



Available online at [www.sciencedirect.com](http://www.sciencedirect.com)

**ScienceDirect**

Comput. Methods Appl. Mech. Engrg. 370 (2020) 113228

**Computer methods  
in applied  
mechanics and  
engineering**

[www.elsevier.com/locate/cma](http://www.elsevier.com/locate/cma)

# What do we hear from a drum? A data-consistent approach to quantifying irreducible uncertainty on model inputs by extracting information from correlated model output data

Troy Butler<sup>a</sup>, Harri Hakula<sup>b,\*</sup>

<sup>a</sup> University of Colorado Denver, Department of Mathematical and Statistical Sciences, 1201 Larimer St, Denver, CO, 80204, USA

<sup>b</sup> Aalto University, Department of Mathematics and System Analysis, P.O. Box 11100, FI-00076 Aalto, Finland

Received 3 March 2020; accepted 18 June 2020

Available online 24 July 2020

## Abstract

In manufacturing processes, controlling system responses with uncertain system inputs, e.g., due to variations in material parameters of critical system sub-components, is a crucial task for performing reliable quality control and verification & validation (V&V) of system design. As a model for a manufacturing process, we consider the production of drums, that is, thin elastic membranes, whose properties are modeled via Dirichlet Laplacian eigenproblems with uncertain diffusion coefficients. Both a quality control and V&V problem are formulated within a data-consistent framework utilizing push-forward and pullback measures. In both problems, the uncertain diffusion coefficients are parameterized for every instance and the corresponding eigen-information defines correlated data streams. Subsequently, the quantities of interest required in the solution to the data-consistent inverse problems are determined by an *a posteriori* analysis of these data streams using feature extraction techniques. While the methodology proposed here is quite general, the specific efficacy of the proposed methodology is comprehensively explored in the numerical results for both the quality control and V&V problems associated with the manufacturing of drums. © 2020 Elsevier B.V. All rights reserved.

MSC: 28A50; 60-04; 60-08

Keywords: Inverse problem; Data-consistent; Push-forward measure; Pullback measure

## 1. Introduction

In 1966 Mark Kac famously posed the question [1], “Can one hear the shape of a drum?” He was in fact referring to the problem of whether the Laplacian operator with Dirichlet boundary conditions could have identical spectra on two distinct planar regions. It is now known that the answer is negative via counterexamples [2]. The interested reader is referred to a review article by Driscoll [3].

In this work, we consider a variant of this famous question, and ask rather simply and succinctly, “what is it we do hear from a drum?” It turns out the answer is quite simple. We hear information. More precisely, by transforming “notes” into the right types of “chords” (i.e., mixtures of notes), we extract unique quantities of interest (QoI) from the sound produced by striking a drum that are sensitive to different variations in the material properties used to

\* Corresponding author.

E-mail addresses: [troy.butler@ucdenver.edu](mailto:troy.butler@ucdenver.edu) (T. Butler), [Harri.Hakula@aalto.fi](mailto:Harri.Hakula@aalto.fi) (H. Hakula).

construct the drum. These QoI are subsequently useful for formulating and solving a specific type of stochastic inverse problem where the goal is to describe what variations in material properties led to the observed variations in the QoI data.

This type of stochastic inverse problem and the methods we describe for constructing QoI from sets of observed data are actually quite general. Below, we provide some level of abstraction to further motivate the approaches taken in this work.

Consider a manufacturing process that is implemented to produce an “ideal” design of some engineered system. Inputs to this process may include either raw or refined materials used to construct system components and potentially many sub-systems and sub-components that may be sourced from other manufacturers (which in turn often rely upon other raw or refined materials). Impurities in raw materials can vary greatly in both composition and concentration and can rarely, if ever, be completely removed by refinement, e.g., see [4,5] where controls of impurities in refined materials are discussed and [6] where impurities are deliberately introduced at certain tolerances to achieve desirable properties. Moreover, sub-systems and sub-components often include machined parts that vary within some “acceptable” range “most of the time” depending on the technology and quality of machinery utilized in their construction, e.g., see [7–9] for analyses on general impacts of machining errors. Consequently, variability in the system outputs related to performance is inevitable. Failure to quantify uncertainties related to these variabilities can subsequently lead to unreliable or unpredictable system performance, e.g., see [10] for a comprehensive framework analyzing the impacts of aleatoric and epistemic sources of uncertainty on model predictions.

Viewing the underlying uncertainty as a random process, we view the parameters (i.e., system inputs) of the engineered system as well as system outputs as random variables/fields with some “exact” or “true” distributions. In other words, variability in system outputs is expected due to variability inherent to the system itself and not necessarily due to the imprecision of measurement equipment used to observe system outputs (i.e., measurement errors). In the language of uncertainty quantification (UQ), we are concerned with uncertainties that are aleatoric (i.e., irreducible) in nature not epistemic (i.e., theoretically reducible). For both quality control purposes and to ensure that the system performs within specifications under a variety of scenarios, it is imperative that these distributions be determined. A first step in solving quality control problems often involves subjecting a sampled set of manufactured systems to a finite number of controlled experiments (e.g., in a laboratory setting) to generate a finite amount of observable data. Then, this data may be used to infer properties related to parameter variability.

Such problems also arise naturally when the simulation model of a manufactured system requires verification and validation (V&V) [11,12]. The V&V process may involve determining the probability of compliance of the manufactured system response within some pre-determined “industry” standards. Such standards may be defined, for instance, using the responses of previously validated simulation models subjected to various input conditions. However, if manufacturing processes or materials are proprietary or involve new technologies and methods, initial descriptions of uncertainty in parameters may be limited to specifications of sets of physically plausible values. Subsequently, a useful first step in V&V is to separate sets of probable parameters from those that are merely physically plausible.

In this work, we make these problems less abstract by formulating them in the context of manufacturing thin elastic membranes that are cut and stretched over a rigid frame. That is, we consider the manufacturing of “drums” where the “tension” of the drums varies spatially and no two drums are exactly alike. Both a quality control and V&V problem are formulated as stochastic inverse problems involving push-forward and pullback measures. The solutions to these stochastic inverse problems are obtained using a recently developed uncertainty quantification (UQ) method based on the disintegration of measures [13]. This method re-weights an initial probability density on parameters in directions informed by the observations. This updated density on parameters is then guaranteed to reconstruct the observed density. We therefore refer to this updated density as a *data-consistent* solution. This is, in a certain sense, a type of probabilistic calibration. However, this approach completely avoids the perturbation analysis and weak nonlinearity assumptions in the QoI map often required by other UQ calibration methods, e.g., see [14]. While a comparison of this type of inverse problem and its solution to other formulations and solutions of inverse problems is provided in [13] using a computationally simple algebraic QoI map, we summarize some key differences with respect to Bayesian-inspired methods [15–17] here since they are some of the most popular means of inferring probabilistic descriptions of model parameters from QoI data.

In typical Bayesian frameworks, a common assumption is that uncertainty in the QoI data is described by an additive noise model. Conceptually, variations in QoI data are then viewed as the result of epistemic rather than

aleatoric uncertainty so that it is at least theoretically possible to reduce uncertainty using improved instrumentation to collect more precise data. The Bayesian posterior is given by a conditional density proportional to the product of a data-likelihood function and prior distribution prescribed on the parameters. The posterior can be interrogated to assess the relative likelihoods that different *fixed estimates* of the parameters could have produced all of the QoI data. Subsequently, a common use of the posterior is to generate point-estimators such as the maximum a posteriori (MAP) or the conditional mean (CM) estimate to describe the Bayes estimator (i.e., the solution) of the inverse problem as a single fixed estimate of the parameters [18]. Given the assumptions of epistemic uncertainty in typical Bayesian frameworks, viewing a single point estimate of the parameters as the solution is actually quite reasonable. In fact, as more data are collected, the formulation of the data-likelihood function used to create the posterior will often cause the posterior to become “spiked” around a single parameter. This phenomenon is explained by the Bernstein–von Mises theorem [19].

Alternative Bayesian formulations exist to address problems where uncertainty in the parameters is viewed as aleatoric instead of epistemic. Perhaps the most popular are hierarchical Bayesian methods [20] that begin with specifying prior distributions from a parametric family (e.g., normal distributions). Hyper-parameters refer to the scalar quantities parameterizing the prior distribution (e.g., means and variances). After assuming additional prior distributions on the hyper-parameters, the Bayesian paradigm is used to generate estimates of the hyper-parameters. Some non-parametric, hierarchical Bayesian methods for modeling aleatoric uncertainty also exist such as those deploying Dirichlet Processes [21,22]. While these are an active research area within the machine/statistical learning community, they are currently more expensive and complex to implement.

In [13], the data-consistent UQ method utilized in this work implicitly assumes that the QoI, defined as the observable system outputs for which data are collected, are specified a priori. A contribution of the current work is the utilization of feature extraction techniques based on principal component analysis to multivariate data that allows us to define an “optimal” QoI map relating system parameters to system outputs. In other words, we *learn*, a posteriori, what QoI are present within the predicted data. Moreover, these QoI possess more optimal geometric properties (compared to the raw observed data) in terms of the both the relationships between QoI components *and the subsequent directions in parameter space informed by each component of a QoI map*. In the recent work of [23], similar feature extraction techniques are used to streamline model validation. However, in that work, the focus is entirely on validation of the predictive capabilities of a model so that only a metric between predicted and observed QoI distribution functions is considered. This is in contrast to this current work that utilizes discrepancies between predicted and observed QoI distributions to aid in the determination of physically probable sets of parameters in the V&V example. Moreover, we utilize a diagnostic from the UQ method to analyze the feasibility of various descriptions of initial parameter uncertainties. This subsequently allows for the specification of the “best” initial design in terms of the capacity of the design to reproduce observable variations in data.

The rest of the paper is organized as follows. Section 2 summarizes the theoretical foundation of the PDE eigenproblem used to model the drums. The formulation and solution of stochastic inverse problems within a data-consistent are then discussed in Section 3. In the interest of scientific reproducibility, the computational framework used to generate all of the subsequent numerical results is briefly summarized in Section 4. The numerical results for both problems are then summarized in Section 5. Concluding remarks follow in Section 6. Finally, for the interested reader, Appendix A discusses system requirements for executing the supplemental documents containing all data sets and scripts that reproduce the relevant figures and table data presented in Section 5 along with additional numerical results omitted in the interest of space.

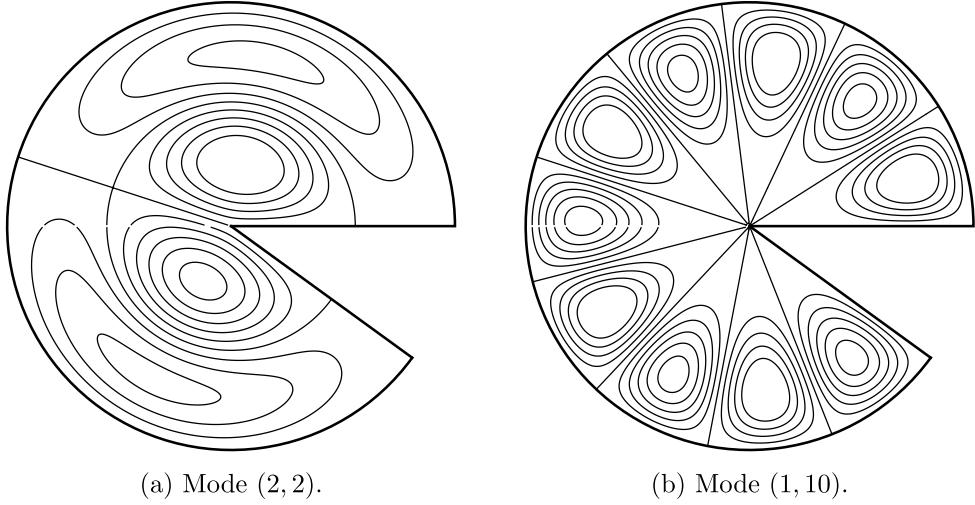
## 2. Preliminaries

In this section, we briefly outline the theoretical foundation of the PDE eigenproblem used to model the drums. For a through overview of the subject we refer the read to Boffi [24]. The actual eigensolver and details of the *hp*-discretization are documented in [25].

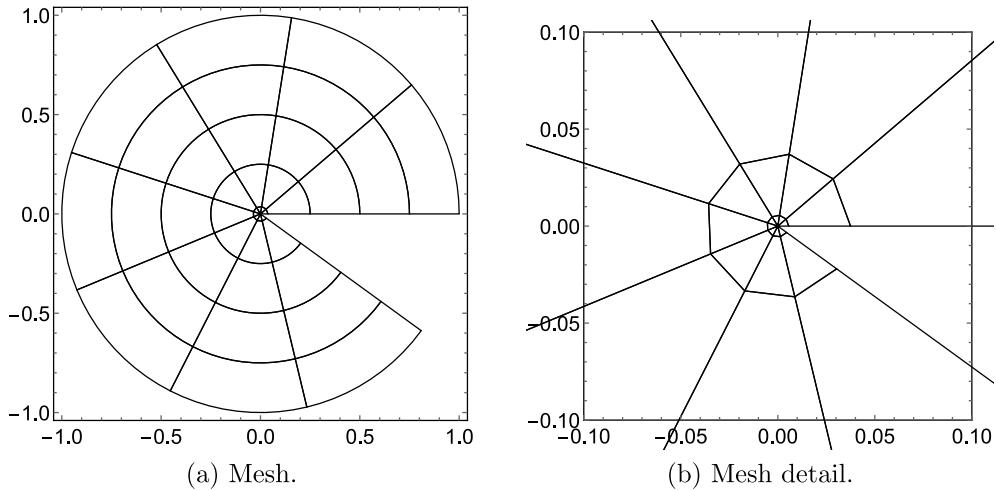
### 2.1. Laplace eigenvalue problem

Let  $\alpha = 5/9$  be a *cut parameter* used to parametrize the computational domain, a cut disk  $\Omega = \mathbb{D}(\mathbf{x}, r \in [0, 1], \theta \in [0, \pi/\alpha])$  centered at the origin  $\mathbf{x} = (0, 0)$ , and consider the Laplace eigenvalue problem:

$$-\Delta\psi = \lambda\psi \text{ in } \Omega, \quad \psi = 0 \text{ on } \partial\Omega. \quad (2.1)$$



**Fig. 1.** Two eigenmodes.



**Fig. 2.** Drum design. Strongly graded  $hp$ -mesh with detail close to the origin.

The eigenvalues and vectors are known explicitly, and are doubly-indexed for  $m, n \in \mathbb{N}$  by

$$\psi_{m,n} = J_{\sigma_n}(j_{m,n}r) \sin(\sigma_n \theta), \quad \lambda_{m,n} = j_{m,n}^2, \quad (2.2)$$

where  $\sigma_n = \alpha n$ ,  $J_{\sigma_n}$  is the first-kind Bessel function of order  $\sigma_n$ , and  $j_{m,n}$  is the  $m$ th positive root of  $J_{\sigma_n}$ . In Fig. 1, contour plots of two different modes are illustrated.

In the simulations, we consider the generalized problem

$$-\nabla \cdot (a(r)\nabla \psi) = \lambda \psi \text{ in } \Omega, \quad \psi = 0 \text{ on } \partial\Omega, \quad (2.3)$$

where we let the diffusion parameter  $a(r)$  vary on the radial direction only. The problem is solved using  $hp$ -FEM on a strongly graded mesh (see Fig. 2) with high polynomial degree  $p = 8$ . With this setup, it is guaranteed that the assumptions of the Babuska–Osborne theorem are satisfied and no mixing of modes occurs when  $a(r) = 1$ , that is, in the reference problem (2.1).

**Table 1**

Mode signatures: Effect of radial diffusion coefficient. Enumerated eigenmodes with signatures and eigenvalues. The eigenvalues are roughly twice as large in the perturbed case.

(a) Reference case: $a(r) = 1$ .								(b) One sample: $a(r) = 1 + r + r^2$ .							
$i$	$m_i$	$n_i$	$\lambda_i$	$i$	$m_i$	$n_i$	$\lambda_i$	$i$	$m_i$	$n_i$	$\lambda_i$	$i$	$m_i$	$n_i$	$\lambda_i$
1	1	1	10.4	11	1	8	65.9	1	1	1	20.5	11	2	4	141
2	1	2	15.8	12	2	4	76.0	2	1	2	31.3	12	1	8	145
3	1	3	22.2	13	1	9	76.9	3	1	3	44.6	13	3	1	157
4	1	4	29.3	14	1	10	88.7	4	1	4	60.1	14	2	5	169
5	1	5	37.3	15	2	5	89.6	5	2	1	72.1	15	1	9	171
6	2	1	40.5	16	3	1	90.4	6	1	5	78.0	16	3	2	187
7	1	6	46.1	17	1	11	101	7	2	2	92.4	17	2	6	199
8	2	2	51.5	18	2	6	104	8	1	6	98.0	18	1	10	200
9	1	7	55.6	19	3	2	107	9	2	3	115	19	3	3	220
10	2	3	63.3	20	1	12	115	10	1	7	120	20	1	11	230

## 2.2. Mode detection

As we let the diffusion parameter  $a(r)$  vary, it is clear that the eigenmodes are not exactly those given by (2.2). However, if the perturbations of  $a(r)$  are sufficiently small, it is reasonable to consider what is the closest mode in the original set to the observed one. The exact eigenmodes have a tensor structure which simplifies the mode detection task considerably. Let  $\hat{\psi}$  be an approximate eigenmode and thus a solution of (2.3). The two indices  $(m, n)$  represent oscillations in the radial and angular directions, respectively. Let us discuss the angular index  $n$  first. We evaluate the mode  $\hat{\psi}$  at two random radii  $r_1$  and  $r_2$  over the whole interval of  $\theta$  and compute the wave number  $n$  using Fourier transformation. If the wave numbers agree, we stop, otherwise a third radius  $r_3$  is selected for tie-breaking. (This happens very rarely.) In the radial direction,  $m$  is detected similarly. Since there is no direct substitute for the Fourier transform, we project the mode to a set of admissible Bessel functions and choose the closest one in the  $L^2$ -sense.

## 2.3. Parameter-dependent mixing of modes

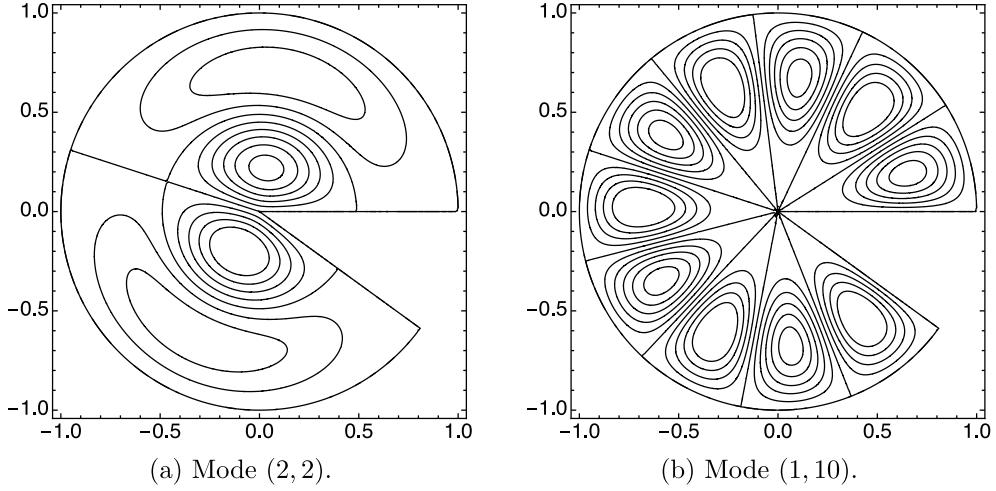
As outlined above, the reference eigenmodes (2.2) have a tensor structure that can be identified with high reliability. In the perturbed case, the eigenvalues naturally change, but also the order of the modes changes. This is referred to as *mixing of modes*. This effect is demonstrated in Table 1, where the first twenty modes are tabulated for the reference case with  $a(r) = 1$  and another perturbed one with  $a(r) = 1 + r + r^2$ . It is clear that not only are the eigenvalues roughly twice as large in the perturbed case, but also the modes with stronger radial oscillations are brought forward in the spectrum. In fact the factor of two is to be expected since in the perturbed case  $a(r) = 1 + r + r^2$  and thus  $\int_{[0,1]} a(r) dr = 11/6$ . In Section 4, we discuss how this mixing of modes is addressed in the analysis of the data (see Fig. 3).

## 3. The stochastic inverse problem

We summarize the terminology, notation, and (both the explicit and implicit) assumptions for solving the stochastic inverse problem considered in this work that is originally explored in [13]. For consistency, we use the notation from the drum problem above although it is relatively straightforward to map this to the notation used in [13]. Then, the stochastic inverse problem and its solution are defined. Finally, we briefly summarize the theory of existence, uniqueness, and stability, as well as some computational approaches for interrogating the solution.

### 3.1. Terminology and notation

Denote by  $\mathbb{A}$  the space of all *physically possible* inputs to the model, which we refer to as parameters. A quantity (or quantities) of interest (QoI) refer to the functional(s) applied to the solution space of the model corresponding to *physically meaningful* scalar (or vector-valued) outputs that are sensitive to perturbations in the parameters.



**Fig. 3.** Two approximated eigenmodes.

Depending on the problem, a QoI map may correspond to either output values we wish to predict or values computed from observable (or simulated) data, e.g., for the purpose of model validation, calibration, or design. The parameter-to-QoI map (often referred to simply as the QoI map) is denoted by  $Q(a(r)) : \mathbb{A} \rightarrow \mathcal{Q}$  to make explicit the dependence on the diffusion parameter  $a(r)$ . Here,  $\mathcal{Q}$  denotes space of all *physically possible* output data that the QoI map can predict.

Since we are interested in solutions to problems defined as densities, assume that  $(\mathbb{A}, \mathcal{B}_{\mathbb{A}}, \mu_{\mathbb{A}})$  and  $(\mathcal{Q}, \mathcal{B}_{\mathcal{Q}}, \mu_{\mathcal{Q}})$  are both measure spaces. Here,  $\mathcal{B}_{\mathbb{A}}$  and  $\mathcal{B}_{\mathcal{Q}}$  denote the Borel  $\sigma$ -algebras inherited from the metric topologies on their respective spaces, and  $\mu_{\mathbb{A}}$  and  $\mu_{\mathcal{Q}}$  denote the dominating measures for which probability densities (i.e., Radon–Nikodym derivatives of probability measures) are defined on each space. It is implicitly assumed that the QoI map is a measurable map between these measure spaces. A practical assumption implicitly made by computations involving finite sampling of the QoI map and standard density approximation techniques is that the QoI map is also piecewise smooth.

### 3.2. The forward and inverse problems

The stochastic inverse problem we consider involves the construction of pullback measures. In other words, this is the direct inverse of a stochastic forward problem involving push-forward measures. At a high-level, the forward problem involves quantifying uncertainties in predictions of QoI data using some initial information to describe uncertainties on parameters. Below, we formally define the forward problem and the subsequent inverse problem that is the focus of this work.

**Definition 3.1 (Forward Problem and Push-forward Measure).** Given an initial (i.e., initially specified) probability measure  $P_{\mathbb{A}}^{\text{init}}$  on  $(\mathbb{A}, \mathcal{B}_{\mathbb{A}})$  the forward problem is the determination of the push-forward probability measure

$$P_{\mathcal{Q}}^{Q(\text{init})}(E) := P_{\mathbb{A}}^{\text{init}}(Q^{-1}(E)), \quad \forall E \in \mathcal{B}_{\mathcal{Q}}$$

on  $(\mathcal{Q}, \mathcal{B}_{\mathcal{Q}})$ .

We often refer to the push-forward of the initial measure as the *predicted* measure.

**Definition 3.2 (Inverse Problem and Consistent Measure).** Given both an observed probability measure  $P_{\mathcal{Q}}^{\text{obs}}$  on  $(\mathcal{Q}, \mathcal{B}_{\mathcal{Q}})$  and an initial probability measure  $P_{\mathbb{A}}^{\text{init}}$  on  $(\mathbb{A}, \mathcal{B}_{\mathbb{A}})$ , the inverse problem is to determine an update to the initial probability measure, denoted by  $P_{\mathbb{A}}^{\text{update}}$ , on  $(\mathbb{A}, \mathcal{B}_{\mathbb{A}})$  that is consistent with the observed data in the sense that

$$P_{\mathbb{A}}^{\text{update}}(Q^{-1}(E)) = P_{\mathcal{Q}}^{\text{obs}}(E), \quad \forall E \in \mathcal{B}_{\mathcal{Q}}. \quad (3.1)$$

From the above two definitions, it is clear that the consistency criterion of (3.1) implies that the push-forward of the updated measure, denoted by  $P_{\mathcal{Q}}^{\text{Q(update)}}$ , matches the observed measure on  $(\mathcal{Q}, \mathcal{B}_{\mathcal{Q}})$ , i.e.,  $P_{\mathbb{A}}^{\text{update}}$  is a pullback of the observed measure. In other words, if we solve a forward problem using  $P_{\mathbb{A}}^{\text{update}}$  and the same QoI map, then we will “predict” what was in fact observed. In this sense, we may interpret the solution to the inverse problem as solving a type of probabilistic calibration problem. Any subsequent forward problem that utilizes this updated measure with a different QoI map is then interpreted as forming a *data-consistent* prediction. This describes a common workflow in engineering design where validation of design parameters using data from a limited set of laboratory experiments is used to predict system performance under different conditions than can be replicated in the laboratory.

### 3.3. Inverse solutions: Theory and numerical interrogation

We assume that the probability measures are absolutely continuous with respect to the dominating measures and admit densities (i.e., Radon–Nikodym derivatives). Let  $\pi_{\mathbb{A}}^{\text{init}}$ ,  $\pi_{\mathcal{Q}}^{\text{Q(init)}}$ ,  $\pi_{\mathcal{Q}}^{\text{obs}}$ , and  $\pi_{\mathbb{A}}^{\text{update}}$  denote the densities corresponding to  $P_{\mathbb{A}}^{\text{init}}$ ,  $P_{\mathcal{Q}}^{\text{Q(init)}}$ ,  $P_{\mathcal{Q}}^{\text{obs}}$ , and  $P_{\mathbb{A}}^{\text{update}}$ , respectfully. With this notation, we make the following *predictability assumption* that ensures the solvability of the inverse problem.

**Assumption 1.** There exists  $C > 0$  such that for a.e.  $q \in \mathcal{Q}$ ,  $\pi_{\mathcal{Q}}^{\text{obs}}(q) \leq C\pi_{\mathcal{Q}}^{\text{Q(init)}}(q)$ .

At a conceptual level, we refer to this as a predictability assumption because it ensures that QoI data we observe to be likely are also predicted to be likely by the push-forward of the initial density. It is important to note that this is an assumption on *both* the initial density and the QoI map. In other words, violation of this assumption may be due to either an initial density that is too restrictive in its support (i.e., physically possible parameters are not considered initially likely) or because of model inadequacy (i.e., missing physics or computational errors make it impossible to predict certain ranges of observable data).

In [13], the predictability assumption, a Disintegration Theorem [26], and Bayes’ rule on generalized contours are used to derive the following closed form expression for the updated density (which we write using the notation of this work):

$$\pi_{\mathbb{A}}^{\text{update}}(a(r)) = \pi_{\mathbb{A}}^{\text{init}}(a(r)) \frac{\pi_{\mathcal{Q}}^{\text{obs}}(Q(a(r)))}{\pi_{\mathcal{Q}}^{\text{Q(init)}}(Q(a(r)))}, \quad a(r) \in \mathbb{A}. \quad (3.2)$$

This updated density solves the inverse problem (i.e., it defines a consistent measure) of **Definition 3.2**. Moreover, up to the choice of an initial density, by the Disintegration theorem it is also unique. Stability (in the total variation metric [27]) with respect to perturbations in either the initial or observed densities also follows. An important takeaway is that the inverse solution is formally determined once the predicted density,  $\pi_{\mathcal{Q}}^{\text{Q(init)}}$ , is constructed. In other words, constructing the density associated with the push-forward measure solving the forward problem in **Definition 3.1** provides *sufficient information* for solving the inverse problem. Moreover, since the predicted density does not depend on the observed density, it is trivial to construct solutions to multiple inverse problems defined by considering various observed densities on the QoI (under the constraint that any such additional observed densities satisfy the predictability assumption). The interested reader is directed to [13] for a thorough discussion.

To make the algorithmic approaches self-evident, we rewrite the updated density as

$$\pi_{\mathbb{A}}^{\text{update}}(a(r)) = \pi_{\mathbb{A}}^{\text{init}}(a(r)) \mathcal{R}(a(r)), \quad \text{with } \mathcal{R}(a(r)) := \frac{\pi_{\mathcal{Q}}^{\text{obs}}(Q(a))}{\pi_{\mathcal{Q}}^{\text{Q(init)}}(Q(a))}. \quad (3.3)$$

The ratio denoted by  $\mathcal{R}(a(r))$  has several practical uses that we outline below.

First, we observe that

$$1 = \int_{\mathbb{A}} \pi_{\mathbb{A}}^{\text{update}} = \int_{\mathbb{A}} \pi_{\mathbb{A}}^{\text{init}}(a(r)) \mathcal{R}(a(r)) = \mathbb{E}_{\text{init}}(\mathcal{R}(a(r))). \quad (3.4)$$

Note that the expectation operator is taken with respect to the initial probability measure. Thus, the expected value of the random variable given by  $\mathcal{R}(a(r))$  defined on the initial probability space of parameters *should be* 1 if the predictability assumption holds. This is turned into a numerical diagnostic as follows. First, replace  $\pi_{\mathcal{Q}}^{\text{Q(init)}}$  with a numerical approximation constructed using a finite number of QoI samples associated with a finite set of parameter

samples drawn from the initial density. Such an approximation may be obtained, e.g., by standard Gaussian kernel density estimation (GKDE) that have well-established rates of convergence [28,29]. Compute and compare to 1 the sample average of  $\mathcal{R}(a(r))$  using this approximation to  $\pi_Q^{Q(\text{init})}$ . As shown in [13] as well as in the numerical examples of Section 5, this is a useful numerical diagnostic for verifying that both the predictability assumption holds and that the approximation of  $\pi_Q^{Q(\text{init})}$  is sufficiently accurate.

Second, the computation of  $\mathcal{R}(a(r))$  on QoI samples provides several algorithmic approaches for interrogating  $\pi_A^{\text{update}}$ . One approach is to use this to re-weight parameter samples coming from the initial density and compute weighted density estimates on the parameter space directly (this is pursued in the first numerical example of Section 5). Alternatively, standard rejection sampling can be performed to sub-select initial parameter samples into a set of independent identically distributed (i.i.d.) samples from the updated density as described in [13]. Specifically, we can generate a set of i.i.d. samples from the initial density, propagate these samples to the output space using any fixed QoI map to define a set of *proposal samples* in the QoI space, and perform rejection sampling using the *target* observed density. The accepted samples, viewed in the QoI space, are i.i.d. samples from the observed density. Subsequently, the accepted samples, viewed in the parameter space, are i.i.d. samples from a consistent solution. Since these are clearly an *update* to the initial sample set, the motivation for the terminology is now self-evident. Finally, we can simply sort the scalar array of  $\mathcal{R}(a(r))$  values and visualize sets of parameters that meet various threshold conditions involving their relative updated likelihoods (this is pursued in the second numerical example of Section 5).

#### 4. A computational framework

In the interest of scientific reproducibility, we provide a high-level overview of the computational framework used to generate and analyze the data for the numerical results presented in the following section. The numerical experiments for generating data related to the drum problem described in Section 2 were run on the Aalto University Triton cluster with the *hp*-FEM eigensolver implemented in Mathematica 12 [25,30]. The analysis of the data within the data-consistent framework described in Section 3 was carried out using Python scripts that utilize the module, *DrumAnalysis*, developed specifically for managing the format of the data sets as well as consolidating the functionality necessary for performing data-consistent inversion. These data sets, Python scripts, and the Python module are provided as supplementary material to this document. In Appendix A, we describe the minimal necessary requirements for running the scripts to re-create the relevant figures and tables presented in Section 5 along with some additional numerical results that are omitted in the interest of space.

##### 4.1. Model computations: Implementation and fingerprinting

In the experiments, observed data are generated using two different models for the diffusion coefficient: (a) linear of the form  $a(r) = a_0 + a_1 r$ , and (b) quadratic  $a(r) = a_0 + a_1 r + a_2 r^2$ . In the linear case, the parameters  $a_i$  were sampled from the prescribed distributions using the routines provided in Mathematica. In the quadratic case, the admissible coefficients were sampled by first interpolating sampled values at three distinct points in the interval  $[0, 1]$  and then verifying that they satisfied the given geometric constraints, for instance, that the graphs of the coefficients stayed within given bounds.

The predicted data associated with the linear case are generated from the same linear model but using different prescribed distributions as discussed in the quality control example of Section 5. The predicted data associated with the quadratic case are generated using several models defined by continuous piecewise-linear splines for  $a(r)$  as discussed in the V&V example of Section 5.

Once the eigenmodes had been computed, each mode was analyzed using techniques described in Section 2.2. Every mode was identified and stored in the data sets using a triple  $(\lambda, m, n)$ , where  $\lambda \in \mathbb{R}$  is the eigenvalue, and  $m \in \mathbb{N}$  and  $n \in \mathbb{N}$  are the radial and angular wavenumbers, respectively.

##### 4.2. Uncertainty quantification: Implementation and QoI construction

The Python module, *DrumAnalysis*, has functions for managing the eigenmode data sets, constructing QoI maps from this data, performing density estimation, and constructing/analyzing  $\mathcal{R}(a(r))$ . Here, we summarize some of the key features of this module. The function `un_mix` utilizes the fingerprinting to ensure that the indices for arrays

of eigenvalues,  $\lambda$ , are aligned with the same indices of wavenumbers,  $(m, n)$ , appearing across all sampled diffusion coefficients. The function `QoI_xform` transforms the un-mixed eigenvalue data into QoI data (this is discussed in the context of the numerical results) for which density estimation can be applied. The function `compute_r` then uses the observed and predicted density estimates to evaluate  $\mathcal{R}(a(r))$  on the sampled QoI data. The module is well-documented for the interested reader who wishes to explore other built-in functionality not discussed here or utilized by the provided scripts.

## 5. Numerical results

We formulate and solve the inverse problem for two examples motivated by distinct engineering problems involving the manufacturing of drums of the shape shown in Fig. 1. Both problems fundamentally involve quantifying uncertainties in the diffusion parameter  $a(r)$  related to the tension of the drumheads (i.e., the elastic membrane that is stretched and tightened along the top of the drum frame). Such uncertainties may be attributed, for instance, to variations in the elastic material properties of individual membranes or to the amount of force used in stretching/securing the membrane to the drum. For simplicity, we ignore any uncertainties due to potential deviations from the drum shape. This is reasonable, for instance, if we assume that either a highly machined and rigid cast is used to construct the drum frames or different membranes are tested upon the same drum.

In general, solutions to either type of problem are theoretically complicated by the fact that  $\mathbb{A}$  is nominally infinite dimensional since it describes a function space for the potential diffusion parameters  $a(r)$ . Moreover, while it may be possible to accurately simulate or experimentally observe many potential eigenmodes, it is not clear which ones form the most useful QoI maps due to high correlation in both their values and sensitivities to perturbations in the parameters. In the context of each of the examples below, we discuss practical constraints placed on  $\mathbb{A}$  that make the space finite-dimensional and amenable to computational approximations.

### 5.1. Example 1: Quality control

Here, we assume that the manufacturing process is setup to stretch elastic membranes over the drum shape in such a way that the resulting diffusion parameter  $a(r)$  should linearly increase from  $r = 0$  to  $r = 1$ . In other words, the manufactured drums should have diffusion parameters of the form  $a(r) = a_0 + a_1 r$ . For simplicity in discussing the numerical results, we refer to the pair  $(a_0, a_1)$  as the actual parameters in the model and seek a properly defined  $\mathbb{A} \subset \mathbb{R}^2$  for which uncertainties are quantified.

#### 5.1.1. Defining the parameter space

Suppose that the ideal drum is given when  $a_0 = 1$  and  $a_1 = 0.15$ , but the tolerances in the manufacturing process can only be controlled to guarantee that  $a_0 \in [0.80, 1.20]$  and  $a_1 \in [0.10, 0.20]$ . Using only this information, we choose  $\pi_{\mathbb{A}}^{\text{init}}(a_0, a_1)$  associated with a uniform probability measure on  $\mathbb{A}$ , i.e.

$$(a_0, a_1) \in \mathbb{A} := [0.80, 1.20] \times [0.10, 0.20], \text{ and } \pi_{\mathbb{A}}^{\text{init}}(a_0, a_1) \equiv 4.$$

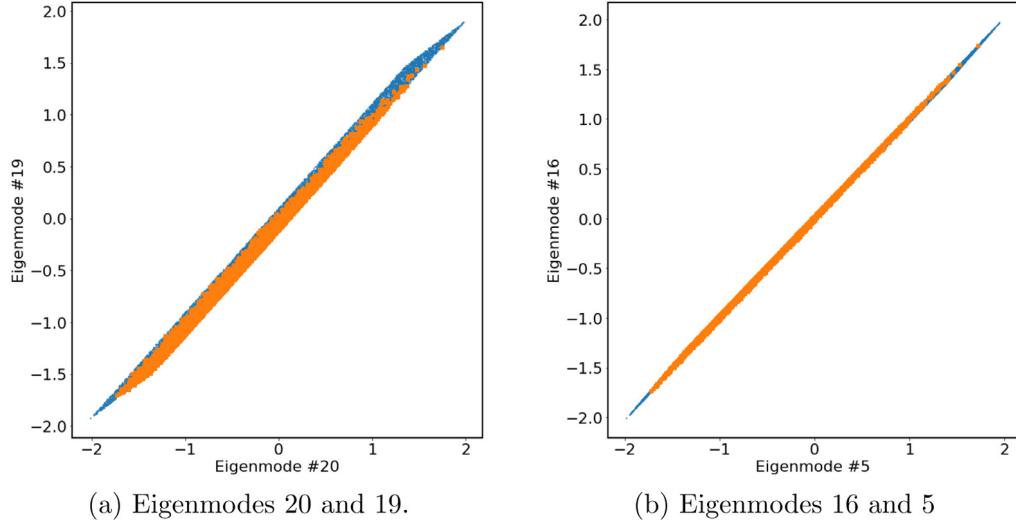
Suppose that the true variability in the parameters is described by the tensor product of densities associated with

$$a_0 \sim \text{Beta}(2, 4) \text{ and } a_1 \sim \text{Beta}(3, 1),$$

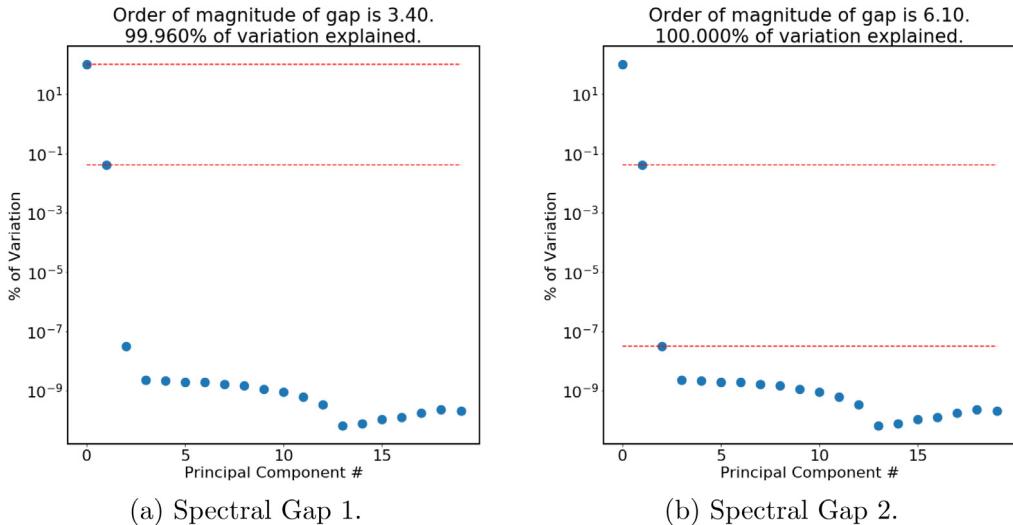
i.e., the true variability of  $(a_0, a_1)$  is described by independent random variables with different Beta distributions. These particular Beta distributions are chosen to exhibit significantly different qualitative behaviors both compared to each other and compared to the initial uniform distribution.

#### 5.1.2. Learning the QoI

Suppose that we can observe up to the first 20 eigenmodes associated with this type of drum, which we simulate using numerical solutions to (2.3). In this problem, we begin with two samples of observable eigenmodes. First, we generate  $1E4$  i.i.d. samples from  $\pi_{\mathbb{A}}^{\text{init}}$  and compute the associated  $1E4$  samples of 20-dimensional vectors defined by these eigenmodes. Second, we generate  $5E3$  i.i.d. samples from each of the Beta distributions described above and compute the associated  $5E3$  samples of the eigenmodes. However, only the eigenmode data associated with these Beta distributed parameter samples are ever used to solve the inverse problem. In other words, we have access



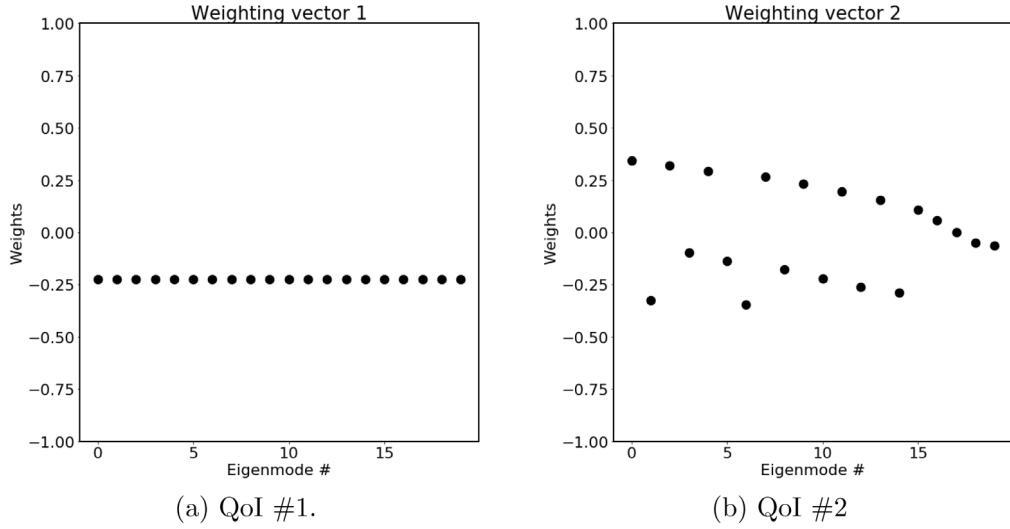
**Fig. 4.** Representative results of high correlation present between eigenmodes (after standard scaling) as parameters are varied. Blue dots are predicted values and orange squares are observed values. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



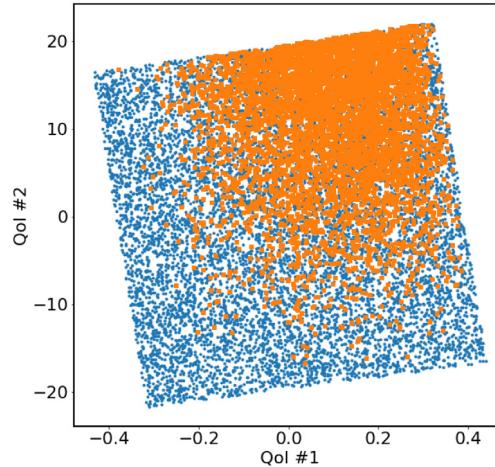
**Fig. 5.** The eigenmode data is effectively summarized by two principal components as illustrated by the combination of spectral gap magnitude and percent of variation (rounded to three places) explained by these principal components.

to three sample sets: (1) the uniform samples from  $\pi_{\mathbb{A}}^{\text{init}}$ , (2) the predicted eigenmodes associated with the uniform samples, and (3) the observed eigenmodes associated with the true distribution of parameters.

In Fig. 4, we provide two representative scatter plots that illustrate the significant positive correlation that exists between the eigenmodes. Given the strong linear relationship between eigenmodes, we choose to apply a principal component analysis (PCA) to this eigenmode data following a standard scaling of the data using the function `QoI_xform` in the `DrumAnalysis` module. Fig. 5 illustrates that the first principal component explains approximately 99.96% of the variation present in the eigenmode data while the addition of the second principal component effectively accounts for the remaining variation in this data. In fact, there is approximately a 6.10 order of magnitude drop between the percent of variation explained by the second and third principal components, which serves as further evidence that, geometrically speaking, the 20-dimensional eigenmode data is accurately described



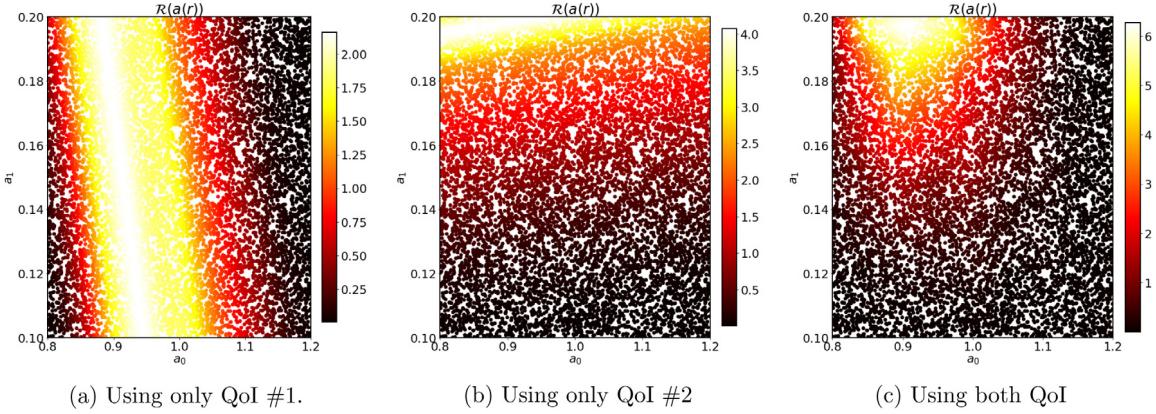
**Fig. 6.** (a) QoI #1 is defined (approximately) by a simple scaled average of all observable eigenmodes. Since this is the first principal component, it illustrates the significant correlation between all eigenmodes. (b) QoI #2 shows the complex linear combination of eigenmodes necessary to describe the orthogonal direction to the first QoI and explain the remaining small amount of variation (see Fig. 5) in the eigenmode data not explained by the first QoI.



**Fig. 7.** The reduction in correlation between the QoI is significant compared to individual eigenmodes (compare to Fig. 4). Blue dots are predicted values and orange squares are observed values. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

as existing inside of a 2-dimensional manifold embedded in  $\mathbb{R}^{20}$ . It should be noted that we *expect* this to happen since the response of each of the eigenmodes is determined by exactly two parameters and the model is deterministic. We take these two principal components as the QoI and label them QoI #1 and QoI #2, respectively.

The weighting vectors defining these QoI are illustrated in Fig. 6. Here, we see that QoI #1 is defined by almost an equal weighting of all the eigenmodes, which is perhaps obvious given the high correlation between the eigenmodes as discussed above. However, QoI #2 is defined by a more complicated linear combination of eigenmodes that is not immediately obvious from any visual inspection of the data. In Fig. 7, we illustrate the predicted and observed samples in the data space defined by these two QoI. Comparing this to the plots in Fig. 4, we observe significantly less correlation between these QoI than between any pair of eigenmodes.



**Fig. 8.** The  $\mathcal{R}(a(r))$  we obtain on parameters  $a_0$  (horizontal axis) and  $a_1$  (vertical axis) if only using (a) QoI #1, (b) QoI #2, or (c) both QoI. The diagnostics  $\mathbb{E}_{\text{init}}(\mathcal{R}(a(r)))$  are approximately 1.00, 1.00, and 1.03, respectively. Here, the QoI numbering is consistent with the labeling in Fig. 6.

### 5.1.3. Solving the inverse problem: Qualitative and quantitative results

We formulate and solve three inverse problems using QoI maps defined by QoI #1, QoI #2, and by using both QoI #1 and QoI #2 to illustrate the geometric relationship between the parameters and each QoI. Using both standard and weighted Gaussian kernel density estimators (GKDE) provided in the `stats` subpackage of `scipy` (with a simple wrapper function provided in the provided module `DrumAnalysis` for convenience), we construct the predicted and observed densities for each QoI map, and compute the corresponding  $\mathcal{R}(a(r))$ . The plots of the updated weighting of the initial samples are in Fig. 8 and of the approximate marginals of the associated updated densities are in Figs. 9–11. Note that the diagnostic,  $\mathbb{E}_{\text{init}}(\mathcal{R}(a(r)))$  (refer to the discussion around (3.4)), is close to 1.00 for each inverse problem solution (see the caption in Fig. 8), which verifies both the predictability assumption for each map as well as the quality of the approximation of the associated predicted densities.

Below, we provide a mixture of qualitative and quantitative analyses of the updated solutions in terms of the marginal PDFs obtained for each parameter. The qualitative analysis is based upon plots of the various marginal PDFs. The quantitative analysis is based upon the distance between various marginal PDFs using the total variation (TV) metric [27], denoted by  $d_{TV}$ , which metrizes a space of probability measures defined on a common measurable space. This metric is sometimes computed by the  $L^1$ -distance between the Radon–Nikodym derivatives (i.e., PDFs) associated to the individual probability measures. Specifically, if  $(\Omega, \mathcal{F}, \mu)$  denotes a measure space for which a family of probability measures, denoted by  $\mathcal{P}$ , is defined and absolutely continuous with respect to dominating measure  $\mu$ , then the TV metric between  $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{P}$  is given by

$$d_{TV}(\mathbb{P}_1, \mathbb{P}_2) := \int_{\Omega} |\pi_1 - \pi_2| \, d\mu,$$

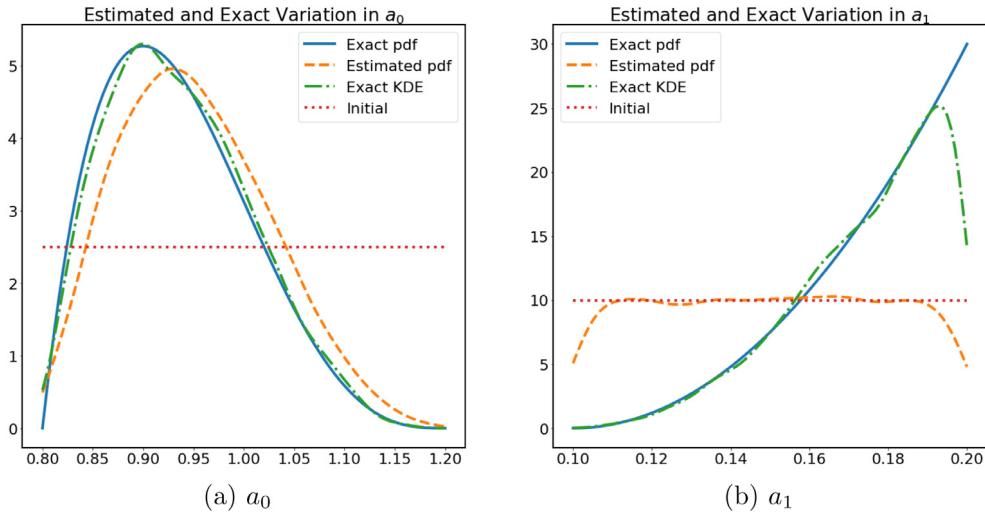
where  $\pi_1$  and  $\pi_2$  denote the Radon–Nikodym derivatives of  $\mathbb{P}_1$  and  $\mathbb{P}_2$ , respectively. Note that

$$d_{TV} : \mathcal{P} \times \mathcal{P} \rightarrow [0, 2].$$

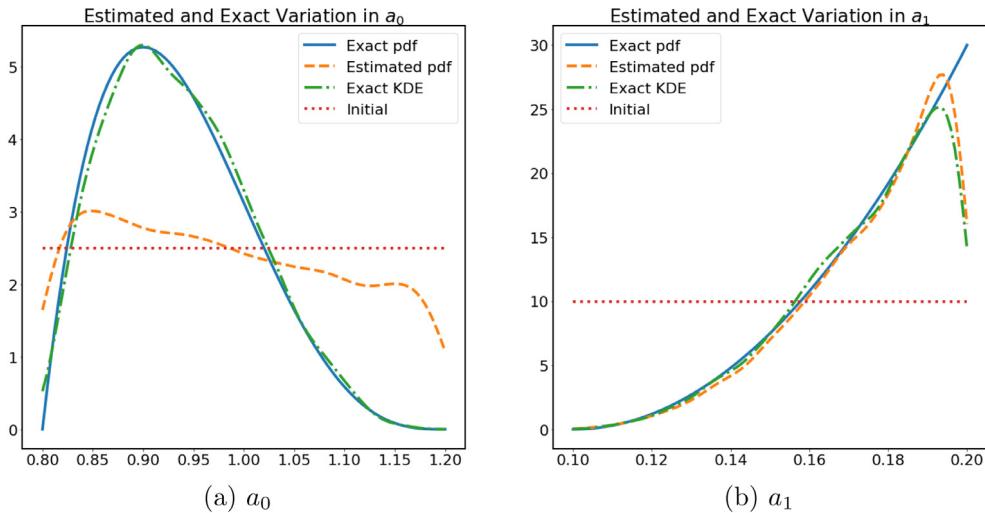
The TV metric computations use the built-in quadrature function available in the `integrate` subpackage of `scipy` with maximum iterations chosen to ensure a relative tolerance of at least  $1.49E-08$  is reached in each case.

In both the figures and tables involving the TV metric, we include a standard KDE approximation of the exact marginal PDF for each parameter computed on the set of  $5E3$  parameter samples used to generate the observed eigenmode data. This provides some perspective on the impact of finite sample error in the observable data on the updated solutions.

Fig. 8 illustrates that QoI #1 is primarily sensitive to changes in the  $a_0$  parameter whereas QoI #2 is primarily sensitive to changes in the  $a_1$  parameter. Subsequently, the updated solutions associated with either using QoI #1 or QoI #2, exclusively, only appears to approximate the true distributions on  $a_0$  or  $a_1$ , respectively, with only minor updates to the initial distributions on  $a_1$  or  $a_0$ , respectively (see Figs. 9 and 10). This is further evidenced by the quantitative results obtained by the TV metric between various marginal PDFs as discussed below.



**Fig. 9.** Results using only QoI #1 from Fig. 6. Exact (in blue) and weighted *kernel density estimates* (KDEs) of the updated (in orange) densities for  $a_0$  (a) and  $a_1$  (b). The weighted KDEs are computed using a Gaussian kernel and weights given by  $\mathcal{R}(a(r))$  evaluated at the predicted QoI samples (see Fig. 7). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



**Fig. 10.** Results using only QoI #2 from Fig. 6. Exact (in blue) and weighted *kernel density estimates* (KDEs) of the updated (in orange) densities for  $a_0$  (a) and  $a_1$  (b). The weighted KDEs are computed using a Gaussian kernel and weights given by  $\mathcal{R}(a(r))$  evaluated at the predicted QoI samples (see Fig. 7). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

From [Table 2](#), we see that the marginal updated PDF for  $a_0$  obtained by using only QoI #1 is closer to the exact PDF and its KDE estimate than the initial PDF (compare 0.203 and 0.163, respectively, to 0.615). However, the distance of this marginal updated PDF for  $a_0$  from the exact PDF and its KDE estimate is more than we would expect from just a KDE estimate of the exact PDF (compare 0.203 and 0.163, respectively, to 0.045). Simultaneously, from [Table 3](#), we see that the marginal updated PDF for  $a_1$  obtained by using only QoI #1 appears to be much closer to the initial PDF than the exact PDF and its KDE estimate (compare 0.052 to 0.769 or 0.733, respectively). Moreover, the distance of this marginal updated PDF for  $a_1$  to the initial PDF is comparable to the distance expected from a KDE estimate of an exact density (compare 0.052 to either 0.068 in this table or to 0.045 in [Table 2](#)).

**Table 2**

The total variation (TV) metrics between different marginal PDFs for  $a_0$  using QoI #1. Each row fixes a PDF  $\pi_1$  and each column fixes a PDF  $\pi_2$  in the computation of  $d_{TV}(\mathbb{P}_1, \mathbb{P}_2)$ .

	$\pi_{a_0}^{\text{update}}$	$\pi_{a_0}^{\text{exact}}$	KDE of $\pi_{a_0}^{\text{exact}}$
$\pi_{a_0}^{\text{init}}$	0.615	0.694	0.679
$\pi_{a_0}^{\text{update}}$	—	0.203	0.163
$\pi_{a_0}^{\text{exact}}$	—	—	0.045

**Table 3**

The total variation (TV) metrics between different marginal PDFs for  $a_1$  using QoI #1. Each row fixes a PDF  $\pi_1$  and each column fixes a PDF  $\pi_2$  in the computation of  $d_{TV}(\mathbb{P}_1, \mathbb{P}_2)$ .

	$\pi_{a_1}^{\text{update}}$	$\pi_{a_1}^{\text{exact}}$	KDE of $\pi_{a_1}^{\text{exact}}$
$\pi_{a_1}^{\text{init}}$	0.052	0.770	0.734
$\pi_{a_1}^{\text{update}}$	—	0.769	0.733
$\pi_{a_1}^{\text{exact}}$	—	—	0.068

**Table 4**

The total variation (TV) metrics between different marginal PDFs for  $a_0$  using QoI #2. Each row fixes a PDF  $\pi_1$  and each column fixes a PDF  $\pi_2$  in the computation of  $d_{TV}(\mathbb{P}_1, \mathbb{P}_2)$ .

	$\pi_{a_0}^{\text{update}}$	$\pi_{a_0}^{\text{exact}}$	KDE of $\pi_{a_0}^{\text{exact}}$
$\pi_{a_0}^{\text{init}}$	0.142	0.694	0.679
$\pi_{a_0}^{\text{update}}$	—	0.561	0.552
$\pi_{a_0}^{\text{exact}}$	—	—	0.045

**Table 5**

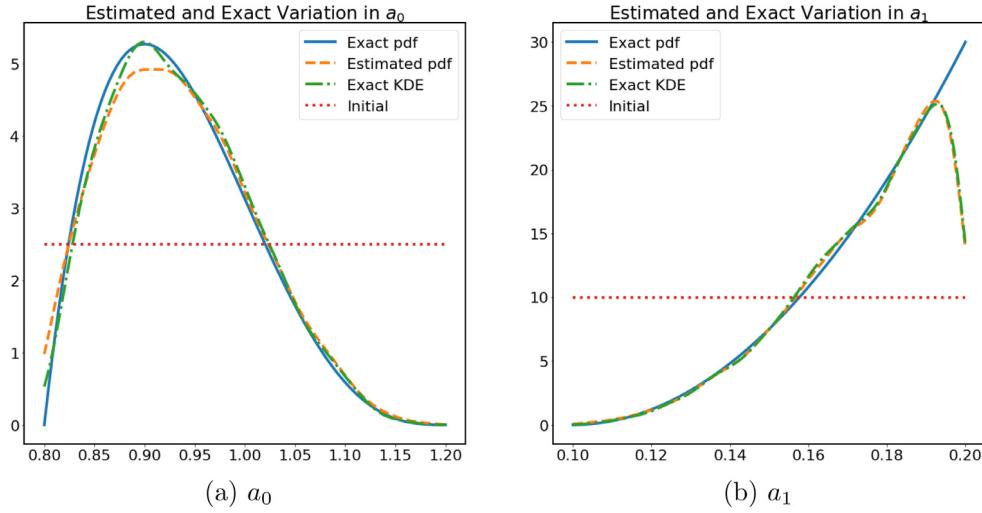
The total variation (TV) metrics between different marginal PDFs for  $a_1$  using QoI #2. Each row fixes a PDF  $\pi_1$  and each column fixes a PDF  $\pi_2$  in the computation of  $d_{TV}(\mathbb{P}_1, \mathbb{P}_2)$ .

	$\pi_{a_1}^{\text{update}}$	$\pi_{a_1}^{\text{exact}}$	KDE of $\pi_{a_1}^{\text{exact}}$
$\pi_{a_1}^{\text{init}}$	0.757	0.770	0.734
$\pi_{a_1}^{\text{update}}$	—	0.072	0.061
$\pi_{a_1}^{\text{exact}}$	—	—	0.068

Almost the opposite is quantitatively observed when using QoI #2 to update the initial PDFs. From [Table 4](#), we see that the marginal updated PDF for  $a_0$  obtained by using only QoI #2 is significantly closer to the initial PDF than either the exact PDF or its KDE estimate than the initial PDF (compare 0.142 to 0.694 or 0.679, respectively). However, this distance from the initial PDF is more than we would expect from a KDE estimate of the initial PDF. Simultaneously, from [Table 5](#), we see that the marginal updated PDF for  $a_1$  obtained by using only QoI #2 appears to be much closer to the exact PDF and its KDE estimate than the initial PDF (compare 0.072 and 0.061, respectively, to 0.757). Moreover, the distance of this marginal updated PDF for  $a_1$  to either the exact PDF or its KDE estimate is similar to other distances observed in these tables of KDE estimates of specified densities.

To summarize the results presented above, QoI #1 produces an updated PDF with a marginal for  $a_0$  that is significantly closer to the exact PDF than the initial, but is still farther away from the exact PDF than expected from a KDE estimate. This QoI #1 also fails to provide any meaningful updates to the initial PDF for  $a_1$ . On the other hand, QoI #2 provides a very reasonable update to  $a_1$  in terms of matching the exact PDF while slightly updating the  $a_0$  parameter. As we see below, the simultaneous use of both QoI produces the best updated PDF to both parameters.

[Fig. 11](#) shows that using both QoI concurrently produces an updated density that exhibits excellent agreement with the true distributions for both parameters.



**Fig. 11.** Results using both QoI from Fig. 6. Exact (in blue) and weighted kernel density estimates (KDEs) of the updated (in orange) densities for  $a_0$  (a) and  $a_1$  (b). The weighted KDEs are computed using a Gaussian kernel and weights given by  $\mathcal{R}(a(r))$  evaluated at the predicted QoI samples (see Fig. 7). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

**Table 6**

The total variation (TV) metrics between different marginal PDFs for  $a_0$  using both QoI. Each row fixes a PDF  $\pi_1$  and each column fixes a PDF  $\pi_2$  in the computation of  $d_{TV}(\mathbb{P}_1, \mathbb{P}_2)$ .

	$\pi_{a_0}^{\text{update}}$	$\pi_{a_0}^{\text{exact}}$	KDE of $\pi_{a_0}^{\text{exact}}$
$\pi_{a_0}^{\text{init}}$	0.642	0.694	0.679
$\pi_{a_0}^{\text{update}}$	—	0.064	0.040
$\pi_{a_0}^{\text{exact}}$	—	—	0.045

**Table 7**

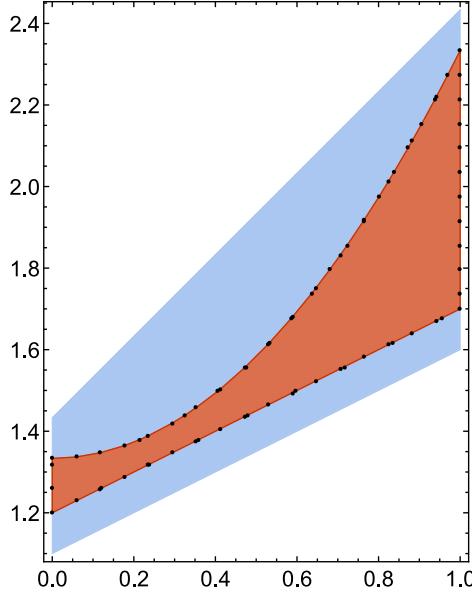
The total variation (TV) metrics between different marginal PDFs for  $a_1$  using both QoI. Each row fixes a PDF  $\pi_1$  and each column fixes a PDF  $\pi_2$  in the computation of  $d_{TV}(\mathbb{P}_1, \mathbb{P}_2)$ .

	$\pi_{a_1}^{\text{update}}$	$\pi_{a_1}^{\text{exact}}$	KDE of $\pi_{a_1}^{\text{exact}}$
$\pi_{a_1}^{\text{init}}$	0.729	0.770	0.734
$\pi_{a_1}^{\text{update}}$	—	0.072	0.013
$\pi_{a_1}^{\text{exact}}$	—	—	0.068

We summarize the TV metric between various (marginal) PDFs for  $a_0$  in Table 6 and for  $a_1$  in Table 7. It is evident that the updated PDFs are both significantly different from the initial PDFs and the distance from the exact PDFs are about what we would expect from a KDE estimate of such PDFs for each parameter.

## 5.2. Example 2: Verification & validation

Suppose a manufacturer of drums seeks verification & validation that their drums are compliant with designs that produce continuous diffusion parameters  $a(r)$  within the large convex-shaped region shown in Fig. 12. We refer to this as the V&V region below. Here, we assume that the manufacturing process stretches elastic membranes over the drum shape in such a way that the resulting diffusion parameter  $a(r)$  is a quadratic function generated in the darker shaded sub-region of Fig. 12, which we refer to as the quadratic region below. It is further assumed that the third-party performing the verification has no knowledge of this quadratic region nor that the curves (along



**Fig. 12.** The V&V region for the diffusion parameter is illustrated by the large orange convex set. The quadratic region containing the true diffusion parameters is illustrated by the blue non-convex subset of the V&V region. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

with their distributions) for the diffusion parameters should be quadratic at all. However, it is straightforward to show that under the sup-norm metric that piecewise-linear continuous splines are dense in the space of continuous diffusion parameters contained within the lightly shaded region. Thus, we assume that the third-party has access to either experimental or simulation data for families of canonical designs involving piecewise-linear continuous splines.

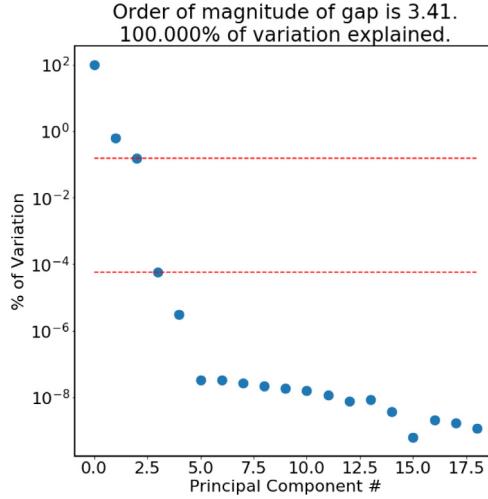
### 5.2.1. Learning possible canonical designs from observed data

As in the previous example, we assume that we can observe up to the first 20 eigenmodes for any drum. However, in contrast to the previous example, we begin with an analysis of *observed* (not predicted) eigenmode data. Specifically, we use simulated data from 5E3 randomly generated quadratic  $a(r)$  functions in the quadratic region of Fig. 12. Performing the standard scaling followed by a PCA on the associated observed eigenmode data, we find that there are three principal components in this data (see Fig. 13). Since any potential QoI is simply a function of the diffusion parameter  $a(r)$ , this implies that variations in  $a(r)$  can be summarized by three variables, which we denote by  $a_0$ ,  $a_1$ , and  $a_2$ .

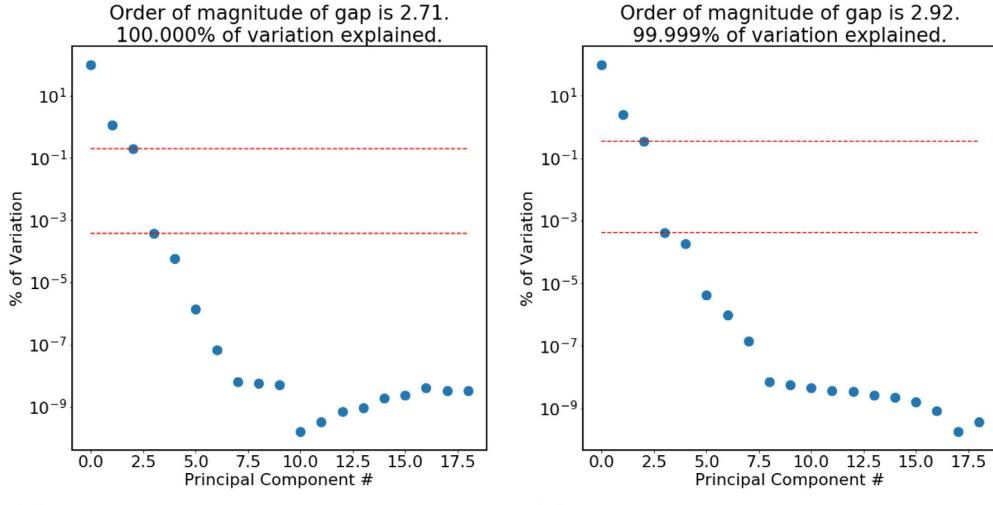
Canonical designs are then restricted to splines with at least three knots. However, in the interest of parsimony, we restrict possible canonical designs to families of splines defined by three knots located at  $k_0 = 0 < k_i < k_1 = 1$  where  $k_i$  is chosen to be either 0.25, 0.5, or 0.75. In other words, two of the three knots are fixed at  $r = 0$  and  $r = 1$  and the location of the interior knot,  $k_i$ , defines which of the three canonical designs are being considered. Since the large lightly shaded region shown in Fig. 12 is convex, the variations of diffusion parameters that belong to any particular canonical design are then described completely by the range of physically allowable values for  $a(k_j)$  for  $j \in \{0, i, 1\}$ , which is simply determined by evaluation of the lower- and upper-bounds for the large lightly shaded region. In the interest of space, we present results only for the canonical designs where  $k_i = 0.5$  or  $k_i = 0.75$ , which prove to be the “best” and “worst” of the three designs, respectively. However, the interested reader can regenerate all these results along with the results for the design with  $k_i = 0.25$  using the provided supplemental Python scripts.

### 5.2.2. Learning the QoI from each canonical design

For each of the canonical designs, we construct a family of 1E4 piecewise-linear splines by generating independent uniform samples over the each of the physically allowable ranges for  $a(k_j)$  for each  $j \in \{0, i, 1\}$ .



**Fig. 13.** Identifying the existence of three QoI using principal component analysis on the observed data set.

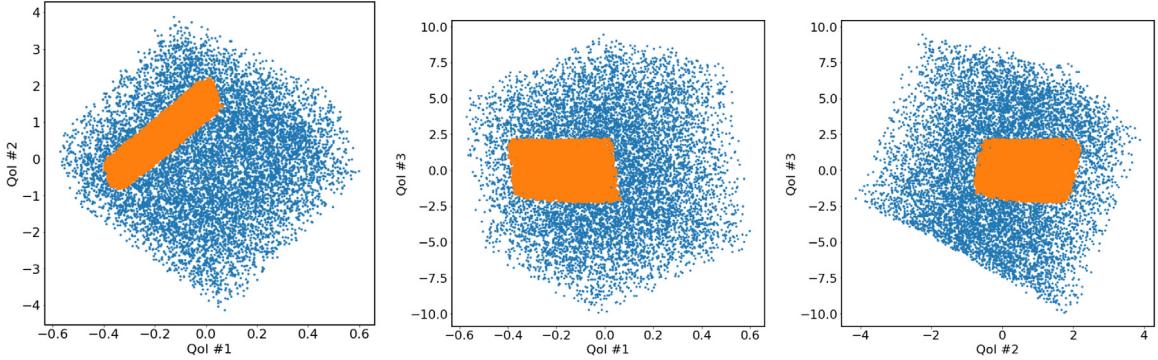


(a) Spectral gap with  $k_i = 0.75$  design. (b) Spectral gap with  $k_i = 0.5$  design.

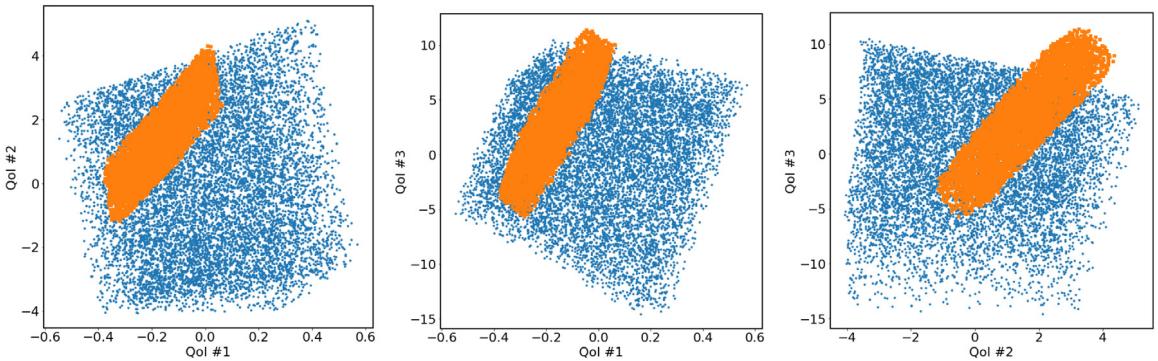
**Fig. 14.** Different QoI defined by all two canonical design sets with different “gaps” in information present in the principal components.

Since the designs can only produce approximations to the true  $a(r)$  variability, we observe slight differences in the PCA for the associated first 20 eigenmodes of each design (see Fig. 14) both between designs and compared to the PCA of the observed data shown in Fig. 13. However, the PCA still leads to the same conclusion that there are three associated QoI defined by the first three principal components associated with each design prediction of eigenmode data.

It is important to note that the QoI maps we learn from the predicted data associated to different canonical design exhibit different capacities to predict the observed data. This is best illustrated by Figs. 15 and 16 where scatter plots show the three different pairs of predicted and observed QoI components associated with the  $k_i = 0.5$  and 0.75 designs, respectively. In these plots, it is evident that the QoI learned from the design with  $k_i = 0.75$  will lead to a violation of the predictability assumption (Assumption 1). However, such qualitative analyses are not always available especially in higher dimensional QoI spaces. Thus, we proceed with solving the inverse problems associated with each design, compute the diagnostic for the inverse problem, and visualize parameter curves to analyze the impact and quality of solutions obtained using either design.



**Fig. 15.** Predicted (blue dots) and observed (orange squares) pairs of QoI data using learned QoI with  $k_i = 0.5$  design.



**Fig. 16.** Predicted (blue dots) and observed (orange squares) pairs of QoI using learned QoI with  $k_i = 0.75$  design.

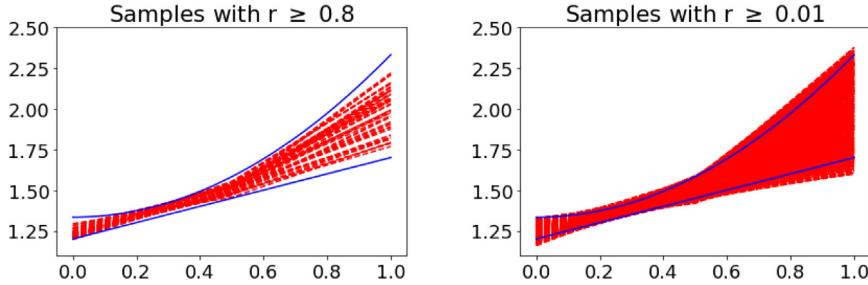
### 5.2.3. Solving the inverse problems

Here, we present results for the solutions of the two inverse problems using the learned three-dimensional QoI maps from the designs with  $k_i = 0.5$  and  $k_i = 0.75$ , respectively. As in the previous example, this requires applying standard density estimation (provided within the module `DrumAnalysis`) to construct the associated predicted and observed densities for each QoI map associated with each design. The associated  $\mathbb{E}_{\text{init}}(\mathcal{R}(a(r)))$  values are estimated to be 0.99 and 0.89 for the solutions associated with the  $k_i = 0.5$  and  $k_i = 0.75$  designs, respectively. This diagnostic immediately indicates that the predictability assumption is violated by the QoI and initial density associated with the  $k + i = 0.75$  design. Thus, that design cannot provide verification that the manufacturer is producing  $a(r)$  curves completely within the V&V region since it cannot explain all of the observed data. However, the other design with  $k_i = 0.5$  can explain all of the observable data and can verify the manufacturer's design. We illustrate the importance of using this diagnostic in making that conclusion since further analysis of the solutions associated with either design is otherwise mostly similar.

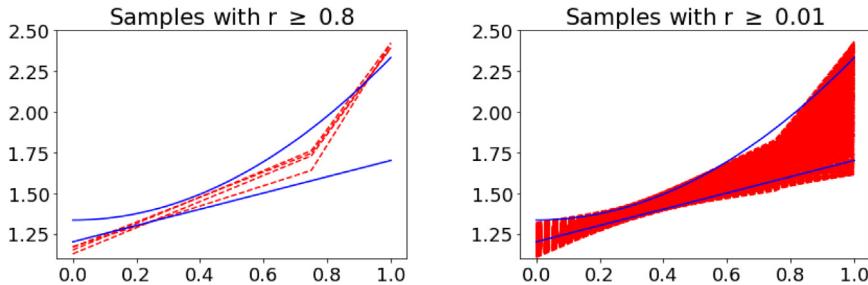
[Fig. 17](#) (resp., [Fig. 18](#)) show all design diffusion curves in the 80th percentile (left) and all design curves deemed “somewhat” likely (1st percentile) as determined by normalized  $\mathcal{R}(a(r))$  values for the design with  $k_i = 0.5$  (resp.,  $k_i = 0.75$ ). Clearly the design with  $k_i = 0.5$  produces more diffusion curves in the 80th percentile compared to the design with  $k_i = 0.75$  and does a better job overall of producing likely diffusion curves that belong or track the quadratic region quite closely. However, the  $k_i = 0.75$  design still does a reasonable job of describing the quadratic region with its likely curves.

To further quantify the performance of these designs, we define the distance of a diffusion curve  $a(r)$  from the quadratic region ( $R$ ) as follows:

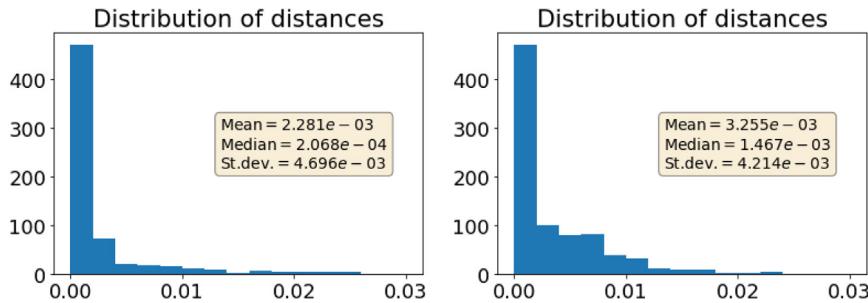
$$d(a(r), R) := \inf_{b(r) \in \mathcal{C}([0,1]) \cap R} \int_{[0,1]} |a(r) - b(r)| \, dr.$$



**Fig. 17.** Likely parameters from cpl3 set with  $r \geq 0.8$  (left) and  $r \geq 0.01$  (right).



**Fig. 18.** Likely parameters from cpl3-right set with  $r \geq 0.8$  (left) and  $r \geq 0.01$  (right).



**Fig. 19.** Statistics on a metric of design fit of likely samples (those with  $\mathcal{R}(a(r)) \geq 0.01$ ) from actual set of  $a(r)$  values. The metric is defined as  $d(x, A)$  where  $A$  is the actual set of  $a(r)$  values and  $d$  is the integral metric on continuous functions.

This distance function is approximated using a uniform regular grid of  $[0, 1]$  and tracking when a given  $a(r)$  (defined by a spline) is either above or below the bounds defining the quadratic region. Then, we compute  $d(a(r), R)$  for all spline approximations to  $a(r)$  in each design with  $\mathcal{R}(a(r)) \geq 0.01$ . In other words, we compute this distance on all “likely” diffusion curves determined by each inverse problem solution. Histograms and statistics of this analysis are shown in Fig. 19. We see that the likely diffusion curves associated with design  $k_i = 0.5$  are, on average, approximately 70% closer to the quadratic region compared to the likely diffusion curves associated with design  $k_i = 0.75$ , with the median distance roughly an order of magnitude less as well. The standard deviations are comparable. However, these types of analyses are not possible when the quadratic region is unknown, which is usually the case in a V&V study such as this. What these analyses do suggest is that even if the predictability assumption is violated, it may still be useful to follow through with the inverse solution for an incorrect design. The analysis of results may lead to insights into choosing different canonical designs that may appear to fit the features that are partially explainable by the incorrect design. The problem of iteratively or adaptively constructing canonical designs is outside the scope of this work though, and we leave it to a future study.

## 6. Conclusions

The manufacturing of drums and the associated quantification of uncertainties has served as the model process throughout our discussion here. The quality control problem has been loosely defined as the process of discovering the distributions of the uncertain inputs, and the V&V problem as the process of separation of sets of probable parameters from those that are merely physically plausible. Both problems have been formulated as stochastic inverse problems, where the system inputs have been discovered by interrogating the outputs using a data-consistent UQ framework based on the disintegration of measures.

The effectiveness of the proposed method has been demonstrated with two examples, where the manufacturing process is modeled via a Dirichlet Laplacian eigenproblem. The diffusion coefficient is parameterized for every instance and the corresponding eigen-information is used in the solution of the inverse problems. Specifically, quantities of interest required in the solution of data-consistent inverse problems are constructed from the correlated eigenvalue data using principal component analysis to extract the information content in each problem. In the V&V example, this is also used to determine the nominal dimension of the parameter space used to define the spatially variable diffusion coefficients. The results from the numerical experiments are exactly aligned with the theoretical considerations and show that the salient aspects of the model are captured accurately.

Indeed, even though we may not hear the shape of the drum, it appears that we can hear the make of the drum.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgment

We acknowledge the computational resources provided by the Aalto Science-IT project.

T. Butler's work is supported by the National Science Foundation, USA under Grant No. DMS-1818941.

### Appendix A. Software dependencies for generating numerical results

The supplemental material allows the interested reader to re-create all of the figures and table data presented in this article. The data files and scripts are available here [31] and the most up-to-date scripts, modules, and data are available at <https://github.com/CU-Denver-UQ/drums-supplemental-material>. Here, we provide a brief description of the data formats and software dependencies required to run the scripts.

All eigenmode data associated with solutions to (2.3) are stored in .mat formats containing descriptive variable names for the various arrays of data in each file. The Python module, DrumAnalysis, and scripts provided as supplemental material to this article were developed using Python 3.7.3 and the following libraries

- `scipy` (version 1.2.1)
- `numpy` (version 1.16.2)
- `sklearn` (version 0.20.3)
- `matplotlib` (version 3.0.3)

The instructions below assume an installation of Python with the appropriate dependencies as described above and that all supplemental material are located within the same path (directory). Then, from a terminal or command prompt, navigate to the appropriate path containing all the files. Use the `python` command to execute either the `Example-1.py` or `Example-2.py` script to generate all of the figures and table data for Sections 5.1 or 5.2, respectively. Running the `Example-1.py` script will prompt the user to enter a maximum number of iterations for the adaptive quadrature used to compute the various TV metrics reported in Section 5.1. Setting this to 1500 will recreate the results in this work, but it may take several minutes to execute. Setting this to 100 will give results leading to the same general conclusions despite the increase in numerical error, but it should execute in a matter of seconds. Of note is that `Example-2.py` will generate additional figures and data that were omitted in this work that use the third design set of parameters mentioned in Section 5.2. This script should take under a minute to run. However, depending on machine specifications and user-edits to either script, they may both take several minutes to execute.

## References

- [1] M. Kac, Can one hear the shape of a drum?, *Amer. Math. Monthly* 73 (part II) (1966) 1–23.
- [2] C. Gordon, D. Webb, S. Wolpert, Isospectral plane domains and surfaces via riemannian orbifolds, *Invent. Math.* 110 (1992) 1–22.
- [3] T.A. Driscoll, Eigenmodes of isospectral drums, *SIAM Rev.* 39 (1) (1997) 1–17.
- [4] W.R. Mateker, J.D. Douglas, C. Cabanatos, I.T. Sachs-Quintana, J.A. Bartelt, E.T. Hoke, A. El Labban, P.M. Beaujuge, J.M.J. Fréchet, M.D. McGehee, Improving the long-term stability of pbdtpp polymer solar cells through material purification aimed at removing organic impurities, *Energy Environ. Sci.* 6 (2013) 2529–2537.
- [5] M. Liu, G.-L. Song, Impurity control and corrosion resistance of magnesium–aluminum alloy, *Corros. Sci.* 77 (2013) 143–150.
- [6] S.-Z. Lu, A. Hellawell, The mechanism of silicon modification in aluminum-silicon alloys: impurity induced twinning, *Metall. Trans. A* 18 (10) (1987) 1721–1733.
- [7] R. Ramesh, M. Mannan, A. Poo, Error compensation in machine tools — a review: Part i: geometric, cutting-force induced and fixture-dependent errors, *Int. J. Mach. Tools Manuf.* 40 (9) (2000) 1235–1256.
- [8] R. Ramesh, M. Mannan, A. Poo, Error compensation in machine tools — a review: Part ii: thermal errors, *Int. J. Mach. Tools Manuf.* 40 (9) (2000) 1257–1284.
- [9] J.V. Abellan-Nebot, J. Liu, F.R. Subirón, J. Shi, State space modeling of variation propagation in multistation machining processes considering machining-induced variations, *J. Manuf. Sci. Eng.* 134 (2) (2012) 021002.
- [10] C.J. Roy, W.L. Oberkampf, A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing, *Comput. Methods Appl. Mech. Engrg.* 200 (25) (2011) 2131–2144.
- [11] R. Rebba, S. Mahadevan, S. Huang, Validation and error estimation of computational models, *Reliab. Eng. Syst. Saf.* 91 (10) (2006) 1390–1397, The Fourth International Conference on Sensitivity Analysis of Model Output (SAMO 2004).
- [12] D. Wu, Z. Lu, Y. Wang, L. Cheng, Model validation and calibration based on component functions of model output, *Reliab. Eng. Syst. Saf.* 140 (2015) 59–70.
- [13] T. Butler, J. Jakeman, T. Wildey, Combining push-forward measures and Bayes' rule to construct consistent solutions to stochastic inverse problems, *SIAM J. Sci. Comput.* 40 (2) (2018) A984–A1011.
- [14] T. Hasselman, G. Lloyd, A top-down approach to calibration, validation, uncertainty quantification and predictive accuracy assessment, *Comput. Methods Appl. Mech. Engrg.* 197 (29) (2008) 2596–2606, Validation Challenge Workshop.
- [15] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, D.B. Rubin, *Bayesian Data Analysis*, third ed., Chapman and Hall/CRC, 2013.
- [16] D. Calvetti, J. Kaipio, E. Somersalo, Inverse problems in the Bayesian framework, *Inverse Problems* 30 (11) (2014) 110–301.
- [17] J.O. Berger, E. Moreno, L.R. Pericchi, M.J. Bayarri, J.M. Bernardo, J.A. Cano, J. De la Horra, J. Martín, D. Ríos-Insúa, B. Betrò, A. Dasgupta, P. Gustafson, L. Wasserman, J.B. Kadane, C. Srinivasan, M. Lavine, A. O'Hagan, W. Polasek, C.P. Robert, C. Goutis, F. Ruggeri, G. Salinetti, S. Sivaganesan, An overview of robust Bayesian analysis, *Test* 3 (1) (1994) 5–124.
- [18] M. Burger, F. Lucka, Maximum posterior estimates in linear inverse problems with log-concave priors are proper bayes estimators, *Inverse Problems* 30 (11) (2014) 114004.
- [19] A.W.v.d. Vaart, *Asymptotic statistics*, Cambridge University Press, Cambridge, UK; New York, NY, USA;, 1998.
- [20] C.K. Wikle, L.M. Berliner, N. Cressie, Hierarchical Bayesian space-time models, *Environ. Ecol. Stat.* 5 (2) (1998) 117–154.
- [21] A.E. Gelfand, A. Kottas, A computational approach for full nonparametric Bayesian inference under dirichlet process mixture models, 2002.
- [22] Y.W. Teh, *Dirichlet Process*, Springer US, Boston, MA, 2010, pp. 280–287.
- [23] L. Li, Z. Lu, A new method for model validation with multivariate output, *Reliab. Eng. Syst. Saf.* 169 (2018) 579–592.
- [24] D. Boffi, Finite element approximation of eigenvalue problems, *Acta Numer.* 19 (2010) 1–120.
- [25] H. Hakula, T. Tuominen, Mathematica implementation of the high order finite element method applied to eigenproblems, *Computing* 95 (1) (2013) 277–301.
- [26] C. Dellacherie, P. Meyer, *Probabilities and Potential*, North-Holland Publishing Co., Amsterdam, 1978.
- [27] A.L. Gibbs, F.E. Su, On choosing and bounding probability metrics, *Internat. Statist. Rev.* (2002) 419–435.
- [28] G.R. Terrell, D.W. Scott, Variable kernel density estimation, *Ann. Statist.* 20 (3) (1992) 1236–1265.
- [29] L. Devroye, L. Györfi, *Nonparametric Density Estimation: The L<sub>1</sub> View*, Wiley, New York, 1985.
- [30] Wolfram Research, Inc. *Mathematica*, Version 12.0, Champaign, IL, 2019.
- [31] T. Butler, CU-Denver-UQ/drums-supplemental-material: 10.5281/ZENODO.3905344, Drums Module and Data version 1.0, Zenodo, 2020.