

Optimal design of large-scale nonlinear Bayesian inverse problems under model uncertainty

Alen Alexanderian^{1,*} , Ruanui Nicholson²
and Noemi Petra³ 

¹ Department of Mathematics, North Carolina State University, Raleigh, NC, United States of America

² Department of Engineering Science, University of Auckland, Auckland, New Zealand

³ Department of Applied Mathematics, University of California, Merced, CA, United States of America

E-mail: alexanderian@ncsu.edu, ruanui.nicholson@auckland.ac.nz and npetra@ucmerced.edu

Received 27 October 2023; revised 7 April 2024

Accepted for publication 8 July 2024

Published 18 July 2024



Abstract

We consider optimal experimental design (OED) for Bayesian nonlinear inverse problems governed by partial differential equations (PDEs) under model uncertainty. Specifically, we consider inverse problems in which, in addition to the inversion parameters, the governing PDEs include secondary uncertain parameters. We focus on problems with infinite-dimensional inversion and secondary parameters and present a scalable computational framework for optimal design of such problems. The proposed approach enables Bayesian inversion and OED under uncertainty within a unified framework. We build on the Bayesian approximation error (BAE) approach, to incorporate modeling uncertainties in the Bayesian inverse problem, and methods for A-optimal design of infinite-dimensional Bayesian nonlinear inverse problems. Specifically, a Gaussian approximation to the posterior at the maximum *a posteriori* probability point is used to define an uncertainty aware OED objective that is tractable to evaluate and optimize. In particular, the OED objective can be computed at a cost, in the number of PDE solves, that does not grow with the dimension of the discretized inversion and secondary parameters. The OED problem is formulated as a binary bilevel PDE constrained optimization problem and a greedy algorithm, which provides a pragmatic approach, is used to find optimal designs. We demonstrate the effectiveness of the proposed

* Author to whom any correspondence should be addressed.

approach for a model inverse problem governed by an elliptic PDE on a three-dimensional domain. Our computational results also highlight the pitfalls of ignoring modeling uncertainties in the OED and/or inference stages.

Keywords: optimal experimental design, sensor placement, Bayesian inverse problems, model uncertainty, Bayesian approximation error

1. Introduction

Models governed by partial differential equations (PDEs) are common in science and engineering applications. Such PDE models often contain parameters that need to be estimated using observed data and the model. This requires solving an inverse problem. The quality of the estimated parameters is influenced significantly by the quantity and quality of the measurement data. Therefore, optimizing the data acquisition process is crucial. This requires solving an optimal experimental design (OED) problem [7, 13, 43]. In the present work, we focus on inverse problems in which measurement data are collected at a set of sensors. In this case, OED amounts to finding an optimal sensor placement. In this context, OED is especially important when only a few sensors can be deployed.

While some parameters in the governing PDEs can be estimated by solving an inverse problem, often there are additional uncertain model parameters that are not being estimated. These parameters might be too costly or impossible to estimate. We call the uncertain parameters that are being estimated in an inverse problem the *inversion parameters* and refer to the additional uncertain parameters as the *secondary parameters*. Such secondary parameters have also been referred to as auxiliary parameters, latent parameters, or nuisance parameters in the literature. When solving inverse problems with secondary parameters with significant uncertainty levels, both the parameter estimation and data acquisition processes need to be aware of such uncertainties. In this article, we present a computational framework for optimal design of nonlinear Bayesian inverse problems governed by PDEs under model uncertainty.

Uncertainties in mathematical models can be divided into two classes: reducible and irreducible [39]. The reducible uncertainties are epistemic uncertainties that can be reduced via statistical parameter estimation. On the other hand, irreducible uncertainties are either aleatoric uncertainties inherent to the model that are impossible to reduce or are epistemic uncertainties that are too costly or impractical to reduce. Also, in some applications we might have access to a probabilistic description of secondary model parameters from previous studies and further reduction of the uncertainty in such parameters may not be worth the additional computational cost. We consider such uncertainties as irreducible as well. In the present work, we focus on the case of irreducible uncertainties.

We consider models of the form

$$\mathbf{y} = \mathcal{G}(m, \xi) + \boldsymbol{\eta}, \quad (1)$$

where \mathbf{y} is a vector of measured data, \mathcal{G} is a PDE-based model, m and ξ are uncertain parameters, and $\boldsymbol{\eta}$ is a random vector that models measurement noise. Herein, m is the inversion parameter which we seek to infer, and ξ is a secondary uncertain model parameter. We assume the uncertainty in ξ to be irreducible. The parameters m and ξ are assumed to be independent random variables that take values in infinite-dimensional real separable Hilbert spaces \mathcal{M} and \mathcal{X} , respectively. Moreover, \mathcal{G} is assumed to be nonlinear in both m and ξ . The methods

presented in this article enable computing optimal experimental designs in such a way that the uncertainty in the secondary parameters is accounted for.

1.1. Related work

In the recent years, there have been numerous research efforts directed at OED in inverse problems governed by PDEs. See the review [1], for a survey of the recent literature in this area. There has also been an increased interest in parameter inversion and design of experiments in systems governed by uncertain forward models; see e.g. [6, 14, 24, 27, 33, 38] for a small sample of the literature addressing inverse problems under uncertainty. Methods for OED in such inverse problems have been studied in [5, 11, 18, 28]. The works [5, 28] concern optimal design of infinite-dimensional Bayesian linear inverse problems governed by PDEs. Specifically, [28] considers design of linear inverse problems governed by PDEs with irreducible sources of model uncertainty. On the other hand, [5] targets OED for linear inverse problems with reducible sources of uncertainty. The efforts [11, 18] focus on inverse problems with finite- and low-dimensional inversion and secondary parameters. These articles devise sampling approaches for estimating the expected information gain in such problems. The setting considered in [18] is that of inverse problems with reducible uncertainties. The approach in [11] employs a small noise approximation that applies to problems with nuisance parameters that have small uncertainty levels.

1.2. Our approach

We focus on Bayesian nonlinear inverse problems governed by PDEs with infinite-dimensional inversion and secondary parameters. Traditionally, when solving such inverse problems all the secondary model parameters are fixed at their nominal values and the focus is on the estimation of the inversion parameters. Considering (1), this amounts to using the *approximate model* $\mathcal{F}(m) := \mathcal{G}(m, \bar{\xi})$, where $\bar{\xi}$ is some nominal value. In section 2, we consider simpler instances of (1) and illustrate the role of model uncertainty in Bayesian inverse problems and the importance of accounting for such uncertainties in parameter inversion and OED. The discussion in section 2 motivates our approach for incorporating secondary uncertainties in nonlinear Bayesian inverse problems and the corresponding OED problems. This is done using the Bayesian approximation error (BAE) approach [24, 26]; see section 3.

Subsequently, we build on methods for optimal design of infinite-dimensional nonlinear inverse problems [1, 4, 44] to derive an uncertainty aware OED objective. Specifically, we follow an A-optimal design strategy where the goal is to obtain designs that minimize average posterior variance. To cope with the non-Gaussianity of the posterior, we rely on a Gaussian approximation to the posterior. This enables deriving approximate measures of posterior uncertainty that are tractable to optimize for infinite-dimensional inverse problems; see section 4. We then present two approaches for formulating and computing the OED objective; see section 5. The first approach uses Monte Carlo trace estimators and the second one formulates the OED problem as an eigenvalue optimization problem. In each case, the OED problem is formulated as a bilevel binary PDE-constrained optimization problem. In both approaches, the cost of computing the OED objective, in terms of the number of PDE solves, is independent of the dimensions of discretized inversion and secondary parameters. This makes these approaches suitable for large-scale applications. In the present work, we rely on a greedy approach to solve the resulting optimization problems. As discussed in section 5, a greedy algorithm is especially suited for the formulation of the OED problem in section 5.2.2, as a binary PDE-constrained eigenvalue optimization problem.

We elaborate the proposed approach in the context of a model inverse problem governed by an elliptic PDE in a three-dimensional domain; see section 6. This inverse problem, which is motivated by heat transfer applications, concerns estimation of a coefficient field on the bottom boundary of the domain, using sensor measurements of the temperature on the top boundary. The secondary uncertain parameter in this inverse problem is the log-conductivity field, which is modeled as a random field on the three-dimensional physical domain. Our computational results in section 7 demonstrate the effectiveness of the proposed strategy in computing optimal sensor placements under model uncertainty. We also systematically study the drawbacks of ignoring the uncertainty in the Bayesian inversion and experimental design stages. These studies illustrate the fact that ignoring uncertainty in OED or inference stages can lead to inferior designs and highly inaccurate results in parameter estimation.

1.3. Contributions

The contributions of this article are as follows: (1) We present an uncertainty aware formulation of the OED problem, uncertainty aware OED objectives along with scalable methods for computing them, and an extensible optimization framework for computing optimal designs. These make OED for nonlinear Bayesian inverse problems governed by PDEs with infinite-dimensional inversion and secondary parameters feasible. Additionally, the proposed approach enables Bayesian inversion and OED under uncertainty within a unified framework. (2) We elaborate the proposed approach for an inverse problem governed by an elliptic PDE, on a three-dimensional domain, with infinite-dimensional inversion and secondary parameters. This is used to elucidate the implementation of our proposed approach for design of inverse problems governed by PDEs under uncertainty. (3) We present comprehensive computational studies that illustrate the effectiveness of the proposed approach and also the importance of accounting for modeling uncertainties in both parameter inversion and experimental design stages. (4) By considering simpler instances of (1), in section 2, we present a systematic study of the role of model uncertainty in Bayesian inverse problems and the importance of accounting for such uncertainties in parameter inversion and OED. That study also reveals a connection between the BAE-based approach taken in the present work and the method for OED in linear inverse problems under uncertainty in [5].

The developments in this article point naturally to a number of extensions of the presented methods. We discuss such issues in section 8, where we present our concluding remarks, and discuss potential limitations of the presented approach and opportunities for future extensions.

2. Motivation and overview

In this section, we motivate our approach for OED under uncertainty and set the stage for the developments in the rest of the article. After a brief coverage of requisite notation and preliminaries in section 2.1, we begin our discussion in section 2.2 by considering a simple form of (1) where \mathcal{G} is linear in m and ξ . This facilitates an intuitive study of the role of model uncertainty in Bayesian inverse problems and OED. We then consider nonlinear models of varying complexity in section 2.3 to motivate our approach for design of inverse problems governed by nonlinear models of the form (1).

2.1. Preliminaries

In this article, we consider inversion and secondary parameters that take values in infinite-dimensional Hilbert spaces. For a Hilbert space \mathcal{H} , we denote the corresponding inner product

by $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and the associated induced norm by $\|\cdot\|_{\mathcal{H}}$; i.e. $\|\cdot\|_{\mathcal{H}} := \langle \cdot, \cdot \rangle_{\mathcal{H}}^{1/2}$. For Hilbert spaces \mathcal{H}_1 and \mathcal{H}_2 , $\mathcal{L}(\mathcal{H}_1, \mathcal{H}_2)$ denotes the space of bounded linear transformations from \mathcal{H}_1 to \mathcal{H}_2 . The space of bounded linear operators on a Hilbert space \mathcal{H} is denoted by $\mathcal{L}(\mathcal{H})$, and the subspace of bounded selfadjoint operators is denoted by $\mathcal{L}^{\text{sym}}(\mathcal{H})$. We let $\mathcal{L}^{\text{sym}+}(\mathcal{H})$ denote the set of bounded positive selfadjoint operators. The subspace of trace-class operators in $\mathcal{L}(\mathcal{H})$ is denoted by $\mathcal{L}_1(\mathcal{H})$, and the subspace of selfadjoint trace-class operators is denoted by $\mathcal{L}_1^{\text{sym}}(\mathcal{H})$. Also, the sets of positive and strictly positive selfadjoint trace-class operators are denoted by $\mathcal{L}_1^{\text{sym}+}(\mathcal{H})$ and $\mathcal{L}_1^{\text{sym}++}(\mathcal{H})$, respectively.

Throughout the article, $\mathcal{N}(a, \mathcal{C})$ denotes a Gaussian measure with mean a and covariance operator \mathcal{C} . For a Gaussian measure on an infinite-dimensional Hilbert space \mathcal{H} , the covariance operator \mathcal{C} is required to be in $\mathcal{L}_1^{\text{sym}+}(\mathcal{H})$. Herein, we consider non-degenerate Gaussian measures; i.e. we assume that $\mathcal{C} \in \mathcal{L}_1^{\text{sym}++}(\mathcal{H})$. For further details on Gaussian measures, we refer to [15, 40]. Also, when considering measures on Hilbert spaces, we equip these spaces with their associated Borel sigma algebra. Throughout the article, for notational convenience, we suppress this choice of the sigma-algebra in our notations.

The adjoint of a linear transformation $\mathcal{A} \in \mathcal{L}(\mathcal{H}_1, \mathcal{H}_2)$, where \mathcal{H}_1 and \mathcal{H}_2 are (real) Hilbert spaces, is denoted by \mathcal{A}^* . Recall that $\mathcal{A}^* \in \mathcal{L}(\mathcal{H}_2, \mathcal{H}_1)$ and

$$\langle \mathcal{A}v_1, v_2 \rangle_{\mathcal{H}_2} = \langle v_1, \mathcal{A}^*v_2 \rangle_{\mathcal{H}_1}, \quad \text{for all } v_1 \in \mathcal{H}_1, v_2 \in \mathcal{H}_2.$$

We also recall the following basic result regarding affine transformations of Gaussian random variables. Let X be an \mathcal{H}_1 -valued Gaussian random variable with law $\mathcal{N}(a, \mathcal{C})$, $\mathcal{A} \in \mathcal{L}(\mathcal{H}_1, \mathcal{H}_2)$, and $b \in \mathcal{H}_2$. Then, the random variable $\mathcal{A}X + b$ is an \mathcal{H}_2 -valued Gaussian random variable with law $\mathcal{N}(\mathcal{A}a + b, \mathcal{A}\mathcal{C}\mathcal{A}^*)$; see [15] for details.

2.2. Linear models

Consider the model

$$\mathbf{y} = \mathcal{S}m + \mathcal{T}\xi + \boldsymbol{\eta}, \tag{2}$$

where $\mathbf{y} \in \mathbb{R}^d$ denotes measurement data, $m \in \mathcal{M}$ is the inversion parameter, $\xi \in \mathcal{X}$ is the secondary uncertain parameter, and $\boldsymbol{\eta}$ is the measurement noise vector. (The spaces \mathcal{M} and \mathcal{X} are as described in the introduction.) This type of model, which was considered in [5], may correspond to inverse problems governed by linear PDEs with uncertainties in source terms or boundary conditions. We assume $\mathcal{S} \in \mathcal{L}(\mathcal{M}, \mathbb{R}^d)$ and $\mathcal{T} \in \mathcal{L}(\mathcal{X}, \mathbb{R}^d)$. Moreover, we assume that $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma}_{\text{noise}})$ and that $\boldsymbol{\eta}$ is independent of m and ξ . We consider a Gaussian prior law $\mu_{\text{pr}} = \mathcal{N}(m_{\text{pr}}, \mathcal{C}_{\text{pr}})$ for m , and for the purpose of this illustrative example, let the *secondary model uncertainty* ξ be distributed according to a Gaussian $\mu_{\xi} = \mathcal{N}(\bar{\xi}, \mathcal{C}_{\xi})$. Also, we assume $\boldsymbol{\Gamma}_{\text{noise}} = \sigma^2 \mathbf{I}$, with σ^2 denoting the noise level.

2.2.1. Incorporating model uncertainty in the inverse problem. For a fixed realization of ξ , estimating m from (2) is a standard problem. This also follows the common practice of fixing additional model parameters at some nominal values before solving the inverse problem. However, it is possible to account for the model uncertainty in this process. To see this, suppose we fix ξ at $\bar{\xi}$ and consider the approximate model $\mathcal{F}(m) = \mathcal{S}m + \mathcal{T}\bar{\xi}$. Note that

$$\underbrace{\mathcal{S}m + \mathcal{T}\xi}_{\text{accurate model}} = \underbrace{\mathcal{S}m + \mathcal{T}\bar{\xi}}_{\mathcal{F}(m)} + \underbrace{\mathcal{T}(\xi - \bar{\xi})}_{\text{error } \boldsymbol{\varepsilon}}.$$

In this case, the approximation error $\varepsilon = \mathcal{T}(\xi - \bar{\xi})$ has a Gaussian law, $\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathcal{T}\mathcal{C}_\xi\mathcal{T}^*)$. Hence, we may rewrite (2) in terms of the approximate model \mathcal{F} as follows:

$$\mathbf{y} = \mathcal{F}(m) + \boldsymbol{\nu}, \quad (3)$$

where $\boldsymbol{\nu} = \varepsilon + \boldsymbol{\eta}$ denotes the *total error*. In the present setting, $\boldsymbol{\nu} \sim \mathcal{N}(\mathbf{0}, \mathcal{T}\mathcal{C}_\xi\mathcal{T}^* + \boldsymbol{\Gamma}_{\text{noise}})$. Note that we have incorporated the uncertainty due to ξ in the error covariance matrix.⁴ The present procedure for incorporating the secondary model uncertainty into the inverse problem is a special case of the Bayesian approximation error (BAE) approach (see section 3).

Since we have Gaussian prior and noise models and \mathcal{F} in (3) is affine, the posterior is also Gaussian with analytic formulas for its mean and covariance operator; see e.g. [40]. In particular, the posterior covariance operator is given by

$$\mathcal{C}_{\text{post}} = (\mathcal{S}^* \boldsymbol{\Gamma}_\nu^{-1} \mathcal{S} + \mathcal{C}_{\text{pr}}^{-1})^{-1} = \left(\mathcal{S}^* (\mathcal{T}\mathcal{C}_\xi\mathcal{T}^* + \boldsymbol{\Gamma}_{\text{noise}})^{-1} \mathcal{S} + \mathcal{C}_{\text{pr}}^{-1} \right)^{-1}. \quad (4)$$

Considering, for example, the A-optimality criterion $\text{tr}(\mathcal{C}_{\text{post}})$, (4) illustrates the manner in which the uncertainty due to use of an approximate model impacts the posterior uncertainty.

2.2.2. Interplay between measurement error and approximation error. Note that $\mathcal{T}\mathcal{C}_\xi\mathcal{T}^* \in \mathcal{L}^{\text{sym}+}(\mathbb{R}^d)$. This operator admits a spectral decomposition $\mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^\top$, where \mathbf{V} is an orthogonal matrix of eigenvectors and $\boldsymbol{\Lambda}$ is a diagonal matrix with the eigenvalues on its diagonal. Thus, the total error covariance matrix can be written as $\boldsymbol{\Gamma}_\nu = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^\top + \sigma^2\mathbf{I} = \sum_j (\lambda_j + \sigma^2)\mathbf{v}_j\mathbf{v}_j^\top$. Therefore, the modes for which $\lambda_j \ll \sigma^2$ may be ignored. To observe the impact of model uncertainty on individual observations, we note that

$$\text{Var}\{\nu_i\} = \mathbf{e}_i^\top \boldsymbol{\Gamma}_\nu \mathbf{e}_i = \sum_j (\lambda_j + \sigma^2) (\mathbf{e}_i^\top \mathbf{v}_j)^2, \quad i \in \{1, \dots, d\},$$

where \mathbf{e}_i 's are the standard basis vectors in \mathbb{R}^d . Thus, we see that only large eigenvalues contribute significantly to the total error in the i th measurement.

We can also consider the interplay between the spectral representation of the error covariance and the posterior covariance operator. Specifically, $\mathcal{C}_{\text{post}} = (\mathcal{S}^* \boldsymbol{\Gamma}_\nu^{-1} \mathcal{S} + \mathcal{C}_{\text{pr}}^{-1})^{-1} = \mathcal{C}_{\text{pr}}^{1/2} (\tilde{\mathcal{S}}^* \boldsymbol{\Gamma}_\nu^{-1} \tilde{\mathcal{S}} + I)^{-1} \mathcal{C}_{\text{pr}}^{1/2}$ with $\tilde{\mathcal{S}} = \mathcal{S}\mathcal{C}_{\text{pr}}^{1/2}$. Thus, letting $\mathbf{E}_j = \mathbf{v}_j\mathbf{v}_j^\top$, $j = 1, \dots, d$, we can write

$$\mathcal{C}_{\text{post}} = \mathcal{C}_{\text{pr}}^{1/2} \left[\sum_{j=1}^d (\lambda_j + \sigma^2)^{-1} \tilde{\mathcal{S}}^* \mathbf{E}_j \tilde{\mathcal{S}} + I \right]^{-1} \mathcal{C}_{\text{pr}}^{1/2}.$$

Note that in the directions corresponding to very large eigenvalues the measurements will have a negligible impact on posterior uncertainty.

The above discussion indicates that model uncertainty cannot be ignored in the inverse problem, especially when the model uncertainty overwhelms measurement noise. The latter is also important in design of experiments. Specifically, some of the measurements might be completely useless due to large amount of model uncertainty associated to the corresponding measurements. This information needs to be accounted for in the OED problem to ensure only measurements that are helpful in reducing posterior uncertainty are selected.

⁴ If the approximate model \mathcal{F} was defined by fixing ξ at a point different from $\bar{\xi}$, then the error term $\boldsymbol{\nu}$ would have nonzero mean.

2.2.3. Connection to post-marginalization. In general, the BAE approach involves pre-marginalization over the secondary model uncertainties. This can be related to the idea of post-marginalization in the case of linear Gaussian inverse problems. Namely, if we consider ξ as a reducible uncertainty that is being estimated along with m and μ_ξ as the corresponding prior law, then (4) is the covariance operator of the marginal posterior law of m . To see this, we first note that

$$(\mathcal{T}\mathcal{C}_\xi\mathcal{T}^* + \mathbf{\Gamma}_{\text{noise}})^{-1} = \mathbf{\Gamma}_{\text{noise}}^{-1} - \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{T} \left(\mathcal{C}_\xi^{-1} + \mathcal{T}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{T} \right)^{-1} \mathcal{T}^* \mathbf{\Gamma}_{\text{noise}}^{-1}. \quad (5)$$

This relation can be derived by following a similar calculation as the one in [40, p 536].⁵ Then, we substitute (5) in (4) to obtain

$$\mathcal{C}_{\text{post}} = \left[\mathcal{C}_{\text{pr}}^{-1} + \mathcal{S}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{S} - \mathcal{S}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{T} \left(\mathcal{C}_\xi^{-1} + \mathcal{T}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{T} \right)^{-1} \mathcal{T}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{S} \right]^{-1}.$$

This $\mathcal{C}_{\text{post}}$ is the same as the marginal posterior covariance operator of m as noted in [5, equation (2.4)]. Thus, for a linear Gaussian inverse problem with reducible secondary uncertainty, the posterior covariance operator obtained using the BAE approach equals the marginal posterior covariance operator of m obtained following joint estimation of m and ξ .

2.3. Nonlinear models

The linear model (2) can be generalized in the following ways:

$$\text{additive model, linear in } m, \text{ nonlinear in } \xi : \mathcal{G}(m, \xi) = \mathcal{S}m + \mathcal{T}(\xi); \quad (6)$$

$$\text{additive model, nonlinear in } m, \text{ linear in } \xi : \mathcal{G}(m, \xi) = \mathcal{S}(m) + \mathcal{T}\xi; \quad (7)$$

$$\text{additive model, nonlinear in both } m \text{ and } \xi : \mathcal{G}(m, \xi) = \mathcal{S}(m) + \mathcal{T}(\xi); \quad (8)$$

$$\text{nonadditive model, nonlinear in both } m \text{ and } \xi. \quad (9)$$

While the cases (6)–(8) might be of independent interest, our focus in this article is on the general case (9). However, items (6)–(8) do serve to illustrate some of the key challenges.

In the case of (6), one can repeat the steps leading to (3), except ε will not be Gaussian anymore and therefore the distribution of ν will not be known analytically. In that case, one may obtain a Gaussian approximation to ε either by fitting a Gaussian to ε or by using a linear approximation of $\mathcal{T}(\xi)$. Then, one may obtain a Gaussian posterior, where one also relies on the linearity of \mathcal{S} . On the other hand, in the case of (7), the total error ν will be Gaussian as before, but due to nonlinearity of $\mathcal{S}(m)$ the posterior will not be Gaussian. The latter leads to one of the fundamental challenges in OED of nonlinear inverse problem—defining a suitable OED objective whose optimization is tractable. The more complicated cases of (8)–(9) inherit the challenges corresponding to the previous cases.

In the rest of this article, we build on the BAE approach to incorporate the uncertainty in ξ in inverse problems governed by nonlinear models of the type (9). The uncertainty in ξ will be assumed irreducible, and in general, ξ will not be assumed to follow a Gaussian law. All that we require is the ability to generate samples of ξ . Following the BAE approach, we approximate the approximation error ε with a Gaussian. This enables incorporating the model uncertainty in the data likelihood; see section 3. To cope with non-Gaussianity of the posterior, we rely on

⁵ Note that (5) can be viewed as a special form of the Sherman–Morrison–Woodbury formula involving Hilbert space operators.

a Gaussian approximation to the posterior, to derive an uncertainty aware OED objective that is tractable to evaluate and optimize for infinite-dimensional inverse problems; see section 4 for the definition of the OED objective and section 5 for computational methods.

3. Infinite-dimensional Bayesian inverse problems under uncertainty

We consider the inverse problem of inferring a parameter m from a model of the form (1), where $\mathcal{G} : \mathcal{M} \times \mathcal{X} \rightarrow \mathbb{R}^d$ is a parameter-to-observable map that in general is nonlinear in both arguments. We focus on problems where \mathcal{G} is defined as a composition of an observation operator and a PDE solution operator. The model \mathcal{G} is assumed to be Fréchet differentiable in m , at $\xi = \bar{\xi}$, where $\bar{\xi} \in \mathcal{X}$ is a nominal value. As before, we employ a Gaussian noise model, $\eta \sim \mathcal{N}(\mathbf{0}, \Gamma_{\text{noise}})$, a Gaussian prior law $\mu_{\text{pr}} = \mathcal{N}(m_{\text{pr}}, \mathcal{C}_{\text{pr}})$ for m , and assume η is independent of m and ξ . The prior induces the Cameron–Martin space $\mathcal{E} = \text{range}(\mathcal{C}_{\text{pr}}^{1/2})$, which is endowed with the inner product [15, 17]

$$\langle a, b \rangle_{\mathcal{E}} = \langle \mathcal{C}_{\text{pr}}^{-1/2} a, \mathcal{C}_{\text{pr}}^{-1/2} b \rangle_{\mathcal{M}}, \quad a, b \in \mathcal{E},$$

where $\langle \cdot, \cdot \rangle_{\mathcal{M}}$ is the inner product on the parameter space \mathcal{M} .

To account for model uncertainty (due to ξ) in the inverse problem, we rely on the BAE approach [24, 26], which we explain next. We fix the secondary parameter to $\bar{\xi}$ and consider the approximate (also known as inaccurate, reduced order, or surrogate) model

$$\mathcal{F}(m) = \mathcal{G}(m, \bar{\xi}). \quad (10)$$

As mentioned before, this is typically what is done in practice where the secondary model parameters are fixed at some (possibly well-justified) nominal values. In the BAE approach, we quantify and incorporate errors due to the use of this approximate model in the Bayesian inverse problem analogously to what was done in section 2.2. Namely, we consider

$$\mathbf{y} = \mathcal{F}(m) + \underbrace{\mathcal{G}(m, \xi) - \mathcal{F}(m)}_{\text{approximation error } \varepsilon} + \eta = \mathcal{F}(m) + \nu(m, \xi), \quad (11)$$

where $\nu(m, \xi) = \varepsilon(m, \xi) + \eta$ is the total error. In the BAE framework, the approximation error ε , is approximated as a conditionally Gaussian random variable. That is, the distribution of $\varepsilon|m$ is assumed to be Gaussian. In the present work, we employ the so-called *enhanced error model* [10, 24, 27], that ignores the correlation between ε and m and approximates the law of ε as a Gaussian $\varepsilon \sim \mathcal{N}(\varepsilon_0, \Gamma_{\varepsilon})$ with

$$\begin{aligned} \varepsilon_0 &= \int_{\mathcal{M}} \int_{\mathcal{X}} \varepsilon(m, \xi) \mu_{\text{pr}}(dm) \mu_{\xi}(d\xi), \\ \Gamma_{\varepsilon} &= \int_{\mathcal{M}} \int_{\mathcal{X}} (\varepsilon(m, \xi) - \varepsilon_0)(\varepsilon(m, \xi) - \varepsilon_0)^T \mu_{\text{pr}}(dm) \mu_{\xi}(d\xi). \end{aligned} \quad (12)$$

In general, the approximation errors can be (highly) correlated with the parameters [24, 34]. However, ignoring the correlation between ε and m (i.e. employing the enhanced error model) is typically viewed as a conservative (safe) approximation as it is analogous to approximating the conditional distribution of $\varepsilon|m$ with the marginal distribution of ε , which cannot reduce variance [25, section 3.4]. On the other hand, employing the enhanced error model can significantly reduce the costs associated with computing the mean and covariance operator of ε [24] which are in practice computed via sampling; see section 5.

With these approximations, and by our assumption on the measurement noise (which is independent of the parameters), the total error ν is modeled by a Gaussian $\mathcal{N}(\varepsilon_0, \Gamma_{\nu})$, where

$\Gamma_\nu = \Gamma_\varepsilon + \Gamma_{\text{noise}}$. Using this approximate noise model along with the approximate model \mathcal{F} we arrive at the following data likelihood:

$$\Pi_{\text{like}}^{\text{BAE}}(\mathbf{y}|m) \propto \exp \left\{ -\frac{1}{2} (\mathbf{y} - \mathcal{F}(m) - \boldsymbol{\varepsilon}_0)^\top \Gamma_\nu^{-1} (\mathbf{y} - \mathcal{F}(m) - \boldsymbol{\varepsilon}_0) \right\}. \quad (13)$$

With the prior measure in place, and using this BAE-based data likelihood, we can state the Bayes formula [40],

$$\frac{d\mu_{\text{post}}^y}{d\mu_{\text{pr}}} \propto \Pi_{\text{like}}^{\text{BAE}}(\mathbf{y}|m).$$

To make computations involved in design of large-scale inverse problems tractable, we rely on a local Gaussian approximation to the posterior. Namely, we use $\hat{\mu}_{\text{post}}^y = \mathcal{N}(m_{\text{MAP}}^y, \mathcal{C}_{\text{post}}^y)$, where m_{MAP}^y is the maximum *a posteriori* probability (MAP) point and $\mathcal{C}_{\text{post}}^y$ is an approximate posterior covariance operator, described below. The MAP point is given by

$$\begin{aligned} m_{\text{MAP}}^y = \arg \min_{m \in \mathcal{E}} \mathcal{J}(m) := & \frac{1}{2} (\mathbf{y} - \mathcal{F}(m) - \boldsymbol{\varepsilon}_0)^\top \Gamma_\nu^{-1} (\mathbf{y} - \mathcal{F}(m) - \boldsymbol{\varepsilon}_0) \\ & + \frac{1}{2} \langle m - m_{\text{pr}}, m - m_{\text{pr}} \rangle_{\mathcal{E}}. \end{aligned} \quad (14)$$

For the approximate posterior covariance operator $\mathcal{C}_{\text{post}}^y$, we use

$$\mathcal{C}_{\text{post}}^y = (\mathcal{F}_m(m_{\text{MAP}}^y)^* \Gamma_\nu^{-1} \mathcal{F}_m(m_{\text{MAP}}^y) + \mathcal{C}_{\text{pr}}^{-1})^{-1}, \quad (15)$$

where $\mathcal{F}_m(m_{\text{MAP}}^y)$ is the Fréchet derivative of \mathcal{F} evaluated at m_{MAP}^y .

Note that the true posterior μ_{post}^y is equivalent to the prior measure. It is also possible to show that the Gaussian approximation is equivalent to μ_{pr} , as well. This fact, which is made precise below, is important in justifying the use of this Gaussian approximation for defining an approximate measure of posterior uncertainty. Namely, to define a notion of uncertainty reduction, it is important that our surrogate for the posterior measure is absolutely continuous with respect to our reference measure, which is given by the prior. The equivalence of $\hat{\mu}_{\text{post}}^y$ to μ_{pr} may be inferred from the more general developments in [35]. However, we present an accessible argument that applies to the specific problem setup under study in the present work. We first present the following result.

Proposition 3.1. *Consider the linearized forward model*

$$\mathcal{L}(m) = \mathcal{F}(m_{\text{MAP}}) + \mathcal{F}_m(m_{\text{MAP}})(m - m_{\text{MAP}}), \quad m \in \mathcal{H}, \quad (16)$$

where we have suppressed the dependence of m_{MAP} to \mathbf{y} for notational convenience. Define the data model

$$\mathbf{y} = \mathcal{L}(m) + \boldsymbol{\nu}, \quad (17)$$

where $\boldsymbol{\nu} \sim \mathcal{N}(\boldsymbol{\varepsilon}_0, \Gamma_\nu)$. Consider the Bayesian inverse problem of estimating m using (17) and the prior $\mu_{\text{pr}} = \mathcal{N}(m_{\text{pr}}, \mathcal{C}_{\text{pr}})$. The corresponding posterior measure is given by $\hat{\mu}_{\text{post}}^y$.

Proof. Using the theory of Bayesian linear inverse problems in a Hilbert space [40], the solution of the linear inverse problem under study yields a Gaussian posterior $\mu_{\text{post}}^{\text{lin}} = \mathcal{N}(m_{\text{MAP}}^{\text{lin}}, \mathcal{C}_{\text{post}})$. The covariance operator $\mathcal{C}_{\text{post}}$ is as in (15) and $m_{\text{MAP}}^{\text{lin}}$ is found by minimizing

$$\mathcal{J}_L(m) := \frac{1}{2} (\mathbf{y} - \mathcal{L}(m) - \boldsymbol{\varepsilon}_0)^\top \Gamma_\nu^{-1} (\mathbf{y} - \mathcal{L}(m) - \boldsymbol{\varepsilon}_0) + \frac{1}{2} \langle m - m_{\text{pr}}, m - m_{\text{pr}} \rangle_{\mathcal{E}}, \quad (18)$$

over the Cameron–Martin space \mathcal{E} . To complete the proof, we show $m_{\text{MAP}}^{\text{lin}} = m_{\text{MAP}}$. Note that \mathcal{J}_L is a strictly convex quadratic functional with a unique global minimizer. We consider the Euler–Lagrange equation for the present optimization problem. Note that the Fréchet derivative of \mathcal{L} is given by $\mathcal{L}_m = \mathcal{F}_m(m_{\text{MAP}})$. It is straightforward to see

$$\frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \mathcal{J}_L(m + \varepsilon\tilde{m}) = \langle \tilde{m}, \mathcal{F}_m(m_{\text{MAP}})^* \Gamma_\nu^{-1} (y - \mathcal{F}(m) - \varepsilon_0) \rangle_{\mathcal{M}} + \langle \tilde{m}, m - m_{\text{pr}} \rangle_{\mathcal{E}},$$

for every $\tilde{m} \in \mathcal{E}$. Thus, recalling $\mathcal{L}(m_{\text{MAP}}) = \mathcal{F}(m_{\text{MAP}})$, we note that for every $\tilde{m} \in \mathcal{E}$,

$$\begin{aligned} \frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \mathcal{J}_L(m_{\text{MAP}} + \varepsilon\tilde{m}) &= \langle \tilde{m}, \mathcal{F}_m(m_{\text{MAP}})^* \Gamma_\nu^{-1} (y - \mathcal{F}(m_{\text{MAP}}) - \varepsilon_0) \rangle_{\mathcal{M}} + \langle \tilde{m}, m_{\text{MAP}} - m_{\text{pr}} \rangle_{\mathcal{E}} \\ &= \frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \mathcal{J}(m_{\text{MAP}} + \varepsilon\tilde{m}) = 0. \end{aligned}$$

The last equality follows from the fact that m_{MAP} is a minimizer of \mathcal{J} in (14). Hence, m_{MAP} is the unique global minimizer of \mathcal{J}_L . Therefore, $m_{\text{MAP}}^{\text{lin}} = m_{\text{MAP}}$. \square

Proposition 3.1 shows that $\hat{\mu}_{\text{post}}^y$ is the posterior measure corresponding to the linearized Bayesian inverse problem considered in the result. Therefore, by construction, $\hat{\mu}_{\text{post}}^y$ is equivalent to μ_{pr} .

4. A-optimal experimental design under uncertainty

We consider inverse problems in which measurement data are collected at a set of sensors. In this case, the OED problem seeks to find an optimal placement of sensors. Specifically, we formulate the OED problem as that of selecting an optimal subset from a set of candidate sensor locations, which is a common approach; see e.g. [3, 20, 43]. To make matters concrete, we begin by fixing a set of points $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n_s}\}$ that indicate the candidate sensor locations. We then assign a binary weight w_i to each candidate location \mathbf{x}_i ; a weight of one indicates that a sensor will be placed at the corresponding candidate location. The binary vector $\mathbf{w} \in \{0, 1\}^{n_s}$ thus fully specifies an experimental design in the present setting. Note that specification of a set of candidate sensor locations will in general depend on the specific application at hand. For example, placing sensors in certain parts of the domain might be impractical or impossible. Also, in problems with Dirichlet boundary conditions, placing sensors on or very close to such boundaries will be a waste of resources.

4.1. Design of the Bayesian inverse problem

The design \mathbf{w} enters the formulation of the Bayesian inverse problem through the data likelihood. This requires additional care in the present work because the total error covariance matrix Γ_ν is non-diagonal. We follow the setup in [31] to incorporate \mathbf{w} in the Bayesian inverse problem formulation. For a binary design vector $\mathbf{w} \in \{0, 1\}^{n_s}$, we define the matrix \mathbf{P}_w as submatrix of $\text{diag}(\mathbf{w})$ with rows corresponding to the zero weights removed. Thus, given a generic measurement vector $\mathbf{d} \in \mathbb{R}^{n_s}$, $\mathbf{P}_w \mathbf{d}$ returns the measurements corresponding to active sensors. For $\mathbf{d} \in \mathbb{R}^{n_s}$, we use the notation $\mathbf{d}_w = \mathbf{P}_w \mathbf{d}$.

Next, we describe how a design vector \mathbf{w} enters the Bayesian inverse problem, within the BAE framework. For a given \mathbf{w} , we consider the model $\mathbf{y}_w = \mathbf{P}_w(\mathcal{F}(m) + \boldsymbol{\nu})$. Using this model leads to the following, \mathbf{w} -dependent, data likelihood

$$\Pi_{\text{like}}^{\text{BAE}}(\mathbf{y}|m) \propto \exp \left\{ -\frac{1}{2} (\mathbf{y} - \mathcal{F}(m) - \varepsilon_0)^T \boldsymbol{\Sigma}(\mathbf{w}) (\mathbf{y} - \mathcal{F}(m) - \varepsilon_0) \right\}, \quad (19)$$

with

$$\Sigma(\mathbf{w}) = \mathbf{P}_w^\top \Gamma_{\nu,w}^{-1} \mathbf{P}_w, \quad \text{where } \Gamma_{\nu,w} = \mathbf{P}_w \Gamma_\nu \mathbf{P}_w^\top. \quad (20)$$

Consequently, we obtain the following \mathbf{w} -dependent Gaussian approximation to the posterior, $\hat{\mu}_{\text{post}}^{\mathbf{y}, \mathbf{w}} = \mathcal{N}(m_{\text{MAP}}^{\mathbf{y}, \mathbf{w}}, \mathcal{C}_{\text{post}}^{\mathbf{y}, \mathbf{w}})$, where $m_{\text{MAP}}^{\mathbf{y}, \mathbf{w}}$ is obtained by minimizing

$$\mathcal{J}_w(m; \mathbf{y}) = \frac{1}{2} (\mathbf{y} - \mathcal{F}(m) - \boldsymbol{\varepsilon}_0)^\top \Sigma(\mathbf{w}) (\mathbf{y} - \mathcal{F}(m) - \boldsymbol{\varepsilon}_0) + \frac{1}{2} \langle m - m_{\text{pr}}, m - m_{\text{pr}} \rangle_{\mathcal{E}}, \quad (21)$$

and

$$\mathcal{C}_{\text{post}}^{\mathbf{y}, \mathbf{w}} = \left(\mathcal{F}_m(m_{\text{MAP}}^{\mathbf{y}, \mathbf{w}})^\top \Sigma(\mathbf{w}) \mathcal{F}_m(m_{\text{MAP}}^{\mathbf{y}, \mathbf{w}}) + \mathcal{C}_{\text{pr}}^{-1} \right)^{-1}. \quad (22)$$

Note that in practical computations, the cost function \mathcal{J}_w in (21) is implemented as

$$\mathcal{J}_w(m; \mathbf{y}) = \frac{1}{2} (\mathbf{y}_w - \mathbf{P}_w \mathcal{F}(m) - \mathbf{P}_w \boldsymbol{\varepsilon}_0)^\top \Gamma_{\nu,w}^{-1} (\mathbf{y}_w - \mathbf{P}_w \mathcal{F}(m) - \mathbf{P}_w \boldsymbol{\varepsilon}_0) + \frac{1}{2} \langle m - m_{\text{pr}}, m - m_{\text{pr}} \rangle_{\mathcal{E}},$$

with $\Gamma_{\nu,w}$ as in (20). This amounts to using the data from the active sensors and removing the rows and columns corresponding to inactive sensors from the error covariance matrix Γ_ν .

4.2. The design criterion

In the present work, we follow an A-optimal design strategy, where the goal is to find designs that minimize the average posterior variance. Generally, computing the average posterior variance for a Bayesian nonlinear inverse problem is computationally challenging. We follow the developments in [4] to define a Bayesian A-optimality criterion in the case of nonlinear inverse problems.

Given a data vector $\mathbf{y} \in \mathbb{R}^{n_s}$, an approximate measure of posterior uncertainty is provided by $\text{tr}(\mathcal{C}_{\text{post}}^{\mathbf{y}, \mathbf{w}})$. However, when solving the OED problem data is not available. Indeed, it is the goal of the OED problem to specify how data should be collected. To overcome this, we follow the general approach in Bayesian OED of nonlinear inverse problems, where we consider $\mathbb{E}_{\mathbf{y}}\{\text{tr}(\mathcal{C}_{\text{post}}^{\mathbf{y}, \mathbf{w}})\}$, with $\mathbb{E}_{\mathbf{y}}$ denoting expectation with respect to the set of all likely data. We compute this expectation by using the information available in the Bayesian inverse problem and the information regarding the distribution of the model uncertainty. Namely, following the approach in [4], we use the design criterion

$$\Phi(\mathbf{w}) = \int_{\mathcal{X}} \int_{\mathcal{M}} \int_{\mathbb{R}^{n_s}} \text{tr}(\mathcal{C}_{\text{post}}^{\mathbf{y}, \mathbf{w}}) \Pi_{\text{noise}}(\mathcal{G}(m, \xi) - \mathbf{y}) \, d\mathbf{y} \, \mu_{\text{pr}}(dm) \, \mu_\xi(d\xi), \quad (23)$$

where Π_{noise} is the probability density function (pdf) of the noise distribution $\mathcal{N}(\mathbf{0}, \Gamma_{\text{noise}})$. In practice, $\Phi(\mathbf{w})$ will be computed via sample averaging. Specifically, we use

$$\Phi_{n_d}(\mathbf{w}) = \frac{1}{n_d} \sum_{i=1}^{n_d} \text{tr}(\mathcal{C}_{\text{post}}^{\mathbf{y}_w^i}), \quad (24)$$

where the *training* data samples \mathbf{y}_w^i are given by $\mathbf{y}_w^i = \mathbf{P}_w(\mathcal{G}(m^i, \xi^i) + \boldsymbol{\eta}^i)$, with $\{(m^i, \xi^i, \boldsymbol{\eta}^i)\}_{i=1}^{n_d}$ a sample set from the product space $(\mathcal{M}, \mu_{\text{pr}}) \otimes (\mathcal{X}, \mu_\xi) \otimes (\mathbb{R}^{n_s}, \mathcal{N}(\mathbf{0}, \Gamma_{\text{noise}}))$. In large-scale applications, typically a small n_d can be afforded. However, in this context, typically a modest n_d enables computing good quality optimal designs. This is also demonstrated in the computational results in the present work.

Note that, for each $i \in \{1, \dots, n_d\}$,

$$\text{tr} \left(\mathcal{C}_{\text{post}}^{\mathbf{y}_w^i} \right) = \sum_{j=1}^{\infty} \langle \mathcal{C}_{\text{post}}^{\mathbf{y}_w^i} e_j, e_j \rangle_{\mathcal{M}},$$

where $\{e_j\}_{j=1}^{\infty}$ is a complete orthonormal set in \mathcal{M} . Inserting this in (24) and using the definition of $\mathcal{C}_{\text{post}}^{\mathbf{y}_w^i}$, we have

$$\Phi_{n_d}(\mathbf{w}) = \frac{1}{n_d} \sum_{i=1}^{n_d} \sum_{j=1}^{\infty} \langle c_{ij}, e_j \rangle_{\mathcal{M}}, \quad (25a)$$

where, for $i \in \{1, \dots, n_d\}$ and $j \in \mathbb{N}$,

$$m_{\text{MAP}}^{\mathbf{y}_w^i} = \arg \min_m \mathcal{J}_w(m; \mathbf{y}^i), \quad (25b)$$

$$\left(\mathcal{F}_m(m_{\text{MAP}}^{\mathbf{y}_w^i})^* \Sigma(\mathbf{w}) \mathcal{F}_m(m_{\text{MAP}}^{\mathbf{y}_w^i}) + \mathcal{C}_{\text{pr}}^{-1} \right) c_{ij} = e_j. \quad (25c)$$

Computing the OED objective as defined above is not practical. However, (25) provides insight into the key challenges in computing the OED objective. In the first place, (25b) is a challenging PDE-constrained optimization problem whose solution is the MAP point $m_{\text{MAP}}^{\mathbf{y}_w^i}$. Moreover, upon discretization, (25c) will be a high-dimensional linear system with the discretization of the operator $(\mathcal{F}_m(m_{\text{MAP}}^{\mathbf{y}_w^i})^* \Sigma(\mathbf{w}) \mathcal{F}_m(m_{\text{MAP}}^{\mathbf{y}_w^i}) + \mathcal{C}_{\text{pr}}^{-1})$ as its coefficient matrix. Such systems can be tackled with Krylov iterative methods that require the application of the coefficient matrix on vectors. In section 5, we present two approaches for efficiently computing the OED objective: one approach uses randomized trace estimation and the other utilizes low-rank approximation of $\mathcal{F}_m(m_{\text{MAP}}^{\mathbf{y}_w^i})^* \Sigma(\mathbf{w}) \mathcal{F}_m(m_{\text{MAP}}^{\mathbf{y}_w^i})$. These approaches rely on scalable optimization methods for (25b) as well as adjoint based gradient and Hessian apply computation.

4.3. The optimization problem for finding an A-optimal design

We state the OED problem of selecting the best K sensors, with $K \leq n_s$, as follows:

$$\begin{aligned} & \min_{\mathbf{w} \in \{0,1\}^{n_s}} \Phi_{n_d}(\mathbf{w}) \\ & \text{s.t. } \sum_{\ell=1}^{n_s} w_{\ell} = K. \end{aligned} \quad (26)$$

This is a challenging binary optimization problem. One possibility is to pursue a relaxation strategy [3, 4, 31] to enable gradient-based optimization with design weights $w_i \in [0, 1]$. A practical alternative is to follow a greedy approach to find an approximate solution to (26). Namely, we place sensors one at a time. In each step of a greedy algorithm, we select the sensor that provides the largest decrease in the value of the OED objective Φ_{n_d} . While the solutions obtained using a greedy algorithm are suboptimal in general, greedy approaches have shown good performance in many sensor placement problems; see, e.g. [5, 23, 29, 37]. In the present work, we follow a greedy approach for finding approximate solutions for (26). The effectiveness of this approach is demonstrated in our computational results in section 7.

5. Computational methods

We begin this section with a brief discussion on computing the BAE error statistics in section 5.1. We then detail our proposed methods for computing the OED objective in section 5.2. The greedy approach for computing optimal designs is outlined in section 5.3. Then, we discuss the computational cost of OED objective evaluation, using our proposed methods, as well as the greedy procedure in section 5.4.

5.1. Estimating approximation error statistics

As discussed in section 3, the mean and covariance of the approximation error in (12) will be approximated via Monte Carlo sampling. Specifically, we begin by drawing samples $\{(m^i, \xi^i)\}_{i=1}^{n_s}$ in $(\mathcal{M} \times \mathcal{X}, \mu_{\text{pr}} \otimes \mu_{\xi})$ and compute

$$\widehat{\varepsilon}_0 = \frac{1}{n_{\text{mc}}} \sum_{i=1}^{n_{\text{mc}}} \varepsilon^i, \quad \widehat{\Gamma}_{\varepsilon} = \frac{1}{n_{\text{mc}} - 1} \sum_{i=1}^{n_{\text{mc}}} (\varepsilon^i - \widehat{\varepsilon}_0) (\varepsilon^i - \widehat{\varepsilon}_0)^T,$$

where $\varepsilon^i = \mathcal{G}(m^i, \xi^i) - \mathcal{F}(m^i)$. (27)

Subsequently, we use $\varepsilon_0 \approx \widehat{\varepsilon}_0$ and $\Gamma_{\varepsilon} \approx \widehat{\Gamma}_{\varepsilon}$. The computational cost of this process is $2n_{\text{mc}}$ model evaluations. Note, however, that these model evaluations are done a priori, in parallel, and can be used for both the OED problem and solving the inverse problem. Typically, only a modest sample size is sufficient for approximating the mean and covariance of the model error [10, 33]. Specifically, as noted in section 2, only dominant modes of the error covariance matrix need to be resolved. Note also that (a subset of) the model evaluations in (27) may be reused in generation of training data needed for computation of the OED objective.

5.2. Computation of the OED objective

In this section, we present scalable computational methods for computing the OED objective. Specifically, we present two methods: (i) a method based on the use of randomized trace estimators (section 5.2.1) and (ii) a method based on low-rank spectral decompositions (section 5.2.2). As discussed further below, the latter is particularly suited to our overall approach in the present work. Therefore, we primarily focus on the method based on low-rank spectral decompositions, which we also fully elaborate in the context of a model inverse problem (see sections 6 and 7). However, the first method does have its own merits and is included to provide an alternative approach. The relative benefits and computational complexity of these methods are discussed in section 5.4.

Before we proceed further, we define some notations that simplify the discussions that follow. For $i \in \{1, \dots, n_d\}$, we define

$$\mathcal{H}^i(\mathbf{w}) = \mathcal{F}_m \left(\mathbf{m}_{\text{MAP}}^{\mathbf{y}_w^i} \right)^* \boldsymbol{\Sigma}(\mathbf{w}) \mathcal{F}_m \left(\mathbf{m}_{\text{MAP}}^{\mathbf{y}_w^i} \right) \quad \text{and} \quad \widetilde{\mathcal{H}}^i(\mathbf{w}) = \mathcal{C}_{\text{pr}}^{1/2} \mathcal{H}^i(\mathbf{w}) \mathcal{C}_{\text{pr}}^{1/2}. \quad (28)$$

Note that the operator \mathcal{H}^i is the so-called Gauss–Newton Hessian of the data misfit term in the definition of $\mathcal{J}_w(m; \mathbf{y}^i)$ in (21), evaluated at $\mathbf{m}_{\text{MAP}}^{\mathbf{y}_w^i}$.

5.2.1. Monte Carlo trace estimator approach. We begin by deriving a randomized (Monte Carlo) trace estimator for the posterior covariance operator (22). This is facilitated by the following technical result.

Proposition 5.1. Let $\mathcal{C} \in \mathcal{L}_1^{\text{sym++}}(\mathcal{M})$ and $\mathcal{K} \in \mathcal{L}(\mathcal{M})$, and consider the Gaussian measure $\mu = \mathcal{N}(0, \mathcal{C})$ on \mathcal{M} . Then, $\int_{\mathcal{M}} \langle \mathcal{K}z, z \rangle_{\mathcal{M}} \mu(dz) = \text{tr}(\mathcal{C}^{1/2} \mathcal{K} \mathcal{C}^{1/2})$.

Proof. By the assumptions on \mathcal{C} and \mathcal{K} , we have that both $\mathcal{C}\mathcal{K}$ and $\mathcal{C}^{1/2}\mathcal{K}\mathcal{C}^{1/2}$ are trace-class. Also, by the formula for the expectation of a quadratic form in the infinite-dimensional Hilbert space setting [2, Lemma 1], we have that $\int_{\mathcal{M}} \langle \mathcal{K}z, z \rangle_{\mathcal{M}} \mu(dz) = \text{tr}(\mathcal{C}\mathcal{K})$. To finish the proof, we let $\{e_j\}_{j=1}^{\infty}$ be the (orthonormal) basis of the eigenvectors of \mathcal{C} with corresponding (real, positive) eigenvalues $\{\lambda_j\}_{j=1}^{\infty}$, and note $\text{tr}(\mathcal{C}\mathcal{K}) = \sum_j \langle \mathcal{C}\mathcal{K}e_j, e_j \rangle_{\mathcal{M}} = \sum_j \lambda_j \langle \mathcal{K}e_j, e_j \rangle_{\mathcal{M}} = \sum_j \langle \mathcal{K}\mathcal{C}^{1/2}e_j, \mathcal{C}^{1/2}e_j \rangle_{\mathcal{M}} = \sum_j \langle \mathcal{C}^{1/2}\mathcal{K}\mathcal{C}^{1/2}e_j, e_j \rangle_{\mathcal{M}} = \text{tr}(\mathcal{C}^{1/2}\mathcal{K}\mathcal{C}^{1/2})$. \square

Next, note that we can write $\mathcal{C}_{\text{post}}^{\mathbf{y}_w^i}$ as

$$\mathcal{C}_{\text{post}}^{\mathbf{y}_w^i} = \left(\mathcal{F}_m \left(\mathbf{m}_{\text{MAP}}^{\mathbf{y}_w^i} \right)^* \Sigma(\mathbf{w}) \mathcal{F}_m \left(\mathbf{m}_{\text{MAP}}^{\mathbf{y}_w^i} \right) + \mathcal{C}_{\text{pr}}^{-1} \right)^{-1} = \mathcal{C}_{\text{pr}}^{1/2} \left(\tilde{\mathcal{H}}^i(\mathbf{w}) + I \right)^{-1} \mathcal{C}_{\text{pr}}^{1/2}. \quad (29)$$

Using this and letting $\mathcal{C} = \mathcal{C}_{\text{pr}}$ and $\mathcal{K} = (\tilde{\mathcal{H}}^i(\mathbf{w}) + I)^{-1}$ in proposition 5.1, we have $\text{tr}(\mathcal{C}_{\text{post}}^{\mathbf{y}_w^i}) = \int_{\mathcal{M}} \langle (\tilde{\mathcal{H}}^i(\mathbf{w}) + I)^{-1} z, z \rangle_{\mathcal{M}} \mu(dz)$. Approximating the integral on the right hand side via sampling, we obtain

$$\text{tr}(\mathcal{C}_{\text{post}}^{\mathbf{y}_w^i}) \approx \frac{1}{n_{\text{tr}}} \sum_{j=1}^{n_{\text{tr}}} \left\langle (\tilde{\mathcal{H}}^i(\mathbf{w}) + I)^{-1} z_j, z_j \right\rangle_{\mathcal{M}},$$

where $\{z_j\}_{j=1}^{n_{\text{tr}}}$ are draws from the measure $\mu = \mathcal{N}(0, \mathcal{C}_{\text{pr}})$. This randomized (Monte Carlo) trace estimator allows approximating the OED objective (24), as follows.

$$\Phi_{n_d}^{\text{tr}}(\mathbf{w}) = \frac{1}{n_{\text{tr}} n_d} \sum_{i=1}^{n_d} \sum_{j=1}^{n_{\text{tr}}} \langle c_{ij}, z_j \rangle_{\mathcal{M}}, \quad (30a)$$

where, for $i \in \{1, \dots, n_d\}$ and $j \in \{1, \dots, n_{\text{tr}}\}$,

$$\mathbf{m}_{\text{MAP}}^{\mathbf{y}_w^i} = \arg \min_m \mathcal{J}_w(m; \mathbf{y}^i), \quad (30b)$$

$$(\tilde{\mathcal{H}}^i(\mathbf{w}) + I) c_{ij} = z_j. \quad (30c)$$

The major computational challenge in computing $\Phi_{n_d}^{\text{tr}}(\mathbf{w})$ is solving the MAP estimation problems in (30b). These *inner* optimization problems can be solved efficiently using an inexact Newton-Conjugate Gradient (Newton-CG) method. The required first and second order derivatives can be obtained efficiently using adjoint-based gradient and Hessian apply computation. The Hessian solves in (30c) are also done using (preconditioned) CG. This only requires the action of $\tilde{\mathcal{H}}^i(\mathbf{w})$ on vectors, which can be done using the adjoint method. A detailed study of the computational cost of computing $\Phi_{n_d}^{\text{tr}}(\mathbf{w})$ is provided in section 5.4.

5.2.2. Low-rank approximation approach. Due to the use of finite-dimensional observations, $\tilde{\mathcal{H}}^i$ in (28) has a finite-dimensional range. Also, often $\tilde{\mathcal{H}}^i$ exhibits rapid spectral decay and, in cases where the dimension of measurements is high, the numerical rank of this operator is typically much smaller than its exact rank. Note that the exact rank is bounded by the measurement dimension. Such low-rank structures are due to possible smoothing properties of the forward operator and that of \mathcal{C}_{pr} ; see, e.g. [12]. The following technical result facilitates exploiting such low-rank structures to compute the OED objective.

Proposition 5.2. Let \mathcal{C} and \mathcal{A} be in $\mathcal{L}_1^{\text{sym}+}(\mathcal{M})$. Letting $\{v_k\}_{k=1}^{\infty}$ be the orthonormal basis of eigenvectors of \mathcal{A} with corresponding (real non-negative) eigenvalues $\{\lambda_k\}_{k=1}^{\infty}$, we have

$$\text{tr}\left(\mathcal{C}^{1/2}(I+\mathcal{A})^{-1}\mathcal{C}^{1/2}\right) = \text{tr}(\mathcal{C}) - \sum_{k=1}^{\infty} \frac{\lambda_k}{1+\lambda_k} \|\mathcal{C}^{1/2}v_k\|_{\mathcal{M}}^2. \quad (31)$$

Proof. The result follows by noting that

$$\begin{aligned} \text{tr}(\mathcal{C}) - \text{tr}\left(\mathcal{C}^{1/2}(I+\mathcal{A})^{-1}\mathcal{C}^{1/2}\right) &= \sum_{k=1}^{\infty} \langle \mathcal{C}v_k, v_k \rangle_{\mathcal{M}} - \sum_{k=1}^{\infty} \langle \mathcal{C}(I+\mathcal{A})^{-1}v_k, v_k \rangle_{\mathcal{M}} \\ &= \sum_{k=1}^{\infty} \langle \mathcal{C}v_k, v_k \rangle_{\mathcal{M}} - \sum_{k=1}^{\infty} \frac{1}{1+\lambda_k} \langle \mathcal{C}v_k, v_k \rangle_{\mathcal{M}} \\ &= \sum_{k=1}^{\infty} \frac{\lambda_k}{1+\lambda_k} \langle \mathcal{C}_{\text{pr}}v_k, v_k \rangle_{\mathcal{M}} = \sum_{k=1}^{\infty} \frac{\lambda_k}{1+\lambda_k} \|\mathcal{C}^{1/2}v_k\|_{\mathcal{M}}^2. \end{aligned}$$

□

The relation (31) facilitates approximating $\text{tr}(\mathcal{C}^{1/2}(I+\mathcal{A})^{-1}\mathcal{C}^{1/2})$. Namely, we can truncate the infinite summation in the right-hand side to the first r terms, corresponding to the r dominant eigenvalues of \mathcal{A} . We can also quantify the approximation error due to this truncation as follows.

Proposition 5.3. Let \mathcal{A} and \mathcal{C} be as in proposition 5.2 and let $\{\lambda_k\}_{k=1}^{r+1}$ be the $r+1$ largest eigenvalues of \mathcal{A} . Define $T_r(\mathcal{A}, \mathcal{C}) = \text{tr}(\mathcal{C}) - \sum_{k=1}^r \frac{\lambda_k}{1+\lambda_k} \|\mathcal{C}^{1/2}v_k\|_{\mathcal{M}}^2$. Then,

$$|\text{tr}\left(\mathcal{C}^{1/2}(I+\mathcal{A})^{-1}\mathcal{C}^{1/2}\right) - T_r(\mathcal{A}, \mathcal{C})| \leq \frac{\lambda_{r+1}}{1+\lambda_{r+1}} \text{tr}(\mathcal{C}). \quad (32)$$

Proof. We have

$$\begin{aligned} |\text{tr}\left(\mathcal{C}^{1/2}(I+\mathcal{A})^{-1}\mathcal{C}^{1/2}\right) - T_r(\mathcal{A}, \mathcal{C})| &= \sum_{k=r+1}^{\infty} \frac{\lambda_k}{1+\lambda_k} \|\mathcal{C}^{1/2}v_k\|_{\mathcal{M}}^2 \\ &\leq \frac{\lambda_{r+1}}{1+\lambda_{r+1}} \sum_{k=r+1}^{\infty} \|\mathcal{C}^{1/2}v_k\|_{\mathcal{M}}^2 \\ &= \frac{\lambda_{r+1}}{1+\lambda_{r+1}} \sum_{k=1}^{\infty} \langle \mathcal{C}v_k, v_k \rangle_{\mathcal{M}} = \frac{\lambda_{r+1}}{1+\lambda_{r+1}} \text{tr}(\mathcal{C}). \end{aligned}$$

□

Next, we consider $\tilde{\mathcal{H}}^i(\mathbf{w})$ and let $\{(\lambda_{ik}, v_{ik})\}_{k=1}^r$ be its dominant eigenpairs. (The dependence of eigenvalues and eigenvectors on \mathbf{w} is suppressed, for notational convenience.) Using proposition 5.2 with $\mathcal{C} = \mathcal{C}_{\text{pr}}$ and $\mathcal{A} = \tilde{\mathcal{H}}^i(\mathbf{w})$, we obtain the approximation

$$\begin{aligned} \text{tr}\left(\mathcal{C}_{\text{post}}^{\mathbf{y}_{\mathbf{w}}^i}\right) &= \text{tr}\left(\mathcal{C}_{\text{pr}}^{1/2} \left(I + \tilde{\mathcal{H}}^i(\mathbf{w})\right)^{-1} \mathcal{C}_{\text{pr}}^{1/2}\right) \\ &= \text{tr}\left(\mathcal{C}_{\text{pr}} \left(I + \tilde{\mathcal{H}}^i(\mathbf{w})\right)^{-1}\right) \approx \text{tr}(\mathcal{C}_{\text{pr}}) - \sum_{k=1}^r \frac{\lambda_{ik}}{1+\lambda_{ik}} \|\mathcal{C}_{\text{pr}}^{1/2}v_{ik}\|_{\mathcal{M}}^2. \end{aligned} \quad (33)$$

Observe that only the second term in the right-hand side depends on \mathbf{w} . This term, which provides a measure of uncertainty reduction, can be used to define the OED objective. Specifically, using (24) along with (33), leads to following form of the OED objective:

$$\Phi_{n_d}^{\text{eig}}(\mathbf{w}) = -\frac{1}{n_d} \sum_{i=1}^{n_d} \sum_{k=1}^r \frac{\lambda_{ik}}{1+\lambda_{ik}} \|\mathcal{C}_{\text{pr}}^{1/2} v_{ik}\|_{\mathcal{M}}^2, \quad (34a)$$

where, for $i \in \{1, \dots, n_d\}$,

$$m_{\text{MAP}}^{\mathbf{y}_w^i} = \arg \min_m \mathcal{J}_w(m; \mathbf{y}^i), \quad (34b)$$

$$\tilde{\mathcal{H}}^i(\mathbf{w}) v_{ik} = \lambda_{ik} v_{ik} \quad k \in \{1, \dots, r\}, \quad (34c)$$

$$\langle v_{ik}, v_{il} \rangle_{\mathcal{M}} = \delta_{kl} \quad k, l \in \{1, \dots, r\}. \quad (34d)$$

Note that the numerical rank of $\tilde{\mathcal{H}}^i$ will, in general, depend on i . In (34), we have used a common target rank r for simplicity. In the present setting, this target rank is bounded by the number of active sensors for a given \mathbf{w} . Note also that proposition 5.3 shows how to quantify the error due to the use of truncated spectral decompositions. Namely, in view of (32), the error due to the truncation is bounded by $\left(\frac{1}{n_d} \sum_{i=1}^{n_d} \frac{\lambda_{i,r+1}}{1+\lambda_{i,r+1}} \right) \text{tr}(\mathcal{C}_{\text{pr}})$.

As in the case of the approach outlined in section 5.2.1, the dominant computational challenge in computing $\Phi_{n_d}^{\text{eig}}(\mathbf{w})$ is the solution of the MAP estimation problems (34b). This will be tackled using the same techniques. The eigenvalue problem (34c) can be tackled via Lanczos or randomized approaches. The target rank r , can be selected as the number of active sensors. This is suitable in cases where we have a small number of active sensors. See section 5.4 for further details regarding the computational cost of evaluating $\Phi_{n_d}^{\text{eig}}$.

We point out that if the dominant eigenvalues are simple, the condition $\langle v_{ik}, v_{ik} \rangle_{\mathcal{M}} = 1$ for all $k \in \{1, \dots, r\}$, implies the orthonormality condition (34d). This follows from the fact that the eigenvectors corresponding to distinct eigenvalues of $\tilde{\mathcal{H}}^i$, which belongs to $\mathcal{L}_1^{\text{sym}+}(\mathcal{M})$, are orthogonal. The simplicity assumption on the dominant eigenvalues is observed in many applications where the dominant eigenvalues decay rapidly. Making this simplicity assumption, and letting $s_{ik} = \mathcal{C}_{\text{pr}}^{1/2} v_{ik}$, we may also write the eigenvalue problem above as the following generalized eigenvalue problem

$$\begin{aligned} \mathcal{H}^i(\mathbf{w}) s_{ik} &= \lambda_{ik} \mathcal{C}_{\text{pr}}^{-1} s_{ik}, \\ \langle s_{ik}, s_{ik} \rangle_{\mathcal{E}} &= 1, \end{aligned} \quad (35)$$

where $i \in \{1, \dots, n_d\}$ and $k \in \{1, \dots, r\}$. This formulation is helpful when describing the eigenvalue problem in the weak form, as seen in section 6.2. In particular, the eigenvalue problem (35) will be formulated in terms of the incremental state and adjoint equations and the adjoint based expression for $\tilde{\mathcal{H}}^i$ applies.

5.3. Greedy optimization

We follow a greedy approach for solving the OED problem. That is, we select sensors one at a time: at each step, we pick the sensor that results in the greatest decrease in the OED objective value; see algorithm 1. As mentioned before, we use the approach in section 5.2.2 for computing the OED objective. That is, we use the OED objective $\Phi_{n_d}^{\text{eig}}$, as defined in (34).

Algorithm 1. Greedy approach for solving the OED problem (26) with $\Phi_{n_d} = \Phi_{n_d}^{\text{eig}}$.

Input: The target number of sensors K in (26).

Output: The optimal design vector \mathbf{w} .

```

1:  $\mathbf{w} \leftarrow \mathbf{0}$ 
2:  $\mathcal{I}_{\text{candidate}} \leftarrow \{1, \dots, n_s\}$ 
3:  $\mathcal{I}_{\text{active}} \leftarrow \emptyset$ 
4: for  $\ell = 1$  to  $K$  do
5:   Evaluate  $\Phi_{n_d}^{\text{eig}}(\mathbf{w} + \mathbf{e}_j)$ , for all  $j \in \mathcal{I}_{\text{candidate}}$        $\{\mathbf{e}_j \text{ is the } j\text{th coordinate vector in } \mathbb{R}^{n_s}\}$ 
6:    $i_\ell \leftarrow \arg \min_{j \in \mathcal{I}_{\text{candidate}}} \Phi_{n_d}^{\text{eig}}(\mathbf{w} + \mathbf{e}_j)$ 
7:    $\mathcal{I}_{\text{active}} \leftarrow \mathcal{I}_{\text{active}} \cup \{i_\ell\}$ 
8:    $\mathcal{I}_{\text{candidate}} \leftarrow \mathcal{I}_{\text{candidate}} \setminus \{i_\ell\}$ 
9:    $\mathbf{w} \leftarrow \mathbf{w} + \mathbf{e}_{i_\ell}$ 
10: end for

```

Theoretical justifications behind the use of a greedy approach for sensor placement, in various contexts, have been investigated in a number of works [23, 29, 37]. The solution obtained using the greedy algorithm is known to be suboptimal. However, as observed in practice, the use of a greedy algorithm is a practical approach and is effective in obtaining near optimal sensor placements; see also [5, 30, 44]. We demonstrate the effectiveness of this approach, in our computational results in section 7.

5.4. Computational cost of sensor placement

The ℓ th step of the greedy algorithm requires $n_s - \ell - 1$ OED objective evaluations (cf step 5 of algorithm 1); these can be performed in parallel. It is straightforward to note that placing K sensors using the greedy approach requires a total of $C(K, n_s) := Kn_s - K(K - 1)/2$ OED objective function evaluations. Therefore, the overall cost of computing an optimal design, in terms of the number of PDE solves, using the proposed approach is $C(K, n_s)$ times the number of PDE solves required in each OED objective evaluation. Thus, to provide a complete picture, in this section, we detail the cost of OED objective evaluation using the two approaches discussed in section 5.2. A key aspect of both of these approaches is that the cost of OED objective evaluation, in terms of the number of required PDE solves, is independent of the dimension of the discretized inversion and secondary parameters.

In the following discussion, the rank of the operators $\tilde{\mathcal{H}}^i(\mathbf{w})$ in (28) plays an important role. As mentioned before, the exact rank of these operators is bounded by the number of active sensors. Moreover, if the number of active sensors is high, the numerical rank is typically considerably smaller than the exact rank. We denote the number of active sensors in a given design by $n_{\text{act}} = n_{\text{act}}(\mathbf{w})$. Note that for a binary design vectors \mathbf{w} , $n_{\text{act}}(\mathbf{w}) = \|\mathbf{w}\|_1$.

5.4.1. Cost of evaluating $\Phi_{n_d}^{\text{tr}}$. The most expensive step in evaluating (30) is solving the inner optimization problems for the MAP points $m_{\text{MAP}}^{\mathbf{y}_w^i}$, $i \in \{1, \dots, n_d\}$. In the present work, the inner optimization problem is solved via inexact Gauss–Newton–CG with line search. The cost of each Gauss–Newton iteration is dominated by the CG solves for the search direction. When using the prior covariance operator as a preconditioner, the number of CG iterations is bounded by $\mathcal{O}(n_{\text{act}})$; see e.g. [22]. Each CG step in turn requires two linearized PDE solves (incremental forward/adjoint solves). Hence, the cost of the n_d MAP estimation problems, in

terms of ‘forward-like’ PDE solves, is $\mathcal{O}(2 \times n_d \times n_{\text{newton}} \times n_{\text{act}})$, where n_{newton} is the (average) number of Gauss–Newton iterations. The $2 \times n_d \times n_{\text{tr}}$ solves in (30c) are also done using preconditioned CG. As noted before, each of these solves requires $\mathcal{O}(n_{\text{act}})$ CG iterations. To summarize, the overall cost of evaluating (30) is bounded by

$$\underbrace{\mathcal{O}(2 \times n_d \times n_{\text{newton}} \times n_{\text{act}})}_{\text{cost of MAP point solves in (30b)}} + \underbrace{\mathcal{O}(2 \times n_d \times n_{\text{tr}} \times n_{\text{act}})}_{\text{cost of solves in (30c)}}. \quad (36)$$

It is important to note that MAP estimation problems as well as the linear solves in (30c) can be performed in parallel.

5.4.2. Cost of evaluating $\Phi_{n_d}^{\text{eig}}$. As is the case with (30), the solution of the MAP estimation problems (34b) dominates the computational cost of evaluating $\Phi_{n_d}^{\text{eig}}$. This cost was analyzed for the case of $\Phi_{n_d}^{\text{tr}}$. We next discuss the cost of solving the eigenvalue problems in (34c). We rely on the Lanczos method [19] for these eigenvalue problems.⁶ For each $i \in \{1, \dots, n_d\}$, the Lanczos method requires $\mathcal{O}(n_{\text{act}})$ applications of $\tilde{\mathcal{H}}^i$ on vectors, each costing two PDE solves. Therefore, solving the eigenvalue problems requires a total of $\mathcal{O}(2 \times n_d \times n_{\text{act}})$ PDE solves. Hence, the overall cost of computing $\Phi_{n_d}^{\text{eig}}$ is bounded by

$$\underbrace{\mathcal{O}(2 \times n_d \times n_{\text{newton}} \times n_{\text{act}})}_{\text{cost of MAP point solves in (34b)}} + \underbrace{\mathcal{O}(2 \times n_d \times n_{\text{act}})}_{\text{cost of solves in (34c)}}. \quad (37)$$

To sum up, the computational cost of evaluating both $\Phi_{n_d}^{\text{tr}}$ and $\Phi_{n_d}^{\text{eig}}$ is dominated by the cost of solving the MAP estimation problems. However, by exploiting the low-rank structure of the operators $\tilde{\mathcal{H}}^i$, $\Phi_{n_d}^{\text{eig}}$ is more efficient to compute, as it does not require n_{tr} Hessian solves (compare also the second terms in (36) and (37)). Moreover, computing the traces using the eigenvalues will be, in general, more accurate than a sampling based approach, as long as sufficiently many eigenvalues are used. Note also that when using a greedy approach, n_{act} starts at $n_{\text{act}} = 1$ in the first step and increase by one in each iteration.

However, the idea of using a randomized trace estimator does have some merits. For one thing, typically a small (in order of tens) n_{tr} is sufficient when solving the OED problem. Moreover, $\Phi_{n_d}^{\text{tr}}$ is simpler to implement as it does not require solving eigenvalue problems. Note, however, that if a large n_{tr} is needed, then the cost of Hessian solves in (30c) might exceed the cost of the MAP estimation problems.

6. Model problem

Here, we present the model problem used to study our approach for OED under uncertainty. We consider a nonlinear inverse problem governed by a linear elliptic PDE, in a three-dimensional domain. This problem, which is adapted from [33], is motivated by heat transfer applications. In section 6.1, we detail the description of the Bayesian inverse problem under study. Subsequently, we detail the description of the OED objective $\Phi_{n_d}^{\text{eig}}$, for this specific model problem in section 6.2.

⁶ Another option is the use of randomized methods [21].

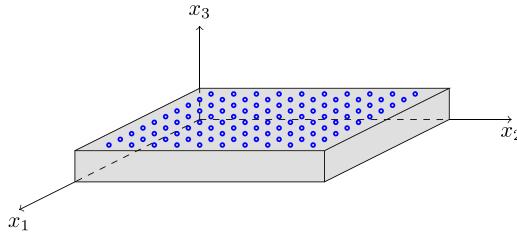


Figure 1. Sketch of the physical domain Ω and location of candidate sensor locations (blue circles).

6.1. The Bayesian inverse problem

We consider the following model

$$\begin{aligned}
 -\nabla \cdot (e^\xi \nabla u) &= 0 \quad \text{in } \Omega, \\
 u &= 0 \quad \text{on } \Gamma_D, \\
 e^\xi \nabla u \cdot \mathbf{n} &= h \quad \text{on } \Gamma_N, \\
 e^\xi \nabla u \cdot \mathbf{n} + e^m u &= 0 \quad \text{on } \Gamma_R.
 \end{aligned} \tag{38}$$

In this problem, $\Omega \subset \mathbb{R}^3$ is bounded domain with sufficiently smooth boundary $\Gamma = \Gamma_D \cup \Gamma_N \cup \Gamma_R$, where Γ_D , Γ_N , and Γ_R are mutually disjoint. In (38), u is the state variable, m is the inversion parameter, and ξ is the secondary uncertain parameter. The Neumann boundary data is $h \in L^2(\Gamma_N)$, \mathbf{n} is the unit-length outward normal for the boundaries Γ_N and Γ_R , and for simplicity we have considered zero volume source, and homogeneous Dirichlet and Robin conditions. In our numerical experiments, Ω is a three-dimensional domain with a unit square base and height such that aspect ratio (base/height) is 100; see figure 1. In the present example, Γ_R is the bottom edge of the domain, Γ_N is the top edge, and Γ_D is the union of the side edges. In our computations, we define the boundary source term h according to $h(\mathbf{x}) = 1 + \sin(4\pi\sqrt{(x_1 - 1)^2 + (x_2 - 1)^2})$.

Note that in this problem, the inversion and secondary parameters are both functions. The inversion parameter belongs to the Hilbert space $\mathcal{M} = L^2(\Gamma_R)$, which is equipped with the $L^2(\Gamma_R)$ -inner product. The secondary parameter ξ takes values in $L^2(\Omega)$. As discussed below, we define ξ as an $L^2(\Omega)$ -valued Gaussian random variable with almost surely continuous realizations, which ensures (almost sure) well-posedness of the problem (38). Below we use $\langle \cdot, \cdot \rangle_\Omega$ and $\langle \cdot, \cdot \rangle_{\Gamma_N}$ to denote the $L^2(\Omega)$ - and $L^2(\Gamma_N)$ -inner products, respectively. Also, with a slight abuse of notation, we denote the $L^2(\Omega)^3$ inner-product with $\langle \cdot, \cdot \rangle_\Omega$ as well.

We use measurements of u on the top boundary, Γ_N , to estimate the inversion parameter m . The measurements are collected at a set of sensor locations as depicted in figure 1. Hence, in the present problem, the parameter-to-observable is defined as the composition of a linear observation operator \mathcal{B} , which extracts solution values at the sensor locations, and the PDE solution operator. Note that this parameter-to-observable map is a nonlinear function of the inversion and secondary parameters.

As detailed in section 3, we assume a Gaussian noise model, and use the BAE approach to account for the secondary uncertainties in the inverse problem. Also, we use a Gaussian

prior law for m and model the secondary uncertainty as a Gaussian random variable.⁷ For both m and ξ , we use covariance operators that are defined as negative powers of Laplacian-like operators [40]; see section 7. Note that with the appropriate choice of the covariance operator, the realizations of ξ will be almost surely continuous on the closure of Ω .

6.2. OED under uncertainty problem formulation

In this section we discuss the precise definition of $\Phi_{n_d}^{\text{eig}}$, for the model inverse problem discussed in section 6.1. The MAP estimation problems in (34b) require minimizing $\mathcal{J}_w(m; y^i)$ defined in (21). For notational convenience, in this section, we use the shorthand m_i^* to denote $m_{\text{MAP}}^{y^i}$. In the present example, the first order optimality conditions for this optimization problem are given by

$$\langle e^\xi \nabla u_i, \nabla \tilde{p} \rangle_\Omega - \langle h, \tilde{p} \rangle_{\Gamma_N} + \langle e^{m_i^*} u_i, \tilde{p} \rangle_{\mathcal{M}} = 0, \quad \forall \tilde{p} \in \mathcal{V}_0, \quad (39)$$

$$\langle e^\xi \nabla p_i, \nabla \tilde{u} \rangle_\Omega + \langle e^{m_i^*} \tilde{u}, p_i \rangle_{\mathcal{M}} + \langle \mathcal{B}^* \Sigma(\mathbf{w}) (\mathcal{B} u_i - y^i + \varepsilon_0), \tilde{u} \rangle_\Omega = 0, \quad \forall \tilde{u} \in \mathcal{V}_0, \quad (40)$$

$$\langle m_i^* - m_{\text{pr}}, \tilde{m} \rangle_{\mathcal{E}} + \langle \tilde{m} e^{m_i^*} u_i, p_i \rangle_{\mathcal{M}} = 0, \quad \forall \tilde{m} \in \mathcal{E}, \quad (41)$$

where $\mathcal{V}_0 = \{v \in H^1(\Omega) : v|_{\Gamma_D} = 0\}$. For the derivation details, we refer to [33]. Note that (39) is the weak form of the state equations (38), (40) and (41) are the weak forms of the adjoint and gradient equations, respectively. Note also that the left hand side of (41) describes the action of the derivative of $\mathcal{J}_w(m; y^i)$ in the direction \tilde{m} .

To specify the eigenvalue problem in (34c), we first discuss the action of the operator \mathcal{H}^i (evaluated at m_i^*) in the direction \tilde{m} . In weak form, this Hessian application satisfies,

$$\langle \mathcal{H}^i(m_i^*) \tilde{m}, \tilde{m} \rangle_{\mathcal{M}} = \langle \tilde{m} e^{m_i^*} u_i, \hat{p}_i \rangle_{\mathcal{M}}, \quad (42)$$

for every $\tilde{m} \in \mathcal{E}$. In (42), for a given m_i^* , u_i solves the state problem (38), \hat{p}_i solves the incremental adjoint problem

$$\langle e^\xi \nabla \hat{p}_i, \nabla \tilde{u} \rangle_\Omega + \langle e^{m_i^*} \tilde{u}, \hat{p}_i \rangle_{\mathcal{M}} + \langle \mathcal{B}^* \Sigma(\mathbf{w}) \mathcal{B} \hat{u}_i, \tilde{u} \rangle_\Omega = 0, \quad \forall \tilde{u} \in \mathcal{V}_0, \quad (43)$$

and \hat{u}_i solves the so-called incremental state problem

$$\langle e^\xi \nabla \hat{u}_i, \nabla \tilde{p} \rangle_\Omega + \langle e^{m_i^*} \hat{u}_i, \tilde{p} \rangle_{\mathcal{M}} + \langle e^{m_i^*} \hat{m} u_i, \tilde{p} \rangle_{\Gamma_R} = 0, \quad \forall \tilde{p} \in \mathcal{V}_0. \quad (44)$$

We next summarize the OED problem of minimizing (34) as a PDE-constrained optimization problem specialized for the model inverse problem in section 6:

$$\min_{\mathbf{w} \in \{0,1\}^{n_d}} -\frac{1}{n_d} \sum_{i=1}^{n_d} \sum_{k=1}^r \frac{\lambda_{ik}}{1 + \lambda_{ik}} \|s_{ik}\|_{\mathcal{M}}^2, \quad (45a)$$

where, for $i \in \{1, \dots, n_d\}$ and $k \in \{1, \dots, r\}$,

$$\langle e^\xi \nabla u_i, \nabla \tilde{p} \rangle_\Omega - \langle h, \tilde{p} \rangle_{\Gamma_N} + \langle e^{m_i^*} u_i, \tilde{p} \rangle_{\mathcal{M}} = 0, \quad \forall \tilde{p} \in \mathcal{V}_0, \quad (45b)$$

⁷ As discussed earlier, the presented framework does not rely on a specific choice of distribution law for ξ . The Gaussian assumption here is merely for computational convenience.

$$\langle e^\xi \nabla p_i, \nabla \tilde{u} \rangle_\Omega + \langle e^{m_i^*} \tilde{u}, p_i \rangle_{\mathcal{M}} + \langle \mathcal{B}^* \Sigma(\mathbf{w}) (\mathcal{B} u_i - \mathbf{y}^i + \boldsymbol{\varepsilon}_0), \tilde{u} \rangle_\Omega = 0, \quad \forall \tilde{u} \in \mathcal{V}, \quad (45c)$$

$$\langle m_i^* - m_{\text{pr}}, \tilde{m} \rangle_{\mathcal{E}} + \langle \tilde{m} e^{m_i^*} u_i, p_i \rangle_{\mathcal{M}} = 0, \quad \forall \tilde{m} \in \mathcal{E}, \quad (45d)$$

$$\langle e^\xi \nabla \hat{p}_{ik}, \nabla \tilde{u} \rangle_\Omega + \langle e^{m_i^*} \tilde{u}, \hat{p}_{ik} \rangle_{\mathcal{M}} + \langle \mathcal{B}^* \Sigma(\mathbf{w}) \mathcal{B} \hat{u}_{ik}, \tilde{u} \rangle_\Omega = 0, \quad \forall \tilde{u} \in \mathcal{V}, \quad (45e)$$

$$\langle \tilde{m} e^{m_i^*} u_i, \hat{p}_{ik} \rangle_{\mathcal{M}} = \lambda_{ik} \langle s_{ik}, \tilde{m} \rangle_{\mathcal{E}}, \quad \forall \tilde{m} \in \mathcal{E}, \quad (45f)$$

$$\langle e^\xi \nabla \hat{u}_{ik}, \nabla \tilde{p} \rangle_\Omega + \langle e^{m_i^*} \hat{u}_{ik}, \tilde{p} \rangle_{\mathcal{M}} + \langle e^{m_i^*} s_{ik} u_i, \tilde{p} \rangle_{\mathcal{M}} = 0, \quad \forall \tilde{p} \in \mathcal{V}, \quad (45g)$$

$$\langle s_{ik}, s_{ik} \rangle_{\mathcal{E}} = 1. \quad (45h)$$

The PDE constraints (45b)–(45d) are the optimality system (39)–(41) characterizing the MAP point $m_i^* = m_{\text{MAP}}^{\mathbf{y}_w^i}$ described in (34b). The equations (45e)–(45g) are the PDE constraints that describe the Hessian apply and eigenvalue problem. Note that we have reformulated the eigenvalue problem according to (35).

7. Computational results

In this section, we numerically study our OED under uncertainty approach, which we apply to the model inverse problem described in section 6. We begin by specifying the Bayesian problem setup and discretization details in section 7.1. Then, we study the impact of secondary model uncertainty on the measurements and compute the BAE error statistics in section 7.2. Finally, in section 7.3, we examine the effectiveness of our approach in computing uncertainty aware designs. We also study the impact of ignoring model uncertainty in (i) experimental design stage and (ii) both experimental design and inference stages. Ignoring uncertainty in the design stage amounts to fixing the secondary parameter ξ at its nominal value (i.e. using the approximate model (10)) and ignoring the approximation error when solving the OED problem. We refer to designs computed in this manner as *uncertainty unaware* designs. Note that in the case of (i), the uncertainty is still accounted for when solving the inverse problem. This study illustrates the importance of accounting for model uncertainty, when computing experimental designs. On the other hand, the study of case (ii) illustrates the pitfalls of ignoring uncertainty in both the optimal design problem and subsequent solution of the inverse problem.

7.1. Problem setup

We consider $n_s = 100$ candidate sensor locations that are arranged in a regular grid on the top boundary Γ_N ; see figure 1. The additive noise in the synthetic measurements has a covariance matrix of the form $\Gamma_{\text{noise}} = \sigma^2 \mathbf{I}$, with $\sigma = 10^{-3}$. This amounts to about one percent noise.

We use a Gaussian prior law $\mu_{\text{pr}} = \mathcal{N}(m_{\text{pr}}, \mathcal{C}_{\text{pr}})$ for m . The prior mean is taken as a constant function $m_{\text{pr}} \equiv 1$, and we use a covariance operator given by the inverse of a Laplacian-like operator. Specifically, we let $\mathcal{C}_{\text{pr}} := \mathcal{A}^{-2}$, with $\mathcal{A}[m] = -\nabla \cdot (\theta \nabla m) + \alpha m$, where we take $\theta = 0.1$ and $\alpha = 1$. To help mitigate undesirable boundary affects that can arise due to the use of PDE-based prior covariance operators, we equip the operator \mathcal{A} with Robin boundary conditions [16]. For illustration, four random draws from the prior distribution are shown in figure 2.

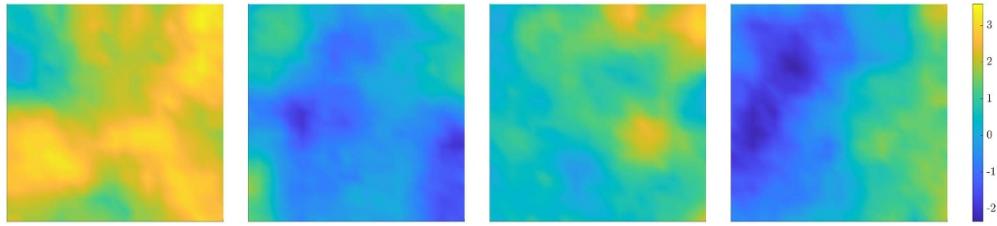


Figure 2. Samples of the primary uncertain parameter m .

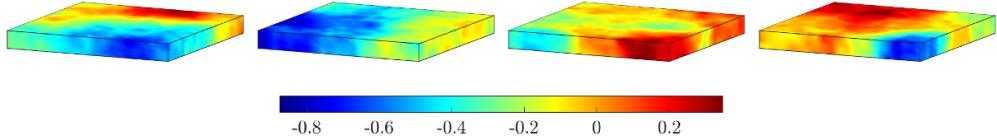


Figure 3. Samples of the secondary uncertain parameter ξ .

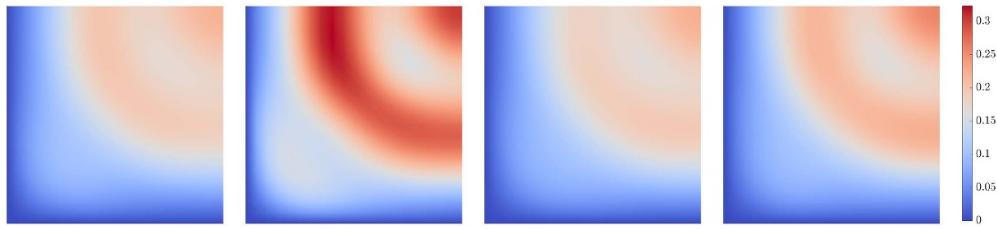


Figure 4. The effect of the secondary uncertain parameter ξ on the state u on top surface for the (fixed) true primary parameter m and four different realizations of ξ .

The law of the secondary parameter ξ is chosen to be a Gaussian measure $\mu_\xi = \mathcal{N}(\bar{\xi}, \mathcal{C}_\xi)$. We set the mean of ξ as the constant function $\bar{\xi} \equiv 0$, and the covariance operator is defined as $\mathcal{C}_\xi := \mathcal{L}^{-2}$, where $\mathcal{L}[\xi] = -\nabla \cdot (\Theta \nabla \xi) + \gamma \xi$ with $\Theta = 0.25 \text{ diag}(1, 1, \frac{1}{100})$ and $\gamma = 50$. This choice of Θ corresponds to a random field with much shorter correlation in the z -direction than in the x - and y -directions, inline with aspect ratio of 100 used in defining the domain Ω . We show four representative samples of ξ in figure 3.

The forward problem (38) is solved using a continuous Galerkin finite element method. We use 9600 tetrahedral elements and 2205 piecewise linear basis functions. As such, the discretized state and adjoint variables, as well as the secondary parameter ξ , have dimension 2205. On the other hand, the discretized inversion parameter m , which is defined on the bottom boundary, is of dimension 441.

7.2. Incorporating the model uncertainty in the inverse problem

We begin by studying the impact of the secondary parameter on the solution of the forward problem. This is illustrated in figure 4. In this experiment, we solve the forward problem for a fixed m and four different realizations of ξ . Note that, although the qualitative behavior of u on Γ_N is similar for the different samples of ξ , there are considerable differences in the values. This indicates that the approximation error due to fixing ξ at the nominal value will have significant variations.

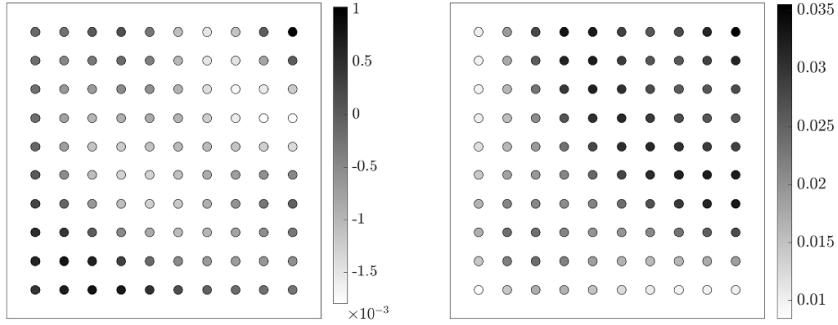


Figure 5. Left: the mean of the approximation error at the candidate sensor locations; right: the marginal standard deviation of the approximation error at these locations (i.e. the square root of diagonal entries of $\widehat{\Gamma}_\varepsilon$ given in (27)).

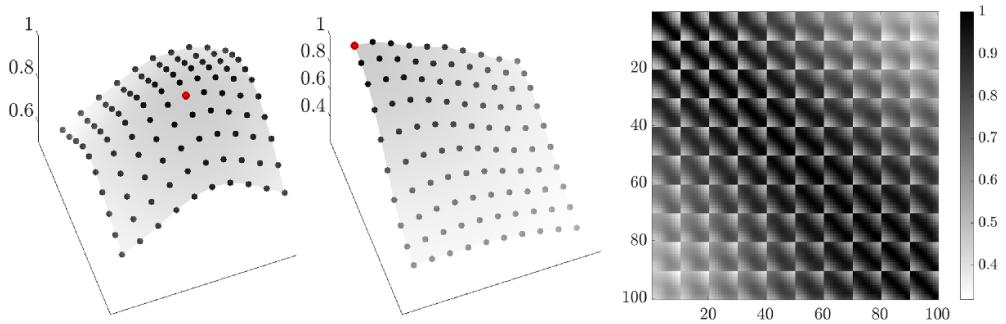


Figure 6. Left: the correlation of a measurement from a sensor near the middle marked by red dot, with the remaining sensors; middle: The same information for the sensor in the top left corner of the sensor grid; right: the correlation matrix of the approximation error.

We compute the sample mean and covariance matrix of the approximation error as in (27) with $n_{mc} = 1000$. The mean and marginal standard deviations are shown in figure 5. To illustrate the correlation structure of the approximation error, we show the correlation of the approximation error at two of the candidate locations with the other candidate locations in figure 6 (left-middle). To provide an overall picture, we report the correlation matrix of the approximation error in figure 6 (right). Note that the approximation errors at the sensor sites are not only larger than the measurement noise, but also they are highly correlated and have a nonzero and non-constant mean. Thus, in the present application, the approximation error due to model uncertainty cannot be ignored and needs to be accounted for.

7.3. Optimal experimental design under uncertainty

We begin by solving the OED problem (45) with $n_d = 5$ training data samples. In figure 7 (left), we show an uncertainty aware optimal sensor placement with 20 sensors. Note that due to the use of a greedy algorithm, we can track the order in which the sensors are picked. To understand the impact of ignoring model uncertainty in the design stage, we also compute an uncertainty unaware design; this is reported in figure 7 (right).

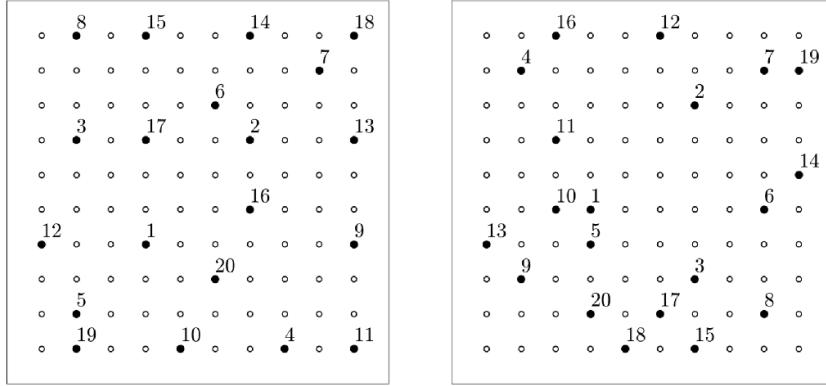


Figure 7. The uncertainty aware (left) and uncertainty unaware (right) sensor placements. The numbers indicate the order in which the sensors are picked in the greedy algorithms for minimizing (34a).

To evaluate the quality of the computed designs, we first study the expected posterior variance and expected relative error of the MAP point. To facilitate this, we draw parameter samples $\{m_i^v\}_{i=1}^{n_v}$ from μ_{pr} and generate validation data samples,

$$\mathbf{d}_w^i = \mathbf{P}_w(\mathcal{G}(m_i^v, \xi_i^v) + \boldsymbol{\eta}_i^v), \quad i \in \{1, \dots, n_v\},$$

where $\{\xi_i^v\}_{i=1}^{n_v}$ and $\{\boldsymbol{\eta}_i^v\}_{i=1}^{n_v}$ are draws from μ_ξ and $\mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma}_{\text{noise}})$, respectively. Then, we consider

$$\bar{V}(\mathbf{w}) = \frac{1}{n_v} \sum_{i=1}^{n_v} \text{tr}(\mathcal{C}_{\text{post}}^{\mathbf{d}_w^i}) \quad \text{and} \quad \bar{E}_{\text{MAP}}(\mathbf{w}) = \frac{1}{n_v} \sum_{i=1}^{n_v} \frac{\|m_{\text{MAP}}^{\mathbf{d}_w^i} - m_i^v\|_{\mathcal{M}}}{\|m_i^v\|_{\mathcal{M}}}. \quad (46)$$

In our numerical experiments we use $n_v = 100$. We compare $\bar{V}(\mathbf{w})$ and $\bar{E}_{\text{MAP}}(\mathbf{w})$ when solving the inverse problem with the computed optimal design versus randomly chosen designs in figure 8, where we consider designs with different numbers of sensors. We also examine the impact of ignoring the model uncertainty in solving the OED problem (see the black dots in figure 8). Note that the uncertainty aware designs outperform the random designs as well as the uncertainty unaware designs. This is most pronounced when the number of sensors is small. This is precisely when optimal placement of sensors is crucial. We also observe that as the number of sensors in the designs increase, the cloud moves left and downward. This is expected. The more sensors we use, the more we can improve the quality of the MAP point and reduce posterior uncertainty. For $n_s \geq 30$, we note that the difference between uncertainty aware and uncertainty unaware designs become small (in the case of $n_s = 30$, they nearly overlap). For $n_d \geq 40$, even the random designs become competitive. All these are expected. As mentioned, earlier, optimal placement of sensors is crucial, when we have access only to a ‘small’ number measurement points. What constitutes ‘small’ is problem dependent.

Next, we illustrate the effectiveness of the computed optimal designs in reducing posterior uncertainty. To this end, we consider the computed optimal designs with 10 sensors. In figure 9, we show the effectiveness of the uncertainty aware optimal design in reducing posterior uncertainty; we also report the posterior standard deviation field, when solving the inverse problem with an uncertainty unaware design. To complement this study, we consider the quality of the MAP points computed using uncertainty aware and uncertainty unaware designs in figure 10.

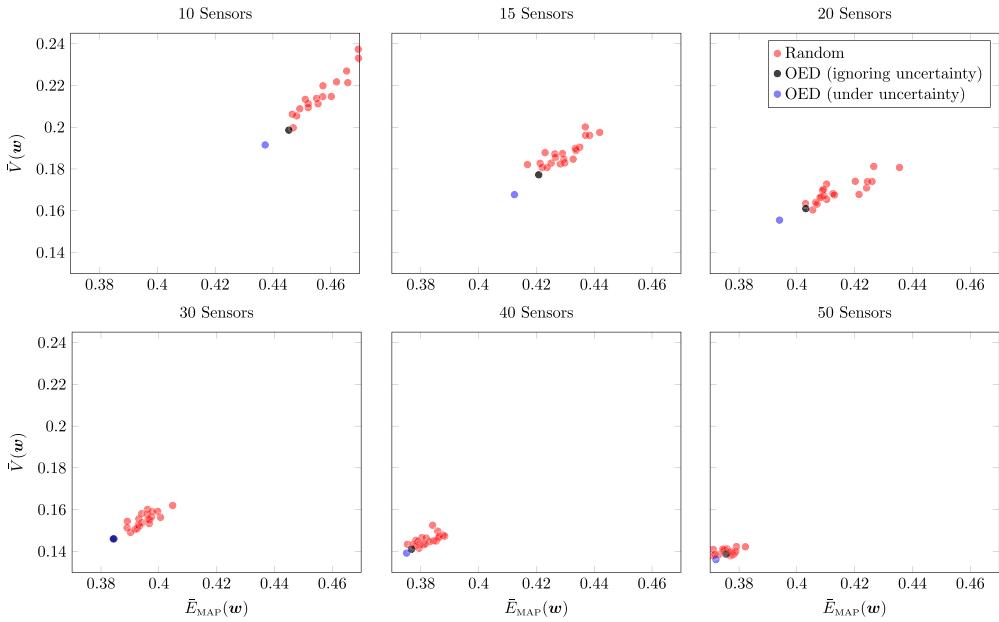


Figure 8. The expected relative error $\bar{E}_{\text{MAP}}(w)$ of the MAP point versus expected posterior variance $\bar{V}(w)$ using random designs (red dots), an uncertainty aware sensor placements (blue dots), and uncertainty unaware sensor placements (black dots) using 10, 15, and 20 sensors.

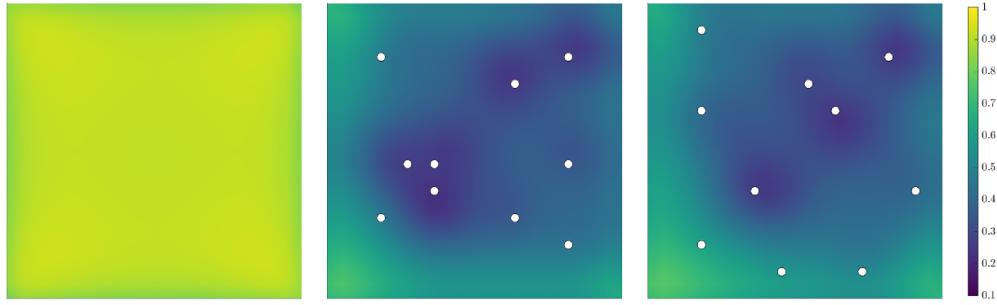


Figure 9. The effectiveness of the computed optimal design with 10 sensors on posterior uncertainty. We report the pointwise prior standard deviation of m (left), and the pointwise posterior standard deviation of m using the uncertainty aware (middle) and uncertainty unaware (right) optimal designs.

Overall, we observe that the uncertainty aware design is more effective in reducing posterior uncertainty and results in a higher quality MAP point. This conclusion is also supported by the results reported in figure 8.

Note that the data used to solve the inverse problem is synthesized by solving the PDE model (1), using our choice of the ‘truth’ inversion parameter (see figure 10 (left)) and a randomly chosen ξ followed by extracting measurements at the sensor sites and adding measurement noise. This simulates the practical situation when field data that corresponds to an unknown choice of ξ is collected.

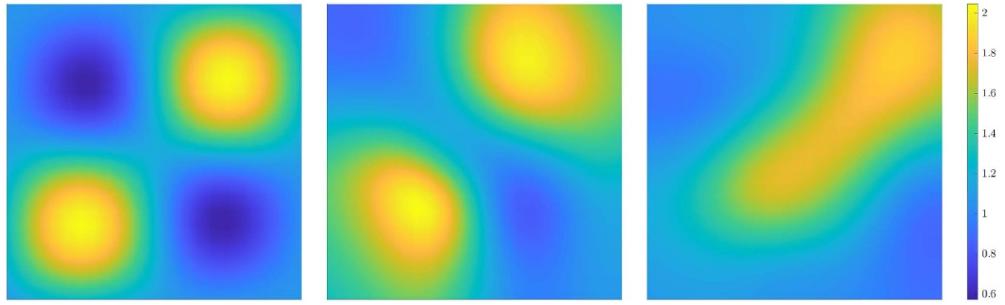


Figure 10. The quality of the MAP points computed using optimal designs with 10 sensors. We show the true parameter field (left), the MAP point computed using the uncertainty aware optimal design (middle) and using the uncertainty unaware design (right).

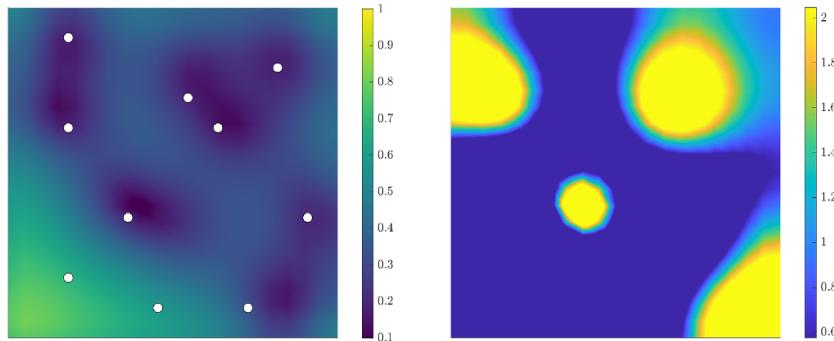


Figure 11. The posterior standard deviation field (left) and MAP point (right) when solving the uncertainty unaware inverse problem using an uncertainty unaware design.

In the above experiments, when examining the performance of uncertainty unaware designs, model uncertainty was still accounted for (following the BAE framework) when solving the inverse problem using these designs. Next, we examine the impact of ignoring model uncertainty in both OED and inference stages. For this experiment, we use the same synthesized data as that used to obtain the results in figure 10 (right). In figure 11 we report the result of solving the inverse problem with an uncertainty unaware design, when ignoring model uncertainty in the inverse problem as well. We note an impressive reduction of uncertainty; see figure 11 (left), which uses the same scale as the standard deviation plots in figure 9. On the other hand, the MAP point computed in this case is of very poor quality; see figure 11 (right), where we have used the same scale as the MAP point plots in figure 10. Intuitively, this indicates that ignoring model uncertainty in both OED and inference stages, in presence of significant modeling uncertainties, can lead to an unfortunate situation where one is highly certain (i.e. low posterior variance) about a very wrong parameter estimate.

As mentioned earlier, we used $n_d = 5$ training data samples when solving the OED problem (45). Clearly, increasing n_d would improve the quality of the design—we would be optimizing a more accurate estimate of the OED objective. However, this comes with increased

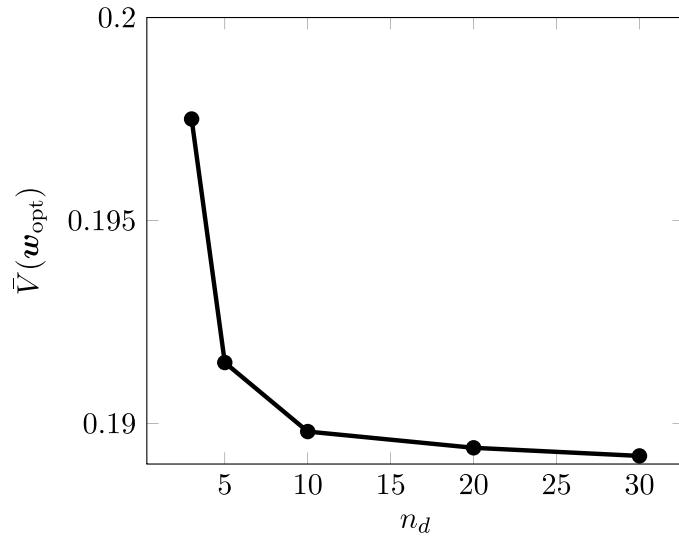


Figure 12. The impact of the number of training data samples n_d on the quality of the computed optimal design with 10 sensors.

computational overhead. In practice, typically a small n_d is effective in obtaining good quality designs. This is demonstrated in the results above (e.g. figure 8). To examine the impact of changing n_d on the quality of the computed optimal design, we perform a brief computational study where we focus on designs with 10 sensors. We compute the expected value of the posterior variance (as in (46)) at the computed optimal design, w_{opt} , with values of $n_d \in \{3, 5, 10, 20, 30\}$. The results are reported in figure 12. As before, $\bar{V}(w_{\text{opt}})$ is computed with $n_v = 100$ validation data samples. Note that going beyond $n_d = 10$ results in diminishing returns. This, however, entails increased computational cost. In practice, studies such as the one conducted here may be necessary to estimate a reasonable choice for n_d . In the present example, we see that $n_d = 5$ is approximately at the ‘elbow’ of the curve in figure 12, making it a reasonable practical choice.

8. Conclusions

In the present work, we addressed OED under uncertainty for Bayesian nonlinear inverse problems governed by PDEs with infinite-dimensional inversion and secondary parameters. We have presented a mathematical framework and scalable computational methods for computing uncertainty aware optimal designs. Our results demonstrate that ignoring the uncertainty in the OED and/or the parameter inversion stages can lead to inferior designs and inaccurate parameter estimation. Hence, it is important to account for modeling uncertainties in Bayesian inversion and OED.

The limitations of the proposed approach are in its reliance on Gaussian approximations for the posterior and the approximation error. The former is a common approach in large-scale Bayesian inverse problems as well as in the BAE literature. The Gaussian approximation to the posterior is suitable if a linearization of the forward model, at the MAP point, is sufficiently

accurate for the set of parameters with significant posterior probability. On the other hand, the Gaussian approximation of the approximation error, which is guided by the BAE approach, might fail to adequately capture the distribution of the approximation error. However, as shown in various studies in the BAE literature, this Gaussian approximation is reasonable in broad classes of inverse problems; see, e.g. [24, 27, 32, 33]. Additionally, in the present work we considered a greedy approach for tackling the binary OED optimization problems. This can become expensive if the number of candidate sensor locations or the desired number of sensors in the computed sensor placements become very large.

The numerical experiments in the present work focus on an academic, albeit application-driven, model problem. The present application was chosen since it exhibits key problem structures seen in broad classes of ill-posed inverse problems. Namely, smoothing properties of the forward operator lead to low-rank structures that can be exploited when developing numerical methods. The present numerical study also builds on and complements the previous study of BAE for the Robin boundary condition inversion problems in [33]. Examining the properties of the present method in more challenging inverse problems with multiple types of modeling uncertainties is a subject for future work.

The discussions in this article point to a number of opportunities for future work. In the first place, we point out that the use of BAE approach to account for modeling uncertainties in Bayesian inverse problems is only one possible application of this approach. In general, BAE has been used in a wide range of applications to account for uncertainty due to the use of approximate forward models. For example, BAE can be used to model the approximation errors due to the use of reduced order models or upscaled models. BAE has also been used to account for the errors due to the use of a mean-field model instead of an underlying high-fidelity stochastic model [38]. Thus, the framework presented for OED under uncertainty in the present work can be extended to OED in inverse problems where the forward model is replaced with a low-fidelity approximate model instead of a computationally intensive high-fidelity model, as long the approximation error can be modeled adequately by a Gaussian.

Another interesting line of inquiry involves replacing the greedy strategy used to find optimal designs with more powerful optimization algorithms. One possibility is to follow a relaxation approach [3, 4, 8, 31] where the design weights are allowed to take values in the interval $[0, 1]$. This enables use of efficient gradient-based optimization algorithms and can be combined with a suitable penalty approach to control the sparsity of the computed sensor placements. An attractive alternative is the approach in [9], which tackles the binary OED optimization problem by replacing it with a related stochastic programming problem. A further line of inquiry is investigating the idea of using fixed MAP points, computed prior to solving the OED problem, as done in [44] for OED problems with no additional uncertainties. This idea, at the expense of further approximations, replaces the bilevel optimization problem with a simpler one, hence reducing computational cost significantly.

Finally, in inverse problems governed by complex models with multiple sources of secondary uncertainty, an *a priori* sensitivity analysis of the Bayesian inverse problem may be necessary to identify secondary modeling uncertainties that are most influential to the solution of the inverse problem. This may be accomplished by suitable adaptations of hyper-differential sensitivity analysis methods for inverse problems [36, 41, 42].

Data availability statement

No new data were created or analysed in this study.

Acknowledgment

The work of N Petra was supported in part by US National Science Foundation Grant DMS #1723211. The work of A Alexanderian was supported in part by US National Science Foundation Grants DMS #1745654 and #2111044.

ORCID iDs

Alen Alexanderian  <https://orcid.org/0000-0002-6371-6618>
 Noemi Petra  <https://orcid.org/0000-0002-9491-0034>

References

- [1] Alexanderian A 2021 Optimal experimental design for infinite-dimensional Bayesian inverse problems governed by PDEs: a review *Inverse Problems* **37** 043001
- [2] Alexanderian A, Gloor P J and Ghattas O 2016 On Bayesian A-and D-optimal experimental designs in infinite dimensions *Bayesian Anal.* **11** 671–95
- [3] Alexanderian A, Petra N, Stadler G and Ghattas O 2014 A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0 -sparsification *SIAM J. Sci. Comput.* **36** A2122–48
- [4] Alexanderian A, Petra N, Stadler G and Ghattas O 2016 A fast and scalable method for A-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems *SIAM J. Sci. Comput.* **38** A243–72
- [5] Alexanderian A, Petra N, Stadler G and Sunseri I 2021 Optimal design of large-scale Bayesian linear inverse problems under reducible model uncertainty: good to know what you don't know *SIAM/ASA J. Uncertain. Quantification* **9** 163–84
- [6] Aravkin A Y and Van Leeuwen T 2012 Estimating nuisance parameters in inverse problems *Inverse Problems* **28** 115016
- [7] Atkinson A C and Donev A N 1992 *Optimum Experimental Designs* (Clarendon Press)
- [8] Attia A and Constantinescu E 2022 Optimal experimental design for inverse problems in the presence of observation correlations *SIAM J. Sci. Comput.* **44** A2808–42
- [9] Attia A, Leyffer S and Munson T S 2022 Stochastic learning approach for binary optimization: Application to Bayesian optimal design of experiments *SIAM J. Sci. Comput.* **44** B395–427
- [10] Babaniyi O, Nicholson R, Villa U and Petra N 2021 Inferring the basal sliding coefficient field for the Stokes ice sheet model under rheological uncertainty *Cryosphere* **15** 1731–50
- [11] Bartuska A, Espanh L and Tempone R 2022 Small-noise approximation for Bayesian optimal experimental design with nuisance uncertainty *Comput. Methods Appl. Mech. Engrg.* **399** 115320
- [12] Bui-Thanh T, Ghattas O, Martin J and Stadler G 2013 A computational framework for infinite-dimensional Bayesian inverse problems. Part I: the linearized case, with application to global seismic inversion *SIAM J. Sci. Comput.* **35** A2494–523
- [13] Chaloner K and Verdinelli I 1995 Bayesian experimental design: a review *Statist. Sci.* **10** 273–304
- [14] Constantinescu E M, Petra N, Bessac J and Petra C G 2020 Statistical treatment of inverse problems constrained by differential equations-based models with stochastic terms *SIAM/ASA J. Uncertain. Quantification* **8** 170–97
- [15] Da Prato G 2006 *An Introduction to Infinite-Dimensional Analysis* (Springer)
- [16] Daon Y and Stadler G 2018 Mitigating the influence of boundary conditions on covariance operators derived from elliptic PDEs *Inverse Probl. Imaging* **12** 1083–102
- [17] Dashti M and Stuart A M 2017 The Bayesian approach to inverse problems *Handbook of Uncertainty Quantification*, ed R Ghanem, D Higdon and H Owhadi (Springer) pp 311–428
- [18] Feng C and Marzouk Y M 2019 A layered multiple importance sampling scheme for focused optimal Bayesian experimental design (arXiv:1903.11187)
- [19] Golub G H and Van Loan C F 2013 Matrix computations *Johns Hopkins Studies in the Mathematical Sciences* 4th edn (Johns Hopkins University Press)
- [20] Haber E, Horesh L and Tenorio L 2008 Numerical methods for experimental design of large-scale linear ill-posed inverse problems *Inverse Problems* **24** 125–37

[21] Halko N, Martinsson P-G and Tropp J A 2011 Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions *SIAM Rev.* **53** 217–88

[22] Isaac T, Petra N, Stadler G and Ghattas O 2015 Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet *J. Comput. Phys.* **296** 348–68

[23] Jagalur-Mohan J and Marzouk Y M 2021 Batch greedy maximization of non-submodular functions: Guarantees and applications to experimental design *J. Mach. Learn. Res.* **22** 252

[24] Kaipio J and Kolehmainen V 2013 Approximate marginalization over modeling errors and uncertainties in inverse problems *Bayesian Theory and Applications* pp 644–72

[25] Kaipio J and Somersalo E 2005 *Statistical and Computational Inverse Problems (Applied Mathematical Sciences* vol 160) (Springer)

[26] Kaipio J and Somersalo E 2007 Statistical inverse problems: discretization, model reduction and inverse crimes *J. Comput. Appl. Math.* **198** 493–504

[27] Kolehmainen V, Tarvainen T, Arridge S R and Kaipio J P 2011 Marginalization of uninteresting distributed parameters in inverse problems-application to diffuse optical tomography *Int. J. Uncertain. Quantification* **1** 1–17

[28] Koval K, Alexanderian A and Stadler G 2020 Optimal experimental design under irreducible uncertainty for linear inverse problems governed by PDEs *Inverse Problems* **36** 075007

[29] Krause A, Singh A and Guestrin C 2008 Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies *J. Mach. Learn. Res.* **9** 235–84

[30] Li F 2019 A combinatorial approach to goal-oriented optimal Bayesian experimental design *Master's Thesis* Massachusetts Institute of Technology

[31] Liu S, Chepuri S P, Fardad M, Maşazade E, Leus G and Varshney P K 2016 Sensor selection for estimation with correlated measurement noise *IEEE Trans. Signal Process.* **64** 3509–22

[32] Mozumder M, Tarvainen T, Arridge S, Kaipio J P, D'Andrea C and Kolehmainen V 2016 Approximate marginalization of absorption and scattering in fluorescence diffuse optical tomography *Inverse Problems Imaging* **10** 227

[33] Nicholson R, Petra N and Kaipio J P 2018 Estimation of the Robin coefficient field in a Poisson problem with uncertain conductivity field *Inverse Problems* **34** 115005

[34] Nicholson R, Petra N, Villa U and Kaipio J P 2023 On global normal linear approximations for nonlinear Bayesian inverse problems *Inverse Problems* **39** 054001

[35] Pinski F J, Simpson G, Stuart A M and Weber H 2015 Kullback–Leibler approximation for probability measures on infinite dimensional spaces *SIAM J. Math. Anal.* **47** 4091–122

[36] Reese W, Hart J, van Bloemen Waanders B, Perego M, Jakeman J and Saibaba A 2024 Hyper-differential sensitivity analysis in the context of Bayesian inference applied to ice-sheet problems *Int. J. Uncertain. Quantification* **14** 1–20

[37] Shulkind G, Horesh L and Avron H 2018 Experimental design for nonparametric correction of misspecified dynamical models *SIAM/ASA J. Uncertain. Quantification* **6** 880–906

[38] Simpson M J, Baker R E, Buenzli P R, Nicholson R and Maclaren O J 2022 Reliable and efficient parameter estimation using approximate continuum limit descriptions of stochastic models *J. Theoret. Biol.* **549** 111201

[39] Smith R C 2013 *Uncertainty Quantification: Theory, Implementation and Applications (Computational Science and Engineering Series* vol 12) (SIAM)

[40] Stuart A M 2010 Inverse problems: a Bayesian perspective *Acta Numer.* **19** 451–559

[41] Sunseri I, Alexanderian A, Hart J and van Bloemen Waanders B 2024 Hyper-differential sensitivity analysis for nonlinear Bayesian inverse problems *Int. J. Uncertain. Quantification* **14** 1–20

[42] Sunseri I, Hart J, van Bloemen Waanders B and Alexanderian A 2020 Hyper-differential sensitivity analysis for inverse problems constrained by partial differential equations *Inverse Problems* **36** 125001

[43] Uciński D 2005 *Optimal Measurement Methods for Distributed Parameter System Identification* (CRC Press)

[44] Wu K, Chen P and Ghattas O 2023 A fast and scalable computational framework for large-scale high-dimensional Bayesian optimal experimental design *SIAM/ASA J. Uncertain. Quantification* **11** 235–61