Chapter 13
Optimal Experiment Design for Large-Scale Inverse Problems with Enhanced Robustness to Model Uncertainties

Wilkins Aquino, Mark Chen, Drew Kouri, Kathryn Maupin, Joshua Mullins, Chandler Smith, Timothy Walsh, and Rebekah White

Abstract In structural dynamics, response data obtained from laboratory-controlled vibration tests or real-time monitoring are essential to inform uncertain model parameters such as structural loads, boundary conditions, or material properties. Often, response data are limited due to cost or physical constraints on the number of sensors that can be placed in an experiment. Optimal experimental design (OED) is a powerful tool for judiciously choosing what experiments will provide the most informative data, i.e. the data that reduces uncertainty in resulting parameter estimates, which is characterized by scalar-valued functions of the associated asymptotic covariance. However, such classical formulations of OED ignore the impact of uncertainties associated with auxiliary parameters – model parameters not inferred from experimental data but whose values impact the accuracy of model predictions. Therefore, this works considers more robust OED strategies that account for uncertainty in the auxiliary parameters when predicting an optimal experimental design. Specifically, we characterize uncertainty in the auxiliary parameters using probability distributions; the classical alphabetical criteria (e.g. I- or D-optimality) are then computed with respect to this distribution. To ensure a scalar-valued OED objective function, minimax and expectation-based approaches are used to minimize worst-case performance or maximize average performance, respectively. We consider two approaches for solving the optimization problem determining the best design: 1) utilizing a greedy algorithm for binary weights and 2) performing gradient-based optimization for fractional weights. To overcome the computational expense associated with the large-scale finite element models (FEM) structural dynamics code, we develop a reduced order model (ROM) to utilize within the robust OED strategies. Specifically, a reduced basis method (based on formal *a posteriori* error estimation concepts) is used to construct the ROM. Using an optimal control inverse problem within the context a ground-based vibration test, we demonstrate the effectiveness of robust OED in comparison to standard OED approaches. Furthermore, we demonstrate how computational intractability inherent to large-scale structural dynamics problems can be overcome through integration of robust OED strategies with reduced order modeling.

Keywords Robust Optimal Experimental Design · Large-Scale Inverse Problems · Experimental Design Under Uncertainty · Reduced Order Modeling

Introduction

Across many engineering applications, accurately representing a physical system of interest of requires large-scale finite-element models (FEM), which can have on the order of millions of degrees of freedom. Furthermore, the parameters comprising such models are often uncertain or even unknown, requiring experimental data for calibration. Within the field of structural dynamics, inverse methods are employed to estimate unknown model parameters such as structural loads or material properties, using vibration response data. One challenge typical to inverse methods is the extraction of comprehensive information regarding the target unknown parameters from, what is often, limited experimental data; for example, laboratory-controlled vibration tests or real-time monitoring, can be costly, limiting the amount of feasible tests. Furthermore, within an

Wilkins Aquino
Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC, 27708
e-mail : wa20@duke.edu

Mark Chen · Drew Kouri · Kathryn Maupin · Joshua Mullins · Chandler Smith · Timothy Walsh · Rebekah White
Sandia National Laboratories, Albuquerque, NM, 87112
e-mail: marchen@sandia.gov; dpkouri@sandia.gov; kmaupin@sandia.gov; jmullin@sandia.gov; chasmit@sandia.gov; tfwalsh@sandia.gov; rebwhit@sandia.gov

Roland Platz et al. (eds.), *Model Validation and Uncertainty Quantification, Vol. 3*, Conference Proceedings of the Society for Experimental Mechanics Series, https://doi.org/10.13052/97887-438-0148-1_13

experiment, there are limitations to the number of sensors that can be placed on a physical system. With such limited data, it is imperative to perform experiments that yield the most information.

Optimal experimental design (OED) is a valuable tool for predicting *a priori* what experiments will be most informative given an experimental budget. An often employed, standard approach to OED, is to relate the informativeness of data to how well that data allows for the estimation (calibration) of the underlying model parameters; here M-estimation (e.g. least squares, maximum likelihood) [1] is used to estimate the model parameters while the corresponding asymptotic covariance provides a measure of uncertainty in the estimate [2]. As a result, scalar-valued functions of the asymptotic covariance are often used to formulate an OED objective function. Examples include the trace (A-optimal) and determinant (D-optimal) of the covariance. Within the context of M-estimation, robust design approaches have been considered. For example [3] formulates designs that are robust to misspecified response functions (given by additive model corrections) and unexpected heteroscedasticity. For nonlinear models, the optimality criterion depends upon the estimated parameters. Thus, uncertainty in the estimated parameters affects the design; approaches accounting for this uncertainty include [4, 5], which leverage minimax strategies, and [6], which leverages expectation-based strategies. Such formulations, however, do not account for uncertainties associated with auxiliary parameters, i.e. model parameters not inferred from experimental data but whose values impact the accuracy of structural dynamics model predictions.

Therefore, this work leverages robust OED strategies to account for uncertainty in the auxiliary parameters when predicting optimal experimental designs. Specifically, we characterize uncertainty in the auxiliary parameters using probability distributions. The OED objective is then computed with respect to this distribution using minimax and expectation-based approaches. The probabilistic representation of and subsequent marginalization over the auxiliary parameters follows from similar approaches in a Bayesian paradigm; see, e.g. [7]. However, here the OED criteria is formulated with frequentist estimates of the asymptotic covariance and risk measures beyond just the expected value are employed. We consider a structural dynamics model acted upon by unknown loads with material model uncertainty, where the experimental design is defined as the location at which sensors measuring vibration response are placed. For such sensor placement problems, the experimental design corresponds to weights of one or zero, indicating whether a sensor is placed or not, respectively; however, combinatorial complexity makes solving such optimization problems intractable. Thus, this work considers 1) utilizing a greedy approach, which selects an optimal sensor location one at time [8] and 2) assigning probability masses (i.e. fractional weights) to the sensor locations, which allows for gradient-based optimization to be employed [9]. To overcome the computational expense associated with large-scale finite element models (FEM), we develop a reduced order model (ROM) to utilize within the robust OED strategies. Specifically, a reduced basis method (based on formal *a posteriori* error estimation concepts) is used to construct the ROM.

The main contributions of this work are as follows:

1. Demonstrate the effectiveness of robust OED in comparison to standard OED approaches for a structural dynamics force estimation inverse problem.
2. Provide a comparison between utilizing a greedy algorithm for binary weights and performing gradient-based optimization on weights representing probability masses.
3. Demonstrate how a reduced order model can enable computational tractability of robust OED strategies for structural dynamics applications.

The remainder of this paper is organized as follows: We begin by providing the relevant background for robust OED generally, highlighting how gradient-based and greedy approaches enable tractable optimization. Following this, we introduce the structural dynamics problem of interest, the wedding cake model, and provide the formulation for the ROM utilized within the OED framework. We then present the computational results, which include verification of the ROM accuracy, verifying the proposed greedy-based method, and comparing the gradient-based optimization strategy to the greedy method in terms of model prediction error. Lastly, final conclusions are drawn.

Background and Methods

First, we present the general framework for robust optimal experimental design (OED). In this work's approach to robust OED, the aim is determine what noisy observational data $\boldsymbol{y}_\epsilon \in \mathscr{Y} \subset \mathbb{R}^{n_y}$ best inform a subset of parameters $\boldsymbol{\theta} = [\theta_1, \ldots, \theta_{n_p}] \in \Theta \subset \mathbb{R}^{n_p}$ of a mathematical model $f(\boldsymbol{\theta}, \boldsymbol{\xi}; \boldsymbol{d}) : \Theta \times \Xi \to \mathscr{Y}$, while accounting for the *irreducible* uncertainty in auxiliary parameters $\boldsymbol{\xi} = [\xi_1, \ldots, \xi_{n_a}] \in \Xi \subset \mathbb{R}^{n_a}$, characterized by the probability distribution $\pi(\boldsymbol{\xi})$. The experimental design vector is denoted $\boldsymbol{d} \in \mathscr{D} \subset \mathbb{R}^{n_d}$ and controls how data is collected; for example, $\boldsymbol{d}$ may represent spatial locations at which response measurements are taken. We assume additive noise, and thus model the data as

$$\boldsymbol{y}_\epsilon = f(\boldsymbol{\theta}, \boldsymbol{\xi}; \boldsymbol{d}) + \boldsymbol{\epsilon}, \tag{1}$$

where $\boldsymbol{\epsilon} \sim \mathscr{N}(\mathbf{0}, \sigma^2 \boldsymbol{I})$, for an identity matrix $\boldsymbol{I} \in \mathbb{R}^{n_y \times n_y}$. Let the asymptotic covariance associated with the M-estimation of $\boldsymbol{\theta}$ be denoted $C(\boldsymbol{d}) \in \mathbb{R}^{n_p \times n_p}$. The OED criterion $\Psi(\boldsymbol{d}, \boldsymbol{\xi})$ provides a scalar-valued representation of the covariance matrix; classical choices include D- and A-optimal criteria, which represent the determinant and trace of the covariance matrix, respectively. Then, the OED objective $U : \mathscr{D} \to \mathbb{R}$ characterizing how "good" a given experimental design $\boldsymbol{d}$ is can be formulated in a robust way as follows:

$$U(\boldsymbol{d}) = \mathscr{R}_{\boldsymbol{\xi}} \left[\Psi(\boldsymbol{d}, \boldsymbol{\xi}) \right], \tag{2}$$

where $\mathscr{R}_{\boldsymbol{\xi}}$ represents a risk measure that allows one to compute an "average" (with respect to $\pi(\boldsymbol{\xi})$) objective. For example, letting $\mathscr{R}_{\boldsymbol{\xi}} := \mathbb{E}_{\boldsymbol{\xi}}$ would result in the following expectation-based robust objective

$$U(\boldsymbol{d}) = \int_{\Xi} \Psi(\boldsymbol{d}, \boldsymbol{\xi}) d\pi(\boldsymbol{\xi}), \tag{3}$$

which we refer to as a risk-neutral approach. Alternatively, one could consider risk measures that penalize tail statistics, such as the average value-at-risk [10]. The optimal design $\boldsymbol{d}^*$ is the one solving the following optimization problem

$$\boldsymbol{d}^* = \arg\min_{\boldsymbol{d} \in \mathscr{D}} U(\boldsymbol{d}). \tag{4}$$

In general, sensor placement problems utilizing binary weights result in (4) being combinatorially complex and thus intractable for expensive high-fidelity models and high-dimensional design spaces. Therefore this work considers two approaches to solving (4). The first is to consider a greedy optimization strategy, which aims to determine optimal sensor locations one (or more generally one batch) at a time. Here, the design is a vector of weights $\boldsymbol{d} = [d_1, \ldots, d_{n_d}]$, where $d_i \in \{0, 1\}$ for all i, and the corresponding optimization problem is given as

$$\boldsymbol{d}^* = \arg\min_{\boldsymbol{d} \in \{0,1\}} U(\boldsymbol{d}), \quad \sum_i d_i = N, \tag{5}$$

where N represents the sensor budget. The design weights d_i are determined sequentially, providing an approach that scales with the number of candidate sensor locations and requires $\mathscr{O}(N n_d)$ computations. Following [8], Algorithm 1 describes the greedy procedure. While not being guaranteed to find the optimal design, greedy algorithms have been shown to perform well in practice [11] and have theoretical justifications; see, e.g. [12, 13].

Algorithm 1: Greedy sensor placement

Input: Target number of sensors N.
Output: Design vector $\boldsymbol{d}$
set $\boldsymbol{d} = \mathbf{0}$, $\mathscr{U} = \{1, \ldots, n_d\}$, and $\mathscr{S} = \emptyset$.
for $k = 1 : n_d$ **do**
    $i = \arg\min_{j \in \mathscr{U} \setminus \mathscr{S}} U(\boldsymbol{d} + \boldsymbol{e}_j)$ $\boldsymbol{e}_j$ is the jth coordinate vector in $\mathbb{R}^{n_d}$
    $\mathscr{S} = \mathscr{S} \cup \{i\}$
    $\boldsymbol{d} = \boldsymbol{d} + \boldsymbol{e}_j$
end

An alternative to utilizing a greedy algorithm is to instead allow the design vector to represent probability masses and solve the following optimization problem

$$\boldsymbol{d}^* = \arg\min_{\boldsymbol{d} \in [0,1]} U(\boldsymbol{d}), \quad \sum_i d_i = 1, \tag{6}$$

utilizing gradient-based optimization. Here, design vector weights reflect a frequency at which each each sensor is utilized; thus, in practice, the larger the weight, the "better" the sensor location. Note that although this framing does not naturally enforce a given experimental budget N, for the problem considered herein, the number of candidate sensor locations is significantly larger than the number of inversion parameters, guaranteeing a sparse design [14]. Furthermore, heuristic approaches can be employed that select the N largest weights and utilize rounding and rescaling to ensure the interpretation of weights as probability masses.

Model problem and reduced-order basis construction

Here, we present the so-called wedding cake model depicted in Figure 1a used to demonstrate the robust OED approach. We also present the reduced order model formulation that enables computationally efficient evaluation of the robust OED objective function. Vibration response is modeled mathematically according to the following linear parameter-to-observable map

$$f(\boldsymbol{\theta} := s, \boldsymbol{\xi}) := H(\boldsymbol{\xi})\boldsymbol{s}, \tag{7}$$

where H is the frequency response function (FRF) matrix, $\boldsymbol{s} \in \mathbb{C}^{N_p}$ are the applied forces we wish to infer from experimental data, and $\boldsymbol{\xi} \in \Xi \subseteq \mathbb{R}^{n_a}$ are auxiliary parameters representing frequency as well as the density, elastic modulus, and Poisson's ratio of the columns on each floor, thus totaling ten auxiliary parameters.

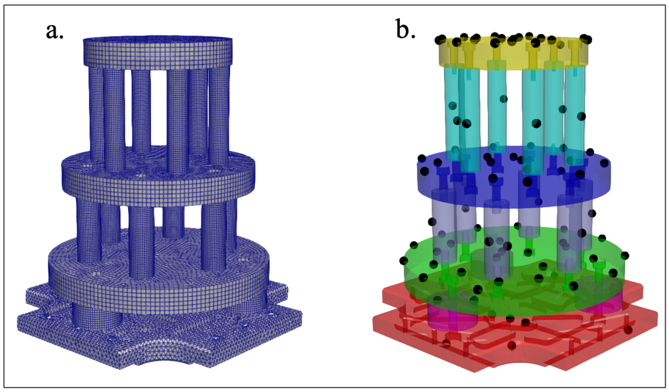

Fig. 1 a) Finite element model of the wedding cake and b) candidate sensor locations superimposed over the model. Uniaxial sensors can be oriented in one of three degrees of freedom.

Irreducible uncertainty in the auxiliary parameters is characterized by a multivariate uniform distribution. We assume a large variance in the uncertainty to encourage significant changes in the modal properties of the structure, in-turn eliciting complex dynamic response behavior.

The high-fidelity (HiFi) model solution is computed using Sandia's Sierra-SD FEM software to solve the finite element solution of the Helmholtz equations [15]. The frequency range of interest is 100Hz to 500Hz spaced at 8Hz intervals, thus totaling 50 frequency lines; on 60 processors, the runtime for one HiFi simulation using 50 frequencies is ~ 10 minutes. Due to the computational costs of evaluating the HiFi model at every sample of the auxiliary parameters, a reduced order model (ROM) is needed to make subsequent experimental design computationally tractable.

We now provide an abbreviated overview of the adaptive reduced basis method of reduced order modeling in the context of approximating the FRF matrix. The formulations and algorithms presented herein are based on the work published in [16]. Consider the following equation for the finite element (FE) approximated Helmholtz problem describing the vibration response of a linear, time-invariant structure with n total degrees of freedom (DoFs),

$$\mathscr{M}(\hat{u}; \boldsymbol{\xi}, \omega) \equiv A(\boldsymbol{\xi})\hat{u} - F = 0, \tag{8}$$

$\hat{u}(\boldsymbol{\xi}) \in \mathscr{U}$ is the solution of the Helmholtz problem representing the vibration response, F is the forcing function, ω is the circular frequency, and $\mathscr{U} \subseteq \mathbb{C}^n$ represents the space of solutions. Note that $A(\boldsymbol{\xi}) \equiv K(\boldsymbol{\xi}) + j\omega C(\boldsymbol{\xi}) - \omega^2 M(\boldsymbol{\xi})$, where K, M, and C are the stiffness, mass, and damping matrices, respectively.

Our goal is to construct an efficient approximation (i.e. avoiding the solution to (8)) $\bar{u} : \Xi \to \mathscr{U}$ that is close to $\hat{u}$ (in some sense). Consider a set of samples $\Theta := \{\boldsymbol{\xi}_i\}_{i=1}^k$, referred to as atoms, drawn from the material parameter set Ξ. According to the reduced basis (RB) method [16], a finite dimensional subspace $\mathscr{U}_k$ for the true solution space $\mathscr{U}$, is given as

$$\mathscr{U}_k \equiv \text{span}(\Phi_k := \{\hat{u}_1, \hat{u}_2, \ldots, \hat{u}_k\}), \tag{9}$$

where $\hat{u}_k := \hat{u}(\boldsymbol{\xi}_k)$ is the solution to the Helmholtz problem given $\boldsymbol{\xi}_k$, and $\Phi_k \in \mathbb{C}^{n \times k}$. We can now use (9) to define a reduced-basis approximation to the model solution as

$$\bar{u}_k(\boldsymbol{\xi}) := \Phi_k a_k(\boldsymbol{\xi}), \tag{10}$$

where $a_k(\boldsymbol{\xi})$ solves the reduced problem

$$\Phi_k^T A(\boldsymbol{\xi}) \Phi_k \, a_k = \Phi_k^T F. \tag{11}$$

The computational savings stem from the small number of basis vectors in Φ_k, where the number of basis vectors k is much smaller than the dimension of the state space n.

In this work, the construction of the subspace $\mathscr{U}_k$ is performed adaptively by sequentially selecting atoms based on the current surrogate solution $\hat{u}$. Given a current set of atoms, we select the next atom $\hat{\boldsymbol{\xi}}_{k+1}$ by identifying the atom corresponding to the largest expected ROM error over the parameter space. To avoid directly computing the ROM solution error, we estimate the error using a reliable *a posteriori* error estimate $\epsilon_u(\boldsymbol{\xi})$. Let the solution error be given by $e(\boldsymbol{\xi}) := \hat{u}(\boldsymbol{\xi}) - \bar{u}(\boldsymbol{\xi})$, and define the residual as $r(\boldsymbol{\xi}) := F(z) - H(\boldsymbol{\xi})\bar{u}(\boldsymbol{\xi})$. We can show that (see [16])

$$\|e(\boldsymbol{\xi})\|_U \leq \frac{\|r(\boldsymbol{\xi})\|}{\gamma(\boldsymbol{\xi})} =: \epsilon_u(\boldsymbol{\xi}), \tag{12}$$

where the inf-sup constant is given as $\gamma(\boldsymbol{\xi}) := \min_i |\lambda_i(\boldsymbol{\xi})|$, with $\lambda_i(\boldsymbol{\xi}), i = 1, ..., n$ being the eigenvalues of $H(\boldsymbol{\xi})$.

There are three major computational hurdles related to training the ROM: evaluating the inf-sup constant, reassembling the operator H, and training the ROM for a wide band of circular frequencies. In this work, to avoid solving a costly eigenvalue problem, we ignore the inf-sup constant and take the error estimator as simply $||r(\boldsymbol{\xi})||_2$. To address the large computational cost of assembling H for all background samples of $\boldsymbol{\xi}$, we exploit the mathematical structure when the operator H is affine in the uncertain parameters; see [16] for details on the affine decomposition. Finally, to avoid computing a ROM for each frequency ω of interest, we treat frequency as an uncertain parameter for the training phase of the ROM only (i.e. $\omega \in \boldsymbol{\xi}$).

Computational Results

Reduced order model verification

First, we establish that the ROM for the wedding cake has adequate accuracy for use in robust OED. Utilizing the framework presented in the previous section, we train six independent ROMs using the reduced basis (RB) method, where each ROM corresponds to one column of the FRF matrix. Note that the ROM depends only on the load position and orientation, and scales linearly with the unknown load amplitude. Hence, we assume a unit load at the DoF corresponding to each unknown load when training each ROM.

In this example problem, the candidate acceleration measurement locations are shown as nodes in Figure 1b. Note that sensors can be oriented in one of three Cartesian directions, resulting in $n_d = 267$ possible sensor DoFs (locations and orientations). Acceleration data is then used to estimate the unknown inference parameters ($\boldsymbol{\theta}$) given by six coincident point loads (three translational and three rotational) that excite the base.

Each ROM requires 40 HiFi model evaluations to achieve an acceptable degree of accuracy. Figure 2 presents the ROM error estimators (see (12)) as a function of atoms (i.e. HiFi model samples); here, we see adequate decay in error to establish the fidelity of the ROM during training. Noting that each HiFi solution is evaluated at one sample of the auxiliary parameters $\boldsymbol{\xi}$ (i.e. one sample of the material properties and a single frequency), a single ROM require approximately 100 minutes (40×2.5 minutes) to train. In total, the six ROMs require approximately 10 hours to train on 60 processors.

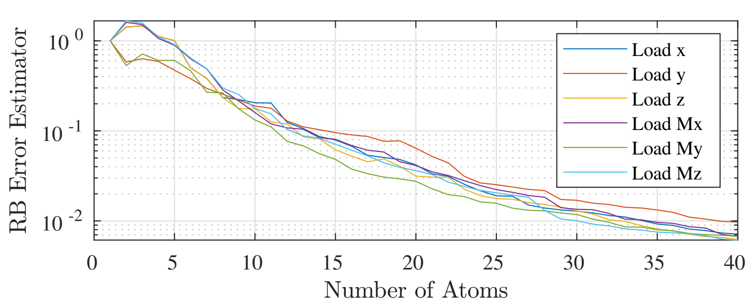

Fig. 2 The reduced basis (RB) training error estimation as a function of atoms (i.e. samples of the HiFi model). Each legend item corresponds to a load.

For verification, we compare the solution to the Helmholtz problem using the trained ROM and the HiFi finite element model. Here, we assume the forces applied to the six DoFs at the model's base are also random variables described by a standard normal distribution to avoid skewing the verification results to one particular set of loads. The force amplitude is held constant across frequencies. The verification set used to compare the trained ROM and the HiFi model is comprised of 50 samples of the applied force distribution and auxiliary parameters. For each sample, the HiFi model and ROM are evaluated across a frequency range of 100 to 500 Hz at 8 Hz intervals. Note that the ROM and HiFi models are compared at the candidate sensor DoFs only.

Figure 3a presents the relative error between the ROM and HiFi solution for each sample of the verification set, where the relative error is taken as,

$$Relative\ Error = \frac{\|\hat{u}(\boldsymbol{\xi}_k) - \bar{u}(\boldsymbol{\xi}_k)\|_2^2}{\|\hat{u}(\boldsymbol{\xi}_k)\|_2^2}, \tag{13}$$

where the l_2-norm is given by

$$\|x(\boldsymbol{\xi}_k)\|_2 = \sqrt{\sum_{i=1}^{n_\omega}\sum_{j=1}^{n_d} |x_j(\boldsymbol{\xi}_k, \omega_i)|^2}, \tag{14}$$

and where $x_j(\boldsymbol{\xi}_k, \omega_i)$ is the HiFi (or ROM) acceleration response evaluated at the j^{th} DoF for the i^{th} frequency corresponding to the random sample $\boldsymbol{\xi}_k$. From 3a, for most random samples, the relative ROM error is low; three samples corresponding to peaks in uncertainty are noted as s-6, s-9, and s-45.

To understand how the relative error as a function of uncertainty samples propagates to response error, Figure 3b presents the root mean square (RMS) of the acceleration response over all DoFs as a function of frequency using the HiFi and ROM models. In Figure 3b, parameter samples corresponding to the three largest relative ROM errors (s-6, s-9, s-45) are used for clarity. From Figure 3b, we see that across significant shifts in resonances between the three model samples, the ROM corresponds well to the HiFi solution.

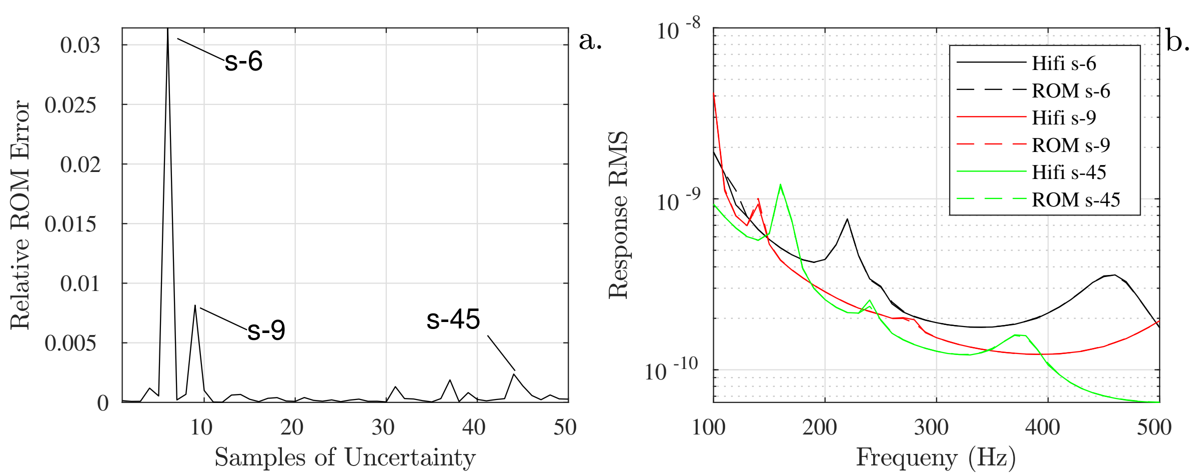

Fig. 3 a) The relative error of the ROM with respect to the HiFi model as a function of the samples of uncertainty, where the samples corresponding to the three largest uncertainties are noted as s-6, s-9, and s-45. b) The RMS responses of the HiFi model and ROM corresponding to s-6, s-9, and s-45.

Robust OED utilizing greedy optimization

With the verified accuracy, we can now evaluate robust OED approaches utilizing the ROM. Here, we focus on strategies utilizing the greedy approach presented in Algorithm 1. In this section, we consider the D-optimal expectation-based (i.e. risk-neutral) and minimax objective functions respectively given by

$$U(\boldsymbol{d}) = \mathbb{E}_{\boldsymbol{\xi}}\left[\Psi\left(\boldsymbol{d}, \boldsymbol{\xi}\right)\right] \approx \frac{1}{n_a}\sum_{i=1}^{n_a} \Psi\left(\boldsymbol{d}, \boldsymbol{\xi}_i\right), \quad \boldsymbol{\xi}_i \sim \pi(\boldsymbol{\xi}) \tag{15}$$

and,

$$U(\boldsymbol{d}) = \max_{\boldsymbol{\xi} \in \Xi} \Psi\left(\boldsymbol{d}, \boldsymbol{\xi}\right) \approx \max_{i=1\ldots n_a} \Psi\left(\boldsymbol{d}, \boldsymbol{\xi}_i\right), \quad \boldsymbol{\xi}_i \sim \pi(\boldsymbol{\xi}) \tag{16}$$

for,

$$\Psi\left(\boldsymbol{d}, \boldsymbol{\xi}\right) = \log\det\left(C(\boldsymbol{d}, \boldsymbol{\xi})\right).$$

The robust experimental design is computed using $1,000$ realizations of the material uncertainty and frequency to compute a corresponding $1,000$ FRF matrix samples. Note that without the ROM, the total runtime for utilizing the HiFi model even with a parallelized greedy approach, is on the order of days. In contrast, utilizing the ROM requires on the order of minutes to compute an optimal design. To verify the performance of a given sensor configuration, we constructed a verification set using the same samples used to solve the robust OED problem. From this verification set, we can evaluate the quality of a given sensor configuration by computing the corresponding objective function of interest.

To begin, we first compare the OED solution using the greedy algorithm to random sensor selection; here, we generate $5,000$ unique random designs with a budget of 10 sensors and evaluate the corresponding robust objectives (see (15) and (16)). Figure 4 presents the histograms of the objective function values alongside the OED solutions using the risk-neutral (a) and minimax objectives (b). From Figure 4, one can see the optimal design determined via the greedy algorithm outperforms random sensor selection.

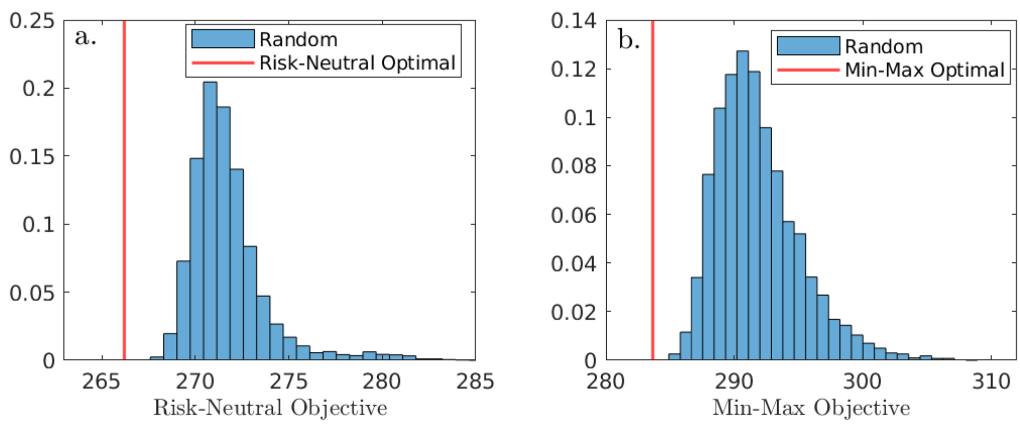

Fig. 4 Histogram comparison of $5,000$ random experimental designs versus the optimal design computed using the greedy algorithm for a) the risk-neutral objective given in (15) and b) the minimax objective given in (16).

Next, we compare the performance of the proposed robust OED approach, which accounts for uncertainties in auxiliary parameters, to standard OED formulations. To do so, we consider that the standard OED formulation utilizes an average (with respect to material uncertainty) FRF matrix. To ensure a fair comparison, the FRF matrices are also evaluated at all the frequencies of interest (i.e. 100 Hz to 500 Hz at 8 Hz intervals).

First, consider the risk-neutral objectives. Figure 5 compares the standard (c) versus robust (d) designs for a budget of $N = 10$ sensors using the risk-neutral objective. From Figure 5c and d, we see the optimal designs differ between robust and standard OED; to further explore this difference, Figure 6a shows the risk-neutral objective evaluated at the $1,000$ FRF matrix validation samples as a function of sensor budget, comparing optimal robust versus standard designs. We see from Figure 6a a negligible difference in robust versus standard risk-neutral designs, indicating uncertainty in the auxiliary parameters does not have a significant impact on the OED problem in this case.

Next, Figure 5 compares the standard (a) versus robust (b) designs for a budget of $N = 10$ sensors using the minimax objective. Unlike the risk-neutral case, we see significant differences in the designs. Furthermore, Figure 6b depicts the minimax objective over $1,000$ FRF matrix samples as a function of sensor budget; here, we see that robust designs result in significantly lower object values, indicating that when tail statistics are of interest, auxiliary uncertainty impacts the OED problem. Note this result indicates that the impact of auxiliary uncertainty on the OED problem depends upon the objective formulation and goals of experimentation. More generally, the impact of such auxiliary uncertainties is dependent upon the model problem being considered, where in some cases (such our risk-neutral formulation) the impact is negligible-however, in other cases not properly accounting for auxiliary uncertainty may have significant impacts on the optimal design.

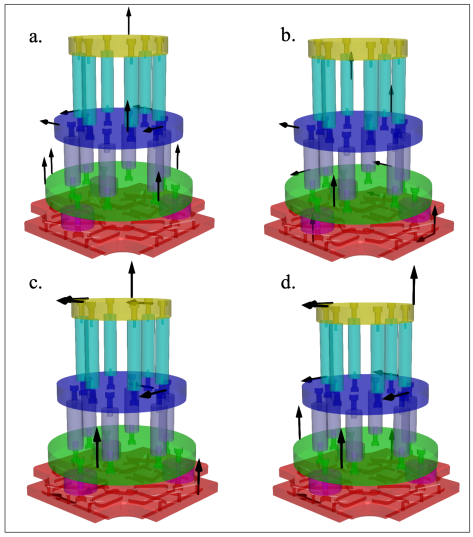

Fig. 5 A comparison of robust versus standard experimental designs on the wedding cake geometry. Note the standard designs utilize the mean auxiliary parameters: a) standard minimax optimal design b) robust minimax optimal design, c) standard risk-neutral optimal design, d) robust risk-neutral optimal design.

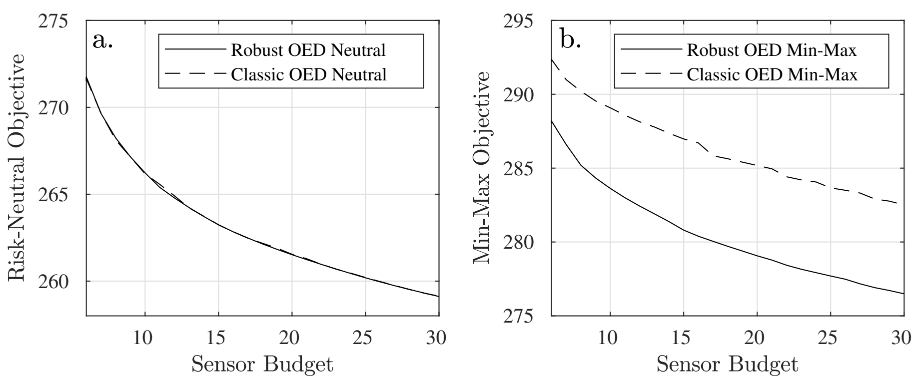

Fig. 6 Comparison of robust versus standard (e.g. classical) designs for the OED objective (averaged over $1,000$ FRF matrix samples) versus the sensor budget: a) risk-neutral objective and b) minimax objective.

Average prediction variance greedy and convex optimization investigation

We now turn our attention to the performance of the convex optimization OED formulation given in (6). Here, we focus on experimental designs that minimize model prediction uncertainty (as opposed to inference parameter uncertainty), where the variance of the prediction error at all model DoFs is used to quantify prediction uncertainty. Prediction error is defined by the difference between the predicted response $f(\tilde{\boldsymbol{\theta}}, \boldsymbol{\xi}; \boldsymbol{d}_j)$ and the model response $f(\boldsymbol{\theta}, \boldsymbol{\xi}; \boldsymbol{d}_j)$ conditioned on the auxiliary parameter uncertainty $\boldsymbol{\xi}$ for all candidate DoFs given by $\boldsymbol{d}_j \in \mathscr{D}$. The predicted response is the model response using the estimated regression parameters $\tilde{\boldsymbol{\theta}}$ obtained from the least squares solution of the inverse problem. Note that $\tilde{\boldsymbol{\theta}}$ is a function of the chosen experiment design and the measurement uncertainty.

Since the estimator is unbiased, the variance of the prediction error at the j^{th} DoF is equivalent to the prediction variance defined as $\sigma^2_{\tilde{y}_j} \equiv \mathbb{E}\left[f(\tilde{\boldsymbol{\theta}}, \boldsymbol{\xi}; \boldsymbol{d}_j)^2\right]$, where the expectation is taken with respect to the measurement noise. We then refer to the average prediction variance as the sample approximated average over the candidate DoFs and over the auxiliary parameter realizations given by

$$\frac{1}{n_a n} \sum_{k=1}^{n_a} \sum_{j=1}^{n} \sigma^2_{\tilde{y}_j}(\boldsymbol{\xi}_k, \boldsymbol{d}), \tag{17}$$

where n is the total number of DoFs, and $\sigma^2_{\tilde{y}_j}$ represents the prediction variance at the j^{th} DoF for the k^{th} sample of $\boldsymbol{\xi}$.

In this section, we consider the gradient-based optimization algorithm where $\Psi(\boldsymbol{d}, \boldsymbol{\xi})$ is given by risk-neutral formulations of the D- and I-optimal criteria. I-optimality seeks to directly minimize the average prediction variance over all candidate experiments. The risk-neutral robust formulation of the I-optimal objective is approximated as

$$U(\boldsymbol{d}) \approx \frac{1}{n_a} \sum_{k=1}^{n_a} \text{Trace}\left(H(\boldsymbol{\xi}_k) C(\boldsymbol{d}, \boldsymbol{\xi}_k) H(\boldsymbol{\xi}_k)^*\right), \quad \boldsymbol{\xi}_k \sim \pi(\boldsymbol{\xi}), \tag{18}$$

where $*$ represents the complex conjugate transpose. In contrast, G-optimality seeks to minimize the worst-case prediction variance. However, in the case of homoscedastic noise and a fixed set of auxiliary parameters, D- and G-optimality are equivalent [17], and hence we consider the objective given in (15).

In this study, we restrict the OED problem to 50 FRF matrix (i.e. auxiliary parameter) samples. To begin, we first compare the optimal designs produced by the convex optimization formulation and the greedy algorithm for the risk-neutral objective. The risk-neutral gradient-based I-optimal design and D-optimal design result in 26 and 22 non-zero sensor weights, respectively, and we select the first 26 sensors of the risk-neutral greedy D-optimal design. Figure 7 compares the optimal designs across criteria and optimization approach employed. In Figure 7, we see the weights corresponding to each candidate sensor location, where the size and color intensity of the marker reflect the magnitude of the sensor weight. Sensors selected by the greedy algorithm are assigned uniform weights. From Figure 7, we see that across varying criteria and optimization approaches, the optimal designs produced are similar.

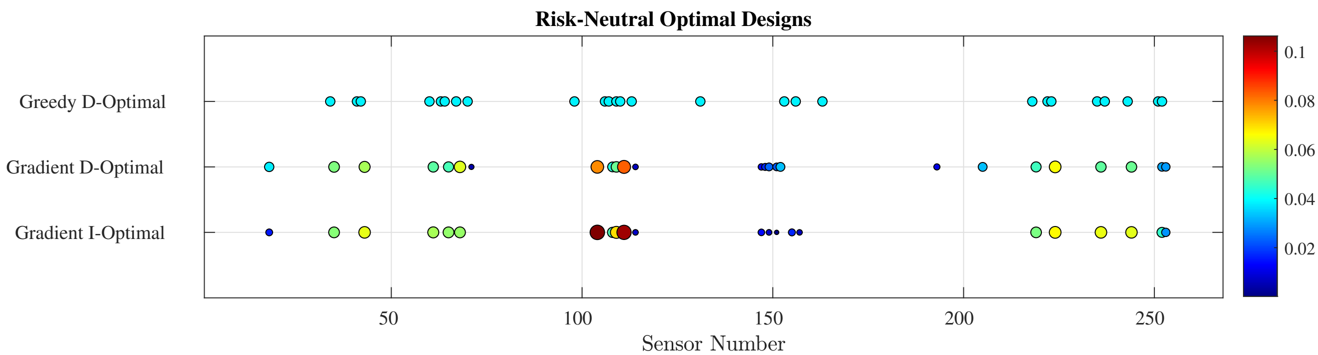

Fig. 7 Optimal sensor designs using D- and I-optimality using convex optimization versus D-optimality computed using a greedy algorithm. Color intensity and marker sizes correspond to the design weight. Non-zero designs of the greedy solution are assigned uniform weights.

Now we proceed to verifying that the prediction variance aligns with the established design criteria. Figure 8 presents the cumulative distribution function (CDF) of the prediction variance across model DoFs averaged over samples of the auxiliary parameters (*Model-Averaged Prediction Variance*) and normalized by the variance of the measurement noise. In this example, the model-averaged prediction variance is computed over the 50 auxiliary parameter samples used to solve the robust OED, and the prediction variance per model sample is sample-approximated using $1,000$ realizations of the estimated regression parameters $\tilde{\boldsymbol{\theta}}$. Each realization of $\tilde{\boldsymbol{\theta}}$ involves generating a sample of the noise-perturbed response at the experiment design under consideration and solving an inverse problem.

From the CDF curves in Figure 8, we observe that the prediction variance aligns with the established design criteria. Specifically, the gradient-based D-optimal design minimizes the maximum model-averaged prediction variance (CDF = 1), while the gradient-based I-optimal design minimizes the average model-averaged prediction variance. In general, there is a notable trade-off between these designs, where for example, minimal average prediction variance comes at the expense of a larger maximum prediction variance and vice-versa. Finally, we observe that the model-averaged prediction variances associated with the greedy D-optimal design tend to exceed those obtained from the gradient-based approaches.

Now, let's assume that the goal of the experiment design is to minimize the average prediction variance for any realization of the auxiliary parameters. To this end, we use the risk-neutral I-optimal OED formulation and compare its performance in terms of average prediction variance to the greedy and gradient-based D-optimal designs. For this study, we consider the average prediction variance given by (17) normalized with respect to the additive measurement noise variance σ^2. The sample-approximated prediction variances are generated using the same sampling method described previously, with the exception that we draw $1,000$ samples of the auxiliary parameters. The corresponding average prediction variances are (i) $4.48\sigma^2$ for gradient-based I-optimality, (ii) $4.62\sigma^2$ for gradient-based D-optimality, and (iii) $4.75\sigma^2$ for greedy-based D-optimality. Only a small improvement in average prediction variance using the I-optimal solution is observed.

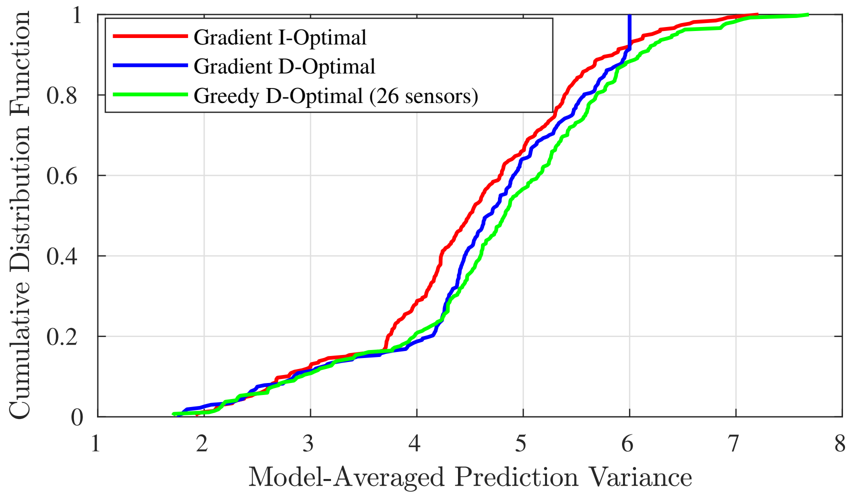

Fig. 8 Cumulative distribution function of the model-averaged prediction variance using the gradient-based D- and I-optimal designs and the greedy-based D-design.

Another trend can be observed if we investigate prediction variance averaged over the DoFs for each auxiliary parameter sample (*DoF-Averaged Prediction Variance*). Note that the DoF-averaged prediction variance is equivalent to the I-criterion. Figure 9 presents the CDF of the DoF-averaged prediction variance for the three design cases, using $1,000$ auxiliary parameter samples. Despite the similarity in average prediction variance between D-optimal designs and the I-optimal design, the D-optimal designs result in long tails in the DoF-averaged prediction variance distribution. In other words, for some realizations of the auxiliary parameters, the DoF-averaged prediction variance using the D-optimal design is significantly worse than the I-optimal design.

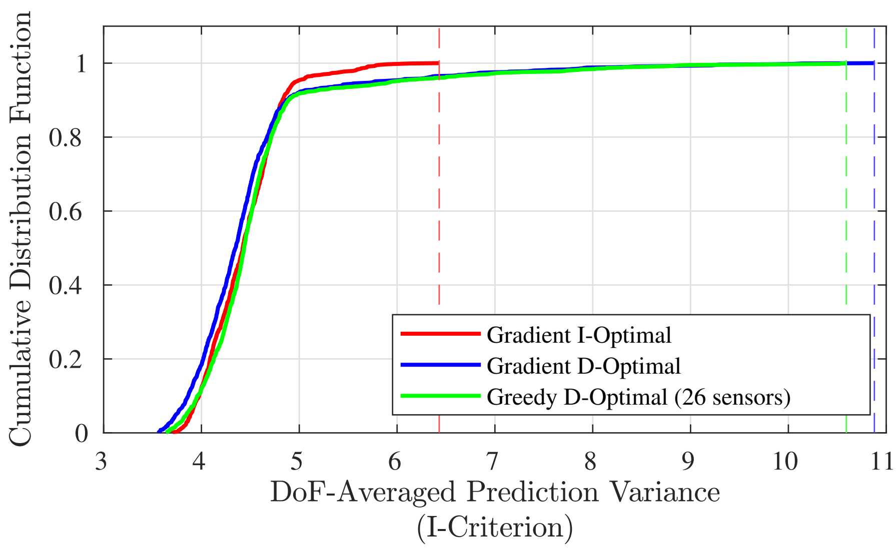

Fig. 9 Cumulative distribution function of the DoF-averaged prediction variance (I-criterion) using the gradient-based D- and I-optimal designs and the greedy-based D-optimal design. Dashed lines indicate maxima for additional clarity.

Finally, we investigate a heuristic for down-selecting optimally selected experiments in order to enforce a sensor budget. The presented gradient-based optimization approach does not naturally incorporate a sensor budget N, so instead we consider one heuristic, which sorts the selected sensors by design weight, selecting the first N largest weights, and then rounding all other weights to zero. In this method, we then scale the non-zero weights so that they sum to one, adhering to their interpretation as probability masses.

Let's now compare the designs (greedy D-optimal versus gradient-based D- and I-optimal) by observing the sample-approximated, DoF-averaged prediction variance (I-criterion) as a function of sensor budget from 6 sensors to 26 sensors. Figure 10 depicts the 25^{th} and 75^{th} quantiles (solid black lines), the median (dashed black line), and the average (solid red line) of the DoF-averaged prediction variance across $1,000$ realizations of $\boldsymbol{\xi}$ using the three design formulations. Note for the gradient-based D-optimal design, we terminated the study with 22 sensors because the optimization algorithm yielded only 22 non-zero weights. Upon inspection, it appears that the performance of the gradient-based I- and D-optimal designs is preserved down to about 19 and 17 sensors, respectively. The greedy-based D-optimal solution increases in prediction variance between 20 and 26 sensors, thus indicating that the prediction variance does not monotonically decrease with additional sensors. Generally speaking, the proposed heuristic appears to be a viable strategy for enforcing a sensor budget, though further investigation is needed.

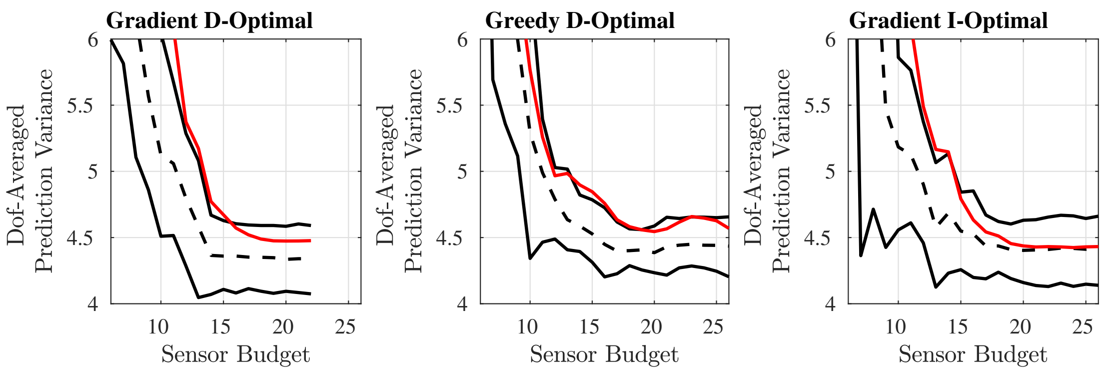

Fig. 10 The 25^{th} and 75^{th} quantiles (solid black line), the median (dashed black line), and the average (solid red line) of the DoF-averaged prediction variance (I-criterion) computed over all auxiliary parameter samples corresponding to the a) gradient D-optimal design, b) greedy D-optimal design, and c) gradient I-optimal design as a function of sensor budget.

Conclusions

This work considered robust OED strategies that account for uncertainty in auxiliary model parameters when estimating an optimal experimental design and applied the approach to a structural dynamics force inversion problem. Uncertainty in auxiliary parameters was characterized probabilistically. The OED objectives were then formulated using minimax and expectation-based (i.e. risk-neutral) approaches, which compute respectively, the worst-case and expected determinant (D-optimal) of the asymptotic covariance or prediction variance (I-optimal). We investigated the performance of the proposed methods on a force inversion problem using measured vibration response in the presence of material uncertainty. A reduced basis method was used to construct a reduced order model of the expensive high-fidelity simulation, enabling computational tractability of the robust OED scheme. We further compared two approaches for optimizing the OED objective: 1) a greedy approach to sequentially determine optimal sensors and 2) a gradient-based approach that assigns probability masses to candidate sensor locations. We found that accounting for uncertainty when designing experiments resulted in a significant improvement in designs when considering the minimax objective function; however, a negligible improvement was noted for risk-neutral objective formulations. We note that such a result is model dependent, and cannot be generalized beyond the context of the presented computational example. Furthermore, in comparing gradient-based and greedy optimization approaches, similar optimal designs resulted. Finally, we found that the risk-neutral I-optimal gradient-based approach was most effective for minimizing the average model prediction uncertainty. In future work, we plan to study the minimax gradient-based optimization formulation and further investigate methods that minimize prediction uncertainty.

Acknowledgments This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, 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-NA0003525.

References

1. Dette, H. and Trampisch, M. "Optimal Designs for Quantile Regression Models". *Journal of the American Statistical Association*, 107(499):1140–1151 (2012) Publisher: [American Statistical Association, Taylor & Francis, Ltd.].
2. Banks, H., Hu, S., and Thompson, W. *Modeling and Inverse Problems in the Presence of Uncertainty*. CRC Press (2014)
3. Kong, L. and Wiens, D.P. "Model-Robust Designs for Quantile Regression". *Journal of the American Statistical Association*, 110(509):233–245 (2015) Publisher: [American Statistical Association, Taylor & Francis, Ltd.].
4. Biedermann, S. and Dette, H. "A note on maximin and bayesian d-optimal designs in weighted polynomial regression". *Mathematical Methods of Statistics*, 12 (2003)
5. Rojas, C.R., Welsh, J.S., Goodwin, G.C., and Feuer, A. "Robust optimal experiment design for system identification". *Automatica*, 43(6):993–1008 (2007)
6. Pronzato, L. and Walter, E. "Robust experiment design via stochastic approximation". *Mathematical Biosciences*, 75(1):103–120 (1985)
7. Alexanderian, A., Petra, N., Stadler, G., and Sunseri, I. "Optimal design of large-scale Bayesian linear inverse problem under reducible model uncertainty: Good to know what you don't know". *SIAM/ASA Journal on Uncertainty Quantification*, 9(1):163–184 (2021)
8. Alexanderian, A. "Optimal Experimental Design for Infinite-dimensional Bayesian Inverse Problems Governed by PDEs: A Review". *Inverse Problems*, 37 (2021)
9. Kouri, D.P., Jakeman, J.D., and Gabriel Huerta, J. "Risk-Adapted Optimal Experimental Design". *SIAM/ASA Journal on Uncertainty Quantification*, 10(2):687–716 (2022) Publisher: Society for Industrial and Applied Mathematics.
10. Chun, S.Y., Shapiro, A., and Uryasev, S. "Conditional value-at-risk and average value-at-risk: Estimation and asymptotics". *Oper. Res.*, 60:739–756 (2012)
11. Jagalur-Mohan, J. and Marzouk, Y.M. "Batch greedy maximization of non-submodular functions: Guarantees and applications to experimental design". *J. Mach. Learn. Res.*, 22:252:1–252:62 (2021)
12. Chamon, L. and Ribeiro, A. "Approximate supermodularity bounds for experimental design" (2017)
13. Nemhauser, G.L., Wolsey, L.A., and Fisher, M.L. "An analysis of approximations for maximizing submodular set functions—I". *Mathematical Programming*, 14(1):265–294 (1978)
14. Silvey, S. *Optimal Design: An introduction to the theory for parameter estimation*. Chapman & Hall (1980)
15. Bunting, G., Crane, N.K., Day, D.M., Dohrmann, C.R., Ferri, B.A., Flicek, R.C., Hardesty, S., Lindsay, P., Miller, S.T., Munday, L., et al. "Sierra Structural Dynamics-Users notes 4.50.". Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States) (2018)
16. Zou, Z., Kouri, D., and Aquino, W. "An adaptive local reduced basis method for solving PDEs with uncertain inputs and evaluating risk". *Computer Methods in Applied Mechanics and Engineering*, 345:302–322 (2019)
17. Kiefer, J. and Wolfowitz, J. "The Equivalence of Two Extremum Problems". *Canadian Journal of Mathematics*, 12(600):363–366 (1960)