

Carl Ehrett¹

Watt Family Innovation Center,
Clemson University,
Clemson, SC 29631
e-mail: cehrett@clemson.edu

D. Andrew Brown

Associate Professor
School of Mathematical and Statistical Sciences,
Clemson University,
Clemson, SC 29631
e-mail: ab7@clemson.edu

Christopher Kitchens

Associate Professor
Department of Chemical and Biomolecular
Engineering,
Clemson University,
Clemson, SC 29631
e-mail: ckitch@clemson.edu

Xinyue Xu

Department of Architectural Engineering,
The Pennsylvania State University,
University Park, PA 16802
e-mail: xkx5055@psu.edu

Roland Platz

Department of Architectural Engineering,
The Pennsylvania State University,
University Park, PA 16802
e-mail: rpx5110@psu.edu

Sez Atamturktur

Harry and Arlene Schell Professor and Head of
the Department of Architectural Engineering,
The Pennsylvania State University,
University Park, PA 16802
e-mail: sez@psu.edu

Simultaneous Bayesian Calibration and Engineering Design With an Application to a Vibration Isolation System

Calibration of computer models and the use of those design models are two activities traditionally carried out separately. This paper generalizes existing Bayesian inverse analysis approaches for computer model calibration to present a methodology combining calibration and design in a unified Bayesian framework. This provides a computationally efficient means to undertake both tasks while quantifying all relevant sources of uncertainty. Specifically, compared with the traditional approach of design using parameter estimates from previously completed model calibration, this generalized framework inherently includes uncertainty from the calibration process in the design procedure. We demonstrate our approach to the design of a vibration isolation system. We also demonstrate how, when adaptive sampling of the phenomenon of interest is possible, the proposed framework may select new sampling locations using both available real observations and the computer model. This is especially useful when a misspecified model fails to reflect that the calibration parameter is functionally dependent upon the design inputs to be optimized. [DOI: 10.1115/1.4050075]

1 Introduction

This paper connects two distinct areas of research concerning computer models of real phenomena. One area is that of computer model calibration, where the goal is to find a posterior distribution of unknown or imperfectly known, parameters by calibrating a computer model using real-world observations of the modeled phenomenon. The second area is that of enlisting a computer model for design, using the model to find settings for controllable system inputs such that the resulting system output is optimized with respect to some design goal. These two problems are structurally similar, both involving finding estimates or distributions of model inputs to achieve some desired effect on model outputs. In the case of calibration, the desired effect is that the model outputs approximate reality, and in the case of design, the desired effect is that the model outputs approximate the optimal achievable outputs. Since calibration and design are typically carried out separately, existing design techniques operate under the assumption that the model is an accurate approximation of the real system of interest. In practice, models used for design typically are known

or suspected to be biased representations of the phenomenon of interest, and often have inputs that require calibration. The goal of the work described here is to provide a unified framework for calibration and design. We refer to this new approach as dual calibration to target outcomes (DCTO). In addition to avoiding the idealization that the model used for design is unbiased, DCTO allows one to focus calibration efforts on regions of interest, prioritizing them over other areas of the model range. For example, one may be more interested in calibrating the model to be accurate in the optimal region of some design variable θ_d than elsewhere. Having a combined framework for calibration and design is especially of interest when those two activities are nontrivially intertwined, as in the case when the value of the calibration parameters is functionally dependent upon the design settings.

Bayesian methods for computer model calibration are developed by Kennedy and O'Hagan [1]. Since their seminal paper, the methodology has seen numerous extensions and refinements [2–7]. Henceforth, we refer to this approach to calibration as “KOH”. Common to the KOH approach is the Bayesian framework in which one places a prior on the calibration parameters θ_c , often pairing it with a Gaussian process (GP) metamodel of the computer model of interest and a GP prior on the model discrepancy $\delta(\cdot)$, and using the available observations y_r of the real system to find a posterior distribution $\theta_c, \delta(\cdot) | y_r$. Such an approach is notable for providing not merely a point estimate of the

¹Corresponding author.

Manuscript received October 30, 2020; final manuscript received January 24, 2021; published online March 9, 2021. Assoc. Editor: Tao Xing.

calibration parameter, but for providing a full posterior distribution quantifying remaining uncertainty about θ_c and about $\delta(\cdot)$.

Herein, we leverage the KOH framework to find a posterior distribution, not only on unknown model parameters but also on controllable design settings. We achieve this via an approach called counterfactual Bayes. In traditional model calibration, one uses Bayes' rule to discover a posterior distribution of calibration parameters using real observations, so that the observations are the source of the Bayesian learning. In a design case, there are no relevant observations. One wants to find design settings that induce the system to behave optimally, but one typically has not observed the system doing so, and therefore there seems to be no relevant source of Bayesian learning that could drive the use of Bayes' rule to discover a posterior distribution of optimal design settings. The idea of counterfactual Bayes is to identify artificial observations, or target outcomes, y_t such that the resulting likelihood is highest in the optimal design region, i.e., target outcomes y_t such that their occurrence is strong evidence that the design settings are optimal. Hence, in addition to calibrating the unknown model parameters against experimental observations, one uses the KOH framework to also find a posterior distribution of design settings given the target outcomes. Given the nature of y_t , this is *defacto* a distribution of optimal design settings for the system. The result retains the benefits of the Bayesian model calibration tools on which it is based, namely, the quantification of remaining uncertainty regarding the optimal design settings. And like KOH, DCTO is especially well-suited to problems that rely on black-box functions.

The combined approach to calibration and design presented here thus contrasts with traditional approaches by providing for quantification of remaining uncertainty in the joint posterior distribution of the calibration and design inputs using an integrated framework incorporating both sets of inputs. The rest of the paper is organized as follows. Section 2 reviews related approaches to calibration and design. Section 3 describes the difficulties involved in extending KOH into a framework that incorporates both design and calibration, illustrating this by considering the failings of a naïve method for combining the two procedures, followed by a description of the proposed DCTO framework for extending KOH. Section 4 considers how DCTO may be useful in the case where sequential sampling is possible. In particular, sequential sampling with DCTO is attractive when the calibration parameter is known or suspected to be functionally dependent upon the design settings. We showcase the application of DCTO with sequential sampling using a synthetic example, comparing its results to that of a more traditional approach of design following calibration. In Sec. 5 we apply DCTO to a dynamic vibration system, using a set of experimental observations simultaneously to calibrate a finite element model and to select gain factor settings to achieve the optimal vibration isolation outcome while demonstrating DCTO's thorough quantification of the relevant uncertainties. Section 6 concludes with a discussion of the results and thoughts about future directions.²

2 Related Work

We may divide optimization approaches for design broadly into three camps [8]. Gradient-based approaches [9] are of limited utility when dealing with black-box functions, where we cannot

evaluate the objective function's derivative. Approximation of the derivative requires additional function evaluations, rapidly inflating the computational cost when each evaluation involves significant expense. Heuristic approaches [10] such as evolutionary algorithms [11–13], particle swarm optimization [14,15], and simulated annealing [16] avoid the need to know or approximate derivatives but often require prohibitively many function evaluations. Furthermore, such methods, like gradient-based approaches, do not inherently provide quantification of remaining uncertainty about optimal design settings and the system outputs at those settings. Methods exist for using heuristic approaches while accommodating and quantifying uncertainties [17,18], but these come at the cost of even further inflating the number of function evaluations required. This problem can be mitigated by relying on a surrogate model, but the resulting uncertainty quantification is accomplished by separate methods that are layered on top of the independent heuristic approach. By contrast, our combined approach to calibration and design includes uncertainty quantification as an intrinsic aspect of the framework.

The third camp is the diverse collection of response surface methodologies (RSMs) [19] used for optimization. RSMs operate by fitting a predictive model to an existing set of model runs, to form a computationally inexpensive metamodel which is then used to explore the model output. The concept of calibration to target outcomes that are built into DCTO is an example of an RSM, using GPs for its metamodel fit. Other popular versions of RSMs include efficient global optimization (EGO) [20,21] and stepwise uncertainty reduction (SUR) [22–28]. EGO and SUR are both designed to facilitate sequential sampling from the system of interest in a search for the global optimum. They differ in their *acquisition functions*, which determine the location of the next sampling location throughout the optimization process. EGO finds the spot that maximizes the expected improvement [20,29], whereas SURs acquisition function seeks to reduce the volume of excursion sets below the current best-known solutions [24]. Furthermore, the acquisition functions employed by EGO and SUR attempt to balance exploitation (proposing a new sample location that optimizes system output) with exploration (proposing a location that promotes learning for subsequent rounds of sampling). Because they rely on sequential sampling, EGO and SUR are of limited utility when one is constrained to rely on a pre-existing set of observations, or in general when the observation locations cannot be chosen purely to suit the goal of optimization. The value of exploration is in guiding future adaptive sampling to avoid settling on merely local optima. In cases where adaptive sampling is not possible, this value cannot be realized, and hence it is preferable to adopt a pure-exploitation result. As a result, although these acquisition functions constitute distributions of sampling locations, by their nature they are not interpretable as distributions of the *optimal* design settings for a given problem, and hence these distributions do not quantify uncertainty regarding the location of that optimum. By contrast, our combined approach (understood as a pure-exploitation method) quantifies the remaining uncertainty regarding the location of the system optimum.

An example of an RSM more closely resembling our approach to design is described by Olalotiti-Lawal and Datta-Gupta [30]. Their approach defines a distribution that is designed to lie both on and near the Pareto front of the objective function and generates a posterior distribution that includes quantified uncertainties via Markov Chain Monte Carlo (MCMC) [31]. However, the posterior distribution in that work is designed by the authors and is not dictated by the model itself; as such, its interpretability is not entirely clear. By contrast, our approach provides a posterior distribution based on the likelihood of the optimal design settings given the (hypothetical) observation of target outcomes y_t , and thus the uncertainty quantified by design using the KOH framework is model-driven and interpretable as uncertainty regarding the optimal values for the design inputs and the resulting system output.

²Note that KOH is usually conceived of as a means of calibrating a computer model with respect to a set of experimental observations. However, the KOH framework, and by extension DCTO, are applicable more generally whenever one has access to both low-fidelity and high-fidelity sources of information and seeks to calibrate the former with respect to the latter. This includes the case in which both the high-fidelity and low-fidelity sources of information are computer models (e.g., with different levels of computational expense). For ease of exposition, we follow the common convention and present DCTO in terms of calibrating a computer model using experimental observations. Nonetheless, in some cases (such as in our discussion of sequential sampling in Sec. 4), the methods discussed may apply more naturally in the context of employing two computer models of varying fidelity.

3 Dual Calibration to Target Outcomes

3.1 Separate Calibration and Design. The version of KOH considered here is that which finds a posterior distribution of a parameter of interest for calibration, θ , using a GP emulator with hyperparameters ϕ_η . Similarly, one may also use a GP prior with hyperparameters ϕ_δ to the model discrepancy between the computer model $\eta(\cdot)$ and the true function $f(\cdot)$ that it represents. In the work described here, we employ stationary GPs with a Gaussian kernel covariance structure $C(\mathbf{x}, \mathbf{x}') = 1/\lambda \times \exp(-\beta(\mathbf{x} - \mathbf{x}')^2)$, so that $\phi_\eta = [\beta, \lambda]$. In our adaptation, $\theta = (\theta_c, \theta_d)$ is partitioned into parameters θ_c to be calibrated and inputs θ_d to be optimized for design purposes. Setting priors on θ and on ϕ_δ , we train the GP emulator on observations η and use MCMC to explore the distribution

$$\pi(\theta, \phi_\eta, \phi_\delta | \mathcal{D}) \propto \pi(\mathcal{D} | \theta, \phi_\eta, \phi_\delta) \times \pi(\theta) \times \pi(\phi_\eta) \times \pi(\phi_\delta) \quad (1)$$

where $\mathcal{D} = (\eta^T, \mathbf{y}^T)^T$ for some observations $\mathbf{y}$.

In a computer calibration problem, $\mathbf{y}$ is a set of observations of the system modeled by $\eta(\cdot)$. When calibrating to target outcomes as in DCTO, by contrast, $\mathbf{y}$ is a set of target outcomes representing the way that one wishes to induce the system to behave (rather than observations one has made of the system in reality). When one wishes to perform design leveraging a simulation model that also requires traditional calibration, then, one might consider combining the two approaches by using Eq. (1) with $\mathbf{y} = (\mathbf{y}_r^T, \mathbf{y}_t^T)^T$, an array containing both real observations $\mathbf{y}_r$ (for calibration) and target outcomes $\mathbf{y}_t$ (for design). However, this approach will not work for two reasons. First, though the matter of whether a given input is considered to be a calibration parameter is dependent upon the nature of the research problem under investigation, typically the inputs to be calibrated are not design settings under researcher control. Second, for successful calibration, one must train one's model on observations of reality rather than on unobserved target outcomes.

Hence, model calibration and system design must be separated. An obvious choice here is to perform KOH calibration first, without involving any target outcomes, and then to use the calibrated model for model-assisted design. Under this approach, with observations $\mathbf{y}_r$ of the system of interest, one would employ the model described in Eq. (1) with $\theta = \theta_c$ (the parameters to be calibrated) and with $\mathcal{D} = \mathcal{D}_c = (\eta^T, \mathbf{y}_r^T)^T$. The result would be a posterior distribution of θ_c and of $\delta(\cdot)$, the systematic discrepancy between the computer model $\eta(\cdot, \cdot)$ and the true system $f(\cdot)$. These can be used to produce estimates $\hat{\theta}_c$ and $\hat{\delta}(\cdot)$ such that $f(\mathbf{z}) \approx \eta(\mathbf{z}, \hat{\theta}_c) + \hat{\delta}(\mathbf{z})$ for all $\mathbf{z}$ in the domain of f . The result is a calibrated model $\eta_c(\mathbf{z}) = \eta(\mathbf{z}, \hat{\theta}_c) + \hat{\delta}(\mathbf{z})$ which can be used for design.

With η_c in hand, one can partition $\mathbf{z}$ into $(\mathbf{x}, \theta_d)$ where θ_d is the set of inputs over which one wishes to optimize, and $\mathbf{x}$ are all other inputs in the operational domain, within which the calibrated model's predictions are reliable. We can write $\eta_c(\mathbf{z})$ as $\eta_c(\mathbf{x}, \theta_d)$. Then one can perform design again using Eq. (1), this time with $\theta = \theta_d$ and $\mathcal{D} = \mathcal{D}_t = (\eta_c^T, \mathbf{y}_t^T)^T$ where $\eta_c = \eta + \hat{\delta} = \eta + (\hat{\delta}(\mathbf{z}_1), \dots, \hat{\delta}(\mathbf{z}_n))^T$. Notice that a single set of simulator runs η can be used both for KOH and for subsequent design. A crucial difference between calibration and design is that for the design step one would not attempt to model any systematic discrepancy between η_c and f , since an estimate of that discrepancy is already included in η_c . For the purposes of Eq. (1), this amounts to setting a degenerate prior on ϕ_δ that is a point mass at 0.

A problem with the above-described approach of performing calibration before a separate design optimization is that relying on static calibration estimates $\hat{\theta}_c$ ignores uncertainty remaining after calibration with respect to the true value of θ_c (if one takes a Bayesian view, one may prefer to consider calibration as finding optimal settings for inducing a model to approximate reality, rather than as finding true values of some parameter; for ease of exposition, we will continue to use the latter phrasing). In order to

produce results that take into account all sources of uncertainty, it is necessary to integrate calibration and design, so that the uncertainty remaining from calibration is propagated through the design process. This can be accomplished either asynchronously (so that the posterior distribution of θ_c is sampled while undertaking design) or, for lower computational overhead, synchronously (so that a single MCMC run is used to perform both calibration and design). In either case, it will be useful to produce an integrated model for the combined tasks of calibration and design which describes the use of both procedures, and which makes clear the relationship between them. This integrated model will also serve to demonstrate the unified framework underlying the combined approach.

3.2 Integrated model. Consider η as having three inputs $(\mathbf{x}, \mathbf{t}_c, \mathbf{t}_d)$ where $\mathbf{t}_c$ denotes the parameters targeted for KOH calibration, $\mathbf{t}_d$ denotes the input settings targeted for design, and $\mathbf{x}$ denotes the remaining controllable inputs. If η can be run quickly, then we use it directly in MCMC. However, if it is computationally expensive, we employ a surrogate by setting a GP prior on η with mean $m_\eta(\mathbf{x}, \mathbf{t}_c, \mathbf{t}_d)$ and covariance function $C_\eta((\mathbf{x}, \mathbf{t}_c, \mathbf{t}_d), (\mathbf{x}', \mathbf{t}_c', \mathbf{t}_d'))$. From here on in this discussion, assume that a GP surrogate is used for η . We model the systematic discrepancy between η and f at the true value of $\mathbf{t}_c = \theta_c$ with another GP prior $\delta(\cdot, \cdot)$ having mean $m_\delta(\mathbf{x}, \mathbf{t}_d)$ and covariance function $C_\delta((\mathbf{x}, \mathbf{t}_d), (\mathbf{x}', \mathbf{t}_d'))$. In addition to systematic discrepancy between η and reality, measurement error ε_r may be included in the model for real observations $\mathbf{y}_r$, and additional Gaussian observation error ε_d may be included for target outcomes $\mathbf{y}_t$.

The purpose of additional observation error ε_d is twofold. Depending on the distribution of ε_c , the target outcomes $\mathbf{y}_t$ may or may not be possible outputs of a model that lacks ε_d . Including ε_d ensures that there is a nonzero probability of an observation falling in the vicinity of the targets. Second, including ε_d and estimating its variance σ_d^2 provides computational benefits. For example, even if the target outcomes are compatible with a model that does not include ε_d , they may (depending on the choice of targets) be extreme outliers to the extent that the relevant likelihoods are small enough to generate significant numerical errors during MCMC. In terms of the interpretation of the model, adding ε_d amounts to supposing that the counterfactual target outcomes were observed with greater than usual observation error, where that additional error is distributed as $N(0, \sigma_d^2)$. Though it is not necessary to assume that ε_c is Gaussian, for simplicity of presentation we assume here that it is distributed as $N(0, \sigma_c^2)$. Finally, we assume that η , δ , ε_c , and ε_d are all mutually independent.

A collection of simulation runs is needed to train the GP code surrogate. Let $(\mathbf{x}_s, \mathbf{t}_{cs}, \mathbf{t}_{ds})$ be the design matrix for the settings of the simulation runs, and let $\mathbf{y}_s$ denote the output of these runs. Similarly, let $\mathbf{y}_r$ be observations made at $\mathbf{x}_r, \mathbf{t}_{dr}$, and let $\mathbf{y}_t$ be target outcomes we wish to observe at $\mathbf{x}_t$. Finally, let $\mathbf{y} = (\mathbf{y}_s^T, \mathbf{y}_r^T, \mathbf{y}_t^T)^T$, and $\mathbf{1}$ a vector of ones. Then it follows that $\mathbf{y} \sim N(\mathbf{m}, \mathbf{C})$, where:

$$\mathbf{m} = \begin{pmatrix} m_s(\mathbf{x}_s, \mathbf{t}_{cs}, \mathbf{t}_{ds}) \\ m_s(\mathbf{x}_r, 1\theta_c^T, \mathbf{t}_{dr}) + m_\delta(\mathbf{x}_r, \mathbf{t}_{dr}) \\ m_s(\mathbf{x}_t, 1\theta_c^T, 1\theta_d^T) + m_\delta(\mathbf{x}_t, 1\theta_d^T) \end{pmatrix} \quad (2)$$

$$\mathbf{C} = \begin{pmatrix} C_{11} & C_{12} & C_{13} \\ C_{21} & C_{22} & C_{23} \\ C_{31} & C_{32} & C_{33} \end{pmatrix} \quad (3)$$

$$C_{11} = C_\eta((\mathbf{x}_s, \mathbf{t}_{cs}, \mathbf{t}_{ds}), (\mathbf{x}_s, \mathbf{t}_{cs}, \mathbf{t}_{ds})) \quad (4)$$

$$C_{21} = C_\eta((\mathbf{x}_s, \mathbf{t}_{cs}, \mathbf{t}_{ds}), (\mathbf{x}_r, 1\theta_c^T, \mathbf{t}_{dr})) \quad (5)$$

$$C_{31} = C_\eta((\mathbf{x}_s, \mathbf{t}_{cs}, \mathbf{t}_{ds}), (\mathbf{x}_t, 1\theta_c^T, 1\theta_d^T)) \quad (6)$$

$$\mathbf{C}_{12} = \mathbf{C}_{21}^T \quad (7)$$

$$\begin{aligned} \mathbf{C}_{22} = & C_\eta((\mathbf{x}_r, 1\theta_c^T, \mathbf{t}_{dr}), (\mathbf{x}_r, 1\theta_c^T, \mathbf{t}_{dr})) \\ & + C_\delta((\mathbf{x}_r, \mathbf{t}_{dr}), (\mathbf{x}_r, \mathbf{t}_{dr})) + \sigma_c^2 \mathbf{I} \end{aligned} \quad (8)$$

$$\mathbf{C}_{32} = C_\eta((\mathbf{x}_r, 1\theta_c^T, \mathbf{t}_{dr}), (\mathbf{x}_t, 1\theta_c^T, 1\theta_d^T)) + C_\delta((\mathbf{x}_r, \mathbf{t}_{dr}), (\mathbf{x}_t, 1\theta_d^T)) \quad (9)$$

$$\mathbf{C}_{13} = \mathbf{C}_{31}^T \quad (10)$$

$$\mathbf{C}_{23} = \mathbf{C}_{32}^T \quad (11)$$

$$\begin{aligned} \mathbf{C}_{33} = & C_\eta((\mathbf{x}_t, 1\theta_c^T, 1\theta_d^T), (\mathbf{x}_t, 1\theta_c^T, 1\theta_d^T)) \\ & + C_\delta((\mathbf{x}_t, 1\theta_d^T), (\mathbf{x}_t, 1\theta_d^T)) + \sigma_c^2 \mathbf{I} + \sigma_d^2 \mathbf{I} \end{aligned} \quad (12)$$

Note that when $\mathbf{y}_t$ and $\mathbf{x}_t$ are empty and $\mathbf{m}$, $\mathbf{C}$ reduce, respectively, to their first two and upper two-by-two block elements, this is simply the KOH framework. Thus, DCTO is an extension of the KOH framework to include design using target outcomes.

A primary benefit of DCTO is that the design process includes quantification of all sources of uncertainty. Performing calibration and then subsequently undertaking design using static estimates for $\hat{\theta}_c$ and $\hat{\delta}$ does not properly account for the uncertainty surrounding the estimates. Another benefit of the combined approach appears in cases in which the model is misspecified in failing to account for functional dependence of θ_c on θ_d . In such cases, one may be interested only or primarily in the value of θ_c at the optimal value of θ_d . If one has the freedom to sample adaptively from the true system, then this freedom can be applied in DCTO to concentrate samples disproportionately in the region of interest. This idea is explored further in Sec. 4.

For DCTO, we employ modularity in the manner of Ref. [32]. A modular analysis intentionally falls short of being a full Bayesian analysis, either for computational benefits, or to quarantine “suspect” aspects of the model, so that the posterior distributions of parameters of interest are robust to model misspecification. The target outcomes $\mathbf{y}_t$ are precisely such a suspect source of Bayesian learning—they are by their nature extreme outliers, and hence are a poor guide both for estimating the hyperparameters of the GP emulator and for estimating the parameter θ_c . To modularize DCTO, we estimate the emulator hyperparameters via maximum likelihood, and we refrain from including $\mathbf{y}_t$ in the updates of θ_c during MCMC. That is, rather than calculating the likelihood of a proposed sample $t_c^{(i+1)}$ at step i of the MCMC using $\mathbf{y} = (\mathbf{y}_s^T, \mathbf{y}_r^T, \mathbf{y}_t^T)^T$, we instead calculate its likelihood using only $\mathbf{y} = (\mathbf{y}_s^T, \mathbf{y}_r^T) \sim N(\mathbf{m}_r, \mathbf{C}_r)$, where $\mathbf{m}_r$ and $\mathbf{C}_r$ are respectively the upper two and upper-left two-by-two components of $\mathbf{m}$ and $\mathbf{C}$. Such modularization ensures that all Bayesian learning of θ_c is based upon the real observations rather than upon $\mathbf{y}_t$.

4 Dependence of θ_c on θ_d

4.1 Background. In many cases of computer model calibration, it is known or suspected that the value of one or more calibration parameters is functionally dependent upon the values of other model inputs [33–35]. If one is interested to understand the functional form of the calibration parameter, then state-aware methods can be used to arrive at such an estimate [33,34,36].

In a case where the calibration parameter is functionally dependent upon the design settings, one might be interested only to know the value of the calibration parameter in the optimal design region. When calibration and design are undertaken simultaneously, as in DCTO, the machinery of state-aware calibration is not needed, and effort is better spent focusing on estimating the fixed calibration parameter value in the region of interest. In such a case, calibration should be founded on observations for which

the design settings are in the optimal design region. This will allow one to calibrate the model using observations taken from the region of design interest, so that the calibration takes on values that are most applicable in that region.

When observations may be made adaptively, other RSM approaches such as EGO [20,21] or SUR [22–28] may be more efficient than the KOH framework for estimating optimal design settings, though the KOH framework offers more interpretable and model-driven uncertainty quantification. Further, RSM approaches, in general, do not include tools to accommodate the case in which a model stands in need of calibration as well as optimization. DCTO provides such a framework for combined calibration and design.

Algorithm 1 DCTO with adaptive sampling

- 1 Set $\mathbf{y} = [\mathbf{y}_r^T, \mathbf{y}_t^T]^T$ where $\mathbf{y}_t$ are the target outcomes and $\mathbf{y}_r = []$ is an empty array.
- 2 Begin MCMC burn-in. Set $i = 1$. Let m be the budget of function evaluations. While $i \leq m$:
 - 2.1 Complete n iterations of MCMC burn-in (where e.g., $n = 100$).
 - 2.2 Draw $\hat{\theta}_d$ from the available size $n \cdot i$ sample of $t_d|\mathbf{y}$.
 - 2.3 Evaluate $f(\mathbf{x}_i, \hat{\theta}_d)$.
 - 2.4 Set $\mathbf{y}_r = [\mathbf{y}_r^T, f(\mathbf{x}_i, \hat{\theta}_d)]^T$.
- 3 Continue burn-in until convergence.
- 4 Draw a sample of the desired size from the posterior distributions of θ_c, θ_d .

Therefore, we now consider under the lens of DCTO the case in which the design settings of the observations of the true system may be chosen adaptively. The use of DCTO with adaptive sampling is potentially of greatest use when it is known or suspected that the calibration parameter is a function $\theta_c(t_d)$ of the design setting t_d , and particularly when interest focuses on learning the optimal design setting θ_d and the corresponding value $\theta_c(\theta_d)$ of the calibration parameter. The process of performing DCTO with adaptive sampling is described in Algorithm 1. When adaptively evaluating the objective function, the locations of the input settings $\mathbf{x}_i$ which are not being optimized for design can be selected to maximize distance from previous observations, or these locations can be predetermined according to a space-filling design over the domain of nondesign inputs. The result of applying this algorithm is that observations are concentrated around the design settings of interest so that the unknown calibration parameter values in those observations are concentrated around the value $\theta_c(\theta_d)$.

4.2 Simulated Example. To demonstrate the use of DCTO with adaptive sampling in a case of functional dependence of the calibration parameter on design settings, we apply the method to a simulated system described by the function of three inputs

$$f_0(x, t_c, t_d) = x/(t_d^{t_c-1} \exp(-0.75t_d) + 1) \quad (13)$$

Figure 1 shows the output of this function for $x = 1$ over the range $(t_c, t_d) \in [1.5, 4.5] \times [0, 5]$. For any value of x and t_c , the optimal (minimizing) value of t_d is $(4/3)(t_c - 1)$. Suppose that the calibration parameter’s “true” value is functionally dependent on the design input, with the relationship

$$\theta_c(t_d) = 2.25 - .75 \frac{\exp\left(40\left(\frac{t_d - 1.5}{.75} - .5\right)\right)}{1 + \exp\left(40\left(\frac{t_d - 1.5}{.75} - .5\right)\right)} \quad (14)$$

which would be unknown in a real application. Figure 2 shows this relationship. Figure 3 shows the locations of the true value of θ_c and the optimal value of θ_d with respect to the function $f(x, \theta_c(t_d), t_d)$ of the design input t_d (where that optimum is




Fig. 1 Example computer model output over the support of the calibration parameter t_c and the design parameter t_d

independent of the value of x). In Fig. 3 it is clear that the true value of θ_c is far from optimal for the purpose of minimizing the objective function output, in the sense that if this value were within our control (which, being a calibration parameter, it is not), we would prefer to place it at the upper end of its support, at 4.5. Thus η showcases the ability of DCTO to perform simultaneously both calibration and design in the case when our “truth-seeking” goals and our design goals are in tension.

We apply DCTO to four versions of the problem. First, we assume that η is free from discrepancy, i.e., that $\eta(x, \theta_c, t_d)$ is an unbiased estimator of the “true” system $f_0(x, t_d)$, described by Eq. (13) above. The other three versions each assume that η suffers from some form of discrepancy. Let f_1, f_2, f_3 denote the “true” systems in these three cases. We set

$$f_1(x, t_d) = \eta(x, \theta_c, t_d)(1 - a(x - .5)(x - 1)/x) \quad (15)$$

$$f_2(x, t_d) = \eta(x, \theta_c, t_d) - a(x - .5)(x - 1)\left(t_d - \frac{4}{3}\right)^2 + b \quad (16)$$

$$f_3(x, t_d) = \eta(x, \theta_c, t_d) + ax t_d + b \quad (17)$$




Fig. 2 True value of the calibration parameter θ_c for each value in the domain of t_d

where a, b are constants that determine how severe the discrepancy is in each case. The function f_1 has a multiplicative discrepancy dependent only on x and a . This discrepancy does not affect the optimal value of t_d . The discrepancies of f_2 and f_3 are both additives. Figure 4 shows the discrepancies for two different versions (corresponding to different settings of (a, b)) of each f_i .

We apply DCTO with and without adaptive sampling to each of seven cases, without using an emulator: the nondiscrepancy case, and the two different versions of each f_i shown in Fig. 4. In each case, we gather 20 “observations” of f_i on a Latin hypercube design over the supports of x and t_d , setting θ_c equal to its “true” value of $\theta_c(t_d)$. After standardizing the response to have mean 0 and standard deviation 1, we add independent and identically distributed $N(0, 0.05)$ noise to the response. An example of the resulting “observations” from nonadaptive DCTO, with noise, appears in Fig. 5. We carry out DCTO using Metropolis-Hastings-within-Gibbs MCMC, drawing 8000 realizations each (discarding the first 4000 as burn-in) of $t_c, t_d, \rho_\delta, \lambda_\delta, \sigma_\delta^2$, where $\phi_\delta = (\rho_\delta^T, \lambda_\delta)^T$. For the adaptive sampling application of DCTO, we begin the MCMC with 0 observations of f_i , making a new observation after every 100 steps of MCMC until we reached the total budget of 20. An example of the resulting difference between the adaptive sampling approach and relying on a space-filling design, with regard to the sampling distribution of our observations of the objective function, appears in Fig. 6. There, one can see that the adaptive sampling approach manages to expend its budget on observations that are near to the region of design interest. This explains the superior performance of adaptive sampling (discussed below) in both design and in calibration (since the value of the calibration parameter is dependent upon that of the design input). This ameliorative effect would likely be even greater in a higher-dimensional case, in which a space-filling design would (due to




Fig. 3 The top plot shows the computer model output at $x = 1$ and the optimal design setting for each value of the calibration parameter t_c . The bottom plot shows the model output at $x = 1, t_c = \theta_c(t_d)$ for each value of the design parameter t_d .

the curse of dimensionality) tend to generate observations even farther from the region of design interest.

In both versions of DCTO, we modularize the analysis by drawing each of $\theta_c, \rho_\delta, \lambda_\delta$ using the likelihood-based only on $(\mathbf{y}_s^T, \mathbf{y}_r^T)^T$ rather than on $(\mathbf{y}_s^T, \mathbf{y}_r^T, \mathbf{y}_t^T)^T$. Convergence was verified visually and by the Gelman-Rubin statistic (≈ 1.01 ; Ref. [37]).

The resulting optimal design settings and a calibration parameter value at the optimum vary in the discrepancy cases, though $\theta_c(\theta_d)$ is near 2.16 in each case. Representative results from performing DCTO with adaptive sampling in each discrepancy case appear in Fig. 7, along with results from applying DCTO non-adaptively (using a space-filling set of observations). A summary




Fig. 4 The i th row shows f_i (the objective function with discrepancy), η (the computer model, identical in each panel), and the discrepancy $f_i - \eta$, all at $x = 0.75$. In each row, a less aggressive version of the discrepancy appears on the left, and a more aggressive on the right. In each plot, the topmost surface is f_i , the middle surface is η , and the bottom surface is the discrepancy $f_i - \eta$.




Fig. 5 Noisy observations of the system and the true system mean, for $f = f_0$ (no discrepancy)

of the results of thirty applications of DCTO both with and without adaptive sampling, for each of the discrepancy cases, appears in Table 1.

The results show superior performance for the adaptive sampling DCTO over DCTO using a space-filling design of experiments for the true phenomenon (or high-fidelity model, in a case of calibrating a low-fidelity model to use for design purposes). The adaptive DCTO posterior means have lower root-mean-square error (RMSEs) in all cases for θ_d , and in all cases except one for θ_c . Though both models suffer from the misspecification of treating as constant a calibration parameter that is more properly understood as functionally dependent upon other model inputs (particularly, e.g., in the case of f_3), the adaptive form of DCTO is consistently more robust to these difficulties. By using the estimate $\hat{\theta}_d$ to sample from the region of interest, DCTO learns from observations such that $\theta_c(\hat{\theta}_d)$ is near to the value $\theta_c(\theta_d)$. This promotes better calibration with respect to the region of interest, and thereby better estimation of the optimal design settings. By relying on DCTO rather than on performing KOH using samples gathered using heuristic optimization methods, or other RSM approaches, we achieve these estimates with quantification of all relevant model-driven uncertainty with respect to the values of θ_c and θ_d .

5 Case study Application: Vibration Isolation Design

The application of DCTO methodology is demonstrated on a vibration isolation design problem. Vibration isolation relies on the balance of inertia, damping, and stiffness properties where, in active vibration isolation, an additional active gain factor enhances the system's damping behavior. To achieve the optimal vibration isolation outcome, the design engineer typically specifies the resonance and isolation frequencies and then balances mass, damping, stiffness, and the gain factor.

5.1 Case Study Problem. The experimental dynamical system studied herein is a one-mass oscillator subjected to passive and active vibration isolation. The system consists of a rigid rectangle frame, a rigid mass held by four identical orthogonally placed leaf springs mounted to the frame, and a voice coil actuator (VCA) for passive and active damping.

In Fig. 8(a) and as a simplified model, a rigid mass m oscillates in the z -direction due to a base point excitation $w(t)$. A damper with the damping coefficient b and a spring element with a stiffness constant k connects the mass to the base point. The damper

and spring provide the system's internal passive damping force, active damping force, and stiffness force $F_b = b[\dot{z}(t) - \dot{w}(t)]$, $F_a = -g\dot{z}(t)$, $F_k = k[z(t) - w(t)]$ with F_a derived from a simple velocity feedback control with the gain factor g .

The inhomogeneous differential equation of the one-mass oscillator's motion in Fig. 8(a) can be written as

$$\ddot{z}(t) + \left[2D_p\omega_0 + \frac{g}{m} \right] \dot{z}(t) + \omega_0^2 z(t) = 2D_p\omega_0 \dot{w}(t) + \omega_0^2 w(t) = \omega_0^2 r(t) \quad (18)$$

using the abbreviation $2D_p\omega_0 = \frac{b}{m}$, and $\omega_0^2 = \frac{k}{m}$ including the damping ratio D_p from passive damping, with $0 < D_p < 1$, and the angular eigenfrequency ω_0 . The term $\omega_0^2 r(t)$ in EQ. (18) is the excitation function, which, in this case, is the linear combination of the damper and spring base point excitation $2D_p\omega_0 \dot{w}(t) + \omega_0^2 w(t)$.

Figure 8(b) depicts the laboratory setup used in this study, in which a rigid frame with mass m_f serves as a base point structure. The frame is fixed by a gliding support assumed to have no friction perpendicular to the z -direction. The frame is constrained by a damper with the damping coefficient b_f and springs with a total stiffness k_f in the z -direction and in the same plane.

In the laboratory application, the frame suspends from a rigid mount via elastic straps vertical to the z -direction, allowing the frame to move freely in the z -direction. The idealized damping b_f and k_f that constrain this movement are relatively small, compared to the b and k of the mass. The frame moves in a translational z -direction because of a time-dependent translational excitation displacement $w(t)$ in the z -direction. As shown in Fig. 8(c), the frame retains two supports that fix a leaf spring at its ends at A and C, with the effective bending length l on sides A-B and B-C, with the rigid mass m in the center position at B. The leaf spring is the practical realization of the spring elements in Figs. 8(a) and 8(b). Its stiffness $k^* = 12EI/l^3$ is a function of the bending stiffness EI , where E is Young's modulus of the leaf spring made from carbon fiber reinforced polymer (CFRP), I is the geometrical moment of inertia, and l is the length of the leaf spring. Two leaf springs are mounted in parallel with length l on each side of A-B and B-C (see Figs. 8(c)). With four leaf springs, the total stiffness becomes $k = 4k^*$. The two supports at A and C in Fig. 8(c) are adjustable




Fig. 6 Design input values for observations made under the adaptive sampling (AS) approach and under a SFD, along with the optimal value θ_d visible as the system global minimum in Fig. 3

along l to tune the leaf spring's bending deflection and therefore its effective stiffness k .

A VCA realizes an electromotive force F_{VCA} as the passive damping and the active force F_b and F_a (Fig. 8(c)). The force

sensor $S_{F_{VCA}}$ at B in Fig. 8(c) measures the sum of forces F_b and F_a acting on the moving mass m . The acceleration sensors $S_{a,z}$ and $S_{a,w}$ measure directly the accelerations of mass and frame, $\ddot{z}$ and $\ddot{w}$. The accelerations are transformed into velocities $\dot{w}$ and $\dot{z}$ by




Fig. 7 Prior and posterior distributions of the calibration parameter θ_c and design parameter θ_d , along with their true/optimal values, for DCTO with AS and with predetermined SFD in each of the cases studied. For each f_b , a and b control the size of the discrepancy as specified in Eqs. (15)–(17).

numerical integration in the SIMULINK-dSpace environment. The masses of $S_{a,z}$, $S_{F_{VCA}}$, and parts of the leaf spring are included in mass m (Table 2).

Figure 8(c) also shows a modal hammer with a force sensor S_F to excite the frame. The hammer creates the impulse force

$$\hat{F}(t_0) = \int_{-\infty}^{\infty} F(t) \cdot \delta(t - t_0) dt \quad (19)$$

including the Dirac-impulse function $\delta(t - t_0)$ that leads to the vibrational response of the frame

$$w(t) = \frac{\hat{F}(t_0)}{m_f \omega_{Df}} \cdot e^{-D_f \omega_{0f} t} \sin \omega_{Df} t \quad (20)$$

Table 1 Posterior root-mean-square error (RMSE) for the calibration variable θ_c and the design variable θ_d for DCTO with AS and a predetermined SFD. The estimator θ_i is the posterior mean of t_i for $i=c, d$. For each t_i , a and b control the size of the discrepancy as specified in Eqs. (15)–(17).

	$\hat{\theta}_c$ RMSE		$\hat{\theta}_d$ RMSE	
Objective	AS	SFD	AS	SFD
f_0 (no discrepancy)	0.188	0.433	0.163	0.479
f_1 , $a = 1.5$	0.233	0.32	0.243	0.414
f_1 , $a = 3.5$	0.188	0.247	0.213	0.393
f_2 , $a = 0.15$, $b = 0.075$	0.221	0.263	0.187	0.348
f_2 , $a = 0.65$, $b = 0.075$	0.228	0.16	0.183	0.206
f_3 , $a = 0.055$, $b = 0$	0.452	0.506	0.182	0.329
f_3 , $a = 0.055$, $b = 0.1$	0.448	0.468	0.167	0.292

in the time domain, with damping ratio D_f , angular eigenfrequency ω_{0f} and damped angular eigenfrequency ω_{Df} of the frame's movement in z -direction. Equation (20) is only valid for low damping $0 < D_f < 1$. This leads to the total vibration response $z(t) = r_0 \left\{ 1 - e^{-D \omega_0 t} [\cos \omega_D t - D \frac{\omega_0}{\omega_D} \sin \omega_D t] \right\}$.

The particular solution r_0 is part of the general excitation function $\omega_0^2 r(t)$ in Eq. (18), which takes the form of an excitation step function $r(t) = r_0 \sigma(t - t_0)$ when multiplied with the unit step function $\sigma(t - t_0)$ as the integral of the Dirac-impulse function $\delta(t - t_0)$ in Eq. (19). From the relation $2D_p \omega_0 \dot{w}(t) + \omega_0^2 w(t) = \omega_0^2 r(t)$ in Eq. (18), it follows that $r_0 = \frac{1}{\omega_0} 2D_p \dot{w}_0 + w_0$ with the velocity $\dot{w}_0$ and displacement w_0 at $t = t_0$ that is derived from Eq. (20).




Fig. 8 Schematic diagram of the test rig for the dynamic vibration system: (a) simplified schematic representation of the one-mass oscillator, (b) one-mass oscillator with an additional frame as the base point, and (c) schematic representation of the real test setup

Table 2 Geometrical, mass, and material values of each component in the vibration isolation test rig, including minimum (min) and maximum (max) of variable properties of the test rig

Category	Property	Variable	Value	Unit
Rigid frame structure	Sum mass	m_f	6.2073	kg
Vibrating rigid mass	Sum mass, min	m	0.7853	kg
	20× add. weights, small	m_{ws}	0.0760	kg
	24× add. weights, large	m_{wl}	0.2880	kg
	Sum mass, max	m	1.1493	kg
Geometry	Leaf spring length, min	l	0.04	m
	Leaf spring length, max	l	0.08	m
	Leaf spring cross section, width	d	0.04	m
	Leaf spring cross section, height	h	0.11	m
Material	Elastic modulus	E	6.2×10^9	N/m ²
	Stiffness CFRP, min	k	25788.1	N/m
	Stiffness CFRP, max	k	206305.0	N/m
VCA	Passive damping coefficient, min	b	16	Ns/m
	Passive damping coefficient, max	b	130	Ns/m
	Passive damping ratio, min	D_p	0.0481	—
	Passive damping ratio, max	D_p	0.628	—
	Active gain factor, min	g	0	Ns/m
	Active gain factor, max	g	95	Ns/m

Table 3 A variety of experiment tests and five-times averaged results

Case	Control parameter x	Design parameter θ_d	Five-times avg. system response y
Variable unit	Mass (kg)	Gain factor (Ns/m)	Overall damping ratio*
1	1.1493	0	0.0523
2	0.9653	0	0.0481
3	0.7853	0	0.0549
4	1.1493	8	0.0798
5	0.9653	8	0.0864
6	0.7853	8	0.0871
7	1.1493	41	0.308
8	0.9653	41	0.264
9	0.7853	41	0.259
10	1.1493	95	0.542
11	0.9653	95	0.527
12	0.7853	95	0.628




Fig. 9 Schematic diagram of the applied impulse force in the time domain (left) and five-times averaged acceleration response of the rigid mass in the time domain (right)

In this demonstration, the design problem is formulated with the gain factor g being the design parameter θ_d . The elastic modulus E of the four-leaf spring is assumed to be poorly known and is assigned as the calibration parameter θ_c . The mass m of the system that needs to be vibration isolated is treated as the control parameter x , and the damping ratio is the design objective y .

5.2 Experimental Observations. For the dual model calibration, 12 operational conditions for the test rig are designed for varying values of the mass m and gain factor g (shown in Table 3). A space-filling experimental design is used because a functional relationship of the calibration parameter upon the design input is not expected in this case. To excite the test rig, an impulse force is applied in the translational z -direction via a modal hammer. The time history response of the hammer excitation is shown in the left panel of Fig. 9. The right panel shows the acceleration response $\ddot{z}(t)$ of the mass, as measured by the acceleration sensor $S_{a,z}$. Since the rigid frame is constrained by a spring of small stiffness in the z -direction, the resulting relatively low resonance frequency of the frame (≈ 1.5 Hz) does not significantly affect the mass vibration with its higher eigenfrequency (>20 Hz) when vibrating in the z -direction. The low-frequency content is filtered out in the measurement chain. The hammer impact is repeated five times, and the impact force and the system response measurements are averaged.

One significant character of an oscillatory system is its damping (i.e., how rapidly a vibration system will decay after the initial excitation). The damping ratio is a dimensionless measure

that describes the damping level, and is calculated as $D = (1 + (2\pi/\delta)^2)^{-1/2}$, where $\delta = \frac{1}{n} \ln(\ddot{z}(t)/\ddot{z}(t+nT))$, $\ddot{z}(t)$ is the first peak value of mass acceleration, $\ddot{z}(t+nT)$ is the $n+1$ th peak value of mass acceleration, and n is the number of peak intervals. δ is the logarithmic decrement, which is used to compute the damping ratio D_p . By following these two equations, the system responses under various numerical simulations are summarized in Table 3.




Fig. 10 The dynamic vibration system: (1) the rigid frame, (2) the leaf springs, (3) the mass oscillator, (4) the damper, (5) the active force, and (6) the spring

Table 4 A partial parameterized input and corresponding numerical results

Case	Control parameter x	Calibration parameter θ_c	Design parameter θ_d	System response y
Variable unit	Mass (kg)	Elastic modulus (N/m ²)	Gain factor (Ns/m)	Overall damping ratio
1	0.9625	54,037,300,000	11.5	0.0979
2	0.8175	58,698,200,000	46.5	0.217
3	0.9525	72,098,300,000	76.5	0.268
4	0.7275	70,350,500,000	80.5	0.3496
5	1.0125	71,515,700,000	73.5	0.2483
6	1.0875	64,233,000,000	10.5	0.0815
...	...	...	...	...
93	1.0575	72,389,600,000	54.5	0.1855
94	0.8775	68,311,300,000	91.5	0.3621
95	0.7675	56,950,400,000	19.5	0.1326
96	0.7525	71,807,000,000	83.5	0.3489
97	0.8125	68,893,900,000	27.5	0.1392
98	1.0525	48,793,800,000	34.5	0.1679

5.3 Numerical Investigation. To fully explore the domain of the control parameter in this dual model calibration problem, a finite element model of the one-mass oscillator is built in ANSYS v. 2018 (Fig. 10). The frame and oscillatory mass are represented by

a linear solid element type C3D8R in ABAQUS. Both the frame and the mass are assigned very high stiffness values to reflect rigid body behavior.³ The rigid frame is constrained in the z -direction of vibration by spring of a small stiffness value, and laterally, by assumed gliding support (see Fig. 8(c)). A passive damping force, an active damping force (that results from the gain factor g), and elastic forces (from the leaf springs) apply on the mass oscillator. Dashpot elements are used for the damper and gain to model velocity-dependent forces. The damper represented by DASHPOT2 element introduces a damping force as a function of the relative velocity between the rigid frame and the mass oscillator, the active damping force due to gain is modeled as a function of the absolute velocity of the mass oscillator through DASHPOT1 element, and the spring is represented by SPRING1 element in ABAQUS. A Latin hypercube sampling is completed with 98 runs for parameters values partially shown in Table 4 for which the damping ratio of the system is calculated.




Fig. 11 The posterior distributions of the calibration and design inputs, respectively, along with their (uniform) priors. The very narrow posterior distribution of gain is concentrated at the minimum of its support.

5.4 Application of DCTO to Vibration Isolation Design.

Since our goal is to minimize the damping ratio, we set our target outcomes y_i to be 0 across a range of oscillator masses. Specifically, we set a grid of size eight over the range of oscillator masses present in the simulation and experimental data, with target outcome 0 for each point in that grid. We define our prior GP surrogate for the FE model using a mean function found via degree-2 polynomial regression on the available FE runs. For the hyperparameters of the surrogate's covariance function, we estimate them as MLEs using the quasi-Newton Broyden–Fletcher–Goldfarb–Shanno (BFGS) method [38]. We perform 10,000 iterations of MCMC using this surrogate and set of target observations, of which the first half are discarded as burn-in. The convergence of the resulting MCMC chains is assessed both visually and using the Gelman-Rubin statistic (≈ 1.01 and 1.001 for calibration and design, respectively), Ref. [37].

The total wall time required for the MCMC to complete DCTO in this case was 94 s (on a laptop with an Intel Core i7-9750H CPU and 16GB of RAM). The posterior distributions of the calibration and design inputs are shown in Fig. 11. Strong Bayesian learning has occurred, particularly for the design input. The posterior distribution of the elastic modulus for the system assigns a high likelihood to the expected value of 6.2×10^{10} , with a posterior mean of 6.188×10^{10} . In comparison with our design results, we also apply the gradient-free non-dominated genetic sorting algorithm II (NSGA-II) [12], to the trained GP model surrogate. We use 100 generations and a population size 50, taking a total of 48 s of computation (wall time). Whereas our method performs both calibration and design, NSGA-II cannot be used for calibration, and so we apply it to a model calibrated with a point estimate (the

³“Overall damping ratio” refers to the combination of active and passive damping.




Fig. 12 The posterior distributions of the model output at three different levels of the operational domain, along with their prior distributions. The posteriors constitute a notable performance increase over the priors.

posterior mean) of elastic modulus from our method's results. The results of NSGA-II agree with our own, in finding the optimal gain setting to be 0.

We also use the surrogate model to estimate the posterior predictive distribution of the system after DCTO. Figure 12 shows the resulting posterior distributions of model output at various levels of oscillator mass, along with the distributions of both experimental and simulator system output. For comparison, the figure also includes the output of the surrogate model using the posterior mean of elastic modulus along with the NSGA-II estimate of optimal gain. Note that the predicted model outputs fall at the bottom of the ranges of observed model outputs across the domain of oscillator masses, implying a successful design outcome for the system has been achieved.

6 Conclusion

DCTO provides a method for generalizing the KOH framework for model calibration to include design. The result secures the benefits of KOH both for calibration and for design. This includes

the ability to quantify uncertainty remaining in the true value of the calibration parameter, the optimal settings for the design input, and the resulting model output. DCTO provides a computationally efficient method of propagating the uncertainties remaining from KOH calibration through the design procedure. In the case when observations of the real system can be carried out sequentially at adaptively chosen locations, DCTO is robust to model misspecification where the calibration parameter is functionally dependent on the value of the design input and the model fails to reflect this. In such a case, if the functional form of the dependence of θ_c on θ_d is of interest, then state-aware calibration should be used. However, if one only wishes to estimate the calibration parameter at the optimal design settings, then DCTO provides a means of doing so. In this application, DCTO with adaptive sampling uses information from both the sequentially performed observations of the real system and from the existing computer model to identify new sampling locations. MCMC methods struggle to converge in the context of high dimensionality; future work on this subject will include the application to DCTO of ongoing research in the area of high-dimensional MCMC [39]. Future work will also include pairing adaptive sampling DCTO with other methodologies for selecting new sampling locations, such as EGO and SUR.

Funding Data

- National Science Foundation (NSF) (Grant Nos. 1633608, CMMI-1934438, EEC-1744497 and OIA-1826715; Funder ID: 10.13039/100000001).
- Department of Education GAANN (Grant No. P200A150310; Funder ID: 10.13039/100000138).

Nomenclature

- $C_\eta(\cdot)$ = covariance of Gaussian process emulator of η
 $C_\delta(\cdot)$ = covariance of Gaussian process model of δ
 $\mathcal{D} = \eta^T, \mathbf{y}^T)^T$ for some observations $\mathbf{y}$
 $f(\cdot)$ = the system of interest
 g = gain of the dynamic vibration system
 k = elastic modulus of leaf spring
 m = mass of oscillator in the dynamic vibration system
 $m_\eta(\cdot)$ = mean of Gaussian process emulator of η
 $m_\delta(\cdot)$ = mean of Gaussian process model of δ
 T = time period of one oscillation of a dynamic vibration system
 t_c = value of calibration parameter used as input in $\eta(\cdot)$
 t_d = value of design variable used as input in $\eta(\cdot)$
 $\mathbf{x}$ = all model inputs in the operational domain of $f(\cdot)$
 $\mathbf{y}_f$ = vector of observations of the system of interest
 $\mathbf{y}_s$ = vector of outputs of the computer model of $f(\cdot)$
 $\mathbf{y}_t$ = vector of target outcomes for the system of interest
 $\mathbf{z}$ = all known and/or controllable inputs of $f(\cdot)$
 β = inverse correlation length for Gaussian process input
 $\delta(\cdot)$ = discrepancy between model and true system
 ε_c = measurement error
 ε_d = discrepancy between optimal system output and $\mathbf{y}_t$
 ζ = Damping ratio of a dynamic vibration system
 $\boldsymbol{\eta}$ = vector of outputs of $\eta(\cdot)$
 $\eta(\cdot)$ = computer model of the system of interest
 θ_c = true value of the parameter to be calibrated
 θ_d = optimal design input
 λ = marginal precision of Gaussian process
 ρ = reparameterization of β
 σ_c^2 = variance of ε_c
 σ_d^2 = variance of ε_d
 ϕ_δ = hyperparameters of Gaussian process model of $\delta(\cdot)$
 ϕ_η = hyperparameters of Gaussian process surrogate for $\eta(\cdot)$

References

- [1] Kennedy, M. C., and O'Hagan, A., 2001, "Bayesian Calibration of Computer Models," *J. R. Stat. Soc. Ser. B*, **63**(3), pp. 425–464.

- [2] Higdon, D., Kennedy, M., Cavendish, J. C., Cafoe, J. A., and Ryne, R. D., 2004, "Combining Field Data and Computer Simulations for Calibration and Prediction," *SIAM J. Sci. Comput.*, **26**(2), pp. 448–466.
- [3] Williams, B., Higdon, D., Gattiker, J., Moore, L., McKay, M., and Keller-McNulty, S., 2006, "Combining Experimental Data and Computer Simulations, With an Application to Flyer Plate Experiments," *Bayesian Anal.*, **1**(4), pp. 765–792.
- [4] Bayarri, M. J., Berger, J. O., Cafoe, J., Garcia-Donato, G., Liu, F., Palomo, J., Parthasarathy, R. J., Paulo, R., Sacks, J., and Walsh, D., 2007, "Computer Model Validation With Functional Output," *Ann. Stat.*, **35**(5), pp. 1874–1906.
- [5] Bayarri, M. J., Berger, J. O., Paulo, R., Sacks, J., Cafoe, J. A., Cavendish, J., Lin, C.-H., and Tu, J., 2007, "A Framework for Validation of Computer Models," *Technometrics*, **49**(2), pp. 138–154.
- [6] Paulo, R., García-Donato, G., and Palomo, J., 2012, "Calibration of Computer Models With Multivariate Output," *Comput. Stat. Data Anal.*, **56**(12), pp. 3959–3974.
- [7] Brynjarsdóttir, J., and O'hagan, A., 2014, "Learning About Physical Parameters: The Importance of Model Discrepancy," *Inverse Probl.*, **30**(11), p. 114007.
- [8] Regis, R. G., and Shoemaker, C. A., 2004, "Local Function Approximation in Evolutionary Algorithms for the Optimization of Costly Functions," *IEEE Trans. Evol. Comput.*, **8**(5), pp. 490–505.
- [9] Nocedal, J., and Wright, S. J., 2006, *Numerical Optimization*, Springer, New York.
- [10] Lee, K. Y., and El-Sharkawi, M. A., 2007, *Modern Heuristic Optimization Techniques: Theory and Applications to Power Systems*, Wiley, Hoboken, NJ.
- [11] 2008, "Multiobjective Optimization: Interactive and Evolutionary Approaches," *LNC3*, J. Branke, K. Deb, K. Miettinen, and R. Slowinski, eds., Vol. 5252, Springer, New York.
- [12] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., 2002, "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, **6**(2), pp. 182–197.
- [13] Kim, M., Hiroyasu, T., Miki, M., and Watanabe, S., 2004, "SPEA2+: Improving the Performance of the Strength Pareto Evolutionary Algorithm 2," *Lecture Notes in Computer Science*, Vol. 3242, Heidelberg, Germany, pp. 742–751.
- [14] Bonyadi, M. R., and Michalewicz, Z., 2017, "Particle Swarm Optimization for Single Objective Continuous Space Problems: A Review," *Evolutionary Computation*, **25**(1), pp. 1–54.
- [15] Mason, K., Duggan, J., and Howley, E., 2017, "Multi-Objective Dynamic Economic Emission Dispatch Using Particle Swarm Optimisation Variants," *Neurocomputing*, **270**, pp. 188–197.
- [16] Robert, C. P., and Casella, G., 2004, "Monte Carlo Optimization," *Monte Carlo Statistical Methods*, 2nd ed., Chap. 5, Springer, New York, pp. 157–204.
- [17] Deb, K., and Gupta, H., 2006, "Introducing Robustness in Multi-Objective Optimization," *Evol. Comput.*, **14**(4), pp. 463–494.
- [18] Zhou, A., Qu, B.-Y., Li, H., Zhao, S.-Z., Suganthan, P. N., and Zhang, Q., 2011, "Multiobjective Evolutionary Algorithms: A Survey of the State of the Art," *Swarm. Evol. Comput.*, **1**(1), pp. 32–49.
- [19] Dean, A., Voss, D., and Draguljić, D., 2017, *Response Surface Methodology*, Springer International Publishing, Cham, pp. 565–614.
- [20] Jones, D. R., Schonlau, M., and Welch, W. J., 1998, "Efficient Global Optimization of Expensive Black-Box Functions," *J. Global Optim.*, **13**(4), pp. 455–492.
- [21] Brochu, E., Cora, V. M., and de Freitas, N., 2010, "A Tutorial on Bayesian Optimization of Expensive Cost Functions, With Application to Active User Modeling and Hierarchical Reinforcement Learning," arXiv preprint [arXiv:1012.2599](https://arxiv.org/abs/1012.2599).
- [22] Geman, D., and Jedynak, B., 1996, "An Active Testing Model for Tracking Roads in Satellite Images," *IEEE Trans. Pattern Anal. Mach. Intell.*, **18**(1), pp. 1–14.
- [23] Villemonteix, J., Vazquez, E., and Walter, E., 2009, "An Informational Approach to the Global Optimization of Expensive-to-Evaluate Functions," *J. Global Optim.*, **44**(4), pp. 509–534.
- [24] Chevalier, C., Bect, J., Ginsbourger, D., Vazquez, E., Picheny, V., and Richet, Y., 2014, "Fast Parallel Kriging-Based Stepwise Uncertainty Reduction With Application to the Identification of an Excursion Set," *Technometrics*, **56**(4), pp. 455–465.
- [25] Picheny, V., 2015, "Multiobjective Optimization Using Gaussian Process Emulators Via Stepwise Uncertainty Reduction," *Stat. Comput.*, **25**(6), pp. 1265–1280.
- [26] Hernández-Lobato, J. M., Gelbart, M. A., Adams, R. P., Hoffman, M. W., and Ghahramani, Z., 2016, "A General Framework for Constrained Bayesian Optimization Using Information-Based Search," *J. Mach. Learn. Res.*, **17**(160), pp. 1–53.
- [27] Picheny, V., Binois, M., and Habbal, A., 2019, "A Bayesian Optimization Approach to Find Nash Equilibria," *J. Global Optim.*, **73**(1), pp. 171–192.
- [28] Binois, M., Picheny, V., Taillardier, P., and Habbal, A., 2020, "The Kalai-Smorodinsky Solution for Many-Objective Bayesian Optimization," *J. Mach. Learn. Res.*, **21**(150), pp. 1–42.
- [29] Mockus, J., Tiesis, V., and Zilinskas, A., 1978, "The Application of Bayesian Methods for Seeking the Extremum," Towards Global Optim. 2, Dixon, L. and Szego, G., eds., Elsevier, Amsterdam, The Netherlands, pp. 117–129.
- [30] Olalotiti-Lawal, F., and Datta-Gupta, A., 2018, "A Multiobjective Markov Chain Monte Carlo Approach for History Matching and Uncertainty Quantification," *J. Pet. Sci. Eng.*, **166**, pp. 759–777.
- [31] Gelfand, A. E., and Smith, A. F. M., 1990, "Sampling-Based Approaches to Calculating Marginal Densities," *J. Am. Stat. Assoc.*, **85**(410), pp. 398–409.
- [32] Liu, F., Bayarri, M. J., and Berger, J. O., 2009, "Modularization in Bayesian Analysis, With Emphasis on Analysis of Computer Models," *Bayesian Anal.*, **4**(1), pp. 119–150.
- [33] Atamturktur, S., and Brown, D. A., 2015, "State-Aware Calibration for Inferring Systematic Bias in Computer Models of Complex Systems," NAFEMS World Congress Proceedings, San Diego, CA, June 21–24, pp. 21–24.
- [34] Atamturktur, S., Hegenderfer, J., Williams, B., Egeberg, M., Lebensohn, R. A., and Unal, C., 2015, "Mechanics of Advanced Materials and Structures: A Resource Allocation Framework for Experiment-Based Validation of Numerical Models," *Mech. Adv. Mater. Struct.*, **22**(8), pp. 641–654.
- [35] Ezzat, A. A., Pourhabib, A., and Ding, Y., 2018, "Sequential Design for Functional Calibration of Computer Models," *Technometrics*, **60**(3), pp. 286–296.
- [36] Brown, D. A., and Atamturktur, S., 2018, "Nonparametric Functional Calibration of Computer Models," *Stat. Sin.*, **28**(2), pp. 721–742.
- [37] Gelman, A., and Rubin, D. B., 1992, "Inference From Iterative Simulation Using Multiple Sequences," *Stat. Sci.*, **7**(4), pp. 457–472.
- [38] Fletcher, R., 2013, *Practical Methods of Optimization*, Wiley, Chichester, UK.
- [39] Saibaba, A. K., Bardsley, J., Brown, D. A., and Alexanderian, A., 2019, "Efficient Marginalization-Based MCMC Methods for Hierarchical Bayesian Inverse Problems," *SIAM/ASA J. Uncertainty Quantif.*, **7**(3), pp. 1105–1131.