

HYPER-DIFFERENTIAL SENSITIVITY ANALYSIS OF UNCERTAIN PARAMETERS IN PDE-CONSTRAINED OPTIMIZATION

Joseph Hart,^{1,*} Bart van Bloemen Waanders,¹ & Roland Hertzog²

¹*Optimization and Uncertainty Quantification, Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM 87123-1320*

²*Technical University Chemnitz, Faculty of Mathematics, 09107 Chemnitz, Germany*

*Address all correspondence to: Joseph Hart, E-mail: joshart@sandia.gov

Many problems in engineering and sciences require the solution of large scale optimization constrained by partial differential equations (PDEs). Though PDE-constrained optimization is itself challenging, most applications pose additional complexity, namely, uncertain parameters in the PDEs. Uncertainty quantification (UQ) is necessary to characterize, prioritize, and study the influence of these uncertain parameters. Sensitivity analysis, a classical tool in UQ, is frequently used to study the sensitivity of a model to uncertain parameters. In this article, we introduce “hyper-differential sensitivity analysis” which considers the sensitivity of the solution of a PDE-constrained optimization problem to uncertain parameters. Our approach is a goal-oriented analysis which may be viewed as a tool to complement other UQ methods in the service of decision making and robust design. We formally define hyper-differential sensitivity indices and highlight their relationship to the existing optimization and sensitivity analysis literatures. Assuming the presence of low rank structure in the parameter space, computational efficiency is achieved by leveraging a generalized singular value decomposition in conjunction with a randomized solver which converts the computational bottleneck of the algorithm into an embarrassingly parallel loop. Two multi-physics examples, consisting of nonlinear steady state control and transient linear inversion, demonstrate efficient identification of the uncertain parameters which have the greatest influence on the optimal solution.

KEY WORDS: *sensitivity analysis, PDE-constrained optimization, randomized linear algebra, low rank approximations*

1. INTRODUCTION

Many critical applications in science and engineering require the analysis of multi-physics phenomena across several spatial and temporal scales. For instance in material science, fusion energy, hydrocarbon extraction, and climate science, the ultimate goal is to solve large scale optimization problems while reconciling uncertainties that arise in constituent models, material properties, boundary conditions, initial conditions, and multi-physics interfaces. Although the desire is to apply modeling, uncertainty quantification, and optimization to arrive at robust solutions, the culmination of such analysis poses a formidable computational challenge rendering many algorithmic strategies ineffective.

In this article, we propose a sensitivity analysis framework to determine the importance of uncertain parameters in the context of optimal solutions. To aid in distinguishing our approach from existing sensitivity analysis techniques, we introduce the term “hyper-differential sensitivity analysis” or HDSA for short. We compute the Fréchet derivative of the solution of a PDE-constrained optimization problem with respect to uncertain parameters. Assuming a low rank structure in the parameter space, a truncated generalized singular value decomposition is employed to efficiently estimate sensitivity indices which enables an identification of inconsequential uncertain parameters and a prioritization of the uncertainties. By identifying low dimensional structure in high dimensional parameter spaces, HDSA is a goal-oriented analysis which considers the sensitivity of the optimal solution.

The complexities of PDE-constrained optimization consists of incorporating the discretization of PDEs as well as implementing complicated linear algebra constructs, including adjoints and Hessians. Furthermore, the large scale nature of these problems requires efficient computational implementation with scalable parallel linear algebra. Considerable research and development has been conducted and the interested reader is referred to a small sampling of the literature [2–14]. Although a range of algorithmic strategies are possible, we use standard methods to solve the underlying PDE-constrained optimization problem, consisting Newton-based solvers, trust region globalization, Tikhonov regularization, finite element discretization of the PDE constraints, and matrix-free operators. The focus of this paper is sensitivity analysis with respect to the solution of PDE-constrained problems and therefore inherits all the associated complexities.

Traditional sensitivity analysis can be divided into two subfields: local sensitivity analysis and global sensitivity analysis [15,16]. There are a plurality of methods within each subfield; we highlight one class of methods from each to provide background for this article. Local sensitivity analysis studies the influence of uncertain parameters on a quantity of interest (for instance, a functional of the PDE solution) at some fixed parameter value. Computing the derivatives of the quantity of interest with respect to the parameters is one measure of local sensitivity. A large derivative indicates that the quantity of interest is sensitive to the parameter. In the context of PDE-constrained optimization, finite difference, direct, and adjoint-based sensitivities are local sensitivities of the objective function with respect to the optimization variables, whereas HDSA is the local sensitivity of the optimal solution with respect to uncertain parameters. The limitation of local sensitivity analysis is that it is only valid in a neighborhood of the fixed parameter value. Global sensitivity analysis [17–21] seeks to alleviate this problem by varying the parameters over a set (typically with an associated probability measure) and measuring the importance of the parameters by averaging over this set. The approach of [19] connects local sensitivity analysis to global sensitivity analysis. Since the derivative is local in the sense that it depends on the user specifying a nominal parameter value, a global approach computes the expected value (in parameter space) of the squared derivative. Both local and global derivative-based (hyper-differential) sensitivity indices are defined in this article, though additional complexities arise in the context of HDSA which do not occur in the traditional sensitivity analysis framework.

Building on the work of Brandes and Griesse* [22] (and related work [23–29]), we define sensitivity indices, provide algorithmic developments, and implement software which enables analysis for large-scale optimization problems with high dimensional uncertain parameter spaces. The evaluation of these sensitivities requires an eigenvalue computation involving the optimality system. To that end, we have reformulated the eigenvalue problem and introduced the use of a randomized eigenvalue solver which allows for parallelization of the underlying matrix-vector products [30]. Our implementation is in C++ with parallel linear algebra constructs. The parallel randomized solver extends the use of the native parallelism for a second level of parallelism. Having developed an efficient computational framework, we introduce global sensitivity indices in the context of PDE-constrained optimization, demonstrate how they may be estimated, and highlight the challenges involved. Finally, we demonstrate our approach on the control of a thermal-fluid flow multi-physics problem and a source inversion problem constrained by Darcy flow and advection diffusion. The main contributions of this paper are: 1) the introduction of local and global hyper-differential sensitivity indices as tools to analyze the sensitivity of optimal solutions in the service of robust design and decision making, 2) the development of C++ software infrastructure to leverage state of the art (matrix free) PDE-constrained optimization and parallel linear algebra, and 3) the formulation of a symmetric generalized eigenvalue problem (to estimate sensitivities) and its numerical solution via randomized methods to achieve (nearly) embarrassingly parallel efficiency.

The article is organized as follows. We first define hyper-differential local and global sensitivities in Section 2. Our algorithmic contributions to accelerate the computation of local sensitivities are given in Section 3, followed by Section 4 which overviews our algorithms, provides an analysis of computational cost, and a strategy for interpreting the sensitivities. Our proposed framework is demonstrated through two applications in Section 5. We provide conclusions and highlight areas of future work in Section 6.

*R. Hertzog was formerly R. Griesse

2. HYPER-DIFFERENTIAL SENSITIVITY ANALYSIS FOR SOLUTIONS OF PDE-CONSTRAINED OPTIMIZATION PROBLEMS

This section provides necessary background and formally defines hyper-differential local and global sensitivities. The implicit function theorem forms the mathematical foundation for the proposed approach. The concept of sensitivities, defined through the implicit function theorem, are present in different portions of the optimization literature with multiple titles and in various contexts. For instance, it is called post optimality analysis, sensitivity analysis, or parametric programming in the operations research literature [31], sensitivity analysis, stability analysis, and perturbation analysis in the nonlinear programming community [32], and parametric sensitivity analysis in PDE-constrained optimal control [24,25]. We follow Brandes and Griesse [22] who introduced such sensitivity analysis for PDE-constrained optimization.

2.1 Local Sensitivity Analysis

Consider the PDE-constrained optimization problem

$$\begin{aligned} & \min_{u,z} J(u, z, \theta) \\ \text{s.t. } & c(u, z, \theta) = 0 \\ & u \in U, z \in Z \end{aligned} \tag{1}$$

where U is the state space, Z is the optimization variable space, $\theta \in \Theta$ are uncertain parameters which are fixed in the optimization problem, $J : U \times Z \times \Theta \rightarrow \mathbb{R}$ is an objective function, $c : U \times Z \times \Theta \rightarrow \Lambda^*$ is the weak form of a PDE, and Λ^* denotes the dual of Λ . It is assumed that U and Λ are reflexive Banach spaces, and Z and Θ are Hilbert spaces. The stronger assumptions on Z and Θ are to enable subsequent analysis with the generalized singular value decomposition (GSVD). The spaces U , Z , Θ , and Λ may be finite or infinite dimensional. In particular, the optimization variable space Z may correspond to Euclidean space or a function space in either a control, design, or inverse problem setting, and the parameter space Θ may represent a finite number of parameters or a functional representation (for instance, a spatially dependent parameter), in which case Θ is infinite dimensional. The PDE represented by c may be time dependent or in steady state. We assume that J and c are twice continuously differentiable with respect to (u, z, θ) and that the Fréchet derivative of c with respect to u is surjective.

Our goal is to develop a computational framework which may be used for large scale PDE systems with infinite (or large finite) dimensional parameter uncertainty. In particular, appealing to the use of HDSA for transient, nonlinear, and multi-physics systems with spatially and/or temporally dependent parameters where it may identify important structure embedded in complex physics.

To analyze (1) as an unconstrained problem, define the Lagrangian $\mathcal{L} : U \times Z \times \Lambda \times \Theta \rightarrow \mathbb{R}$ as

$$\mathcal{L}(u, z, \lambda, \theta) = J(u, z, \theta) + \langle \lambda, c(u, z, \theta) \rangle, \tag{2}$$

where $\lambda \in \Lambda$ is the (unique) Lagrange multiplier, in the context of PDE-constrained optimization it is called the adjoint state. Throughout the article the subscripts J_* and c_* are used to denote the Fréchet derivative of J and c with respect to $*$, respectively. After discretizing, J_* and c_* will denote the gradient of J and Jacobian of c , respectively.

We seek to perform derivative-based analysis, specifically the derivative of the optimal solution with respect to the parameters. Care must be taken since (1) may admit multiple local minima. The following result (Lemma 2.6 in [22]) is foundational for our sensitivity analysis.

Let (u_0, z_0) be a local minimum of (1) when $\theta = \theta_0$ and λ_0 be the unique adjoint state. Assuming the second order sufficient optimality condition holds, see [22] for details, there exists neighborhoods $\mathcal{N}(\theta_0) \subset \Theta$ and $\mathcal{N}(u_0, z_0, \lambda_0) \subset U \times Z \times \Lambda$ and a continuously differentiable function

$$\mathcal{F} : \mathcal{N}(\theta_0) \rightarrow \mathcal{N}(u_0, z_0, \lambda_0)$$

such that for all $\theta \in \mathcal{N}(\theta_0)$, $\mathcal{F}(\theta) = (u_{opt}(\theta), z_{opt}(\theta), \lambda_{opt}(\theta))$ is the unique stationary point of $\mathcal{L}(\cdot, \cdot, \cdot, \theta)$ in $\mathcal{N}(u_0, z_0, \lambda_0)$, that is,

$$\nabla \mathcal{L}_{(u, z, \lambda)}(\mathcal{F}(\theta), \theta) = 0.$$

The Fréchet derivative of \mathcal{F} at θ_0 is given by

$$\mathcal{F}'(\theta_0; z_0) = \mathcal{K}(\theta_0; z_0)^{-1} \mathcal{B}(\theta_0; z_0)$$

where

$$\mathcal{K}(\theta_0; z_0) = \begin{pmatrix} \mathcal{L}_{uu} & \mathcal{L}_{uz} & c_u^* \\ \mathcal{L}_{zu} & \mathcal{L}_{zz} & c_z^* \\ c_u & c_z & 0 \end{pmatrix} \quad \text{and} \quad \mathcal{B}(\theta_0; z_0) = - \begin{pmatrix} \mathcal{L}_{u, \theta} \\ \mathcal{L}_{z, \theta} \\ \mathcal{L}_{\lambda, \theta} \end{pmatrix}, \quad (3)$$

with the second derivatives of \mathcal{L} evaluated at $(u_0, z_0, \lambda_0, \theta_0)$; \star denotes the adjoint of an operator. We only write explicit dependence on $(\theta_0; z_0)$ to simplify notation. The operator \mathcal{K} is the Karush-Kuhn-Tucker (KKT) operator.

This gives a computable expression for the change in the optimal solution when the parameters $\theta_0 \in \Theta$ are perturbed. We emphasize that this is a local result in the sense that the optimization problem (1) may have many stationary points; \mathcal{F} is only defined in the neighborhood of a particular stationary point, $\mathcal{N}(u_0, z_0, \lambda_0)$, ensuring existence of the mapping and its derivative.

The directional derivative of optimal solution in the direction θ is given by the solution, (u, z, λ) , of the linear system

$$\mathcal{K}(\theta_0; z_0) \begin{pmatrix} u \\ z \\ \lambda \end{pmatrix} = \mathcal{B}(\theta_0; z_0) \theta. \quad (4)$$

To determine the sensitivity of the optimization variables to changes in θ , define the projection operator $\mathcal{P} : U \times Z \times \Lambda \rightarrow Z$ by

$$\mathcal{P} \begin{pmatrix} u \\ z \\ \lambda \end{pmatrix} = z.$$

Then the Fréchet derivative of the optimal solution, z , with respect to the parameters θ may be expressed as the linear operator $\mathcal{D}(\theta_0; z_0) : \Theta \rightarrow Z$,

$$\mathcal{D}(\theta_0; z_0) = \mathcal{P} \mathcal{K}(\theta_0; z_0)^{-1} \mathcal{B}(\theta_0; z_0). \quad (5)$$

Note that the sensitivity of the state, or a function of the state, adjoint, or combination of them may be considered by using a different \mathcal{P} . To simplify the presentation, this article focuses on the sensitivity with respect to the optimization variables z .

We introduce the hyper-differential local sensitivity function $\mathcal{S}(\theta_0; z_0) : \Theta \rightarrow \mathbb{R}$ which is defined by

$$\mathcal{S}(\theta_0; z_0) \phi = \left\| \mathcal{D}(\theta_0; z_0) \frac{\phi}{\|\phi\|} \right\| \quad \forall \phi \in \Theta. \quad (6)$$

The scalar $\mathcal{S}(\theta_0; z_0) \phi$ may be interpreted as the magnitude of the change in the optimal z when the parameters are perturbed in the direction ϕ . Throughout the article we will commonly refer to (6) as a local sensitivity for short, but adopt the formal title ‘‘hyper-differential local sensitivity’’ to distinguish this approach from other forms of derivative-based sensitivity analysis.

Example 2.1 illustrates the difference between a ‘‘traditional’’ sensitivity analysis approach versus hyper-differential sensitivity analysis. We argue that the latter is applicable to many engineering and science problems and while it is computationally intensive, the final sensitivity results are in the context of control, design, and inversion goals.

A traditional sensitivity approach fixes z to the optimal z_0 and analyzes the sensitivity of the state (or objective function) to changes in the parameters; there are no KKT solves involved. This approach suffers two drawbacks:

- The optimal z will change when the parameters change, so fixing the optimization variable may lead to sensitivity which occurs because of a poor z .
- There may be parameters for which the state u is sensitive but the optimal z is not, or vice versa. Since the solution of the optimization problem z is our end goal, fixing the z does not give the needed sensitivities.

The following example illustrates differences in these two approaches.

Example 2.1.

Consider the optimization problem

$$\begin{aligned} \min_{u,z} \quad & J(u,z) = (u - 2)^2 + .0005z^2 \\ \text{s.t.} \quad & u = \frac{1}{1 + e^{-\theta_1 z}} + \theta_2 \end{aligned} \quad (7)$$

where $\theta = (\theta_1, \theta_2)$ are uncertain parameters.

Let $u_{opt}(\theta)$, $z_{opt}(\theta)$, as a function of θ , be the optimal solution of (7). Then

$$\frac{\partial z_{opt}}{\partial \theta_1}(0.5, 0.5) = 9.99 \quad \text{and} \quad \frac{\partial z_{opt}}{\partial \theta_2}(0.5, 0.5) = 3.12.$$

If instead, (7) is solved with $\theta = (0.5, 0.5)$ and the derivative of

$$g(\theta_1, \theta_2) = \left(\frac{1}{1 + e^{-\theta_1 z_{opt}(0.5, 0.5)}} + \theta_2 - 2 \right)^2 + .0005z_{opt}(0.5, 0.5)^2$$

is computed with respect to (θ_1, θ_2) we get

$$\frac{\partial g}{\partial \theta_1}(0.5, 0.5) = 0.135 \quad \text{and} \quad \frac{\partial g}{\partial \theta_2}(0.5, 0.5) = 1.03.$$

Hence computing the sensitivity of z_{opt} to θ gives a different conclusion than the sensitivity of the objective function (at the optimal z) to θ . Figure 1 gives some intuition for this result. The constraint is plotted with u as a function of z . Curves are plotted for different values of θ and the solution of (7) for each fixed θ is given by the dot on the curve. In the left (right) panel $\theta_2 = 0.5$ ($\theta_1 = 0.5$) is fixed and θ_1 (θ_2) varies from 0.3 to 0.7. This demonstrates that as θ_1 varies, left panel, the optimal solution z varies significantly while the state u is kept nearly constant; whereas as θ_2 varies, right panel, z is nearly constant while u varies significantly.

In many applications, the parameter space Θ may be a product of sets corresponding to different physical parameters. For instance, there may be parameters corresponding to functions defined on the PDE's computational domain (such as a spatially distributed coefficient in the PDE), functions defined on the boundary (such as a non homogeneous Dirichlet boundary condition), or scalar parameters (such as constant coefficients in the PDE). It is useful to assign a scalar measure of sensitivity for each set of parameters. To do so, consider $\Theta = \Xi_1 \times \Xi_2 \times \dots \times \Xi_T$ as the product of T different parameter sets and define the local set sensitivity index as

$$\mathcal{S}_{\Xi_i} = \max_{\phi \in \Theta} \left\| \mathcal{D}(\theta_0; z_0) \Pi_{\Xi_i} \frac{\phi}{\|\phi\|} \right\| \quad i = 1, 2, \dots, T, \quad (8)$$

where $\Pi_{\Xi_i} : \Theta \rightarrow \Theta$ is the projection operator which maps $\phi = (\xi_1, \xi_2, \dots, \xi_T) \in \Xi_1 \times \Xi_2 \times \dots \times \Xi_T$ to $\Pi_{\Xi_i} \phi = (0, 0, \dots, 0, \xi_i, 0, \dots, 0) \in \Xi_1 \times \Xi_2 \times \dots \times \Xi_T$, i.e. it annihilates all parameter variability except those in Ξ_i . For complex multi-physics problems where T may be large (we consider $T = 5$ as large), the set of indices $\{S_{\Xi_i}\}_{i=1}^T$ may be easily tabulated or visualized to see the relative importance of the different parameter sets. For such complex problems, (6) provides detailed information but it may be difficult to compare the sensitivities of different parameters if, for instance, they have different time and/or space dependencies. The set sensitivity indices (8) provide a simple summary.

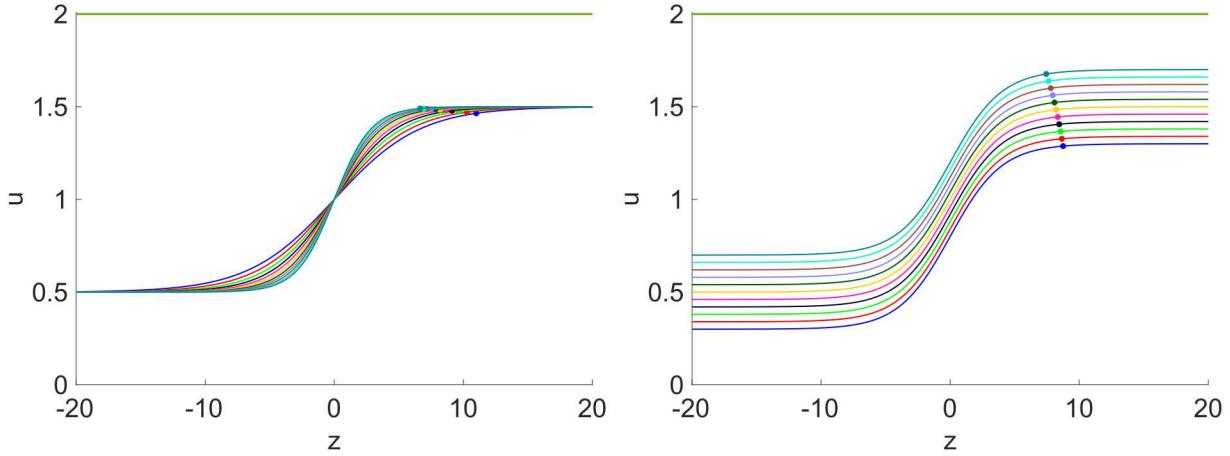


FIG. 1: Plot of the constraint u in (7) as a function of z for different values of θ . Left: varying θ_1 from 0.3 to 0.7 with $\theta_2 = 0.5$ fixed; right: θ_2 varying from 0.3 to 0.7 with $\theta_1 = 0.5$ fixed. Each curve corresponds to a different θ and each dot corresponds to the solution of (7) for that given θ . The green horizontal line is the target $u = 2$.

2.2 Global Sensitivity Analysis

Caution must be exercised when interpreting local sensitivities because of their dependence on the nominal parameter θ_0 . If θ is uncertain and the function \mathcal{F} is nonlinear, then local sensitivities may not be adequate to make inferences about the parameters. To address this issue, we introduce a global sensitivity function. There are two important, and related, points to consider when defining a global sensitivity function:

- The term “global” refers to being global in the parameter space, not in the optimization variables. In general, only local optimality of the optimization problem (1) may be ensured, but we may perform analysis at different θ_0 ’s thus making the analysis global in the parameter space.
- For a fixed θ_0 , our analysis is only valid around the local minimum (u_0, z_0) . It is necessary to account for different local minima which may occur for a fixed θ_0 .

As in classical global sensitivity analysis, let Θ be a random field (or vector if $\Theta = \mathbb{R}^n$), taking values in Θ , which encodes the uncertainty in the parameters. The presence of multiple local minima poses an additional challenge beyond what is typically considered in global sensitivity analysis. To formally define a global sensitivity function, let \mathbf{I} be a random field taking values in $U \times Z$. Realizations of \mathbf{I} are used as initial iterates for an optimization routine. Assume that each (θ_0, I_0) is a realization of (Θ, \mathbf{I}) and is uniquely associated with a particular local minimum $(u_0(\theta_0, I_0), z_0(\theta_0, I_0))$ of (1). This ensures that the global sensitivity function defined below is well-posed. If the initial iterate is set to zero, then \mathbf{I} will equal zero with probability 1. Otherwise, random initial iterates can be used in which case different local minima may exist for fixed parameters and different realizations of \mathbf{I} .

For each (θ_0, I_0) as a realization of (Θ, \mathbf{I}) , there is a unique operator $\mathcal{D}(\theta_0; z_0(\theta_0, I_0))$ corresponding to (5) when $\theta = \theta_0$ and the local minimum is identified by solving (1) with initial iterate $I_0 \in U \times Z$. Then considering (6) as a function of (Θ, \mathbf{I}) yields that $\mathcal{S}(\Theta; z_0(\Theta, \mathbf{I}))$ is a random field. Assume that $\mathcal{S}(\Theta; z_0(\Theta, \mathbf{I}))\phi$ is measurable for each $\phi \in \Theta$. Mimicking ideas in derivative-based global sensitivity analysis [19,33–35], we define the global hyper-differential sensitivity function as $\mathcal{S}^G : \Theta \rightarrow \mathbb{R}$,

$$\mathcal{S}^G \phi = \mathbb{E}_{\Theta, \mathbf{I}} [\mathcal{S}(\Theta; z_0(\Theta, \mathbf{I}))\phi], \quad (9)$$

where $\mathbb{E}_{\Theta, \mathbf{I}} [\cdot]$ denotes the expected value computed with respect to the distribution of (Θ, \mathbf{I}) . We will frequently refer to (9) as a global sensitivity for short. Global set sensitivity indices may be defined in a similar manner by taking the expectation of (8), we omit the details for brevity.

The global sensitivity function S^G may be related to the solution of the optimization problem (1) as follows. Let $\mathcal{C} : \Theta \times U \times Z \rightarrow Z$ be defined by $\mathcal{C}(\theta_0, I_0) = \mathcal{P}\mathcal{F}(\theta_0; z_0(\theta_0, I_0))$, i.e. an operator mapping parameter and initial iterates to optimal solutions. Theorem 1 bounds the expected difference in optimal solutions in terms of the global sensitivity function.

Theorem 1. *Let $\phi \in \Theta$ and assume that*

- $\|\phi\| = 1$,
- *the Hessian of \mathcal{C} with respect to θ , $\mathcal{C}_{\theta,\theta}$, exists and is continuous,*
- $\exists Q > 0$ such that $\|\mathcal{C}_{\theta,\theta}(\theta, I)\| \leq Q \forall (\theta, I) \in \Theta \times U \times Z$,

Then

$$\mathbb{E}_{\Theta, I} [\|\mathcal{C}(\Theta + \delta\phi, I) - \mathcal{C}(\Theta, I)\|] \leq |\delta| S^G \phi + \frac{Q}{2} \delta^2.$$

Proof. Let (θ_0, I_0) be a fixed realization of (Θ, I) . Taylor's theorem implies that

$$\begin{aligned} \|\mathcal{C}(\theta_0 + \delta\phi, I_0) - \mathcal{C}(\theta_0, I_0)\| &\leq \|\mathcal{D}(\theta_0; z_0(\theta_0, I_0))\delta\phi\| + \frac{Q}{2} \delta^2 \\ &= |\delta| S(\theta_0; I_0) \phi + \frac{Q}{2} \delta^2. \end{aligned}$$

Taking the expectation over all (θ_0, I_0) completes the proof. \square

The constant Q in Theorem 1 corresponds to the nonlinearity of the operator \mathcal{C} (or equivalently \mathcal{F} since \mathcal{P} is linear), if \mathcal{C} is approximately linear then Q will be small. Theorem 1 implies that the average change in the optimal solution when the nominal parameters are perturbed in the direction ϕ is bounded by the global sensitivity function acting on ϕ , and a measure of the nonlinearity of \mathcal{C} .

Computing the expected value in S^G is very difficult in practice. As elaborated in Section 4.3, a sparse sampling approach is taken and the variability of the local sensitivity function is used as a heuristic to assess the nonlinearity of \mathcal{C} . This approach is akin to derivative-based global sensitivity analysis [19,33–35] and Morris screening [17,36]. Example 2.2 illustrates a special case where the parameter to optimal solution mapping \mathcal{C} is linear.

Example 2.2.

Consider the following optimization problem

$$\begin{aligned} \min_{u,z} \quad & J(u, z) = \frac{1}{2} \|u - d\|^2 \\ \text{s.t.} \quad & c(u, z, \theta) = A(\theta)u - z \end{aligned} \tag{10}$$

where $d \in U$ is a target state and $A(\theta)$ is linear differential operator and is linear in its parameters. The solution of the PDE is given by

$$u(z; \theta) = A^{-1}(\theta)z,$$

which is a nonlinear function of θ . To perform sensitivity analysis of the PDE solution with respect to θ requires adequate sampling of the parameter space to account for the nonlinearity. On the other hand, the optimal solution is given by

$$z_{opt}(\theta) = A(\theta)d.$$

Hence the optimal z is a linear function of θ even though the PDE solution is a nonlinear function of θ .

Even though the assumptions in Example 2.2 are too restrictive to permit any general conclusions, the example illustrates a case where it is easier to compute sensitivities of the optimal z than sensitivities of the PDE solution. Local sensitivity analysis of the optimal solution is therefore sufficient. In general, local sensitivities should be computed at various samples from parameter space to assess the nonlinearity.

3. COMPUTATIONS OF LOCAL SENSITIVITIES

Having formally defined sensitivities, next we present our approach to efficiently compute local sensitivities. In particular we seek low rank structure in the parameter space through the SVD of an operator arising from the solution of the PDE-constrained optimization problem. Computation of global sensitivities will be considered in Section 4.3.

3.1 Finite Dimensional Approximation

The optimization problem (1) is discretized by defining the finite dimensional subspaces $U_h \subseteq U, Z_h \subseteq Z, \Lambda_h \subseteq \Lambda, \Theta_h \subseteq \Theta$. In practice these subspaces typically arise from a discretization of the PDE, for instance, a finite element discretization. Let $\theta_0 \in \Theta_h$ be the nominal parameters, and $(u_0, z_0, \lambda_0) \in U_h \times Z_h \times \Lambda_h$ be a local minimum of the discretization of (1) with nominal parameters θ_0 .

Letting $\{\phi_1, \phi_2, \dots, \phi_n\}$ be a basis for Θ_h , we define the local hyper-differential sensitivity indices as

$$S_i(\theta_0, z_0) = \mathcal{S}(\theta_0, z_0)\phi_i, \quad i = 1, 2, \dots, n. \quad (11)$$

Global hyper-differential sensitivity indices $S_i^G, i = 1, 2, \dots, n$, may be defined in a similar manner by having (9) act on the basis functions. For simplicity we assume that $\|\phi_i\| = 1$ for $i = 1, 2, \dots, n$.

The indices $S_i(\theta_0, z_0), i = 1, 2, \dots, n$, may be computed directly by solving (1) once, and subsequently computing each $S_i(\theta_0, z_0)$ separately. This approach requires solving n linear systems with coefficient matrix \mathcal{K} (or solving a block system with n right hand sides). However, in many applications \mathcal{B} (in (3)) possesses a low rank structure which may be exploited to accelerate computation.

Assume that $\mathcal{D}(\theta_0; z_0)$ is a compact operator. Letting $\sigma_k, \theta_k, z_k, k = 1, 2, \dots$ denote the singular values and vectors of $\mathcal{D}(\theta_0; z_0)$, i.e. $\mathcal{D}(\theta_0; z_0)\theta_k = \sigma_k z_k$, we have

$$S_i(\theta_0, z_0) = \sqrt{\sum_{k=1}^{\infty} \sigma_k^2(\theta_0, \phi_i)^2}, \quad i = 1, 2, \dots, n. \quad (12)$$

Truncating the series in (12) yields the approximations

$$S_i(\theta_0, z_0) \approx \sqrt{\sum_{k=1}^K \sigma_k^2(\theta_0, \phi_i)^2}, \quad i = 1, 2, \dots, n.$$

The singular values/vectors $\sigma_k, \theta_k, k = 1, 2, \dots, K$ may frequently be computed with far fewer than n applications of the operator $\mathcal{D}(\theta_0; z_0)$. For many problems in practice, n is on the order of hundreds, thousands, or more, whereas K on the order of tens may be sufficient to accurately approximate $S_i, i = 1, 2, \dots, n$. Leveraging the truncated SVD representation of $\mathcal{D}(\theta_0; z_0)$ may reduce the number of linear system solves by an order of magnitude (or more) if such low rank structure exists. The presence (or lack thereof) of low rank structure is easily identified by computing the leading singular values of $\mathcal{D}(\theta_0; z_0)$, so the approximation is easily certified in practice.

In addition to facilitating efficient estimation of the sensitivity indices, the truncated SVD also provides directions of greatest sensitivity in parameter space (the singular vectors). These directions provide an alternative coordinate system in parameter space with gives a useful low dimensional representation. Though different in several ways, these singular vectors and sensitivity indices have a similar intuition as active subspaces [37] and activity scores [38], respectively. Additionally, the singular vectors in the Z indicate which features of the optimal solution (in space and/or time) are most sensitivity to the parameters.

Assuming that the singular values/vectors $\sigma_k, \theta_k, z_k, k = 1, 2, \dots, K$ have been computed, the local set sensitivity index (8) may be estimated by computing the largest singular value of the linear operator

$$\phi \mapsto \sum_{k=1}^K \sigma_k z_k(\theta_0, \Pi_{\Xi_i} \phi).$$

For low rank operators (small K), the computational cost of this estimate is negligible in comparison to computing $\sigma_k, \theta_k, z_k, k = 1, 2, \dots, K$.

For problems where n is large and $\mathcal{D}(\theta_0; z_0)$ does not admit a low rank structure, computing the local sensitivity indices (11) is prohibitive. In such cases, the local set sensitivity indices (8) may be estimated by using a SVD routine to compute the leading singular value of $\mathcal{D}(\theta_0; z_0)\Pi_{\Xi_i}$ for each $i = 1, 2, \dots, T$. This is a significant saving when, for instance, $T = \mathcal{O}(5)$ and $n = \mathcal{O}(1,000)$, which occurs when there are multiple temporally and/or spatially dependent parameters.

3.2 Computing the Truncated SVD

The operator \mathcal{D} is defined on function spaces so the truncated SVD should be computing using the inner products from Θ_h and Z_h . To achieve this, an SVD routine using Euclidean inner products may be applied to the matrix

$$R_{Z_h} D R_{\Theta_h}^{-1}, \quad (13)$$

where R_{Θ_h} and R_{Z_h} are the Cholesky factors of the symmetric positive definite mass matrices (or weighting matrices more generally) $M_{\Theta_h} = R_{\Theta_h}^T R_{\Theta_h} \in \mathbb{R}^{n \times n}$ and $M_{Z_h} = R_{Z_h}^T R_{Z_h} \in \mathbb{R}^{m \times m}$ defined by

$$(M_{\Theta_h})_{i,j} = (\phi_i, \phi_j)_{\Theta_h} \quad (M_{Z_h})_{i,j} = (y_i, y_j)_{Z_h},$$

where $\{y_1, y_2, \dots, y_m\}$ is a basis for Z_h .

The matrix $D = C_{Z_h} \mathcal{D} E_{\Theta_h} \in \mathbb{R}^{m \times n}$ represents the action of \mathcal{D} on coordinates in the discretized spaces, where $E_{\Theta_h} : \mathbb{R}^n \rightarrow \Theta_h$ and $C_{Z_h} : Z_h \rightarrow \mathbb{R}^m$ denote the coordinate transformation operators. Computing the truncated SVD of (13) directly is not scalable because the Cholesky factors R_{Z_h} and R_{Θ_h} will be dense whereas the mass matrices are typically sparse. In what follows, we propose to a symmetric generalized eigenvalue problem instead, and mitigate computational limitations by introducing the use of a randomized generalized eigenvalue solver from [30] which facilitates parallel evaluations of matrix vector products.

3.3 Proposed Reformulation

For notational simplicity, θ and z are used to denote the coordinate representations for parameters and optimization variables throughout this section. The singular values and singular vectors of (13) correspond to the positive eigenvalues and eigenvectors of the Jordan-Wielandt matrix

$$W = \begin{pmatrix} 0 & R_{Z_h} D R_{\Theta_h}^{-1} \\ (R_{\Theta_h}^{-1})^T D^T (R_{Z_h})^T & 0 \end{pmatrix}.$$

This gives the symmetric eigenvalue problem

$$W \begin{pmatrix} z \\ \theta \end{pmatrix} = \alpha \begin{pmatrix} z \\ \theta \end{pmatrix} \quad (14)$$

for which we seek to compute the largest eigenvalues.

To avoid computing the Cholesky factors of the mass matrices, define

$$X = \begin{pmatrix} R_{Z_h}^{-1} & 0 \\ 0 & R_{\Theta_h}^{-1} \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \tilde{z} \\ \tilde{\theta} \end{pmatrix} = X \begin{pmatrix} z \\ \theta \end{pmatrix}.$$

Then we may reformulate (14) as

$$A \begin{pmatrix} \tilde{z} \\ \tilde{\theta} \end{pmatrix} = \alpha B \begin{pmatrix} \tilde{z} \\ \tilde{\theta} \end{pmatrix}. \quad (15)$$

where

$$A = (X^{-1})^T W X^{-1} = \begin{pmatrix} 0 & M_{Z_h} D \\ D^T M_{Z_h} & 0 \end{pmatrix} \quad \text{and} \quad B = (X^{-1})^T X^{-1} = \begin{pmatrix} M_{Z_h} & 0 \\ 0 & M_{\Theta_h} \end{pmatrix},$$

Since A is symmetric and B is symmetric positive definite, a symmetric generalized eigenvalue problem must be solved. Symmetry in our formulation ensures good numerical properties and leverages powerful theoretical results from linear algebra [39]. We propose to solve the generalized eigenvalue problem with a randomized algorithm, the advantages of this approach will be elaborated on in Section 4.

Assume that the K largest eigenvalues $\alpha_k > 0$ and corresponding eigenvectors $(\tilde{z}_k, \tilde{\theta}_k)$, $k = 1, 2, \dots, K$, of (15) have been computed. The corresponding singular values, right singular vectors, and left singular vectors of (13) (using Euclidean inner products) are α_k , $R_{\Theta_h} \tilde{\theta}_k$, and $R_{Z_h} \tilde{z}_k$, $k = 1, 2, \dots, K$, respectively. It appears that the mass matrix Cholesky factor is needed to compute these singular vectors; however, as shown in [22], the coordinate representation of the right and left singular vectors of (5) (using Θ_h and Z_h inner products) are given by

$$\frac{\tilde{\theta}_k}{\sqrt{\tilde{\theta}_k^T M_{\Theta_h} \tilde{\theta}_k}} \quad \text{and} \quad \frac{\tilde{z}_k}{\sqrt{\tilde{z}_k^T M_{Z_h} \tilde{z}_k}},$$

$k = 1, 2, \dots, K$, respectively. Hence we only compute the largest eigenvalues and eigenvectors of (15), and normalize them with the mass matrices.

An alternative formulation solves $D^T M_{Z_h} D \theta = \alpha M_{\Theta_h} \theta$ instead of (15). This system is $n \times n$ positive definite, instead of $(n+m) \times (n+m)$ indefinite, and squares the singular values, which is advantageous for computing singular values greater than one. However, this formulation requires an additional collection of K matrix vector products to compute the left singular vectors z_k , $k = 1, 2, \dots, K$. Both formulations were considered, the results in this article focus on the generalized eigenvalue problem (15) because of its ease computing the left singular vectors.

4. ALGORITHMIC OVERVIEW

Algorithm 1 below provides an overview of hyper-differential sensitivity analysis and highlights the important computational features. The solution to the underlying PDE-constrained problem is encapsulated in Line 3. A core component of Algorithm 1 is the randomized algorithm, in Lines 4-12, used to solve the generalized eigenvalue problem. We begin by motivating the randomized generalized eigenvalue solver and subsequently consider the computational complexity of Algorithm 1.

4.1 Randomized Linear Algebra

Randomized linear algebra has emerged as a powerful tool in scientific computation [40]. The utility of randomized methods is that they permit a reordering of the computation which may better exploit computing architectures. Most traditional algorithms are inherently serial, for instance, constructing Krylov subspaces require serial matrix-vector products since each vector is formed using the previous ones. In contrast, randomized methods may require a comparable number of matrix-vector products which can be computed in parallel with minimal communication overhead.

We adopt the randomized generalized eigenvalue solver from [30]. It is well suited for estimating the largest eigenvalues. The user specifies the desired number of eigenvalues p and an oversampling factor L . Then $p+L$ matrix vector products are computed (in parallel) by applying the coefficient matrix to independently generated random vectors. The resulting vectors form an approximation of the subspace of eigenvectors corresponding to the largest eigenvalues. Inexpensive computation may be done in this subspace to estimate the eigenvalues and eigenvectors.

Algorithm 1 inputs an integer N specifying the number of local sensitivities to compute, an integer K specifying the number of singular pairs to compute for each local sensitivity, and an integer L specifying the oversampling factor. Large values of N may be necessary to accurately estimate global sensitivities; however, relatively small values of N , for instance $\mathcal{O}(10)$, are frequently sufficient to capture important features. The user should choose N based on their available computational resources. Ideally, the N local sensitivities will provide similar parameter inferences.

Algorithm 1: Hyper-differential Sensitivity Analysis Algorithm

Input: number of parameter samples N , number of singular pairs K , oversampling factor L

1: for j from 1 to N (embarrassingly parallel loop)

2: sample $(\bar{\theta}^j, \bar{I}^j)$ from the distribution of (Θ, \mathbf{I})

3: solve (1) with $\theta = \bar{\theta}^j$ and initial iterate \bar{I}^j and store solution u_{opt}^j, z_{opt}^j

4: for i from 1 to $2K + L$ (embarrassingly parallel loop)

5: draw a standard normal sample $\begin{pmatrix} \tilde{z}^i \\ \tilde{\theta}^i \end{pmatrix}$ and compute $y_i = B^{-1}A\begin{pmatrix} \tilde{z}^i \\ \tilde{\theta}^i \end{pmatrix} \in \mathbb{R}^{m+n}$

6: end

7: compute the decomposition $QR = [y_1, y_2, \dots, y_{2K+L}] \in \mathbb{R}^{(m+n) \times (2K+L)}$ with B inner products

8: for i from 1 to $2K + L$ (embarrassingly parallel loop)

9: compute $Aq_i \in \mathbb{R}^{(m+n)}$ where $Q = [q_1, q_2, \dots, q_{2K+L}] \in \mathbb{R}^{(m+n) \times (2K+L)}$

10: end

11: form $T = Q^T A Q \in \mathbb{R}^{(2K+L) \times (2K+L)}$

12: compute the eigenvalue decomposition $T = VEV^T$ with ordering $E_{1,1} \geq E_{2,2} \geq \dots \geq E_{2K+L}$

13: form $\begin{pmatrix} \tilde{z}_k \\ \tilde{\theta}_k \end{pmatrix} = Qv_k \in \mathbb{R}^{(m+n)}$, $k = 1, 2, \dots, K$, where $V = [v_1, v_2, \dots, v_{2K+L}]$

14: store singular values $\sigma_k^j = E_{k,k}$, $k = 1, 2, \dots, K$

15: compute and store singular vectors $z_k^j = \frac{\tilde{z}_k}{\tilde{z}_k^T M_{Z_h} \tilde{z}_k}$ and $\theta_k^j = \frac{\tilde{\theta}_k}{\tilde{\theta}_k^T M_{\Theta_h} \tilde{\theta}_k}$, $k = 1, 2, \dots, K$

16: end

Return: $u_{opt}^j, z_{opt}^j, \sigma_k^j, z_k^j, \theta_k^j, k = 1, 2, \dots, K, j = 1, 2, \dots, N$

If the local sensitivities differ significantly (the operator \mathcal{F} is highly nonlinear) then the user should exercise caution making inferences with them. The “optimal” choice for K is not clear a-priori; however, the singular values returned from Algorithm 1 certify the choice of K (by assessing the low rank structure). The user should start with small values for K , for instance $K = 4$ is used in Section 5, and increase it if necessary. If there is not sufficient decay in the first K singular values then Lines 4-12 may be repeated to augment the existing computation. The oversampling factor L scales the cost versus accuracy. Typically $L < 20$ (or even $L < 10$) is sufficient, see [30,40] and references therein. In Algorithm 1, $2K + L$ random vectors are used (see Lines 4-6) because the eigenvalues of A correspond to positive and negative pairs of the desired singular values.

The outer loop initialized in Line 1 of Algorithm 1 is embarrassingly parallel. The most computationally intensive portions within this loop are Line 3 (solving the PDE-constrained optimization problem), Line 5 (applying A), and Line 9 (applying A). All of the other calculations are simple linear algebra which may be executed quickly using standard libraries. Solving the PDE-constrained optimization problem in Line 3 involves a host of complexities including trust region and/or line search globalization, finite element discretization, solving large systems of (possibly nonlinear) equations, and adjoint calculations to evaluate gradients and/or hessians. In the scope of this work, we assume that efficient solvers are available for this end but emphasize its complexity. Lines 5 and 9 are computationally intensive because applying A requires applying \mathcal{K}^{-1} twice, which involves potentially many PDE solves.

The benefit of the randomized generalized eigenvalue solver is that the $2K + L$ application of A in Lines 5 and 9 may be parallelized. The loops are embarrassingly parallel since the random vectors are independent; however, barriers are necessary after each loop because the data from each matrix vector product must be shared across processors in order to form Q and T (Lines 7 and 11). There may be parallel inefficiency if the execution time to apply A differs significantly across the samples.

4.2 Algorithmic Complexity

To assess the computational complexity of Algorithm 1 we count the number of large scale linear system solves. For simplicity our assessment focuses on Lines 4-15 as the computational complexity of Line 3 is a question of PDE-constrained optimization and the outer loop initialized at Line 1 simply scales the cost of Lines 2-15 by a factor N . The computational cost of Lines 4-15 is approximately equal to the cost of applying \mathcal{K}^{-1} four times (twice in Lines 5 and 9). Assuming that each application of \mathcal{K}^{-1} requires an average of s_{CG} iterations of conjugate gradient, this amounts to $4s_{CG}$ times the cost of each matrix vector product $\mathcal{K}v$. Since \mathcal{K} is evaluated at the optimal solution we may use the existing solution for the primal and adjoint PDE solves, so computing $\mathcal{K}v$ requires a state sensitivity solve (linear system involving the state Jacobian of the constraint) and a adjoint sensitivity solve (linear system involving the state Jacobian of the constraint transposed). The total computational cost will be approximately $4(2K + L)s_{CG}$ state sensitivity solves plus $4(2K + L)s_{CG}$ adjoint sensitivity solves. The factor of $(2K + L)$ may be mitigated through the parallelism of the randomized solver.

4.3 Interpretation

Estimating (9) accurately through sampling based approaches may require extensive computational effort (a large N in Algorithm 1). However, useful information may be obtained by computing local sensitivities at a sparse collection of samples from the parameter space. This section details practical considerations computing, visualizing, and interpreting samples of local sensitivities. The proposed approach is demonstrated in Section 5.

Assume that Algorithm 1 has been executed yielding optimal solutions and local sensitivities at N (chosen based upon the user's computational budget) different realizations of (Θ, \mathbf{I}) . For each $j = 1, 2, \dots, N$, we have,

- the state u_{opt}^j and optimization variables z_{opt}^j determined by solving (1) with initial iterate \bar{I}^j when $\theta = \bar{\theta}^j$,
- the K leading singular triples of (5), $(\sigma_k^j, \theta_k^j, z_k^j)$, $k = 1, 2, \dots, K$, in coordinate representation.

The spectrum of (5) summarizes its low rank structure (or lack thereof). The singular values σ_k^j , $k = 1, 2, \dots, K$, $j = 1, 2, \dots, N$, may be visualized in a scatter plot which readily reveals the structure of the leading singular values for different samples in parameter space. Ideally we will see a decay in the singular values for each j , indicating a low rank structure. If such a low rank structure exists, the K singular values and singular vectors may be used to analyze the sensitivities. Otherwise, a larger value of K is needed. Our approach is designed to exploit low rank structure which we are implicitly assuming is present.

The local sensitivity index (12) is approximated for the i^{th} parameter basis function using the j^{th} sample with

$$\hat{S}_i^j = \sqrt{\sum_{k=1}^K (\sigma_k^j)^2 \left(\left(M_{\Theta_h} \theta_k^j \right)_i \right)^2} \quad (16)$$

for $i = 1, 2, \dots, m$, $j = 1, 2, \dots, N$; $\left(M_{\Theta_h} \theta_k^j \right)_i$ denotes the i^{th} entry of the vector $M_{\Theta_h} \theta_k^j$. Then \hat{S}_i^j , $i = 1, 2, \dots, m$, $j = 1, 2, \dots, N$, may be visualized in a scatter plot which reveals the relative influence of the parameters and the variability of their local sensitivity indices over the samples $\bar{\theta}^j, \bar{I}^j$, $j = 1, 2, \dots, N$. This variability provides a heuristic to measure the nonlinearity of the parameter to optimal solution mapping. We hope to find a similar low rank structure and local sensitivity indices for each parameter sample which indicates desirable structure in the problem. If the local sensitivities vary significantly over the (sparse) sampling of parameter space, this indicates strong nonlinearities which will mandate greater computational effort.

As discussed in Subsection 3.1, the local set sensitivity indices may be estimated as a by-product using the singular values and vectors. They may be easily visualized in a manner similar to the local sensitivity indices. This is particularly useful when the visualization of the local sensitivity indices becomes cumbersome because of temporal and/or spatial dependencies of the parameters.

Along with the spectrum and parameter sensitivity information above, we may also visualize the resulting changes in the optimal solution when the parameters are perturbed. In particular, z_k^j is the change in the optimal solution if the parameter vector $\bar{\theta}^j$ is perturbed in the direction θ_k^j . The singular vectors z_k^j may be visualized by overlaying them on a plot, or plotting statistical quantities computed from the sample $\{z_k^j\}_{j=1}^N$. In both cases, it informs us which features of the optimal solution are most sensitive to the parametric uncertainty.

5. NUMERICAL RESULTS

In this section we demonstrate our proposed hyper-differential sensitivity analysis on two examples. The first is a control problem for a steady state nonlinear multi-physics system modeling a high pressure chemical vapor deposition (CVD) reactors, see [41,42]. The second example is an inverse problem for a transient multi-physics system modeling subsurface contaminant transport.

Our HDSA software is implemented in the Rapid Optimization Library (ROL) [43] of Trilinos [44], a collection of C++ libraries for scientific computation. The implementation is based on C++, Trilinos parallel constructs, and special PDE-constrained solver interfaces that generalize the solution procedure to a range of PDE-based models. ROL consists of state-of-the-art Newton Krylov based optimization methods with both reduced and full space solution methods (trust region and line search globalization). The HDSA implementation is matrix free, has three levels of parallelism (two embarrassingly parallel loops and parallel linear algebra constructs), and may be easily adapted to a variety of applications. The parallel and matrix free design achieves scalable performance.

5.1 Control of Thermal Fluids

In this subsection we consider control of the steady state nonlinear multi-physics Boussinesq flow equations in two spatial dimensions, a model for a CVD reactor. Reactant gases are injected in the top of a reactor and flow downwards to create an epitaxial film on the bottom. Vorticities created by buoyancy-driven convection inhibit some gases from reaching the bottom of the reactor. Thermal fluxes are controlled on the side walls of the reactor in order to minimize the vorticity. Formally, consider the control problem,

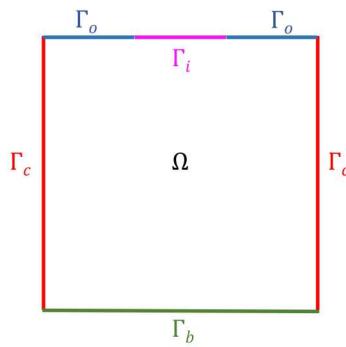


FIG. 2: Computation domain and boundaries for (17).

$$\begin{aligned}
& \min_{v,p,T,z} \frac{1}{2} \int_{\Omega} (\nabla \times v)^2 dx + \frac{\gamma}{2} \int_{\Gamma_c} z^2 dx \\
& \text{s.t.} \\
& -\epsilon(\theta) \nabla^2 v + (v \cdot \nabla) v + \nabla p + \eta(\theta) T g = 0 \quad \text{in } \Omega \\
& \nabla \cdot v = 0 \quad \text{in } \Omega \\
& -\kappa(\theta) \Delta T + v \cdot \nabla T = 0 \quad \text{in } \Omega \\
& T = 0 \quad \text{and} \quad v = v_i \quad \text{on } \Gamma_i \\
& \kappa(\theta) \nabla T \cdot n = 0 \quad \text{and} \quad v = v_o \quad \text{on } \Gamma_o \\
& T = T_b(\theta) \quad \text{and} \quad v = 0 \quad \text{on } \Gamma_b \\
& \kappa(\theta) \nabla T \cdot n + \nu(\theta)(z - T) = 0 \quad \text{and} \quad v = 0 \quad \text{on } \Gamma_c
\end{aligned} \tag{17}$$

where $\Omega = (0, 1) \times (0, 1)$, $\Gamma_i = [1/3, 2/3] \times \{1\}$, $\Gamma_o = [0, 1/3] \times \{1\} \cup [2/3, 1] \times \{1\}$, $\Gamma_b = [0, 1] \times \{0\}$, $\Gamma_c = \{0, 1\} \times [0, 1]$, and n denotes the outward pointing normal vector to the boundary. Figure 2 depicts the domain and boundaries.

The state consists of horizontal and vertical velocities $v = (v_1, v_2)$, the pressure p , and the temperature T ; the control z is a function defined on the left and right boundaries of Ω , denote their union Γ_c . The deterministic inflow and outflow conditions v_i and v_o are given by

$$v_i(x) = -4 \left(x - \frac{1}{3} \right) \left(\frac{2}{3} - x \right)$$

and

$$v_o(x) = \begin{cases} 2 \left(\frac{1}{3} - x \right) x & \text{if } x \in [0, \frac{1}{3}] \\ 2 \left(x - \frac{2}{3} \right) (1 - x) & \text{if } x \in [\frac{2}{3}, 1] \end{cases}.$$

Uncertainties enter the the model through the finite dimensional vector (coming from a spatial discretization) $\theta \in \mathbb{R}^m$ where $m = 2m_b + 2m_\ell + 2m_r + 3$. In particular, the boundary term T_b is defined through the sum

$$T_b(x, \theta) = 1 + 0.2 \sum_{k=1}^{m_b} \left(\theta_k \frac{\sin(\pi k x)}{k} + \theta_{m_b+k} \frac{\cos(\pi k x)}{k} \right). \tag{18}$$

In addition, the boundary term $\nu(\theta)$, a function defined on Γ_c , is expressed as a sum in the form of (18) with $2m_\ell$ parameters in the sum defining the left boundary term and $2m_r$ parameters in the sum defining the right boundary term. Along with these $2m_b + 2m_\ell + 2m_r$ parameters which define T_b and ν , there are three parameters, $\theta_{m-2}, \theta_{m-1}, \theta_m$, which appear in the uncertain scalar quantities ϵ, η , and κ , as

$$\begin{aligned}
\epsilon &= \frac{1}{Re} = \frac{1 + 0.05\theta_{m-2}}{100}, & \kappa &= \frac{1}{RePr} = \frac{1 + 0.05\theta_{m-1}}{72} \\
\text{and } \eta &= \frac{Ge}{Re^2} = 1 + 0.05\theta_m,
\end{aligned}$$

where Re , Ge , and Pr are the Reynolds number, Grashof number, and Prandtl number respectively. Each θ_k , $k = 1, 2, \dots, m$, is assumed to be independent and uniformly distributed on $[-1, 1]$. The initial iterate is taken to be zero for all samples; similar results were found using a random initial iterate. The parameter weighting matrix M_p is the identity since the uncertain boundary conditions T_b and ν are discretized by orthogonal global basis functions.

The PDE is discretized with finite elements on a 99x99 rectangular mesh. The velocity and pressure are represented with the Q2-Q1 Taylor-Hood finite element pair and the temperature is represented with the Q2 finite element.

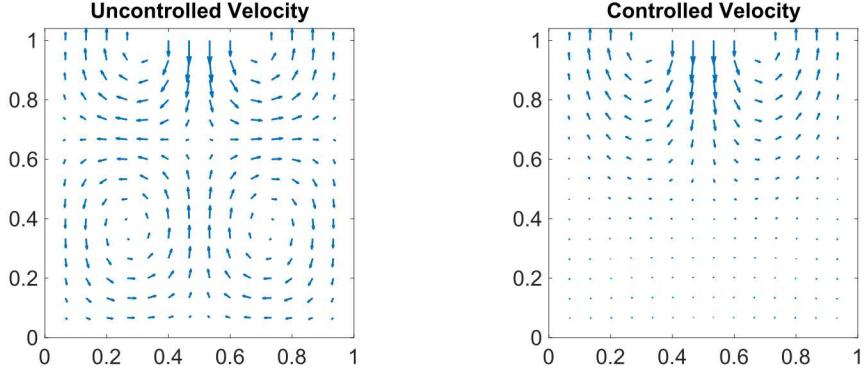


FIG. 3: Uncontrolled (left) and controlled (right) velocity field.

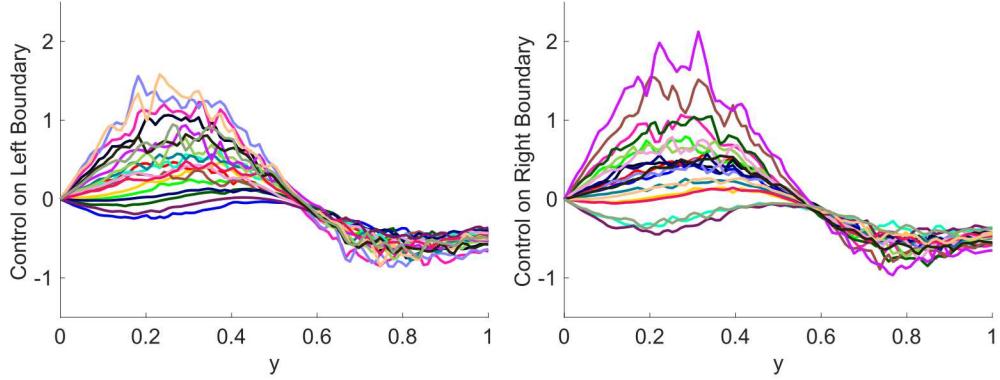


FIG. 4: Control solutions for (17) corresponding to 20 different parameter samples. The left and right panels are the controllers on the left and right boundaries, respectively. Each curve is a control solution for a given parameter sample.

We take $m_b = m_\ell = m_r = 25$ which yields a total of $m = 153$ uncertain parameters. The deterministic control problem is solved using the full space composite step algorithm in ROL with the control penalty $\gamma = 0.01$. Figure 3 displays the uncontrolled (left) and controlled (right) velocity field with parameters $\theta_k = 0$, $k = 1, 2, \dots, 153$. The undesired vorticities are observed in the uncontrolled velocity field and are reduced by the control strategy.

Local sensitivities are evaluated at $N = 20$ samples from parameter space. Figure 4 displays the optimal control solutions for these 20 samples. The left and right panels display the controller on the left and right boundary, respectively. Each curve corresponds to the control solution for a different parameter sample. There is significant variability in the solutions which indicates a strong dependence of the controller on the uncertain parameters. Our objective in the hyper-differential sensitivity analysis is to determine which parameters cause the greatest changes in the controller so that uncertainty quantification and robust optimization may focus on them rather than the full set of 153 parameters.

We compute the leading $K = 4$ singular triples of (5) and follow the approach presented in Section 4.3 to analyze them. An oversampling factor of $L = 8$ is used. Figure 5 shows the leading 4 singular values from each of the 20 parameter samples. Each vertical slice in Figure 5 gives the 4 singular values for the fixed parameter sample. We observe that there are 2 dominant singular triples (thus validating that $K = 4$ is sufficiently large) and that the singular values do not vary significantly over the different parameter samples.

The sensitivity indices (16) and displayed in Figure 6. There are 20 circles in each vertical slice of Figure 6 corresponding to the local sensitivity index for a fixed parameter over the 20 samples. We observe several interesting features:

- The local sensitivity analysis yields similar results for each parameter sample.

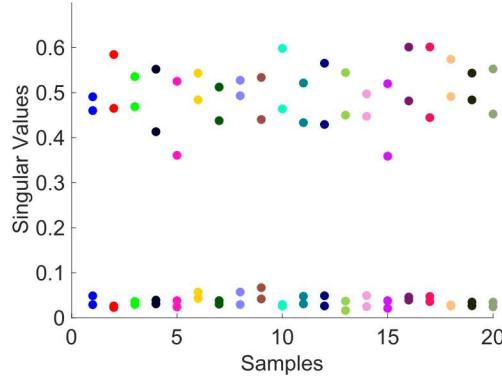


FIG. 5: Leading 4 singular values of (5) at 20 different parameter samples. Each vertical slice corresponds to the leading 4 singular values for a fixed sample.

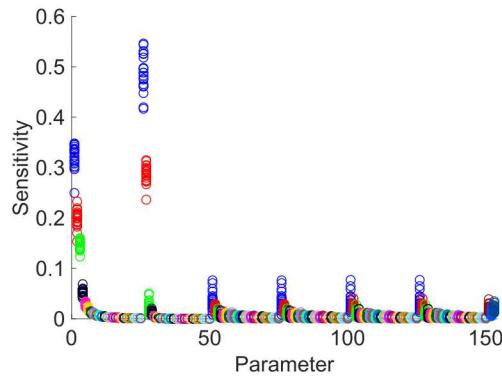


FIG. 6: Local sensitivities (16) for the 153 uncertain parameters in (17). The 20 circles in each vertical slice indicates the sensitivity index for a fixed parameter as it varies over the 20 parameter samples. The repeating color scheme indicates the grouping of parameters as they correspond to sine and cosine components of each boundary condition.

- Only around 10% of the uncertain parameters exhibit significant influence on the control strategy.
- The bottom boundary condition, T_b , has the greatest influence on the control strategy.
- The cosine components of T_b are more important in the first two frequencies, the sine component is more important in the third frequency.

To complement these parameter sensitivities, Figure 7 displays the singular vectors, z_k , $k = 1, 2$, corresponding to the leading parameter perturbations θ_k , $k = 1, 2$. The top left and top right panels of Figure 7 show the first singular vector on the left and right boundaries, respectively; the bottom left and bottom right panels are the second singular vector on the left and right boundaries, respectively. The singular vectors z_1 and z_2 may be interpreted as the change in the control strategy if the parameters are perturbed according to the singular vectors θ_1 and θ_2 , respectively. We observe that the control strategy will change more near the bottom of the domain, an unsurprising result since the greatest sensitivity is in the bottom boundary condition.

The computation was performed using 80 compute nodes, each containing 16 processors. By taking 20 parameter samples and 16 random vectors in the eigenvalue solver, the embarrassingly parallel loop in the eigenvalue solver distributed the computation over all 1280 processors, using 4 processors in parallel for each matrix vector product. This reduces the overall execution time to approximately one matrix vector product in Line 5 and one matrix vector product in Line 9 of Algorithm 1, each leveraging parallel linear algebra with 4 processors. Since the KKT solve (4) dominates the computational cost, the total execution time for computing 20 local sensitivities is approximately equal

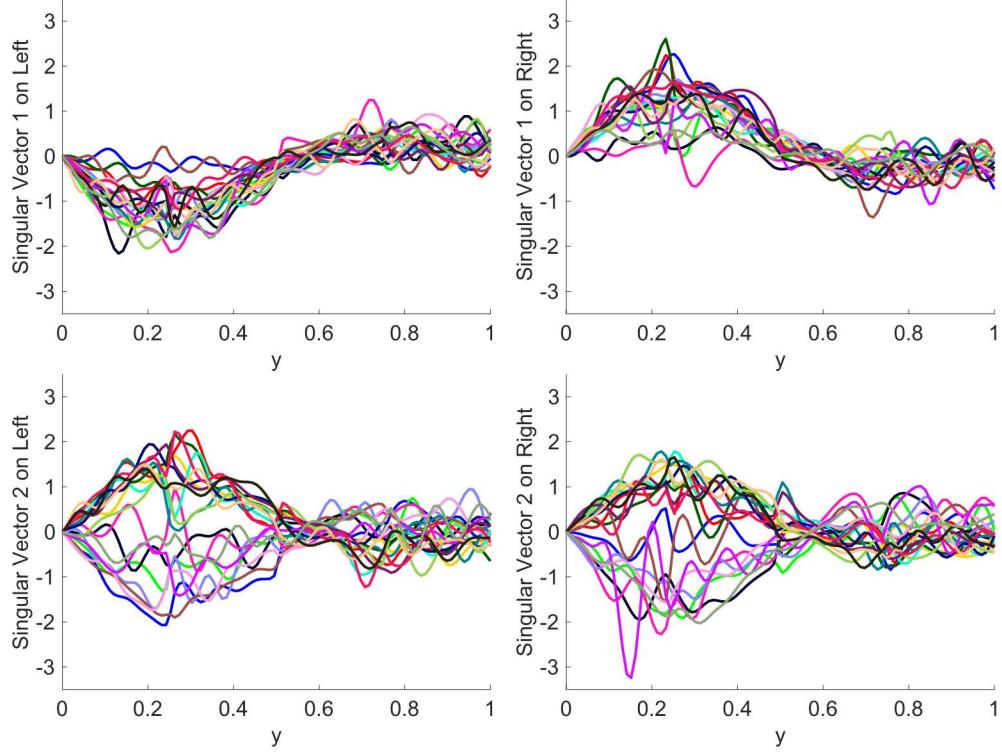


FIG. 7: First two singular vectors z_k , $k = 1, 2$, of (5), i.e. $\mathcal{D}\theta_k = \sigma_k z_k$. The top (bottom) row shows the first (second) singular vector on the left and right boundaries, respectively. Each curve corresponds to a different parameter sample.

to 4 KKT solves.

5.2 Inversion of Permeability

In this subsection we consider an inverse problem constrained by the transient advection diffusion equation and Darcy's equation in two spatial dimensions. This emulates a model for subsurface contaminant transport. We seek to invert for a contaminant source given sparse noisy measurements of the contaminant. Consider the following inverse formulation:

$$\min_{p,c,z} \frac{1}{2} \sum_{i=1}^T \sum_{j=1}^M (\mathcal{P}_{i,j} c - d_{i,j})^2 + \frac{\alpha}{2} \int_{\Omega} z^2 dx \quad (19)$$

s.t.

$$\nabla \cdot (-\kappa(\theta) \nabla p) = 0 \quad \text{in } \Omega$$

$$\frac{\partial c}{\partial t} + \nabla \cdot (-\epsilon(\theta) \nabla c) + (-\kappa(\theta) \nabla p) \cdot \nabla c = \chi_{[0.01, 0.02]}(t) z \quad \text{in } \Omega \text{ for } t > 0$$

$$p = \psi(\theta) \quad \text{on } \Gamma_L \cup \Gamma_R$$

$$-\kappa(\theta) \nabla p \cdot n = 0 \quad \text{on } \Gamma_B \cup \Gamma_T$$

$$\nabla c \cdot n = 0 \quad \text{on } \Gamma \text{ for } t > 0$$

$$u = 0 \quad \text{in } \Omega \text{ for } t = 0 \quad (20)$$

where $\Omega = (0, 1)^2$ with boundary $\Gamma = \Gamma_L \cup \Gamma_B \cup \Gamma_R \cup \Gamma_T$, $\Gamma_L = \{0\} \times (0, 1)$, $\Gamma_B = (0, 1) \times \{0\}$, $\Gamma_R = \{1\} \times (0, 1)$, $\Gamma_T = (0, 1) \times \{1\}$, and n denotes the outward pointing normal vector to the boundary.

The state consists of the transient contaminant concentration c and the steady state pressure p ; the stationary source z is a function defined on Ω which appears on the right hand side of the advection diffusion equation along with the characteristic function $\chi_{[.01,.02]}(t)$ which equals 1 on the time interval $[.01, .02]$ and 0 otherwise. Contaminant concentration data is collected at $M = 121$ sensors which are uniformly distributed on an 11×11 grid in $[0, 1]^2$ and $T = 40$ instances in time. Synthetic data $d_{i,j}$ (at the i^{th} time instance and j^{th} sensor location) is generated by solving the governing PDEs (with double the mesh resolution used in the inverse problem) with a specified source term (given in Figure 9) and adding independent Gaussian noise whose standard deviation is 3% of the observed concentration. The observation operator $\mathcal{P}_{i,j}$ maps the PDE solution to the contaminant concentration at the i^{th} time instance and j^{th} sensor location.

Uncertainties enter the the model through the finite dimensional vector (coming from a spatial discretization) $\theta \in \mathbb{R}^m$ where $m = 256 + 1 + 16 + 16 = 289$ which parameterizes perturbations of the scalar and spatially dependent parameters in the PDE, in particular, the permeability field κ , left and right Dirichlet boundary condition ψ , and diffusion coefficient ϵ . These parameters are fixed to nominal values in order to solve the inverse problem and we consider the sensitivity of the source estimate to perturbations of the parameters.

The nominal value of the permeability field, denote it by $\bar{\kappa}$, is shown in Figure 8 alongside the Darcy pressure generated by solving Darcy's equation with the nominal permeability field $\bar{\kappa}$ and nominal boundary condition defined by

$$\bar{\Psi}(x, y) = \begin{cases} 10 + 2 \cos(2\pi y) & \text{if } x \in \Gamma_L \\ 12 + \cos(2\pi y) + .5 \cos(4\pi y) & \text{if } x \in \Gamma_R \end{cases}.$$

The nominal scalar diffusion coefficient in the advection diffusion equation is $\bar{\epsilon} = 1$.

To study the sensitivity of the source estimate to uncertainties in the permeability field, boundary condition, and diffusion coefficient, we represent them as

$$\begin{aligned} \kappa &= \bar{\kappa} \left(1 + a \sum_{k=1}^{(L+1)^2} \theta_k \phi_k \right) \\ \psi &= \bar{\Psi} \left(1 + a \sum_{k=1}^{2(L+1)} \theta_{(L+1)^2+k} \eta_k \right) \\ \epsilon &= \bar{\epsilon} \left(1 + a \theta_{(L+1)^2+2(L+1)+1} \right) \end{aligned}$$

where $a = 0.2$ represents the level of uncertainty (20%), $L = 15$ is an integer specifying a spatial discretization, and ϕ_k and η_k are linear finite element basis functions defined on a rectangular mesh with $L + 1$ equally spaced nodes in each spatial dimension (ϕ_k is defined on $[0, 1]^2$ and η_k is defined on $[0, 1]$). The parameter weighting matrix M_p is block diagonal with its blocks given by the finite element mass matrices arising from discretization of κ and ψ , and the scalar identity for the 1×1 block corresponding to ϵ .

The advection-diffusion equation is discretized on the time mesh $t_i = \frac{.11i}{44}$, $i = 0, 2, \dots, 44$, and solved using backward Euler time stepping. Both the pressure and contaminant concentration are spatially discretized with a 2,601 degrees of freedom linear finite element approximation on a rectangular grid. The spatially dependent parameters are discretized on a coarser mesh (256 degrees of freedom) to enforce smoothness of the perturbations. The data $d_{i,j}$ is generated by solving on a mesh with 10,201 degrees of freedom. A regularization coefficient $\alpha = 0.0005$ is used and the optimization problem is solved using a truncated CG trust region solver in ROL.

The contaminant source used to generate synthetic data is given in the left panel of Figure 9 and the source estimated by solving the inverse problem is given in the right panel of Figure 9. The location and magnitude of the source is recovered well, albeit, the estimate is noisy as a result of the sparse noisy data and uninformed regularization.

The time evolution of the contaminant concentration is given in Figure 10 where three time instances are shown ($t = 0.02$, $t = 0.06$, and $t = 0.10$) from left to right. The observed noisy data is given on the top row and the estimated contaminant concentration (using the estimated source) is given on the bottom row.

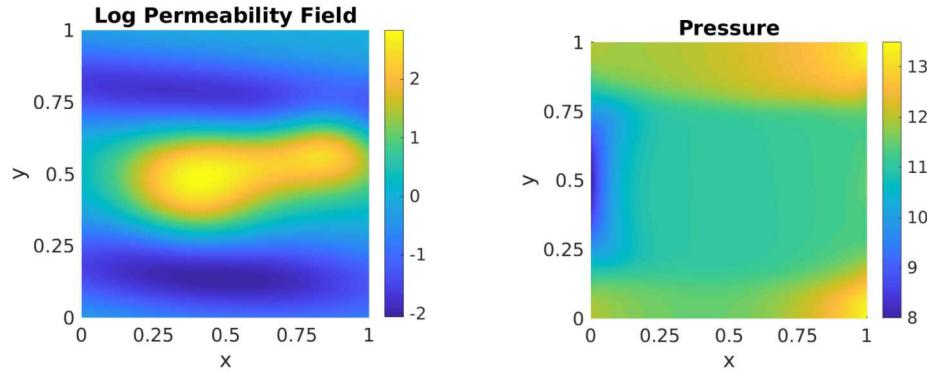


FIG. 8: Left: the log (base 10) of the nominal permeability field $\bar{\kappa}$; right: the solution of the pressure equation with nominal parameters $\bar{\kappa}$ and $\bar{\psi}$.

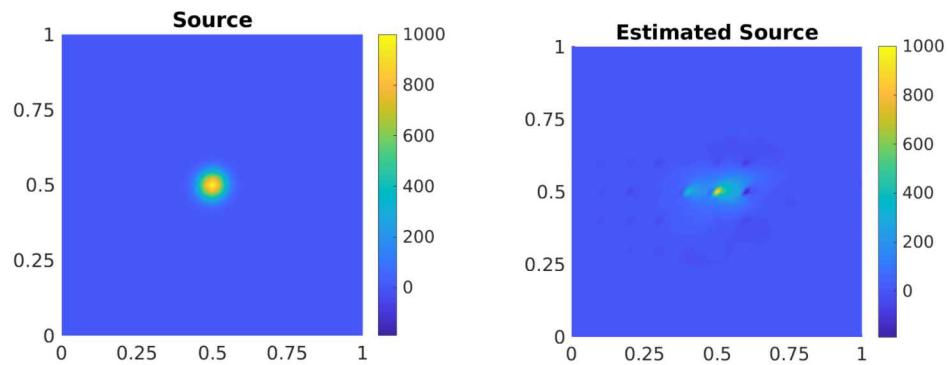


FIG. 9: Left: the “true” source term used to generate synthetic data; right: the estimated source term computed by solving (19) with sparse noisy data.

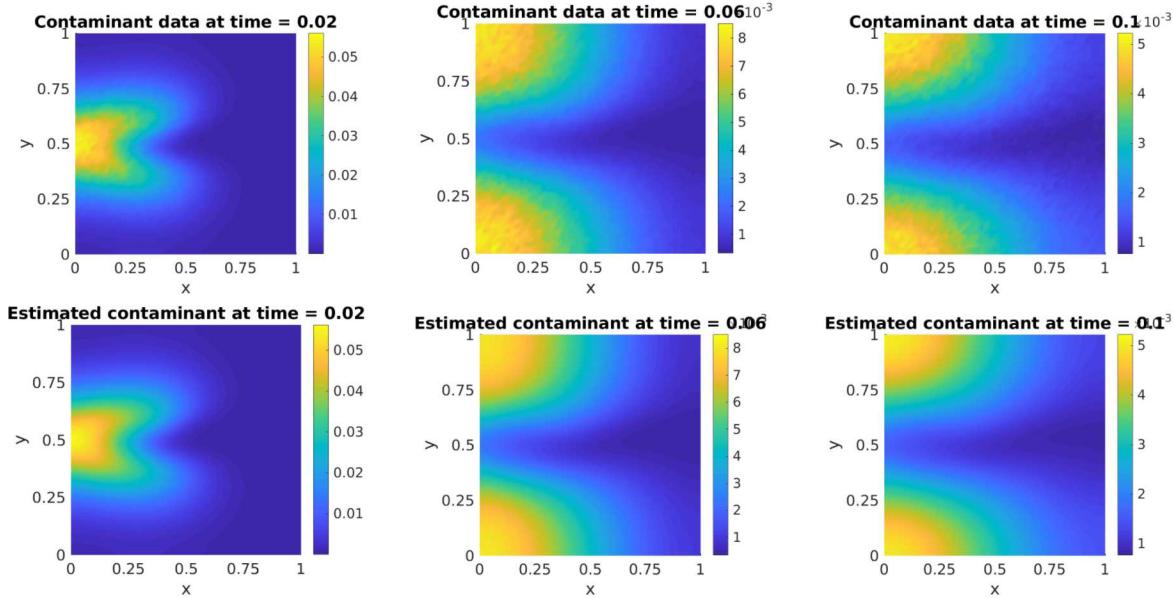


FIG. 10: Time evolution of the contaminant concentration; from left to right, the solutions at times 0.02, 0.06, and 0.10. The top row corresponds to the “true” state polluted by noise; the bottom row corresponds to the estimated state using the estimated source in the right panel of Figure 9.

In this example, we only consider local sensitivity analysis (around the parameter nominal estimates). In the case of inverse problems, we interpret the sensitivity indices as the bias created in our source estimation as a result of potentially misspecifying other parameters (permeability, boundary conditions, diffusion). The sensitivity indices indicate where we should invest efforts in characterizing the parameters in order to ensure a high fidelity source estimation, the leading singular value provides an error bound on the source estimation, and the source singular vectors indicate which features of the estimate will change given the uncertainty in the parameters.

We compute the leading $K = 12$ singular triples of (5) with an oversampling factor of $L = 8$ is used. Figure 11 shows the leading 12 singular values with an order of magnitude decay from first to last.

The sensitivity indices are shown in Figure 12. The left, center, and right panels gives the sensitivities for the parameterization of the permeability field κ , boundary condition ψ , and diffusion coefficient ϵ . The center of the basis functions (ϕ 's and η 's) correspond to the spatial location of the sensitivity indices (dots) in the left and center

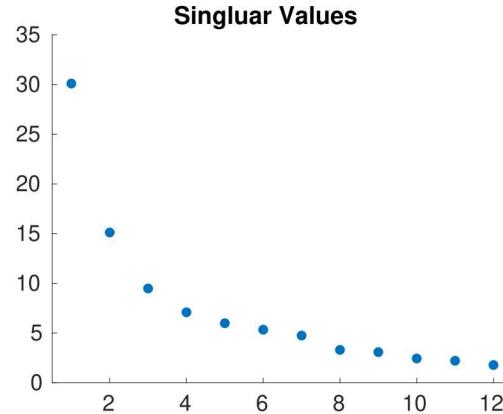


FIG. 11: The 12 leading singular values for the source estimation inverse problem sensitivities.

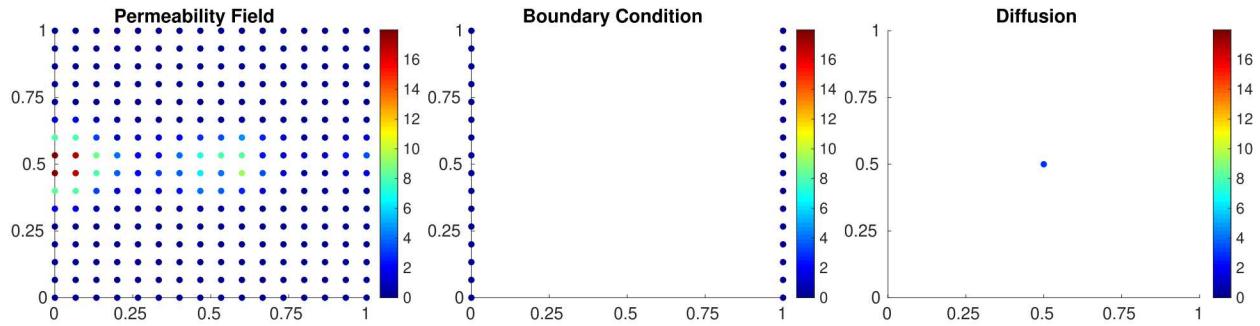


FIG. 12: Sensitivity indices for the parameterization of the permeability field (left), pressure equation Dirichlet boundary conditions (center), and diffusion coefficient (right).

panels, the diffusion coefficient is a scalar and hence only one sensitivity index (dot) is present.

We see greater sensitivity in the permeability field which is primarily localized in the region around $x = 0.05$ and $y = 0.50$. We also observe a slight asymmetry which is consistent with the asymmetry observed in the state solution. The diffusion coefficient sensitivity is smaller than the largest permeability sensitivities while the pressure equation Dirichlet boundary condition sensitivities are nearly zero. The set sensitivity indices for the permeability field, diffusion coefficient, and left and right Dirichlet boundary condition are given in Table 1. The set sensitivity index for the permeability field is larger than any individual sensitivity index, indicating that a unit norm perturbation of the entire field (perturbing multiple basis functions) will generate a greater change in the source estimate than any unit norm perturbation on a single basis element, this result is unsurprising. The leading left singular vector (omitted for conciseness), indicates that such a perturbation will result in a source estimate with greater contaminant upstream in the flow direction.

Set of Parameters	Permeability	Left BC	Right BC	Diffusion
Set Sensitivity Index	30.09	2.94	4.86×10^{-6}	4.83×10^{-6}

TABLE 1: Estimated set sensitivity indices for the parameters corresponding to the permeability field, left Dirichlet boundary condition, right Dirichlet boundary condition, and diffusion coefficient.

The computation was performed using 32 processors and 32 random vectors in the eigenvalue solver. This enables the eigenvalue solver to execute all 32 matrix vector products simultaneously and reduce the overall execution time to approximately 4 KKT solves.

6. CONCLUSION

We have introduced hyper-differential sensitivity analysis (HDSA) arguing for a goal-oriented paradigm for sensitivity analysis in the context of PDE-constrained optimization. We have shown through an analytic example that HDSA is different from traditional sensitivity analysis. The complication of this approach is that directional derivatives are required through the optimality conditions, which includes partial differential equations as constraints. The high computational requirements are addressed through efficient adjoint-based methods for the inner optimization problem and randomized generalized eigenvalue solvers to estimate the sensitivity indices. We explore three levels of parallelism including the underlying linear algebra constructs, parallel randomized algorithms, and global sensitivity sampling procedures. Two numerical examples demonstrate the flexibility of HDSA in that uncertainties in both control and inverse problems can be addressed, in addition to demonstrating the extensibility to different sets of PDEs (steady state and transient). HDSA provides a computational tractable approach to prioritize large numbers of uncertain parameters relevant to the optimal solution of PDE-constrained optimization problems. A global sensitivity strategy has been developed with intuition through a theoretical bound that depends on the underlying nonlinearity. The numerical examples demonstrate the challenge managing uncertainties in an optimization problem and how HDSA provides

critical insight that can guide subsequent analysis such as field measurements, laboratory experiments, physics development, and robust optimization.

HDSA facilitates new exploration by asking and answering questions that augment other UQ methodologies in the service of optimal design and decision making. A particular focus is given to multi-physics applications and the computational efficiency needed to explore high dimensional parameter spaces which contain spatial and/or temporally dependent parameters corresponding to a variety of physical quantities. The abstraction and generality of the proposed method permits various extensions and applications. For instance, we may consider parameters in the objective function such as weights, data, or user specified algorithmic parameters. The parallelism of our method, matrix free software design, and underlying Trilinos constructs facilitates computationally scalability. By combining mathematical abstraction, efficient software infrastructure, and exploitation of low rank structure (when present), HDSA enables us to explore a variety of UQ questions in the context of optimization constrained by PDEs.

ACKNOWLEDGMENTS

The authors are grateful to Arvind Saibaba for helpful discussions facilitating this work. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. SAND2019-XXXX.

REFERENCES

1. Ghanem, R., Higdon, D., and Owhadi, H. (Eds.), *Handbook of Uncertainty Quantification*, Springer, 2016.
2. Vogel, C.R., Sparse matrix computations arising in distributed parameter identification, *SIAM J. Matrix Anal. Appl.*, pp. 1027–1037, 1999.
3. Ascher, U.M. and Haber, E., Grid refinement and scaling for distributed parameter estimation problems, *Inverse Problems*, 17:571–590, 2001.
4. Haber, E. and Ascher, U.M., Preconditioned all-at-once methods for large, sparse parameter estimation problems, *Inverse Problems*, 17:1847–1864, 2001.
5. Vogel, C.R., *Computational Methods for Inverse Problems*, SIAM Frontiers in Applied Mathematics Series, 2002.
6. L. T. Biegler, O. Ghattas, M.H. and van Bloemen Waanders, B. (Eds.), *Large-Scale PDE-Constrained Optimization*, Vol. 30, Springer-Verlag Lecture Notes in Computational Science and Engineering, 2003.
7. Biros, G. and Ghattas, O., Parallel Lagrange-Newton-Krylov-Schur methods for PDE-constrained optimization. Parts I-II, *SIAM J. Sci. Comput.*, 27:687–738, 2005.
8. Laird, C.D., Biegler, L.T., van Bloemen Waanders, B., and Bartlett, R.A., Time dependent contaminant source determination for municipal water networks using large scale optimization, *ASCE J. Water Res. Mgt. Plan.*, pp. 125–134, 2005.
9. Hintermuller, M. and Vicente, L.N., Space mapping for optimal control of partial differential equations, *SIAM J. Opt.*, 15:1002–1025, 2005.
10. Hazra, S.B. and Schulz, V., Simultaneous pseudo-timestepping for aerodynamic shape optimization problems with state constraints, *SIAM J. Sci. Comput.*, 28:1078–1099, 2006.
11. Biegler, L.T., Ghattas, O., Heinkenschloss, M., Keyes, D., and van Bloemen Waanders, B. (Eds.), *Real-Time PDE-Constrained Optimization*, Vol. 3, SIAM Computational Science and Engineering, 2007.
12. Borzi, A., High-order discretization and multigrid solution of elliptic nonlinear constrained optimal control problems, *J. Comp. Applied Math*, 200:67–85, 2007.
13. Hinze, M., Pinnau, R., Ulbrich, M., and Ulbrich, S., *Optimization with PDE Constraints*, Springer, 2009.
14. Biegler, L., Biros, G., Ghattas, O., Heinkenschloss, M., Keyes, D., Mallick, B., Marzouk, Y., Tenorio, L., van Bloemen Waanders, B., and Willcox, K. (Eds.), *Large-Scale Inverse Problems and Quantification of Uncertainty*, John Wiley and Sons, 2011.

15. Iooss, B. and Saltelli, A. Introduction to sensitivity analysis. In: Ghanem, R., Higdon, D., and Owhadi, H. (Eds.), *Handbook for Uncertainty Quantification*, pp. 1103–1122. Springer, 2016.
16. Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., and Tarantola, S., *Global sensitivity analysis: the primer*, Wiley, 2008.
17. Iooss, B. and Lemaître, P. A review on global analysis methods. In: Dellino, G. and Meloni, C. (Eds.), *Uncertainty management in simulation-optimization of complex systems*, chapter 5, pp. 543–501. Springer, 2015.
18. Borgonovo, E., A new uncertainty importance measure, *Reliability Eng. Sys. Safety*, 92:771–784, 2007.
19. Kucherenko, S. and Iooss, B. *Handbook of Uncertainty Quantification*, chapter Derivative-based global sensitivity measures. Springer, 2016.
20. Prieur, C. and Tarantola, S. Variance-based sensitivity analysis: Theory and estimation algorithms. In: Ghanem, R., Higdon, D., and Owhadi, H. (Eds.), *Handbook for Uncertainty Quantification*, pp. 1217–1239. Springer, 2016.
21. Song, E., Nelson, B.L., and Staum, J., Shapley effects for global sensitivity analysis: Theory and computation, *SIAM/ASA J. Uncertain. Quantif.*, 4:1060–1083, 2016.
22. Brandes, K. and Griesse, R., Quantitative stability analysis of optimal solutions in PDE-constrained optimization, *Journal of Computational and Applied Mathematics*, 206:908–926, 2007.
23. Büskens, C. and Griesse, R., Parametric sensitivity analysis of perturbed PDE optimal control problems with state and control constraints, *Journal of Optimization Theory and Applications*, 131(1):17–35, 2006.
24. Griesse, R., Parametric sensitivity analysis in optimal control of a reaction-diffusion system – part II: practical methods and examples, *Optimization Methods and Software*, 19(2):217–242, 2004.
25. Griesse, R., Parametric sensitivity analysis in optimal control of a reaction diffusion system. I. solution differentiability, *Numerical Functional Analysis and Optimization*, 25(1-2):93–117, 2004.
26. Griesse, R. Stability and sensitivity analysis in optimal control of partial differential equations, 2007.
27. Griesse, R. and Vexler, B., Numerical sensitivity analysis for the quantity of interest in PDE-Constrained optimization, *SIAM Journal on Scientific Computing*, 29(1):22–48, 2007.
28. Griesse, R. and Volkwein, S., Parametric sensitivity analysis for optimal boundary control of a 3D reaction-difusion system, In: Pillo, G.D. and Roma, M. (Eds.), *Nonconvex Optimization and its Applications*, Vol. 83. Springer, Berlin, 2006.
29. Griesse, R. and Walther, A., Parametric sensitivities for optimal control problems using automatic differentiation, *Optimal Control Applications and Methods*, 24:297–314, 2003.
30. Saibaba, A.K., Lee, J., and Kitanidis, P.K., Randomized algorithms for generalized Hermitian eigenvalue problems with application to computing Karhunen-Loève expansion, *Numerical Linear Algebra with Applications*, 23:314–339, 2016.
31. Murthy, P.R., *Operations Research*, New Age International Publishers, 2nd edition, 2007.
32. Bonnans, J.F. and Shapiro, A., Optimization problems with perturbations: A guided tour, *SIAM Review*, 40(2):228–264, 1998.
33. Rakovec, O., Hill, M.C., Clark, M.P., Weerts, A.H., Teuling, A.J., and Uijlenhoet, R., Distributed evaluation of local sensitivity analysis (delsa), with application to hydrologic models, *Water Resources Research*, 50:409–426, 2014.
34. Sobol', I. and Kucherenko, S., Derivative based global sensitivity measures and the link with global sensitivity indices, *Math. Comp. Simul.*, 79:3009–30017, 2009.
35. Sobol', I. and Kucherenko, S., A new derivative based importance criterion for groups of variables and its link with the global sensitivity indices, *Comput. Phys. Comm.*, 181:1212–1217, 2010.
36. Morris, M., Factorial sampling plans for preliminary computational experiments, *Technometrics*, 33:161–174, 1991.
37. Constantine, P.G., *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*, SIAM, 2015.
38. Constantine, P.G. and Diaz, P., Global sensitivity metrics from active subspaces, *Reliability Engineering & System Safety*, 162:1–13, 2017.
39. Fassbender, H. and Kressner, D., Structured eigenvalue problems, *GAMM-Mitteilungen*, 29(2):297–318, 2006.
40. Halko, N., Martinsson, P.G., and Tropp, J.A., Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Review*, 53(2):217–288, 2011.
41. Ito, K. and Ravindran, S.S., Optimal control of thermally convected fluid flows, *SIAM J. on Scientific Computing*, 19:1847–1869, 1998.

42. Kouri, D.P. and Ridzal, D. *Inexact Trust-Region Methods for PDE-Constrained Optimization*, pp. 83–121. Springer New York, New York, NY, 2018.
43. Kouri, D.P., von Winckel, G., and Ridzal, D. ROL: Rapid Optimization Library.
44. Heroux, M., Bartlett, R., Howle, V., Hoekstra, R., Hu, J., Kolda, T., Lehoucq, R., Long, K., Pawlowski, R., Phipps, E., Salinger, A., Thornquist, H., Tuminaro, R., Willenbring, J., and Williams, A., An overview of Trilinos, Tech. Rep. SAND2003-2927, Sandia National Laboratories, 2003.