

Goal-Oriented Optimal Sensor Placement for PDE-Constrained Inverse Problems in Crisis Management^{*)}

Marco MATTUSCHKA^{1)*}, Noah AN DER LAN¹⁾, Max VON DANWITZ¹⁾, Daniel WOLFF²⁾, Alexander POPP^{1),2)}

¹⁾ *Institute for the Protection of Terrestrial Infrastructures, German Aerospace Center (DLR), Sankt Augustin, Germany*

²⁾ *Institute for Mathematics and Computer-Based Simulation (IMCS), University of the Bundeswehr Munich, Neubiberg, Germany*

* *Corresponding Author: marco.mattuschka@dlr.de*

This paper presents a novel framework for goal-oriented optimal static sensor placement and dynamic sensor steering in PDE-constrained inverse problems, utilizing a Bayesian approach accelerated by low-rank approximations. The framework is applied to airborne contaminant tracking, extending recent dynamic sensor steering methods to complex geometries to improve computational efficiency. A C-optimal design criterion is employed to strategically place sensors, minimizing uncertainty in predictions. Numerical experiments validate the approach's effectiveness for source identification and monitoring, highlighting its potential for real-time decision-making in crisis management scenarios.

Keywords: airborne contaminant transport, advection-diffusion equation, large-scale inverse problems, optimal experimental design, dynamic sensor steering.



Copyright © 2026 The Author(s).

Published by IPPT PAN. This work is licensed under the Creative Commons Attribution License CC BY 4.0 (<https://creativecommons.org/licenses/by/4.0/>).

1. INTRODUCTION

The growing capabilities of unmanned systems (UxS), such as unmanned aerial vehicles (UAVs, drones), autonomous underwater vehicles (AUVs), and unmanned ground vehicles (e.g., mobile robots) have made these systems indispensable tools in crisis management situations (see Fig. 1 for examples). Such autonomous sensor platforms facilitate the collection of valuable information in regions where manned missions would be too dangerous or simply impossible due

^{*)} The results of this study were presented at GAMM 2025 – 95th Annual Meeting of the International Association of Applied Mathematics and Mechanics, April 7–11, 2025, Poznań, Poland.



FIG. 1. Unmanned systems operated as autonomous sensor platforms at the German Aerospace Center (DLR). Images: DLR, CC BY-NC-ND 3.0.

to inaccessibility for humans. However, during the routing of a sensor platform in the complex environment of a crisis situation, the question arises as to at which locations (additional) measurements provide information gain and thereby constitute added value for decision-makers. This very question leads to the need for goal-oriented optimal sensor placement and, in the dynamic case, to a sensor steering problem. The task of selecting optimal measurement locations to infer system knowledge falls within the classical field of optimal experimental design (OED). Goal-oriented optimal experimental design generalizes this concept by determining the optimal design not for the parameters themselves, but for a derived quantity of interest (QoI). This approach offers two major advantages. First, it enables a more application-specific formulation, as the optimal design targets the QoI directly rather than the entire parameter space of the inverse problem. Second, it substantially reduces computational cost, since the QoI typically has far fewer degrees of freedom in its discretized representation.

This contribution addresses the problem in the mathematical setting of inverse problems constrained by partial differential equations (PDEs). A selected sensor steering strategy is developed and applied to the specific challenge of mapping airborne contaminant dispersion in the region of interest using discrete sensor measurements. By specifying the region of interest in both space and time, the problem naturally falls within the framework of goal-oriented optimal sensor placement.

Current methods for contaminant source identification and spread prediction rely heavily on an informative sensor placement. The selection of measuring points (sensor positions) is crucial; however, many existing studies focus solely

on stationary sensors. This work bridges this gap by incorporating recent advances in sensor selection and experimental design to derive an algorithm for optimal sensor steering. Our research focuses on developing a systematic approach to select sensor positions that maximize the accuracy of contaminant source identification and prediction. By integrating current methodological advancements, our aim is to provide a practical solution for rescue workers and first responders, enabling informed decision-making in high-stakes situations. Whereas the numerical examples presented in this work focus on the specific application of airborne contaminant transport, the goal-oriented optimal sensor placement strategy is independent of the underlying physical model and can therefore be easily transferred to other crisis management applications.

The starting point for this work is the recent publication by WOGGRIN *et al.* [23], which pioneers a dynamic sensor steering method in the context of airborne contaminant transport. In the present work, we extend this approach to a significantly more complex geometry. Moreover, the inverse problem solution follows a more advanced approach that uses a Laplacian-like trace class operator as prior information within a Bayesian inverse problem framework [16, 20]. To achieve approximate real-time capability, low-rank approximations of the Hessian matrix are precomputed in an offline phase [13, 15], enabling efficient problem solving in the online phase using a preconditioned inexact Newton-conjugate gradient (CG) solver [19].

In the calculation of an optimal (stationary) sensor layout, ALEXANDERIAN *et al.* [3] used a reduced model for the contaminant transport to determine an A-optimal design that minimizes the average point-wise posterior variance of the inferred parameter vector. Following extensions proposed in [5] and [18], we focus on a goal-oriented design, i.e., the uncertainty associated with the prediction of the contaminant concentration in a specific region and time is minimized. To demonstrate the feasibility of sensor steering, a relaxed optimality criterion compared to the A-optimal criterion is chosen for this test case. Specifically, we use the C-optimal criterion, which focuses on minimizing the posterior variance of a particular linear combination of the inversion parameters. This approach eliminates the need to estimate the trace of the full covariance matrix and allows us to directly assess the impact of the covariance matrix on the parameter of interest. Alternatively, a D-optimal goal-oriented design in infinite dimensions maximizes the expected information gain [3, 24]. For a broader perspective on optimal experimental design for infinite-dimensional Bayesian inverse problems governed by PDEs, the interested reader is referred to [1]. The remainder of this paper is organized as follows. Section 2 provides background and mathematical formulations of the forward problem of contaminant transport, the inverse problem of source identification, as well as sensor positioning strategies and goal-oriented optimization. The combination of methodological

developments into an algorithm for goal-oriented optimal sensor placement and steering is described in [Sec. 3](#). [Section 4](#) introduces dynamic sensor steering, showing how a mobile sensor can be guided in real time to maximize information about the contaminated field. Numerical results are presented in [Sec. 5](#) for three test cases of goal-oriented optimal sensor placement: (a) identifying an instantaneous contaminant source in a user-defined area of interest, (b) monitoring an area of special interest over a predefined time window, and (c) steering a moving sensor. Finally, [Sec. 6](#) offers a conclusion and outlook.

2. BACKGROUND

2.1. FORWARD PROBLEM: CONTAMINANT DISTRIBUTION EVOLUTION

A mathematical description of the transport of a substance (contaminant) concentration u in a bounded open domain $\Omega \subseteq \mathbb{R}^n$ for $n \in \{2, 3\}$, as shown in [Fig. 2](#), is given by the following equation:

$$\begin{aligned} r_{\mathcal{K}}(u) &:= u_t - \kappa \Delta u + \mathbf{v} \cdot \nabla u = 0 && \text{in } (0, T) \times \Omega, \\ \kappa \nabla u \cdot \mathbf{n} &= 0 && \text{in } (0, T) \times (\Gamma_+ \cup \Gamma_0), \\ u &= 0 && \text{in } (0, T) \times \Gamma_-, \\ u(0, \cdot) &= m && \text{in } \Omega. \end{aligned} \quad (1)$$

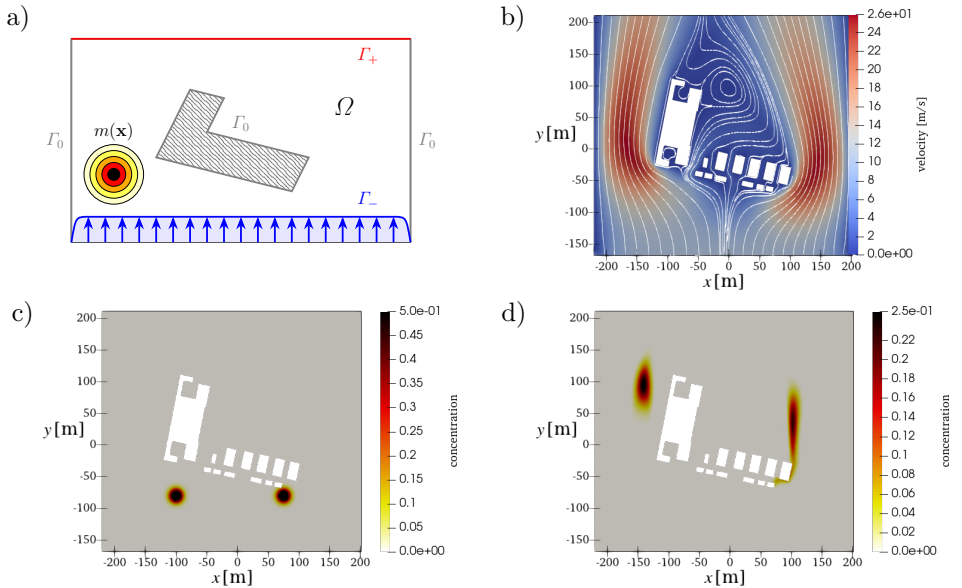


FIG. 2. Forward simulation of airborne contaminant transport on a campus geometry. Computational domain Ω with highlighted inflow (Γ_-), outflow (Γ_+), and characteristic (Γ_0) boundaries as well as the initial contaminant source $u(0, \cdot) = m(\cdot)$ (a). Estimated wind vector field \mathbf{v} (b), initial condition (c), and simulated concentration at $t = 10$ s (d).

The parameter-dependent forward problem shown in Eq. (1) is formulated for realizations of the parameter m . A visualization of the contaminant dispersion is provided in Fig. 2.

The underlying transport process is governed by a wind vector field \mathbf{v} , which is assumed to be sufficiently smooth, bounded – i.e., $\mathbf{v} \in L^\infty(\Omega, \mathbb{R}^n)$ – and divergence-free – i.e., $\nabla \cdot \mathbf{v} = 0$. The example wind vector field used hereinafter is shown in Fig. 2b. Based on the orientation of the wind vector relative to the outward-pointing boundary normal \mathbf{n} , the domain boundary $\partial\Omega$ is partitioned into three disjoint subsets: the outflow boundary $\Gamma_+ \subset \partial\Omega$, where $\mathbf{v} \cdot \mathbf{n} > 0$; the characteristic (or tangential) boundary $\Gamma_0 \subset \partial\Omega$, where $\mathbf{v} \cdot \mathbf{n} = 0$; and the inflow boundary $\Gamma_- \subset \partial\Omega$, where $\mathbf{v} \cdot \mathbf{n} < 0$, following the convention in [10].

2.2. INVERSE PROBLEM: SOURCE IDENTIFICATION

Whenever measurements of the concentration at discrete locations and times are available, an obvious question is whether the initial condition can be reconstructed on the basis of the given measurements. The respective inverse problem is illustrated in Fig. 3, see also [20]. As the function space for the initial condition, we consider an admissible subset of square-integrable functions. For instance, we define $D := H_{\Gamma_-}^{1,2}(\Omega) := \{m \in H^{1,2}(\Omega) \mid m|_{\Gamma_-} = 0\}$, where $H^{1,2}(\Omega)$ denotes the Sobolev space of square-integrable functions whose weak derivatives $\partial_i u$ exist and are square-integrable. The gradient is denoted as $\nabla u = (\partial_1 u, \partial_2 u)$, and the norm on $H^{1,2}(\Omega)$ is defined as

$$\|u\|_{H^{1,2}(\Omega)} := \left(\int_{\Omega} u^2 + \|\nabla u\|_2^2 dx \right)^{1/2}.$$

In this setting, the estimation of the initial value leads to an optimization problem, which will be addressed in the following. The first step is to describe sensor measurements within this formulation.

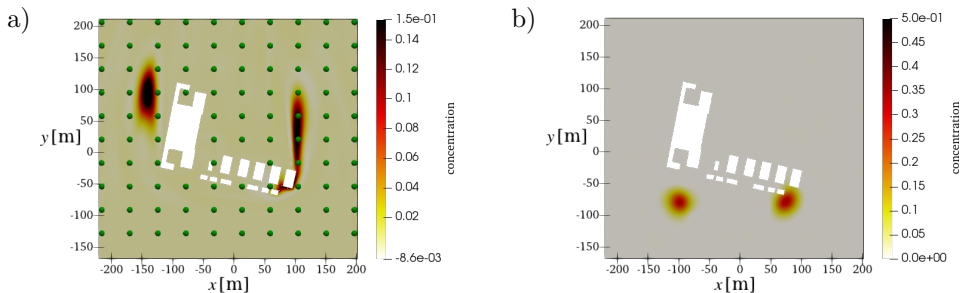


FIG. 3. Inverse problem. Measurements at 96 equidistantly spaced sensor positions (a) and reconstructed initial condition (solution of the inverse problem (b)).

To incorporate sensor measurements, let $T_0 > 0$ be fixed and consider space-time points $(t_i^{\text{obs}}, x_i^{\text{obs}})$, $i = 1, \dots, q$, with $t_i^{\text{obs}} \in [T_0, T)$ and interior points $x_i^{\text{obs}} \in \Omega = \overline{\Omega} \setminus \partial\Omega$. For the transient problem in Eq. (1) on a smooth domain Ω , the diffusion term induces strong smoothing on the parameter field m . By standard parabolic regularity (see, e.g., [11]), $u(t, \cdot)$ embeds in the space of continuous functions $C^0(\overline{\Omega})$ on the closure of Ω for any $t > 0$. Since continuity is only required near the observation points, define $\overline{\Omega}_o := \bigcup_{i=1}^q B_{r_0}(x_i^{\text{obs}}) \subset \Omega$, where $r_0 > 0$ is chosen such that $u \in C^0([T_0, T] \times \overline{\Omega}_o)$ admits a continuous representation. This ensures the formulation remains valid for less regular domains Ω . Finally, we define a well-posed and bounded space-time observation operator $\mathcal{B} : C^0([T_0, T] \times \overline{\Omega}_o) \rightarrow \mathbb{R}^q$ by $u \mapsto \sum_{i=1}^q \delta_{(t_i^{\text{obs}}, x_i^{\text{obs}})}(u) \mathbf{e}_i = (u(t_i^{\text{obs}}, x_i^{\text{obs}}))_{i=1}^q$, where \mathbf{e}_i is the standard basis of \mathbb{R}^q . Using this observation operator, the final parameter-to-observable map $\mathcal{F} : D \rightarrow \mathbb{R}^q$ is defined by

$$\mathcal{F}(m) := \mathcal{B} \circ \mathcal{K}(m), \quad \text{with } \mathcal{K}(m) := u \text{ such that } r_{\mathcal{K}}(u) = m. \quad (2)$$

Here, \mathcal{K} is the parameter-to-state map, mapping the parameter space D to the state space, often referred to as the ‘model’ in this context [8]. An example problem in which \mathcal{K} maps the initial condition u_0 to the solution of $u(t = 10 \text{ s}, \cdot)$ is illustrated in Fig. 2.

The next step is to model sensor noise, which is usually present in real-world measurements. To this end, it is assumed that the observations $\mathbf{d} = \mathcal{F}(m) + \boldsymbol{\epsilon}$ contain centered additive Gaussian noise $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \Gamma_{\text{noise}})$ due to measurement uncertainties. For the sake of simplicity, it is further assumed that the sensor noise at the different sensor positions is uncorrelated and of equal magnitude, represented by the diagonal matrix $\Gamma_{\text{noise}} = \text{diag}(\sigma^2, \dots, \sigma^2)$. Moreover, the conformity of the simulation with the measured values, also called misfit, $\mathbf{y} = \mathcal{F}(m) - \mathbf{d}$, is evaluated in the following norm:

$$\|\mathcal{F}(m) - \mathbf{d}\|_{\Gamma_{\text{prior}}^{-1}}^2 = 1/\sigma^2 \sum_{i=1}^q \left(u(t_i^{\text{obs}}, x_i^{\text{obs}}) - d_i \right)^2.$$

In summary, the inverse problem consists of using available measurements \mathbf{d} to infer the values of the unknown parameter field m . Alternatively, the maximum a-posteriori point m_{map} can be characterized as the solution of the minimization problem:

$$m_{\text{map}} = \arg \min_{m \in D} J(m) := \frac{1}{2} \|\mathcal{F}(m) - \mathbf{d}\|_{\Gamma_{\text{noise}}^{-1}}^2 + \frac{1}{2} \|m - m_{\text{prior}}\|_{\Gamma_{\text{prior}}^{-1}}^2, \quad (3)$$

with the prior information encoded as a Tikhonov regularization term.

Unfortunately, the system is heavily under-determined under realistic conditions, as sensor measurements are only available at a few locations, whereas the initial condition is to be reconstructed for the entire domain. In order to transform this into a well-posed problem, some prior knowledge is needed and a formulation as a Bayesian inverse problem provides a suitable framework. In this setting, a Gaussian prior $\mathcal{N}(m_{\text{prior}}, \Gamma_{\text{prior}})$ with mean m_{prior} and covariance Γ_{prior} is chosen for parameter regularization. Then, the posterior density of m satisfies by Bayes' theorem: $\pi_{\text{post}}(m|\mathbf{d}) \propto \pi_{\text{like}}(\mathbf{d}|m) \pi_{\text{prior}}(m)$. Here, $\pi_{\text{like}}(\mathbf{d}|m) \propto \exp(-\frac{1}{2} \|\mathcal{F}(m) - \mathbf{d}\|_{\Gamma_{\text{noise}}^{-1}}^2)$ is the likelihood function under the observational Gaussian noise $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \Gamma_{\text{noise}})$. Due to the linearity of \mathcal{F} , the posterior distribution is again a Gaussian distribution $\mathcal{N}(m_{\text{map}}, \Gamma_{\text{post}})$ with covariance and mean:

$$\begin{aligned} \Gamma_{\text{post}} &= (\mathcal{F}^* \Gamma_{\text{noise}}^{-1} \mathcal{F} + \Gamma_{\text{prior}}^{-1})^{-1}, \\ m_{\text{map}} &= \Gamma_{\text{post}} (\mathcal{F}^* \Gamma_{\text{noise}}^{-1} \mathbf{d} + \Gamma_{\text{prior}}^{-1} m_{\text{prior}}). \end{aligned} \quad (4)$$

Here, the formal adjoint operator $\mathcal{F}^* : \mathbb{R}^q \rightarrow D^*$ is required for this framework. For the mapping $\mathcal{F} : D \rightarrow \mathbb{R}^q$ between Hilbert spaces, the formal adjoint operator \mathcal{F}^* is characterized by the relation $\langle \mathcal{F}(m), \mathbf{y} \rangle_{\mathbb{R}^q} = \langle m, \mathcal{F}^*(\mathbf{y}) \rangle_{L^2(\Omega)}$ for all $\mathbf{y} \in \mathbb{R}^q$ and $m \in D$, and its existence follows from Riesz's representation theorem [4]. The posterior covariance or Hessian $\mathcal{H} := \mathcal{F}^* \Gamma_{\text{noise}}^{-1} \mathcal{F} + \Gamma_{\text{prior}}^{-1}$ of the objective function J , see Eq. (3), contains a wealth of information about the system. In line with the Bayesian framework [20], the covariance or inverse Hessian matrix can be employed to predict the uncertainty of the system and is of particular significance for optimal sensor placement in the following section.

The mean value m_{map} is a reliable estimate for the initial value and thus represents the solution of the inverse problem Fig. 3. In order to calculate m_{map} from Eq. (4), a further specification of the adjoint operator \mathcal{F}^* is necessary. By applying integration by parts to the weak form of Eq. (1), the adjoint state p can be derived and satisfies the following equation:

$$\begin{aligned} -p_t - \kappa \Delta p - \text{div}(p\mathbf{v}) &= -\frac{1}{\sigma^2} \sum_{i=1}^q y_i \delta_{(t_i^{\text{obs}}, x_i^{\text{obs}})} && \text{in } (0, T) \times \Omega, \\ (\mathbf{v}p + \kappa \nabla p) \cdot \mathbf{n} &= 0 && \text{in } (0, T) \times (\Gamma_+ \cup \Gamma_0), \\ p &= 0 && \text{in } (0, T) \times \Gamma_-, \\ p(T, \cdot) &= 0 && \text{in } \Omega, \end{aligned} \quad (5)$$

for given $\mathbf{y} \in \mathbb{R}^q$. Finally, the adjoint operator \mathcal{F}^* can be explicitly determined, resulting in $\mathcal{F}^* \mathbf{y} = p(0, \cdot)$.

2.3. GOAL-ORIENTED OPTIMAL EXPERIMENTAL DESIGN: SENSOR PLACEMENT

So far, a model for the forward problem and an estimate of the initial condition, given a fixed sensor configuration, have been derived. However, the question of how such a sensor arrangement ought to be chosen remains unanswered. As mentioned in the previous section, the covariance of the posterior, or more specifically, the Hessian matrix, plays a crucial role in developing an indicator for the uncertainty in the system. In classical Bayesian optimal experimental design (OED), a design functional Φ serves as a quantitative criterion for the uncertainty described by the covariance matrix and is used to guide the selection of an optimal experimental setup. Specifically, the optimization problem

$$\min_{\mathbf{w} \in \mathbb{W}} \left(\Phi[\Gamma_{\text{post}}(\mathbf{w})] + \mathcal{R}(\mathbf{w}) \right) \quad (6)$$

is solved, where \mathbb{W} denotes the set of all valid sensor configurations, and \mathcal{R} is a suitable regularization term. As in [2], a finite set of candidate sensor placements $(t_i^{\text{obs}}, x_i^{\text{obs}}) \in [T_0, T] \times \Omega$ for $1 \leq i \leq q$ is considered. An example of a spatial grid with 96 sensors can be seen in Fig. 3. For this set of candidate locations, a weight vector $\mathbf{w} \in [0, 1]^q$ is defined with the i -th candidate entry corresponding to the i -th location in space and time. In fact, the weight vector determines which measurements are realized or taken into account. In the case of stationary sensors located at positions $\mathbf{x}_s^{\text{obs}}$, measurements collected at these spatial points over the entire time horizon, denoted as $(\cdot, \mathbf{x}_s^{\text{obs}})$, are constantly weighted and included in the misfit. Hence, the number of independent entries in the weight vector reduces to the number of possible stationary sensor positions. If we consider a mobile sensor, we have a trajectory $\gamma : \{t_0^{\text{obs}}, \dots, t_s^{\text{obs}}\} \rightarrow \{\mathbf{x}_0^{\text{obs}}, \dots, \mathbf{x}_s^{\text{obs}}\}$. For points on the trajectory $(t_i^{\text{obs}}, \mathbf{x}_i^{\text{obs}})$ the sensor weight is set to $\mathbf{w}_i = 1$. All weights away from the trajectory are set to 0. To adjust the forward model to the chosen sensor configuration, we consider the diagonal matrix $W \in \mathbb{R}^{q \times q}$ with $W_{ii} = \mathbf{w}_i$. If we denote the parameter-to-observable map for all sensor positions as \mathcal{F} , then for each design $\mathbf{w} \in \mathbb{W}$, we have $\mathcal{F}(\mathbf{w}) = W\mathcal{F}$. Taking this further and using W to modify the noise matrix, the influence of a selected sensor layout is also captured in the likelihood function

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

In consequence, the posterior covariance and mean also depend on the sensor layout via (cf. Eq. (4)):

$$\begin{aligned} \Gamma_{\text{post}}(\mathbf{w}) &= (\mathcal{F}^* W^{1/2} \Gamma_{\text{noise}}^{-1} W^{1/2} \mathcal{F} + \Gamma_{\text{prior}}^{-1})^{-1}, \\ m_{\text{map}}(\mathbf{w}) &= \Gamma_{\text{post}}(\mathcal{F}^* \Gamma_{\text{noise}}^{-1} W \mathbf{d} + \Gamma_{\text{prior}} m_{\text{prior}}). \end{aligned} \quad (7)$$

In Bayesian optimal experimental design, the objective is to construct experiments that minimize the posterior uncertainty in the unknown parameter field m . This approach belongs to the broader class of inverse uncertainty quantification methods. Goal-oriented optimal experimental design extends this framework by additionally incorporating forward uncertainty quantification into the design objective. This is achieved by introducing a goal QoI, denoted by ρ . Following [5, Sec. 3], we consider a QoI defined as the action of a bounded linear operator \mathcal{P} on the parameter field, i.e., $\rho = \mathcal{P}m$. Due to the linearity of \mathcal{P} , the prior distribution of ρ is Gaussian. In particular, $\rho \sim \mathcal{N}(\rho_{\text{prior}}, \Sigma_{\text{prior}})$, with mean $\rho_{\text{prior}} = \mathcal{P}m_{\text{prior}}$ and covariance $\Sigma_{\text{prior}} = \mathcal{P}\Gamma_{\text{prior}}\mathcal{P}^*$. The resulting Bayesian inverse problem is well-defined, and the posterior distribution of the QoI given measurements \mathbf{d} is also Gaussian: $\pi_{\text{post}}(\rho \mid \mathbf{d}) \sim \mathcal{N}(\rho_{\text{post}}, \Sigma_{\text{post}})$. The posterior mean and covariance are given by:

$$\Sigma_{\text{post}} = \mathcal{P}\Gamma_{\text{post}}\mathcal{P}^* \quad \text{and} \quad \rho_{\text{post}} = \mathcal{P}m_{\text{map}},$$

where Γ_{post} and m_{map} are the derived quantities from Eq. (7).

3. DISCRETIZATION, PRECONDITIONING, AND SPARSIFICATION

3.1. FINITE ELEMENT DISCRETIZATION

To solve the PDE problems (Eq. (1) and Eq. (5)) numerically, a finite element discretization is employed using n_{dof} Lagrange basis functions $\mathcal{V}_h = \text{span}\{\phi_1, \dots, \phi_{n_{\text{dof}}}\}$. Moreover, we find an identification between a vector in $\mathbb{R}^{n_{\text{dof}}}$ and finite elements $I : \mathbb{R}^{n_{\text{dof}}} \rightarrow \mathcal{V}_h$ via $I(a) = \sum_{i=1}^{n_{\text{dof}}} a_i \phi_i$. This leads to discretized versions of the parameter-to-observable map $\mathcal{F}_h : \mathbb{R}^{n_{\text{dof}}} \rightarrow \mathbb{R}^q$ defined by $\mathcal{F}_h(m_h) = \mathcal{B}(u_h)$, where u_h solves Eq. (1) in a weak form, and its adjoint operator $\mathcal{F}_h^* : \mathbb{R}^q \rightarrow \mathbb{R}^{n_{\text{dof}}}$ given by $\mathcal{F}_h^* \mathbf{y} = p_h(0, \cdot)$, where p_h solves Eq. (5) in a weak form. The identification is an isometry, i.e., $\langle I(a), I(b) \rangle_{L^2(\Omega)} = \langle a, b \rangle_M =: a^T M b$, where the mass matrix $M_{ji} := \int_{\Omega} \phi_i(\mathbf{x}) \phi_j(\mathbf{x}) d\mathbf{x}$, $M \in \mathbb{R}^{n_{\text{dof}} \times n_{\text{dof}}}$ is used to define the corresponding scalar product. For further details of the finite element discretization, we refer to [20] and [21].

3.2. PRECONDITIONING OF THE DISCRETE INVERSE PROBLEM

To solve the discrete inverse problem, the prior distribution needs to be carefully chosen. A Laplacian-like operator of trace class $\mathcal{A} := (\eta I - \gamma \Delta)$, with Robin boundary condition, $\gamma \nabla m \cdot \mathbf{n} + \beta m = 0$ in $(0, T) \times \partial\Omega$, is applied with the constant β as proposed in [9]. This operator serves as a suitable covariance operator for the prior, e.g., $\Gamma_{\text{prior}} = \mathcal{A}^{-2} = (\eta I - \gamma \Delta)^{-2}$. In addition,

its discrete counterpart is given by the mapping $\Gamma_{\text{prior,h}} : \mathbb{R}^{n_{\text{dof}}} \rightarrow \mathbb{R}^{n_{\text{dof}}}$ via $\Gamma_{\text{prior,h}} = (M^{-1}A)^{-2} = A^{-1}MA^{-1}M := R^{-1}M$, with matrix representation $A_{ij} = \int \phi_i(x)\mathcal{A}\phi_j(x) dx$.

Combining this covariance operator with an appropriate prior mean m_{prior} (in our applications, e.g., $m_{\text{prior}} = 0$) renders the inverse problem well-posed and its solution can be found by solving the following equation for m_{map} :

$$\mathcal{H}_h(\mathbf{w})m_{\text{map}} = \mathcal{F}_h^* \Gamma_{\text{noise}}^{-1} (W\mathbf{d}) + \Gamma_{\text{prior,h}}^{-1} m_{\text{prior}}, \quad (8)$$

for the discrete version of the Hessian, that is, $\mathcal{H}_h(\mathbf{w}) = \mathcal{F}_h^* W^{1/2} \Gamma_{\text{noise}}^{-1} W^{1/2} \mathcal{F}_h + \Gamma_{\text{prior,h}}^{-1}$. Since directly determining the Hessian matrix is computationally expensive for large-scale problems (requiring $O(n_{\text{dof}})$ -PDE solutions), an iterative CG method is employed. This approach requires only the action of the Hessian-vector on a given vector $m_k \in \mathbb{R}^{n_{\text{dof}}}$ at each iteration. Specifically, the Hessian action is computed by the following steps: first, solve the forward equation $\mathbf{d} = \mathcal{F}_h(m_k) = \mathcal{B}(u_h)$, then, solve the corresponding adjoint equation $\mathcal{F}_h^*(W^{1/2} \Gamma_{\text{noise}}^{-1} W^{1/2} \mathbf{d}) = p_h(0, \cdot)$, next, compute $\tilde{m}_k = \Gamma_{\text{prior,h}} m_k$, and finally obtain the Hessian action as $\mathcal{H}_h(\mathbf{w})m_k = p_h(0, \cdot) + \tilde{m}_k$.

Since two PDE solutions have to be determined per iteration, a reduced model of the Hessian matrix is precomputed in advance so that the inverse problem can be solved quickly. In addition, using the Cholesky decomposition of the prior covariance, $\Gamma_{\text{prior,h}}^{-1} = (M^{-1}A)(M^{-1}A)^*$, one obtains the preconditioned Hessian matrix as:

$$(\mathcal{A}_h^{-1})^* \mathcal{H}_h(\mathbf{w}) \mathcal{A}_h^{-1} = (\mathcal{F}_h \circ \mathcal{A}_h^{-1})^* W^{1/2} \Gamma_{\text{noise}}^{-1} W^{1/2} (\mathcal{F}_h \circ \mathcal{A}_h^{-1}) + I, \quad (9)$$

where $\mathcal{A}_h^{-1} = (A^{-1}M)$. This preconditioned system $\mathcal{F} \circ \mathcal{A}^{-1}$ has rapidly decaying eigenvalues. Therefore, we follow [14, 20] in constructing a low-rank approximation of the prior-preconditioned misfit part of the Hessian, i.e., $\mathcal{H}_h^{\text{misfit}}(\mathbf{w}) := (\mathcal{F}_h \circ \mathcal{A}_h^{-1})^* W^{1/2} \Gamma_{\text{noise}}^{-1} W^{1/2} (\mathcal{F}_h \circ \mathcal{A}_h^{-1})$ by solving the symmetric eigenvalue problem [13, 15]:

$$\mathcal{H}_h^{\text{misfit}}(\mathbf{w})v_i = \lambda_i M \Gamma_{\text{prior,h}}^{-1} v_i = \lambda_i R v_i$$

for an orthogonal basis $V_r = (v_1, \dots, v_r) \in \mathbb{R}^{n_{\text{dof}} \times r}$ and $\lambda_1 \geq \dots \geq \lambda_r$ using the scalar product induced by $M \Gamma_{\text{prior,h}}^{-1}$, i.e., $\langle a, b \rangle_{M \Gamma_{\text{prior,h}}^{-1}} := a^T M \Gamma_{\text{prior,h}}^{-1} b$. Applying the Sherman–Morrison–Woodbury formula, we write:

$$\mathcal{A}_h (\mathcal{H}_h(\mathbf{w}))^{-1} \mathcal{A}_h = (\tilde{\mathcal{H}}_h^{\text{misfit}}(\mathbf{w}) - I)^{-1} \approx I + V_r D_r(\mathbf{w}) V_r^T, \quad (10)$$

where $D_r = \text{diag}(\lambda_1/(1 + \lambda_1), \dots, \lambda_r/(1 + \lambda_r))$ is a low-rank approximation of the Hessian. Detailed information on this approach can be found in [2] and [14]. Using this approximation, the solution of Eq. (8) can be determined with a preconditioned Newton-CG method, see [19, 20].

3.3. SPARSIFICATION OF SENSOR LAYOUTS AND OPTIMALITY CRITERIA

Obviously, the trace of the posterior covariance is minimized when every sensor weight is set to 1, which corresponds to using all available information to reduce the level of uncertainty. Thus, to derive a sparse sensor configuration, a penalty term must be introduced into the optimization problem (Eq. (6)). A common choice for the regularization term in Eq. (6) is the ℓ_1 -norm, which leads to a convex minimization problem with a unique minimizer. Specifically, the regularization term is defined as:

$$\mathcal{R}(\mathbf{w}) := \alpha \|\mathbf{w}\|_1 = \alpha \mathbf{1}^\top \mathbf{w}, \quad (11)$$

where $\alpha > 0$ is the regularization parameter and $\mathbf{1}$ is a vector of ones. Finally, a binary sensor configuration $\{0, 1\}^q$ is obtained by considering only sensor locations with weights above a chosen threshold.

It should be noted that the resulting binary vector is not the optimal binary weight vector. Obtaining such an optimal solution requires solving the ℓ_0 -regularized problem considered in [2]. The key idea of that approach is to employ a regularization term based on the number of nonzero entries in \mathbf{w} , i.e.,

$$\mathcal{R}(\mathbf{w}) := \alpha \sum_{w_i \neq 0} 1.$$

To approximate this discontinuous penalty, a sequence of strictly convex functions $f_j(\mathbf{w})$ is introduced such that $f_j(\mathbf{w}) \rightarrow \sum_{w_i \neq 0} 1$ as $j \rightarrow \infty$. At each iteration, the optimization problem

$$\min_{\mathbf{w} \in \mathbb{W}} \left(\Phi[\Gamma_{\text{post}}(\mathbf{w})] + f_j(\mathbf{w}) \right)$$

is solved and the optimizer obtained at iteration j is used as the initial value for iteration $j+1$. While effective, this procedure entails substantially higher computational cost. Moreover, as observed empirically in [2, Subsec. 6.1.4], the sensor configurations obtained using this ℓ_0 -approximation strategy offer only marginal improvement compared with those produced by ℓ_1 -regularization followed by a simple thresholding step. Since a highly efficient method is required for the sensor steering presented in Subsec. 5.3, we therefore employ an ℓ_1 -regularization approach.

The A-optimality criterion for sensor placement minimizes the integrated point-wise posterior variance in linear Bayesian inverse problems: $\int_{\Omega} \text{Var}[m(x)] dx = \text{tr}(\Gamma_{\text{post}}(\mathbf{w}))$. To compute the discrete variance field $\sigma^2(x) := \text{Var}[m(x)]$, one can extract the diagonal of the inverse Hessian matrix in the finite element basis, assigning each diagonal entry to its corresponding node. However, this exact

calculation is computationally expensive ($O(n_{\text{dof}})$), so the reduced-order model (Eq. (10)) is used for visualization. Figure 4b shows the variance field $\sigma^2(x)$ of m_{map} obtained with a C-optimal sparse sensor layout.




FIG. 4. Region of interest P defining the QoI (a) and point-wise variance $\sigma^2(x)$ of m_{map} obtained with five selected sensors (green spheres (b)).

Moreover, the prior covariance is computed using approximate random sampling [20]. Finally, computing the trace of the inverse Hessian is costly. Therefore, we relax A-optimality to a C-optimal design, which requires only the evaluation of the Hessian's action on a fixed vector $c \in \mathbb{R}^{n_{\text{dof}}}$ [1]. In this case, the design criterion becomes $\Phi(\mathbf{w}) = c^T \Gamma_{\text{post}}(\mathbf{w})c$.

4. GOAL-ORIENTATION AND SENSOR STEERING

4.1. GOAL-ORIENTED OPTIMAL EXPERIMENTAL DESIGN

In the next step, the presented method is adjusted to achieve goal-oriented C-optimal experimental designs for stationary sensors. To this end, an operator $\mathcal{P} : D \rightarrow \mathbb{R}$ is first selected so that the design, namely the sensor placement, is optimized to observe the initial conditions in a region $P \subset \Omega$, as shown in Fig. 4. This simplification via forward uncertainty quantification, defined by the operator \mathcal{P} , leads to a C-optimal design for the problem in Eq. (7), as described in [1].

The indicator function of a subset, that is the function $\mathbf{1}_A : \Omega \rightarrow \{0, 1\}$, which for a given subset A of Ω , takes the value 1 at points in A and the value 0 at points outside of A , is used to define our operator \mathcal{P} . If an optimal design for the initial condition m is desired, the operator \mathcal{P} does not depend on the solution u , or more specifically, it does not depend on the operator \mathcal{F} . Concretely, the operator for the QoI is given by $\mathcal{P}(m) := \langle m, \mathbf{1}_P \rangle_{L^2(\Omega)} = \int_P m(\mathbf{x}) dx$. By

identifying the dual space of $L^2(\Omega)$ and defining the adjoint operator, it follows that $\mathcal{P}^* = \mathbf{1}_P$, and the design function for the optimal experimental design takes the form $\text{tr}[\Gamma_{\text{post}}(\mathbf{w})] = \langle \mathbf{1}_P, \Gamma_{\text{post}}(\mathbf{w})\mathbf{1}_P \rangle_{L^2(\Omega)}$. In a finite element setting,

the function $\mathbf{1}_P$ is represented by a vector $c \in \mathbb{R}^{n_{\text{dof}}}$ via the standard projection, i.e., $\mathbf{1}_P - c_h \perp \mathcal{V}_h$ into the finite element space $\mathcal{V}_h = \{\phi_1, \dots, \phi_{n_{\text{dof}}}\}$. According to [1, Subsec. 4.1], the design function is given by $\Phi(\mathbf{w}) = c_h^T M \mathcal{H}_h^{-1}(\mathbf{w}) c_h$. Together with a suitable regularization term, e.g., $\alpha \|\mathbf{w}\|_{L^1}$, an optimal design can be obtained by minimizing the objective $c_h^T M \mathcal{H}_h^{-1}(\mathbf{w}) c_h + \alpha \|\mathbf{w}\|_{L^1}$ which is illustrated in Fig. 4. If this procedure is generalized to determine an optimal sensor placement for the contaminant concentration over a specific spatial region and time interval, the operator \mathcal{P} must be extended accordingly. Specifically, by defining \mathcal{P} over a subset of space-time, e.g., $[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P$, the operator \mathcal{P} is then constructed as:

$$\mathcal{P}(m) = \int_0^T \int_{\Omega} \mathbf{1}_{[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P} \mathcal{K}(m)(t, \mathbf{x}) dt d\mathbf{x} = \int_{T_0^{\text{QoI}}}^{T_{\text{final}}^{\text{QoI}}} \int_P u(t, \mathbf{x}) dt d\mathbf{x}.$$

Again, the adjoint of \mathcal{P} is needed. For this, we calculate

$$\begin{aligned} \mathcal{P}^*(1)(\widehat{m}) &= \langle \mathcal{K}(\widehat{m}), \mathbf{1}_{[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P} \rangle_{L^2(\Omega)} \\ &= \langle \widehat{m}, \mathcal{K}^*(\mathbf{1}_{[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P}) \rangle_{L^2(\Omega)}, \end{aligned}$$

and therefore the map c is obtained by $c := \mathcal{P}^*(1) = \mathcal{K}^*(\mathbf{1}_{[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P})$, since \widehat{m} was arbitrary. More precisely, the map c satisfies the following PDE:

$$\begin{aligned} -c_t - \kappa \Delta c - \text{div}(c\mathbf{v}) &= \mathbf{1}_{[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P} && \text{in } (0, T) \times \Omega, \\ (\mathbf{v}c + \kappa \nabla c) \cdot \mathbf{n} &= 0 && \text{in } (0, T) \times (\Gamma_+ \cup \Gamma_0), \\ c &= 0 && \text{in } (0, T) \times \Gamma_-, \\ c(T, \cdot) &= 0 && \text{in } \Omega. \end{aligned} \tag{12}$$

In Fig. 5a an example of a constant-over-time QoI $\mathbf{1}_{[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P}$ is shown. The corresponding map $c(t=0)$ is visualized in Fig. 5b. Returning to the definition, the objective function for the goal-oriented sensor design reads:

$$\begin{aligned} \Phi(\mathbf{w}) &= \mathcal{P} \mathcal{H}^{-1}(\mathbf{w}) \mathcal{P}^*(1) = \mathcal{P}(\mathcal{H}^{-1}(\mathbf{w})c) \\ &= \langle \mathcal{K}(\mathcal{H}^{-1}(\mathbf{w})c), \mathbf{1}_{[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P} \rangle_{L^2(\Omega)} \\ &= \langle \mathcal{H}^{-1}(\mathbf{w})c, \mathcal{K}^*(\mathbf{1}_{[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P}) \rangle_{L^2(\Omega)} = \langle c, \mathcal{H}^{-1}(\mathbf{w})c \rangle_{L^2(\Omega)}. \end{aligned}$$

Thus, the time-dependent case is reduced to finding an optimal design for the initial conditions of the transported QoI c , which must be found and coincides with the stationary case.

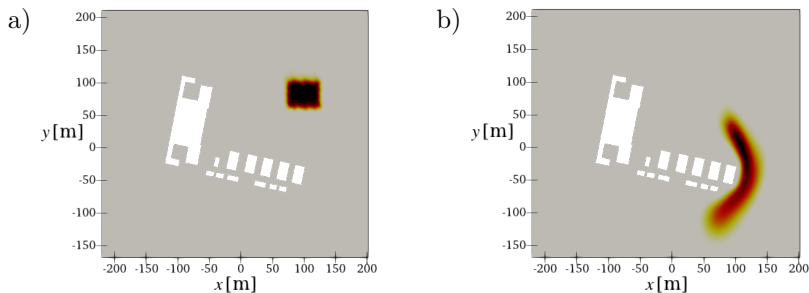


FIG. 5. Illustration of an example QoI $\mathbf{1}[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P$ (kept constant over time (a)) and solution $\mathcal{K}^*(\mathbf{1}[T_0^{\text{QoI}}, T_{\text{final}}^{\text{QoI}}] \times P)(t = 0, \cdot)$ of transport problem induced by adjoint operator \mathcal{K}^* (b).

For the numerical evaluation of this objective function and its gradient, which are required to minimize the regularized design function using the limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm with bound constraints (L-BFGS-B) solver, we start again with the inverse low-rank approximation of the Hessian from Eq. (10) and proceed to compute the design function as follows:

$$\Phi(\mathbf{w}) = \langle c, \mathcal{H}^{-1}(\mathbf{w})c \rangle_{L^2} \approx c_h^T \mathcal{A}_h^{-1} (I - V_r D_r V_r^T) \mathcal{A}_h^{-1} c_h.$$

For a shorter notation, we introduce $\hat{q}_h := (I - V_r D_r V_r^T) \mathcal{A}_h^{-1} c_h$, $q_h := \mathcal{A}_h^{-1} \hat{q}_h$ and obtain $\Phi(\mathbf{w}) \approx c_h^T q_h$. Hence, the calculation of the trace reduces to a projection in the low-rank subspace and solutions of an elliptic problem \mathcal{A} , respectively \mathcal{A}_h , for which very fast solving strategies exist. To calculate the derivative, we follow [2] and conclude for this simplified case

$$\frac{\partial}{\partial \mathbf{w}_i} \Phi(\mathbf{w}) = (\mathcal{F}^i(q))^2 \approx (\mathcal{F}_h^i \circ \mathcal{A}_h^{-1}(\hat{q}))^2. \quad (13)$$

So, this calculation can be replaced by a surrogate model for the preconditioned forward operator. In principle, this procedure can be extended to a stronger optimality criteria, such as A- or D-optimal designs.

4.2. DYNAMIC SENSOR STEERING BASED ON GOAL-ORIENTED OPTIMAL SENSOR PLACEMENT

A method will now be presented which is capable of dynamically controlling a sensor in such a way that a greater knowledge of the actual contaminant concentration can be generated. We assume that some knowledge about the concentration is already available from certain stationary sensors, i.e., the true contaminant field already possesses an appreciable concentration at these sensor locations, so that the inverse problem can be solved. This scenario can be seen in Fig. 6. We then define the QoI so that its center point is at the maximum of the reconstructed initial condition. The optimum design is then calculated on




FIG. 6. Dynamic sensor steering test case with stationary sensor (green) and mobile sensor (purple trajectory): a) true concentration field at $u(\cdot, t = 2.2\text{s})$ during the second steering time step, b) maximum-a-posteriori estimate m_{map} of the inverse problem and algorithmically selected zone of interest (red square), c) optimal sensor design, d) target position of the mobile sensor (red sphere), true concentration field $u(\cdot, t = 7\text{s})$ shown in background of (c) and (d) to visualize transport problem dynamics.

this basis and the sensor is steered to the position with the highest weight \mathbf{w} . After the next measurement is collected, the procedure is repeated from the beginning. This method is schematically shown in Fig. 7. In this way, we obtain a trajectory $\gamma : \{t_0^{\text{obs}}, \dots, t_s^{\text{obs}}\} \rightarrow \{\mathbf{x}_0^{\text{obs}}, \dots, \mathbf{x}_s^{\text{obs}}\}$ for the steered sensor.

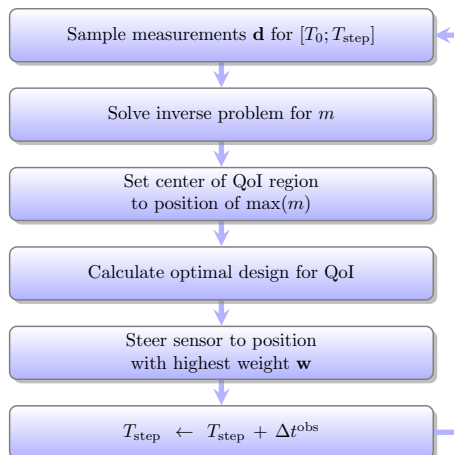


FIG. 7. Algorithm for dynamic sensor steering based on goal-oriented sensor placement [22, 23].

5. NUMERICAL RESULTS OF OPTIMAL EXPERIMENTAL DESIGNS

In order to simulate scenarios on real-world domains, we use a highly automated process for grid generation. Building imprints as obstacles for two-dimensional contaminant transport are imported directly from OpenStreetMap (OSM), and locally refined triangular meshes are generated for the region of interest [7]. The forward model is implemented with stabilized linear Lagrange finite elements in the software framework `FEniCs` [6]. The `FEniCs` extension `hIPPYlib` (Inverse Problem PYthon library [20]) is used for the implementation of the inverse problem.

In the three following inverse problem OED examples, we use the forward simulation of Eq. (1) illustrated in Fig. 2, as the ground truth. Two radially symmetric functions:

$$m_{\mathbf{x}_s}(\mathbf{x}; \mathbf{x}_s, r) = \min \{0.5, \exp(-\ln(\epsilon) \|\mathbf{x} - \mathbf{x}_s\|_2^2 / r^2)\}, \quad \epsilon = 0.001,$$

describe the initial concentration field defined as:

$$\begin{aligned} u_0(\mathbf{x}) &= m_{\mathbf{x}_s}(\mathbf{x}; \mathbf{x}_s = [-100 \text{ m}, -80 \text{ m}], r = 25 \text{ m}) \\ &+ m_{\mathbf{x}_s}(\mathbf{x}; \mathbf{x}_s = [75 \text{ m}, -80 \text{ m}], r = 25 \text{ m}). \end{aligned} \quad (14)$$

The initial concentration field is transported by the vector field \mathbf{v} . For the considered test cases, we estimate the stationary wind vector field as the solution of the incompressible Navier–Stokes equations with wind entering the given geometry from the south at $\mathbf{v} = 10 \text{ m s}^{-1}$. This condition is realized using a Dirichlet boundary condition. In the inner boundaries that represent the imprints of the buildings, a no-slip condition is applied. The remaining edges correspond to free boundary conditions. For the chosen Reynolds number of 50, we obtain the laminar wind field visualized in Fig. 2. Moreover, the diffusion coefficient is selected as $\kappa = 1 \text{ m}^2 \text{ s}^{-1}$, resulting in a transport problem with a moderate Peclet number. Finally, the time step size for the implicit Euler time-stepping scheme is set to 0.05 s. In the parametrization of the prior, the constants were chosen as $\eta = 8$ and $\gamma = 800$, yielding the operator $\mathcal{A} := 8I - 800\Delta$.

In order to make this problem computationally feasible, reduced-order models (ROMs) of the forward and adjoint operators are derived. Considering the forward operator $\mathcal{F}_h : \mathbb{R}^{n_{\text{dof}}} \rightarrow \mathbb{R}^q$, it is observed that it constitutes a linear mapping from a high-dimensional to a lower-dimensional space. Thus, a singular value decomposition is performed to construct a ROM; see also [13, 15]. The decomposition provides singular values $\lambda_1 \geq \dots \geq \lambda_r$, an L^2 -orthogonal basis $U_r = (u_1, \dots, u_r) \in \mathbb{R}^{n_{\text{dof}} \times r}$ and an orthogonal basis $V_r = (v_1, \dots, v_r) \in \mathbb{R}^{q \times r}$. During the online phase, for example when the precomputed ROM is used for sensor steering, the selected initial condition is projected so that only matrix-vector multiplications are required:

$$\mathcal{F}_h(m_h) \approx V_r D_r U_r M m_h,$$

where $m_h \in \mathbb{R}^{n_{\text{dof}}}$ and $D_r = \text{diag}(\lambda_1, \dots, \lambda_r)$.

As discussed in Eq. (9) and Eq. (13), it is also viable to construct a ROM directly for the preconditioned forward operator. The singular values of the operators \mathcal{F}_h and $\mathcal{F}_h \circ \mathcal{A}$ are compared in Fig. 8a. It is observed that the singular values of the preconditioned operator decay faster and therefore the number of computations that have to be performed to construct a ROM with acceptable accuracy is reduced. Fig. 8b shows that the ROM approximates the forward operator fairly well. Furthermore, the computed singular value decomposition is reused to approximate the adjoint operator $\mathcal{F}_h^* : \mathbb{R}^q \rightarrow \mathbb{R}^{n_{\text{dof}}}$ with a ROM as well, namely, $\mathcal{F}_h^*(y) \approx M U_r D_r V_r y$, for $y \in \mathbb{R}^q$.

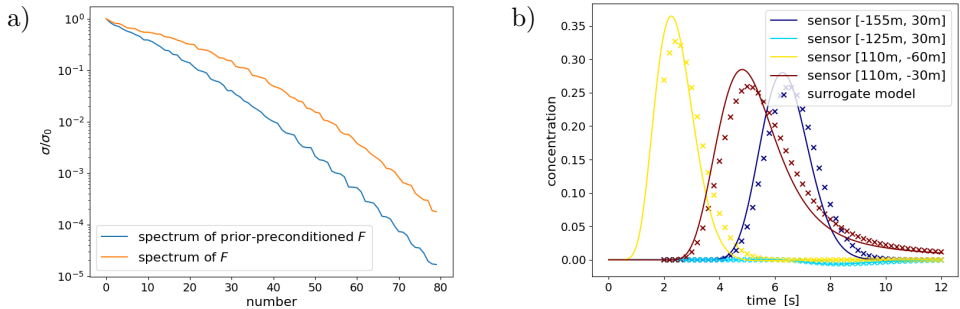


FIG. 8. Reduced-order modeling. Decay of singular values of the \mathcal{F}_h and of preconditioned $\mathcal{F}_h \circ \mathcal{A}$ (a), comparison of the ROM with the forward model \mathcal{F}_h evaluated at the sensor positions (b).

To assess the benefits of the reduced model in the context of forward evaluations, we first performed a single evaluation of the full-order model, which required approximately 1 s on our hardware. To simulate a scenario relevant to optimal sensor placement, the reduced model was used to evaluate the full-order model at 96 spatial positions over 90 time steps, resulting in a total of $q = 8640$ measurements. In total, we computed 200 spectral values. On average, each evaluation of the reduced model took 6.25 ms, yielding a relative speedup of approximately $160\times$ compared to the full-order model.

5.1. OED 1. SENSOR CONFIGURATION TO RECONSTRUCT INITIAL CONDITION IN A CRITICAL AREA

As the first example of a goal-oriented optimal experimental design, we address the problem of identifying an optimal sensor layout to recover the initial condition in a defined subset of the computational domain $P_1 := \{(x, y) \in \Omega \mid 75 \leq x \leq 125, -100 \leq y \leq -60\}$. In a practical application, P_1 might represent a critical area of a chemical plant site where hazardous material is stored.

The inverse problem is posed under the assumption that only stationary sensors are used. These sensors sample the concentration at a rate of 5 Hz, beginning at $T_0 = 2$ s. Measurements taken after 12 s are not taken into account. A noise variance of $\sigma^2 = (0.005)^2$ is assumed, resulting in a signal-to-noise ratio of approximately $\text{SNR} \approx 100$. Moreover, a regularization parameter of $\alpha = 0.1$ is applied to obtain a sparse sensor configuration, see Eq. (11). The selected domain where the QoI is inferred is indicated in Fig. 4a, along with the optimal sensor configuration (Fig. 4b). Moreover, the point-wise variance, which represents the uncertainty in the reconstruction, is also illustrated in Fig. 4. The solution to the inverse problem represented by m_{map} is visualized in Fig. 9. The numerical result demonstrates a reconstruction quality in P_1 comparable to that achieved using the full configuration with 96 sensors (Fig. 3), despite using five optimally selected sensors in OED 1.

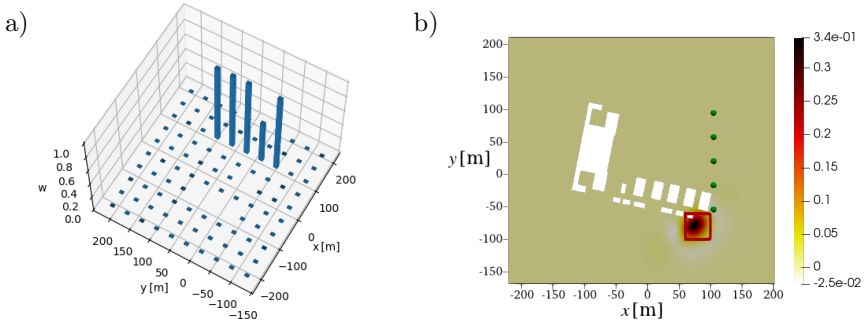


FIG. 9. OED 1. Weights of the optimal sensor configuration for monitoring P_1 (a) and reconstructed initial condition ((b), ground truth shown in Fig. 2).

5.2. OED 2. CONFIGURATION TO MONITOR CONCENTRATION EVOLUTION IN A CRITICAL AREA

In the second scenario, our aim is to secure a specific area for a given time period. To achieve this, the QoI is defined to depend on the state u . To create a meaningful scenario, we shift $P_2 := \{(x, y) \in \Omega \mid 75 \leq x \leq 125, 60 \leq y \leq 100\}$ upwards. The goal of OED 2 is to ensure that concentration values can be predicted correctly in region P_2 during the time interval from 5 s to 12 s. This setup results in the following operator:

$$\mathcal{P}_2(m) = \int_{T_0^{\text{QoI}}=5\text{ s}}^{T_{\text{final}}^{\text{QoI}}=12\text{ s}} \int_{P_2} \mathcal{K}(m)(t, \mathbf{x}) dt d\mathbf{x} = \int_{T_0^{\text{QoI}}=5\text{ s}}^{T_{\text{final}}^{\text{QoI}}=12\text{ s}} \int_{P_2} u(t, \mathbf{x}) dt d\mathbf{x}.$$

The sensor weights \mathbf{w} in Fig. 10 are computed using a regularization parameter of $\alpha = 1.0$. In this case as well, the source relevant to the QoI is reconstructed accurately using the optimized sensor configuration. The reconstructed initial

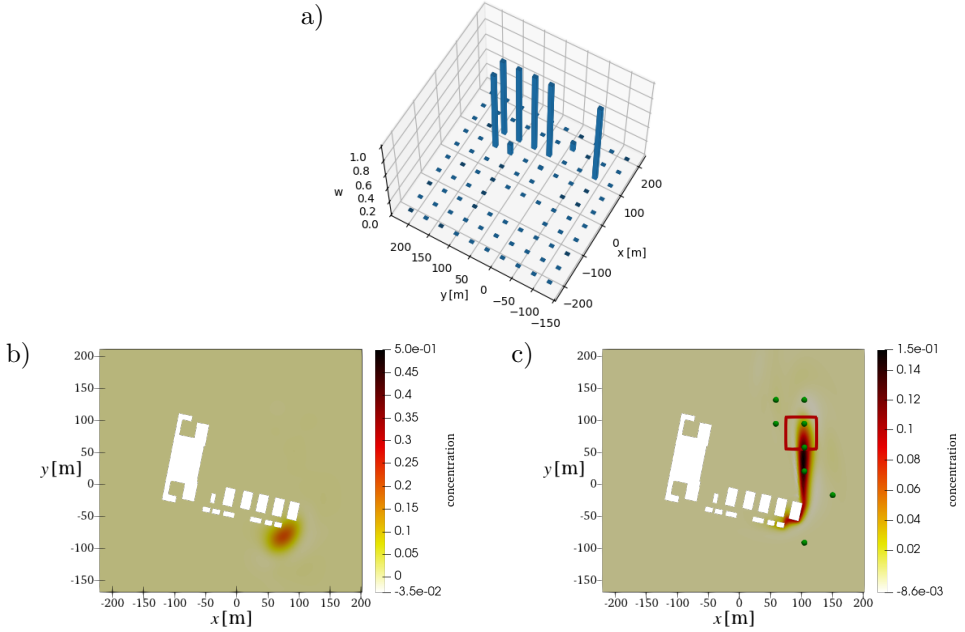


FIG. 10. OED 2. Weights of the optimal configuration for monitoring the concentration evolution in \mathcal{P}_2 (a), reconstructed initial condition (b), predicted concentration for $T = 10$ s (c).

condition and the corresponding prediction are shown in Fig. 10. As illustrated in Fig. 11, the reduction in uncertainty is concentrated primarily in the region relevant to QoI compared with the complete sensor configuration.

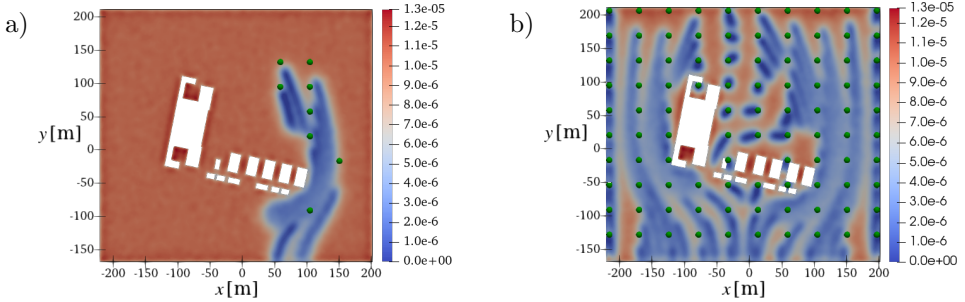


FIG. 11. OED 2. Point-wise variance σ^2 as a measure of uncertainty in the inferred parameter for the optimal configuration for \mathcal{P}_2 (a) compared to the full sensor grid (b).

5.3. OED 3. DYNAMIC SENSOR STEERING FOR SOURCE IDENTIFICATION

Finally, the sensor steering method described in Subsec. 4.2 is tested in a numerical application case with $\kappa = 10 \text{ m}^2 \text{ s}^{-1}$. To steer the sensor, a much finer sensor grid consisting of 1511 possible sensor locations is used. In this setup, the moving sensor is allowed to take one step on this grid per cycle, which

needs 0.2s and thus corresponds to a measurement frequency of 5 Hz. This corresponds to a speed of approximately 40 ms^{-1} for the moving sensor. To demonstrate the capabilities of the sensor steering approach, we placed a single stationary sensor just behind one of the obstacles. However, due to the transport characteristics in this region, information solely from this stationary sensor results in an inaccurate reconstruction of the source, which grossly underestimates the degree of contamination further from the buildings. In addition to the stationary sensor, measurements from a mobile sensor are available. The measurement process begins at time $T_0 = 2 \text{ s}$, with data collected at a frequency of 5 Hz. The state u at $T_{\text{step}} = T_0 = 2 \text{ s}$ is shown in Fig. 12a. At this point, the stationary sensor receives very limited information and is unable to provide an accurate source estimate. However, computing the optimal sensor design based

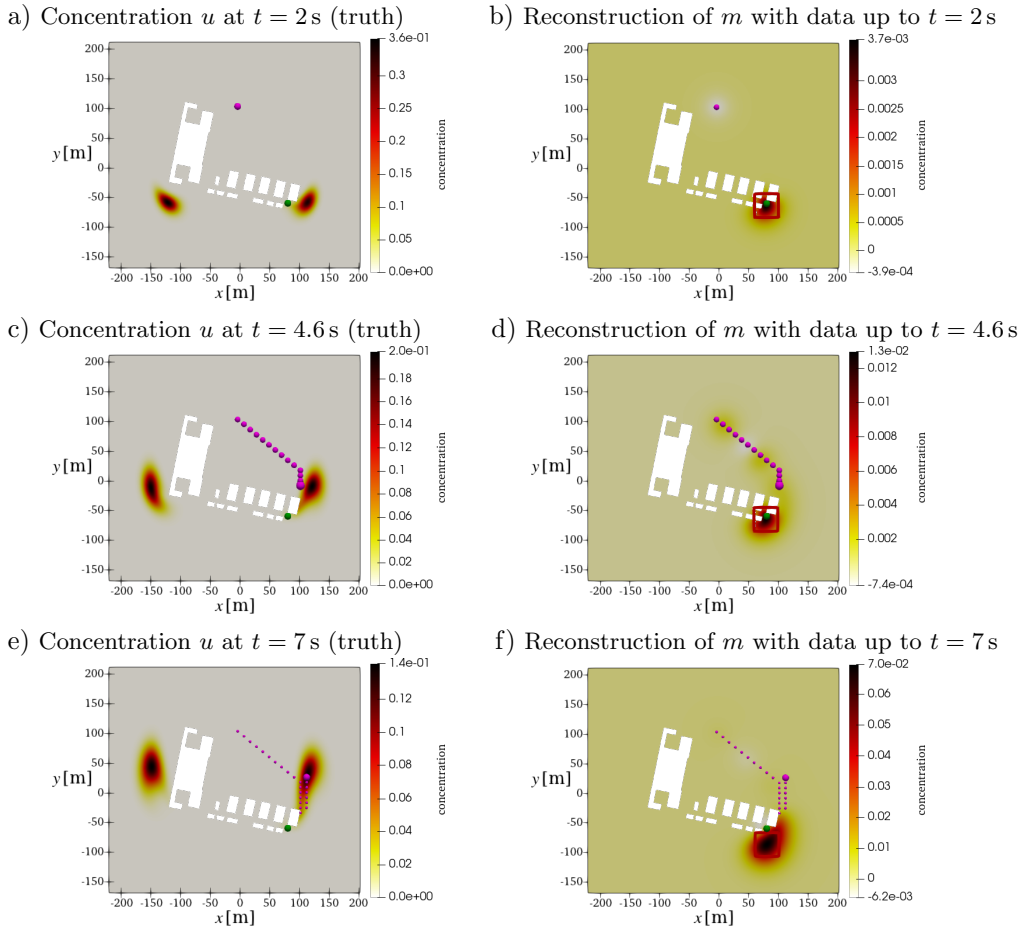


FIG. 12. OED 3. Data fusion of stationary sensor (marked in green) and mobile sensor (marked in purple). Mobile sensor steered to maximize information gain (trajectory marked with small purple spheres).

on the current QoI, defined as the integral over a square measuring 40 m on each side, centered on the maximal point of the reconstructed initial condition, yields favorable estimates for informative measurement positions. The mobile sensor is then directed toward the location associated with the highest weight \mathbf{w} , as determined by the C-optimal design criterion, computed over the observation period $[T_{\text{step}}, T_{\text{step}} + 2 \text{ s}]$ with the same sampling rate of 5 Hz, wherein we take as sensor weights for times prior to T_{step} , the actual past locations of the sensor. In the subsequent time steps, illustrated in Fig. 12c, the sensor continues to move toward regions of increasing concentration. In Fig. 12d, at $T_{\text{step}} = 4.6 \text{ s}$, the mobile sensor reaches the core of the contaminant. Finally, Fig. 12e and Fig. 12f demonstrate that the mobile sensor continues to accurately track the contaminant in further time steps. Comparing the performance of the stationary sensor alone with the combination of a stationary sensor and dynamically steered one, we find, as depicted in Fig. 13, that the mobile sensor achieves substantially improved reconstruction accuracy after just 7 s, whereas the stationary sensor fails to produce a reliable estimate even after 12 s.

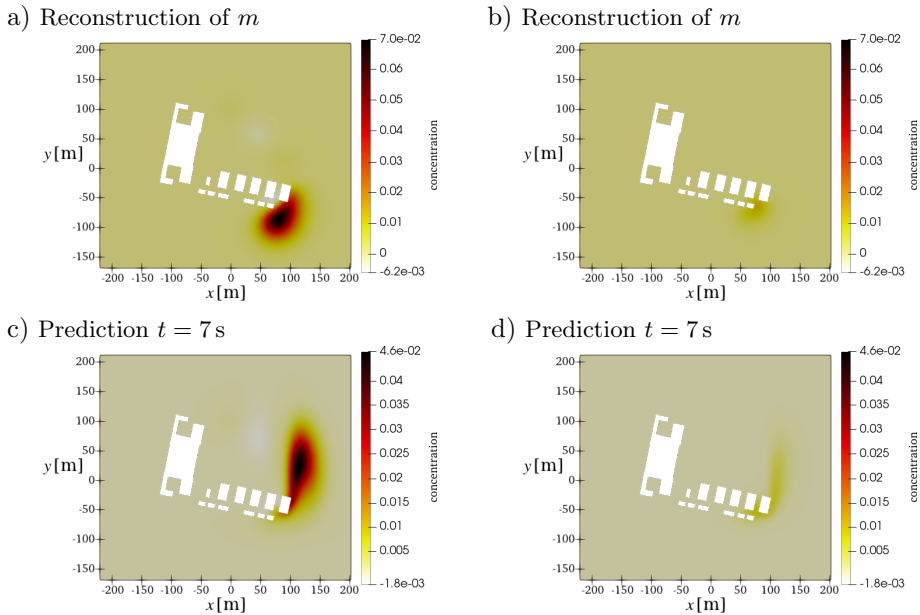


FIG. 13. OED 3. Comparison of predictions with mobile sensor and stationary sensor (a,c) and stationary sensor alone (b,d).

6. CONCLUSION AND OUTLOOK

This paper investigated a novel approach for goal-oriented optimal static sensor placement and dynamical sensor steering for PDE-constrained problems. Adopting previous work by WOGGRIN *et al.* [22, 23] on dynamic sensor steering,

we leveraged a Bayesian approach for the solution of the inverse problem, accelerated by offline low-rank approximations of the Hessian matrix and an online preconditioned inexact Newton-CG method. The resulting framework was then applied to a more complex geometry extracted from real-world map data. We showcased the strengths of the proposed workflow in three application cases in the field of airborne contaminant transport. In the first example, we derived an optimal placement of stationary sensors to recover the initial condition inside a spatially constrained rectangular region. The results showed that our proposed method only required five sensors to reconstruct the initial condition locally with accuracy comparable to the full configuration of 96 sensors. In our second example, we extended the QoI in the sense that a RoI was monitored not only at a specific time instance, but also over a fixed time period. From a practical point of view, this corresponds to the goal of securing a specific area for a given time period. Using only eight optimally placed sensors, the evolution of the concentration was accurately reconstructed and the uncertainty was minimized in the area of interest. Lastly, we investigated a dynamic sensor steering problem. Here, we showed that while we were still able to roughly predict the general shape of the initial condition with an unfavorably placed stationary sensor, adding a mobile sensor significantly improved the reconstruction while simultaneously reducing the required measurement time to one-third compared to the stationary case. This proves that the presented method is able to successfully handle the complexity of moving sensors and steer them to achieve fast and reliable reconstruction of unknown (in practice) initial condition.

While we believe this work to be an important step toward the optimal steering of unmanned sensor platforms in crisis situations, there still remain several points for improvement and further investigation of the proposed algorithm is needed. One potential area for improvement concerns the mathematical formulation and solution of the inverse problem. In many applications, it is reasonable to assume that the initial condition is sparse. Incorporating this additional knowledge into the solution procedure is expected to speed up computations and further improve real-time capabilities of the method [17]. Moreover, we plan to extend the sensor steering methodology to a reinforcement learning approach, where the position and size of the QoI at each step is determined by an agent that was previously trained based on trial-and-error interactions with the forward model [12].

ACKNOWLEDGEMENTS

Daniel Wolff and Alexander Popp gratefully acknowledge funding by Digitalization and Technology Research Center of the Bundeswehr (dtec.bw) (project RISK.twin). dtec.bw is funded by the European Union – NextGenerationEU.

REFERENCES

1. ALEXANDERIAN A., Optimal experimental design for infinite-dimensional Bayesian inverse problems governed by PDEs: A review, *Inverse Problems*, **37**(4): 043001, 2021, <https://doi.org/10.1088/1361-6420/abe10c>.
2. ALEXANDERIAN A., PETRA N., STADLER G., GHATTAS O., A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0 -sparsification, *SIAM Journal on Scientific Computing*, **36**(5): A2122–A2148, 2014, <https://doi.org/10.1137/130933381>.
3. ALEXANDERIAN A., SAIBABA A.K., Efficient D-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems, *SIAM Journal on Scientific Computing*, **40**(5): A2956–A2985, 2018, <https://doi.org/10.1137/17M115712X>.
4. ALT H.W., *Lineare Funktionalanalysis*, Springer, Berlin, Heidelberg, 2012, <https://doi.org/10.1007/978-3-642-22261-0>.
5. ATTIA A., ALEXANDERIAN A., SAIBABA A.K., Goal-oriented optimal design of experiments for large-scale Bayesian linear inverse problems, *Inverse Problems*, **34**(9): 095009, 2018, <https://doi.org/10.1088/1361-6420/aad210>.
6. BARATTA I.A. *et al.*, DOLFINx: The next generation FEniCS problem solving environment, *Zenodo*, 2023, <https://doi.org/10.5281/zenodo.10447666>.
7. BONARI J., KÜHN L., VON DANWITZ M., POPP A., Towards real-time urban physics simulations with digital twins, [in:] *2024 28th International Symposium on Distributed Simulation and Real Time Applications (DS-RT)*, Urbino, Italy, pp. 18–25, 2024, <https://doi.org/10.1109/DS-RT62209.2024.00013>.
8. CVETKOVIĆ N., LIE H.C., BANSAL H., VEROY K., Choosing observation operators to mitigate model error in Bayesian inverse problems, *SIAM/ASA Journal on Uncertainty Quantification*, **12**(3): 723–758, 2024, <https://doi.org/10.1137/23M1602140>.
9. DAON Y., STADLER G., Mitigating the influence of the boundary on PDE-based covariance operators, *Inverse Problems and Imaging*, **12**(5): 1083–1102, 2018, <https://doi.org/10.3934/ipi.2018045>.
10. ELMAN H.C., SU T., A low-rank solver for the stochastic unsteady Navier–Stokes problem, *Computer Methods in Applied Mechanics and Engineering*, **364**: 112948, 2020, <https://doi.org/10.1016/j.cma.2020.112948>.
11. EVANS L.C., *Partial Differential Equations*, Graduate Studies in Mathematics, Vol. 19, American Mathematical Society, Providence, Rhode Island, 2nd ed., 2022.
12. FRICKE C., WOLFF D., KEMMERLING M., ELGETI S., Investigation of reinforcement learning for shape optimization of 2D profile extrusion die geometries, *Advances in Computational Science and Engineering*, **1**(1): 1–35, 2023, <https://doi.org/10.3934/acse.2023001>.
13. HALKO N., MARTINSSON P.G., TROPP J.A., Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Review*, **53**(2): 217–288, 2011, <https://doi.org/10.1137/090771806>.
14. ISAAC T., PETRA N., STADLER G., GHATTAS O., Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet, *Journal of Computational Physics*, **296**: 348–368, 2015, <https://doi.org/10.1016/j.jcp.2015.04.047>.

15. LIBERTY E., WOOLFE F., MARTINSSON P.-G., ROKHLIN V., TYGERT M., Randomized algorithms for the low-rank approximation of matrices, *Proceedings of the National Academy of Sciences*, **104**(51): 20167–20172, 2007, <https://doi.org/10.1073/pnas.0709640104>.
16. PETRA N., STADLER G., *Model variational inverse problems governed by partial differential equations*, ICES REPORT 11-05, The Institute for Computational Engineering and Sciences, The University of Texas at Austin, 2011, <https://math.nyu.edu/~stadler/papers/PetraStadler11.pdf>.
17. PIEPER K., WALTER D., Linear convergence of accelerated conditional gradient algorithms in spaces of measures, *ESAIM: Control, Optimisation and Calculus of Variations*, **27**: 38, 2021, <https://doi.org/10.1051/cocv/2021042>.
18. SPANTINI A., CUI T., WILLCOX K., TENORIO L., MARZOUK Y., Goal-oriented optimal approximations of Bayesian linear inverse problems, *SIAM Journal on Scientific Computing*, **39**(5): S167–S196, 2017, <https://doi.org/10.1137/16M1082123>.
19. STEIHAUG T., Local and superlinear convergence for truncated iterated projections methods, *Mathematical Programming*, **27**(2): 176–190, 1983, <https://doi.org/10.1007/BF02591944>.
20. VILLA U., PETRA N., GHATTAS O., hIPPYlib: An extensible software framework for large-scale inverse problems governed by PDEs: Part I: Deterministic inversion and linearized Bayesian inference, *ACM Transactions on Mathematical Software*, **47**(2): 16, 2021, <https://doi.org/10.1145/3428447>.
21. VON DANWITZ M., BONARI J., FRANZ P., KÜHN L., MATTUSCHKA M., POPP A., Contaminant dispersion simulation in a digital twin framework for critical infrastructure protection, [in:] *9th European Congress on Computational Methods in Applied Sciences and Engineering*, CIMNE, 2024, <https://doi.org/10.23967/eccomas.2024.301>.
22. WOGRIN S., *Model Reduction for Dynamic Sensor Steering: A Bayesian Approach to Inverse Problems*, Master’s Thesis, Massachusetts Institute of Technology, Computation for Design and Optimization Program, 2008, <http://hdl.handle.net/1721.1/43739>.
23. WOGRIN S., SINGH A., ALLAIRE D., GHATTAS O., WILLCOX K., From data to decisions: A real-time measurement–inversion–prediction–steering framework for hazardous events and health monitoring, [in:] Darema F., Blasch E.P., Ravela S., Aved A.J. [Eds], *Handbook of Dynamic Data Driven Applications Systems*, pp. 195–227, Springer, Cham, 2023, https://doi.org/10.1007/978-3-031-27986-7_8.
24. WU K., CHEN P., GHATTAS O., An offline-online decomposition method for efficient linear Bayesian goal-oriented optimal experimental design: Application to optimal sensor placement, *SIAM Journal on Scientific Computing*, **45**(1): B57–B77, 2023, <https://doi.org/10.1137/21M1466542>.

*Received July 2, 2025; revised November 24, 2025; accepted January 18, 2026;
available online February 9, 2026; version of record April 20, 2026.*