Example 1					
REF	8	117	15	0	250
CEM	11	101	10	4	228
BAE	5	57	12	0	124
Example 2					
REF	6	54	9	0	120
CEM	8	95	12	0	206
BAE	5	97	20	0	204




Figure 11. Anisotropic conductivity. Far left: the true Robin coefficient β_{true} . Centre left to far right: the reference MAP estimate $\beta_{\text{MAP}}^{\text{REF}}$, the conventional error model MAP estimate $\beta_{\text{MAP}}^{\text{CEM}}$ and the approximation error model MAP estimate $\beta_{\text{MAP}}^{\text{BAE}}$, respectively. The cross sections are shown in figures 12 and 13.

conventional error model with reduced model and the approximation error model we need about 20. In the second example the reference case requires the retention of about 20 eigenvalues, the conventional error model with reduced model requires about 30, and the reduced model with approximation error model we need 25. We note that these low numerical ranks are substantially smaller than the 961 degrees of freedom of the parameter (i.e. we see a compression of the parameter dimension of about 30). Hence the approximate posterior covariance matrix along with draws from the posterior can be cheaply computed using (23) and (24), respectively.

In figures 15 and 16 we show several eigenvectors corresponding to the dominant eigenvalues of the prior-preconditioned data misfit Hessians corresponding to the three cases.




Figure 12. Anisotropic conductivity. The cross sections of figure 11 along the line q - q^* . The description otherwise as in figure 9.




Figure 13. Anisotropic conductivity. The cross sections of figure 11 along the line q - q^* . The description otherwise as in figure 9.




Figure 14. Semi-log plots of the eigenvalues of the prior-preconditioned misfit Hessians for the isotropic case (left) and the anisotropic case (right).

The most dominant eigenvalues can be interpreted as the modes in the basal Robin coefficient for which the data contains the most information about. The first few eigenvectors of all three cases are fairly similar (note the sign change for the first eigenvector is the approximation error model case). However as the level of oscillation in the eigenvectors increases we see that the differences between corresponding eigenvectors between the three models increase.

We can further assess the uncertainty in the estimates by analyzing the full pointwise posterior variances, i.e. the diagonal of the posterior covariance matrices. Such analysis also provides insight into how much the variance (from the prior to the posterior) is reduced by taking the data into consideration. In figures 17 and 18 we show the diagonal of the prior covariance




Figure 15. Eigenvectors of the reference prior-preconditioned misfit Hessian (top row), conventional error model prior-preconditioned misfit Hessian (middle row), and approximation error model prior-preconditioned misfit Hessian (bottom row) computed at $\beta_{\text{MAP}}^{\text{REF}}$, $\beta_{\text{MAP}}^{\text{CEM}}$, and $\beta_{\text{MAP}}^{\text{BAE}}$, respectively. From left to right: the eigenvectors corresponding to the first (i.e. the largest), the third, the fifth and the tenth eigenvalues.

matrix and the three posterior covariance matrices, for the first and second numerical examples, respectively.

Most information, and consequently the greatest reduction in uncertainty, is directly below the measurement locations and in areas where the parameter field attains relatively high values. The reduction in variance in the vicinity of the measurement locations is typical for inverse problems. On the other hand, it is evident that the reduction in variance is greatest where the parameter achieves higher values. This is due to the fact that for higher values of the parameter the Robin boundary condition begins to behave like a Dirichlet boundary condition, meaning the potential is less free to vary which leads to less variance in the inferred parameter in these regions.

Finally, the centre images in figures 17 and 18 further illustrate the extent to which the conventional error model leads to overly optimistic (narrow) confidence intervals. This feature is especially evident in the anisotropic case (figure 18), where the posterior variance using the conventional error model is significantly smaller than the reference posterior variance.

6. Discussion

In this paper, we considered the problem of inferring the distributed (basal) Robin coefficient from surface measurements under an unknown random conductivity field. The forward




Figure 16. Eigenvectors of the reference prior-preconditioned misfit Hessian (top row), conventional error model prior-preconditioned misfit Hessian (middle row), and approximation error model prior-preconditioned misfit Hessian (bottom row) computed at $\beta_{\text{MAP}}^{\text{REF}}$, $\beta_{\text{MAP}}^{\text{CEM}}$ and $\beta_{\text{MAP}}^{\text{BAE}}$, respectively. From left to right: the eigenvectors corresponding to the first (i.e. the largest), the third, the fifth and the tenth eigenvalues.




Figure 17. The reduction in variance for the isotropic case. Far left: the diagonal of the prior covariance matrix, $\mathbf{\Gamma}_\beta$, as outlined in section 2. Centre left: the diagonal of the reference approximate posterior covariance matrix, $\mathbf{\Gamma}_{\text{post}}^{\text{REF}}$. Centre right: the diagonal of the conventional error model approximate posterior covariance matrix, $\mathbf{\Gamma}_{\text{post}}^{\text{CEM}}$. Far right: the diagonal of the approximation error model approximate posterior covariance matrix, $\mathbf{\Gamma}_{\text{post}}^{\text{BAE}}$.




Figure 18. The reduction in variance for the anisotropic case. Far left: the diagonal of the prior covariance matrix, $\mathbf{\Gamma}_\beta$, as outlined in section 2. Centre left: the diagonal of the reference approximate posterior covariance matrix, $\mathbf{\Gamma}_{\text{post}}^{\text{REF}}$. Centre right: the diagonal of the conventional error model approximate posterior covariance matrix, $\mathbf{\Gamma}_{\text{post}}^{\text{CEM}}$. Far right: the diagonal of the approximation error model approximate posterior covariance matrix, $\mathbf{\Gamma}_{\text{post}}^{\text{BAE}}$.

Table 2. Statistics of the mesh. The first column (Mesh used for) relates what the mesh is used for; the second (#Nodes) third (#Els) and fourth (#Param) columns give the number of total number of FEM nodes used in the entire volume, tetrahedral elements used in the entire volume, and the number nodes on the domain of the parameter, β , respectively.

	Mesh used for	#Nodes	#Els	#Param
Example 1				
	Data synthesis	28 611	150 000	2601
	Inversion	6727	32 400	961
Example 2				
	Data synthesis	132 651	750 000	2601
	Inversion	29 791	162 000	961

model at hand was the (anisotropic) Poisson equation with mixed boundary conditions. To account for model errors that stem from the uncertainty in the conductivity coefficient in the underlying PDE, we carry out approximative marginalization over the conductivity. In this process, we approximate the related modelling errors and uncertainties as normal, which is also referred to as the Bayesian approximation error (BAE) approach.

The uncertainty analysis presented here relies on a local linearization of the parameter-to-observable maps at the MAP point estimates, leading to a normal (Gaussian) approximation of the parameter posterior density, which is also referred to as the Laplace approximation. We considered two cases of the conductivity field, an isotropically smooth field and an anisotropically smooth (horizontal strata) one. The results indicate that fixing the conductivity as an incorrect but otherwise well justified (distributed parameter) field can result in infeasible and misleading posterior estimates in the sense that the true parameter is not supported by the posterior model. On the other hand, carrying out approximative marginalization does provide feasible estimates with both isotropically and anisotropically smooth unknown conductivities.

The computational feasibility (in large-scale) distributed Robin coefficient problems is provided by the adjoint method and being able to avoid the simultaneous estimation of the conductivity, which in contrast to the Robin coefficient, is a random field in the entire domain.

Future work will be concentrated in two distinct direction. Firstly from a qualitative stand point it would be of interest to consider how the quality of the estimate changes when the

number and location of measurements changes. Secondly, in the simplified model (with fixed $a = a_*$) there are no parameters distributed in the domain, warranting an investigation into the computational feasibility of other numerical methods such as boundary element methods (BEM) for solving the the simplified model.

Acknowledgments

This work was partially supported by NSF grants ACI-1550547 and 1654311.

ORCID iDs

Ruanui Nicholson  <https://orcid.org/0000-0002-9256-5513>

References

- [1] Kaup P G and Santosa F 1995 Nondestructive evaluation of corrosion damage using electrostatic measurements *J. Nondestruct. Eval.* **14** 127–36
- [2] Inglese G 1997 An inverse problem in corrosion detection *Inverse Problems* **13** 977–94
- [3] Alessandrini G, Piero L D and Rondi L 2003 Stable determination of corrosion by a single electrostatic boundary measurement *Inverse Problems* **19** 973–84
- [4] Chantasiriwan S 1999 Inverse heat conduction problem of determining time-dependent heat transfer coefficient *Int. J. Heat Mass Transfer* **42** 4275–85
- [5] Divo E, Kassab A J, Kapat J S and Chyu M K 2005 Retrieval of multidimensional heat transfer coefficient distributions using an inverse BEM-based regularized algorithm: numerical and experimental results *Eng. Anal. Bound. Elem.* **29** 150–60
- [6] Chaabane S and Jaoua M 1999 Identification of Robin coefficients by the means of boundary measurements *Inverse Problems* **15** 1455–68
- [7] Fasino D and Inglese G 1999 An inverse Robin problem for Laplace’s equation: theoretical results and numerical methods *Inverse Problems* **15** 41–8
- [8] Chaabane S, Ferchichi J and Kunisch K 2004 Differentiability properties of the L1-tracking functional and application to the Robin inverse problem *Inverse Problems* **20** 1083–98
- [9] Jin B 2007 Conjugate gradient method for the Robin inverse problem associated with the Laplace equation *Int. J. Numer. Methods Eng.* **71** 433–53
- [10] Ma Y-B 2015 Newton method for estimation of the Robin coefficient *J. Nonlinear Sci. Appl.* **8** 660–9
- [11] Bui-Thanh T, Ghattas O, Martin J and Stadler G 2013 A computational framework for infinite-dimensional Bayesian inverse problems part I: the linearized case, with application to global seismic inversion *SIAM J. Sci. Comput.* **35** A2494–523
- [12] Petra N, Martin J, Stadler G and Ghattas O 2014 A computational framework for infinite-dimensional Bayesian inverse problems part II: stochastic Newton MCMC with application to ice sheet flow inverse problems *SIAM J. Sci. Comput.* **36** A1525–55
- [13] Stuart A M 2010 Inverse problems: a Bayesian perspective *Acta Numer.* **19** 451–559
- [14] Gunzburger M 2003 *Perspectives in Flow Control and Optimization* (Philadelphia, PA: SIAM)
- [15] Hinze M, Pinnau R, Ulbrich M and Ulbrich S 2009 *Optimization with PDE Constraints* (New York: Springer)
- [16] Bui-Thanh T and Ghattas O 2014 An analysis of infinite dimensional Bayesian inverse shape acoustic scattering and its numerical approximation *SIAM/ASA J. Uncertain. Quantification* **2** 203–22
- [17] Alexanderian A, Petra N, Stadler G and Ghattas O 2016 A fast and scalable method for a-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems *SIAM J. Sci. Comput.* **38** A242–72
- [18] Kaipio J and Somersalo E 2005 *Statistical and Computational Inverse Problems* (New York: Springer)

- [19] Kaipio J and Somersalo E 2007 Statistical inverse problems: discretization, model reduction and inverse crimes *J. Comput. Appl. Math.* **198** 493–504
- [20] Nissinen A, Heikkinen L M, Kolehmainen V and Kaipio J P 2009 Compensation of errors due to discretization, domain truncation and unknown contact impedances in electrical impedance tomography *Inverse Problems* **20** 105504
- [21] Kolehmainen V, Tarvainen T, Arridge S R and Kaipio J P 2011 Marginalization of uninteresting distributed parameters in inverse problems: application to diffuse optical tomography *Int. J. Uncertain. Quantification* **1** 1–17
- [22] Mozumder M, Tarvainen T, Arridge S, Kaipio J P, D’Andrea C and Kolehmainen V 2016 Approximate marginalization of absorption and scattering in fluorescence diffuse optical tomography *Inverse Problems Imaging* **10** 227–46
- [23] Kaipio J and Kolehmainen V 2013 Approximate marginalization over modeling errors and uncertainties in inverse problems *Bayesian Theory and Applications* (Oxford: Oxford University Press) ch 32, pp 644–72
- [24] Daon Y and Stadler G 2016 Mitigating the influence of the boundary on PDE-based covariance operators *Inverse Problems Imaging* **12** 1083–102
- [25] Lindgren F, Rue H and Lindström J 2011 An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach *J. R. Stat. Soc. B* **73** 423–98
- [26] Rue H and Martino S 2007 Approximate Bayesian inference for hierarchical Gaussian Markov random field models *J. Stat. Plan. Inference* **137** 3177–92
- [27] Bekas C, Kokiopoulou E and Saad Y 2007 An estimator for the diagonal of a matrix *Appl. Numer. Math.* **57** 1214–29
- [28] Bekas C, Curioni A and Fedulova I 2009 Low cost high performance uncertainty quantification *Proc. of the 2nd Workshop on High Performance Computational Finance*
- [29] Lin L, Lu J, Ying L, Car R and Weinan E 2009 Fast algorithm for extracting the diagonal of the inverse matrix with application to the electronic structure analysis of metallic systems *Commun. Math. Sci.* **7** 755–77
- [30] Tang J M and Saad Y 2012 A probing method for computing the diagonal of a matrix inverse *Numer. Linear Algebr. Appl.* **19** 485–501
- [31] Calvetti D, Kaipio J P and Somersalo E 2006 Aristotelian prior boundary conditions *Int. J. Math. Comput. Sci.* **1** 63–81
- [32] Roininen L, Huttunen J M J and Lasanen S 2014 Whittle-Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography *Inverse Problems Imaging* **8** 561–86
- [33] Dashti M, Law K J H, Stuart A M and Voss J 2013 MAP estimators and their consistency in Bayesian nonparametric inverse problems *Inverse Problems* **29** 095017
- [34] Helin T and Burger M 2015 Maximum *a posteriori* probability estimates in infinite-dimensional bayesian inverse problem *Inverse Problems* **31** 085009
- [35] Tarantola A 2004 *Inverse Problem Theory and Methods for Model Parameter Estimation* (Philadelphia, PA: SIAM)
- [36] Flath P H, Wilcox L C, Akçelik V, Hill J, van Bloemen Waanders B and Ghattas O 2011 Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations *SIAM J. Sci. Comput.* **33** 407–32
- [37] Vogel C R 2002 *Computational Methods for Inverse Problems* (Philadelphia, PA: SIAM)
- [38] Golub G H and Loan C F V 2013 *Matrix Computations* 4th edn (Baltimore, MD: The Johns Hopkins University Press)
- [39] Frangos M, Marzouk Y, Willcox K and van Bloemen Waanders B 2010 Surrogate, reduced-order modeling: a comparison of approaches for large-scale statistical inverse problems *Large-Scale Inverse Problems and Quantification of Uncertainty* (New York: Wiley) ch 7, pp 123–49
- [40] Nissinen A, Kolehmainen V and Kaipio J P 2011 Reconstruction of domain boundary and conductivity in electrical impedance tomography using the approximation error approach *Int. J. Uncertain. Quantification* **1** 203–22
- [41] Arridge S R, Kaipio J P, Kolehmainen V, Schweiger M, Somersalo E, Tarvainen T and Vauhkonen M 2006 Approximation errors and model reduction with an application in optical diffusion tomography *Inverse Problems* **22** 175–96
- [42] Kaipio J, Lehtikoinen A, Voutilainen A and Finsterle S 2007 Approximation errors and truncation of computational domains with application to geophysical tomography *Inverse Problems* **1** 371–89

- [43] Huttunen J and Kaipio J 2007 Approximation error analysis in nonlinear state estimation with an application to state-space identification *Inverse Problems* **23** 2141–58
- [44] Nissinen A, Heikkinen L M and Kaipio J P 2008 The Bayesian approximation error approach for electrical impedance tomography: experimental results *Meas. Sci. Technol.* **19** 015501
- [45] Lipponen A, Seppanen A and Kaipio J P 2011 Nonstationary approximation error approach to imaging of three-dimensional pipe flow: experimental evaluation *Meas. Sci. Technol.* **22** 104013
- [46] Koponen J, Huttunen T, Tarvainen T and Kaipio J 2014 Bayesian approximation error approach in full-wave ultrasound tomography *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **61** 1627–37
- [47] Calvetti D, Ernst O and Somersalo E 2014 Dynamic updating of numerical model discrepancy using sequential sampling *Inverse Problems* **30** 110301
- [48] Calvetti D, Dunlop M M, Somersalo E and Stuart A M 2018 Iterative updating of model error for Bayesian inversion *Inverse Problems* **34** 025008
- [49] Tröltzsch F 2010 *Optimal Control of Partial Differential Equations, Methods and Applications* (Providence, RI: American Mathematical Society)
- [50] Borz A and Schulz V 2012 *Computational Optimization of Systems Governed by Partial Differential Equations* (Philadelphia, PA: SIAM)
- [51] Gilbarg D and Trudinger N S 2015 *Elliptic Partial Differential Equations of Second Order* (New York: Springer)
- [52] Grisvard P 2011 *Elliptic Problems in Nonsmooth Domains* vol 69 (Philadelphia, PA: SIAM)
- [53] Petra N, Zhu H, Stadler G, Hughes T J and Ghattas O 2012 An inexact Gauss–Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model *J. Glaciol.* **58** 889–903
- [54] Eisenstat S C and Walker H F 1996 Choosing the forcing terms in an inexact Newton method *SIAM J. Sci. Comput.* **17** 16–32