Understanding the Effects of Model Uncertainty in Robust Design With Computer Experiments

Daniel W. Apley[1]
Associate Professor

Jun Liu
Graduate Student

Department of Industrial Engineering and Management Sciences,
Northwestern University,
Evanston, IL, 60208-3119

Wei Chen
Associate Professor
Department of Mechanical Engineering,
Northwestern University,
Evanston, IL, 60208-3111

The use of computer experiments and surrogate approximations (metamodels) introduces a source of uncertainty in simulation-based design that we term model interpolation uncertainty. Most existing approaches for treating interpolation uncertainty in computer experiments have been developed for deterministic optimization and are not applicable to design under uncertainty in which randomness is present in noise and/or design variables. Because the random noise and/or design variables are also inputs to the metamodel, the effects of metamodel interpolation uncertainty are not nearly as transparent as in deterministic optimization. In this work, a methodology is developed within a Bayesian framework for quantifying the impact of interpolation uncertainty on the robust design objective, under consideration of uncertain noise variables. By viewing the true response surface as a realization of a random process, as is common in kriging and other Bayesian analyses of computer experiments, we derive a closed-form analytical expression for a Bayesian prediction interval on the robust design objective function. This provides a simple, intuitively appealing tool for distinguishing the best design alternative and conducting more efficient computer experiments. We illustrate the proposed methodology with two robust design examples—a simple container design and an automotive engine piston design with more nonlinear response behavior and mixed continuous-discrete design variables. [DOI: 10.1115/1.2204974]

Keywords: model uncertainty, interpolation uncertainty, metamodel, robust design, Bayesian prediction interval, computer experiments

1 Introduction

Simulation-based engineering analyses (e.g., finite element analysis (FEA), computational fluid dynamics) are increasingly important tools for product design [1–3]. This is particularly true when the need to reduce product development time and cost necessitates early-stage simulation analyses before actual production or prototype experimentation is possible. On a parallel front, with increasing expectations for high quality, reliable, low cost products, the issue of robust product design under uncertainty continues to grow in importance [4–15]. In robust design, key system response variables are viewed as functions of a set of design variables (or "control" variables) specified as part of the decision making in product design and a set of uncontrollable "noise" variables that involve uncertainty. If we assign probability distributions to the noise variables, then the response variables are also viewed as probabilistic. Broadly speaking, the robust design objective is to select the design variables in order to optimize certain probabilistic characteristics of the objective responses (typically involving the mean μ and standard deviation σ of each response) [16], while ensuring that the constraint responses satisfy certain conditions with some desired probability [6,17–20].

In the simulation arena, increases in computing power are often offset by increased accuracy (e.g., finer mesh sizes in FEA) and increased simulation scope/complexity (e.g., multidisciplinary simulation), so that computational expense remains a barrier. For example, finite-element based automotive structural simulations to assess crash worthiness may take days to execute a single run [21]. The inevitable result is that it generally becomes impossible to conduct enough simulation runs to thoroughly cover the entire input variable space and provide a globally accurate understanding of the functional relationship between the response and the inputs. We must instead attempt to optimize the design based on the output of simulation runs conducted at only a limited number of input variable combinations (also known as sampling sites).

A common strategy in this situation is to fit some form of surrogate model (also called "metamodel" [22]) to the samples and to rely on that for optimization purposes. Comprehensive reviews of metamodeling applications in mechanical and aerospace systems can be found in Refs. [23–25]. Typical choices for metamodel structure are polynomial functions [26], kriging models [27,28], and radial basis function networks [29,30]. Examples of recent studies on enhancing metamodeling techniques for engineering design applications include the work from Refs. [31–40]. Although the majority of prior work on metamodel-based design is for deterministic optimization, a number of researchers [21,41,42] have used fitted metamodels for robust design optimization with noise variables.

The obvious drawback of simply using a fitted metamodel to optimize the design is that it may be quite far from the true response surface. In essence, it ignores metamodel uncertainty, i.e., the uncertainty that results from not knowing the output of the computer simulation code, except at a finite set of sampling sites. Kennedy and O'Hagan [43] referred to this type of uncertainty as *code uncertainty*. We adopt the term *interpolation uncertainty*, because we believe it is more descriptive: The uncertainty disappears precisely at the sampling sites because the simulation code is repeatable, and the uncertainty grows as we stray further from the sampling sites and attempt to interpolate the response between sampling sites using some appropriately fitted metamodel.

[1]Corresponding author.

Contributed by the Design Automation Committee of ASME for publication in the JOURNAL OF MECHANICAL DESIGN. Manuscript received September 16, 2005; final manuscript received December 19, 2005. Review conducted by Shapour Azarm.

Widespread efforts have been made in the statistical community to quantify interpolation uncertainty in computer experiments [28,43,44]. In one of the most popular approaches, the response surface is viewed as a realization of a Gaussian random process (GRP), and Bayesian methods are used to interpolate/extrapolate the response surface by calculating its posterior distribution given the sampling sites. This provides a direct means of representing the effects of interpolation uncertainty in *deterministic* optimization problems, in which the objective function is simply the output response of the computer simulation. Bayesian analyses with GRPs have been used in deterministic optimization to develop sequential sampling strategies for the purposes of either efficiently arriving at the optimum [45,46] or capturing local irregularities in the global response behavior [39,40].

Although Jin et al. [47] have argued that metamodel inaccuracy has an even larger effect in *robust design* optimization than in deterministic optimization, very few existing approaches for treating interpolation uncertainty in computer experiments apply in the former situation. The reason is that in robust design optimization, the effects of interpolation uncertainty on the objective function are not nearly as transparent as in deterministic optimization, because some of the metamodel inputs are random (the noise variables). Consider a robust design objective function that is a weighted combination $c_1\mu+c_2\sigma$ of the response mean and standard deviation (or equivalently, $\mu+c\sigma$ with weighting constant $c=c_2/c_1$), which might be appropriate if the design objective was to make the response as small as possible ("smaller-the-better") under consideration of the uncertainty in the noise variables. The mean μ and variance σ^2 of the response must be obtained by integrating the *uncertain* response surface over the full range of values of the noise variables. Moreover, unlike in deterministic optimization, we cannot eliminate the effects of interpolation uncertainty with validation runs in robust design optimization, because this would require a large number of simulations covering the full range of values for the noise variables at the selected design sites. Due to these difficulties, existing metamodel-based design approaches are either for deterministic optimization problems or, in the case of robust design optimization with noise variables, they simply rely on a fitted global metamodel while ignoring interpolation uncertainty.

The aforementioned difficulties pose a number of challenges and research issues. The fundamental question is how should we quantify the effects of interpolation uncertainty on the *robust design objective function with uncertain noise variables*, and subsequently, how can we use this information for design decision making, e.g., to distinguish the best design alternative. Because the computer experiment drives the accuracy of a metamodel, another fundamental question is how the quantification of interpolation uncertainty can be used to guide the computer experiment in determining whether or not additional samples are needed and, if so, how to select the additional sampling site(s)? A final issue is the potential computational challenges of the proposed method, especially in problems involving a large number of design and noise variables.

To address these issues, in this work we develop an approach using Bayesian prediction intervals (PIs) for quantifying the effects of interpolation uncertainty on the objective function in robust design. To better illustrate the concepts and introduce the proposed methods, a running design example (described in Sec. 2) is used throughout the paper. A brief background on Bayesian analysis of GRPs is presented next (Sec. 3). In Sec. 4 we derive closed-form analytical expressions for the prediction intervals on the response mean, the response variance, and the robust design objective function. The prediction intervals are useful in many contexts, for example to evaluate a candidate design or compare two or more candidate designs. Based on the prediction intervals, in Sec. 5 we provide a method for deciding whether to terminate or continue simulation. The stopping criterion is objective oriented and is intended to terminate when further simulation will not lead to a significantly better design. In Sec. 6 we discuss simple graphical methods for guiding the selection of additional sampling sites. An automotive engine piston design problem is presented in Sec. 7 to illustrate the applicability of our method for problems with nonlinear function behavior and mixed continuous-discrete design variables. In Sec. 8 we discuss the handling of various computational issues. Conclusions are summarized in Sec. 9.

2 An Illustrative Design Example

Consider the following simple design optimization example from Ref. [48], in which the task is to design an open box that will be used to ferry a specified quantity of gravel across a river. The design objective is to minimize the total cost $y(d,W)=80d^{-2}+2dW+d^2W$ (refer to Ref. [48] for details), where the first term is the cost of ferrying, the remaining two terms are the total material cost, and the variables are as follows. The noise variable W represents the material cost per unit, and the design variable d represents the length and width of the square box. In order to represent the material cost uncertainty, we assign to W a normal probability distribution $W\sim N(\mu_w,\sigma_w^2)$ with mean $\mu_w=10$ \$/m^2 and variance $\sigma_w^2=2$ \$^{2}/m^4. The response is plotted in Fig. 1(*a*).

We cannot minimize cost directly, because of the uncertainty in W. As a *robust* design objective (robust with respect to uncertainty in W), we will instead minimize the objective function $f(d)\equiv\mu(d)+c\sigma(d)$, where $\mu(d)$ and $\sigma(d)$ represent the mean and standard deviation of the response (total cost) with respect to variations in W. Suppose we choose $c=3$ as our desired relative weighting of μ and σ (choice of c is discussed in Sec. 4). If the true response surface in Fig. 1(*a*) were known, this would be a rather trivial design problem. Imagine that instead of having the analytical expression for $y(d,w)$, however, we must rely on a computationally expensive simulation to evaluate the response and a fitted metamodel to optimize the design. Fig. 1(*b*) shows the simulation output at four sampling sites, as well as one particular fitted metamodel (the details of which are discussed in the following section) denoted by $\hat{y}(d,w)$. If we ignored interpolation uncertainty and treated the fitted model as the true response surface, we could calculate the response mean and variance as a function of d via

$$\mu(d)=\int \hat{y}(d,w)p_w(w)dw, \tag{1}$$

and

$$\sigma^2(d)=\int [\hat{y}(d,w)-\mu(d)]^2p_w(w)dw \tag{2}$$

for use in the objective function $\mu(d)+3\sigma(d)$. Here, $p_w(w)$ denotes the probability density function of w. Figure 1(*d*) shows $\mu(d)$, $\sigma(d)$, and $f(d)$ calculated from Eqs. (1) and (2), and Fig. 1(*c*) shows the corresponding values calculated by substituting the true response surface for $\hat{y}(d,w)$ in Eqs. (1) and (2).

Ignoring metamodel uncertainty clearly has a large effect on the design optimization. From Fig. 1(*d*), the optimal design that would result from assuming that the true response surface coincides with $\hat{y}(d,w)$ is $d=0.8$ (we have defined the feasible design space as $0.8\leq d\leq 2.5$). In contrast, the optimal design that would result if we knew the true response surface is $d=1.3$ (see Fig. 1(*c*)). In the subsequent sections, we propose a method for quantifying the effects of interpolation uncertainty on the robust design objective function f, to help designers avoid selecting a design that is far from optimal because of insufficient computer simulation and the resulting interpolation uncertainty.

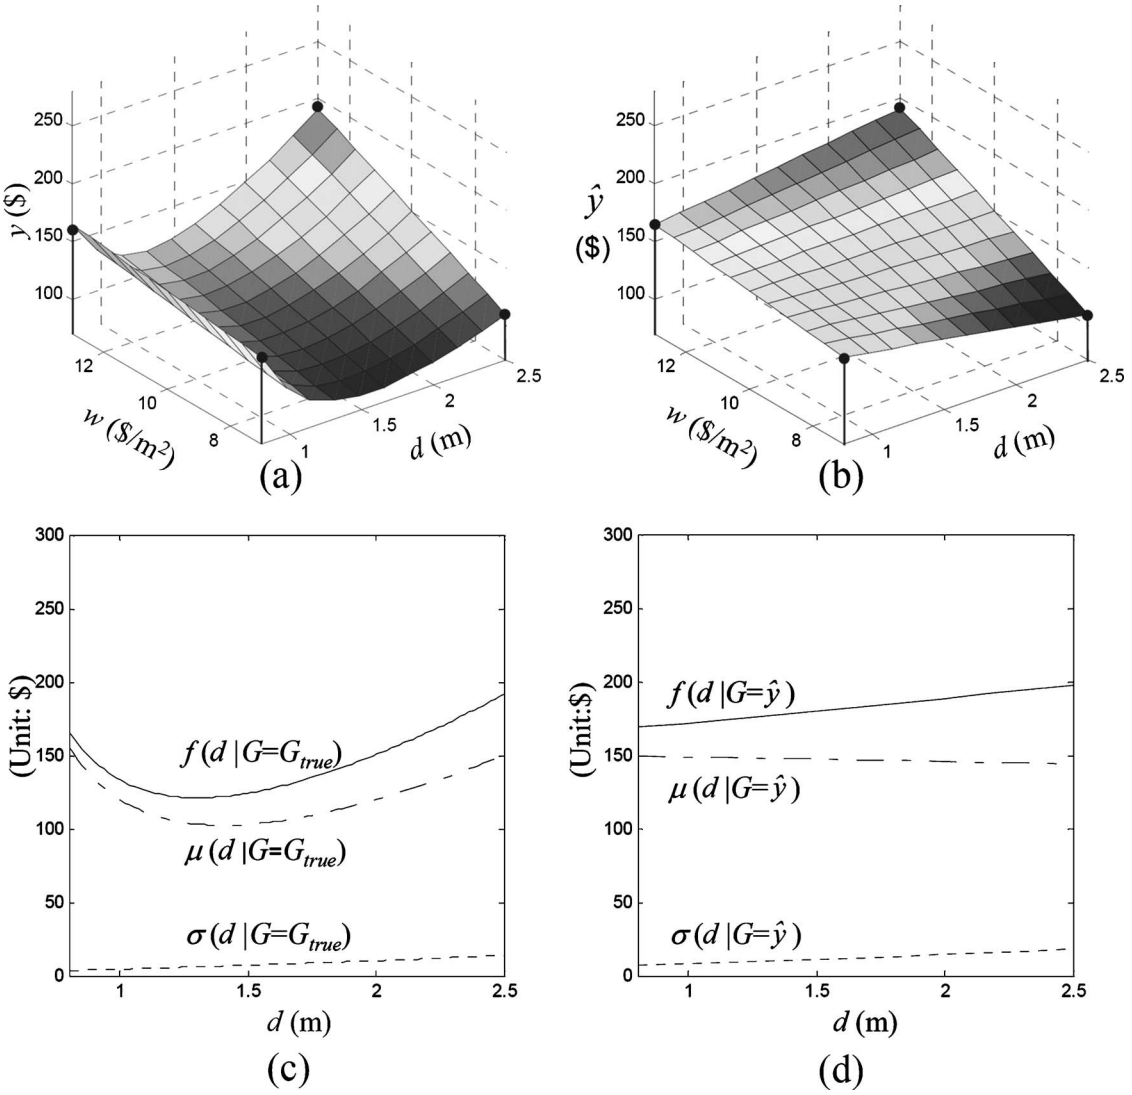

Fig. 1 (*a*) True response surface (treated as unknown), the four bullets indicate simulation output at four sampling sites. (*b*) Fitted surrogate model $\hat{y}$ based on the four sampling sites (*c*) Mean μ, standard deviation σ, and robust design objective function *f* calculated from the true response surface. (*d*) Mean, standard deviation, and objective function calculated from the fitted model $\hat{y}$.

3 Review of Gaussian Random Processes and Their Bayesian Analysis

We provide a brief review of Bayesian analysis of GRPs for representing computer simulation response surfaces in this section. More comprehensive discussions can be found in Refs. [28,43,44,49]. Let the vectors $\boldsymbol{d}=[d_1,d_2,\ldots,d_{q_d}]^T$ and $\boldsymbol{w}=[w_1,w_2,\ldots,w_{q_w}]^T$ denote the sets of design and noise variables, and let $y(\boldsymbol{x})$ be a scalar response variable, where $\boldsymbol{x}=[\boldsymbol{d}^T,\boldsymbol{w}^T]^T$ denotes the vector of inputs. Suppose we run a computer simulation at a set of N separate input sampling sites $\mathbf{S}^N=\{\mathbf{s}_1,\mathbf{s}_2,\ldots,\mathbf{s}_N\}$ to obtain a vector of N responses $\boldsymbol{y}^N=[y(\boldsymbol{s}_1),y(\boldsymbol{s}_2),\ldots,y(\boldsymbol{s}_N)]^T$. Figures 1(*a*) and 1(*b*) show the response at $N=4$ sampling sites. Given the observed values $\boldsymbol{y}^N$, two natural questions are how can we predict (in this case, interpolate/extrapolate) the response as we move away from the specific sampling sites in $\mathbf{S}^N$, and how can we reasonably quantify the prediction (interpolation) uncertainty?

Bayesian analysis of GRPs, which one might view as a generalization of the kriging method that was popularized in the field of geostatistics [50], is the most widely investigated method for these purposes. The basic idea is to view the actual response surface $y(\boldsymbol{x})$ as a realization of a Gaussian random process $G(\boldsymbol{x})$ with prior mean function $E[G(\boldsymbol{x})]=\boldsymbol{h}^T(\boldsymbol{x})\boldsymbol{\beta}$ and prior covariance function $\mathrm{Cov}[G(\boldsymbol{x}),G(\boldsymbol{x}')]=\alpha^2R_{\phi}(\boldsymbol{x},\boldsymbol{x}')$ (these are intended to be prior to actually running the computer simulation). Here, $\boldsymbol{\beta}=[\beta_1,\beta_2,\ldots,\beta_k]^T$ is a vector of unknown parameters, $\boldsymbol{h}^T(\boldsymbol{x})=[h_1(\boldsymbol{x}),h_2(\boldsymbol{x}),\ldots,h_k(\boldsymbol{x})]^T$ is a vector of known functions of $\boldsymbol{x}$ (e.g., linear or quadratic), α^2 is the variance of $G(\cdot)$, and for any two input vectors $\boldsymbol{x}$ and $\boldsymbol{x}'$, $R_{\phi}(\boldsymbol{x},\boldsymbol{x}')$ is the correlation coefficient between $G(\boldsymbol{x})$ and $G(\boldsymbol{x}')$, parameterized in terms of a vector $\boldsymbol{\phi}$ of parameters. One of the most common choices for correlation function is

$$R_{\phi}(\boldsymbol{x},\boldsymbol{x}')=\prod_{i=1}^{q}\exp[-\phi_i(x_i-x_i')^2] \tag{3}$$

where $q=q_d+q_w$. Because $\boldsymbol{\phi}$ determines the rate at which the correlation between $G(\boldsymbol{x})$ and $G(\boldsymbol{x}')$ decays as $\boldsymbol{x}$ and $\boldsymbol{x}'$ move further apart, $\boldsymbol{\phi}$ and α relate to the smoothness of the response surface and the extent to which it might differ from a parametric surface of the form $\boldsymbol{h}^T(\boldsymbol{x})\boldsymbol{\beta}$. A noninformative prior distribution is often assigned to $\boldsymbol{\beta}$. Although we may view α^2 and $\boldsymbol{\phi}$ as hyper-

parameters and assign to them a prior distribution as well, it is more common (for computational purposes) to estimate them from the data and then simply view the estimates as given values. We discuss this in Sec. 8(*a*).

For notational purposes, let $Y(\cdot)$ denote the response surface $y(\cdot)$ when viewed as a realization of the random process and write $Y(G,\boldsymbol{d},\boldsymbol{W})=G(\boldsymbol{d},\boldsymbol{W})=G(\boldsymbol{x})$. This allows us to express the response $Y(\cdot)$ as a function of the design $\boldsymbol{d}$ and two sources of uncertainty—the noise $\boldsymbol{W}$ and the interpolation uncertainty represented by $G(\cdot)$. The random process $G(\cdot)$, together with a prior mean and covariance function and a prior distribution for $\boldsymbol{\beta}$, are used to represent our prior uncertainty in the response surface $Y(\cdot)$. Given the observed values $\boldsymbol{y}^N$, it is quite straightforward to use Bayesian methods to calculate the *posterior* distribution of $G(\cdot)$. The posterior distribution of $G(\cdot)$ provides a direct quantification of the effects of interpolation uncertainty on the predicted response. For a noninformative prior on $\boldsymbol{\beta}$, the posterior distribution is Gaussian with posterior mean and covariance given by [27,51]

$$\hat{y}(\boldsymbol{x}) \equiv E_{G|\boldsymbol{w},\boldsymbol{y}^N}[Y(G,\boldsymbol{d},\boldsymbol{w})|\boldsymbol{w},\boldsymbol{y}^N] = \boldsymbol{h}^T(\boldsymbol{x})\hat{\boldsymbol{\beta}} + \boldsymbol{r}^T(\boldsymbol{x})\mathbf{R}^{-1}(y^N - \mathbf{H}\hat{\boldsymbol{\beta}}), \tag{4}$$

and

$$\begin{aligned}\mathrm{Cov}[Y(G,\boldsymbol{d},\boldsymbol{w}),Y(G,\boldsymbol{d}',\boldsymbol{w}')|\boldsymbol{y}^N] &\equiv E_{G|\boldsymbol{w},\boldsymbol{w}'\boldsymbol{y}^N}\{[Y(G,\boldsymbol{d},\boldsymbol{w}) - \mu(\boldsymbol{d},\boldsymbol{w}|\boldsymbol{y}^N)][Y(G,\boldsymbol{d}',\boldsymbol{w}') \\ &\quad - \mu(\boldsymbol{d}',\boldsymbol{w}'|\boldsymbol{y}^N)]|\boldsymbol{w},\boldsymbol{w}',\boldsymbol{y}^N\} \\ &= \alpha^2\{R_\phi(\boldsymbol{x},\boldsymbol{x}') - \boldsymbol{r}^T(\boldsymbol{x})\mathbf{R}^{-1}\boldsymbol{r}(\boldsymbol{x}') + [\boldsymbol{h}(\boldsymbol{x}) \\ &\quad - \mathbf{H}^T\mathbf{R}^{-1}\boldsymbol{r}(\boldsymbol{x})]^T[\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}]^{-1}[\boldsymbol{h}(\boldsymbol{x}') - \mathbf{H}^T\mathbf{R}^{-1}\boldsymbol{r}(\boldsymbol{x}')]\}\end{aligned} \tag{5}$$

where $\mathbf{H}$ is an $N\times k$ matrix whose ith row is $\boldsymbol{h}^T(\boldsymbol{s}_i)$, $\mathbf{R}$ is an $N\times N$ matrix whose ith row, jth column element is $R_\phi(\boldsymbol{s}_i,\boldsymbol{s}_j)$, $\boldsymbol{r}(\boldsymbol{x})$ is an $N\times 1$ vector whose ith element is $R_\phi(\boldsymbol{x},\boldsymbol{s}_j)$, and $\hat{\boldsymbol{\beta}} = [\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}]^{-1}\mathbf{H}^T\mathbf{R}^{-1}\boldsymbol{y}^N$ is the posterior mean of $\boldsymbol{\beta}$. The subscripts on the expectation operator $EG|w,y^N[\cdot]$ indicate that the expectation is with respect to the posterior distribution of $G(\cdot)$ given $\boldsymbol{W}=\boldsymbol{w}$ and $\boldsymbol{y}^N$. Note that conditioning the expectation on $\boldsymbol{W}$ treats it as a fixed quantity. The preceding results, as well as the vast majority of the literature on Bayesian analysis of computer experiments, have only considered deterministic inputs $\boldsymbol{x}$. The distinction that we will make in subsequent sections between the deterministic $\boldsymbol{d}$ and the random $\boldsymbol{W}$ that characterizes robust design problems is a nontrivial one.

The usual interpretation of Eqs. (4) and (5) is the following: For any site $\boldsymbol{x}$ in the input space, $\hat{y}(\boldsymbol{x})$ is the best prediction of the response $Y(\boldsymbol{x})$ at that site, and the prediction error variance is given by Eq. (5) with $\boldsymbol{x}'=\boldsymbol{x}$. We may view $\hat{y}(\boldsymbol{x})$ as the fitted surface. One can easily verify that for any sampling site $\boldsymbol{s}_i = [\boldsymbol{d}_i^T,\boldsymbol{w}_i^T]^T$ at which we have already run a simulation, the prediction reduces to $\hat{y}(\boldsymbol{s}_i)=y(\boldsymbol{s}_i)$ and the prediction error variance is zero, as we would desire with perfectly repeatable computer experiments.

The reasons why GRPs and Bayesian analyses are so popular for representing and analyzing computer experiment response surfaces becomes apparent when one considers three fundamental criteria that any metamodel should possess: (i) It should be convenient to work with, allowing tractable analyses. (ii) The structure of the fitted surface $\hat{y}(\boldsymbol{x})$ should be sufficiently versatile to represent a wide variety of true response surfaces. (iii) It should provide an inherent mechanism for quantitatively assessing the interpolation uncertainty in the fitted surface. The GRP approach satisfies all three criteria. The flexibility of the structure of $\hat{y}(\boldsymbol{x})$ in Eq. (4) stems in part from the generality of the $\boldsymbol{h}(\boldsymbol{x})$ term. As pointed out in Ref. [28,44], many of the most common model structures used in metamodeling and approximation theory can be viewed as special cases of the Bayesian approach, depending on the choice for $\boldsymbol{h}(x)$ and $\boldsymbol{r}(x)$. These include polynomial functions [26], kriging models [27,28], and radial basis function networks [29,30].

If the form of the fitted metamodel $\hat{y}(\boldsymbol{x})$ ends up resembling the rather familiar polynomial or radial basis function models, one might question why it is couched in the more abstract language of GRPs. The reason relates to the third criterion of quantitatively assessing the interpolation uncertainty. Standard statistical methods for assessing uncertainty in models fitted to random experimental data (see, e.g., Ref. [52]) are inapplicable with computer experiments, because all of the common estimates of the variance of the model prediction error are meaningless if the response is perfectly repeatable, as is the case with computer experiments.

The GRP approach, on the other hand, is ideally suited for this situation. The random process G should be viewed primarily as a means of representing our prior knowledge of what we might expect the response surface to look like, before any simulations are run. When this prior knowledge is integrated with the observed response values at the sampling sites within the context of a Bayesian analysis, it dictates our posterior expectations of what the response surface might look like between sampled sites. To illustrate, we continue the container design example by calculating the posterior mean and covariance from Eqs. (4) and (5) based on the four design sites shown in Fig. 1(*b*) using prior parameters $\boldsymbol{\phi}=[\phi_d,\phi_w]=[2,0.1]$ and $\alpha=35$, $\boldsymbol{h}(\boldsymbol{x})=[1,d,w,dw]^T$, and a noninformative prior for $\boldsymbol{\beta}$ (Sec. 8.1 discusses estimation of the prior parameters). The posterior mean $\hat{y}(d,w)$ is that shown in Fig. 1(*b*). A number of different GRP realizations were generated from this *posterior* distribution, of which three representative realizations (call them G_1, G_2, and G_3) are plotted in Figs. 2(*a*)–2(*c*). Refer to Sec. 8.2 for details on how each GRP realization is generated.

Each of the three different GRP realizations reflect what we might expect the complete true response surface to look like, based on only the four sampling sites and our limited prior knowledge. If one were to generate and plot a large number of GRP realizations, instead of just the three shown in Figs. 2(*a*)–2(*c*), the differences between them would directly reflect interpolation uncertainty. Specifically, if we consider a fixed point (d,w) in the input space and look at the response $y(d,w)$ evaluated at that point for each GRP realization, the mean and variance of the set of response values would be given by Eqs. (4) and (5). It is in this sense that the GRP approach allows us to quantify the effects of interpolation uncertainty.

4 Derivation of Prediction Intervals for the Response Mean and Variance and the Objective Function

In this section, we derive a form of Bayesian PI for the mean and variance of the response, which we propose as a direct measure of the impact of interpolation uncertainty on the response mean and variance. This will lead to a PI on the robust design objective function for objective functions that are a weighted combination of the response mean and standard deviation as in Eq. (8) below. *Hypothetically*, if we had a complete response surface $G(\boldsymbol{x})$ available, then in analogy with Eqs. (1) and (2) we could calculate the response mean and variance (with respect to variations in the noise variables $\boldsymbol{W}$ over their joint distribution $p_{\boldsymbol{w}}(\cdot)$) directly via

$$\mu(\boldsymbol{d}|G) \equiv E_{\boldsymbol{W}|G}[Y(G,\boldsymbol{d},\boldsymbol{W})|G] \equiv \int G(\boldsymbol{d},\boldsymbol{w})p_{\boldsymbol{w}}(\boldsymbol{w})d\boldsymbol{w}, \tag{6}$$

and

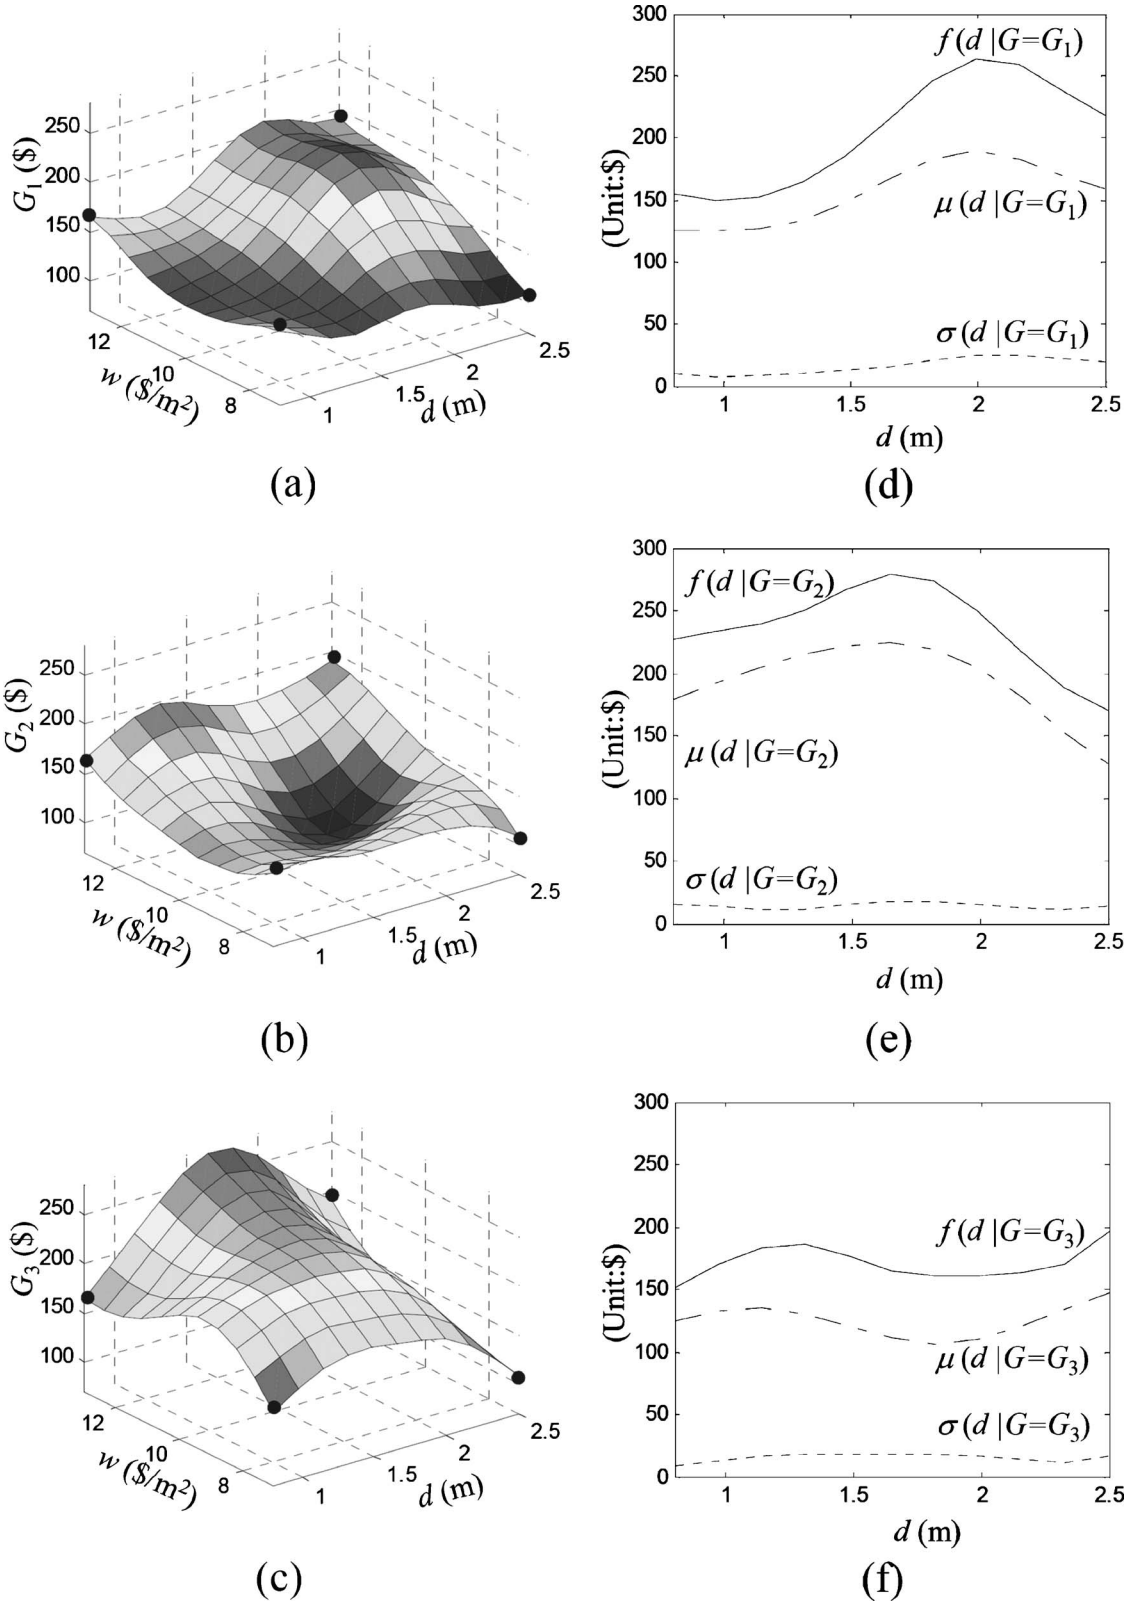

Fig. 2 (*a*)–(*c*) **Three realizations of the posterior Gaussian random process *G*, each of which represents a potential response surface consistent with the four simulated output in Fig. 1(*b*). (*d*)–(*f*) Mean, standard deviation, and objective function calculated from the potential response surfaces in (*a*)–(*c*). Prior parameters:** $[\phi_d, \phi_w]=[2, 0.1]$ **and** $\alpha=35$.

$$\sigma^2(\boldsymbol{d}|G) \equiv E_{\boldsymbol{W}|G}\{[Y(G,\boldsymbol{d},\boldsymbol{W}) - \mu(\boldsymbol{d}|G)]^2|G\}$$
$$\equiv \int [G(\boldsymbol{d},\boldsymbol{w}) - \mu(\boldsymbol{d}|G)]^2 p_w(\boldsymbol{w})d\boldsymbol{w} \quad (7)$$

The notation throughout will be that $p(\cdot)$ denotes a probability density, and subscripts indicate which random quantity the distribution is with respect to and/or conditioned on. Note that the quantities plotted in Fig. 1(*c*) are $\mu(\boldsymbol{d}|G=G_{\text{true}})$ and $\sigma(\boldsymbol{d}|G=G_{\text{true}})$, where G_{true} denotes the unknown true response surface, and the corresponding quantities in Fig. 1(*d*) are $\mu(d|G=\hat{y})$ and $\sigma(d|G=\hat{y})$.

For the remainder of the paper, the robust design objective that we consider is to minimize the objective function

$$f(\boldsymbol{d}|G) \equiv \mu(\boldsymbol{d}|G) + c\sigma(\boldsymbol{d}|G) \quad (8)$$

for some specified constant c (e.g., $c=3$). For simplicity, we consider unconstrained optimization. The approach can be readily extended to the more practical constrained optimization scenario as we briefly discuss in Sec. 9. The objective function (8) assumes that we would like the response to be as small a possible (in some probabilistic sense). For situations in which we would like the response to be as large as possible, we could instead attempt to maximize an objective function of the form $\mu(\boldsymbol{d}|G)-c\sigma(\boldsymbol{d}|G)$. The treatment would be similar to what we describe in this section.

We might view the minimization of (8) as simply attempting to ensure that both the mean and the standard deviation of the response are small. Alternatively, we might view $\mu(\boldsymbol{d}|G)+c\sigma(\boldsymbol{d}|G)$ as an approximate probabilistic upper bound on the response. For example, if the response is approximately normally distributed, then the probability that $Y(G,\boldsymbol{d},\boldsymbol{W})|G \leqslant \mu(\boldsymbol{d}|G)+c\sigma(\boldsymbol{d}|G)$ is controlled by choice of c ($c=3$ translates to a probability of 0.9987). Minimizing (8) therefore corresponds to minimizing the worst-case value for (i.e., a probabilistic upper bound on) the response under consideration of the uncertainty in $\boldsymbol{W}$. Regardless of which viewpoint one takes, the choice of c should reflect the designer's risk aversion. A choice of $c=0$ represents complete risk neutrality, because this would result in simply minimizing the response mean $\mu(\boldsymbol{d}|G)$ with complete disregard of the response variance. A large c represents high risk aversion. For many design problems the "optimal" design will obviously depend strongly on the desired level of risk aversion.

Of course, we cannot substitute $G=G_{\text{true}}$ into Eqs. (6)–(8) to calculate $\mu(\boldsymbol{d}|G=G_{\text{true}})$, $\sigma(\boldsymbol{d}|G=G_{\text{true}})$, nor the objective function $f(\boldsymbol{d}|G=G_{\text{true}})$, because we do not have the complete response surface G_{true}. Instead, we propose to use prediction intervals for these quantities in order to guide the design process. The Bayesian PI approach involves viewing the quantities $\mu(\boldsymbol{d}|G)$, $\sigma(\boldsymbol{d}|G)$, and $f(\boldsymbol{d}|G)$ as random, being functions of the random surface G. Because the (posterior) distribution of G is used to represent interpolation uncertainty, the resulting (posterior) distribution of $f(\boldsymbol{d}|G)$, for example, represents the effects of interpolation uncertainty on the objective function. This is illustrated in Fig. 2, the bottom panels of which show $\mu(\boldsymbol{d}|G)$, $\sigma(\boldsymbol{d}|G)$, and $f(\boldsymbol{d}|G)$ for the three realizations of the random G shown in the top panels. The quantities differ randomly between the three figures, because they are obtained by integrating the random surface G. Note that $\mu(\boldsymbol{d}|G)$, $\sigma(\boldsymbol{d}|G)$, and $f(\boldsymbol{d}|G)$ are random because of their functional dependence on G and not because of any dependence on $\boldsymbol{W}$. The random effects of $\boldsymbol{W}$ have already been averaged out via the integration in Eqs. (6) and (7).

This suggests a straightforward approach for quantifying the effects of interpolation uncertainty on $\mu(\boldsymbol{d}|G)$, $\sigma(\boldsymbol{d}|G)$, and $f(\boldsymbol{d}|G)$: Calculate their posterior distributions (or approximate posterior distributions) by considering the posterior distribution of G, similar to the treatment in [53,54]. Based on this, we can then calculate PIs (similar to confidence intervals) for the quantities. Rather than producing a single point estimate of the unknown objective function $f(\boldsymbol{d}|G)$ (which could be quite inaccurate because of interpolation uncertainty), the PI approach produces a range of possible values that takes into account the level of interpolation uncertainty.

In order to derive the PI, first consider the posterior distribution of $\mu(\boldsymbol{d}|G)$. The integral in Eq. (6) can be viewed as a linear transformation of the Gaussian random process G. Consequently, $\mu(\boldsymbol{d}|G)$ is exactly Gaussian. Its posterior distribution is therefore determined by its mean and standard deviation, which we denote by $\mu_\mu(\boldsymbol{d})$ and $\sigma_\mu(\boldsymbol{d})$. Using Eq. (6) and basic results on the mean and variance of linear transformations of random processes [55], we have (see Ref. [53] also)

$$\mu_\mu(\boldsymbol{d}) = E_{G|\boldsymbol{y}^N}[\mu(\boldsymbol{d}|G)|\boldsymbol{y}^N] = E_{G|\boldsymbol{y}^N}\left[\int G(\boldsymbol{d},\boldsymbol{w})p_w(\boldsymbol{w})d\boldsymbol{w}|\boldsymbol{y}^N\right]$$
$$= \int E_{G|\boldsymbol{y}^N}[G(\boldsymbol{d},\boldsymbol{w})|\boldsymbol{y}^N]p_w(\boldsymbol{w})d\boldsymbol{w}$$
$$= \int \hat{y}(\boldsymbol{d},\boldsymbol{w})p_w(\boldsymbol{w})d\boldsymbol{w} \quad (9)$$

$$\sigma^2_\mu(\boldsymbol{d}) = E_{G|\boldsymbol{y}^N}\{[\mu(\boldsymbol{d}|G) - \mu_\mu(\boldsymbol{d})]^2|\boldsymbol{y}^N\}$$
$$= \int\int \text{Cov}[Y(G,\boldsymbol{d},\boldsymbol{w}),Y(G,\boldsymbol{d},\boldsymbol{w}')|\boldsymbol{y}^N]p_w(\boldsymbol{w})d\boldsymbol{w}p_w(\boldsymbol{w}')d\boldsymbol{w}' \quad (10)$$

where $\hat{y}(\boldsymbol{d},\boldsymbol{w})$ and $\text{Cov}[Y(G,\boldsymbol{d},\boldsymbol{w}),Y(G,\boldsymbol{d},\boldsymbol{w}')|\boldsymbol{y}^N]$ are given by Eqs. (4) and (5). In the Appendix we show that a closed-form analytical expression for $\mu_\mu(\boldsymbol{d})$ and $\sigma_\mu(\boldsymbol{d})$ can be obtained from Eqs. (9) and (10) for many common choices of $\boldsymbol{h}(x)$ and $\boldsymbol{r}(x)$. If z_p denotes the $1-p$ percentile of the standard normal distribution, then a $1-p$ PI for $\mu(\boldsymbol{d}|G)$ is

$$\mu(\boldsymbol{d}|G) \in \mu_\mu(\boldsymbol{d}) \pm z_{p/2}\sigma_\mu(\boldsymbol{d}) \quad (11)$$

We refer to $\mu_\mu(\boldsymbol{d})+z_{p/2}\sigma\mu(\boldsymbol{d})$ and $\mu_\mu(\boldsymbol{d})-z_{p/2}\sigma_\mu(\boldsymbol{d})$ as the upper and lower limits of the PI for $\mu(\boldsymbol{d}|G)$. For all examples in the remainder of the paper, we will use $z_{p/2}=2.0$, which corresponds roughly to a 95% PI.

Similarly, Eq. (7) shows that $\sigma^2(\boldsymbol{d}|G)$ is a linear transformation of the random process $(G(\boldsymbol{d},\boldsymbol{w})-\mu(\boldsymbol{d}|G))^2$. Although $\sigma^2(\boldsymbol{d}|G)$ will not be exactly Gaussian (it will have a positive skewness because of the squaring operation), we can still evaluate its mean and variance similar to Eqs. (9) and (10) after determining the mean and covariance functions of the random process $(G(\boldsymbol{d},\boldsymbol{w})-\mu(\boldsymbol{d}|G))^2$. Details of the calculations are provided in the appendix. From the mean and variance of $\sigma^2(\boldsymbol{d}|G)$, the Appendix also describes how to calculate the approximate mean and standard deviation of $\sigma(\boldsymbol{d}|G)$, which we denote by $\mu_\sigma(\boldsymbol{d})$ and $\sigma_\sigma(\boldsymbol{d})$, respectively. An approximate $1-p$ PI for $\sigma(\boldsymbol{d}|G)$ is then given by

$$\sigma(\boldsymbol{d}|G) \in \mu_\sigma(\boldsymbol{d}) \pm z_{p/2}\sigma_\sigma(\boldsymbol{d}) \quad (12)$$

An approximate PI on the objective function $f(\boldsymbol{d}|G)=\mu(\boldsymbol{d}|G)+c\sigma(\boldsymbol{d}|G)$ can be obtained as follows. The Gaussian approximation used in the Appendix for $\sigma(\boldsymbol{d}|G)$ leads naturally to an approximate Gaussian posterior distribution for $f(\boldsymbol{d}|G)$. Its mean and standard deviation (denoted by $\mu_f(\boldsymbol{d})$ and $\sigma_f(\boldsymbol{d})$) can be readily obtained from those of $\mu(\boldsymbol{d}|G)$ and $\sigma(\boldsymbol{d}|G)$ in Eqs. (9) and (10), (A1), and (A2) and the covariance between $\mu(\boldsymbol{d}|G)$ and $\sigma(\boldsymbol{d}|G)$. Specifically

$$\mu_f(\boldsymbol{d}) = \mu_\mu(\boldsymbol{d}) \pm c\mu_\sigma(\boldsymbol{d}) \quad (13)$$

and

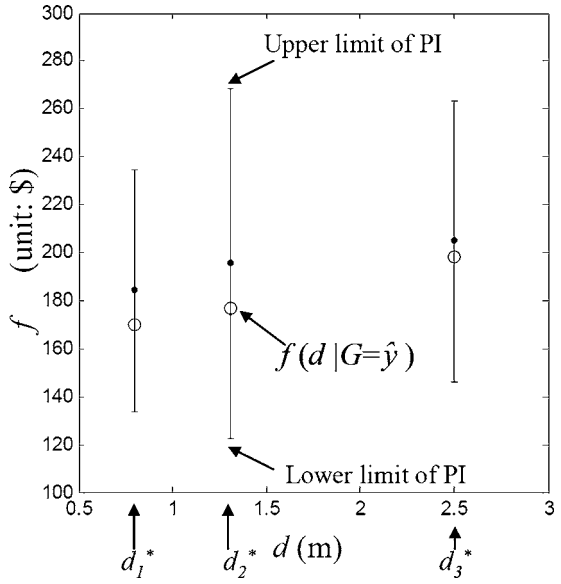

Fig. 3 Prediction intervals for the objective function values for three candidate designs, d_1^*, d_2^*, and d_3^*, based on the four simulation outputs in Fig. 1(*b*). The dot at the center of each bar indicated the center $\mu_f(\boldsymbol{d})$ of the prediction interval, and the circles indicate the objective function values obtained from the fitted model $\hat{y}$.

$$\sigma_f^2(\boldsymbol{d}) = \sigma_\mu^2(\boldsymbol{d}) + c^2\sigma_\sigma^2(\boldsymbol{d}) + 2c\mathrm{Cov}[\mu(\boldsymbol{d}|G), \sigma(\boldsymbol{d}|G)] \quad (14)$$

The Appendix describes how to calculate a closed-form analytical expression for $\mathrm{Cov}(\mu(\boldsymbol{d}|G), \sigma(\boldsymbol{d}|G))$ based on the assumption that $\mu(\boldsymbol{d}|G)$ and $\sigma(\boldsymbol{d}|G)$ are approximately jointly Gaussian. An approximate $1-p$ PI for $f(\boldsymbol{d}|G)$ becomes

$$f(\boldsymbol{d}|G) \in \mu_f(\boldsymbol{d}) \pm z_p\sigma_f(\boldsymbol{d}) \quad (15)$$

which provides a direct measure of the effects of interpolation uncertainty on the objective function.

To illustrate the use of these PIs for understanding the effects of interpolation uncertainty, we continue the container design example. After simulating the response at the four sampling sites shown in Fig. 1(*b*), fitting the metamodel $\hat{y}$ without considering interpolation uncertainty leads to selecting $d_1^*=0.80$ as the optimal design (see Fig. 1(*d*)) and a corresponding objective function of $f(d_1^*|G=\hat{y})=170$. Considering d_1^* as a potentially optimal design, we would naturally be interested in the extent to which interpolation uncertainty will affect the objective function at d_1^*. The (95%) PI obtained from Eq. (15) for this design is $f(d_1^*|G) \in [134, 233]$. The interpretation is that the true objective value may actually be as small as 134 or as large as 233, even though $f(d_1^*|G=\hat{y})=170$. This level of uncertainty is consistent with the values of $f(d_1^*|G)$ for the three realizations of G shown in Figs. 2(*d*)–2(*f*).

In order to have a better idea of how "robust" the current optimal design d_1^* is with respect to interpolation uncertainty, we can compare its PI with that of other candidate optimal designs. For example, suppose $d_2^*=1.31$ and $d_3^*=2.50$ are two additional candidate designs under consideration. Figure 3 shows the PIs at all three candidate designs. Because there is significant overlap between the PIs for $f(d_1^*|G)$, $f(d_2^*|G)$, and $f(d_3^*|G)$, none of the three candidate designs emerges as the clear winner or loser in the presence of the interpolation uncertainty. Even though $f(d|G=\hat{y})$ (values for which are shown as open circles in Fig. 3) is the smallest for d_1^*, the overlapping PIs imply that based on our current information (only four samples), any of the three designs could perhaps have the smallest true objective function value $f(d|G=G_{\mathrm{true}})$. One conclusion is that if we want to guarantee that our design is close to optimal, we must conduct simulations at additional sampling sites to reduce interpolation uncertainty. We elaborate on this in the following section. Note that the PIs shown in Figs. 3–6 are not necessarily symmetric about $f(d|G=\hat{y})$, which will often differ from the center $\mu_f(d)$ of the PIs.

5 Guidelines for Terminating Versus Continuing Simulation

One observation from Fig. 3 is that when relatively few simulation runs ($N=4$ in this example) have been conducted, the interpolation uncertainty is high, which is reflected in the wide PIs in Fig. 3. The only way to reduce the interpolation uncertainty is to run the simulations at additional sampling sites. This raises the question of how do we know when interpolation uncertainty is

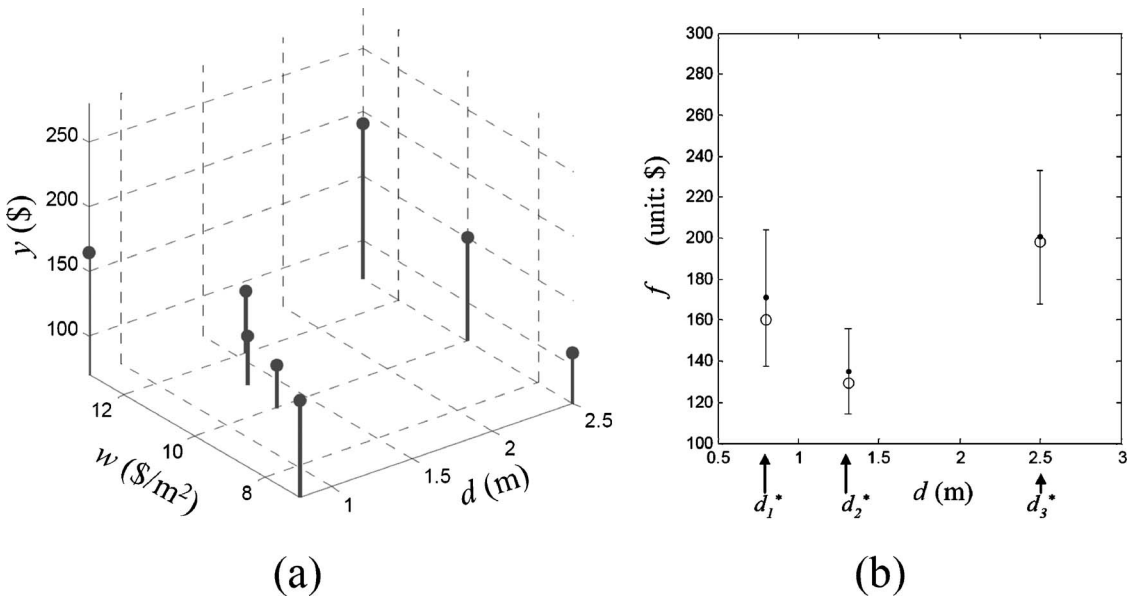

Fig. 4 (*a*) Simulation output at the first four sampling sites plus four additional sites. (*b*) Updated prediction intervals at the three candidate designs based on all eight simulation outputs in (*a*).

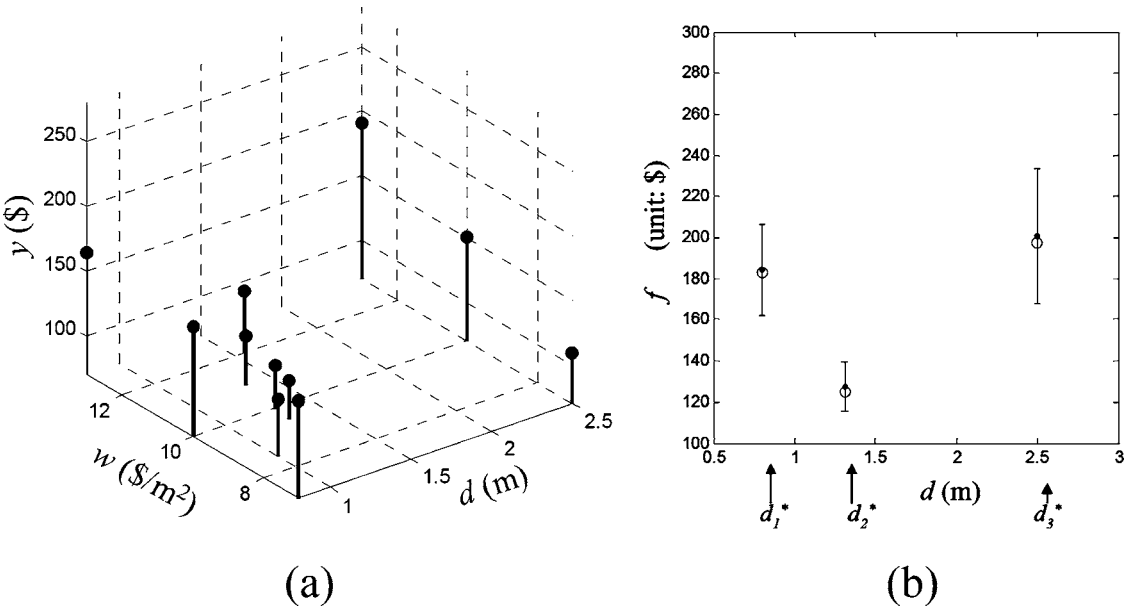

Fig. 5 (*a*) Three more simulation outputs added to Fig. 4(*a*). (*b*) Updated prediction intervals at the three candidate designs based on all the eleven outputs in (*a*).

small enough that we can terminate simulation and finalize the design? The PIs that we derived in Sec. 4 provide a criterion for helping to decide when to terminate. To illustrate this, we continue the container design example with the discrete set of candidate designs (d_1^*, d_2^*, and d_3^*) shown in Fig. 3. More general situations, in which the feasible design space is continuous, will be discussed in later sections.

Because the PIs overlap in Fig. 3, we clearly must conduct further computer experimentation to reduce interpolation uncertainty. The response values at four additional sampling sites, chosen in a somewhat ad hoc manner, are shown in Fig. 4(*a*). Based on the entire set of eight simulation runs, we updated the posterior mean and covariance of $G(\cdot)$ from Eqs. (4) and (5) and then calculated new PIs using Eq. (15). The updated PIs for the three candidate designs are shown in Fig. 4(*b*). Comparing this with Fig. 3, it is clear that interpolation uncertainty has been reduced substantially, especially at d_2^* and d_3^*. Because the lower limit of the PI for $f(d_3^*|G)$ now falls above the upper limit of the PI for $f(d_2^*|G)$, we can conclusively rule out d_3^* as inferior to d_2^* (with 95% confidence). Note that after updating our knowledge of the response surface from the additional sampling sites, the fitted model $\hat{y}$ (not shown) has also been updated to better reflect the true response surface. This is somewhat evident from the fact that $f(d_2^*|G=\hat{y})$ is now smaller than $f(d_1^*|G=\hat{y})$, just as it is for $G=G_{\text{true}}$ (see Fig. 1(*c*)). Because the PIs at d_1^* and d_2^* still overlap, however, interpolation uncertainty is still too large to confidently conclude that d_2^* is superior to d_1^*. Consequently, although we have gotten closer to the final solution, we still need to conduct additional simulation to further reduce interpolation uncertainty in the vicinity of d_1^* and d_2^*.

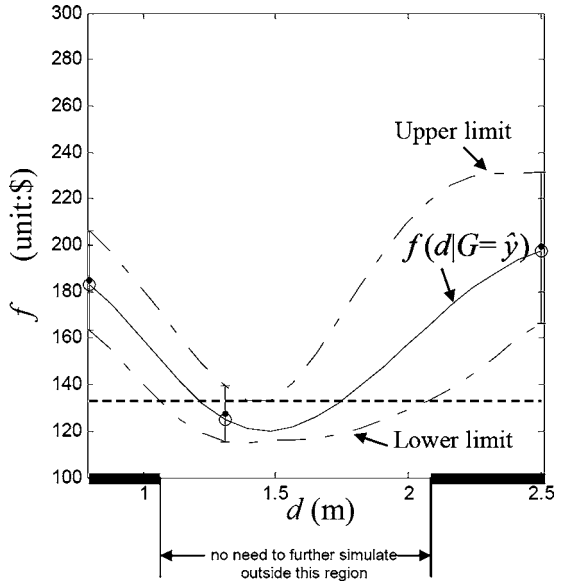

Fig. 6 Prediction interval plot based on the eleven simulation outputs in Fig. 5(*a*)

For these purposes, three additional sampling sites were added, and the simulation results at the entire set of 11 sampling sites are shown in Fig. 5(*a*). We have not added any additional sampling sites in the vicinity of d_3^*, because this design has already been ruled out as inferior. After updating the posterior distribution of G based on the 11 sampling sites, the PI for $f(d_2^*|G)$ is now clearly separated from those of the other candidates, as shown in Fig. 5(*b*). Specifically, the upper limit of the PI for $f(\boldsymbol{d}|G)$ at d_2^* falls well below the lower limits of the PIs at d_1^* or d_3^*. Consequently, if these are the only three design candidates, we can terminate simulation and conclude that d_2^* is the optimal one (with 95% confidence).

We reiterate the important point that it is impossible to eliminate the effects of interpolation uncertainty via validation runs at even the single final selected design, because this would require a large number of simulation runs at sampling sites spread densely over the entire noise variable domain. Consequently, the PIs derived in this section are critical for understanding the effects of interpolation uncertainty and appropriately assessing whether additional sampling and computer simulation is necessary.

6 Prediction Interval Plots for Guiding the Simulation

The graphical methods illustrated in the previous section are primarily relevant for simple problems with only a relatively small number of design possibilities. In this section, we extend the

PI concepts to more general problems with continuous design spaces. For situations in which one decides that additional simulation is necessary, we also discuss how to efficiently guide the selection of the additional sampling sites based on the previous simulation results and the obtained PIs.

We illustrate the approach with the same container design example, except that the feasible design space is now $0.8 \leq d \leq 2.5$, as opposed to the three candidate values considered in the previous section. Suppose we have run simulations at the same eleven sites shown in Fig. 5(a). Figure 6 shows the PIs analogous to those shown in Fig. 5(b), except that we have calculated and plotted them for the entire range of feasible design values. We refer to this as a PI plot. Note that the three PIs shown in Fig. 5(b) have been added to Fig. 6 to illustrate that they are consistent.

Figure 6 implies that we need not further consider design values larger than $d=2.08$ or small than $d=1.06$ (the darkened region of the d axis), because they have essentially been ruled out as nonoptimal. Specifically, the upper boundary $\mu_f(d)+2\sigma_f(d)$ of the PI on the robust design objective function f is minimized at $d \approx 1.45$, and the minimum value is roughly $f=133$. Consequently, if we select $d=1.45$ as the design, we are guaranteed (with 95% confidence) that the objective function will be no larger than 133. Extending this value across the plot as the dashed horizontal line, we see that the *lower* limit $\mu_f(d)-2\sigma_f(d)$ of the PI is always larger than 133 for $d>2.08$ and $d<1.06$. Therefore, we are guaranteed (with 95% confidence) that any design with $d>2.08$ or $d<1.06$ will have an objective function larger than 133 and, therefore, larger than for the design $d=1.45$. We conclude that we can rule out the entire region of the design space outside the interval [1.06, 2.08] when selecting additional sampling sites. This will result in a much more efficient simulation strategy than if we cover the entire input space with uniformly dense sampling sites.

Note that the same concept was used to rule out d_3^* in the simple discrete design situation shown in Fig. 4(b). For the continuous design situation depicted in Fig. 6, we can say that the PI at $d=1.45$ does not overlap with the PI for any d outside the interval [1.06, 2.08].

Note that the PI width in Fig. 6 is not reduced to zero at any of the sampled design sites. This represents a fundamental difference between deterministic design and design under uncertainty. The PIs shown in Fig. 6 are PIs on the objective function $f(\boldsymbol{d}|G)=\mu(\boldsymbol{d}|G)+c\sigma(\boldsymbol{d}|G)$, as opposed to PIs on the response $Y(G,\boldsymbol{d},\boldsymbol{w})$. Because $\mu(\boldsymbol{d}|G)$ and $\sigma(\boldsymbol{d}|G)$ represent the response mean and variance with respect to variation in $\boldsymbol{W}$, they depend on the true response surface over a *range* of values of $\boldsymbol{W}$. Hence, we cannot eliminate uncertainty in $f(\boldsymbol{d}|G)$ by conducting a simulation experiment at a single site $\boldsymbol{x}=[\boldsymbol{d}^T,\boldsymbol{w}^T]^T$. The only way to completely eliminate uncertainty in $f(\boldsymbol{d}|G)$ at any design site $\boldsymbol{d}$ is to conduct an exhaustive set of simulations over the entire set of $\{\boldsymbol{d},\boldsymbol{w}\}$ values as $\boldsymbol{w}$ varies over its distribution range. In contrast, for deterministic optimization problems, we can always eliminate the effects of interpolation uncertainty at any single design site $\boldsymbol{d}$ (for example, the final chosen design) by conducting a single simulation run at the site $\boldsymbol{x}=\boldsymbol{d}$.

7 Case Study: Robust Design of an Engine Piston

To better demonstrate the general application of the proposed PI methodology, in this section we consider the automotive piston design case study previously analyzed in Ref. [56]. The goal is to design an engine piston in order to robustly minimize piston slap noise. Engine noise is one of the key contributors to customer dissatisfaction with vehicles. Piston slap noise is the engine noise resulting from piston secondary motion, which can be simulated using computationally intensive computer experiments (multibody dynamics). The response variable y is the sound power level of the piston slap noise. Previous results [56] indicate that the skirt profile (SP) and pin offset (PO) are two piston geometry design variables that have major impact on the response. In addition, the piston-to-bore clearance (CL) is an uncontrollable noise factor that also strongly affects the response. Therefore, we will consider these two design variables $\boldsymbol{d}=[\mathrm{SP\ PO}]^{\mathrm{T}}$ and a single noise variable $W=\mathrm{CL}$ with a feasible design space of $\{\mathrm{SP} \in [1,3]$ (dimensionless), $\mathrm{PO} \in [0.5,1.3]$ (mm)$\}$ and CL assumed to follow a Gaussian distribution with mean 50 (μm) and standard deviation 11 (μm). We also consider the objective function $f(\mathrm{SP},\mathrm{PO}|G)=\mu(\mathrm{SP},\mathrm{PO}|G)+3\sigma(\mathrm{SP},\mathrm{PO}|G)$ (Eq. (8) with $c=3$), where $\mu(\mathrm{SP},\mathrm{PO}|G)$ and $\sigma(\mathrm{SP},\mathrm{PO}|G)$ are as defined in Eqs. (6) and (7). The true response surface G represents the output (sound power level) of the computationally intensive simulation code as a function of the three input variables $\{\boldsymbol{d},w\}=\{\mathrm{SP},\mathrm{PO},\mathrm{CL}\}$.

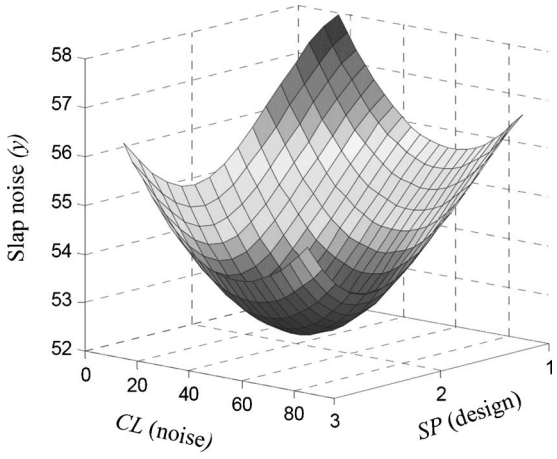

Fig. 7 True response surface for the piston design example

Because of the computational expense, the true response surface is of course unknown. The design process illustrated in the subsequent paragraphs is based on information obtained from a limited number of sampling sites, under consideration of the interpolation uncertainty in metamodel. In order to gauge the effectiveness of the approach, however, we constructed what will serve as our "true" response surface as follows. A large set of 200 sampling sites were used to fit an elaborate, relatively accurate kriging metamodel passing through all 200 simulated output responses. The much smaller set of sampled response values used to guide the design process in the following example were then obtained by sampling from this true response surface. The true response surface is plotted in Fig. 7 as a function of the noise variable (CL) and the single design variable (SP). For the purpose of illustrating the surface with a 3D plot, the second design variable (PO) was held fixed at 1.3 in Fig. 7. It should be noted that the true response surface is only used for generating response values at sampling sites and for gauging the effectiveness of the PI approach. Otherwise, we treat it as completely unknown.

Figure 7 illustrates that the response is highly nonlinear in CL and SP. In the first stage of the example, the response values at ten sampling sites were generated and used to find the posterior response mean and covariance functions from Eqs. (4) and (5). The prediction intervals for the objective function were then calculated as previously described and were plotted as a function of the two design variables in the 3D PI plots shown in Fig. 8(a). The wide gaps between the upper and lower limits of the PIs clearly indicate that additional simulation is needed to distinguish good designs from unacceptable ones. Hence, eight additional sites were sampled (for a total of 18 sites), and the resulting PI plots are shown in Fig. 8(b). The number of additional sampling sites was chosen rather arbitrarily. The upper and lower PI limits in Fig. 8(b) are substantially tighter than in Fig. 8(a) and are now tight enough to rule out some clearly inferior designs using the same reasoning discussed in Sec. 6. For example, designs with PO close to 0.5 can be ruled out with approximately 95% confidence (the

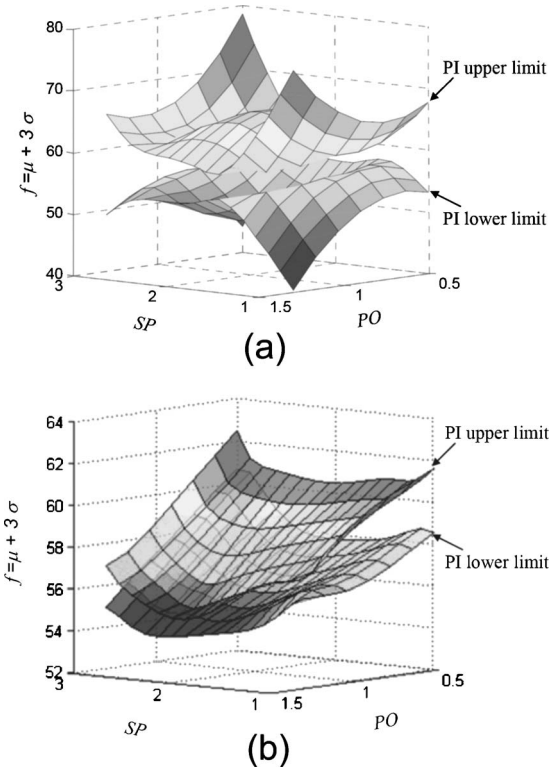

Fig. 8 PI plots of robust design objective after: (*a*) 10 initial simulation runs and (*b*) 18 simulation runs

lower limits of their PIs fall above the PI upper limits for many other designs). Hence, all further simulation can focus on the more promising designs. Note that as more and more sampling sites are added, the upper and lower PI limits will become tighter and tighter (Fig. 9 shows a one-dimensional counterpart after 27 sampling sites).

This can be more clearly visualized if we consider the design variable PO to be discrete and construct three separate 2D plots for the three discrete values PO={0.5,0.9,1.3}, as shown in Fig. 10. Note that the PI plots in Fig. 10 are precisely those in Fig. 8(*b*) for the three discrete values of PO. From Fig. 10 we can see that a portion of design region with PO=1.3, as well as the entire design region with PO=0.9 or 0.5, can be ruled out as inferior. If we are only considering these three values for PO, the potentially optimal design region then reduces to the much simpler 1D segment shown in Fig. 10(*a*). The dashed curves in Fig. 10 represent

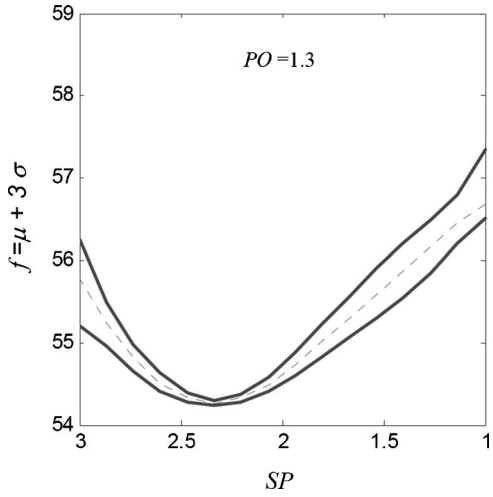

Fig. 9 PI plot after 27 simulation runs. The robust design solution [$PO^*=1.3$, $SP^*=2.33$] has an objective value within the PI of [54.24, 54.29].

the true unknown objective function $f(\mathrm{SP},\mathrm{PO}|G=G_{\mathrm{true}})$, with G_{true} being that plotted in Fig. 7. Note that the PIs shown in Fig. 10 all contain $f(\mathrm{SP},\mathrm{PO}|G=G_{\mathrm{true}})$, although the true objective function is sometimes closer to the boundary of the PI than to the center. This is all that we can expect, of course, given the definition of a prediction interval.

Ruling out the inferior design regions (the shaded regions in Fig. 10) and conducting nine additional simulations in the still-potentially optimal design region $\{1.5\leqslant \mathrm{SP}\leqslant 2.9,\ \mathrm{PO}=1.3\}$, we arrive at the PI plot shown in Fig. 9. Figure 9 indicates that we can now narrow the optimal design down to a small neighborhood around $\boldsymbol{d}^*=(\mathrm{PO}^*,\mathrm{SP}^*)=(1.3,2.333)$. More specifically, if we terminate simulation at this point and choose $\boldsymbol{d}^*$ as the optimal design, we are guaranteed (with 95% confidence) that the true objective function $f(\mathrm{SP}^*,\mathrm{PO}^*|G_{\mathrm{true}})$ at the chosen design will fall within its PI range [54.24, 54.29]. If we decide that this range is negligibly small, we would adopt $\boldsymbol{d}^*=(1.3,2.333)$ as the final design. For comparison, the optimal design based on the true response surface G_{true} is (PO=1.3, SP=2.334).

8 Discussion

8.1 Selecting $\boldsymbol{\phi}$ and $\boldsymbol{\alpha}$. The prior distribution parameters $\boldsymbol{\phi}$ and α play an important role in quantifying the posterior uncertainty in the response surface G and therefore the uncertainty in

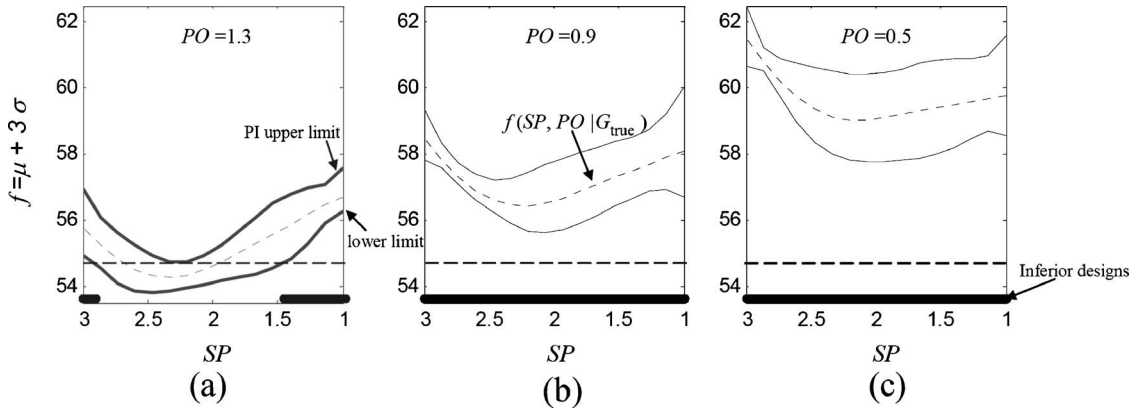

Fig. 10 Three 2D cross sections of the PI plot in Fig. 8(*b*). All dark shaded regions of the design space can be ruled out (with 95% confidence) as inferior to the unshaded design region shown in panel (*a*).

the objective function $f(\boldsymbol{d}|G)$. In particular, $\boldsymbol{\phi}$ and α influence the width of the PI. Because prior knowledge of appropriate values for these parameters would generally be limited, we recommend estimating them directly from the sampling data. One method would be to use a full Bayesian approach in which we assume hyperprior distributions for the prior parameters $\boldsymbol{\phi}$, α, and $\boldsymbol{\beta}$ and attempt to calculate the full posterior distribution $G(\boldsymbol{d},\boldsymbol{w})$ marginalized over the posterior distribution of $\boldsymbol{\phi}$, α, and $\boldsymbol{\beta}$. Although this would provide the most accurate representation of interpolation uncertainty, the full posterior distribution of G would be far too complex to calculate analytically for even the simplest choice of hyperpriors. The only recourse would be to use Monte Carlo methods to evaluate the posterior distribution of G. Using Monte Carlo simulation, Martin and Simpson [57] and Martin [58] demonstrated that for some typical situations the posterior distribution of G appears close to a t distribution and that Eq. (5) underestimates the posterior variance of G.

Although Monte Carlo methods are often feasible, for reasons discussed in Sec. 8.2 they are far too computationally intensive to be used for the problem addressed in this paper. In light of this, we recommend the same approximation that has been used in the vast majority of prior work on Bayesian analysis of computer experiments (e.g, Refs. [27,28,43]). The approximation, which is sometimes referred to as an empirical Bayes approach [59], is to first calculate suitable estimates of the prior parameters $\boldsymbol{\phi}$, α, and $\boldsymbol{\beta}$ and then simply treat the estimates as the true values in all subsequent Bayesian computations. One might use maximum likelihood estimates (MLEs) or assume hyperpriors for $\boldsymbol{\phi}$, α, and $\boldsymbol{\beta}$ and then use their posterior means or modes as the estimates. Martin and Simpson [57] recommended using the posterior means rather than the MLEs or posterior modes, because $\boldsymbol{\phi}$ and α often have skewed distributions. The MLEs are simpler because they can be readily obtained by maximizing the log-likelihood function, which is equivalent to minimizing the function $(\mathbf{y}_N - \boldsymbol{H}\boldsymbol{\beta})^T\Sigma^{-1}(\mathbf{y}_N - \boldsymbol{H}\boldsymbol{\beta}) - \log[\det(\Sigma)]$, where the ith-row, jth-column element of the $N \times N$ matrix Σ is $\alpha^2 R_{\boldsymbol{\phi}}(\boldsymbol{s}_i,\boldsymbol{s}_j)$, and all other quantities are as defined in Sec. 3. Because the posterior mean $\hat{\beta}$ defined in Sec. 3 cannot be calculated until we have values for $\boldsymbol{\phi}$ and α available, some other suitable estimate for $\boldsymbol{\beta}$ should be substituted into the log-likelihood function. For example, one may use the MLE of $\boldsymbol{\beta}$, which is equivalent to jointly maximizing the log likelihood over $\boldsymbol{\phi}$, α, and $\boldsymbol{\beta}$.

8.2 Monte Carlo Simulation Versus Analytical Expressions. As discussed in Sec 8.1, in order to obtain estimates of the prior parameters $\boldsymbol{\phi}$, α, and $\boldsymbol{\beta}$, one approach is to calculate the posterior distribution of $\boldsymbol{\phi}$, α, and $\boldsymbol{\beta}$ and then use the posterior mean or mode as estimates. Although for certain assumed prior distributions, an analytical expression for the posterior distribution might be tractable, one would often be forced to use Monte Carlo simulation to estimate the posterior distribution. Markov Chain Monte Carlo (MCMC) methods are particularly useful in this regard [59]. Martin and Simpson [57] used the Metropolis-Hastings MCMC method for such purposes.

For estimating the *prior parameters*, the computational burden of Monte Carlo methods is quite reasonable. In contrast, for estimating the *prediction intervals* for constructing PI plots such as those shown in Fig. 8, Monte Carlo methods (either Markov chain or conventional) would involve prohibitively large computational expense. In order to illustrate the reason, consider how one might attempt to use Monte Carlo simulation to estimate the posterior mean $\mu_f(\boldsymbol{d})$ and standard deviation $\sigma_f(\boldsymbol{d})$ for use in the PI of Eq. (15): Consider the most optimistic scenario in which one already has available an expression for the posterior distribution of G, given the response values at a set of sampling sites. One would first discretize the entire (perhaps high dimensional) input $\{\boldsymbol{d},\boldsymbol{w}\}$ variable space. For each replicate of the Monte Carlo simulation, one would implement the following two steps: Step (1) Generate one realization of the discretized random surface G over the entire space of discretized input values (similar to those shown in Figs. 2(*a*)–2(*c*). Step (2) For each value of $\boldsymbol{d}$ in the discretized design variable space, numerically integrate (using Eqs. (6) and (7)) and the realization of G from step (1) over the entire discretized noise variable space to obtain one realization of $\mu(\boldsymbol{d}|G)$, $\sigma(\boldsymbol{d}|G)$, and $f(\boldsymbol{d}|G)$. Steps (1) and (2) would then be repeated for some large number of Monte Carlo replicates (e.g., 10,000 or 100,000), each one corresponding to a different realization of G, to obtain the posterior distribution of $f(\boldsymbol{d}|G)$ empirically. From this, we could estimate $\mu_f(\boldsymbol{d})$ and $\sigma_f(\boldsymbol{d})$ and calculate a PI on $f(\boldsymbol{d}|G)$ for each value of $\boldsymbol{d}$.

To better appreciate the computational challenges, consider step (1), and suppose we had ten input variables with the range for each one discretized into 20 points (a relatively course discretization). Because the entire input space is discretized into 20^{10} points, in order to generate one realization of G in step (1), we would have to generate 20^{10} random variables. Moreover, because the 20^{10} random variables are not independent of each other, we would first have to calculate their $20^{10} \times 1$ mean vector and $20^{10} \times 20^{10}$ covariance matrix using Eqs. (4) and (5) Clearly, this is computationally prohibitive, especially when one considers that it must be repeated for a large number of Monte Carlo replicates. Note that this is how the realizations of G shown in Fig 2 were generated, albeit for a much lower-dimensional problem.

One might consider replacing the numerical integration in step (2) with an "inner" Monte Carlo simulation loop in which we generate a large number of observations of $\boldsymbol{W}$ and empirically estimate $\mu(\boldsymbol{d}|G)$ and $\sigma(\boldsymbol{d}|G)$ for each specific realization of G from step (1). Oakley and O'Hagan [54] discuss such an approach, as well as some strategies for improving the computational efficiency. For high-dimensional $\boldsymbol{w}$, this would typically be more computationally efficient than the numerical integration. However, the result would be a nested Monte Carlo simulation in which the inner simulation must be repeated for each value of $\boldsymbol{d}$ in the design variable space, as well as for each replicate of the outer simulation loop (one replicate corresponding to one realization of G). Even with the computational improvements suggested in Ref. [54], the nested Monte Carlo simulation is still far too computationally expensive for high-dimensional design spaces.

If the numerical integration is replaced by a Monte Carlo simulation in the nested simulation scenario described above, one might also consider combining the inner and outer simulations. This, however, is not possible given the nature of our prediction interval formulation. By design, the prediction interval distinguishes the effects of noise variable uncertainty (inner simulation) and interpolation uncertainty (outer simulation).

For these reasons, Monte Carlo methods to estimate the prediction intervals are impractical, and analytical expressions become more important. An additional advantage of analytical expressions is that they are better suited for design optimization algorithms. For example, we can differentiate the expressions to facilitate a gradient search procedure. Monte Carlo sampling methods to evaluate design objective functions often result in jitter and slow (or non) convergence.

9 Conclusions

On two parallel fronts, robust design under uncertainty and design optimization using computer experiments have both received a great deal of attention recently. When the two are combined, however, the "optimal" robust design solution may not be as robust as it ought to be, because of the effects of metamodel interpolation uncertainty. Moreover, the interpolation uncertainty cannot be eliminated with a final validation run, as it can in deterministic optimization problems. We have proposed a Bayesian methodology with GRPs to represent the response surface in order to clearly reveal the impact of interpolation uncertainty and to directly quantify its effect on the robust design objective func-

tion. This is particularly important for robust design using metamodels fitted to the output of computer experiments, because the conventional approach of estimating prediction uncertainty due to the random error present in physical experiments is completely invalid with computer experiments. The Bayesian prediction intervals that we have derived can be used by designers as a simple intuitive tool to quantify the effects of interpolation uncertainty, evaluate and compare robust design candidates, and guide a more efficient simulation. One unique contribution is that our closed-form expressions for the prediction intervals obviate the need for numerical integration and/or Monte Carlo simulation, which we have argued is computationally infeasible in all but the simplest design problems, due to the curse of dimensionality.

Because we have only considered interpolation uncertainty, one might view our treatment of model uncertainty as being relative to a base line that assumes the computer simulation is a perfect representation of reality. The approach is intended to ensure that the chosen design is close to what one would choose as optimal in the hypothetical scenario in which one were able to run an infinite number of simulations. Without an accurate assessment of the error between reality and the output of the simulation code, it is obviously impossible to ensure that the chosen design will be even close to optimal when implemented in the physical world. A proper treatment of this would involve a great many factors specific to the application and the code and would naturally require some form of physical experimentation. If such were available, however, we believe that the Bayesian methodology that we have utilized can be conveniently extended to encompass these other forms of uncertainty. Bayesian statistical methods are often well suited for analyzing data obtained via physical experiments. It is typically straightforward to integrate uncertainty from multiple levels in a hierarchical Bayesian analysis, with the posteriors from one level becoming the priors at the next level. For guidelines on defining and quantitatively representing various forms of uncertainty in computer experimentation and on their integrated analysis within a Bayesian framework, we refer the reader to the comprehensive discussion in Kennedy and O'Hagan [43].

It should be noted that although a low-dimensional design example was used to illustrate the approach throughout the paper, the closed-form expressions for the PIs and the concept of using PIs to guide computer simulations and to distinguish the best design alternative are applicable to high-dimensional problems. The primary challenge would lie in projecting the high-dimensional PI information down to meaningful three-dimensional PI plots. Precisely how to do this is the subject of ongoing work.

Furthermore, we have only treated unconstrained optimization, for which we only need to consider uncertainty in the objective response. Typical design problems also have (i) constraint responses that are functions of the inputs and that must be satisfied with a specified probability (when they are functionally dependent on uncertain noise variables); and (ii) design variables that are themselves uncertain, due perhaps to manufacturing variation. We believe that our approach can be readily extended to both of these scenarios. Uncertain design variables can be treated simply by decomposing each uncertain design variables into two components: A design component that represents the desired value for the design variable, and a random noise component that represents the difference between the actual design variable and its specified desired value. Regarding responses that must satisfy a constraint with some specified probability: Constraint responses that are normally viewed as random because they depend on the random $\boldsymbol{W}$ can simply be viewed as random because of their dependence on both $\boldsymbol{W}$ and G. When selecting a design that ensures the constraint response is satisfied with a specified probability, we would consider the distribution of $\boldsymbol{W}$ and the posterior distribution of G given the samples.

Throughout the paper, the implication has been that one can use the prediction intervals to help ensure that the effect of interpolation uncertainty is minimized, to the extent that it does not alter the choice of optimal design. This would often involve running additional simulations in stages, guidelines and examples for which were presented in Secs. 5–7. If the effects of interpolation uncertainty are not negligible, but no additional simulation is possible because of limited resources, then one would hope the final design is robust to interpolation uncertainty, much like it is robust to noise uncertainty. One simple strategy for achieving this is to use the upper boundary of the prediction interval on $f(\boldsymbol{d}|G)$ (e.g., the upper curve in Fig. 6) as the objective function in optimization. The upper boundary of the prediction interval represents what we might view as the worst-case value (at the specified confidence level for the prediction interval) of $\mu(\boldsymbol{d}|G)+c\sigma(\boldsymbol{d}|G)$, under direct consideration of interpolation uncertainty. As future research, we are also exploring methods of considering interpolation uncertainty on par with noise uncertainty for decision making in robust design.

Acknowledgment

This work was supported in part by NSF Grants No. DMI 0354824 and No. DMI 0522662.

Nomenclature

$\boldsymbol{d}$ = vector of design variables
$\boldsymbol{W}$ = vector of random noise variables
$\boldsymbol{w}$ = a specific value or realization of $\boldsymbol{W}$
$\boldsymbol{x}=\{\boldsymbol{d},\boldsymbol{w}\}$ = vector of design and noise variables together
$y(\boldsymbol{x})$ = response surface as a function of $\boldsymbol{x}$
$Y(\boldsymbol{x})$ = response surface viewed as random
N = number of simulation runs conducted
$\boldsymbol{S}^N$ = set of N sampled sites
$\boldsymbol{y}^N$ = vector of N simulated response values corresponding to $\mathbf{S}^N$
$\hat{y}(\boldsymbol{x})$ = best Bayesian prediction of $Y(\boldsymbol{x})$
$p_w(w)$ = pdf of $\boldsymbol{W}$ evaluated at $\boldsymbol{w}$
$G(\boldsymbol{x})$ = Gaussian random process (GRP) to represent uncertainty in the response surface
$\boldsymbol{h}^T(\boldsymbol{x})\boldsymbol{\beta}$ = prior mean function of $G(\boldsymbol{x})$
$\alpha^2 R_\phi(\boldsymbol{x},\boldsymbol{x})$ = prior covariance function of $G(\boldsymbol{x})$
$\mu(\boldsymbol{d}|G)$ = mean of $Y(\boldsymbol{d},\boldsymbol{W})$ with respect to $\boldsymbol{W}$, given a specific realization of G
$\sigma^2(\boldsymbol{d}|G)$ = variance of $Y(\boldsymbol{d},\boldsymbol{W})$ with respect to $\boldsymbol{W}$, given a specific realization of G
$f(\boldsymbol{d}|G)$ = robust design objective function (e.g., $\mu(\boldsymbol{d}|G)+c\sigma(\boldsymbol{d}|G)$)
$\mu_\mu(\boldsymbol{d})$ = mean of $\mu(\boldsymbol{d}|G)$ with respect to the random process G
$\sigma_\mu(\boldsymbol{d})$ = standard deviation of $\mu(\boldsymbol{d}|G)$ with respect to the random process G
$\mu_\sigma(\boldsymbol{d})$ = mean of $\sigma(\boldsymbol{d}|G)$ with respect to the random process G
$\sigma_\sigma(\boldsymbol{d})$ = standard deviation of $\sigma(\boldsymbol{d}|G)$ with respect to the random process G
$\mu_f(\boldsymbol{d})$ = mean of $f(\boldsymbol{d}|G)$ with respect to the random process G
$\sigma_f(\boldsymbol{d})$ = standard deviation of $f(\boldsymbol{d}|G)$ with respect to the random process G
PI = prediction interval

Appendix: Derivation of the PI for the Robust Design Objective Function

As discussed in Sec. 4, to obtain the PI for the objective function $f(\boldsymbol{d}|G)=\mu(\boldsymbol{d}|G)+c\sigma(\boldsymbol{d}|G)$, we must calculate $\mu_\mu(\boldsymbol{d})$, $\sigma^2_\mu(\boldsymbol{d})$, $\mu_\sigma(\boldsymbol{d})$, $\sigma^2_\sigma(\boldsymbol{d})$, and $\mathrm{Cov}(\mu(\boldsymbol{d}|G),\sigma(\boldsymbol{d}|G))$. The mean and variance of $\mu(\boldsymbol{d}|G)$ are given by Eqs. (9) and (10). For Gaussian noise, a

polynomial $\boldsymbol{h}(\boldsymbol{x})$, and the form of $\boldsymbol{r}(x)$ given by Eq. (3), the terms $\hat{y}(\boldsymbol{d},\boldsymbol{w})$, $\mathrm{Cov}[Y(G,\boldsymbol{d},\boldsymbol{w}),Y(G,\boldsymbol{d},\boldsymbol{w}')|\boldsymbol{y}^N]$ and $p_{\boldsymbol{w}}(\boldsymbol{w})$ can all be written in terms of exponentials of quadratic forms of $\boldsymbol{w}$, multiplied by polynomials in $\boldsymbol{w}$. Consequently, Eqs. (9) and (10) can be analytically integrated to obtain closed-form expressions, instead of evaluating them using numerical integration or Monte Carlo simulation. For example, suppose we have a single sampling site and $\boldsymbol{h}(\boldsymbol{x})=1$ and $\boldsymbol{r}(\mathbf{x})=\exp[-\phi_d(d-d_1)^2-\phi_w(w-w_1)^2]$ In this case, $\hat{y}(\boldsymbol{d},\boldsymbol{w})=\hat{\boldsymbol{\beta}}+\exp[-\phi_d(d-d_1)^2-\phi_W(w-w_1)^2]\boldsymbol{R}^{-1}(\boldsymbol{y}^1-\hat{\boldsymbol{\beta}})$, where only $\boldsymbol{r}^T(\boldsymbol{x})$ is a function of $\boldsymbol{w}$, while $\hat{\boldsymbol{\beta}}$, $\boldsymbol{R}^{-1}$ and $\boldsymbol{y}^N$ are constant with respect to $\boldsymbol{w}$. $\boldsymbol{y}^1$ is the simulation output at sampling site (d_1,w_1) and $[\phi_d,\phi_w]$ are the prior parameters. Thus from Eq. (9), one can obtain the analytical expression $\mu_{\mu}(d)=\hat{\boldsymbol{\beta}}+\exp[-\phi_d(d-d_1)^2]\ \exp[-\phi_w w_1{}^2/(2\phi_w+1)]\mathbf{R}^{-1}(\boldsymbol{y}^1-\hat{\boldsymbol{\beta}}/\sqrt{2\phi_w+1}$. The analytical expressions for general polynomial $\boldsymbol{h}(\boldsymbol{x})$ and $\boldsymbol{r}(\mathbf{x})$ of Eq. (3) can be obtained in a similar, albeit more tedious, manner.

To simplify notation, define $S\equiv\sigma^2(\boldsymbol{d}|G)$ and $K(\boldsymbol{w},\boldsymbol{w}')\equiv\mathrm{Cov}[Y(G,\boldsymbol{d},\boldsymbol{w}),Y(G,\boldsymbol{d}',\boldsymbol{w}')|\boldsymbol{y}^N]$, and omit the term "$|\boldsymbol{y}^N$" in the following derivation, although it should be understood that the posterior distributions are all conditioned on the simulation output $\boldsymbol{y}^N$. The mean of S is

$$\mu_S(\boldsymbol{d})=E_G[\sigma^2(\boldsymbol{d}|G)]=E_G\{E_{\boldsymbol{w}|G}[(Y(G,\boldsymbol{d},\boldsymbol{w})-\mu(\boldsymbol{d}|G))^2|G]\}.$$

After expanding the squared term, interchanging the order of the expectation operators, and some tedious algebra, we obtain

$$\mu_S=E_{\boldsymbol{w}}[\hat{y}(\boldsymbol{d},\boldsymbol{w})-\mu_{\mu}(\boldsymbol{d})]^2+\int K(\boldsymbol{w},\boldsymbol{w})p_{\boldsymbol{w}}(\boldsymbol{w})d\boldsymbol{w}-\int\int K(\boldsymbol{w},\boldsymbol{w}')p_{\boldsymbol{w}}(\boldsymbol{w})d\boldsymbol{w}p_{\boldsymbol{w}}(\boldsymbol{w}')d\boldsymbol{w}' \tag{A1}$$

We can similarly obtain the variance of S as

$$\begin{aligned}\mathrm{Var}(S)=&2\int\int[K(\boldsymbol{w},\boldsymbol{w}')]^2p_{\boldsymbol{w}}(\boldsymbol{w})d\boldsymbol{w}p_{\boldsymbol{w}}(\boldsymbol{w}')d\boldsymbol{w}'\\&-4\int\int\int K(\boldsymbol{w},\boldsymbol{w}')K(\boldsymbol{w},\boldsymbol{w}'')p_{\boldsymbol{w}}(\boldsymbol{w})d\boldsymbol{w}p_{\boldsymbol{w}}(\boldsymbol{w}')d\boldsymbol{w}'\\&\times p_{\boldsymbol{w}}(\boldsymbol{w}'')d\boldsymbol{w}''+\int\int\int\int[K(\boldsymbol{w},\boldsymbol{w}'')K(\boldsymbol{w}',\boldsymbol{w}''')\\&+K(\boldsymbol{w},\boldsymbol{w}''')K(\boldsymbol{w}',\boldsymbol{w}'')]p_{\boldsymbol{w}}(\boldsymbol{w})d\boldsymbol{w}p_{\boldsymbol{w}}(\boldsymbol{w}')d\boldsymbol{w}'\\&\times p_{\boldsymbol{w}}(\boldsymbol{w}'')d\boldsymbol{w}''p_{\boldsymbol{w}}(\boldsymbol{w}''')d\boldsymbol{w}'''+4\int\int[\hat{y}(\boldsymbol{d},\boldsymbol{w})\\&-\mu_{\mu}(\boldsymbol{d})][\hat{y}(\boldsymbol{d},\boldsymbol{w}')-\mu_{\mu}(\boldsymbol{d})][K(\boldsymbol{w},\boldsymbol{w}')\\&-2\int K(\boldsymbol{w}',\boldsymbol{w}'')p_{\boldsymbol{w}}(\boldsymbol{w}'')d\boldsymbol{w}''+\int\int K(\boldsymbol{w}'',\boldsymbol{w}''')p_{\boldsymbol{w}}(\boldsymbol{w}'')\\&\times d\boldsymbol{w}p_{\boldsymbol{w}}(\boldsymbol{w}''')d\boldsymbol{w}''']p_{\boldsymbol{w}}(\boldsymbol{w})d\boldsymbol{w}p_{\boldsymbol{w}}(\boldsymbol{w}')d\boldsymbol{w}'\end{aligned} \tag{A2}$$

Because $S\equiv[\sigma(\boldsymbol{d}|G)]^2$, if we approximate $\sigma(\boldsymbol{d}|G)$ as Gaussian (which is a better approximation than assuming $\sigma^2(\boldsymbol{d}|G)$ is normal) we can express $\mu_{\sigma}(\boldsymbol{d})$ and $\sigma^2_{\sigma}(\boldsymbol{d})$ in terms of μ_S and $\mathrm{Var}(S)$ via

$$\mu_{\sigma}(\boldsymbol{d})=[\mu_S^2-\mathrm{Var}(S)/2]^{1/4}, \tag{A3}$$

and

$$\sigma^2_{\sigma}(\boldsymbol{d})=\mu_S-[\mu_S^2-\mathrm{Var}(S)/2]^{1/2}. \tag{A4}$$

The covariance between $\mu(\boldsymbol{d}|G)$ and $S\equiv\sigma^2(\boldsymbol{d}|G)$ can be calculated from the identity

$$\mathrm{Cov}[\mu(\boldsymbol{d}|G),S]=E_G[\mu(\boldsymbol{d}|G)S]-[\mu_{\mu}(\boldsymbol{d})\mu_S].$$

By substituting in the definition of S, expanding the squared terms, and interchanging the order of expectations, we obtain

$$\begin{aligned}\mathrm{Cov}[\mu(\boldsymbol{d}|G),S]=&\mu_{\mu}(\boldsymbol{d})\{E_{\boldsymbol{w}}[\hat{y}(\boldsymbol{d},\boldsymbol{w})]^2+E_{\boldsymbol{w}}[K(\boldsymbol{w},\boldsymbol{w})]-[\mu_{\mu}(\boldsymbol{d})]^2\\&-3E_{\boldsymbol{w}}E_{\boldsymbol{w}'}[K(\boldsymbol{w},\boldsymbol{w}')]\}+2E_{\boldsymbol{w}}E_{\boldsymbol{w}'}[\hat{y}(\boldsymbol{d},\boldsymbol{w})K(\boldsymbol{w},\boldsymbol{w}')]\\&-\mu_{\mu}(\boldsymbol{d})\mu_S\end{aligned} \tag{A5}$$

From Eq. (A5), if we further approximate $\mu(\boldsymbol{d}|G)$ and $\sigma(\boldsymbol{d}|G)$ as jointly Gaussian, then from the relation between $\mathrm{Cov}[\mu(\boldsymbol{d}|G),\sigma(\boldsymbol{d}|G)]$ and $\mathrm{Cov}[\mu(\boldsymbol{d}|G),S]$, we obtain

$$\mathrm{Cov}[\mu(\boldsymbol{d}|G),\sigma(\boldsymbol{d}|G)]=\mathrm{Cov}(\mu(\boldsymbol{d}|G),S)/[2\mu_{\sigma}(\boldsymbol{d})] \tag{A6}$$

For the same reasons stated above, Eqs. (A1)–(A6) can all be evaluated analytically. Hence, $\mu_f(\boldsymbol{d})$ and $\sigma^2_f(\boldsymbol{d})$ can be calculated from Eqs. (13) and (14), and the prediction interval of $f(\boldsymbol{d}|G)$ is given by Eq. (15).

References

[1] Chang, K. H., Choi, K. K., Wang, J., Tsai, C. S., and Hardee, E., 1998, "A Multilevel Product Model for Simulation-based Design of Mechanical Systems," Concurr. Eng. Res. Appl., **6**(2), pp. 131–144.

[2] Parry, J., Bornoff, R. B., Stehouwer, P., Driessen, L. T., and Stinstra, E., 2004, "Simulation-Based Design Optimization Methodologies Applied to CFD," IEEE Trans. Compon., Hybrids, Manuf. Technol., **27**(2), pp. 391–397.

[3] Horstemeyer, M. F., and Wang, P., 2003, "Cradle-to-Grave Simulation-Based Design Incorporating Multiscale Microstructure-Property Modeling: Reinvigorating Design With Science," J. Comput.-Aided Mater. Des., **10**(1), pp. 13–34.

[4] Parkinson, A., 1995, "Robust Mechanical Design Using Engineering Models," ASME J. Mech. Des., **117**(Sp. Iss. B), pp. 48–54.

[5] Chen, W., Allen, J. K., Tsui, K. L., and Mistree, F., 1996, "A Procedure for Robust Design: Minimizing Variations Caused by Noise Factors and Control Factors," ASME J. Mech. Des., **118**(4), pp. 478–485.

[6] Du, X. P., and Chen, W., 2000, "Towards a Better Understanding of Modeling Feasibility Robustness in Engineering Design," ASME J. Mech. Des., **122**(4), pp. 385–394.

[7] Thornton, A. C., 2001, "Optimism vs. Pessimism: Design Decisions in the Face of Process Capability Uncertainty," ASME J. Mech. Des., **123**(3), pp. 313–321.

[8] Suri, R., and Otto, K., 2001, "Manufacturing System Robustness Through Integrated Modeling," ASME J. Mech. Des., **123**(4), pp. 630–636.

[9] Kalsi, M., Hacker, K., and Lewis, K., 2001, "A Comprehensive Robust Design Approach for Decision Trade-Offs in Complex Systems Design," ASME J. Mech. Des., **123**(1), pp. 1–10.

[10] Hernandez, G., Allen, J. K., Woodruff, G. W., Simpson, T. W., Bascaran, E., Avila, L. F., and Salinas, F., 2001, "Robust Design of Families of Products With Production Modeling and Evaluation," ASME J. Mech. Des., **123**(2), pp. 183–190.

[11] McAllister, C. D., and Simpson, T. W., 2003, "Multidisciplinary Robust Design Optimization of an Internal Combustion Engine," ASME J. Mech. Des., **125**(1), pp. 124–130.

[12] Du, X. P., and Chen, W., 2004, "Sequential Optimization and Reliability Assessment Method for Efficient Probabilistic Design," ASME J. Mech. Des., **126**(2), pp. 225–233.

[13] Gunawan, S., and Azarm, S., 2004, "Non-Gradient Based Parameter Sensitivity Estimation for Single Objective Robust Design Optimization," ASME J. Mech. Des., **126**(3), pp. 395–402.

[14] Koch, P. N., Yang, R. J., and Gu, L., 2004, "Design for Six Sigma Through Robust Optimization," Struct. Multidiscip. Optim., **26**(3–4), pp. 235–248.

[15] Al-Widyan, K., and Angeles, J., 2005, "A Model-Based Formulation of Robust Design," ASME J. Mech. Des., **127**(3), pp. 388–396.

[16] Chen, W., Wiecek, M. M., and Zhang, J., 1999, "Quality Utility—A Compromise Programming Approach to Robust Design," ASME J. Mech. Des., **121**(2), pp. 179–187.

[17] Du, X. P., Sudjianto, A., and Chen, W., 2004, "An Integrated Framework for Optimization Under Uncertainty Using Inverse Reliability Strategy," ASME J. Mech. Des., **126**(4), pp. 562–570.

[18] Tu, J., Choi, K. K., and Park, Y. H., 2001, "Design Potential Method for Robust System Parameter Design," AIAA J., **39**(4), pp. 667–677.

[19] Zou, T., Mahadevan, S., Mourelatos, Z., and Meernik, P., 2002, "Reliability Analysis of Automotive Body-Door Subsystem," Materialwiss. Werkstofftech., **78**(3), pp. 315–324.

[20] Youn, B. D., Choi, K. K., and Park, Y. H., 2003, "Hybrid Analysis Method for Reliability-Based Design Optimization," ASME J. Mech. Des., **125**(2), pp. 221–232.

[21] Yang, R. J., Akkerman, A., Anderson, D. F., Faruque, O. M., and Gu, L., 2000, "Robustness Optimization for Vehicular Crash Simulations," Rep. Sci. Res. Inst., **2**(6), pp. 8–13.

[22] Kleijnen, J. P. C., 1987, *Statistical Tools for Simulation Practitioners*, Marcel Dekker, New York.

[23] Simpson, T. W., Peplinski, J., Koch, P. N., and Allen, J. K., 2001, "Metamodels for Computer Based Engineering Design: Survey and Recommendations," Eng. Comput., **17**(2), pp. 129–150.
[24] Barthelemy, J.-F. M., and Haftka, R. T., 1993, "Approximation Concepts for Optimum Structural Design—A Review," Struct. Optim., **5**, pp. 129–144.
[25] Sobieszczanski-Sobieski, J., and Haftka, R. T., 1997, "Multidisciplinary Aerospace Design Optimization: Survey of Recent Developments," Struct. Optim., **14**, pp. 1–23.
[26] Box, G. E. P., Hunter, W. G., and Hunter, J. S., 1978, *Statistics for Experimenters*, Wiley, New York.
[27] Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P., 1989, "Design and Analysis of Computer Experiments," Stat. Sci., **4**(4), pp. 409–435.
[28] Currin, C., Mitchell, T., Morris, M. D., and Ylvisaker, D., 1991, "Bayesian Prediction of Deterministic Functions, With Applications to the Design and Analysis of Computer Experiments," J. Am. Stat. Assoc., **86**(416), pp. 953–963.
[29] Hardy, R. L., 1971, "Multiquadratic Equations of Topography and Other Irregular Surfaces," J. Geophys. Res., **76**, pp. 1905–1915.
[30] Dyn, N., Levin, D., and Rippa, S., 1986, "Numerical Procedures for Surface Fitting of Scattered Data by Radial Basis Functions," SIAM (Soc. Ind. Appl. Math.) J. Sci. Stat. Comput., **7**(2), pp. 639–659.
[31] Mullur, A. A., and Messac, A., 2005, "Extended Radial Basis Functions: More Flexible and Effective Metamodeling," AIAA J., **43**(6), pp. 1306–1315.
[32] Meckesheimer, M., Barton, R. R., Simpson, T., Limayen, F., and Yannou, B., 2001, "Metamodeling of Combined Discrete/Continuous Responses," AIAA J., **39**(10), pp. 1950–1959.
[33] Meckesheimer, M., Booker, A. J., Barton, R. R., and Simpson, T. W., 2002, "Computationally Inexpensive Metamodel Assessment Strategies," AIAA J., **40**(10), pp. 2053–2060.
[34] Martin, J. D., and Simpson, T. W., 2005, "Use of Kriging Models to Approximate Deterministic Computer Models," AIAA J., **43**(4), pp. 853–863.
[35] Sasena, M. J., Papalambros, P., and Goovaerts, P., 2002, "Exploration of Metamodeling Sampling Criteria for Constrained Global Optimization," Eng. Optimiz., **34**(3), pp. 263–278.
[36] Wang, G. G., 2003, "Adaptive Response Surface Method using inherited Latin Hypercube Design points," ASME J. Mech. Des., **125**(2), pp. 210–220.
[37] Romero, V. J., Swiler, L. P., and Giunta, A. A., 2004, "Construction of Response Surfaces Based on Progressive-Lattice-Sampling Experimental Designs With Application to Uncertainty Propagation," Struct. Safety, **26**(2), pp. 201–219.
[38] Perez, V. M., Renaud, J. E., and Watson, L. T., 2004, "An Interior-Point Sequential Approximate Optimization Methodology," Struct. Multidiscip. Optim., **27**(5), pp. 360–370.
[39] Farhang Mehr, A., and Azarm, S., 2005, "Bayesian Meta-Modeling of Engineering Design Simulations: A Sequential Approach With Adaptation to Irregularities in the Response Behavior," Int. J. Numer. Methods Eng., **62**(15), pp. 2104–2126.
[40] Lin, Y., Mistree, F., Allen, J. K., Tsui, K.-L., and Chen, V., 2004, "Sequential Metamodeling in Engineering Design," *AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Albany, NY. Paper No. AIAA-2004-4304.
[41] Gu, L., Yang, R. J., Tho, C. H., Makowski, M., Faruque, O., and Li, Y., 2001, "Optimization and Robustness for Crashworthiness of Side Impact," Int. J. Veh. Des., **26**(4), pp. 348–360.
[42] Chen, W., Jin, R., and Sudjianto, A., 2005, "Analytical Global Sensitivity Analysis and Uncertainty Propagation for Robust Design," J. Quality Technol., in press.
[43] Kennedy, M. C., and O'Hagan, A., 2001, "Bayesian Calibration of Computer Experiments," J. R. Stat. Soc. Ser. B. Methodol., **63**, pp. 425–464.
[44] Handcock, M. S., and Stein, M. L., 1993, "A Bayesian Analysis of Kriging," Technometrics, **35**, pp. 403–410.
[45] Jones, D. R., Schonlau, M., and Welch, W. J., 1998, "Efficient Global Optimization of Expensive Black-Box Functions," J. Global Optim., **13**, pp. 455–492.
[46] Pacheco, J. E., Amon, C. H., and Finger, S., 2003, "Bayesian Surrogates Applied to Conceptual Stages of the Engineering Design Process," ASME J. Mech. Des., **125**(4), pp. 664–672.
[47] Jin, R., Du, X., and Chen, W., 2003, "The Use of Metamodeling Techniques for Optimization Under Uncertainty," Struct. Multidiscip. Optim., **25**(2), pp. 99–116.
[48] Duffin, R. J., Peterson, E. L., and Zener, C., 1967, *Geometric Programming*, Wiley, New York, p. 5.
[49] Steinberg, D. M., and Bursztyn, D., 2004, "Data Analytic Tools for Understanding Random Field Regression Models," Technometrics, **46**(4), pp. 411–420.
[50] Matheron, G., 1963, "Principles of Geostatistics," Econ. Geol., **58**, pp. 1246–1266.
[51] O'Hagan, A., 1992, "Some Bayesian Numerical Analysis," in *Bayesian Statistics 4*, J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, eds. Oxford University Press, New York, pp. 345–363.
[52] Montgomery, D. C., Peck, E. A., and Vining, G. G., 2001, *Introduction to Linear Regression Analysis*, 3rd ed., Wiley, New York.
[53] Haylock, R. G., and O'Hagan, A., 1996, "On Inference for Outputs of Computationally Expensive Algorithms With Uncertainty on the Inputs," in *Bayesian Statistics 5* J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith, eds. Oxford University Press, New York, pp. 629–37.
[54] Oakley, J., and O'Hagan, A., 2002, "Bayesian Inference for the Uncertainty Distribution of Computer Model Outputs," Biometrika, **89**(4), pp. 769–784.
[55] Van Trees, H. L., 1968, *Detection, Estimation, and Modulation Theory, Part I*, Wiley, New York, p. 177.
[56] Jin, R., Chen, W., and Sudjianto, A., 2004, "Analytical Metamodel-Based Global Sensitivity Analysis and Uncertainty Propagation for Robust Design," SAE Trans. Journal of Materials and Manufacturing, paper 2004-01–0429.
[57] Martin, J. D., and Simpson, T. W., 2004, "A Monte Carlo Simulation of the Kriging Model," *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Albany, NY.
[58] Martin, J. D., 2005, "A Methodology for Evaluating System-level Uncertainty in the Conceptual Design of Complex Multidisciplinary Systems," Ph.D. thesis, The Pennsylvania State University.
[59] Carlin, B. P., and Louis, T. A., 2000, *Bayes and Empirical Bayes Methods for Data Analysis*, 2nd ed., Chapman & Hall/CRC, New York.