

HYPER-DIFFERENTIAL SENSITIVITY ANALYSIS FOR NONLINEAR BAYESIAN INVERSE PROBLEMS

Isaac Sunseri,^{1,2,*} Alen Alexanderian,¹ Joseph Hart,³ & Bart van Bloemen Waanders³

¹North Carolina State University, Raleigh, North Carolina, 27607, USA

²Pennsylvania State University, State College, Pennsylvania, 16802, USA

³Sandia National Laboratories, Albuquerque, New Mexico, 87185, USA

*Address all correspondence to: Isaac Sunseri, North Carolina State University, Raleigh, North Carolina, 27607, USA, E-mail: ixs5268@psu.edu

Original Manuscript Submitted: 8/10/2022; Final Draft Received: 6/3/2023

We consider hyper-differential sensitivity analysis (HDSA) of nonlinear Bayesian inverse problems governed by partial differential equations (PDEs) with infinite-dimensional parameters. In previous works, HDSA has been used to assess the sensitivity of the solution of deterministic inverse problems to additional model uncertainties and also different types of measurement data. In the present work, we extend HDSA to the class of Bayesian inverse problems governed by PDEs. The focus is on assessing the sensitivity of certain key quantities derived from the posterior distribution. Specifically, we focus on analyzing the sensitivity of the MAP point and the Bayes risk and make full use of the information embedded in the Bayesian inverse problem. After establishing our mathematical framework for HDSA of Bayesian inverse problems, we present a detailed computational approach for computing the proposed HDSA indices. We examine the effectiveness of the proposed approach on an inverse problem governed by a PDE modeling heat conduction.

KEY WORDS: Bayesian inverse problems, post-optimality-sensitivity analysis, model uncertainty, design of experiments

1. INTRODUCTION

Many natural phenomena can be described by systems of partial differential equations (PDEs). The governing PDEs, however, often include parameters that are unknown and challenging to measure directly. This gives rise to inverse problems, in which one uses the PDE model and measurement data to estimate the unknown model parameters. In this article we consider Bayesian inverse problems [1,2], whose solution is a posterior distribution that is informed by our prior knowledge and the data measurements. Specifically, we focus on Bayesian inverse problems governed by PDEs with infinite-dimensional parameters.

In addition to the parameters being estimated, the governing PDEs typically contain parameters that are uncertain but needed for a full model specification. For clarity, we refer to the parameters being estimated by the inverse problem as *inversion parameters* and call the additional model parameters the *auxiliary parameters*. Another source of uncertainty in the inverse problem arises from the parameters specifying the experimental conditions, such as the locations of measurement devices or their accuracy. We call these the *experimental parameters*. Throughout the article we will refer to the union of auxiliary and experimental parameters as *complementary parameters*. Our goal in this article is to develop methods for assessing the sensitivity of the solution of a Bayesian inverse problem with respect to perturbations of complementary parameters; see Section 2 for a simple illustrative example.

Understanding the sensitivity of an inverse problem's solution to complementary parameters is important. These parameters may differ from their measured or estimated values, which in turn will result in a solution different from the one we would obtain if we had access to perfect measurements and true auxiliary parameters. Determining the sensitivity of the solution to perturbations in these parameters can inform our modeling assumptions and experimental

design practices. Specifically, this can guide a goal oriented prioritization of resources and focus efforts on obtaining accurate values for the important auxiliary parameters. Moreover, if one has data that are informative to the important auxiliary parameters, the inverse problem may be redesigned to include these parameters in the set of inversion parameters.

The present work builds on previous efforts in hyper-differential sensitivity analysis (HDSA) [3–12]. Traditional HDSA uses the derivative of the solution of an optimization problem with respect to complementary parameters to define sensitivity indices. These indices measure how much the solution of the optimization problem changes when the complementary parameters are perturbed. Specifically, in our previous work [3], which targets deterministic inverse problems, we define two types of sensitivity indices: pointwise sensitivities, and generalized sensitivities. Pointwise indices measure the sensitivity of the solution to perturbation in specific complementary parameters. Generalized indices measure the maximum possible change in the solution with respect to any unit perturbation of groups of complementary parameters. These indices provide a framework to study the sensitivity of the inverse problem solution; see [3] for more details.

In the present work, we extend HDSA to the class of Bayesian inverse problems. This allows us to study the change in the posterior distribution with respect to perturbations of the complementary parameters. We focus on HDSA of Bayesian inverse problems governed by PDEs with infinite-dimensional parameters. Section 3 provides a brief overview of the inverse problems under study. HDSA of a Bayesian inverse problem is difficult, because the solution of such problems is a statistical distribution. In general, the posterior distribution is difficult to approximate; this makes assessing the sensitivity of the posterior to complementary parameters challenging. A tractable approach is to instead focus on certain key aspects of the posterior distribution. Namely, we consider specific quantities of interest (QoIs) derived from the posterior distribution to perform HDSA on; we call such quantities the *HDSA QoIs*.

A first possibility, which we consider in this article, is to assess the sensitivity of the maximum *a posteriori* probability (MAP) point to the complementary parameters. This builds on the developments in [3]. Of greater difficulty is obtaining sensitivities of a measure of the posterior uncertainty. A natural setting for defining such measures is provided by the theory of optimal experimental design (OED) [13–18]. OED aims at finding experiments that minimize posterior uncertainty or, more generally, optimize the statistical quality of the estimated parameters. This is done by optimizing certain design criteria. Examples include the A-optimality criterion, which quantifies the average posterior variance, or the Bayesian D-optimality criterion, measuring the expected information gain; see, e.g., [16]. In the present work, we consider the Bayes risk, which has been used previously in OED for PDE-constrained inverse problems [19–21], as an HDSA QoI. Our motivations for considering the Bayes risk as an HDSA QoI are twofold. First, the Bayes risk is defined as an average error of the MAP estimator; see Section 4.1. Thus, HDSA of the Bayes risk builds further on methods for HDSA of the MAP point. Moreover, it is well-known [16,22] that the Bayes risk, with respect to the L^2 loss function, reduces to the A-optimality criterion, in the case of the Gaussian linear Bayesian inverse problem. Hence, up to a linearization, the Bayes risk may be considered as a proxy for the average posterior variance. Bayes risk is also a common utility function in decision theory.

In addition to the need for assessing the sensitivity of posterior uncertainty, there are other key differences between methods for HDSA of deterministic and Bayesian inverse problems. In the first place, in a deterministic formulation, some experimental measurements are needed before conducting sensitivity analysis. In contrast, our proposed framework does not require experimental measurements *a priori*, but uses the information encoded in the Bayesian inverse problem to generate likely data realizations. Moreover, HDSA of Bayesian inverse problems with respect to experimental parameters requires additional care. These parameters inform how data are collected and need to be accounted for properly in the data likelihood as well as the generated data samples. These issues are detailed throughout the article.

The contributions of this article are as follows:

- We develop a mathematical framework to assess the sensitivity of the Bayes risk and the MAP point in non-linear Bayesian inverse problems with respect to complementary parameters; see Section 4.
- We present a scalable computational framework for computing the HDSA indices for the MAP point and the Bayes risk; see Section 5. In that section, we also detail the computational cost of the various components of the proposed approach.

- We present comprehensive numerical results for a model problem of heat flow across a conductive surface that examine the effectiveness and efficiency of the proposed approach; see Section 6 for the description of the model under study and Section 7 for our computational results.

2. MOTIVATING EXAMPLE

We consider a simple example to motivate the problem considered in this work, which is to conduct sensitivity analysis on the solution of Bayesian inverse problems. Consider the heat equation,

$$\frac{\partial u(x, t)}{\partial t} = \exp(m) \frac{\partial^2 u(x, t)}{\partial x^2}, \quad x \in (0, \pi), t \in (0, 1), \quad (1a)$$

$$u(0, t) = u(\pi, t) = 0, \quad t \in (0, 1), \quad (1b)$$

$$u(x, 0) = \sin(x) + \exp(\theta) \sin(2x), \quad x \in (0, \pi). \quad (1c)$$

We focus on the inverse problem of estimating the inversion parameter m using measurements of u at the final time. In this problem, θ is an uncertain auxiliary parameter. For simplicity, we let m and θ be scalars here. Following a Bayesian framework, we endow m with a Gaussian prior $m \sim \mathcal{N}(1.3, 0.1)$. The problem (1) can be solved analytically, and the solution is given by

$$u(x, t) = \exp(-\exp(m)t) \sin(x) + \exp(-4\exp(m)t + \theta) \sin(2x). \quad (2)$$

We collect data measurements with additive Gaussian noise at the final time $t = 1$ at six evenly spaced points between $x = 0$ and $x = 2\pi$ depicted in Fig. 1. The Gaussian noise is unbiased with a standard deviation of 0.004 in this example. By Bayes rule, the posterior probability density function (pdf) is proportional to the product of the likelihood and prior pdfs. Because the auxiliary parameter is uncertain, we want to understand how the posterior distribution changes as the auxiliary parameter is perturbed. We view the posterior pdf of m solved at the nominal value of $\theta = 10$ and a perturbed value of $\theta = 11$ in Fig. 1.

We notice two primary changes in the posterior pdf as the auxiliary parameter is perturbed. First, there is a shift in the location of the distribution's peak, which equates to a change in the MAP point. Second, there is a change in the spread or variance of the distribution which equates to a changing of the posterior uncertainty. We can see that perturbations of auxiliary parameters can have significant impact upon both of these posterior quantities.




FIG. 1: Left: the solution of the PDE at the final time with noisy data measurements. Right: the prior pdf of m and its posterior pdf for both a nominal and perturbed value of the auxiliary parameter θ .

3. BAYESIAN INVERSE PROBLEMS AND COMPLEMENTARY PARAMETERS

We let θ_a and θ_e denote the auxiliary and experimental parameters, and call the augmented parameter vector $\theta = [\theta_a^\top \ \theta_e^\top]^\top$ the complementary parameters. Note that it is possible to have some auxiliary parameters that are functions (see [3]) but to keep the presentation simple, we consider finite-dimensional complementary parameters. The precise definition of the experimental parameters is, in general, application dependent, but in what follows we consider a particular case. Our goal is to provide a comprehensive framework for analyzing the sensitivity of the solution of the inverse problems under study to perturbations in θ .

We focus on Bayesian inverse problems governed by PDEs, with infinite-dimensional parameters. In our description of the computational methods, we will be working with discretized formulations. However, since the underlying inverse problem is formulated in infinite dimensions, we will begin by formulating the Bayesian inverse problem in a function space setting. This facilitates the discussion of the requirements on the prior distribution as well as the discretization of the Bayesian inverse problem.

We assume that the governing PDE (the state equation), represented abstractly by

$$v(u, m, \theta_a) = 0, \quad (3)$$

has a unique solution u for a given m and fixed auxiliary parameters θ_a . The inversion parameter m belongs to an infinite-dimensional Hilbert space $\mathcal{M}$ that is equipped with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{M}}$ and the induced norm $\| \cdot \|_{\mathcal{M}}$. The state variable u belongs to an infinite-dimensional reflexive Banach space U . In the present work, $\mathcal{M} = L^2(\Omega)$ where Ω is a suitable physical domain and $\langle \cdot, \cdot \rangle_{\mathcal{M}}$ is the standard L^2 inner product.

To infer m , we solve a Bayesian inverse problem that uses observed measurements along with information known about the governing system of PDEs. We assume that (noisy) measurement data are related to m according to the following model:

$$\mathbf{y} = \mathbf{F}(m, \theta_a) + \boldsymbol{\eta}(\theta_e), \quad (4)$$

where $\mathbf{y}$ is a vector of n_y experimental measurements, $\mathbf{F}$ the parameter-to-observable map that takes in the inversion parameter m and maps it to a vector of measurements, and $\boldsymbol{\eta}$ a vector that models additive Gaussian noise, $\boldsymbol{\eta} \sim \mathcal{N}(0, \mathbf{\Gamma}_{\text{noise}}(\theta_e))$. Evaluating $\mathbf{F}(m, \theta_a)$ requires solving the state equation (3) followed by application of an observation operator $\mathcal{O}$ which evaluates the state u at the n_y sensor locations. In the present work we let θ_e parameterize the noise levels of the sensors. This can correspond to situations where the experimental error at various sensors can be controlled either by repeated measurements or by choice of the measurement device, or possibly recalibration of existing devices. Our model for the experimental parameters is detailed in Section 6.

The Bayesian inverse problem setup. To solve an inverse problem with the data model (4) we define the data likelihood pdf $\pi_{\text{like}}(\mathbf{y}|m; \theta)$, which describes the distribution of data measurements $\mathbf{y}$, given a particular inversion parameter m . Given our assumption of an additive Gaussian noise model, we have $\mathbf{y}|m \sim \mathcal{N}(\mathbf{F}(m), \mathbf{\Gamma}_{\text{noise}})$ and, thus,

$$\pi_{\text{like}}(\mathbf{y}|m; \theta) \propto \exp\left(-\frac{1}{2}(\mathbf{F}(m, \theta_a) - \mathbf{y}(\theta_e))^\top \mathbf{\Gamma}_{\text{noise}}^{-1}(\theta_e)(\mathbf{F}(m, \theta_a) - \mathbf{y}(\theta_e))\right). \quad (5)$$

Note that in Eq. (5) the auxiliary parameters appear in the parameter-to-observable map $\mathbf{F}$ while we assume the experimental parameters for our problem appear only in the data measurements themselves and the noise covariance matrix.

In a Bayesian paradigm, we model our uncertainty regarding the inversion parameter by modeling m as a random variable. Accordingly, we endow the inversion parameter m with a prior distribution that reflects our knowledge of m *a priori*. In the present work, we let the prior distribution law be a Gaussian $\mu_{\text{pr}} = \mathcal{N}(m_{\text{pr}}, \mathcal{C}_{\text{pr}})$, with mean m_{pr} and covariance operator $\mathcal{C}_{\text{pr}}$. We let $\mathcal{C}_{\text{pr}} = \mathcal{A}^{-2}$ where $\mathcal{A}$ is a Laplace-like differential operator; see, e.g., [2, 18, 23]. The Gaussian prior measure is meaningful since $\mathcal{A}^{-2}$ is a trace class operator which guarantees bounded variance and almost surely pointwise well-defined samples. The prior measure induces the Cameron-Martin space $\mathcal{E} = \text{range}(\mathcal{C}_{\text{pr}}^{1/2})$, which is endowed with the following inner product:

$$\langle x, y \rangle_{\mathcal{E}} = \langle \mathcal{A}x, \mathcal{A}y \rangle. \quad (6)$$

We assume $m_{\text{pr}} \in \mathcal{E}$.

The definitions of the prior measure and the data likelihood complete the description of the Bayesian inverse problem. The solution of this inverse problem is the posterior measure $\mu_{\text{post}}^{\mathbf{y}}$, which describes the probability law of m conditioned on experimental measurements $\mathbf{y}$. We will often denote the posterior measure as μ_{post} for notational simplicity when no confusion arises from doing so. The Bayes formula takes the following form in the infinite-dimensional Hilbert space setting [2]:

$$\frac{d\mu_{\text{post}}}{d\mu_{\text{pr}}} \propto \pi_{\text{like}}(\mathbf{y}|m; \boldsymbol{\theta}). \quad (7)$$

Also, for a fixed $\boldsymbol{\theta}$, the maximum *a posteriori* probability (MAP) estimator of m is found by solving

$$m^*(\boldsymbol{\theta}) := \underset{m \in \mathcal{E}}{\operatorname{argmin}} J(m, \boldsymbol{\theta}), \quad (8)$$

where

$$J(m, \boldsymbol{\theta}) := \frac{1}{2} (\mathbf{F}(m, \boldsymbol{\theta}_a) - \mathbf{y}(\boldsymbol{\theta}_e))^\top \boldsymbol{\Gamma}_{\text{noise}}^{-1}(\boldsymbol{\theta}_e) (\mathbf{F}(m, \boldsymbol{\theta}_a) - \mathbf{y}(\boldsymbol{\theta}_e)) + \frac{1}{2} \langle m - m_{\text{pr}}, m - m_{\text{pr}} \rangle_{\mathcal{E}}. \quad (9)$$

Discretization. In the present work, we follow a continuous Galerkin finite element discretization and let $\mathbf{m}$ and $\mathbf{u}$ be the discretizations of their continuous counterparts m and u . We let n_m be the dimension of the discretized parameter. The discretized space is $\mathbb{R}^{n_m}$ equipped with the inner product,

$$\langle \mathbf{a}, \mathbf{b} \rangle_{\mathbf{M}} = \mathbf{a}^\top \mathbf{M} \mathbf{b}, \quad \mathbf{a}, \mathbf{b} \in \mathbb{R}^{n_m}, \quad (10)$$

where $\mathbf{M}$ is the finite element mass matrix, and the norm $\|\cdot\|_{\mathbf{M}}$ induced by this inner product. Note that when working with linear operators on $(\mathbb{R}^{n_m}, \|\cdot\|_{\mathbf{M}})$ or linear transformations between $(\mathbb{R}^{n_m}, \|\cdot\|_{\mathbf{M}})$ and $(\mathbb{R}^n, \|\cdot\|)$, where $\|\cdot\|$ is the Euclidean inner product, the adjoint operators need to be defined appropriately; see [23] for this and further details on discretization of different components of infinite-dimensional Bayesian inverse problems. In the remainder of this article, we present the proposed methods in the discretized setting.

4. HDSA FOR NONLINEAR BAYESIAN INVERSE PROBLEMS

In this section, we outline our framework for HDSA of nonlinear Bayesian inverse problems.

4.1 The HDSA QoIs

As discussed in the Introduction, we consider two HDSA QoIs for a Bayesian inverse problem: (i) the MAP point, which is obtained by minimizing Eq. (9), and (ii) the Bayes risk. For a fixed vector $\boldsymbol{\theta}$ of complementary parameters, the Bayes risk is defined by

$$\Psi_{\text{risk}}(\boldsymbol{\theta}) = \int_{\mathcal{M}} \int_{\mathbb{R}^{n_y}} \|\mathbf{m}^*(\boldsymbol{\theta}) - \mathbf{m}\|_{\mathbf{M}}^2 \pi_{\text{like}}(\mathbf{y}|\mathbf{m}; \boldsymbol{\theta}) d\mathbf{y} d\mu_{\text{pr}}^{n_m}(\mathbf{m}). \quad (11)$$

Note that here we have expressed the Bayes risk for the discretized version of the Bayesian inverse problem, and $\mathbf{m}^*$ is the discretized MAP point. The discretized prior measure, which we denote by $\mu_{\text{pr}}^{n_m}$, should be defined appropriately, as described in [23].

In practice, Bayes risk is approximated via sample averaging. Namely, we draw n_s samples $\{\mathbf{m}_1, \dots, \mathbf{m}_{n_s}\}$ from the prior distribution to compute data samples $\{\mathbf{y}_1, \dots, \mathbf{y}_{n_s}\}$ with the forward data model,

$$\mathbf{y}_i = \mathbf{F}(\mathbf{m}_i, \boldsymbol{\theta}_a) + \boldsymbol{\eta}_i(\boldsymbol{\theta}_e), \quad i = 1, \dots, n_s, \quad (12)$$

where $\boldsymbol{\eta}_i$ are draws from noise distribution $\mathcal{N}(0, \boldsymbol{\Gamma}_{\text{noise}}(\boldsymbol{\theta}_e))$. We can then write the approximate Bayes risk as

$$\widehat{\Psi}_{\text{risk}}(\boldsymbol{\theta}) = \frac{1}{n_s} \sum_{i=1}^{n_s} \|\mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}) - \mathbf{m}_i\|_{\mathbf{M}}^2. \quad (13)$$

4.2 Sensitivity Operator of Bayes Risk

We assess the sensitivity of Bayes risk by computing the partial derivative of $\widehat{\Psi}_{\text{risk}}(\boldsymbol{\theta})$ with respect to θ_j , the j th component of $\boldsymbol{\theta}$, at a set of nominal complementary parameter values $\boldsymbol{\theta}^*$:

$$D_j^{\text{R}} := \frac{\partial}{\partial \theta_j} \widehat{\Psi}_{\text{risk}}(\boldsymbol{\theta}^*) = \frac{2}{n_s} \sum_{i=1}^{n_s} \frac{\partial}{\partial \theta_j} (\mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}^*))^\top \mathbf{M} \mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}^*) - \frac{\partial}{\partial \theta_j} (\mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}^*))^\top \mathbf{M} \mathbf{m}_i. \quad (14)$$

We define the discretized sensitivity operator of the approximate Bayes risk as

$$\mathbf{D}^{\text{R}} = [D_1^{\text{R}} \quad D_2^{\text{R}} \quad \dots \quad D_{n_\theta}^{\text{R}}], \quad (15)$$

where n_θ denotes the dimension of the complementary parameter vector. Note that $\mathbf{D}^{\text{R}} \tilde{\boldsymbol{\theta}}$ can be interpreted as the sensitivity of the approximate Bayes risk with respect to a perturbation of the complementary parameters in the direction $\tilde{\boldsymbol{\theta}}$.

To compute the derivative of the approximate Bayes risk, we need $(\partial \mathbf{m}^*)/\partial \boldsymbol{\theta}(\mathbf{y}_i, \boldsymbol{\theta}^*)$, $i = 1, \dots, n_s$, which measure the sensitivity of the MAP points (for each data sample $\mathbf{y}_i$) to the complementary parameters.[†] For clarity, we denote the discretized cost functional in Eq. (9) by $\mathbf{J}$. As discussed in [3], under mild regularity assumptions [4,5,24] using the implicit function theorem, we obtain

$$\mathbf{D}^{\text{M}_i} = \frac{\partial \mathbf{m}^*}{\partial \boldsymbol{\theta}}(\mathbf{y}_i, \boldsymbol{\theta}^*) = -\mathbf{H}_i^{-1} \mathbf{B}_i, \quad (16)$$

where $\mathbf{H}_i = (\partial^2 \mathbf{J} / \partial \mathbf{m}^2)(\mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}^*), \boldsymbol{\theta}^*)$ and $\mathbf{B}_i = (\partial^2 \mathbf{J} / \partial \mathbf{m} \partial \boldsymbol{\theta})(\mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}^*), \boldsymbol{\theta}^*)$ are evaluated at the solution $\mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}^*)$ with fixed nominal parameters $\boldsymbol{\theta}^*$. By averaging these computed sensitivities over the number of data samples n_s , we can simultaneously measure both the average MAP point and Bayes risk sensitivities.

It is important to note the significance of this process. In a deterministic formulation [3], computing the sensitivities of the inverse problem solution requires data measurements. That is, some experimental measurements would be needed before conducting sensitivity analysis. In contrast, the method proposed here does not require experimental measurements and can be computed *a priori* by using the information encoded in the Bayesian inverse problem to generate likely data realizations. This makes the methodology applicable to a broad range of problems where data are not available at the time of performing HDSA.

4.3 Sensitivity Indices

Given a sensitivity operator, we define scalar sensitivity indices that measure the magnitude of the change in the solution with respect to a particular perturbation of the complementary parameters. We first group related complementary parameters together into K subsets. For example, we group data measurements corresponding to the same state variable together; scalar auxiliary parameters form their own group (of size 1), while all parameters defining the discretization of an uncertain function may form another group. Let Θ_k be the inner product space containing the k th set of parameters for $k = 1, \dots, K$ and let $\{\mathbf{b}_k^1, \mathbf{b}_k^2, \dots, \mathbf{b}_k^{n_k}\}$ be a basis for Θ_k of dimension n_k . We then define $\{\mathbf{e}_k^j\}$ as the basis of $\Theta = \Theta_1 \times \Theta_2 \times \dots \times \Theta_K$ for $k = 1, \dots, K$ and $j = 1, \dots, n_k$ where

$$\mathbf{e}_k^j = [0_1 \quad \dots \quad 0_{k-1} \quad \mathbf{b}_k^j \quad 0_{k+1} \quad \dots \quad 0_K]^\top. \quad (17)$$

We define pointwise sensitivity indices to measure the sensitivity of the individual MAP points, average MAP point sensitivities, and the sensitivity of the approximate Bayes risk, respectively, according to

$$S_{k,i}^j = \frac{\|\mathbf{D}^{\text{M}_i} \mathbf{e}_k^j\|_{\mathbf{M}}}{\|\mathbf{e}_k^j\|_{\Theta}}, \quad S_k^j = \frac{1}{n_s} \sum_{i=1}^{n_s} S_{k,i}^j, \quad \text{and} \quad \mathbb{S}_k^j = \frac{\|\mathbf{D}^{\text{R}} \mathbf{e}_k^j\|}{\|\mathbf{e}_k^j\|_{\Theta}}. \quad (18)$$

[†] As discussed further in Section 5, we only need to compute the action of this sensitivity operator on vectors.

Note that $S_{k,i}^j$ and $\mathbb{S}_k^j$, respectively, measure the change in $\mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}^*)$ and the Bayes risk to a perturbation of the k th parameter in the j th direction $\mathbf{b}_k^j$.

We would also like to determine the importance of the K parameter subgroups relative to one another. To do so, we define generalized sensitivity indices that provide a single measure of sensitivity for each parameter subgroup. Let $\mathbf{T}_k : \Theta \rightarrow \Theta$ be a selection operator that zeros out components of $\boldsymbol{\theta}$ not in Θ_k . We define the generalized sensitivity of the MAP points and approximate Bayes risk to the k th subgroup of complementary parameters according to

$$S_{k,i} = \max_{\boldsymbol{\theta} \in \Theta} \frac{\|\mathbf{D}^{\mathbf{M}_i} \mathbf{T}_k \boldsymbol{\theta}\|_{\mathbf{M}}}{\|\boldsymbol{\theta}\|_{\Theta}} \quad \text{and} \quad \mathbb{S}_k = \max_{\boldsymbol{\theta} \in \Theta} \frac{|\mathbf{D}^{\mathbf{R}} \mathbf{T}_k \boldsymbol{\theta}|}{\|\boldsymbol{\theta}\|_{\Theta}}, \quad (19)$$

respectively. We also define average generalized MAP point sensitivity indices by

$$S_k = \frac{1}{n_s} \sum_{i=1}^{n_s} S_{k,i}, \quad k = 1, \dots, K. \quad (20)$$

The generalized sensitivities in Eq. (19) measure the maximum change that can be observed in the HDSA QoIs to a norm-1 perturbation of the k th parameter subgroup. We can interpret this as a “worst case scenario” sensitivity because it measures the maximum change in the solution. More importantly, the generalized sensitivities provide a single measure of sensitivity for each parameter subgroup that can be used to compare their relative importance, despite their potentially diverse range of physical characteristics. Note that the parameter groupings should be specified by the user and are problem dependent. In the model problem considered in Section 6 we allow scalar auxiliary parameters to each consist of their own subgroup while the experimental parameters, corresponding to noise in the data measurements, are grouped together. It is important to note that if a subgroup consists of a single scalar parameter, its pointwise and generalized sensitivities will be identical. We direct the reader to [3] for additional details on the construction of these sensitivities.

To compare the MAP point and Bayes risk sensitivities it is important to note that each sensitivity is endowed with specific units. If we were only concerned with a single HDSA QoI, this would not matter because we would be primarily concerned with the relative differences between sensitivities of that measure. When comparing the sensitivities of the MAP point to Bayes risk, however, we must normalize with respect to the QoI to compare the sensitivities to each other reasonably. To do so, we divide the sensitivities with respect to the MAP point by the average norm of the computed MAP points, $(1/n_s) \sum_{i=1}^{n_s} \|\mathbf{m}^*(\mathbf{y}_i, \boldsymbol{\theta}^*)\|_{\mathbf{M}}$, and the sensitivities with respect to Bayes risk by the computed value of Bayes risk, $\hat{\Psi}_{\text{risk}}(\boldsymbol{\theta}^*)$.

4.4 An Illustrative Example

Here, we present an illustrative example to demonstrate how our sensitivity analysis framework can be used practically to improve the posterior distribution.

Consider a function $u : [0, 1] \rightarrow \mathbb{R}$ defined as

$$u(x) = \theta_1 x m^3 + (0.1) \theta_2. \quad (21)$$

Here, $x \in [0, 1]$ is the spatial coordinate, $m \in \mathbb{R}$ is the inversion parameter, and $\boldsymbol{\theta} = [\theta_1 \ \theta_2]^\top$ is the vector of auxiliary parameters. We let the “ground-truth” values of m and $\boldsymbol{\theta}$ be $m_{\text{true}} = 2$ and $\boldsymbol{\theta}_{\text{true}} = [1 \ 1]^\top$, respectively. To simulate what happens in practical computations, rather than using the truth $\boldsymbol{\theta}$, we assume a nominal value of $\boldsymbol{\theta}^* = [1.4 \ 1.4]^\top$ for the auxiliary parameter vector. We assume that the value of u is observed at two spatial locations $x = 0.5$ and $x = 1.0$. Using these data, we formulate a Bayesian inverse problem for m with a mean zero normal prior distribution whose variance is 5, and use a pointwise noise variance of 0.25.

Before collecting data, we analyze the sensitivity of the inverse problem by computing the average MAP point sensitivity and the Bayes risk sensitivity using $n_s = 10,000$ samples from the prior. This yields sensitivity indices

$$S_1 = 0.43, \quad S_2 = 0.05, \quad \mathbb{S}_1 = 0.03, \quad \text{and} \quad \mathbb{S}_2 = 0.0001.$$

Observe that both the MAP point and Bayes risk sensitivities are much larger for θ_1 compared to θ_2 . This is consistent with intuition based on Eq. (21) since $\theta_1 x m^3 = \mathcal{O}(1)$ and $(0.1)\theta_2 = \mathcal{O}(10^{-1})$. Furthermore, θ_1 scales the influence of m on u , whereas θ_2 shifts the value of y independent of m .

Next, we assume that (noisy) data are generated by evaluating $u(x, m_{\text{true}}, \theta_{\text{true}})$ at the observation locations. We solve the inverse problem using these data and the nominal auxiliary parameter vector $\theta^* = [1.4 \ 1.4]^\top$. Based on the sensitivity information computed prior to observing data, we conclude that it is advantageous to improve our estimate of θ_1 . Thus, we also solve the inverse problem with an updated value of $\theta_1 = 1$; i.e., we consider the auxiliary parameter vector $\theta^{(1)} = [1 \ 1.4]^\top$. For comparison, we also solve the inverse problem by updating the “unimportant” auxiliary parameter θ_2 , where we use $\theta^{(2)} = [1.4 \ 1]^\top$. Figure 2 displays the three posterior PDFs.

We observe that using the nominal auxiliary parameters yields a poor estimate of m with the true value lying outside the posterior distribution’s support. Correcting θ_2 yields a negligible impact. In contrast, correcting the “important” auxiliary parameter θ_1 significantly improves the posterior distribution.

5. COMPUTATIONAL METHODS

In this section, we present computational methods to implement the framework proposed in Section 4.

5.1 Computing the Sensitivity Indices

We can write the sensitivity operator of the approximate Bayes risk with respect to the complementary parameters (15) as follows:

$$\mathbf{D}^R = \frac{2}{n_s} \sum_{i=1}^{n_s} (\mathbf{D}^{M_i})^\top \mathbf{M}(\mathbf{m}^*(\mathbf{y}_i, \theta^*) - \mathbf{m}_i) = \frac{2}{n_s} \sum_{i=1}^{n_s} -\mathbf{B}_i^\top \mathbf{H}_i^{-\top} \mathbf{M}(\mathbf{m}^*(\mathbf{y}_i, \theta^*) - \mathbf{m}_i). \quad (22)$$

Note that, as before, the subscript i on the operators $\mathbf{D}^{M_i} \in \mathbb{R}^{n_m \times n_\theta}$, $\mathbf{B}_i \in \mathbb{R}^{n_m \times n_\theta}$, and $\mathbf{H}_i \in \mathbb{R}^{n_m \times n_m}$ indicates the dependence on the i th data sample [see Eq. (16)].

To compute matrix-free actions of $\mathbf{H}$ and $\mathbf{B}$ to vectors, we use a discretized formal Lagrangian approach. We note that this method is utilized to both compute sensitivity indices and solve for the MAP point. Note also that, for notational convenience, we suppress the index i in the Hessian and mixed derivative operators when describing the adjoint based expressions describing their application to vectors. We begin by defining the discrete Lagrangian as




FIG. 2: Posterior PDFs corresponding to solving the inverse problem with nominal auxiliary parameters, $\theta^* = [1.4 \ 1.4]^\top$, an update of θ_1 , $\theta^{(1)} = [1 \ 1.4]^\top$, and an update of θ_2 , $\theta^{(2)} = [1.4 \ 1]^\top$. The true parameter value $m_{\text{true}} = 2$ is indicated by the vertical line.

$$\mathcal{L}(\mathbf{u}, \mathbf{m}, \mathbf{p}; \boldsymbol{\theta}) = J(\mathbf{m}, \boldsymbol{\theta}) - \langle \mathbf{p}, v(\mathbf{u}, \mathbf{m}, \boldsymbol{\theta}_a) \rangle_{\mathbf{M}}, \quad (23)$$

where $v(\mathbf{u}, \mathbf{m}, \boldsymbol{\theta}_a)$ is the discretized form of the PDE v , and $\mathbf{p}$ is the adjoint variable. Next, we use variational derivatives to compute the action of the discretized gradient of the cost function. We let $\mathcal{L}_{\mathbf{p}}[\hat{\mathbf{p}}]$ denote the variational derivative of Eq. (23) with respect to $\mathbf{p}$, acting on $\hat{\mathbf{p}}$, with the input arguments suppressed for brevity. A similar notation is used for the variational derivatives with respect to $\mathbf{u}$ and $\mathbf{m}$. We can also compute the action of the Hessian by constructing a meta-Lagrangian,

$$\mathcal{L}^H(\mathbf{u}, \mathbf{m}, \mathbf{p}, \hat{\mathbf{u}}, \hat{\mathbf{m}}, \hat{\mathbf{p}}; \boldsymbol{\theta}) = \mathcal{L}_{\mathbf{p}}[\hat{\mathbf{p}}] + \mathcal{L}_{\mathbf{u}}[\hat{\mathbf{u}}] + \mathcal{L}_{\mathbf{m}}[\hat{\mathbf{m}}]. \quad (24)$$

By computing variational derivatives of the meta-Lagrangian, we can evaluate the action of the discretized Hessian to vectors. The basic steps of this solution process are outlined in Algorithm 1, and we direct the reader to [25,26] for additional details.

Next, we discuss computing the action of the mixed derivative operator $\mathbf{B}^\top$. We follow a similar approach as one used to compute the action of the Hessian using the meta-Lagrangian $\mathcal{L}^H$. Namely, we differentiate the meta-Lagrangian with respect to $\boldsymbol{\theta}$ to obtain

$$\mathcal{L}_{\boldsymbol{\theta}}^H(\mathbf{u}, \mathbf{m}, \mathbf{p}, \hat{\mathbf{u}}, \hat{\mathbf{m}}, \hat{\mathbf{p}}; \boldsymbol{\theta})[\tilde{\boldsymbol{\theta}}] = \tilde{\boldsymbol{\theta}}^\top \mathbf{B}^\top \hat{\mathbf{m}}, \quad (25)$$

where $\hat{\mathbf{u}}$ and $\hat{\mathbf{p}}$ satisfy the incremental state and adjoint equations, respectively.

Note that we can also compute the action of $\mathbf{B}$ by reversing the order of differentiation, deriving through the Lagrangian by $\boldsymbol{\theta}$ and the meta-Lagrangian by $\mathbf{m}$, which will result in modified incremental equations. These adjoint based methods provide a computationally efficient method to evaluate the sensitivity operators $\mathbf{D}^M$ and $\mathbf{D}^R$.

To compute the discretized sensitivity operator $\mathbf{D}^R$, we must first generate data samples $\mathbf{y}_i$ for $i = 1, \dots, n_s$ and then evaluate Eq. (22), which involves a nontrivial computational cost. We also compute sensitivities of the MAP point, efficiently reusing PDE solves whenever applicable. This process is summarized in Algorithm 2. Note that for clarity, we have separated the processes of data generation and sensitivity operator computation in the algorithm. Recall that the second subscript in sensitivity indices $S_{k,i}$ denotes dependence of the index upon the i th data sample.

5.2 Computational Costs

Here we discuss the areas of high computational cost in Algorithm 2. To gain computational efficiency, we rely on the following key tools from PDE-constrained optimization: inexact Newton-CG for MAP estimation, adjoint methods gradient and Hessian computation, and low-rank approximations for efficient computation of inverse Hessian applies [23,27–29]. By combining methods that make maximum use of the problem structure, we ensure that the computational complexity of our approach, in the terms of the number of PDE solves, does not scale with the dimension of the discretized inversion parameter.

Generating data samples. We solve the forward problem n_s times and use the resulting solutions to generate data.

MAP point solves. We solve the inverse problem n_s times (line 9) using an inexact Newton-CG method with backtracking Armijo line search. Each Newton step requires 2 PDE solves to compute the gradient and an additional $2I$ PDE solves to compute the Hessian apply where I is the number of iterations required by the CG solver to find

Algorithm 1: Compute the gradient $g(\mathbf{m})$ and action of the Hessian $\mathbf{H}$ in the direction $\hat{\mathbf{m}}$

- 1: Solve the state equation $\mathcal{L}_{\mathbf{p}} = 0$ for the state variable $\mathbf{u}$
 - 2: Solve the adjoint equation $\mathcal{L}_{\mathbf{u}} = 0$ for the adjoint variable $\mathbf{p}$
 - 3: Evaluate $g(\mathbf{m})^\top = \mathcal{L}_{\mathbf{m}}$
 - 4: Solve the incremental state equation $\mathcal{L}_{\mathbf{p}}^H = 0$ for the incremental state variable $\hat{\mathbf{u}}$
 - 5: Solve the incremental adjoint equation $\mathcal{L}_{\mathbf{u}}^H = 0$ for the incremental adjoint variable $\hat{\mathbf{p}}$
 - 6: Evaluate the Hessian apply $\mathbf{H}(\mathbf{m})[\hat{\mathbf{m}}] = \mathcal{L}_{\mathbf{m}}^H$
-

Algorithm 2: Compute the sensitivity indices

```

1: % Data sample generation
2: for  $i = 1$  to  $n_s$  do
3:   Draw prior sample  $\mathbf{m}_i$ 
4:   Solve the forward equation  $\mathbf{v}(\mathbf{u}_i, \mathbf{m}_i, \boldsymbol{\theta}_a^*) = 0$  for  $\mathbf{u}_i$ 
5:   Synthesize data samples  $\mathbf{y}_i = \mathcal{O}\mathbf{u}_i + \boldsymbol{\eta}_i(\boldsymbol{\theta}_e^*)$  { $\mathcal{O}$  observes  $\mathbf{u}$  at measurement locations}
6: end for

7: % Computation of the Bayes risk sensitivities
8: for  $i = 1$  to  $n_s$  do
9:   Solve the discretized inverse problem for  $\mathbf{m}_i^*(\mathbf{d}_i, \boldsymbol{\theta}^*)$ 
10:  Solve  $-\mathbf{H}_i \mathbf{z}_i = \mathbf{M}(\mathbf{m}_i^* - \mathbf{m}_i)$  for  $\mathbf{z}_i$ 
11:  Compute  $\mathbf{r}_i = \mathbf{B}_i^\top \mathbf{z}_i$ 
12: end for
13:  $\mathbf{D}^R = (2/n_s) \sum_{i=1}^{n_s} \mathbf{r}_i$ 
14: Compute  $\mathbb{S}_k^j$  and  $\mathbb{S}_k$  for all  $k = 1, \dots, K$  and  $j = 1, \dots, n_\theta$ , see Eqs. (18) and (19)

15: % Computation of the average MAP point sensitivities
16: for  $i = 1$  to  $n_s$  do
17:   for  $j = 1$  to  $n_\theta$  do
18:     Compute  $S_{k,i}^j$  for  $k = 1, \dots, K$ , see Eq. (18)
19:   end for
20:   Compute  $S_{k,i}$  for  $k = 1, \dots, K$ , see Eq. (19)
21: end for
22: Compute averaged pointwise sensitivities  $S_k^j = (1/n_s) \sum_{i=1}^{n_s} S_{k,i}^j$ , see Eq. (18)
23: Compute averaged generalized sensitivities  $S_k = (1/n_s) \sum_{i=1}^{n_s} S_{k,i}$ , see Eq. (20)

```

an appropriate search direction. Letting I denote a bound on the number of CG iterations over L (outer) Newton iterations, the total cost is $2L + \mathcal{O}(2LI)$ PDE solves. This cost in PDE solves multiplied by the number of samples n_s becomes quite significant. However, since the samples drawn from the prior are independent of each other, these computations can be performed in parallel. We also note that we initialize the MAP point solves with the prior samples used to generate data samples.

Evaluating inverse Hessian applies. We now address the problem of repeated application of the inverse Hessian, which is required to compute both Bayes risk and MAP point sensitivities in lines 10, 18, and 20. We note that if one only wishes to compute Bayes risk sensitivities, this will not require repeated use of the same Hessian inverse, and line 10 can be evaluated with (preconditioned) CG. Assuming that MAP point sensitivities are also desired, we can offset this cost by utilizing a low-rank approximation for the Hessian of the data-misfit term in definition of the cost function $\mathbf{J}$. Specifically, the Hessian $\mathbf{H}_i$ can be written as $\mathbf{H}_i = \mathbf{H}_{\text{misfit}}^{(i)} + \boldsymbol{\Gamma}_{\text{pr}}^{-1}$ where $\mathbf{H}_{\text{misfit}}^{(i)}$ is the data-misfit Hessian and $\boldsymbol{\Gamma}_{\text{pr}}$ is the discretized prior covariance operator. In ill-posed inverse problems, the prior-preconditioned data-misfit Hessian $\boldsymbol{\Gamma}_{\text{pr}}^{1/2} \mathbf{H}_{\text{misfit}}^{(i)} \boldsymbol{\Gamma}_{\text{pr}}^{1/2}$ often admits a low-rank approximation. As detailed in [23], this low-rank approximation, which can be computed efficiently using the Lanczos method, enables fast Hessian inverse applies. After computing this low-rank approximation, application of the Hessian inverse can be approximated by matrix-vector products. The computational cost of this process is $\mathcal{O}(2r) + 2$ PDE solves, where r is the rank of the desired approximation.

Computing Bayes risk sensitivities. The sensitivity operator of Bayes risk is a vector, so we built this operator directly before computing indices. We begin this discussion by noting that we can solve the state and adjoint equations around the MAP point once for each data sample, and reuse these solves for each Hessian $\mathbf{H}_i$ and mixed derivative operator $\mathbf{B}_i$ or $\mathbf{B}_i^\top$ apply. Each Hessian apply requires two additional PDE solves (in addition to the forward and adjoint solves) for the incremental state and incremental adjoint equations. These incremental equation solves can be

reused to compute the application of $\mathbf{B}_i^\top$, while $\mathbf{B}_i$ apply requires two more PDE solves for the modified incremental equations.

Computing MAP point sensitivities. The greatest computational cost in estimating the MAP point sensitivities comes in the repeated application of $\mathbf{D}^{\mathbf{M}_i}$ to standard basis vectors e_j (line 18) to compute n_θ pointwise sensitivity indices for all n_s data samples. As mentioned previously, this cost is significantly reduced by precomputing a low-rank approximation that allows for fast Hessian inverse applications. It is also important to note that we reuse the inverse problem solves from computing the Bayes risk sensitivities in computing the MAP point sensitivities and we do not require any additional inverse problem solves here. Due to these various computational savings, we can estimate the MAP point sensitivities through sample averaging at a significantly reduced cost.

The discussed computational costs are summarized in Table 1 for clarity. We remark that for the problem considered in the present work n_θ is not very large. For problems with a large number of complementary parameters, computing a suitable low-rank approximation of $\mathbf{B}$ may be helpful to reduce the cost of computing many MAP point sensitivities. We plan to investigate this in our future work.

6. MODEL PROBLEM

In this section we consider a model inverse problem, involving heat flow across a conductive surface, that will be used to study our HDSA framework. We begin by describing the forward problem in Section 6.1 followed by the setup of the Bayesian inverse problem in Section 6.2.

6.1 Forward Model

Consider the problem of inferring the log-conductivity field of a medium from measurements of temperature. Focusing on a cross section, we consider the problem in two space dimensions. The forward problem is governed by the following elliptic PDE, modeling steady state heat conduction on a unit square domain Ω with boundary $\partial\Omega = \cup_{i=1}^4 \Gamma_i$, where $\Gamma_1, \Gamma_2, \Gamma_3$, and Γ_4 denote the bottom, right, top, and left edges of Ω respectively,

$$-\nabla \cdot (e^m \nabla u) = f \quad \text{in } \Omega, \quad (26a)$$

$$e^m \nabla u \cdot n = 0 \quad \text{on } \Gamma_1 \cup \Gamma_3, \quad (26b)$$

$$e^m \nabla u \cdot n = \beta(T_{\text{amb}} - u) \quad \text{on } \Gamma_2, \quad (26c)$$

$$e^m \nabla u \cdot n = s \quad \text{on } \Gamma_4. \quad (26d)$$

In this model, the inversion parameter $m(\mathbf{x})$ is a function representing the log of the heat conductivity of the non-homogeneous two-dimensional surface. We let $u(\mathbf{x})$ denote the temperature, $f(\mathbf{x})$ the heat source in the domain, β the heat transfer coefficient of the medium, T_{amb} the ambient temperature of the medium, and $s(x_2)$ a boundary heat source function representing heat entering the domain from the left boundary. In this model problem, Eq. (26) are dimensionless and we let $T_{\text{amb}} = 22$ and consider the heat transfer coefficient β to be an uncertain auxiliary parameter with a nominal value of $\beta = 1$.

The boundary heat source $s(x_2)$ is modeled as follows:

TABLE 1: Computational costs summary

Computation	Significant cost per sample (n_s)
Data generation	1 PDE solve
Inverse problem solves	$2L + \mathcal{O}(2LI)$ PDE solves for L Newton steps each having at most I CG iterations
Hessian inverse approximation	$\mathcal{O}(2r) + 2$ PDE solves where r is the rank of the desired approximation
Bayes risk sensitivities	2 PDE solves
MAP point sensitivities	$2n_\theta$ PDE solves

$$s(x_2) = s_1 \exp\left(-\left(\frac{x_2 - s_3}{s_2}\right)^2\right), \quad (27)$$

with auxiliary parameters s_1, s_2 , and s_3 fixed at nominal values $s_1 = 30, s_2 = 0.1$, and $s_3 = 0.65$. The auxiliary parameters consist of the amplitude, spread, and location of the boundary heat source, respectively. The heat source in the domain $f(\mathbf{x})$ is modeled as

$$f(\mathbf{x}) = f_1 \exp\left[-\frac{1}{2}(\mathbf{x} - \mathbf{w})^\top \mathbf{C}_1(\mathbf{x} - \mathbf{w})\right] + f_2 \exp\left[-\frac{1}{2}(\mathbf{x} - \mathbf{z})^\top \mathbf{C}_2(\mathbf{x} - \mathbf{z})\right], \quad (28)$$

with

$$\mathbf{C}_1 = \begin{bmatrix} \frac{\cos^2(\gamma_1)}{\sigma_{x_1}^2} + \frac{\sin^2(\gamma_1)}{\sigma_{x_2}^2} & \frac{\sin(2\gamma_1)}{2\sigma_{x_2}^2} - \frac{\sin(2\gamma_1)}{2\sigma_{x_1}^2} \\ \frac{\sin(2\gamma_1)}{2\sigma_{x_2}^2} - \frac{\sin(2\gamma_1)}{2\sigma_{x_1}^2} & \frac{\sin^2(\gamma_1)}{\sigma_{x_1}^2} + \frac{\cos^2(\gamma_1)}{\sigma_{x_2}^2} \end{bmatrix} \text{ and } \mathbf{C}_2 = \begin{bmatrix} \frac{\cos^2(\gamma_2)}{\sigma_{x_1}^2} + \frac{\sin^2(\gamma_2)}{\sigma_{x_2}^2} & \frac{\sin(2\gamma_2)}{2\sigma_{x_2}^2} - \frac{\sin(2\gamma_2)}{2\sigma_{x_1}^2} \\ \frac{\sin(2\gamma_2)}{2\sigma_{x_2}^2} - \frac{\sin(2\gamma_2)}{2\sigma_{x_1}^2} & \frac{\sin^2(\gamma_2)}{\sigma_{x_1}^2} + \frac{\cos^2(\gamma_2)}{\sigma_{x_2}^2} \end{bmatrix}. \quad (29)$$

In this formulation, f_1 and f_2 control the amplitude of the heat sources, $\mathbf{w}$ and $\mathbf{z}$ control the centers of the two sources, γ_1 and γ_2 their respective tilt angles, and σ_{x_1} and σ_{x_2} the spread of the heat sources in the x_1 and x_2 directions, respectively. For this problem we fix these parameters at the following nominal values: $f_1 = 100, f_2 = 105, \mathbf{w} = (0.8, 0.25), \mathbf{z} = (0.5, 0.8), \gamma_1 = -\pi/4, \gamma_2 = 0.15, \sigma_{x_1} = 0.8, \sigma_{x_2} = 0.1$. We consider the amplitude, center point, and angle of each bar to be uncertain and thus let $f_1, f_2, w_1, w_2, z_1, z_2, \gamma_1$, and γ_2 be the auxiliary parameters for the right-hand side heat source $f(\mathbf{x})$. Figure 3 depicts this heat source in the domain.

We note that this model problem has been kept intentionally simple to aid in the interpretation and understanding of the complicated algorithmic methodology. Even so, this example is motivated by many uncertainties surrounding additive manufacturing processes (such as powered bed laser fusion) that cause high residual stresses and even defects in the final parts. Variability in the powder material, boundary conditions, rasterization patterns, and laser power result in uneven heat distribution with problematic microcrystallographic structures and inhomogeneous material properties. Although the underlying physics for additive manufacturing is more complicated, our model problem conceptually demonstrates the ability of our approach to provide insight into a complicated application area.

6.2 Prior Measure and State Solution

In many inverse problems a “true solution” is chosen to synthesize data and evaluate the accuracy of the proposed methodology. We do not have any such “true solution” here and instead we compute data from samples of the prior




FIG. 3: The heat source function $f(\mathbf{x})$

distribution. We specify the Bayesian prior on m as a Gaussian random field on Ω with mean m_{pr} and covariance operator $\mathcal{C}_{\text{pr}}$. We model the prior mean as a sinusoidal function:

$$m_{\text{pr}}(\mathbf{x}) = 1.5 \sin(2\pi x_1) \cos(2\pi x_2) + 2. \quad (30)$$

Furthermore, we let the covariance operator $\mathcal{C}_{\text{pr}}$ be the inverse of a squared elliptic differential operator $\mathcal{A}$, where $m = \mathcal{A}^{-1}s$ satisfies

$$\alpha \int_{\Omega} (\Phi \nabla m) \cdot \nabla q + m q \, d\mathbf{x} = \int_{\Omega} s q \, d\mathbf{x}, \quad (31)$$

for all $q \in H^1(\Omega)$, with $\alpha = 5$, and $\Phi = 0.01$. This formulation of the prior covariance ensures that $\mathcal{C}_{\text{pr}}$ is trace class and provides a computationally convenient formulation. For more details see [23].

Measurements are collected on an evenly spaced 5×5 grid of observation locations depicted in Fig. 4. We consider the standard deviation of the noise in each data measurement to be our uncertain experimental parameters. Additive Gaussian noise models “error” in our data and we assume the measurements are uncorrelated, with nominal standard deviations of $\sigma = 0.1$; thus $\mathbf{\Gamma}_{\text{noise}}(\theta_e) = \text{diag}(\sigma_1(\theta_{e,1})^2, \dots, \sigma_{n_y}(\theta_{e,n_y})^2)$. Although we allow the measurement standard deviations to take the same nominal value, we consider each standard deviation individually when computing sensitivities of the solution.

Perturbing the noise standard deviation will also result in a perturbation of the noise realization $\eta_i \sim \mathcal{N}(0, \mathbf{\Gamma}_{\text{noise}})$, directly proportional to the multiplicative perturbation of σ_i . Therefore, the experimental parameters θ_e enter the inverse problem through the cost function (9), both in the noise covariance matrix $\mathbf{\Gamma}_{\text{noise}}(\theta_e)$ and the data measurements $y(\theta_e)$ which depend on the noise realizations. The solution of the governing PDE system detailed in Eq. (26) at the nominal parameter values with m fixed at the prior mean is depicted in Fig. 4.

7. RESULTS

We focus on the model Bayesian inverse problem described in Section 6. Following Algorithm 2 to evaluate our Bayesian hyper-differential sensitivities, we take samples from the prior distribution on m and push them through the forward mapping to generate noisy data. Each data sample is then used to solve Eq. (8), giving a unique MAP point reconstruction for each sample. To illustrate this process, we present three prior samples and their corresponding MAP point reconstructions in Fig. 5.

Each MAP point is attempting to estimate the above prior sample from noisy data. This example is illustrative in that it gives us some insight into Bayes risk, which measures the average difference in norm between the prior samples (top) and the inferred MAP points (bottom).




FIG. 4: The state solution of the governing system of partial differential equations with the experimental sensor locations indicated by filled black circles




FIG. 5: Top: three samples from the log-conductivity prior. Bottom: inverse problem MAP point estimates solved using data generated by the above prior sample.

In Section 7.1 we detail how perturbations of the complementary parameters are modeled. Following this we present and discuss the significance of the generalized sensitivities of the complementary parameters as well as the pointwise sensitivities of the experimental parameters with respect to Bayes risk (Section 7.2) and the MAP point (Section 7.3). We note that the pointwise sensitivities of the auxiliary parameters are identical to their generalized sensitivities as each auxiliary parameter is scalar valued in this model problem.

7.1 Modeling Parameter Perturbations

Suppose ρ is an uncertain scalar parameter. We model our uncertainty in ρ as

$$\rho = \tilde{\rho}(1 + a\theta), \quad (32)$$

where $\tilde{\rho}$ is the nominal value, a is a scaling coefficient quantifying our degree of uncertainty, and $\theta \in [-1, 1]$ defines a perturbation of $\tilde{\rho}$. Perturbations of vector valued complementary parameters, such as data measurements, are modeled as componentwise scalar perturbations as in Eq. (32).

In this particular model problem we use a perturbation scaling coefficient of $a = 0.05$ for each auxiliary parameter, which represents our uncertainty in that parameter's estimate being 5% of the parameter's nominal value. For the experimental parameters we instead use a scaling coefficient of $a = 1$ to represent that our uncertainty in the standard deviation of the data noise is the full quantity of the standard deviation.

7.2 Sensitivities of Bayes Risk

The approximate Bayes risk is computed via sample averaging as detailed in Eq. (13). We present the generalized sensitivities of each complementary parameter with respect to Bayes risk in Fig. 6. We study the effect of the sample size on the computed sensitivities by comparing generalized sensitivities for Bayes risk computed from ten groups of 20 samples, ten groups of 100 samples, and ten groups of 500 samples, each taken randomly from a group of 3000 precomputed samples.

First let us discuss the spread of the samples.




FIG. 6: Generalized sensitivities of Bayes risk to complementary parameters computed for ten groups of samples of size 20, 100, and 500

We can see that while the groups of 100 samples can sometimes vary significantly, such as in the case of γ_2 , they generally capture the rankings of the parameters relative to one another correctly. Thus, if we are primarily concerned with determining the relative importance of the parameters compared to each other we may conclude that 100 samples provide sensitivity estimates that suit our needs.

Next, we note that Bayes risk is most sensitive to the tilt angle of the second domain heat source γ_2 . We observe that as the tilt angles γ_1 and γ_2 are changed, they can overlap in the domain interior, causing a large increase in heat where the overlap occurs and will result in a significant change in f . One possible reason γ_2 is so important is that even a relatively small perturbation will result in increased or decreased overlap of these bars in the domain. Of secondary importance are the heat amplitude (f_2) and center in the x_1 direction (z_1) of the second domain heat source, heat transfer coefficient (β), and the standard deviations of data noise (σ). This sensitivity information can then be used by an experimenter to inform their experimental design choices for this problem. To accurately estimate the Bayes risk for this problem as a measure of posterior uncertainty, it is more important to invest resources in ensuring that the parameters γ_2 , f_2 , z_1 , β , and σ are more accurately estimated than the other complementary parameters. Specifically, we can interpret these sensitivities as “a 5% perturbation in the scalar auxiliary parameters or a norm-1 perturbation in the experimental parameters (σ) will result in a perturbation of Bayes risk proportional to the sensitivity.”

While these sensitivities appear to be very small, we note that the problem is highly diffusive and steady state. Both of these factors are likely making the problem highly insensitive to perturbations of complementary parameters. This itself showcases the benefits of using HDSA. For such an insensitive problem, it would be extremely difficult to gather any kind of intuition or conclusion as to the relative importance of various parameters *a priori*. With our framework however, we can rigorously determine the relative importance of uncertain parameters before any physical experimentation is done, even for highly insensitive problems, which is valuable to experimenters who seek to efficiently allocate experimental resources.

Next we study the pointwise sensitivities of Bayes risk to the experimental parameters, the standard deviation of noise in the data measurements, presented in Fig. 7. By perturbing the noise, we model perturbations of each collected data measurement in a way that we can experimentally control through sensor accuracy.

We first note the scale of sensitivities presented in Fig. 7. Although these sensitivities are very small (on the order of 10^{-5} or 10^{-6}), this is not entirely unexpected given the scale of the generalized sensitivities presented in Fig. 6.




FIG. 7: Pointwise sensitivities of Bayes risk to experimental parameters computed from 1000 samples. Each bar represents the sensitivity of perturbing the standard deviation of the data noise at that particular sensor location.

Indeed, we would expect that perturbing a single data measurement's noise would not result in a very large change in Bayes risk. We can see that the sensors grouped around small values of x_1 and large values of x_2 are most important with respect to Bayes risk. Thus, we can conclude that the data measured at these sensors is the most important to collect accurately for the purposes of estimating our measure of posterior uncertainty. We observe that these sensors are located in the region where the state solution depicted in Fig. 4 is largest. This is also the region near the boundary source term s . We also note that the sensors at (0.9, 0.1) and (0.9, 0.9) are relatively important, which are located in the areas where the state solution is smallest. These results provide information that may not be obvious *a priori* and helps practitioners understand what parameters and sensor measurements the solution is most sensitive to.

7.3 Sensitivities of the MAP Point

We now study the averaged generalized sensitivities of the MAP point. As was done previously, we study the effect of the sample size on the generalized sensitivities. This is done by computing generalized MAP point sensitivities for 3000 data samples. We then randomly select and average ten groups of 20 sensitivities, ten groups of 100 sensitivities, and ten groups of 500 sensitivities, which are plotted in Fig. 8.

In this case we can see that even groups of 20 sensitivities produce little variation in the averaged sensitivity measure. Thus, we can conclude that for this application, using a sample average of just 20 sensitivities provides sufficient accuracy for our purposes. Furthermore, we notice that the generalized sensitivities of the MAP point are significantly greater in magnitude than those computed for Bayes risk. For this problem, it appears that the MAP point is more sensitive to perturbations in the complementary parameters than the posterior uncertainty is. We see that the MAP point has the greatest sensitivity to γ_2 , β , z_1 , and f_2 . It is interesting to note that for Bayes risk, σ and β had the second and fifth greatest sensitivity, respectively. In contrast, the sensitivity rankings of these two parameters have switched places with respect to the MAP point.

Finally, we examine the pointwise sensitivities of the MAP point to the experimental parameters depicted in Fig. 9. Each pointwise sensitivity is computed as an average of 20 sensitivities computed from different data samples. We compared these pointwise sensitivities with those computed from an average of 1000 sensitivities, and as our study on sample size in Fig. 8 would indicate, there was minimal difference.

We can again see that the sensors grouped around small values of x_1 and large values of x_2 are most important with respect to the MAP point. We observe that for this problem the sensors with greatest importance to the MAP point coincide closely with those sensors that are important for Bayes risk.




FIG. 8: Averaged generalized sensitivities of the MAP point to complementary parameters computed for ten groups of sensitivities of size 20, 100, and 500




FIG. 9: Pointwise sensitivities of the MAP point to experimental parameters computed as an average of 20 sample sensitivities

8. CONCLUSION

In this article we take foundational steps in applying HDSA to large-scale nonlinear Bayesian inverse problems. In particular, we focus on HDSA of the MAP point and Bayes risk to the auxiliary and experimental parameters and present efficient methods for computing the corresponding HDSA indices. Performing HDSA is important as it reveals the auxiliary parameters the inverse problem is most sensitive to. Moreover, HDSA with respect to measurement data helps identify the measurements that are important to the solution of the inverse problem, and can guide the design of experiments by investing resources to obtain good quality data from important measurement points.

It is also important to note that the Bayesian formulation allows for the computation of HDSA indices prior to conducting experiments. Namely, we use the information encoded in the Bayesian inverse problem to obtain likely realizations of measurement data, which are used to compute the Bayes risk sensitivities and average MAP point sensitivities. This is a key factor that makes this approach attractive for HDSA of Bayesian inverse problems, while minimizing experimental costs.

While the steady-state heat conduction model presented in Section 6 is an academic model problem, it has many features that are seen in real applications. We found that the tilt angle, heat amplitude, and center in the horizontal direction of volume heat source as well as the heat transfer coefficient and data noise were the parameters that both the Bayes risk and the MAP point were most sensitive to. We also determined which sensors provide the most informative data and found that for this problem the Bayes risk is generally less sensitive to perturbations of the complementary parameters than the MAP point is. Such observations can be instrumental in areas such as additive manufacturing. By applying the proposed methods to additive manufacturing problems, one can determine *a priori* which experimental factors the inverse problem solution will be most sensitive to and thereby guide the calibration of equipment tolerances with this information.

The MAP point is a key point estimator for the inversion parameters and performing HDSA on this quantity provides valuable insight regarding the sensitivity of the inverse problem to complementary parameters. On the other hand, Bayes risk provides a measure of the statistical quality of the estimated parameters, and is a common utility function in decision theory. Additionally, up to a linearization, Bayes risk can be considered as a proxy for posterior uncertainty. These considerations, coupled with the fact that the methods for HDSA of Bayes risk build on methods for HDSA of the MAP point, made Bayes risk a suitable HDSA QoI in the first steps towards HDSA of Bayesian inverse problems.

In our future work, we plan to investigate HDSA of different quantities such as average posterior variance or expected information gain. Suitable approximations of the posterior, such as a Laplace approximation, can be considered, to mitigate the high cost of HDSA of such quantities in large-scale nonlinear inverse problems. Another interesting line of inquiry is to use HDSA within the context of OED under uncertainty [30,31]. HDSA can reveal model uncertainties that the OED criterion is most sensitive to and thus must be accounted for in the optimal design process. On the other hand, model uncertainties the design criterion is less sensitive to may be fixed at some nominal values, hence reducing the complexity of OED under uncertainty problems.

ACKNOWLEDGMENTS

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under Contract No. DE-NA-0003525. This work was supported by the US Department of Energy, Office of Advanced Scientific Computing Research, Field Work Proposal 20-023231.

The work of I. Sunseri and A. Alexanderian was supported in part by the National Science Foundation under Grant No. DMS-1745654. Additionally, the work of A. Alexanderian was also supported in part by the National Science Foundation under Grant No. DMS-2111044.

REFERENCES

1. Tarantola, A., *Inverse Problem Theory and Methods for Model Parameter Estimation*, Philadelphia: Society for Industrial and Applied Mathematics, 2005.
2. Stuart, A.M., Inverse Problems: A Bayesian Perspective, *Acta Numer.*, **19**:451–559, 2010.
3. Sunseri, I., Hart, J., van Bloemen Waanders, B., and Alexanderian, A., Hyper-Differential Sensitivity Analysis for Inverse Problems Constrained by Partial Differential Equations, *Inverse Probl.*, **36**(12):125001, 2020.

4. Hart, J., van Bloemen Waanders, B., and Herzog, R., Hyper-Differential Sensitivity Analysis of Uncertain Parameters in PDE-Constrained Optimization, *Int. J. Uncertainty Quantif.*, **10**(3):225–248, 2020.
5. Brandes, K. and Griesse, R., Quantitative Stability Analysis of Optimal Solutions in PDE-Constrained Optimization, *J. Comput. Appl. Math.*, 206(2):908–926, 2007.
6. Griesse, R. and Walther, A., Parametric Sensitivities for Optimal Control Problems Using Automatic Differentiation, *Optim. Control Appl. Methods*, **24**:297–314, 2003.
7. Griesse, R. and Vexler, B., Numerical Sensitivity Analysis for the Quantity of Interest in PDE-Constrained Optimization, *SIAM J. Sci. Comput.*, **29**(1):22–48, 2007.
8. Griesse, R., Stability and Sensitivity Analysis in Optimal Control of Partial Differential Equations, Habilitation Thesis, Faculty of Natural Sciences, Karl-Franzens University, 2007.
9. Griesse, R., Parametric Sensitivity Analysis in Optimal Control of a Reaction Diffusion System—Part I: Solution Differentiability, *Numer. Funct. Anal. Optim.*, **25**(1-2):93–117, 2004.
10. Griesse, R., Parametric Sensitivity Analysis in Optimal Control of a Reaction-Diffusion System—Part II: Practical Methods and Examples, *Optimi. Methods Software*, **19**(2):217–242, 2004.
11. Büskens, C. and Griesse, R., Parametric Sensitivity Analysis of Perturbed PDE Optimal Control Problems with State and Control Constraints, *J. Optim. Theory Appl.*, **131**(1):17–35, 2006.
12. Griesse, R. and Volkwein, S., Parametric Sensitivity Analysis for Optimal Boundary Control of a 3D Reaction-Diffusion System, in *Nonconvex Optimization and Its Applications*, Vol. 83, G.D. Pillo and M. Roma, Eds., Berlin: Springer, 2006.
13. Atkinson, A.C. and Donev, A.N., *Optimum Experimental Designs*, Oxford: Oxford University Press, 1992.
14. Pázman, A., *Foundations of Optimum Experimental Designs*, Dordrecht, the Netherlands: D. Reidel Publishing Co., 1986.
15. Pukelsheim, F., *Optimal Design of Experiments*, New York: John Wiley & Sons, 1993.
16. Chaloner, K. and Verdinelli, I., Bayesian Experimental Design: A Review, *Stat. Sci.*, **10**(3):273–304, 1995.
17. Uciński, D., *Optimal Measurement Methods for Distributed Parameter System Identification*, Boca Raton, FL: CRC Press, 2005.
18. Alexanderian, A., Optimal Experimental Design for Infinite-Dimensional Bayesian Inverse Problems Governed by PDEs: A Review, *Inverse Probl.*, **37**:043001, 2021.
19. Haber, E., Horesh, L., and Tenorio, L., Numerical Methods for Experimental Design of Large-Scale Linear Ill-Posed Inverse Problems, *Inverse Probl.*, **24**(055012):125–137, 2008.
20. Haber, E., Horesh, L., and Tenorio, L., Numerical Methods for the Design of Large-Scale Nonlinear Discrete Ill-Posed Inverse Problems, *Inverse Probl.*, **26**(2):025002, 2010.
21. Horesh, L., Haber, E., and Tenorio, L., *Optimal Experimental Design for the Large-Scale Nonlinear Ill-Posed Problem of Impedance Imaging*, New York: Wiley, pp. 273–290, 2010.
22. Alexanderian, A., Gloor, P.J., and Ghattas, O., On Bayesian A- and D-Optimal Experimental Designs in Infinite Dimensions, *Bayesian Anal.*, **11**(3):671–695, 2016.
23. Bui-Thanh, T., Ghattas, O., Martin, J., and Stadler, G., 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**(6):A2494–A2523, 2013.
24. Ambrosetti, A. and Prodi, G., *A Primer of Nonlinear Analysis*, Cambridge: Cambridge University Press, 1995.
25. Villa, U., Petra, N., and Ghattas, O., hIPPYlib: An Extensible Software Framework for Large-Scale Inverse Problems Governed by PDEs: Part I: Deterministic Inversion and Linearized Bayesian Inference, *ACM Trans. Math. Software (TOMS)*, **47**(2):1–34, 2021.
26. Gunzburger, M.D., *Perspectives in Flow Control and Optimization*, Philadelphia: SIAM, 2003.
27. Petra, N., Martin, J., Stadler, G., and Ghattas, O., A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems, Part II: Stochastic Newton MCMC with Application to Ice Sheet Flow Inverse Problems, *SIAM J. Sci. Comput.*, **36**(4):A1525–A1555, 2014.
28. Flath, H.P., Wilcox, L.C., Akçelik, V., Hill, J., van Bloemen Waanders, B., and Ghattas, O., Fast Algorithms for Bayesian Uncertainty Quantification in Large-Scale Linear Inverse Problems Based on Low-Rank Partial Hessian Approximations, *SIAM J. Sci. Comput.*, **33**(1):407–432, 2011.

29. Martin, J., Wilcox, L.C., Burstedde, C., and Ghattas, O., A Stochastic Newton MCMC Method for Large-Scale Statistical Inverse Problems with Application to Seismic Inversion, *SIAM J. Sci. Comput.*, **34**(3):A1460–A1487, 2012.
30. Alexanderian, A., Petra, N., Stadler, G., and Sunseri, I., Optimal Design of Large-Scale Bayesian Linear Inverse Problems under Reducible Model Uncertainty: Good to Know What You Don't Know, *SIAM/ASA J. Uncertainty Quantif.*, **9**(1):163–184, 2021.
31. Koval, K., Alexanderian, A., and Stadler, G., Optimal Experimental Design under Irreducible Uncertainty for Inverse Problems Governed by PDEs, *Inverse Probl.*, **36**(7):075007, 2020.