

PAPER

Iterative updating of model error for Bayesian inversion

To cite this article: Daniela Calvetti *et al* 2018 *Inverse Problems* **34** 025008

View the [article online](#) for updates and enhancements.

Iterative updating of model error for Bayesian inversion

Daniela Calvetti¹, Matthew Dunlop², Erkki Somersalo¹ 
and Andrew Stuart²

¹ Department of Mathematics, Applied Mathematics and Statistics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106, United States of America

² California Institute of Technology, Computing & Mathematical Sciences, 1200 E California Boulevard, Pasadena, California, CA 91125, United States of America

E-mail: dx57@case.edu, mdunlop@caltech.edu, ejs49@case.edu
and astuart@caltech.edu

Received 10 July 2017, revised 10 December 2017

Accepted for publication 20 December 2017

Published 18 January 2018



CrossMark

Abstract

In computational inverse problems, it is common that a detailed and accurate forward model is approximated by a computationally less challenging substitute. The model reduction may be necessary to meet constraints in computing time when optimization algorithms are used to find a single estimate, or to speed up Markov chain Monte Carlo (MCMC) calculations in the Bayesian framework. The use of an approximate model introduces a discrepancy, or modeling error, that may have a detrimental effect on the solution of the ill-posed inverse problem, or it may severely distort the estimate of the posterior distribution. In the Bayesian paradigm, the modeling error can be considered as a random variable, and by using an estimate of the probability distribution of the unknown, one may estimate the probability distribution of the modeling error and incorporate it into the inversion. We introduce an algorithm which iterates this idea to update the distribution of the model error, leading to a sequence of posterior distributions that are demonstrated empirically to capture the underlying truth with increasing accuracy. Since the algorithm is not based on rejections, it requires only limited full model evaluations.

We show analytically that, in the linear Gaussian case, the algorithm converges geometrically fast with respect to the number of iterations when the data is finite dimensional. For more general models, we introduce particle approximations of the iteratively generated sequence of distributions; we also prove that each element of the sequence converges in the large particle limit under a simplifying assumption. We show numerically that, as in the linear case, rapid convergence occurs with respect to the number of iterations.

Additionally, we show through computed examples that point estimates obtained from this iterative algorithm are superior to those obtained by neglecting the model error.

Keywords: model discrepancy, discretization error, particle approximation, importance sampling, electrical impedance tomography, Darcy flow

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

1. Introduction

The traditional way of describing an inverse problem is to define a forward map relating the unknown to an observed quantity, and to look for an estimate of the unknown when the data is corrupted by noise. In this description, it is often tacitly assumed that an underlying ‘truth’ exists, and the noiseless data arises from applying the forward map on this true value. On the other hand, it is commonly acknowledged that a mathematical model does not coincide with the reality, and therefore part of the noise must be attributed to the model discrepancy, or the mismatch between the model and the reality. Modeling this discrepancy is an active research topic in statistics—see [6, 7, 24] and the references therein; it is also a closely related to the concept of the ‘inverse crime’, a procedure of testing a computational method with data that has been generated by the same model that is used to solve the inverse problem [13, 23].

Common sources of modeling errors in inverse problems include:

- (i) model reduction—a complex, computationally intensive model is replaced by a simpler, less demanding model;
- (ii) parametric reduction—in a model depending on poorly known parameters, some of them are frozen to fixed values, assuming that the solution is not sensitive to them;
- (iii) unknown geometry—a computational domain of unknown shape is approximated by a standard geometry.

Including the modeling error into the computations in the traditional deterministic setting may not be straightforward. Recasting the inverse problem via Bayesian inference provides tools to carry this out in a natural statistical fashion. The present article introduces and analyzes a Bayesian methodology for model error estimation which demonstrably leads to improved estimates of the true unknown function generating the data.

1.1. Background

In this article, we consider the problem of estimating an unknown quantity u based on indirect observations. In the Bayesian framework, the prior belief about the quantity u is encoded in the prior probability distribution $\mathbb{P}(u)$, and given the observed data b , the posterior distribution $\mathbb{P}(u | b)$ follows from Bayes’ formula,

$$\mathbb{P}(u | b) \propto \mathbb{P}(b | u)\mathbb{P}(u),$$

where \propto denotes proportionality up to a scaling constant depending on b but not on u , and the distribution $\mathbb{P}(b | u)$ of b is the likelihood. To construct the likelihood, a forward model from u to b needs to be specified. A commonly used model, assuming additive observation noise that is independent of the unknown u , is

$$b = F(u) + \varepsilon$$

where $\varepsilon \sim \pi_{\text{noise}}(\cdot)$ is a realization of random noise, and F is a mapping defined typically on the parameter space for u , which is often infinite-dimensional (u is a function) or high dimensional (a function has been discretized to obtain u). Under these assumptions, the likelihood is equal to the distribution of ε with mean shifted by $F(u)$, i.e.

$$\mathbb{P}(b \mid u) = \pi_{\text{noise}}(b - F(u)).$$

To estimate u , or derived quantities based on it, numerical approximations of integrals with respect to the posterior distribution are required, and a common approach is to use sampling methods such as Markov chain Monte Carlo (MCMC). This requires a large number of repeated evaluations of the forward map F , which is often expensive to evaluate numerically. A particular instance that we have in mind is the situation where evaluation of F requires numerical solution of a partial differential equation. If computational resources or time are an issue, an attractive approach is to trade off the accuracy of evaluations with the computational cost by adjusting the resolution of the mesh that the PDE is solved upon. Denoting by f an approximation to F on a coarse mesh, a model for the data can be written as

$$b = f(u) + m + \varepsilon$$

where $m = F(u) - f(u)$ denotes the model error induced by moving from the accurate model F to the approximate one. If we ignore the fact that m depends on u , and instead model it as additive independent noise, the conditional likelihood $\mathbb{P}(b \mid u, m)$ is then given by

$$\mathbb{P}(b \mid u, m) = \pi_{\text{noise}}(b - f(u) - m);$$

evaluations of this map then only require evaluation of the approximate map f . Furthermore, the likelihood $\mathbb{P}(b \mid u)$ can be found by marginalizing out the model error m . However, the marginalization requires the distribution of m which is not known. As suggested in [22], the Bayesian approach provides a natural approximate solution to this problem: by using the prior distribution of u and the model error mapping $M(u) = F(u) - f(u)$, one can generate a sample of model errors to estimate the model error distribution. This approach, referred to as the *enhanced error model*, has been shown to produce more accurate point estimates than those that come from neglecting the model error (the *conventional error model*), see, e.g. [2, 3, 17, 23] for static inverse problems, and [19–21] for extensions to dynamic inverse problems.

In [8], the enhanced error model was developed further using the observation that the posterior distribution based on the error model contains refined information about the unknown u beyond the point estimate determined by the enhanced error model; as a consequence the model error distribution can be updated by pushing forward the distribution under the model error mapping M . When the data are particularly informative, posterior samples may differ significantly from prior samples, and this should produce a much better approximation to the distribution of the model error, potentially yielding a better approximation of the posterior distribution of u , and at the very least providing point estimates of higher accuracy. The procedure can be iterated to produce a sequence of approximate posterior distributions that have the potential to yield a sequence of point estimates of improved accuracy; they may also approximate the true posterior distribution with increasing accuracy. In this article, we address this approach in a systematic way, with particular focus on convergence of the iterative updating.

The effect of model error and model discrepancy in Bayesian inference is a widely studied topic. Early works focus primarily on code uncertainty—the uncertainty that arises due to expense of forward model evaluations meaning that it is only practical to compute outputs of the model for a limited finite number of inputs. A review of work in this direction is given in [31], including the problems of optimal choice of inputs at which to evaluate the forward model, and how to predict the output for inputs away from the computed ones. In [24]

the numerous sources of uncertainty within computational inverse problems are discussed, including those arising from model inadequacy and code uncertainty. The authors model this error as a function independent of the approximate model, which can be justified in certain cases. Hierarchical Gaussian priors are placed upon the model and the model error, and the model and error are then linked by imposing correlations between the hyperparameters. The technique has subsequently been developed further in [7], and used in, for example, the context of model validation [4] and uncertainty quantification [18]. More recent work in probabilistic numerical methods [6, 12] provides a unifying perspective on this body of work, linking it to earlier research connecting numerical algorithms to Bayesian statistics [16].

1.2. Our contribution

- We develop further the iterative updating of the posterior probability densities based on repeated updating of the model error distribution, leading to an approximation of the posterior probability density. While the approximation error is defined through the computationally expensive accurate model, the posterior approximation we introduce relies primarily on the computationally inexpensive approximate model, and a limited number of evaluations of the accurate model.
- In the case where the models are linear and the noise and prior distributions are Gaussian, we show that the means and covariances of the resulting sequence of posterior distributions converge to a limit geometrically fast.
- For more general models and prior/noise distributions we introduce particle approximations to allow the algorithm to be implemented numerically, and show convergence of these approximations in the large particle limit.
- We illustrate numerically the effectiveness of the algorithms in multiple different settings, showing the advantage over the conventional and enhanced error models.

1.3. Outline

The iterative approach of updating the posterior distributions is introduced in section 2. In section 3 we focus on the particular case where the forward model is linear, and the noise and prior distributions are Gaussian. The assumptions imply that the approximate posterior distributions are also Gaussian, and can therefore be characterized by their means and covariances. We identify conditions guaranteeing the convergence of the sequence of approximate posteriors to a non-degenerate limit as the number of iterations tends to infinity. In section 4 we discuss different sampling methods which may be used to implement the algorithm in practice. In particular we focus on particle methods that require a finite number of full model evaluations in order to estimate the modeling error and posterior distribution, and show convergence to the correct sequence of approximate distributions in the large particle limit. Finally, in section 5, we provide numerical illustrations of the behavior and performance of the algorithm for three different forward models. Section 6 contains the conclusions and discussion.

2. Problem formulation

We start by introducing the main ingredients of the iterative algorithm: the accurate and approximate models are defined in section 2.1, along with some examples whose details will be discussed later on. The enhanced error model [22] is reviewed in section 2.2, prompting the question of how to update the density of the modeling error. In section 2.3 we provide

an iterative algorithm for doing this, in the case where all measures involved have Lebesgue densities; a more general algorithm is provided in the appendix for cases such as those arising when the measures are defined on infinite-dimensional function spaces.

2.1. Accurate versus approximate model

Let X, Y be two Banach spaces representing the parameter and data spaces, respectively. Let

$$F : X \rightarrow Y, \quad u \mapsto b$$

denote a reference forward model, referred to as the *accurate model*, and let the *approximate model* be denoted by

$$f : X \rightarrow Y, \quad u \mapsto b.$$

We write the observation model using the accurate model,

$$b = F(u) + \varepsilon, \tag{1}$$

and equivalently, using the approximate model, as

$$\begin{aligned} b &= f(u) + (F(u) - f(u)) + \varepsilon \\ &= f(u) + m + \varepsilon, \end{aligned} \tag{2}$$

where m represents the modeling error,

$$m = F(u) - f(u) = M(u).$$

In light of the above observation, we may view the data as coming from the approximate model, with an error term which reflects both observational noise ε and modeling error m . The main problem addressed in this paper is how to model the probability distribution of the model error. We study this question with the goal of providing computations to estimate the unknown u from b using only the less expensive model f , and not the true model F , without creating modeling artifacts.

We conclude this subsection by giving two examples of approximate models, both of which may be relevant in applications.

Example 2.1 (Linearization). In electrical impedance tomography (EIT), the goal is to estimate the conductivity distribution inside a body from a finite set of current/voltage measurements at the boundary, as discussed in more detail in section 5.2. We denote by F the differentiable non-linear forward model, mapping the appropriately parametrized conductivity distribution to the voltage measurements. We define the approximate model through the linearization,

$$f(u) = F(u_0) + DF(u_0)(u - u_0),$$

where $DF(u_0)$ is the Jacobian of the forward map, and u_0 is a fixed parameter value, representing, e.g. a constant conductivity background.

Example 2.2 (Coarse mesh). In the EIT model, the accurate model represents the forward model computed with a FEM grid fine enough to guarantee that the numerical solution approximates the solution of the underlying PDE within a required precision. To speed up computations, we introduce the reduced model f as the forward model based on FEM built on a coarse grid. We assume that both computational grids are built on an underlying independent discretization of the conductivity σ , appropriately parametrized, and the FEM stiffness matri-

ces that require integration over elements are computed by evaluating σ in the Gauss points of the elements, respectively.

2.2. The enhanced error model

We start by reviewing the basic ideas of the enhanced error model and, for simplicity, assume here that X and Y are Euclidean spaces and that all probability distributions are expressible in terms of Lebesgue densities. We assume that u is an X -valued random variable with a given *a priori* density,

$$u \sim \pi_{\text{prior}}(u).$$

Furthermore, we assume that the additive noise ε is a Y -valued random variable, independent of u , with the density

$$\varepsilon \sim \pi_{\text{noise}}(\varepsilon).$$

In view of the model (1), we may write the likelihood as

$$\pi(b | u) = \pi_{\text{noise}}(b - F(u)),$$

and the posterior density is, according to Bayes' formula, given by

$$\pi(u | b) \propto \pi_{\text{prior}}(u) \pi_{\text{noise}}(b - F(u)).$$

If, instead, we want to use the approximate model $f(u)$ for the data, the modeling error m needs to be taken into account. The idea of the *enhanced error model* in [22] is the following: given the prior distribution $\mu_0(du) = \pi_{\text{prior}}(u)du$, with no other information about u , the probability distribution of $m = M(u)$ is obtained as a push-forward of the prior distribution:

$$m \sim M^{\#} \mu_0.$$

To obtain a computationally efficient formalism, the distribution of m is approximated by a Gaussian distribution sharing the mean and covariance of the push-forward measure,

$$m \sim \mathcal{N}(\bar{m}, \Sigma),$$

where in practice the mean and covariance may be estimated numerically by sampling the modeling error. The Gaussian approximation is particularly convenient if the additive noise is Gaussian,

$$\varepsilon \sim \mathcal{N}(0, \Gamma),$$

as it leads to the approximate likelihood model

$$\pi(b | u) \propto \exp\left(-\frac{1}{2}(b - f(u) - \bar{m})^T(\Gamma + \Sigma)^{-1}(b - f(u) - \bar{m})\right),$$

and, consequently, to the posterior model

$$\pi(u | b) \propto \pi_{\text{prior}}(u) \exp\left(-\frac{1}{2}(b - f(u) - \bar{m})^T(\Gamma + \Sigma)^{-1}(b - f(u) - \bar{m})\right).$$

Note that the enhanced error model can be interpreted as a form of variance inflation: if the model error is neglected as in the conventional error model, then the covariance matrix in the likelihood would be smaller in the sense of positive-definite quadratic forms, since Σ is non-negative definite. This is to be expected as the lack of accuracy in the model contributes

additional uncertainty to the problem; we do not wish to be over-confident in an inaccurate model. In the next section we explain how this idea may be built upon to produce the algorithms studied in this paper.

2.3. The iterative algorithm

In this subsection we generalize the preceding enhanced error model in two ways: (i) we iterate the construction of the model error, updating its probability distribution by pushing forward by M the measure μ_ℓ , the posterior distribution of u when the model error distribution is computed as the pushforward under M of the measure $\mu_{\ell-1}$; in this iterative method we choose μ_0 to be the prior and so the first step is analogous to what is described in the previous subsection; (ii) we do not invoke a Gaussian approximation of the model error, leaving open other possibilities for practical implementation. We describe the resulting algorithm here in the case where Lebesgue densities exist, and refer to the appendix for its formulation in a more abstract setting.

Algorithm (Lebesgue densities). Let μ_ℓ denote the posterior distribution at stage ℓ , with density π_ℓ , so that $\mu_\ell(\mathrm{d}u) = \pi_\ell(u) \mathrm{d}u$. Denote $\pi_\ell(\mathbf{b} | u)$ the likelihood at stage ℓ .

1. Set $\pi_0(u) = \pi_{\text{prior}}(u)$ and $\ell = 0$.
2. Given μ_ℓ , assume $m \sim M^\# \mu_\ell$. Assuming that u and m are mutually independent, we have

$$\pi(\mathbf{b} | u, m) = \pi_{\text{noise}}(\mathbf{b} - f(u) - m),$$

and by marginalization,

$$\begin{aligned} \pi_{\ell+1}(\mathbf{b} | u) &= \int_Y \pi_{\text{noise}}(\mathbf{b} - f(u) - m) (M^\# \mu_\ell)(\mathrm{d}m) \\ &= \int_X \pi_{\text{noise}}(\mathbf{b} - f(u) - M(z)) \pi_\ell(z) \mathrm{d}z. \end{aligned}$$

Hence using Bayes' theorem, update the posterior distribution:

$$\pi_{\ell+1}(u) \propto \pi_{\text{prior}}(u) \int_X \pi_{\text{noise}}(\mathbf{b} - f(u) - M(z)) \pi_\ell(z) \mathrm{d}z. \quad (3)$$

3. Set $\ell \mapsto \ell + 1$ and go to 2. □

We can give an explicit expression for the above density $\pi_\ell(u)$:

$$\pi_\ell(u) \propto \pi_{\text{prior}}(u) \int_X \cdots \int_X \left(\prod_{i=1}^{\ell} \pi_{\text{noise}}(\mathbf{b} - f(z_{i+1}) - M(z_i)) \pi_{\text{prior}}(z_i) \right) \mathrm{d}z_1 \cdots \mathrm{d}z_\ell$$

where we define $z_{\ell+1} = u$.

The above algorithm can be generalized to the case when no Lebesgue densities exist, such as will be the case on infinite dimensional Banach spaces, see the appendix.

3. The linear Gaussian case

We analyze the convergence of the general algorithm in the case where both the accurate model F and the approximate model f are linear, and the noise and prior distributions are Gaussian. With these assumptions the posterior distributions forming the sequence remain Gaussian, and are hence characterized by the sequences of means and covariances. Convergence properties

of the iteration can be established by using the updating formulas for these sequences. Though the iteration is not immediately implementable, since the full matrix for the accurate model F is required to calculate each covariance matrix, the explicit convergence results give insight into the implementable variants of the algorithm introduced in section 4, as well as for non-linear forward maps.

In section 3.1, we first describe how the posterior density evolves, establishing the iterations that the means and covariances satisfy. In section 3.2 we show that if the model error is sufficiently small, the sequences of posterior means and covariances converge geometrically fast. Moreover, despite the repeated incorporation of the data into the posterior, the limiting distribution does not become singular. Additionally, in section 3.3 we show that in finite dimensions the assumption of small model error is not needed in order to establish convergence of the covariance.

3.1. Evolution of the Posterior distribution

When the noise distribution π_{noise} and prior distribution π_{prior} are Gaussian, the measure μ_ℓ is Gaussian at each stage ℓ , and we can write down expressions for the evolution of its mean and covariance. Let X, Y be separable Hilbert spaces and $A_\star : X \rightarrow Y$ a linear operator. Assume that the data $b \in Y$ arise from A_\star via

$$b = A_\star u + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \Gamma)$$

where $\Gamma : Y \rightarrow Y$ is a symmetric positive definite covariance operator. Let $A : X \rightarrow Y$ denote an approximation to A_\star so that the expression for b may be rewritten

$$b = Au + (A_\star - A)u + \varepsilon.$$

We define the model error operator $M : X \rightarrow Y$ by $M = A_\star - A$, so that Mu represents the (unknown) model error. We assume that the model error is Gaussian, with unknown mean w and covariance Σ . Additionally we assume that it is independent of the observation noise. The data is now given by

$$b = Au + \hat{\varepsilon}, \quad \hat{\varepsilon} \sim \mathcal{N}(w, \Sigma + \Gamma).$$

Let $\mu_0 = \mathcal{N}(m_0, C_0)$ denote the prior distribution on u . We first estimate w and Σ by pushing forward μ_0 by the model error operator M :

$$M^\# \mu_0 = \mathcal{N}(Mm_0, MC_0M^*) \equiv \mathcal{N}(\hat{w}_1, \hat{\Sigma}_1).$$

Then, assuming for now that the model error has this distribution, the resulting posterior distribution on u can be calculated as $\mu_1 = \mathcal{N}(m_1, C_1)$, where

$$C_1 = (A^*(\Gamma + MC_0M^*)^{-1}A + C_0^{-1})^{-1},$$

$$m_1 = C_1(A^*(\Gamma + MC_0M^*)^{-1}(b - Mm_0) + C_0^{-1}m_0).$$

As described previously, in order to obtain a better approximation of the model error, the above step can be repeated so that the measure pushed forward to approximate the model error is closer to the posterior distribution. Therefore, in the next step, we approximate the distribution of the model error by

$$M^\# \mu_1 = \mathcal{N}(Mm_1, MC_1M^*) \equiv \mathcal{N}(\hat{w}_2, \hat{\Sigma}_2),$$

and make a further posterior approximation. Iterating this process, we see that the posterior mean and covariance evolve via

$$\begin{aligned} C_{\ell+1} &= (A^*(\Gamma + MC_\ell M^*)^{-1}A + C_0^{-1})^{-1}, \\ m_{\ell+1} &= C_{\ell+1}(A^*(\Gamma + MC_\ell M^*)^{-1}(b - Mm_\ell) + C_0^{-1}m_0). \end{aligned} \quad (4)$$

We wish to show that these sequences are convergent.

We may write the above iteration in an equivalent form as

$$\begin{aligned} C_{\ell+1} &= C_0 - C_0 A^* (\Gamma + MC_\ell M^* + AC_0 A^*)^{-1} AC_0, \\ m_{\ell+1} &= m_0 + C_0 A^* (\Gamma + MC_\ell M^* + AC_0 A^*)^{-1} (b - Am_0 - Mm_\ell), \end{aligned}$$

using results from [25, 27], assuming that C_0 is trace-class and Γ is trace-class or white. This form has the advantage that the unbounded operator C_0^{-1} does not appear, and so we need not worry about its domain of definition. In what follows we simply assume that this equivalent expression for the evolution of the means and covariances is valid.

In the following subsections, we consider two different cases. In the first one, we limit the data into a finite dimensional space Y but let the space X to be a Hilbert space. The convergence of the algorithm is demonstrated under certain restrictive conditions: the modeling error needs to be small enough. In the second case, we also limit the unknown u to a finite dimensional space, and show that in this case the convergence proof can be obtained without the restrictions needed in the former case. We emphasize that, although we establish convergence of the iteration in various settings, the limiting distribution does not coincide with the true posterior distribution found in the absence of model error. Nonetheless our numerical experiments will show that the limiting distribution can significantly improve point estimates of the underlying value used to generate the data.

3.2. Convergence in infinite dimensions

We introduce a scalar parameter δ controlling the accuracy of the approximation A of A_* , writing δM in place of M . By writing explicitly the dependence of the mean and covariance on δ , we have

$$\begin{aligned} C_{\ell+1}(\delta) &= C_0 - C_0 A^* (\Gamma + \delta^2 MC_\ell(\delta) M^* + AC_0 A^*)^{-1} AC_0, \\ m_{\ell+1}(\delta) &= m_0 + C_0 A^* (\Gamma + \delta^2 MC_\ell(\delta) M^* + AC_0 A^*)^{-1} (b - Am_0 - \delta Mm_\ell(\delta)). \end{aligned}$$

Let $\mathcal{L}(X)$ denote the space of bounded linear operators on X , equipped with the operator norm. Let $S_+(X) \subseteq \mathcal{L}(X)$ denote the set of positive bounded linear operators on X , and $\overline{S_+(X)}$ the set of non-negative bounded linear operators on X . We may write the iteration for $C_\ell(\delta)$ as

$$C_{\ell+1}(\delta) = \mathcal{F}(C_\ell(\delta), \delta) \quad (5)$$

where $\mathcal{F} : \overline{S_+(X)} \times \mathbb{R} \rightarrow \overline{S_+(X)}$ is given by

$$\mathcal{F}(B, \delta) = C_0 - C_0 A^* (\Gamma + \delta^2 MBM^* + AC_0 A^*)^{-1} AC_0.$$

We show that under certain assumptions, for all δ sufficiently small, $\mathcal{F}(\cdot, \delta)$ has a unique stable fixed point $C(\delta)$. The assumptions we make are as follows.

Assumptions 3.1.

- (i) $C_0 \in S_+(X)$ and is trace-class.
- (ii) Y is finite dimensional, and $\Gamma \in S_+(Y)$.
- (iii) A_* , $A : X \rightarrow Y$ are bounded.

We first establish the following result concerning convergence of the sequence of covariance operators:

Proposition 3.2. *Let assumptions 3.1 hold. Then there is a $\beta > 0$ such that for all $\delta < 1/\beta$, a unique $C(\delta) \in S_+(X)$ exists with*

$$C(\delta) = \mathcal{F}(C(\delta), \delta).$$

Moreover, $C(\delta)$ is a stable fixed point of $\mathcal{F}(\cdot, \delta)$, and there is a constant α_1 such that

$$\|C_\ell(\delta) - C(\delta)\|_{\mathcal{L}(X)} \leq \alpha_1(\beta\delta)^{2\ell} \quad \text{for all } \ell \geq 1.$$

In particular, for $\delta < 1/\beta$, the sequence $\{C_\ell(\delta)\}_{\ell \geq 1}$ converges geometrically fast.

From this geometric convergence of $\{C_\ell(\delta)\}_{\ell \geq 1}$ we can deduce that also the means $m_\ell(\delta)$ converge: define the maps $G_\ell : X \times \mathbb{R} \rightarrow X$, $\ell \geq 0$, by

$$G_\ell(m, \delta) = m_0 + C_0 A^* (\Gamma + \delta^2 M C_\ell(\delta) M^* + A C_0 A^*)^{-1} (b - A m_0 - \delta M m)$$

so that the update for $m_\ell(\delta)$ is given by

$$m_{\ell+1}(\delta) = G_\ell(m_\ell(\delta), \delta).$$

Define also the limiting map $G : X \times \mathbb{R} \rightarrow X$ by

$$G(m, \delta) = m_0 + C_0 A^* (\Gamma + \delta^2 M C(\delta) M^* + A C_0 A^*)^{-1} (b - A m_0 - \delta M m).$$

Then we have the following result:

Proposition 3.3. *Let assumptions 3.1 hold, and let β be as in proposition 3.2. Then for all $\delta < 1/\beta$, there exists a unique $m(\delta) \in X$ with*

$$m(\delta) = G(m(\delta), \delta).$$

Moreover, there is an $\alpha_2 > 0$ such that

$$\|m_\ell(\delta) - m(\delta)\|_X \leq \alpha_2(\beta\delta)^\ell \quad \text{for all } \ell \geq 1.$$

Hence, for $\delta < 1/\beta$, the sequence $\{m_\ell(\delta)\}_{\ell \geq 1}$ converges geometrically fast.

To prove the above propositions we first prove the following lemma. In the proof, the following notation is used: given symmetric non-negative linear operators $B_1, B_2 \in \mathcal{L}(X)$, we write $B_1 \leq B_2$ to mean that $B_2 - B_1$ is non-negative.

Lemma 3.4. *Let assumptions 3.1 hold. Then the family of operators $K(B, \delta) : Y \rightarrow X$ given by*

$$K(B, \delta) = C_0 A^* (\Gamma + \delta^2 M B M^* + A C_0 A^*)^{-1}$$

is bounded uniformly over $B \in \overline{S_+(X)}$ and $\delta \in \mathbb{R}$.

Proof. We have that

$$\Gamma + AC_0A^* \leq \Gamma + Q + AC_0A^*,$$

for any $Q \in \overline{S_+(Y)}$, which implies that

$$(\Gamma + Q + AC_0A^*)^{-1} \leq (\Gamma + AC_0A^*)^{-1},$$

and consequently

$$\|(\Gamma + Q + AC_0A^*)^{-1}\|_{\mathcal{L}(Y)} \leq \|(\Gamma + AC_0A^*)^{-1}\|_{\mathcal{L}(Y)}.$$

By choosing $Q = \delta^2MBM^*$ the claim follows. \square

In what follows we will denote $K_{\max} = \sup \{\|K(B, \delta)\|_{\mathcal{L}(Y,X)} \mid B \in \overline{S_+(X)}, \delta \in \mathbb{R}\}$. Furthermore, we define the parameter $\beta = K_{\max}\|M\|_{\mathcal{L}(Y)}$.

Remark 3.5. The conclusion of the above lemma does not hold in general unless the space Y is finite dimensional. For example, let $X = Y$ be infinite dimensional, and suppose that A , Γ and C_0 are symmetric, positive and simultaneously diagonalizable, with normalized eigenbasis $\{\varphi_j\}$ satisfying

$$A\varphi_j = a_j\varphi_j, \quad \Gamma\varphi_j = \gamma_j\varphi_j, \quad C_0\varphi_j = \lambda_j\varphi_j.$$

Then choosing $B \in \overline{S_+(X)}$ to be the zero operator, we have

$$\begin{aligned} K(B, \delta)\varphi_j &= \frac{\lambda_j a_j}{\gamma_j + a_j^2 \lambda_j} \varphi_j \\ &= \frac{1}{\frac{\gamma_j}{a_j \lambda_j} + a_j} \varphi_j. \end{aligned}$$

Assume that $ja_j \rightarrow 0$, and set $\gamma_j = a_j \lambda_j / j$. Then $\|K(B, \delta)\varphi_j\| \geq Cj \rightarrow \infty$, with $C > 0$, and so the uniform boundedness cannot hold.

This does not necessarily mean that the conclusions of propositions 3.2 and 3.3 are false if Y is not assumed finite dimensional, however the proofs we provide of these propositions rely strongly on the conclusions of lemma 3.4.

Proof of proposition 3.2. We first show that for δ sufficiently small, the map $\mathcal{F}(\cdot, \delta)$ is a contraction on $\overline{S_+(X)}$. To do this, we look at the Fréchet derivative of the map, which may be calculated explicitly. For $B \in \overline{S_+(X)}$ and $V \in \mathcal{L}(X)$, we have

$$\begin{aligned} D_B \mathcal{F}(B, \delta)V &= \delta^2 C_0 A^* (\Gamma + \delta^2 MBM^* + AC_0 A^*)^{-1} MVM^* (\Gamma + \delta^2 MBM^* + AC_0 A^*)^{-1} AC_0 \\ &= \delta^2 K(B, \delta) MVM^* K(B, \delta)^* \end{aligned}$$

where $K(B, \delta)$ is as defined in lemma 3.4. The norm of the derivative can be estimated as

$$\begin{aligned}
\|D_{\mathbf{B}}\mathcal{F}(\mathbf{B}, \delta)\|_{\mathcal{L}(X) \rightarrow \mathcal{L}(X)} &= \sup_{\|V\|_{\mathcal{L}(X)}=1} \|D_{\mathbf{B}}\mathcal{F}(\mathbf{B}, \delta)V\|_{\mathcal{L}(X)} \\
&\leq \sup_{\|V\|_{\mathcal{L}(X)}=1} \delta^2 \|K(\mathbf{B}, \delta)\|_{\mathcal{L}(Y, X)}^2 \|M\|_{\mathcal{L}(Y)}^2 \|V\|_{\mathcal{L}(X)} \\
&\leq (\beta\delta)^2
\end{aligned}$$

by the estimate of lemma 3.4. Since the above bound is uniform in \mathbf{B} , we may use the mean value theorem to deduce that for all $\mathbf{B}_1, \mathbf{B}_2 \in \overline{S_+(X)}$,

$$\|\mathcal{F}(\mathbf{B}_1, \delta) - \mathcal{F}(\mathbf{B}_2, \delta)\|_{\mathcal{L}(X)} \leq (\beta\delta)^2 \|\mathbf{B}_1 - \mathbf{B}_2\|_{\mathcal{L}(X)}$$

and so $\mathcal{F}(\cdot, \delta)$ is a contraction for $\delta < 1/\beta$. The set $\overline{S_+(X)}$ is a complete subset of the space $\mathcal{L}(X)$, and so by the Banach fixed-point theorem there exists a unique $\mathbf{C}(\delta) \in \overline{S_+(X)}$ such that

$$\mathbf{C}(\delta) = \mathcal{F}(\mathbf{C}(\delta), \delta),$$

and we have $\mathbf{C}_\ell(\delta) \rightarrow \mathbf{C}(\delta)$. Moreover,

$$\begin{aligned}
\|\mathbf{C}_\ell(\delta) - \mathbf{C}(\delta)\|_{\mathcal{L}(X)} &= \|\mathcal{F}(\mathbf{C}_{\ell-1}(\delta), \delta) - \mathcal{F}(\mathbf{C}(\delta), \delta)\|_{\mathcal{L}(X)} \\
&\leq (\beta\delta)^2 \|\mathbf{C}_{\ell-1}(\delta) - \mathbf{C}(\delta)\|_{\mathcal{L}(X)},
\end{aligned}$$

and recursively,

$$\|\mathbf{C}_\ell(\delta) - \mathbf{C}(\delta)\|_{\mathcal{L}(X)} \leq (\beta\delta)^{2\ell} \|\mathbf{C}_1(\delta) - \mathbf{C}_0\|_{\mathcal{L}(X)} = \alpha_1 (\beta\delta)^{2\ell}.$$

We finally show that we actually have $\mathbf{C}(\delta) \in S_+(X)$ and so the covariance does not become degenerate in the limit. We denote $\mathbf{C}_{\text{post}} = \mathcal{F}(\mathbf{C}_{\text{post}}, 0)$ the exact posterior covariance in the absence of model error, noting that $\mathbf{C}_{\text{post}} \in S_+(X)$ as we assume $\mathbf{C}_0 \in S_+(X)$ and $\Gamma \in S_+(Y)$. From a similar argument as in the proof of lemma 3.4, we have

$$\begin{aligned}
0 < \mathbf{C}_{\text{post}} &= \mathbf{C}_0 - \mathbf{C}_0 \mathbf{A}^* (\Gamma + \mathbf{A} \mathbf{C}_0 \mathbf{A}^*)^{-1} \mathbf{A} \mathbf{C}_0 \\
&\leq \mathbf{C}_0 - \mathbf{C}_0 \mathbf{A}^* (\Gamma + \delta^2 \mathbf{M} \mathbf{C}(\delta) \mathbf{M}^* + \mathbf{A} \mathbf{C}_0 \mathbf{A}^*)^{-1} \mathbf{A} \mathbf{C}_0 \\
&= \mathbf{C}(\delta),
\end{aligned}$$

which gives the result. □

Proof of proposition 3.3. We may express G_ℓ in the form of an affine mapping,

$$G_\ell(m, \delta) = H_\ell(\delta)m + g_\ell(\delta)$$

where $H_\ell(\delta)$ and $g_\ell(\delta)$ are given by

$$\begin{aligned}
H_\ell(\delta) &= -\delta \mathbf{C}_0 \mathbf{A}^* (\Gamma + \delta^2 \mathbf{M} \mathbf{C}_\ell(\delta) \mathbf{M}^* + \mathbf{A} \mathbf{C}_0 \mathbf{A}^*)^{-1} \mathbf{M} \\
&= -\delta \mathbf{K}(\mathbf{C}_\ell(\delta), \delta) \mathbf{M}, \\
g_\ell(\delta) &= m_0 + \mathbf{C}_0 \mathbf{A}^* (\Gamma + \delta^2 \mathbf{M} \mathbf{C}_\ell(\delta) \mathbf{M}^* + \mathbf{A} \mathbf{C}_0 \mathbf{A}^*)^{-1} (b - \mathbf{A} m_0) \\
&= m_0 + \mathbf{K}(\mathbf{C}_\ell(\delta), \delta) (b - \mathbf{A} m_0),
\end{aligned}$$

respectively. From the estimates of lemma 3.4, we obtain the uniform bounds

$$\begin{aligned}\|H_\ell(\delta)\|_{\mathcal{L}(X)} &\leq \beta\delta < 1, \\ \|g_\ell(\delta)\|_X &\leq \|m_0\|_X + \|K(C_\ell(\delta), \delta)\|_{\mathcal{L}(Y)}\|b - Am_0\|_X \\ &= \|m_0\|_X + \beta\|b - Am_0\|_X = L.\end{aligned}$$

From the convergence of the sequence $\{C_\ell(\delta)\}_{\ell \geq 1}$ in the previous proposition, we see that $\{H_\ell(\delta)\}_{\ell \geq 1}$ and $\{g_\ell(\delta)\}_{\ell \geq 1}$ also converge, the limits being denoted by $H(\delta)$ and $g(\delta)$, respectively. Explicitly,

$$\begin{aligned}H(\delta) &= -\delta K(C(\delta), \delta)M, \\ g(\delta) &= m_0 + K(C(\delta), \delta)(b - Am_0).\end{aligned}$$

Moreover, since $B \mapsto K(B, \delta)$ is Fréchet differentiable, this convergence occurs at the same rate as the convergence of $\{C_\ell(\delta)\}_{\ell \geq 1}$.

Next we show that $\{m_\ell(\delta)\}_{\ell \geq 1}$ remains bounded for sufficiently small δ . From the bounds above, we have

$$\begin{aligned}\|m_\ell(\delta)\|_X &\leq \|H_{\ell-1}(\delta)\|_{\mathcal{L}(X)}\|m_{\ell-1}(\delta)\|_X + \|g_{\ell-1}(\delta)\|_X \\ &\leq \beta\delta\|m_{\ell-1}\|_X + L,\end{aligned}$$

and therefore, by repeatedly applying the estimate, we obtain

$$\begin{aligned}\|m_\ell(\delta)\|_X &\leq (\beta\delta)^\ell \|m_0\|_X + L \sum_{j=0}^{\ell-1} (\beta\delta)^j \\ &\leq (\beta\delta)^\ell \|m_0\|_X + \frac{L}{1 - \beta\delta},\end{aligned}$$

which provides a uniform bound for $\delta < 1/\beta$.

To prove the convergence, we write first for $i \geq 1$ the estimate

$$\begin{aligned}\|m_{i+1}(\delta) - m_i(\delta)\|_X &= \|H_i(\delta)m_i(\delta) + g_i(\delta) - H_{i-1}(\delta)m_{i-1}(\delta) - g_{i-1}(\delta)\|_X \\ &\leq \|H_i(\delta)m_i(\delta) - H_i(\delta)m_{i-1}(\delta)\|_X \\ &\quad + \|H_i(\delta)m_{i-1}(\delta) - H_{i-1}(\delta)m_{i-1}(\delta)\|_X \\ &\quad + \|g_i(\delta) - g_{i-1}(\delta)\|_X \\ &\leq \|H_i(\delta)\|_{\mathcal{L}(X)}\|m_i(\delta) - m_{i-1}(\delta)\|_X \\ &\quad + \|H_i(\delta) - H_{i-1}(\delta)\|_{\mathcal{L}(X)}\|m_{i-1}(\delta)\|_X \\ &\quad + \|g_i(\delta) - g_{i-1}(\delta)\|_X,\end{aligned}$$

and further, by the geometric convergence of the sequences $\{H_i(\delta)\}$ and $\{g_i(\delta)\}$, and the uniform boundedness, for some $\gamma > 0$,

$$\|m_{i+1}(\delta) - m_i(\delta)\|_X \leq \beta\delta\|m_i(\delta) - m_{i-1}(\delta)\|_X + \gamma(\beta\delta)^{2i}.$$

By repeatedly applying the estimate, we arrive at

$$\begin{aligned}\|m_{i+1}(\delta) - m_i(\delta)\|_X &\leq (\beta\delta)^{i+1}\|m_0\|_X + \gamma \sum_{j=0}^i (\beta\delta)^j ((\beta\delta)^2)^{i-j} \\ &= (\beta\delta)^{i+1}\|m_0\|_X + \gamma \frac{((\beta\delta)^2)^{i+1} - (\beta\delta)^{i+1}}{(\beta\delta)^2 - \beta\delta}.\end{aligned}$$

From this bound, it follows that $\{m_\ell(\delta)\}_{\ell \geq 1}$ is a Cauchy sequence: for $k > \ell$, we have

$$\begin{aligned}\|m_k(\delta) - m_\ell(\delta)\|_X &\leq \sum_{i=\ell}^{k-1} \|m_{i+1}(\delta) - m_i(\delta)\|_X \\ &\leq \|m_0\|_X \sum_{i=\ell}^{k-1} (\beta\delta)^{i+1} + \gamma \sum_{i=\ell}^{k-1} \frac{((\beta\delta)^2)^{i+1} - (\beta\delta)^{i+1}}{(\beta\delta)^2 - \beta\delta} \\ &= \|m_0\|_X \frac{(\beta\delta)^{\ell+1} - (\beta\delta)^{k+1}}{1 - \beta\delta} + \gamma \frac{(\beta\delta)^{\ell+1} - (\beta\delta)^{k+1}}{(\beta\delta - 1)((\beta\delta)^2 - \beta\delta)} \\ &\quad + \gamma \frac{((\beta\delta)^2)^{\ell+1} - ((\beta\delta)^2)^{k+1}}{(\delta^2 - 1)(\delta^2 - \beta\delta)}\end{aligned}$$

which tends to zero as $k, \ell \rightarrow \infty$, provided δ is small enough. Thus the sequence $\{m_\ell(\delta)\}$ converges, and we denote the limit by $m(\delta)$. Taking the limit as $k \rightarrow \infty$ in the above inequality, we have

$$\begin{aligned}\|m_\ell(\delta) - m(\delta)\|_X &\leq \frac{(\beta\delta)^{\ell+1}}{\beta\delta - 1} + \gamma \frac{(\beta\delta)^{\ell+1}}{(\beta\delta - 1)((\beta\delta)^2 - \beta\delta)} + \gamma \frac{((\beta\delta)^2)^{\ell+1}}{(\delta^2 - 1)((\beta\delta)^2 - \beta\delta)} \\ &= \mathcal{O}((\beta\delta)^\ell)\end{aligned}$$

for all $\ell \geq 1$, and it follows that $\{m_\ell(\delta)\}_{\ell \geq 1}$ converges geometrically with rate $\beta\delta$.

To show that the limit $m(\delta)$ is indeed a fixed point of $G(\cdot, \delta)$, we first we note that

$$\begin{aligned}\|G_\ell(m_\ell(\delta), \delta) - G(m(\delta), \delta)\|_X &= \|H_\ell(\delta)m_\ell(\delta) + g_\ell(\delta) - H(\delta)m(\delta) - g(\delta)\|_X \\ &\leq \|H_\ell(\delta) - H(\delta)\|_{\mathcal{L}(X)} \|m_\ell(\delta)\|_X \\ &\quad + \|H(\delta)\|_{\mathcal{L}(X)} \|m_\ell(\delta) - m(\delta)\|_X + \|g_\ell(\delta) - g(\delta)\|_X \\ &\rightarrow 0,\end{aligned}$$

as $\ell \rightarrow \infty$, and so it follows that

$$m(\delta) = \lim_{\ell \rightarrow \infty} m_{\ell+1}(\delta) = \lim_{\ell \rightarrow \infty} G_\ell(m_\ell(\delta), \delta) = G(m(\delta), \delta).$$

All that remains is to check that the fixed point is unique. Supposing that $h(\delta)$ is another fixed point, it follows that

$$\begin{aligned}\|m(\delta) - h(\delta)\|_X &= \|G(m(\delta), \delta) - G(h(\delta), \delta)\|_X \\ &\leq \|H(\delta)\|_{\mathcal{L}(X)} \|m(\delta) - h(\delta)\|_X \\ &\leq \beta\delta \|m(\delta) - h(\delta)\|_X.\end{aligned}$$

Hence, for $\delta < 1/\beta$ we must have that $m(\delta) = h(\delta)$, and the result follows. \square

3.3. Convergence in finite dimensions

Above we had to assume that the prior distribution was sufficiently close to the posterior in order to guarantee convergence; in finite dimensions we may drop this assumption and still get convergence of the covariances. Here we assume that X and Y are Euclidean spaces, $X = \mathbb{R}^N$ and $Y = \mathbb{R}^J$.

We use the following convention: for symmetric matrices B_1, B_2 , we write $B_1 \geq B_2$, or $B_1 > B_2$ to indicate that $B_1 - B_2$ is non-negative definite, or positive definite, respectively. We denote by $S_+^N \subseteq \mathbb{R}^{N \times N}$ the set of positive definite $N \times N$ matrices. In this section, the adjoint operators are denoted as transposes.

We start by showing that the iterative updating formula (4) gives a convergent sequence. However, instead of the covariance matrices, it is more convenient to work with the precision matrices, defined as $B_\ell := C_\ell^{-1}$. Observe that formula (4) can be written as

$$C_{\ell+1}^{-1} = A^T(\Gamma + MC_\ell M^T)^{-1}A + C_0^{-1},$$

which motivates the following result.

Proposition 3.6. *Let $B_0 = C_0^{-1} \in S_+^N$ be a positive definite precision matrix, and let the sequence $\{B_\ell\}_{\ell \geq 0}$ be generated iteratively by the formula $B_{\ell+1} = R(B_\ell)$, where $R : S_+^N \rightarrow S_+^N$ is given by*

$$R(B) = A^T(\Gamma + MB^{-1}M^T)^{-1}A + B_0.$$

Then the sequence $\{B_\ell\}_{\ell \geq 0}$ is increasing in the sense of quadratic forms, and there exists a positive definite $B \in S_+^N$ such that $B_\ell \uparrow B$.

Consequently, the sequence of covariances $\{C_\ell\}_{\ell \geq 0}$ defined by (4) satisfies $C_\ell \downarrow C := B^{-1}$.

Proof. We first show that $B_{\ell+1} \geq B_\ell$ for all ℓ using induction. Write

$$r(B) = A^T(\Gamma + MB^{-1}M^T)^{-1}A,$$

so that $R(B) = r(B) + B_0$. If $B > 0$, then $r(B) \geq 0$, proving that $B_1 = r(B_0) + B_0 \geq B_0$.

Now assume that $B_\ell \geq B_{\ell-1}$. Then we have

$$\begin{aligned} B_{\ell+1} - B_\ell &= R(B_\ell) - R(B_{\ell-1}) \\ &= r(B_\ell) - r(B_{\ell-1}) \\ &= A^T \left((\Gamma + MB_\ell^{-1}M^T)^{-1} - (\Gamma + MB_{\ell-1}^{-1}M^T)^{-1} \right) A. \end{aligned}$$

To prove the claim, it suffices to show that the bracketed difference is non-negative definite. Consider the difference

$$(\Gamma + MB_{\ell-1}^{-1}M^T) - (\Gamma + MB_\ell^{-1}M^T) = M(B_{\ell-1}^{-1} - B_\ell^{-1})M^T \geq 0$$

by the induction assumption. Therefore,

$$(\Gamma + MB_{\ell-1}^{-1}M^T)^{-1} \leq (\Gamma + MB_{\ell}^{-1}M^T)^{-1},$$

which implies the desired non-negative definiteness.

To prove that the sequence $\{B_{\ell}\}_{\ell \geq 0}$ is bounded as quadratic forms, denote the quadratic form as

$$Q_{\ell}(u, v) = u^T B_{\ell} v, \quad u, v \in \mathbb{R}^N.$$

Since

$$(\Gamma + MB_{\ell-1}^{-1}M^T)^{-1} \leq \Gamma^{-1},$$

we have

$$Q_{\ell}(u, u) \leq (Au)^T \Gamma^{-1} Au + u^T B_0 u,$$

proving the boundedness of the sequence.

In particular, it follows that for each $u \in \mathbb{R}^N$,

$$Q_{\ell}(u, u) \rightarrow u^T B u, \quad \text{as } \ell \rightarrow \infty$$

for some symmetric positive definite matrix $B \in \mathbb{R}^{N \times N}$. The fact that the matrix entries of B_{ℓ} converge to the corresponding entries of B follows from the polar identity,

$$u^T B_{\ell} v = \frac{1}{4} (Q_{\ell}(u+v) - Q_{\ell}(u-v)),$$

with u, v being the canonical basis vectors. This completes the proof. \square

4. The general case

In general, the sequence of distributions π_{ℓ} are not Gaussian, and so it is considerably harder to analyze the convergence as we did in the previous section. In this section we consider how the algorithm may be implemented in practice, and in particular, how to produce an approximation to the posterior distribution using a finite number of full model evaluations. This approximate distribution can be used for generating samples using only approximate model evaluations, leading to a significantly lower computational cost over sampling using the true posterior based on the full model.

In section 4.1 we outline the general framework for sampling from the approximate posterior sequence and updating the density, making use of particle approximations. In section 4.2 we reformulate the iteration (3) in terms of operators on the set of probability measures, and provide results on properties of these operators. Convergence in the large particle limit is shown, using the new formulation of the update. In section 4.3 a particular rejection sampling method, based on a Gaussian mixture proposal, is studied. Importance sampling is then considered in section 4.4 and similar convergence is shown.

4.1. Updating the densities via sampling

We consider the algorithm in section 2.3, and in particular, address the question of how to generate a sequence of approximate samples from the iteratively defined densities μ_ℓ given by (3). We shall use particle approximations to do this. Assume that

$$\mathcal{S}_\ell = \{(u_\ell^1, w_\ell^1), (u_\ell^2, w_\ell^2), \dots, (u_\ell^N, w_\ell^N)\}, \quad \ell = 0, 1, 2, \dots$$

is the current approximate sample of the unknowns with relative weights w_ℓ^j . For $\ell = 0$, the sample is obtained by independent sampling from the prior, and $w_0^j = 1/N$. We then compute the modeling error sample,

$$\mathcal{M}_\ell = \{m_\ell^1, m_\ell^2, \dots, m_\ell^N\}, \quad \ell = 0, 1, 2, \dots$$

by defining

$$m_\ell^j = M(u_\ell^j).$$

Consider now the model (2). Assuming that the modeling error is independent of the unknown u , we may write a conditional likelihood model,

$$\pi(b | u, m) \propto \pi_{\text{noise}}(b - f(u) - m).$$

Let $\nu_\ell(m)$ denote the probability density of the modeling error based on our current information. Then, the updated likelihood model based on the approximate model is

$$\pi_{\ell+1}(b | u) = \int \pi(b | u, m) \nu_\ell(m) dm,$$

and, using a Monte Carlo integral approximation, postulating that the realizations m_ℓ^j inherit the weights of the sample points u_ℓ^j , we obtain

$$\pi_{\ell+1}(b | u) \approx \sum_{j=1}^N w_\ell^j \pi(b | u, m_\ell^j).$$

The current approximation for the posterior density is

$$\pi_{\ell+1}(u | b) \propto \pi_{\text{prior}}(u) \sum_{j=1}^N w_\ell^j \pi(b | u, m_\ell^j),$$

suggesting an updating scheme for $\mathcal{S}_\ell \rightarrow \mathcal{S}_{\ell+1}$:

- (a) Draw an index k_j by replacement from $\{1, 2, \dots, N\}$, using the probabilities w_ℓ^j ;
- (b) Draw the sample point

$$u_{\ell+1}^j \sim \pi_{\text{prior}}(u) \pi(b | u, m_\ell^{k_j}). \quad (6)$$

Part (b) above is straightforward, in particular, if the model is Gaussian and f is linear, such as in the linearized model for EIT, since the measure (6) is then a Gaussian. We will demonstrate the effectiveness of this approach in section 5.2. Otherwise we may consider other sampling methods such as importance sampling; this is what is done in the following subsections.

4.2. A convergence result for particle approximations

In this section, we rewrite the updating formula in terms of mappings of measures, and analyze the convergence of the particle approximation under certain limited conditions.

Let μ_ℓ denote the current approximation of the posterior density for u . The updated likelihood based on the modeling error is

$$\pi_{\ell+1}(b | u) \propto \int_X \pi_{\text{noise}}(b - f(u) - M(z)) \mu_\ell(dz),$$

and therefore, the updating, by Bayes' formula, is given by

$$\mu_{\ell+1}(du) \propto \mu_{\text{prior}}(du) \int_X \pi_{\text{noise}}(b - f(u) - M(z)) \mu_\ell(dz) = P\mu_\ell(du). \quad (7)$$

Furthermore, we write the normalization formally as an operator,

$$L\mu = \frac{\mu}{\mu(1)}, \quad \mu(1) = \int_X \mu(du).$$

The model updating algorithm can therefore be written concisely as

$$\mu_{\ell+1} = LP\mu_\ell, \quad \mu_0 = \mu_{\text{prior}}.$$

Let $M(X)$ denote the set of finite measures on X . Denote by $P(X)$ the set of probability measures on X , and for $p \in (0, 1)$ denote by $M_p(X)$ the set of finite measures with total mass lying in the interval $[p, p^{-1}]$.

Let μ and ν denote two random $M(X)$ -valued measures, i.e. $\mu^\omega, \nu^\omega \in M(X)$ for $\omega \in \Omega$, where Ω is a probability space. Denoting by \mathbb{E} the expectation, we define the distance between random measures through

$$d(\mu, \nu)^2 = \sup_{\|\varphi\|_\infty=1} \mathbb{E} |\mu(\varphi) - \nu(\varphi)|^2,$$

where the functions φ are continuous over X . For non-random measures, the definition coincides with the total variation distance.

In the following two lemmas, which we need for the large particle convergence result that follows them, we make the following assumption about the model, prior and noise.

Assumption 4.1. *There exists $\kappa \in (0, 1)$ such that for all $z \in X$,*

$$\kappa \leq \int_X \pi_{\text{prior}}(u) \pi_{\text{noise}}(b - f(u) - M(z)) du \leq \kappa^{-1}.$$

Observe that this assumption holds if, for example, π_{noise} is Gaussian and both f, F are bounded. It also holds if f is linear rather than bounded, as will be established later. Under this assumption, we show the following results concerning the mappings P and L defined before.

Lemma 4.2. *Let assumptions 4.1 hold. Then $P : P(X) \rightarrow M_\kappa(X)$, and*

$$d(P\mu, P\nu) \leq \kappa^{-1} d(\mu, \nu).$$

Proof. First note that the assumption implies that $\kappa \leq (P\mu)(1) \leq \kappa^{-1}$, and so P does indeed map $P(X)$ into $M_\kappa(X)$. Exchanging the order of integration, we see for any bounded measurable φ ,

$$(P\mu)(\varphi) = \int_X \underbrace{\left(\int_X \pi_{\text{prior}}(u) \pi_{\text{noise}}(b - f(u) - M(z)) \varphi(u) du \right)}_{=: \psi(z)} \mu(dz)$$

and so

$$|(P\mu)(\varphi) - (P\nu)(\varphi)|^2 = |\mu(\psi) - \nu(\psi)|^2.$$

Using the assumption we see that $\|\varphi\|_\infty \leq 1$ implies that $\|\psi\|_\infty \leq \kappa^{-1}$, and so

$$\begin{aligned} d(P\mu, P\nu)^2 &\leq \sup_{\|\psi\|_\infty \leq \kappa^{-1}} \mathbb{E} |\mu(\psi) - \nu(\psi)|^2 \\ &\leq \sup_{\|\psi\|_\infty \leq 1} \kappa^{-2} \mathbb{E} |\mu(\psi) - \nu(\psi)|^2 \\ &= \kappa^{-2} d(\mu, \nu)^2, \end{aligned}$$

implying the claim. \square

A similar result for the mapping L can be obtained.

Lemma 4.3. *Let assumptions 4.1 hold. Then it follows that $L : M_\kappa(X) \rightarrow P(X)$, and furthermore, for $\mu, \nu \in M_\kappa(X)$, we have*

$$d(L\mu, L\nu) \leq 2\kappa^{-2} d(\mu, \nu).$$

Proof. The proof is essentially identical to that of lemma 5.17 in [15], with 1 in place of g . We skip the details here. \square

We use the above results to analyze the convergence of particle approximations of the measures. We introduce the sampling operator $S^N : P(X) \rightarrow P(X)$,

$$S^N \mu = \frac{1}{N} \sum_{j=1}^N \delta_{u^j}, \quad u^1, \dots, u^N \sim \mu \text{ i.i.d.}$$

and we have

$$(S^N \mu)(\varphi) = \frac{1}{N} \sum_{j=1}^N \varphi(u^j), \quad u^1, \dots, u^N \sim \mu \text{ i.i.d.}$$

Observe that $S^N \mu$ is a random measure, as it depends on the sample. It is shown in [15], lemma 5.15, that the operator S^N satisfies

$$\sup_{\mu \in P(X)} d(S^N \mu, \mu) \leq \frac{1}{\sqrt{N}}.$$

Define the sequence of particle approximations $\{\mu_\ell^N\}_{\ell \geq 0}$ to $\{\mu_\ell\}_{\ell \geq 0}$ by

$$\begin{aligned} \mu_0^N &= S^N \mu_0, \\ \mu_{\ell+1}^N &= S^N L P \mu_\ell^N. \end{aligned} \tag{8}$$

in light of the previous lemmas, now we prove the following result regarding convergence of this approximation as $N \rightarrow \infty$:

Proposition 4.4. *Let assumptions 4.1 hold. Define $\{\mu_\ell\}_{\ell \geq 0}$, $\{\mu_\ell^N\}_{\ell \geq 0}$ as above. Then, for each ℓ ,*

$$d(\mu_\ell^N, \mu_\ell) \leq \frac{1}{\sqrt{N}} \sum_{k=0}^{\ell} (2\kappa^{-3})^k.$$

In particular, $d(\mu_\ell^N, \mu_\ell) \rightarrow 0$ as $N \rightarrow \infty$.

Proof. The triangle inequality for $d(\cdot, \cdot)$ yields

$$e_\ell := d(\mu_\ell^N, \mu_\ell) \leq d(S^N LP\mu_{\ell-1}^N, LP\mu_{\ell-1}^N) + d(LP\mu_{\ell-1}^N, LP\mu_{\ell-1}),$$

and applying the bounds given by the previous lemmas, we obtain

$$\begin{aligned} e_\ell &\leq \frac{1}{\sqrt{N}} + 2\kappa^{-2} d(P\mu_{\ell-1}^N, P\mu_{\ell-1}) \\ &\leq \frac{1}{\sqrt{N}} + 2\kappa^{-3} d(\mu_{\ell-1}^N, \mu_{\ell-1}) \\ &= \frac{1}{\sqrt{N}} + 2\kappa^{-3} e_{\ell-1}. \end{aligned}$$

The result follows since $e_0 = d(S^N \mu_0, \mu_0) \leq 1/\sqrt{N}$. □

4.3. Particle approximation with Gaussian densities

In this section, we consider the particle approximation when the approximate model is linear, while the accurate model need not be. This is the situation in the computed examples that will be discussed later.

Suppose that the approximate model f is linear, $f(u) = Au$, and the noise and prior distributions are Gaussian,

$$\pi_{\text{noise}} = \mathcal{N}(0, \Gamma), \quad \pi_{\text{prior}} = \mathcal{N}(m_0, C_0).$$

Then the measure (6) is a Gaussian mixture:

$$\pi_{\text{prior}}(u)\pi(b | u, m_\ell^j) = \frac{1}{N} \sum_{j=1}^N \mathcal{N}(u | p_\ell^j, C)$$

where the means and covariance are given by

$$\begin{aligned} C &= (A^T \Gamma^{-1} A + C_0)^{-1}, \\ p_\ell^j &= C(A^T \Gamma^{-1}(b - m_\ell^j) + C_0^{-1} m_0). \end{aligned}$$

The collection of samples \mathcal{S}_ℓ can then be evolved via the following algorithm.

Algorithm (Linear approximate model).

1. Set $\ell = 0$. Define the covariance operator $C = (A^T \Gamma^{-1} A + C_0)^{-1}$. Draw an initial ensemble of particles $\{u_\ell^j\}_{j=1}^N$ from the prior measure $\mu_0(du) = \pi_{\text{prior}}(u)du$, and define the collection $\mathcal{S}_\ell = \{u_\ell^1, u_\ell^2, \dots, u_\ell^N\}$.
2. Define the means $p_\ell^j = C(A^T \Gamma^{-1}(b - M(u_\ell^j)) + C_0^{-1} m_0)$, $j = 1, \dots, N$.

3. For each $j = 1, \dots, N$
 - (i) Sample k_j uniformly from the set $\{1, \dots, N\}$
 - (ii) Sample $u_{\ell+1}^j \sim \mathcal{N}(p_{\ell}^{k_j}, \mathbf{C})$
4. Set $\mathcal{S}_{\ell+1} = \{u_{\ell+1}^1, u_{\ell+1}^2, \dots, u_{\ell+1}^N\}$.
5. Set $\ell \mapsto \ell + 1$ and go to 2.

Note that the theory applies to the above example when the accurate model F is bounded:

Lemma 4.5. *Let $\pi_{\text{prior}} = \mathcal{N}(m_0, \mathbf{C}_0)$ and $\pi_{\text{noise}} = \mathcal{N}(0, \Gamma)$ be Gaussian. Assume that $f(\cdot) = \mathbf{A}\cdot$ is linear and F is bounded. Then assumption 4.1 holds.*

Proof. Let $F_0 > 0$ be such that $\|F(z)\|_{\Gamma} \leq F_0$ for all $z \in X$. We have via Cauchy-Schwarz and Young's inequalities

$$\begin{aligned}
 \pi_{\text{noise}}(0) &\geq \pi_{\text{noise}}(b - f(u) - M(z)) \\
 &= \pi_{\text{noise}}(b - \mathbf{A}(u - z) - F(z)) \\
 &= D_1 \exp\left(-\frac{1}{2}\|(b + \mathbf{A}z) - \mathbf{A}u - F(z)\|_{\Gamma}^2\right) \\
 &= D_1 \exp\left(-\frac{1}{2}(\|(b + \mathbf{A}z) - \mathbf{A}u\|_{\Gamma}^2 - 2\langle(b + \mathbf{A}z) - \mathbf{A}u, F(z)\rangle_{\Gamma} + \|F(z)\|_{\Gamma}^2)\right) \\
 &\geq D_1 \exp(-\|(b + \mathbf{A}z) - \mathbf{A}u\|_{\Gamma}^2 - F_0^2) \\
 &= D_2 \exp(-\|(b + \mathbf{A}z) - \mathbf{A}u\|_{\Gamma}^2).
 \end{aligned}$$

Integrating with respect to the prior distribution, we therefore have that for any $z \in X$,

$$\begin{aligned}
 \pi_{\text{noise}}(0) &\geq \int_X \pi_{\text{prior}}(u) \pi_{\text{noise}}(b - f(u) - M(z)) \, du \\
 &\geq D_3 \int_X \exp\left(-\|(b + \mathbf{A}z) - \mathbf{A}u\|_{\Gamma}^2 - \frac{1}{2}\|u - m_0\|_{\mathbf{C}_0}^2\right) \, du \\
 &= D_3 \sqrt{\det\left((\sqrt{2}\mathbf{A}^{\top}\Gamma^{-1}\mathbf{A} + \mathbf{C}_0^{-1})^{-1}\right)}
 \end{aligned}$$

which gives the result. \square

We consider this setup later numerically in sections 5.2 and 5.3 in the contexts of electrical impedance tomography and groundwater flow problems respectively. In these cases the accurate model F is not bounded and so we cannot deduce convergence in the large particle limit; however in practice we may modify the map F to be bounded by composing with a thresholding map, corresponding to an additional prior belief that the conductivity/permeability fields are bounded uniformly below by positive constants.

Remark 4.6. For more general models, one could use a method such as rejection sampling in order to produce exact samples from the measure (6). A suitable proposal distribution for this rejection sampling could be, for example, a Gaussian mixture with appropriately chosen means and covariances [9].

Two natural candidates for non-Gaussian priors, that retain some of the simplicity of the Gaussian models without being as limited, are:

- (i) Hierarchical, conditionally Gaussian prior models,

$$\pi_{\text{prior}}(\mathbf{u} \mid \theta) \sim \mathcal{N}(\mu_{\theta}, \mathbf{C}_{\theta}),$$

where the mean and covariance depend on a hyperparameter vector θ that follows a hyperprior distribution,

$$\theta \sim \pi_{\text{hyper}}.$$

The hypermodels allow the introduction of sparsity promoting priors, similar to total variation; [10, 11].

- (ii) Gaussian mixtures, which allow a fast sampling from non-Gaussian distributions through a local approximation by Gaussian or other simple distributions [33].

4.4. Importance sampling and convergence

In this section we consider an approximate sampling based updating scheme of the probability densities using importance sampling. This method effectively turns a collection of prior samples into samples from the posterior by weighting them appropriately, using the fact that the posterior is absolutely continuous with respect to the prior.

Assume that at stage ℓ of the approximation scheme, we have a collection of N particles and the corresponding weights, $\mathcal{S}_{\ell} = \{(u_{\ell}^j, w_{\ell}^j)\}_{j=1}^N$. The associated particle approximation μ_{ℓ}^N of the probability distribution acting on a test function φ is

$$\mu_{\ell}^N(\varphi) = \sum_{j=1}^N w_{\ell}^j \varphi(u_{\ell}^j).$$

We evolve this distribution by acting on it with P and L . By the definition (7) of P , we first get an approximation

$$\begin{aligned} P\mu_{\ell}^N(\mathbf{d}u) &= \left(\sum_{j=1}^N w_{\ell}^j \pi_{\text{noise}}(b - f(u) - M(u_{\ell}^j)) \right) \mu_{\text{prior}}(\mathbf{d}u) \\ &=: g_{\ell}(u) \mu_{\text{prior}}(\mathbf{d}u). \end{aligned}$$

To generate an updated sample based on this approximation, we use independent sampling to draw $u_{\ell+1}^j \sim \mu_0$, $j = 1, \dots, N$, and define the particle approximation by

$$\mu_{\ell+1}^N(\varphi) = \sum_{j=1}^N w_{\ell+1}^j \varphi(u_{\ell+1}^j), \quad w_{\ell+1}^j = \frac{g_{\ell}(u_{\ell+1}^j)}{\sum_{j=1}^N g_{\ell}(u_{\ell+1}^j)}.$$

Denoting by $T^N : \mathcal{P}(X) \rightarrow \mathcal{P}(X)$ the importance sampling step, consisting of independent sampling and weighting, we may define an iterative algorithm symbolically as

$$\begin{aligned} \mu_0^N &= T^N \mu_0, \\ \mu_{\ell+1}^N &= T^N L P \mu_{\ell}^N. \end{aligned} \tag{9}$$

Explicitly, the algorithm can be described as follows.

Algorithm (Importance sampling).

1. Set $\ell = 0$. Draw an initial ensemble of particles $\{u_\ell^j\}_{j=1}^N$ from the prior measure $\mu_0(\mathrm{d}u) = \pi_{\text{prior}}(u)\mathrm{d}u$, and initialize the weights $w_\ell^j = 1/N$ for each $j = 1, \dots, N$. Define the collection $\mathcal{S}_\ell = \{(u_\ell^1, w_\ell^1), (u_\ell^2, w_\ell^2), \dots, (u_\ell^N, w_\ell^N)\}$.
2. Define

$$g_\ell(u) = \sum_{j=1}^N w_\ell^j \pi_{\text{noise}}(b - f(u) - M(u_\ell^j)).$$

3. Sample $u_{\ell+1}^j \sim \mu_0$, $j = 1, \dots, N$ i.i.d. and define the weights

$$w_{\ell+1}^j = \frac{g_\ell(u_{\ell+1}^j)}{\sum_{j=1}^N g_\ell(u_{\ell+1}^j)}, \quad j = 1, \dots, N.$$

4. Set $\mathcal{S}_{\ell+1} = \{(u_{\ell+1}^1, w_{\ell+1}^1), (u_{\ell+1}^2, w_{\ell+1}^2), \dots, (u_{\ell+1}^N, w_{\ell+1}^N)\}$.
5. Set $\ell \mapsto \ell + 1$ and go to 2.

As in the previous section, we establish a convergence result for $N \rightarrow \infty$ only under the restrictive condition of assumption 4.1. We recall the following result from [1]:

Lemma 4.7. *Let $\mu \in \mathcal{P}(X)$ be absolutely continuous with respect to the prior measure μ_0 ,*

$$\mu(\mathrm{d}u) \propto g(u)\mu_0(\mathrm{d}u),$$

where $\mu_0(g^2) < \infty$. Define the quantity $\rho \geq 1$ by $\rho = \mu_0(g^2)/\mu_0(g)^2$. Then

$$d(T^N \mu, \mu) \leq 2\sqrt{\frac{\rho}{N}}.$$

By assumption 4.1, there exists $\kappa \in (0, 1)$ such that $\kappa \leq g_\ell(\varepsilon) \leq \kappa^{-1}$ for all $\varepsilon \in Y$, implying that

$$\frac{\mu_0(g_\ell^2)}{\mu_0(g_\ell)^2} \leq \kappa^{-4}.$$

In particular, by applying the above lemma to the measure $\hat{\mu}_{\ell+1} = LP\mu_\ell$, we see that

$$d(T^N LP\mu_\ell, LP\mu_\ell) \leq \frac{2\kappa^{-2}}{\sqrt{N}}.$$

We are ready to prove the following proposition establishing the convergence of the particle approximations as $N \rightarrow \infty$:

Proposition 4.8. *Let assumptions 4.1 hold for the noise distribution, and let $\{\mu_\ell\}_{\ell \geq 0}$ be the sequence of the model error approximations of the posterior, and $\{\mu_\ell^N\}_{\ell \geq 0}$ a sequence of importance sampling approximations obtained as above. Then, for each ℓ ,*

$$d(\mu_\ell^N, \mu_\ell) \leq \frac{2\kappa^{-2}}{\sqrt{N}} \sum_{k=0}^{\ell-1} (2\kappa^{-3})^k + \frac{(2\kappa^{-3})^\ell}{\sqrt{N}}.$$

In particular, $d(\mu_\ell^N, \mu_\ell) \rightarrow 0$ as $N \rightarrow \infty$.

Proof. From the triangle inequality for $d(\cdot, \cdot)$, we have

$$e_\ell := d(\mu_\ell^N, \mu_\ell) \leq d(T^N LP \mu_{\ell-1}^N, LP \mu_{\ell-1}^N) + d(LP \mu_{\ell-1}^N, LP \mu_{\ell-1}).$$

The bounds derived above yield

$$\begin{aligned} e_\ell &\leq \frac{2\kappa^{-2}}{\sqrt{N}} + 2\kappa^{-2} d(P \mu_{\ell-1}^N, P \mu_{\ell-1}) \\ &\leq \frac{2\kappa^{-2}}{\sqrt{N}} + 2\kappa^{-3} d(\mu_{\ell-1}^N, \mu_{\ell-1}) \\ &= \frac{2\kappa^{-2}}{\sqrt{N}} + 2\kappa^{-3} e_{\ell-1}. \end{aligned}$$

Since we have $e_0 = d(T^N \mu_0, \mu_0) = d(S^N \mu_0, \mu_0) \leq 1/\sqrt{N}$, the result follows. \square

Remark 4.9. In theory, the importance sampling method described above can be used with very weak assumptions on the forward maps and prior/noise distributions. However in practice it may be ineffective if the posterior is significantly far from the prior, such as when the size of the observational noise is small. To overcome this issue, one could instead consider Sequential Monte Carlo or sequential importance sampling methods to evolve prior samples into posterior samples by introducing a sequence of intermediate measures [5, 26].

5. Numerical illustrations

In this section, we demonstrate the convergence properties established in the preceding sections by means of computed examples. Furthermore, we demonstrate the enhanced reconstructions obtained by modelling error as advocated in this paper. The first example is a linear inverse source problem, elucidating the geometric convergence in the linear Gaussian case. The second example is the EIT problem with linearized approximate model with a coarse FEM mesh, allowing for straightforward particle updates. In the last example we consider the problem of recovering the permeability field in the steady state Darcy flow model, again with a linearized approximate model.

5.1. Inverse source problem

As a proof of concept, we start by considering a simple one-dimensional inverse source problem. Let $\Omega = (0, 1)$ and define $X = L^2(\Omega)$. Given $u \in X$, let $p = P(u) \in H_0^1(\Omega)$ be the solution to the Laplace equation,

$$\begin{cases} -p'' = u & x \in \Omega \\ p = 0 & x \in \partial\Omega. \end{cases}$$

The inverse problem is to estimate the source u from pointwise observations of p . Therefore, define the observation operator $\mathcal{O} : H_0^1(\Omega) \rightarrow \mathbb{R}^J$ by

$$\mathcal{O}(u) = (u(q_1), \dots, u(q_J))$$

for some set of points $\{q_1, \dots, q_J\} \subseteq \Omega$. We define the exact forward operator $A_\star^{\text{exact}} = \mathcal{O} \circ P$. For numerical simulations, the exact forward model A_\star^{exact} is approximated by a high fidelity

Table 1. The approximation errors of the approximate posterior means and covariances for various approximation levels with and without the inclusion of the modeling error correction. The matrix norms are the Frobenius norms.

n	$\ m_L - m_{\text{post}}\ $	$\ m_{\text{post}}^n - m_{\text{post}}\ $	$\ C_L - C_{\text{post}}\ $	$\ C_{\text{post}}^n - C_{\text{post}}\ $
4	0.1906	0.2986	0.0676	0.0381
5	0.0455	0.0739	0.0170	0.0095
6	0.0111	0.0182	0.0043	0.0024
7	0.0028	0.0045	0.0011	0.0006
8	0.0007	0.0011	0.00026	0.00015
9	0.0002	0.0003	0.00006	0.00004

proxy, A_* , obtained by approximating the solution p through a finite difference solution on a fine mesh. The coarse mesh approximation of A_*^{exact} , used in the inverse model, is denoted by $A = A_n$. In our computed example, we use $2^{10} - 1 = 1023$ equally spaced interior points for A_* , while the coarse mesh model A_n is computed with of $2^n - 1$ equally spaced interior points, $n < 10$.

We let $q_j = j/16$, $j = 1, \dots, 15 = J$ be equally spaced observation points, and to generate the simulated data, we corrupt the high fidelity data with a small amount of white noise, $\varepsilon \sim \mathcal{N}(0, \Gamma)$, where we set $\Gamma = 10^{-8}I_J$. The prior is chosen to be a standard Brownian motion, specifically we take $\mu_0 = \mathcal{N}(0, C_0)$ with

$$C_0 = (-\Delta)^{-1}, \quad \Delta = \frac{d^2}{dx^2},$$

$$\mathcal{D}(-\Delta) = \left\{ u \in H^2(\Omega) \mid u(0) = 0, \quad \frac{du}{dx}(1) = 0 \right\},$$

and the true source used for data generation is drawn from the prior. Numerically the precision operator C_0^{-1} is implemented as the finite difference Laplacian matrix. We perform $L = 30$ iterations in each simulation.

The posterior mean and covariance, $m_{\text{post}}, C_{\text{post}}$, corresponding to the high fidelity model, and the corresponding mean and covariance, $m_{\text{post}}^n, C_{\text{post}}^n$, based on the approximate model are given by

$$m_{\text{post}} = m_0 + C_0 A_*^T (\Gamma + A_* C_0 A_*^T)^{-1} (b - A_* m_0),$$

$$C_{\text{post}} = C_0 - C_0 A_*^T (\Gamma + A_* C_0 A_*^T)^{-1} A_* C_0,$$

$$m_{\text{post}}^n = m_0 + C_0 A_n^T (\Gamma + A_n C_0 A_n^T)^{-1} (b - A_n m_0),$$

$$C_{\text{post}}^n = C_0 - C_0 A_n^T (\Gamma + A_n C_0 A_n^T)^{-1} A_n C_0,$$

respectively. The approximate posterior mean and covariances (4) obtained by the modeling error approach, after ℓ iterations, are denoted by m_ℓ and C_ℓ , respectively.

Table 1 shows the approximation errors arising from both approximations of C_{post} and m_{post} with different discretization levels. The table shows that the modeling error approach produces a better approximation of the posterior mean than the model ignoring the modeling error, while the approximate covariances are slightly less accurate as approximations of the posterior covariance than those found without allowing for the modeling error correction. These experiments confirm our assertion at the start of the paper, namely that allowing for model error can result in improved point estimates (here the posterior mean) but that the




Figure 1. The trace of the errors $\|m_\ell - m_L\|$ and $\|C_\ell - C_L\|$, illustrating their convergence rates. From left to right, top to bottom, the approximation level n is increased from 4 to 9.

Table 2. The convergence rates quantified in terms of the slope of the logarithmic plot of mean and covariance deviation from the limit values m_L and C_L before the deviations plateau, indicating that the algorithm has converged.

n	Slope of $\log \ m_\ell - m_L\ $	Slope of $\log \ C_\ell - C_L\ $	$\ A_n - A_\star\ $
4	-1.80	-3.60	0.101
5	-3.13	-6.26	0.0706
6	-4.57	-9.00	0.0491
7	-6.30	-11.9	0.0335
8	-7.73	-14.8	0.0219
9	-9.34	-14.9	0.0127

iteration introduced does not converge to the true posterior distribution (as evidenced by the error in the covariance at fixed n and large L .)

To demonstrate the convergence rate, figure 1 shows the mean and covariance errors for various approximation levels as functions of the number of iterations. The plots, as well as the tabulated values, table 2, of the logarithmic slopes of the approximation errors verify the geometric convergence rates, with their dependence on the approximation level. Observe that the logarithm of the convergence rate for the covariance, a quadratic quantity, is twice that of the mean.

5.2. Electrical impedance tomography (EIT)

In this section, we revisit the modeling error due to coarse discretization of a PDE model in the context of electrical impedance tomography (EIT). Let $\Omega \subset \mathbb{R}^d$, $d = 2, 3$, denote a bounded connected set with boundary $\partial\Omega$, and let $\sigma : \Omega \rightarrow \mathbb{R}$ be a function modeling the electric conductivity in Ω , $0 < \sigma_m \leq \sigma \leq \sigma_M < \infty$. We assume that S electrodes are attached to the boundary $\partial\Omega$, and we model them as open disjoint surface patches $e_s \subset \partial\Omega$, $1 \leq s \leq S$. Assuming that an electric current I_s is injected through e_s into the body modeled by Ω ,

the electric voltage potential v in Ω , and the electrode voltages V_s on the electrodes can be found as a solution of the complete electrode model (CEM) boundary value problem [32],

$$\begin{cases} \nabla \cdot (\sigma \nabla v) = 0 & x \in \Omega \\ \sigma \frac{\partial v}{\partial n} = 0 & x \in \partial\Omega \setminus \bigcup_{s=1}^S e_s \\ v + z_s \sigma \frac{\partial v}{\partial n} = V_s & x \in e_s, 1 \leq s \leq S \\ \int_{e_s} \sigma \frac{\partial v}{\partial n} dS = I_s & 1 \leq s \leq S. \end{cases}$$

Here, the parameters $z_s > 0$ are the presumably known contact impedances, and the currents satisfy the Kirchhoff's law, or conservation of charge condition,

$$\sum_{s=1}^S I_s \in \mathbb{R}_0^S = \left\{ V \in \mathbb{R}^L \mid \sum_{s=1}^S V_s = 0 \right\}.$$

The solution of the boundary value problem is the unique solution $(v, V) \in H^1(\Omega) \times \mathbb{R}_0^S$ of the weak form variational problem

$$\mathcal{B}((w, W), (v, V)) = \sum_{s=1}^S I_s W_s = \langle (w, W), b_I \rangle, \quad \text{for all } (w, W) \in H^1(\Omega) \times \mathbb{R}_0^S,$$

where $b_I = (0, I) \in H^1(\Omega) \times \mathbb{R}_0^S$, and

$$\mathcal{B}((w, W), (v, V)) = \int_{\Omega} \sigma \nabla w \cdot \nabla v dx + \sum_{s=1}^S \frac{1}{z_s} \int_{e_s} (w - W_s)(v - V_s) dS.$$

To discretize the problem, assume that Ω is approximated by the union of triangular or tetrahedral elements, the mesh containing n_f nodes ('f' for fine), and let $\{\psi_j\}_{j=1}^{n_f}$ denote a nodal-based piecewise polynomial Lagrange basis. Further, let $\{\phi_s\}_{s=1}^{S-1}$ denote a basis of \mathbb{R}_0^S . We define the basis functions $\bar{\psi}_j \in H^1(\Omega) \times \mathbb{R}_0^S$ as

$$\bar{\psi}_j = (\psi_j, 0), \quad 1 \leq j \leq n_f, \quad \bar{\psi}_{n_f+s} = (0, \phi_s), \quad 1 \leq s \leq S-1.$$

We approximate the potential-voltage pair (v, V) as

$$(v, V) = \sum_{j=1}^{n_f+S-1} \alpha_j \bar{\psi}_j,$$

and discretize the forward problem by choosing $(w, W) = \bar{\psi}_k$, to arrive at the Galerkin approximation,

$$\sum_{j=1}^{n_f+S-1} \mathcal{B}(\bar{\psi}_k, \bar{\psi}_j) \alpha_j = \langle \bar{\psi}_k, b_I \rangle, \quad 1 \leq k \leq n_f + S - 1. \quad (10)$$

Further, to parametrize the conductivity, we define a discretization of Ω by triangular or tetrahedral elements, independent of the discretization above, with K nodes, and denote by $\{\eta_j\}_{j=1}^K$ the nodal-based piecewise polynomial Lagrange basis functions. We then parametrize the conductivity by writing

$$\sigma(x) = \sigma_0 \exp \left(\sum_{j=1}^K u_j \eta_j(x) \right), \quad x \in \Omega,$$

where $\sigma_0 > 0$ is a fixed background conductivity. The matrix $[\mathcal{B}(\bar{\psi}_k, \bar{\psi}_j)]$ defining the system (10) is parametrized by the vector $u \in \mathbb{R}^K$, and we write the equation in matrix form concisely as

$$A_u^{n_f} \alpha = b(I),$$

where we have indicated explicitly the dependency on the discretization by the number n_f of nodes. Solving this system for α , extracting the last $S - 1$ components α_{n_f+s} , $1 \leq s \leq S - 1$, and representing the voltage in terms of the basis functions ϕ_s defines the forward map

$$u \mapsto V = \sum_{s=1}^{S-1} \alpha_{n_f+s} \phi_s = R_u^{n_f} I, \quad \text{where } \alpha = (A_u^{n_f})^{-1} b(I),$$

where $R_u^{n_f} \in \mathbb{R}^{S \times S}$ is the resistance matrix. We repeat the calculation for a full frame of $S - 1$ linearly independent current patterns, $I^1, \dots, I^{S-1} \in \mathbb{R}_0^S$, obtaining the full frame of voltage patterns V^1, \dots, V^{S-1} . Finally, the voltage patterns are stacked together in a vector, constituting the forward model for the observation,

$$V = \begin{bmatrix} V^1 \\ \vdots \\ V^{S-1} \end{bmatrix} = F^{n_f}(u), \quad F^{n_f} : \mathbb{R}^K \rightarrow \mathbb{R}^{S(S-1)}.$$

To guarantee satisfactory accuracy of the forward model, the discretization needs to be fine enough, in particular to capture the singularities of the voltage potential v at the electrode edges. To demonstrate the modeling error effect, we construct a forward map defined over a coarser FEM mesh with n_c nodes ('c' for coarse), $n_c < n_f$, and denote the corresponding forward map by

$$V = F^{n_c}(u), \quad F^{n_c} : \mathbb{R}^K \rightarrow \mathbb{R}^{S(S-1)}, \quad n_c < n_f.$$

Observe that the discretization of u is independent of the FEM mesh, and is not changed when passing to a coarser computational mesh. In our computed examples, we use a piecewise linear Lagrange basis to represent both u and v over the different meshes. The three meshes that we base our simulations, generated with the mesh generator described in [28], are shown in figure 2. The number of electrodes is $S = 16$.

We assign the Whittle–Matérn prior [29, 30] for the vector u defining the conductivity so that

$$\zeta \lambda^{-1} (-\lambda^2 L_g + I_K) u \sim \mathcal{N}(0, I_K), \quad (11)$$

where $L_g \in \mathbb{R}^{K \times K}$ is the graph Laplacian defined on the conductivity mesh, $\lambda > 0$ is a correlation length parameter, $\zeta > 0$ is amplitude scaling, and I_K is the identity matrix. In figure 3, three independently drawn realizations of the conductivity distributions are shown. The values of the model parameters are indicated in the figure caption.

We generate the data using the fine scale model $F = F^{n_f}$, and using the conventional error model, i.e. ignoring the modeling error, compute a MAP estimate u_{MAP} using the forward map $f = F^{n_c}$ in the inverse solver. The estimate is based on a simple Gauss-Newton iteration. The additive noise covariance in this simulation is $\Gamma = \gamma^2 I_{S(S-1)}$ with $\gamma = 10^{-3} V_{\text{max}}$, where V_{max} is the maximum of all noiseless electrode voltages over the full frame of $S - 1$ voltage patterns. The noise level is assumed to be low so that the modeling error is the predominant part in the uncertainty. In figure 4, we show the conductivity distribution that was used to generate the synthetic data with the model F , the conventional error model MAP estimate based on the




Figure 2. Triangular meshes used in the numerical simulations. The number of electrodes is $L = 16$, and they are indicated by red nodal points in the plot. The mesh for representing the conductivity distribution (left) has $K = 733$ vertices and 1364 elements. The coarse mesh for the forward solver (middle) has $n_c = 877$ vertices and 1592 elements, and the fine scale mesh (right) consist of $n_f = 2418$ vertices and 4562 elements.




Figure 3. Conductivities $\sigma = \sigma_0 \exp(u)$ corresponding to three independent draws of u from the prior density. The parameter values used here are $\lambda = 0.2$ and $\zeta = 1/15$. The background conductivity is $\sigma_0 = 1.5$. The radius of the disc is unity, and the units are arbitrary.

coarse mesh model f , as well as the enhanced error model estimate. In the latter, the modeling error mean and covariance are estimated from a sample of 1 500 random draws from the prior of u .

Observe that in the reconstruction based on the conventional error model, the true inclusions are completely overshadowed by the boundary artifacts that are concentrated around the edges of the electrodes. This is to be expected, since the basis functions in the coarse FEM mesh do not capture the voltage singularities at the electrode edges, and the inverse solution compensates the modeling error with elevated conductivity at the edges to mitigate the singularity. In agreement with previously published results, the enhanced error model produces a solution without modeling error artifacts.

The computation of the MAP estimate, regardless of the error model, requires repeated linearization of the forward map. The re-evaluation of the Jacobian may be time consuming, and therefore it is tempting to replace the coarse mesh FEM model with a linearized approximation around the background conductivity σ_0 corresponding to $u = 0$,




Figure 4. Left: the true conductivity used to generate the test data using the finely discretized FEM forward model. Center: a Gauss–Newton-based MAP estimate based on the coarsely discretized FEM forward model, using the conventional error model that ignores the modeling error. Right: the MAP estimate computed by using the enhanced error model, in which the modeling error mean and covariance are estimated from 1500 random draws from the prior.

$$f(u) = F^{n_c}(0) + DF^{n_c}(0)u.$$

The solution of the inverse problem with the linearized model and Gaussian prior is particularly straightforward, requiring a solution of a linear system. We iterate the posterior updating algorithm, generating samples $\mathcal{S}_\ell = \{u_\ell^1, \dots, u_\ell^N\}$, $\ell = 0, 1, 2, \dots$ using the modeling error updating scheme. In figure 5, we plot the conductivities corresponding to the posterior means,

$$\bar{\sigma}_\ell = \sigma_0 \exp(\bar{u}_\ell), \quad \bar{u}_\ell = \frac{1}{N} \sum_{j=1}^N u_\ell^j,$$

as well as the marginal variances of the parameters u_ℓ , that is,

$$\text{var}_\ell = \text{diag} \left(\frac{1}{N} \sum_{j=1}^N (u_\ell^j - \bar{u}_\ell)(u_\ell^j - \bar{u}_\ell)^\top \right).$$

The sample size here was $N = 5000$.

Finally, we consider the convergence of the iterated densities π_ℓ towards the posterior density by means of the Kullback–Leibler divergence, which we approximate using the particles drawn from π_ℓ ,

$$\begin{aligned} D_{\text{KL}}(\pi_\ell \parallel \pi_{\text{post}}) &= \int \pi_\ell(u) \log \left(\frac{\pi_\ell(u)}{\pi_{\text{post}}(u)} \right) du \approx \frac{1}{N} \sum_{j=1}^N \log \left(\frac{\pi_\ell(u_\ell^j)}{\pi_{\text{post}}(u_\ell^j)} \right) \\ &= \frac{1}{N} \sum_{j=1}^N \log \left(\frac{\pi_\ell(b | u_\ell^j)}{\pi(b | u_\ell^j)} \right) - \log \left(\frac{\pi(b)}{\pi_\ell(b)} \right), \end{aligned}$$

the second term corresponding to the normalization factors of the true and approximate posteriors. Observe that to evaluate the posterior density, the fine mesh model needs not to be evaluated anew, since the fine mesh evaluations are already computed for the modeling error




Figure 5. Sample means of the conductivity \bar{u}_ℓ for $\ell = 1, 2, 3, 4$ (upper row), and sample marginal variances of the components of the vectors u at the same iterations.

sample. The sample-based approximation of the KL divergence is straightforward to compute up to the normalizing constants.

Figure 6 shows the sample-based estimates of the Kullback–Leibler divergence for $\ell = 1, 2, \dots, \ell_{\max} = 5$. To subtract the unknown normalization offset, we plot the differences

$$\Delta D_{\text{KL}}(\pi_\ell \| \pi_{\text{post}}) = D_{\text{KL}}(\pi_\ell \| \pi_{\text{post}}) - D_{\text{KL}}(\pi_{\ell_{\max}} \| \pi_{\text{post}}). \quad (12)$$

The figure shows also the relative error of the sample mean approximating the true conductivity,

$$e_r(\bar{u}_\ell) = \frac{\|\sigma - \bar{\sigma}_\ell\|}{\|\sigma\|}.$$

As in the previous subsection, the numerical results demonstrate that the approach to model error advocated in this paper leads to improved estimates of the true value used to generate the data.

5.3. Steady state Darcy flow

In the last computed example, we consider the inverse problem of estimating the permeability distribution in porous medium from a discrete set of pressure measurements. More precisely, let the computational domain be $\Omega = (0, 1)^2$, and define $X = L^\infty(\Omega)$. For a given $u \in X$, called *log-permeability* define the pressure field $p = P(u) \in H_0^1(\Omega)$ to be the solution to the steady-state Darcy equation with Dirichlet boundary conditions,

$$\begin{cases} -\nabla \cdot (e^u \nabla p) = g & x \in \Omega \\ p = 0 & x \in \partial\Omega \end{cases} \quad (13)$$

for some fixed and presumably known source term $g \in H^{-1}(\Omega)$.

We define now the observation operator $\mathcal{O} : H_0^1(\Omega) \rightarrow \mathbb{R}^J$ by

$$\mathcal{O}_j(p) = \frac{1}{2\pi\varepsilon} \int_\Omega p(x) e^{-\frac{1}{2\varepsilon^2}(x-q_j)^2} dx, \quad j = 1, \dots, J,$$




Figure 6. Left: the Kullback–Leibler divergence (12) difference estimated by using a sample of 5000 realizations drawn from the approximate posterior densities π_ℓ . Right: the relative distance of the sample mean from the true conductivity over five iterations.

for some set of points $\{q_1, \dots, q_J\} \subseteq \Omega$. The noiseless data consist of smoothed observations at the points $\{q_1, \dots, q_J\}$, converging to point observations as $\varepsilon \rightarrow 0$. Note that each \mathcal{O}_j is a bounded linear functional on $H_0^1(\Omega)$. In what follows, we choose $\varepsilon = 0.02$, $g(x_1, x_2) = 100 \sin(\pi x_1) \sin(\pi x_2)$, and let $\{q_1, \dots, q_{25}\}$ be a uniformly spaced grid of 25 points in Ω . The accurate model is then defined by the composition $F = \mathcal{O} \circ P$. As in the EIT example, the approximate model is defined through linearization,

$$f(u) = F(u_0) + DF(u_0)u.$$

for some fixed $u_0 \in X$; in this example we choose $u_0 = 0$. To construct the linear model, the derivative may be computed inexpensively using an adjoint method. The computation of the full Jacobian $DF(u_0)$ requires $J + 1$ numerical solutions of a PDE of the form (13), and needs to be performed only once.

In this example, we generate three different data sets corresponding to different noise levels: the noiseless data generated by using the non-linear model is perturbed by additive observational noise drawn from normal distribution $\mathcal{N}(0, \Gamma_i)$, where $\Gamma_i = 10^{-i-1}I$, and $i = 1, 2, 3$. The true log-permeability u^\dagger , defined as the sum of two unnormalized Gaussian densities, is shown in figure 7. In the same figure, the computed pressure field is shown, with the observation points indicated by black dots. Each data set is generated using a uniform mesh of 128×128 points, while in the inverse computations, we use a reduced model with a uniform mesh of 64×64 points. We perform 10 iterations of the posterior updating algorithm, using as few as 100 particles in the particle approximation model.

As in the EIT simulations, we choose a Whittle–Matérn prior distribution for the vector u defining the permeability as given in (11). We make the choices $\lambda = 0.1$ and $\zeta = 1$, and note that here L_g corresponds to the the finite-difference Laplacian on the reduced mesh using the standard 5-point stencil.

In figure 8 the conditional means arising from the different error models and data sets are shown. In this example the conventional and enhanced error models have very similar performance in terms of inferring the conditional mean. They are both able to infer the geometry of the log-permeability field, particularly when the observational noise is small, however they fail to obtain the magnitude. The iterative algorithm proposed in this article is able to obtain both the geometry and magnitude with good accuracy in a small number of iterations. Figure 9 shows the evolution of the size of the error between the conditional mean at iteration ℓ and the




Figure 7. (Left) The true log-permeability used to generate the data. (Right) The true pressure and the observation points $\{q_j\}_{j=1}^{25}$.

true log-permeability field. In all cases the error has converged in 4 or 5 iterations, similarly to what was observed in the EIT experiments. As in the previous two subsections, the numerical results demonstrate that the approach to model error advocated in this paper leads to improved accuracy of the point estimates of the true parameter underlying the data.

5.4. Computational cost considerations

We consider the cost of generating a collection of S (effectively) independent samples from the posterior and from the three approximations considered. Denote G_+ , G_- the costs of evaluating the accurate and approximate models respectively. Let N denote the number of particles in the approximations, and L the number of iterations in the iterative algorithm. Let $\rho_j \geq 1$ denote the reciprocal of the integrated autocorrelation time of each MCMC chain. Finally let Q denote the cost of producing a sample from a given Gaussian distribution.

In table 3 we show, up to constant scalings, the cost of generating these S samples. In the first column we focus on the case where the approximate model f is linear and the noise and prior are Gaussian. For the conventional and enhanced error models, the approximate posterior is then Gaussian, and for the iterative error model it is a Gaussian mixture. After these approximations have been constructed, independent samples may be directly sampled. The cost of constructing the approximation is more expensive for the iterative error model than the enhanced error model by a factor of L , the number of iterations. However, from the results earlier in this section, there appears to be little increase in accuracy from choosing $L > 2$.

In the second column we show the cost of generating the samples when f is nonlinear and π_{noise} is not necessarily Gaussian. The posterior approximations are then no longer Gaussian, and so we use MCMC methods to perform sampling. We consider cases where the prior is Gaussian, as then we may use known dimension robust MCMC methods to generate the samples. We focus on the cost of the pCN method [14] in which each proposed state is constructed using a draw from the prior distribution, and accepted according to its relative likelihood. Typically, when the data is very informative, states in the MCMC chain will be strongly correlated due to smaller jumps in proposed states, and so the reciprocal integrated autocorrelation times ρ_j will be very large. This includes situations when the observational noise is small, which is when accounting for model error is important.




Figure 8. (Left column) Conditional mean arising from conventional error model. (Middle column) Conditional mean arising from enhanced error model. (Right column) Conditional mean arising from iterative error model, iteration 10. From top to bottom, observational noise standard deviation is 10^{-2} , 10^{-3} , 10^{-4} respectively.

For the approximate posteriors, after constructing the approximations, only evaluations of the approximate model f are required to generate samples. Conversely, to sample the full posterior distribution no approximation needs to be constructed, but the accurate model F must be evaluated to generate each sample. Thus which method is the most appropriate to use is going to depend on the balancing of the number of samples required, the cost of evaluating F versus f , and the desired accuracy of approximation. Finally it should also be noted that the construction of the particle approximations may be parallelized in a natural way, whereas standard MCMC methods are inherently sequential.

6. Conclusions

Ill-posedness is a characteristic feature of inverse problems, and therefore, special attention needs to be paid to model uncertainties and model discrepancies that manifest themselves as highly correlated noise, deviating the measured data from the value predicted by the forward model. The modeling error is particularly detrimental when the quality of the data is good,




Figure 9. Convergence of the error $\|u_\ell - u^\dagger\|$ between the conditional mean and the truth as the number of iterations increases. From left to right observational noise standard deviation is 10^{-2} , 10^{-3} , 10^{-4} respectively.

Table 3. Comparisons of costs of producing samples from different approximations of the posterior distribution, as well as the full posterior. We assume in all cases that π_{prior} is Gaussian.

	f linear, π_{noise} Gaussian	General case
Conventional	SQ	$SQ\rho_1G_-$
Enhanced	$N(G_+ + G_-) + SQ$	$N(G_+ + G_-) + SQ\rho_2G_-$
Iterative	$N(G_+ + G_-)L + SQ$	$N(G_+ + G_-)L + SQ\rho_3G_-$
Full	—	$SQ\rho_4G_+$

and the exogenous noise does not mask the modeling errors that may become the predominant component of the noise. Quantification of the uncertainty due to the modeling errors is therefore an important part of successfully solving the inverse problem. Modeling error depends on the unknown that is the target of the inverse problem, and therefore, the Bayesian framework provides a natural basis for attacking the problem: the unknown of interest, modeled as a random variable, can be used in a natural way to define the modeling error as a random variable, thus allowing a statistical interpretation of the modeling error. In this article we introduce, and study the properties of, an iterative method of refining the statistical description of the modeling error as our information about the unknown increases.

From the implementational point of view, two cases in which the refinement of the modeling error distribution can be computed are identified. When the model is linear and the distributions are Gaussian, a fairly straightforward updating strategy of the posterior estimate is found, and convergence of this iteration can be shown. For non-linear inverse problems, a linearized approximate model leads to a tractable iterative algorithm based on particle approximations of the posterior, and as demonstrated in the numerical experiments, the computed point estimates can be very good, significantly improving on estimates which ignore model error. However, as pointed out in the article, the limiting approximate probability density obtained by the iterative algorithm is not identical to the Bayesian posterior density, although it may be close to it. Regarding both the point estimate and the posterior it is important to recognize that while the approximation error approach does require a number of evaluations of the expensive forward model, unlike traditional MCMC algorithms no rejections occur. Thus the methodology has potential to compute point estimates more economically than conventional non-Bayesian approaches such as Tikhonov regularization; and it also holds the potential to produce reasonable posterior distributions at considerably lower cost than MCMC using the fully accurate Bayesian posterior. One of the future directions of research is to see how the approximation process proposed in this article can be effectively used to produce an estimate of the true posterior density.

Acknowledgments

The work of D Calvetti is partially supported by NSF grant DMS-1522334. E Somersalo's work is partly supported by the NSF grants DMS-1312424 and DMS-1714617. The research of AM Stuart was partially supported by the EPSRC programme grant EQUIP, by AFOSR Grant FA9550-17-1-0185 and ONR Grant N00014-17-1-2079. M Dunlop was partially supported by the EPSRC MASDOC Graduate Training Program. Both M Dunlop and AM Stuart are supported by DARPA funded program Enabling Quantification of Uncertainty in Physical Systems (EQUIPS), contract W911NF-15-2-0121.

Appendix. Abstract formulation of algorithm

Let (Z, \mathcal{Z}) be a measurable space, and given $A \in \mathcal{Z}$ define the indicator function $\mathbb{I}_A : Z \rightarrow \mathbb{R}$ by

$$\mathbb{I}_A(z) = \begin{cases} 1 & z \in A \\ 0 & z \notin A. \end{cases}$$

Given two measures μ, ν on (Z, \mathcal{Z}) , let $\mu * \nu$ denote their convolution, i.e. the measure on (Z, \mathcal{Z}) given by

$$(\mu * \nu)(A) = \int_{Z \times Z} \mathbb{I}_A(u+v) \mu(du) \nu(dv)$$

for any $A \in \mathcal{Z}$. Note that if we have $u \sim \mu$ and $v \sim \nu$ independently, then $u+v \sim \mu * \nu$.

Algorithm (General). Let μ_0 denote the prior distribution on u and \mathbb{Q}_0 the distribution of the noise ε . Given $v \in Y$, define $T_v : Y \rightarrow Y$ to be the translation operator $T_v(y) = y + v$. Set $\ell = 0$.

1. Given μ_ℓ , assume $m \sim M^\# \mu_\ell$ independently of ε , so $m + \varepsilon \sim \mathbb{Q}_0^{(\ell+1)} := M^\# \mu_\ell * \mathbb{Q}_0$. The likelihood is given by

$$b \mid u \sim \mathbb{Q}_u^{(\ell+1)} := T_{f(u)}^\# \mathbb{Q}_0^{(\ell+1)}.$$

Assume that $\mathbb{Q}_u^{(\ell+1)} \ll \mathbb{Q}_0^{(\ell+1)}$, so that we have Radon–Nikodym density

$$\frac{d\mathbb{Q}_u^{(\ell+1)}}{d\mathbb{Q}_0^{(\ell+1)}}(b) = \exp(-\Phi^{(\ell+1)}(u; b)).$$

Bayes' theorem gives the posterior distribution

$$\mu_{\ell+1}(du) \propto \exp(-\Phi^{(\ell+1)}(u; b)) \mu_0(du). \quad (\text{A.1})$$

2. Set $\ell \mapsto \ell + 1$ and go to 1. □

The above iteration could be written more directly as

$$\mu_{\ell+1}(du) \propto \frac{d[T_{f(u)}^\#(M^\# \mu_\ell * \mathbb{Q}_0)]}{d[M^\# \mu_\ell * \mathbb{Q}_0]}(b) \mu_0(du)$$

though the expression (A.1) makes links with previous work on non-parametric Bayesian inverse problems clearer.

ORCID iDs

Erkki Somersalo  <https://orcid.org/0000-0001-5099-3512>

References

- [1] Agapiou S, Papaspiliopoulos O, Sanz-Alonso D and Stuart A M 2017 Importance sampling: computational complexity and intrinsic dimension *Stat. Sci.* **32** 405–31
- [2] 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** 175–95
- [3] Banasiak R, Ye Z and Soleimani M 2012 Improving three-dimensional electrical capacitance tomography imaging using approximation error model theory *J. Electromagn. Waves Appl.* **26** 411–21
- [4] Bayarri M J, Berger J O, Paulo R, Sacks J, Cafeo J A, Cavendish J, Lin C-H and Tu J 2007 A framework for validation of computer models *Technometrics* **49** 138–54
- [5] Beskos A, Jasra A, Muzaffer E and Stuart A M 2015 Sequential Monte Carlo methods for Bayesian elliptic inverse problems *Stat. Comput.* **25** 727–37
- [6] Briol F X, Oates C J, Girolami M and Osborne M A 2016 Probabilistic integration: a role for statisticians in numerical analysis? (arXiv:1512.00933)
- [7] Brynjarsdottir J and O’Hagan A 2014 Learning about physical parameters: the importance of model discrepancy *Inverse Problems* **30** 114007
- [8] Calvetti D, Ernst O and Somersalo E 2014 Dynamic updating of numerical model discrepancy using sequential sampling *Inverse Problems* **30** 114019
- [9] Calvetti D and Somersalo E 2007 *Introduction to Bayesian Scientific Computing—Ten Lectures on Subjective Computing* (New York: Springer)
- [10] Calvetti D and Somersalo E 2008 Hypermodels in the Bayesian imaging framework *Inverse Problems* **24** 034013
- [11] Calvetti D, Hakula H, Pursiainen S and Somersalo E 2009 Conditionally Gaussian hypermodels for cerebral source localization *SIAM J. Imaging Sci.* **2** 879–909
- [12] Cockayne J, Oates C, Sullivan T and Girolami M 2016 Probabilistic numerical methods for partial differential equations and Bayesian inverse problems (arXiv:1605.07811)
- [13] Colton D and Kress R 1992 *Inverse Acoustic and Electromagnetic Scattering Theory* (Berlin: Springer) (new edition: 1998)
- [14] Cotter S L, Roberts G O, Stuart A M and White D 2013 MCMC methods for functions: modifying old algorithms to make them faster *Stat. Sci.* **28** 424–46
- [15] Dashti M and Stuart A M 2016 The Bayesian approach to inverse problems *Handbook of Uncertainty Quantification* (New York: Springer)
- [16] Diaconis P 1988 Bayesian numerical analysis *Stat. Decis. Theory Relat. Top.* IV **1** 163–75
- [17] Heino J, Somersalo E and Kaipio J P 2005 Statistical compensation of geometric mismodeling in optical tomography *Opt. Express* **13** 296–308
- [18] Higdon D, Kennedy M, Cavendish J C, Cafeo J A and Ryne R D 2004 Combining field data and computer simulations for calibration and prediction *SIAM J. Sci. Comput.* **26** 448–66
- [19] Huttunen J M J, Kaipio J P and Somersalo E 2007 Approximation errors in nonstationary inverse problems *Inverse Problems Imaging* **1** 77–93
- [20] Huttunen J M J and Kaipio J P 2007 Approximation error analysis in nonlinear state estimation with an application to state-space identification *Inverse Problems* **23** 2141
- [21] Huttunen J M J, Lehtikoinen A, Hämäläinen J and Kaipio J P 2010 Importance sampling approach for the nonstationary approximation error method *Inverse Problems* **26** 125003
- [22] Kaipio J and Somersalo E 2004 *Statistical and Computational Inverse Problems* (New York: Springer)
- [23] Kaipio J and Somersalo E 2007 Statistical inverse problems: discretization, model reduction and inverse crimes *J. Comput. Appl. Math.* **198** 493–504
- [24] Kennedy M C and O’Hagan A 2001 Bayesian calibration of computer models *J. R. Stat. Soc. B* **63** 425–64

- [25] Lehtinen M S, Päivärinta L and Somersalo E 1989 Linear inverse problems for generalised random variables *Inverse Problems* **5** 599–612
- [26] MacEachern S N, Clyde M and Liu J S 1999 Sequential importance sampling for nonparametric Bayes models: the next generation *Can. J. Stat.* **27** 251–67
- [27] Mandelbaum A 1984 Linear estimators and measurable linear transformations on a Hilbert space *Z. Wahrscheinlichkeitstheor. Verwandte Geb.* **65** 385–97
- [28] Persson P and Strang G 2004 A simple mesh generator in MATLAB *SIAM Rev.* **46** 329–45
- [29] Roininen L, Lehtinen M, Lasanen S, Orispää M and Markkanen M 2011 Correlation priors *Inverse Problems Imaging* **5** 167–84
- [30] Roininen L, Huttunen J M and Lasanen S 2014 Whittle–Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography *Inverse Problems Imaging* **8** 561–86
- [31] Sacks J, Welch W J, Mitchell T J and Wynn H P 1989 Design and analysis of computer experiments *Stat. Sci.* **4** 409–23
- [32] Somersalo E, Cheney M and Isaacson D 1992 Existence and uniqueness for electrode models for electric current computed tomography *SIAM J. Appl. Math.* **52** 1023–40
- [33] West M 1993 Approximating posterior distributions by mixture *J. R. Stat. Soc. B* **55** 409–22