

# A-OPTIMAL DESIGN OF EXPERIMENTS FOR INFINITE-DIMENSIONAL BAYESIAN LINEAR INVERSE PROBLEMS WITH REGULARIZED $\ell_0$ -SPARSIFICATION\*

ALEN ALEXANDERIAN<sup>†</sup>, NOEMI PETRA<sup>†</sup>, GEORG STADLER<sup>†</sup>, AND  
OMAR GHATTAS<sup>†‡§</sup>

**Abstract.** We present an efficient method for computing A-optimal experimental designs for infinite-dimensional Bayesian linear inverse problems governed by partial differential equations (PDEs). Specifically, we address the problem of optimizing the location of sensors (at which observational data are collected) to minimize the uncertainty in the parameters estimated by solving the inverse problem, where the uncertainty is expressed by the trace of the posterior covariance. Computing optimal experimental designs (OEDs) is particularly challenging for inverse problems governed by computationally expensive PDE models with infinite-dimensional (or, after discretization, high-dimensional) parameters. To alleviate the computational cost, we exploit the problem structure and build a low-rank approximation of the parameter-to-observable map, preconditioned with the square root of the prior covariance operator. The availability of this low-rank surrogate, relieves our method from expensive PDE solves when evaluating the optimal experimental design objective function, i.e., the trace of the posterior covariance, and its derivatives. Moreover, we employ a randomized trace estimator for efficient evaluation of the OED objective function. We control the sparsity of the sensor configuration by employing a sequence of penalty functions that successively approximate the  $\ell_0$ -“norm”; this results in binary designs that characterize optimal sensor locations. We present numerical results for inference of the initial condition from spatio-temporal observations in a time-dependent advection-diffusion problem in two and three space dimensions. We find that an optimal design can be computed at a cost, measured in number of forward PDE solves, that is independent of the parameter and sensor dimensions. Moreover, the numerical optimization problem for finding the optimal design can be solved in a number of interior-point quasi-Newton iterations that is insensitive to the parameter and sensor dimensions. We demonstrate numerically that  $\ell_0$ -sparsified experimental designs obtained via a continuation method outperform  $\ell_1$ -sparsified designs.

**Key words.** Optimal experimental design, A-optimal design, Bayesian inference, sensor placement, ill-posed inverse problems, low-rank approximation, randomized trace estimator, randomized SVD.

**AMS subject classifications.** 62K05, 35Q62, 62F15, 35R30, 35Q93, 65C60.

**1. Introduction.** Recent advances in theory [32] and numerical algorithms (e.g., [8]) are enabling efficient solution of infinite-dimensional Bayesian inverse problems. This opens the door to consideration of the upstream question: how do we place sensors to *optimally* infer model parameters for large-scale problems? Here we present an efficient method for such optimal experimental design (OED) problems. Specifically, we consider Bayesian linear inverse problems governed by PDEs whose solution is the posterior probability law for a parameter field. The numerical solution of such inverse problems is challenging due to the infinite (or, when discretized, large) dimension of the parameters, ill-posedness of the inverse problem, and expensive-to-compute PDE models. The Bayesian inverse problem is by itself very challenging, but it is merely a subproblem within the OED problem, and must be solved repeatedly when using conventional OED methods. Hence, it is essential to make maximum use

<sup>†</sup>Institute for Computational Engineering & Sciences, The University of Texas at Austin, Austin, TX, USA

<sup>‡</sup>Department of Mechanical Engineering, The University of Texas at Austin, Austin, TX, USA

<sup>§</sup>Department of Geological Sciences, The University of Texas at Austin, Austin, TX, USA

\*This work was partially supported by NSF grant ARC-0941678; DOE grants DE-FC02-13ER26128, DE-SC0010518 DE-FC02-11ER26052, DE-11018096, and DE-FG02-09ER25914; and AFOSR grant FA9550-12-1-0484.

of the problem structure to realize efficient algorithms that are *scalable*, i.e., their performance—measured in number of governing (forward) PDE solves—is independent of the discretized parameter and sensor dimensions, and the discretization of the governing PDE.

When formulating an OED problem, a basic question is the precise meaning of what constitutes the *design*. The present work concerns computation of optimal sensor locations where observational data will be collected. A subsequent question concerns the definition of an *optimal* design, which leads to the choice of the design criterion. For Bayesian inverse problems, a natural choice is to seek a design that minimizes the average posterior variance of the inversion parameters, leading to the Bayesian A-optimal design criterion.

Standard references for optimal experimental design include [2,6,28,30,33]. While most of the classical texts concern problems with small or moderate parameter dimension and focus mainly on well-posed inverse problems, there has been recent interest in optimal design for large-scale ill-posed linear [16,18] and nonlinear [17,20] inverse problems. The numerical methods in the present paper are closest to those in [16,18], where the authors consider finite-dimensional linear inverse problems, and develop a framework to control the mean square error of the regularized Tikhonov estimates. This leads to a design criterion that seeks to minimize the sum of the estimation bias and the variability of the estimator around its mean, which can be related to A-optimal designs. These contributions use  $\ell_1$ -penalties to control the sparsity of the design and, in [18], a low-rank singular value decomposition (SVD) of the parameter-to-observable map is used; the present paper builds on and extends both of these ideas. Further recent work that employs a Bayesian formulation includes [21,22], where the authors use a decision theoretic design criterion, generalized polynomial chaos surrogates, and stochastic optimization to tackle nonlinear inverse problems, albeit in low to moderate parameter dimension.

In this paper, we devise scalable numerical methods for computing A-optimal designs for infinite-dimensional Bayesian linear inverse problems governed by (time-dependent) PDEs. As suggested in [18], having a low-rank SVD *surrogate* of the parameter-to-observable map relieves the OED method of repeated PDE solves. In the Bayesian context, we can improve on this idea and further exploit problem structure; namely, we construct a low-rank SVD representation of the parameter-to-observable map *preconditioned* by the square root of the prior covariance operator [8,14]. This preconditioning amounts to filtering through the prior the information gained from the data about the model parameters. In the case of smoothing priors usually used in infinite-dimensional problems, this preconditioning results in faster spectral decay and thus allows for a more efficient low-rank approximation. The remaining steps in the solution of the OED problem use this low-rank surrogate and thus do not require additional PDE solves. As a result of a consistent discretization of the problem, i.e., one that respects the infinite-dimensional Hilbert-space structure, the numerical rank of the prior-preconditioned parameter-to-observable map is bounded with respect to the discretized parameter dimension.

We consider a finite number of candidate locations for the placement of sensors; the optimal configuration is a sparse subset of these locations. To each candidate sensor location we assign a non-negative number that weights the observation from that sensor. Finding an optimal design then amounts to choosing an optimal weight vector; a weight of 0 indicates absence of a sensor and a weight of 1 corresponds to a sensor being placed at that location. For computational convenience, we allow the

weights to take on any value in  $[0, 1]$ , and use a sparsifying penalty to control the number of nonzero weights, and thus, the number of allocated sensors. One option for such a penalty is the  $\ell_1$ -norm; see, e.g., [16, 18, 20]. This approach, however, does not lead to a binary (i.e., 0–1) design. Motivated by continuation methods used in topology optimization [4, 5], we propose to solve a sequence of OED problems with penalty functions that successively approximate the  $\ell_0$ -“norm”. This, in contrast to an  $\ell_1$ -penalty approach, does result in a binary design. As a test problem for our OED method, we consider a forward problem in the form of a time-dependent advection diffusion model, in which we infer the probability law of the initial condition. For this problem, we demonstrate the success of our continuation method, and show that the weights found by this continuation approach improve over optimal designs obtained via an  $\ell_1$ -penalty approach.

For typical infinite-dimensional Bayesian inverse problems, the performance of our OED method is insensitive to the number of candidate sensor locations. This is due to the fact that although the dimension of the observations increases with the number of candidate locations, the amount of independent information that can be gained from nearby sensors is typically limited. Thus, increasing the number of sensors beyond a certain point does not significantly increase the numerical rank of the prior-preconditioned parameter-to-observable map. As a consequence, the number of PDE solves required to compute a low-rank SVD surrogate for the (prior-preconditioned) parameter-to-observable map is bounded as the number of candidate sensor locations increases. Moreover, in our computational results we find that the numerical optimization problem to compute the optimal design can be solved in a number of (quasi-Newton) iterations that is independent of the number of candidate locations.

The large-scale nature of the Bayesian inverse problems we target necessitates the use of randomized methods in linear algebra. In particular, we utilize randomized trace estimators [3, 23] to estimate the trace of the posterior covariance operator, and randomized SVD [19] to compute a low-rank surrogate of the prior-preconditioned parameter-to-observable map. Moreover, for computing the application of matrix square roots, as needed in our method, we employ matrix-free iterative methods [11].

The structure of this paper is as follows. After presenting the requisite background material in Section 2, we formulate the optimal design problem in infinite dimensions in Section 3. Then, we detail the components of our OED method in Section 4. Section 5 provides a description of our model problem, namely the inference of the initial condition in a time-dependent advection-diffusion equation. We present a comprehensive numerical study in Section 6. Finally, in Section 7, we draw conclusions, and discuss limitations and possible extensions of our method.

**2. Background.** In this section, we provide the background material required for the formulation and numerical solution of optimal experimental design problems in the context of infinite-dimensional Bayesian inverse problems. In Section 2.1, we present the Bayesian inverse problem in an infinite-dimensional Hilbert space setting, adopting the framework in [32]. In Section 2.2, we describe a discretization that is consistent with the infinite-dimensional inference problem formulation; this presentation follows [8]. Finally, in Section 2.3, we briefly comment on the randomized SVD and randomized trace estimators, which are used in our numerical method.

**2.1. Bayesian inversion in Hilbert spaces.** We begin our discussion by first considering a deterministic inverse problem. Given finite-dimensional observations

$\mathbf{d} \in \mathbb{R}^q$ , we seek the model parameter  $m$  that solves

$$\min_{m \in \mathcal{H}} \mathcal{J}(m) := \frac{1}{2} \|\mathcal{F}(m) - \mathbf{d}\|_{\Gamma}^2 + \mathcal{R}(m). \quad (2.1)$$

The function  $\mathcal{F} : \mathcal{H} \rightarrow \mathbb{R}^q$  is the parameter-to-observable map and  $\mathcal{R}$  denotes a regularization term. In the applications targeted in this paper  $\mathcal{H} = L^2(\mathcal{D})$ , where  $\mathcal{D} \subset \mathbb{R}^d$  is a bounded domain (with  $d = 2, 3$ ) and an evaluation of  $\mathcal{F}$  involves the solution of a PDE, followed by the application of an observation operator. Note that the solution of a deterministic inverse problem can be thought of as a *point estimate* of  $m$ . To obtain a full probabilistic description of the parameter  $m$ , we are led to a Bayesian formulation of the problem, whose solution is a posterior probability law for  $m$ .

In this paper, we consider a parameter  $m$  which is modeled as a random-field (random function). To be precise, letting  $(\Omega, \Sigma, \mathbb{P})$  be an appropriate probability space,  $m : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$  is a function such that for each  $\mathbf{x} \in \mathcal{D}$ ,  $m(\mathbf{x}, \cdot)$ , is a real-valued random variable; thus, we can view  $m$  as an indexed collection of random variables,  $\{m(\mathbf{x})\}_{\mathbf{x} \in \mathcal{D}}$ , where, following the common practice, we suppress the dependence on  $\omega$ . On the other hand, for each  $\omega \in \Omega$ ,  $m(\cdot, \omega) : \mathcal{D} \rightarrow \mathbb{R}$  is a real-valued function. We consider the case where  $m(\cdot, \omega) \in \mathcal{H}$ , and thus, we can also view  $m$  as a random-variable,  $m : (\Omega, \Sigma, \mathbb{P}) \rightarrow (\mathcal{H}, \mathfrak{B}(\mathcal{H}))$ , where  $\mathfrak{B}(\mathcal{H})$  denotes the Borel  $\sigma$ -algebra on  $\mathcal{H}$ . Recall that the law of  $m$  is a probability measure  $\mu$  on  $(\mathcal{H}, \mathfrak{B}(\mathcal{H}))$  given by  $\mu(E) = \mathbb{P}(m \in E)$  for  $E \in \mathfrak{B}(\mathcal{H})$ .

The infinite-dimensional Bayesian inverse problem can then be formulated as using observations to update our knowledge of the law of  $m$ , as a probability measure on  $(\mathcal{H}, \mathfrak{B}(\mathcal{H}))$ . Since, in contrast to the finite-dimensional case, there is no Lebesgue measure on  $\mathcal{H}$ , the infinite-dimensional Bayes formula is given by

$$\frac{d\mu_{\text{post}}}{d\mu_0} \propto \pi_{\text{like}}(\mathbf{d}|m). \quad (2.2)$$

Here,  $\frac{d\mu_{\text{post}}}{d\mu_0}$  denotes the Radon-Nikodym derivative [34] of the posterior measure  $\mu_{\text{post}}$  with respect to  $\mu_0$ , and  $\pi_{\text{like}}(\mathbf{d}|m)$  denotes the data likelihood. Conditions under which the posterior measure is well defined and (2.2) holds are given in detail in [32]. Note that we consider a finite-dimensional observation vector, motivated by the fact that in practice data are available only at a finite number of sensor locations and a finite number of points in time.

In the present work, we consider the Gaussian-linear case, i.e., the parameter-to-observable map  $\mathcal{F}$  is linear. Moreover, we assume an additive noise model,  $\mathbf{d} = \mathcal{F}m + \boldsymbol{\eta}$ , where  $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Gamma}_{\text{noise}})$  is a centered Gaussian on  $\mathbb{R}^q$ ; the latter implies

$$\pi_{\text{like}}(\mathbf{d}|m) \propto \exp \left\{ -\frac{1}{2} (\mathcal{F}m - \mathbf{d})^T \mathbf{\Gamma}_{\text{noise}}^{-1} (\mathcal{F}m - \mathbf{d}) \right\}. \quad (2.3)$$

We use a Gaussian prior,  $\mu_0 = \mathcal{N}(m_0, \mathcal{C}_0)$ , where  $m_0 \in \mathcal{H}$  is sufficiently regular, and  $\mathcal{C}_0$  is an appropriate covariance operator, i.e.,  $\mathcal{C}_0$  must be symmetric, positive, and of trace-class. We define the covariance operator as the inverse of an elliptic differential operator. A common alternative choice for statistical inverse problems is to specify a covariance function between any two spatial points, which results in a dense covariance matrix. For large-scale problems, however, the construction and “inversion” of such a dense covariance matrix can be infeasible. On the contrary,

specifying the covariance as the inverse of an elliptic differential operator allows to build on existing fast solvers for elliptic equations. As detailed in [8, 32], the PDE solution operator used as covariance operator  $\mathcal{C}_0$  must be sufficiently smoothing and have bounded Green's functions. For example, the biharmonic operator has bounded Green's functions in two and three space dimensions. Therefore, we choose  $\mathcal{C}_0 = \mathcal{A}^{-2}$ , with  $\mathcal{A}$  a Laplacian-like operator in the sense of Assumption 2.9 in [32]. This choice also allows efficient applications of the square root operator  $\mathcal{C}_0^{1/2} = \mathcal{A}^{-1}$ , as required below. The elliptic PDE corresponding to  $\mathcal{A}$  written in weak form is as follows: For  $s \in \mathcal{H} = L^2(\mathcal{D})$ , the solution  $m = \mathcal{A}^{-1}s$  satisfies

$$\int_{\mathcal{D}} \alpha \nabla m \cdot \nabla p + \beta m p \, d\mathbf{x} = \int_{\mathcal{D}} s p \, d\mathbf{x}, \quad \text{for all } p \in H^1(\mathcal{D}), \quad (2.4)$$

with  $\alpha, \beta > 0$  controlling the variation and the correlation length. Due to the present choice of the prior and the noise model, and the linearity of  $\mathcal{F}$ , the posterior measure is a Gaussian,  $\mathcal{N}(m_{\text{post}}, \mathcal{C}_{\text{post}})$  with [32, Section 6.4],

$$\mathcal{C}_{\text{post}} = (\mathcal{F}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{F} + \mathcal{C}_0^{-1})^{-1}, \quad m_{\text{post}} = \mathcal{C}_{\text{post}} (\mathcal{F}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{d} + \mathcal{C}_0^{-1} m_0), \quad (2.5)$$

where  $\mathcal{F}^* : \mathbb{R}^q \rightarrow \mathcal{H}$  is the adjoint of  $\mathcal{F}$ .

**2.2. Discretization of the infinite-dimensional Bayesian inverse problem.** In this section, we describe the discretization of the Bayesian inverse problem (2.2) in the Gaussian linear case. We consider a finite-dimensional subspace  $V_h \subset L^2(\mathcal{D})$  given by  $V_h = \text{span}\{\phi_1, \dots, \phi_n\}$ , where  $\{\phi_j\}_{j=1}^n$  are continuous Lagrange nodal basis functions. Given  $m_h \in V_h$ , we denote by  $\mathbf{m}$ , the vector of its coordinates in  $V_h$ ; i.e., for  $m_h = \sum_{j=1}^n m_j \phi_j$ , we have  $\mathbf{m} = (m_1, \dots, m_n)^T$ . After this discretization, we replace the task of inferring the parameter  $m \in L^2(\mathcal{D})$  with that of inferring the coefficients for the finite-element approximation  $m_h$  of  $m$ .

Following [8], we state the finite-dimensional Bayesian inverse problem such that it is consistent with the corresponding inference problem in  $L^2(\mathcal{D})$ . Consequently, we work in  $\mathbb{R}^n$ , with the weighted inner product,  $\langle \cdot, \cdot \rangle_{\mathbf{M}}$  given by  $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbf{M}} = \langle \mathbf{M}\mathbf{x}, \mathbf{y} \rangle$ , where  $\langle \cdot, \cdot \rangle$  denotes the Euclidean inner product, and  $\mathbf{M}$  is the (symmetric positive definite) finite-element mass matrix. It is convenient to introduce the notation  $\mathbb{R}_{\mathbf{M}}^n$  for  $\mathbb{R}^n$  when endowed with the  $\langle \cdot, \cdot \rangle_{\mathbf{M}}$  inner product. Note that the mapping  $m_h \mapsto \mathbf{m}$  is a Hilbert-space isomorphism between  $V_h$  (with  $L^2$ -inner product) and  $\mathbb{R}_{\mathbf{M}}^n$ . For a linear operator  $\mathbf{A} : \mathbb{R}_{\mathbf{M}}^n \rightarrow \mathbb{R}_{\mathbf{M}}^n$ , the adjoint operator is given by  $\mathbf{A}^* = \mathbf{M}^{-1} \mathbf{A}^T \mathbf{M}$ . A linear operator  $\mathbf{A}$  on  $\mathbb{R}_{\mathbf{M}}^n$  is self-adjoint if  $\mathbf{A} = \mathbf{A}^*$ ; for convenience, we refer to such operators as  $\mathbf{M}$ -symmetric. In the sequel, we will encounter linear mappings  $\mathbf{A}_1 : \mathbb{R}_{\mathbf{M}}^n \rightarrow \mathbb{R}^q$  and  $\mathbf{A}_2 : \mathbb{R}^r \rightarrow \mathbb{R}_{\mathbf{M}}^n$ , where  $\mathbb{R}^q$  and  $\mathbb{R}^r$  are endowed with the Euclidean inner product; the corresponding adjoints are given by [8],  $\mathbf{A}_1^* = \mathbf{M}^{-1} \mathbf{A}_1^T$  and  $\mathbf{A}_2^* = \mathbf{A}_2^T \mathbf{M}$ .

For the discretized problem, the density for the prior (as a measure over the space  $\mathbb{R}_{\mathbf{M}}^n$ ) is characterized by

$$\pi_{\text{prior}}(\mathbf{m}) \propto \exp \left\{ -\frac{1}{2} \|\mathbf{A}(\mathbf{m} - \mathbf{m}_0)\|_{\mathbf{M}}^2 \right\}, \quad (2.6)$$

with  $\mathbf{A} = \mathbf{M}^{-1} \mathbf{L}$ , where  $\mathbf{L} = \alpha \mathbf{K} + \beta \mathbf{M}$  and  $\mathbf{K}$  is the finite-element stiffness matrix. Note that  $\mathbf{A}$  is  $\mathbf{M}$ -symmetric; moreover, it follows from the above definition that  $\mathbf{\Gamma}_{\text{prior}} = \mathbf{A}^{-2}$ . Subsequently, the posterior is a Gaussian  $\mathcal{N}(\mathbf{m}_{\text{post}}, \mathbf{\Gamma}_{\text{post}})$  with

$$\mathbf{m}_{\text{post}} = \mathbf{\Gamma}_{\text{post}} \left( \mathbf{F}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{d} + \mathbf{\Gamma}_{\text{prior}}^{-1} \mathbf{m}_0 \right), \quad \mathbf{\Gamma}_{\text{post}} = \left( \mathbf{F}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{F} + \mathbf{\Gamma}_{\text{prior}}^{-1} \right)^{-1},$$

where  $\mathbf{F} : \mathbb{R}_M^n \rightarrow \mathbb{R}^q$  is the discretization of the parameter-to-observable map  $\mathcal{F}$ .

Next, we summarize how to draw samples from and to compute the variance of the discretized posterior measure; here we follow [8], where details of these operations are provided. For a posterior measure  $\mathcal{N}(\mathbf{m}_{\text{post}}, \mathbf{\Gamma}_{\text{post}})$  generating samples requires a decomposition of  $\mathbf{\Gamma}_{\text{post}}$  in the form,  $\mathbf{\Gamma}_{\text{post}} = \mathbf{Q}\mathbf{Q}^*$ . Then, to generate a sample  $\boldsymbol{\nu}$  from  $\mathcal{N}(\mathbf{m}_{\text{post}}, \mathbf{\Gamma}_{\text{post}})$ , we draw a realization  $\mathbf{z}$  from  $\mathcal{N}(0, \mathbf{I})$  and compute  $\boldsymbol{\nu}$  as

$$\boldsymbol{\nu} = \mathbf{m}_{\text{post}} + \mathbf{Q}\mathbf{M}^{-1/2}\mathbf{z}. \quad (2.7)$$

The discretized covariance is given by  $\text{Cov}\{m_i, m_j\} = \mathbf{e}_i^T \mathbf{\Gamma}_{\text{post}} \mathbf{M}^{-1} \mathbf{e}_j$ , for  $i, j = 1, \dots, n$ . In particular,

$$\text{Var}\{m_i\} = \mathbf{e}_i^T \mathbf{\Gamma}_{\text{post}} \mathbf{M}^{-1} \mathbf{e}_i, \quad i = 1, \dots, n, \quad (2.8)$$

where  $\text{Var}\{m_i\}$  denotes the variance of  $m_i$ .

**2.3. Randomized linear algebra algorithms.** One of the major components of our method is a low-rank SVD surrogate for the prior-preconditioned parameter-to-observable map. To compute such low-rank surrogates, we use a randomized SVD [19]. This choice is motivated by the flexibility and the robustness of the method, and by the fact that, as opposed to Krylov subspace methods, it only requires independent matrix-vector products. This aspect is particularly useful for the large-scale problems we target, in which the matrix-vector applications involve expensive PDE solves; see also [7]. Randomized SVD methods can be made very accurate with negligible probability of failure [19].

Computing A-optimal designs requires minimizing the trace of large dense covariance matrices, which, in our target problems, usually have a rapidly decaying spectrum and the eigenvalues are clustered around 0. For such matrices, which are defined implicitly through their applications to vectors, randomized trace estimators provide a reasonably accurate approximation of the trace with a small number of random vectors (see e.g., [31]). These estimators involve only matrix-vector products, which makes them suitable for large scale problems. In particular, randomized trace estimators approximate the trace of a matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$ , via Monte-Carlo estimates of the form  $\text{tr}(\mathbf{A}) \approx \frac{1}{N_{\text{tr}}} \sum_{i=1}^{N_{\text{tr}}} [\mathbf{z}^{(i)}]^T \mathbf{A} \mathbf{z}^{(i)}$ , where the trial vectors  $\mathbf{z}^{(i)}$  are random  $n$ -vectors. A well-known example is the Hutchinson estimator [23], which uses random vectors  $\mathbf{z}^{(i)}$  with  $\pm 1$  entries, each with a probability of 1/2. Another possibility, used in this paper, is the Gaussian trace estimator, which uses Gaussian random vectors with independent and identically distributed (i.i.d.) standard normal entries. For a description and analysis of different trace estimators, we refer to [3].

**3. A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems.** In this section, we formulate the A-optimal design problem for infinite-dimensional Bayesian linear inverse problems. The extension of the A-optimal design criterion to the infinite-dimensional setting is described in Section 3.1. In Section 3.2, we specify the mathematical definition of a *design*, and describe how the design is introduced in the Bayesian inverse problem. Finally, in Section 3.3, we formulate the resulting OED optimization problem.

**3.1. A-optimal design in Hilbert spaces.** In a finite-dimensional inference problem, an A-optimal design minimizes the average posterior variance of the inference parameters [33]. In the linear case, this is accomplished by minimizing the trace of the posterior *covariance matrix*. Since the present work concerns inference problems with

a random field as the inference parameter, we first extend the notion of an A-optimal design to the infinite-dimensional Hilbert space setting.

The average posterior variance of  $m$  over the physical domain  $\mathcal{D}$  is given by

$$\frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} c_{\text{post}}(\mathbf{x}, \mathbf{x}) d\mathbf{x},$$

where  $|\mathcal{D}|$  denotes the Lebesgue measure of the domain  $\mathcal{D}$ , and  $c_{\text{post}}$  is the *covariance function* of  $m$ :

$$c_{\text{post}}(\mathbf{x}, \mathbf{y}) = \mathbb{E} \left\{ \left( m(\mathbf{x}) - m_{\text{post}}(\mathbf{x}) \right) \left( m(\mathbf{y}) - m_{\text{post}}(\mathbf{y}) \right) \right\}, \quad \mathbf{x}, \mathbf{y} \in \mathcal{D},$$

where  $\mathbb{E} \{ \cdot \}$  denotes the expectation operator. Note that  $c_{\text{post}}$  is related to the *covariance operator*  $\mathcal{C}_{\text{post}}$ , i.e.,

$$[\mathcal{C}_{\text{post}}u](\mathbf{x}) = \int_{\mathcal{D}} c_{\text{post}}(\mathbf{x}, \mathbf{y})u(\mathbf{y}) d\mathbf{y}, \quad u \in L^2(\mathcal{D}).$$

The covariance operator  $\mathcal{C}_{\text{post}}$  is positive, symmetric, and of trace-class, and thus has real eigenvalues,  $\{\lambda_i\}_{i=1}^{\infty}$ , and a complete orthonormal set of eigenvectors,  $\{e_i\}_{i=1}^{\infty}$ . Mercer's Theorem [24, 26] then provides a decomposition of the covariance function  $c_{\text{post}}$  from the spectral decomposition of  $\mathcal{C}_{\text{post}}$  through  $c_{\text{post}}(\mathbf{x}, \mathbf{y}) = \sum_j \lambda_j e_j(\mathbf{x})e_j(\mathbf{y})$ , where the convergence of the infinite sum is uniform and absolute in  $\mathcal{D} \times \mathcal{D}$ . From this representation, one obtains,

$$\int_{\mathcal{D}} c_{\text{post}}(\mathbf{x}, \mathbf{x}) d\mathbf{x} = \sum_j \lambda_j = \text{tr}(\mathcal{C}_{\text{post}}).$$

Thus, we formulate the A-optimal design problem as that of minimizing the trace of the covariance *operator*,  $\text{tr}(\mathcal{C}_{\text{post}})$ . This shows that the definition of an A-optimal design in finite dimensions, namely the minimization of the average variance of the estimates, extends naturally to infinite-dimensional Bayesian inverse problems involving a random-field as the inversion parameter. We emphasize that the trace of the posterior covariance operator in the A-optimal design criterion is well defined due to the proper choice of the prior measure for infinite-dimensional inference problems.

**3.2. Introducing the design in the Bayesian inverse problem.** Let us first specify the notion of a design in the context of our target applications, namely optimal sensor placement. We use a finite-dimensional design space, that is, we fix a set of points  $\mathbf{x}_i$ ,  $i = 1, \dots, N_s$  as the set of candidate sensor locations, and associate to each  $\mathbf{x}_i$  a non-negative weight  $w_i \in \mathbb{R}$ ; see also [16, 18, 33]. The OED problem is then formulated as an optimization problem over the weight vector,  $\mathbf{w} = (w_1, \dots, w_{N_s})^T$ . The points  $\mathbf{x}_i$ ,  $i = 1, \dots, N_s$  can be thought of as a discretization of the *sensor domain*, which is a subset of  $\mathcal{D}$ .

The design weights can have different interpretations. For example, in classical formulations such as in [33],  $w_i$  define a probability mass function, i.e.,  $w_i \geq 0$  and  $\sum w_i = 1$ ; one may then interpret large weights as ones indicating promising locations for placing sensors. If the inversion is based on a repeatable experiment, weights can also be used to specify the number of experiments performed at each sensor location to control the observation noise; see, e.g., [16]. In many inverse problems, however, experiments cannot be repeated or the mathematical model is not an exact



representation of the physical phenomenon underlying the observations, such that the observation error cannot be controlled. Thus, we prefer weight vectors containing zeros and ones only, indicating absence or presence of sensors over the candidate sensor grid. Unfortunately, solving optimization problems for vectors with binary components is a difficult combinatorial problem. Hence, we employ a relaxation of the problem and consider weights  $w_i \in [0, 1]$ . In Section 4.5, we devise a method of recovering the desired 0–1 structure using sparsifying penalties combined with a continuation procedure.

Next, we describe the process of introducing the design vector  $\mathbf{w}$  into the Bayesian inverse problem. Since the design guides the collection of data, the weight vector  $\mathbf{w}$  enters the inference problem (2.2) through the data likelihood (2.3). The  $\mathbf{w}$ -weighted data-likelihood is given by,

$$\pi_{\text{like}}(\mathbf{d}|m; \mathbf{w}) \propto \exp \left\{ -\frac{1}{2}(\mathcal{F}m - \mathbf{d})^T \mathbf{W}^{1/2} \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{W}^{1/2} (\mathcal{F}m - \mathbf{d}) \right\}, \quad (3.1)$$

where  $\mathbf{W} \in \mathbb{R}^{q \times q}$  is a diagonal matrix with weights on its diagonal. For time-dependent problems, with observations collected at sensor locations at discrete points in time, we have  $\mathcal{F} : \mathcal{H} \rightarrow \mathbb{R}^q$  with  $q = N_s N_\tau$ , where  $N_s$  and  $N_\tau$  are the number of candidate sensors and observation times, respectively. Here,  $\mathbf{W}$  is a block-diagonal matrix having  $N_\tau$  blocks, where each block is an  $N_s \times N_s$  diagonal matrix with  $\mathbf{w}$  on its diagonal. The posterior covariance operator is thus given by

$$\mathbf{C}_{\text{post}}(\mathbf{w}) = (\mathcal{F}^* \mathbf{W}^{1/2} \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{W}^{1/2} \mathcal{F} + \mathbf{C}_0^{-1})^{-1}. \quad (3.2)$$

From this point on, we work with the discretization of the infinite-dimensional problem, where we follow the discretization strategy described in Section 2.2. In particular, we have the following discretized posterior covariance,

$$\mathbf{\Gamma}_{\text{post}}(\mathbf{w}) = (\mathbf{F}^* \mathbf{W}^{1/2} \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{W}^{1/2} \mathbf{F} + \mathbf{\Gamma}_{\text{prior}}^{-1})^{-1}. \quad (3.3)$$

The mean  $\mathbf{m}_{\text{post}}$  of the discretized posterior measure coincides with the maximum a posteriori probability (MAP) estimate, given as the solution of the minimization problem,

$$\min_{\mathbf{m} \in \mathbb{R}^n} \frac{1}{2} \left\langle \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{W}^{1/2} (\mathbf{d} - \mathbf{F}\mathbf{m}), \mathbf{W}^{1/2} (\mathbf{d} - \mathbf{F}\mathbf{m}) \right\rangle + \frac{1}{2} \left\langle \mathbf{\Gamma}_{\text{prior}}^{-1} (\mathbf{m} - \mathbf{m}_0), \mathbf{m} - \mathbf{m}_0 \right\rangle_{\mathbf{M}}.$$

The inverse of the Hessian  $\mathbf{H}(\mathbf{w})$  of the above functional coincides with  $\mathbf{\Gamma}_{\text{post}}(\mathbf{w})$  defined in (3.3). In the following, we call the Hessian of the misfit term,

$$\mathbf{H}_{\text{misfit}}(\mathbf{w}) = \mathbf{F}^* \mathbf{W}^{1/2} \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{W}^{1/2} \mathbf{F},$$

the *misfit Hessian*; thus,  $\mathbf{H}(\mathbf{w}) = \mathbf{H}_{\text{misfit}}(\mathbf{w}) + \mathbf{\Gamma}_{\text{prior}}^{-1}$ . In this paper, we consider the case of independent observations, and hence  $\mathbf{\Gamma}_{\text{noise}}$  is a diagonal matrix. This assumption turns  $\mathbf{W}^{1/2} \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{W}^{1/2}$  into a diagonal matrix with diagonal entries  $w_i / \sigma_i^2$ . For simplicity of the presentation, we further assume that  $\mathbf{\Gamma}_{\text{noise}}$  is a constant multiple of identity, i.e.,  $\mathbf{\Gamma}_{\text{noise}} = \sigma \mathbf{I}$ , and set  $\sigma = 1$ . Thus, the misfit Hessian takes the form  $\mathbf{H}_{\text{misfit}} = \mathbf{F}^* \mathbf{W} \mathbf{F}$ . The algorithms presented below can be easily modified to accommodate a general diagonal noise covariance matrices  $\mathbf{\Gamma}_{\text{noise}}$ .



**3.3. The OED problem.** Finally, we can formulate the Bayesian A-optimal experimental design problem as optimization problem for the weight vector  $\mathbf{w}$ . Following the discussions in the previous sections, the OED objective function is the trace of the posterior covariance operator (3.2), which, after discretization is given by (3.3). Additionally, we use a penalization to control the sparsity of the design. Hence, the optimal design vector is the solution to the following optimization problem:

$$\begin{aligned} \min_{\mathbf{w} \in \mathbb{R}^{N_s}} \quad & \text{tr}[\mathbf{\Gamma}_{\text{post}}(\mathbf{w})] + \gamma\Phi(\mathbf{w}), \\ \text{subject to} \quad & 0 \leq w_i \leq 1, \quad i = 1, \dots, N_s, \end{aligned} \tag{3.4}$$

where  $\Phi : \mathbb{R}_+^{N_s} \rightarrow [0, \infty)$  is a penalty function and  $\gamma \geq 0$  controls the sparsity of the design. We note that the function  $\mathbf{w} \mapsto \text{tr}[\mathbf{\Gamma}_{\text{post}}(\mathbf{w})] = \text{tr}[\mathbf{H}(\mathbf{w})^{-1}]$  is strictly convex due to strict convexity of  $\mathbf{X} \mapsto \text{tr}(\mathbf{X}^{-1})$  on the cone of symmetric positive definite matrices (see [28, p. 82]) and the fact that  $\mathbf{w}$  enters linearly in  $\mathbf{H}(\mathbf{w})$ . Therefore, if the penalty function  $\Phi$  is convex, (3.4) has a unique solution. An example of a convex penalty is given by  $\Phi(\mathbf{w}) = \mathbf{1}^T \mathbf{w}$ , i.e., an  $\ell_1$ -penalty, whose sparsening property has been used extensively in compressive sensing [9, 13] and has also been adapted to OED for inverse problems [16, 18]. In this paper, we will use a continuation approach involving a family  $\{\Phi_\varepsilon\}_{\varepsilon>0}$  of penalty functions, which approximate the  $\ell_0$ -“norm”. This allows us to find binary optimal design vectors.

We remark that there exists an alternative interpretation of the A-optimal design criterion in the Gaussian linear case considered here. Namely, minimizing the trace of the posterior covariance is equivalent to minimizing the average mean square error (MSE) of the posterior mean, where the average is with respect to the prior measure. This average MSE is also referred to as the Bayes risk of the posterior mean. MSE is a concept in frequentist inference, in which the posterior mean is interpreted as an estimator for the unknown parameter. This frequentist point of view of A-optimal design is used in [16, 18]. For completeness of our presentation, we detail this relation between average MSE and the trace of posterior covariance in Appendix B.

**4. Numerical solution of the OED problem (3.4).** We begin this section with deriving a decomposition of the misfit Hessian in terms of contributions from different sensors. Then, in Section 4.2, we present an approximation of  $\text{tr}(\mathbf{\Gamma}_{\text{post}})$  using a randomized trace estimator, and derive expressions for the gradient of the resulting OED objective function. Subsequently, in Sections 4.3 and 4.4, we present algorithmic components that allow efficient realization of these computations. Finally, in Section 4.5, we discuss methods to control the sparsity of the design, i.e., the number of allocated sensors.

**4.1. Decomposition of the misfit Hessian.** The misfit Hessian plays an important role in the derivative computation of the OED objective function with respect to the sensor location weights. Therefore, we first derive a decomposition of  $\mathbf{H}_{\text{misfit}}$  as a weighted sum of terms corresponding to individual sensor locations.

We consider a linear parameter-to-observable map  $\mathbf{F}$ , which involves a time-dependent PDE. Parameter-to-observable maps with stationary equations are included as special case in the discussion below by considering a single time step. We consider observations at the candidate sensor locations  $\mathbf{x}_1, \dots, \mathbf{x}_{N_s}$  in  $\mathcal{D}$ . For each sensor location, the time evolution of the observation is discretized using Lagrange elements in time (we use piecewise linear elements in this paper) for the nodal time instances  $\tau_1, \dots, \tau_{N_\tau}$  ( $\tau_i \in [0, T]$  for  $i = 1, \dots, N_\tau$ ). These observation times are independent of the time steps used in the integration of the PDE.

The parameter-to-observable map  $\mathbf{F}$  takes a parameter vector  $\mathbf{m} \in \mathbb{R}^n$  and maps it to the space-time observation vector  $\bar{\mathbf{d}} \in \mathbb{R}^{N_s N_\tau}$ :

$$\mathbf{F} : \mathbf{m} \xrightarrow{\mathbf{S}} \bar{\mathbf{u}} \xrightarrow{\mathbf{B}} \bar{\mathbf{d}}. \quad (4.1)$$

Here,  $\mathbf{S}$  is the discretized PDE solution operator,  $\bar{\mathbf{u}} \in \mathbb{R}^{n(N_t+1)}$  is the space-time solution vector, and  $\mathbf{B}$  is the space-time observation operator. We target designs where the sensor locations coincide for all time observations. Thus, the sensor weight matrix  $\mathbf{W}$  can be written as  $\mathbf{W} = \sum_{j=1}^{N_s} w_j \mathbf{E}_j$ , where  $\mathbf{E}_j$  is an  $N_s N_\tau \times N_s N_\tau$  block-diagonal matrix, with  $N_\tau$  blocks, with each block equal to  $\mathbf{e}_j \otimes \mathbf{e}_j = \mathbf{e}_j \mathbf{e}_j^T$ ; here  $\mathbf{e}_j$  denotes the  $j^{\text{th}}$  coordinate vector in  $\mathbb{R}^{N_s}$ . Thus, the misfit Hessian for a weight vector  $\mathbf{w}$  can be decomposed as  $\mathbf{H}_{\text{misfit}}(\mathbf{w}) = \mathbf{F}^* \mathbf{W} \mathbf{F} = \sum_{j=1}^{N_s} w_j \mathbf{F}^* \mathbf{E}_j \mathbf{F}$ , where the matrices  $\mathbf{F}^* \mathbf{E}_j \mathbf{F}$ ,  $j = 1, \dots, N_s$  are the *atoms* corresponding to the different sensor locations. This decomposition, which is also used in [33], reveals the identity

$$\frac{\partial \mathbf{H}_{\text{misfit}}(\mathbf{w})}{\partial w_j} = \mathbf{F}^* \mathbf{E}_j \mathbf{F}. \quad (4.2)$$

Next, we approximate the OED objective function using trace estimators.

**4.2. The OED objective function and its derivative.** We consider the problem (3.4) and recall that  $\Gamma_{\text{post}}(\mathbf{w}) = \mathbf{H}(\mathbf{w})^{-1}$ , where the Hessian is an  $\mathbf{M}$ -symmetric linear mapping on  $\mathbb{R}_M^n$ . For the numerical solution of the OED problem, we approximate the trace of  $\mathbf{H}(\mathbf{w})^{-1}$  using a randomized trace estimator (see Section 2.3). The trace estimator-based OED objective functional is

$$\Theta(\mathbf{w}) := \frac{1}{N_{\text{tr}}} \sum_{i=1}^{N_{\text{tr}}} \left\langle \mathbf{z}^{(i)}, \mathbf{H}(\mathbf{w})^{-1} \mathbf{z}^{(i)} \right\rangle_{\mathbf{M}}, \quad (4.3)$$

where  $\mathbf{z}^{(i)} = \mathbf{M}^{-1/2} \mathbf{y}^{(i)}$ ,  $i = 1, \dots, N_{\text{tr}}$ , with  $\mathbf{y}^{(i)}$  appropriately chosen random vectors<sup>1</sup>. Therefore, we consider the following OED optimization problem in our numerical computations,

$$\begin{aligned} \min_{\mathbf{w} \in \mathbb{R}^{N_s}} \quad & \Theta(\mathbf{w}) + \gamma \Phi(\mathbf{w}), \\ \text{subject to} \quad & 0 \leq w_i \leq 1, \quad i = 1, \dots, N_s. \end{aligned} \quad (4.4)$$

Since we will use gradient-based optimization methods to solve (4.4), we need to compute the gradient of  $\Theta(\mathbf{w})$  with respect to  $\mathbf{w}$ . For  $j \in \{1, \dots, N_s\}$ , we obtain

$$\begin{aligned} \frac{\partial \Theta(\mathbf{w})}{\partial w_j} &= \frac{1}{N_{\text{tr}}} \sum_{i=1}^{N_{\text{tr}}} \left\langle \mathbf{z}^{(i)}, \frac{\partial \mathbf{H}(\mathbf{w})^{-1}}{\partial w_j} \mathbf{z}^{(i)} \right\rangle_{\mathbf{M}} \\ &= -\frac{1}{N_{\text{tr}}} \sum_{i=1}^{N_{\text{tr}}} \left\langle \mathbf{z}^{(i)}, \mathbf{H}(\mathbf{w})^{-1} \frac{\partial \mathbf{H}(\mathbf{w})}{\partial w_j} \mathbf{H}(\mathbf{w})^{-1} \mathbf{z}^{(i)} \right\rangle_{\mathbf{M}}. \end{aligned}$$

The  $\mathbf{M}$ -symmetry of  $\mathbf{H}(\mathbf{w})^{-1}$ , the fact that the prior does not depend on  $\mathbf{w}$ , and denoting  $\mathbf{q}^{(i)} = \mathbf{H}^{-1} \mathbf{z}^{(i)}$  yields

$$\frac{\partial \Theta(\mathbf{w})}{\partial w_j} = -\frac{1}{N_{\text{tr}}} \sum_{i=1}^{N_{\text{tr}}} \left\langle \mathbf{q}^{(i)}, \frac{\partial \mathbf{H}_{\text{misfit}}(\mathbf{w})}{\partial w_j} \mathbf{q}^{(i)} \right\rangle_{\mathbf{M}}, \quad j = 1, \dots, N_s. \quad (4.5)$$

<sup>1</sup>In the present work, we rely on Gaussian trace estimators. See Appendix A for a justification of trace estimation in the context of weighted inner products.

Using (4.2) and denoting  $\tilde{\mathbf{d}}^{(i)} = (\mathbf{d}^{(i),1}, \dots, \mathbf{d}^{(i),N_\tau})^T := \mathbf{F}\mathbf{q}^{(i)}$ , with  $\mathbf{d}^{(i)} \in \mathbb{R}^{N_s}$  the spatial observations corresponding to time  $\tau_i$ , each of the summands in (4.5) can be computed as

$$\begin{aligned} \left\langle \mathbf{q}^{(i)}, \frac{\partial \mathbf{H}_{\text{misfit}}(\mathbf{w})}{\partial w_j} \mathbf{q}^{(i)} \right\rangle_{\mathbf{M}} &= \left\langle \mathbf{F}\mathbf{q}^{(i)}, \mathbf{E}_j \mathbf{F}\mathbf{q}^{(i)} \right\rangle \\ &= \sum_{\ell=1}^{N_\tau} [\mathbf{d}^{(i),\ell}]^T (\mathbf{e}_j \otimes \mathbf{e}_j) [\mathbf{d}^{(i),\ell}] = \sum_{\ell=1}^{N_\tau} [d_j^{(i),\ell}]^2. \end{aligned} \quad (4.6)$$

To summarize, the evaluation of  $\Theta(\mathbf{w})$  requires  $N_{\text{tr}}$  multiplications of vectors with  $\mathbf{H}(\mathbf{w})^{-1}$ . The computation of the derivative of  $\Theta(\mathbf{w})$  with respect to  $\mathbf{w}$  additionally requires  $N_{\text{tr}}$  evaluations of the parameter-to-observable map  $\mathbf{F}$ . In the Sections 4.3 and 4.4 below, we discuss the efficient realization of these computations.

**4.3. Low-rank approximation of the prior-preconditioned parameter-to-observable map.** As shown above, the repeated application of  $\mathbf{H}(\mathbf{w})^{-1}$  is necessary to compute  $\Theta(\mathbf{w})$  and the gradient of  $\Theta(\mathbf{w})$  with respect to  $\mathbf{w}$ . In a numerical optimization algorithm to solve (4.4), these computations are required in each iteration. Despite the use of a trace estimator, this is computationally demanding and can render OED for large-scale Bayesian inverse problems infeasible.

As a remedy, we construct a surrogate model for the parameter-to-observable map that can be used to efficiently compute the application of  $\mathbf{H}(\mathbf{w})^{-1}$  to vectors. The surrogate construction exploits the fact that for a large class of infinite-dimensional inverse problems,  $\mathbf{F}$  can be well approximated by a low-rank operator due to properties of the underlying PDE and the limited number of observations; see for instance [8,14]. Thus, we can compute a low-rank SVD surrogate for  $\mathbf{F}$  *upfront*, and use this surrogate in all subsequent computations involving  $\mathbf{F}$  to find an optimal experimental design.

Since only parameters consistent with the data *and* the prior have a significant posterior probability, only those parameters can influence the OED. Thus, it suffices to compute a surrogate of the *prior-preconditioned* parameter-to-observable map,  $\tilde{\mathbf{F}} := \mathbf{F}\mathbf{\Gamma}_{\text{prior}}^{1/2}$ . The smoothing property of the priors usually employed in infinite-dimensional Bayesian inversion<sup>2</sup> results in  $\tilde{\mathbf{F}}$  having faster decaying singular values than  $\mathbf{F}$ . Thus, the number of forward and adjoint PDE solves required to construct an accurate low-rank SVD surrogate for  $\tilde{\mathbf{F}}$  is usually smaller than that for  $\mathbf{F}$ . For the model problem in Section 6.1.1, this results in a significant speedup for the construction of the low-rank surrogate.

We employ randomized SVD (see Section 2.3) to compute the low-rank surrogate of  $\tilde{\mathbf{F}}$ . This algorithm only requires the applications of  $\tilde{\mathbf{F}}$  and  $\tilde{\mathbf{F}}^*$  to a set of independent random vectors, which is convenient and can often be implemented efficiently in large-scale computations. In the next section, we show how the low-rank surrogate of  $\tilde{\mathbf{F}}$  can be used to compute  $\Theta(\mathbf{w})$  and its gradient with respect to  $\mathbf{w}$ .

**4.4. Efficient computation of  $\Theta(\mathbf{w})$  and its derivatives.** First, we discuss the efficient computation of  $\mathbf{H}(\mathbf{w})^{-1}\mathbf{z}$  for  $\mathbf{z} \in \mathbb{R}^n$ , as required to evaluate  $\Theta(\mathbf{w})$  and its derivative. Note that the inverse of the Hessian  $\mathbf{H}(\mathbf{w})$  can be written as [8,25],

$$\mathbf{H}(\mathbf{w})^{-1} = \mathbf{\Gamma}_{\text{prior}}^{1/2} (\tilde{\mathbf{H}}_{\text{misfit}}(\mathbf{w}) + \mathbf{I})^{-1} \mathbf{\Gamma}_{\text{prior}}^{1/2}, \quad (4.7)$$

<sup>2</sup>Note that a smoothing property is necessary to render the infinite-dimensional Bayesian inverse problem well-posed, [32].

where, for a weight vector  $\mathbf{w} \in \mathbb{R}^{N_s}$ ,  $\tilde{\mathbf{H}}_{\text{misfit}}(\mathbf{w}) = \tilde{\mathbf{F}}^* \mathbf{W} \tilde{\mathbf{F}}$  is the *prior-preconditioned* misfit Hessian. Assuming a rank- $r$  surrogate  $\tilde{\mathbf{F}}_r$  for  $\tilde{\mathbf{F}}$  is available,  $\tilde{\mathbf{H}}_{\text{misfit},r}(\mathbf{w}) = \tilde{\mathbf{F}}_r^* \mathbf{W} \tilde{\mathbf{F}}_r$  is the resulting approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}(\mathbf{w})$  used below. Note that for every weight vector  $\mathbf{w}$ , the rank of  $\tilde{\mathbf{H}}_{\text{misfit},r}(\mathbf{w})$  is less than or equal to  $r$ . One approach to apply  $(\tilde{\mathbf{H}}_{\text{misfit},r} + \mathbf{I})^{-1}$  to a vector is to first compute the spectral decomposition  $\tilde{\mathbf{H}}_{\text{misfit},r}(\mathbf{w}) = \mathbf{V}_r \mathbf{\Lambda}_r \mathbf{V}_r^*$ , where  $\mathbf{V}_r$  is the matrix containing the eigenvectors of  $\tilde{\mathbf{H}}_{\text{misfit},r}(\mathbf{w})$  as its columns, and  $\mathbf{\Lambda}_r$  is the diagonal matrix with  $r$  largest eigenvalues on the diagonal.<sup>3</sup> Then, by the Sherman-Morrison-Woodbury formula [15], we have  $(\mathbf{V}_r \mathbf{\Lambda}_r \mathbf{V}_r^* + \mathbf{I})^{-1} = \mathbf{I} - \mathbf{V}_r \mathbf{D}_r \mathbf{V}_r^*$ , where  $\mathbf{D}_r$  is a diagonal matrix with  $\lambda_i/(1 + \lambda_i)$ ,  $i = 1, \dots, r$ , on its diagonal.<sup>4</sup> Defining

$$\hat{\mathbf{q}} = (\mathbf{I} - \mathbf{V}_r \mathbf{D}_r \mathbf{V}_r^*) \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{z} \quad (4.8)$$

allows us to compute

$$\mathbf{q} = \mathbf{H}(\mathbf{w})^{-1} \mathbf{z} \approx \mathbf{\Gamma}_{\text{prior}}^{1/2} \hat{\mathbf{q}}, \quad (4.9)$$

where the approximation is only due to the use of  $\tilde{\mathbf{F}}_r$  instead of  $\tilde{\mathbf{F}}$ . To compute  $\mathbf{F}\mathbf{q}$ , as needed in (4.6), we use

$$\mathbf{F}\mathbf{q} = \mathbf{F}\mathbf{\Gamma}_{\text{prior}}^{1/2} (\tilde{\mathbf{H}}_{\text{misfit}}(\mathbf{w}) + \mathbf{I})^{-1} \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{z} \approx \tilde{\mathbf{F}}_r \hat{\mathbf{q}}, \quad (4.10)$$

where the approximation is again due to the use of the low-rank surrogate for  $\tilde{\mathbf{F}}$ . We now summarize the procedure for computing  $\Theta(\mathbf{w})$  and its gradient. Note that once  $\tilde{\mathbf{F}}_r$  is available, our method does not require further forward or adjoint PDE solves.

We close this subsection with two remarks concerning the computations involved in Algorithm 1. First, note that while the low-rank SVD surrogate for  $\tilde{\mathbf{F}}$  relieves the OED method of repeated PDE solves, we still require the repeated applications of  $\mathbf{\Gamma}_{\text{prior}}^{1/2}$ . Due to our choice of the prior as the inverse of a squared elliptic operator, this amounts to an elliptic solve, for which optimal complexity solvers (e.g., multigrid) are available.

Next, note that we require the application of  $\mathbf{M}^{-1/2}$  to random vectors for the trace estimator in the OED method; see (4.3). The same operation, which transforms vectors from the Euclidean space  $\mathbb{R}^n$  to  $\mathbb{R}_{\mathbf{M}}^n$  (see Section 2.2) via a mapping that preserves the inner products, is also needed to draw samples from the posterior distribution (see (2.7)). Since the explicit computation of the matrix square root is expensive in high dimensions, we utilize an iterative algorithm to compute the application of the inverse matrix square root to vectors [11]. The method relies on the approximation of the square root via orthogonal polynomials on an interval containing the spectrum of the matrix; since the convergence of this method is fastest for matrices with clustered eigenvalues, we apply it to the mass matrix  $\mathbf{M}$ , symmetrically preconditioned with the lumped mass matrix,  $\mathbf{M}_l$ . The eigenvalues of the resulting matrix,  $\tilde{\mathbf{M}} = \mathbf{M}_l^{-1/2} \mathbf{M} \mathbf{M}_l^{-1/2}$ , are clustered around 1, resulting in fast convergence when applying  $\tilde{\mathbf{M}}^{-1/2}$  to vectors as discussed in [11]. Then, instead of  $\mathbf{M}^{-1/2}$  we

<sup>3</sup>Note that thanks to the low-rank surrogate of the prior-preconditioned parameter-to-observable map  $\tilde{\mathbf{F}}_r$ , this computation does not require the solution of PDEs.

<sup>4</sup>Note that there exist alternatives to computing a spectral decomposition of  $\tilde{\mathbf{H}}_{\text{misfit},r}(\mathbf{w})$  and then using the Sherman-Morrison-Woodbury formula. For instance, one can use a Krylov method to compute  $\mathbf{H}(\mathbf{w})^{-1} \mathbf{z}$ .

---

**Algorithm 1** Algorithm for computing  $\Theta(\mathbf{w})$  and its gradient  $\mathbf{g} = \nabla\Theta(\mathbf{w})$  with random vectors  $\mathbf{z}^{(i)}$ ,  $i = 1, \dots, N_{\text{tr}}$  for the trace estimator.

---

Compute  $\tilde{\mathbf{H}}_{\text{misfit},r}(\mathbf{w}) = \mathbf{V}_r \mathbf{\Lambda}_r \mathbf{V}_r^*$       {low-rank approx to  $\tilde{\mathbf{H}}_{\text{misfit}}(\mathbf{w}) = \tilde{\mathbf{F}}^* \mathbf{W} \tilde{\mathbf{F}}$ }  
 Set  $\mathbf{D}_r = \text{diag}\left(\frac{\lambda_1}{1+\lambda_1}, \dots, \frac{\lambda_r}{1+\lambda_r}\right)$   
**for**  $i = 1$  to  $N_{\text{tr}}$  **do**  
   Compute  $\hat{\mathbf{q}}^{(i)} = \left(\mathbf{I} - \mathbf{V}_r \mathbf{D}_r \mathbf{V}_r^*\right) \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{z}^{(i)}$       {as in (4.8)}  
   Compute  $\mathbf{q}^{(i)} = \mathbf{\Gamma}_{\text{prior}}^{1/2} \hat{\mathbf{q}}^{(i)}$       {compute  $\mathbf{q}^{(i)} = \mathbf{H}^{-1} \mathbf{z}^{(i)}$  as in (4.9)}  
**end for**  
 Compute  $\Theta(\mathbf{w}) = \frac{1}{N_{\text{tr}}} \sum_{i=1}^{N_{\text{tr}}} \left\langle \mathbf{z}^{(i)}, \mathbf{q}^{(i)} \right\rangle_{\mathbf{M}}$       {evaluation of  $\Theta(\mathbf{w})$ }  
 $\mathbf{g} = \mathbf{0}$       {initialize the gradient vector}  
**for**  $i = 1$  to  $N_{\text{tr}}$  **do**  
    $\bar{\mathbf{d}} = \tilde{\mathbf{F}}_r \hat{\mathbf{q}}^{(i)}$       {compute  $\bar{\mathbf{d}} = \mathbf{F} \mathbf{q}^{(i)}$  as in (4.10)}  
   **for**  $j = 1$  to  $N_s$  **do**  
      $g_{i,j} = \sum_{\ell=1}^{N_r} (d_j^\ell)^2$       { $g_{i,j} = \left\langle \mathbf{q}^{(i)}, \frac{\partial \mathbf{H}_{\text{misfit}}}{\partial w_j} \mathbf{q}^{(i)} \right\rangle_{\mathbf{M}}$  as in (4.6)}  
   **end for**  
    $\mathbf{g} = \mathbf{g} - \frac{1}{N_{\text{tr}}} \mathbf{g}_i$   
**end for**

---

use the matrix  $\mathbf{L} = \mathbf{M}_l^{-1/2} \tilde{\mathbf{M}}^{-1/2}$  as the isomorphism between  $\mathbb{R}^n$  and  $\mathbb{R}_{\mathbf{M}}^n$ . Note, in particular, that for all  $\mathbf{x}$  and  $\mathbf{y}$  in  $\mathbb{R}^n$

$$\langle \mathbf{L}\mathbf{x}, \mathbf{L}\mathbf{y} \rangle_{\mathbf{M}} = \mathbf{x}^T \tilde{\mathbf{M}}^{-\frac{1}{2}} \mathbf{M}_l^{-\frac{1}{2}} \mathbf{M} \mathbf{M}_l^{-\frac{1}{2}} \tilde{\mathbf{M}}^{-\frac{1}{2}} \mathbf{y} = \mathbf{x}^T \tilde{\mathbf{M}}^{-\frac{1}{2}} \tilde{\mathbf{M}} \tilde{\mathbf{M}}^{-\frac{1}{2}} \mathbf{y} = \langle \mathbf{x}, \mathbf{y} \rangle,$$

which shows that  $\mathbf{L}$  is a Hilbert space isomorphism between  $\mathbb{R}^n$  and  $\mathbb{R}_{\mathbf{M}}^n$ .

**4.5. Sparsity-enforcing penalty functions.** Here, we discuss the penalty term  $\gamma\Phi(\cdot)$  in (3.4). If  $\Phi(\cdot)$  favors sparse solutions, the degree of sparsity depends on the choice of  $\gamma$ . However,  $\gamma$  only provides an indirect control of the sparsity of the design and in practice adjusting  $\gamma$  and resolving (3.4) might be necessary to obtain a design with a specified (or close to a specified) number of sensors.

Next, we discuss two possible choices for the penalty function  $\Phi(\cdot)$ . An  $\ell_1$ -penalty function  $\Phi(\mathbf{w}) = \mathbf{1}^T \mathbf{w}$  (see also [16–18]), is known to result in sparse solution vectors  $\mathbf{w}$ . While this choice is convenient due to the convexity of  $\Phi$ , it results in a weight vector  $\mathbf{w}$  whose components can take on any value in  $[0, 1]$ . Unfortunately, the direct interpretation of weights  $w_i \notin \{0, 1\}$  for the placement of sensors is unclear. A practical approach to use such a solution vector is to place sensors in locations where the corresponding weight does not vanish. An alternative to  $\ell_1$ -sparsification is to use non-convex functions for  $\Phi$  that directly lead to binary weight vectors. This approach is discussed next.

To obtain binary weight vectors, we select penalties that successively approximate the  $\ell_0$ -“norm”,<sup>5</sup> which, for  $\mathbf{x} \in \mathbb{R}^n$ ,  $\|\mathbf{x}\|_0$  is defined as the number of non-zero elements in  $\mathbf{x}$ . One choice for such a family of penalizations is  $\Phi_q(\mathbf{w}) = \|\mathbf{w}\|_q^q$ , for  $q < 1$ . Since

---

<sup>5</sup>Note that  $\|\mathbf{x}\|_0$  is in fact not a *norm*.

these penalization functions are not Lipschitz continuous at the origin, we use an alternative family of regularizations, namely, for  $\varepsilon > 0$ ,

$$\Phi_\varepsilon(\mathbf{w}) := \sum_{i=1}^{N_s} f_\varepsilon(w_i), \quad \text{where } f_\varepsilon(w) = \begin{cases} \frac{w}{\varepsilon}, & 0 \leq w \leq \frac{1}{2}\varepsilon, \\ p_\varepsilon(w), & \frac{1}{2}\varepsilon < w \leq 2\varepsilon, \\ 1, & 2\varepsilon < w \leq 1. \end{cases} \quad (4.11)$$

Here,  $p_\varepsilon(\cdot)$  is the uniquely defined third order polynomial that makes  $f_\varepsilon(x) : [0, 1] \rightarrow [0, 1]$  continuously differentiable. In Figure 4.1, we show these penalty functions for different values of  $\varepsilon$ ; note how  $f_\varepsilon(\cdot)$  approximates the  $\ell_0$ -“norm” as  $\varepsilon \rightarrow 0$ .




FIG. 4.1. Graphs of  $f_\varepsilon(\cdot)$  as defined in (4.11) for  $\varepsilon = 1/2^i$ ,  $i = 1, \dots, 4$ .

To cope with potential multiple local minima due to the non-convexity of the penalties  $\Phi_\varepsilon$ , we use a continuation procedure with respect to  $\varepsilon$ . We fix the regularization parameter  $\gamma$  in (3.4) and solve the OED problem with  $\ell_1$ -penalty, resulting in a weight vector  $\mathbf{w}^0$ . We then choose a decreasing sequence of positive real numbers  $\{\varepsilon_i\}_{i \geq 1}$ , and compute a new weight vector  $\mathbf{w}^i$  by solving (3.4) with penalty function  $\Phi_{\varepsilon_i}(\mathbf{w})$  using  $\mathbf{w}^{i-1}$  as initialization. In practice, once  $\varepsilon_i$  is sufficiently small, the solution  $\mathbf{w}^i$  is a 0–1 vector that remains unchanged as  $\varepsilon$  is further decreased.

A similar continuation strategy is used in topology optimization [4, 5], where the target is to design an “optimal structure” by deciding on an optimal distribution of material in a physical domain  $\mathcal{D}$ . In this application, one seeks a *density* function  $\rho : \mathcal{D} \rightarrow \{0, 1\}$  that characterizes the absence or presence of material in  $\mathcal{D}$ . One of the main approaches to solving topology optimization problems is to relax the requirement of  $\rho \in \{0, 1\}$  to  $0 \leq \rho \leq 1$  and to solve a sequence of optimization problems with successively steeper penalty functions to approach a 0–1 solution  $\rho$ —similar our  $\Phi_\varepsilon$ -continuation approach outlined above.

In the rest of this paper, we refer to the designs obtained via the above continuation method as  $\Phi_\varepsilon$ -sparsified designs. Optimal designs obtained by solving (3.4) with  $\ell_1$ -penalty (followed by thresholding to obtain a 0–1 weight vector) are referred to as  $\ell_1$ -sparsified designs. In the numerical experiments presented in Section 6, we compare  $\ell_1$ - and  $\Phi_\varepsilon$ -sparsified designs.

**5. Model problem setup.** Here, we present the model problem used to study the methods for OED presented in this paper. We consider a time-dependent advection-diffusion equation in which we seek to infer an unknown initial condition from point measurements; see also [1, 14, 29].

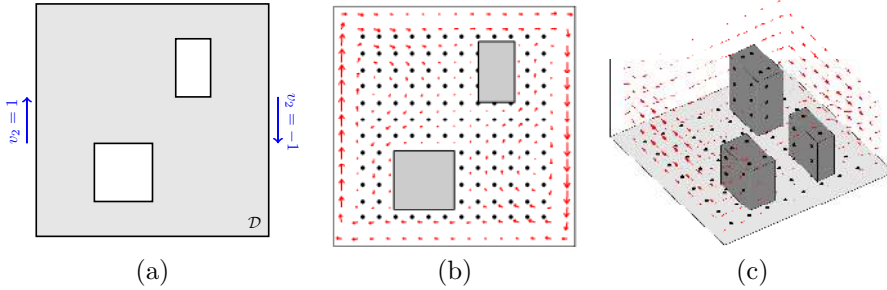


FIG. 5.1. (a) The computational domain  $\mathcal{D}$  for the two-dimensional model problem is  $[0, 1]^2$  with two rectangular regions (representing buildings) removed. (b) The velocity field  $\mathbf{v}$  (red arrows) and the candidate sensor locations (black dots) for the two-dimensional problem. (c) The computational domain  $\mathcal{D}$  for the three-dimensional model problem is  $[0, 1]^2 \times [0, 0.5]$  with the buildings (gray blocks) removed; shown are also the velocity field (red arrows) and the candidate sensor locations (black dots).

**5.1. The parameter-to-observable map.** The PDE in the parameter-to-observable map models diffusive transport in a domain  $\mathcal{D} \subset \mathbb{R}^d$ , which is depicted in Figure 5.1(a) for  $d = 2$  and 5.1(c) for  $d = 3$ . The domain boundaries  $\partial\mathcal{D}$  include the outer edges/faces as well as the internal boundaries of the rectangles/hexahedra, which model buildings. The parameter-to-observable map maps an initial condition  $m \in L^2(\mathcal{D})$  to spatial and temporal point observations through the solution of the advection-diffusion equation  $u(\mathbf{x}, t)$  given by

$$\begin{aligned} u_t - \kappa \Delta u + \mathbf{v} \cdot \nabla u &= 0 && \text{in } \mathcal{D} \times (0, T), \\ u(\cdot, 0) &= m && \text{in } \mathcal{D}, \\ \kappa \nabla u \cdot \mathbf{n} &= 0 && \text{on } \partial\mathcal{D} \times (0, T). \end{aligned} \quad (5.1)$$

Here,  $\kappa > 0$  is the diffusion coefficient and  $T > 0$  is the final time. In our numerical experiments, we use  $\kappa = 0.001$  for the two-dimensional problem and  $\kappa = 0.003$  for the three-dimensional problem. The velocity field  $\mathbf{v}$ , shown in Figures 5.1(b), (c) for the two- and the three-dimensional problems, respectively, is computed by solving the following steady-state Navier-Stokes equation with the side walls driving the flow:

$$\begin{aligned} -\frac{1}{\text{Re}} \Delta \mathbf{v} + \nabla q + \mathbf{v} \cdot \nabla \mathbf{v} &= 0 && \text{in } \mathcal{D}, \\ \nabla \cdot \mathbf{v} &= 0 && \text{in } \mathcal{D}, \\ \mathbf{v} &= \mathbf{g} && \text{on } \partial\mathcal{D}. \end{aligned} \quad (5.2)$$

Here,  $q$  is pressure,  $\text{Re}$  is the Reynolds number, set to  $\text{Re} = 50$  in the present examples. The Dirichlet boundary data  $\mathbf{g} \in \mathbb{R}^d$  is given by  $\mathbf{g} = \mathbf{e}_2$  on the left wall of the domain,  $\mathbf{g} = -\mathbf{e}_2$  on the right wall, and  $\mathbf{g} = \mathbf{0}$  everywhere else (see, e.g., Figure 5.1(a) for  $d = 2$ ). To illustrate the physics of the forward problem for a given initial condition, we show snapshots of the time evolution of the state variable  $u$  in Figure 5.2; for snapshots of the time evolution for the two-dimensional model problem we refer to [29].

The observation operator  $\mathcal{B}$  then extracts the values of  $u$  on a set of sensor locations  $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_s}\} \subset \mathcal{D}$  at times  $\{\tau_1, \dots, \tau_{N_\tau}\} \subset [0, T]$ . To summarize, the infinite-dimensional linear parameter-to-observable map  $\mathcal{F}$  maps the initial condition  $m$  to  $q = N_s N_\tau$  observations as follows: First, we solve the time-dependent advection-diffusion equation (5.1), resulting in the space-time solution  $u = u(m)$ . Then, the






FIG. 5.2. Volume renderings of the time evolution of  $u$  for the three-dimensional model problem. The initial condition  $m$  is shown in the left, and the middle and right images correspond to snapshots at times  $t = 1$  and  $t = 2$ .

observation operator extracts point values from  $u$  at the measurement locations and times, i.e., evaluates  $\mathcal{B}u$ . The corresponding discrete parameter-to-observable map  $\mathbf{F}$  is obtained by discretizing  $m$  and  $u$  using, for instance, finite elements.

**5.2. The Bayesian inverse problem.** Following the Bayesian setup in Section 2.1, we utilize a Gaussian prior measure  $\mu_0 = \mathcal{N}(m_0, \mathcal{C}_0)$ , with  $\mathcal{C}_0 = \mathcal{A}^{-2}$  where  $\mathcal{A}$  is an elliptic differential operator as described in Section 2.1, and use an additive Gaussian noise model. Therefore, the solution of the Bayesian inverse problem is the posterior measure,  $\mu_{\text{post}} = \mathcal{N}(m_{\text{post}}, \mathcal{C}_{\text{post}})$  with  $m_{\text{post}}$  and  $\mathcal{C}_{\text{post}}$  as given in (2.5). The posterior mean  $m_{\text{post}}$  is characterized as the minimizer of

$$\mathcal{J}(m) := \frac{1}{2} \|\mathcal{B}u(m) - \mathbf{d}\|_{\Gamma_{\text{noise}}^{-1}}^2 + \frac{1}{2} \|\mathcal{A}(m - m_0)\|_{L^2(\mathcal{D})}^2,$$

which can also be interpreted as the regularized functional to be minimized in deterministic inversion. Next, we specify the action of the adjoint  $\mathcal{F}^*$  of the parameter-to-observable map. Given an observation vector  $\bar{\mathbf{d}} \in \mathbb{R}^q$  as defined in (4.1),  $\mathcal{F}^*\bar{\mathbf{d}}$  is computed by solving the *adjoint equation* (see [1, 14, 29]) for the adjoint variable  $p = p(\mathbf{x}, t)$ ,

$$\begin{aligned} -p_t - \nabla \cdot (p\mathbf{v}) - \kappa\Delta p &= -\mathcal{B}^*\bar{\mathbf{d}} && \text{in } \mathcal{D} \times (0, T), \\ p(\cdot, T) &= 0 && \text{in } \mathcal{D}, \\ (\mathbf{v}p + \kappa\nabla p) \cdot \mathbf{n} &= 0 && \text{on } \partial\mathcal{D} \times (0, T), \end{aligned} \quad (5.3)$$

and setting  $\mathcal{F}^*\bar{\mathbf{d}} = p(\cdot, 0)$ . Note that (5.3) is a final value problem, which has to be solved backwards in time.

**5.3. Discretization and solvers.** The discretization of the forward and adjoint problems uses linear triangular/tetrahedral continuous Galerkin finite elements for two/three space dimensions, and the implicit Euler method in time. The discrete adjoint equation is obtained as the adjoint of the discretized forward equation, i.e., we follow a discretize-then-optimize approach. Due to the large diffusion parameter  $\kappa$  used, no stabilization of the advection-diffusion equation is needed. A factorization of the matrix in each implicit Euler time step is computed upfront, such that every time step (of the forward as well as the adjoint equation) only requires triangular solves.

The OED optimization problems are solved using MATLAB's built-in interior-point solver `fmincon`, to which we supply the OED objective function evaluation and its derivative, which we compute using the methods presented above. The optimization algorithm approximates second derivatives using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) [27] method.

**6. Numerical results.** In this section, we summarize our numerical experiments. In Section 6.1, we provide a detailed numerical study of our method and compare different sparsifications for the two-dimensional model problem. Then, in Section 6.2, we use the three-dimensional model problem to test the scalability and performance of our OED method.

**6.1. The two-dimensional model problem.** This section contains numerical experiments for the two-dimensional model problem. We use  $N_s = 122$  candidate locations for placing sensors as shown in Figure 5.1(b), and use 1864 triangles to discretize the domain. Linear finite elements are used for the parameter  $m$  and the state variable  $u$  resulting in  $n = 1012$  spatial degrees of freedom (and thus inversion parameters). The final time is  $T = 4$  and the time interval is discretized using  $N_t + 1$  implicit Euler time steps; unless otherwise specified,  $N_t = 64$ . Observations at the sensor locations are taken at  $N_\tau = 19$  equally spaced points in the time interval  $[1, 4]$ . The prior covariance is specified as described in Section 2.1 with  $\alpha = 8 \times 10^{-3}$  and  $\beta = 10^{-2}$ ; these values result in a Gaussian prior that is not overly restrictive or smoothing. Next, we study the low-rank approximations of the forward operator and the misfit Hessian.

**6.1.1. Low-rank approximation of  $\mathbf{F}$ ,  $\tilde{\mathbf{F}}$  and  $\mathbf{H}_{\text{misfit}}$ .** We use a randomized SVD to compute a low-rank surrogate for the parameter-to-observable map  $\mathbf{F} : \mathbb{R}^n \rightarrow \mathbb{R}^q$  and its prior-preconditioned counterpart  $\tilde{\mathbf{F}}$ , where  $q = N_s \times N_\tau$ . This directly leads to a low-rank approximation of the prior-preconditioned misfit Hessian,  $\tilde{\mathbf{H}}_{\text{misfit}}(\mathbf{w}) = \tilde{\mathbf{F}}^* \mathbf{W} \tilde{\mathbf{F}}$ , where the diagonal weight matrix  $\mathbf{W}$  depends on the design. Since the diagonal entries of  $\mathbf{W}$  may vanish, the numerical rank of  $\tilde{\mathbf{H}}_{\text{misfit}}(\mathbf{w})$  is less than or equal than that of  $\tilde{\mathbf{F}}^* \tilde{\mathbf{F}}$ , which corresponds to assigning unit weights to all sensors; thus, it suffices to examine the accuracy of the approximation of  $\tilde{\mathbf{H}}_{\text{misfit}} = \tilde{\mathbf{F}}^* \tilde{\mathbf{F}}$  only.

First, we study the low-rank approximation of the (prior-preconditioned) parameter-to-observable map. In Figure 6.1 (left), we show the normalized singular values of  $\tilde{\mathbf{F}}$  and of  $\mathbf{F}$ . The eigenvalues decay rapidly and thus both mappings can be accurately approximated by low-rank matrices. Note the faster decay of the singular values of  $\tilde{\mathbf{F}}$  compared to those of  $\mathbf{F}$ , which shows that due to the preconditioning with  $\mathbf{\Gamma}_{\text{prior}}^{1/2}$ ,  $\tilde{\mathbf{F}}$  can be approximated more efficiently with a low-rank operator than  $\mathbf{F}$ .

Next, we investigate the influence of spatial and temporal discretization on the prior-preconditioned misfit Hessian; the right plot in Figure 6.1 shows the eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  for different spatial and temporal discretizations. Note that the spectra lie almost on top of each other. This indicates that the discretizations  $\tilde{\mathbf{H}}_{\text{misfit}}$  converge to their infinite-dimensional counterpart,  $\tilde{\mathcal{H}}_{\text{misfit}} : L^2(\mathcal{D}) \rightarrow L^2(\mathcal{D})$  and that all discretizations studied in Figure 6.1 resolve the dominant physics of the problem.

Finally, we study how the number of candidate sensor locations affects the spectral decay of  $\tilde{\mathbf{F}}$ . Increasing the number of sensors has the potential to increase the information in the data and thus the numerical rank of  $\tilde{\mathbf{F}}$ . In Figure 6.2, we plot the singular values of  $\tilde{\mathbf{F}}$  for different sensor grids. We observe convergence of the singular value curves as the sensor grids are refined. This is a consequence of the correlation of the information that can be gained from neighboring sensors, which is due to the diffusion term in (5.1).

**6.1.2. A-optimal designs with  $\ell_1$ -sparsification.** We now apply our method to compute  $\ell_1$ -sparsified A-optimal designs. The weights  $\mathbf{w}$  found with  $\gamma = 6 \times 10^1$  are shown in Figure 6.3 (left). As discussed in Section 4.5, sensors are then placed

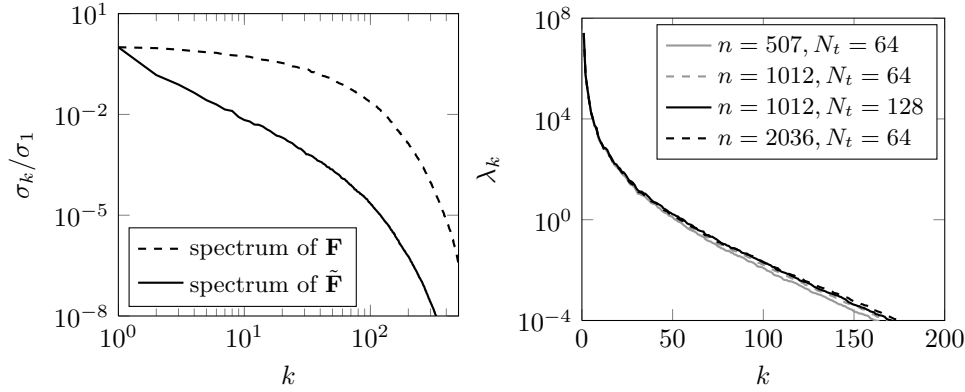


FIG. 6.1. Shown on the left are the normalized singular values of  $\tilde{\mathbf{F}}$  and  $\mathbf{F}$ . The right plot depicts the spectrum of the prior-preconditioned misfit Hessian with unit sensor weights  $\tilde{\mathbf{H}}_{\text{misfit}} = \tilde{\mathbf{F}}^* \tilde{\mathbf{F}}$  for discretizations with different numbers  $n$  of spatial and  $N_t$  of temporal unknowns. To compute these spectra, we used a rank  $r = 500$  approximation  $\tilde{\mathbf{F}}_r$  to diminish the influence of the low-rank approximation for  $\tilde{\mathbf{F}}$ .




FIG. 6.2. Singular values of  $\tilde{\mathbf{F}}$  for different values of  $N_s$ . The results correspond to a discretization of the problem with parameter dimension  $n = 1012$  and  $N_t = 2^6$  time steps.

at locations corresponding to non-vanishing weights. To account for numerical errors introduced by the interior point method used to solve (3.4), non-vanishing weights are defined as weights  $w_i$  that satisfy  $w_i / \sum_j w_j > 4 \times 10^{-3}$ ; the resulting sensor locations are shown in Figure 6.3 (right). Next, we study the behavior of our method for computing  $\ell_1$ -sparsified designs.

Since we rely on a low-rank surrogate  $\tilde{\mathbf{F}}_r$  for the prior-preconditioned parameter-to-observable map  $\tilde{\mathbf{F}}$  to avoid repeated PDE solves in the OED method, we first examine the influence of this truncation on the optimal design. For that purpose, we fix  $\gamma = 5 \times 10^1$  and solve the A-optimal design problem using approximations  $\tilde{\mathbf{F}}_r$  with  $r = 10, 15, 20, 30, 40, 60, 80$ . In Figure 6.4, we plot the optimal objective value  $\Theta(\mathbf{w}_{\text{opt}})$  computed using low-rank surrogates  $\tilde{\mathbf{F}}_r$  with different  $r$ . The convergence of the objective value shows that very similar results are found for low-rank surrogates  $\tilde{\mathbf{F}}_r$  with  $r \geq 40$ . This shows that even a significant compression of the (preconditioned) parameter-to-observable map has little influence on the optimal design.

Next, we study the effect of increasing the parameter dimension (through mesh refinements) on the number of interior-point iterations required to solve the OED




FIG. 6.3. Shown on the left are the weights found by solving the optimization problem (3.4) with  $\ell_1$ -sparsification. The right image shows the optimal sensor locations (black dots) obtained by placing sensors at locations with non-vanishing weights. Also shown is the advective flow field  $\mathbf{v}$  (red arrows).




FIG. 6.4. Left: Value of  $\Theta(\mathbf{w}_{opt})$  for an  $\ell_1$ -sparsified  $A$ -optimal design obtained with rank- $r$  approximations  $\tilde{\mathbf{F}}_r$ . Right: Number of interior-point quasi Newton iterations and function evaluations for different discretizations (and thus different parameter dimensions).

optimization problem (4.4). As observed in Section 6.1.1, increasing the parameter dimension does not increase the numerical rank of  $\tilde{\mathbf{F}}$  once the mesh is sufficiently fine. This, and the use of a Newton-type optimization method results in a nearly constant number of interior-point iterations required to solve (4.4), as shown in the table in Figure 6.4.

Finally, we study how the performance of the numerical optimization encountered in our OED method is affected by increasing the number  $N_s$  of candidate sensor locations; note that  $N_s$  is also the dimension of the weight vector  $\mathbf{w}$ , the unknown in the OED optimization problem (4.4). The results shown in Table 6.1 show that the number of interior-point quasi-Newton iterations and the number of evaluations of  $\Theta(\cdot)$  do not increase significantly as the number of sensor candidate locations increases. We again attribute this to the use of a Newton-type method, since Newton's method satisfies—under certain assumptions—a mesh independence property [12].

Regarding the complexity of our method with respect to the number  $N_s$  of candidate sensor locations, we draw the following conclusions from Figure 6.2 and Table 6.1: First, the number of PDE solves required to compute the low-rank surrogate  $\tilde{\mathbf{F}}_r$  is

TABLE 6.1

Number of interior-point BFGS iterations and evaluations of  $\Theta(\cdot)$  required for the solution of (4.4) with  $\ell_1$ -sparsification, for different numbers  $N_s$  of sensor candidate locations. The optimization iteration was terminated when the norm of the gradient was decreased by a factor of  $10^4$ .

| $N_s$       | 33  | 58  | 85  | 122 | 172 | 216 | 264 | 340 |
|-------------|-----|-----|-----|-----|-----|-----|-----|-----|
| Iterations  | 58  | 68  | 68  | 81  | 58  | 68  | 75  | 72  |
| Func. eval. | 149 | 165 | 141 | 158 | 136 | 137 | 122 | 160 |

bounded with respect to  $N_s$ , and second, the number of optimization iterations to solve (4.4) is insensitive to the number of candidate sensor locations, and thus the dimension of the weight vector  $\mathbf{w}$ .

**6.1.3. A-optimal designs with  $\Phi_\varepsilon$ -sparsification.** Next, we present  $\Phi_\varepsilon$ -sparsified designs obtained by a continuation procedure with respect to  $\varepsilon$ . Optimal designs for  $\gamma = 0.05$  and  $\gamma = 0.06$  are shown on the left in Figure 6.5. Note that the optimal sensor locations obtained with  $\gamma = 0.06$  is not a subset of the locations obtained with  $\gamma = 0.05$ , i.e., the designs are not nested.

For the continuation procedure, we use the values  $\varepsilon_i = (2/3)^i$ ,  $i = 1, \dots, 10$ . To elucidate the convergence of the weights to the desired binary structure, the weights corresponding to several values of  $\varepsilon_i$  are shown on the right in Figure 6.5. Note that the behavior of the weights is not monotone as  $\varepsilon$  is decreased. A possible explanation for this behavior is that, as  $\varepsilon$  decreases, weights at neighboring locations can get merged into a single weight at a new location.

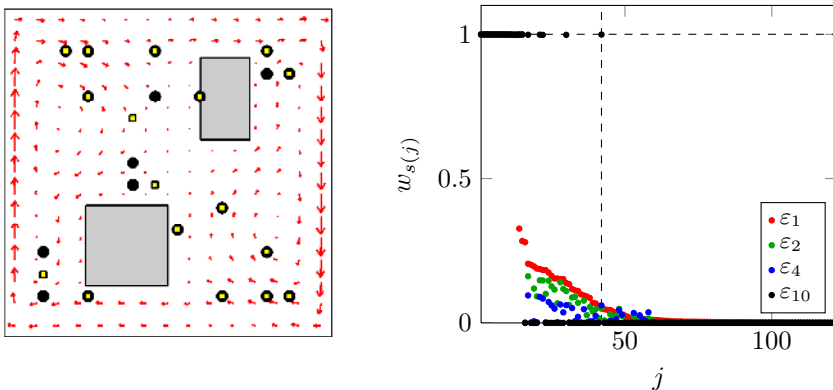


FIG. 6.5. Left: Sensor locations using  $\Phi_\varepsilon$ -sparsified A-optimal design for  $\gamma = 0.05$  (black dots) and for  $\gamma = 0.06$  (yellow squares). Right: Convergence of the weights to a 0-1 structure as  $\varepsilon \rightarrow 0$ . The weights are ordered based on the values obtained with  $\varepsilon_1$ , and shown is the evolution of the weights as  $\varepsilon$  decreases. Here,  $s(j)$  denotes the descending ordering obtained with  $\varepsilon_1$ .

To illustrate the decreased variance for A-optimal designs, in Figure 6.6 we compare the pointwise posterior standard deviation obtained with  $\gamma = 0.05$  with a uniform and two randomly generated designs employing the same number of sensors. This comparison shows that the A-optimal design provides a notable reduction in the standard deviation (and thus the variance) fields. Next, we compare the performance of sensor placements obtained with  $\ell_1$ - and  $\Phi_\varepsilon$ -sparsification.

**6.1.4. Comparing  $\ell_1$ - and  $\Phi_\varepsilon$ -sparsifications.** We observe that  $\ell_1$ - and  $\Phi_\varepsilon$ -sparsifications lead to different designs for the same number of sensors. Thus, a




FIG. 6.6. Shown is the pointwise standard deviation of the posterior for a  $\Phi_\varepsilon$ -sparsified A-optimal design with 20 sensors (a), for a manually chosen uniform design (b), and two randomly generated designs with the same number of sensors (c), (d). The white dots indicate the sensor locations. Red and blue correspond to regions with large and small standard deviation, respectively. We find that compared to the optimal design, the uniform design and the two random designs result in 7%, 36% and 26% increase in the average variance, respectively.

naturally arising question is (1) *which sparsification results in the better placement of sensors?* Moreover, (2) *how much do optimal designs improve over randomly chosen designs?* To answer these questions, we compute optimal designs based on different sparsification strategies and compare with randomly generated designs. We use trace estimators with 100 Gaussian random vectors chosen differently in each OED problem, but report the exact value of  $\text{tr}(\mathbf{\Gamma}_{\text{post}}(\mathbf{w})) = \text{tr}(\mathbf{H}(\mathbf{w})^{-1})$ .

Our results are summarized in Figure 6.8. We compute  $\ell_1$ -sparsified designs for various values of  $\gamma$  and report the value of  $\text{tr}(\mathbf{H}(\mathbf{w})^{-1})$ . Similarly, we report  $\text{tr}(\mathbf{H}(\mathbf{w})^{-1})$  for 0–1 optimal designs computed via  $\Phi_\varepsilon$ -sparsification. Additionally, we compute  $\text{tr}(\mathbf{H}(\mathbf{w})^{-1})$  for a collection of random sensor configurations. We find that  $\Phi_\varepsilon$ -sparsified designs consistently outperform  $\ell_1$ -sparsified designs and that both improve significantly over randomly chosen designs. Another observation from Figure 6.8 is the diminishing returns as the number of sensors is increased; using more than 20 sensors only results in negligible decrease of the OED objective function value.

**6.1.5. Influence of trace estimation.** To assess the accuracy of randomized trace estimation for a typical posterior covariance matrix, we compare the exact trace with results from randomized trace estimation with different numbers of random vectors. We find that trace estimators based on 1, 5, 10, 20, 100 vectors estimate the exact trace with an average error of about 15%, 7%, 5%, 2% and 1.5%, respectively. This is consistent with the experiments reported in [3] showing that randomized trace estimation is reasonably accurate with a small number of random vectors, but that many random vectors may be needed to obtain very accurate approximations.

A more relevant question in the context of OED is how the use of trace estimators to compute optimal sensor locations influences the resulting designs. To study this issue, we compute optimal designs with fixed  $\gamma$ , but different trace estimators. The results are shown in Figure 6.7. Note that trace estimation in the OED objective function does have an influence on the designs. More accurate trace estimation decreases the variation in both the number of sensors of the designs and the value of the trace of the posterior covariance. However, as can be seen from Figure 6.8 optimal designs computed using trace estimation (based on 100 random vectors) consistently improve over random designs with respect to reducing the *exact* trace of the posterior covariance. We thus conclude that the use of trace estimation does not have a significant impact on the quality of the optimal designs.




FIG. 6.7. Sensitivity of  $\Phi_\varepsilon$ -sparsified designs with respect to the trace estimator with fixed  $\gamma = 0.05$ . The left and the middle image visualize the frequency of a candidate location being part of the optimal design with a trace estimator based on a single random vector (left) and 100 random vectors (middle). The results are based on 30 different realizations of the respective trace estimators. The darker the candidate location, the more often the corresponding sensor was part of the optimal design. Note the decreased variation in the designs in the middle plot, which is due to the increased accuracy of the trace estimator. The right plot depicts the number of sensors and the value of the trace of the posterior covariance for 30 different realizations of trace estimators with a single random vector (empty dots) and with 100 random vectors (filled dots). Note that the number of sensors varies for the different realizations and that the more accurate trace estimator results in less scattering with respect to the number of sensors and the value of the trace.




FIG. 6.8. Comparison of performance for different designs. Shown is the value of  $\text{tr}(\mathbf{H}^{-1})$  versus the number of sensors for random designs (red dots),  $\ell_1$ -sparsified designs (cyan dots) and  $\Phi_\varepsilon$ -sparsified designs (black dots).

**6.2. The three-dimensional model problem.** The main target of this three dimensional model problem is to study the applicability of our OED method to large-scale inverse problems. The computational domain used is depicted in Figure 5.1(c), where 101 candidate sensor locations are shown as black dots. Note that we allow sensors on the ground and on the sides of the buildings. Observations are collected at  $N_\tau = 12$  equally spaced points in the time interval  $[1, T]$ , where  $T = 4$ . The parameter and the state variable  $u$  in the advection-diffusion equation are discretized using linear finite elements on a tetrahedral mesh with 10,652 spatial degrees of freedom, and 64 implicit Euler time steps are used for the time integration. Thus, the dimension of the inversion parameter is  $n = 10,652$ .






FIG. 6.9. *Top row: Sensor locations (black dots) obtained with A-optimal design with  $\Phi_\varepsilon$ -sparsification (left) and with random design (right). Bottom row: Standard deviation fields of the posterior for an A-optimal design with  $\Phi_\varepsilon$ -sparsification (left) and random design (right). Here, red indicates regions with large standard deviation, and blue corresponds to regions with small standard deviation.*

In our computations, we found a rank-100 SVD surrogate for the prior-preconditioned parameter-to-observable map adequate since the 100th singular value of  $\tilde{\mathbf{F}}$  is about four orders of magnitude smaller than the largest singular value. This translates to eight orders of magnitude difference between the first and the 100th eigenvalue of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . Due to the large parameter space dimension we do not compute  $\mathbf{M}^{-1/2}$  explicitly, but use the algorithm for the iterative application of  $\mathbf{M}^{-1/2}$  to vectors from [11]; see also Section 4.4. We observe that the algorithm is well converged after 10 applications of  $\tilde{\mathbf{M}}$ —the relative difference between the iterative application of  $\tilde{\mathbf{M}}^{-1/2}$  based on 10 applications of  $\tilde{\mathbf{M}}$  compared with that based on 500 applications of  $\tilde{\mathbf{M}}$  is approximately  $8 \times 10^{-6}$ . For the solution of the OED optimization problem with  $\gamma = 1.2 \times 10^{-1}$  and  $\Phi_\varepsilon$ -sparsification, the continuation procedure presented in Section 6.1.2 is used. To compute the initial  $\ell_1$ -sparsified weight vector  $\mathbf{w}^0$ , 18 interior-point iterations were necessary. This initialization was followed by 10 continuation steps using the penalties  $\Phi_{\varepsilon_i}$ , with  $\varepsilon_i = (2/3)^i$ ; each step required the solution of an auxiliary OED problem, which amounted to a total of 554 interior-point quasi-Newton iterations to arrive at a 0–1 design vector. The iteration for each of these auxiliary OED problems was terminated when the relative residual dropped below  $10^{-4}$  or a maximum of 150 iterations was reached. The number of quasi-Newton iterations in the continuation procedure can likely be decreased by incorporating the continuation procedure into the interior-point algorithm rather than solving an independent interior point problem for each  $\varepsilon_i$ .

The resulting optimal design is shown in top left image of Figure 6.9; note that several sensors are placed on the top of buildings; these sensors help reduce the variance in the upper parts of the domain. To illustrate the effectiveness of the optimal design to reduce the variance, in Figure 6.9 we show a volume rendering of the point-

wise posterior standard deviation obtained with the optimal design, and compare it with a rendering obtained with a randomly generated design that has the same number of sensors.

**7. Concluding remarks.** We have developed a structure-exploiting efficient numerical method for computing  $\mathbf{A}$ -optimal experimental designs in infinite-dimensional Bayesian linear inverse problems governed by PDEs. Numerical experiments for the inversion of the initial condition in an advection-diffusion equation indicate that an optimal design can be computed at a cost, measured in forward PDE solves, that is independent of the parameter and candidate sensor dimensions. Moreover, we find that experimental designs obtained with regularized  $\ell_0$ -sparsification are consistently superior to designs obtained with  $\ell_1$ -sparsification, and significantly improve over uniform and random designs.

One limitation of the OED method we have presented is that it relies on linearity of the parameter-to-observable map and assumes Gaussian prior and noise distributions. However, the methods presented here are also applicable in situations where a nonlinear parameter-to-observable map can be well approximated by a linearization over the set of parameters that have significant posterior probability. Moreover, using Gaussian noise and prior distributions is common, particularly in infinite-dimensional inference problems. The computational efficiency of our methods depends on low-rank approximations of the preconditioned parameter-to-observable map, which rely on properties of the associated forward and observation operators. As a consequence of ill-posedness, many parameter-to-observable operators admit such low-rank approximations. Finally, our relaxation of the design problem using a continuous weight vector in combination with sparsification provides just indirect control of the number of sensors through the value of  $\gamma$ . However, this approach makes the combinatorial OED problem of optimal sensor placement computationally tractable.

Possible extensions of the present work include consideration of (1) alternative optimal experimental design criteria that are meaningful in infinite dimensions and (2) nonlinear parameter-to-observable maps. OED with nonlinear parameter-to-observable maps is particularly challenging as, in general, the posterior is non-Gaussian, the (linearization of the) parameter-to-observable map in general depends on the state, parameter, and data variables and thus its low-rank approximation cannot be computed a priori, optimal designs might not be unique even if the sparsifying penalty is convex, and the misfit Hessian depends on (usually unavailable) observations.

**Acknowledgments.** We would like to thank James Martin for providing us with an implementation of the algorithm in [11], which was used to compute the application of the inverse square root of the mass matrix.

**Appendix A. The mass-weighted trace estimator.** The following basic result justifies the form of the randomized trace estimator for an  $\mathbf{M}$ -symmetric matrix as defined in section 2.2. The proof given below adapts the arguments given in [3] regarding Gaussian trace estimators for symmetric linear mappings on the standard Euclidean inner product space.

**PROPOSITION A.1.** *Let  $\mathbf{A}$  be an  $\mathbf{M}$ -symmetric linear mapping on  $\mathbb{R}_{\mathbf{M}}^n$ ,  $n \geq 1$ . Suppose  $\mathbf{y}$  is a random  $n$ -vector with i.i.d.  $\mathcal{N}(0, 1)$ -entries, and let  $\mathbf{z} = \mathbf{M}^{-1/2}\mathbf{y}$ . Then,  $T(\mathbf{A}) := \langle \mathbf{z}, \mathbf{A}\mathbf{z} \rangle_{\mathbf{M}}$  is an unbiased estimator for  $\text{tr}(\mathbf{A})$ ; that is,  $\mathbb{E}\{T(\mathbf{A})\} = \text{tr}(\mathbf{A})$ .*

*Proof.* Since  $\mathbf{A}$  is  $\mathbf{M}$ -symmetric, it admits a spectral decomposition,  $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^*$ , with  $\mathbf{V}$  a matrix with  $\mathbf{M}$ -orthogonal eigenvectors of  $\mathbf{A}$  as its columns (i.e.,  $\mathbf{V}^T\mathbf{M}\mathbf{V} = \mathbf{I}$ ) and  $\mathbf{\Lambda}$  a diagonal matrix with real eigenvalues  $\{\lambda_i\}_{i=1}^n$  of  $\mathbf{A}$  on its diagonal. Then,

$$\mathbb{E}\{T(\mathbf{A})\} = \mathbb{E}\{\mathbf{z}^T\mathbf{M}\mathbf{V}\mathbf{\Lambda}\mathbf{V}^*\mathbf{z}\} = \mathbb{E}\{\mathbf{y}^T\mathbf{M}^{1/2}\mathbf{V}\mathbf{\Lambda}\mathbf{V}^T\mathbf{M}^{1/2}\mathbf{y}\} = \mathbb{E}\{\mathbf{q}^T\mathbf{\Lambda}\mathbf{q}\},$$

with  $\mathbf{q} = \mathbf{V}^T\mathbf{M}^{1/2}\mathbf{y}$ ; note that above we also used  $\mathbf{V}^* = \mathbf{V}^T\mathbf{M}$ . It is straightforward to show that  $\mathbf{q} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . Therefore,

$$\mathbb{E}\{T(\mathbf{A})\} = \mathbb{E}\{\mathbf{q}^T\mathbf{\Lambda}\mathbf{q}\} = \sum_{i=1}^n \lambda_i \mathbb{E}\{q_i^2\} = \sum_{i=1}^n \lambda_i = \text{tr}(\mathbf{A}),$$

where the penultimate equality follows from the fact that  $\mathbb{E}\{q_i^2\} = 1$ , which is the case because  $q_i^2$  is the square of standard normal random variable and is thus  $\chi^2$  distributed with one degree of freedom.  $\square$

**Appendix B. Relation between the trace of posterior covariance and the expected MSE of the posterior mean.** Here, we show that

$$\int_{\mathbb{R}^n} \int_{\mathbb{R}^q} \|\mathbf{m}_{\text{post}}^{\mathbf{d}} - \mathbf{m}\|_{\mathbf{M}}^2 d\mu_{\mathbf{d}|\mathbf{m}}(\mathbf{d}) d\mu_0(\mathbf{m}) = \text{tr}(\mathbf{\Gamma}_{\text{post}}), \quad (\text{B.1})$$

where, according to our choice of the noise model,  $\mu_{\mathbf{d}|\mathbf{m}} = \mathcal{N}(\mathbf{F}\mathbf{m}, \mathbf{\Gamma}_{\text{noise}})$  and  $\mu_0$  is the (discretized) prior measure over  $\mathbb{R}_{\mathbf{M}}^n$ ,  $\mu_0 = \mathcal{N}(\mathbf{m}_0, \mathbf{\Gamma}_{\text{prior}})$ . While (B.1) is known [10] in the context of inference problems in  $n$ -dimensional Euclidean space, below we include a proof of this result adjusted to our choice of mass-weighted inner product, which results from a consistent discretization of the infinite-dimensional Bayesian inverse problem. We use the notation  $\mathbf{m}_{\text{post}} = \mathbf{m}_{\text{post}}^{\mathbf{d}}$  to make the dependence of the posterior mean to data explicit. For brevity, we write (B.1) as

$$\mathbb{E}_{\mu_0} \left\{ \mathbb{E}_{\mathbf{d}|\mathbf{m}} \left\{ \|\mathbf{m}_{\text{post}}^{\mathbf{d}} - \mathbf{m}\|_{\mathbf{M}}^2 \right\} \right\} = \text{tr}(\mathbf{\Gamma}_{\text{post}}). \quad (\text{B.2})$$

Suppose  $\mu$  is a Gaussian measure with mean  $\mathbf{z}_0$  and covariance  $\mathbf{Q}$  defined on a Hilbert space  $V$ . In the proof of the following result we use the fact that  $\int_V \|\boldsymbol{\xi}\|^2 d\mu(\boldsymbol{\xi}) = \text{tr}(\mathbf{Q}) + \|\mathbf{z}_0\|^2$ . Moreover if  $\mathbf{z}$  is Gaussian random variable with probability law  $\mu$  and  $\mathbf{A} : V \rightarrow W$  is a linear transformation with  $W$  a Hilbert space, then  $\mathbf{A}\mathbf{z} + \mathbf{b}$  is also Gaussian with law  $\mathcal{N}(\mathbf{A}\mathbf{z}_0 + \mathbf{b}, \mathbf{A}\mathbf{Q}\mathbf{A}^*)$ , where  $\mathbf{A}^*$  is the adjoint of  $\mathbf{A}$ .

**PROPOSITION B.1.** *Consider a Bayesian linear inverse problem with a Gaussian prior  $\mu_0 = \mathcal{N}(\mathbf{m}_0, \mathbf{\Gamma}_{\text{prior}})$  on  $\mathbb{R}_{\mathbf{M}}^n$  with additive Gaussian noise model as described above. Then, (B.2) holds.*

*Proof.* An elementary calculation shows that,

$$\begin{aligned} \mathbb{E}_{\mathbf{d}|\mathbf{m}} \left\{ \|\mathbf{m}_{\text{post}}^{\mathbf{d}} - \mathbf{m}\|_{\mathbf{M}}^2 \right\} &= \\ \mathbb{E}_{\mathbf{d}|\mathbf{m}} \left\{ \|\mathbf{m}_{\text{post}}^{\mathbf{d}} - \mathbb{E}_{\mathbf{d}|\mathbf{m}} \{ \mathbf{m}_{\text{post}}^{\mathbf{d}} \}\|_{\mathbf{M}}^2 \right\} &+ \|\mathbb{E}_{\mathbf{d}|\mathbf{m}} \{ \mathbf{m}_{\text{post}}^{\mathbf{d}} \} - \mathbf{m}\|_{\mathbf{M}}^2, \end{aligned} \quad (\text{B.3})$$

Let us consider the first term on the right-hand-side of (B.3); denote  $S\mathbf{d} = \mathbf{m}_{\text{post}}^{\mathbf{d}} - \mathbb{E}_{\mathbf{d}|\mathbf{m}} \{ \mathbf{m}_{\text{post}}^{\mathbf{d}} \}$ . Recalling that  $\mathbf{d}|\mathbf{m} \sim \mathcal{N}(\mathbf{F}\mathbf{m}, \mathbf{\Gamma}_{\text{noise}})$ , we note that for any fixed  $\mathbf{m}$ , we have  $S\mathbf{d} = \mathbf{\Gamma}_{\text{post}}\mathbf{F}^*\mathbf{\Gamma}_{\text{noise}}^{-1}(\mathbf{d} - \mathbf{F}\mathbf{m})$ . Therefore, the law of the random variable

$S : (\mathbb{R}^q, \mu_{d|m}) \rightarrow \mathbb{R}_M^n$  is  $\mu_S = \mathcal{N}(\mathbf{0}, \Gamma_{\text{post}} \mathbf{H}_{\text{misfit}} \Gamma_{\text{post}})$ , where  $\mathbf{H}_{\text{misfit}} = \mathbf{F}^* \Gamma_{\text{noise}}^{-1} \mathbf{F}$ . Therefore,

$$\mathbb{E}_{d|m} \left\{ \left\| \mathbf{m}_{\text{post}}^d - \mathbb{E}_{d|m} \left\{ \mathbf{m}_{\text{post}}^d \right\} \right\|_{\mathbf{M}}^2 \right\} = \int_{\mathbb{R}^n} \|\boldsymbol{\xi}\|_{\mathbf{M}}^2 d\mu_S(\boldsymbol{\xi}) = \text{tr}(\Gamma_{\text{post}}^2 \mathbf{H}_{\text{misfit}}), \quad (\text{B.4})$$

where we used that  $\text{tr}(\Gamma_{\text{post}} \mathbf{H}_{\text{misfit}} \Gamma_{\text{post}}) = \text{tr}(\Gamma_{\text{post}}^2 \mathbf{H}_{\text{misfit}})$ . Next, we consider the expectation over  $\mu_0$  of the second term in (B.3). We let  $T : (\mathbb{R}_M^n, \mu_0) \rightarrow \mathbb{R}_M^n$  be defined by  $T\mathbf{m} = \mathbf{m} - \mathbb{E}_{d|m} \left\{ \mathbf{m}_{\text{post}}^d \right\}$ . Then,

$$T\mathbf{m} = \Gamma_{\text{post}} (\Gamma_{\text{post}}^{-1} - \mathbf{H}_{\text{misfit}}) \mathbf{m} - \Gamma_{\text{post}} \Gamma_{\text{prior}}^{-1} \mathbf{m}_0 = \Gamma_{\text{post}} \Gamma_{\text{prior}}^{-1} (\mathbf{m} - \mathbf{m}_0). \quad (\text{B.5})$$

Hence, the law of  $T$  is given by  $\mu_T = \mathcal{N}(\mathbf{0}, \mathbf{C})$  with  $\mathbf{C} = (\Gamma_{\text{post}} \Gamma_{\text{prior}}^{-1}) \Gamma_{\text{prior}} (\Gamma_{\text{post}} \Gamma_{\text{prior}}^{-1})^* = \Gamma_{\text{post}} \Gamma_{\text{prior}}^{-1} \Gamma_{\text{post}}$ . Therefore,

$$\mathbb{E}_{\mu_0} \left\{ \left\| \mathbb{E}_{d|m} \left\{ \mathbf{m}_{\text{post}}^d \right\} - \mathbf{m} \right\|_{\mathbf{M}}^2 \right\} = \int_{\mathbb{R}^n} \|\boldsymbol{\xi}\|_{\mathbf{M}}^2 d\mu_T(\boldsymbol{\xi}) = \text{tr}(\Gamma_{\text{post}}^2 \Gamma_{\text{prior}}^{-1}). \quad (\text{B.6})$$

Hence, using (B.3), (B.4), and (B.6), we have

$$\begin{aligned} \mathbb{E}_{d|m} \left\{ \left\| \mathbf{m}_{\text{post}}^d - \mathbf{m} \right\|_{\mathbf{M}}^2 \right\} &= \text{tr}(\Gamma_{\text{post}}^2 \mathbf{H}_{\text{misfit}}) + \text{tr}(\Gamma_{\text{post}}^2 \Gamma_{\text{prior}}^{-1}) \\ &= \text{tr}(\Gamma_{\text{post}}^2 (\mathbf{H}_{\text{misfit}} + \Gamma_{\text{prior}}^{-1})) = \text{tr}(\Gamma_{\text{post}}^2), \end{aligned}$$

where, for the last equality we used  $\mathbf{H}_{\text{misfit}} + \Gamma_{\text{prior}}^{-1} = \Gamma_{\text{post}}^{-1}$ .  $\square$

## REFERENCES

- [1] VOLKAN AKÇELİK, GEORGE BIROS, ANDREI DRAGANESCU, OMAR GHATTAS, JUDITH HILL, AND BART VAN BLOEMAN WAANDERS, *Dynamic data-driven inversion for terascale simulations: Real-time identification of airborne contaminants*, in Proceedings of SC2005, Seattle, 2005.
- [2] ANTHONY C. ATKINSON AND ALEXANDER N. DONEV, *Optimum Experimental Designs*, Oxford, 1992.
- [3] HAIM AVRON AND SIVAN TOLEDO, *Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix*, Journal of the ACM (JACM), 58 (2011), p. 17.
- [4] MARTIN P. BENDSØE, *Optimization of structural topology, shape and material*, Springer Verlag, Berlin, Heidelberg, New York, 1995.
- [5] MARTIN P. BENDSØE AND OLE SIGMUND, *Topology Optimization. Theory, methods and Applications*, Springer Verlag, Berlin, Heidelberg, New York, 2003.
- [6] HANS GEORG BOCK, STEFAN KÖRKEL, AND JOHANNES P. SCHLÖDER, *Parameter estimation and optimum experimental design for differential equation models*, in Model Based Parameter Estimation, Hans Georg Bock, Thomas Carraro, Willi Jäger, Stefan Körkel, Rolf Rannacher, and Johannes P. Schlöder, eds., vol. 4 of Contributions in Mathematical and Computational Sciences, Springer Berlin Heidelberg, 2013, pp. 1–30.
- [7] TAN BUI-THANH, CARSTEN BURSTEDDE, OMAR GHATTAS, JAMES MARTIN, GEORG STADLER, AND LUCAS C. WILCOX, *Extreme-scale UQ for Bayesian inverse problems governed by PDEs*, in SC12: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 2012.
- [8] TAN BUI-THANH, OMAR GHATTAS, JAMES MARTIN, AND GEORG 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 (2013), pp. A2494–A2523.
- [9] EMMANUEL J. CANDÈS, JUSTIN ROMBERG, AND TERENCE TAO, *Stable signal recovery from incomplete and inaccurate measurements*, Communications on Pure and Applied Mathematics, 59 (2006), pp. 1207–1223.
- [10] KATHRYN CHALONER AND ISABELLA VERDINELLI, *Bayesian experimental design: A review*, Statistical Science, 10 (1995), pp. 273–304.

- [11] JIE CHEN, MIHAI ANITESCU, AND YOUSEF SAAD, *Computing  $f(A)b$  via least squares polynomial approximations*, SIAM Journal on Scientific Computing, 33 (2011), pp. 195–222.
- [12] PETER DEUFLHARD, *Newton methods for nonlinear problems*, vol. 35 of Springer Series in Computational Mathematics, Springer-Verlag, Berlin, 2004.
- [13] DAVID L. DONOHO, *Compressed sensing*, IEEE Transactions on Information Theory, 52 (2006), pp. 1289–1306.
- [14] H. PEARL FLATH, LUCAS C. WILCOX, VOLKAN AKÇELIK, JUDY HILL, BART VAN BLOEMEN WAANDERS, AND OMAR 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 (2011), pp. 407–432.
- [15] GENE H. GOLUB AND CHARLES F. VAN LOAN, *Matrix Computations*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, third ed., 1996.
- [16] ELDAD HABER, LIOR HORESH, AND LUIS TENORIO, *Numerical methods for experimental design of large-scale linear ill-posed inverse problems*, Inverse Problems, 24 (2008), pp. 125–137.
- [17] ELDAD HABER, LIOR HORESH, AND LUIS TENORIO, *Numerical methods for the design of large-scale nonlinear discrete ill-posed inverse problems*, Inverse Problems, 26 (2010), p. 025002.
- [18] ELDAD HABER, ZHUOJUN MAGNANT, CHRISTIAN LUCERO, AND LUIS TENORIO, *Numerical methods for A-optimal designs with a sparsity constraint for ill-posed inverse problems*, Computational Optimization and Applications, (2012), pp. 1–22.
- [19] NATHAN HALKO, PER-GUNNAR MARTINSSON, AND JOEL A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288.
- [20] LIOR HORESH, ELDAD HABER, AND LUIS TENORIO, *Optimal experimental design for the large-scale nonlinear ill-posed problem of impedance imaging*, Large-Scale Inverse Problems and Quantification of Uncertainty, (2010), pp. 273–290.
- [21] XUN HUAN AND YOUSSEF M. MARZOUK, *Simulation-based optimal Bayesian experimental design for nonlinear systems*, Journal of Computational Physics, 232 (2013), pp. 288–317.
- [22] ———, *Gradient-based stochastic optimization methods in Bayesian experimental design*, International Journal for Uncertainty Quantification, (2014). to appear.
- [23] MICHAEL F. HUTCHINSON, *A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines*, Communications in Statistics-Simulation and Computation, 19 (1990), pp. 433–450.
- [24] PETER D. LAX, *Functional Analysis*, John Wiley & Sons, New-York, Chicester, Brisbane, Toronto, 2002.
- [25] JAMES MARTIN, LUCAS C. WILCOX, CARSTEN BURSTEDDE, AND OMAR GHATTAS, *A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion*, SIAM Journal on Scientific Computing, 34 (2012), pp. A1460–A1487.
- [26] JAMES MERCER, *Functions of positive and negative type, and their connection with the theory of integral equations*, Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character, (1909), pp. 415–446.
- [27] JORGE NOCEDAL AND STEPHEN J. WRIGHT, *Numerical Optimization*, Springer Verlag, Berlin, Heidelberg, New York, second ed., 2006.
- [28] ANDREJ PÁZMAN, *Foundations of Optimum Experimental Design*, D. Reidel Publishing Co., 1986.
- [29] NOEMI PETRA AND GEORG STADLER, *Model variational inverse problems governed by partial differential equations*, Tech. Report 11-05, The Institute for Computational Engineering and Sciences, The University of Texas at Austin, 2011.
- [30] FRIEDRICH PUKELSHEIM, *Optimal Design of Experiments*, John Wiley & Sons, New-York, 1993.
- [31] FARBOD ROOSTA-KHORASANI AND URI ASCHER, *Improved bounds on sample size for implicit matrix trace estimators*. Preprint, 2013.
- [32] ANDREW M. STUART, *Inverse problems: A Bayesian perspective*, Acta Numerica, 19 (2010), pp. 451–559.
- [33] DARIUSZ UCIŃSKI, *Optimal measurement methods for distributed parameter system identification*, CRC Press, Boca Raton, 2005.
- [34] DAVID WILLIAMS, *Probability with Martingales*, Cambridge University Press, 1991.