ABSTRACT

SUNSERI, ISAAC PAUL. Design and Sensitivity Analysis of Inverse Problems Governed by Partial Differential Equations (Under the direction of Dr. Alen Alexanderian).

Many applications and physical phenomena can be described by systems of partial differential equations (PDEs). These PDEs may be coupled together and include multiple state variables of differing units and scales, as well as many physical model parameters. Inverse problems arise when these physical model parameters are uncertain and we wish to estimate them using a set of measurement data that typically corresponds to the PDE states. Using the measurement data and the PDE model that describes the physical system being considered, an optimization problem is solved to estimate one or more unknown parameters, which are called the inversion parameters. Other model parameters besides the inversion parameter are referred to as auxiliary parameters, while parameters specifying the experimental conditions, such as types of measurements or noise levels, are experimental parameters. In this work, the author explores several methods to better understand PDE-constrained inverse problems and design effective experiments for these problems while considering additional uncertainties in the auxiliary and experimental parameters.

Optimal experimental design (OED) seeks to determine the arrangement of sensors, measurement times, and other experimental conditions that optimize some desired design criterion. Hyper-differential sensitivity analysis (HDSA) is a form of post-optimal sensitivity analysis that quantifies the change in the solution of an optimization problem as auxiliary or experimental parameters are perturbed. By combining these techniques in this dissertation, the author develops a comprehensive method to study inverse problems and design effective experiments for them. This dissertation addresses the following topics in the context of PDE-constrained inverse problems: (i) OED that takes into account additional, reducible model uncertainties (ii) HDSA of nonlinear deterministic inverse problems, (iii) HDSA of nonlinear Bayesian inverse problems.

© Copyright 2022 by Isaac Paul Sunseri

All Rights Reserved

Design and Sensitivity Analysis of Inverse Problems Governed by Partial Differential Equations

by
Isaac Paul Sunseri

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Mathematics

Raleigh, North Carolina
2022

APPROVED BY:

Arvind Krishna Saibaba

Pierre Gremaud

Bart Van Bloemen Waanders
External Member

Murthy Guddati

Alen Alexanderian
Chair of Advisory Committee

DEDICATION

to Haley

BIOGRAPHY

Isaac Paul Sunseri graduated summa cum laude with a Batchelor of Sciences degree in Mathematics from North Carolina State University in 2017. Isaac earned a Master of Sciences degree in Applied Mathematics from North Carolina State University in December 2020 and will receive a Doctorate of Philosophy degree in Applied Mathematics from North Carolina State University in May 2022. During his graduate studies Isaac served as the student coordinator of the Applied Math Graduate Student Seminar at NC State University, co-president of the Graduate Instructor Tools & Support group at NC State, was a graduate fellow in a National Science Foundation research training group program, received the Franke-Norris-Griggs Teaching Award from the NC State department of mathematics, and completed an internship with Sandia National Laboratories.

ACKNOWLEDGMENTS

I want to thank Alen Alexanderian for the great amount of time and energy he invested in me. Alen has always taken time out of his week to provide excellent advice, encouragement, and assistance to my work and has been supportive of my research and professional development. I am grateful to Arvind Saibaba, Pierre Gremaud, Murthy Guddati and Bart Van Bloemen Waanders for serving on my committee and their helpful feedback and support. Furthermore, I am thankful for my internship mentors Bart Van Bloemen Waanders and Joseph Hart for their great amount of assistance and support they provided during and after my internship with Sandia National Laboratories. I would also like to thank and acknowledge the help and work of Noemi Petra and Georg Stadler with the work presented in chapter 2 of this dissertation.

I would also like to acknowledge my sources of financial support. My work was supported in part by the National Science Foundation (NSF) award DMS-1745654. The work in Chapters 3 and 4 was performed, in part, at the Center for Integrated Nanotechnologies, an Office of Science User Facility operated for the U.S. Department of Energy. Sandia National Laboratories is a multi-mission 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 National Nuclear Security Administration under contract DE-NA0003525. The views expressed in Chapters 3 and 4 do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

TABLE OF CONTENTS

LIST OF TABLES . . . . **vii**

LIST OF FIGURES . . . . **viii**

Chapter 1 Introduction **1**
- 1.1 Outline of the dissertation 2
- 1.2 Contributions 3

Chapter 2 Optimal Experimental Design Under Uncertainty **4**
- 2.1 Introduction 4
- 2.2 Bayesian inverse problems governed by models with reducible uncertainties . 7
  - 2.2.1 Notation 8
  - 2.2.2 Marginals of Gaussian measures 8
  - 2.2.3 Bayesian inverse problem setup 10
  - 2.2.4 The discretized problem 11
  - 2.2.5 Optimal experimental design 12
- 2.3 Marginalized Bayesian A-optimality 14
  - 2.3.1 Alternative form of the posterior 14
  - 2.3.2 The marginalized A-optimality criterion 16
  - 2.3.3 Computing optimal designs 18
- 2.4 Computational methods 18
  - 2.4.1 Efficient computation of mOED objective and its gradient 18
  - 2.4.2 Sparsity control 21
  - 2.4.3 Greedy sensor placement 21
- 2.5 Model problem setup 21
  - 2.5.1 Parameter-to-observable map 22
  - 2.5.2 Prior laws of m and b 23
  - 2.5.3 Discretization 24
  - 2.5.4 Illustrating the impact of the secondary uncertainty 25
- 2.6 Computational results 26
  - 2.6.1 Comparison of sparsification algorithms 26
  - 2.6.2 Studying the posterior uncertainty 28
  - 2.6.3 Study of MAP points 29
- 2.7 Conclusion 30

Chapter 3 Hyper-Differential Sensitivity Analysis for Deterministic Inverse Problems **32**
- 3.1 Introduction 32
- 3.2 Preliminaries 35
  - 3.2.1 Inverse Problems 36
  - 3.2.2 Design of Experiments 37
- 3.3 Hyper-differential Sensitivity Analysis for Inverse Problems 38
  - 3.3.1 Mathematical Formulation 38

3.3.2 Sensitivity Indices 39
3.3.3 Interpretation of Experimental Parameter HDSA for Linear Inverse Problems 41
3.4 Computational Considerations 43
3.5 Model Problem 44
3.6 Computational Results 47
3.6.1 Generalized Sensitivity Results 48
3.6.2 HDSA with Respect to Experimental Parameters 50
3.6.3 HDSA with Respect to Auxiliary Parameters 54
3.6.4 Discussion 56
3.7 Conclusion 57
3.8 Appendix: Details for the model inverse problem (3.7) 58

Chapter 4 Hyper-Differential Sensitivity Analysis for Bayesian Inverse Problems 60
4.1 Introduction 60
4.2 An Illustrative Example 63
4.3 Bayesian inverse problems and complementary parameters 64
4.4 HDSA for Nonlinear Bayesian Inverse Problems 67
4.4.1 The HDSA QoIs 67
4.4.2 Sensitivity operator of Bayes risk 67
4.4.3 Sensitivity Indices 68
4.5 Computational Methods 69
4.5.1 Computing the sensitivity indices 69
4.5.2 Computational Costs 72
4.6 Model Problem 73
4.6.1 Forward Model 74
4.6.2 Prior Measure and State Solution 75
4.7 Results 77
4.7.1 Modeling Parameter Perturbations 78
4.7.2 Sensitivities of Bayes Risk 79
4.7.3 Sensitivities of the MAP Point 81
4.8 Conclusion 83

Chapter 5 Conclusion 85

References 87

LIST OF TABLES

Table 4.1 Computational costs summary 73

LIST OF FIGURES

Figure 2.1 Left: Sketch of domain $\mathcal{D}$ and velocity field $\boldsymbol{v}$ in (2.24). The red dot indicates the location $\boldsymbol{x}_0 = (0.5, 0.35)$ where the source term (2.25) is centered. Right: the "truth" source term m and five samples from the prior distribution of m shown in cyan and various shades of orange, respectively. 23

Figure 2.2 Shown in each row are snapshots of the concentration at times $t = 0, 0.4, 0.6, 1$ (from left to right). For the primary parameter m entering on the right hand side of eq. (2.24a), the "truth" parameter shown in fig. 2.1 (right) is used. For the secondary parameter b, i.e., the initial condition, two different realizations from the distribution of b are used. Note that a different colorbar is used for the initial conditions than for the other snapshots. 26

Figure 2.3 Left: mOED objective values (y-axis) plotted against number of sensors (x-axis) for the greedy (red dots) and the ℓ_0-sparsification approaches (blue dots). Right: Number of mOED objective evaluations required to converge for computing greedy (red) and ℓ_0-sparsified (blue) designs. 27

Figure 2.4 Shown are A-optimal designs with 20 sensors (filled squares) using mOED (left) and OED without marginalization (center), i.e., the design obtained with OED neglecting secondary uncertainties. Inactive sensors are shown as empty squares. On the right, the marginal posterior standard deviation field (i.e., square root of the diagonal of $\boldsymbol{\Gamma}_{\text{post,m}}(\boldsymbol{w})$ in eq. (2.14)) is shown for the two designs. 28

Figure 2.5 Comparison of MAP estimates computed with mOED and OED without marginalization. Shown are the MAP estimates computed using sensor placements obtained via mOED (black solid line), OED with the secondary parameter b set to the "truth" (blue solid line), and OED with b taken as realizations from corresponding prior distribution (red dotted lines). 29

Figure 2.6 *Left:* Relative error in the MAP estimate (x-axis) and reduction in the objective (y-axis) for mOED (black dot), OED with the secondary parameter b set to the "truth" (blue dot), and OED with b taken as different realizations of b (red dots). *Right:* The distribution of the errors with various realizations of the noise in the data. Note that the x-axis is cut at 2 due to the long tail of the error distribution corresponding to OED with b taken as different realizations of b. In this study, we used 200 samples of the secondary parameter, and 500 samples of measurement noise. 30

Figure 3.1 Concentration sensor, pressure sensor, and source locations. 46

Figure 3.2 Left: True Permeability Right: Reconstructed Solution. 47

Figure 3.3 Bar graph of generalized parameter sensitivities. 49

Figure 3.4 Perturbation of the left pressure Dirichlet boundary condition. . . . 50
Figure 3.5 Concentration sensitivities at times 0.01, 0.04, 0.07, 0.10, 0.13, and 0.1567. 51
Figure 3.6 A spatial distribution of the time evolution of contaminant sensitivities. Each subplot has the same horizontal axis range depicting time from 0 to .16. Each vertical axis subplot has the same range depicting sensitivity from 0 to .2832 53
Figure 3.7 Pressure sensitivities 54
Figure 3.8 Left: Source sensitivities, Right: Darcy Velocity Field 55
Figure 3.9 Pressure Dirichlet boundary condition sensitivities 56

Figure 4.1 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. 64
Figure 4.2 The heat source function $f(\boldsymbol{x})$ 75
Figure 4.3 The state solution of the governing system of partial differential equations with the experimental sensor locations indicated by filled black circles. 77
Figure 4.4 Top: three samples from the log-conductivity prior. Bottom: inverse problem MAP point estimates solved using data generated by the above prior sample. 78
Figure 4.5 Generalized sensitivities of Bayes Risk to complementary parameters computed with sample sizes of 20, 100, and 500. 79
Figure 4.6 Point-wise 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. 81
Figure 4.7 Averaged generalized sensitivities of the MAP point to complementary parameters computed with sample sizes of 20, 100, and 500. . . 82
Figure 4.8 Point-wise sensitivities of the MAP point to experimental parameters computed as an average of 20 sample sensitivities. 83

CHAPTER

1

INTRODUCTION

Throughout scientific and engineering fields, partial differential equations are used to model a wide variety of physical systems. These models are used to predict outcomes, design experiments, and estimate unknown quantities. In practice, these models frequently contain a number of uncertain parameters, some of which may need to be estimated. Using data measurements of the PDE states, along with knowledge of the governing model, an optimization problem can be solved to infer the unknown parameters of interest. This type of problem is referred to as a PDE-constrained inverse problem and is widely studied [17, 24, 25, 50, 53, 81, 86]. Inverse problems have many challenges that must be carefully considered including ill-posedness, uncertainty in the parameters, noisy-data, and experimental designs. In addition to the unknown parameter(s) of interest that are inferred by the inverse problem, there are often many other uncertain parameters that must be fixed to fully specify the inverse problem. These additional uncertain parameters are of concern because incorrectly estimating and fixing these parameters may adversely affect the quality of the solution, or else lead experimenters to make incorrect conclusions if they are unaware of these uncertainties.

This dissertation explores the design of PDE-constrained inverse problems through the lens of optimal experimental design (OED) [5, 6, 8, 11, 13, 15, 37, 38, 46, 47, 55, 59, 60, 67, 84] and sensitivity analysis [41, 42, 44, 48, 49, 71, 77] which take additional model uncertainties into account. Optimal experimental design seeks to determine a set of sensor locations, measurement times, or other experimental conditions that optimize some measure of the solution's quality. Sensitivity analysis is a broad field of study and herein hyper-differential sensitivity analysis (HDSA) of PDE-constrained inverse problems is considered. HDSA is a form of derivative-based sensitivity analysis that quantifies the impact of perturbations in the auxiliary and experimental parameters on the solution of the inverse problem. This is useful as it allows for comparison of relative importance between various uncertain model and experimental parameters which may be used to inform the design of inverse problems. In particular, this information can be used to perform parameter subset selection, give physical insight, inform sensor calibration, and prioritize data collection and model parameter estimation techniques.

1.1 Outline of the dissertation

This section provides an overview of the work discussed in each chapter of this dissertation.

Chapter 2: Optimal Experimental Design Under Uncertainty. Traditional OED methods may not consider additional model uncertainties and this can lead to suboptimal design choices. In this chapter we consider designs for linear inverse problems that minimize the posterior uncertainty in the inversion parameters, while accounting for additional reducible uncertainties. To accomplish this we develop efficient computational methods to compute a marginalized A-optimality criterion. By considering optimal experimental design under uncertainty, we show that this approach produces quantifiable improvements in the resulting inverse problem solution.

Chapter 3: Hyper-Differential Sensitivity Analysis for Deterministic Inverse Problems. There are many sources of uncertainty in inverse problems which include both experimental conditions, such as sensor measurements, and model uncertainties that must be estimated and fixed to fully specify the inverse problem. We wish to understand how sensitive the inverse problem solution is to these uncertain parameters, as this can aid experimenters in allocating resources to most accurately estimate and measure the quantities of greatest importance. To this end, we use a local derivative-based sensitivity analysis approach to evaluate the sensitivity of the solution with respect to uncertain parameters. We demonstrate this methodology on a high-dimensional multi-physics problem involving estimation of log-permeability in a porous medium flow problem.

Chapter 4: Hyper-Differential Sensitivity Analysis for Bayesian Inverse Problems. The work of this chapter extends that of chapter 3 to Bayesian inverse problems. This is an important and necessary step to making HDSA a more widely used approach. The solution of a Bayesian inverse problem is a posterior distribution for the inversion parameterrs. This immediately raises several computational and theoretical challenges such as “how does one compute derivative-based sensitivities of a distribution?” We evaluate the sensitivity of the posterior distribution through the MAP point and Bayes risk, and present a detailed analysis of the computational cost and efficiency of this approach. While this approach is specific to using Bayes risk as a measure of posterior uncertainty, it is foundational in providing a framework for evaluating the sensitivity of posterior distributions through other measures of posterior uncertainty.

Chapter 5: Conclusion. Chapter 5 concludes the dissertation offering a high-level synthesis of the methods detailed herein and proposes several possible areas of future work.

1.2 Contributions

The primary contributions of this dissertation are:

- We provide a formulation for OED under reducible model uncertainty for high-dimensional Bayesian linear inverse problems that uses a marginalized A-optimality criterion to appropriately account for these uncertainties; see chapter 2.

- We provide a mathematical and computationally efficient framework for HDSA of deterministic nonlinear inverse problems governed by PDEs with respect to auxiliary and experimental parameters; see chapter 3.

- We extend HDSA to Bayesian nonlinear inverse problems governed by PDEs by studying the sensitivity of various measures of the posterior distribution including the MAP point and the Bayes risk; see chapter 4.

- We demonstrate the effectiveness and interpretation of the methods explored in this dissertation both through simple motivating examples and high dimensional inverse model problems; see chapters 2-4.

CHAPTER

2

OPTIMAL EXPERIMENTAL DESIGN UNDER UNCERTAINTY

The work presented in this chapter is based upon a collaborative publication [7]. The authors of this article are cited alphabetically. The author of this dissertation provided significant contributions to the mathematical theory, computational results, and writing of the article.

2.1 Introduction

An inverse problem uses measurement data and a mathematical model to estimate a set of uncertain model parameters. An experimental design specifies the strategy for collecting measurement data. For example, in inverse problems where measurement data are collected using sensors, an experimental design specifies the placement of the sensors. This is the setting considered in this chapter. Optimal experimental design (OED) [11, 84] refers to the task of determining an experimental setup such that the measurements are most informative about the underlying parameters. This is particularly important in situations where experiments are costly or time-consuming, and thus only a small number of measurements can be collected. In addition to the parameters estimated by the inverse problem, the governing mathematical models often involve simplifications, approximations, or modeling assumptions, resulting in additional uncertainty. These additional uncertainties must be taken into account in the experimental design process; failing to do so could result in suboptimal designs.

We distinguish between two types of uncertainties: *reducible* and *irreducible* [76]. Reducible uncertainties, also referred to as epistemic uncertainties, are those that can be reduced through parameter inference. In contrast, irreducible uncertainties, also known as aleatoric uncertainties, are inherent to the model and are impractical or impossible to reduce through parameter inference. In this chapter, we aim at computing optimal experimental designs in the presence of *reducible* model uncertainty.

In what follows, we consider the model

$$\boldsymbol{y} = \mathcal{E}(m, b) + \boldsymbol{\eta}, \tag{2.1}$$

where $\boldsymbol{y}$ is a vector of measurement data, (m, b) a pair of uncertain parameter vectors or functions, $\mathcal{E}$ a model that maps (m, b) to measurements, and $\boldsymbol{\eta}$ a random vector that models additive measurement errors. Herein, m is the parameter of primary interest, which we seek to infer, and b represents additional uncertain parameters. We assume m and b are elements of infinite-dimensional Hilbert spaces. Furthermore, we assume that the uncertainty in b is reducible. Thus, we can formulate an inverse problem to estimate both m and b. However, when designing experiments to solve the inverse problem, our main interest is reducing the uncertainty in m. We achieve this by finding sensor placements that minimize the posterior uncertainty in m, while taking into account the uncertainty in b. This results in an OED problem in which we minimize the marginal posterior uncertainty in m.

In this chapter, we focus on the case of a model that is linear in m and b and is of the form:

$$\mathcal{E}(m, b) = \mathcal{F}m + \mathcal{G}b. \tag{2.2}$$

Here, $\mathcal{F}$ and $\mathcal{G}$ are bounded linear transformations from suitably defined Hilbert spaces to the space of measurement data. This models, for example, a linear inverse problem with uncertain volume or boundary terms. The mathematical foundations for Bayesian inversion and design of experiments in this context are discussed in section 2.2.

Examples for secondary uncertainties are initial conditions, boundary conditions that are introduced into a model due to the necessity to truncate a computational domain, or unknown forcing or source terms in a real world system that are only incorporated approximately in the mathematical model. When designing experiments, failure to properly account for these secondary uncertainties may result in suboptimal experimental designs. For instance, not taking into account secondary uncertainties in the mathematical model may result in sensors being located close to uncertain sources, resulting in observations that can provide biased information on the parameter of primary interest. On the other hand, if one aims at finding designs that are optimal for both primary and secondary uncertain parameters, the design is likely to be suboptimal for inference of the primary parameter.

Such designs would require optimizing design criteria for primary and secondary uncertain parameters, which would be more expensive. Furthermore, such an approach might require introducing a weight coefficient to balance the importance of the uncertainty in the primary and secondary parameters, further complicating matters.

Related work. In many inverse problems, one has model uncertainties in addition to the inversion parameters. A robust parameter inversion strategy must account for such additional model uncertainties; see [10, 22, 52–54, 64, 65] for a small sample of the literature addressing such issues. This work is about A-optimal experimental design for Bayesian linear inverse problems governed by partial differential equations (PDEs) with model uncertainties. For a review of the literature on optimal design of inverse problems governed by computationally intensive models, we refer to [3]. Here, we mainly review related work on optimal design of linear inverse problems.

The articles [5, 37, 39] present methods for large-scale ill-posed linear inverse problems. Specifically, the present chapter builds on [5], which focuses on A-optimal experimental design of infinite-dimensional Bayesian linear inverse problems.

Recent work also considers A-optimal design of infinite-dimensional Bayesian linear inverse problems with model uncertainties [56]. The key difference to the current chapter is that [56] considers OED for inverse problems governed by models with *irreducible* uncertainties and formulates the OED problem as one of optimization under uncertainty. In contrast, in this work we consider OED under *reducible* model uncertainties and propose a formulation that aims at minimizing the marginal posterior variance of the primary parameters. By combining primary and secondary uncertainties, the problem considered in this work can formally be written as goal-oriented OED problem, as studied in [12]. However, taking a model uncertainty perspective and considering infinite-dimensional primary and secondary uncertain parameters require a tailored approach that distinguishes primary and secondary uncertainties.

Other related efforts include [28, 45, 74]. In [28], the authors present an adaptive A-optimal design strategy for linear dynamical systems. OED for linear inverse problems with linear equality and inequality constraints is addressed in [74]. This results in OED with an effectively nonlinear inverse problem for which the authors propose an approach based on Bayes risk minimization. In [45], the authors present an approach for A-optimal design of infinite-dimensional Bayesian linear inverse problems using ideas from randomized subspace iteration and reweighted ℓ_1-minimization.

Contributions. This chapter makes the following contributions to the state-of-the-art in OED for large-scale linear inverse problems. (i) We provide a mathematical formulation of OED under reducible model uncertainty and show how the OED problem can be reformulated to take advantage of the often low dimensionality of the measurement space (see 2.3); in particular, our formulation eliminates the need for trace estimation in the discretized parameter space. (ii) We develop a scalable computational framework for solving the class of OED problems under study (see section 2.4). Specifically, the computational complexity of our methods, in terms of the number of PDE solves, does not grow with the dimension of the discretized primary and secondary parameters. (iii) We present illustrative numerical results, in context of a contaminant transport inverse problem (see section 2.5 and section 2.6) where we seek to estimate an unknown source term, but have additional uncertainty in the initial state. Our numerical experiments examine different aspects of our proposed framework, and elucidate the importance of incorporating additional model uncertainties in the OED problem.

The latter avoids trace estimation in the infinite-dimensional parameter space, which becomes high-dimensional upon discretization of the Bayesian inverse problem.

Chapter overview. We begin by presenting some preliminaries on Bayesian inverse problems with additional model uncertainties in section 2.2. In section 2.3, we present a mathematical framework for marginalized A-optimal design of experiments. enable efficient computation of marginalized A-optimal designs. Section 2.5 describes the model inverse problem used to illustrate our framework. Specifically, we illustrate the efficacy of the proposed strategy in the context of a contaminant transport problem. In that inverse problem, we seek to infer an uncertain source term, but also have (reducible) uncertainty in the initial state. We present our computational results in section 2.6. Finally, we present concluding remarks in section 2.7.

2.2 Bayesian inverse problems governed by models with reducible uncertainties

After introducing basic notation in section 2.2.1, we present preliminaries regarding Gaussian measures on Hilbert spaces in section 2.2.2. Next, we outline the setup of Bayesian linear inverse problems with additional reducible model uncertainties in infinite-dimensions (section 2.2.3) as well as in discretized form (section 2.2.4). We discuss basics on optimal design of such inverse problems in section 2.2.5.

2.2.1 Notation

Herein we consider a probability space $(\Omega, \mathfrak{A}, \mathbb{P})$, where Ω is a sample space, $\mathfrak{A}$ a sigma-algebra on Ω, and $\mathbb{P}$ is a probability measure. Given a Hilbert space $\mathscr{X}$, we denote by $\mathfrak{B}(\mathscr{X})$ the Borel sigma-algebra on $\mathscr{X}$. A Gaussian measure on $(\mathscr{X}, \mathfrak{B}(\mathscr{X}))$, with mean $\bar{z} \in \mathscr{X}$ and covariance operator $\mathcal{C} : \mathscr{X} \to \mathscr{X}$, is denoted by $\mathcal{N}(\bar{z}, \mathcal{C})$. We also recall that for a random variable $Z : (\Omega, \mathfrak{A}, \mathbb{P}) \to (\mathscr{X}, \mathfrak{B}(\mathscr{X}))$, its law is a Borel measure $\mathcal{L}_Z$ on $\mathscr{X}$, that satisfies $\mathcal{L}_Z(A) = \mathbb{P}(Z \in A)$ for every $A \in \mathfrak{B}(\mathscr{X})$ [88]. Also, for a linear transformation $T : \mathscr{X} \to \mathscr{Y}$, where $\mathscr{Y}$ is another Hilbert space, we denote the adjoint by T^*.

2.2.2 Marginals of Gaussian measures

Here we discuss some preliminaries regarding Gaussian measures and Gaussian random variables taking values in Hilbert spaces. First we record the following known result about the law of a linear transformation of a Hilbert space-valued Gaussian random variable, which we prove for completeness.

Lemma 2.2.1. *Let $\mathscr{X}$ and $\mathscr{Y}$ be infinite-dimensional Hilbert spaces. Suppose $Z : \Omega \to \mathscr{X}$ is a Gaussian random variable with law $\mu = \mathcal{N}(\bar{z}, \mathcal{C})$. Consider the random variable $Y = TZ$, where $T : \mathscr{X} \to \mathscr{Y}$ is a bounded linear transformation. Then, $Y : (\Omega, \mathfrak{A}, \mathbb{P}) \to (\mathscr{Y}, \mathfrak{B}(\mathscr{Y}))$ is a Gaussian random variable with law $\nu = \mathcal{N}(T\bar{z}, T\mathcal{C}T^*)$.*

Proof. Using [70, Proposition 1.18], we know that the law of the random variable $T : (\mathscr{X}, \mathfrak{B}(\mathscr{X}), \mu) \to (\mathscr{Y}, \mathfrak{B}(\mathscr{Y}))$ is given by $\mu \circ T^{-1} = \mathcal{N}(T\bar{z}, T\mathcal{C}T^*) = \nu$. To complete the proof we show $\mathcal{L}_Y = \nu$. Namely, for every $A \in \mathfrak{B}(\mathscr{Y})$,

$$\mathcal{L}_Y(A) = \mathbb{P}(Y \in A) = \mathbb{P}(TZ \in A) = \mathbb{P}(Z \in T^{-1}(A)) = \mu(T^{-1}(A)) = \nu(A).$$

□

Consider a Hilbert space $\mathscr{V} = \mathscr{V}_1 \times \mathscr{V}_2$, where $\mathscr{V}_1$ and $\mathscr{V}_2$ are real, separable, infinite-dimensional Hilbert spaces with inner products $\langle \cdot, \cdot \rangle_1$ and $\langle \cdot, \cdot \rangle_2$, respectively. An element $z \in \mathscr{V}$ is of the form $z = (z_1, z_2)$ with $z_1 \in \mathscr{V}_1$ and $z_2 \in \mathscr{V}_2$, respectively. We assume that $\mathscr{V}$ is equipped with the natural inner product

$$\langle\!\langle x, y \rangle\!\rangle = \langle x_1, y_1 \rangle_1 + \langle x_2, y_2 \rangle_2, \quad x, y \in \mathscr{V}.$$

Let $Z : (\Omega, \mathcal{F}, \mathbb{P}) \to (\mathscr{V}, \mathcal{B}(\mathscr{V}), \mu)$ be a Gaussian random variable with law $\mu = \mathcal{N}(\bar{z}, \mathcal{C})$. The marginal laws of Z can be defined analogously to the finite-dimensional setting, as shown next. This shows that the familiar marginalization results for Gaussian random variables remain meaningful in infinite dimensions.

We denote realizations of Z by $z = (z_1, z_2) = (\Pi_1 z, \Pi_2 z) \in \mathscr{V}$, where Π_1 and Π_2 denote linear projection operators onto $\mathscr{V}_1$ and $\mathscr{V}_2$, respectively. The following result concerns the law of $\Pi_i Z$, $i = 1, 2$, i.e., marginal laws of Z.

Lemma 2.2.2. *$Z_i = \Pi_i Z$ has a Gaussian law μ_i with mean $\bar{z}_i = \Pi_i \bar{z}$ and covariance operator $\mathcal{C}_{ii}$, which satisfies*

$$\langle \mathcal{C}_{ii} u, v \rangle_i = \int_{\mathscr{V}_i} \langle s - \bar{z}_i, u \rangle_i \langle s - \bar{z}_i, v \rangle_i \, \mu_i(ds), \quad i = 1, 2, \text{ for all } u, v \in \mathscr{V}_i. \tag{2.3}$$

Proof. By lemma 2.2.1, $\Pi_i Z$ has a Gaussian law $\mu_i = \mathcal{N}(\Pi_i \bar{z}, \mathcal{C}_{ii})$ with $\mathcal{C}_{ii} = \Pi_i \mathcal{C} \Pi_i^*$. It remains to show that $\mathcal{C}_{ii}$ satisifes eq. (2.3). Without loss of generality, we assume $\bar{z} \equiv 0$ and show the result for $i = 1$. By definition of the covariance operator $\mathcal{C}$ of μ, we have $\langle\!\langle \mathcal{C} a, b \rangle\!\rangle = \int_{\mathscr{V}} \langle\!\langle z, a \rangle\!\rangle \langle\!\langle z, b \rangle\!\rangle \, \mu(dz)$, for $a, b \in \mathscr{V}$. Therefore, for arbitrary $u, v \in \mathscr{V}_1$, we have

$$\langle \mathcal{C}_{11} u, v \rangle_1 = \langle\!\langle \mathcal{C} \Pi_1^* u, \Pi_1^* v \rangle\!\rangle = \int_{\mathscr{V}} \langle\!\langle z, (u, 0) \rangle\!\rangle \langle\!\langle z, (v, 0) \rangle\!\rangle \mu(dz) = \int_{\mathscr{V}_1} \langle s, u \rangle_1 \langle s, v \rangle_1 \mu_1(ds).$$

$\square$

In the present work, $\mathscr{V}_1 = L^2(\mathcal{T})$ and $\mathscr{V}_2 = L^2(\mathcal{D})$ with $\mathcal{T}$ and $\mathcal{D}$ bounded open sets in $\mathbb{R}^{d_i}$ with $d_i \in \{1, 2, 3\}$, for $i = 1, 2$. In this case, realizations of $\Pi_1 Z$ and $\Pi_2 Z$ are square-integrable functions on $\mathcal{T}$ and $\mathcal{D}$, respectively. Thus, we can also view $\Pi_i Z$ as a random field. Consider, e.g., $Z_1 = \Pi_1 Z$. This marginalized random field has mean $\bar{z}_1(x)$ and the following covariance function (kernel):

$$c_{11}(x, y) := \int_{\Omega} (Z_1(x, \omega) - \bar{z}_1(x))(Z_1(y, \omega) - \bar{z}_1(y)) \, \mathbb{P}(d\omega).$$

As expected, the (marginal) covariance operator $\mathcal{C}_{11}$ can be written as an integral operator with kernel c_{11}. To show this, we use eq. (2.3) and again, for simplicity, assume $\bar{z} \equiv 0$. Note that

$$\begin{aligned}
\langle \mathcal{C}_{11} u, v\rangle_1 = \int_{\mathscr{V}_1} \langle s, u\rangle_1 \langle s, v\rangle_1 \mu_1(ds) &= \int_\Omega \langle Z_1(\omega), u\rangle_1 \langle Z_1(\omega), v\rangle_1 \, \mathbb{P}(d\omega) \\
&= \int_\Omega \int_\mathcal{T} \int_\mathcal{T} Z_1(x,\omega) Z_1(y,\omega) u(x) v(y) \, dx \, dy \, \mathbb{P}(d\omega) \\
&= \int_\mathcal{T} \left[\int_\mathcal{T} \left(\int_\Omega Z_1(x,\omega) Z_1(y,\omega) \, \mathbb{P}(d\omega) \right) v(y) \, dy \right] u(x) \, dx \\
&= \int_\mathcal{T} \left[\int_\mathcal{T} c_{11}(x,y) v(y) \, dy \right] u(x) \, dx,
\end{aligned}$$

where we used Fubini's theorem to change the order of the integrals. From this, we deduce

$$[\mathcal{C}_{11} v](\cdot) = \int_\mathcal{T} c_{11}(\cdot, y) v(y) \, dy.$$

In finite dimensions, we recover the following well-known [82] result, which we prove here for completeness:

Lemma 2.2.3. *Consider a Gaussian random vector*

$$\boldsymbol{Z} = \begin{bmatrix} \boldsymbol{Z}_1 \\ \boldsymbol{Z}_2 \end{bmatrix} \sim \mathcal{N}(\bar{\boldsymbol{z}}, \mathbf{C}) = \mathcal{N}\left(\begin{bmatrix} \bar{\boldsymbol{z}}_1 \\ \bar{\boldsymbol{z}}_2 \end{bmatrix}, \begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} \\ \mathbf{C}_{21} & \mathbf{C}_{22} \end{bmatrix} \right),$$

where $\boldsymbol{Z}_1$ and $\boldsymbol{Z}_2$ denote subsets of entries of $\boldsymbol{Z}$ and the mean and covariance matrix are partitioned consistent with partitioning of $\boldsymbol{Z}$. Then, the marginals of $\boldsymbol{Z}$ are Gaussian, with $\boldsymbol{Z}_1 \sim \mathcal{N}(\bar{\boldsymbol{z}}_1, \mathbf{C}_{11})$ and $\boldsymbol{Z}_2 \sim \mathcal{N}(\bar{\boldsymbol{z}}_2, \mathbf{C}_{22})$.

Proof. Note that $\boldsymbol{Z}_1 = \mathbf{P}\boldsymbol{Z}$ with $\mathbf{P} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \end{bmatrix}$, where $\mathbf{I}$ is the identity matrix of dimension equal to that of $\boldsymbol{Z}_1$ and $\mathbf{0}$ the zero matrix of the same size as $\boldsymbol{Z}_2$. Thus, $\boldsymbol{Z}_1 \sim \mathcal{N}\big(\mathbf{P}\bar{\boldsymbol{z}}, \mathbf{P}\mathbf{C}\mathbf{P}^\mathsf{T}\big) = \mathcal{N}(\bar{\boldsymbol{z}}_1, \mathbf{C}_{11})$. Showing the statement about the law of $\boldsymbol{Z}_2$ is analogous. $\square$

Remark 2.2.1. *For additional details on marginals of Gaussian random vectors see [14, 27, 36, 78, 82].*

2.2.3 Bayesian inverse problem setup

We consider a Bayesian linear inverse problem for $\theta = (m, b) \in \mathscr{V} = \mathscr{V}_1 \times \mathscr{V}_2$ and where the forward model is of the form eq. (2.2). We assume Gaussian priors for the primary and secondary parameters, which we denote by m and b, respectively, and for simplicity of the presentation assume no prior correlation between m and b.[1] Thus, the prior law of (m, b)

---

[1] The presented framework can be modified to allow for prior correlations between m and b.

is the product measure $\mu_{\text{pr}} = \mu_{\text{pr,m}} \otimes \mu_{\text{pr,b}}$, with $\mu_{\text{pr,m}}$ and $\mu_{\text{pr,b}}$ each Gaussian measures on $\mathscr{V}_1$ and $\mathscr{V}_2$, i.e., $\mu_{\text{pr,m}} = \mathcal{N}(m_{\text{pr}}, \Gamma_{\text{pr,m}})$ and $\mu_{\text{pr,b}} = \mathcal{N}(b_{\text{pr}}, \Gamma_{\text{pr,b}})$. Note that $\mu_{\text{pr}} = \mathcal{N}(\theta_{\text{pr}}, \Gamma_{\text{pr}})$ with $\theta_{\text{pr}} = (m_{\text{pr}}, b_{\text{pr}})$ and $\Gamma_{\text{pr}} = \Gamma_{\text{pr,m}} \times \Gamma_{\text{pr,b}}$, where

$$(\Gamma_{\text{pr,m}} \times \Gamma_{\text{pr,b}})(u_1, u_2) = (\Gamma_{\text{pr,m}} u_1, \Gamma_{\text{pr,b}} u_2), \quad (u_1, u_2) \in \mathscr{V}.$$

The inverse problem under study considers inference of m and b using measurement data $\boldsymbol{y} \in \mathbb{R}^{n_\text{d}}$ and the model

$$\boldsymbol{y} = \mathcal{F}m + \mathcal{G}b + \boldsymbol{\eta}. \tag{2.4}$$

The measurement noise vector $\boldsymbol{\eta}$ is assumed to be independent of (m, b), and we assume a Gaussian noise model, $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Gamma}_{\text{noise}})$. Under these assumptions, the posterior is the Gaussian measure $\mu_{\text{post}}^{\boldsymbol{y}} = \mathcal{N}(\theta_{\text{post}}, \Gamma_{\text{post}})$ with [62]

$$\Gamma_{\text{post}}^{-1} = \mathcal{E}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{E} + \Gamma_{\text{pr}}^{-1}, \quad \theta_{\text{post}} = \Gamma_{\text{post}}(\mathcal{E}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \boldsymbol{y} + \Gamma_{\text{pr}}^{-1} \theta_{\text{pr}}). \tag{2.5}$$

Note that $\mathcal{E}^*$ denotes the adjoint of the linear transformation $\mathcal{E}$. Specifically, $\mathcal{E}^*$ satisfies $\mathcal{E}^* \boldsymbol{y} = (\mathcal{F}^* \boldsymbol{y}, \mathcal{G}^* \boldsymbol{y}) \in \mathscr{V}$, for $\boldsymbol{y} \in \mathbb{R}^{n_\text{d}}$.

2.2.4 The discretized problem

Let $\boldsymbol{m}$ and $\boldsymbol{b}$ be discretized versions of m and b. Recall that we consider a parameter space $\mathscr{V}$ of the form $\mathscr{V} = L^2(\mathcal{T}) \times L^2(\mathcal{D})$. The discretized parameter space is $\mathscr{V}_n = \mathbb{R}^{n_\text{m}} \times \mathbb{R}^{n_\text{b}} \cong \mathbb{R}^n$, where n_m and n_b are the dimensions of the discretized parameters $\boldsymbol{m}$ and $\boldsymbol{b}$, respectively, and $n = n_\text{m} + n_\text{b}$. An element $\boldsymbol{u} \in \mathscr{V}_n$, $\boldsymbol{u} = (\boldsymbol{u}_1, \boldsymbol{u}_2)$ with $\boldsymbol{u}_1 \in \mathbb{R}^{n_\text{m}}$ and $\boldsymbol{u}_2 \in \mathbb{R}^{n_\text{b}}$, can be represented as $\boldsymbol{u} = [\boldsymbol{u}_1^\mathsf{T} \quad \boldsymbol{u}_2^\mathsf{T}]^\mathsf{T}$. The discretized parameter space is endowed with the inner product

$$\langle \boldsymbol{u}, \boldsymbol{v} \rangle_{\mathbb{M}} = \boldsymbol{u}_1^\mathsf{T} \mathbf{M}_1 \boldsymbol{v}_1 + \boldsymbol{u}_2^\mathsf{T} \mathbf{M}_2 \boldsymbol{v}_2 = \boldsymbol{u}^\mathsf{T} \mathbf{M} \boldsymbol{v}, \quad \boldsymbol{u}, \boldsymbol{v} \in \mathscr{V}_n,$$

with $\mathbf{M} = \left[\begin{smallmatrix} \mathbf{M}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_2 \end{smallmatrix}\right]$, and where the "weight" matrices $\mathbf{M}_1$ and $\mathbf{M}_2$ are defined based on the method used to discretize the L^2-inner products on $L^2(\mathcal{T})$ and $L^2(\mathcal{D})$, respectively; see section 2.6 for examples.

The discretized forward operator is defined by

$$\mathcal{E}\boldsymbol{\theta} = \begin{bmatrix} \mathbf{F} & \mathbf{G} \end{bmatrix} \begin{bmatrix} \boldsymbol{m} \\ \boldsymbol{b} \end{bmatrix} = \mathbf{F}\boldsymbol{m} + \mathbf{G}\boldsymbol{b},$$

where $\mathbf{F}$ and $\mathbf{G}$ are discretizations of $\mathcal{F}$ and $\mathcal{G}$ in eq. (2.4). The respective marginal priors are $\mathcal{N}(\boldsymbol{m}_{\text{pr}}, \mathbf{\Gamma}_{\text{pr,m}})$ and $\mathcal{N}(\boldsymbol{b}_{\text{pr}}, \mathbf{\Gamma}_{\text{pr,b}})$, and the prior covariance is $\mathbf{\Gamma}_{\text{pr}} = \left[\begin{smallmatrix} \mathbf{\Gamma}_{\text{pr,m}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Gamma}_{\text{pr,b}} \end{smallmatrix}\right]$.

Using eq. (2.5), the posterior covariance operator satisfies

$$\boldsymbol{\Gamma}_{\text{post}}^{-1} = \begin{bmatrix} \boldsymbol{\Gamma}_{\text{pr,m}}^{-1} + \mathbf{F}^* \boldsymbol{\Gamma}_{\text{noise}}^{-1} \mathbf{F} & \mathbf{F}^* \boldsymbol{\Gamma}_{\text{noise}}^{-1} \mathbf{G} \\ \mathbf{G}^* \boldsymbol{\Gamma}_{\text{noise}}^{-1} \mathbf{F} & \boldsymbol{\Gamma}_{\text{pr,b}}^{-1} + \mathbf{G}^* \boldsymbol{\Gamma}_{\text{noise}}^{-1} \mathbf{G} \end{bmatrix}.$$

Computing the inverse of the block matrix on the right is facilitated by the well-known formula for the inverse of a such matrices [61, Theorem 2.1(ii)]. Specifically, we can show that the covariance operator of the marginal posterior law of $\boldsymbol{m}$ is given by

$$\boldsymbol{\Gamma}_{\text{post,m}} = \big(\boldsymbol{\Gamma}_{\text{pr,m}}^{-1} + \mathbf{F}^* \boldsymbol{\Gamma}_{\text{noise}}^{-1} \mathbf{F} - \mathbf{F}^* \boldsymbol{\Gamma}_{\text{noise}}^{-1} \mathbf{G} (\boldsymbol{\Gamma}_{\text{pr,b}}^{-1} + \mathbf{G}^* \boldsymbol{\Gamma}_{\text{noise}}^{-1} \mathbf{G})^{-1} \mathbf{G}^* \boldsymbol{\Gamma}_{\text{noise}}^{-1} \mathbf{F}\big)^{-1}. \tag{2.6}$$

Note also that for

$$\mathbf{F} : (\mathbb{R}^{n_\text{m}}, \langle \cdot, \cdot \rangle_{\mathbf{M}_1}) \to (\mathbb{R}^{n_\text{d}}, \langle \cdot, \cdot \rangle_{\mathbb{R}^{n_\text{d}}}) \quad \text{and} \quad \mathbf{G} : (\mathbb{R}^{n_\text{b}}, \langle \cdot, \cdot \rangle_{\mathbf{M}_2}) \to (\mathbb{R}^{n_\text{d}}, \langle \cdot, \cdot \rangle_{\mathbb{R}^{n_\text{d}}}),$$

where $\langle \cdot, \cdot \rangle_{\mathbb{R}^{n_\text{d}}}$ denotes the Euclidean inner product on $\mathbb{R}^{n_\text{d}}$, the respective adjoint operators are defined by (cf. e.g., [18])

$$\mathbf{F}^* = \mathbf{M}_1^{-1} \mathbf{F}^\mathsf{T} \quad \text{and} \quad \mathbf{G}^* = \mathbf{M}_2^{-1} \mathbf{G}^\mathsf{T}. \tag{2.7}$$

The optimal design approach we follow consists of minimizing the average posterior variance in $\boldsymbol{m}$ by minimizing the trace of the marginal posterior covariance operator defined in eq. (2.6). We call the resulting OED criterion the *marginalized A-optimality criterion.* In section 2.3, we derive an alternative expression for the marginal posterior covariance operator, which is useful in applications which only allow low or moderate dimensional measurements.

2.2.5 Optimal experimental design

We formulate the sensor placement problem using the approach in [5, 39]. We assume $\boldsymbol{x}_i$, $i = 1, \dots, n_\text{d}$, represent a fixed set of candidate sensor locations. The goal is to select an optimal subset of these locations. We assign a non-negative weight $w_i \in \mathbb{R}$ to each $\boldsymbol{x}_i$, $i = 1, \dots, n_\text{d}$. An experimental design is specified by the vector $\boldsymbol{w} = [w_1, w_2, \dots, w_{n_\text{d}}]^\mathsf{T}$.

As detailed in [5,39], binary weight vectors are desirable to decide whether or not to place a sensor in each of the candidate locations. However, solving an OED problem with binary weights is challenging due to its combinatorial complexity. Thus, as in [5], we relax the problem by considering weights $w_i \in [0, 1]$, $i = 1, \dots, n_\text{d}$. Binary weights are obtained using sparsifying penalty functions, as discussed further in section 2.4.2. An alternative approach to obtaining binary weights, which can be suitable for some problems, is a greedy strategy;

see section 2.4.3.

The vector $\boldsymbol{w}$ is introduced into the Bayesian inverse problem through the data likelihood [5]. We assume uncorrelated measurements; i.e., the noise covariance is diagonal, $\boldsymbol{\Gamma}_{\text{noise}} = \text{diag}(\sigma_1^2, \sigma_2^2, \ldots, \sigma_{n_\text{d}}^2)$, with σ_j^2 the noise level at the jth sensor. For $\boldsymbol{w} \in \mathbb{R}^{n_\text{d}}$, we define the diagonal weight matrix $\mathbf{W} = \text{diag}(w_1, w_2, \ldots, w_{n_\text{d}})$ and the matrix $\mathbf{W}_\sigma$ as follows:

$$\mathbf{W}_\sigma := \text{diag}\Big(\frac{w_1}{\sigma_1^2}, \frac{w_2}{\sigma_2^2}, \ldots, \frac{w_{n_d}}{\sigma_{n_d}^2}\Big) = \sum_{j=1}^{n_\text{d}} w_j \sigma_j^{-2} \boldsymbol{e}_j \boldsymbol{e}_j^\top, \tag{2.8}$$

where $\boldsymbol{e}_j$ is the jth coordinate vector in $\mathbb{R}^{n_\text{d}}$. The $\boldsymbol{w}$-dependent MAP estimator and posterior covariance operator are then given by [5]

$$\boldsymbol{\theta}_{\text{MAP}}(\boldsymbol{w}) = \boldsymbol{\Gamma}_{\text{post}}(\boldsymbol{w})\big(\mathbf{E}^*\mathbf{W}_\sigma \boldsymbol{y} + \boldsymbol{\Gamma}_{\text{pr}}^{-1}\boldsymbol{\theta}_{\text{pr}}\big) \quad \text{and} \quad \boldsymbol{\Gamma}_{\text{post}}(\boldsymbol{w}) = (\mathbf{E}^*\mathbf{W}_\sigma\mathbf{E} + \boldsymbol{\Gamma}_{\text{pr}}^{-1})^{-1}. \tag{2.9}$$

Optimal experimental design (OED) is the problem of finding a design that, within constraints on the number of sensors allowed, minimizes the posterior uncertainty in the estimated parameters. This is done by minimizing certain design criteria that quantify the posterior uncertainty [20, 84]. In this chapter, we use the A-optimal design criterion which is given by $\mathsf{tr}\big[\boldsymbol{\Gamma}_{\text{post}}(\boldsymbol{w})\big]$; this criterion quantifies the average posterior variance of the parameter $\boldsymbol{\theta}$. Using this approach for eq. (2.9), the OED objective is given by the sum of the average posterior variance of the primary and secondary parameters. The primary parameter being the main focus of parameter estimation, we seek sensor placements that minimize the uncertainty in the primary parameter, while being aware of the uncertainty in the secondary parameters. This is done by finding designs that minimize the average posterior variance of the primary parameters, quantified according to the corresponding marginalized posterior distribution. We call such designs marginalized A-optimal designs, which are the subject of section 2.3.

Note that ignoring the uncertainty in the secondary parameter and fixing $\boldsymbol{b}$ to some nominal value $\boldsymbol{b}_0$, results in the affine forward model $\mathbf{E}_0\boldsymbol{m} = \mathbf{F}\boldsymbol{m} + \mathbf{G}\boldsymbol{b}_0$. In this case, the posterior law of $\boldsymbol{m}$ is $\mathcal{N}\big(\boldsymbol{m}_{\text{MAP}}^0, \boldsymbol{\Gamma}_{\text{post}}^0\big)$ with

$$\begin{aligned} \boldsymbol{m}_{\text{MAP}}^0(\boldsymbol{w}) &= \boldsymbol{\Gamma}_{\text{post}}^0(\boldsymbol{w})\big(\mathbf{F}^*\mathbf{W}_\sigma(\boldsymbol{y} - \mathbf{G}\boldsymbol{b}_0) + \boldsymbol{\Gamma}_{\text{pr,m}}^{-1}\boldsymbol{m}_{\text{pr}}\big) \quad \text{and} \\ \boldsymbol{\Gamma}_{\text{post}}^0(\boldsymbol{w}) &= (\mathbf{F}^*\mathbf{W}_\sigma\mathbf{F} + \boldsymbol{\Gamma}_{\text{pr,m}}^{-1})^{-1}, \end{aligned} \tag{2.10}$$

and an A-optimal design $\boldsymbol{w}$ is one that minimizes the classical A-optimality criterion

$$\psi(\boldsymbol{w}) := \mathsf{tr}\big[(\mathbf{F}^*\mathbf{W}_\sigma\mathbf{F} + \boldsymbol{\Gamma}_{\text{pr,m}}^{-1})^{-1}\big]. \tag{2.11}$$

Notice that the optimal design does not depend on the choice of $\boldsymbol{b}_0$. More importantly, such an optimal design is completely unaware of the uncertainty in $\boldsymbol{b}$.

2.3 Marginalized Bayesian A-optimality

In this section, we present our formulation of the marginalized A-optimality criterion. We first derive a reformulation of the marginalized posterior covariance that facilitates an efficient computational procedure for computing marginalized A-optimal designs; see section 2.3.1. Then, we present the definition of the marginalized A-optimality criterion, in section 2.3.2, and prove its convexity. Finally, the formulation of the optimization problem for finding marginalized A-optimal designs is discussed in section 2.3.3.

2.3.1 Alternative form of the posterior

Computing optimal designs based on the marginalized posterior covariance operator eq. (2.6) entails computing traces of operators defined on the discretized parameter spaces. The corresponding expressions also include inverses of operators of dimensions n_{m} and n_{b}; see eq. (2.6). The discretized parameter dimensions are typically large and depend on the computational grids used for discretization. In many large scale inverse problems, the dimension n_{d} of the measurement vector $\boldsymbol{y}$ is considerably smaller than the dimension of the discretized uncertain parameters. Also, in our approach, this measurement dimension is fixed a priori. Here we derive an alternative expression for the posterior covariance operator eq. (2.9) that facilitates exploiting this problem structure. In particular, this allows reformulating the marginalized A-optimality criterion in terms of an operator defined on the measurement space, which can then be computed directly (see section 2.4). This is in contrast to previous works such as [5,28,37,39,45] that use randomized trace estimation (in the discretized parameter space) to compute the OED objective.

Theorem 2.3.1. *The following relation holds.*

$$(\mathbf{E}^*\mathbf{W}_\sigma\mathbf{E} + \boldsymbol{\Gamma}_{\mathrm{pr}}^{-1})^{-1} = \boldsymbol{\Gamma}_{\mathrm{pr}} - \boldsymbol{\Gamma}_{\mathrm{pr}}\mathbf{E}^*(\mathbf{I} + \mathbf{W}_\sigma\mathbf{E}\boldsymbol{\Gamma}_{\mathrm{pr}}\mathbf{E}^*)^{-1}\mathbf{W}_\sigma\mathbf{E}\boldsymbol{\Gamma}_{\mathrm{pr}}. \tag{2.12}$$

Proof. First, we need to show that $\mathbf{I} + \mathbf{W}_\sigma\mathbf{E}\boldsymbol{\Gamma}_{\mathrm{pr}}\mathbf{E}^*$ is invertible. To do this, we show that $\mathbf{W}_\sigma\mathbf{E}\boldsymbol{\Gamma}_{\mathrm{pr}}\mathbf{E}^*$ has non-negative eigenvalues. Note that $\boldsymbol{\Gamma}_{\mathrm{pr}} = \boldsymbol{\Gamma}_{\mathrm{pr}}^* = \mathbf{M}^{-1}\boldsymbol{\Gamma}_{\mathrm{pr}}^{\mathsf{T}}\mathbf{M}$. Moreover, we have that $\mathbf{E}^* = \mathbf{M}^{-1}\mathbf{E}^{\mathsf{T}}$. Thus, we have $(\mathbf{E}\boldsymbol{\Gamma}_{\mathrm{pr}}\mathbf{E}^*)^{\mathsf{T}} = (\mathbf{E}^*)^{\mathsf{T}}\boldsymbol{\Gamma}_{\mathrm{pr}}^{\mathsf{T}}\mathbf{E}^{\mathsf{T}} = \mathbf{E}\mathbf{M}^{-1}\boldsymbol{\Gamma}_{\mathrm{pr}}^{\mathsf{T}}\mathbf{M}\mathbf{E}^* = \mathbf{E}\boldsymbol{\Gamma}_{\mathrm{pr}}\mathbf{E}^*$. That is, $\mathbf{E}\boldsymbol{\Gamma}_{\mathrm{pr}}\mathbf{E}^*$ is symmetric; it is also clearly positive semidefinite.

To show that $\mathbf{W}_\sigma\mathbf{E}\boldsymbol{\Gamma}_{\mathrm{pr}}\mathbf{E}^*$ has non-negative eigenvalues, we recall a basic result from linear algebra: if $\mathbf{A}$ and $\mathbf{B}$ are two square matrices, $\mathbf{AB}$ and $\mathbf{BA}$ have the same eigenvalues; see e.g., [66, page 249]. Applying this result with $\mathbf{A} = \mathbf{W}_\sigma^{1/2}\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*$ and $\mathbf{B} = \mathbf{W}_\sigma^{1/2}$, we have that $\mathbf{W}_\sigma^{1/2}\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*\mathbf{W}_\sigma^{1/2}$ and $\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*$ have the same eigenvalues. Therefore, since $\mathbf{W}_\sigma^{1/2}\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*\mathbf{W}_\sigma^{1/2}$ is symmetric positive semidefinite, it follows that $\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*$ has non-negative eigenvalues. This implies that $\mathbf{I} + \mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*$ is invertible. The relation eq. (2.12) is now seen as follows:

$$
\begin{aligned}
&(\mathbf{E}^*\mathbf{W}_\sigma\mathbf{E} + \mathbf{\Gamma}_{\text{pr}}^{-1})(\mathbf{\Gamma}_{\text{pr}} - \mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*(\mathbf{I} + \mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*)^{-1}\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}})\\
&\quad = \mathbf{E}^*\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}} - \mathbf{E}^*\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*(\mathbf{I} + \mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*)^{-1}\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}} + \mathbf{I} - \mathbf{E}^*(\mathbf{I} + \mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*)^{-1}\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\\
&\quad = \mathbf{I} + \mathbf{E}^*\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}} - \mathbf{E}^*(\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^* + \mathbf{I})(\mathbf{I} + \mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\mathbf{E}^*)^{-1}\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}}\\
&\quad = \mathbf{I} + \mathbf{E}^*\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}} - \mathbf{E}^*\mathbf{W}_\sigma\mathbf{E}\mathbf{\Gamma}_{\text{pr}} = \mathbf{I}.
\end{aligned}
$$

□

Notice that this result is well known in the case $\mathbf{W}_\sigma = \mathbf{\Gamma}_{\text{noise}}^{-1}$. The challenge here is to account for the possibility of a singular $\mathbf{W}_\sigma$. Note that the expression in the left hand side of eq. (2.12) involves the inverse of an $n \times n$ matrix, where $n = n_{\text{m}} + n_{\text{b}}$, whereas the expression on the right hand side involves the inverse of an $n_{\text{d}} \times n_{\text{d}}$ matrix. It is also worth noting that the proof of theorem 2.3.1 can be simplified by the use of the Sherman–Morrison–Woodbury formula. Above, we chose to present a direct linear algebra argument instead, for clarity.

We introduce the following notations, which will be used in the remainder of this chapter.

$$
\mathbf{Q}(\boldsymbol{w}) := (\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\mathbf{W}_\sigma, \quad \text{where} \quad \mathbf{C} := \mathbf{F}\mathbf{\Gamma}_{\text{pr,m}}\mathbf{F}^* + \mathbf{G}\mathbf{\Gamma}_{\text{pr,b}}\mathbf{G}^*. \tag{2.13}
$$

Next, we present tractable representations for the posterior mean and covariance operator in a (discretized) Bayesian linear inverse problem, as formulated in section 2.2.4. Recall that the primary parameter is $\boldsymbol{m}$ and the secondary parameter is $\boldsymbol{b}$.

Theorem 2.3.2. *The posterior law of* $\begin{bmatrix}\boldsymbol{m}\\ \boldsymbol{b}\end{bmatrix}$ *is* $\mathcal{N}\left(\begin{bmatrix}\boldsymbol{m}_{\text{MAP}}\\ \boldsymbol{b}_{\text{MAP}}\end{bmatrix}, \begin{bmatrix}\mathbf{\Gamma}_{\text{post,m}}(\boldsymbol{w}) & \mathbf{\Gamma}_{\text{post,mb}}(\boldsymbol{w})\\ \mathbf{\Gamma}^*_{\text{post,mb}}(\boldsymbol{w}) & \mathbf{\Gamma}_{\text{post,b}}(\boldsymbol{w})\end{bmatrix}\right)$, *where*

$$
\begin{aligned}
\mathbf{\Gamma}_{\text{post,m}}(\boldsymbol{w}) &= \mathbf{\Gamma}_{\text{pr,m}} - \mathbf{\Gamma}_{\text{pr,m}}\mathbf{F}^*\mathbf{Q}(\boldsymbol{w})\mathbf{F}\mathbf{\Gamma}_{\text{pr,m}},\\
\mathbf{\Gamma}_{\text{post,b}}(\boldsymbol{w}) &= \mathbf{\Gamma}_{\text{pr,b}} - \mathbf{\Gamma}_{\text{pr,b}}\mathbf{G}^*\mathbf{Q}(\boldsymbol{w})\mathbf{G}\mathbf{\Gamma}_{\text{pr,b}},\\
\mathbf{\Gamma}_{\text{post,mb}}(\boldsymbol{w}) &= -\mathbf{\Gamma}_{\text{pr,m}}\mathbf{F}^*\mathbf{Q}(\boldsymbol{w})\mathbf{G}\mathbf{\Gamma}_{\text{pr,b}},\\
\boldsymbol{m}_{\text{MAP}}(\boldsymbol{w}) &= \mathbf{\Gamma}_{\text{post,m}}(\boldsymbol{w})(\mathbf{F}^*\mathbf{W}_\sigma\boldsymbol{y} + \mathbf{\Gamma}_{\text{pr,m}}^{-1}\boldsymbol{m}_{\text{pr}}) + \mathbf{\Gamma}_{\text{post,mb}}(\boldsymbol{w})(\mathbf{G}^*\mathbf{W}_\sigma\boldsymbol{y} + \mathbf{\Gamma}_{\text{pr,b}}^{-1}\boldsymbol{b}_{\text{pr}}),\\
\boldsymbol{b}_{\text{MAP}}(\boldsymbol{w}) &= \mathbf{\Gamma}_{\text{post,b}}(\boldsymbol{w})(\mathbf{G}^*\mathbf{W}_\sigma\boldsymbol{y} + \mathbf{\Gamma}_{\text{pr,b}}^{-1}\boldsymbol{b}_{\text{pr}}) + \mathbf{\Gamma}^*_{\text{post,mb}}(\boldsymbol{w})(\mathbf{F}^*\mathbf{W}_\sigma\boldsymbol{y} + \mathbf{\Gamma}_{\text{pr,m}}^{-1}\boldsymbol{m}_{\text{pr}}).
\end{aligned}
\tag{2.14}
$$

Proof. Recall that the discretized forward operator $\mathbf{E}$ can be represented in a block matrix form $\mathbf{E} = \begin{bmatrix}\mathbf{F} & \mathbf{G}\end{bmatrix}$. Using this and the expression for $\mathbf{\Gamma}_{\text{post}}$ given in theorem 2.3.1, we obtain

$$\begin{aligned}
\mathbf{\Gamma}_{\text{post}}(\boldsymbol{w}) &= \mathbf{\Gamma}_{\text{pr}} - \mathbf{\Gamma}_{\text{pr}} \begin{bmatrix}\mathbf{F}^* \\ \mathbf{G}^*\end{bmatrix} \left(\mathbf{I} + \mathbf{W}_\sigma \begin{bmatrix}\mathbf{F} & \mathbf{G}\end{bmatrix} \mathbf{\Gamma}_{\text{pr}} \begin{bmatrix}\mathbf{F}^* \\ \mathbf{G}^*\end{bmatrix}\right)^{-1} \mathbf{W}_\sigma \begin{bmatrix}\mathbf{F} & \mathbf{G}\end{bmatrix} \mathbf{\Gamma}_{\text{pr}} \\
&= \mathbf{\Gamma}_{\text{pr}} - \mathbf{\Gamma}_{\text{pr}} \begin{bmatrix}\mathbf{F}^* \\ \mathbf{G}^*\end{bmatrix} (\mathbf{I} + \mathbf{W}_\sigma(\mathbf{F}\mathbf{\Gamma}_{\text{pr,m}}\mathbf{F}^* + \mathbf{G}\mathbf{\Gamma}_{\text{pr,b}}\mathbf{G}^*))^{-1} \mathbf{W}_\sigma \begin{bmatrix}\mathbf{F} & \mathbf{G}\end{bmatrix} \mathbf{\Gamma}_{\text{pr}} \\
&= \begin{bmatrix}\mathbf{\Gamma}_{\text{pr,m}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Gamma}_{\text{pr,b}}\end{bmatrix} - \begin{bmatrix}\mathbf{\Gamma}_{\text{pr,m}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Gamma}_{\text{pr,b}}\end{bmatrix} \begin{bmatrix}\mathbf{F}^* \\ \mathbf{G}^*\end{bmatrix} \mathbf{Q}(\boldsymbol{w}) \begin{bmatrix}\mathbf{F} & \mathbf{G}\end{bmatrix} \begin{bmatrix}\mathbf{\Gamma}_{\text{pr,m}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Gamma}_{\text{pr,b}}\end{bmatrix} \\
&= \begin{bmatrix}\mathbf{\Gamma}_{\text{pr,m}} - \mathbf{\Gamma}_{\text{pr,m}}\mathbf{F}^*\mathbf{Q}(\boldsymbol{w})\mathbf{F}\mathbf{\Gamma}_{\text{pr,m}} & -\mathbf{\Gamma}_{\text{pr,m}}\mathbf{F}^*\mathbf{Q}(\boldsymbol{w})\mathbf{G}\mathbf{\Gamma}_{\text{pr,b}} \\ -\mathbf{\Gamma}_{\text{pr,b}}\mathbf{G}^*\mathbf{Q}(\boldsymbol{w})\mathbf{F}\mathbf{\Gamma}_{\text{pr,m}} & \mathbf{\Gamma}_{\text{pr,b}} - \mathbf{\Gamma}_{\text{pr,b}}\mathbf{G}^*\mathbf{Q}(\boldsymbol{w})\mathbf{G}\mathbf{\Gamma}_{\text{pr,b}}\end{bmatrix}.
\end{aligned} \tag{2.15}$$

This establishes the representation of the posterior covariance operator. The expressions for $\boldsymbol{m}_{\text{MAP}}(\boldsymbol{w})$ and $\boldsymbol{b}_{\text{MAP}}(\boldsymbol{w})$ can be obtained using eq. (2.9) and eq. (2.15). $\square$

Using lemma 2.2.3 in conjunction with theorem 2.3.2, the marginal posterior laws of $\boldsymbol{m}$ and $\boldsymbol{b}$ are given by $\mathcal{N}(\boldsymbol{m}_{\text{MAP}}(\boldsymbol{w}), \mathbf{\Gamma}_{\text{post,m}}(\boldsymbol{w}))$ and $\mathcal{N}(\boldsymbol{b}_{\text{MAP}}(\boldsymbol{w}), \mathbf{\Gamma}_{\text{post,b}}(\boldsymbol{w}))$, respectively. Since we target the primary parameter $\boldsymbol{m}$, we focus on the corresponding marginal posterior law $\mathcal{N}(\boldsymbol{m}_{\text{MAP}}(\boldsymbol{w}), \mathbf{\Gamma}_{\text{post,m}}(\boldsymbol{w}))$. The marginal covariance operator $\mathbf{\Gamma}_{\text{post,m}}(\boldsymbol{w})$ will be used to define the marginalized A-optimality criterion (see below). Also, note that the expression for $\boldsymbol{m}_{\text{MAP}}$ in eq. (2.14) is the sum of two terms: the first is the familiar expression for the posterior mean if $\boldsymbol{b}$ was fixed to $\boldsymbol{b} = \mathbf{0}$; the second reflects the impact of the uncertainty in $\boldsymbol{b}$.

2.3.2 The marginalized A-optimality criterion

The marginalized A-optimal design (mOED) criterion is given by

$$\Phi(\boldsymbol{w}) := \mathsf{tr}(\mathbf{\Gamma}_{\text{post,m}}(\boldsymbol{w})) = \mathsf{tr}(\mathbf{\Gamma}_{\text{pr,m}}) - \mathsf{tr}(\mathbf{\Gamma}_{\text{pr,m}}\mathbf{F}^*\mathbf{Q}(\boldsymbol{w})\mathbf{F}\mathbf{\Gamma}_{\text{pr,m}}). \tag{2.16}$$

Next, we show the convexity of the mOED objective. Before proving this, we consider a slightly more general result. Below, $S^{\mathbf{M}}_{++}$ denotes the cone of self-adjoint and positive definite operators on $\mathbb{R}^n$ equipped with the weighted inner product $\langle \cdot, \cdot \rangle_{\mathbb{M}}$.

Theorem 2.3.3. *Let the function* $f : \mathbb{R}^{n_s}_{\geq 0} \to \mathbb{R}$ *be given by*

$$f(\boldsymbol{w}) = \mathsf{tr}(\mathbf{R}\mathbf{\Gamma}_{\text{post}}(\boldsymbol{w})\mathbf{R}^*),$$

where $\mathbf{R}$ *is an* $n \times n$ *matrix and* $\mathbf{R}^*$ *denotes its adjoint with respect to* $\langle\cdot,\cdot\rangle_{\mathbb{M}}$*. Then, the function* f *is convex.*

Proof. Let $\mathbf{A}(\boldsymbol{w}) = \boldsymbol{\Gamma}_{\text{post}}(\boldsymbol{w})^{-1}$, and note that $\mathbf{A}(\boldsymbol{w}) \in S_{++}^{\mathbf{M}}$ for all $\boldsymbol{w} \in \mathbb{R}_{\geq 0}^{n_s}$. First we show the function $G(\mathbf{A}) = \mathsf{tr}(\mathbf{R}\mathbf{A}^{-1}\mathbf{R}^*)$ is convex on $S_{++}^{\mathbf{M}}$. Consider the restriction of G to a line, $\mathbf{S} + t\mathbf{B}$, where $\mathbf{S} \in S_{++}^{\mathbf{M}}$ and $\mathbf{B}$ is self-adjoint; we consider values of t for which $\mathbf{S} + t\mathbf{B} \in S_{++}^{\mathbf{M}}$. Let $\mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^*$ be the spectral decomposition of $\mathbf{V} = \mathbf{S}^{-1/2}\mathbf{B}\mathbf{S}^{-1/2}$; here $\boldsymbol{\Lambda}$ is a diagonal matrix with the eigenvalues $\{\lambda_i\}_{i=1}^n$ of $\mathbf{V}$ on its diagonal and $\mathbf{U}$ is a matrix with the corresponding eigenvectors $\{\boldsymbol{u}_i\}_{i=1}^n$ as its columns. Letting $\mathbf{L} = \mathbf{S}^{-1/2}\mathbf{R}^*$, we note

$$\begin{aligned} G(\mathbf{S} + t\mathbf{B}) &= \mathsf{tr}(\mathbf{R}\mathbf{S}^{-1/2}(\mathbf{I} + t\mathbf{S}^{-1/2}\mathbf{B}\mathbf{S}^{-1/2})^{-1}\mathbf{S}^{-1/2}\mathbf{R}^*) \\ &= \mathsf{tr}(\mathbf{L}\mathbf{L}^*(\mathbf{I} + t\mathbf{V})^{-1}) = \sum_{i=1}^{n} \left\langle \mathbf{L}\mathbf{L}^*(\mathbf{I} + t\mathbf{V})^{-1}\boldsymbol{u}_i, \boldsymbol{u}_i \right\rangle_{\mathbb{M}} = \sum_{i=1}^{n} (1 + t\lambda_i)^{-1} \langle \mathbf{L}^*\boldsymbol{u}_i, \mathbf{L}^*\boldsymbol{u}_i \rangle_{\mathbb{M}}. \end{aligned}$$

Thus, $G(\mathbf{S} + t\mathbf{B})$ is a linear combination of convex functions with non-negative coefficients, $\langle \mathbf{L}^*\boldsymbol{u}_i, \mathbf{L}^*\boldsymbol{u}_i \rangle_{\mathbb{M}} \geq 0$, and is thus convex. This shows that G is convex on $S_{++}^{\mathbf{M}}$. It remains to show that $f(\boldsymbol{w}) = G(\mathbf{A}(\boldsymbol{w}))$ is convex. Recall that $\mathbf{A}(\boldsymbol{w}) = \boldsymbol{\Gamma}_{\text{pr}}^{-1} + \mathbf{E}^*\mathbf{W}_\sigma\mathbf{E}$; thus $\mathbf{A}$ is affine in $\boldsymbol{w}$ and therefore, for $\alpha \in [0, 1]$,

$$\begin{aligned} f(\alpha\boldsymbol{w} + (1-\alpha)\boldsymbol{v}) = G(\mathbf{A}(\alpha\boldsymbol{w} + (1-\alpha)\boldsymbol{v})) &= G(\alpha\mathbf{A}(\boldsymbol{w}) + (1-\alpha)\mathbf{A}(\boldsymbol{v})) \\ &\leq \alpha G(\mathbf{A}(\boldsymbol{w})) + (1-\alpha)G(\mathbf{A}(\boldsymbol{v})) = \alpha f(\boldsymbol{w}) + (1-\alpha)f(\boldsymbol{v}). \end{aligned}$$

□

Corollary 2.3.1. *The function* $\Phi : \mathbb{R}_{\geq 0}^{n_d} \to \mathbb{R}$*, defined in eq.* (2.16)*, is convex.*

Proof. Using eq. (2.15), we can write $\Phi(\boldsymbol{w})$ as

$$\Phi(\boldsymbol{w}) = \mathsf{tr}(\mathbf{R}\boldsymbol{\Gamma}_{\text{post}}(\boldsymbol{w})\mathbf{R}^*) \quad \text{with} \quad \mathbf{R} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

Thus, the convexity of $\Phi(\boldsymbol{w})$ can be concluded from Proposition 2.3.3. □

Consider the marginalized A-optimality criterion $\Phi(\boldsymbol{w})$ in eq. (2.16). Since the prior covariance operator is independent of $\boldsymbol{w}$, minimizing $\Phi(\boldsymbol{w})$ is equivalent to minimizing

$$\Psi(\boldsymbol{w}) := -\mathsf{tr}(\mathbf{F}\boldsymbol{\Gamma}_{\text{pr,m}}^2\mathbf{F}^*\mathbf{Q}(\boldsymbol{w})). \tag{2.17}$$

This is the objective function we use in finding a marginalized A-optimal design. Henceforth, we refer to this objective function as the mOED objective or the mOED criterion.

2.3.3 Computing optimal designs

Here we describe the optimization problem for computing mOEDs. The ultimate goal is to find a binary optimal design vector that minimizes the mOED objective Ψ, defined in eq. (2.17). That is, letting $\mathcal{X} = \{0,1\}^{n_{\mathrm{d}}}$, we would like to solve

$$\min_{\boldsymbol{w}\in\mathcal{X}} \Psi(\boldsymbol{w}), \quad \text{s.t.} \sum_{i=1}^{n_{\mathrm{d}}} w_i = N, \tag{2.18}$$

where N is a desired number of sensors. However, as mentioned above, solving such a binary optimization problem can be intractable due to its combinatorial complexity. One possible way to find an approximate solution to (2.18) is via a greedy procedure, i.e., place sensors one-by-one. This method does not require derivatives of the objective with respect to weights. Generally, greedy approaches result in suboptimal solutions, which, in practice, are often quite good. Computational details of this approach are discussed in section 2.4.3. We also compare, in section 2.6.1, the performance of the greedy approach against the approach described next.

As an alternative to the greedy approach, one can consider a relaxation of the problem and allow for design weights in the interval $[0,1]$. Binary weights are then obtained using sparsifying penalty functions. Specifically, we consider an optimization problem of the form

$$\min_{\boldsymbol{w}\in\mathcal{W}} \Psi(\boldsymbol{w}) + \gamma P(\boldsymbol{w}), \tag{2.19}$$

where $\mathcal{W} = [0,1]^{n_{\mathrm{d}}}$, $\Psi(\boldsymbol{w})$ is the mOED objective, $\gamma > 0$ is a penalty parameter, and $P(\boldsymbol{w})$ is a penalty function. Minimization of (2.19) usually requires gradients of the objective. Key computational aspects are discussed in the next section where we outline computational methods for tackling the mOED problem.

2.4 Computational methods

In this section, we present a computational framework for computing mOEDs.

2.4.1 Efficient computation of mOED objective and its gradient

Consider the objective function $\Psi(\boldsymbol{w})$ defined in eq. (2.17). We note that the argument of the trace in eq. (2.17) is an operator defined on $\mathbb{R}^{n_{\mathrm{d}}\times n_{\mathrm{d}}}$, where n_{d} is the number of candidate sensor locations (i.e., the dimension of the measurement vector). This objective function can be computed as follows:

$$\Psi(\boldsymbol{w}) = -\sum_{i=1}^{n_\mathrm{d}} \boldsymbol{e}_i{}^\mathsf{T}\mathbf{D}\mathbf{Q}(\boldsymbol{w})\boldsymbol{e}_i = -\sum_{i=1}^{n_\mathrm{d}} \boldsymbol{e}_i^\mathsf{T}\mathbf{D}\boldsymbol{q}_i, \quad \text{where } \mathbf{D} = \mathbf{F}\boldsymbol{\Gamma}^2_{\mathrm{pr,m}}\mathbf{F}^*, \tag{2.20}$$

$\boldsymbol{q}_i = \mathbf{Q}(\boldsymbol{w})\boldsymbol{e}_i$ with $\mathbf{Q}(\boldsymbol{w})$ given in eq. (2.13), and $\boldsymbol{e}_i$ is the ith standard basis vector in $\mathbb{R}^{n_\mathrm{d}}$, $i = 1, \ldots, n_\mathrm{d}$. Note that

$$\boldsymbol{q}_i = (\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\mathbf{W}_\sigma\boldsymbol{e}_i = \sigma_i^{-2}w_i(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\boldsymbol{e}_i. \tag{2.21}$$

To derive the expression for the gradient of Ψ, we first need the following derivative:

$$\frac{\partial}{\partial w_j}\mathbf{Q}(\boldsymbol{w}) = -\sigma_j^{-2}\left(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C}\right)^{-1}(\boldsymbol{e}_j\boldsymbol{e}_j^\mathsf{T})\mathbf{C}\left(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C}\right)^{-1}\mathbf{W}_\sigma + \sigma_j^{-2}\left(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C}\right)^{-1}\boldsymbol{e}_j\boldsymbol{e}_j^\mathsf{T}.$$

Thus,

$$\begin{aligned}
\frac{\partial\Psi}{\partial w_j} &= -\frac{\partial}{\partial w_j}\mathsf{tr}(\mathbf{Q}(\boldsymbol{w})\mathbf{D}) \\
&= \mathsf{tr}\left[\sigma_j^{-2}\left(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C}\right)^{-1}\boldsymbol{e}_j\boldsymbol{e}_j^\mathsf{T}\mathbf{C}\left(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C}\right)^{-1}\mathbf{W}_\sigma\mathbf{D}\right] - \mathsf{tr}\left[\sigma_j^{-2}\left(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C}\right)^{-1}\boldsymbol{e}_j\boldsymbol{e}_j^\mathsf{T}\mathbf{D}\right] \\
&= \sigma_j^{-2}\boldsymbol{e}_j^\mathsf{T}\mathbf{C}(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\mathbf{W}_\sigma\mathbf{D}(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\boldsymbol{e}_j - \sigma_j^{-2}\boldsymbol{e}_j^\mathsf{T}\mathbf{D}(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\boldsymbol{e}_j \\
&= \sum_{i=1}^{n_\mathrm{d}} w_i\sigma_i^{-2}\sigma_j^{-2}\boldsymbol{e}_j^\mathsf{T}\mathbf{C}(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\boldsymbol{e}_i\boldsymbol{e}_i^\mathsf{T}\mathbf{D}(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\boldsymbol{e}_j - \sigma_j^{-2}\boldsymbol{e}_j^\mathsf{T}\mathbf{D}(\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\boldsymbol{e}_j,
\end{aligned}$$

where we have used the cyclic property of the trace and the definition of $\mathbf{W}_\sigma$ in eq. (2.8). Letting $\boldsymbol{y}_i = (\mathbf{I} + \mathbf{W}_\sigma\mathbf{C})^{-1}\boldsymbol{e}_i$, $i = 1, \ldots, n_\mathrm{d}$, and substituting in the above expression, leads to

$$\frac{\partial\Psi}{\partial w_j} = \sum_{i=1}^{n_\mathrm{d}} w_i\sigma_i^{-2}\sigma_j^{-2}(\boldsymbol{e}_j^\mathsf{T}\mathbf{C}\boldsymbol{y}_i)\boldsymbol{e}_i^\mathsf{T}\mathbf{D}\boldsymbol{y}_j - \sigma_j^{-2}\boldsymbol{e}_j^\mathsf{T}\mathbf{D}\boldsymbol{y}_j, \quad j = 1, \ldots, n_\mathrm{d}. \tag{2.22}$$

Note that the vectors $\boldsymbol{q}_i$ in eq. (2.21) and vectors $\boldsymbol{y}_i$ in the definition of the gradient are related according to $\boldsymbol{q}_i = w_i\sigma_i^{-2}\boldsymbol{y}_i$, $i = 1, \ldots, n_\mathrm{d}$.

The matrices $\mathbf{C}$ and $\mathbf{D}$ in eqs. (2.21) and (2.22) are of size $n_\mathrm{d} \times n_\mathrm{d}$. As mentioned previously, in many cases, the measurement dimension n_d is considerably smaller than the dimension of the discretized primary and secondary parameters. This case typically arises in inverse problems governed by PDEs, where the dimension of the discretized parameters grow upon grid refinements, while the measurement dimension n_d is fixed a priori.

The matrices $\mathbf{C}$ and $\mathbf{D}$ can be built in a precomputation step, as outlined in algorithm 1. The computational cost to build $\mathbf{C}$ and $\mathbf{D}$ is $3n_\mathrm{d}$ forward and $2n_\mathrm{d}$ adjoint PDE solves. Once the matrices $\mathbf{C}$ and $\mathbf{D}$ are computed, the OED objective and gradient evaluation can

Algorithm 1 Computing matrices $\mathbf{C}$ in eq. (2.13) and $\mathbf{D}$ in eq. (2.20) needed for mOED objective and gradient evaluation.

```
for $i = 1$ to $n_\mathrm{d}$ do
    Compute $\boldsymbol{a}_i = \boldsymbol{\Gamma}_{\mathrm{pr,m}} \mathbf{F}^* \boldsymbol{e}_i$
    Compute $\boldsymbol{d}_i = \mathbf{F} \boldsymbol{\Gamma}_{\mathrm{pr,m}} \boldsymbol{a}_i$    {columns of $\mathbf{D} = \mathbf{F}\boldsymbol{\Gamma}_{\mathrm{pr,m}}^2 \mathbf{F}^*$}
    Compute $\boldsymbol{c}_i = \mathbf{F}\boldsymbol{a}_i + \mathbf{G}\boldsymbol{\Gamma}_{\mathrm{pr,b}} \mathbf{G}^* \boldsymbol{e}_i$    {columns of $\mathbf{C} = \mathbf{F}\boldsymbol{\Gamma}_{\mathrm{pr,m}} \mathbf{F}^* + \mathbf{G}\boldsymbol{\Gamma}_{\mathrm{pr,b}} \mathbf{G}^*$}
end for
Build $\mathbf{C} = [\boldsymbol{c}_1 \ \cdots \ \boldsymbol{c}_{n_\mathrm{d}}]$ and $\mathbf{D} = [\boldsymbol{d}_1 \ \cdots \ \boldsymbol{d}_{n_\mathrm{d}}]$
```

be performed without further PDE solves and require only linear algebra operations; see algorithm 2. The cost of evaluating the objective function is dominated by the cost of steps 1–3, which amount to computing $\mathbf{Y} = (\mathbf{I} + \mathbf{W}_\sigma \mathbf{C})^{-1}$; this can be done in $\mathcal{O}(n_\mathrm{d}^3)$ arithmetic operations, by precomputing an LU factorization of $\mathbf{I} + \mathbf{W}_\sigma \mathbf{C}$ and then performing triangular solves to compute columns of $\mathbf{Y}$. We also need the matrix-matrix product $\mathbf{DY}$ (see step 5 of algorithm 2), which requires an additional $\mathcal{O}(n_d^3)$ operations. The additional effort in computing the gradient is dominated by one matrix-matrix product, $\mathbf{CY}$, amounting to $\mathcal{O}(n_\mathrm{d}^3)$ arithmetic operations.

Algorithm 2 Computing $\Psi(\boldsymbol{w})$ and its gradient $\nabla\Psi(\boldsymbol{w})$.

```
Input: Design vector $\boldsymbol{w}$
Output: $\Psi = \Psi(\boldsymbol{w})$ and $\nabla\Psi = \nabla\Psi(\boldsymbol{w})$
1: /* evaluation of the objective function */
2: for $i = 1$ to $n_\mathrm{d}$ do
3:     Solve the system $(\mathbf{I} + \mathbf{W}_\sigma \mathbf{C})\boldsymbol{y}_i = \boldsymbol{e}_i$
4: end for
5: Compute $\Psi = -\sum_{i=1}^{n_\mathrm{d}} w_i \sigma_i^{-2} \boldsymbol{e}_i^\mathsf{T} \mathbf{D}\mathbf{Y} \boldsymbol{e}_i$    {$\mathbf{Y} = [\boldsymbol{y}_1 \ \boldsymbol{y}_2 \ \cdots \ \boldsymbol{y}_{n_\mathrm{d}}]$}
6: /* evaluation of the gradient */
7: for $j = 1$ to $n_\mathrm{d}$ do
8:     Compute $\frac{\partial \Psi}{\partial w_j} = \sum_{i=1}^{n_\mathrm{d}} w_i \sigma_i^{-2} \sigma_j^{-2} (\boldsymbol{e}_j^\mathsf{T} \mathbf{C}\mathbf{Y} \boldsymbol{e}_i)(\boldsymbol{e}_i^\mathsf{T} \mathbf{D}\mathbf{Y} \boldsymbol{e}_j) - \sigma_j^{-2} \boldsymbol{e}_j^\mathsf{T} \mathbf{D}\mathbf{Y} \boldsymbol{e}_j$
9: end for
```

2.4.2 Sparsity control

Here we discuss several options for choosing the penalty function $P(\boldsymbol{w})$ in eq. (2.19). A straightforward choice for $P(\boldsymbol{w})$ is the ℓ_1-norm of $\boldsymbol{w}$; see e.g., [37, 39]. As is well-known, the ℓ_1-penalty promotes sparsity, but not necessarily a binary structure, in the computed design vectors. Another option is to solve a sequence of optimization problems where penalty functions approximating ℓ_0-"norm" (the number of nonzero elements in a vector) are used. An example is the so-called regularized ℓ_0-sparsification approach proposed in [5]; in this approach, which we use in this chapter, a continuation approach is used, and a sequence of optimization problems, with non-convex penalty functions approaching the ℓ_0-norm, are solved. A related approach is the use of reweighted ℓ_1-minimization, as done in [45]. Solving optimization problems with continuous weights, combined with a suitable penalty method, enables the use of powerful gradient-based optimization algorithms to explore the set of admissible designs. The effectiveness of such approaches in obtaining optimal sensor placements has been demonstrated in a number of previous works; see e.g., [5, 37, 39, 45].

2.4.3 Greedy sensor placement

An alternative approach for finding sparse mOEDs is to use a greedy strategy. Greedy approaches have been used successfully in many sensor placement applications to obtain designs that, while suboptimal, provide near optimal performance; see e.g., [21,51,57,75]. In a greedy approach, we place sensors one at a time: in each step, we select a sensor that provides the largest decrease in the design criterion. A greedy approach can be attractive due to its simplicity and the fact that it does not require the gradient of the design criterion. However, the computational complexity of greedy sensor placement, in terms of function evaluations, scales with the number of candidate sensor locations and the number of the sensors in the optimal design. Note that the computational cost, in terms of function evaluations, of finding a greedy sensor placement (in its most basic form) with K sensors is

$$C(K, n_{\mathrm{d}}) = K n_{\mathrm{d}} - (K-1)K/2. \tag{2.23}$$

2.5 Model problem setup

To illustrate our approach for computing optimal designs under reducible uncertainty, we consider a linear inverse problem governed by a time-dependent advection-diffusion equation with two sources of uncertainty: the parameter of primary interest is a time-dependent scalar-valued function $m = m(t)$, which models the time amplitude of a source entering on the right hand side of the equation. The second uncertain parameter is the spatially distributed initial condition $b = b(\boldsymbol{x})$. Specifically, we consider:

$$u_t - \kappa\Delta u + \boldsymbol{v} \cdot \nabla u = \delta(\boldsymbol{x})m(t) \quad \text{in } \mathcal{D} \times \mathcal{T}, \tag{2.24a}$$

$$u(\cdot, 0) = b(\boldsymbol{x}) \quad \text{in } \mathcal{D}, \tag{2.24b}$$

$$\kappa\nabla u \cdot \boldsymbol{n} = 0 \quad \text{on } \partial\mathcal{D} \times \mathcal{T}. \tag{2.24c}$$

Here, $\mathcal{D}$ is a bounded open set in $\mathbb{R}^2$, the time interval $\mathcal{T} = (0, T)$, where $T > 0$ is a final time, $\kappa > 0$ is the diffusion coefficient, and $\boldsymbol{v}$ is a given velocity field. Note that the solution $u(\boldsymbol{x}, t)$, which can be interpreted as concentration, depends affinely on m and b. In our numerical experiments, $\kappa = 0.001$ and $\mathcal{D}$ is a unit square with two cutouts as shown in fig. 2.1 (left). If (2.24a) models the flow of a contaminant in a region, the cutouts could represent buildings, for instance. The velocity field $\boldsymbol{v}$ (shown in fig. 2.1) is obtained by solving Navier-Stokes equations with no-outflow boundary conditions and non-zero tangential boundary conditions as in [5]. The function δ in the source term is given by a mollified delta-function:

$$\delta(\boldsymbol{x}) = \left(\frac{1}{2\pi L} e^{-\frac{1}{2L^2}\|\boldsymbol{x}-\boldsymbol{x}_0\|^2}\right), \tag{2.25}$$

where the "correlation length" L is 0.05 in our experiments, and $\boldsymbol{x}_0 = (0.5, 0.35)$ as indicated by the red dot in fig. 2.1 (left).

2.5.1 Parameter-to-observable map

The parameter-to-observable map maps the time evolution of the right hand side amplitude, $m \in L^2(\mathcal{T})$ and the initial condition $b \in L^2(\mathcal{D})$ to point measurements of the solution of the advection-diffusion equation eq. (2.24). To write the parameter-to-observable map in the form eq. (2.2), we define the continuous linear operators $\mathcal{S}_1$ and $\mathcal{S}_2$ as follows: $\mathcal{S}_1$ maps m to the PDE solution u, with $b = 0$, and $\mathcal{S}_2$ maps b to the PDE solution u, with $m = 0$.

Then, the solution to the initial-boundary value problem eq. (2.24) can be written as $u = \mathcal{S}_1 m + \mathcal{S}_2 b$; see [83, p.152]. Next, let $\mathcal{B}$ be a linear observation operator that extracts the values of $u(\boldsymbol{x}, t)$ on a set of sensor locations $\{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_{n_\text{d}}\} \in \mathcal{D}$, and takes an average of u over the time interval $[0.95, 0.99]$. Then $\mathcal{F} = \mathcal{B}\mathcal{S}_1$ and $\mathcal{G} = \mathcal{B}\mathcal{S}_2$ map the primary inference parameter m and the additional uncertain parameter b to measurement $\boldsymbol{y} \in \mathbb{R}^{n_\text{d}}$:

$$\mathcal{F} : m(t) \xmapsto{\mathcal{S}_1} u(\boldsymbol{x}, t) \xmapsto{\mathcal{B}} \boldsymbol{y}, \qquad \mathcal{G} : b(\boldsymbol{x}) \xmapsto{\mathcal{S}_2} u(\boldsymbol{x}, t) \xmapsto{\mathcal{B}} \boldsymbol{y}. \tag{2.26}$$

The corresponding discrete parameter-to-observable maps $\mathbf{F}$ and $\mathbf{G}$ are obtained through

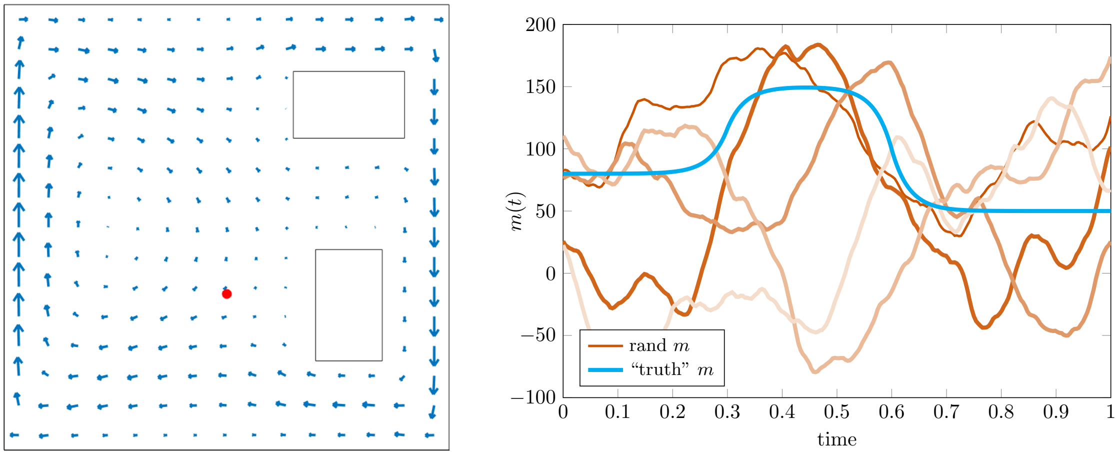

Figure 2.1: Left: Sketch of domain $\mathcal{D}$ and velocity field $\boldsymbol{v}$ in (2.24). The red dot indicates the location $\boldsymbol{x}_0 = (0.5, 0.35)$ where the source term (2.25) is centered. Right: the “truth” source term m and five samples from the prior distribution of m shown in cyan and various shades of orange, respectively.

discretization using, for instance, finite elements.

Computations of derivatives of an objective that involves the parameter-to-observable map requires the adjoint operators $\mathcal{F}^*$ and $\mathcal{G}^*$. These can be derived using the formal Lagrangian method, resulting in the following adjoint equations [83]. Given a vector of observations $\boldsymbol{y} \in \mathbb{R}^{n_\text{d}}$, we first solve the adjoint equation (see [1, 5]) for the adjoint variable $p = p(\boldsymbol{x}, t)$

$$-p_t - \nabla \cdot (p\boldsymbol{v}) - \kappa \Delta p = -\mathcal{B}^* \boldsymbol{y} \quad \text{in } \mathcal{D} \times \mathcal{T}, \tag{2.27a}$$

$$p(\cdot, T) = 0 \quad \text{in } \mathcal{D}, \tag{2.27b}$$

$$(\boldsymbol{v}p + \kappa \nabla p) \cdot \boldsymbol{n} = 0 \quad \text{on } \partial\mathcal{D} \times \mathcal{T}, \tag{2.27c}$$

and obtain the action of the adjoint operators as $\mathcal{F}^*\boldsymbol{y} = -\int_{\mathcal{D}} f(\boldsymbol{x}) p(\boldsymbol{x}, \cdot) d\boldsymbol{x}$ and $\mathcal{G}^*\boldsymbol{y} = -p(\cdot, 0)$.

2.5.2 Prior laws of m and b

To complete the definition of the Bayesian inverse problem, we specify the prior laws for m and b. We assume both to be Gaussian random fields, and thus it is sufficient to specify the mean and covariance operator. For the primary parameter m, which is a function of time only, we choose the mean to be the constant function $m_{\text{pr}} \equiv 65$, and specify the covariance operator $\Gamma_{\mathrm{pr,m}}$ according to

$$[\Gamma_{\mathrm{pr,m}} z](t) = \int_{\mathcal{T}} c(s,t) z(t)\, dt, \quad z \in L^2(\mathcal{T}),$$

where we chose the Matérn-3/2 covariance kernel

$$c(s,t) = \sigma^2 \left(1 + \frac{\sqrt{3}|s-t|}{\ell}\right) \exp\left(-\frac{\sqrt{3}|s-t|}{\ell}\right). \tag{2.28}$$

This covariance function ensures that draws from the prior law of m are (almost surely) continuously differentiable; see, e.g., [40, 58, 87]. In our numerical experiments, we use the parameters $\sigma = 80$ and $\ell = 0.17$ in (2.28). Samples from the resulting distribution are shown in fig. 2.1 (right).

The realizations of the secondary parameter b are functions defined over the spatial domain $\mathcal{D}$. For the distribution of b we choose a Gaussian with mean $b_{\mathrm{pr}} \equiv 50$, and a Laplacian-like covariance operator of the form $(-\epsilon\Delta + \alpha I)^{-2}$ [62], with $\epsilon = 4.5 \times 10^{-3}$ and $\alpha = 2.2 \times 10^{-1}$. We equip the Laplace operator with homogeneous Robin boundary conditions with a constant coefficient. We do this to mitigate undesired boundary effects that can arise when PDE operators are used to define covariance operators [23, 73].

2.5.3 Discretization

We discretize the forward problem using linear finite elements on triangular meshes in space and use the implicit Euler method in time. This guides the discretization of the primary and secondary uncertainties m and b. Specifically, the discretized uncertain source terms is the vector $\boldsymbol{m}$ whose entries are the values of m at the time-steps used by the forward solver. We discretize the $L^2(\mathcal{T})$ inner product using quadrature. That is, for $f, g \in L^2(\mathcal{T})$,

$$\langle f, g \rangle_1 = \int_{\mathcal{T}} f(t) g(t)\, dt \approx \sum_{j=1}^{n_{\mathrm{m}}} \nu_j f(t_j) g(t_j) = \boldsymbol{f}^{\mathsf{T}} \mathbf{M}_1 \boldsymbol{g} =: \langle \boldsymbol{f}, \boldsymbol{g} \rangle_{\mathrm{M}_1},$$

where $\{\nu_j\}_{j=1}^{n_{\mathrm{m}}}$ are quadrature weights, $\boldsymbol{f}$ and $\boldsymbol{g}$ are vectors (in $\mathbb{R}^{n_{\mathrm{m}}}$) of function values at the time-steps, and $\mathbf{M}_1 = \mathrm{diag}(\nu_1, \nu_2, \ldots, \nu_{n_{\mathrm{m}}})$. In the present work, we use the composite trapezoid rule to discretize the $L^2(\mathcal{T})$ inner product.

The uncertain initial state b is discretized using finite element Lagrange nodal basis functions, $\varphi_1(\boldsymbol{x}), \ldots, \varphi_{n_{\mathrm{b}}}(\boldsymbol{x})$. This leads to the discretization $b(\boldsymbol{x}) \approx b_h(\boldsymbol{x}) = \sum_{j=1}^{n_{\mathrm{b}}} b_i \varphi_i(\boldsymbol{x})$. The discretized initial state is given by the vector $\boldsymbol{b}$ of finite-element coefficients. This finite element method is also used to discretize the PDE operator $(-\epsilon\Delta + \alpha I)$, which is the square root of the covariance operator of the distribution of $\boldsymbol{b}$. The covariance operator is thus defined as the square of the finite element operator, corresponding to a mixed discretization of the 4th-order covariance operator [18]. Also, note that the discretized $L^2(\mathcal{D})$-inner product is given by $\langle \boldsymbol{u}, \boldsymbol{v} \rangle_{\mathrm{M}_2} = \boldsymbol{u}^\mathsf{T} \mathbf{M}_2 \boldsymbol{v},$ for $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^{n_\mathrm{b}}$, where $\mathbf{M}_2$ is the finite-element mass matrix.

In the numerical experiments below, we use a discretization with $n_\mathrm{m} = 257$ time steps and $n_\mathrm{b} = 1{,}529$ spatial degrees of freedom. The "truth" primary parameter m is shown in fig. 2.1 (right), and the "truth" secondary parameter b is given by a random draw from the prior law of b, depicted in fig. 2.2 (top left). For computing solutions for the inverse problem, we synthesize data using "truth" parameters b and m, and add Gaussian noise with standard deviation $\sigma_\mathrm{noise} = 0.25$ to each data point. That is, we assume $\mathbf{\Gamma}_\mathrm{noise} = \sigma^2_\mathrm{noise}\mathbf{I}$, with $\sigma_\mathrm{noise} = 0.25$. Notice that the sensor measurements obtained from the model range approximately in the interval $[51, 54]$; see e.g., fig. 2.2 (top right). Thus, a noise standard deviation of 0.25 is significant compared to the variations of model output at the sensors.

2.5.4 Illustrating the impact of the secondary uncertainty

To depict the impact of the secondary uncertainty on the solution of the forward problem, in fig. 2.2 we show snapshots of the solution of the state equation. Here, we use two random draws from the prior distribution of b, i.e., the secondary uncertainty, as initial conditions. Recall that the initial condition used for the first row is also used as "truth" secondary parameter. For the primary uncertainty, the time evolution of the right hand side source, the "truth" parameter (see fig. 2.1 (right)) is used. Note that even at the final snapshot, around which measurements are taken for inference, distinct differences caused by the different initial conditions are visible. This indicates that the uncertainty in the initial state cannot be ignored.

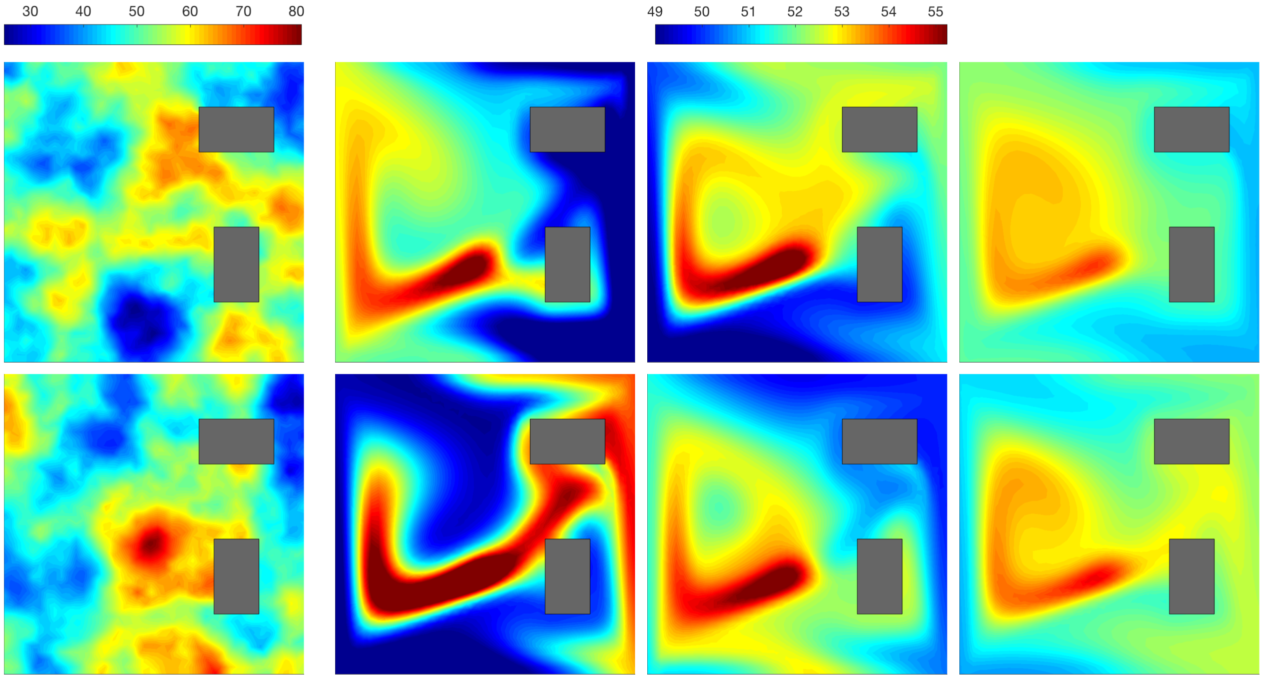

Figure 2.2: Shown in each row are snapshots of the concentration at times $t = 0, 0.4, 0.6, 1$ (from left to right). For the primary parameter m entering on the right hand side of eq. (2.24a), the "truth" parameter shown in fig. 2.1 (right) is used. For the secondary parameter b, i.e., the initial condition, two different realizations from the distribution of b are used. Note that a different colorbar is used for the initial conditions than for the other snapshots.

2.6 Computational results

In this section, we present numerical results for the model problem described in section 2.5. In section 2.6.1, we compare the performance of regularized ℓ_0-sparsification and greedy approaches for computing mOEDs. Then, in sections 2.6.2 and 2.6.3, we demonstrate the importance of taking the additional model uncertainty into account for computing sensor placements.

2.6.1 Comparison of sparsification algorithms

Here, we compare the two different approaches to obtain binary mOEDs discussed in section 2.4. As discussed in section 2.4.2, when using ℓ_0-sparsification we solve a sequence of optimization problems with non-convex penalty functions using a gradient-based optimization algorithm. Here, we use MATLAB's interior point quasi-Newton solver provided by the `fmincon` function, which we supply with routines implementing the mOED objective and its gradient. In contrast, the greedy approach only requires the mOED objective. As can be seen in fig. 2.3 (left), the greedy and the ℓ_0-sparsified designs perform similary. While in this figure the ℓ_0-sparsification finds slightly lower objective values, we have also observed tests where the objective values are identical or the greedy approach is slightly better.

It is also important to consider the computational cost of these algorithms. We do so by recording the number of mOED objective function evaluations required by the two algorithms in fig. 2.3 (right). Note that the cost of greedy sensor placement scales with the number of sensors in the optimal design, see also eq. (2.23). The cost of the ℓ_0-sparsification, in terms of function evaluations, remains nearly constant. Of course, the regularized ℓ_0-sparsification method requires gradients additionally to objective evaluations. However, as discussed in section 2.4.1, the additional cost of computing the gradient is small compared to the cost of mOED objective function evaluation. Therefore, the number of objective function evaluations is a reasonable measure to compare the cost of the two algorithms.

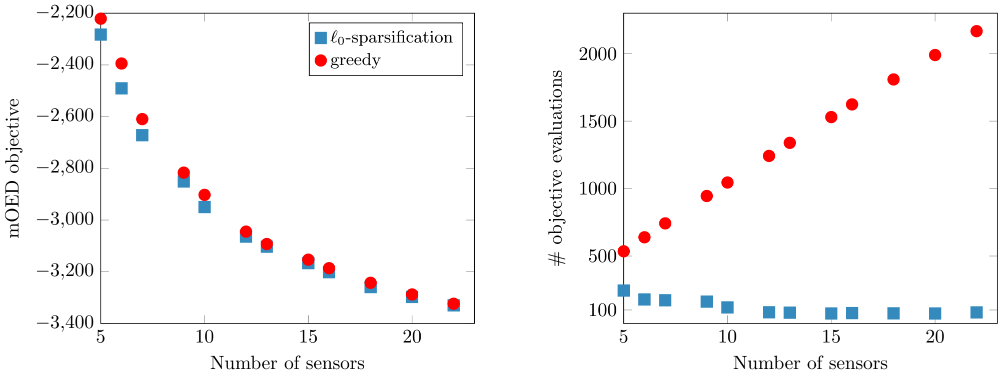

Figure 2.3: Left: mOED objective values (y-axis) plotted against number of sensors (x-axis) for the greedy (red dots) and the ℓ_0-sparsification approaches (blue dots). Right: Number of mOED objective evaluations required to converge for computing greedy (red) and ℓ_0-sparsified (blue) designs.

In the remainder of this section, where we compare the performance of designs obtained with and without marginalization, we use the greedy approach to find optimal designs. This is motivated by the fact that the greedy approach facilitates computing (near) optimal designs with a desired number of sensors, while the ℓ_0-sparsification approach only provides indirect control on the number of sensors by changing the penalty parameter γ.

2.6.2 Studying the posterior uncertainty

Next, we compare the performance of designs obtained by performing mOED against those using OED with no marginalization in terms of the resulting marginal posterior uncertainty. Note that designs obtained without marginalization, which we simply refer to as OED, minimize the classical A-optimality criterion ψ in eq. (2.11) whereas designs with marginalization minimize the mOED criterion in eq. (2.17).

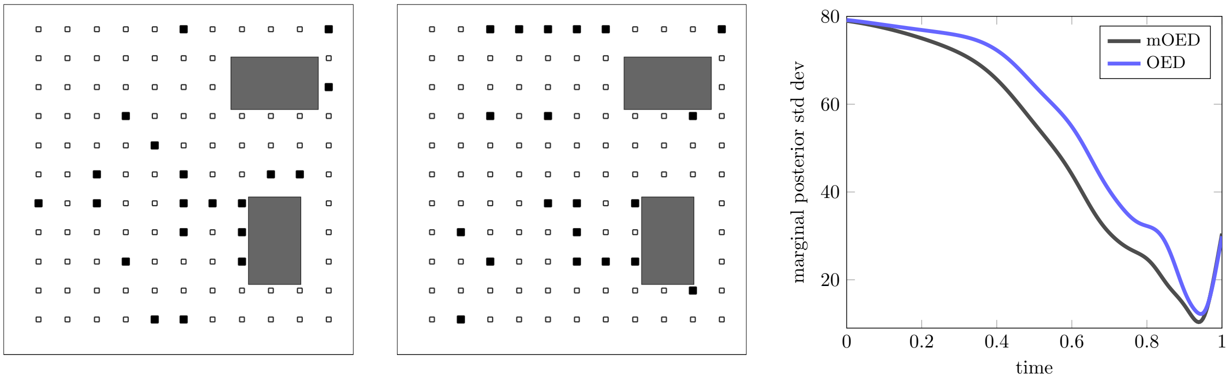

Figure 2.4: Shown are A-optimal designs with 20 sensors (filled squares) using mOED (left) and OED without marginalization (center), i.e., the design obtained with OED neglecting secondary uncertainties. Inactive sensors are shown as empty squares. On the right, the marginal posterior standard deviation field (i.e., square root of the diagonal of $\mathbf{\Gamma}_{\text{post,m}}(\boldsymbol{w})$ in eq. (2.14)) is shown for the two designs.

Figure 2.4 shows two designs with 20 sensors, one taking into account the secondary uncertainty through marginalization, and one assuming that there is no secondary uncertainty. On the right panel of fig. 2.4, the pointwise standard deviation of the marginalized posterior distribution are shown for the two sensor placements. The following conclusions can be drawn. First, note that mOED is superior, with respect to the marginalized posterior variance, to the design computed without taking the secondary uncertainty into account. Of course, this is by construction of mOEDs. However, the difference is significant and exists for all times $t \in \mathcal{T}$. Second, since measurements are taken around the final time, the uncertainty is more reduced for later times. However, close to the final time T, the uncertainty increases again as there is not enough time for the concentration field to propagate to and be picked up by sensors.

2.6.3 Study of MAP points

Next, we compare MAP points computed with the mOED and OED designs shown in fig. 2.4. Note that the MAP point for mOED does not depend on a realization of the secondary parameter (see eq. (2.14)), while it does for OED without marginalization (see eq. (2.10)). In fig. 2.5, we show the MAP point for the mOED, which recovers features from the "truth" parameter but resorts to the prior mean when little information can be gathered from observations.

As mentioned above, we need a realization of the secondary parameter b when computing the MAP point using the classical OED. If we knew the "truth" b, the additional uncertainty would vanish and the problem reduces to an inverse (and OED) problem with fully specified model as, e.g., in [5]. The corresponding MAP point, shown in blue in fig. 2.5, slightly improved compared to the MAP point from the mOED formulation. However, in general the "truth" secondary parameter is unknown, and we only know its distribution. If random draws from the secondary parameter distribution are used in the MAP computation, the model error is underestimated and the corresponding MAP points may be poor. This can be seen in fig. 2.5, where MAP points obtained with random draws from the distribution of b are shown in red.

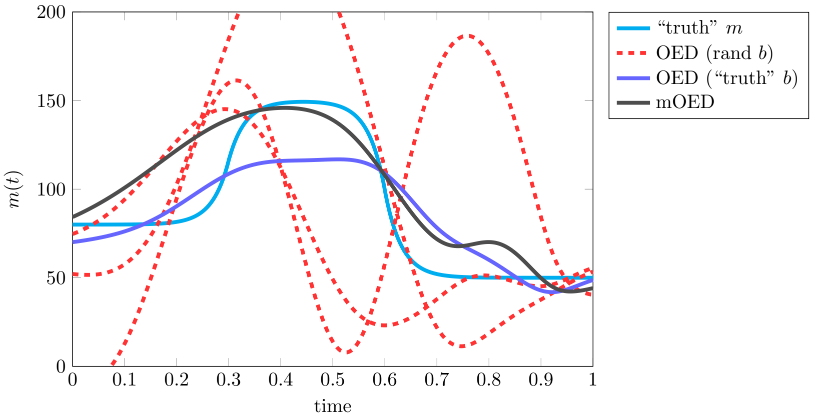

Figure 2.5: Comparison of MAP estimates computed with mOED and OED without marginalization. Shown are the MAP estimates computed using sensor placements obtained via mOED (black solid line), OED with the secondary parameter b set to the "truth" (blue solid line), and OED with b taken as realizations from corresponding prior distribution (red dotted lines).

The above discussed difference between mOED and OED without marginalization is summarized in fig. 2.6. On the left, we plot the relative $L^2(\mathcal{T})$-error between the MAP point and the "truth" primary parameter versus the mOED objective. Using OED with random draws for b result in MAP points that tend to be further from the "truth" parameter than the mOED MAP point. If the "truth" secondary parameter is used in the computation of the MAP point using OED, the reconstruction is slightly better than the result of mOED. It can also be seen that the mOED objective is independent from draws of the secondary parameter, as also discussed above. The results in fig. 2.6 (left) depend on the noise realizations in the synthetic data. In fig. 2.6 (right), we show the probability density function of the error between the MAP point and the "truth" primary parameter for random observation noise. As can be seen, it is slightly more likely to obtain a better MAP point when using OED with the "truth" parameter than with mOED. However, it can clearly be seen that mOED MAP points significantly outperform OED MAP points with random realizations from the prior distribution of b.

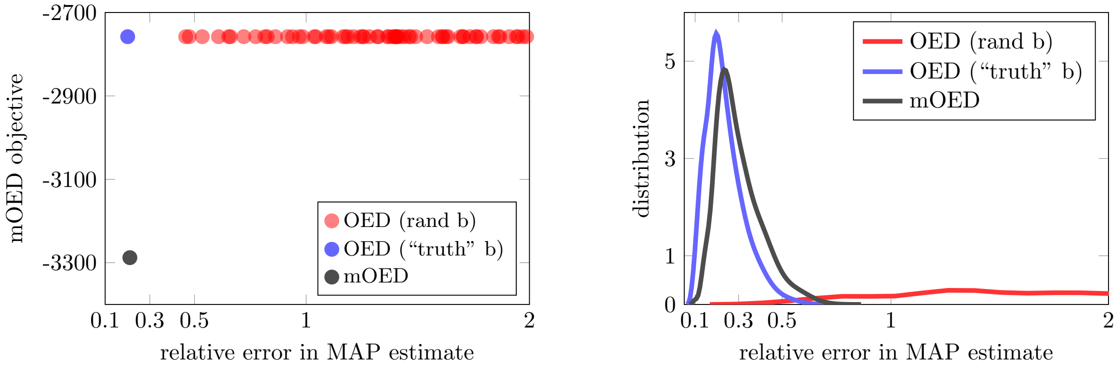

Figure 2.6: *Left:* Relative error in the MAP estimate (x-axis) and reduction in the objective (y-axis) for mOED (black dot), OED with the secondary parameter b set to the "truth" (blue dot), and OED with b taken as different realizations of b (red dots). *Right:* The distribution of the errors with various realizations of the noise in the data. Note that the x-axis is cut at 2 due to the long tail of the error distribution corresponding to OED with b taken as different realizations of b. In this study, we used 200 samples of the secondary parameter, and 500 samples of measurement noise.

2.7 Conclusion

In this chapter, we have considered linear inverse problems with reducible model uncertainty and presented a mathematical and computational framework for computing marginalized A-optimal sensors placements. Our results show that it is important to take into account additional sources of model uncertainty for the optimal design and the inverse problem in general. The designs computed by minimizing the marginalized A-optimality criterion are superior compared to classical A-optimal designs, in terms of the quality of the estimated primary parameters: the marginalized optimal designs result in optimal uncertainty reduction as well as more accurate MAP estimates. The overall conclusions support the claim made in this chapter's title, namely that in the context of design of inverse problems, it is good to know what you don't know. This information should be used when computing optimal designs.

An important direction for future work is design of nonlinear inverse problems under model uncertainty. A related direction is a sensitivity analysis framework for detecting sources of model uncertainty that are most important to the solution of the inverse problem. This would enable incorporating only the most important sources of model uncertainty in the OED problem, hence reducing the computational complexity of the problem. For deterministic inverse problems, first steps in this direction are presented in chapter 3.

CHAPTER

3

HYPER-DIFFERENTIAL SENSITIVITY ANALYSIS FOR DETERMINISTIC INVERSE PROBLEMS

The work presented in this chapter is based upon a collaborative publication [80]. The author of this dissertation is listed as the first author of that paper and was responsible for the majority of the writing and results therein.

3.1 Introduction

Rapid advances in numerical algorithms and computing infrastructure have made it feasible to simulate complex multiphysics systems governed by systems of partial differential equations (PDEs) on high resolution computational grids. Inverse problems arise when some model parameters cannot be determined directly, but rather are estimated using measurements of the model state variables. The states may correspond to different physical quantities with varying data volumes and measurement fidelities, and measurements are typically sparse and noisy due to budget and hardware limitations.

Inverse problems governed by systems with complex physics involve various sources of uncertainty. This includes the uncertainty in the parameters being estimated, uncertainty in measurement data, and uncertainty in parameters in governing PDEs that are not the focus of the parameter estimation, but are needed for a full model specification. For clarity, we refer to the model parameters being estimated as *inversion parameters* and to the other model parameters besides the inversion parameter as *auxiliary parameters*. Additionally, we refer to the parameters specifying the experimental conditions, such as types of measurements or measurement noise levels, as *experimental parameters*. The auxiliary parameters and the experimental parameters are needed for the formulation of the inverse problem. We call the combination of auxiliary and experimental parameters the *complementary parameters*. This chapter is about understanding and quantifying the impact and relative importance of the perturbations in complementary parameters to the solution of an inverse problem.

For illustration, let us consider a subsurface flow problem, in which we seek to invert for the log-permeability field using a tracer test. The forward model we consider is given by the mass conservation, constrained by Darcy's law, resulting in a linear elliptic PDE governing the pressure, and a time-dependent PDE governing diffusion and transport of the tracer. The inversion parameter here is the log-permeability field. The auxiliary parameters include the source terms (e.g., the tracer injection), boundary conditions, and coefficients (e.g., the diffusion coefficient) in the governing PDE system. The measurements correspond to the two states: pressure and concentration. The experimental parameters correspond to noise in these measurements.

We propose a general framework to assess the relative importance of complementary parameters in determining the solution of the inverse problem. To do so, we build upon previous work in [43], and a series of related articles [16, 19, 29–34], that introduced hyper-differential sensitivity analysis (HDSA) for PDE-constrained optimization. HDSA computes the Fréchet derivative of the solution of the inverse problem with respect to complementary parameters. We use this derivative to define hyper-differential sensitivities of the inverse problem solution with respect to the complementary parameters. These sensitivities describe the change in the solution of the inverse problem with respect to perturbations of a given parameter. We also define generalized sensitivity indices that determine maximum (worst case) changes in the inverse problem solution with respect to perturbations in a set of complementary parameters.

By providing sensitivity information on the experimental parameters, our framework provides vital information for effective data collection. For instance, by discovering the sensor measurements the inverse problem solution is most sensitive to, we can identify the sensors where higher fidelity measurements are desired. This can be achieved by designing sensors with improved error tolerances, or in problems where this is possible, repeating these sensor measurements to reduce the associated measurement noise. Furthermore, we demonstrate that HDSA can be used to compare the relative importance of different types of sensors. As such, this process complements experimental design which is used to design optimal sensor placement for data measurements. By calibrating the measurement fidelities in a given experimental design (e.g., a sensor network), one can make the most out of the measurements for effective parameter estimation. Therefore, the proposed framework can be combined with an optimal experimental design (OED) problem [11,67,84] to: (i) identify an optimal set of experiments; and (ii) calibrate the fidelities of the experiments or further prune the specified set of experiments, based on the sensitivity analysis results. HDSA provides a systematic framework to distinguish between measurements obtained from different sensor types and understand the relative importance of spatial and temporal sensor distributions. While HDSA is not intended to replace OED, it augments it by providing unique insights into the influence of various sensors in large-scale multiphysics applications.

Another important application of the proposed framework is guiding OED under uncertainty. In practical applications, an OED problem must be found in such a way that it is robust with respect to uncertainty in auxiliary model parameters; see e.g., [56]. Performing sensitivity analysis of the inverse problem solution with respect to auxiliary parameters informs the sources of model uncertainty one needs to focus on when solving an OED under uncertainty problem. By focusing on sources of model uncertainty the inverse problem solution is most sensitive to, our framework can significantly reduce the complexity of an OED under uncertainty problem. Furthermore, for applications in which multiple experiments may be designed to target calibration of different auxiliary parameters, the proposed sensitivities may identify where experimental efforts should be invested to calibrate the most influential parameters through a sequence of different experiments.

Additionally, in complex physics systems, typically the influence of various sources of model uncertainty on the solution behavior is not clear a priori. The proposed sensitivity analysis framework provides important insight about the governing model.

In contrast to traditional sensitivity analysis, where one quantifies the contribution of auxiliary parameters to variability in model output, our proposed framework provides a *goal oriented sensitivity analysis* approach by quantifying the impact of perturbations in auxiliary parameters on estimation of unknown model parameters. This enables determining which auxiliary parameters need to be specified more accurately. In fact, it might be that some auxiliary parameters should be estimated along with the inversion parameters, if possible.

Lastly, these sensitivities provide a computationally efficient low order approach to uncertainty quantification for large-scale systems. For instance, if the solution of the inverse problem must be determined in real time to inform critical decision making, coupling the estimated solution with a notion of uncertainty contributed by errors in the complementary parameters provides real time uncertainty estimation which is critical for making informed decisions.

This chapter considers local sensitivity analysis for deterministic variational inverse problems. We consider connections to statistical formulations and global sensitivity analysis in Section 3.7. The contributions of this chapter are as follows:

- We define HDSA with respect to experimental parameters. This provides a systematic approach to compare different sensor types and reveals information about the relative importance of distributed (spatially and temporally) sensor measurements, neither of which can be easily determined by traditional OED.
- Theoretical results are presented for linear inverse problems to provide intuition and demonstrate properties of the sensitivities with respect to experimental parameters.
- We build upon previous work [43] to develop a more comprehensive mathematical framework for HDSA of nonlinear multiphysics inverse problems. This is done, in particular, by maturing the idea of generalized sensitivities as a tool for systematically comparing the importance of auxiliary parameters (which may be of differing physical characteristics and scales) alongside the novel development of HDSA for experimental parameters.
- Comprehensive numerical results, in a large-scale subsurface flow application, demonstrate the interpretation and use of HDSA for nonlinear multiphysics inverse problems.

The remainder of the chapter is organized as follows. Section 3.2 outlines the basic principles of inverse problems and design of experiments. Section 3.3 provides the mathematical formulation of hyper-differential sensitivities and their interpretation for inverse problems constrained by multiphysics. The computational considerations, implementation, and cost analysis of HDSA is detailed in Section 3.4. Section 3.5 presents a large scale, multiphysics model problem, which is then used to construct sensitivity results that are detailed in Section 3.6. Concluding remarks and notes on potential areas of future work are highlighted in Section 3.7.

3.2 Preliminaries

In this section, we recall background material on inverse problems and design of experiments, which are augmented by the proposed sensitivity analysis in subsequent sections.

3.2.1 Inverse Problems

In this chapter, we are concerned with ill-posed inverse problems governed by PDEs. Specifically, we seek to estimate a parameter m, henceforth called the inversion parameter, using data $\boldsymbol{y}$ and a model of the form

$$F(m) + \boldsymbol{\eta} = \boldsymbol{y}.$$

Here F is the parameter-to-observable map and $\boldsymbol{\eta}$ represents measurement noise. Evaluating $F(m)$ requires solving the governing PDEs and extracting the solution at measurement points.

Due to ill-posedness and availability of only sparse noisy measurements, we are led to variational formulations with suitable regularizations. Specifically, we consider an optimization problem of the following form:

$$\begin{aligned} &\min_{u,m} J(u, m, \theta_e) \\ \text{s.t.} \quad & v(u, m, \theta_a) = 0 \\ & u \in U, m \in \mathcal{M}. \end{aligned} \tag{3.1}$$

Here, U is an infinite dimensional reflexive Banach space containing the state, $\mathcal{M}$ is a possibly infinite dimensional Hilbert space, J is a regularized data misfit cost functional (we make this precise below), and v represents the constraining PDE system. The experimental parameters, θ_e, represent uncertainty in the data, while θ_a are the auxiliary parameters contained in the system of PDEs. Generally, solving this optimization problem produces parameter estimates that are consistent with measurement data and the model. For the remainder of the chapter, we refer to (3.1) as the inverse problem.

We mention that an alternative approach to address ill-posed inverse problems is to consider a Bayesian formulation [62]. In this approach, the inversion parameter m is modeled as a random variable, and the goal is to find a distribution law for m that is consistent with measurement data, the model, and a prior distribution of m that models our prior knowledge/beliefs about m. In the current chapter, we restrict our attention to deterministic formulation of inverse problems, as described above.

We assume that the PDE represented by v is uniquely solvable for any admissible m and θ_a. This allows us to formulate (3.1) in *reduced space* [2]. Letting $\mathcal{A}(m, \theta_a)$ denote the solution operator for the PDE, i.e. $v(\mathcal{A}(m, \theta_a), m, \theta_a) = 0$ for all m and θ_a, we define the reduced objective function $\hat{J}(m, \theta_e, \theta_a) = J(\mathcal{A}(m, \theta_a), m, \theta_e)$. In this chapter we focus on objective functions defined as a linear combination of data misfit and regularization, yielding a general form for the inverse problem

$$\min_{m\in\mathcal{M}} \hat{J}(m,\theta_e,\theta_a) := \frac{1}{2}\|\mathcal{Q}\mathcal{A}(m,\theta_a)-\boldsymbol{y}(\theta_e)\|^2+\alpha\mathcal{R}(m), \tag{3.2}$$

where $\boldsymbol{y}(\theta_e)$ is a vector of measured data (with uncertainty parameterized by θ_e), $\mathcal{Q}$ an observation operator that maps the PDE solution to a set of observation locations, $\mathcal{R}$ is a regularization operator, and α is a regularization parameter. Traditional approaches to solving inverse problems fix $\theta = (\theta_e, \theta_a)$ to a best estimate and solve (3.2) by optimizing over m. Analyzing the influence of θ on the solution of (3.2) is the focus of this chapter.

Besides ill-posedness, such inverse problems are difficult to solve for a number of other reasons. These include having noisy observations, expensive forward PDE solves, tuning multiple experimental and modeling parameters, and optimization in infinite (or large finite) dimensional spaces. Common optimization methods used to tackle such problems include quasi-Newton, inexact Newton-CG, Gauss-Newton, and truncated CG trust region solvers. These optimization problems often require efficient gradient and Hessian computation through adjoint state methods, and repeated large scale linear system solves with Krylov iterative methods. We direct the interested reader to a number of classical inverse problem references [17,24,25,50,53,81,86].

3.2.2 Design of Experiments

An important aspect of solving an inverse problem is the collection of informative measurement data, which is guided by *design of experiments*. In our target inversion, this generally corresponds to specifying the locations of the sensors used to collect measurement data and is known as an optimal experimental design (OED) problem [11,67,84]. OED for inverse problems governed by differential equations has received significant attention in recent years; see e.g., [5,6,8,13,15,37,38,46,47,55,59,60]. An OED problem is typically formulated with a statistical formulation of the inverse problem in mind. An optimal design is one that optimizes the statistical quality of the estimated parameters. Examples include maximizing the expected information gain, leading to a D-optimal design problem, or minimization of average posterior variance, leading to a Bayesian A-optimal design problem.

OED is a powerful tool that is used on a wide variety of problems. It is also a very challenging problem both from mathematical and computational point of view, especially when it comes to nonlinear inverse problems governed by PDEs; see e.g., [6,46]. The developments in this chapter are closely related to OED: while we do not directly solve an OED problem, we address the following relevant questions: (i) which measurements is the solution of an inverse problem most sensitive to? And (ii) which measurement types are most influential to the solution of the inverse problem? The latter is tied to important questions typically not addressed in OED literature: how should multi-purpose sensors that can take different types of measurements be deployed, and how should different sensor types be designed and deployed in an existing experimental design?

3.3 Hyper-differential Sensitivity Analysis for Inverse Problems

This section is devoted to our proposed framework for hyper-differential sensitivity analysis (HDSA) of PDE-constrained inverse problems. In Subsection 3.3.1, we detail the mathematical formulation of the operator mapping complementary parameters to solutions of the PDE-constrained inverse problem, and its Fréchet derivative. In Subsection 3.3.2 the hyper-differential sensitivities are defined, as well as the generalized sensitivity index which is used to compare the importance of sets of complementary parameters with different physical characteristics. Subsection 3.3.3 presents an analytical result for linear inverse problems which connects the sensitivities to the trace of the covariance in the solution of the inverse problem.

3.3.1 Mathematical Formulation

HDSA differs from traditional sensitivity analysis in that it determines the sensitivity of the solution of an optimization problem rather than simply a model (which is typically a constraint in the optimization problem). We seek to perform HDSA on (3.2) to determine the sensitivity of the optimal m to uncertainty in complementary (both experimental and auxiliary) parameters θ which are fixed when solving (3.2).

HDSA uses the derivative of the solution of (3.2) with respect to θ. To formally define HDSA, we assume that $\hat{J}$ is twice continuously differentiable with respect to (m, θ) and that $m^\star$ is a local minimum of (3.2) for specified complementary parameters $\theta = \theta^\star$. Assuming that the Hessian of $\hat{J}$ with respect to m, evaluated at $(m^\star, \theta^\star)$, is positive definite [16, 43], we can apply the implicit function theorem [9, p. 38] to $\hat{J}_m$ (the Fréchet derivative of $\hat{J}$ with respect to m), to define a continuously differentiable mapping $\mathcal{F}$ from a neighborhood of $\theta^\star$ to a neighborhood of $m^\star$,

$$\mathcal{F} : \mathcal{N}(\theta^\star) \to \mathcal{N}(m^\star)$$

such that

$$\hat{J}_m(\mathcal{F}(\theta), \theta) = 0, \qquad \text{for all} \qquad \theta \in \mathcal{N}(\theta^\star),$$

i.e., $\mathcal{F}$ maps complementary parameters to stationary points of (3.2). The Fréchet derivative of $\mathcal{F}$ with respect to θ, evaluated at $\theta^\star$, is given by

$$\mathcal{D} := \mathcal{F}_\theta(\theta^\star) = -\mathcal{H}^{-1}\mathcal{B}, \tag{3.3}$$

where $\mathcal{H}$ is the Hessian of $\hat{J}$ with respect to m, evaluated at $m^\star$ and $\theta^\star$, i.e., $\mathcal{H} := \hat{J}_{m,m}(m^\star, \theta^\star)$, and $\mathcal{B}$ is the Fréchet derivative of $\hat{J}_m$ with respect to θ, evaluated at $m^\star$ and $\theta^\star$, i.e., $\mathcal{B} := \hat{J}_{m,\theta}(m^\star, \theta^\star)$.

An intuitive interpretation of (3.3) is that once (3.2) has been solved to optimality for the specified $\theta^\star$, we take a perturbation with respect to θ ($\mathcal{B}$) and apply a Newton step ($-\mathcal{H}^{-1}$) to update the solution of the inverse problem. We interpret $\mathcal{D}\overline{\theta}$ as the sensitivity of the solution of the inverse problem when the complementary parameters are perturbed in the direction $\overline{\theta}$. Note that upon discretization, applying the inverse of $\mathcal{H}$ to a vector requires a large-scale linear solve, which requires many PDE solves.

3.3.2 Sensitivity Indices

Given a sensitivity operator, we define scalar sensitivity indices that measure the magnitude of the change in the inverse problem solution with respect to perturbations 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); and the parameters defining the discretization of a function-valued auxiliary parameter may form another group.

In general, the complementary parameters $\theta = (\theta_e, \theta_a) \in \Theta$ take values in a possibly infinite dimensional space $\Theta = \Theta_1 \times \Theta_2 \times ... \times \Theta_K$, which is a product of K Hilbert spaces. The first ℓ parameter spaces contain the experimental parameters $\theta_e \in \Theta_1 \times ... \times \Theta_\ell$, while the remainder contain the auxiliary parameters $\theta_a \in \Theta_{\ell+1} \times ... \times \Theta_K$. The product space Θ is equipped with the inner product

$$\langle \theta, \phi \rangle_\Theta = \langle \theta_1, \phi_1 \rangle_{\Theta_1} + \cdots + \langle \theta_K, \phi_K \rangle_{\Theta_K}, \qquad \text{for} \qquad \theta, \phi \in \Theta.$$

We are particularly interested in cases where the uncertain parameters have varying physical characteristics, such as different units or differences in their spatial or temporal dependence.

To better understand spatial and temporal patterns of importance for a particular parameter or data source, we define pointwise sensitivity indices in space and time, and later define generalized sensitivities which remove these units. From here on, we use $\boldsymbol{\Theta}$ and $\boldsymbol{\Theta}_k$ to denote the discretizations of the possibly infinite dimensional spaces Θ and Θ_k. To respect spatiotemporal structure in discrete data, we use weighted norms corresponding to spatial and/or temporal discretizations. For instance, if $\boldsymbol{\theta}_k$ models perturbations of spatiotemporal data measurements then

$$\|\boldsymbol{\theta}_k\|_{\boldsymbol{\Theta}_k} = \sqrt{\frac{1}{n_t n_s}\sum_{i=1}^{n_t n_s}(\theta_k^i)^2}$$

where θ_k^i denotes the ith component of the vector $\boldsymbol{\theta}_k \in \mathbb{R}^{n_t n_s}$, n_s and n_t denote the number of spatial and temporal points, respectively.

Upon discretization of (3.1), we let $\{\boldsymbol{b}_k^1, \boldsymbol{b}_k^2, \ldots, \boldsymbol{b}_k^{n_k}\}$ denote a basis for each parameter space $\boldsymbol{\Theta}_k$, where n_k is the dimension of $\boldsymbol{\Theta}_k$. We define a basis for $\boldsymbol{\Theta}$ as $\{\boldsymbol{e}_k^i\}$ for $k = 1, \ldots, K$ and $i = 1, \ldots, n_k$ where

$$\boldsymbol{e}_k^i = \begin{pmatrix}\mathbf{0}_1 & \ldots & \mathbf{0}_{k-1} & \boldsymbol{b}_k^i & \mathbf{0}_{k+1} & \ldots & \mathbf{0}_K\end{pmatrix}^\top.$$

We define the pointwise sensitivity indices for $k = 1, \ldots, K$ and $i = 1, \ldots, n_k$ as,

$$S_k^i = \frac{\|\mathbf{D}\boldsymbol{e}_k^i\|_{\mathbf{M}}}{\|\boldsymbol{e}_k^i\|_{\boldsymbol{\Theta}}}, \tag{3.4}$$

where $\mathbf{D}$ is the discretized sensitivity operator (3.3), and $\|\cdot\|_{\mathbf{M}}$ is the norm discretized consistently with respect to the norm in $\mathcal{M}$. The pointwise sensitivities measure the change in the solution of the inverse problem with respect to a perturbation of the kth parameter in direction $\boldsymbol{b}_k^i$. Thus, high sensitivity indicates that errors in the parameter will cause a significant change in the reconstructed solution. This leads to an interpretation of the sensitivities as quantifying the importance of accurately measuring or modeling the parameter.

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 which provide a single measure of sensitivity for each parameter subgroup. Let $\mathbf{T}_k : \Theta \to \Theta$ be a selection operator that zeros out components of $\boldsymbol{\theta}$ not in Θ_k. We define the generalized sensitivity of the kth subgroup of complementary parameters as

$$S_k = \max_{\boldsymbol{\theta}\in\boldsymbol{\Theta}} \frac{\|\mathbf{D}\mathbf{T}_k\boldsymbol{\theta}\|_{\mathbf{M}}}{\|\boldsymbol{\theta}\|_{\boldsymbol{\Theta}}}. \tag{3.5}$$

The generalized sensitivities measure the maximum change that can be observed in the solution to a norm-1 perturbation of the kth 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 which 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 3.5 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. This is because (3.4) and (3.5) are both equivalent to the largest singular value of $\mathbf{M}^{1/2}\mathbf{D}\mathbf{T}_k$ when $n_k = 1$.

3.3.3 Interpretation of Experimental Parameter HDSA for Linear Inverse Problems

Sensitivity of auxiliary parameters, θ_a in the notation of this chapter, is a natural concept with a clear physical interpretation. The sensitivity of the solution of the inverse problem to experimental parameters is less intuitive, so we present an analytic result in Proposition 3.3.1 to provide intuition. For conciseness and clarity in this subsection, we consider only uncertainty in the experimental parameters.

Assume that $\tilde{\boldsymbol{y}} = (\tilde{y}_1, \tilde{y}_2, \ldots, \tilde{y}_n)^\top$ is a vector of noisy data which may be modeled by a linear parameter-to-observable map $\mathcal{QA}$ acting on an unknown parameter m. Estimating this unknown m gives rise to a linear inverse problem. To apply HDSA with respect to experimental parameters, we model the data as

$$y_i = \tilde{y}_i(1 + \theta_e^i) \qquad i = 1, 2, \ldots, n,$$

where $(\theta_e^1, \theta_e^2, \ldots, \theta_e^n)^\top = \boldsymbol{\theta_e} \sim N(0, \boldsymbol{\Sigma})$ is a perturbation of the nominal value $\tilde{y}_i$, i.e. a noise model. Assuming uncorrelated observations, the noise covariance matrix is $\boldsymbol{\Sigma} = \operatorname{diag}(\sigma_1^2, \ldots, \sigma_n^2)$. An estimate of the inversion parameter can be obtained by solving

$$\min_m \hat{J}(m) := \frac{1}{2}\|\mathcal{QA}m - \boldsymbol{y}\|^2_{\boldsymbol{\Sigma}^{-1}} + \frac{\alpha}{2}\|m\|^2_R \tag{3.6}$$

with $\boldsymbol{\theta_e} = 0$, i.e. solving a linear least squares problem with the data $\tilde{\boldsymbol{y}}$. The norm in the regularization term is weighted by regularization operator $R : \mathcal{M} \to \mathcal{M}$ which we assume is a self-adjoint, strictly positive linear operator on $\mathcal{M}$ [86].

The estimator $m^\star$ obtained from solving (3.6) is a random variable, due to the random noise in the data. HDSA provides the sensitivity of its solution with respect to $\boldsymbol{\theta_e} = (\theta_e^1, \theta_e^2, \ldots, \theta_e^n)$. This may be interpreted as the sensitivity of the least squares estimate with respect to the data, a metric to assess the relative importance of the observations. Such sensitivity information can be used to inform sensor designs and measurement tolerances. Proposition 3.3.1 relates the variance of $m^\star$ (with respect to randomness in $\boldsymbol{\theta_e}$) to the pointwise sensitivities with respect to data measurements S^i as defined in equation (3.4). Here, we omit the subscript ($K = 1$) on the pointwise sensitivities, as we only consider one type of parameter in this section.

Proposition 3.3.1. $\mathsf{tr}(\mathrm{Cov}(m^\star)) = \sum_{i=1}^n (S^i)^2$, *where* tr *denotes the trace of a linear operator and* S^i *is defined as in* (3.4) *when* $\Theta = \mathbb{R}^n$ *is equipped with the* $\mathbf{\Sigma}^{-1}$ *weighted norm.*

Proof. Computing the Fréchet derivative of the objective $\hat{J}$ in (3.6), setting it equal to zero, and solving for m yields the solution of the inverse problem,

$$m^\star = (\mathcal{A}^*\mathbf{W}\mathcal{A} + \alpha R)^{-1}\mathcal{A}^*\mathbf{W}\boldsymbol{y} = (\mathcal{A}^*\mathbf{W}\mathcal{A} + \alpha R)^{-1}\mathcal{A}^*\mathbf{W}(\tilde{\boldsymbol{y}} + \tilde{\mathbf{Y}}\boldsymbol{\theta_e}),$$

where $\mathcal{A}^*$ denotes the adjoint of the linear operator $\mathcal{A}$ and $\tilde{\mathbf{Y}} = \mathrm{diag}\Big(\tilde{y_1}, \tilde{y_2}, \ldots, \tilde{y_n}\Big)$. Computing the Fréchet derivative of $m^\star$ with respect to $\boldsymbol{\theta_e}$ (which coincides with (3.3)) yields

$$\mathcal{D} = (\mathcal{A}^*\mathbf{W}\mathcal{A} + \alpha R)^{-1}\mathcal{A}^*\mathbf{W}\tilde{\mathbf{Y}}.$$

The covariance of the estimator $m^\star$ (with respect to randomness in $\boldsymbol{\theta_e}$) is

$$\mathrm{Cov}(m^\star) = \mathrm{Cov}(\mathcal{D}\boldsymbol{\theta_e}) = \mathcal{D}\mathbf{\Sigma}\mathcal{D}^*.$$

Therefore,

$$\begin{aligned} \mathsf{tr}(\mathrm{Cov}(m^\star)) = \mathsf{tr}(\mathcal{D}\mathbf{\Sigma}\mathcal{D}^*) = \mathsf{tr}(\mathcal{D}^*\mathcal{D}\mathbf{\Sigma}) &= \sum_{i=1}^n \langle \boldsymbol{e}_i, \mathcal{D}^*\mathcal{D}\mathbf{\Sigma}\boldsymbol{e}_i\rangle \\ &= \sum_{i=1}^n (\sigma_i \|\mathcal{D}\boldsymbol{e}_i\|_{\mathcal{M}})^2 = \sum_{i=1}^n (\sigma_i \|\boldsymbol{e}_i\|_{\mathbf{\Sigma}^{-1}} S^i)^2 = \sum_{i=1}^n (S^i)^2, \end{aligned}$$

where $\boldsymbol{e}_i$ is the ith canonical unit vector in $\mathbb{R}^n$. □

This result provides useful intuition into the sensitivity indices and indicates that the variance of the inverse problem solution is scaled by the magnitude of the sensitivities with respect to data.

Remark 3.3.1. *Notice that Proposition 3.3.1 generalizes naturally to Bayesian linear inverse problems. With a Gaussian prior and likelihood, the solution of the Bayesian linear inverse problem is a Gaussian posterior. With the assumption* $\boldsymbol{\theta}_e \sim N(0, \mathbf{\Sigma})$ *on the measurement noise, and an appropriately chosen prior covariance, the maximum a posteriori (MAP) point is equivalent to the solution* $m^\star$ *of the deterministic linear inverse problem. Taking the trace of the MAP point's covariance we again conclude Proposition 3.3.1. Note that in the*

Bayesian setting, we consider the average variance of the MAP estimator as a measure of robustness of this point estimator for the inversion parameter. This is different than the average posterior uncertainty in the parameter given by the trace of the posterior covariance operator.

3.4 Computational Considerations

We begin our discussion with a simple illustrative example. To highlight the dimensions of the discretized operators and gain insight into the computational complexity of HDSA, we consider a discretized inverse problem with only auxiliary parameters:

$$\begin{aligned} &\min_{\boldsymbol{m}} \hat{J}(\boldsymbol{m}) = \frac{1}{2}\|\mathbf{Q}\boldsymbol{u} - \boldsymbol{y}\|_{\mathbf{W}}^2 + \frac{\alpha}{2}\|\boldsymbol{m}\|_{\mathbf{R}}^2 \\ &\text{where } \ \mathbf{L}(\boldsymbol{m})\boldsymbol{u} = \mathbf{V}\boldsymbol{\theta}. \end{aligned} \tag{3.7}$$

where $\mathbf{Q} \in \mathcal{R}^{d\times n}$ is an observation operator, $\boldsymbol{u} \in \mathcal{R}^n$ the state vector, $\boldsymbol{y} \in \mathcal{R}^d$ is the vector of experimental observations, $\mathbf{W} \in \mathcal{R}^{d\times d}$ a symmetric weight matrix, $\alpha > 0$ a regularization coefficient, $\boldsymbol{m} \in \mathcal{R}^p$ the discretized inversion parameter, $\mathbf{R} \in \mathcal{R}^{p\times p}$ a symmetric positive definite regularization operator, $\mathbf{L}(\boldsymbol{m}) \in \mathcal{R}^{n\times n}$ a discretized differential operator, $\boldsymbol{\theta} \in \mathcal{R}^k$ the vector of auxiliary parameters, and $\mathbf{V} \in \mathcal{R}^{n\times k}$. Note that the discretized state dimension n corresponds to the number of degrees of freedom in the mesh (typically large), and for problems with distributed parameters, p will also have a comparable dimension to n (frequently equal). The dimension of the auxiliary parameters, k, can also be large, potentially larger than n when there are multiple distributed auxiliary parameters.

In practice, we compute the action of the gradient and Hessian of $\hat{J}$ using a formal Lagrangian approach where each application of the Hessian requires two linear PDE solves (inverting the matrix $\mathbf{L}(m)$). To compute the action of $\mathbf{D}$ (the discretization of (3.3)) on a vector, we require 2 linear solves to apply the matrix $\mathbf{B}$ to a vector, and then $2I$ linear solves to apply $\mathbf{H}^{-1}$ to the resulting vector, where I is the number of iterations needed by an iterative linear solver. We direct the reader to 3.8, where we demonstrate the adjoint method used to compute the gradient and Hessian of the reduced objective function $\hat{J}$, as well as the operator $\mathbf{B}$.

In general, the sensitivity indices (3.4) and generalized sensitivity indices (3.5) may be computed in a variety of ways. The efficiency of different approaches depends upon (i) the dimension of the parameter space, (ii) the computational cost of the PDE solves, (iii) the structure of the Fréchet derivative $\mathcal{D}$, and (iv) the available computational resources.

The computational bottleneck when computing (3.4) and (3.5) is repeatedly inverting $\mathcal{H}$

(a large linear system solve). In general, we are interested in systems of nonlinear PDEs. For such systems, each application of $\mathcal{H}^{-1}$ requires $2I$ linearized PDE solves, where I is the number of iterations required by the linear solver (such as conjugate gradient). Because HDSA is post-optimality analysis, we do not require solving nonlinear systems repeatedly as in the inverse problem, but rather solving PDEs which are linearized about the solution of the inverse problem.

As introduced in [43], a randomized generalized eigenvalue problem may be formulated to estimate the truncated generalized singular value decomposition (GSVD) of $\mathcal{D}$. When the parameter dimension is large and Fréchet derivative $\mathcal{D}$ is low rank, both (3.4) and (3.5) may be efficiently estimated by using the truncated GSVD and leveraging the parallelism of randomized methods. We refer the reader to [43] for additional details.

If $\mathcal{D}$ is not low rank but the parameter dimension and cost per PDE solve is mild, we may compute (3.4) and (3.5) directly by applying $\mathcal{D}$ to each basis function in Θ. This does not exploit structure as in the GSVD approach, but it is embarrassingly parallel making it feasible for moderate parameter dimensions.

If $\mathcal{D}$ is not low rank and the parameter dimension or cost per PDE solver prohibits computing (3.4) and (3.5) directly, we may still compute (3.5) using a GSVD. Each generalized sensitivity index corresponds to the leading singular value of $\mathcal{D}$ acting on a projection operator. Since the number of generalized sensitivities are typically small, they may be estimated by using randomized solvers to compute the leading singular value. By exploiting parallelism, this may be done with a modest number of linearized PDE solvers regardless of the spectral decay in $\mathcal{D}$.

3.5 Model Problem

In this section, we present a multiphysics model problem which is then used in Section 3.6 to compute hyper-differential sensitivities and demonstrate the usefulness and flexibility of HDSA. As a motivating example, we consider the problem of identifying the permeability field of a porous subsurface region with a tracer substance flowing through the domain. We consider a unit square domain Ω with boundary $\Gamma = \cup_{i=0}^{3}\Gamma_i$, where Γ_0, Γ_1, Γ_2, and Γ_3 denote the bottom, right, top, and left edges of Ω, respectively.

We model subsurface flow of a fluid through a porous medium with Darcy's Law and consider transport of the tracer through the medium governed by the advection diffusion equation:

$$
\begin{aligned}
-\nabla \cdot (e^m \nabla p) &= 0 \quad \text{in } \Omega && \text{(3.8a)}\\
c_t - \nabla \cdot \Big(\epsilon \nabla c\Big) + \nabla \cdot \Big(\boldsymbol{v} c\Big) &= g \quad \text{in } [0,T] \times \Omega && \text{(3.8b)}\\
p &= p_1 \quad \text{on } \Gamma_1 && \text{(3.8c)}\\
p &= p_2 \quad \text{on } \Gamma_3 && \text{(3.8d)}\\
\nabla p \cdot n &= 0 \quad \text{on } \Gamma_0 \cup \Gamma_2 && \text{(3.8e)}\\
\nabla c \cdot n &= 0 \quad \text{on } [0,T] \times \{\Gamma_0 \cup \Gamma_1 \cup \Gamma_2 \cup \Gamma_3\} && \text{(3.8f)}\\
c(0,\cdot) &= 0 \quad \text{in } \Omega && \text{(3.8g)}
\end{aligned}
$$

Here p denotes the pressure field, m the log-permeability field of the medium, $\boldsymbol{v} = -e^m \nabla p$ the Darcy velocity, $c(t,x)$ the tracer concentration, ϵ the diffusivity constant, and g the source term of the injected tracer. In the present example, we used $\epsilon = 0.025$. For simplicity of notation, the constant fluid viscosity and constant porosity of the medium have been removed from the model. The Dirichlet pressure boundary conditions (on left and right boundaries) are described by non-zero fuctions p_1 and p_2. We let p_1 be greater in magnitude than p_2, as this pressure difference will drive fluid flow from right to left through the domain,

$$
\begin{aligned}
p_1(y) &= 15 + \cos(2\pi y) + \frac{1}{2}\cos(4\pi y),\\
p_2(y) &= 10 + 2\cos(2\pi y).
\end{aligned}
$$

The tracer source is described by

$$
g(x,y) = \sum_{k=1}^{16} 10 e^{-100((x-v_k)^2 + (y-w_k)^2)}
$$

where the source injection locations (v_k, w_k) are arranged in a 4×4 grid as depicted by the diamonds in Figure 3.1.

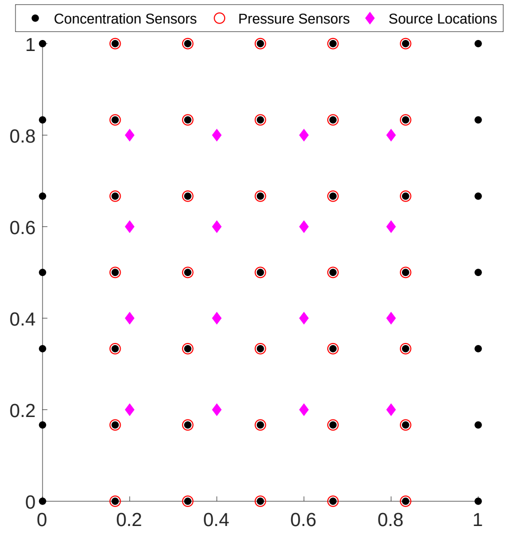

Figure 3.1: Concentration sensor, pressure sensor, and source locations.

We seek to solve an inverse problem to reconstruct the log-permeability field m, using pressure and concentration measurements. Let $\mathcal{Q}$ denote the observation operator and $\boldsymbol{y} \in \mathbb{R}^n$ be a vector of n_p pressure measurements and n_c concentration measurements at n_t measurement times, giving a total of $n = n_p + n_c n_t$ data points,

$$\boldsymbol{y} = \begin{bmatrix} p_1 & p_2 & \dots & p_{n_p} & c_1 & c_2 & \dots & c_{n_c n_t} \end{bmatrix}.$$

The observation (sensor) locations are depicted in Figure 3.1.

We consider the inverse problem

$$\min_m \hat{J}(m) := \frac{1}{2}\|\mathcal{Q}\mathcal{A}(m) - \boldsymbol{y}\|_{\mathbf{W}}^2 + \frac{\alpha}{2}\int_\Omega \|\nabla m\|_2^2 \, dx$$

where $\mathcal{A}$ is the solution operator for (3.8),

$$\mathbf{W} = \begin{pmatrix} \frac{1}{\overline{p}^2\sigma^2}\mathbf{I}_{n_p} & 0 \\ 0 & \frac{1}{\overline{c}^2\sigma^2}\mathbf{I}_{n_c n_t} \end{pmatrix}$$

is a data misfit weight matrix and α is the regularization coefficient. We used $\alpha = 3 \times 10^{-2}$ in our numerical experiments; this was chosen based on numerical experimentations seeking a regularization coefficient that is large enough to mitigate ill-posedness and at the same time produces a reasonable parameter reconstruction. The weight matrix divides each measurement by the measurement noise σ and the average of its data type ($\overline{p}$ and $\overline{c}$

respectively) to ensure the two data types, which are on different scales, have equivalent importance in the data misfit term.

We synthesize data for this problem with additive Gaussian noise that perturbs the data with a standard deviation of 3% of the true value, i.e. $\theta_e^i \sim \mathcal{N}(0, \sigma)$, with $\sigma = 0.03$. Note that we assume pressure and concentration sensors have the same proportional measurement error, $\sigma_p = \sigma_c = 0.03$.

The inverse problem is solved on a 55×55 finite element spacial discretization with 48 time steps, while the data is generated from a forward PDE solve with a 109×109 finite element spacial discretization and 98 time steps. We use a completely uninformed, constant 0 initial guess, with 49 concentration sensors and 35 pressure sensors arrayed throughout the domain as depicted in Figure 3.1 to solve the inverse problem.

3.6 Computational Results

Using the model presented in Section 3.5, we solve the inverse problem and compute sensitivity indices to determine which parameters and data sources are most important. Figure 3.2 depicts the true log permeability field we seek to reconstruct through the inverse problem, and the reconstructed solution found by solving the inverse problem with a truncated CG trust region solver. Note that since the optimization problem is non-convex we are only guaranteed to find a local minimizer.

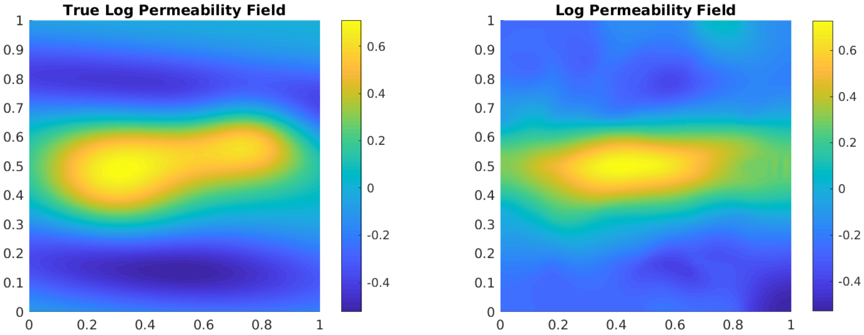
True Log Permeability Field

Log Permeability Field

Figure 3.2: Left: True Permeability Right: Reconstructed Solution.

In Subsection 3.6.1 we display and interpret generalized sensitivity indices for the pressure data, concentration data, tracer source term, diffusion coefficient, and the left and right pressure Dirichlet boundary conditions. In Subsection 3.6.2 we analyze the pointwise sensitivities with respect to the experimental parameters (pressure and concentration data), and in Subsection 3.6.3 we analyze the pointwise sensitivities with respect to auxiliary parameters (source term, pressure Dirichlet boundary conditions, and diffusion coefficient). General discussion of the importance and interpretation of the sensitivities is presented in Subsection 3.6.4.

We model uncertain parameters as a nominal value times a parameterized perturbation. When our parameters of interest are constants, such as data measurements or modeling coefficients, we can model uncertain parameters as

$$d = \tilde{d}(1 + a\theta),$$

where d is the parameter of interest, $\tilde{d}$ is the nominal value, a is a scaling coefficient, and $\theta \in [-1, 1]$ the parameterization of the perturbation. In practice, we compute the sensitivity with $\theta = 0$, which corresponds to computing the sensitivity at the nominal parameter value $\tilde{d}$. Extensions to global sensitivity analysis may consider sampling θ in $[-1, 1]$. The scaling coefficient a is problem dependent, and should be set based on prior knowledge of the level of uncertainty in the parameter of interest. When the parameter is a spatially and/or temporally distributed, we model uncertainty in the function using a linear combination of basis functions

$$f(x) = \tilde{f}(x)\Big(1 + a\sum_{i=1}^{L}\theta_i\phi_i(x)\Big)$$

where $\tilde{f}$ its nominal estimate, L the dimension of the discretized basis, and $\{\phi_i\}$ are basis functions. The results in this chapter take ϕ_i as linear finite element basis functions defined on a coarser mesh than the PDE is solved on (to enforce smoothness in perturbations). For this model problem we let $a = 0.05$ for the experimental parameters and $a = 0.2$ for the auxiliary parameters which indicates 5% uncertainty in the data (experimental parameters) and 20% uncertainty in the auxiliary parameters.

3.6.1 Generalized Sensitivity Results

We calculate the generalized sensitivities for each parameter type, which allows for comparison between their relative importance. These are plotted in Figure 3.3.

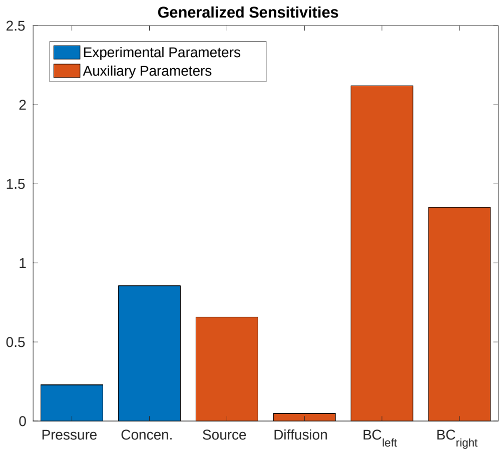

Figure 3.3: Bar graph of generalized parameter sensitivities.

From Figure 3.3, we can see that for this specific model problem it is most important to accurately measure the left and right boundary conditions. This makes sense intuitively, as the boundary conditions drive the fluid flow and the problem is advection dominated. We can also tell that in terms of data collection, it is more important to accurately measure concentration than pressure and that measuring the diffusion coefficient with a high degree of accuracy is relatively unimportant.

We consider the left boundary condition as an example and illustrate the interpretation of its generalized sensitivity index. Such principles of interpretation may be extended to any other generalized sensitivity indices but are omitted for conciseness. Figure 3.4 displays the true left boundary condition and the perturbation of the left boundary condition corresponding to the generalized sensitivity index (the argument of the maximization in (3.5). The perturbation plotted in Figure 3.4 is the unit norm perturbation that results in the maximum change in the inverse problem's solution. Thus, the generalized sensitivity index $S_5 = 2.12$ indicates that this unit norm perturbation will result in a change of about 2.12 in the norm of the solution of the inverse problem. Scaling by the norm of the solution of the inverse problem gives an interpretation that the unit norm perturbation shown in Figure 3.4 results in approximately a 14% change in the solution of the inverse problem. With this interpretation, a user may associate a level of uncertainty in the boundary condition and the resulting change in the estimated permeability field to determine if further calibration is needed.

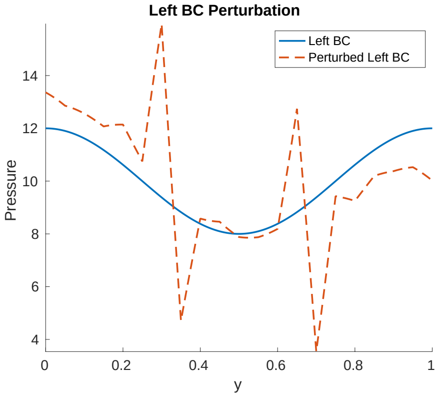

Figure 3.4: Perturbation of the left pressure Dirichlet boundary condition.

3.6.2 HDSA with Respect to Experimental Parameters

In this section, we turn to the pointwise hyper-differential sensitivities (3.4) to study the spatial and temporal dependence within the experimental parameters. Using the reconstructed log-permeability field, we compute the sensitivities of the solution with respect to both pressure and contaminant measurements at each sensor location, and for concentration, each time step. Figure 3.5 shows the spatial distribution of contaminant sensitivities (depicted by colored points using the right colorbar scale) at informative time snapshots, overlaid atop the tracer concentration field (depicted by a greyscale concentration map using the left colorbar scale). By overlaying these plots, we are able to study how the sensitivities relate to the tracer advection.

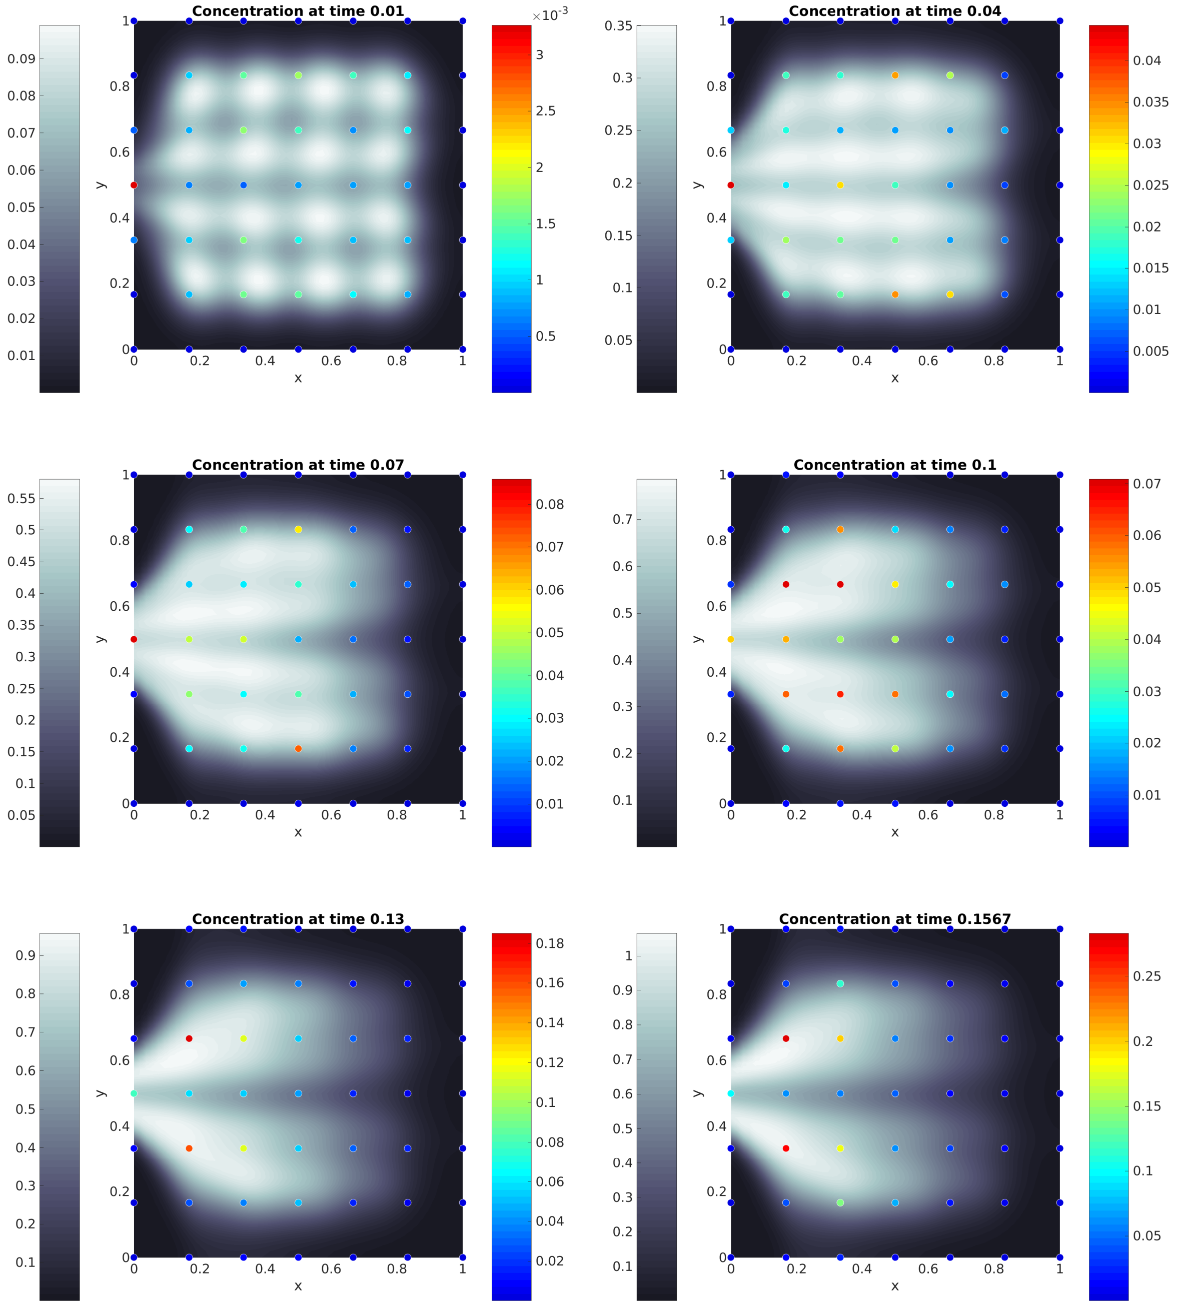

Figure 3.5: Concentration sensitivities at times 0.01, 0.04, 0.07, 0.10, 0.13, and 0.1567.

Note that both the sensitivity and concentration colorbar scales change in each plot, which is to allow the reader to visualize the results more clearly. We make the following observations about the concentration sensitivities. (1) As a general trend, the sensitivity of concentration increases in time. This is because the continuous source injection increases the amount of tracer in the domain as time progresses, making concentration sensors increasingly important. (2) We also see that as the mass of high tracer concentration (depicted by bright white in the color map) moves, sensors that observe this change in mass have increased importance while the mass moves toward or away from the sensor and then decrease in importance after the mass has moved passed. This phenomenon is particularly noticeable from the sensors in the high permeability channel at $y = 0.5$. As the tracer mass moves from right to left at $y = 0.5$ the sensor's importance increases following the back edge of the mass, and then decreases after the mass moves past. (3) We also notice that the sensor on the left boundary at $y = .5$ is highly important early in time and the sensors at $(0.2, 0.33)$ and $(0.2, 0.66)$ are more important later in time. This is because the majority of the tracer is getting advected toward those sensors which are therefore observing a large amount of tracer flow.

To further study the concentration sensitivities, Figure 3.6 depicts the time evolution of each concentration sensitivity in an array of plots. Each plot describes the time evolution of a single contaminant sensitivity and they are spatially arranged to correspond to the sensor locations they depict (compare with Figure 3.1 for their spatial association).

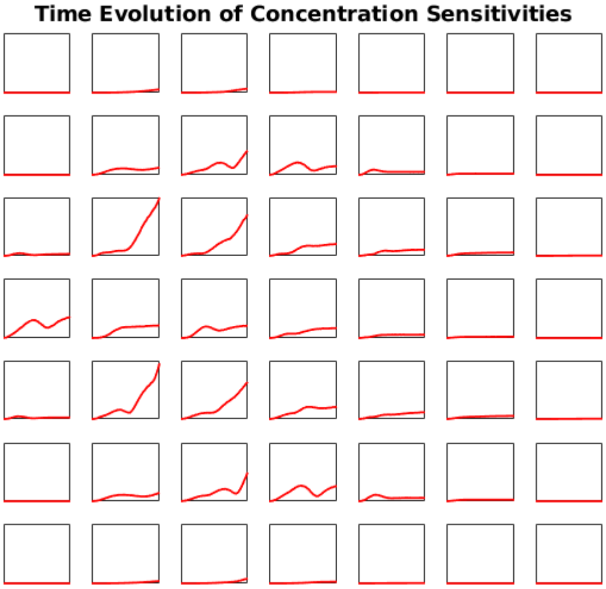

Figure 3.6: A spatial distribution of the time evolution of contaminant sensitivities. Each subplot has the same horizontal axis range depicting time from 0 to .16. Each vertical axis subplot has the same range depicting sensitivity from 0 to .2832

From Figure 3.6 we can see that some of the sensitivities decrease in importance over time, or have a range of time during which they decrease before they begin to increase again. Sensors for which the tracer permanently moves away experience a long-term decrease in importance. The decrease in sensitivity at a sensor location, followed by a subsequent increase is likely caused by the movement of tracer through the high permeability region at $y = 0.5$. As the tracer empties out of this central region, the concentration changes from being a single mass to being two distinct masses, one in the upper portion and the other in the lower portion of the domain. This splitting of the concentration mass, when observed by a nearby sensor, likely causes a minor disturbance in the general trend of the sensitivity.

Because this problem is advection dominated, the movement of the tracer through the domain has a large impact on the interpretation of the pressure sensitivities as well. Figure 3.7 shows the spatial distribution of pressure sensitivities (depicted by colored points using the right colorbar scale) overlaid atop the pressure field (depicted by a greyscale pressure map using the left colorbar scale).

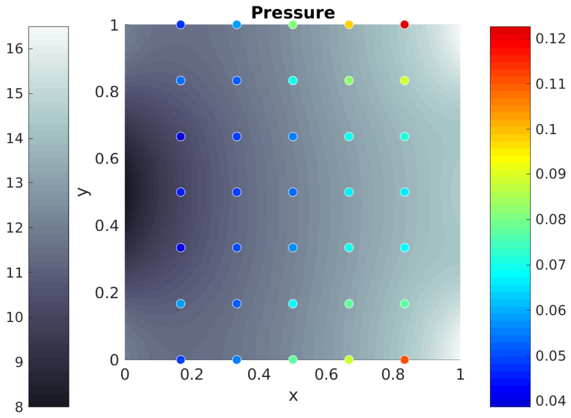

Figure 3.7: Pressure sensitivities

First, notice that the pressure sensitivities are larger than the concentration sensitivities at early time steps, but are eventually surpassed by the steadily increasing concentration sensitivities over time. Also observe that the sensors with highest pressure sensitivity, are in the upper right and lower right corners of the domain. Because the tracer moves from right to left, the majority of information about the tracer advection is found on the left side of the domain. The right side of the domain, and particularly the low permeability regions at the top and bottom, have a lack of tracer flow information. Thus, the inverse problem relies heavily on the pressure measurements in these regions to reconstruct the permeability field, which corresponds to the higher pressure sensitivities in this region.

3.6.3 HDSA with Respect to Auxiliary Parameters

In this section, we study the pointwise hyper-differential sensitivities with respect to auxiliary parameters: the source term and pressure Dirichlet boundary conditions. The diffusion coefficient is also an auxiliary parameter, however according to the generalized sensitivities it is relatively unimportant so we will not investigate it further.

We begin by analyzing the sensitivities with respect to the source term. Figure 3.8 depicts the sensitivities of the source at each injection well, next to a plot of the Darcy velocity field. The source uncertainty is discretized by taking a 3×3 mesh locally around each well. This models our uncertainty in the rate and distribution of the injected tracer at each well due to hardware limitations, while the location of the injected tracer is known.

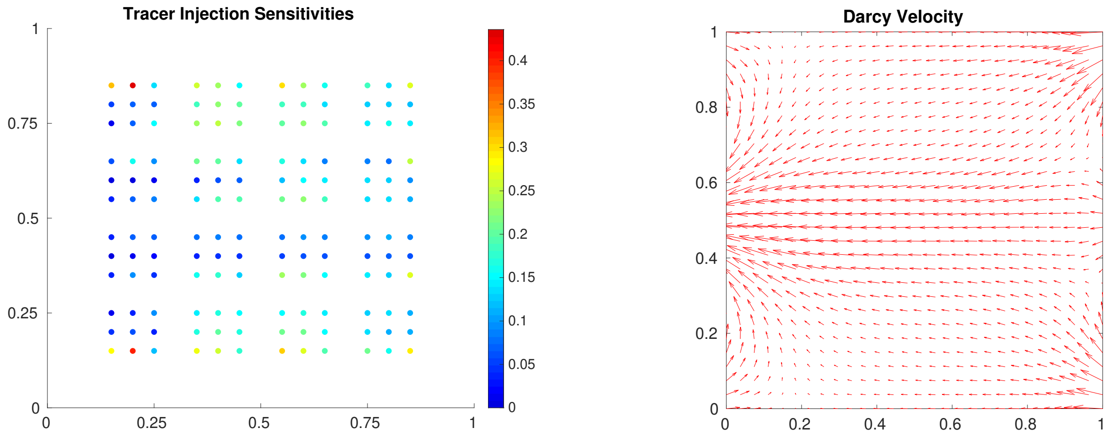

Figure 3.8: Left: Source sensitivities, Right: Darcy Velocity Field

From Figure 3.8 we can see that the areas of high source sensitivity generally occur in regions with low Darcy velocity. This is likely because in these regions, the problem is diffusion dominated and the tracer is not advected away from the injection site very quickly. Thus, if the source injection is perturbed, the tracer will stay in that region and slowly diffuse, affecting concentration measurements in that area for may time steps. If a source injection is perturbed in a region of high Darcy velocity, the tracer will be pulled away and mix with the rest of the tracer moving through the domain. Thus for this problem, the source injections have highest sensitivity in regions that are diffusion dominated.

According to the generalized sensitivities, the boundary conditions have the largest relative impact of any uncertain parameter on the solution. To further understand the influence of the boundary conditions on the physical systems in the model, we consider the hyper-differential sensitivities of the boundary conditions. Figure 3.9 depicts the sensitivities of the solution with respect to the pressure Dirichlet boundary conditions, discretized by 21 equally spaced nodes on each boundary.

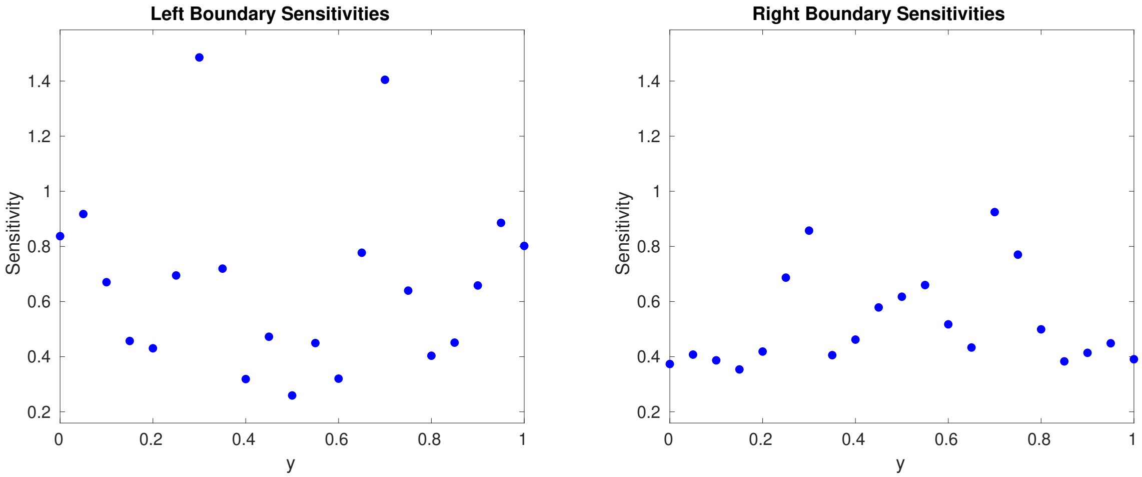

Figure 3.9: Pressure Dirichlet boundary condition sensitivities

From the boundary condition sensitivities in Figure 3.9 we can see that on both the left and right boundary, there is a heightened sensitivity around $y = 0.3$ and $y = 0.7$. This corresponds to the area between the high permeability region through the middle of the domain, and the low permeability regions above and below it. A perturbation of the boundary conditions near the low permeability region is going to have minimal effect because the low permeability region is going to keep the Darcy velocity small relative to the rest of the domain regardless. Similarly, a perturbation of the boundary conditions in the high permeability region is going to have little effect because the high permeability is going to keep the Darcy velocity relatively high in that area. A perturbation in the pressure boundary conditions will have maximal impact in the thin region between the high and low permeability regions. This perturbation can cause the region of moderate Darcy velocity to become a region of either high or low Darcy velocity relative to the other regions of the domain, significantly impacting the advection flow.

3.6.4 Discussion

It is evident from our analysis that the hyper-differential sensitivities can provide a wealth of information about how the solution of the inverse problem depends on the interactions within the governing physics systems. These observations are not readily apparent without the sensitivities, which emphasizes their usefulness in understanding the inverse problem. In particular, the sensitivities with respect to experimental parameters can be used to determine where and when to place expensive, high fidelity sensors in an experimental design, and where less accurate and more cost effective sensors can be uesd.

In addition to providing information about the underlying physics in a model, these observations allow experimenters to determine how to design experiments, by prioritizing the measurement and estimation of all complementary parameters considered. For this problem, the generalized sensitivities inform us that tracer concentration is more important to measure accurately than pressure, which informs the design of sensors and data collection techniques. We also learn that accurately estimating the pressure Dirichlet boundary conditions is highly important, while the diffusivity coefficient is relatively unimportant. This information informs the model specification and how these parameters should be estimated and considered in the model.

3.7 Conclusion

In this chapter, we have presented a mathematical framework for hyper-differential sensitivity analysis in the context of inverse problems constrained by multiphysics systems of partial differential equations. The mathematical formulation involves derivative based local sensitivity analysis of the solution of an inverse problem with respect to perturbations of parameters. This framework is general and can be applied to a wide variety of inverse problems. The usefulness of HDSA is most apparent in the context of complicated multiphysics systems with many uncertain parameters. By introducing sensitivity analysis with respect to experimental parameters and maturing the generalized hyper-differential sensitivity indices, we have enabled analysis of the relative importance of both auxiliary model parameters and experimental parameters, such as data sources. Studying hyper-differential sensitivities provides new insights into both understanding the underlying physical systems of a model, and designing experiments to solve inverse problems. In particular, comparing the relative importance of spatially and temporally distributed measurements with various sensor types provides unique insights that are difficult to attain using traditional experimental design methodologies. HDSA compliments experimental design by providing a systematic way to compare multiphysics parameters and data sources.

HDSA is an emerging technology and as such there are several important considerations to be studied in future work. One questions is “how robust are the sensitivities of experimental parameters to perturbations of design?” In practice, one may not be able to place a sensor exactly where a design indicates it should be. If a sensor’s location is perturbed within a local area, how will this affect the magnitude of the sensitivity, and that of its neighboring sensors? Ideally, perturbing the location of a sensor slightly will have a minimal impact on the sensitivity at that sensor and its neighbors, indicating that the hyper-differential sensitivities are robust to perturbations of the design, but this has yet to be rigorously verified.

Moreover, HDSA of inverse problems requires availability of measurement data. For the computational results in this chapter, we assumed that we had some set of experimental data, but in many applications we would like to compute sensitivities a priori, before data is collected. In such cases, one could generate training data by applying the forward model to a sample of the inversion parameters drawn from a prior distribution, giving rise to a distribution of the sensitivities.

3.8 Appendix: Details for the model inverse problem (3.7)

First, we illustrate the computation of gradient and Hessian of $\hat{J}$. To facilitate this, we introduce the Lagrangian

$$\mathcal{L}(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{\theta}) = \hat{J}(\boldsymbol{m}) + \boldsymbol{\lambda}^\top (\mathbf{L}(\boldsymbol{m})\boldsymbol{u} - \mathbf{V}\boldsymbol{\theta}),$$

where $\boldsymbol{\lambda}$ is a Lagrange multiplier (vector). To compute the gradient of $\hat{J}$, using the so called formal Lagrange approach, we consider the variations of $\mathcal{L}$ with respect to $\boldsymbol{\lambda}$, $\boldsymbol{u}$ and $\boldsymbol{m}$. Note that

$$\mathcal{L}_{\boldsymbol{\lambda}}(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{\theta}) = \mathbf{L}(\boldsymbol{m})\boldsymbol{u} - \mathbf{V}\boldsymbol{\theta} \quad \text{and} \quad \mathcal{L}_{\boldsymbol{u}}(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{\theta}) = \frac{\partial \hat{J}}{\partial \boldsymbol{u}} + \boldsymbol{\lambda}^\top \mathbf{L}(\boldsymbol{m}).$$

Setting these variations equal to 0 results in the state and the adjoint equations:

$$\mathbf{L}(\boldsymbol{m})\boldsymbol{u} = \mathbf{V}\boldsymbol{\theta} \quad \text{and} \quad \mathbf{L}(\boldsymbol{m})^\top \boldsymbol{\lambda} = -\mathbf{Q}^\top \mathbf{W}(\mathbf{Q}\boldsymbol{u} - \boldsymbol{y}).$$

Then, the gradient $g(\boldsymbol{m})$ of $\hat{J}$ satisfies, $g(\boldsymbol{m})^T = \mathcal{L}_{\boldsymbol{m}}(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{\theta}) = \frac{\partial \hat{J}}{\partial \boldsymbol{m}} + \boldsymbol{\lambda}^\top \frac{\partial \mathbf{A}}{\partial \boldsymbol{m}} \boldsymbol{u}$; therefore,

$$g(\boldsymbol{m}) = \alpha \mathbf{R}\boldsymbol{m} + \mathbf{C}^\top \boldsymbol{\lambda} \quad \text{with} \quad \mathbf{C} = \frac{\partial}{\partial \boldsymbol{m}} \big(\mathbf{L}(\boldsymbol{m})\boldsymbol{u}\big).$$

To compute the action of the Hessian $\mathbf{H}(\boldsymbol{m})$ of $\hat{J}$ (at $\boldsymbol{m}$) on a vector $\widehat{\boldsymbol{m}}$, we differentiate through the directional derivative $\langle g(\boldsymbol{m}), \widehat{\boldsymbol{m}} \rangle$. This is faciliated by introducing the "meta-Lagrangian":

$$\begin{aligned}\mathcal{L}^H(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{\theta}, \widehat{\boldsymbol{u}}, \widehat{\boldsymbol{\lambda}}; \widehat{\boldsymbol{m}}) = \widehat{\boldsymbol{m}}^\top \big[\mathbf{C}^\top \boldsymbol{\lambda} + \alpha \mathbf{R}\boldsymbol{m}\big] + \widehat{\boldsymbol{\lambda}}^\top \big[\mathbf{L}(\boldsymbol{m})\boldsymbol{u} - \mathbf{V}\boldsymbol{\theta}\big] + \\ \widehat{\boldsymbol{u}}^\top \big[\mathbf{L}(\boldsymbol{m})^\top \boldsymbol{\lambda} + \mathbf{Q}^\top \mathbf{W}(\mathbf{Q}\boldsymbol{u} - \boldsymbol{y})\big].\end{aligned}$$

The Lagrange multipliers $\widehat{\boldsymbol{u}}$ and $\widehat{\boldsymbol{\lambda}}$ are refered to as the incremental state and adjoint variables; see e.g., [69]. Letting the variations of $\mathcal{L}^H$ with respect to $\widehat{\boldsymbol{\lambda}}$ and $\widehat{\boldsymbol{u}}$ vanish gives

$$\mathbf{L}(\boldsymbol{m})\widehat{\boldsymbol{u}} = -\mathbf{C}\widehat{\boldsymbol{m}}, \qquad \text{(incremental state equation)} \tag{3.9}$$

$$\mathbf{L}(\boldsymbol{m})^\top\widehat{\boldsymbol{\lambda}} = -\mathcal{L}_{\boldsymbol{um}}\widehat{\boldsymbol{m}} - \mathcal{L}_{\boldsymbol{uu}}\widehat{\boldsymbol{u}}. \qquad \text{(incremental adjoint equation)} \tag{3.10}$$

The Hessian apply is then given by, $[\mathbf{H}(\boldsymbol{m})\widehat{\boldsymbol{m}}]^\top = \mathcal{L}^H_{\boldsymbol{m}}\widehat{\boldsymbol{m}}$, resulting in

$$\mathbf{H}(\boldsymbol{m})\widehat{\boldsymbol{m}} = \mathcal{L}_{\boldsymbol{mm}}\widehat{\boldsymbol{m}} + \mathcal{L}_{\boldsymbol{mu}}\widehat{\boldsymbol{u}} + \mathbf{C}^\top\widehat{\boldsymbol{\lambda}}. \tag{3.11}$$

Note that in the above equations

$$\mathcal{L}_{\boldsymbol{uu}} = \mathbf{Q}^\top\mathbf{W}\mathbf{Q}, \quad \mathcal{L}_{\boldsymbol{mm}} = \alpha\mathbf{R} + \frac{\partial}{\partial \boldsymbol{m}}(\mathbf{C}^T\boldsymbol{\lambda}), \quad \text{and} \quad \mathcal{L}_{\boldsymbol{mu}} = \mathcal{L}_{\boldsymbol{um}}^\top = \frac{\partial}{\partial \boldsymbol{m}}(\mathbf{L}(\boldsymbol{m})^T\boldsymbol{\lambda})$$

We summarize the compuation of $\mathbf{H}(\boldsymbol{m})\widehat{\boldsymbol{m}}$ in Algorithm 3. Note that the cost associates with Algorithm 3 is two (linear) PDE solves. Also, by replacing the expressions for the incremental state and adjoint variables in the expression for the Hessian apply, we can write the (reduced) Hessian as:

Algorithm 3 Computation of $\mathbf{H}(\boldsymbol{m})\widehat{\boldsymbol{m}}$ for a given $\widehat{\boldsymbol{m}}$.

```
solve the incremental state equation (3.9) for $\widehat{\boldsymbol{u}}$
Solve the incremental adjoint equation (3.10) for $\widehat{\boldsymbol{\lambda}}$
Evaluate $\mathbf{H}(\boldsymbol{m})\widehat{\boldsymbol{m}}$ accordng to (3.11).
```

$$\mathbf{H} = \mathbf{C}^\top\mathbf{L}(\boldsymbol{m})^{-\top}\mathcal{L}_{\boldsymbol{uu}}\mathbf{L}(\boldsymbol{m})^{-1}\mathbf{C} + \mathcal{L}_{\boldsymbol{mm}} - \mathcal{L}_{\boldsymbol{mu}}\mathbf{L}(\boldsymbol{m})^{-1}\mathbf{C} - \mathbf{C}^\top\mathbf{L}(\boldsymbol{m})^{-\top}\mathcal{L}_{\boldsymbol{um}}.$$

Letting $\mathbf{B}$ be the (discretized) Fréchet derivative of the gradient with respect to $\boldsymbol{\theta}$,

$$\mathbf{B} = g_{\boldsymbol{\theta}}(\boldsymbol{m}) = -\mathbf{C}^\top\mathbf{L}(\boldsymbol{m})^{-\top}\mathbf{Q}^\top\mathbf{W}\mathbf{Q}\mathbf{L}(\boldsymbol{m})^{-1}\mathbf{V} + \mathcal{L}_{\boldsymbol{mu}}\mathbf{L}(\boldsymbol{m})^{-1}\mathbf{V}.$$

This is the discretized version of the operator $\mathcal{B}$ in (3.3).

CHAPTER

4

HYPER-DIFFERENTIAL SENSITIVITY ANALYSIS FOR BAYESIAN INVERSE PROBLEMS

The work presented in this chapter is based upon the collaborative publication [79]. The author of this dissertation is listed as the first author of that paper and was responsible for the majority of the writing and results therein.

4.1 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 chapter we consider Bayesian inverse problems [62, 81], whose solution is a posterior distribution that is informed by both 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 location of measurement devices or their accuracy. We call these the *experimental parameters*. Throughout the chapter we will refer to the union of auxiliary and experimental parameters as *complementary parameters*. Our goal in this chapter is to develop methods for assessing the sensitivity of the solution of a Bayesian inverse problems with respect to perturbations of complementary parameters; see Section 4.2 for a simple illustrative example.

Understanding the sensitivity of an inverse problem 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 is informative to the important auxiliary parameters, the inverse problem may be redesigned to include these parameters in the set of inversion parameters.

The present chapter builds on the work in chapter 2 and other efforts in hyper-differential sensitivity analysis (HDSA) [16, 19, 29–34, 43, 80]. 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 chapter 3 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 chapter 3 for more details.

In the present chapter, we extend HDSA to the class of Bayesian inverse problems, which 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 4.3 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 chapter, is to assess the sensitivity of the maximum a posteriori probability (MAP) point to the complementary parameters. This builds directly on the developments in chapter 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) [3, 11, 20, 67, 72, 84]. Recall that in OED, one seeks 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., [20]. In the present chapter, we consider the Bayes risk, which has been used previously in OED for PDE-constrained inverse problems [37, 38, 46]. Our motivations for using the Bayes risk as an HDSA QoI are two-fold. First, the Bayes risk is defined as an average error of the MAP estimator; see Section 4.4.1. Thus, HDSA of the Bayes risk builds further on methods for HDSA of the MAP point. Moreover, it is well-known [4, 20] that the Bayes risk, with respect to the L^2 loss function, reduces to the A-optimality criterion, in the case of 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.

The contributions of this chapter are as follows:

- We develop a mathematical framework to assess the sensitivity of the Bayes risk and the MAP point in nonlinear Bayesian inverse problems with respect to complementary parameters; see Section 4.4.
- We present a scalable computational framework for computing the HDSA indices for the MAP point and the Bayes risk; see Section 4.5. 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 4.6 for the description of the model under study and Section 4.7 for our computational results.

4.2 An Illustrative Example

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

$$\frac{\partial y(x,t)}{\partial t} = \exp(m)\frac{\partial^2 y(x,t)}{\partial x^2}, \qquad x \in (0,\pi), t \in (0,1) \tag{4.1a}$$

$$y(0,t) = y(\pi,t) = 0, \qquad t \in (0,1) \tag{4.1b}$$

$$y(x,0) = \sin(x) + \exp(\theta)\sin(2x), \qquad x \in (0,\pi). \tag{4.1c}$$

In this problem, m is the inversion parameter and 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 $m \sim \mathcal{N}(1.3, .1)$. The problem (4.1) can be solved analytically and the solution is given by,

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

We have collected data measurements with additive Gaussian noise at the final time $t = 1$ at six evenly spaced points between $x = 0$ and $x = \pi$ depicted in Figure 4.1. The Gaussian noise is unbiased with a standard deviation of 26 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 $\theta = -.29$ and a perturbed value of $\theta = -.28$ in Figure 4.1, using data generated from the true parameter value of $\theta = -.3$.

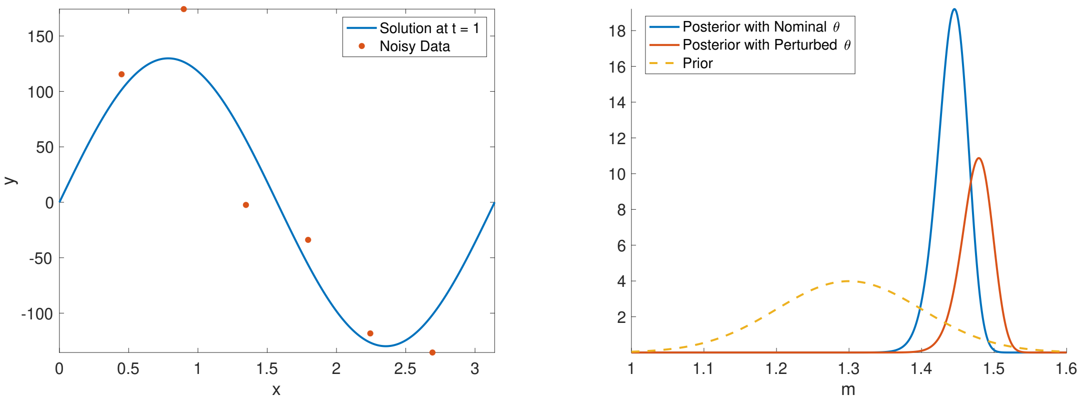

Figure 4.1: Left: the solution of the PDE at the final time with noisy data measurements, Right: the prior pdf of m and its posterior pdf for both a nominal and perturbed value of the auxiliary parameter θ.

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.

4.3 Bayesian inverse problems and complementary parameters

We let $\boldsymbol{\theta}_a$ and $\boldsymbol{\theta}_e$ denote the auxiliary and experimental parameters, and call the augmented parameter vector $\boldsymbol{\theta} = \begin{bmatrix} \boldsymbol{\theta}_a \\ \boldsymbol{\theta}_e \end{bmatrix}$ the complementary parameters. Note that it is possible to have some auxiliary parameters that are functions (see chapter 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 $\boldsymbol{\theta}$.

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

$$v(u, m, \boldsymbol{\theta}_a) = 0, \tag{4.2}$$

has a unique solution u for a given m and fixed auxiliary parameters $\boldsymbol{\theta}_a$. 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 chapter, $\mathcal{M} = L^2(\Omega)$ 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 is related to m according to the following model:

$$\boldsymbol{y} = \boldsymbol{F}(m, \boldsymbol{\theta}_a) + \boldsymbol{\eta}(\boldsymbol{\theta}_e), \tag{4.3}$$

where $\boldsymbol{y}$ is a vector of n_y experimental measurements, $\boldsymbol{F}$ 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, $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma}_{\text{noise}}(\boldsymbol{\theta}_e))$. Evaluating $\boldsymbol{F}(m, \boldsymbol{\theta}_a)$ requires solving the state equation (4.2) followed by application of an observation operator $\mathcal{O}$ which evaluates the state u at the n_y sensor locations. In the present chapter we let $\boldsymbol{\theta}_e$ 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 4.6.

The Bayesian inverse problem setup. To solve an inverse problem with the data model (4.3) we define the data likelihood pdf $\pi_{\text{like}}(\boldsymbol{y}|m; \boldsymbol{\theta})$, which describes the distribution of data measurements $\boldsymbol{y}$, given a particular inversion parameter m. Given our assumption of an additive Gaussian noise model, we have $\boldsymbol{y}|m \sim \mathcal{N}(\boldsymbol{F}(m), \boldsymbol{\Gamma}_{\text{noise}})$, and thus,

$$\pi_{\text{like}}(\boldsymbol{y}|m; \boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2}(\boldsymbol{F}(m, \boldsymbol{\theta}_a) - \boldsymbol{y}(\boldsymbol{\theta}_e))^\top \boldsymbol{\Gamma}_{\text{noise}}^{-1}(\boldsymbol{\theta}_e)(\boldsymbol{F}(m, \boldsymbol{\theta}_a) - \boldsymbol{y}(\boldsymbol{\theta}_e)) \right). \tag{4.4}$$

Note that in (4.4) the auxiliary parameters appear in the parameter-to-observable map $\boldsymbol{F}$ 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 chapter, 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., [3,18,62]. 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 $\mathscr{E} = \mathrm{range}(\mathcal{C}_{\mathrm{pr}}^{1/2})$, which is endowed with the following inner product,

$$\langle x, y\rangle_{\mathscr{E}} = \langle \mathcal{A}x, \mathcal{A}y\rangle.$$

We assume $m_{\mathrm{pr}} \in \mathscr{E}$.

The definition of the prior measure and the data likelihood completes the description of the Bayesian inverse problem. The solution of this inverse problem is called the posterior measure $\mu_{\mathrm{post}}^{\boldsymbol{y}}$, which describes the probability law of m conditioned on experimental measurements $\boldsymbol{y}$. 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 [62]

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

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

$$m^*(\boldsymbol{\theta}) \coloneqq \underset{m\in\mathscr{E}}{\mathrm{argmin}}\ J(m,\boldsymbol{\theta}), \tag{4.5}$$

where

$$J(m,\boldsymbol{\theta}) \coloneqq \frac{1}{2}\big(\boldsymbol{F}(m,\boldsymbol{\theta}_a)-\boldsymbol{y}(\boldsymbol{\theta}_e)\big)^\top \boldsymbol{\Gamma}_{\mathrm{noise}}^{-1}(\boldsymbol{\theta}_e)\big(\boldsymbol{F}(m,\boldsymbol{\theta}_a)-\boldsymbol{y}(\boldsymbol{\theta}_e)\big)+\frac{1}{2}\langle m-m_{\mathrm{pr}}, m-m_{\mathrm{pr}}\rangle_{\mathscr{E}}. \tag{4.6}$$

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

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

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 $(\mathbb{R}^{n}, \|\cdot\|_{\mathbf{M}})$ and $(\mathbb{R}^{n}, \|\cdot\|)$, where $\|\cdot\|$ is the Euclidean inner product, the adjoint operators need to be defined appropriately; see [18] for this and further details on discretization of different components of infinite-dimensional Bayesian inverse problems. In the remainder of this chapter, we present the proposed methods in the discretized setting.

4.4 HDSA for Nonlinear Bayesian Inverse Problems

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

4.4.1 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 (4.6) and (ii) the Bayes risk. The Bayes risk, for a fixed vector $\boldsymbol{\theta}$ of complementary parameters, is defined by

$$\Psi_{\text{risk}}(\boldsymbol{\theta}) = \int_{\mathcal{M}} \int_{\mathbb{R}^d} \|\boldsymbol{m}^*(\boldsymbol{\theta}) - \boldsymbol{m}\|^2_{\mathbf{M}} \pi_{\text{like}}(\boldsymbol{y}|\boldsymbol{m};\boldsymbol{\theta}) d\boldsymbol{y} \mu_{\text{pr}}^{n_m}(d\boldsymbol{m}). \tag{4.7}$$

Note that here we expressed the Bayes risk for the discretized version of the Bayesian inverse problem, and $\boldsymbol{m}^*$ 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 [18].

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

$$\boldsymbol{y_i} = \boldsymbol{F}(\boldsymbol{m}_i, \boldsymbol{\theta}_a) + \boldsymbol{\eta}_i(\boldsymbol{\theta}_e), \qquad i = 1, \dots, n_s,$$

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

$$\widehat{\Psi}_{\text{risk}}(\boldsymbol{\theta}) = \frac{1}{n_s} \sum_{i=1}^{n_s} \|\boldsymbol{m}^*(\boldsymbol{y}_i, \boldsymbol{\theta}) - \boldsymbol{m}_i\|^2_{\mathcal{M}}. \tag{4.8}$$

4.4.2 Sensitivity operator of Bayes risk

To compute the sensitivity operator of the approximate Bayes risk estimator, we first discretize (4.8) and denote its discretization by $\widehat{\boldsymbol{\Psi}}_{\text{risk}}(\boldsymbol{\theta})$. We assess the sensitivity of Bayes risk by differentiating $\widehat{\boldsymbol{\Psi}}_{\text{risk}}(\boldsymbol{\theta})$ with respect to θ_j, the jth component of $\boldsymbol{\theta}$, and evaluate at a set of nominal complementary parameter values $\bar{\boldsymbol{\theta}}$:

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

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

$$\mathbf{D}^{\mathrm{R}} = \begin{bmatrix} D_1^{\mathrm{R}} & D_2^{\mathrm{R}} & \dots & D_{n_\theta}^{\mathrm{R}} \end{bmatrix}, \tag{4.9}$$

where n_θ denotes the dimension of the complementary parameter vector. Note that $\mathbf{D}^{\mathrm{R}}\tilde{\boldsymbol{\theta}}$ 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 $\frac{\partial \boldsymbol{m}^*}{\partial \theta}(\boldsymbol{y}_i, \bar{\boldsymbol{\theta}})$, $i = 1, \dots, n_s$, which measures the sensitivity of the MAP points (for each data sample $\boldsymbol{y}_i$) to the complementary parameters [1]. For clarity, we denote the discretized cost functional by $\boldsymbol{J}$. As discussed in chapter 3,

$$\mathbf{D}^{\mathrm{M}} = \frac{\partial \boldsymbol{m}^*}{\partial \boldsymbol{\theta}} = -\left(\frac{\partial^2 \boldsymbol{J}}{\partial \boldsymbol{m}^2}\right)^{-1} \frac{\partial^2 \boldsymbol{J}}{\partial \boldsymbol{m} \partial \boldsymbol{\theta}} = -\mathbf{H}^{-1}\mathbf{B}, \tag{4.10}$$

where $\mathbf{H}$ and $\mathbf{B}$ are evaluated at the solution $\boldsymbol{m}^*$ with fixed nominal parameters $\bar{\boldsymbol{\theta}}$. 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 [80], the sensitivities of the inverse problem solution $\boldsymbol{m}^*$ require data measurements to compute. 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 is not available at the time of performing HDSA.

4.4.3 Sensitivity Indices

As detailed in section 3.3.2, we define sensitivity indices to provide a scalar which measures the magnitude of the change in the solution with respect to a particular perturbation of the complementary parameters. In section 3.3.2 we defined pointwise and generalized sensitivities for the MAP point, which work exactly the same in the context of this chapter. Herein we also define the pointwise sensitivity of the approximate Bayes risk as

---

[1] As discussed further in the Section 4.5, we only need to compute the action of this sensitivity operator on vectors

$$\mathbb{S}_k^j = \frac{|\mathbf{D}^{\mathrm{R}} \boldsymbol{e}_k^j|}{\|\boldsymbol{e}_k^j\|_\Theta} \tag{4.11}$$

These pointwise sensitivities measure the change in the Bayes risk to a perturbation of the kth parameter in the jth direction $\boldsymbol{b}_k^j$. We also define the generalized sensitivity of the kth subgroup of complementary parameters with respect to the approximate Bayes risk as,

$$\mathbb{S}_k = \max_{\boldsymbol{\theta} \in \Theta} \frac{|\mathbf{D}^{\mathrm{R}} \mathbf{T}_k \boldsymbol{\theta}|}{\|\boldsymbol{\theta}\|_\Theta}. \tag{4.12}$$

This formulation uses the selection operator $\mathbf{T}_k$ defined in 3.3.2. The generalized sensitivities measure the maximum change that can be observed in the Bayes risk with respect to a norm-1 perturbation of the kth parameter subgroup.

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 in a reasonable fashion. To do so, we divide the sensitivities with respect to the MAP point by the average norm of the computed MAP points, $\frac{1}{n_s}\sum_{i=1}^{n_s} \|\boldsymbol{m}^*(\boldsymbol{y}_i, \bar{\boldsymbol{\theta}})\|_{\mathbf{M}}$, and the sensitivities with respect to Bayes risk by the computed value of Bayes risk, $\widehat{\boldsymbol{\Psi}}_{\text{risk}}(\bar{\boldsymbol{\theta}})$.

4.5 Computational Methods

In this section, we present computational methods to implement the framework proposed in Section 4.4.

4.5.1 Computing the sensitivity indices

With the discretization of $\boldsymbol{m}$ and its prior, we can write the sensitivity operator of the approximate Bayes risk with respect to the complementary parameters (4.9) as,

$$\mathbf{D}^{\mathrm{R}} = \frac{2}{n_s} \sum_{i=1}^{n_s} (\mathbf{D}^{\mathrm{M}_i})^\top \mathbf{M} \left(\boldsymbol{m}^*(\boldsymbol{y}_i, \bar{\boldsymbol{\theta}}) - \boldsymbol{m}_i \right) = \frac{2}{n_s} \sum_{i=1}^{n_s} -\mathbf{B}_i^\top \mathbf{H}_i^{-\top} \mathbf{M} \left(\boldsymbol{m}^*(\boldsymbol{y}_i, \bar{\boldsymbol{\theta}}) - \boldsymbol{m}_i \right). \tag{4.13}$$

We let the subscript i on the operators $\mathbf{D}^{\mathrm{M}_i} \in \mathbb{R}^{n_m \times n_\theta}$, $\mathbf{B}_i \in \mathbb{R}^{n_m \times n_\theta}$, and $\mathbf{H}_i \in \mathbb{R}^{n_m \times n_m}$ indicate the dependence on the ith data sample.

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 as well as solve for the MAP point. We begin by defining the discrete Lagrangian as

$$\mathcal{L}(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{p}; \boldsymbol{\theta}) = \boldsymbol{J}(\boldsymbol{m}, \boldsymbol{\theta}) - \langle \boldsymbol{p}, \boldsymbol{v}(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{\theta}_a) \rangle_{\mathbb{M}} \tag{4.14}$$

where $\boldsymbol{v}(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{\theta}_a)$ is the discretized form of the PDE v, and $\boldsymbol{p}$ is the adjoint variable. Next, we use variational derivatives to compute the action of the discretized gradient of the cost function. We let $\mathcal{L}_{\boldsymbol{p}}[\hat{\boldsymbol{p}}]$ denote the variational derivative of (4.14) with respect to $\boldsymbol{p}$, acting on $\hat{\boldsymbol{p}}$, with the input arguments suppressed for brevity. A similar notation is used for the variational derivatives with respect to $\boldsymbol{u}$ and $\boldsymbol{m}$. We can also compute the action of the Hessian by constructing a meta-Lagrangian,

$$\mathcal{L}^H(\boldsymbol{u}, \boldsymbol{m}, \boldsymbol{p}, \hat{\boldsymbol{u}}, \hat{\boldsymbol{m}}, \hat{\boldsymbol{p}}; \boldsymbol{\theta}) = \mathcal{L}_{\boldsymbol{p}}[\hat{\boldsymbol{p}}] + \mathcal{L}_{\boldsymbol{u}}[\hat{\boldsymbol{u}}] + \mathcal{L}_{\boldsymbol{m}}[\hat{\boldsymbol{m}}]. \tag{4.15}$$

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 4, and we direct the reader to [35, 85] for additional details.

Algorithm 4 Compute the gradient g($\boldsymbol{m}$) and action of the Hessian $\mathbf{H}$ in the direction $\hat{\boldsymbol{m}}$

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

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 mata-Lagrangian with respect to $\boldsymbol{\theta}$ to obtain

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

where $\hat{\boldsymbol{u}}$ and $\hat{\boldsymbol{p}}$ satisfy the incremenal state and 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 $\boldsymbol{m}$, which will result in modified incremental equations. These adjoint based methods provide a computationally efficient method to evaluate the sensitivity operators $\mathbf{D}^{\mathrm{M}}$ and $\mathbf{D}^{\mathrm{R}}$.

To compute the discretized sensitivity operator $\mathbf{D}^{\mathrm{R}}$, we must first generate data samples $\boldsymbol{y}_i$ for $i = 1, \dots, n_s$ and then evaluate (4.13) which requires non-trivial computational cost. We also compute sensitivities of the MAP point, efficiently reusing PDE solves whenever applicable. This process is summarized in Algorithm 5. Note that for clarity, we have separated the processes of data generation and sensitivity operator computation in the algorithm. Furthermore, in Algorithm 5 the second subscript in sensitivity indices $S_{k,i}$ denotes dependence of the index upon the ith data sample.

Algorithm 5 Compute the sensitivity indices

```
% Data sample generation
for $i = 1$ to $n_s$ do
    Draw prior sample $\boldsymbol{m}_i$
    Solve the forward equation $\boldsymbol{v}(\boldsymbol{u}_i, \boldsymbol{m}_i, \bar{\boldsymbol{\theta}}_a) = \mathbf{0}$ for $\boldsymbol{u}_i$
    Synthesize data samples $\boldsymbol{y}_i = \mathcal{O}\boldsymbol{u}_i + \boldsymbol{\eta}_i(\bar{\boldsymbol{\theta}}_e)$ {$\mathcal{O}$ observes $\boldsymbol{u}$ at measurement locations}
end for
% Computation of the Bayes risk sensitivities
for $i = 1$ to $n_s$ do
    Solve the discretized inverse problem for $\boldsymbol{m}_i^*(\boldsymbol{d}_i, \bar{\boldsymbol{\theta}})$
    Solve $-\mathbf{H}_i \boldsymbol{z}_i = \mathbf{M}(\boldsymbol{m}_i^* - \boldsymbol{m}_i)$ for $\boldsymbol{z}_i$
    Compute $\boldsymbol{r}_i = \mathbf{B}_i^\top \boldsymbol{z}_i$
end for
$\mathbf{D}^{\mathrm{R}} = \frac{2}{n_s}\sum_{i=1}^{n_s} \boldsymbol{r}_i$
Compute $\mathbb{S}_k$ and $\mathbb{S}_k^j$ for all $k = 1, \dots, K$ and $j = 1, \dots, n_\theta$, see (4.11)
% Computation of the average MAP point sensitivities
for $i = 1$ to $n_s$ do
    Compute $S_{k,i}$ for $k = 1, \dots, K$
    for $j = 1$ to $n_\theta$ do
        Compute $S_{k,i}^j$ for $k = 1, \dots, K$
    end for
end for
Compute averaged generalized sensitivities $S_k \approx \bar{S}_k = \frac{1}{n_s}\sum_{i=1}^{n_s} S_{k,i}$, see (4.12)
Compute averaged pointwise sensitivities $S_k^j \approx \bar{S}_k^j = \frac{1}{n_s}\sum_{i=1}^{n_s} S_{k,i}^j$, see (4.12)
```

4.5.2 Computational Costs

Here we discuss the areas of high computational cost in Algorithm 5. To gain computational efficiency, we rely on some 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 [18, 26, 63, 68]. In particular, 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 PDEs solves, does not scale with the dimension of the discretized inversion parameter.

Generate 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 conjugate gradient line search algorithm with Armijo backtracking. 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 an appropriate search direction. Thus the total cost is $2L + 2LI$ PDE solves where L is the number of Newton steps taken. 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, 17, and 19. 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 PCG. Assuming that MAP point sensitivities are also desired, we can offset this cost by computing a low-rank approximation with the Lancoz method to apply the Hessian inverse efficiently, as detailed in [18]. After computing this low-rank approximation, application of the Hessian inverse can be approximated by matrix-vector products. The computational cost of the Lanczos method is $2r + 2$ 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 2 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 2 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 $\mathbf{D}^{\mathrm{M}}$ to standard basis vectors $\boldsymbol{e}_i$ (line 19) to compute n_θ pointwise sensitivity indices for all n_s data samples. As mentioned previously, this cost is significantly reduced by pre-computing 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.

Table 4.1: Computational costs summary

Computation	**Significant Cost per Sample (n_s)**
Data Generation	1 PDE solve
Inverse Problem Solves	$2L + 2LI$ PDE solves for L Newton steps and I PCG iterations
Hessian Inverse Approximation	$2r + 2$ PDE solves where r is the rank of the desired approximation
Bayes risk sensitivities	2 PDE solves
MAP point sensitivities	$2n_\theta$ PDE solves

The discussed computational costs are summarized in Table 4.1 for clarity. We remark that for the problem considered in the current chapter 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.

4.6 Model Problem

In this section, we present 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 4.6.1 followed by the setup of the Bayesian inverse problem in Section 4.6.2.

4.6.1 Forward Model

Consider the problem of infering 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 Γ_1, Γ_2, Γ_3, and Γ_4 denote the bottom, right, top, and left edges of Ω respectively,

$$
\begin{aligned}
-\nabla \cdot (e^m \nabla u) &= f && \text{in } \Omega, && \text{(4.16a)}\\
e^m \nabla u \cdot n &= 0 && \text{on } \Gamma_1 \cup \Gamma_3, && \text{(4.16b)}\\
e^m \nabla u \cdot n &= \beta(T_{\text{amb}} - u) && \text{on } \Gamma_2, && \text{(4.16c)}\\
e^m \nabla u \cdot n &= s && \text{on } \Gamma_4. && \text{(4.16d)}
\end{aligned}
$$

In this model, the inversion parameter $m(\boldsymbol{x})$ is a function representing the log of the heat conductivity of the non-homogeneous two-dimensional surface. We let $u(\boldsymbol{x})$ denote the temperature, $f(\boldsymbol{x})$ 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, the equations in (4.16) 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,

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

with auxiliary parameters $s_1, s_2,$ and s_3 fixed at nominal values $s_1 = 30, s_2 = .1,$ and $s_3 = .65$. The auxiliary parameters consist of the amplitude, spread, and location of the boundary heat source respectively. The heat source in the domain $f(\boldsymbol{x})$ is modeled as,

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

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

In this formulation, f_1 and f_2 control the amplitude of the heat sources, $\boldsymbol{w}$ and $\boldsymbol{z}$ 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: $f_1 = 100, f_2 = 105, \boldsymbol{w} = (.8, .25), \boldsymbol{z} = (.5, .8), \gamma_1 = -\pi/4, \gamma_2 = .15, \sigma_{x_1} = .8, \sigma_{x_2} = .1$. We consider the amplitude, center point, and angle of each bar to be uncertain and thus let $f_1, f_5, w_1, w_2, z_1, z_2, \gamma_1$, and γ_2 be the auxiliary parameters for the right hand side heat source $f(\boldsymbol{x})$. Figure 4.2 depicts this heat source in the domain.

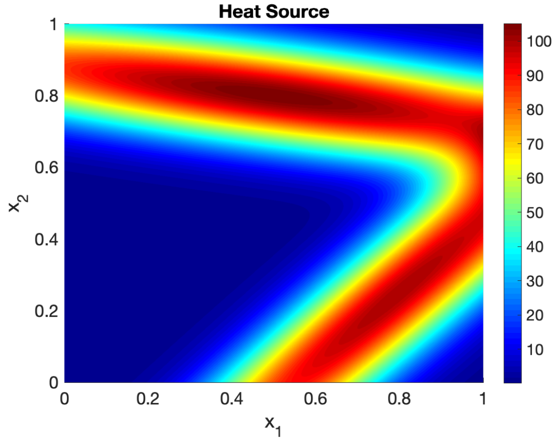

Figure 4.2: The heat source function $f(\boldsymbol{x})$

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 final parts. Variability in the powder material, boundary conditions, rasterization patterns, and laser power result in uneven heat distribution with problematic micro-crystallographic 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.

4.6.2 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. Note that we do not have any such "true solution" here and instead we compute data from samples of the prior 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:

$$m_{pr}(\boldsymbol{x}) = 1.5\sin(2\pi x_1)\cos(2\pi x_2) + 2.$$

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

$$\alpha \int_{\Omega} (\Phi \nabla m) \cdot \nabla q + mq \; d\boldsymbol{x} = \int_{\Omega} sq \; d\boldsymbol{x}$$

for all $q \in H^1(\Omega)$, with $\alpha = 5$, and $\Phi = .01$. 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 [18].

Measurements are collected on an evenly spaced 5×5 grid of observation locations depicted in Figure 4.3. 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 $\sigma = .1$, thus $\boldsymbol{\Gamma}_{\text{noise}} = \sigma^2\mathbf{I}$. 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 $\eta_i \sim \mathcal{N}(0, \boldsymbol{\Gamma}_{\text{noise}})$, directly proportional to the multiplicative perturbation of σ_i. Therefore the experimental parameters $\boldsymbol{\theta}_e$ enter the inverse problem through the cost function (4.6), both in the noise covariance matrix $\boldsymbol{\Gamma}_{\text{noise}}(\boldsymbol{\theta}_e)$ and the data measurements $y(\boldsymbol{\theta}_e)$ which depend on the noise realizations.

The solution of the governing PDE system detailed in (4.16) at the nominal parameter values with m fixed at the prior mean is depicted in Figure 4.3.

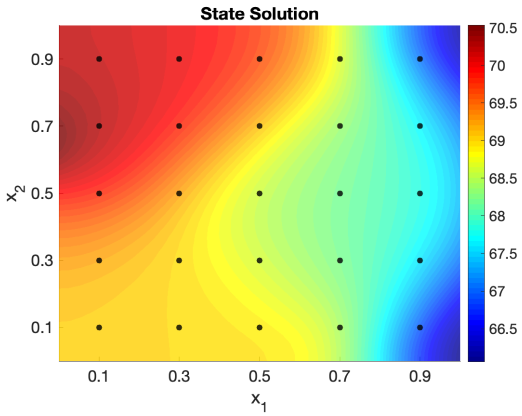

Figure 4.3: The state solution of the governing system of partial differential equations with the experimental sensor locations indicated by filled black circles.

4.7 Results

Using the model of heat flow across a conductive surface from Section 4.6, we solve an inverse problem to estimate the posterior distribution. Following Algorithm 5 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 (4.5), 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 Figure 4.4.

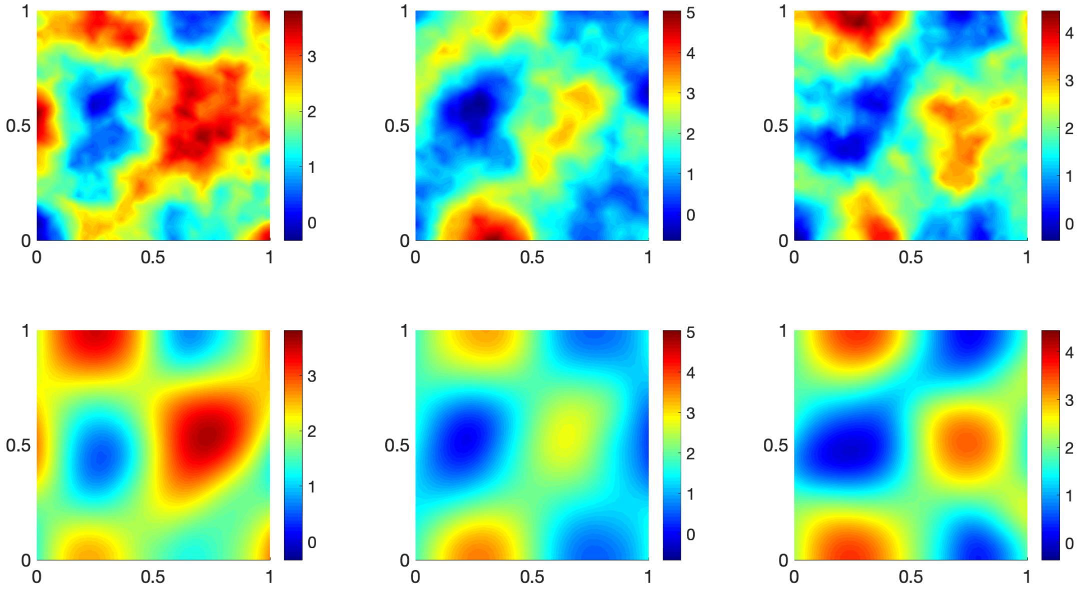

Figure 4.4: Top: three samples from the log-conductivity prior. Bottom: inverse problem MAP point estimates solved using data generated by the above prior sample.

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).

In Section 4.7.1 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 4.7.2) and the MAP point (Section 4.7.3). 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.

4.7.1 Modeling Parameter Perturbations

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

$$\rho = \tilde{\rho}(1 + a\theta) \tag{4.17}$$

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 (4.17).

In this particular model problem we use a perturbation scaling coefficient of $a = .05$ 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.

4.7.2 Sensitivities of Bayes Risk

The approximate Bayes risk is computed as a sample average as detailed in (4.8). We present the generalized sensitivities of each complementary parameter with respect to Bayes risk in Figure 4.5. 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 pre-computed samples.

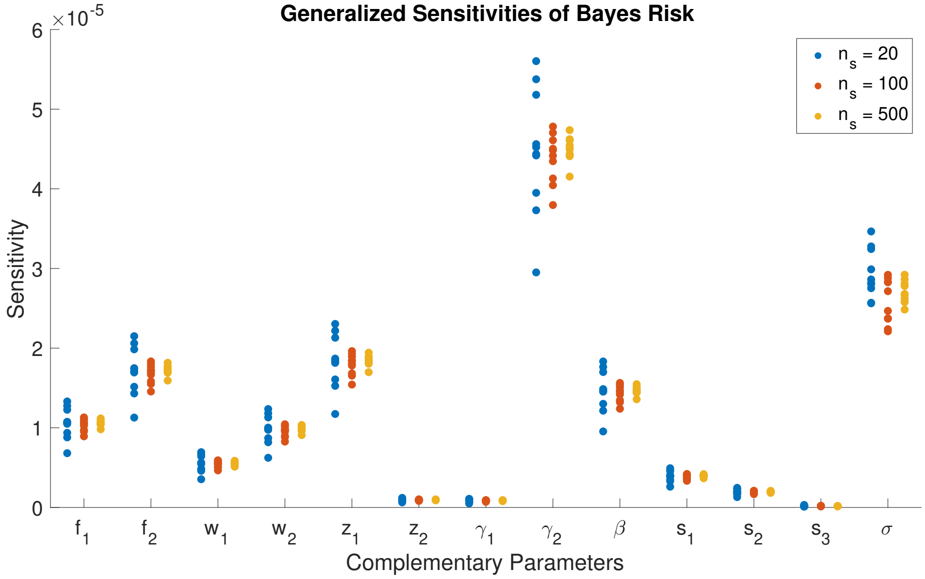

Figure 4.5: Generalized sensitivities of Bayes Risk to complementary parameters computed with sample sizes of 20, 100, and 500.

First let us discuss the spread of the samples. 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 provides 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 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_5) 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 ($\boldsymbol{\sigma}$). 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 $\gamma_2, f_5, z_1, \beta$, and $\boldsymbol{\sigma}$ are 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 ($\boldsymbol{\sigma}$) 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 in of 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 Figure 4.6. By perturbing the noise, we model perturbations of each collected data measurement in a way that we can experimentally control through sensor accuracy.

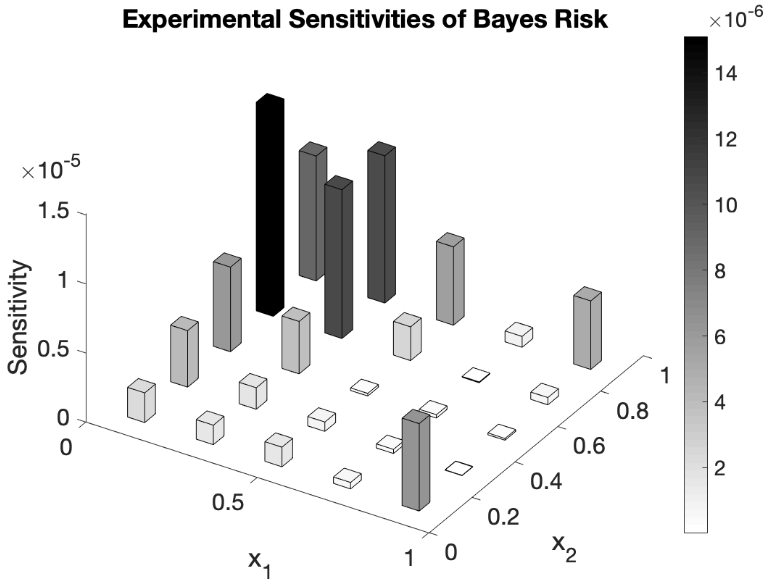

Figure 4.6: Point-wise sensitivities of Bayes Risk to experimental parameters computed from 1000 samples. Each bar represents the sensitivity of perturbing the standard deviation of the data noise at that particular sensor location.

We first note the scale of sensitivities presented in Figure 4.6. 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 Figure 4.5. 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 that the state solution depicted in Figure 4.3 is largest. This is also the region near the boundary source term s. We also note that the sensors at (.9,.1) and (.9,.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.

4.7.3 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 Figure 4.7.

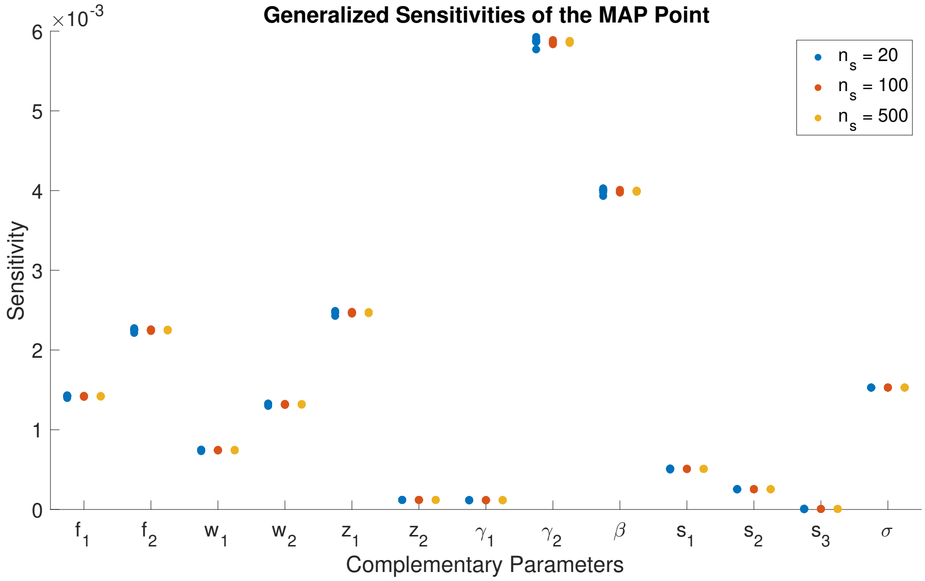

Figure 4.7: Averaged generalized sensitivities of the MAP point to complementary parameters computed with sample sizes of 20, 100, and 500.

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 greatest sensitivity to γ_2, β, z_1, and f_5. It is interesting to note that for Bayes risk, $\boldsymbol{\sigma}$ 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 Figure 4.8. 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 Figure 4.7 would indicate, there was minimal difference.

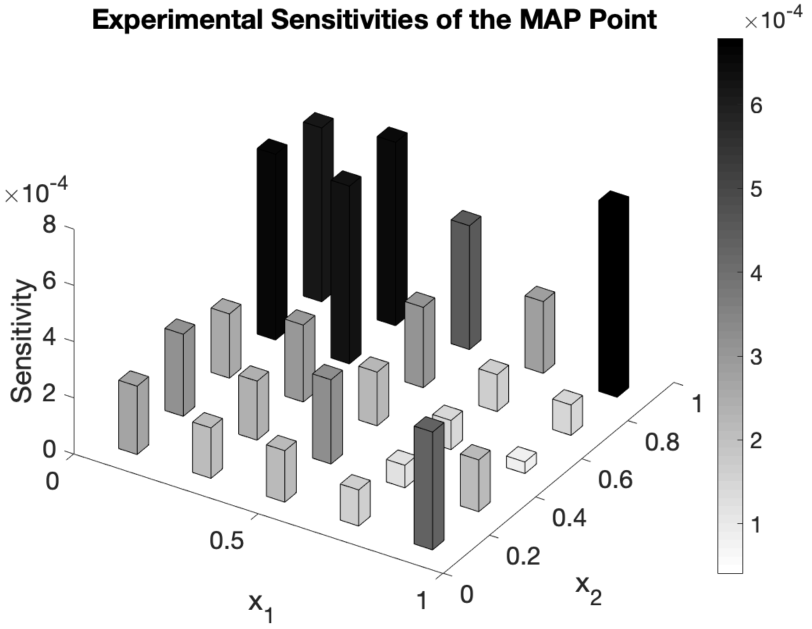

Figure 4.8: Point-wise sensitivities of the MAP point to experimental parameters computed as an average of 20 sample sensitivities.

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.

4.8 Conclusion

In this chapter we take foundational steps in applying hyper-differential sensitivity analysis (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 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 4.6 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 MAP point, made Bayes risk a suitable HDSA QoI in 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 optimal experimental design (OED) under uncertainty [7, 56]. 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 problem.

CHAPTER

5

CONCLUSION

The goal of the work provided herein is to introduce new methods that augment traditional approaches to understand and design PDE-constrained inverse problems. In particular, the work in chapters 2-4 seeks to understand how to account for and quantify the effect of uncertain model and experimental parameters, which are abundant in practical applications.

Chapter 2 studies optimal experimental designs of linear inverse problems and takes into account reducible model uncertainties by optimizing over a marginalized A-optimality criterion. Our experiments indicate that minimizing the marginalized A-optimality criterion offers a significant increase in the quality of the inferred inversion parameters, both in parameter estimation and uncertainty reduction. Chapter 3 examines inverse problems through the lens of post-optimal sensitivity analysis. The hyper-differential sensitivity analysis framework (HDSA) proposed therein offers a robust method of understanding the sensitivity of the solution to a large variety of uncertain parameters including those present in the governing model itself as well as the experimental conditions. While the work in chapter 3 pertains to deterministic inverse problems, chapter 4 extends this work to Bayesian inverse problems which presents many new mathematical and computational challenges. This work measures posterior uncertainty through Bayes risk which is equivalent to the average posterior variance in linear inverse problems and is used as a measure of solution quality for nonlinear inverse problems. Each of these works shows that additional uncertainties present in the model and experimental conditions can have significant effects upon the inverse problem solution and should be taken into account when solving these problems.

There are are various lines of future investigation stemming from this work. While chapter 2 presents optimal experimental design under uncertainty (OEDUU) for linear inverse problems, this work should be extended to nonlinear inverse problems as well. Chapters 2 and 4 are very synergistic and can be utilized together to efficiently explore high-dimensional OEDUU problems. OEDUU problems with a large number of uncertain auxiliary parameters become very computationally expensive. In this case, we can leverage the Bayesian HDSA methodology to determine which uncertain parameters the inverse problem is most sensitive to. One can then consider only these important parameters as uncertain for the OED computation, fixing the remaining parameters at nominal values and drastically reducing the computational expense of the OEDUU problem. While this outlines a clear framework for utilizing HDSA and OEDUU together, it may be possible to connect the two methods in other ways. Two additional interesting directions of investigation include studying how experimental sensitivities can be used to inform the OED problem and examining how the sensor placement problem itself, rather than the solution of the inverse problem, is sensitive to perturbations of the complementary parameters. The work in chapter 4 is a foundational work in the direction of HDSA for Bayesian inverse problems. While this dissertation presents a computationally efficient method for computing sensitivities of Bayes risk, there are many other measures of posterior uncertainty, such as the various OED criteria, which HDSA can be extended to.

References

[1] V. Akçelik, G. Biros, A. Dr***u***ag***u***anescu, O. Ghattas, J. Hill, and B. van Bloeman Waanders. Dynamic data-driven inversion for terascale simulations: Real-time identification of airborne contaminants. In *Proceedings of SC2005*, Seattle, 2005.

[2] V. Akçelik, G. Biros, O. Ghattas, J. Hill, D. Keyes, and B. van Bloemen Waanders. *Parallel Processing for Scientific Computing*, chapter Parallel Algorithms for PDE-Constrained Optimization. SIAM, 2006.

[3] A. Alexanderian. Optimal experimental design for Bayesian inverse problems governed by PDEs: A review. *Inverse Problems*, 2021.

[4] A. Alexanderian, P. J. Gloor, and O. Ghattas. On Bayesian A- and D-Optimal Experimental Designs in Infinite Dimensions. *Bayesian Analysis*, 11(3):671–695, 2016.

[5] A. Alexanderian, N. Petra, G. Stadler, and O. Ghattas. A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0-sparsification. *SIAM Journal on Scientific Computing*, 36(5):A2122–A2148, 2014.

[6] A. Alexanderian, N. Petra, G. Stadler, and O. Ghattas. A fast and scalable method for a-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems. *SIAM Journal on Scientific Computing*, 38(1), 2016.

[7] A. Alexanderian, N. Petra, G. Stadler, and I. Sunseri. Optimal design of large-scale Bayesian linear inverse problems under reducible model uncertainty: good to know what you don't know. *SIAM/ASA Journal on Uncertainty Quantification*, 9(1):163–184, 2021.

[8] A. Alexanderian and A. K. Saibaba. Efficient D-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems. *SIAM Journal on Scientific Computing*, 40(5):A2956–A2985, 2018.

[9] A. Ambrosetti and G. Prodi. *A Primer of Nonlinear Analysis.* Cambridge University Press, 1995.

[10] A. Y. Aravkin and T. Van Leeuwen. Estimating nuisance parameters in inverse problems. *Inverse Problems*, 28(11):115016, 2012.

[11] A. C. Atkinson and A. N. Donev. *Optimum Experimental Designs.* Oxford, 1992.

[12] A. Attia, A. Alexanderian, and A. K. Saibaba. Goal-oriented optimal design of experiments for large-scale Bayesian linear inverse problems. *Inverse Problems*, 34(9):095009, 2018.

[13] I. Bauer, H. G. Bock, S. Körkel, and J. P. Schlöder. Numerical methods for optimum experimental design in DAE systems. *Journal of Computational and Applied Mathematics*, 120(1-2):1–25, 2000. SQP-based direct discretization methods for practical optimal control problems.

[14] R. N. Bhattacharya and E. C. Waymire. *A basic course in probability theory*, volume 69. Springer, 2007.

[15] H. G. Bock, S. Körkel, and J. P. Schlöder. Parameter estimation and optimum experimental design for differential equation models. In H. G. Bock, T. Carraro, W. Jäger, S. Körkel, R. Rannacher, and J. P. Schlöder, editors, *Model Based Parameter Estimation*, volume 4 of *Contributions in Mathematical and Computational Sciences*, pages 1–30. Springer Berlin Heidelberg, 2013.

[16] K. Brandes and R. Griesse. Quantitative stability analysis of optimal solutions in PDE-constrained optimization. *Journal of Computational and Applied Mathematics*, 2006.

[17] T. Bui-Thanh. A gentle tutorial on statistical inversion using the Bayesian paradigm. Technical report, The University of Texas at Austin, 2012.

[18] T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler. A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494–A2523, 2013.

[19] C. Büskens and R. Griesse. Parametric sensitivity analysis of perturbed PDE optimal control problems with state and control constraints. *Journal of Optimization Theory and Applications*, 131(1):17–35, 2006.

[20] K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995.

[21] L. Chamon and A. Ribeiro. Approximate supermodularity bounds for experimental design. In *Advances in Neural Information Processing Systems*, pages 5403–5412, 2017.

[22] E. M. Constantinescu, N. Petra, J. Bessac, and C. G. Petra. Statistical treatment of inverse problems constrained by differential equations-based models with stochastic terms. *SIAM/ASA Journal on Uncertainty Quantification*, 8(1):170–197, 2020.

[23] Y. Daon and G. Stadler. Mitigating the influence of boundary conditions on covariance operators derived from elliptic PDEs. *Inverse Problems and Imaging*, 12(5):1083–1102, 2018.

[24] M. Dashti and A. M. Stuart. The Bayesian approach to inverse problems. *Handbook of Uncertainty Quantification*, 2017.

[25] H. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems.* Kluwer Academic Publishers, 1996.

[26] H. P. Flath, L. C. Wilcox, V. Akçelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas. Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial hessian approximations. *SIAM Journal on Scientific Computing*, 33(1):407–432, 2011.

[27] B. Flury. *A first course in multivariate statistics.* Springer Science & Business Media, 2013.

[28] J. Fohring and E. Haber. Adaptive A-optimal experimental design for linear dynamical systems. *SIAM-ASA Journal on Uncertainty Quantification*, 4:1138–1159, 01 2016.

[29] R. Griesse. Parametric sensitivity analysis in optimal control of a reaction-diffusion system – part II: practical methods and examples. *Optimization Methods and Software*, 19(2):217–242, 2004.

[30] R. Griesse. Parametric sensitivity analysis in optimal control of a reaction diffusion system. I. solution differentiability. *Numerical Functional Analysis and Optimization*, 25(1-2):93–117, 2004.

[31] R. Griesse. Stability and sensitivity analysis in optimal control of partial differential equations. Habilitation Thesis, Faculty of Natural Sciences, Karl-Franzens University, 2007.

[32] R. Griesse and B. Vexler. Numerical sensitivity analysis for the quantity of interest in PDE-Constrained optimization. *SIAM Journal on Scientific Computing*, 29(1):22–48, 2007.

[33] R. Griesse and S. Volkwein. Parametric sensitivity analysis for optimal boundary control of a 3D reaction-diffusion system. In G. D. Pillo and M. Roma, editors, *Nonconvex Optimization and its Applications*, volume 83. Springer, Berlin, 2006.

[34] R. Griesse and A. Walther. Parametric sensitivities for optimal control problems using automatic differentiation. *Optimal Control Applications and Methods*, 24:297–314, 2003.

[35] M. D. Gunzburger. *Perspectives in Flow Control and Optimization.* SIAM, 2003.

[36] A. Gut. *An intermediate course in probability.* Springer Texts in Statistics. Springer, New York, second edition, 2009.

[37] E. Haber, L. Horesh, and L. Tenorio. Numerical methods for experimental design of large-scale linear ill-posed inverse problems. *Inverse Problems*, 24(055012):125–137, 2008.

[38] E. Haber, L. Horesh, and L. Tenorio. Numerical methods for the design of large-scale nonlinear discrete ill-posed inverse problems. *Inverse Problems*, 26(2):025002, 2010.

[39] E. Haber, Z. Magnant, C. Lucero, and L. Tenorio. Numerical methods for A-optimal designs with a sparsity constraint for ill-posed inverse problems. *Computational Optimization and Applications*, pages 1–22, 2012.

[40] M. S. Handcock and M. L. Stein. A Bayesian analysis of kriging. *Technometrics*, 35(4):403–410, 1993.

[41] J. Hart and P. A. Gremaud. An approximation theoretic perspective of Sobol' indices with dependent variables. *International Journal for Uncertainty Quantification*, 8(6), 2018.

[42] J. Hart and P. A. Gremaud. Robustness of the Sobol' indices to distributional uncertainty. *International Journal for Uncertainty Quantification*, 9(5), 2019.

[43] J. Hart, B. van Bloemen Waanders, and R. Herzog. Hyper-differential sensitivity analysis of uncertain parameters in PDE-constrained optimization. *International Journal for Uncertainty Quantification*, 10(3):225–248, 2020.

[44] J. L. Hart, A. Alexanderian, and P. A. Gremaud. Efficient computation of Sobol' indices for stochastic models. *SIAM Journal on Scientific Computing*, 39(4):A1514–A1530, 2017.

[45] E. Herman, A. Alexanderian, and A. K. Saibaba. Randomization and reweighted ℓ_1-minimization for A-optimal design of linear inverse problems. *SIAM Journal on Scientific Computing*, accepted, 2020. `https://arxiv.org/abs/1906.03791`.

[46] L. Horesh, E. Haber, and L. Tenorio. *Optimal Experimental Design for the Large-Scale Nonlinear Ill-Posed Problem of Impedance Imaging*, pages 273–290. Wiley, 2010.

[47] X. Huan and Y. M. Marzouk. Simulation-based optimal Bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232(1):288–317, 2013.

[48] B. Iooss and P. Lemaître. A review on global sensitivity analysis methods. In *Uncertainty management in simulation-optimization of complex systems*, pages 101–122. Springer, 2015.

[49] B. Iooss and A. Saltelli. Introduction to sensitivity analysis. *Handbook of uncertainty quantification*, pages 1103–1122, 2017.

[50] K. Ito and B. Jin. *Inverse problems: Tikhonov theory and algorithms.* World Scientific, 2015.

[51] J. Jagalur-Mohan and Y. Marzouk. Batch greedy maximization of non-submodular functions: Guarantees and applications to experimental design. *Journal of Machine Learning Research*, 22(252):1–62, 2021.

[52] J. Kaipio and V. Kolehmainen. Approximate marginalization over modeling errors and uncertainties in inverse problems. *Bayesian Theory and Applications*, pages 644–672, 2013.

[53] J. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems.* Springer, 2005.

[54] V. Kolehmainen, T. Tarvainen, S. R. Arridge, and J. P. Kaipio. Marginalization of uninteresting distributed parameters in inverse problems-application to diffuse optical tomography. *International Journal for Uncertainty Quantification*, 1(1), 2011.

[55] S. Körkel, E. Kostina, H. G. Bock, and J. P. Schlöder. Numerical methods for optimal control problems in design of robust optimal experiments for nonlinear dynamic processes. *Optimization Methods & Software*, 19(3-4):327–338, 2004. The First International Conference on Optimization Methods and Software. Part II.

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

[57] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9:235–284, 2008.

[58] F. Lindgren, H. Rue, and J. Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498, 2011.

[59] Q. Long, M. Motamed, and R. Tempone. Fast Bayesian optimal experimental design for seismic source inversion. *Computer Methods in Applied Mechanics and Engineering*, 291:123–145, 2015.

[60] Q. Long, M. Scavino, R. Tempone, and S. Wang. Fast estimation of expected information gains for Bayesian experimental designs based on Laplace approximations. *Computer Methods in Applied Mechanics and Engineering*, 259:24–39, 2013.

[61] T.-T. Lu and S.-H. Shiou. Inverses of 2×2 block matrices. *Computers & Mathematics with Applications*, 43(1-2):119–129, 2002.

[62] A. M. Stuart. Inverse problems: A Bayesian perspective. *Acta Numerica*, 19:451–559, 05 2010.

[63] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas. A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012.

[64] J. B. Nagel. *Bayesian techniques for inverse uncertainty quantification.* PhD thesis, ETH Zurich, 2017.

[65] R. Nicholson, N. Petra, and J. P. Kaipio. Estimation of the Robin coefficient field in a Poisson problem with uncertain conductivity field. *Inverse Problems*, 34(11):115005, 2018.

[66] J. M. Ortega. *Matrix theory.* The University Series in Mathematics. Plenum Press, New York, 1987. A second course.

[67] A. Pázman. *Foundations of Optimum Experimental Designs.* D. Reidel Publishing Co., 1986.

[68] N. Petra, J. Martin, G. Stadler, and O. Ghattas. A computational framework for infinite-dimensional Bayesian inverse problems, part ii: Stochastic Newton MCMC with application to ice sheet flow inverse problems. *SIAM Journal on Scientific Computing*, 36(4):A1525–A1555, 2014.

[69] N. Petra and G. Stadler. Model variational inverse problems governed by partial differential equations. Technical Report 11-05, The Institute for Computational Engineering and Sciences, The University of Texas at Austin, 2011.

[70] G. D. Prato. *An Introduction to Infinite-Dimensional Analysis*. Springer-Verlag Berlin Heidelberg, 2006.

[71] C. Prieur and S. Tarantola. Variance-based sensitivity analysis: Theory and estimation algorithms. *Handbook of uncertainty quantification*, pages 1217–1239, 2017.

[72] F. Pukelsheim. *Optimal Design of Experiments*. John Wiley & Sons, New-York, 1993.

[73] L. Roininen, J. M. Huttunen, and S. Lasanen. Whittle-Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography. *Inverse Probl. Imaging*, 8(2):561, 2014.

[74] L. Ruthotto, J. Chung, and M. Chung. Optimal experimental design for inverse problems with state constraints. *SIAM Journal on Scientific Computing*, 40(4):B1080–B1100, 2018.

[75] G. Shulkind, L. Horesh, and H. Avron. Experimental design for nonparametric correction of misspecified dynamical models. *SIAM/ASA Journal on Uncertainty Quantification*, 6(2):880–906, 2018.

[76] R. C. Smith. *Uncertainty quantification: Theory, implementation, and applications*, volume 12 of *Computational Science and Engineering Series*. SIAM, 2013.

[77] I. Sobol and S. Kucherenko. Derivative based global sensitivity measures. *Procedia-Social and Behavioral Sciences*, 2(6):7745–7746, 2010.

[78] I. Sunseri and A. Alexanderian. On marginals of Gaussian random vectors. Technical report, North Carolina State University, 2018.

[79] I. Sunseri, A. Alexanderian, J. Hart, and B. van Bloemen Waanders. Hyper differential sensitivity analysis for nonlinear Bayesian inverse problems. *Submitted*, 2022.

[80] I. Sunseri, J. Hart, B. van Bloemen Waanders, and A. Alexanderian. Hyper-differential sensitivity analysis for inverse problems constrained by partial differential equations. *Inverse Problems*, 36(12):125001, 2020.

[81] A. Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation.* Society for Industrial and Applied Mathematics, 2005.

[82] Y. L. Tong. *The multivariate normal distribution.* Springer Science & Business Media, 2012.

[83] F. Tröltzsch. *Optimal Control of Partial Differential Equations: Theory, Methods and Applications*, volume 112 of *Graduate Studies in Mathematics.* American Mathematical Society, 2010.

[84] D. Uciński. *Optimal Measurement Methods for Distributed Parameter System Identification.* CRC Press, Boca Raton, 2005.

[85] U. Villa, N. Petra, and O. Ghattas. hIPPYlib: An extensible software framework for large-scale inverse problems governed by PDEs: Part i: Deterministic inversion and linearized Bayesian inference. *ACM Transactions on Mathematical Software (TOMS)*, 47(2):1–34, 2021.

[86] C. Vogel. *Computational Methods for Inverse Problems.* Society for Industrial and Applied Mathematics, 2002.

[87] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

[88] D. Williams. *Probability with martingales.* Cambridge university press, 1991.