


PAPER

Data-consistent inversion for stochastic input-to-output maps^{*}

To cite this article: Troy Butler *et al* 2020 *Inverse Problems* **36** 085015

View the [article online](#) for updates and enhancements.



IOP | ebooksTM

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Data-consistent inversion for stochastic input-to-output maps*

Troy Butler^{1,3} , T Wildey² and Tian Yu Yen¹

¹ Department of Mathematical and Statistical Sciences, University of Colorado Denver, Denver, CO 80202, United States of America

² Sandia National Laboratories, Center for Computing Research, Albuquerque, NM 87185, United States of America

E-mail: troy.butler@ucdenver.edu and tmwilde@sandia.gov

Received 28 December 2019, revised 8 April 2020

Accepted for publication 1 May 2020

Published 20 August 2020



Abstract

Data-consistent inversion is a recently developed measure-theoretic framework for solving a stochastic inverse problem involving models of physical systems. The goal is to construct a probability measure on model inputs (i.e., parameters of interest) whose associated push-forward measure matches (i.e., is consistent with) a probability measure on the observable outputs of the model (i.e., quantities of interest). Previous implementations required the map from parameters of interest to quantities of interest to be deterministic. This work generalizes this framework for maps that are stochastic, i.e., contain uncertainties and variation not explainable by variations in uncertain parameters of interest. Generalizations of previous theorems of existence, uniqueness, and stability of the data-consistent solution are provided while new theoretical results address the stability of marginals on parameters of interest. A notable aspect of the algorithmic generalization is the ability to query the solution to generate independent identically distributed samples of the parameters of interest without requiring knowledge of the so-called stochastic parameters. This work therefore extends the applicability of the data-consistent inversion framework to a much wider class of problems. This includes those based on purely experimental and field data where only a subset of conditions are either controllable or can be

*The views expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

³Author to whom any correspondence should be addressed.

documented between experiments while the underlying physics, measurement errors, and any additional covariates are either uncertain or not accounted for by the researcher. Numerical examples demonstrate application of this approach to systems with stochastic sources of uncertainties embedded within the modeling of a system and a numerical diagnostic is summarized that is useful for determining if a key assumption is verified among competing choices of stochastic maps.

Keywords: uncertainty quantification, inverse problem, push-forward measure, pullback measure, data consistent

(Some figures may appear in colour only in the online journal)

1. Introduction

As computational models are increasingly used to aid engineering design [12, 22], shape public policy [7], and predict the future behavior of physical and biological systems [1, 20], quantifying the uncertainties in model inputs (i.e., parameters of interest) that influence solution characteristics is becoming more critical. Such quantification of uncertainties may occur by formulating and solving an inverse problem using data associated with observable model outputs [i.e., quantities of interest (QoI)] to infer likely model inputs. Uncertainties in observable data often necessitate the formulation and solution of a *stochastic* inverse problem where the solution on the space of parameters takes the form of a probability measure. With a probability measure, or its associated probability density function (PDF), we may assess the relative likelihoods of parameters, estimate the probability of important events (such as those associated with catastrophic system failure), or determine the robustness of system designs.

Uncertainties in observable data for some QoI may be due to a variety of epistemic and aleatoric factors, i.e., lack of precise knowledge of the system or fundamentally random properties of causal factors. Of particular interest in this work is the quantification of aleatoric factors due to variability in parameters of interest. Consider a manufacturing process for an engineered system relying upon various materials and sub-components sourced from other manufacturers. In raw materials, concentrations and locations of impurities can vary substantially, which can be reduced through refinement but never eliminated. Moreover, the production of various sub-components may only be produced to be within certain tolerances. These are sources of aleatoric uncertainty that will impact parameters for a computational model of the system. These subsequently impact both system performance and the QoI computed from the model. Viewing the underlying uncertainty impacting system inputs as a random process, the associated parameters of interest and QoI are random variables with uncertain distributions that we wish to determine.

Bayesian methods [2, 8, 13, 17] are some of the most popular means of inferring probabilistic descriptions of model parameters from QoI data. We first introduce some basic notation prior to comparing the extension of the data-consistent approach—which is the focus of this work—to Bayesian methods. First, we use λ to denote parameters of interest and ξ to denote stochastic parameters. We denote by $Q(\lambda)$ (or simply Q) the QoI map defined only on the space of parameters of interest and $\hat{Q}(\lambda, \xi)$ (or simply $\hat{Q}$) to denote the QoI map defined on the joint space of parameters of interest and stochastic parameters. In other words, Q comes from a deterministic (usually physics-based) computational model where a single choice of λ is mapped to a single QoI whereas $\hat{Q}$ maps λ to a new random variable whose output depends upon the ξ parameter.

In a typical Bayesian framework, one of the initial assumptions is that data, denoted by d , are obtained from a map $\hat{Q}$ defined by $Q + \xi$, where ξ is a random variable following a given distribution, usually assumed to be Gaussian. The random variable ξ is used to describe the uncertainty associated with measurement errors that can theoretically be reduced using improved instrumentation to collect more precise data. The Bayesian solution for the parameters of interest λ , known as a posterior, is given by a conditional density, $\pi(\lambda|d) \propto \pi(\lambda)\pi(d|\lambda)$, where $\pi(\lambda|d)$ is the posterior, $\pi(\lambda)$ is a prior density on parameters of interest, and $\pi(d|\lambda)$ is the data-likelihood function whose form is often determined by products of the noise distribution evaluated at residuals between Q and d . The posterior is used to assess the relative likelihoods that a fixed estimate for the parameters of interest could have produced all of the data d . Subsequently, it is common to use one of two point-estimators, 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 of interest [4]. This use of a single point estimate of the parameters of interest as a solution to the inverse problem is actually quite reasonable within this framework since as more data are collected (i.e., as the dimension of the data vector given by d grows), the posterior will often become concentrated (or ‘spiked’) around a single parameter of interest, which is a phenomenon explained by the Bernstein–von Mises theorem [31]. In other words, this typical formulation treats the uncertainty in parameters of interest as epistemic rather than aleatoric.

When parameters in the model are instead viewed as sources of aleatoric uncertainty (i.e., they possess some natural, irreducible variability described by a distribution as in this work), then hierarchical Bayesian methods provide an alternative to the regular Bayesian framework outlined above [32]. Hierarchical Bayesian methods commonly specify prior distributions from a parametric family of distributions, where so-called hyper-parameters are introduced as random variables into the inference process to describe the uncertainty in the parametric family of distributions used for the prior. These approaches and the software packages that implement them (e.g., pyMC3 [23]) consequently require an increase in the number of samples that must be computed to obtain accurate inferences because samples must be drawn from a ‘hierarchy’ of distributions. Other non-parametric, hierarchical Bayesian methods for modeling aleatoric uncertainty exist (namely Dirichlet processes) but are typically more costly than regular parametric Bayesian methods.

An alternative approach for constructing non-parametric estimates of densities on parameters of interest, referred to in this work as data-consistent inversion, was recently developed in [5]. In this data-consistent approach, any prior knowledge of aleatoric uncertainty on the parameters of interest is used to construct an *initial* density on these parameters. This initial density serves a different purpose than the prior used in Bayesian or hierarchical Bayesian approaches, so we do not refer to it as a prior density in this work to avoid confusion. Subsequently, assuming the variation in QoI data is due to this aleatoric uncertainty, the map Q is used to formulate a *prediction* of the observed data distribution using the push-forward of the initial density. Then, the ratio between the observed and predicted densities, evaluated on the map Q , serves to *update* the initial density only in directions informed by the data. This updated density on the parameters of interest has the property that its associated push-forward through the map Q exactly reconstructs the observed density. In other words, the solution is a pullback of the observed density through Q that is regularized by the initial density only in directions in the parameter space to which Q is not sensitive. In this work, we generalize the theoretical and algorithmic approaches for sampling the parameters of interest from this data-consistent solution when the map is given by $\hat{Q}$. Moreover, the algorithm we provide generates samples for any of the specified parameters of interest associated with the

Table 1. Summary of notation and terminology for data-consistent inversion with a deterministic map.

Notation	Description
$(\Lambda, \mathcal{B}_\Lambda)$	Measurable space of parameters of interest.
$(\mathcal{D}, \mathcal{B}_\mathcal{D})$	Measurable space of observable data.
$\mu_\Lambda, \mu_\mathcal{D}$	Dominating measure on spaces Λ and $\mathcal{D}$, respectively.
Q	Quantity of interest (QoI) map between $(\Lambda, \mathcal{B}_\Lambda)$ and $(\mathcal{D}, \mathcal{B}_\mathcal{D})$.
$\mathbb{P}_{\text{init}}, \pi_{\text{init}}$	Initial probability measure and associated density assumed on Λ .
$\mathbb{P}_{\text{obs}}, \pi_{\text{obs}}$	Probability measure and associated density based on observations in $\mathcal{D}$.
$\mathbb{P}_{\text{predict}}, \pi_{\text{predict}}$	Push-forward probability measure and associated density of $\mathbb{P}_{\text{init}}$ propagated through Q .
$r(\lambda)$	Ratio of $\pi_{\text{obs}}(Q(\lambda))$ to $\pi_{\text{predict}}(Q(\lambda))$ used to update the initial density π_{init} .
$\mathbb{P}_{\text{update}}, \pi_{\text{update}}$	Updated probability measure and associated density for $(\Lambda, \mathcal{B}_\Lambda)$ using data-consistent inversion.

data-consistent solution without requiring an explicit representation of $\hat{Q}$ or knowledge of any of the stochastic parameters ξ .

The rest of this paper is outlined as follows. In section 2, we first review the basic theory of the data-consistent framework. In section 3, we discuss the extension of the data-consistent framework to uncertain stochastic maps, providing a rigorous theoretical framework for extending data-consistent inversion including theory of existence, uniqueness, and stability of solutions. We then provide an algorithm for sampling directly from data-consistent marginal distributions in section 4 that does not require knowledge of either the stochastic components of the map or their distributions. Finally, in section 5, we demonstrate this extended framework using numeric examples which illustrate the flexibility of the data-consistent approach in handling stochastic uncertainties within model maps. Concluding remarks follow in section 6.

2. Overview of data-consistent inversion

2.1. Notation, terminology and assumptions

In this section, we summarize the data-consistent framework for deterministic maps. For ease of reference, the notation and terminology used in this summary is consolidated in table 1.

Let Λ denote the space of model inputs characterizing physical properties (e.g., a diffusion coefficient or initial/boundary data), which we refer to as either the physical parameters or parameters of interest. In this work, we assume $\Lambda \subset \mathbb{R}^P$ and let $\mathcal{B}_\Lambda$ denote the Borel σ -algebra on $\mathbb{R}^P$ restricted to Λ so that $(\Lambda, \mathcal{B}_\Lambda)$ defines a measurable space of physical parameters. Let $Q : \Lambda \rightarrow \mathcal{D}$ denote the quantity of interest (QoI) map from physical parameters to the space of observable model output data denoted by $\mathcal{D} \subset \mathbb{R}^D$, which we refer to as the data space. Denote by $\mathcal{B}_\mathcal{D}$ the Borel σ -algebra on $\mathbb{R}^D$ restricted to $\mathcal{D}$ so that $(\mathcal{D}, \mathcal{B}_\mathcal{D})$ defines a measurable data space. We assume that Q is a measurable piecewise-smooth map from $(\Lambda, \mathcal{B}_\Lambda)$ to $(\mathcal{D}, \mathcal{B}_\mathcal{D})$. This implies that the *contour* σ -algebra, defined by

$$\mathcal{C}_\Lambda := \{Q^{-1}(A) : A \in \mathcal{B}_\mathcal{D}\},$$

is a sub- σ -algebra of $\mathcal{B}_\Lambda$. We assume that some initial guess of uncertainty on $(\Lambda, \mathcal{B}_\Lambda)$ is given in the form of a probability measure denoted by $\mathbb{P}_{\text{init}}$. Let $\mathbb{P}_{\text{predict}}$ denote the corresponding push-forward measure on $(\mathcal{D}, \mathcal{B}_\mathcal{D})$ defined by

$$\mathbb{P}_{\text{predict}}(A) := \mathbb{P}_{\text{init}}(Q^{-1}(A)), \forall A \in \mathcal{B}_{\mathcal{D}}.$$

Since $\mathbb{P}_{\text{predict}}$ is a solution to a forward uncertainty quantification (UQ) problem that may occur before any data are collected on $\mathcal{D}$, we refer to $\mathbb{P}_{\text{predict}}$ as the predicted probability measure. We assume that once observable model output data are collected and analyzed that uncertainty in these data are represented in the form of a probability measure $\mathbb{P}_{\text{obs}}$ on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$.

If the probability measures are absolutely continuous with respect to dominating measures, then probability density functions (i.e., Radon–Nikodym derivatives, which we sometimes refer to simply as densities or PDFs) associated with each probability measure can be substituted in the analysis. We denote such dominating measures on the parameter and data spaces by μ_{Λ} and $\mu_{\mathcal{D}}$, respectively. In this work, we assume that μ_{Λ} and $\mu_{\mathcal{D}}$ are given by the P - and D -dimensional Lebesgue measures, respectively. We now formalize the definition of the type of inverse problem considered in this work and what is meant by consistency of the solution we seek.

Definition 1 (inverse problem and consistent measure). Assume we are given both

- an initial probability measure $\mathbb{P}_{\text{init}}$ on $(\Lambda, \mathcal{B}_{\Lambda})$ that is absolutely continuous with respect to μ_{Λ} and admits a density π_{init} , and
- a probability measure $\mathbb{P}_{\text{obs}}$ on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ that is absolutely continuous with respect to $\mu_{\mathcal{D}}$ and admits a density π_{obs} .

The inverse problem is to *update* the initial probability measure in such a way that the updated measure, denoted by $\mathbb{P}_{\text{update}}$, admits a density π_{update} that is consistent in the sense that

$$\mathbb{P}_{\text{update}}(Q^{-1}(A)) = \int_{Q^{-1}(A)} \pi_{\text{update}} d\mu_{\Lambda} = \int_A \pi_{\text{obs}} d\mu_{\mathcal{D}} = \mathbb{P}_{\text{obs}}(A), \forall A \in \mathcal{B}_{\mathcal{D}}. \quad (1)$$

We refer to $\mathbb{P}_{\text{update}}$ as the *solution* to the inverse problem.

2.2. Existence and uniqueness of solutions

The following assumption on the push-forward of the initial density, π_{predict} , guarantees the existence and uniqueness of such an updated measure. It is also practically necessary for implementing a simple rejection sampling algorithm to draw samples from the updated measure, as discussed below.

Assumption 1 (predictability assumption). There exists $C > 0$ such that $\pi_{\text{obs}}(q) \leq C\pi_{\text{predict}}(q)$ for a.e. $q \in \mathcal{D}$.

We refer to this assumption as the *predictability assumption* since it implies that any output event with non-zero *observed* probability has a non-zero *predicted probability*. While the initial density serves a different purpose than the prior used in a Bayesian setting, we may nonetheless draw an analogy to the convention in Bayesian methods of choosing a prior to be as uninformative as possible.

In [5], a disintegration theorem [10] is used to prove the following theorem where $\mu_{\Lambda, q}$ denotes the disintegration of the dominating measure, μ_{Λ} . Conceptually, it is convenient to think of this theorem as a nonlinear version of Fubini's theorem where the disintegrated measure, $\mu_{\Lambda, q}$, is analogous to using lower-dimensional Lebesgue measures in the iterated portions of the integral.

Theorem 2 (existence and uniqueness [5]). *The probability measure $\mathbb{P}_{\text{update}}$ on $(\Lambda, \mathcal{B}_{\Lambda})$ defined by*

$$\mathbb{P}_{\text{update}}(A) = \int_{\mathcal{D}} \left(\int_{A \cap Q^{-1}(q)} \frac{\pi_{\text{init}}(\lambda) \pi_{\text{obs}}(Q(\lambda))}{\pi_{\text{predict}}(Q(\lambda))} d\mu_{\Lambda, q}(\lambda) \right) d\mu_{\mathcal{D}}(q), \quad \forall A \in \mathcal{B}_{\Lambda} \quad (2)$$

is a consistent solution to the inverse problem in the sense of (1) and is uniquely determined for a given initial probability measure $\mathbb{P}_{\text{init}}$ on $(\Lambda, \mathcal{B}_{\Lambda})$.

The probability density of the consistent solution is given by

$$\pi_{\text{update}}(\lambda) = \pi_{\text{init}}(\lambda) \frac{\pi_{\text{obs}}(Q(\lambda))}{\pi_{\text{predict}}(Q(\lambda))}, \quad \lambda \in \Lambda. \quad (3)$$

Since we assume π_{init} and π_{obs} are given, it is evident that the updated density (3) is immediately obtained once the predicted density, π_{predict} , is constructed. We usually rewrite the updated density in the form,

$$\pi_{\text{update}}(\lambda) = \pi_{\text{init}}(\lambda) r(\lambda), \quad \text{with } r(\lambda) = \frac{\pi_{\text{obs}}(Q(\lambda))}{\pi_{\text{predict}}(Q(\lambda))}, \quad (4)$$

where the ratio $r(\lambda)$ is interpreted as providing the necessary re-weighting (i.e., update) to the initial density to produce the desired updated density. Below, we provide a few remarks to give both insight and intuition into the structure of this updated measure and density defining the solution to the inverse problem.

Remark 2.1. $\mathbb{P}_{\text{update}}$ is a *pullback* measure of $\mathbb{P}_{\text{obs}}$ and $\mathbb{P}_{\text{obs}}$ is the push-forward measure of $\mathbb{P}_{\text{update}}$.

Remark 2.2. The predictability assumption implies that rejection sampling can be applied to a set of independent identically distributed (i.i.d.) samples from π_{init} to produce a set of i.i.d. samples from π_{update} . Specifically, in [5], evaluation of $r(\lambda)$ on an i.i.d. set of samples from π_{init} is used to formulate the rejection ratio for each sample in the set. While we omit the details here, a generalization of this approach for stochastic maps is summarized in algorithm 1 in section 4.2.

Remark 2.3. The data-consistent updated density is a fundamentally different object than the typical Bayesian posterior. Whereas the Bayesian posterior weights the prior distribution by the ratio of a data likelihood function to a constant, the updated density weights the initial density by the ratio of the observed density to the predicted density (i.e., $r(\lambda)$). The impact of $r(\lambda)$ is to fundamentally update the *structure* of the initial density only in certain local directions in parameter space. Specifically, $r(\lambda)$ is a fixed constant when restricting λ to the same (generalized) contour—defined by a singleton in the sub- σ -algebra $\mathcal{C}_{\Lambda}$ —so that $Q(\lambda)$ is a fixed constant. Consequently, evaluation of π_{update} at points along the same contour are given by evaluation of π_{init} re-scaled by the same constant, but this constant will likely change between contours. Even supposing that the likelihood function and the observed density are taken to be the same function, the data consistent and Bayesian approaches produce very different solutions due to the differences in re-weighting the prior/initial density, as demonstrated by example 7.2 in [5].

3. A data-consistent inversion framework for stochastic maps

In some scenarios, the QoI map from physical parameter space to data space is more appropriately modeled as a stochastic process to account for processes impacting the data that are not accounted for in (or perhaps are unexplainable by) the physics-based computational model. In

Table 2. Summary of additional notation and terminology for extending the data-consistent framework to stochastic maps.

Notation	Description
$(\Xi, \mathcal{B}_\Xi)$	Measurable space of stochastic parameters.
$\mathbb{P}_{\text{stoch}}, \pi_{\text{stoch}}$	Probability measure and associated density for $(\Xi, \mathcal{B}_\Xi)$.
$(\Lambda \times \Xi, \mathcal{B}_\Lambda \times \mathcal{B}_\Xi)$	Joint measurable space of parameters of interest and stochastic parameters.
$\hat{Q}$	A stochastic QoI map from Λ to $\mathcal{D}$ viewed as a deterministic QoI map from the joint space $\Lambda \times \Xi$ to the data space $\mathcal{D}$.
$\hat{\mathbb{P}}_{\text{init}}, \hat{\pi}_{\text{init}}$	Initial probability measure and associated density assumed on joint space $\Lambda \times \Xi$.
$\hat{\mathbb{P}}_{\text{predict}}, \hat{\pi}_{\text{predict}}$	Push-forward probability measure and associated density of $\hat{\mathbb{P}}_{\text{init}}$ propagated through $\hat{Q}$.
$\hat{r}(\lambda, \xi)$	Ratio of $\pi_{\text{obs}}(\hat{Q}(\lambda, \xi))$ to $\hat{\pi}_{\text{predict}}(\hat{Q}(\lambda, \xi))$ used to update the initial density $\hat{\pi}_{\text{init}}$.
$\hat{\mathbb{P}}_{\text{update}}, \hat{\pi}_{\text{update}}$	Updated probability measure and associated density for $(\Lambda \times \Xi, \mathcal{B}_\Lambda \times \mathcal{B}_\Xi)$ using data-consistent inversion.
$\hat{\mathbb{P}}_{\text{update}, \Lambda}, \hat{\pi}_{\text{update}, \Lambda}$	Updated marginal probability measure and associated density for $(\Lambda, \mathcal{B}_\Lambda)$ using generalization of data-consistent inversion.

these cases, we say that the QoI map is stochastic, and in this section we extend and interpret the data-consistent inverse methodology for stochastic maps. We prove several stability results with a particular emphasis on the (marginal) updated PDF on the physical parameter space. A numerical algorithm is also provided for constructing and sampling from this updated PDF on physical parameters.

3.1. Extended framework and existence/uniqueness

3.1.1. Notation. The extension of the data-consistent framework to stochastic maps introduces several new terms that we summarize in table 2 for ease of reference.

In the rest of this work, we assume that variability in the observable data used to construct $\mathbb{P}_{\text{obs}}$ are due to additional sources of uncertainty beyond any assumed aleatoric uncertainty in physical parameters. It is relatively common in the UQ literature, and especially in Bayesian frameworks, to consider additional sources of uncertainty in observable data as attributed to measurement errors represented by additive noise models on the QoI (e.g., see [9, 18, 19, 29]). More recently, the idea of model inadequacy (e.g., due to missing or parameterized physics at various spatial or temporal scales) has gained more attention as a way to describe additional uncertainty that is embedded within the model (e.g., see [21] and the references therein). We additionally assume that it is possible to model these additional sources of uncertainty using random vectors

$$\xi(\omega) : \Omega \rightarrow \Xi \subset \mathbb{R}^S$$

where $(\Omega, \mathcal{F}, \mathbb{P})$ defines the probability space on the random outcomes, denoted by ω , belonging to the sample space Ω . We refer to these random variables as the *stochastic parameters* to clearly distinguish $\xi(\omega) \in \Xi$ and its impact on predicted data from that of a physical parameter $\lambda \in \Lambda$. In other words, for a fixed $\lambda \in \Lambda$, the predicted datum associated with the output of the QoI map is defined probabilistically over sets of possible output values.

Generally, we omit direct reference to the sample space and work directly in the output space of the stochastic parameters. Therefore, we use the simpler notation $\xi \in \Xi$ to refer to a particular realization of the stochastic parameters that impacts a predicted QoI datum. Subsequently, we let $\mathcal{B}_\Xi$ denote the Borel σ -algebra of $\mathbb{R}^S$ restricted to Ξ , and $\mathbb{P}_{\text{stoch}}$ denotes the




Figure 1. The solid (blue) curve represents a deterministic observable map, $Q(\lambda)$. The red dots indicate noisy observations representing a sample from the distributions plotted vertically on the dotted lines. These distributions indicate the relative likelihoods of obtaining an observable datum given a fixed value of input parameter to the deterministic observable map.

push-forward probability measure of $\mathbb{P}$ on $(\Xi, \mathcal{B}_\Xi)$, i.e.,

$$\mathbb{P}_{\text{stoch}}(A) := \mathbb{P}(\xi^{-1}(A)), \quad \forall A \in \mathcal{B}_\Xi.$$

We again assume that there exists a dominating measure μ_Ξ on $(\Xi, \mathcal{B}_\Xi)$ so that the PDF of $\mathbb{P}_{\text{stoch}}$ can be defined by its Radon–Nikodym derivative, which we denote by π_{stoch} . Finally, we replace the QoI map notation with $\hat{Q}(\lambda, \xi)$ to denote the measurable map from the measurable product space $(\Lambda \times \Xi, \mathcal{B}_\Lambda \times \mathcal{B}_\Xi)$ to the observable data space $(\mathcal{D}, \mathcal{B}_\mathcal{D})$.

3.1.2. Conceptual example part I: the maps and spaces. We use a simple example to provide context for some of this notation. Suppose $\Lambda = [0, 9]$ and an initial deterministic model defines the QoI map,

$$Q(\lambda) = (\lambda - 1)(\lambda - 5)(\lambda - 7), \quad (5)$$

illustrated by the blue curve in figure 1, but that observable data are polluted by measurement errors, which we represent using an additive noise model so that

$$\hat{Q}(\lambda, \xi) = Q(\lambda) + \xi, \quad \xi \sim N(0, \sigma^2). \quad (6)$$

Here, we take $\sigma = 7$. Then, for any fixed λ , $\hat{Q}(\lambda, \cdot)$ is a random variable following a $N(Q(\lambda), 7^2)$ distribution as illustrated by the two distributions plotted along the vertical lines associated with two different physical parameters. The red dots indicate observable data for the two different physical parameters, which come from two particular realizations of the stochastic parameters. Note that neither of these data points are on the curve defined by $Q(\lambda)$, yet both are in the range defined by $Q(\Lambda)$. Thus, it is at least possible to use the map Q to invert such observable data. However, doing so would clearly result in inaccurate *point* estimates of the physical parameters that led to such observations.

In figure 2, we plot $\hat{Q}(\lambda, \xi)$ over the product space $\Lambda \times \Xi$ (where, for illustrative purposes, we truncate Ξ by $\pm 4\sigma$). The blue curve illustrates the mapping $\hat{Q}(\lambda, 0) = Q(\lambda)$. The red dots are the same data points as before, and while they are not on the blue curve, they do exist on the *surface* defined by $\hat{Q}$ over the product space. Thus, applying $\hat{Q}^{-1}$ to either of these data points produces a *contour* in $\Lambda \times \Xi$ that contains the true physical parameter value responsible for producing the corresponding datum. In other words, the map $\hat{Q}$ is a completely deterministic and measurable real-valued function whose pre-images exist as sets in the product space $\Lambda \times \Xi$.




Figure 2. The blue curve $Q(\lambda)$ and sampled points from figure 1 are embedded onto a deterministic surface defined over the input space of physical parameters for the model and stochastic parameters defining the additive noise of model observables.

3.1.3. Existence and uniqueness. We let $\hat{\mathbb{P}}_{\text{init}}$ and $\hat{\pi}_{\text{init}}$ denote an initial (joint) probability measure and its PDF, respectively, on $(\Lambda \times \Xi, \mathcal{B}_\Lambda \times \mathcal{B}_\Xi)$. For simplicity in notation, we let π_{init} and π_{stoch} denote the marginals of $\hat{\pi}_{\text{init}}$ on Λ and Ξ , respectively.

Let $\hat{\mathbb{P}}_{\text{predict}}$ and $\hat{\pi}_{\text{predict}}$ denote the push-forward of the initial probability measure and its PDF, respectively, through the map $\hat{Q}$. We assume that the form of the predictability assumption holds with $\hat{\pi}_{\text{predict}}$ replacing π_{predict} . With this notation and updated predictability assumption, the following result for stochastic maps is a direct extension of theorem 2 where $\mu_{\Lambda \times \Xi, q}$ now denotes a disintegration of the dominating product measure $\mu_{\Lambda \times \Xi} := \mu_\Lambda \times \mu_\Xi$. Otherwise, the proof is identical to the proof of theorem 2.

Theorem 3 (existence and uniqueness for stochastic maps). *The probability measure $\hat{\mathbb{P}}_{\text{update}}$ on $(\Lambda \times \Xi, \mathcal{B}_\Lambda \times \mathcal{B}_\Xi)$ defined by*

$$\hat{\mathbb{P}}_{\text{update}}(A) = \int_{\mathcal{D}} \left(\int_{A \cap \hat{Q}^{-1}(q)} \hat{\pi}_{\text{init}}(\lambda, \xi) \frac{\pi_{\text{obs}}(\hat{Q}(\lambda, \xi))}{\hat{\pi}_{\text{predict}}(\hat{Q}(\lambda, \xi))} d\mu_{\Lambda \times \Xi, q}(\lambda, \xi) \right) d\mu_{\mathcal{D}}(q), \quad (7)$$

for all $A \in \mathcal{B}_\Lambda \times \mathcal{B}_\Xi$ is a consistent solution to the inverse problem in the sense that $\mathbb{P}_{\text{obs}}$ is a push-forward of $\hat{\mathbb{P}}_{\text{update}}$ (or, equivalently, that $\hat{\mathbb{P}}_{\text{update}}$ is a pullback of $\mathbb{P}_{\text{obs}}$) and is uniquely determined for a given initial probability measure $\hat{\mathbb{P}}_{\text{init}}$ on $(\Lambda \times \Xi, \mathcal{B}_\Lambda \times \mathcal{B}_\Xi)$.

The PDF of this consistent solution is given by

$$\hat{\pi}_{\text{update}}(\lambda, \xi) = \hat{\pi}_{\text{init}}(\lambda, \xi) \frac{\pi_{\text{obs}}(\hat{Q}(\lambda, \xi))}{\hat{\pi}_{\text{predict}}(\hat{Q}(\lambda, \xi))}, \quad (\lambda, \xi) \in \Lambda \times \Xi. \quad (8)$$

As before, this updated PDF is immediately obtained once the predicted density, $\hat{\pi}_{\text{predict}}$, is constructed, and we again rewrite this updated PDF in the form

$$\hat{\pi}_{\text{update}}(\lambda, \xi) = \hat{\pi}_{\text{init}}(\lambda, \xi) \hat{r}(\lambda, \xi), \quad \text{with } \hat{r}(\lambda, \xi) := \frac{\pi_{\text{obs}}(\hat{Q}(\lambda, \xi))}{\hat{\pi}_{\text{predict}}(\hat{Q}(\lambda, \xi))}. \quad (9)$$




Figure 3. The four steps of sampling from the updated PDF. Left: samples from the joint initial PDF. Right: samples from the updated density in the joint parameter space obtained using rejection sampling.

We emphasize that this updated PDF is updating the initially assumed *joint* structure on the product space. Subsequently, the ratio of the observed density to this predicted density, denoted by $\hat{r}$, serves to update both the joint density as well as the marginal densities, π_{stoch} and π_{init} , simultaneously.

3.1.4. Conceptual example part II: the measures. We now return to the simple example from section 3.1.2 (with Q and $\hat{Q}$ defined by (5) and (6)) to both illustrate these ideas and demonstrate one approach to (1) sample from the updated PDF and (2) verify the samples come from a consistent solution. Suppose that

$$\pi_{\text{init}} \sim \text{Beta}(2, 5).$$

Here, by $\text{Beta}(2, 5)$, we denote a standard $\text{Beta}(\alpha, \beta)$ distribution that is transformed from $[0, 1]$ to $\Lambda = [0, 9]$ by a simple scaling of inputs. To approximate $\hat{\pi}_{\text{predict}}$, we first generate $1E + 4$ independent identically distributed (i.i.d.) samples from $\hat{\pi}_{\text{init}}$ (see the left plot of figure 3), evaluate $\hat{Q}$ on each of these samples (see the left plot figure 4), and then use a standard Gaussian kernel density estimator (GKDE) to approximate $\hat{\pi}_{\text{predict}}$ (shown as the blue dashed-dotted curve in figure 5). Suppose we observe data that follows a normal distribution with an observed mean and standard deviation of 8 and 4, respectively, so that $\pi_{\text{obs}} \sim N(8, 4^2)$ (shown as the red curve in figure 5). Clearly, the predictability assumption is satisfied, so we use a straightforward rejection sampling algorithm on the initial $1E + 4$ samples to generate a set of approximately $3E + 3$ i.i.d. samples from the updated density (shown in the plot of figure 3).

To verify that this set of accepted i.i.d. samples does in fact come from a consistent solution, we propagate these back through the map $\hat{Q}$ (shown in the right plot of figure 4). We again use a GKDE to estimate the density, and verify that it reconstructs the observed density (compare the black dashed and red curves in figure 5). Numerical diagnostics also verify that the sample average and standard deviation of the push-forward of these updated samples are within 2% of the observed values of 8 and 4, respectively.

Following construction or sampling of the joint updated density on $\Lambda \times \Xi$, it is straightforward to analyze and compare the *marginals* of this updated density even if the joint updated density has complex structure. For example, we can simply use the individual components of the i.i.d. samples of $\hat{\pi}_{\text{update}}$ obtained by rejection sampling as above to construct i.i.d. sets of samples for the marginals.

In the plots of initial and updated samples in Λ shown in figure 3, we also summarize on the top and right of the axes the associated marginal densities on Λ and Ξ , respectively. We may use a mixture of qualitative and quantitative analyses to interpret the specific updates to the initially




Figure 4. The four steps of sampling from the updated PDF. The surface represents the map $\hat{Q}$. Left: evaluation of the initial samples using $\hat{Q}$. Right: evaluation of the updated samples using $\hat{Q}$.




Figure 5. Visualizing the predicted density obtained by evaluation of $\hat{Q}$ on the initial samples (blue dash-dotted curve), the observed density (red solid curve), and the consistency of the updated density is verified by comparing its push-forward (black dashed curve) to the observed density.

assumed uncertainties in the physical and stochastic parameters. For example, qualitatively it appears that the initial uni-modal beta distribution on Λ is updated into a bi-modal non-parametric distribution whereas the updated distribution on Ξ is still approximately normal but with a shift in mean and variance. This is numerically verified by comparing the sample mean and standard deviation of the initial set of ξ -samples (which are approximately 0 and 6.9, respectively) to the sample mean and standard deviation of the updated set of ξ -samples (which are approximately -2.4 and 6).

3.2. Stability results: theory and interpretations

We now consider the stability of the updated joint probability measure with respect to perturbations in the initial, observed, or predicted probability measures. We use the total-variation (TV) metric, 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 $\mathcal{P}$, are defined and absolutely continuous with respect to dominating measure μ ,

then the TV metric between $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{P}$ is given by

$$d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_2) := \int_{\Omega} |\pi_1 - \pi_2| \, d\mu, \quad (10)$$

where π_1 and π_2 denote the Radon–Nikodym derivatives of $\mathbb{P}_1$ and $\mathbb{P}_2$, respectively.

We assume that any perturbations in either the initial or observed measures still lead to probability measures that are absolutely continuous with respect to the dominating measure on the corresponding measure space, and that the predictability assumption is never violated. With this assumption, we use the TV metric in the form given by (10) to state the stability results. First, we provide a general definition of stability used in this work.

Definition 4 (stability). Let $\mathbb{P}_\alpha$ be a probability measure on $(\Omega_\alpha, \mathcal{F}_\alpha, \mu_\alpha)$ and $\mathbb{P}_\beta$ be a probability measure on $(\Omega_\beta, \mathcal{F}_\beta, \mu_\beta)$ that depends upon $\mathbb{P}_\alpha$ (e.g., through a transformation or functional dependency). Let $\widetilde{\mathbb{P}}_\alpha$ denote a perturbation to $\mathbb{P}_\alpha$ and $\widetilde{\mathbb{P}}_\beta$ the corresponding perturbation to $\mathbb{P}_\beta$. We say that $\mathbb{P}_\beta$ is *stable* with respect to perturbations in $\mathbb{P}_\alpha$ if for all $\epsilon > 0$ there exists $\delta > 0$ such that

$$d_{\text{TV}}(\mathbb{P}_\alpha, \widetilde{\mathbb{P}}_\alpha) < \delta \Rightarrow d_{\text{TV}}(\mathbb{P}_\beta, \widetilde{\mathbb{P}}_\beta) < \epsilon. \quad (11)$$

With this definition of stability, the following corollary extends several results originally proven in [5] for deterministic maps.

Corollary 5. *The consistent solution $\widehat{\mathbb{P}}_{\text{update}}$ is stable with respect to perturbations in (i) $\widehat{\mathbb{P}}_{\text{init}}$, (ii) $\mathbb{P}_{\text{obs}}$, and (iii) $\widehat{\pi}_{\text{predict}}$.*

Proof. The stability with respect to perturbations in $\mathbb{P}_{\text{obs}}$ follows from theorem 4.1 and corollary 4.3 in [5] after making straightforward changes in notation to account for the disintegration over the joint space $\Lambda \times \Xi$ (details omitted here). The stability with respect to perturbations in $\widehat{\mathbb{P}}_{\text{init}}$ and $\widehat{\pi}_{\text{predict}}$ follows from theorems 4.5 and 5.1, respectively, in [5] after making similar changes in notation. $\square$

A practical interpretation of the stability results (i) and (ii) in corollary 5 is that small errors in the specification of $\widehat{\mathbb{P}}_{\text{init}}$ or $\mathbb{P}_{\text{obs}}$ (i.e., in setting up the inverse problem) lead to small errors in the solution given by $\widehat{\mathbb{P}}_{\text{update}}$. Stability result (iii) in corollary 5 is interpreted best in terms of numerical errors in approximations. Specifically, even if the specifications of $\widehat{\mathbb{P}}_{\text{init}}$ or $\mathbb{P}_{\text{obs}}$ are exact, the PDF of the predicted measure, $\widehat{\pi}_{\text{predict}}$, is often numerically approximated (e.g., by propagating a finite set of samples from $\widehat{\mathbb{P}}_{\text{init}}$ and then applying a GKDE as in the conceptual example above), which leads to a numerical approximation of the consistent solution. Viewing these numerical approximations as perturbations, we interpret this last stability result as stating that as numerical errors in the approximation of the predicted density are made small, the errors in the numerical approximation to the updated density are also small.

4. The marginal updated measure on physical parameters

From (8), specification of $\widehat{\pi}_{\text{init}}$ is required to form $\widehat{\pi}_{\text{update}}$, which implicitly assumes knowledge of π_{stoch} . However, it is possible to generate i.i.d. samples of the random vector $(\lambda, \xi) \sim \widehat{\pi}_{\text{update}}$ without direct knowledge of $\widehat{\pi}_{\text{init}}$ using rejection sampling on a set of i.i.d. samples generated according to this (potentially unknown) joint initial density.

In practice, we may fail to have knowledge of either π_{stoch} or the particular ξ -components of the joint initial sample sets even if π_{stoch} is known. For example, in models exhibiting chaos

(such as Lorenz models), the impact of rounding errors on computationally predicted data may be represented as a stochastic parameter. In such cases, the specification of π_{stoch} and knowledge of the exact sample taken from π_{stoch} may prove to be an intractable proposition.

In less extreme cases where either the exact π_{stoch} or a good approximation of it is known, we may still fail to have knowledge of the stochastic samples that impact predicted data. Moreover, it may be unreasonable to assume that the stochastic parameters are independent of the physical parameters so that even if knowledge of π_{stoch} is available, the joint structure of $\hat{\pi}_{\text{init}}$ may remain unknown. In general settings where a physical model is used in a laboratory setting to obtain predicted data, certain subsets of physical parameters may be specified/controlled between experiments while stochastic parameters representing various sources of error or uncertainty in the experimental setup or data may not be known. For a specific example, consider the use of a wave tank to simulate storm surge in a scaled physical model of a coastal community. It is well-established (see, e.g., [3]) that storm surge is sensitive to the bathymetry and Manning's coefficient of roughness, which are impacted by sediment transport (e.g., due to shipping channels or agricultural runoff). Repeated wave tank experiments can be run with various values of bathymetry and Manning's coefficient of roughness specified by the experimenter to create an ensemble of predicted storm surge data. This represents a specification of the samples of physical parameters coming from an implicitly defined π_{init} . However, assuming knowledge of π_{stoch} exists to describe measurement errors (e.g., due to splashing effects and instrument errors), we are unlikely to know the exact measurement errors for any given experiment that impact the recorded maximum wave height data. Moreover, the measurement errors are likely to be higher for more turbulent flows suggesting some conditional dependence of the stochastic parameters on the physical parameters. In the examples described above, we are likely to only have information on the physical parameter components of the sample set generated from an (unknown) $\hat{\pi}_{\text{init}}$. This motivates the material below, which

- (a) Proves stability of the marginal updated probability measure on physical parameters with respect to perturbations in the (unknown) $\hat{\mathbb{P}}_{\text{update}}$ and subsequently with respect to perturbations in both $\hat{\mathbb{P}}_{\text{init}}$ and $\mathbb{P}_{\text{obs}}$; and
- (b) Provides a simple algorithm for sampling directly from the marginal updated distribution on physical parameters that does not require any knowledge of the stochastic parameters.

4.1. Stability of the marginal

First, we prove a fundamental result that the TV metric between two joint distributions on the same product space is always greater than the TV metric between *any* of the marginals. This result is subsequently used to prove stability results for the marginal updated probability measure on physical parameters.

Theorem 6. *Let $(\Omega_n, \mathcal{F}_n, \mathbb{P}_n)$ denote an n -dimensional probability space, $m < n$, and $\mathbb{P}_m$ denote the marginal probability measure of $\mathbb{P}_n$ on any m -dimensional probability subspace denoted by $(\Omega_m, \mathcal{F}_m, \mathbb{P}_m)$. Also, let μ_n and μ_m denote the dominating (product) measures on $(\Omega_n, \mathcal{F}_n)$ and $(\Omega_m, \mathcal{F}_m)$, respectively. Then, if $\widetilde{\mathbb{P}}_n$ denotes any perturbation to $\mathbb{P}_n$ and $\widetilde{\mathbb{P}}_m$ denotes the corresponding perturbation to the marginal probability measure $\mathbb{P}_m$,*

$$d_{\text{TV}}(\mathbb{P}_m, \widetilde{\mathbb{P}}_m) \leq d_{\text{TV}}(\mathbb{P}_n, \widetilde{\mathbb{P}}_n). \quad (12)$$

In the proof below, we use the following notation

- $\pi_n, \widetilde{\pi}_n, \pi_m$, and $\widetilde{\pi}_m$ denote the PDFs of $\mathbb{P}_n, \widetilde{\mathbb{P}}_n, \mathbb{P}_m$, and $\widetilde{\mathbb{P}}_m$, respectively;

- $k := n - m$ and $(\Omega_k, \mathcal{F}_k)$ denotes the k -dimensional measurable subspace with dominating (k -dimensional product) measure μ_k such that $\Omega_n = \Omega_m \times \Omega_k$ and $\mathcal{F}_n$ is generated by the completion of the product σ -algebra between $\mathcal{F}_m$ and $\mathcal{F}_k$.

Proof. Using the notation conventions listed above along with the definition of marginal probability measures and Fubini's theorem, we have

$$\begin{aligned} d_{\text{TV}}(\mathbb{P}_m, \widetilde{\mathbb{P}}_m) &= \int_{\Omega_m} |\pi_m - \widetilde{\pi}_m| \, d\mu_m = \int_{\Omega_m} \left| \int_{\Omega_k} (\pi_n - \widetilde{\pi}_n) \, d\mu_k \right| \, d\mu_m \\ &\leq \int_{\Omega_m} \int_{\Omega_k} |\pi_n - \widetilde{\pi}_n| \, d\mu_k \, d\mu_m = \int_{\Omega_n} |\pi_n - \widetilde{\pi}_n| \, d\mu_n = d_{\text{TV}}(\mathbb{P}_n, \widetilde{\mathbb{P}}_n). \end{aligned}$$

□

Recalling definition 4, an immediate consequence of theorem 6 is that every marginal probability measure is stable with respect to perturbations in the corresponding joint probability measure. Let $\widehat{\mathbb{P}}_{\text{update}, \Lambda}$ and $\widehat{\pi}_{\text{update}, \Lambda}$ denote the *marginal* updated probability measure and density, respectively, on the physical parameter space $(\Lambda, \mathcal{B}_\Lambda)$ induced by $\widehat{\mathbb{P}}_{\text{update}}$.

Corollary 7. *The marginal updated probability measure $\widehat{\mathbb{P}}_{\text{update}, \Lambda}$ on physical parameter space $(\Lambda, \mathcal{B}_\Lambda)$ is stable with respect to perturbations in $\widehat{\mathbb{P}}_{\text{init}}$, $\mathbb{P}_{\text{obs}}$, and $\widehat{\mathbb{P}}_{\text{predict}}$.*

Proof. This follows from combining corollary 3.1 and theorem 4.1. □

4.2. Sampling directly from the marginal

Algorithm 1 summarizes how to generate an i.i.d. sample set from an approximate marginal updated PDF on physical parameters without requiring knowledge of the stochastic parameter samples involved in these computations. The algorithm is based on a fundamental fact that if

$$\left\{ (X_1^{(i)}, X_2^{(i)}, \dots, X_n^{(i)}) \right\}_{i=1}^K$$

is an i.i.d. set of K samples from an n -dimensional joint distribution, then for $m < n$,

$$\left\{ (X_1^{(i)}, X_2^{(i)}, \dots, X_m^{(i)}) \right\}_{i=1}^K$$

is an i.i.d. set of K samples from the m -dimensional marginal distribution defined by the first m -components of the probability space.

The algorithm is both straightforward and easy to implement given two arrays of sample sets (an input array and the corresponding output array) and a PDF on the data space. The only knowledge that is required to use this algorithm are the physical parameter components from an initial sample set, the corresponding set of predicted output data, and specification of an observed PDF. Knowledge of either π_{stoch} (or, in fact, $\widehat{\pi}_{\text{init}}$) or the actual values of the ξ -components of this initial sample set are not required. Moreover, knowledge of how the stochastic parameters impact the map is *not* required if we are given the set of output samples

$$\left\{ \widehat{Q}(\lambda^{(i)}, \xi^{(i)}) \right\}_{i=1}^N$$

in the algorithm directly. In other words, we do not actually need to specify the map $\widehat{Q}$ explicitly as long as we are able to query the desired output data of the model (or experiment).

Algorithm 1. Sampling from the marginal updated PDF.**Required inputs:**

- List of the λ -components from an i.i.d. sample set

$$\{(\lambda^{(i)}, \xi^{(i)})\}_{i=1}^N \sim \hat{\pi}_{\text{init}}.$$

- Corresponding predicted sample set

$$\{\hat{Q}^{(i)}\}_{i=1}^N := \{\hat{Q}(\lambda^{(i)}, \xi^{(i)})\}_{i=1}^N.$$

- The observed data PDF π_{obs} .

Pre-processing computations:

1. Use $\{\hat{Q}^{(i)}\}_{i=1}^N$ to construct $\widetilde{\pi}_{\text{predict}} \approx \hat{\pi}_{\text{predict}}$ (e.g., using GKDE).
2. Compute

$$\{\hat{r}^{(i)}\}_{i=1}^N := \left\{ \frac{\pi_{\text{obs}}(\hat{Q}^{(i)})}{\widetilde{\pi}_{\text{predict}}(\hat{Q}^{(i)})} \right\}_{i=1}^N.$$

3. Estimate $M = \max_{\lambda} \hat{r}(\lambda, \xi)$ with $M \approx \max_{1 \leq i \leq N} \hat{r}^{(i)}$.

Rejection sampling:

Set $K = 0$.

for $i = 1, \dots, N$ **do**

Generate a random number $y^{(i)} \sim U([0, 1])$;

Compute the ratio $\eta^{(i)} = \hat{r}^{(i)} / M$;

if $y^{(i)} < \eta^{(i)}$ **then**

Accept and set $K = K + 1$ and $\lambda_{\text{accept}}^{(K)} := \lambda^{(i)}$;

else

Reject $\lambda^{(i)}$;

end

end

Output: $\{\lambda_{\text{accept}}^{(k)}\}_{k=1}^K$.

This is a particularly useful feature when the model is a computational black box for which we set only some inputs defined as physical parameters without setting or knowing other parameter values that may either be chosen automatically by the code or defined by unknown round-off errors that impact the model outputs and are treated as the stochastic parameters. Similarly, this algorithm can also be applied to experimentally obtained data where knowledge of certain physical parameter values between experiments is available (e.g., due to knowledge of the experimental setup) while knowledge of other aspects of the model or data acquisition, such as missing physics or measurement errors, can be treated as unknown stochastic parameters with unspecified distributions.

4.3. Comments and variation of algorithm 1

Algorithm 1 will generate at least one accepted sample by design. Specifically, any sample giving the computed bound M utilized in the algorithm will subsequently have a normalized rejection ratio $\eta^{(i)} = 1$ and automatically be accepted. Such samples are then interpreted as the ‘most likely’ of all initial samples considered by the algorithm. However, if an exact M is known or an over-estimated value of M is used in the algorithm, then it is theoretically possible

that no samples are accepted. In either case, the final number of samples that are accepted, denoted by K , may be undesirably small for a given number of input samples, denoted by N . While the development and analysis of more efficient sampling schemes is outside the scope of this work, it is the topic of ongoing research, and we provide some comments in the conclusions about this topic.

A simple alternative to algorithm 1 is to remove rejection sampling entirely and simply return the array of ratios given by $\{r(\hat{Q}^{(i)})\}_{i=1}^N$ for all N samples. These ratios provide useful insight into both the marginal updated PDF on physical parameters and in predictive UQ analyses. For instance, we can use a weighted GKDE on the physical parameter samples $\{\lambda^{(i)}\}_{i=1}^N$ with weights given by the ratios to estimate the marginal PDFs on either the entire physical parameter space or on individual physical parameters. Alternatively, if there are QoI labeled as ‘prediction’ QoI (i.e., quantities for which there are no current observable data), we may propagate the larger set of physical parameter samples and use the associated ratios to compute weighted GKDE estimates of the distributions for these QoI. In the numerical examples of section 5, we use either rejection sampling or weighted GKDEs to analyze the marginal updated PDF on physical parameters.

An important feature of algorithm 1 (or the alternative discussed above) is that while samples are *initially* generated in the parameter space using Monte Carlo sampling, all subsequent computations take place in the data space defined by the range of the QoI map. Specifically, only the corresponding push-forward samples, $\{\hat{Q}^{(i)}\}_{i=1}^N$, are used to estimate the predicted density and ratios that are used in the remainder of the algorithm. Consequently, the algorithm scales very well as the dimension of the parameter space increases if the dimension of the data space stays fixed and the variance does not increase. On the other hand, the use of a simple GKDE to estimate the predicted density negatively impacts how this algorithm scales with respect to the dimension of the data space. Such issues of scalability are explored for the deterministic variant of this algorithm in [5], which can be viewed as a special case of algorithm 1 that omits a stochastic parameter. In this paper, the final numerical example of section 5 demonstrates the ability of the algorithm to solve high-dimensional stochastic inverse problems by computing a reference distributions using this special case of algorithm 1 from a 100-dimensional space for the parameters of interest.

5. Examples

Below, we illustrate how the extension to the data-consistent framework outlined above is used to solve inverse problems with QoI that incorporate additional stochastic sources of uncertainties. We first consider a linear ‘wobbly plate’ example that is motivated by common types of measurement devices used in a variety of earth science applications. We then consider a variation of an example presented in [5] regarding inferences into stochastic parameters of a Karhunen–Loève expansion used to define a diffusion coefficient in an elliptic differential equation.

When reading these examples, it is useful to keep in mind the following take home messages. First, it is possible to solve the stochastic inverse problem using the *wrong* model/QoI map as long as the predictability assumption holds. However, the predictability assumption can be numerically verified (as demonstrated in the wobbly plate example). Such a numerical diagnostic can prove useful in determining if a model or QoI map along with variations in the physical parameters are unable to explain the observable data. For instance, in [6], the predictability assumption is checked for a sequence of approximate maps converging to




Figure 6. Schematic of ‘wobbly-plate’ example setup. Two sensors are placed below the wobbly-plate to measure its height relative to the sensor.

an exact map. Moreover, in [6], it is shown that the error in the sequence of updated densities associated with the sequence of approximate maps is bounded by the error in the approximate maps. Second, when a model/QoI map is modified to account for other sources of uncertainty/variation impacting the observational data, the inferences we draw from the updated PDF on the physical parameters are changed. Specifically, we generally avoid ‘over-updating’ the initial PDFs on the physical parameters when we account for other sources of uncertainty impacting the observational data. Remarks on this are provided both within the examples and in the conclusions.

5.1. ‘Wobbly plate’

5.1.1. Motivation. Monitoring the evolution of landscapes is important in many environmental and civil engineering applications. In erosion analysis, such monitoring is often done using a micro-erosion meter (MEM) or simple variants of such devices [28]. For instance, MEMs have been used to measure erosion in cave limestone [25], building stone decay [30], and most commonly in the erosion of shore platforms [27] that can significantly impact a coastal areas susceptibility to storm surge. A typical MEM has three legs that form a triangular base on which either a flat plate or other fixture mimicking a flat surface is placed along with a gauge that measures the movement of the plate as erosion occurs (e.g., see [28] and the figures and references within). For predicting volcanic eruptions, a common type of monitoring device is a tiltmeter that is placed on the side of the volcano [11, 26]. When the main chamber of the volcano fills with magma, swelling on the side of the volcano occurs which is recorded by the change in angle of the tiltmeter.

Whether it is a tiltmeter, an MEM, or other similar earth monitoring system devices, the natural variation in the physical parameters describes a significant amount of variation in the measurements over time. However, uncertainties in the actual deployment of the devices can impact the inferences drawn from these measurements. For instance, any initial slope or roughness to a surface impacts the measurements as well as any miss-calibration of the device that is either initially present or may occur over time.

To describe how the approach in this paper can be used to model and quantify the impact of such uncertainties on physical parameters, we consider the problem of determining the distribution of slope parameters in a ‘wobbly plate’ using height measurements obtained by accurate laser readings. This is depicted in figure 6. We study how measurement uncertainties either due to instrumentation error or misplacement of the lasers impacts results.




Figure 7. (Left) shows initial description of slopes (λ_1, λ_2) and the target distribution of slopes $(\lambda_1, \lambda_2) \sim U([0.85, 1.45] \times [1.6, 1.85])$. (Right) shows density estimates (using GKDE on the data-sets) of the simulated height measurement data at locations $\mathbf{x}_A$ and $\mathbf{x}_B$.

5.1.2. Model setup. Consider a square plate that is ‘wobbling’ randomly around a point in space centered just above the origin. The height, y , of the wobbling plate at any location $(x_1, x_2) \in \mathbb{R}^2$ is given by

$$y = y_0 + \lambda_1 x_1 + \lambda_2 x_2$$

where y_0 is the height of the plate above the origin and (λ_1, λ_2) are the slopes of the plate at a particular snapshot of time.

Suppose we are interested in estimating the distribution of the slope parameters $\lambda = (\lambda_1, \lambda_2)$ of the wobbly-plate using two accurate measurement devices (e.g., as given by laser devices) placed at locations $\mathbf{x}_A$ and $\mathbf{x}_B$ to take repeated measurements of the height of the plate (see figure 6 for an illustration). This implies a QoI map $Q: \Lambda \rightarrow \mathbb{R}^2$:

$$\mathbf{q} = Q(\lambda) = \mathbf{y}_0 + \mathbf{X}\lambda$$

where $\mathbf{q}$ is a vector of the height measurements at location $\mathbf{x}_A$ and $\mathbf{x}_B$, $\mathbf{y}_0$ is the vector (y_0, y_0) representing the height of the plate above the origin, λ is a vector of the slope parameters of interest, $(\lambda_1, \lambda_2)^T$, and $\mathbf{X}$ is a 2×2 matrix with the coordinates of the measuring instruments $\mathbf{X} = \begin{bmatrix} \mathbf{x}_A \\ \mathbf{x}_B \end{bmatrix}$.

5.1.3. Data consistent update with error-free setup. A data-generating uniform distribution of slope parameters λ in the box $[0.85, 1.45] \times [1.6, 1.85]$ is used to simulate 250 i.i.d. observed QoI data assuming the laser locations are setup exactly at $\mathbf{x}_A = (0.6, 0.6)$ and $\mathbf{x}_B = (0.8, 0.6)$, with $y_0 = 3$. The data-generating distribution of λ (the target) and the simulated distributions for the QoI (computed using standard Gaussian kernel density estimates) are illustrated in figure 7. The goal is to use the QoI distributions to update a sample set of λ drawn from a different initial distribution that is consistent with these observations.

Suppose the initial distribution of λ is uniform on $[0, 2] \times [0, 2]$, i.e., $\pi_{\text{init}} \sim U([0, 2] \times [0, 2])$ represents the initially assumed distribution over all physically plausible slope parameters. Using algorithm 1, updated samples (i.e., accepted samples) of slope parameters are computed. In the absence of any other uncertainties, the QoI map is 1–1. Subsequently, the updated sample set of λ is in general agreement with the ‘target’ domain associated with the data-generating distribution of parameters as illustrated in the left plot of figure 8. There




Figure 8. The data consistent update in the absence of additional stochastic uncertainties.




Figure 9. (Left) shows specified locations $\mathbf{x}_A$ and $\mathbf{x}_B$ and the actual perturbed locations of the measurement instruments. (Right) shows a Gaussian kernel density estimation of the simulated height data with stochastic noise in the locations $\mathbf{x}_A$ and $\mathbf{x}_B$.

is a slight over-estimation of this region that is due to the over-smoothing effect of the QoI densities from using GKDE.

5.1.4. Competing error models and a numerical diagnostic. Now consider a situation where the data are polluted by stochastic sources of uncertainty not explained by the original QoI map Q but instead by some map $\hat{Q}$. We explicitly address the predictability assumption in this case where we assume competing models are used to account for the presence of additional stochastic sources of uncertainty.

To illustrate these concepts, we simulate height data using the same target distribution of λ described previously but by perturbing the locations $\mathbf{x}_A$ and $\mathbf{x}_B$ by stochastic noise $\xi_A, \xi_B \sim N(\mathbf{0}, \delta^2 \mathbf{I})$, with $\delta = 0.075$. This models the situation where the setup of the device may contain errors (e.g., due to manufacturing imperfections or improper field deployment). Figure 9 shows the simulated locations and resulting height data.

Knowledge of the device and how errors inherent to design setup impact observations leads to the following map $\hat{Q}(\lambda, \xi)$:

$$\mathbf{q} = \hat{Q}(\lambda, \xi) = \mathbf{y}_0 + (\mathbf{X} + \xi)\lambda$$

where ξ is now the matrix $\xi = \begin{bmatrix} \xi_A \\ \xi_B \end{bmatrix}$. In other words, the stochastic parameters are *embedded* in the map.

Alternatively, in much of the uncertainty quantification literature, it is common to assume that measurement data are polluted by additive noise. In other words, noise in the measurements of the heights at points $\mathbf{x}_A$ and $\mathbf{x}_B$, may instead be applied so that the map $\widehat{Q}(\lambda, \xi)$ is defined as

$$\mathbf{q} = \widehat{Q}(\lambda, \xi) = Q(\lambda) + \xi, \quad \xi \sim N(\mathbf{0}, \sigma^2 \mathbf{I}).$$

Competing hypothesized models of $\widehat{Q}$ may also be the result of differing expert opinions. Below, we compare the data consistent updates using the additive noise model versus the embedded location noise model above to illustrate (1) the potential of this framework for distinguishing the capabilities of competing models for constructing data-consistent updates, and (2) the impact on inferences drawn on physical parameters.

First, we discuss a useful numerical diagnostic. If the predictability assumption is satisfied, then constructing $\hat{r}(\lambda, \xi)$ results in the formal construction of the density $\hat{\pi}_{\text{update}}$. Subsequently, we have that

$$\mathbb{E}(\hat{r}(\lambda, \xi)) = \int_{\Lambda \times \Xi} \hat{r}(\lambda, \xi) \hat{\pi}_{\text{init}}(\lambda, \xi) = \int_{\Lambda \times \Xi} \hat{\pi}_{\text{update}}(\lambda, \xi) = 1.$$

The above is true for any $\widehat{Q}$ and $\hat{\pi}_{\text{init}}$ that satisfy the predictability assumption. This integral is estimated in practice using a Monte-Carlo technique by re-using the initial *predicted* samples generated in the accept-reject algorithm to construct estimated samples of $\hat{r}(\lambda, \xi)$. In other words, we may estimate this integral using only knowledge of the initially predicted output quantities. When the sample average of $\hat{r}(\lambda, \xi)$ deviate significantly from 1 (i.e., more so than expected from finite-sampling error), this indicates a violation of the predictability assumption (see [5] for more details on this diagnostic for deterministic maps). Thus, sample averages of $\hat{r}(\lambda, \xi)$ prove to be useful numerical diagnostics in evaluating the capability of a map $\widehat{Q}$ to solve the inverse problem.

For the additive noise model, we propose an initial distribution for the noise $\xi \sim N(\mathbf{0}, 0.0825^2 \mathbf{I})$. For the location noise model, we propose an initial distribution for the noise terms $\xi_A, \xi_B \sim N(\mathbf{0}, 0.0825^2 \mathbf{I})$.

Figure 10 shows a comparison of the data consistent updated (i.e., accepted) samples from applying algorithm 1 with the additive noise model (left) or location noise model (right) using the same initial samples. Both updated sample sets include a wider range of physical parameters that are consistent with the data than in the previous example where the data (and map) had no additional stochastic uncertainties. This is expected since the new maps can predict likely QoI data using wider ranges of physical parameter values than the map Q . In other words, the introduction of stochastic sources of uncertainty in a map will in general lead to a reduction in the degree of updating of physical parameters. This is analogous to avoiding over-fitting a model.

Despite the qualitative similarity either map produces in the *range* of updated physical parameter samples consistent with the new data, substantial quantitative and qualitative differences exist that are explainable by the diagnostic. In the case where the stochastic noise is modeled correctly (location uncertainty), $E(\hat{r}(\lambda, \xi)) \approx 0.981$, whereas when modeled incorrectly (additive noise), $E(\hat{r}(\lambda, \xi)) \approx 0.846$. Thus, for this situation, the additive noise model appears to violate the predictability assumption and is overall insufficient for describing data with variations partially due to location uncertainties in the device setup. The impact of this on




Figure 10. Shows data consistent updates using the additive noise model (left) and the location noise model (right) for simulated data where $(\lambda_1, \lambda_2) \sim U([0.75, 1.45] \times [1.6, 1.85])$ and observations are polluted with noise in location measurements.

the updated sample is also qualitatively apparent where the physical parameters in the ‘target’ region responsible for the data is significantly under-sampled by the additive noise model map.

5.2. An elliptic PDE with uncertain diffusion

5.2.1. Motivation. In reservoir engineering and other subsurface applications, a partial differential equation (PDE) is often used to model the flow of oil, water, and other gases or contaminants in the subsurface (e.g., see [24]). The spatially heterogeneous structure of porosity, geological layering of rock types, and other subsurface characteristics impact many of the coefficients (i.e., the physical parameters) appearing in such PDEs. To model the spatial heterogeneity of these physical parameters, random fields with a specified spatial covariance are often used to allow for computationally tractable and accurate representations as truncated Karhunen–Loève (KL) expansions [14, 16]. In the context of this work, the physical parameters refer to the coefficients of a truncated KL expansion. We demonstrate that by including a stochastic parameter *embedded* in the diffusion coefficient defined by a low-order truncated KL expansion, it is possible to obtain inferences about physical parameters that are more aligned with using higher-order truncated KL expansions without incurring the additional cost of constructing such expansions. Moreover, as mentioned in section 4.3, the computations involving the higher-order truncated KL expansions use algorithm 1, which demonstrates the ability of the algorithm to be applied to high-dimensional input spaces.

5.2.2. Model and inverse problem. Consider the following incompressible flow model, which describes the pressure field p in terms of a permeability field $K(\lambda; x)$ and simple boundary conditions:

$$\begin{cases} -\nabla \cdot (K(\lambda; x) \nabla p) = 0, & x \in \Omega = (0, 1)^2 \\ p = 1, & x_1 = 0 \\ p = 0, & x_1 = 1 \\ K \nabla p \cdot \mathbf{n} = 0, & x_2 = 0 \text{ and } x_2 = 1 \end{cases} \quad (13)$$

We represent the permeability field $K(\lambda; x)$ by using a transformation of the KL expansion of a Gaussian process. In particular, let $K(\lambda; x) = \exp(Y + \bar{Y})$, where $\bar{Y}$ is the mean of the Gaussian




Figure 11. The first four eigenfunctions computed on a 10×10 mesh.

process. Then define Y using the KL expansion as

$$Y(\lambda; x) = \sum_{i=1}^{\infty} \lambda_i \cdot \sqrt{\eta_i} \cdot f_i(x) \quad (14)$$

where the λ_i are independent identically distributed $N(0, 1)$ random variables and (η_i, f_i) are the eigenvalue–eigenfunction pairs associated with the exponential covariance function $C(x, y) = \exp(|x_1 - y_1|/3 + |x_2 - y_2|/2)$. We use a stochastic finite element approach [14, 16] to numerically estimate the eigen-pairs by forming a generalized eigenvalue problem. Specifically, since the correlation lengths in $C(x, y)$ are large relative to the size of Ω , we perform all computations for the KL expansion and representations of the variability of the diffusion coefficient on a triangulation of a 10×10 mesh using piecewise-linear continuous functions, which is sufficient for maintaining the accuracy of the lower-order eigenfunctions considered in this example (see figure 11).

We refer to the 10×10 mesh as the *parameter mesh* to emphasize that it is chosen specifically to represent the variability in the diffusion coefficient. For simplicity, we abuse notation slightly and refer to the numerical approximations to η_i and f_i as the eigenvalues and eigenfunctions associated to $C(x, y)$ where

$$f_i(x) := \sum_{j=1}^{\text{DOF}} f_{i,j} \varphi_j(x),$$

and $\varphi_j(x)$ defines the j th ‘tent’ function (i.e., basis function) at the j th degree of freedom (DOF) for the finite element space defined on this parameter mesh.




Figure 12. Shows three sample permeability fields (rows) for three KL expansion approximations: a baseline generated by truncating after four terms (left), a reference generated by truncating after 100 terms (middle), and an expansion with four terms plus the additional stochastic parameters ξ_j (right).

Numerically computing the first 100 pairs of eigenvalues and eigenfunctions reveals that approximately 90% of the variance in the random fields is maintained by truncating at the fourth term⁴. In other words, the parameter mesh is used to represent how variations in the so-called physical parameters $\lambda_1, \dots, \lambda_4$ map to variations in $K(\lambda; x)$.

We then use a triangulation of a 20×20 mesh to solve (13) using standard piecewise-linear finite elements. This is referred to as the *computational mesh*. The computational mesh is chosen finer than the parameter mesh to ensure that the length-scales for which the diffusion parameters vary are sufficiently resolved and any deterministic numerical errors are subsequently neglected in the solution of the PDE. For example, see the left column of plots in figure 12 for samples of the permeability field that exhibit variability on the order of a magnitude over a fraction of the length scale of the physical domain.

The goal of the inverse problem is to use an observed distribution of the pressure field p at the point $(0.25, 0.5)$ to obtain an updated distribution of the physical parameters $\lambda_1, \dots, \lambda_4$ where the initial distribution is given by $N(0, I)$ where $0 \in \mathbb{R}^4$ and I is the 4×4 identity matrix. As this is a toy problem, the observed distribution in this case is simply specified to be $N(\mu, \sigma^2)$ with $\mu = 0.68$ and $\sigma^2 = 10E - 4$. In more realistic settings, an observed distribution may be specified as part of an engineering goal assuming some control on the physical parameters is possible, e.g., using direct intervention in the design of an engineered or physical system.

⁴This was also numerically verified by recomputing the first 100 pairs of eigenvalues and eigenfunctions on a 50×50 mesh where the maximum difference in the L^2 -norm between any eigenfunctions was 0.5.

The solution to the inverse problem is thus interpreted as describing the variabilities in physical parameters that lead to this particular observed distribution.

5.2.3. Representing the stochastic parameter. While the truncated KL expansion maintains a significant amount of the variation initially modeled by the diffusion parameters $K(\lambda; x)$, approximately 10% of the variability in this true physical parameter has been removed by the truncation. Thus, we do not expect that all of the variability in observed data is due to only the variations in the truncated KL expansion. However, if we solve the stochastic inverse problem using only the truncated KL expansion, then, by construction, the updated distribution on $\lambda_1, \dots, \lambda_4$ must necessarily propagate forward to the specified observed distribution. We therefore seek a stochastic parameter to account for variability in model outputs that are not entirely due to the physical parameters retained by the truncated KL expansion. Here, we seek a random function $\delta(\xi; x)$ defined on the parameter mesh such that

$$\delta(\xi; x) \approx \sum_{i=N+1}^{\infty} \lambda_i \cdot \sqrt{\eta_i} \cdot f_i(x).$$

In other words, we define a random function $\delta(\xi; x)$ parameterized by a stochastic parameter ξ which approximates the variability in the diffusion coefficient not accounted for by the truncation of the KL expansion.

A simple approach for constructing such a stochastic parameter is to ignore any correlation structure of the remaining terms of the KL expansion and construct a residual spatial variance model. Specifically, we define stochastic parameter $\xi = (\xi_j) \in \mathbb{R}^{\text{DOF}}$ where $\xi_j \sim N(0, \sigma_j^2)$ is defined at each of the DOF of the parameter mesh (i.e., $\text{DOF} = 100$) with magnitude of the variance at the j th degree of freedom, σ_j^2 , approximated by

$$\sigma_j^2 = \sqrt{1 - \sum_{i=1}^4 \eta_i f_i(x_j)}$$

which requires minimal additional computation since the eigenvalues, η_i , and eigenfunctions, f_i , have already been computed for $1 \leq i \leq 4$. This produces a computationally tractable log-permeability field $\hat{Y}(\lambda)$ as a function of both physical and stochastic parameters:

$$\hat{Y}(\lambda, \xi; x) = \sum_{i=1}^4 \lambda_i \cdot \sqrt{\eta_i} f_i(x) + \sum_{j=1}^{\text{DOF}} \xi_j \varphi_j(x). \quad (15)$$

In other words, $\delta(\xi; x)$ is given by the right-most term above.

5.2.4. Defining a reference. The approximations and constructions described above are dependent on an appropriate choice of truncation term and parameter mesh used for the KL expansion. In general, such choices are functions of the correlation lengths of the covariance function $C(x, y)$. For the purposes of this example, we simply show that the choice of approximations and incorporation of a stochastic term embedded in the model provides some clear qualitative benefit to the representation of the permeability field and the inferences drawn by solving the stochastic inverse problem. However, this requires defining some point of reference. Here, the point of reference is the solution obtained when the KL expansion is truncated at the 100th term.




Figure 13. Shows the data consistent update for each parameter λ_i of the KL expansion using either the deterministic, reference, or stochastic QoI.

In figure 12, we show the variability of three different permeability fields obtained when truncating the KL expansion at the fourth term (left column), 100th term (middle column, defining the reference fields), and when the stochastic parameter is sampled and added to the lower-order truncation (right column). Plots are generated on the computational mesh, i.e., permeability fields are generated on the parameter mesh and then projected onto the computational mesh since this is the form of the fields that appears in the numerical solution of the PDE. Each row has a fixed vector for the $\lambda_1, \dots, \lambda_4$ values. We observe finer-scale features in the permeability fields in the middle and right columns than in the left column. Qualitatively, the addition of the stochastic term appears to capture some of the finer-scale features present in the reference fields.

5.2.5. The updated PDFs. In what follows, we refer to the QoI map obtained by solving (13) using the four term KL expansion as the *deterministic* QoI, using the 100 term KL expansion as the *reference* QoI, and using the four term KL expansion with the additional stochastic parameter impacting the diffusion coefficient the *stochastic* QoI. We summarize all of the updated PDFs obtained on the physical parameters, $\lambda_1, \dots, \lambda_4$, using the various QoI maps in figure 13.




Figure 14. Shows initial predictions (left) and updated predictions (right) of the QoI compared to the observed density for each of the KL approximations (deterministic, reference, and stochastic).

In figure 14, we show the push-forward of these updated densities through each of the maps to verify that each QoI map separately solved the stochastic inverse problem correctly (i.e., the predictability assumption holds for all the maps).

For λ_1 , λ_2 , and λ_4 , the updated marginal PDFs are qualitatively similar across all maps. Moreover, they are not significantly different from the initial PDFs for each of these parameters suggesting that the various QoI maps are not sensitive to these particular parameters. However, for the third physical parameter, λ_3 , the updated marginal PDFs are all significantly different from the initial PDF for all QoI maps. This suggests two things. First, all the QoI maps are sensitive to this parameter. Second, a significant amount of the variation in the observed PDF is explained by the distribution of this parameter. This is explained by a quick review of the model, the type of measurement defined by the QoI, and the eigenfunctions associated with each of the physical parameters. First, the boundary conditions in (13) imply that the pressure drops from 1 to 0 as we move from the left-boundary to the right-boundary. Second, the measurement location of p is biased towards the left-half of the domain ($x_1 = 0.25$) and is in the center of the vertical portion of the domain ($x_2 = 0.5$). Combined, we would expect this measurement to be more sensitive to perturbations of permeability field values that primarily impact the variation of permeability as we move in the x_1 -direction. It is evident from the eigenfunctions in figure 11 that such perturbations occur by varying λ_3 .

Further examination of the updated marginal PDFs for λ_3 also reveal that the deterministic QoI leads to an ‘over-updating’ of the initial PDF compared to the update obtained from the reference or stochastic QoI. This is due to the fact that each QoI map creates an updated PDF that must be consistent with the observations in the sense that their associated push-forwards will match the observed PDF (as seen in figure 14). For the deterministic map, this requires that *all* variations in the observations must be explained by variations in the physical parameters $\lambda_1, \dots, \lambda_4$. For the reference and stochastic maps, variations in other parameters (additional physical parameters for the reference map and the stochastic parameters for the stochastic map) can explain some of the variability in the observations. While the focus of this work is not on *constructing* stochastic maps but rather extending the data-consistent framework to apply to stochastic maps, this still highlights an important point about the utility of considering stochastic maps within this framework. Specifically, by accounting for variations in the

output data associated with a (potentially deterministic) model that are not due to variations in a (potentially limited) set of physical parameters, we can obtain inferences on the physical parameters that are more aligned with higher-fidelity references. Moreover, it may be possible to achieve such ‘enhanced’ inferences at a fraction of the computational cost required in constructing such references. However, we leave the general problem of identifying potential sources of uncertainty in models and accounting for those with proper stochastic parameter choices to future works as such modeling is heavily application/domain specific.

6. Conclusions and future work

Extending the data-consistent framework to stochastic maps greatly widens the class of problems that may be solved with this approach. In practice, it is rarely the case that a QoI map from parameters of interest to the data space is known exactly. Nonetheless, by incorporating and modeling the uncertainty in the QoI map, it is still possible to obtain an informative update of the uncertainties in parameters of interest. The theoretical analysis described in sections 3 and 4 ensure that such a data-consistent update exists in the form of a joint distribution whose marginal can be sampled from directly. In addition, the update is stable as long as the assumption of predictability is satisfied.

In section 5, we demonstrate some of the flexibility provided by this methodology. With the data consistent approach, the researcher has a great deal of choice in how to model the additional variation or noise that is observed in the data. Such variation may be modeled as additive noise in the resulting observations, but may also enter into the model in a variety of other more complex embeddings, such as specific errors in an experimental setup (e.g., example 5.1) or general approximations of random processes (e.g., example 5.2). This is reminiscent of the kind of flexibility provided by mixed effect models, where statisticians use information about the structure of experiments to improve the quality of their inferences except that the data-consistent framework is not restricted to linear models. Future research may reveal applications for the data-consistent approach in modeling aleatoric-type uncertainty in similar contexts.

One notable feature of this data-consistent framework is that the additional variation need not be modeled explicitly as long as parameters of interest can be paired with corresponding predicted data values. This is particularly useful in several scenarios including, but not limited to, experimental settings where only certain covariates can be controlled/measured between trials and also large-scale computational models where complex stochastic processes are employed unbeknownst to the user to compute model predictions.

We show in section 5.2 that the algorithm is easily applied to a parameter space of 100-dimensions for computing the reference results. As mentioned in section 4.3, the algorithm scales well with increasing parameter dimension for a fixed data space and variance of the predicted density since the computations required to determine the ratios in algorithm 1 occur in the data space. However, as the dimension of the data space increases, we may require significantly more samples to compute accurate estimates of the predicted density with a GKDE. When obtaining a QoI sample is expensive (either experimentally or computationally), it may only be possible to generate a relatively low number of predicted samples. There are several on-going and future research efforts we are pursuing to address this issue. As algorithms are developed and analyzed, we will incorporate these into the open-source Python library *vert* BETvert [15]. A future *vert* BETvert project is to develop and encode criteria for automatically choosing an algorithm based on the dimensions of spaces and the number of QoI samples available either experimentally or due to a computational budget.

Acknowledgments

T Butler and TY Yen's work is supported by the National Science Foundation (DMS-1818941). T Wildey's work was supported by the Office of Science Early Career Research Program.

ORCID iDs

Troy Butler  <https://orcid.org/0000-0001-7177-3105>

References

- [1] Aqvist J, Medina Llanos C and Samuelsson J E 1994 A new method for predicting binding affinity in computer-aided drug design *Protein Eng.* **7** 385–91
- [2] Berger J O *et al* 1994 An overview of robust Bayesian analysis *Test* **3** 5–124
- [3] Bunya S *et al* 2010 A high resolution coupled riverine flow, tide, wind, wind wave and storm surge model for southern Louisiana and Mississippi: part I—model development and validation *Mon. Weather Rev.* **138** 345–77
- [4] Burger M and Lucka F 2014 Maximum *a posteriori* estimates in linear inverse problems with log-concave priors are proper Bayes estimators *Inverse Problems* **30** 114004
- [5] Butler T, Jakeman J and Wildey T 2018 Combining push-forward measures and Bayes' rule to construct consistent solutions to stochastic inverse problems *SIAM J. Sci. Comput.* **40** A984–1011
- [6] Butler T, Jakeman J and Wildey T 2018 Convergence of probability densities using approximate models for forward and inverse problems in uncertainty quantification *SIAM J. Sci. Comput.* **40** A3523–48
- [7] Calder M *et al* 2018 Computational modelling for decision-making: where, why, what, who and how *R. Soc. Open Sci.* **5** 172096
- [8] Calvetti D, Kaipio J and Somersalo E 2014 Inverse problems in the Bayesian framework *Inverse Problems* **30** 110301
- [9] Cotter S L, Dashti M and Stuart A M 2010 Approximation of Bayesian inverse problems *SIAM J. Numer. Anal.* **48** 322–45
- [10] Dellacherie C and Meyer P A 1978 *Probabilities and Potential* (Amsterdam: North-Holland)
- [11] Dzurisin D 2003 A comprehensive approach to monitoring volcano deformation as a window on the eruption cycle *Rev. Geophys.* **41** 1004
- [12] Finger S and Dixon J R 1989 A review of research in mechanical engineering design. Part i: Descriptive, prescriptive, and computer-based models of design processes *Res. Eng. Des.* **1** 51–67
- [13] Gelman A, Carlin J B, Stern H S, Dunson D B, Vehtari A and Rubin D B 2013 *Bayesian Data Analysis* 3rd edn (London: Chapman and Hall/CRC)
- [14] Ghanem R and Spanos P 1991 *Stochastic Finite Elements: A Spectral Approach* (New York, NY: Springer)
- [15] Graham L, Mattis S, Scott W, Butler T, Pilosov M and McDougall D 2016 *BET: Butler, Estep, Tavenner Method v2.0.0*
- [16] Huang S P, Quek S T and Phoon K K 2001 Convergence study of the truncated Karhunen–Loeve expansion for simulation of stochastic processes *Int. J. Numer. Methods Eng.* **52** 1029–43
- [17] Jaynes E T 1998 *Probability Theory: The Logic of Science* (Cambridge: Cambridge University Press)
- [18] Kennedy M C and Anthony O'H 2001 Bayesian calibration of computer models *J. R. Stat. Soc. B* **63** 425–64
- [19] Marzouk Y M, Najm H N and Rahn L A 2007 Stochastic spectral methods for efficient Bayesian solution of inverse problems *J. Comput. Phys.* **224** 560–86
- [20] Moosavi L, Mahyuddin N, Ghafar N, Zandi M and Bidi M 2018 Numerical prediction of thermal and airflow conditions of a naturally ventilated atrium and validation of cfd models *J. Renew. Sustain. Energy* **10** 065101
- [21] Morrison R, Oliver T and Moser R 2018 Representing model inadequacy: a stochastic operator approach *SIAM/ASA J. Uncertain. Quantif.* **6** 457–96
- [22] Panchal J H, Kalidindi S R and McDowell D L 2013 Key computational modeling issues in integrated computational materials engineering *Comput.-Aided Des.* **45** 4–25

- [23] Salvatier J, Wiecki T V and Fonnesbeck C 2016 Probabilistic programming in Python using pyMC3 *PeerJ Comput. Sci.* **2** e55
- [24] Sarma P, Durlowsky L J, Aziz K and Chen W H 2006 Efficient real-time reservoir management using adjoint-based optimal control and model updating *Comput. Geosci.* **10** 3–36
- [25] Smith D I, Greenaway M A, Moses C and Spate A P 1995 Limestone weathering in eastern australia. Part 1: erosion rates *Earth Surf. Process. Landf.* **20** 451–63
- [26] Sparks R S J, Biggs J and Neuberg J W 2012 Monitoring volcanoes *Science* **335** 1310–1
- [27] Stephenson W J and Kirk R M 1998 Rates and patterns of erosion on inter-tidal shore platforms, Kaikoura Peninsula, South Island, New Zealand *Earth Surf. Process. Landf.* **23** 1071–85
- [28] Stephenson W J and Finlayson B L 2009 Measuring erosion with the micro-erosion meter—contributions to understanding landform evolution *Earth-Sci. Rev.* **95** 53–62
- [29] Stuart A M 2010 Inverse problems: a Bayesian perspective *Acta Numer.* **19** 451–559
- [30] Trudgill S T, Viles H A, Inkpen R, Moses C, Gosling W, Yates T, Collier P, Smith D I and Cooke R U 2001 Twenty-year weathering remeasurements at St Paul’s Cathedral, London *Earth Surf. Process. Landf.* **26** 1129–42
- [31] van der Vaart A W 1998 *Asymptotic Statistics* (Cambridge: Cambridge University Press)
- [32] Wikle C K, Mark Berliner L and Noel C 1998 Hierarchical Bayesian space–time models *Environ. Ecol. Stat.* **5** 117–54