

PAPER

Optimal experimental design under irreducible uncertainty for linear inverse problems governed by PDEs

To cite this article: Karina Koval *et al* 2020 *Inverse Problems* **36** 075007

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



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Optimal experimental design under irreducible uncertainty for linear inverse problems governed by PDEs

Karina Koval¹, Alen Alexanderian² and Georg Stadler^{1,3} 

¹ Courant Institute of Mathematical Sciences, New York University, New York, NY, United States of America

² Department of Mathematics, North Carolina State University, Raleigh, NC, United States of America

E-mail: koval@cims.nyu.edu, alexanderian@ncsu.edu and stadler@cims.nyu.edu

Received 15 December 2019, revised 30 March 2020

Accepted for publication 16 April 2020

Published 17 June 2020



CrossMark

Abstract

We present a method for computing A-optimal sensor placements for infinite-dimensional Bayesian linear inverse problems governed by PDEs with irreducible model uncertainties. Here, irreducible uncertainties refers to uncertainties in the model that exist in addition to the parameters in the inverse problem, and that cannot be reduced through observations. Specifically, given a statistical distribution for the model uncertainties, we compute the optimal design that minimizes the expected value of the posterior covariance trace. The expected value is discretized using Monte Carlo leading to an objective function consisting of a sum of trace operators and a binary-inducing penalty. Minimization of this objective requires a large number of PDE solves in each step. To make this problem computationally tractable, we construct a composite low-rank basis using a randomized range finder algorithm to eliminate forward and adjoint PDE solves. We also present a novel formulation of the A-optimal design objective that requires the trace of an operator in the observation rather than the parameter space. The binary structure is enforced using a weighted regularized ℓ_0 -sparsification approach. We present numerical results for inference of the initial condition in a subsurface flow problem with inherent uncertainty in the flow fields and in the initial times.

Keywords: optimal design, inverse problems, model uncertainty, optimization under uncertainty, model reduction, subsurface flow

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

³Author to whom any correspondence should be addressed.

1. Introduction

Many problems in the sciences and engineering require inference of unknown/uncertain parameters from indirect observations and a mathematical model that relates the parameters to these observations. Having access to informative data is integral to accurate parameter inference. However, the cost of data collection or physical restrictions often limit how much data one can collect. Even if one has access to large stores of data, processing large amounts of it can be expensive; moreover, data can be full of redundancies—poor experimental design choices can limit the information data contains about the parameters. A natural question to ask is: how can we design experimental conditions for data collection to optimally reconstruct/infer parameters of interest? Addressing this requires solving an optimal experimental design (OED) problem.

Inverse problems stemming from real-world applications often contain uncertainties in the governing model, in addition to the uncertain inversion parameters. We focus on *irreducible* model uncertainties, i.e., uncertainties that—for all practical purposes—cannot be reduced via parameter estimation. Examples of such irreducible uncertainties arise in inversion of the initial concentration of a contaminant in groundwater flow. In this problem, one usually only has rough estimates of the true groundwater velocity field. Additionally, one might not exactly know the exact time at which the contaminant has been released. Unlike uncertainties that could be reduced with observations and parameter estimation, these uncertainties are difficult to reduce. However, they should be taken into account when computing experimental designs. This article is about design of experiments for inverse problems governed by models containing such irreducible uncertainties.

We focus on infinite-dimensional Bayesian inverse problems in which the parameter-to-observable map is linear; see section 2 for a brief overview. The irreducible uncertainty may enter nonlinearly into the model and we assume some knowledge about its distribution. This knowledge could be based, for instance, on historic data, or it could be the posterior distribution obtained as the solution of another Bayesian inverse problem. To accommodate such an irreducible additional uncertainty in OED, we extend the notion of A-optimal design to A-optimal design under (irreducible) uncertainty, which we define as the design that minimizes the expected value of the average posterior variance; see section 3.

In the present work, we restrict the idea of experimental design to that of choosing locations for placing data collecting sensors, though our approach can be adapted to more general experimental design problems. We formulate the OED problem as that of finding an optimal subset of locations for sensor placement from a pre-specified network of candidate sensor locations. We assign to each candidate location a non-negative *design weight* indicating its importance. We seek binary designs and interpret a candidate sensor location with a design weight of 1 as a location where a sensor should be placed. To circumvent the computational challenges of binary optimization, we relax the binary condition and add a sparsity-inducing (or binary-inducing) penalty to our optimal design objective; see section 3.3.

Literature survey and challenges. Classical references for optimal experimental design problems include [6, 25, 27, 32]. OED for inverse problems governed by computationally intensive models has been subject to intense research activity in the past couple of decades; see, e.g., [4, 7, 16, 19–21, 23, 29, 34]. Our focus is on ill-posed infinite-dimensional Bayesian *linear* inverse problems. We build on previous work [3, 15, 17], which developed efficient methods for computing A-optimal designs for high- or infinite-dimensional linear inverse problems using either a Bayesian or a frequentist approach. Other approaches for computing optimal experimental designs for high/infinite-dimensional linear inverse problems are explored, for example, in [2, 5, 24]. In [24], the authors propose a measure-based OED formulation that does not choose sensor locations from a finite number of candidate locations but allows sensors to

be placed anywhere on a closed subset of the domain. The articles [2, 5] explore an alternate OED criterion—D-optimality—for infinite-dimensional Bayesian linear inverse problems.

We extend previous approaches by taking into account irreducible uncertainty in the model. This results in a challenging optimization under uncertainty (OUU) problem [9, 22, 30], as we now describe. In linear inverse problems, the average posterior variance of the inversion parameters—the A-optimal criterion—is defined by the trace of the posterior covariance operator $\mathcal{C}_{\text{post}}$. This operator is high-dimensional (upon discretization), dense, and expensive-to-apply. Specifically, applying $\mathcal{C}_{\text{post}}$ to vectors requires many PDE solves. This covariance operator is a function of the vector of experimental design weights and the random variables characterizing the model uncertainty. The resulting OED under uncertainty (OEDUU) problem is challenging because the OEDUU objective is the expected value of $\text{tr} [\mathcal{C}_{\text{post}}]$, with expectation taken with respect to the model uncertainty. It is worth noting that optimizing $\text{tr} [\mathcal{C}_{\text{post}}]$ even for a single instance of $\mathcal{C}_{\text{post}}$ is in itself a challenging problem [3].

As mentioned above, our goal is finding binary design vectors, which are in general difficult to compute. Our approach for this considers a relaxation of the problem and uses an adaptation of the regularized ℓ_0 -sparsification approach outlined in [3]. Note that there are different options to control sparsity of designs; see e.g., [11, 17, 34].

Our approach. We follow a sample average approximation (SAA) approach for the OEDUU problem, where we approximate the expectation in the OEDUU objective by sample averaging over the model uncertainty. Corresponding to each realization of the irreducible model uncertainty, we have a realization of the parameter-to-observable map (forward operator) that defines a specific instance of $\mathcal{C}_{\text{post}}$; see section 4.1. Computing traces of these operators is challenging due to their high-dimensionality (upon discretization) and the high cost (in terms of PDE solves) of applying the covariance operators to vectors.

To mitigate the computational cost of OEDUU, we present a novel formulation of the OED criterion in the observation space. Thus, we only require computing traces of operators defined on the observation space which in many infinite-dimensional inverse problems has a smaller dimension than that of the discretized parameter space; see section 4.2. However, computing the resulting OEDUU objective and its gradient still requires many PDE solves, and these computations are repeated in each step of an optimization algorithm used for solving the OEDUU problem. Hence, it is imperative to exploit problem structure to compute low-rank approximations of the forward operator samples to eliminate frequent PDE solves from the optimization iterations. To do so, we employ randomized matrix methods to compute a low rank basis that *jointly* approximates the range space for all or subsets (clusters) of forward operator samples; see section 5.

We explore the effectiveness and performance of our methods for a realistic groundwater initial condition inversion problem; see sections 6–8. In this application, the (high-dimensional) irreducible uncertainties stem from: an unknown groundwater flow field (as a result of uncertainty in the subsurface permeability field), and from an uncertain observation time.

Contributions. The contributions of this article are as follows: (1) we propose a mathematical formulation for OED under irreducible model uncertainty in infinite-dimensional Bayesian linear inverse problems. (2) We present a novel OED objective formulation in the observation space that avoids trace estimation in high-dimensional discretized parameter spaces. This formulation also applies to A-optimal OED without additional model uncertainty. (3) We develop an efficient and practical reduced order modeling framework for OEDUU that eliminates PDE solves from the optimization process. (4) We present a comprehensive set of numerical experiments that illustrate the proposed approach and demonstrate its effectiveness.

Limitations. The present work also has limitations. (1) Our formulation is restricted to linear parameter-to-observation maps, Gaussian priors, and additive Gaussian noise. Further work is needed to extend it to nonlinear inverse problems. One possible extension is to use a Gaussian approximation of the posterior distribution as in [4]. (2) We use Monte Carlo sampling to approximate the irreducible uncertainty. If the aim is a highly accurate approximation of this uncertainty, a large number of samples might be needed.

2. Background

In this section, we review relevant material required for the formulation of OED problems under uncertainty for infinite-dimensional Bayesian inverse problems. We also summarize previous work which this paper builds on.

2.1. Infinite-dimensional Bayesian linear inverse problems

We begin our discussion by formulating a prototypical Bayesian linear inverse problem, which is the primary focus of the present work. A detailed treatment of general nonlinear Bayesian inverse problems can be found for instance in [13].

Given finite-dimensional observations, $\mathbf{d} \in \mathbb{R}^d$, we seek to infer an unknown parameter, m , which is related to the data through

$$\mathbf{d} = \mathcal{F}m + \eta. \quad (1)$$

We consider the case where m is an element of an infinite-dimensional Hilbert space \mathcal{H} , e.g., $\mathcal{H} = L^2(\mathcal{D})$ with \mathcal{D} being a bounded domain in \mathbb{R}^2 or \mathbb{R}^3 . Here, $\mathcal{F} : \mathcal{H} \rightarrow \mathbb{R}^d$ is a continuous linear parameter-to-observable map (forward model). In our target applications, computing $\mathcal{F}m$ for a given m involves solving a partial differential equation (PDE) followed by application of an observation operator. In (1), we assume $\eta \sim \mathcal{N}(\mathbf{0}, \Gamma_{\text{noise}})$.

Following a Bayesian approach, we model the parameter m as a random variable and impose a prior probability law for m . The prior law is a probabilistic description of our prior knowledge about the parameter. We use a Gaussian prior $\mathcal{N}(m_{\text{pr}}, \mathcal{C}_{\text{pr}})$ where the mean m_{pr} is a sufficiently regular element of \mathcal{H} and the covariance operator \mathcal{C}_{pr} is a trace-class operator defined through the inverse of a differential operator. More precisely, we let $\mathcal{C}_{\text{pr}} = \mathcal{A}^{-2}$ where $\mathcal{A} = -\rho\Delta + \delta\mathcal{I}$; here, Δ is the Laplace operator, $\delta > 0$ controls the magnitude of the variance and $\rho > 0$ controls the correlation length. This choice ensures that the prior covariance operator is trace-class in two and three space dimensions; see [13] for details.

The solution of the Bayesian inverse problem is the posterior probability law for the parameter, which is conditioned on measurement data. For a linear inverse problem with a Gaussian prior and an additive Gaussian noise model, which is what we assume, it is well known [13] that the posterior distribution is also Gaussian, namely $\mathcal{N}(m_{\text{post}}, \mathcal{C}_{\text{post}})$ with

$$\mathcal{C}_{\text{post}} = (\mathcal{F}^* \Gamma_{\text{noise}}^{-1} \mathcal{F} + \mathcal{C}_{\text{pr}}^{-1})^{-1} \quad \text{and} \quad m_{\text{post}} = \mathcal{C}_{\text{post}} (\mathcal{F}^* \Gamma_{\text{noise}}^{-1} \mathbf{d} + \mathcal{C}_{\text{pr}}^{-1} m_{\text{pr}}). \quad (2)$$

2.2. Bayesian linear inverse problem with model uncertainty

We are interested in the design of experiments for inverse problems governed by PDEs with uncertain parameters representing the irreducible uncertainty. Note that these uncertainties are

in addition to the uncertainty in the inversion parameter. In such cases, the forward operator \mathcal{F} is a function of uncertain parameters. We will formalize this below.

Let (Ω, \mathcal{G}, P) be a probability space, where Ω is a sample space, \mathcal{G} is a suitable σ -algebra on Ω , and P is a probability measure. We let $\xi = \xi(\omega)$ denote a random variable, defined on (Ω, \mathcal{G}, P) , that models the irreducible uncertainty in \mathcal{F} ; for example, ξ could be a random vector whose entries define uncertain parameters in the governing PDEs, or ξ could be a function-valued random variable (i.e., a random field coefficient). In this case, the forward operator is a random variable $\mathcal{F} : \Omega \rightarrow \mathcal{L}(\mathcal{H}, \mathbb{R}^d)$, where $\mathcal{L}(\mathcal{H}, \mathbb{R}^d)$ is the space of linear transformations from \mathcal{H} to \mathbb{R}^d . For each $\omega \in \Omega$, we have a realization $\mathcal{F}(\omega) = \mathcal{F}(\xi(\omega))$ of the forward model. Thus, the solution of the corresponding Bayesian linear inverse problem depends on ξ and is given by $\mathcal{N}(m_{\text{post}}(\xi), \mathcal{C}_{\text{post}}(\xi))$, with

$$\mathcal{C}_{\text{post}}(\xi) = (\mathcal{F}(\xi)^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{F}(\xi) + \mathcal{C}_{\text{pr}}^{-1})^{-1} \quad \text{and} \quad m_{\text{post}}(\xi) = \mathcal{C}_{\text{post}} (\mathcal{F}(\xi)^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{d} + \mathcal{C}_{\text{pr}}^{-1} m_{\text{pr}}). \quad (3)$$

2.3. A-optimal design of infinite-dimensional inverse problems

We focus on A-optimal design of Bayesian linear inverse problems. That is, we seek designs (sensor placements) that result in minimized average posterior variance. In the present infinite-dimensional formulation, this amounts to minimizing the trace of $\mathcal{C}_{\text{post}}$; see [3] for details. The additional challenge for inverse problems with uncertainties in the governing model is that the covariance operator itself depends on the random variable ξ that defines the uncertain parameters in the model. In the next section, we formulate the A-optimal design problem as that of optimization under uncertainty. We refer to this problem as optimal experimental design under uncertainty (OEDUU).

3. A-optimal design of experiments for linear Bayesian inverse problems with irreducible model uncertainty

We begin with a discussion of experimental design for infinite-dimensional linear inverse problems with uncertain forward models in section 3.1. In section 3.2, we present our formulation of Bayesian A-optimality for such inverse problems, and in section 3.3, we formulate the optimization problem for finding A-optimal designs in the inverse problems under study.

3.1. Design in the Bayesian inverse problem with uncertain forward models

In OED for inverse problems we are interested in determining how to collect measurement data to optimize the parameter inference. The definition of an ‘experimental design’ is problem specific. For example, in inverse problems in tomography, the design could correspond to choosing a subset of angles for an x-ray source to hit an object. On the other hand, in the inverse problem of identifying the source of a contaminant, the design corresponds to the placement of sensors that are used to measure the contaminant concentration.

We consider a finite generic set of candidate experiments denoted by $x_i \in X$, $i = 1, \dots, d$, where X is a problem specific set of admissible experiments. More concretely, in the tomography example, x_i s are measurement angles and X corresponds to the set of all possible measurement angles; and in the subsurface flow example, x_i s indicate sensor locations and X is the physical domain in which the sensors can be placed. We assign a nonnegative weight w_i to each choice of x_i . Thus, an experimental design is specified by a design vector

$\mathbf{w} = (w_1, w_2, \dots, w_d)$. Ideally, we would like binary design vectors: in the case $w_i = 1$, we will collect the measurement corresponding to x_i , and if $w_i = 0$, the corresponding experiment will not be performed (or the data not be collected). However, an OED problem of finding binary optimal design vectors has combinatorial complexity—an extremely challenging problem. To cope with this, we relax the binary assumption on the weights and allow the weights to take values in the interval $[0, 1]$. We then seek to enforce a binary structure on the computed weights using a suitable penalty method, as discussed further below.

To incorporate a generic design \mathbf{w} into the Bayesian inverse problem, we define a diagonal matrix $\mathbf{W} \in \mathbb{R}^{d \times d}$, which, in a generic OED problem, contains the weights $\mathbf{w} = (w_1, w_2, \dots, w_d)$ on the diagonal. In a sensor placement problems for an inverse problem governed by a time-dependent PDE, as in the example used in section 6, the goal is to select an optimal subset of s candidate sensor locations, which collect measurements at r observation times. When a sensor location is chosen, we assume that it collects measurements for all r observation times. In this setup, the design vector \mathbf{w} has dimension s , and the vector of measurement data has dimension $d = sr$. Thus, $\mathbf{W} \in \mathbb{R}^{sr \times sr}$ is a block diagonal matrix with each $s \times s$ diagonal block being a diagonal matrix with the sensor weights on the diagonal.

We incorporate the design vector \mathbf{w} in the Bayesian inverse problem by considering a *weighted forward operator* $\mathcal{F}(\xi; \mathbf{w}) := \mathbf{W}^{\frac{1}{2}} \mathcal{F}(\xi)$. As before, ξ is the random variable that models uncertainty in the governing PDEs. The posterior law of m now depends on the design \mathbf{w} as follows:

$$\begin{aligned} \mathcal{C}_{\text{post}}(\xi, \mathbf{w}) &= \left(\mathcal{F}^*(\xi) \mathbf{W}^{\frac{1}{2}} \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{W}^{\frac{1}{2}} \mathcal{F}(\xi) + \mathcal{C}_{\text{pr}}^{-1} \right)^{-1} \quad \text{and} \\ m_{\text{post}}(\xi, \mathbf{w}) &= \mathcal{C}_{\text{post}} \left(\mathcal{F}(\xi)^* \mathbf{W}^{\frac{1}{2}} \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{d} + \mathcal{C}_{\text{pr}}^{-1} m_{\text{pr}} \right). \end{aligned} \quad (4)$$

As discussed later, if the noise covariance $\mathbf{\Gamma}_{\text{noise}}^{-1}$ is diagonal, the expression for the posterior covariance operator simplifies and it is not necessary to consider the square root of \mathbf{W} .

3.2. A-optimal design under uncertainty

We focus on A-optimal design of linear inverse problems. Following [3], for a fixed realization of ξ , the A-optimal design is one which minimizes the average posterior variance, i.e., the design which minimizes $\text{tr} [\mathcal{C}_{\text{post}}(\xi, \mathbf{w})]$. We extend the notion of A-optimal designs to Bayesian inverse problems with uncertain model parameters by formulating the OED problem as that of minimizing the expected value of $\text{tr} [\mathcal{C}_{\text{post}}(\xi(\omega), \mathbf{w})]$. Thus, the OED criterion we consider is given by

$$\phi(\mathbf{w}) := \int_{\Omega} \text{tr} [\mathcal{C}_{\text{post}}(\xi(\omega), \mathbf{w})] P(d\omega). \quad (5)$$

Note that taking the perspective of optimization under uncertainty, minimizing (5) amounts to a risk-neutral design. Other scalarization methods are possible, e.g., risk-averse objectives, which place emphasis on avoiding particularly poor designs [22, 31].

3.3. The OED problem

An effective solution method for finding OEDs must provide the user a mechanism to strike a balance between the competing goals of minimizing posterior uncertainty and using as few sensors as possible. We address this by using a sparsifying penalty function to promote sparse (and eventually, binary) optimal design vectors. Accordingly, we formulate the optimization

problem for finding an OED as follows:

$$\min_{\mathbf{w} \in [0,1]^s} \phi(\mathbf{w}) + \gamma \psi(\mathbf{w}), \quad (6)$$

where $\phi(\mathbf{w})$ is the OED criterion defined in (5), $\psi : \mathbb{R}_+^s \rightarrow [0, \infty)$ is a sparsity-inducing penalty function, and γ controls the degree of sparsity.

A simple choice for ψ is the ℓ_1 -norm, $\psi(\mathbf{w}) = \mathbf{1}^T \mathbf{w}$, where $\mathbf{1}$ is the vector of all ones. While using an ℓ_1 -norm penalty leads to sparse designs, it does not yield binary design vectors. To obtain binary designs, using an ℓ_1 -norm penalty can be combined with thresholding to decide where the sensors should be placed, but as numerically studied in [3], the resulting designs are suboptimal.

In the present work, we follow the regularized ℓ_0 -sparsification procedure introduced in [3], which typically results in binary optimal design vectors. This approach involves a continuation procedure in which a sequence of minimization problems with penalty functions that successively approximate the ℓ_0 -‘norm’ are solved. For $\mathbf{w} \in \mathbb{R}^s$, we denote the ℓ_0 -‘norm’ as $\|\mathbf{w}\|_{\ell_0}$ and define it as the number of non-zero entries in \mathbf{w} . The procedure is initialized with $\mathbf{w}_{\varepsilon(0)}$, which is obtained by solving (6) with an ℓ_1 penalty. Then, at step i of the continuation procedure, a new solution $\mathbf{w}_{\varepsilon(i)}$ is obtained by solving

$$\min_{\mathbf{w} \in [0,1]^s} \phi(\mathbf{w}) + \gamma \psi_{\varepsilon(i)}(\mathbf{w}) \quad (7)$$

using the previous $\mathbf{w}_{\varepsilon(i-1)}$ as an initial guess; the penalty function $\psi_{\varepsilon(i)}(\mathbf{w})$ is chosen from a family of continuously differentiable penalty functions that approach the ℓ_0 -‘norm’ as $i \rightarrow \infty$. This procedure uses non-convex functions and thus one cannot guarantee uniqueness of solutions beyond the ℓ_1 -norm initialization. However, numerical experiments in [3] and in section 8 show that the method performs well in practical examples.

We find that using the ℓ_0 sparsification approach, there may be a large disparity between the sparsity (and the values of the objectives) for the ℓ_1 solution, $\mathbf{w}_{\varepsilon(0)}$, and for the binary weight vector \mathbf{w}_* obtained at the end of the continuation procedure. That is, for a fixed γ , we found that $\|\mathbf{w}_{\varepsilon(0)}\|_{\ell_1} \ll \|\mathbf{w}_*\|_{\ell_0}$. This results in the continuation procedure being rather sensitive to the choice of the sequence $\varepsilon(i)$. As a remedy, we use a rescaling to help mitigate this issue, as follows. Note we can scale the ℓ_1 penalty linearly with a scaling factor $\alpha > 0$ through $\|\alpha \mathbf{w}\|_{\ell_1} = \alpha \|\mathbf{w}\|_{\ell_1}$. However, this scaling has no effect on the ℓ_0 norm $\|\alpha \mathbf{w}\|_{\ell_0} = \|\mathbf{w}\|_{\ell_0}$. Since the first step in the continuation strategy in [3] involves solving the minimization problem with an ℓ_1 penalty to obtain $\mathbf{w}_{\varepsilon(0)}$, we can control the sparsity of $\mathbf{w}_{\varepsilon(0)}$ by using an α -scaled ℓ_1 penalty. More specifically, $\alpha \in (0, 1)$ produces a less sparse initial guess for the continuation procedure than $\alpha \geq 1$. Once an initial guess $\mathbf{w}_{\varepsilon(0)}$ is obtained, we use the following family of scaled sparsity-inducing penalty functions:

$$\psi_{\varepsilon}(\alpha \mathbf{w}) := \sum_{i=1}^s f_{\varepsilon}(\alpha w_i), \quad (8)$$

where, for $\varepsilon > 0$,

$$f_{\varepsilon}(\alpha w) = \begin{cases} \alpha w / \varepsilon & \text{if } w < \frac{\varepsilon}{2\alpha}, \\ c_3(\alpha w)^3 + c_2(\alpha w)^2 + c_1(\alpha w) + c_0 & \text{if } \frac{\varepsilon}{2\alpha} \leq w < \frac{2\varepsilon}{\alpha}, \\ 1 & \text{if } w \geq \frac{2\varepsilon}{\alpha}; \end{cases}$$

as in [3], the constants c_3, c_2, c_1 and c_0 are chosen such that $f_\varepsilon(\alpha w)$ is continuously differentiable on $[0, 1]$.

To summarize, we choose a decreasing sequence $\{\varepsilon(i)\}_{i=1}^\infty$ such that $\varepsilon(1) < 1$ and $\varepsilon(i) \rightarrow 0$ as $i \rightarrow \infty$. For a fixed γ , we initialize the procedure with $\mathbf{w}_{\varepsilon(0)}$, the solution to (7) with an α -weighted ℓ_1 penalty, and in each subsequent step i of the procedure we minimize

$$\phi(\mathbf{w}) + \gamma \psi_{\varepsilon(i)}(\alpha \mathbf{w}), \quad (9)$$

using $\mathbf{w}_{\varepsilon(i-1)}$ as the initial iterate, until we converge to a binary solution, \mathbf{w}_* . In practice, the choice of α is problem-specific and we heuristically choose it such that $\|\mathbf{w}_{\varepsilon(0)}\|_{\ell_1} \approx \|\mathbf{w}_*\|_{\ell_1}$, where $\mathbf{w}_{\varepsilon(0)}$ is the non-binary weight vector obtained using the penalty function $\|\alpha \mathbf{w}\|_{\ell_1}$, and \mathbf{w}_* is the binary weight-vector obtained following the continuation procedure with initial guess $\mathbf{w}_{\varepsilon(0)}$. This might require solving the problem with a few choices of α .

4. Finite-dimensional approximation of OED objective and its gradient

In this section, we describe the discretization of the OED problem. This includes the discretization of the operators defining the OED objective and gradient, and the approximation of the expected value in the OED objective in (6) (see section 4.1). We also present novel efficient-to-evaluate expressions for the OED objective and its gradient by taking traces of operators on the measurement space (see section 4.2). The latter provides computational advantages in cases where the measurement dimension is significantly smaller than the discretized parameter dimension.

4.1. The discretized OED problem

Henceforth, we assume that the forward operator has been discretized in space and, if applicable, in time. That is, we consider $\mathbf{F} : (\Omega, \mathcal{G}, P) \rightarrow \mathcal{L}(V_h, \mathbb{R}^d)$ where V_h is a finite-dimensional subspace of \mathcal{H} given by the span of n finite-element nodal basis functions $\{\varphi_i\}_{i=1}^n$. For each realization of ξ , which parameterizes model uncertainty, $\mathbf{F}(\xi)$ is a linear transformation from V_h to the measurement space \mathbb{R}^d . The discretized inversion parameter is given by $m_h = \sum_{i=1}^n m_i \varphi_i$. Thus, instead of inferring the probability law for our random function m , we focus on characterizing the posterior distribution for the vector of coefficients, $\mathbf{m} = (m_1 \ m_2 \ \dots \ m_n)^T$.

The discretized parameter space is \mathbb{R}^n equipped with the so called mass-weighted inner product that approximates the $L^2(\mathcal{D})$ inner product. The latter is the Euclidean inner product weighted by the finite element mass matrix; see [10] for details. Thus, the discretized forward operator is $\mathbf{F}(\xi) : \mathbb{R}^n \rightarrow \mathbb{R}^d$. In this setting, the adjoint $\mathbf{F}^*(\xi)$ of $\mathbf{F}(\xi)$, is given by $\mathbf{F}^*(\xi) = \mathbf{M}^{-1} \mathbf{F}^T(\xi)$ where \mathbf{M} is the mass matrix.

In what follows, we assume the observations are corrupted by uncorrelated Gaussian noise with a constant variance of σ^2 ; that is, $\mathbf{\Gamma}_{\text{noise}} = \sigma^2 \mathbf{I}$. The posterior distribution of the discretized parameter \mathbf{m} is then given by $\mathcal{N}(\mathbf{m}_{\text{post}}(\xi, \mathbf{w}), \mathbf{\Gamma}_{\text{post}}(\xi, \mathbf{w}))$ with

$$\begin{aligned} \mathbf{\Gamma}_{\text{post}}(\xi, \mathbf{w}) &= (\sigma^{-2} \mathbf{F}^*(\xi) \mathbf{W} \mathbf{F}(\xi) + \mathbf{\Gamma}_{\text{pr}}^{-1})^{-1} \quad \text{and} \\ \mathbf{m}_{\text{post}}(\xi, \mathbf{w}) &= \mathbf{\Gamma}_{\text{post}} \left(\sigma^{-2} \mathbf{F}(\xi)^* \mathbf{W}^{\frac{1}{2}} \mathbf{d} + \mathbf{\Gamma}_{\text{pr}}^{-1} \mathbf{m}_{\text{pr}} \right). \end{aligned} \quad (10)$$

Here, $\mathbf{\Gamma}_{\text{pr}}$ and \mathbf{m}_{pr} are the discretizations of \mathcal{C}_{pr} and m_{pr} , respectively.

We follow a sample average approximation (SAA) approach for solving (6). Thus, the expectation in the OED objective is approximated via sample averaging. Possibilities include

Monte Carlo (MC) sampling, quasi-Monte Carlo, or quadrature. In the present work, we rely on MC. Let $\xi_i, i = 1, \dots, N$, be realizations of ξ , and let $\mathbf{F}_i = \mathbf{F}(\xi_i)$. We approximate (5) with

$$\phi(\mathbf{w}) \approx \bar{\phi}_N(\mathbf{w}) := \frac{1}{N} \sum_{i=1}^N \text{tr} \left[(\sigma^{-2} \mathbf{F}_i^* \mathbf{W} \mathbf{F}_i + \mathbf{\Gamma}_{\text{pr}}^{-1})^{-1} \right]. \quad (11)$$

Note that the sample set $\{\xi_i\}_{i=1}^N$ will be fixed in the optimization problem.

Additionally, we note that the map $\mathbf{w} \mapsto \bar{\phi}_N(\mathbf{w})$ is convex. Since $\mathbf{\Gamma}_{\text{post}}(\xi, \mathbf{w})$ is a self-adjoint positive definite operator for any fixed $\xi \in \Omega$ and $\mathbf{w} \in \mathbb{R}^d$, the convexity of $\bar{\phi}_N$ follows from (1) the strict convexity of $\mathbf{X} \mapsto \text{tr} [\mathbf{X}^{-1}]$ on the cone of self-adjoint positive definite operators (see [25]), (2) the mapping $\mathbf{w} \mapsto [\mathbf{\Gamma}_{\text{post}}(\xi, \mathbf{w})]^{-1}$ being affine (for fixed ξ), and (3) the finite sum of convex functions being convex. To argue strict convexity of $\bar{\phi}_N$, we need an additional assumption on the affine map $\mathbf{w} \mapsto [\mathbf{\Gamma}_{\text{post}}(\xi, \mathbf{w})]^{-1}$, that is, we require that for any $\mathbf{w}_1, \mathbf{w}_2 \in \mathbb{R}^d$ ($\mathbf{w}_1 \neq \mathbf{w}_2$), $\mathbf{F}_i^* \mathbf{W}_1 \mathbf{F}_i \neq \mathbf{F}_i^* \mathbf{W}_2 \mathbf{F}_i$, for at least one $i \in \{1, \dots, N\}$. This will not hold, for example, if no information about the parameter \mathbf{m} could be learned from data obtained at one or more sensor locations.

4.2. Efficient computation of OED objective and its gradient

In (11), each term in the summation requires computing the trace of the inverse of an operator whose dimension is determined by the discretized parameter dimension n . For inverse problems governed by PDEs in two and three space dimensions, n is typically very large. Thus, computing these traces directly is infeasible and methods based on low-rank spectral decomposition or randomized trace estimation must be employed [3]. Here, we outline an alternate strategy and present a reformulation of the OED objective that involves computing traces of operators defined on the measurement space. In problems where the measurement dimension d is considerably smaller than the discretized parameter dimension n , this approach provides significant computational savings, in particular in combination with low-rank approximation of the (preconditioned) parameter-to-observable map as in [3], which we generalize in section 5 to accommodate additional model error. Moreover, while n grows upon mesh refinement, d remains fixed due to finite-dimensionality of the observations.

The following result facilitates our proposed reformulation of the A-optimal OED objective. We state the result for a generic forward operator \mathbf{F} .

Proposition 1. *The following relation holds:*

$$(\sigma^{-2} \mathbf{F}^* \mathbf{W} \mathbf{F} + \mathbf{\Gamma}_{\text{pr}}^{-1})^{-1} = \mathbf{\Gamma}_{\text{pr}} - \sigma^{-2} \mathbf{\Gamma}_{\text{pr}} \mathbf{F}^* (\mathbf{I} + \sigma^{-2} \mathbf{W} \mathbf{\Gamma}_{\text{pr}} \mathbf{F}^*)^{-1} \mathbf{W} \mathbf{F} \mathbf{\Gamma}_{\text{pr}}. \quad (12)$$

Proof. The result follows by the Sherman–Morrison–Woodbury identity ([14], 2.1.4), which states that for matrices $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{n \times d}$, $\mathbf{A} + \mathbf{U} \mathbf{V}^T = \mathbf{A}^{-1} - \mathbf{A}^{-1} \mathbf{U} (\mathbf{I} + \mathbf{V}^T \mathbf{A}^{-1} \mathbf{U})^{-1} \mathbf{V}^T \mathbf{A}^{-1}$ provided $(\mathbf{I} + \mathbf{V}^T \mathbf{A}^{-1} \mathbf{U})$ and \mathbf{A} are invertible. Setting $\mathbf{A} := \mathbf{\Gamma}_{\text{pr}}^{-1}$, $\mathbf{U} := \sigma^{-2} \mathbf{F}^*$ and $\mathbf{V}^T := \mathbf{W} \mathbf{F}$, it is thus sufficient to prove the invertibility of the matrix $\mathbf{I} + \mathbf{W} \mathbf{F} \mathbf{\Gamma}_{\text{pr}} \mathbf{F}^*$.

With no loss of generality, we assume $\sigma^2 = 1$ for simplicity. In the present setup, we have $\mathbf{\Gamma}_{\text{pr}}^* = \mathbf{M}^{-1} \mathbf{\Gamma}_{\text{pr}}^T \mathbf{M} = \mathbf{\Gamma}_{\text{pr}}$, where \mathbf{M} is symmetric positive definite. Moreover, we have that $\mathbf{F}^* = \mathbf{M}^{-1} \mathbf{F}^T$ and thus

$$(\mathbf{F} \mathbf{\Gamma}_{\text{pr}} \mathbf{F}^*)^T = (\mathbf{F}^*)^T \mathbf{\Gamma}_{\text{pr}}^T \mathbf{F}^T = \mathbf{F} \mathbf{M}^{-1} \mathbf{\Gamma}_{\text{pr}}^T \mathbf{M} \mathbf{F}^* = \mathbf{F} \mathbf{\Gamma}_{\text{pr}} \mathbf{F}^*.$$

Hence, $\mathbf{F}\Gamma_{\text{pr}}\mathbf{F}^*$ is symmetric; it is also clearly positive semidefinite. Since \mathbf{W} and $\mathbf{F}\Gamma_{\text{pr}}\mathbf{F}^*$ are both symmetric positive semidefinite matrices, their product $\mathbf{W}\mathbf{F}\Gamma_{\text{pr}}\mathbf{F}^*$ has nonnegative eigenvalues. Hence, $\mathbf{I} + \mathbf{W}\mathbf{F}\Gamma_{\text{pr}}\mathbf{F}^*$ is invertible, which ends the proof. \square

Using this result, we have, for all $i \in \{1, \dots, N\}$,

$$\mathbf{\Gamma}_{\text{post}}(\xi_i, \mathbf{w}) = \mathbf{\Gamma}_{\text{pr}} - \sigma^{-2}\mathbf{F}_i\mathbf{F}_i^*\mathbf{S}_i^{-1}(\mathbf{w})\mathbf{W}\mathbf{F}_i\mathbf{\Gamma}_{\text{pr}}, \quad (13)$$

where $\mathbf{S}_i(\mathbf{w}) := (\mathbf{I} + \sigma^{-2}\mathbf{W}\mathbf{F}_i\mathbf{\Gamma}_{\text{pr}}\mathbf{F}_i^*)$. Using properties of the trace,

$$\text{tr} [\mathbf{\Gamma}_{\text{post}}(\xi_i, \mathbf{w})] = \text{tr} [\mathbf{\Gamma}_{\text{pr}}] - \text{tr} [\sigma^{-2}\mathbf{S}_i^{-1}(\mathbf{w})\mathbf{W}\mathbf{F}_i\mathbf{\Gamma}_{\text{pr}}^2\mathbf{F}_i^*]. \quad (14)$$

Denoting

$$\mathbf{K}(\xi_i, \mathbf{w}) = \sigma^{-2}\mathbf{S}_i^{-1}(\mathbf{w})\mathbf{W}\mathbf{F}_i\mathbf{\Gamma}_{\text{pr}}^2\mathbf{F}_i^*,$$

the discretized OED objective $\bar{\phi}_N(\mathbf{w})$ in (11) can thus be rewritten as

$$\bar{\phi}_N(\mathbf{w}) = \text{tr} [\mathbf{\Gamma}_{\text{pr}}] + \phi_N(\mathbf{w}), \quad (15)$$

where

$$\phi_N(\mathbf{w}) := -\frac{1}{N} \sum_{i=1}^N \text{tr} [\mathbf{K}(\xi_i, \mathbf{w})] = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^d \langle \mathbf{e}_j, \mathbf{K}(\xi_i, \mathbf{w}) \mathbf{e}_j \rangle. \quad (16)$$

Here $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product.

Since $\mathbf{\Gamma}_{\text{pr}}$ is independent of \mathbf{w} , we can neglect that term in the optimization and focus on minimizing $\phi_N(\mathbf{w})$. The optimization problem for finding an OED is then,

$$\min_{\mathbf{w} \in [0,1]^s} \phi_N(\mathbf{w}) + \gamma\psi(\mathbf{w}). \quad (17)$$

Note that since $\bar{\phi}_N$ (and hence ϕ_N) is convex, the above objective is convex as long as the penalty function is convex.

We next derive the gradient of $\phi_N(\mathbf{w})$. We begin by defining some notations to facilitate this derivation. Let us denote $\mathbf{W}^\sigma := \sigma^{-2}\mathbf{W}$. Note that

$$\mathbf{W}^\sigma = \sum_{l=1}^s w_l \mathbf{E}_l \quad \text{with} \quad \mathbf{E}_l = \sigma^{-2} \mathbf{I}_r \otimes \mathbf{e}_l \mathbf{e}_l^T,$$

where \mathbf{e}_l the l th coordinate vector in \mathbb{R}^s and \mathbf{I}_r the $r \times r$ identity matrix. Next, we consider the partial derivatives of $\mathbf{S}_i(\mathbf{w})$ which appears in the definition of $\mathbf{K}(\xi_i, \mathbf{w})$ in (16). Note that with the notation we just introduced

$$\mathbf{K}(\xi_i, \mathbf{w}) = \mathbf{S}_i^{-1}(\mathbf{w})\mathbf{W}^\sigma\mathbf{F}_i\mathbf{\Gamma}_{\text{pr}}^2\mathbf{F}_i^*, \quad \text{with} \quad \mathbf{S}_i(\mathbf{w}) = (\mathbf{I} + \mathbf{W}^\sigma\mathbf{F}_i\mathbf{\Gamma}_{\text{pr}}\mathbf{F}_i^*), \quad i = 1, \dots, N.$$

It is straightforward to see that

$$\frac{\partial \mathbf{S}_i(\mathbf{w})}{\partial w_k} = \mathbf{E}_k \mathbf{F}_i \mathbf{\Gamma}_{\text{pr}} \mathbf{F}_i^*, \quad k = 1, \dots, s, i = 1, \dots, N. \quad (18)$$

Then, for $k = 1, \dots, s$,

$$\begin{aligned}
\frac{\partial \phi_N(\mathbf{w})}{\partial w_k} &= -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^d \langle \mathbf{e}_j, \left(\mathbf{S}_i^{-1}(\mathbf{w}) \frac{\partial \mathbf{W}^\sigma}{\partial w_k} - \mathbf{S}_i^{-1}(\mathbf{w}) \frac{\partial \mathbf{S}_i(\mathbf{w})}{\partial w_k} \mathbf{S}_i^{-1}(\mathbf{w}) \mathbf{W}^\sigma \right) \mathbf{F}_i \mathbf{\Gamma}_{\text{pr}}^2 \mathbf{F}_i^* \mathbf{e}_j \rangle \\
&= -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^d \langle \mathbf{e}_j, \mathbf{S}_i^{-1}(\mathbf{w}) \mathbf{E}_k (\mathbf{I} - \mathbf{F}_i \mathbf{\Gamma}_{\text{pr}} \mathbf{F}_i^* \mathbf{S}_i^{-1}(\mathbf{w}) \mathbf{W}^\sigma) \mathbf{F}_i \mathbf{\Gamma}_{\text{pr}}^2 \mathbf{F}_i^* \mathbf{e}_j \rangle \\
&= -\frac{1}{N} \sum_{i=1}^N \text{tr} [\mathbf{S}_i^{-1}(\mathbf{w}) \mathbf{E}_k \mathbf{A}_i(\mathbf{w}) \mathbf{B}_i], \tag{19}
\end{aligned}$$

where $\mathbf{A}_i(\mathbf{w}) := \mathbf{I} - \mathbf{F}_i \mathbf{\Gamma}_{\text{pr}} \mathbf{F}_i^* \mathbf{S}_i^{-1}(\mathbf{w}) \mathbf{W}^\sigma$ and $\mathbf{B}_i := \mathbf{F}_i \mathbf{\Gamma}_{\text{pr}}^2 \mathbf{F}_i^*$, $i = 1, \dots, N$.

5. Model reduction

Evaluating the objective function in (17) and its partial derivatives in (19) requires many discretized PDE solves. Specifically, computation of the trace in the OED objective function (17) requires d applications of $\mathbf{K}(\xi_i, \mathbf{w})$ for each $i = 1, \dots, N$; see (16). Each application of $\mathbf{K}(\xi_i, \mathbf{w})$ involves a forward and adjoint PDE solve as well as the inverse of the operator $\mathbf{S}_i(\mathbf{w})$, and each application of $\mathbf{S}_i(\mathbf{w})$ requires two PDE solves. Additionally, evaluating the partial derivatives $\frac{\partial \phi_N(\mathbf{w})}{\partial w_k}$ for each $k = 1, \dots, s$ requires d applications of $\mathbf{S}_i^{-1}(\mathbf{w}) \mathbf{E}_k \mathbf{A}_i(\mathbf{w}) \mathbf{B}_i$ for $i = 1, \dots, N$, each of which require four PDE solves and two applications of $\mathbf{S}_i^{-1}(\mathbf{w})$. Since computing the objective and its gradient is required in each step of an optimization procedure, the resulting large number of PDE solves can become computationally infeasible.

In this section, we propose a method for replacing PDE solves in the OED objective and gradient computation with a reduced model to make the optimization computationally tractable. The upfront computation of this reduced order model (ROM) for \mathbf{F}_i for $i = 1, \dots, N$ requires upfront forward and adjoint PDE solves. Once the reduced order models are computed, we can solve the optimization problem for all choices of γ without requiring additional PDE solves. In section 5.1, we exploit the problem structure to find low-dimensional subspaces of the observation and parameter spaces that capture the effective action for each forward operator sample \mathbf{F}_i , $i = 1, \dots, N$ (preconditioned by the prior). As discussed in section 5.2, this can be made more efficient by clustering the samples of uncertain parameters $\{\xi_i\}_{i=1}^N$ such that the corresponding forward operator samples (\mathbf{F}_i 's) in each cluster share similar features. Then, low-dimensional bases are computed for each cluster.

5.1. Composite low-rank basis

The article [3], which concerns OED with no model uncertainty, proposes computing a low-rank approximation to the prior-preconditioned forward map $\tilde{\mathbf{F}}_i := \mathbf{F}_i \mathbf{\Gamma}_{\text{pr}}^{\frac{1}{2}}$ in terms of a low-rank singular value decomposition (SVD); in that article only one copy of $\tilde{\mathbf{F}}_i$ is considered, due to lack of model uncertainty. The prior-preconditioned forward operator is commonly low-rank due to properties of the inverse problem, the limited number of observations and the smoothing properties of the prior. Following this procedure directly would require computing and storing the left and right singular vectors for each $\tilde{\mathbf{F}}_i$, $i = 1, \dots, N$. While each individual map may require a small number of vectors to approximate its effective domain and range, computing and storing such a low-rank approximation for every $\tilde{\mathbf{F}}_i$ individually could become infeasible. Additionally, there may be some overlap in the singular vectors required to approximate

Algorithm 1. Composite randomized range finder algorithm.

```

1: procedure CRF( $[\tilde{\mathbf{F}}_1, \dots, \tilde{\mathbf{F}}_N], [\tilde{\mathbf{F}}_1^*, \dots, \tilde{\mathbf{F}}_N^*]$ ,  $\mu$ )
2:   Given  $N$  Gaussian random matrices  $\Omega_i \in \mathbb{R}^{d \times r}$  and  $N$  Gaussian random matrices
    $\widehat{\Omega}_i \in \mathbb{R}^{n \times r}$ ,  $i = 1, \dots, N$ 
3:   Compute  $\mathbf{Y}_i = \tilde{\mathbf{F}}_i \Omega_i$  and  $\widehat{\mathbf{Y}}_i = \tilde{\mathbf{F}}_i^* \widehat{\Omega}_i$ 
4:   Set  $\mathbf{Y} = [\mathbf{Y}_1, \dots, \mathbf{Y}_N]$  and  $\widehat{\mathbf{Y}} = [\widehat{\mathbf{Y}}_1, \dots, \widehat{\mathbf{Y}}_N]$ 
5:   Compute SVDs of  $\mathbf{Y} = \mathbf{U} \Sigma \mathbf{V}^T$  and of  $\widehat{\mathbf{Y}} = \widehat{\mathbf{U}} \widehat{\Sigma} \widehat{\mathbf{V}}^T$  with singular values  $\sigma_j$ 
   and  $\widehat{\sigma}_l$  ( $j = 1, \dots, d$ ,  $l = 1, \dots, n$ ) respectively
6:   Set  $k = \max \left\{ \max \left( j \text{ s.t. } \frac{\sigma_j}{\sigma_1} \leq \mu \right), \max \left( l \text{ s.t. } \frac{\widehat{\sigma}_l}{\widehat{\sigma}_1} \leq \mu \right) \right\}$ 
7:   Set  $\mathbf{Q}$  to be the first  $k$  columns of  $\mathbf{U}$  and  $\widehat{\mathbf{Q}}$  to be the first  $k$  columns of  $\widehat{\mathbf{U}}$ 
8: return  $\mathbf{Q}$  and  $\widehat{\mathbf{Q}}$ 

```

each forward operator. Thus, we propose a method for finding spaces that capture the effective composite action of $\tilde{\mathbf{F}}_i$, $i = 1, \dots, N$.

Accordingly, we seek to find two matrices, $\mathbf{Q} \in \mathbb{R}^{d \times k}$ and $\widehat{\mathbf{Q}} \in \mathbb{R}^{n \times k}$, with $k \in \mathbb{N}$ as small as possible, such that

$$\begin{aligned} [\tilde{\mathbf{F}}_1, \dots, \tilde{\mathbf{F}}_N] &\approx \mathbf{Q} \mathbf{Q}^T [\tilde{\mathbf{F}}_1, \dots, \tilde{\mathbf{F}}_N], \\ [\tilde{\mathbf{F}}_1^*, \dots, \tilde{\mathbf{F}}_N^*] &\approx \widehat{\mathbf{Q}} \widehat{\mathbf{Q}}^T [\tilde{\mathbf{F}}_1^*, \dots, \tilde{\mathbf{F}}_N^*]. \end{aligned} \quad (20)$$

While the composite spaces spanned by \mathbf{Q} and $\widehat{\mathbf{Q}}$ do not give us precise bounds for $\|\tilde{\mathbf{F}}_i - \mathbf{Q} \mathbf{Q}^T \tilde{\mathbf{F}}_i \widehat{\mathbf{Q}} \widehat{\mathbf{Q}}^T\|$ for each individual $\tilde{\mathbf{F}}_i$ ($i = 1, \dots, N$), in practice we find that

$$\tilde{\mathbf{F}}_i \approx \mathbf{Q} \mathbf{Q}^T \tilde{\mathbf{F}}_i \widehat{\mathbf{Q}} \widehat{\mathbf{Q}}^T, \quad i = 1, \dots, N. \quad (21)$$

We find \mathbf{Q} and $\widehat{\mathbf{Q}}$ via the randomized range finder algorithm (RRF) ([18], algorithm 4.1). The idea is to simultaneously compute a basis for the subspaces of \mathbb{R}^d and \mathbb{R}^n that capture the action of each $\tilde{\mathbf{F}}_i$ and $\tilde{\mathbf{F}}_i^*$ ($i = 1, \dots, N$) respectively. To do so, we choose a $d \times r$ Gaussian random matrix⁴ Ω_i and an $n \times r$ Gaussian random matrix $\widehat{\Omega}_i$ and compute $\mathbf{Y}_i = \tilde{\mathbf{F}}_i \Omega_i$ and $\widehat{\mathbf{Y}}_i = \tilde{\mathbf{F}}_i^* \widehat{\Omega}_i$ for each i . The SVDs of $[\mathbf{Y}_1, \dots, \mathbf{Y}_N]$ and $[\widehat{\mathbf{Y}}_1, \dots, \widehat{\mathbf{Y}}_N]$ are computed and truncated up to a specified tolerance to obtain \mathbf{Q} and $\widehat{\mathbf{Q}}$ respectively. The algorithm is summarized in algorithm 1.

To fully eliminate the PDEs from the optimization problem, the small inner matrices $\mathbf{Q}^T \tilde{\mathbf{F}}_i \widehat{\mathbf{Q}}$ in (21) must be computed and stored for each $i = 1, \dots, N$, which requires solving k more PDEs for each forward operator sample. If desired, these additional PDE solves can be avoided at the cost of additional error by modifying the single-pass approach presented in [18]. Specifically, a matrix \mathbf{B}_i approximating $\mathbf{Q}^T \tilde{\mathbf{F}}_i \widehat{\mathbf{Q}}$ for each $i = 1, \dots, N$ can be found using a minimal residual method to approximately satisfy the relations

$$\mathbf{B}_i \widehat{\mathbf{Q}}^T \Omega_i = \mathbf{Q}^T \mathbf{Y}_i \quad \text{and} \quad \mathbf{B}_i^T \mathbf{Q}^T \widehat{\Omega}_i = \widehat{\mathbf{Q}}^T \widehat{\mathbf{Y}}_i. \quad (22)$$

⁴ By Gaussian random matrix, $\Omega \in \mathbb{R}^{d \times k}$, we mean that entries of Ω are independent standard normal random variables.

Algorithm 2. Composite randomized range finder algorithm with clustering.

```

1: procedure CRRFWC( $\tilde{\mathbf{F}}_1, \dots, \tilde{\mathbf{F}}_N, [\tilde{\mathbf{F}}_1^*, \dots, \tilde{\mathbf{F}}_N^*], \mu, l, \mathbf{m}$ )
2:    $[c_1, \dots, c_N] = \text{CLUSTER}([\tilde{\mathbf{F}}_1\mathbf{m}, \dots, \tilde{\mathbf{F}}_N\mathbf{m}], l, \text{tol})$ 
3:   Initialize  $\mathbf{Q} = []$  and  $\hat{\mathbf{Q}} = []$ 
4:   for  $i = 1, \dots, l$  do
5:     Initialize  $\tilde{\mathbf{F}} = []$  and  $\tilde{\mathbf{F}}^* = []$ ,
6:     for  $j = 1, \dots, N$  do
7:       if  $c_j == i$  then append  $\tilde{\mathbf{F}}_j$  to  $\tilde{\mathbf{F}}$  and append  $\tilde{\mathbf{F}}_j^*$  to  $\tilde{\mathbf{F}}^*$ 
8:       Compute  $[\mathbf{Q}_i, \hat{\mathbf{Q}}_i] = \text{CRF}[\tilde{\mathbf{F}}, \tilde{\mathbf{F}}^*, \mu]$ 
9:       Append to  $\mathbf{Q}$ ,  $\mathbf{Q} = [\mathbf{Q}, \mathbf{Q}_i]$  and to  $\hat{\mathbf{Q}}$ ,  $\hat{\mathbf{Q}} = [\hat{\mathbf{Q}}, \hat{\mathbf{Q}}_i]$ 
10:    for  $i = 1, \dots, N$  do
11:       $\mathbf{Q}_{c_i} \mathbf{F} \mathbf{Q}_{c_i} = \mathbf{Q}_{c_i} \tilde{\mathbf{F}}_i \hat{\mathbf{Q}}_{c_i}$ 
12:    return  $[\mathbf{Q}_1, \dots, \mathbf{Q}_l], [\hat{\mathbf{Q}}_1, \dots, \hat{\mathbf{Q}}_l]$  and  $[\mathbf{Q}_{c_1} \mathbf{F} \mathbf{Q}_{c_1}, \dots, \mathbf{Q}_{c_N} \mathbf{F} \mathbf{Q}_{c_N}]$ 

```

5.2. Clustering

To reduce the amount of PDE solves needed to compute the inner matrices and the amount of basis vectors stored, we follow the ideas presented in [26] and break up the sample space Ω into l clusters. To do this, we use a standard k -means clustering algorithm where we define the distance measure between two samples ξ_i and ξ_j as the Euclidean distance between the observations they produce for an instance of the inversion parameter \mathbf{m} . Specifically, for an $\mathbf{m} \in \mathbb{R}^n$, we define the distance between two samples as

$$d(\xi_i, \xi_j; \mathbf{m}) = \|\mathbf{F}_i \mathbf{m} - \mathbf{F}_j \mathbf{m}\|_2. \quad (23)$$

Naturally, this distance depends on the choice of \mathbf{m} . One could pick \mathbf{m} as a random draw from the prior distribution or one could pick a suitable parameter \mathbf{m} , which is an interesting design problem itself. For the model problem used in this paper, we choose the parameter as a sum of radial basis functions with different centers and magnitudes.

We compute matrices \mathbf{Q}_p and $\hat{\mathbf{Q}}_p$ of rank k_p using algorithm 1 for each cluster $p = 1, \dots, l$ to approximate the effective action of all the forward operators in the cluster. This preliminary clustering step allows us to reduce the amount of PDE solves needed when computing the inner matrices in algorithm 1. In addition to k_p being smaller than k due to the clusters containing less samples, sample parameters ξ_i and ξ_j corresponding to the same cluster p produce more similar data than parameters in different clusters. This means that forward operators corresponding to the same cluster have overlapping range spaces, thus requiring less basis vectors to cover them.

The updated composite low-rank basis algorithm given is provided in algorithm 2, where the CLUSTER procedure refers to a standard k -means clustering procedure which takes as inputs the synthetic observations obtained for an instance of the inversion parameter $([\tilde{\mathbf{F}}_1\mathbf{m}, \dots, \tilde{\mathbf{F}}_N\mathbf{m}])$, the number of desired clusters l , and a tolerance tol , and outputs the clusters $[c_1, \dots, c_N]$ for each forward operator $\tilde{\mathbf{F}}_i$ ($i = 1, \dots, N$).

6. Subsurface flow example: the inverse problem

This section is devoted to the description of the example inverse problem used to illustrate our methods. The subsurface flow forward and inverse problems are described in this section. The sources of irreducible uncertainty, which enter in this problem, are discussed in section 7, and optimal designs taking into account the irreducible uncertainty are the topic of section 8. The

inverse problem we consider seeks to estimate an uncertain initial contaminant concentration field using sensor measurements of contaminant concentration recorded at a discrete set of observation times.

6.1. The forward problem

We consider the transport of a contaminant in a rectangular domain $D := [0, a] \times [0, b] \subset \mathbb{R}^2$.⁵ The evolution of the contaminant's concentration, $u(\mathbf{x}, t)$, in groundwater flow is modeled by a time-dependent advection-diffusion equation

$$\begin{aligned} u_t - \kappa \Delta u + \mathbf{v} \cdot \nabla u &= 0 && \text{in } D \times (T_0, T_1), \\ u(\cdot, T_0) &= m && \text{in } D, \\ (-\kappa \nabla u + u \mathbf{v}) \cdot \mathbf{n} &= 0 && \text{in } \Gamma_l \times (T_0, T_1), \\ \kappa \nabla u \cdot \mathbf{n} &= 0 && \text{in } \partial D \setminus \Gamma_l \times (T_0, T_1), \end{aligned} \tag{24}$$

where $\kappa > 0$ is a known diffusion coefficient, \mathbf{v} is the advection velocity field, $T_0 < T_1$ are the initial and the final time, respectively, and $m \in L_2(D)$ is the initial concentration field. Here, $\Gamma_l := \{0\} \times [0, b]$ is the left boundary of the domain. We assume Γ_l is impermeable, as modeled by the zero total flux condition at that boundary. We want the contaminant to be able to leave the domain, so we allow it to advect freely through the remaining portion of the boundary, $\partial D \setminus \Gamma_l$; this is modeled by imposing a homogeneous pure Neumann condition. For a more detailed explanation and a derivation of a model for two-dimensional flow in an aquifer, see e.g., ([8], section 5.3).

A major source of uncertainty in the governing equation (24) is the velocity field \mathbf{v} , which is an irreducible model uncertainty considered herein. Moreover, the time interval $[T_0, T_1]$ might be uncertain as we can only estimate how long ago a contaminant has been released. Our model for these irreducible uncertainties is detailed in section 7. Before that, we detail the Bayesian inverse problem for fixed advection velocity \mathbf{v} and time interval $[T_0, T_1]$.

6.2. Bayesian inversion for initial state

For the Bayesian inversion of the initial concentration m given a fixed velocity field \mathbf{v}_i and a fixed time interval $[T_0, T_1]$, we impose a Laplacian-like prior. We define the prior operator on the domain D as $\mathcal{C}_{\text{pr}} := \mathcal{A}^{-2} := (-\rho \Delta + \delta I)^{-2}$. To reduce the variance near the boundary of the domain resulting from combining the differential operator \mathcal{A} with Neumann boundary conditions, we instead impose Robin boundary conditions [12, 28]. Accordingly, given $s \in L^2(D)$, the weak solution m of $\mathcal{A}m = s$ satisfies

$$\rho \int_D \nabla m \cdot \nabla p \, d\mathbf{x} + \delta \int_D m p \, d\mathbf{x} + \beta \int_{\partial D} m p \, ds = \int_D s p \, d\mathbf{x}, \quad \text{for all } p \in H^1(D), \tag{25}$$

where $\beta = (\rho/1.42)\sqrt{\delta/\rho}$ as proposed in [28]. The parameters ρ and δ control the correlation length and variance of the covariance operator, and were set to $\rho = 0.008$ and $\delta = 0.02$ for our simulations as these parameters lead to prior samples which have realistic smoothness and correlation length.

⁵This two-dimensional setting can be understood as a top-down view of the evolution of an initial concentration in a horizontal slice of an aquifer or a slice resulting from averaging the properties of a thin 3-dimensional domain.

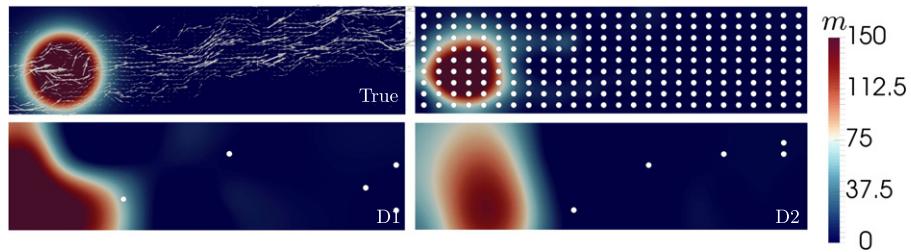


Figure 1. Discrete MAP points, or \mathbf{m}_{post} , at the ‘true’ velocity field and time interval $[-1, 16]$ obtained for different sensor locations. The top left figure shows the ‘true’ initial concentration, \mathbf{m} , we wish to characterize in (24), overlaid with the velocity field \mathbf{v}^* . The top right figure shows the MAP point obtained using all 234 possible sensor and thus depicts the best we can do for our particular problem formulation. On the bottom, we show MAP points obtained with two different designs (D1 and D2) consisting of 5 sensors.



Figure 2. Shown is the pointwise variance for the prior (middle) and posterior using the designs D1 and D2 shown in the bottom row of figure 1.

We choose $s = 234$ candidate locations (these locations are shown in the top right graphic in figure 1) where we can place sensors. Setting $T_1 = 16$, we assume that at each candidate location, we can take $r = 5$ concentration measurements at equally spaced observation times $\tau_i = 7, 9, 11, 13, 15$. Measurements are time-averaged concentrations over intervals $[\tau_i - 0.5, \tau_i + 0.5]$ for each $i = 1, \dots, 5$. This leads to a vector of possible (i.e., if sensors were to be placed on all possible locations) observations $\mathbf{d} \in \mathbb{R}^{sr}$.

As discussed in section 3.1, for a fixed velocity field \mathbf{v} , time interval $[T_0, T_1]$, and design \mathbf{w} , under the assumption of Gaussian prior and additive Gaussian noise, the posterior distribution is also Gaussian. In particular, it is fully characterized by its mean \mathbf{m}_{post} , and posterior covariance matrix, $\mathcal{C}_{\text{post}}$. Different sensor placements lead to different MAP points and different updates to the prior covariance matrix, as can be seen in figures 1 and 2. The advection velocity field used for these results is also shown in figure 1. We illustrate the effect of different design choices on the posterior pointwise variance in figure 2 setting $T_0 = -1$ and $T_1 = 16$. In figure 1 we show the MAP points (\mathbf{m}_{post}) obtained for these same choices of design.

6.3. Discretization and implementation

To discretize the forward and adjoint operators as well as the prior operator, we use the hIPPYlib [33] framework. We use implicit Euler for time-stepping using 250 timesteps in the interval $(T_0, T_1]$. A streamline upwind Petrov Galerkin (SUPG) discretization in space is used to stabilize the discretization of the advective term. The spatial discretization uses 3750 triangles resulting in 1976 degrees of freedom. To accelerate the linear solves in the implicit Euler

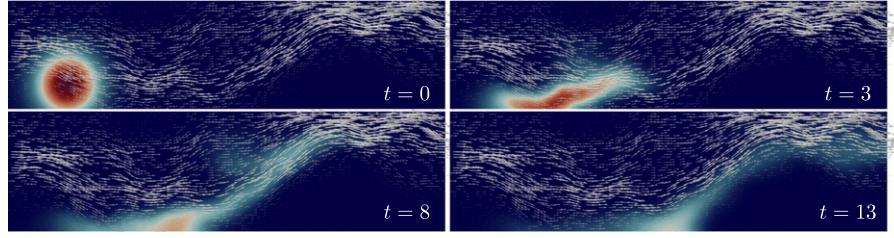


Figure 3. Snapshots of the forward propagation of an initial concentration ($t = 0$). The white arrows depict the velocity field used for the simulation.

steps, we build sparse LU factorizations of the spatial forward and adjoint operators and reuse them throughout the implicit Euler iterations. The evolution for one choice of velocity field and initial condition is shown in figure 3. All other solver components detailed in sections 7 and 8 also use hIPPYlib for the PDE discretization and Python for the numerical linear algebra.

7. Subsurface flow example: irreducible uncertainties

In this section, we characterize the irreducible uncertainties we must take into account when computing optimal sensor locations for the inverse problem presented in section 6.2. The sources of irreducible uncertainty are the advection velocity \mathbf{v} and the initial time T_0 in (24). As explained further below, the uncertainty in the velocity field (26) stems from the log-permeability field $\theta(\mathbf{x})$ of the aquifer being uncertain. Therefore, the random variable ξ , introduced in section 2.2, that parameterizes model uncertainty is given by $\xi = (\theta, T_0)$.

7.1. Characterizing the irreducible uncertainty in the subsurface flow velocity

We will use a steady state inverse problem, different from the inverse problem discussed in section 6.2, which is the main target of this work, to characterize a distribution of the velocity field \mathbf{v} and obtain samples from this distribution. Darcy's law describes the flow of a fluid through a medium in terms of the physical properties of the medium and the pressure gradient. Using Darcy's law, the background velocity field \mathbf{v} in (24) is described by

$$\mathbf{v}(\mathbf{x}) = -e^{\theta(\mathbf{x})} \nabla p(\mathbf{x}), \quad (26)$$

where $\theta(\mathbf{x})$ is the log-permeability field of the aquifer and $p(\mathbf{x})$ denotes the pressure of the groundwater transporting the contaminant through the medium.

The equation governing the pressure field is obtained using Darcy's law along with mass conservation and assuming incompressibility (of the fluid carrying the contaminant). This results in a linear elliptic PDE that can be written in the following dual-mixed form:

$$\begin{aligned} -\mathbf{v} - e^{\theta(\mathbf{x})} \nabla p &= 0 && \text{in } D, \\ \nabla \cdot \mathbf{v} &= 0 && \text{in } D, \\ p &= p_0 && \text{on } \Gamma_l \cup \Gamma_r, \\ -\mathbf{v} \cdot \mathbf{n} &= 0 && \text{on } \partial D \setminus (\Gamma_l \cup \Gamma_r). \end{aligned} \quad (27)$$

Here, $\Gamma_l := \{0\} \times [0, b]$ and $\Gamma_r := \{a\} \times [0, b]$ denote the left and right domain boundaries, respectively. The Dirichlet boundary conditions are prescribed as $p_0 \equiv 0$ on Γ_l and $p_0 \equiv 1$

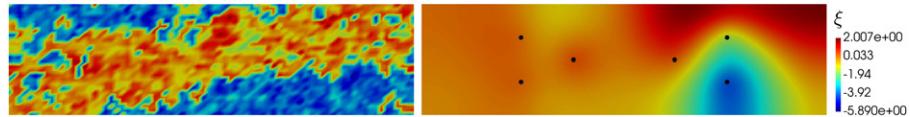


Figure 4. Left: the ‘true’ log permeability field used to generate synthetic pressure data to estimate the log-permeability field. This ‘true’ log-permeability field is a scaled version of the 71st slice of the log-permeability field data from the SPE10 model [1]. Specifically, letting $l(x)$ be the SPE log-permeability data, we use $\theta_{\text{true}}(x) = 0.5 [l(x) - \min(l(x))] + \min(l(x))$. Right: the maximum *a posteriori* estimate of the permeability field obtained after solving the Bayesian inverse problem governed by (27) using six pressure observations. The observation locations are shown as black dots.

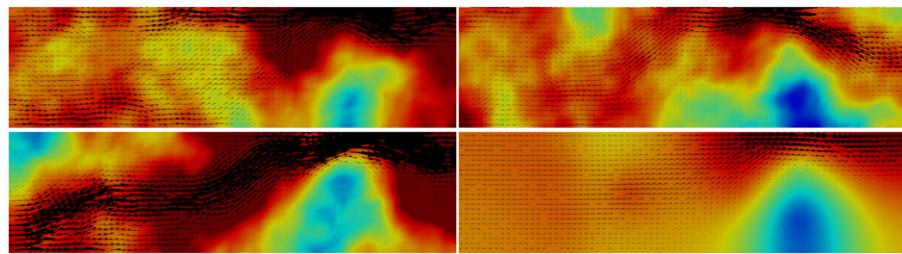


Figure 5. Three log-permeability field samples θ_i and their corresponding velocity fields $v(\theta_i)$ (shown as black arrows) obtained from the posterior distribution π_θ and the log-permeability field and velocity field corresponding to the MAP point (lower right).

on Γ_r . We mention that using the mixed form (27) when deriving the weak formulation and finite element discretization ensures mass conservation over the elements in the numerical solution. Also, porosity and fluid viscosity are omitted from (26) and (27) because they are assumed constant and thus can be absorbed in the remaining terms in the equations through scaling.

As mentioned earlier, the uncertainty in the velocity field (26) is due to uncertainty in the log-permeability field $\theta(x)$. One way to obtain a statistical distribution for the uncertain permeability field is to solve a Bayesian inverse problem governed by the forward problem described in (27). This is described next.

As synthetic data for estimating θ , we use six noisy pressure observations obtained by solving (27) using a scaled version of the 71st slice of the SPE10 permeability model (shown in figure 4). We use a bi-Laplacian prior for θ (see [10]) and assume a Gaussian noise model. The Bayesian inverse problem of estimating θ using pressure measurements is a nonlinear inverse problem that in general requires a sampling algorithm. We compute an approximate solution to that Bayesian inverse problem using a Laplace approximation to the posterior, which is a Gaussian approximation of the posterior centered at the maximum *a posteriori* probability (MAP) point; see e.g., [10].

To generate samples of the uncertain velocity field, we proceed as follows: we generate log-permeability field realizations by drawing samples from the approximate posterior for θ ; subsequently, we solve the pressure equation to obtain the corresponding pressure fields, which are then used to compute velocity field samples, $v(\theta_i)$, using Darcy’s Law (26). For illustration, four different Darcy velocity field samples, together with the permeability fields, are shown in figure 5.

Table 1. Comparison of the number of basis vectors needed to compute \mathbf{Q} and $\widehat{\mathbf{Q}}$ for different numbers N of velocity samples (columns) and tolerances $\mu = 0.002, 0.0001$ (rows). Here, we use algorithm 2 with one cluster.

$\mu \setminus N$	50	100	300	500
0.002	350	410	469	494
0.0001	881	985	1051	1082

7.2. Uncertainty in the initial time

The second source of irreducible uncertainty we consider is the initial time T_0 . Receiving measurements of the concentration, one usually does not know exactly *when* the contaminant has been introduced and thus we might only have an estimate for T_0 . This uncertainty of the initial time should be taken into account when computing optimal designs, i.e., optimal designs should be tailored to a range of initial times T_0 . This uncertainty in the initial time is an additional irreducible uncertainty we take into account, and we target designs that are optimized for $T_0 \in [-1, 1]$ and assume $T_0 \sim \mathcal{U}(-1, 1)$.

8. Subsurface flow example: optimal experimental design under uncertainty

In this section we present numerical results for optimal sensor placement under uncertainty and compare results obtained taking into account the irreducible model uncertainty with results that are computed for a fixed realization of the irreducible uncertainty.

8.1. Setup of the OEDUU problem

As explained above, we solve the optimal experimental design under uncertainty (OEDUU) problem using sample average approximation. In our numerical tests, we use 100 Monte Carlo samples of the irreducible model uncertainty $\xi = (\theta, T_0)$, i.e., $N = 100$ in (17). In table 1, we show how many reduced model basis vectors are needed for different tolerances. The moderate growth in the number of needed ROM basis vectors indicates that 100 Monte Carlo samples capture a reasonable amount of the overall uncertainty.

As discussed in section 5, minimization of the OEDUU objective (17) is made tractable by elimination of the PDE solves, throughout the optimization iterations, via computation of a composite reduced order basis using algorithm 2. In table 2, we study the accuracy and dimension of the joint basis for different numbers of clusters and choices of μ , which control the accuracy of the reduced models.

Our clustering algorithm sorts the uncertain model parameters into bins based on a distance measure (23) that requires choosing a suitable initial concentration \mathbf{m} . Since contaminants in groundwater typically originate from a few localized sources, we choose the initial concentration \mathbf{m} , in the definition of the distance measure (23), to be the sum of three radial basis functions with varying centers, spreads and magnitudes. From table 2, we note that for 100 samples of the irreducible uncertainty ξ , clustering is not needed as the resulting dimension of \mathbf{Q} and $\widehat{\mathbf{Q}}$ in (21) is rather small. For smaller choices of μ , such as $\mu = 10^{-4}$, or larger numbers of samples, clustering becomes more important to reduce memory usage and compute time.

We find that the joint basis obtained using 1 cluster and $\mu = 0.002$ is sufficient for our purposes, both for approximation of the OEDUU objective ϕ_N in (17) and approximation of $\widehat{\mathbf{F}}_i$ for each forward operator sample. Thus, this choice of parameters is used henceforth.

Table 2. Comparison of ROMs obtained using algorithm 2 with $\mu = 0.002, 0.0001$ and different numbers of clusters. Using $N = 100$ sample velocity fields and initial times, $\phi_N(\mathbf{w})$ in objective (17) is computed for $\mathbf{w} = [1, \dots, 1]$, i.e., the design which includes all sensor locations. The relative error in column four is computed as $\frac{|\phi_N(\mathbf{w}) - \phi_N^t(\mathbf{w})|}{|\phi_N^t(\mathbf{w})|}$, where the true value $\phi_N^t(\mathbf{w})$ was approximated using a reduced basis computed using algorithm 2 with $\mu = 10^{-6}$. The third column shows how many basis vectors k were needed in the algorithm to obtain the desired accuracy as described by line 6 in algorithm 2.

	μ	Number basis vectors	Error
1 Cluster	0.002	410	6.876×10^{-4}
	0.0001	985	2.335×10^{-6}
4 Clusters	0.002	283, 267, 271, 279	8.348×10^{-4}
	0.0001	667, 751, 710, 731	7.943×10^{-6}

To obtain sparse and binary optimal designs, we solve a sequence of minimization problems of the form (17). Following the sparsification approach outlined in section 3.3, we use penalty functions of the form (8) with $\alpha = 0.1$ and $\varepsilon(i) = \left(\frac{2}{3}\right)^i, i = 1, 2, 3, \dots$. For our problem setup, the weights converge to binary values in approximately 20 iterations of the sparsification procedure. In our computations, we use the Broyden–Fletcher–Goldfarb–Shannon (BFGS) method available in python’s `scipy` library. We supply this minimization algorithm with the objective function and its gradient as described in section 4.

8.2. Solving the OEDUU problem

Here, we demonstrate that solving the OEDUU problem with our proposed approach is effective in producing designs for which the expected value of the posterior uncertainty (5) is small. In particular, we numerically verify that the SAA (11) is a reasonable approximation. This is done by solving the OEDUU problem with $N = 100$ SAA samples, and comparing the expected value of the objective obtained using the optimal experimental designs under uncertainty, which we will refer to as *uncertainty-aware designs*, with that obtained with *deterministic designs*. By deterministic designs we mean optimal designs obtained using a single sample from the irreducible uncertainty (velocity field and initial time); this amounts to minimizing (17) with $N = 1$.

The results are shown in figure 6. To obtain designs with different numbers of sensors, (17) was solved for different regularization parameters γ , which indirectly controls the sparsity of the optimal designs. To approximate the expectation shown on the y -axis, we use a Monte Carlo approximation of the expectation of the trace update $\mathbb{E}[\text{tr}(\mathbf{K}(\xi, \mathbf{w}))]$ with 100 samples from the irreducible uncertainty that are drawn independently from the SAA samples used to compute the uncertainty-aware designs. Additionally, to avoid bias due to the model reduction, we use more accurate reduced models, specifically we use algorithm 2 with $\mu = 10^{-4}$.

If the sampling as well as the model reduction errors are small enough, we would expect that the uncertainty-aware designs reduce the expected trace more than deterministic designs. As can be seen in figure 6, this is the case most of the time, i.e., we find good binary minimizers of the objective (5). Only very few deterministic designs are superior to the uncertainty-aware designs, which is likely due to sampling error.

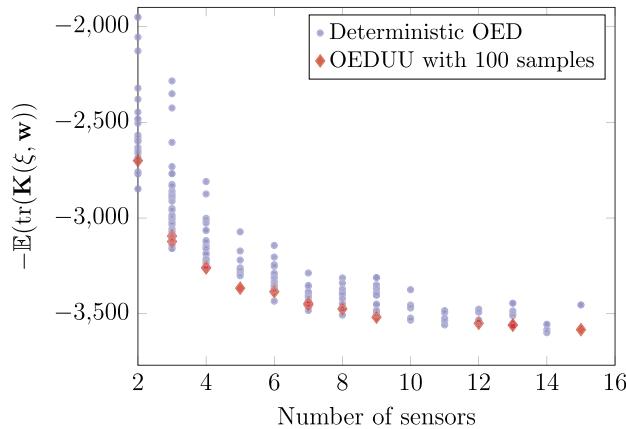


Figure 6. Trace reduction ϕ_N in objective (16) for deterministic OED (light blue dots) and OEDUU (red diamonds). The irreducible uncertainty for the uncertainty-aware designs is approximated using 100 samples from the irreducible uncertainty (velocity field and initial times). The deterministic designs are computed for 20 samples from the irreducible uncertainty. The x -axis shows the number of sensors for each design. The expectation in the objective, which is shown on the y -axis, is approximated using 100 samples that are chosen independently from the samples used to compute the designs.

8.3. Quality of the computed optimal designs

Ultimately, our goal is to choose the sensors for data collection which will allow us to learn the most about the initial concentration for a fixed velocity field and initial time T_0 . Here we demonstrate numerically that the designs computed via OEDUU perform better (produce more informative data) on average than deterministic designs computed for realizations of the irreducible model uncertainty.

Here, we study the effectiveness of the same uncertainty-aware designs used above not in expectation, but for 100 individual realizations ξ_i of the irreducible uncertainty, which again differ from the SAA samples used to compute the uncertainty-aware designs. Results are shown in figure 7, where now we show percentiles for the trace updates for individual ξ_i .

We observe that the designs obtained using OEDUU tend to have a smaller mean in the update trace than the designs obtained using individual ξ_i , particularly when we only use few sensors. Additionally, the 25th–75th and 2nd–98th percentiles show that poorly performing designs are less likely when one accounts for the uncertainty in the design computation. This is again, particularly the case for designs with small numbers of sensors. We believe that the reason for diminishing benefit of computing designs under uncertainty for larger number of sensors is that at each sensor location, five measurements in time are used. Thus, most information about the initial condition that can be recovered is already available from a rather small number of sensors and thus different designs play a less important role. Here, the diffusion contained in the governing PDE plays an important role as well, as it limits the resolution of the initial condition reconstruction that can be obtained from observations. This is also reflected by the decreasing gain of using more than five sensors.

The benefit of computing uncertainty-aware designs is greater for inverse problems where less information can be gained at each sensor location. This can be seen in figure 8. To obtain these results, we use the same inverse problem and ROM setup as described in sections 6–8.1 but restrict the times at which we can measure data to $\tau_i = 12, 15$, effectively reducing the

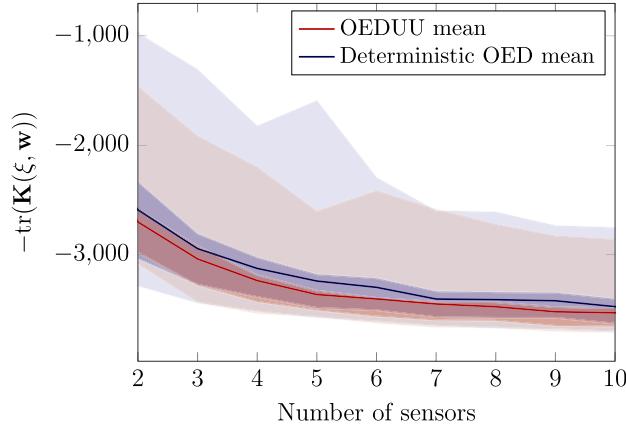


Figure 7. Comparing the mean for the trace in the deterministic objective, ϕ_N with $N = 1$ in (16), using deterministic designs and designs taking into account the uncertainty. The uncertainty-aware designs and the deterministic designs are the same as used in figure 6. Each design is used to evaluate the deterministic objective for 100 realizations of the irreducible uncertainty. The sample mean is plotted as a solid line and the shaded regions depict the 25th–75th and 2nd–98th percentile envelopes (for both the mean and the envelopes, red is used for uncertainty-aware designs and blue for deterministic designs).

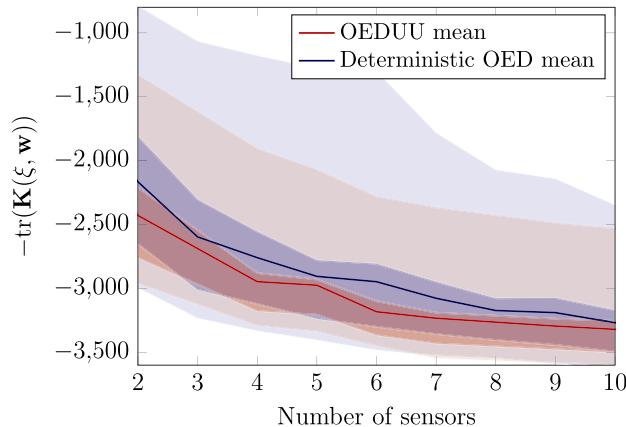


Figure 8. Same as figure 7, but with only two time observations at $\tau_i = 12, 15$ (rather than five). In this regime with less observations, the uncertainty-aware design outperforms deterministic designs more significantly also for larger numbers of sensors.

amount of information gathered at sensors and increasing the importance of careful sensor placement. Comparing figures 7 and 8, we can see that for ‘harder’ inverse problems, i.e., those with less-informative data, the difference between uncertainty-aware designs and deterministic designs is more significant.

9. Discussion and conclusions

We have developed a mathematical formulation and numerical scheme for computing A-optimal experimental designs for infinite-dimensional Bayesian linear inverse problems

governed by PDEs with irreducible model uncertainty. The proposed measurement space approach replaces trace estimation in an infinite-dimensional (high-dimensional upon discretization space) with trace estimation in the (finite-dimensional) measurement space. The computation of a *joint* reduced basis capturing the action of the forward operators for different random samples allows efficient computation of optimal designs. Numerical experiments for the inversion of initial concentration of a contaminant in groundwater indicate that, on average, designs that take the model uncertainty into account are superior to those that do not. This superiority is less pronounced as more sensors are being used.

We have used a Monte Carlo approach for dealing with uncertainty, which can require many samples for adequate resolution. A possible extension of our work is to consider alternate approaches for approximating the uncertainty, e.g., Taylor expansions of the uncertainty or stochastic approximation (SA). In the latter, the samples for the irreducible uncertainty are not chosen *a priori*, but are varied during the optimization. This avoids potential bias of the design towards the chosen samples but leads to several additional challenges in the optimization. Another important research question is the generalization of our approach to nonlinear inverse problems. One possibility is to follow the framework in [4], which is based on Gaussian approximations at the MAP point. Clearly, nonlinear problems under uncertainty become computationally rather expensive.

One aspect of OED under model uncertainty that is not explored in the present work is that of dealing with *reducible* sources of model uncertainty, i.e., additional uncertainties that could be reduced through observational data. However, the focus might be on estimation of primary parameters of interest and not on estimation of this (secondary) reducible uncertainty. Hence, the design should be chosen to focus mainly on the primary parameter and only on the secondary parameters to the extent that it aids inference of the primary parameters.

Acknowledgments

KK was supported in part by the Research Training Group in Modeling and Simulation funded by the National Science Foundation (NSF) via Grant RTG/DMS #1646339. GS acknowledges partial support from the NSF Grant #1723211. The authors would like to thank Noemi Petra and Benjamin Peherstorfer for helpful discussions.

References

- [1] Christie M A and Blunt M J 2001 *Tenth SPE comparative solution project: A comparison of upscaling techniques* SPE reservoir simulation symposium (London: Society of Petroleum Engineers)
- [2] Alexanderian A, Gloor P J, Ghattas O et al 2016 On Bayesian A-and D-optimal experimental designs in infinite dimensions *Bayesian Analysis* **11** 671–95
- [3] Alexanderian A, Petra N, Stadler G and Ghattas O 2014 A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0 -sparsification *SIAM J. Sci. Comput.* **36** A2122–48
- [4] Alexanderian A, Petra N, Stadler G and Ghattas O 2016 A fast and scalable method for A-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems *SIAM J. Sci. Comput.* **38** A243–72
- [5] Alexanderian A and Saibaba A K 2018 Efficient D-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems *SIAM J. Sci. Comput.* **40** A2956–85
- [6] Atkinson A C and Donev A N 1992 *Optimum Experimental Designs* (Oxford: Clarendon Press)
- [7] Attia A, Alexanderian A and Saibaba A K 2018 Goal-oriented optimal design of experiments for large-scale bayesian linear inverse problems *Inverse Problems* **34** 095009

[8] Bear J 2018 *Modeling Phenomena of Flow and Transport in Porous Media* vol 31 (Berlin: Springer)

[9] Beyer H-G and Sendhoff B 2007 Robust optimization—a comprehensive survey *Comput. Methods Appl. Mech. Eng.* **196** 3190–218

[10] Bui-Thanh T, Ghattas O, Martin J and Stadler G 2013 A computational framework for infinite-dimensional Bayesian inverse problems: Part I. The linearized case, with application to global seismic inversion *SIAM J. Sci. Comput.* **35** A2494–523

[11] Candes E J, Wakin M B and Boyd S P 2008 Enhancing sparsity by reweighted ℓ_1 minimization *J. Fourier Anal. Appl.* **14** 877–905

[12] Daon Y and Stadler G 2018 Mitigating the influence of the boundary on PDE-based covariance operators *Inverse Probl. Imag.* **12** 1083–102

[13] Dashti M and Stuart A M 2017 *The Bayesian Approach to Inverse Problems* (Berlin: Springer) pp 311–428

[14] Golub G H and Van Loan C F 2012 *Matrix Computations* vol 3 (Baltimore, MD: John Hopkins University Press)

[15] Haber E, Horesh L and Tenorio L 2008 Numerical methods for experimental design of large-scale linear ill-posed inverse problems *Inverse Problems* **24** 055012

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

[17] Haber E, Magnant Z, Lucero C and Tenorio L 2012 Numerical methods for A-optimal designs with a sparsity constraint for ill-posed inverse problems *Comput. Optim. Appl.* **1**–22

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

[19] Horesh L, Haber E and Tenorio L 2010 *Optimal Experimental Design for the Large-Scale Nonlinear Ill-Posed Problem of Impedance Imaging* (New York, NY: Wiley) pp 273–90

[20] Huan X and Marzouk Y M 2013 Simulation-based optimal Bayesian experimental design for nonlinear systems *J. Comput. Phys.* **232** 288–317

[21] Körkel S, Kostina E, Bock H G and Schlöder J P 2004 Numerical methods for optimal control problems in design of robust optimal experiments for nonlinear dynamic processes *Optim. Methods Software* **19** 327–38 The First Int. Conf. on Optimization Methods and Software. Part II

[22] Kouri D P and Shapiro A 2018 Optimization of PDEs with uncertain inputs *Frontiers in PDE-Constrained Optimization* (Berlin: Springer) pp 41–81

[23] Long Q, Scavino M, Tempone R and Wang S 2013 Fast estimation of expected information gains for Bayesian experimental designs based on Laplace approximations *Comput. Methods Appl. Mech. Eng.* **259** 24–39

[24] Neitzel I, Pieper K, Vexler B and Walter D 2019 A sparse control approach to optimal sensor placement in PDE-constrained parameter estimation problems (arXiv:1905.01696)

[25] Pázman A 1986 *Foundations of Optimum Experimental Design* 14 (Berlin: Springer)

[26] Peherstorfer B, Butnaru D, Willcox K and Bungartz H-J 2014 Localized discrete empirical interpolation method *SIAM J. Sci. Comput.* **36** A168–92

[27] Pukelsheim F 1993 *Optimal Design of Experiments* (New York, NY: Wiley)

[28] Roininen L, Huttunen J M J and Lasanen S 2014 Whittle–Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography *Inverse Probl. Imag.* **8** 561

[29] Ruthotto L, Chung J and Chung M 2018 Optimal experimental design for inverse problems with state constraints *SIAM J. Sci. Comput.* **40** B1080–100

[30] Sahinidis N V 2004 Optimization under uncertainty: state-of-the-art and opportunities *Comput. Chem. Eng.* **28** 971–83

[31] Shapiro A, Dentcheva D and Ruszczynski A 2014 *Lectures on Stochastic Programming: Modeling and Theory* (Society for Industrial and Applied Mathematics)

[32] Uciński D 2005 *Optimal Measurement Methods for Distributed Parameter System Identification* (Boca Raton: CRC Press)

[33] Villa U, Petra N and Ghattas O 2019 hIPPYlib: an extensible software framework for large-scale inverse problems governed by PDEs: Part I. Deterministic inversion and linearized bayesian inference (arXiv:1909.03948)

[34] Yu J and Anitescu M 2019 *Multidimensional Sum-Up Rounding for Integer Programming in Optimal Experimental Design* (Mathematical Programming) pp 1–40