THE ROMES METHOD FOR STATISTICAL MODELING OF REDUCED-ORDER-MODEL ERROR

MARTIN DROHMANN AND KEVIN CARLBERG

Abstract. This work presents a technique for statistically modeling errors introduced by reduced-order models. The method employs Gaussian-process regression to construct a mapping from a small number of computationally inexpensive 'error indicators' to a distribution over the true error. The variance of this distribution can be interpreted as the (epistemic) uncertainty introduced by the reduced-order model. To model normed errors, the method employs existing rigorous error bounds and residual norms as indicators; numerical experiments show that the method leads to a near-optimal expected effectivity in contrast to typical error bounds. To model errors in general outputs, the method uses dual-weighted residuals—which are amenable to uncertainty control—as indicators. Experiments illustrate that correcting the reduced-order-model output with this surrogate can improve prediction accuracy by an order of magnitude; this contrasts with existing 'multifidelity correction' approaches, which often fail for reduced-order models and suffer from the curse of dimensionality. The proposed error surrogates also lead to a notion of 'probabilistic rigor', i.e., the surrogate bounds the error with specified probability.

Primary 65G99, Secondary 65Y20, 62M86

1. Introduction

As computing power increases, computational models of engineered systems are being employed to answer increasingly complex questions that guide decision making, often in time-critical scenarios. It is becoming essential to rigorously quantify and account for both aleatory and epistemic uncertainties in these analyses. Typically, the high-fidelity computational model can be viewed as providing a (costly-to-evaluate) mapping between system *inputs* (e.g., uncertain parameters, decision variables) and system *outputs* (e.g., outcomes, measurable quantities). For example, data assimilation employs collected sensor data (outputs) to update the distribution of uncertain parameters (inputs) of the model; doing so via Bayesian inference requires sampling from the posterior distribution, which can entail thousands of forward model simulations. The computational resources (e.g., weeks on a supercomputer) required for large-scale simulations preclude such high-fidelity models from being feasibly deployed in such scenarios.

To avoid this bottleneck, analysts have turned to surrogate models that approximate the input–output map of the high-fidelity model, yet incur a fraction of their computational cost. However, to be rigorously incorporated in uncertainty-quantification (UQ) contexts, it is critical to quantify the *additional uncertainty* introduced by such an approximation. For example, Bayesian inference aims to sample from the posterior distribution

$$\mathrm{P}[\boldsymbol{\mu}|\bar{s}] \propto \mathrm{P}[\boldsymbol{\mu}]\mathrm{P}[\bar{s}|\boldsymbol{\mu}], \tag{1}$$

where $\boldsymbol{\mu} \in \mathcal{P} \subset \mathbb{R}^{n_\mu}$ denote system inputs, $\bar{s} \in \mathbb{R}$ denotes the measured output,[1] $\mathrm{P}[\boldsymbol{\mu}]$ represents the prior, and $\mathrm{P}[\bar{s}|\boldsymbol{\mu}]$ denotes the likelihood function. Typically, the measured output is modeled as $\bar{s} = s(\boldsymbol{\mu}) + \varepsilon$, where $s : \mathcal{P} \to \mathbb{R}$ denotes the outputs predicted by the high-fidelity model for inputs $\boldsymbol{\mu}$, and ε is a random variable representing measurement noise. Sampling from this posterior distribution (e.g., via Markov-chain Monte–Carlo or importance sampling) is costly, as each sample requires at least one evaluation of the high-fidelity input–output map $\boldsymbol{\mu} \mapsto s$ that appears in the likelihood function.

When a surrogate model is employed, the measured output becomes $\bar{s} = s_{\mathrm{surr}}(\boldsymbol{\mu}) + \delta_s(\boldsymbol{\mu}) + \varepsilon$, where $s_{\mathrm{surr}} : \mathcal{P} \to \mathbb{R}$ denotes the output predicted by the surrogate model, and $\delta_s : \mathcal{P} \to \mathbb{R}$ represents the *surrogate-model output error* or bias. In this case, posterior sampling requires only evaluations of the *surrogate-model* input–output map $\boldsymbol{\mu} \mapsto s_{\mathrm{surr}}$—which is computationally inexpensive—as well as evaluation of the surrogate-model error $\delta_s(\boldsymbol{\mu})$, which is not precisely known in practice. As such, it can be considered a source of *epistemic uncertainty*, as it can be reduced in principle by employing the original high-fidelity model (or a higher fidelity surrogate model). The goal of this work is to construct a statistical model of this surrogate-model error $\widetilde{\delta}_s(\boldsymbol{\mu})$ that is 1) cheaply computable, 2) exhibits low variance (i.e., introduces minimal epistemic uncertainty), and 3) whose distribution can be numerically validated.

Various approaches have been developed for different surrogate models to quantify the surrogate error $\delta_s(\boldsymbol{\mu})$. Surrogate models can be placed into three categories [19]: 1) data fits, 2) lower-fidelity models, and 3) reduced-order models. Data fits employ supervised machine-learning methods (e.g., Gaussian processes, polynomial interpolation [21]) to directly model the high-fidelity input–output map. Within this class of surrogates, it is possible to statistically model the error for stochastic-process data fits, as a prediction for inputs $\boldsymbol{\mu}$ yields a mean $s_{\mathrm{surr}}(\boldsymbol{\mu})$ and a mean-zero distribution $\delta_s(\boldsymbol{\mu})$ that can be associated with epistemic uncertainty. While such models are (unbeatably) fast to query and non-intrusive to implement,[2] they suffer from the curse of dimensionality and lack access to the underlying model's physics, which can hinder predictive robustness.

Lower-fidelity models simply replace the high-fidelity model with a 'coarsened' model obtained by neglecting physics, coarsening the mesh, or employing lower-order finite elements, for example. While such models remain physics based, they often realize only modest computational savings. For such problems, 'multifidelity correction' methods have been developed, primarily in the optimization context. These techniques model the mapping $\boldsymbol{\mu} \mapsto \delta_s$ using a data-fit surrogate; they either enforce 'global' zeroth-order consistency between the corrected surrogate prediction and the high-fidelity prediction at training points [23, 29, 33, 39, 35], or 'local' first- or second-order consistency at trust-region centers [2, 19]. Such approaches tend to work well when the surrogate-model error exhibits a lower variance than the high-fidelity response [35] and the input-space dimension is small.

Reduced-order models (ROMs) employ a projection process to reduce the state-space dimensionality of the high-fidelity computational model. Although intrusive to implement, such physics-based surrogates often lead to more significant computational gains than lower-fidelity models, and higher robustness than data fits. For such models, error analysis has been limited

[1] This work considers one output for notational simplicity. All concepts can be straightforwardly extended to multiple outputs. The numerical experiments treat the case of multiple outputs.

[2] Their construction requires only black-box evaluations of the input–output map of the high-fidelity model.

	ROM	data fits	multifidelity correction	ROM + ROMES
non–intrusive	×	✓	✓	×
output-error correction	×	N/A	✓	✓
rigorous error bounds	✓	×	×	(✓)*
tight error bounds	(✓)†	×	×	✓

* probabilistically rigorous
† good effectivity can only be obtained with very intrusive methods.

TABLE 1. Features of different surrogate models

primarily to computing rigorous *a posteriori* error bounds $\Delta_s(\boldsymbol{\mu})$ satisfying $|\delta_s(\boldsymbol{\mu})| \leq \Delta_s(\boldsymbol{\mu})$ [10, 25, 42]. Especially for nonlinear problems, however, these error bounds are often highly ineffective, i.e., they overestimate the actual error by orders of magnitude [17]. To overcome this shortcoming and obtain tighter bounds, the ROM must be equipped with complex machinery that both increases the computational burden [47, 30] and is intrusive to implement (e.g., reformulate the discretization of the high-fidelity model [45, 48]). Further, *rigorous* bounds are not directly useful for uncertainty quantification (UQ) problems, where a statistical error model that is unbiased, has low variance, and is stochastic is more useful. Recent work [35, Section IV.D] has applied multifidelity correction to ROMs. However, the method did not succeed because the ROM error is often a highly oscillatory function of the inputs and therefore typically exhibits a *higher* variance than the high-fidelity response.

In this paper, we introduce the ROM Error Surrogates (ROMES) method that aims to combine the utility of multifidelity correction with the computational efficiency and robustness of reduced-order modeling. Table 1 compares the proposed approach with existing surrogate-modeling techniques. Similar to the multifidelity-correction approach, we aim to model the ROM error δ_s using a data-fit surrogate. However, as directly approximating the mapping $\boldsymbol{\mu} \mapsto \delta_s$ is ineffective for ROMs, we instead exploit the following key observation: ROMs often generate a small number of physics-based, cheaply computable *error indicators* $\boldsymbol{\rho} : \mathcal{P} \to \mathbb{R}^q$ that correlate strongly with the true error $\delta_s(\boldsymbol{\mu})$. Examples of indicators include the residual norm, dual-weighted residuals, and the rigorous error bounds discussed above. To this end, ROMES approximates the *low-dimensional, well-behaved* mapping $\boldsymbol{\rho}(\boldsymbol{\mu}) \mapsto \delta_s(\boldsymbol{\mu})$ using Gaussian-process regression, which is a stochastic-process data-fit method. Note that ROMES constitutes a generalization of the multifidelity correction approach, as the inputs (or features) of the error model can be any user-defined error indicator—they need not be the system inputs $\boldsymbol{\mu}$. Figure 1 depicts the propagation of information for the proposed method.

In addition to constructing an error surrogate for the system outputs, ROMES can also be used to construct a statistical model for the norm of the error in the system state. Further, ROMES can be used to generate error bounds with 'probabilistic rigor', i.e., an error bound that overestimates the error with a specified probability.

Next, Section 2 introduces the problem formulation and provides a general but brief introduction to model reduction. In Section 3, we introduce the ROMES method, including its objectives, ingredients, and some choices of these ingredients for particular errors. Section 4 briefly summarizes the Gaussian-process kernel method [41] and the relevance vector machine

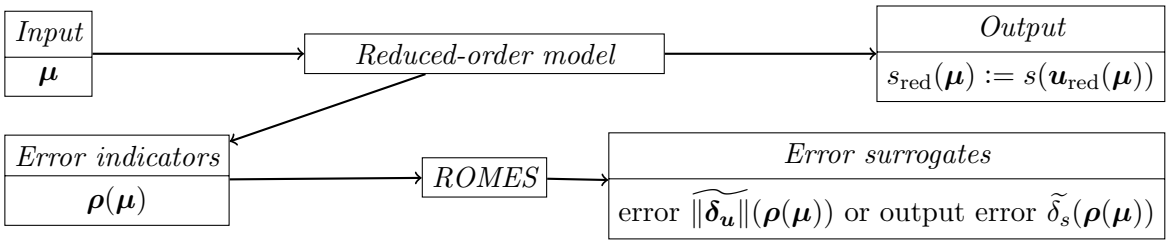

FIGURE 1. ROMES method. The output quantities of interest can be 'corrected' by adding the ROM error surrogate to the ROM output prediction, i.e., $s(\boldsymbol{\mu}) \approx s_{\text{corr}}(\boldsymbol{\mu}) := s_{\text{red}}(\boldsymbol{\mu}) + \widetilde{\delta}_s(\boldsymbol{\rho}(\boldsymbol{\mu}))$.

[44], which are the two machine-learning algorithms we employ to construct the ROMES surrogates. However, the ROMES methodology does not rely on these two techniques, as any supervised machine learning algorithm that generates a stochastic process can be used, as long as it generates a statistical model that meets the important conditions described in Section 3.1. Section 5 analyses the performance of the method when applied with the reduced-basis method to solve Poisson's equation in two dimensions using nine system inputs. The method is also compared with existing rigorous error bounds for normed errors, and with the multifidelity correction approach for errors in general system outputs.

For additional information on the reduced-basis method, including the algorithms to generate the reduced-basis spaces and the computation of error bounds, we refer to the supplementary Section A.

2. Problem formulation

This section details aspects of the high-fidelity and reduced-order models that are important for the ROMES surrogates. We begin with a formulation of the high-fidelity model in Section 2.1 and the reduced-order model in Section 2.2. Finally, we elaborate on the errors introduced by the model-reduction process and possible problems with their quantification in Section 2.3.

2.1. High-fidelity model.

Consider solving a parameterized systems of equations

$$\boldsymbol{r}(\boldsymbol{u}; \boldsymbol{\mu}) = 0, \tag{2}$$

where $\boldsymbol{u} : \mathcal{P} \to \mathbb{R}^n$ denotes the state implicitly defined by Eq. (2), $\boldsymbol{\mu} \in \mathcal{P} \subset \mathbb{R}^{n_{\boldsymbol{\mu}}}$ denotes the system inputs, and $\boldsymbol{r} : \mathbb{R}^n \times \mathbb{R}^{n_{\boldsymbol{\mu}}} \to \mathbb{R}^n$ denotes the residual operator. This model is appropriate for stationary problems, e.g., those arising from the finite-element discretization of elliptic PDEs. For simplicity, assume we are interested in computing a single output

$$s(\boldsymbol{\mu}) := g(\boldsymbol{u}(\boldsymbol{\mu})) \tag{3}$$

with $s : \mathbb{R}^{n_{\boldsymbol{\mu}}} \to \mathbb{R}$ and $g : \mathbb{R}^n \to \mathbb{R}$.

When the dimension n of the high-fidelity model is 'large', computing the system outputs s by first solving Eq. (2) and subsequently applying Eq. (3) can be prohibitively expensive. This is particularly true for many-query problems arising in UQ such as Bayesian inference, which may require thousands of input–output map evaluations $\boldsymbol{\mu} \mapsto s$.

2.2. **Reduced-order model.** Model-reduction techniques aim to reduce the burden of solving Eq. (2) by employing a projection process. First, they execute a computationally expensive *offline stage* (e.g., solving Eq. (2) for a training set $\boldsymbol{\mu} \in \mathcal{P}_{\text{train}} \subset \mathcal{P}$) to construct 1) a low-dimensional trial basis (in matrix form) $\boldsymbol{V} \in \mathbb{R}^{n\times p}$ with $p \ll n$ that (hopefully) captures the behavior of the state $\boldsymbol{u}$ throughout the parameter domain $\mathcal{P}$, and 2) an associated test basis $\boldsymbol{W} \in \mathbb{R}^{n\times p}$. Then, during the computationally inexpensive *online stage*, these methods approximately solve Eq. (2) for arbitrary $\boldsymbol{\mu} \in \mathcal{P}$ by searching for solutions in the trial subspace $\operatorname{range}(\boldsymbol{V}) \subset \mathbb{R}^n$ and enforcing orthogonality of the residual $\boldsymbol{r}$ to the test subspace $\operatorname{range}(\boldsymbol{W}) \subset \mathbb{R}^n$:

$$\boldsymbol{W}^t \boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}; \boldsymbol{\mu}) = 0. \tag{4}$$

Here, the state is approximated as $\boldsymbol{u}_{\text{red}}(\boldsymbol{\mu}) := \boldsymbol{V}\hat{\boldsymbol{u}}(\boldsymbol{\mu})$ and the reduced state $\hat{\boldsymbol{u}}(\boldsymbol{\mu}) \in \mathbb{R}^p$ is implicitly defined by Eq. (4). The ROM-predicted output is then $s_{\text{red}}(\boldsymbol{\mu}) := g\left(\boldsymbol{u}_{\text{red}}(\boldsymbol{\mu}); \boldsymbol{\mu}\right)$.

When the residual operator is nonlinear in the state or non-affine in the inputs, additional complexity-reduction approximations such as empirical interpolation [6, 26], collocation [31, 3, 43], discrete empirical interpolation [16, 22, 18], or gappy proper orthogonal decomposition (POD) [13, 14] are required to ensure that computing the low-dimensional residual $\boldsymbol{W}^t\boldsymbol{r}$ incurs an n-independent operation count. In this case, the residual is approximated as $\tilde{\boldsymbol{r}} \approx \boldsymbol{r}$ and the reduced-order equations become

$$\boldsymbol{W}^t \tilde{\boldsymbol{r}}(\boldsymbol{V}\hat{\boldsymbol{u}}; \boldsymbol{\mu}) = 0. \tag{5}$$

When the output operator is nonlinear and the vector $\partial g/\partial \boldsymbol{u}$ is dense, approximations in the output calculation are also required to ensure an n-independent operation count.

Section A describes in detail the construction of a reduced-order model using the reduced-basis method applied to a parametrically coercive, affine, linear, elliptic PDE.

2.3. **Reduced-order-model error bounds.** One is typically interested in quantifying two types of error incurred by model reduction: the state-space error $\boldsymbol{\delta}_{\boldsymbol{u}}(\boldsymbol{\mu}) := \boldsymbol{u}(\boldsymbol{\mu}) - \boldsymbol{u}_{\text{red}}(\boldsymbol{\mu}) \in \mathbb{R}^n$ and the output error $\delta_s(\boldsymbol{\mu}) := s(\boldsymbol{\mu}) - s_{\text{red}}(\boldsymbol{\mu}) \in \mathbb{R}$. In particular, many ROMs are equipped with computable, rigorous error bounds for these quantities [37, 24, 36, 11, 17]:

$$\Delta_{\boldsymbol{u}}(\boldsymbol{\mu}) \geq \|\boldsymbol{\delta}_{\boldsymbol{u}}(\boldsymbol{\mu})\|, \qquad \Delta_s(\boldsymbol{\mu}) \geq |\delta_s(\boldsymbol{\mu})| \tag{6}$$

In cases, where the norm of the residual operator can be estimated tightly, lower bounds also exist:

$$\Delta_{\boldsymbol{u}}^{\text{LB}}(\boldsymbol{\mu}) \leq \|\boldsymbol{\delta}_{\boldsymbol{u}}(\boldsymbol{\mu})\|, \qquad \Delta_s^{\text{LB}}(\boldsymbol{\mu}) \leq |\delta_s(\boldsymbol{\mu})|. \tag{7}$$

The performance of an upper bound is usually quantified by its *effectivity*, i.e., the factor by which it overestimates the true error

$$\eta^{\boldsymbol{u}}(\boldsymbol{\mu}) := \frac{\Delta_{\boldsymbol{u}}(\boldsymbol{\mu})}{\|\boldsymbol{\delta}_{\boldsymbol{u}}(\boldsymbol{\mu})\|} \geq 1, \qquad \eta^s(\boldsymbol{\mu}) := \frac{\Delta_s(\boldsymbol{\mu})}{|\delta_s(\boldsymbol{\mu})|} \geq 1. \tag{8}$$

The closer these values are to 1, the 'tighter' the bound. For coercive PDEs, these effectivities can be controlled by choosing a tight lower bound of the coercivity constant.

While this can be easily accomplished for stationary, linear problems, it is difficult to find tight lower bounds in almost all other cases. In fact, the resulting bounds often overestimate the error by orders of magnitude [42, 17]. Because effectivity is critically important in practice, various efforts have been undertaken to improve the tightness of the bounds. Huynh et

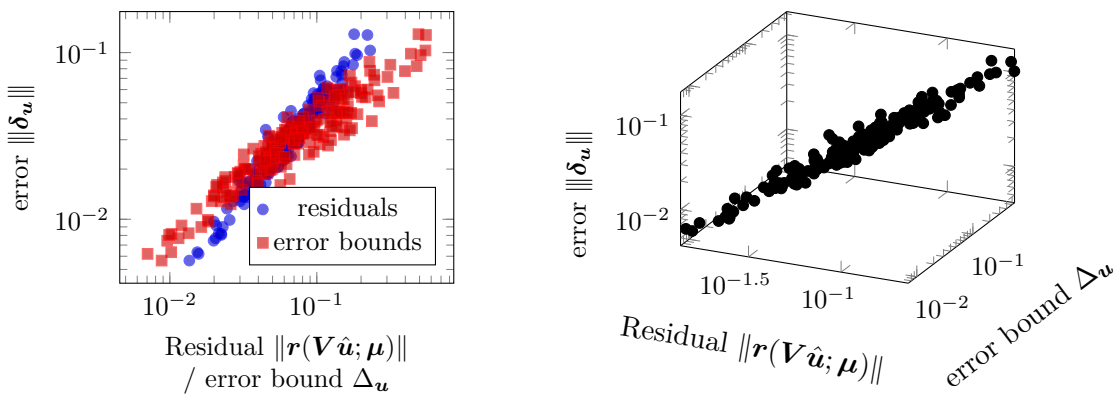

FIGURE 2. Relationship between RB error bounds $\Delta_{\boldsymbol{u}}$, residual norms $\|\boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}};\boldsymbol{\mu})\|$, and the true state-space errors $|||\boldsymbol{\delta_u}|||$, visualized by evaluation of 200 random sample points in the input space. Here, $|||\cdot|||$ denotes the energy norm defined in Section 5.1.

al. [30] developed the successive constraint method for this purpose; the method approximates the coercivity-constant lower bounds by solving small linear programs online, which depend on additional expensive offline computations. Alternatively, Refs. [45, 48] reformulate the entire discretization of time-dependent problems using a space–time method that improves the error bounds by incorporating solutions to dual problems. Another approach [47] aims to approximate the coercivity constant by eigenvalue analysis of the reduced system matrices. These methods all bloat the offline and the online computation time and often incur intrusive changes to the high-fidelity-model implementation.

Regardless of effectivity, rigorous bounds satisfying inequalities (6) are not directly useful for quantifying the epistemic uncertainty incurred by employing the ROM. Rather, a *statistical model* that reflects our knowledge of these errors would be more appropriate. For such a model, the mean of the distribution would provide an expected error; the variance would provide a notion of epistemic uncertainty. The most straightforward way to achieve this would be to model the error as a uniform probability distribution on an interval whose boundaries correspond to the lower and upper bounds. Unfortunately, such an approach leads to wide, uninformative intervals when the bounds suffer from poor effectivity; this will be demonstrated in the numerical experiments.

Instead, we exploit the following observation: error bounds tend to be strongly correlated with the true error. Figure 2 depicts this observed structure for a reduced-basis ROM applied to an elliptic PDE (see Section 5.1 for details). On a logarithmic scale, the true error exhibits a roughly linear dependence on both the bound and the residual norm, and the variance of the data is fairly constant. As will be shown in Section 5.3, employing a multifidelity correction approach wherein the error is modeled as a function of the inputs $\boldsymbol{\mu}$ does not work well for this example, both because the input-space dimension is large (nine) and the error is a highly oscillatory function of these inputs.

Therefore, we propose constructing a stochastic process that maps such *error indicators* to a *random variable for the error*. For this purpose, we employ Gaussian-process regression. The approach leverages one strength of ROMs compared to other surrogate models: ROMs generate strong 'physics-based' error indicators (e.g., error bounds) in addition to output predictions. The next section describes the proposed method.

3. The ROMES method

The objective of the ROMES method is to construct a statistical model of the deterministic, but generally unknown ROM error $\delta(\boldsymbol{\mu})$ with $\delta : \mathcal{P} \to \mathbb{R}$ denoting an $\mathbb{R}$-valued error that may represent the norm of the state-space error $\|\boldsymbol{\delta_u}\|$, the output error δ_s, or its absolute value $|\delta_s|$, for example. The distribution of the random variable representing the error should reflect our (epistemic) uncertainty about its value. We assume that we can employ a set of training points $\delta(\boldsymbol{\mu}_n)$, $n = 1, \dots, N$ to construct this model.

3.1. **Statistical model.** Define a probability space $(\Omega, \mathcal{F}, P)$. We seek to approximate the deterministic mapping $m : \boldsymbol{\mu} \mapsto d(\delta(\boldsymbol{\mu}))$ by a stochastic mapping $\tilde{m} : \boldsymbol{\rho}(\boldsymbol{\mu}) \mapsto \tilde{d}$ with $\tilde{d} : \Omega \to \mathbb{R}$ a real-valued random variable. Here, $d : \mathbb{R} \to \mathbb{R}$ is an invertible *transformation function* (e.g., logarithm) that can be specified to facilitate construction of the statistical model. We can then interpret the statistical model of the error as a random variable $\tilde{\delta} : \boldsymbol{\rho} \mapsto d^{-1}(\tilde{m}(\boldsymbol{\rho}))$.

Three ingredients must be selected to construct this mapping $\tilde{m}$: 1) the error indicators $\boldsymbol{\rho}$, 2) the transformation function d, and 3) the methodology for constructing the statistical model from the training data. We will make these choices such that the stochastic mapping satisfies the following conditions:

(1) The indicators $\boldsymbol{\rho}(\boldsymbol{\mu})$ are *cheaply computable* and *low dimensional* given any $\boldsymbol{\mu} \in \mathcal{P}$. In practice, they should also incur a reasonably small implementation effort, e.g., not require modifying the underlying high-fidelity model.

(2) The mapping $\tilde{m}$ exhibits *low variance*, i.e., $\mathrm{E}\left[(\tilde{m}(\boldsymbol{\rho}(\boldsymbol{\mu})) - \mathrm{E}\left[\tilde{m}(\boldsymbol{\rho}(\boldsymbol{\mu}))\right])^2\right]$ is 'small' for all $\boldsymbol{\mu} \in \mathcal{P}$. This ensures that little additional epistemic uncertainty is introduced.

(3) The mapping $\tilde{m}$ is *validated*:

$$\omega_{\text{validation}}(\omega) \approx \omega, \quad \forall \omega \in [0, 1), \tag{9}$$

where $\omega_{\text{validation}}(\omega)$ is the frequency with which validation data lie in the ω-confidence interval predicted by the statistical model

$$\omega_{\text{validation}}(\omega) := \frac{\operatorname{card}\left(\{\boldsymbol{\mu} \in \mathcal{P}_{\text{validation}} \mid d(\delta(\boldsymbol{\mu})) \in C_\omega(\boldsymbol{\mu})\}\right)}{\operatorname{card}\left(\mathcal{P}_{\text{validation}}\right)}. \tag{10}$$

Here, the validation set $\mathcal{P}_{\text{validation}} \subset \mathcal{P}$ should not include any of the points $\boldsymbol{\mu}_n$, $n = 1, \dots, N$ employed to train the error surrogate, and the confidence interval $C_\omega(\boldsymbol{\mu}) \subset \mathbb{R}$, which is centered at the mean of $\tilde{m}(\boldsymbol{\rho}(\boldsymbol{\mu}))$, is defined for all $\boldsymbol{\mu} \in \mathcal{P}$ such that

$$\mathrm{P}[\tilde{m}(\boldsymbol{\rho}(\boldsymbol{\mu})) \in C_\omega(\boldsymbol{\mu})] = \omega. \tag{11}$$

In essence, validation assesses whether or not the data do indeed behave as random variables with probability distributions predicted by the statistical model.

The next section describes the proposed methodology for selecting indicators $\boldsymbol{\rho}$ and transformation function d. For constructing the mapping $\tilde{m}$ from training points, we will employ the two supervised machine learning algorithms described in Section 4: the Gaussian process (GP) kernel method and the relevance vector machine (RVM). Note that these are merely guidelines for model construction, as there are usually no strong analytical tools to prove that the mapping behaves according to a certain probability distribution. Therefore, any choice must be computationally validated according to condition 3 above.

3.2. **Choosing indicators and transformation function.** The class of multifidelity-correction algorithms can be cast within the framework proposed in Section 3.1. In particular, when a stochastic process is used to model additive error, these methods are equivalent to the proposed construction with ingredients $\delta = \delta_s$, $\boldsymbol{\rho} = \boldsymbol{\mu}$, and $d = \mathrm{id}_{\mathbb{R}}$ with $\mathrm{id}_{\mathbb{R}}(x) = x$, $\forall x \in \mathbb{R}$ the identity function over $\mathbb{R}$. However, as previously discussed, the mapping $\boldsymbol{\mu} \mapsto \delta_s$ can be highly oscillatory and non-smooth for reduced-order models; further, this approach is infeasible for high-dimensional input spaces, i.e., $n_{\boldsymbol{\mu}}$ large. This was shown by Ng and Eldred [35, Section IV.D]; we also demonstrate this in the numerical experiments of Section 5.3.

Note that all indicators and errors proposed in this section should be scaled (e.g., via linear transformations) in practice such that they exhibit roughly the same range. This 'feature scaling' task is common in machine-learning and is particularly important when the ROMES surrogate employs multiple indicators.

3.2.1. *Normed and compliant outputs.* As discussed in Section 2.3, many ROMs are equipped with bounds for normed errors. Further, there is often a strong, well-behaved relationship between such error bounds and the normed error (see Figure 2). In the case of compliant outputs, the error is always non-negative, i.e., $\delta_s = |\delta_s|$ (see Section A.3), so we can treat this error as a normed error.

To this end, we propose employing error bounds as indicators for the errors in the compliant output $|\delta_s|$ and in the state $\|\boldsymbol{\delta_u}\|$. However, because the bound effectivity often lies in a small range (even for a large range of errors) [42], employing a logarithmic transformation is appropriate. To see this, consider a case where the effectivity η of the error bound, defined as

$$\eta(\boldsymbol{\mu}) := \frac{\Delta(\boldsymbol{\mu})}{\delta(\boldsymbol{\mu})} \geq 1, \quad \forall \boldsymbol{\mu} \in \mathcal{P}, \tag{12}$$

lies within a range $\eta_1 \leq \eta(\boldsymbol{\mu}) \leq \eta_2$, $\forall \boldsymbol{\mu} \in \mathcal{P}$. Then the relationship between the error bound and the true error is

$$\frac{\Delta(\boldsymbol{\mu})}{\eta_1} \geq \delta(\boldsymbol{\mu}) \geq \frac{\Delta(\boldsymbol{\mu})}{\eta_2} \tag{13}$$

$$\log \Delta(\boldsymbol{\mu}) - \log \eta_1 \geq \log \delta(\boldsymbol{\mu}) \geq \log \Delta(\boldsymbol{\mu}) - \log \eta_2 \tag{14}$$

for all $\boldsymbol{\mu} \in \mathcal{P}$. In this case, one would expect an affine model mapping $\log \Delta(\boldsymbol{\mu})$ to $\log \delta(\boldsymbol{\mu})$ with constant Gaussian noise to accurately capture the relationship. So, employing a logarithmic transformation permits the use of simpler surrogates that assume a constant variance in the error variable. Therefore, we propose employing $\boldsymbol{\rho} = \log \Delta$ and $d = \log$ for statistical modeling of normed errors and compliant outputs.

A less computationally expensive candidate for an indicator is simply the logarithm of the residual norm $\boldsymbol{\rho} = \log r$, where r is the Euclidean norm of the residual vector

$$r(\boldsymbol{\mu}) := \|\boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}(\boldsymbol{\mu}); \boldsymbol{\mu})\|_2 . \tag{15}$$

For more information on the efficient computation of (15), we refer to Section A.3.1. One would expect this choice of indicator to produce a similar model to that produced by the logarithm of the error bound: the error bound is often equal to the residual norm divided by a (costly-to-compute) coercivity-constant lower bound (see Section A.3). Further, employing the residual norm leads to a model that is less sensitive to strong variations in this approximated lower bound.

Returning to the example depicted in Figure 2, the relationship between the error bound and the energy norm of the state error in log-log space is roughly linear, and the variance is relatively small. The same is true of the relationship between the (computationally cheaper) residual norm and the true error. As expected, these relationships can be accurately modeled as a stochastic process with a linear mean function and constant variance (more details in Section 5.1). Here, strong candidates for ROMES error indicators include $\boldsymbol{\rho}_1(\boldsymbol{\mu}) := r(\boldsymbol{\mu}), \boldsymbol{\rho}_2(\boldsymbol{\mu}) := \Delta(\boldsymbol{\mu})$, and $\boldsymbol{\rho}_3(\boldsymbol{\mu}) := (r(\boldsymbol{\mu}), \Delta(\boldsymbol{\mu}))$. In the experiments in Section 5, we will consider the first choice, which is the least expensive and intrusive option, yet leads to excellent results. For cases where the data are less well behaved, more error indicators can be included, e.g., linear combinations of inputs or the output prediction.

Unfortunately, this set of ROMES ingredients is not applicable to errors in *general* outputs of interest because the logarithmic transformation function assumes strictly positive errors. The next section presents a strategy for handling this general case.

Remark 3.1 (Log-Normal distribution)**.** *In the case where* $d = \log$*, the error models* $\widetilde{\delta}(\boldsymbol{\mu})$*,* $\boldsymbol{\mu} \in \mathcal{P}$ *are random variables with log-normal distribution. If one is interested in the most probable error, one might think to use the expected value of* $\widetilde{\delta}$*. However, the maximum of the probability distribution function of a log-normally distributed random variable is defined by its* mode*, which is less than the expected value. We therefore use* $\mathrm{mode}(\widetilde{\delta})$ *if scalar values for the estimation of the output error or the reduced state error are required.*

3.2.2. *General outputs.* This section describes the ROMES ingredients we propose for modeling the error δ_s in a general output $s(\boldsymbol{\mu}) := g(\boldsymbol{u}(\boldsymbol{\mu}))$. Dual-weighted-residual approaches are commonly adopted for approximating general output errors in the context of *a posteriori* adaptivity [20, 4, 7, 40, 46, 32], model-reduction adaptivity [12], and model-reduction error estimation [38, 11, 45, 34]. The latter references compute adjoint solutions in order to improve the accuracy of ROM output-error bounds. The computation of these adjoint solutions entails a low-dimensional linear solve; thus, they are efficiently computable and can potentially serve as error indicators for the ROMES method.

The main idea of dual-weighted-residual error estimation is to approximate the output error to first-order using the solution to a dual problem. For notational simplicity in this section, we drop dependence on the inputs $\boldsymbol{\mu}$.

To begin, we approximate the output arising from the (unknown) high-fidelity state $\boldsymbol{u}$ to first order about the ROM-approximated state $\boldsymbol{V}\hat{\boldsymbol{u}}$:

$$g(\boldsymbol{u}) \approx g(\boldsymbol{V}\hat{\boldsymbol{u}}) + \frac{\partial g}{\partial \boldsymbol{u}}(\boldsymbol{V}\hat{\boldsymbol{u}})(\boldsymbol{u} - \boldsymbol{V}\hat{\boldsymbol{u}}) \tag{16}$$

with $g : \mathbb{R}^n \to \mathbb{R}$ and $\frac{\partial g}{\partial \boldsymbol{u}} : \mathbb{R}^n \to \mathbb{R}^{1\times n}$. Similarly, we can approximate the residual to first order about the approximated state as

$$0 = \boldsymbol{r}(\boldsymbol{u}) \approx \boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}) + \frac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}(\boldsymbol{V}\hat{\boldsymbol{u}})(\boldsymbol{u} - \boldsymbol{V}\hat{\boldsymbol{u}}), \tag{17}$$

where $\boldsymbol{r} : \mathbb{R}^n \to \mathbb{R}^n$ with $\frac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}} : \mathbb{R}^n \to \mathbb{R}^{n\times n}$. Solving for the error yields

$$(\boldsymbol{u} - \boldsymbol{V}\hat{\boldsymbol{u}}) \approx -\left[\frac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}(\boldsymbol{V}\hat{\boldsymbol{u}})\right]^{-1} \boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}). \tag{18}$$

Substituting (18) in (16) leads to

$$g(\boldsymbol{u}) - g(\boldsymbol{V}\hat{\boldsymbol{u}}) \approx \boldsymbol{y}^T \boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}), \tag{19}$$

where the dual solution $\boldsymbol{y} \in \mathbb{R}^n$ satisfies

$$\frac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}^t (\boldsymbol{V}\hat{\boldsymbol{u}})\boldsymbol{y} = -\frac{\partial g}{\partial \boldsymbol{u}}^t (\boldsymbol{V}\hat{\boldsymbol{u}}). \tag{20}$$

Approximation (19) is first-order accurate; therefore, it is exact in the case of linear outputs and a linear residual operator. In the general nonlinear case, this approximation is accurate in a neighborhood of the ROM-approximated state $\boldsymbol{V}\hat{\boldsymbol{u}}$.

Because we would like to avoid high-dimensional solves, we approximate $\boldsymbol{y}$ as the reduced-order dual solution $\boldsymbol{y}_{\text{red}} := \boldsymbol{Y}\hat{\boldsymbol{y}} \in \mathbb{R}^n$, where $\hat{\boldsymbol{y}}$ satisfies

$$\boldsymbol{Y}^T \frac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}^t (\boldsymbol{V}\hat{\boldsymbol{u}})\boldsymbol{Y}\hat{\boldsymbol{y}} = \boldsymbol{Y}^T \frac{\partial g}{\partial \boldsymbol{u}}^t (\boldsymbol{V}\hat{\boldsymbol{u}}), \tag{21}$$

and $\boldsymbol{Y} \in \mathbb{R}^{n\times p_y}$ with $p_y \ll n$ is a reduced basis (in matrix form) for the dual system. Section A.3.2 provides details on the construction of $\boldsymbol{Y}$ for elliptic PDEs. Substituting the approximation $\boldsymbol{y}_{\text{red}}$ for $\boldsymbol{y}$ in (19) yields a cheaply computable error estimate

$$g(\boldsymbol{u}) - g(\boldsymbol{V}\hat{\boldsymbol{u}}) \approx \boldsymbol{y}_{\text{red}}^t \boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}). \tag{22}$$

This relationship implies that one can construct an accurate, cheaply computable ROMES model for general-output error $\delta = \delta_s = g(\boldsymbol{u}) - g(\boldsymbol{V}\hat{\boldsymbol{u}})$ by employing indicators $\boldsymbol{\rho} = \boldsymbol{y}_{\text{red}}^t \boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}})$ and transformation function $d = \text{id}_{\mathbb{R}}$ the identity function over $\mathbb{R}$.

Remark 3.2 (Uncertainty control for dual-weighted-residual error indicators)**.** *The accuracy of the reduced-order dual solution can be controlled by changing p_y—the dimension of the dual basis $\boldsymbol{Y}$. In general, one would expect an increase in p_y to lead to a lower-variance ROMES surrogate at the expense of a higher dimensional dual problem* (21)*. The experiments in Section 5.6 highlight this uncertainty-control attribute.*

3.3. **Probabilistically rigorous error bounds.** Clearly, the ROMES surrogate does not strictly bound the error, even when error bounds are used as indicators. That is, the mean probability of overestimation is generally less than one

$$c := \frac{1}{|\mathcal{P}|} \int_{\boldsymbol{\mu}\in\mathcal{P}} \mathrm{P}[\tilde{m}(\boldsymbol{\rho}(\boldsymbol{\mu})) > d(\delta(\boldsymbol{\mu}))] d\boldsymbol{\mu} < 1. \tag{23}$$

This frequency of overestimation depends on the probability distribution of the random variable $\tilde{m}(\boldsymbol{\rho})$. Using the machine learning methods proposed in the next section, we infer normally distributed random variables

$$\tilde{m}(\boldsymbol{\rho}) \sim \mathcal{N}(\nu(\boldsymbol{\rho}), \overline{\sigma}^2(\boldsymbol{\rho})) \tag{24}$$

with mean $\nu(\boldsymbol{\rho})$ and variance $\overline{\sigma}^2(\boldsymbol{\rho})$. If the model is perfectly validated, then the mean probability of overestimation is $c = 0.5$. However, knowledge about the distribution of the random variable can be used to control the overestimation frequency. In particular, the modified surrogate

$$\tilde{m}^c(\boldsymbol{\rho}) := \tilde{m}(\boldsymbol{\rho}) + m_{\mathrm{LB}}(\boldsymbol{\rho}, c) \tag{25}$$

enables *probabilistic rigor*: it bounds the error with mean specified probability c assuming the model is perfectly validated. Here, m_{LB} fulfills

$$\mathrm{P}[X > m_{\mathrm{LB}}(\boldsymbol{\rho}, c)] = c, \qquad \text{for } X \sim \mathcal{N}(0, \overline{\sigma}^2(\boldsymbol{\rho})). \tag{26}$$

This value can be computed as

$$m_{\mathrm{LB}}(\boldsymbol{\rho}) = \sqrt{2}\overline{\sigma}(\boldsymbol{\rho})\mathrm{erf}^{-1}(2c-1) \tag{27}$$

where erf^{-1} is the inverse of the *error function*.

4. Gaussian processes

This section describes the two methods we employ to construct the stochastic mapping $\tilde{m} : \boldsymbol{\rho}(\boldsymbol{\mu}) \mapsto \tilde{d}$:

- *Gaussian process kernel regression* (i.e., kriging) [41] and
- the *relevance vector machine* (RVM) [44].

Both methods are examples of supervised learning methods that generate a stochastic process from a set of N training points for independent variables $\mathbf{x} := (\boldsymbol{x}_n)_{n=1}^N$ and a dependent variable $\mathbf{y} := (y_n)_{n=1}^N$. Using these training data, the methods generate predictions $\tilde{y}(\boldsymbol{x}_m^*; \theta^{\mathrm{ML}})$, $m = 1, \ldots, M$ associated with a set of M prediction points $\mathbf{x}^* := (\boldsymbol{x}_m^*)_{m=1}^M$. Here, θ^{ML} denotes hyperparameters that are inferred using a Bayesian approach; the predictions are random variables with a multivariate normal distribution.

In the context of ROMES, the independent variables correspond to error indicators $\boldsymbol{x}_n = \boldsymbol{\rho}(\boldsymbol{\mu}_n)$ with $\boldsymbol{\mu}_n \in \mathcal{P}$, $n = 1, \ldots, N$ and the dependent variable corresponds to the (transformed) reduced-order-model error such that $y_n = d(\delta(\boldsymbol{\mu}_n))$, $n = 1, \ldots, N$. To make this paper as self-contained as possible, the following sections briefly present and compare the two approaches.

4.1. GP kernel method.

A Gaussian process is defined as a collection of random variables such that any finite number of them has a joint Gaussian distribution. The GP kernel method constructs this Gaussian process via Bayesian inference using the training data and a specified kernel function. To begin, the prior distribution is set to

$$\tilde{y}_{\mathrm{prior}}(\underline{\mathbf{x}}) \sim \mathcal{N}\left(0, \mathbf{K}\left(\underline{\mathbf{x}}, \underline{\mathbf{x}}\right) + \sigma^2 \mathbf{I}_{N+M}\right) \tag{28}$$

with $\underline{\mathbf{x}} := (\underline{\boldsymbol{x}}_i)_{i=1}^{N+M} = (\mathbf{x}, \mathbf{x}^*)$. Here, the GP kernel assumes that the covariance between any two points can be described analytically by a kernel k with additive noise $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{M+N})$

such that

$$\mathbf{K}(\underline{\mathbf{x}}, \underline{\mathbf{x}}) = \left(k(\underline{\boldsymbol{x}}_i, \underline{\boldsymbol{x}}_j)\right)_{1 \le i,j \le N+M}. \tag{29}$$

In this work, we employ the most commonly used squared-exponential-covariance kernel

$$k(\boldsymbol{x}_i, \boldsymbol{x}_j) = \exp\left(-\frac{\|\boldsymbol{x}_i - \boldsymbol{x}_j\|_2^2}{2l^2}\right), \tag{30}$$

which induces high correlation between geometrically nearby points. Here, $l \in \mathbb{R}$ is the 'width' hyperparameter.

Assuming the predictions are generated as independent samples from the stochastic process,[3] the GP kernel method then generates predictions for each point $\boldsymbol{x}^* \in \mathbf{x}^*$. These predictions correspond to random variables with posterior distributions $\tilde{y}(\boldsymbol{x}^*; \theta) \sim \mathcal{N}(\nu(\boldsymbol{x}^*), \overline{\sigma}^2(\boldsymbol{x}^*))$ with

$$\nu(\boldsymbol{x}^*) = \mathbf{K}(\boldsymbol{x}^*, \mathbf{x}) \left(\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma^2 \mathbf{I}_N\right)^{-1} \mathbf{y} \tag{31}$$

$$\overline{\sigma}^2(\boldsymbol{x}^*) = \Sigma(\boldsymbol{x}^*) + \sigma^2 \tag{32}$$

$$\Sigma(\boldsymbol{x}^*) = \mathbf{K}(\boldsymbol{x}^*, \boldsymbol{x}^*) - \mathbf{K}(\boldsymbol{x}^*, \mathbf{x}) \left(\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma^2 \mathbf{I}_N\right)^{-1} \mathbf{K}(\mathbf{x}, \boldsymbol{x}^*). \tag{33}$$

More details on the derivation of these expressions can be found in Ref. [41, ch 2.2].

The hyperparameters $\theta := (l^2, \sigma^2)$ can be set to the maximum-likelihood values θ^{ML} computed as the solution to an optimization problem

$$\theta^{\mathrm{ML}} = \arg\max_{\theta} \mathcal{L}(\theta) \tag{34}$$

with the log-likelihood function defined as

$$\mathcal{L}(l^2, \sigma^2) = -\frac{1}{2}\mathbf{y}^t \left(\mathbf{K}(\mathbf{x}, \mathbf{x}; l^2) + \sigma^2 \mathbf{I}_N\right)^{-1} \mathbf{y} - \frac{1}{2}\log\left|\mathbf{K}(\mathbf{x}, \mathbf{x}; l^2) + \sigma^2 \mathbf{I}_N\right| - \frac{N}{2}\log 2\pi. \tag{35}$$

For details on the derivation of the log likelihood function and problem (34), we refer to Ref. [41, ch 5.4].

Remark 4.1. *The noise component σ^2 of posterior covariance $\overline{\sigma}^2(\boldsymbol{x}^*)$ accounts for uncertainty in the assumed GP structure. It plays a crucial role for the ROMES method: it accounts for the non-uniqueness of the mapping $\boldsymbol{\rho} \mapsto \delta$, as it is possible for $\delta(\boldsymbol{\mu}_i) \neq \delta(\boldsymbol{\mu}_j)$ even if $\boldsymbol{\rho}(\boldsymbol{\mu}_i) = \boldsymbol{\rho}(\boldsymbol{\mu}_j)$. In particular, this noise component represents the 'information loss' incurred by employing the error indicators in lieu of the system inputs as independent variables in the GP. Therefore, this component can be interpreted as the* inherent uncertainty in the error due to the non-uniqueness of the mapping $\boldsymbol{\rho} \mapsto \delta$.

On the other hand, the remaining term $\Sigma(\boldsymbol{x}^*)$ of the posterior variance quantifies the uncertainty in the mean prediction. This decreases as the number of training points increases. Therefore, $\Sigma(\boldsymbol{x}^*)$ can be interpreted as the uncertainty due to a lack of training data.

For example, the multifidelity-correction approach employs $\boldsymbol{\rho} = \boldsymbol{\mu}$ and therefore should be characterized by $\sigma^2 = 0$, as the mapping $\boldsymbol{\mu} \mapsto \delta$ is unique. However, due to the high-dimensional nature of the system-input space $\mathcal{P}$ in many problems, the uncertainty due to lack

[3] Typically in the GP literature, predictions at all points $\mathbf{x}^*$ are generated simultaneously as a single sample from the Gaussian process. In this work, we treat all predictions as arising from independent samples of the GP.

of training $\Sigma(\boldsymbol{x}^*)$ can be very large unless many training points are employed. On the other hand, the ROMES method aims to significantly reduce $\Sigma(\boldsymbol{x}^*)$ by employing a small number of indicators, albeit at the cost of a nonzero σ^2.

In light of this remark, we will employ two different types of ROMES models: one that includes the uncertainty due to a lack of training data (i.e., variance $\overline{\sigma}^2(\boldsymbol{x}^*)$), and one that neglects this uncertainty (i.e., variance σ^2).

4.2. **Relevance vector machine (RVM) method.** The RVM [44] is based on a parameterized discretization of the predictive random variable

$$\tilde{y}(\boldsymbol{x}) = \sum_{k=1}^{K} w_k \phi_k(\boldsymbol{x}) + \epsilon = \boldsymbol{\phi}(\boldsymbol{x})^t \mathbf{w} + \epsilon, \tag{36}$$

with specified basis functions $\boldsymbol{\phi}(\boldsymbol{x}) := [\phi_1(\boldsymbol{x}) \ \cdots \ \phi_K(\boldsymbol{x})]^t \in \mathbb{R}^K$, a corresponding set of random variables $\mathbf{w} := [w_1 \ \cdots \ w_K]^t \in \mathbb{R}^K$, with $w_k \sim \mathcal{N}(0, \beta_k^2)$ for $k = 1, \ldots, K$ and noise $\epsilon \sim \mathcal{N}(0, \sigma^2)$. The hyperparameters $\boldsymbol{\beta} = [\beta_1 \ \cdots \ \beta_K]^t \in \mathbb{R}^K$ define the prior probability distribution, and are usually chosen by a likelihood maximization over the training samples. Radial basis functions

$$\phi_k^{RBF}(\boldsymbol{x}) = \exp\left(-\frac{1}{r^2} \|\bar{\boldsymbol{x}}_k - \boldsymbol{x}\|_2^2\right), \quad k = 1, \ldots, K \tag{37}$$

constitute the most common choice for basis functions. For the ROMES method, we often expect a linear relationship between the indicators and true errors, likely with a small-magnitude high-order-polynomial deviation. Therefore, we also consider Legendre polynomials [1, Ch.8]

$$\phi_k^{Leb}(\boldsymbol{x}) = P_k(\boldsymbol{x}), \quad k = 1, \ldots, K. \tag{38}$$

Note that both sets of basis functions are dependent on the training data: while the centering points $\bar{\boldsymbol{x}}_k$, $k = 1, \ldots, K$ in the radial basis functions can be chosen arbitrarily, they are typically chosen to be equal to the training points. The domain of the Legendre polynomials, on the other hand, is nominally $[-1, 1]$; therefore the independent variables must be appropriately scaled to ensure the range of training and prediction points $(\mathbf{x}, \mathbf{x}^*)$ is included in this interval.

The RVM method also employs a Bayesian approach to construct the model from training data. In particular, the vector of hyperparameters $\boldsymbol{\beta}$ affects the variance of the Gaussian random variables $\mathbf{w}$. If these hyperparameters are computed by a maximum-likelihood or a similar optimization algorithm, large values for these hyperparameters identify insignificant components that can be removed. Therefore, in the ROMES context, the RVM can be used to filter out the least significant error indicators. Apart from this detail, the RVM can be considered a special case of the GP kernel method with kernel

$$k(\boldsymbol{x}_i, \boldsymbol{x}_j) = \sum_{k=1}^{K} \frac{1}{\beta_k} \phi_k(\boldsymbol{x}_i) \phi_k(\boldsymbol{x}_j). \tag{39}$$

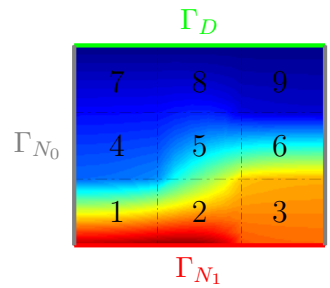

FIGURE 3. Domain and sample solution $\boldsymbol{u}(\boldsymbol{\mu})$ for the thermal-block problem.

5. Numerical experiments

This section analyzes the performance of the ROMES method on Poisson's equation with nine system inputs, using the reduced-basis method to generate the reduced-order model. First, Section 5.1 introduces the test problem. Section 5.2 discusses implementation and validation of the ROMES models. Section 5.3 compares the ROMES method to the multifidelity-correction approach characterized by employing the model inputs as error indicators. Section 5.4 compares the ROMES stochastic error estimate to the error bound given by the reduced-basis method. Section 5.5 compares the performance of the two machine-learning algorithms: the Gaussian process kernel method and the relevance vector machine. Finally, Section 5.6 considers non-compliant and multiple output functionals, which ROMES handles via dual-weighted-residual error indicators.

5.1. **Problem setup.** Consider a finite-element model of heat transport on a square domain $\Omega := \cup_{i=1}^{9} \Omega_i$ composed of nine parameterized materials. The block is cooled along the top boundary to a reference temperature of zero, a nonzero heat flux is specified on the bottom boundary, and the leftmost boundary is adiabatic. The compliant output functional for this problem is defined as the integral over the Neumann domain Γ_{N_1}

$$\bar{g}(u(\boldsymbol{\mu})) = \int_{\Gamma_{N_1}} u(\boldsymbol{\mu}) \mathrm{d}x, \quad \boldsymbol{\mu} \in \mathcal{P}, \tag{40}$$

where the parameter domain is set to $\mathcal{P} = [0.1, 10]^9$ and u is the continuous representation of the finite-element solution. The state variable $u(\boldsymbol{\mu}) \in X = H_0^1 := \left\{ w \in H^1(\Omega) \mid w|_{\Gamma_D} = 0 \right\}$ fulfills the weak form of the parameterized Poisson's equation: find $u(\boldsymbol{\mu}) \in X$, such that

$$a(u(\boldsymbol{\mu}), v) = f(v) \quad \text{for all } v \in X. \tag{41}$$

Here, the bilinear form $a : X \times X \to X$ and the functional $f : X \to X$ are defined as

$$a(u, v) := \int_{\Omega} b(x; \boldsymbol{\mu}) \nabla u(\boldsymbol{\mu}) \cdot \nabla v \, \mathrm{d}x, \qquad f(v) := \int_{\Gamma_{N_1}} v \, \mathrm{d}x \tag{42}$$

with boundary conditions

$$\nabla b(x; \boldsymbol{\mu}) u(\boldsymbol{\mu}) \cdot n = 0 \quad \text{on } \Gamma_{N_0}, \qquad \nabla b(x; \boldsymbol{\mu}) u(\boldsymbol{\mu}) \cdot n = 1 \quad \text{on } \Gamma_{N_1}. \tag{43}$$

We define the coefficient function $b : \Omega \times \mathcal{P} \to \mathbb{R}$ as

$$b(x; \boldsymbol{\mu}) = \sum_{i=1}^{9} \mu_i \mathbb{1}_{\Omega_i}(x), \tag{44}$$

where μ_i denotes the ith component of the parameter vector $\boldsymbol{\mu}$, and the indicator function $\mathbb{1}_A(x) = 1$ if $x \in A$ and is zero otherwise. Figure 3 depicts the composition of the domain and the location of the boundary conditions.

By replacing the infinite-dimensional function space X with the (finite) n-dimensional finite-element space $X_\mathrm{h} \subset X$ in problem (41), one can compute the parameter-dependent state function $u_\mathrm{h}(\boldsymbol{\mu}) \in X_\mathrm{h}$ represented by vectors containing the function's degrees of freedom $\boldsymbol{u}(\boldsymbol{\mu}) \in \mathbb{R}^n$ (see Section A.1). In the experiments, the domain is discretized by triangular finite elements, which results in a finite-element space X_h of dimension $n = 10^4$. The high-fidelity output (in the notation of Section 2.1) is then $g(\boldsymbol{u}(\boldsymbol{\mu})) := \bar{g}(u_\mathrm{h}(\boldsymbol{\mu}))$, $\boldsymbol{\mu} \in \mathcal{P}$.

As described in Section A.4, we employ a greedy algorithm[4] to generate a reduced-basis space $X_\mathrm{red} \subset X_\mathrm{h}$ of dimension $p \ll n$. The algorithm employs a training set of 100 randomly selected points (i.e., $\mathrm{card}\,(\mathcal{P}_\mathrm{greedy}) = 100$ in Section A.4), until the maximum computed error bound in the training set is less than 1; it stops after $p = 11$ iterations.

Replacing X_h with X_red in Eq. (41) leads to reduced state functions $u_\mathrm{red}(\boldsymbol{\mu}) \in X_\mathrm{red}$ for all $\boldsymbol{\mu} \in \mathcal{P}$. As before, these solutions can be represented by vectors $\boldsymbol{u}_\mathrm{red}(\boldsymbol{\mu}) := \boldsymbol{V}\hat{\boldsymbol{u}}(\boldsymbol{\mu}) \in \mathbb{R}^n$, where $\boldsymbol{V} \in \mathbb{R}^{n \times p}$ is the discrete representation of a basis for the function space X_red.

In the following, we analyze two types of error: (i) the energy norm of the state-space error $|||\boldsymbol{\delta_u}||| = |||\boldsymbol{u} - \boldsymbol{u}_\mathrm{red}||| := a(u_\mathrm{h} - u_\mathrm{red}, u_\mathrm{h} - u_\mathrm{red})$ and (ii) the output error $\delta_s = g(u_\mathrm{h}) - g(u_\mathrm{red})$. Because the output functional in this case is compliant (i.e., $g = f$ and a is symmetric), the output error is always non-negative; see Eq. (70) of Section A.3. For more details regarding the finite-element discretization, the reduced-order-model generation, and error bounds, consult Section A of the Supplementary Materials.

5.2. **ROMES implementation and validation.** We first compute ROMES surrogates for the two errors $|||\boldsymbol{\delta_u}|||$ and (compliant) δ_s. As proposed in Section 3.2.1, the three ROMES ingredients we employ are: 1) log-residual-norm error indicators $\boldsymbol{\rho}(\boldsymbol{\mu}) = \log(r(\boldsymbol{\mu}))$, 2) a logarithmic transformation function $d = \log$, and 3) both the GP kernel and the RVM supervised machine-learning methods. To train the surrogates, we compute $|||\boldsymbol{\delta_u}(\boldsymbol{\mu})|||$, $\delta_s(\boldsymbol{\mu})$, and $\boldsymbol{\rho}(\boldsymbol{\mu})$ for $\boldsymbol{\mu} \in \bar{\mathcal{P}} \subset \mathcal{P}$ with $\mathrm{card}\left(\bar{\mathcal{P}}\right) = 2000$. The first $N = 100$ points comprise the ROMES training set $\{\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_N\} =: \mathcal{P}_\mathrm{learn} \subset \bar{\mathcal{P}}$ and the following 1900 points define a validation set $\mathcal{P}_\mathrm{validation} \subset \bar{\mathcal{P}}$; note that the validation set was not used to construct the error surrogates. Reported results relate to statistics computed over this validation set.

For the kernel method, we employ the squared exponential covariance kernel (30). For the RVM method, we choose Legendre polynomials P_k as basis functions, as we expect a linear relationship between the indicators and true errors (see Section 4.2). Because Legendre polynomials are defined on the interval $[-1, 1]$, we must transform and scale this domain to span the possible range of indicator values. For this purpose, we apply the heuristic of setting the domain of the polynomials to be 20% larger than the interval bounded by the smallest

[4] All reduced-basis computations are conducted with the reduced-basis library RBMatlab (`http://www.morepas.org/software/`).

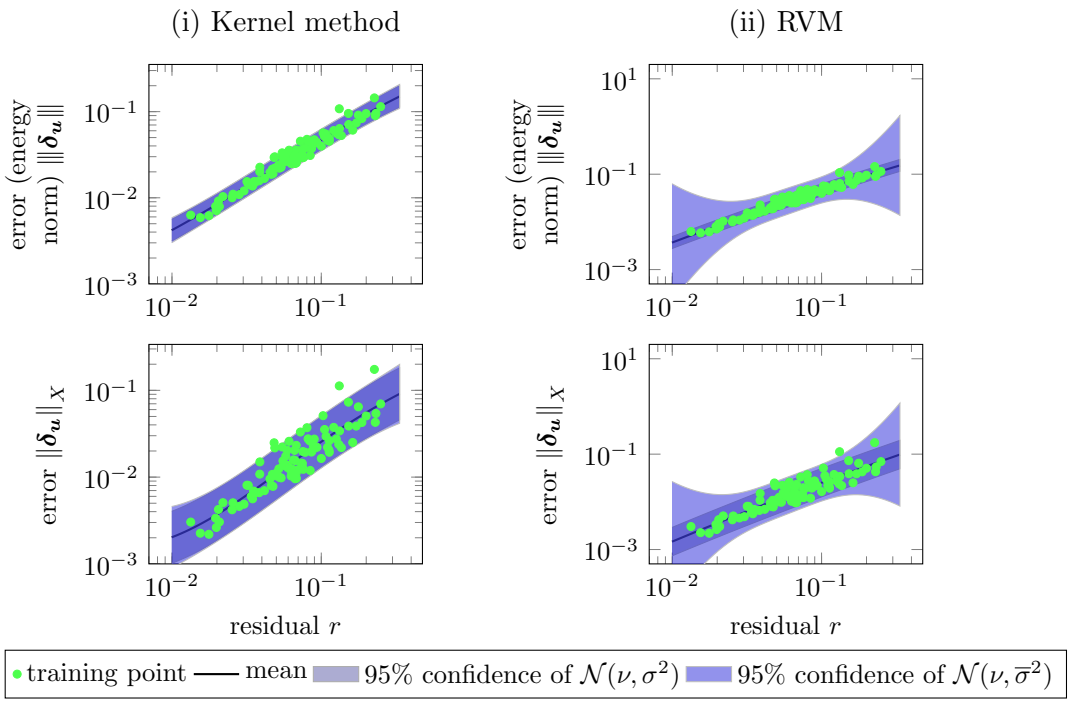

FIGURE 4. Visualization of ROMES surrogates ($\delta = |||\boldsymbol{\delta_u}|||$ and $\|\boldsymbol{\delta_u}\|_X$, $\boldsymbol{\rho} = \log r$, $d = \log$), computed using $N = 100$ training points and the (i) GP kernel method and (ii) RVM.

and largest indicator values:

$$[\boldsymbol{\rho}_{\min} - 0.1(\boldsymbol{\rho}_{\max} - \boldsymbol{\rho}_{\min}), \boldsymbol{\rho}_{\max} + 0.1(\boldsymbol{\rho}_{\max} - \boldsymbol{\rho}_{\min})]\,, \tag{45}$$

where $\boldsymbol{\rho}_{\min} = \min_{\boldsymbol{\mu}\in\mathcal{P}_{\text{learn}}} \boldsymbol{\rho}(\boldsymbol{\mu})$ and $\boldsymbol{\rho}_{\max} = \max_{\boldsymbol{\mu}\in\mathcal{P}_{\text{learn}}} \boldsymbol{\rho}(\boldsymbol{\mu})$. We include Legendre polynomials of orders 0 to 4; however, the RVM method typically discards the higher order polynomials due to the near-linear relation between indicators and errors.

Figure 4 depicts the ROMES surrogate $\widetilde{|||\boldsymbol{\delta_u}|||}$ generated by both machine-learning methods using all 100 training points. For comparison, we also create ROMES surrogates $\widetilde{\|\boldsymbol{\delta_u}\|_X}$ for errors in the parameter-independent norm $\|\cdot\|_X$ of the state space $X = H_0^1$. In addition to the expected mean of the inferred surrogate, the figure displays two 95%-confidence intervals for the prediction (see Remark 4.1):

- The darker shaded interval corresponds to the confidence interval arising from the inherent uncertainty in the error due to the non-uniqueness of the mapping $\boldsymbol{\rho} \mapsto |||\boldsymbol{\delta_u}|||$, i.e., the inferred variance σ^2 of Eq. (32).
- The lighter shaded interval also includes the 'uncertainty in the mean' due to a lack of training data, i.e., Σ of Eq. (32). With an increasing number of training points, this area should be indistinguishable from the darker one.

All ROMES models find a linear trend between the indicators and the errors, where the variance is slightly larger for the parameter-independent norm. This larger variance can be attributed to the larger range of the coercivity constants the parameter-independent norm (see Section A.3). For this example, however, both ROMES are functional. In the following examples, we focus on the energy norm only.

Note that the 'uncertainty in the mean' is dominant for the RVM surrogate. This can be explained as follows: the high-order polynomials have values close to zero near the mean of the data. As such, the training data are not very informative for the coefficients of these polynomials. This results in a large inferred variance for those coefficients. Section 5.5 further compares the two machine-learning methods; due to its superior performance, we now proceed with the kernel method.

We now assess the validity of the Gaussian-process assumptions underlying the ROMES surrogates $\widetilde{|||\boldsymbol{\delta}_{\boldsymbol{u}}|||}$ and $\widetilde{\delta}_s$, i.e., Condition 3 of Section 3.1. From the discussion in Remark 4.1, we know if the underlying GP model form is correct, then as the number of training points increases, the uncertainty about the mean decreases and the set $\{D(\boldsymbol{\mu}) \mid \boldsymbol{\mu} \in \mathcal{P}_{\text{validation}}\}$ with

$$D(\boldsymbol{\mu}) := d\left(|||\boldsymbol{\delta}_{\boldsymbol{u}}(\boldsymbol{\mu})|||\right) - \mathrm{E}\left[d\left(\widetilde{|||\boldsymbol{\delta}_{\boldsymbol{u}}|||}(\boldsymbol{\rho}(\boldsymbol{\mu}))\right)\right] = d\left(|||\boldsymbol{\delta}_{\boldsymbol{u}}(\boldsymbol{\mu})|||\right) - \nu\left(\boldsymbol{\rho}(\boldsymbol{\mu})\right) \tag{46}$$

should behave like samples from the distribution $\mathcal{N}(0, \sigma^2)$. Figure 5 reports this validation test and verifies that this condition does indeed hold for a sufficiently large number of training points.

Further, we can validate the inferred confidence intervals as proposed in Eq. (9). The table within Figure 5 reports $\omega_{\text{validation}}(\omega)$ (see Eq. (10)), which represents the frequency of observed predictions in the validation set that lie within the inferred confidence interval C_ω. We declare the ROMES model to be validated, as $\omega_{\text{validation}}(\omega) \approx \omega$ for several values of ω as the number of training points increases.

The results for the ROMES surrogate $\widetilde{\delta}_s$ are very similar to those presented in Figure 5 and will be further discussed in Section 5.3. Note that the inferred Gaussian process is well-converged with a moderately sized training set consisting of only $N = 35$ points.

5.3. **Output error: comparison with multifidelity correction.** As discussed in Section 3.2, multifidelity-correction methods construct a surrogate $\widetilde{\delta_{s,\text{MF}}}$ of the output error using the system inputs as error-surrogate inputs, i.e., $\delta = \delta_s$, $\boldsymbol{\rho} = \boldsymbol{\mu}$, and $d = \mathrm{id}_{\mathbb{R}}$. In this section, we construct this multifidelity correction surrogate using the same GP kernel method as ROMES. Ref. [35] demonstrated that this error surrogate fails to improve the 'corrected output' when the low-fidelity model corresponds to a reduced-order model. We now verify this result and show that—in contrast to the multifidelity correction approach—the ROMES surrogate $\widetilde{\delta}_s$ constructed via the GP kernel method with $\delta = \delta_s$, $\boldsymbol{\rho} = \log r$, and $d = \log$ yields impressive results: on average, the output 'corrected' by the ROMES surrogate reduces the error by an order of magnitude, and the Gaussian-process assumptions can be validated. The validation quality improves as the number of training points increases, but a moderately sized set of only $N = 20$ training points leads to a converged surrogate.

The reason multifidelity correction fails for most reduced-order models is twofold. First, the mapping $\boldsymbol{\mu} \mapsto \delta_s$ can be highly oscillatory in the input space. This behavior arises from the fact the the reduced-order model error is zero at the (greedily-chosen) ROM training points but

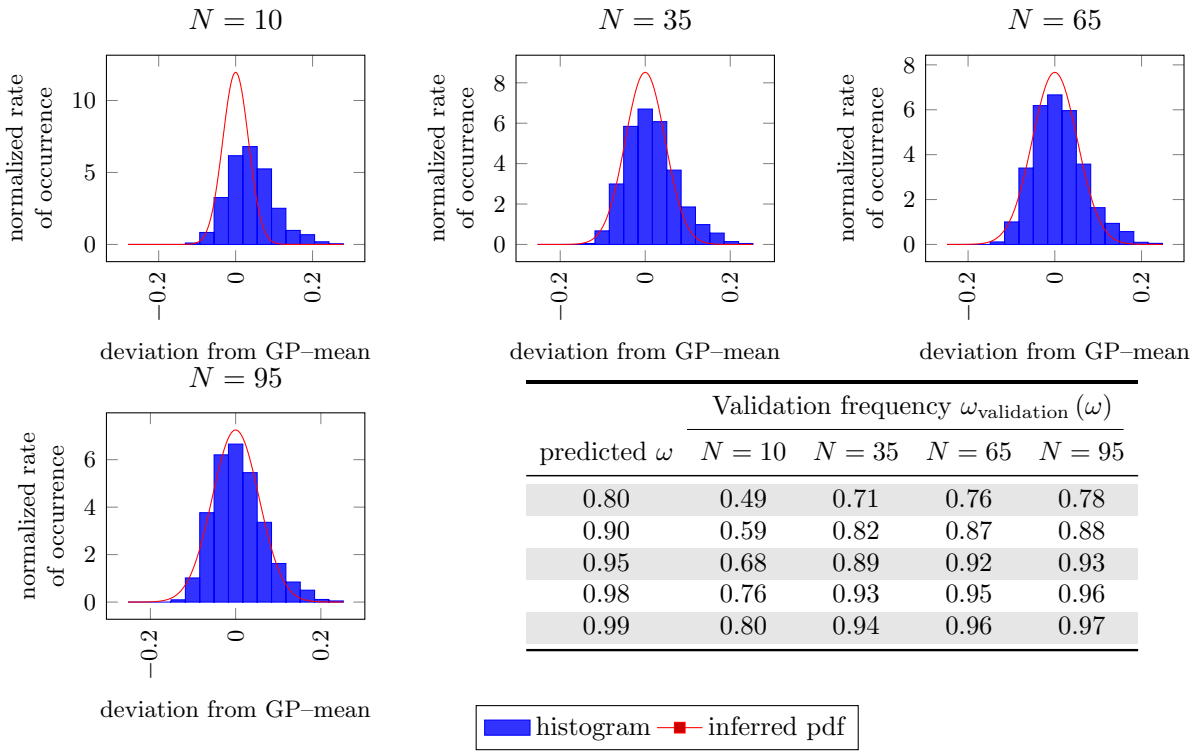

predicted ω	Validation frequency $\omega_{\text{validation}}$ (ω) $N = 10$	$N = 35$	$N = 65$	$N = 95$
0.80	0.49	0.71	0.76	0.78
0.90	0.59	0.82	0.87	0.88
0.95	0.68	0.89	0.92	0.93
0.98	0.76	0.93	0.95	0.96
0.99	0.80	0.94	0.96	0.97

FIGURE 5. Gaussian-process validation for the ROMES surrogate (GP kernel, $\delta = |||\boldsymbol{\delta_u}|||$, $\boldsymbol{\rho} = \log r$, $d = \log$) with a varying number of training points N. The histogram corresponds to samples of $D(\boldsymbol{\mu})$ and the red curve depicts the probability distribution function $\mathcal{N}(0, \sigma^2)$. The table reports how often the actual error lies in the inferred confidence intervals.

grows (and can grow quickly) away from these points. Such complex behavior requires a large number of error-surrogate training points to accurately capture. In addition, the number of system inputs is often large (in this case nine); this introduces curse-of-dimensionality difficulties in modeling the error. Figure 6(ii) visualizes this problem. The depicted mapping between the first two parameter components μ_1, μ_2 and the output error $\delta_s(\boldsymbol{\mu})$ displays no structured behavior. As a result, there is no measurable improvement of the corrected output $s_{\text{red}} + \widetilde{\delta_{s,\text{MF}}}$ over the evaluation of the ROM output s_{red} alone.

In order to quantify the performance of the error surrogates, we introduce a normalized *expected improvement*

$$I(\widetilde{\delta}, \boldsymbol{\mu}) := \left| \frac{\delta_s(\boldsymbol{\mu}) - \text{mode}\left(\widetilde{\delta}(\boldsymbol{\rho}(\boldsymbol{\mu}))\right)}{\delta_s(\boldsymbol{\mu})} \right|. \tag{47}$$

If this value is less than one, then the exected corrected output $s_{\text{red}} + \widetilde{\delta}$ is more accurate than the ROM output s_{red} itself for point $\boldsymbol{\mu} \in \mathcal{P}$, i.e., the additive error surrogate *improves* the

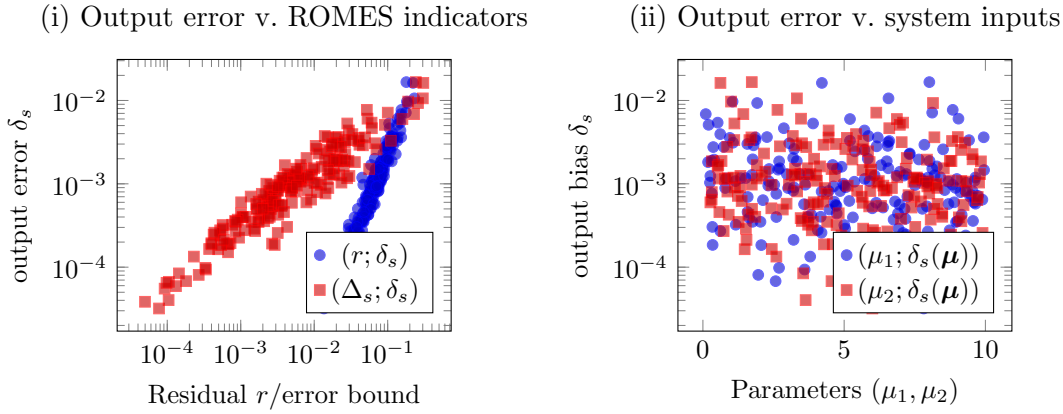

FIGURE 6. Relationship between (i) ROMES error indicators and the compliant-output error and (ii) the first two parameter components and the (compliant) output error, visualized by evaluation of 200 random sample points in the input space. Clearly, the observed structure in the former relationship is more amenable to constructing a Gaussian process.

prediction of the ROM. On the other hand, values above one indicate that the error surrogate *worsens* the ROM prediction.

Figure 7 reports the mean, median, standard deviations, and extrema for the expected improvement (47) evaluated for all validation points $\mathcal{P}_{\text{validation}}$ and a varying number of training points. Here, we also compare with the performance of the error surrogate $\widetilde{\delta_{\text{uni}}}$, which is defined as a uniform distribution on the interval $\left[\Delta_s^{\text{LB}}(\boldsymbol{\mu}), \Delta_s(\boldsymbol{\mu})\right]$, where $\Delta_s^{\text{LB}}(\boldsymbol{\mu})$ and $\Delta_s(\boldsymbol{\mu})$ are the the lower and upper bounds for the output error, respectively. Note that $\widetilde{\delta_{\text{uni}}}$ does not require training data, as it is based purely on error bounds.

The expected improvement for the ROMES output-error surrogate $I(\widetilde{\delta_s}, \boldsymbol{\mu})$ as depicted in Figure 7(i) is approximately 0.2 on average, which constitutes an improvement of nearly an order of magnitude. Further, the maximum expected improvement almost always remains below 1; this implies that the corrected ROM output is almost always more accurate than the ROM output alone.

On the other hand, the expected improvement generated by the error surrogate $\widetilde{\delta_{\text{uni}}}$ is always greater than one, which means that its correction always increases the error. This arises from the fact that the center of the interval $\left[\Delta_s^{\text{LB}}(\boldsymbol{\mu}), \Delta_s(\boldsymbol{\mu})\right]$ is a poor approximation for the true error.

In addition, Figure 7(ii) shows that the expected improvement produced by the multifidelity-correction surrogate $I\left(\widetilde{\delta_{s,\text{MF}}}, \boldsymbol{\mu}\right)$ is often far greater than one. This shows that the multifidelity-correction approach is not well suited for this problem. Presumably, with (far) more training points, these results would improve.

Again, we can validate the Gaussian-process assumptions underlying the error surrogates. For $N = 100$ training points, Figure 8 compares a histogram of deviation of the true error from the surrogate mean to the inferred probability density function. The associated table

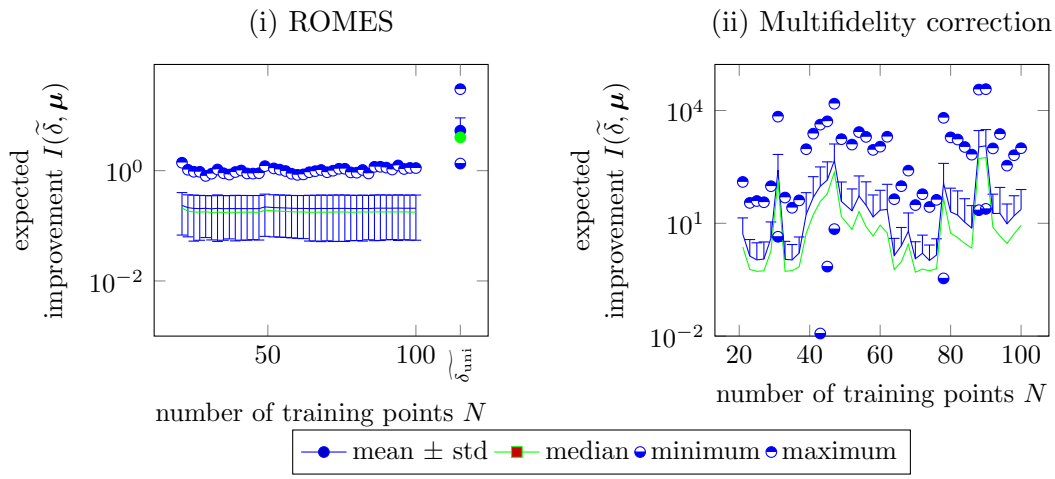

FIGURE 7. Expected improvement $I(\widetilde{\delta}, \boldsymbol{\mu})$ for a varying number of training points N: (i) ROMES (GP kernel, compliant $\delta = \delta_s$, $\boldsymbol{\rho} = \log r$, $d = \log$) with uniform distribution based on reduced-basis error bounds $\widetilde{\delta_{\text{uni}}}$ and (ii) multifidelity correction (GP kernel, compliant $\delta = \delta_s$, $\boldsymbol{\rho} = \boldsymbol{\mu}$, and $d = \text{id}_{\mathbb{R}}$). (1: no improvement, > 1: error worsened, < 1: error improved).

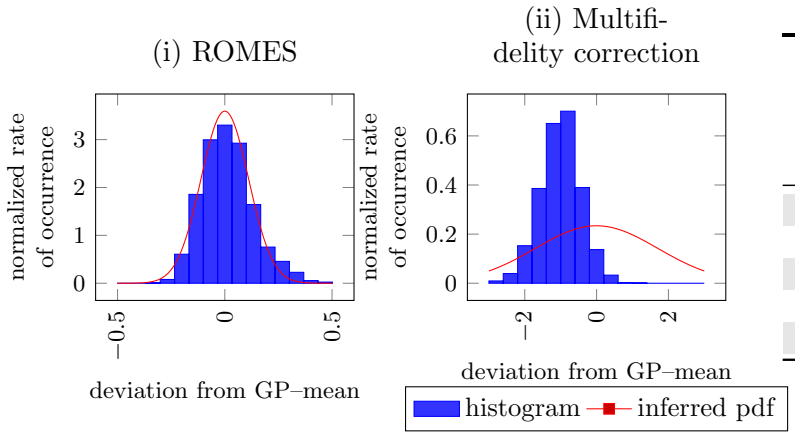

predicted ω	Validation frequency $\omega_{\text{validation}}(\omega)$ (i)	(ii)
0.80	0.78	0.98
0.90	0.89	1.00
0.95	0.93	1.00
0.98	0.96	1.00
0.99	0.97	1.00

FIGURE 8. Gaussian-process validation for the ROMES surrogate (GP kernel, compliant $\delta = \delta_s$, $\boldsymbol{\rho} = \log r$, $d = \log$) and multifidelity-correction surrogate (GP kernel, compliant $\delta = \delta_s$, $\boldsymbol{\rho} = \boldsymbol{\mu}$, and $d = \text{id}_{\mathbb{R}}$) using $N = 100$ training points The histogram corresponds to samples of $D(\boldsymbol{\mu})$ and the red curve depicts the probability distribution function $\mathcal{N}(0, \sigma^2)$. The table reports how often the actual error lies in the inferred confidence intervals. Clearly, this validation test fails for the multilfidelity-correction surrogate.

reports how often the validation data lie in inferred confidence intervals. We observe that the confidence intervals are valid for the ROMES surrogate, but are not for the multifidelity-correction surrogate, as the bins do not align with the inferred distribution. Figure 9 depicts the convergence of these confidence-interval validation metrics as the number of training points increases. As expected (see Remark 4.1) the ROMES observed confidence intervals more closely align with the confidence intervals arising from the *inherent uncertainty* (i.e., σ^2) as the number of training points increases, as this effectively decreases the *uncertainty due to a lack of training*. In addition, only a moderate number of training points (around 20) is required to generate a reasonably converged ROMES surrogate. On the other hand the multifidelity-correction surrogate exhibits no such convergence when fewer than 100 training points are used.

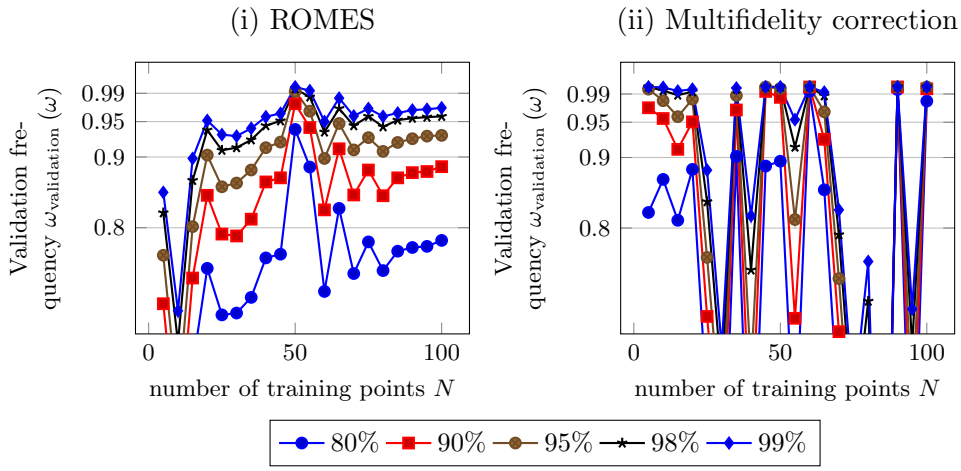

FIGURE 9. Gaussian-process validation for the ROMES surrogate (GP kernel, compliant $\delta = \delta_s$, $\boldsymbol{\rho} = \log r$, $d = \log$) and multifidelity-correction surrogate (GP kernel, compliant $\delta = \delta_s$, $\boldsymbol{\rho} = \boldsymbol{\mu}$, and $d = \mathrm{id}_{\mathbb{R}}$) and a varying number of training points N. The plots depict how often the actual error lies in the inferred confidence intervals.

5.4. **Reduced-basis error bounds.** In this section, we compare the reduced-basis error bound $\Delta_{\boldsymbol{u}}^{\boldsymbol{\mu}}$ (64) with the *probabilistically rigorous* ROMES surrogates $\widetilde{|||\boldsymbol{\delta_u}|||}^{c}$ (8) with rigor values of $c = 0.5$ and $c = 0.9$ as introduced Section 3.3.[5] The ROMES surrogate is constructed with the GP kernel method and ingredients $\delta = |||\boldsymbol{\delta_u}|||$, $\boldsymbol{\rho} = \log r$, and $d = \log$. As discussed in Section 2.3 the error-bound effectivity (8) is important to quantify the performance of these bounds; a value of 1 is optimal, as it implies no over-estimation.

[5]Note that $c = 0.5$ implies no modification to the original ROMES surrogate, as $\widetilde{|||\boldsymbol{\delta_u}|||}^{0.5} = \widetilde{|||\boldsymbol{\delta_u}|||}$ (see Eqs. (25)–(27)).

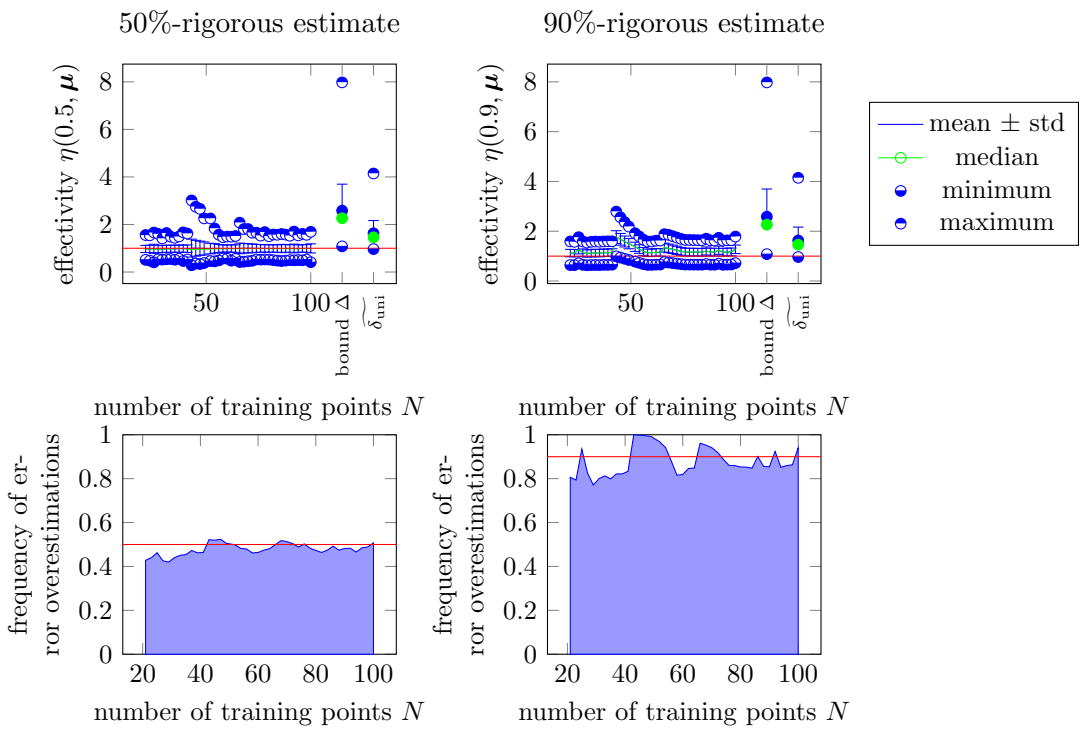

FIGURE 10. Validation of the probabilistically rigorous ROMES surrogates $\widetilde{|||\boldsymbol{\delta_u}|||}^c$ (GP kernel, $\delta = |||\boldsymbol{\delta_u}|||$, $\boldsymbol{\rho} = \log r$, $d = \log$) and comparison with RB error upper bound $\Delta_{\boldsymbol{u}}^{\boldsymbol{\mu}}$ and uniform distribution based on reduced-basis error bounds $\widetilde{\delta_{\text{uni}}}$. The top plots compare statistics of the effectivities $\eta(c, \boldsymbol{\mu})$ with $c = 0.5$ and $c = 0.9$ of the probabilistically rigorous ROMES surrogates with the RB error-bound surrogates. The bottom plots compare the frequency of error overestimation $c_{\text{validation}}$ with the desired value c (red line).

As the probabilistically rigorous ROMES surrogates $\widetilde{|||\boldsymbol{\delta_u}|||}^c$ are stochastic processes, we can measure their (most common) effectivity as

$$\eta(c, \boldsymbol{\mu}) := \frac{\operatorname{mode}\left(\widetilde{|||\boldsymbol{\delta_u}|||}^c(\boldsymbol{\rho}(\boldsymbol{\mu}))\right)}{|||\boldsymbol{\delta_u}(\boldsymbol{\mu})|||}. \tag{48}$$

The top plots of Figure 10 report the mean, median, standard deviation, and extrema of the effectivities $\eta(0.5, \boldsymbol{\mu})$ and $\eta(0.9, \boldsymbol{\mu})$ for all validation points $\boldsymbol{\mu} \in \mathcal{P}_{\text{validation}}$. Again, we compare with $\widetilde{\delta_{\text{uni}}}$, which is a uniform distribution on an interval whose endpoints correspond to the lower and upper bounds for the error $|||\boldsymbol{\delta_u}(\boldsymbol{\mu})|||$. We also compare with the corresponding statistics for the effectivity of the RB error bound $\Delta_{\boldsymbol{u}}^{\boldsymbol{\mu}}$. The lower bound for the coercivity constant that is needed in the RB error bound $\Delta_{\boldsymbol{u}}^{\boldsymbol{\mu}}$ is chosen as the minimum over all parameter components $\alpha_{\text{LB}}(\boldsymbol{\mu}) = \min_{i \in \{1,\dots,9\}} \mu_i$. This simple choice is effective because the example is affinely parameter dependent and linear [37, Ch. 4.2].

We observe that the ROMES surrogate yields better results than both the error bound $\Delta_{\boldsymbol{u}}^{\boldsymbol{\mu}}$ (which produces effectivities roughly between one and eight) and the uniform distribution $\widetilde{\delta_{\text{uni}}}$ (which produces mode effectivities roughly between one and four). The 50%-rigorous ROMES surrogate has an almost perfect mean effectivity of 1 as desired. The 90%-rigorous surrogate has a higher mean effectivity as expected; however, it is only slightly higher. Furthermore, the effectivities of the ROMES surrogate exhibit a much smaller variance[6] than both $\Delta_{\boldsymbol{u}}^{\boldsymbol{\mu}}$ and $\widetilde{\delta_{\text{uni}}}$. Finally, a moderate number (around 20) of training samples is sufficient to obtain well-converged surrogates.

The bottom plots of Figure 10 report the frequency of error overestimation

$$c_{\text{validation}} := \frac{\text{card}\left(\{\boldsymbol{\mu} \in \mathcal{P}_{\text{validation}} \mid \text{median}\left(\tilde{m}\left(\boldsymbol{\rho}(\boldsymbol{\mu})\right)\right) > d(\delta(\boldsymbol{\mu}))\}\right)}{\text{card}\left(\mathcal{P}_{\text{validation}}\right)} \tag{49}$$

for the probabilistically rigorous ROMES surrogates (i.e., $\tilde{m} = \widetilde{|||\boldsymbol{\delta_u}|||}^{c}$) as the number of training points increases to show that $c_{\text{validation}} \approx c$, which validates the rate of error overestimation (see Eq. (23)). Note that the overestimation frequency $c_{\text{validation}}$ converges to its predicted value c, which demonstrates that the rigor of the ROMES estimates can in fact be controlled.

5.5. **Comparison of machine-learning algorithms.** This section compares in detail the ROMES surrogates generated using the two machine-learning methods introduced in Section 4. Recall that Figure 4 visualizes both surrogates. We observe that both methods work well overall and generate well-converged surrogates with a modest number of training samples. As previously mentioned, the GP kernel leads to a smaller inferred variance due to more accurate and localized estimates of the mean. The RVM is characterized by global basis functions that preclude it from accurately resolving localized features of the mean and lead to large uncertainty about the mean (see Figure 4). On the other hand, the confidence intervals computed with the RVM are (slightly) better validated.

Figure 11(i) and (ii) report the validation test, i.e., the frequency of deviations from the inferred mean $D(\boldsymbol{\mu})$ (46) with a training sample containing $N = 80$ training points. We observe a smaller inferred variance σ^2 for the GP kernel method, which implies that the mean is identified more accurately. In both cases, the validation samples align well with the probability density function of the inferred distribution $\mathcal{N}(0, \sigma^2)$.

The confidence intervals of this inferred distribution can be validated, and they turn out to be (slightly) more realistic for the RVM method. The table within Figure 11 shows that the kernel method results are usually optimistic, i.e., the actual confidence intervals are smaller than predicted. This effect can be corrected, however, by adding $\Sigma(\boldsymbol{x}^*)$ as an indicator of the uncertainty of the mean as discussed in Remark 4.1. However, doing so for the RVM prediction yields extremely wide confidence intervals due to the significant uncertainty about the RVM mean (see Figure 4).

As the inferred variance is larger for the relevance vector machine, this also affects the performance of effectivity and improvement measures for the error surrogates. Figure 12 depicts statistics of $\eta(0.5, \boldsymbol{\mu})$ computed with both the methods, and we observe that all statistical measures are significantly better for the kernel method estimate.

[6]The higher variance apparent between 45 and 53 training points can be explained by the fact that the minimization algorithm for the log–likelihood function stops after it exceeds the maximum number of iterations.

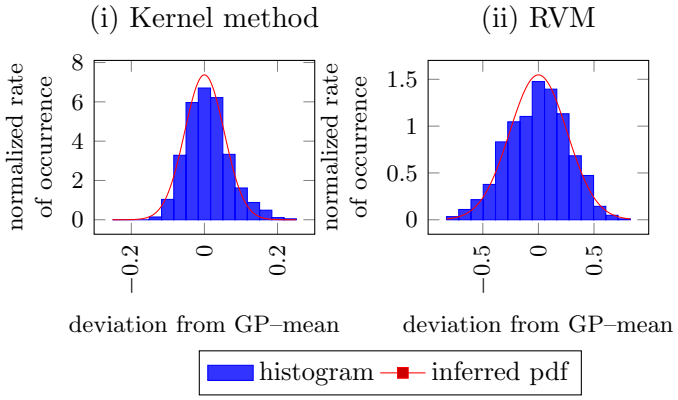

	Validation frequency $\omega_{\text{validation}}(\omega)$			
	σ^2		$\overline{\sigma}^2$	
predicted ω	(i)	(ii)	(i)	(ii)
0.80	0.78	0.78	0.84	0.93
0.90	0.88	0.89	0.93	0.98
0.95	0.93	0.95	0.95	1.00
0.98	0.96	0.97	0.97	1.00
0.99	0.97	0.99	0.99	1.00

FIGURE 11. Gaussian-process validation for the ROMES surrogate using both the GP kernel method and the RVM ($\delta = |||\boldsymbol{\delta_u}|||$, $\boldsymbol{\rho} = \log r$, $d = \log$) using $N = 80$ training points. The histogram corresponds to samples of $D(\boldsymbol{\mu})$ and the red curve depicts the probability distribution function $\mathcal{N}(0, \sigma^2)$. The table reports how often the actual error lies in the inferred confidence intervals.

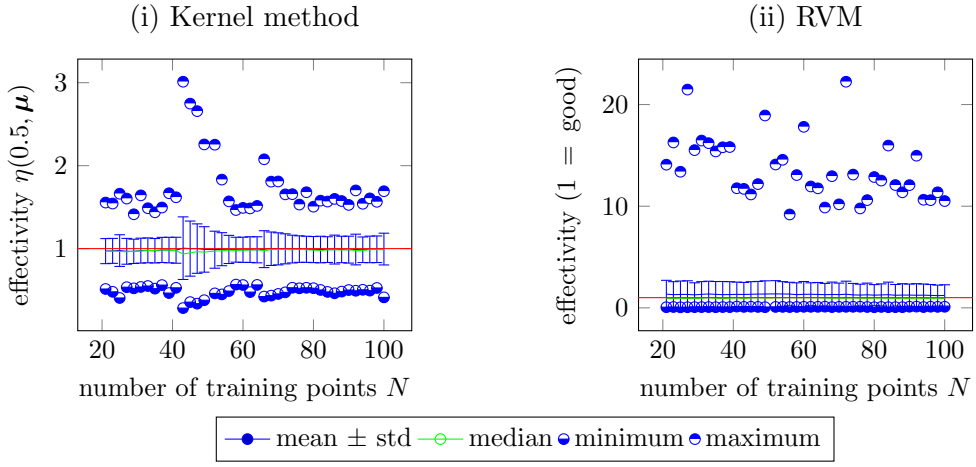

FIGURE 12. Comparison of the effectivity $\eta(0.5, \boldsymbol{\mu})$ of ROMES surrogates ($\delta = |||\boldsymbol{\delta_u}|||$, $\boldsymbol{\rho} = \log r$, $d = \log$) with the GP computed via (i) the GP kernel and (ii) the relevance vector machine method.

We conclude that while both methods produce feasible ROMES surrogates, the GP kernel method produces consistently better results. In particular, the lower inferred variance implies that a lower amount of *epistemic uncertainty* is introduced by the error surrogate (See condition 2 from Section 3.1). This is critically important for many UQ tasks. Therefore, we recommend the GP kernel method to construct ROMES surrogates.

5.5.1. *Dependence on reduced-basis size.* To assess the generalizability of the ROMES method, we apply it to a ROM of higher fidelity, i.e., larger p. We construct two ROMES surrogates: one for the state-space error $\widetilde{|||\boldsymbol{\delta_u}|||}$ (GP kernel, $\delta = |||\boldsymbol{\delta_u}|||$, $\boldsymbol{\rho} = \log r$, $d = \log$) and one for the compliant-output error $\widetilde{\delta_s}$ (GP kernel, $\delta = \delta_s$, $\boldsymbol{\rho} = \log r$, $d = \log$).

We increase the reduced-basis size by decreasing the maximum error over the training set from 1.0 to 1.0×10^{-3} and applying the greedy method. The resulting reduced-basis dimension is $p = 62$. Figure 13 reports the error data and ROMES surrogates. Comparing the leftmost plot of Figure 13 with Figure 4(i) and the rightmost plot of Figure 13 with Figure 6(i) reveals that while the errors are several orders of magnitude smaller for the current experiments, the data (and the resulting ROMES surrogates) exhibit roughly the same structure as before.

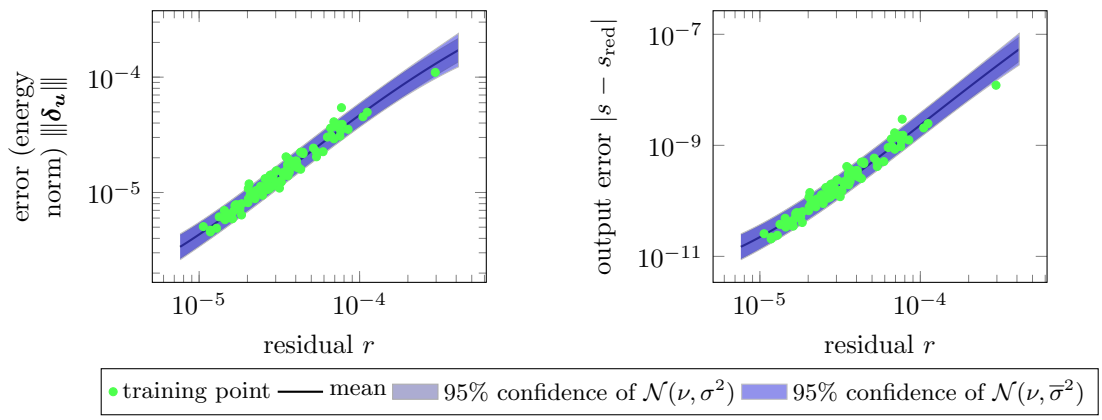

FIGURE 13. Visualization of ROMES surrogates ($\delta = |||\boldsymbol{\delta_u}|||$ and δ_s, $\boldsymbol{\rho} = \log r$, $d = \log$), computed using $N = 100$ training points and the GP kernel method for a higher dimensional ROM with $p = 62$.

Figure 14 reports the performance of these surrogates. Comparing the leftmost plot of Figure 14 with Figure 12(i) shows that the state-space error surrogate exhibits nearly identical convergence for the larger- and smaller-dimension reduced-order models. As in the experiments of Section 5.4, the value of $\operatorname{mode}(\eta(0.5, \boldsymbol{\mu}))$ is close to 1 in the mean. Comparing the rightmost plot of Figure 14 with Figure 7(i) shows that the expected improvement for the output-error correction with the surrogate $\widetilde{\delta_s}$ is around 0.2 in the mean for both the larger- and smaller-dimension reduced-order models. However, for the larger-dimension reduced-order model, more training points are required to reduce the occurrence of improvement factors larger than 1. This is an artifact of the low errors already produced by the larger-dimension ROM itself (i.e., small denominator in Eq. (47)).

We therefore conclude that the ROMES method is applicable to ROMs of different fidelity.

5.6. **Multiple and non-compliant outputs.** Finally, we assess the performance of ROMES on a model with multiple and non-compliant output functionals as discussed in Section 3.2.2.

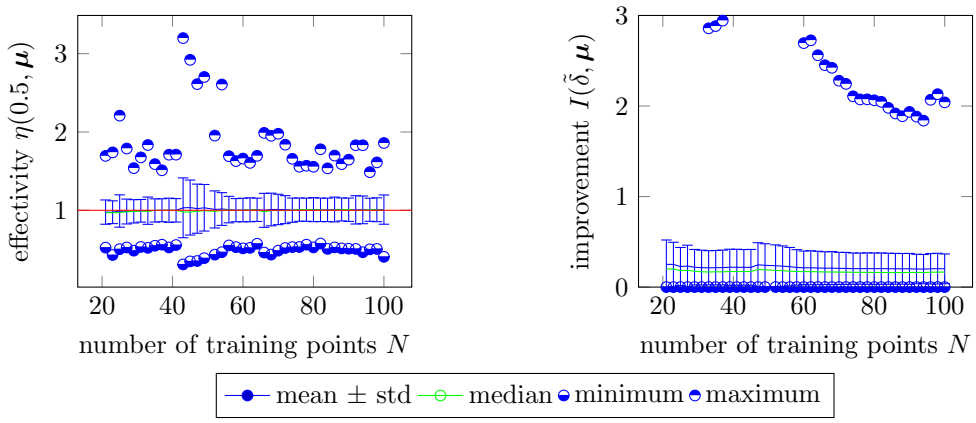

FIGURE 14. Effectivity $\eta(0.5, \boldsymbol{\mu})$ of ROMES surrogate ($\delta = |||\boldsymbol{\delta_u}|||$, $\boldsymbol{\rho} = \log r$, $d = \log$) and expected improvement $I(\widetilde{\delta}, \boldsymbol{\mu})$ of ROMES surrogate (GP kernel, compliant $\delta = \delta_s$, $\boldsymbol{\rho} = \log r$, $d = \log$) for a higher dimensional ROM with $p = 62$ and a varying number of training points N.

For this experiment, we set two outputs to be temperate measurements at points x_1 and x_2:

$$s_i(\boldsymbol{\mu}) := g_i(\boldsymbol{u}(\boldsymbol{\mu})) := \bar{g}_i(u(\boldsymbol{\mu})) = \int_\Omega \delta_{\text{Dirac}}(x - x_i) u(x_i; \boldsymbol{\mu}) \, \mathrm{d}x = u(x_i; \boldsymbol{\mu}), \quad i = 1, 2. \tag{50}$$

where δ_{Dirac} denotes the Dirac delta function. In this case, we construct a separate ROMES surrogates for each output error $\widetilde{\delta_{s_1}}$ and $\widetilde{\delta_{s_2}}$. As previously discussed, we use dual-weighted residuals as indicators $\boldsymbol{\rho}_i(\boldsymbol{\mu}) = \boldsymbol{y}_{\text{red},i}(\boldsymbol{\mu})^t \boldsymbol{r}(\boldsymbol{u}_{\text{red}}; \boldsymbol{\mu})$, $i = 1, 2$ and no transformation $d \equiv \text{id}_{\mathbb{R}}$. This necessitates the computation of approximate dual solutions, for which dual reduced-basis spaces must be generated in the offline stage. The corresponding finite element problem can be found in Eq. (78), where Eq. (50) above provides the right-hand sides. The algebraic problems can be inferred from Eq. (79), where the discrete right-hand sides are canonical unit vectors because the points x_1 and x_2 coincide with nodes of the finite-element mesh.

Like the primal reduced basis, the dual counterpart can be generated with a greedy algorithm that minimizes the approximation error for the reduced dual solutions.

To assess the ability for *uncertainty control* with the dual-weighted-residual indicators (see Remark 3.2) we generate three dual reduced bases of increasing fidelity: 1) error tolerance of 1 (basis sizes p_y of 10 and 11), 2) error tolerance of 0.5 (basis sizes p_y of 15 and 17), 3) error tolerance of 0.1 (basis sizes p_y of 20 and 23).

To train the surrogates, we compute $\delta_{s_1}(\boldsymbol{\mu})$, $\delta_{s_2}(\boldsymbol{\mu})$, $\boldsymbol{\rho}_1(\boldsymbol{\mu})$ (of varying fidelity), $\boldsymbol{\rho}_2(\boldsymbol{\mu})$ (of varying fidelity), for $\boldsymbol{\mu} \in \bar{\mathcal{P}} \subset \mathcal{P}$ with $\text{card}\left(\bar{\mathcal{P}}\right) = 500$. The first $T = 100$ points define the training set $\mathcal{P}_{\text{learn}} \subset \bar{\mathcal{P}}$ and the following 400 points constitute the validation set $\mathcal{P}_{\text{validation}} \subset \bar{\mathcal{P}}$.

Figure 15 depicts the observed relationship between indicators $\boldsymbol{\rho}_1(\boldsymbol{\mu})$ (of different fidelity) and the error in the first output $\delta_{s_1}(\boldsymbol{\mu})$. Note that as the dual-basis size p_y increases, the output error exhibits a nearly exact linear dependence on the dual-weighted residuals. This

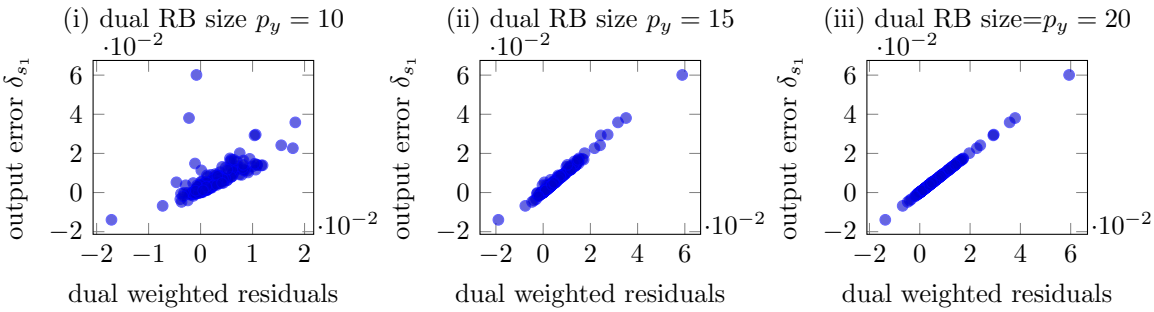

FIGURE 15. Relationship between between dual-weighted-residual indicators $\boldsymbol{\rho}_1 = \boldsymbol{y}_{\mathrm{red},1}(\boldsymbol{\mu})^t \boldsymbol{r}\,(\boldsymbol{u}_{\mathrm{red}}; \boldsymbol{\mu})$ and errors in the (non-compliant) first output δ_{s_1}.

is expected, as the residual operator is linear in the state. Therefore, the RVM with a linear polynomial basis produces the best (i.e., minimum variance) results for the ROMES surrogates in this case.

Figure 16 reflects the necessity of employing a large enough dual reduced basis to compute the dual-weighted-residual error indicators. For a small dual reduced basis, there is almost no improvement in the mean, and only a slight improvement in the median; in some cases, the 'corrected' outputs are actually less accurate. However, the most accurate dual solutions yield a mean and median error improvement of two orders of magnitude. This illustrates the ability and utility of *uncertainty control* when dual-weighted residuals are used as error indicators.

Table 17 reports validation results for the inferred confidence intervals. While the validation results are quite good (and appear to be converging to the correct values), they are not as accurate as those obtained for the compliant output.

6. Conclusions and outlook

This work presented the ROMES method for statistically modeling reduced-order-model errors. In contrast to rigorous error bounds, such statistical models are useful for tasks in uncertainty quantification. The method employs supervised machine learning methods to construct a mapping from existing, cheaply computable ROM error indicators to a *distribution* over the true error. This distribution reflects the epistemic uncertainty introduced by the ROM. We proposed ROMES ingredients (supervised-learning method, error indicators, and transformation function) that yield low-variance, numerically validated models for different types of ROM errors.

For normed outputs, the ROMES surrogates led to effectivities with low variance and means close to the optimal value of one, as well as a notion of probabilistic rigor. This is in contrast to existing ROM error bounds, which exhibited mean effectivities close to ten; this improvement will likely be more pronounced for more complex (e.g., nonlinear, time dependent) problems. Further, the ROMES surrogates were computationally less expensive than the error bounds, as the coercivity-constant lower bound was not required.

For general outputs, the ROMES surrogate allowed the ROM output to be corrected, which yielded a near 10x accuracy improvement. Further, the uncertainty in this error could be *controlled* by modifying the dimension of the dual reduced basis. On the other hand, existing

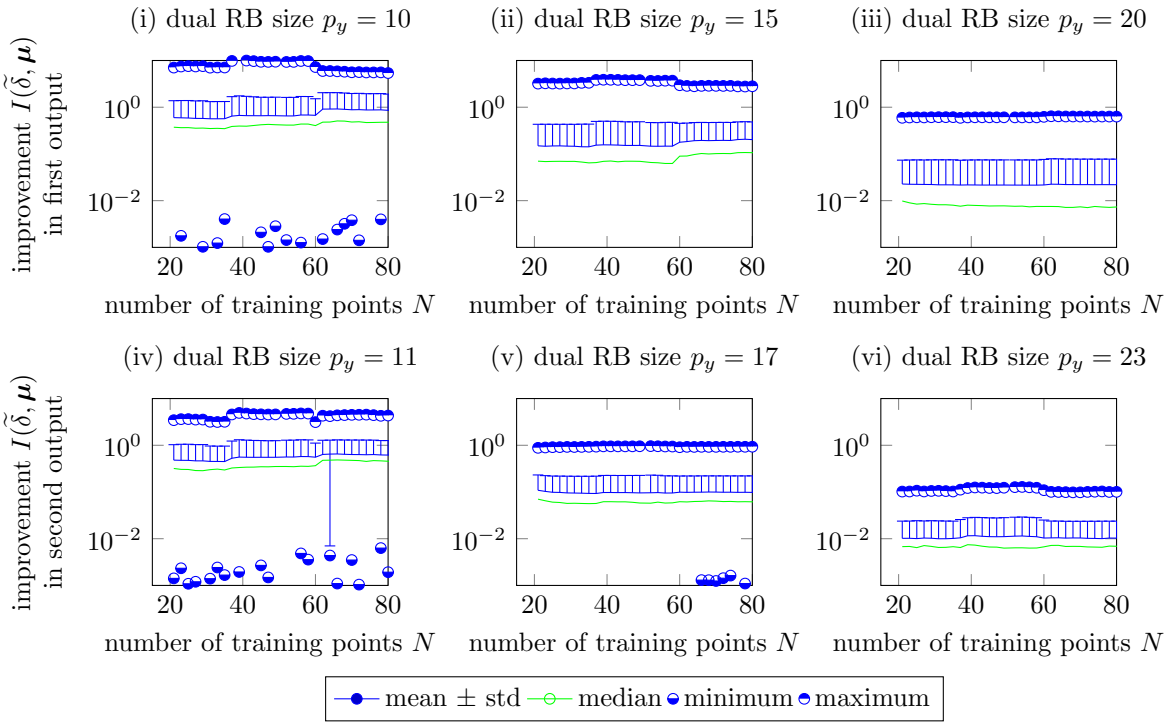

FIGURE 16. Expected improvement $I(\widetilde{\delta}, \boldsymbol{\mu})$ for ROMES surrogate (RVM, $\delta = \delta_s$, $\boldsymbol{\rho}_i = \boldsymbol{y}_{\mathrm{red},i}(\boldsymbol{\mu})^t \boldsymbol{r}(\boldsymbol{u}_{\mathrm{red}}; \boldsymbol{\mu})$, $i = 1, 2$, $d = \mathrm{id}_{\mathbb{R}}$) for a varying number of training points T and different dual reduced-basis-space dimensions. Compare with Figure 7 (1: no improvement, > 1: error worsened, < 1: error improved).

	Validation frequency $\omega_{\mathrm{validation}}$ (ω)							
	first output				second output			
predicted ω	$N = 29$	$N = 53$	$N = 76$	$N = 100$	$N = 29$	$N = 53$	$N = 76$	$N = 100$
0.8	1.00	0.83	0.82	0.82	1.00	0.82	0.87	0.88
0.9	1.00	0.87	0.86	0.86	1.00	0.87	0.91	0.92
0.95	1.00	0.89	0.89	0.89	1.00	0.89	0.93	0.93
0.98	1.00	0.90	0.91	0.90	1.00	0.92	0.94	0.95
0.99	1.00	0.92	0.92	0.91	1.00	0.94	0.95	0.96

FIGURE 17. Gaussian-process validation for the ROMES surrogates (RVM, $\delta = \delta_s$, $\boldsymbol{\rho}_i = \boldsymbol{y}_{\mathrm{red},i}(\boldsymbol{\mu})^t \boldsymbol{r}(\boldsymbol{u}_{\mathrm{red}}; \boldsymbol{\mu})$, $i = 1, 2$, $d = \mathrm{id}_{\mathbb{R}}$). The table reports how often the actual error lies in the inferred confidence intervals. The largest dual reduced-basis space dimensions ($p_y = 20$ and $p_y = 23$) are used to compute the error indicators. .

approaches (i.e., multifidelity correction) that employ system inputs (not error indicators) as inputs to the error model did not lead to improved output predictions. This demonstrated the ability of ROMES to mitigate the curse of dimensionality: although the problem was characterized by nine system inputs, only one error indicator was necessary to construct a low-variance, validated ROMES surrogate.

We foresee the combination of ROMs with ROMES error surrogates to be powerful in UQ applications, especially when the number of system inputs is large. Future work entails integrating and analyzing the ROM/ROMES combination for specific UQ problems, e.g., Bayesian inference, as well as automating the procedure for selecting ROMES ingredients for different problems. Future work will also involve integrating the ROMES surrogates into the greedy algorithm for the reduced-basis-space selection; this has the potential to improve ROM quality due to the near-optimal expected effectivities of the error surrogates.

Acknowledgments

We thank Khachik Sargsyan for his support in the understanding, selection, and development of the machine-learning algorithms. This research was supported in part by an appointment to the Sandia National Laboratories Truman Fellowship in National Security Science and Engineering, sponsored by Sandia Corporation (a wholly owned subsidiary of Lockheed Martin Corporation) as Operator of Sandia National Laboratories under its U.S. Department of Energy Contract No. DE-AC04-94AL85000. We also acknowledge support by the Department of Energy Office of Advanced Scientific Computing Research under contract 10-014804.

Appendix A. Reduced-basis method for parametric elliptic PDEs with affine parameter dependence

This section summarizes the reduced-basis method for generating a reduced-order model for parametric elliptic PDEs with affine parameter dependence. See Ref. [37] for additional details.

A.1. **Parametric elliptic problem.** We start with the discretized version of the elliptic problem, where the state resides in a function space $X_\mathrm{h} \subset X$ of finite dimension $n := \dim(X_\mathrm{h})$, e.g., a finite-element space. For details on the derivation of such a finite element discretization from the analytical formulation of an elliptic PDE and its convergence properties, we refer to Ref. [9].

The problem reads as follows: For any point in the input space $\boldsymbol{\mu} \in \mathcal{P}$, compute a quantity of interest

$$s(\boldsymbol{\mu}) := \bar{g}(u_\mathrm{h}(\boldsymbol{\mu})) \tag{51}$$

with $u_\mathrm{h} \in X_\mathrm{h}$ fulfilling

$$a(u_\mathrm{h}(\boldsymbol{\mu}), v_\mathrm{h}; \boldsymbol{\mu}) = f(v_\mathrm{h}; \boldsymbol{\mu}) \text{ for all } v_\mathrm{h} \in X_\mathrm{h}. \tag{52}$$

Here, for all inputs $\boldsymbol{\mu} \in \mathcal{P}$, $a(\cdot, \cdot; \boldsymbol{\mu}) : X_\mathrm{h} \times X_\mathrm{h} \to X_\mathrm{h}$ is a bilinear form on a Hilbert space X_h and is symmetric, continuous, and coercive, fulfilling

$$\sup_{u \in X} \sup_{v \in X} \frac{a(u, v; \boldsymbol{\mu})}{\|u\|_X \|v\|_X} < \gamma(\boldsymbol{\mu}) < \infty \quad \text{and} \quad \inf_{u \in X, \|u\|_X = 1} a(u, u; \boldsymbol{\mu}) > \alpha(\boldsymbol{\mu}) \tag{53}$$

for constants $\gamma(\boldsymbol{\mu})$ and $\alpha(\boldsymbol{\mu}) > 0$. In addition, we have $f(\boldsymbol{\mu}) \in X_\mathrm{h}^*$.

Given a basis $\{\varphi_i\}_{i=1}^n$ spanning X_h, the solution can be expressed as a linear combination $u_\mathrm{h}(\boldsymbol{\mu}) = \sum_{i=1}^n u_i(\boldsymbol{\mu})\varphi_i \in X_\mathrm{h}$ with the state vector $\boldsymbol{u} := [u_1 \ \cdots \ u_n]^t : \mathcal{P} \to \mathbb{R}^n$ containing the solution degrees of freedom. Then, problem (52) can be solved algebraically by solving the system of linear equations

$$\boldsymbol{r}(\boldsymbol{u};\boldsymbol{\mu}) := \mathbf{A}(\boldsymbol{\mu})\boldsymbol{u}(\boldsymbol{\mu}) - \mathbf{f}(\boldsymbol{\mu}) = 0, \tag{54}$$

with $\mathbf{A} : \mathcal{P} \to \mathbb{R}^{n\times n}$ given by $\mathbf{A}_{ij}(\boldsymbol{\mu}) = a\left(\varphi_i, \varphi_j; \boldsymbol{\mu}\right)$ and $\mathbf{f} : \mathcal{P} \to \mathbb{R}^n$ given by $\mathbf{f}_i(\boldsymbol{\mu}) = f(\varphi_i;\boldsymbol{\mu})$. Similarly, the output can be expressed as a function of the state degrees of freedom:

$$g(\boldsymbol{u}(\boldsymbol{\mu})) := \bar{g}\left(\sum_{i=1}^n u_i(\boldsymbol{\mu})\varphi_i\right). \tag{55}$$

The key concept underlying projection-based model-reduction techniques is to find a lower dimensional subspace $X_\mathrm{red} \subset X_\mathrm{h}$ of dimension $p := \dim(X_\mathrm{red}) \ll n$, such that (54) can be solved more efficiently. For the reduced-basis method, this problem-dependent *reduced-basis space* is obtained through identification of a few training points $\mathcal{P}_\mathrm{train} := \{\tilde{\boldsymbol{\mu}}_i\}_{i=1}^p \subset \mathcal{P}$ for which the full-order equations (54) are solved. These functions then span the reduced basis space $X_\mathrm{red} = \mathrm{span}\,\{u(\tilde{\boldsymbol{\mu}}_i)\}_{i=1}^p$ from which we generate an orthonormal basis $\{\psi_i\}_{i=1}^p$. We refer to subsection A.4 for more details on the selection process of these parameters. Then, these basis vectors can be expressed in terms of the original finite-element degrees of freedom, i.e.,

$$\psi_j = \sum_{i=1}^n \varphi_i \boldsymbol{V}_{ij}, \quad j = 1,\dots,p. \tag{56}$$

This transformation provides the entries of the trial basis $\boldsymbol{V}$. As we consider only Galerkin projection (due to the symmetric and coercive nature of the PDE), we set $\boldsymbol{W} = \boldsymbol{V}$, which leads to a reduced-order problem of the form (4):

$$\boldsymbol{V}^t\boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}};\boldsymbol{\mu}) = \boldsymbol{V}^t\mathbf{A}(\boldsymbol{\mu})\boldsymbol{V}\hat{\boldsymbol{u}}(\boldsymbol{\mu}) - \boldsymbol{V}^t\mathbf{f}(\boldsymbol{\mu})\boldsymbol{V} = 0. \tag{57}$$

The reduced quantity of interest is then given by $s_\mathrm{red} = g(\boldsymbol{V}\hat{\boldsymbol{u}})$.

A.2. **Offline–online decomposition.** The low-dimensional operators in Eq. (57) can be computed efficiently assuming that the bilinear forms a and the functional f are *affinely parameter dependent*, i.e., they can be written as

$$a(\cdot,\cdot;\boldsymbol{\mu}) = \sum_{q=1}^{Q_a} \sigma_a^q(\boldsymbol{\mu})a^q(\cdot,\cdot) \quad \text{and} \quad f(\cdot;\boldsymbol{\mu}) = \sum_{q=1}^{Q_f} \sigma_f^q(\boldsymbol{\mu})f^q(\cdot). \tag{58}$$

with parameter-independent bilinear forms $a^q : X_h \times X_h \to X_h, q = 1,\dots,Q_a$ and functionals $f^q : X_h \to X_h, q = 1,\dots,Q_f$, and coefficient functions $\sigma_a^q : \mathcal{P} \to \mathbb{R}, q = 1,\dots,Q_a$ and $\sigma_f^q : \mathcal{P} \to \mathbb{R}, q = 1,\dots,Q_f$.

In this affinely parameter dependent case, the parameter-dependent reduced quantities $\boldsymbol{V}^t\mathbf{A}(\boldsymbol{\mu})\boldsymbol{V}$ and $\boldsymbol{V}^t\mathbf{f}(\boldsymbol{\mu})\boldsymbol{V}$ can be quickly assembled via a linear combination of their parameter-independent components:

$$\boldsymbol{V}^t\mathbf{A}(\boldsymbol{\mu})\boldsymbol{V} = \sum_{q=1}^{Q_a} \sigma_a^q(\boldsymbol{\mu})\boldsymbol{V}^t\mathbf{A}^q\boldsymbol{V} \tag{59}$$

$$\boldsymbol{V}^t\mathbf{f}\boldsymbol{V} = \sum_{q=1}^{Q_f} \sigma_f^q(\boldsymbol{\mu})\boldsymbol{V}^t\mathbf{f}^q\boldsymbol{V}. \tag{60}$$

Here the matrices $\mathbf{A}^q \in \mathbb{R}^{n\times n}, q = 1,\dots,Q_a$ are given by $\mathbf{A}^q_{ij} = a^q(\varphi_i,\varphi_j)$ and the vectors $\mathbf{f}^q \in \mathbb{R}^n, q = 1,\dots,Q_f$ are given by $\mathbf{f}^q_i = f(\varphi_i)$.

The parameter-independent components

$$\boldsymbol{V}^t\mathbf{A}^q\boldsymbol{V} \in \mathbb{R}^{p\times p},\ q = 1,\dots,Q_a \qquad \text{and} \qquad \boldsymbol{V}^t\mathbf{f}^q\boldsymbol{V} \in \mathbb{R}^p,\ q = 1,\dots,Q_f \tag{61}$$

must be computed only once, during the so-called *offline stage*. All parameter-dependent computations are carried out in the *online stage* in an efficient manner, i.e., with computational complexities that only depend on the reduced-basis dimension p. We refer to Ref. [37] for more details on the offline/online decomposition. Note that in the non-affine case, low-dimensional operators can be approximated using the *empirical interpolation method* [5, 28, 15].

A.3. **Error bounds.** The state error can be measured either in the problem–dependent energy norm or in the norm given by the underlying Hilbert space X_h. The energy norm for the above problem is defined as

$$|||u||| := a(u,u;\boldsymbol{\mu}) \tag{62}$$

and is equivalent to the X_h-norm

$$\alpha \left\|u\right\|^2_{X_\mathrm{h}} \le |||u|||^2 \le \gamma \left\|u\right\|^2_{X_\mathrm{h}}. \tag{63}$$

The bounds for the state errors $\Delta^{\boldsymbol{\mu}}_{\boldsymbol{u}}$ and $\Delta_{\boldsymbol{u}}$ for the errors $|||\boldsymbol{\delta_u}|||$ and $\left\|\boldsymbol{\delta_u}\right\|_{X_\mathrm{h}}$ are defined as follows:

$$\Delta^{\boldsymbol{\mu}}_{\boldsymbol{u}}(\boldsymbol{\mu}) := \frac{r(\boldsymbol{\mu})}{\sqrt{\alpha_{\mathrm{LB}}(\boldsymbol{\mu})}} \ge |||\boldsymbol{\delta_u}(\boldsymbol{\mu})||| \text{ and} \tag{64}$$

$$\Delta_{\boldsymbol{u}}(\boldsymbol{\mu}) := \frac{r(\boldsymbol{\mu})}{\alpha_{\mathrm{LB}}(\boldsymbol{\mu})} \quad \ge \left\|\boldsymbol{\delta_u}(\boldsymbol{\mu})\right\|_{X_\mathrm{h}}, \tag{65}$$

where $\alpha_{\mathrm{LB}}(\boldsymbol{\mu}) \le \alpha(\boldsymbol{\mu})$ is a lower bound for the coercivity constant. This lower bound can easily computed if the bilinear form a is symmetric and *parametrically coercive* as defined in [37, Ch.4.2], i.e., there is a point $\bar{\boldsymbol{\mu}} \in \mathcal{P}$, such that

$$(u,v)_{X_h} = a(u,v;\bar{\boldsymbol{\mu}}), \tag{66}$$

the parameter independent bilinear forms are coercive

$$a^q(w,w) \ge 0 \qquad \forall \boldsymbol{\mu} \in \mathcal{P}, q = 1,\dots,Q_a \tag{67}$$

and their coefficients strictly positive

$$\sigma_a^q(\boldsymbol{\mu}) > 0 \qquad \forall \boldsymbol{\mu} \in \mathcal{P}, q = 1,\dots,Q_a. \tag{68}$$

In this special case, the lower coercivity constant can be chosen as

$$\alpha_{\mathrm{LB}} = \min_{q=1,\dots,Q_a} \frac{\sigma_a^q(\boldsymbol{\mu})}{\sigma_a^q(\bar{\boldsymbol{\mu}})}. \tag{69}$$

A similar bound can be generated for errors in *compliant* output functionals $\bar{g} = f$ (and a symmetric):

$$\Delta_s(\boldsymbol{\mu}) := \frac{r(\boldsymbol{\mu})^2}{\alpha_{\mathrm{LB}}(\boldsymbol{\mu})} \geq \delta_s(\boldsymbol{\mu}) := s(\boldsymbol{\mu}) - s_{\mathrm{red}}(\boldsymbol{\mu}) > 0. \tag{70}$$

Analogously to the upper bounds, lower bounds for the errors can be computed as

$$\Delta_{\boldsymbol{u}}^{LB,\boldsymbol{\mu}}(\boldsymbol{\mu}) := \frac{r(\boldsymbol{\mu})}{\sqrt{\gamma_{\mathrm{UB}}(\boldsymbol{\mu})}} \leq |||\boldsymbol{\delta}_{\boldsymbol{u}}(\boldsymbol{\mu})|||, \tag{71}$$

$$\Delta_{\boldsymbol{u}}^{\mathrm{LB}}(\boldsymbol{\mu}) := \frac{r(\boldsymbol{\mu})}{\gamma_{\mathrm{UB}}(\boldsymbol{\mu})} \quad \leq \|\boldsymbol{\delta}_{\boldsymbol{u}}(\boldsymbol{\mu})\|_{X_{\mathrm{h}}} \text{ and} \tag{72}$$

$$\Delta_s^{\mathrm{LB}}(\boldsymbol{\mu}) := \frac{r(\boldsymbol{\mu})^2}{\gamma_{\mathrm{UB}}(\boldsymbol{\mu})} \quad \leq \delta_s(\boldsymbol{\mu}). \tag{73}$$

where $\gamma_{\mathrm{UB}} \geq \gamma$ is an upper bound for the continuity constant γ.

A.3.1. *Offline–online decomposition of the residual norm.* In order to compute the error bounds efficiently, the residual norm $r(\boldsymbol{\mu}) := \|\boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}(\boldsymbol{\mu});\boldsymbol{\mu})\|_2$ must be computed efficiently through an offline/online decomposition:

$$\begin{aligned}
\|\boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}(\boldsymbol{\mu});\boldsymbol{\mu}\|_2^2 &= (\boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}(\boldsymbol{\mu});\boldsymbol{\mu}))^t \boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{u}}(\boldsymbol{\mu});\boldsymbol{\mu}) \\
&= \hat{\boldsymbol{u}}(\boldsymbol{\mu}) \left(\boldsymbol{V}^t \mathbf{A}^t(\boldsymbol{\mu})\mathbf{A}(\boldsymbol{\mu})\boldsymbol{V}\right) \hat{\boldsymbol{u}}(\boldsymbol{\mu}) - 2\left(\boldsymbol{V}^t \mathbf{f}^t(\boldsymbol{\mu})\mathbf{A}(\boldsymbol{\mu})\boldsymbol{V}\right)\hat{\boldsymbol{u}}(\boldsymbol{\mu}) + \left(\boldsymbol{V}^t\mathbf{f}^t(\boldsymbol{\mu})\mathbf{f}(\boldsymbol{\mu})\boldsymbol{V}\right) \\
&= \sum_{1\leq q,q'\leq Q_a} \sigma_a^q(\boldsymbol{\mu})\sigma_a^{q'}(\boldsymbol{\mu}) \;\; \hat{\boldsymbol{u}}(\boldsymbol{\mu}) \left(\boldsymbol{V}^t \left(\mathbf{A}^q\right)^t \mathbf{A}^{q'}\boldsymbol{V}\right)\hat{\boldsymbol{u}}(\boldsymbol{\mu}) \\
&\quad -2\sum_{q=1}^{Q_f}\sum_{q'=1}^{Q_a} \sigma_f^q(\boldsymbol{\mu})\sigma_a^{q'}(\boldsymbol{\mu}) \;\; \left(\boldsymbol{V}^t \left(\mathbf{f}^q\right)^t \mathbf{A}^{q'}\boldsymbol{V}\right)\hat{\boldsymbol{u}}(\boldsymbol{\mu}) \\
&\quad + \sum_{1\leq q,q'\leq Q_f} \sigma_f^q(\boldsymbol{\mu})\sigma_f^{q'}(\boldsymbol{\mu}) \;\; + \left(\boldsymbol{V}^t \left(\mathbf{f}^q\right)^t \mathbf{f}^{q'}\boldsymbol{V}\right)
\end{aligned} \tag{74}$$

All low-dimensional and parameter-independent matrices

$$\left(\boldsymbol{V}^t \left(\mathbf{A}^q\right)^t \mathbf{A}^{q'}\boldsymbol{V}\right) \in \mathbb{R}^{p\times p}, \qquad 1 \leq q, q' \leq Q_a \tag{75}$$

$$\left(\boldsymbol{V}^t \left(\mathbf{f}^q\right)^t \mathbf{A}^{q'}\boldsymbol{V}\right) \in \mathbb{R}^{p}, \qquad 1 \leq q, q' \leq Q_f, Q_a \tag{76}$$

$$\left(\boldsymbol{V}^t \left(\mathbf{f}^q\right)^t \mathbf{f}^{q'}\boldsymbol{V}\right) \in \mathbb{R}, \qquad 1 \leq q, q' \leq Q_f \tag{77}$$

can be pre-computed during the *offline stage*.

A.3.2. *Dual weighted error estimates.* As discussed in Section 3.2.2, the ROMES surrogate for modeling errors in general (non-compliant) outputs $\bar{g}$ requires a dual solution $\boldsymbol{y}(\boldsymbol{\mu}) \in \mathbb{R}^n$. In the present context, the associated problem is: Find $y_{\mathrm{h}}(\boldsymbol{\mu}) \in X_h$ fulfilling

$$a(v_{\mathrm{h}}, y_{\mathrm{h}}(\boldsymbol{\mu}); \boldsymbol{\mu}) = -\bar{g}(v_{\mathrm{h}}) \qquad \text{for all } v_{\mathrm{h}} \in X_h. \tag{78}$$

Analogously to the discussion for the primal problem, we require a reduced-basis space $X_{\mathrm{red}}^{\bar{g}} \subset X_h$ for this dual problem. Algebraically, this leads to a reduced-basis matrix $\boldsymbol{Y} \in \mathbb{R}^{n \times p_y}$ associated for a reduced dual solution $\hat{\boldsymbol{y}} \in \mathbb{R}^{p_y}$ such that

$$\boldsymbol{Y}^T \boldsymbol{r}_g(\boldsymbol{Y}\hat{\boldsymbol{y}}(\boldsymbol{\mu}); \boldsymbol{\mu}) = 0, \qquad \text{with} \qquad \boldsymbol{r}_g(\boldsymbol{y}; \boldsymbol{\mu}) := \mathbf{A}(\boldsymbol{\mu})\boldsymbol{y} + \mathbf{g}. \tag{79}$$

Here, $\mathbf{g} \in \mathbb{R}^n$ is given by $(\mathbf{g})_i = \bar{g}(\varphi_i)$. Then, this error can be bounded as

$$\Delta_s(\boldsymbol{\mu}) := \frac{r(\boldsymbol{\mu}) r_g(\boldsymbol{\mu})}{\alpha_{\mathrm{LB}}(\boldsymbol{\mu})} \geq |\delta_s(\boldsymbol{\mu})|, \tag{80}$$

with the dual residual norm $r_g(\boldsymbol{\mu}) := \|\boldsymbol{r}_g(\boldsymbol{Y}\hat{\boldsymbol{y}}(\boldsymbol{\mu}); \boldsymbol{\mu})\|_2$.

A.4. **Basis generation with greedy algorithm.** We have so far withheld the mechanism by which the reduced basis spaces $X_{\mathrm{red}}, X_{\mathrm{red}}^g \subset X_{\mathrm{h}}$ are constructed. This section provides a brief overview on this topic, and to ease exposition, we will focus on the generation of the primal reduced-basis space only. As previously mentioned, the reduced-basis space is constructed from p chosen input-space points $\mathcal{P}_{\mathrm{train}} \subset \mathcal{P}$ with $\operatorname{card}(\mathcal{P}_{\mathrm{train}}) = p$. As the reduced-basis solutions should provide accurate approximations for all points $\boldsymbol{\mu} \in \mathcal{P}$, the ultimate goal is to find a reduced-basis space of low dimension $p \ll n$ that minimizes some distance measure between itself and the manifold $\mathcal{M} := \{u(\boldsymbol{\mu}) \mid \boldsymbol{\mu} \in \mathcal{P}\}$. Two possible candidates for such a distance measure are the maximum projection error

$$\operatorname{dist}_1(X_{\mathrm{red}}, \mathcal{M}) := \max_{u \in \mathcal{M}} \inf_{v \in X_{\mathrm{red}}} \|u - v\| \tag{81}$$

and the maximum reduced state error

$$\operatorname{dist}_2(X_{\mathrm{red}}, \mathcal{M}) := \max_{u(\boldsymbol{\mu}) \in \mathcal{M}} \|u(\boldsymbol{\mu}) - u_{\mathrm{red}}(\boldsymbol{\mu})\|. \tag{82}$$

The optimally achievable distance of a reduced-basis space from the manifold $\mathcal{M}$ is known as the Kolmogorov N-width and is given by

$$d_N(\mathcal{M}) := \inf_{\tilde{X} \subset X, \dim(\tilde{X}) = N} \operatorname{dist}_1(\tilde{X}, \mathcal{M}). \tag{83}$$

While the Kolmogorov N-width is a theoretical measure and is only known for a few simple manifolds, it is possible to construct a reduced-basis space with a myopic *greedy algorithm* that often comes close to this optimal value. The greedy algorithm works as follows: First, the manifold of 'interesting solutions' is discretized by defining a finite subset $\mathcal{M}_{\mathrm{greedy}} := \{u(\boldsymbol{\mu}) \mid \boldsymbol{\mu} \in \mathcal{P}_{\mathrm{greedy}}\} \subset \mathcal{M}$, with $\operatorname{card}(\mathcal{P}_{\mathrm{greedy}}) \geq p$ for which the distance measures (81) and (82) can be computed in theory. However, as computing these distance measures exactly requires knowledge of the full-order solution u, the method substitutes these distance measures by error bounds $\max_{\boldsymbol{\mu} \in \mathcal{P}_{\mathrm{greedy}}} \Delta_{\boldsymbol{u}}(\boldsymbol{\mu})$. This allows for the construction of a sequence of reduced-basis spaces $X_1 \subset X_2 \subset \cdots \subset X_p =: X_{\mathrm{red}}$ by choosing the first subspace arbitrarily, as

$$X_1 = \operatorname{span}\{u(\boldsymbol{\mu}_1)\} \tag{84}$$

with a randomly chosen parameter $\boldsymbol{\mu}_1 \in \mathcal{P}$. Then, the following spaces can be chosen iteratively by computing the parameter that maximizes the error bound

$$\boldsymbol{\mu}^k_{\max} := \arg \max_{\boldsymbol{\mu} \in \mathcal{P}_{\text{greedy}}} \Delta^k_{\boldsymbol{u}}(\boldsymbol{\mu}) \tag{85}$$

and adding the corresponding full-order solution to the reduced basis spaces

$$X_{k+1} = X_k \cup \operatorname{span}\left\{ u(\boldsymbol{\mu}^k_{\max}) \right\}, \tag{86}$$

where $\Delta^k_{\boldsymbol{u}}(\boldsymbol{\mu})$ denotes the bound for model-reduction errors $\|u(\boldsymbol{\mu}) - u_{\text{red}}(\boldsymbol{\mu})\|$ computed using reduced-basis space X_k. This algorithm works well in practice and has recently been verified by theoretical convergence results: Refs. [8, 27] proved that the distances of these heuristically constructed reduced spaces converge at a rate similar to that of the theoretical Kolmogorov N-width if this N-width converges algebraically or exponentially fast.

In practice, the performance of this greedy algorithm depends strongly on both the effectivity and computational cost of the error bounds, as less expensive error bounds allow for a larger number of points in $\mathcal{M}_{\text{greedy}}$. Therefore, we expect that employing the ROMES surrogates in lieu of the rigorous bounds may lead to performance gains for the greedy algorithm; this constitutes an interesting area of future investigation.

References

[1] M. Abramowitz and I.A. Stegun, eds., *Handbook of mathematical functions*, vol. 55, National Bureau of Standards Applied Mathatics Series, 1972.

[2] N.M. Alexandrov, R.M. Lewis, C.R. Gumbert, L.L. Green, and P.A. Newman, *Approximation and model management in aerodynamic optimization with variable-fidelity models*, AIAA Journal of Aircraft, 38 (2001), pp. 1093–1101.

[3] P. Astrid, S. Weiland, K. Willcox, and T. Backx, *Missing point estimation in models described by proper orthogonal decomposition*, IEEE Transactions on Automatic Control, 53 (2008), pp. 2237–2251.

[4] I Babuška and A Miller, *The post-processing approach in the finite element method—part 1: Calculation of displacements, stresses and other higher derivatives of the displacements*, International Journal for numerical methods in engineering, 20 (1984), pp. 1085–1109.

[5] M. Barrault, Y. Maday, N.C. Nguyen, and A.T. Patera, *An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations*, C. R. Math. Acad. Sci. Paris Series I, 339 (2004), pp. 667–672.

[6] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera, *An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations*, Comptes Rendus Mathématique Académie des Sciences, 339 (2004), pp. 667–672.

[7] Roland Becker and Rolf Rannacher, *Weighted a posteriori error control in finite element methods*, vol. preprint no. 96-1, Universitat Heidelberg, 1996.

[8] Peter Binev, Albert Cohen, Wolfgang Dahmen, Ronald DeVore, Guergana Petrova, and Przemyslaw Wojtaszczyk, *Convergence rates for greedy algorithms in reduced basis methods*, SIAM J. Math. Anal., 43 (2011), pp. 1457–1472.

[9] Dietrich Braess, *Finite elements: Theory, fast solvers, and applications in solid mechanics*, Cambridge University Press, 2007.

[10] A. Buffa, Y. Maday, Anthony T. Patera, Christophe Prud'homme, and Gabriel Turinici, *A priori convergence of the greedy algorithm for the parametrized reduced basis*, ESAIM-Math. Model. Numer. Anal., 46 (2012), pp. 595–603. Special Issue in honor of David Gottlieb.

[11] C. Canuto, T. Tonn, and K. Urban, *A-posteriori error analysis of the reduced basis method for non-affine parameterized nonlinear pde's*, SIAM J. Numer. Anal, 47 (2009), pp. 2001–2022.

[12] K. Carlberg, *Adaptive h-refinement for reduced-order models*, arXiv preprint 1404.0442, (2014).

[13] K. Carlberg, C. Bou-Mosleh, and C. Farhat, *Efficient non-linear model reduction via a least-squares Petrov–Galerkin projection and compressive tensor approximations*, International Journal for Numerical Methods in Engineering, 86 (2011), pp. 155–181.

[14] K. Carlberg, C. Farhat, J. Cortial, and D. Amsallem, *The GNAT method for nonlinear model reduction: effective implementation and application to computational fluid dynamics and turbulent flows*, Journal of Computational Physics, 242 (2013), pp. 623–647.

[15] S. Chaturantabut and D.C. Sorensen, *Discrete empirical interpolation for nonlinear model reduction*, SIAM J. Sci. Comp., 32 (2010), pp. 2737–2764.

[16] S. Chaturantabut and D. C. Sorensen, *Nonlinear model reduction via discrete empirical interpolation*, SIAM Journal on Scientific Computing, 32 (2010), pp. 2737–2764.

[17] M. Drohmann, B. Haasdonk, and M. Ohlberger, *Reduced basis approximation for nonlinear parametrized evolution equations based on empirical operator interpolation*, SIAM J. Sci Comp, 34 (2012), pp. A937–A969.

[18] ———, *Reduced basis approximation for nonlinear parametrized evolution equations based on empirical operator interpolation*, SIAM Journal on Scientific Computing, 34 (2012), pp. A937–A969.

[19] M. S. Eldred, A. A. Giunta, S. S. Collis, N. A. Alexandrov, and R. M. Lewis, *Second-order corrections for surrogate-based optimization with model hierarchies.*, in In Proceedings of the 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Albany, NY, Aug 2004.

[20] Donald Estep, *A posteriori error bounds and global error control for approximation of ordinary differential equations*, SIAM Journal on Numerical Analysis, 32 (1995), pp. 1–48.

[21] Alexander I. J. Forrester, Andras Sobester, and Andy J. Keane, *Engineering Design via Surrogate Modelling - A Practical Guide.*, Wiley, 2008.

[22] D. Galbally, K. Fidkowski, K. Willcox, and O. Ghattas, *Non-linear model reduction for uncertainty quantification in large-scale inverse problems*, International Journal for Numerical Methods in Engineering, 81 (2009), pp. 1581–1608.

[23] Shawn E Gano, John E Renaud, and Brian Sanders, *Hybrid variable fidelity optimization by using a kriging-based scaling function*, AIAA Journal, 43 (2005), pp. 2422–2433.

[24] M.A. Grepl, Y. Maday, N.C. Nguyen, and A.T. Patera, *Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations*, M2AN, Math. Model. Numer. Anal., 41 (2007), pp. 575–605.

[25] M.A. Grepl and A.T. Patera, *A posteriori error bounds for reduced-basis approximations of parametrized parabolic partial differential equations*, M2AN, Math. Model. Numer. Anal., 39 (2005), pp. 157–181.

[26] M. A. Grepl, Y. Maday, N. C. Nguyen, and A. T. Patera, *Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations*, ESAIM-Mathematical Modelling and Numerical Analysis (M2AN), 41 (2007), pp. 575–605.

[27] B Haasdonk, *Convergence rates of the POD-greedy method*, Tech. Report 23, SimTech Preprint 2011, University of Stuttgart, 2011.

[28] B. Haasdonk, M. Ohlberger, and G. Rozza, *A reduced basis method for evolution schemes with parameter-dependent explicit operators*, Electronic Transactions on Numerical Analysis, 32 (2008), pp. 145–161.

[29] Deng Huang, TT Allen, WI Notz, and RA Miller, *Sequential kriging optimization using multiple-fidelity evaluations*, Structural and Multidisciplinary Optimization, 32 (2006), pp. 369–382.

[30] D.B.P. Huynh, G. Rozza, S. Sen, and A.T. Patera, *A successive constraint linear optimization method for lower bounds of parametric coercivity and inf-sup stability constants*, C. R. Math. Acad. Sci. Paris Series I, 345 (2007), pp. 473–478.

[31] P. A. LeGresley, *Application of Proper Orthogonal Decomposition (POD) to Design Decomposition Methods*, PhD thesis, Stanford University, 2006.

[32] James Ching-Chieh Lu, *An a posteriori error control framework for adaptive precision optimization using discontinuous Galerkin finite element method*, PhD thesis, Massachusetts Institute of Technology, 2005.

[33] Andrew March and Karen Willcox, *Provably convergent multifidelity optimization algorithm not requiring high-fidelity derivatives*, AIAA journal, 50 (2012), pp. 1079–1089.

[34] M. Meyer and H.G. Matthies, *Efficient model reduction in non-linear dynamics using the Karhunen-Loève expansion and dual-weighted-residual methods*, Computational Mechanics, 31 (2003), pp. 179–191.

[35] Leo Wai-Tsun Ng and Michael Eldred, *Multifidelity uncertainty quantification using non-intrusive polynomial chaos and stochastic collocation*, in Structures, Structural Dynamics, and Materials and Co-located Conferences, American Institute of Aeronautics and Astronautics, Apr. 2012, pp. –.

[36] N. C. Nguyen, G. Rozza, and A.T. Patera, *Reduced basis approximation and a posteriori error estimation for the time-dependent viscous burgers' equation*, Calcolo, 46 (2009), pp. 157–185.

[37] A.T. Patera and G. Rozza, *Reduced Basis Approximation and a Posteriori Error Estimation for Parametrized Partial Differential Equations*, MIT, 2007. Version 1.0, Copyright MIT 2006-2007, to appear in (tentative rubric) MIT Pappalardo Graduate Monographs in Mechanical Engineering.

[38] C. Prud'homme, D.V. Rovas, K. Veroy, L. Machiels, Y. Maday, A.T. Patera, and G. Turinici, *Reliable real-time solution of parametrized partial differential equations: Reduced-basis output bound methods*, J. Fluids Engineering, 124 (2002), pp. 70–80.

[39] Dev Rajnarayan, Alex Haas, and Ilan Kroo, *A multifidelity gradient-free optimization method and application to aerodynamic design*, in 12th AIAA/ISSMO multidisciplinary analysis and optimization conference, Victoria, British Columbia, AIAA, vol. 6020, 2008.

[40] Rolf Rannacher, *The dual-weighted-residual method for error control and mesh adaptation in finite element methods*, MAFELEAP, 99 (1999), pp. 97–115.

[41] C.E. Rasmussen and C.K.I. Williams, *Gaussian Processes for Machine Learning*, Adaptative computation and machine learning series, University Press Group Limited, 2006.

[42] G. Rozza, D.B.P. Huynh, and A.T. Patera, *Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations*, Arch. Comput. Meth. Eng., 15 (2007), pp. 229–275.

[43] D. Ryckelynck, *A priori hyperreduction method: an adaptive approach*, Journal of Computational Physics, 202 (2005), pp. 346–366.

[44] Michael E. Tipping, *Sparse bayesian learning and the relevance vector machine*, J. Mach. Learn. Res., 1 (2001), pp. 211–244.

[45] K. Urban and A.T. Patera, *An improved error bound for reduced basis approximation of linear parabolic problems*, Math. Comp, (2013). (in print).

[46] D.A. Venditti and D.L. Darmofal, *Adjoint error estimation and grid adaptation for functional outputs: Application to quasi-one-dimensional flow*, Journal of Computational Physics, 164 (2000), pp. 204–227.

[47] D. Wirtz, D. Sorensen, and B. Haasdonk, *A posteriori error estimation for deim reduced nonlinear dynamical systems*, SIAM Journal on Scientific Computing, 36 (2014), pp. A311–A338.

[48] M. Yano and A.T. Patera, *A spacetime variational approach to hydrodynamic stability theory*, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science, 469 (2013).

(Martin Drohmann) Sandia National Laboratories, 7100 East Ave, Livermore, California 94550
E-mail address, M. Drohmann: `mdrohma@@sandia.gov`

(Kevin Carlberg) Sandia National Laboratories, 7100 East Ave, Livermore, California 94550
E-mail address, K. Carlberg: `ktcarlb@@sandia.gov`