RESEARCH PROGRESS REPORT

MICHAEL PILOSOV

Abstract. Over the past several years, a measure-theoretic framework for formulating and solving stochastic inverse problems has been developed and analyzed. In this framework, we begin with a model, model inputs (called parameters), and a quantity of interest map that defines the types of model outputs used to formulate the stochastic inverse problem. The solution to the stochastic inverse problem is a (non-parametric) probability measure on the space of parameters. A non-intrusive sample based algorithm has also been developed and analyzed to approximate this probability measure. More recent work has shown that a particular geometric property (the skewness) of the quantity of interest map is correlated to the accuracy in the finite sample approximation of the probability measure. In this work, we provide a more in-depth numerical investigation into how skewness impacts the accuracy and convergence rates of such approximate probability measures. We provide an algorithmic procedure for computing a metric on a space of probability measures, and we demonstrate that while convergence rates are unaffected by the choice of map, the absolute error is directly proportional to the skewness of the map.

Contents

1. Introduction

The description and reduction of uncertainties in physics-based computational models is of great interest to the community of computational mathematicians. The process of parameter tuning after the model's initial development starts with the modeler specifying bounds using domain-specific knowledge. However, exact values of parameters within these bounds are often uncertain. These uncertainties lead to uncertainties in the observable model outputs, i.e., the quantities of interest (QoIs). This is problematic for the predictive capabilities of the model. The reduction and quantification of these uncertainties in the parameters can be informed by a process of exploratory sampling which is used to solve a stochastic inverse problem (SIP).

The choice of observable model outputs used in the formulation and solution of the stochastic inverse problem lead to different solutions defined as probabilities measures or densities. Not only may the supports of these densities differ, but our ability to approximate these densities with finite sampling will be affected as well. Here, we study how a quantifiable geometric property of the QoI maps relates to the number of samples we need to sufficiently approximate the probability density.

This work builds on results presented in [7] and in [5], where an end-to-end algorithmic framework was developed for the analysis and reduction of uncertainties in computational models. In [7], a geometric quantity referred to as the skewness of a QoI map, was introduced and linked to the solution process, though no efficient method to compute the quantity was provided. This was expanded upon and remedied in [9], where efficient computational methods (relying on sampled Jacobians of the QoI map) were defined and developed to quantify both skewness and another geometric property of the QoI map. There, we defined a multi-objective optimization problem for selecting a QoI map based on these properties. In doing so, the framework for *simulation-based optimal experimental* was defined. This framework was shown to be effective in its infancy on a hurricane storm surge example in [11].

In this work, we specifically study the impact skewness has on the errors and convergence of solutions to the SIP using the BET software [12]. BET is an open-source Python package that uses random (or regular) sampling to compute solutions to the SIP. We expand on the analysis of [8], which provided a description of errors in approximating the probabilities of events using solutions. Since errors caused by skewness can cancel out globally, it is not clear what the impact can be on the solution provided by BET. To discuss errors and convergence of solutions to a SIP, we describe how we metrize a space of probability measures. We also provide an algorithm for computing the metric used in this work.

We begin the paper with a brief review of the measure-theoretic background necessary for the development of the problem and solution method given in [5]. In Section 3, an algorithm is provided for the definition and solution of the SIP from Section 2 and introduces notation to describe the errors involved in the process. A precise definition of the geometric property we study in this work is also given in this section. Next, we describe several commonly used metrics on spaces of probability measures in Section 4. Then, in Section 5,

we handle the theoretical foundations of the numerical implementation of equipping our space with such a metric. Finally we provide the numerical results that are the novel contributions of this work in Section 6 before making some concluding remarks.

2. Measure-Theoretic Inversion: A Review

2.1. **Notation & Terminology.** We begin by assuming a (deterministic) model, denoted by

$$\mathcal{M}(u, \lambda) = 0,$$

is specified that relates state variables u to model inputs (parameters) denoted by the vector λ. The components of λ may include parameters in either the model operator (e.g. a diffusion coefficient) or input data (e.g. the frequency of a sinusoidal source, initial, or boundary data). Let Λ denote the set of all possible input parameters. The solution operator of the model $\mathcal{M}$ then defines a map taking $\lambda \in \Lambda$ to a solution denoted $u(\lambda)$ which is assumed to be unique.

In an experimental setting, we often cannot fully observe $u(\lambda)$, but rather we define some finite set of observable scalar quantities that are mathematically modeled by functionals of the solution. For example, in experiments of heat diffusion in a media, we typically only record the temperature at some small number of pre-specified points in space-time where measurement devices can be positioned. We refer to the defined set of observables and their associated functionals as *quantities of interest* (QoI), and the set of functionals we consider defines a (vector-valued) QoI map, which we denote by Q. Monte-Carlo integration appears throughout the uncertainty quantification literature and is an integral part of implementing analytical results computationally.

Here, we motivate the reduction of a quantity called *skewness* in pursuit of optimizing the geometry of set-valued solutions to stochastic inverse problems with respect to their ability to be well-approximated by Monte-Carlo integration. However, the results hold for any attempt to approximate densities defined on sets induced by random samples, and thus may be of interest to the larger research community. We metrize the space of probability measures to demonstrate that the number of samples required to approximate densities using uniform i.i.d. sampling is proportional to the skewness of the map used for inversion, though the convergence rate of the algorithm used to solve the stochastic inverse problem is unaffected.

Since the solution to the model, $u(\lambda)$, depends on λ, so do the QoI, and we adopt the notation that

$$Q(\lambda) := Q(u(\lambda))$$

to make this dependence on model parameters explicit. Furthermore, this convention also expresses the reality of the limitations of an experimental setting, where we may be able to control λ in order to observe $Q(\lambda)$ without observing $u(\lambda)$ in its entirety.

The outputs of the QoI map Q are what we refer to as the *data* and the range of the QoI defines the *data space* $\mathcal{D}_Q$, i.e.

$$\mathcal{D}_Q = Q(\Lambda).$$

Here, we use $\mathcal{D}_Q$ to emphasize that the data space depends on the choice of QoI map Q.

We let $\mathcal{Q}$ denote the set of possible QoI maps for which it is possible to collect experimental data. For example, suppose we may record only a single temperature measurement at any of ten locations in space-time. Then $\mathcal{Q}$ is defined by ten possible QoI maps. If we can record any two such measurements, then $\mathcal{Q}$ is defined by $\binom{10}{2} = 45$ possible maps. Note that $\mathcal{Q}$ could easily be uncountable, for example if we were not limited to the spatial locations (or time) at which we could record temperature measurements. However, for simplicity, we will only discuss problems where $\mathcal{Q}$ is finite.

2.2. **The Stochastic Inverse Problem.** To properly summarize the stochastic inverse problem (SIP) and desired solution, we define several measure/probability spaces and refer to the schematic given in Figure 1—borrowed from [8]—in order to illustrate the steps and spaces required in the formulation and solution of the SIPs we consider herein. For a more extensive review, we refer the reader to [1], [4], [5], and [8].

$$\overbrace{\underbrace{(\Lambda, \mathcal{B}_\Lambda, \mu_\Lambda) \xmapsto{Q} (\mathcal{D}, \mathcal{B}_\mathcal{D}, \mu_\mathcal{D}) \xmapsto{P_\mathcal{D}} (\mathcal{D}, \mathcal{B}_\mathcal{D}, P_\mathcal{D})}_{\text{(S2): Solution to SIP Satisfying Eq. (2.2)}} \xmapsto{Q^{-1}} (\Lambda, \mathcal{C}_\Lambda, P_\Lambda) \xmapsto{\{P_\ell\}_{\ell \in \mathcal{L}}} (\Lambda, \mathcal{B}_\Lambda, P_\Lambda)}^{\text{(S1): Stochastic Inverse Problem (SIP)}} \tag{2.1}$$

$$\underbrace{\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad}_{\text{(S3): Unique Solution to SIP by Eq. (2.5) and Ansatz}}$$

FIGURE 1. The first step (S1) defines (i) the formulation of the SIP by specification of the model, (ii) the measure spaces of parameters and (iii) observable outputs, and (iv) the probability measure on the latter. The second step (S2) defines a unique solution to the SIP on the space Λ equipped with the contour σ-algebra $\mathcal{C}_\Lambda$ using the definition of the push-forward measure. In (S3), the Disintegration Theorem and and Ansatz are applied to define a unique solution on the space of interest $(\Lambda, \mathcal{B}_\Lambda)$ equipped with a probability measure P_Λ.

The initial measure/probability spaces involved in the formulation of the SIP are summarized in step (S1) of Fig. 1. We assume $\Lambda \subset \mathbb{R}^n$ is equipped a (volume) measure, μ_Λ, on the Borel σ-algebra $\mathcal{B}_\Lambda$, defining the measure space $(\Lambda, \mathcal{B}_\Lambda, \mu_\Lambda)$. Assume that Q is at least piecewise differentiable, which implies the measurability of the QoI map, and that the space $\mathcal{D}_Q$ induced by Q is equipped with the Borel σ-algebra $\mathcal{B}_{\mathcal{D}_Q}$. The "push-forward" measure $\mu_{\mathcal{D}_Q}$ on $(\mathcal{D}_Q, \mathcal{B}_{\mathcal{D}_Q})$ is defined as

$$\mu_{\mathcal{D}_Q}(A) = \int_A d\mu_{\mathcal{D}_Q} = \int_{Q^{-1}(A)} d\mu_\Lambda = \mu_\Lambda\left(Q^{-1}(A)\right) \quad \forall\, A \in \mathcal{B}_{\mathcal{D}_Q}. \tag{2.2}$$

This defines the measure space $(\mathcal{D}_Q, \mathcal{B}_{\mathcal{D}_Q}, \mu_{\mathcal{D}_Q})$. When referring to properties of the data space that are not unique to the choice of map used to induce $\mathcal{D}_Q$, we will drop the subscript notation and assume the dependence is understood, as expressed in Fig. 1.

Specifying a probability measure $P_{\mathcal{D}_Q}$ on $(\mathcal{D}_Q, \mathcal{B}_{\mathcal{D}_Q})$ to model the uncertainty in data—the final step in (S1)—leads to the following SIP: determine a probability measure P_Λ on $(\Lambda, \mathcal{B}_\Lambda)$ such that the push-forward measure of P_Λ matches $P_{\mathcal{D}_Q}$. In other words, determine a P_Λ satisfying

$$P_\Lambda\left(Q^{-1}(E)\right) = P_{\mathcal{D}_Q}(E) \; \forall \, E \in \mathcal{B}_{\mathcal{D}_Q}. \tag{2.3}$$

We call any such solution P_Λ to Eq. (2.3) a (measure-theoretic) solution to the SIP. This equation implies that any solution is uniquely determined on the induced contour σ-algebra

$$\mathcal{C}_\Lambda = \left\{Q^{-1}(E) : E \in \mathcal{B}_{\mathcal{D}_Q}\right\} \subset \mathcal{B}_\Lambda, \tag{2.4}$$

This is summarized as step (S2) of Fig. 1.

However, for sets $A \in \mathcal{B}_\Lambda \setminus \mathcal{C}_\Lambda$, more information is required than is provided in Eq. (2.3) in order to determine $P_\Lambda(A)$. By the Implicit Function Theorem, if $Q \in C^1(\Lambda)$ and we let $q \in \mathcal{D}$ be a fixed datum, $Q^{-1}(q)$ exists as a $(n-d)$-dimensional manifold (possibly piecewise-defined) that we refer to as a *generalized contour* [5]. These generalized contours can be indexed by a manifold (possibly piecewise-defined) of dimension d called a *transverse parameterization* that intersects each contour once and only once. Transverse parameterizations exist but are in general not unique [5]. We let $\mathcal{L}$ denote any particular transverse parameterization. Each $\ell \in \mathcal{L}$ corresponds to a unique generalized contour $\mathcal{C}_\ell \in \Lambda$ and each point $\lambda \in \Lambda$ belongs to a unique $\mathcal{C}_\ell \in \Lambda$. Thus, a transverse parameterization defines a bijection between the manifold $\mathcal{L}$ and the partitioning of Λ into generalized contours.

The induced σ-algebra $\mathcal{C}_\Lambda$ and this bijection can then be used to define the measurable space $(\mathcal{L}, \mathcal{B}_\mathcal{L})$. We denote the projection map $\pi_\mathcal{L} : \Lambda \to \mathcal{L}$, and $\{\mathcal{C}_\ell\}_{\ell \in \mathcal{L}}$ denotes the family of generalized contours indexed by $\mathcal{L}$. It is then possible to define the associated family of measurable spaces $\{(\mathcal{C}_\ell, \mathcal{B}_{\mathcal{C}_\ell})\}_{\ell \in \mathcal{L}}$. A Disintegration Theorem (cite) is then leveraged to define a unique decomposition for any P_Λ defined on $(\Lambda, \mathcal{B}_\Lambda)$ as a (marginal) probability measure $P_\mathcal{L}$ on $(\mathcal{L}, \mathcal{B}_\mathcal{L})$ and a family of (conditional) probability measures $\{P_\ell\}_{\ell \in \mathcal{L}}$ on $\{(\mathcal{C}_\ell, \mathcal{B}_{\mathcal{C}_\ell})\}_{\ell \in \mathcal{L}}$ such that

$$P_\Lambda(A) = \int_{\pi_\mathcal{L}(A)} \left(\int_{\pi_\mathcal{L}^{-1}(\ell) \cap A} dP_\ell(\lambda) \right) dP_\mathcal{L}(\ell), \; \forall \, A \in \mathcal{B}_\Lambda \tag{2.5}$$

The uniqueness of a probability measure P_Λ on $(\Lambda, \mathcal{C}_\Lambda)$ satisfying Eq. (2.3) implies that the marginal $P_\mathcal{L}$ is unique for any particular specification of $P_\mathcal{D}$ on $(\mathcal{D}, \mathcal{B}_\mathcal{D})$. The disintegration of Eq. (2.5) implies that a specification of a family of conditional probability measures $\{P_\ell\}_{\ell \in \mathcal{L}}$ gives us a unique solution to the SIP on $(\mathcal{C}_\ell, \mathcal{B}_{\mathcal{C}_\ell})$. However, the conditional measures cannot be determined by observations of $q \in \mathcal{D}$. We follow the work of [5] and adopt the *standard ansatz* determined by the disintegration of the volume measure μ_Λ

to compute probabilities of events contained within contour events. The standard ansatz is given by

$$P_\ell = \mu_{\mathcal{C}_\ell}/\mu_{\mathcal{C}_\ell}(\mathcal{C}_\ell),\ \forall\ \ell \in \mathcal{L}, \tag{2.6}$$

where $\mu_{\mathcal{C}_\ell}$ is the disintegrated volume measure on generalized contour $\mathcal{C}_\ell$. Thus, we have defined a unique solution to the SIP on $(\Lambda, \mathcal{B}_\Lambda)$, completing step (S3) in Fig. 1.

3. Numerical Approximations and Analysis

In practice, we must rely on a finite numerical approximation of the often uncountable number of events in the σ-algebras $\mathcal{B}_{\mathcal{D}_Q}$, $\mathcal{C}_\Lambda$, and $\mathcal{B}_\Lambda$. Since $\mathcal{C}_\Lambda \subset \mathcal{B}_\Lambda$, events in both of these σ-algebras can be approximated simultaneously. Thus, there are two primary sources of approximation error: (1) partitioning the parameter space Λ to approximate events in $\mathcal{C}_\Lambda$, and $\mathcal{B}_\Lambda$, and (2) partitioning the data space $\mathcal{D}_Q$ to approximate events in $\mathcal{B}_{\mathcal{D}_Q}$.

3.1. **Algorithm.** We present a non-intrusive sample-based algorithm initially introduced in [5] and further analyzed in [6] that is structured in four stages (written as four independent for-loops) that are linked to the stages in Fig. 1. We direct the interested reader to [6] for more detailed information and analysis of this algorithm, e.g., on the requirement of a sampler being "$\mathcal{B}_\Lambda$-consistent" to ensure convergence.

The first two stages correspond to formulating the discretized version of the SIP given in step (S1) in Fig. 1. We first discretize the probability space $(\mathcal{D}, \mathcal{B}_\mathcal{D}, P_\mathcal{D})$. Then, we simultaneously discretize the measure space $(\Lambda, \mathcal{B}_\Lambda, \mu_\Lambda)$ and construct a simple-function approximation to the map Q. These stages introduce the primary sources of error, and the third and fourth stages may be thought of as solving the discretized SIP exactly.

The third stage then identifies the collection of Voronoi cells in Λ that approximate the contour events defined by $Q^{-1}(D_i)$ for $i = 1, \ldots, M$. This allows us to formulate the consistent solution to the SIP on $(\Lambda, \mathcal{C}_\Lambda, P_\Lambda)$ as illustrated in step (S2) of Fig. 1. Finally, the fourth stage constructs a discretized approximation to step (S3) in Fig. 1 and uses a discrete version of the ansatz to approximate the probability of $\mathcal{V}_j$ for $j = 1, \ldots, N$. This results in an approximate probability measure $P_{\Lambda,M,N}$ which produces the same probability estimates for events A and $A \setminus \left\{\lambda^{(j)}\right\}_{j=1}^N$, which are identical almost everywhere with respect to μ_Λ.

Note that Algorithm 1 is independent of the methods by which the samples $\left\{\lambda^{(j)}\right\}_{j=1}^N$ were generated or sets in $\{D_i\}_{i=1}^M$ are chosen. A thorough discussion of the choices involved in making such decisions is beyond the scope of this work, though we touch briefly on the discretization of $\mathcal{D}$ below.

Algorithm 1: Numerical Approximation of the Inverse Density

Choose a discretization partition $\{D_i\}_{i=1}^M$ of $\mathcal{D}$.
for $i = 1, \dots, M$ **do**
  Compute $p_{\mathcal{D},i} = P_{\mathcal{D}}(D_i)$.
Choose samples $\left\{\lambda^{(j)}\right\}_{j=1}^N \subset \Lambda$, which implicitly defines a Voronoi partition $\{\mathcal{V}_j\}_{j=1}^N$ of Λ.
for $j = 1, \dots, N$ **do**
  Compute quantity of interest vectors $Q_j = Q(\lambda^{(j)})$.
  Bin Q_j in the partition $\{D_i\}_{i=1}^M$ and let $\mathcal{O}_j = \{i : Q_j \in D_i\}$.
  Compute approximations $V_j \approx \mu_\Lambda(\mathcal{V}_j)$.
for $i = 1, \dots, M$ **do**
  Compute $\mathcal{C}_i = \{j : Q_j \in D_i\}$.
for $j = 1, \dots, N$ **do**
  Compute $p_{\Lambda,j} = \left(V_j / \sum_{k \in \mathcal{C}_{\mathcal{O}_j}} V_k\right) p_{\mathcal{D},\mathcal{O}_j}$.
For any $A \in \mathcal{B}_\Lambda$, compute

$$P_{\Lambda,M,N}(A) = \sum_{j=1}^N p_{\Lambda,j} \chi_{\mathcal{V}_j}(A) \tag{3.1}$$

3.2. **Descriptions of Error.** We begin by assuming that $P_{\mathcal{D}}$ is absolutely continuous with respect to $\mu_{\mathcal{D}}$, which allows us to describe $P_{\mathcal{D}}$ with a density $\rho_{\mathcal{D}}$. Then, for any partition $\{D_i\}_{i=1}^M$ of $\mathcal{D}$,

$$P_{\mathcal{D}}(D_i) = \int_{D_i} \rho_{\mathcal{D}} \, \mu_{\mathcal{D}}, \qquad \text{for } i = 1, \dots, M.$$

We often use Monte Carlo approximations to compute the approximations $p_{\mathcal{D},i} = P_{\mathcal{D}}(D_i)$ in the first for-loop in Algorithm 1. These samples are generated on $\mathcal{D}$ and do not require numerical solutions to the model. We therefore assume that for any discretization of $\mathcal{D}$, these approximations can be made sufficiently accurate and neglect the error in this computation. We denote the exact solution to the SIP associated with this partitioning of $\mathcal{D}$ by $P_{\Lambda,M}$. Approximate solutions to the SIP given in the final for-loop of Algorithm 1 are denoted by $P_{\Lambda,M,N,h}$. Here, the h is in reference to a mesh or other numerical parameter that determines the accuracy of the numerical solution $u_h(\lambda^{(j)}) \approx u(\lambda^{(j)})$, and subsequently the accuracy in the computations of $Q_j = Q(\lambda^{(j)})$ in Algorithm 1.

We assume that h is tunable so that

$$\lim_{h \downarrow 0} P_{\Lambda,M,N,h} = P_{\Lambda,M,N}.$$

In [8], the focus was on proving the convergence of $P_{\Lambda,M,N,h}(A) \to P_\Lambda(A)$ for some $A \in \mathcal{B}_\Lambda$ and on estimating the error in $P_{\Lambda,M,N,h}(A)$. There, as well as in [7], adjoint-based a posteriori estimates in the computed QoI are combined with a statistical analysis to both estimate and bound the error in $P_{\Lambda,M,N,h}(A)$. In [8], adjoints were used to compute both error and derivative estimates of $Q(\lambda^{(j)})$ to improve the accuracy in $P_{\Lambda,M,N,h}(A)$.

However, no work has to date fully explored the *convergence rates* of Algorithm 1. Furthermore, no work has yet to establish that these rates are independent of the choice of QoI map despite other studies establishing that the absolute error is very much affected by the geometric properties of the QoI maps [3].

In order to study convergence, we need to define a notion of distance on the space of probability measures on Λ, which we denote by $\mathcal{P}_\Lambda$. There are many choices available to us and we discuss several useful metrics on $\mathcal{P}_\Lambda$ in Section 4. However, for now, let d represent a metric on $\mathcal{P}_\Lambda$.

Then, we have by repeated application of the triangle inequality that

$$d(P_{\Lambda,M,N,h}, P_\Lambda) \leq \underbrace{d(P_{\Lambda,M,N,h}, P_{\Lambda,M,N})}_{\text{(E1)}} + \underbrace{d(P_{\Lambda,M,N}, P_{\Lambda,M})}_{\text{(E2)}} + \underbrace{d(P_{\Lambda,M}, P_\Lambda)}_{\text{(E3)}}. \tag{3.2}$$

The term (E1) describes the effect of the error in the numerically evaluated Q_j on the solution to the SIP. The term (E2) describes the effect of finite sampling error in Λ on the solution to the SIP and (E3) describes the effect of discretization error of $P_\mathcal{D}$ on the solution to the SIP.

3.3. **Skewness and Accuracy.** In [7], the concept of skewness in a QoI map Q was introduced, quantified, and related to the accuracy in solving the SIP with a finite number of samples. Essentially, skewness is a geometric property that describes how the right angles in generalized rectangles belonging to $\mathcal{B}_\mathcal{D}$ are transformed by Q^{-1}. An a priori analysis demonstrated that the number of samples from a *regular uniform grid* in Λ required to approximate the μ_Λ-measure of $Q^{-1}(E)$ to a desired level of accuracy was proportional to the skewness of Q raised to the $(d-1)$ power where d is the dimension of $\mathcal{D}$. This is a version of the so-called curse-of-dimension. Since the numerical solution of the SIP relies fundamentally on the approximation of such $Q^{-1}(E)$, it was assumed that skewness impacted the probability measure computed using Algorithm 1 in a similar way.

Skewness was explored further in [9] in the context of optimal experimental design. There, an additional geometric property of Q related to the *precision* in the solution of the associated SIP was introduced and quantified. Under the same assumption that skewness impacted accuracy of the numerical solution to the SIP, a multi-criteria optimization problem was formulated and solved to determine a $Q \in \mathcal{Q}$ that balanced accuracy and precision in the SIP solution.

Since no previous study has considered the numerical convergence rates of Algorithm 1, the impact of skewness on such rates is unclear. Moreover, it is not clear that skewness causes errors to pollute the entire numerical solution of the SIP since local approximation errors in the μ_Λ-measure of contour events have a tendency to cancel out over Λ. This is the primary focus of the numerical results in Section 6. Here, for completeness, we define skewness below and refer the interested reader to [7, 9] for more details.

Definition 1. For any $Q \in \mathcal{Q}$, $\lambda \in \Lambda$, and a specified row vector $\mathbf{j_k}$ of the Jacobian $J_{\lambda,Q}$, we define

$$S_Q(J_{\lambda,Q}, \mathbf{j_k}) := \frac{|\mathbf{j_k}|}{|\mathbf{j_k^\perp}|}. \tag{3.3}$$

We define the **local skewness** of a map Q at a point λ as

$$S_Q(\lambda) = \max_{1 \le k \le d} S_Q(J_{\lambda,Q}, \mathbf{j_k}). \tag{3.4}$$

Definition 2. The **average** *(or **expected**)* **skewness** is defined as

$$\overline{S_Q} = \frac{1}{\mu_\Lambda(\Lambda)} \int_\Lambda S_Q(\lambda)\, d\mu_\Lambda \tag{3.5}$$

In [9], it was shown that $S_Q(\lambda)$ can be efficiently computed using a singular value decomposition of the Jacobian $J_{\lambda,Q}$. In general, we approximate $\overline{S_Q}$ with Monte-Carlo approximations.

4. The Metrization of a Space of Probability Measures

Throughout this section, let Ω denote a metric space with Borel σ-algebra $\mathcal{B}$, and $\mathcal{P}$ the space of all probability measures on $(\Omega, \mathcal{B})$. Let η and ν be two probability measures in $\mathcal{P}$ with associated densities given by f and g, respectively, with respect to a $\sigma-$finite dominating measure μ. Such a measure μ can always be constructed, e.g., take $\mu = (\eta + \nu)/2$. While there are many metrics we can define on $\mathcal{P}$—each with its own qualitative and practical benefits—we borrow from the overview by Gibbs and Su [10] of some of the most commonly used ones among probabilists.

Definition 3. Discrepancy metric

$$d_D(\eta, \nu) := \sup_{\text{all closed balls } B \in \mathcal{B}} |\eta(B) - \nu(B)| \tag{4.1}$$

The discrepancy metric recognizes the underlying topology of Ω. We can think of it as a "worst case" difference between measures. It is similar to the Kolmogorov-Smirnov statistic used to measure the distance between empirical and continuous one-dimensional cumulative distribution functions. Unfortunately, it has limitations with respect to computational implementations owing to the domain over which the supremum is being taken.

Definition 4. (Lévy-)Prokhorov metric

$$d_P(\eta, \nu) := \inf\left\{\varepsilon > 0 \,:\, \eta(B) \le \nu(B^\varepsilon) + \varepsilon \quad \forall\, B \in \mathcal{B}\right\}, \text{ where } B^\varepsilon = \left\{x : \inf_{y \in B} d(x, y) \le \varepsilon\right\} \tag{4.2}$$

Note that although not intuitive, this does satisfy the symmetric requirement of metrics. It should be evident from the definition that this is not an easy quantity to compute nor implement. However, it is theoretically

relevant as it metrizes weak convergence on any separable metric space. It has been shown that $d_P(\eta, \nu)$ is the minimum distance between random variables distributed according to η, ν [10].

Definition 5. Total variation distance

$$d_T(\eta, \nu) := \sup_{A \in \mathcal{B}} |\eta(A) - \nu(A)| = \frac{1}{2} \max_{|h| \leq 1} \left| \int h \, d\eta - \int h \, d\nu \right| \tag{4.3}$$

where $h : \Omega \to \mathbb{R}$ satisfies $|h| \leq 1$.

For a countable state space Ω, we have

$$d_T(\eta, \nu) = \frac{1}{2} \sum_{x \in \Omega} |\eta(x) - \nu(x)| \,, \tag{4.4}$$

which is half the L^1-norm between the two measures. Note that this metric takes values in $[0, 1]$. We note that an equivalent definition is given by

$$d_T(\eta, \nu) := \frac{1}{2} \int_\Omega |f - g| \, d\mu. \tag{4.5}$$

Definition 6. Hellinger distance

$$d_H(\eta, \nu) := \frac{1}{\sqrt{2}} \left[\int_\Omega (\sqrt{f} - \sqrt{g})^2 \, d\mu \right]^{1/2}. \tag{4.6}$$

Note this definition is independent of the dominating measure μ (this may be more easily seen in Eq. (4.8)). Also note that while many authors do not include the leading coefficient, we do so to normalize the metric to $[0, 1]$ purely for aesthetic purposes. We also have that for a countable state space Ω,

$$d_H(\eta, \nu) := \frac{1}{\sqrt{2}} \left[\sum_{\omega \in \Omega} \left(\sqrt{\eta(\omega)} - \sqrt{\nu(\omega)} \right)^2 \right]^{1/2}. \tag{4.7}$$

Notice that by properties of probability measures, Eq. (4.6) is equivalent to

$$d_H(\eta, \nu) = \frac{1}{\sqrt{2}} \left[2 \left(1 - \int_\Omega \sqrt{fg} \, d\mu \right) \right]^{1/2}. \tag{4.8}$$

While the above error analysis presented in Sec. 3.2 can be undertaken using any of the metrics mentioned above, we choose the Hellinger Distance in the numerical results of Sec. 6 for several reasons. First, it is computationally efficient to implement since it does not contain any maximums or supremums over σ-algebra $\mathcal{B}$. While the same can be said for the Total Variation distance given in Eq. (4.5), the Hellinger distance bounds the former as follows:

$$d_H^2(\eta, \nu) \leq d_{TV}(\eta, \nu) \leq \sqrt{2} \, d_H(\eta, \nu). \tag{4.9}$$

Thus, the Hellinger distance and its square also provide a useful insight into the distance between measures defined by the Total Variation metric.

Second, and perhaps most importantly, since the spaces Λ we are considering are generally bounded and finite, the Hellinger metric metrizes weak convergence (see Thm. 6 in [10]). The latter property is of notable importance because the QoI maps we study are indeed (component-wise) functionals on the space of model inputs Λ. Thus, convergence of a sequence of probability measures under the Hellinger metric implies that the QoI's will also converge component-wise in $\mathbb{R}$. In other words, convergence in the Hellinger metric implies the convergence of the sampled QoI map to the exact QoI map since the map is a linear functional of the probability measure. In other words, if $P_{\Lambda,M,N,h}$ converges to either $P_{\Lambda,M,N}$, $P_{\Lambda,M}$, or P_{Λ} using the Hellinger metric, this implies that the error converged to zero in the numerically computed $Q(\lambda^{(j)})$. Thus, convergence in the Hellinger metric implies, in a sense, convergence of the numerical method used to construct the QoI map. Furthermore, recall that weak convergence $P_n \to P$ is defined to mean

$$\int P_n f \to \int P f$$

for bounded Lipschitz functions f. Taking $f = \chi_A$, this leads to the following implication:

$$P_{\Lambda,M,N} \to P_{\Lambda} \implies P_{\Lambda,M,N}(A) \to P_{\Lambda}(A) \quad \forall A \in \mathcal{B}_{\Lambda}.$$

It is a combination of computational ease of implementation and theoretical implications that motivates the choice of the Hellinger distance as the metric used in the numerical results of Section 6.

5. Numerical Implementations and Software Contributions

The measures computed from Algorithm 1 are defined on a set of samples $S = \left\{\lambda^{(j)}\right\}_{j=1}^{N}$ which implicitly define a Voronoi-cell partition $\left\{\mathcal{V}^{(j)}\right\}_{j=1}^{N}$ of the parameter space Λ. We let $\mathcal{B}_{\Lambda,N}$ denote the *computational algebra* generated by $\left\{\mathcal{V}^{(j)}\right\}_{j=1}^{N}$, i.e., using standard measure theory notation,

$$\mathcal{B}_{\Lambda,N} = \sigma\left(\left\{\mathcal{V}^{(j)}\right\}_{j=1}^{N}\right).$$

Clearly, $\mathcal{B}_{\Lambda,N} \subset \mathcal{B}_{\Lambda}$ and the events $A \in B_{\Lambda,N}$ represent the $A \in \mathcal{B}_{\Lambda}$ for which we can "easily" compute probabilities and make inferences. While Algorithm 1 ultimately defines a probability measure implicitly on $(\Lambda, \mathcal{B}_{\Lambda})$, computationally this is almost never done and the measures are only interrogated on the computational algebra associated with the set of samples.

Different sets $S_k = \left\{\lambda^{(i)}\right\}_{i=1}^{N_k}$, where the $\lambda^{(i)}$'s and N_k's may be completely different for each k, will lead to different measures computed from Algorithm 1. Each S_k induces a computational algebra which we index using the notation $\mathcal{B}_k$ for simplicity, where it is understood that $\mathcal{B}_k = \mathcal{B}_{\Lambda,N_k}$.

This poses an immediate problem with respect to a computational approach to computing d_H: how do we compare measures P_{Λ,M,N_1} and P_{Λ,M,N_2} which may be defined on completely different computational algebras (even if $N_1 = N_2$)? See Figure 2 for an illustration of such a scenario.

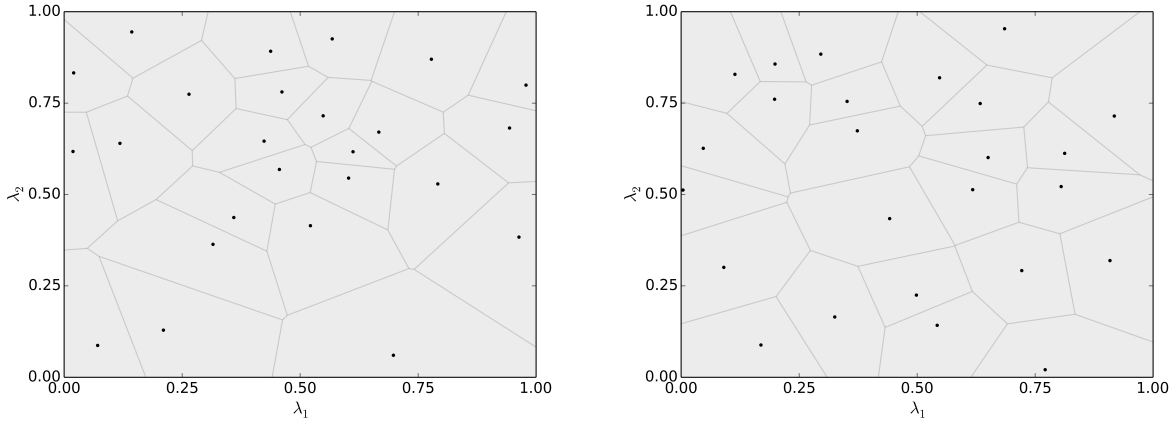

FIGURE 2. Two different Voronoi partitions induced by $N_1 = N_2 = 25$ uniform i.i.d. random samples.

The proof of the following Lemma describes how to "computationally extend" *any* probability measure defined on a computational algebra to a full σ-algebra $\mathcal{B}$, which we exploit in Algorithm 2.

Lemma 1. *Let μ be a measure on $(\Lambda, \mathcal{B}_\Lambda)$, $\left\{\mathcal{V}^{(j)}\right\}_{j=1}^N$ be a partition of Λ, and $\mathcal{B}_{\Lambda,N}$ the computational algebra generated by $\left\{\mathcal{V}^{(j)}\right\}_{j=1}^N$. Assume $\mu(\mathcal{V}^{(j)}) > 0 \; \forall \; j = 1, \ldots, N$. Then, there exists a probability measure η on $(\Lambda, \mathcal{B}_\Lambda)$ such that $\eta(A) = \eta_N(A) \; \forall \; A \in \mathcal{B}_{\Lambda,N}$.*

In the proof below, we use η_N and μ to construct a type of "discrete" Radon-Nikodym derivative of η. This is motivated by the formal structure of solutions given by Algorithm 1.

Proof. Let

$$f_N(\lambda) = \sum_{j=1}^{N} \frac{\eta_N(\mathcal{V}^{(j)})}{\mu(\mathcal{V}^{(j)})} \chi_{\mathcal{V}^{(j)}}(\lambda). \tag{5.1}$$

Then, for any $A \in \mathcal{B}_\Lambda$, define

$$\eta(A) = \int_A f_N(\lambda) \, d\mu. \tag{5.2}$$

We verify that η is a probability measure on $(\Lambda, \mathcal{B}_\Lambda)$ and that $\eta(A) = \eta_N(A) \; \forall \; A \in \mathcal{B}_{\Lambda,N}$ below:

(i) [Positive] Let $A \in \mathcal{B}_\Lambda$.

$$\begin{aligned}\eta(A) &= \int_A f_N(\lambda)\, d\mu \\ &= \int \chi_A \sum_{j=1}^{N} \frac{\eta_N(\mathcal{V}^{(j)})}{\mu(\mathcal{V}^{(j)})} \chi_{\mathcal{V}^{(j)}}(\lambda)\, d\mu \\ &= \sum_{j=1}^{N} \left(\frac{\eta_N(\mathcal{V}^{(j)})}{\mu(\mathcal{V}^{(j)})} \int \chi_{A \cap \mathcal{V}^{(j)}}(\lambda)\, d\mu \right) \\ &= \sum_{j=1}^{N} \left(\frac{\eta_N(\mathcal{V}^{(j)})}{\mu(\mathcal{V}^{(j)})} \mu\left(A \cap \mathcal{V}^{(j)}\right) \right) \geq 0\end{aligned}$$

(ii) [Definite]

$$\eta(\varnothing) = \int_\varnothing f_N(\lambda)\, d\mu = \int \chi_\varnothing f_N(\lambda)\, d\mu = \mu(\varnothing) = 0$$

(iii) [Countably Additive] Let $\{A_k\}_{k=1}^{\infty} \subset \mathcal{B}_\Lambda$.

$$\begin{aligned}\eta(\cup_k A_k) &= \int_{\cup_k A_k} f_N(\lambda)\, d\mu = \int \chi_{\cup_k A_k} f_N(\lambda)\, d\mu \\ &= \int \left(\sum_k \chi_{A_k} \right) f_N(\lambda)\, d\mu = \sum_k \int \chi_{A_k} f_N(\lambda)\, d\mu \\ &= \sum_k \int_{A_k} f_N(\lambda)\, d\mu = \sum_k \eta(A_k)\end{aligned}$$

Finally, let $A \in \mathcal{B}_{\Lambda,N} \subset \mathcal{B}_\Lambda$. Then there exists some $j^* \in \{1, 2, \cdots, N\}$ such that $\mathcal{V}^{(j^*)} = A$. We have that

$$\begin{aligned}\eta(A) &= \int_A f_N(\lambda)\, d\mu = \int \sum_{j=1}^{N} \left(\frac{\eta_N(\mathcal{V}^{(j)})}{\mu(\mathcal{V}^{(j)})} \chi_{\mathcal{V}^{(j)}} \right) \chi_{\mathcal{V}^{(j^*)}}\, d\mu \\ &= \int \frac{\eta_N(\mathcal{V}^{(j^*)})}{\mu(\mathcal{V}^{(j^*)})} \chi_{\mathcal{V}^{(j^*)}}\, d\mu = \frac{\eta_N(\mathcal{V}^{(j^*)})}{\mu(\mathcal{V}^{(j^*)})} \mu(\mathcal{V}^{(j^*)}) \\ &= \eta_N(\mathcal{V}^{(j^*)}) = \eta_N(A).\end{aligned}$$

□

We note that in practice, $\chi_{\mathcal{V}^{(j)}}(\lambda)$ requires the use of nearest-neighbor computations, but otherwise evaluation of Eq. (5.1) is straightforward to compute.

Since we now have a way to extend probability measures defined on $(\Lambda, \mathcal{B}_{\Lambda,N})$ to probability measure on $(\Lambda, \mathcal{B}_\Lambda)$, we can use simple Monte-Carlo approximation schemes to the Hellinger distance between two probability measures defined on two separate computational algebras. This is demonstrated in Algorithm 2.

Algorithm 2: Hellinger Discretization

Let $(\Lambda, B_{\Lambda,N_1}, \eta_{N_1})$ and $(\Lambda, B_{\Lambda,N_2}, \eta_{N_2})$ be given.
Construct f_{N_1} and f_{N_2} and corresponding η_1, η_2 using Eq. (5.1) and Eq. (5.2), respectively.
Use Monte Carlo sampling to approximate

$$d_H^2(\eta_1, \eta_2) = \int_\Lambda \sqrt{f_{N_1}(\lambda)} - \sqrt{f_{N_2}(\lambda)}\, d\mu$$

.

6. Numerical Results

We established in Section 4 the definition of metrics on the space $\mathcal{P}$ of probability measures on some arbitrary space $(\Omega, \mathcal{B})$ and justified our choice for the Hellinger distance for its implications for the convergence of QoI. We want an estimated probability measure $\hat{\mu}_\Lambda$ on the parameter space to converge with respect to d_H to some reference measure μ_Λ (either some known prior distribution taken as truth or another approximation deemed to be sufficiently resolved for the given application or computational budget).[1]

All of our experiments follow the same structure:

[0-a] Select $Q \in \mathcal{Q}$ and define $P_{\mathcal{D}_Q}$ as a uniform distribution centered on a reference QoI value $Q(\lambda_{ref})$ for λ_{ref} taken as the midpoint of Λ. Note that $P_{\mathcal{D}_Q}$ is exactly discretized with $M = 1$ sample, so that

$$P_{\Lambda,1} = P_\Lambda.$$

[0-b] Create a regular grid of samples in $\Lambda = [0,1]^n$ using $N_{\text{ref},i}$ equispaced points in each dimension. Define $\bar{N} := \prod N_{\text{ref},i}$. Since n is small in the numerical examples shown here, we chose $N_{\text{ref},i} = 200\ \forall\ i$ in each example.

[0-c] Use Algorithm 1 to construct a reference solution $P_{\Lambda,\bar{N}} \approx P_\Lambda$.

[1] Generate $\left\{S_k^{(T)}\right\}_{T=1}^{50}$ sets of uniform i.i.d. random samples where $N_k = 25, 50, 100, 200, \ldots, 6400$, and T represents the number of repeated trials of a sample size N_k.

[2] Solve the SIPs using Algorithm 1 to construct $\left\{P_{\Lambda,M,N}^{(T)}\right\}_{T=1}^{50}$.

[3] Use $1E5$ i.i.d. random samples in the Monte Carlo step of Algorithm 2 to estimate $\left\{d_H^2(P_{\Lambda,M,N}^{(T)}, P_{\Lambda,\bar{N}})\right\}_{T=1}^{50}$.

[4] Average over all trials T for each N to estimate the *expected* Hellinger distance for N samples and analyze convergence to $P_{\Lambda,\bar{N}}$.

[5] Repeat steps [0-a]–[4] for each $Q \in \mathcal{Q}$ under consideration.

To isolate the effect of skewness on our ability to approximate sets with finite sampling, we choose our maps so that they preserve the sizes of sets between Λ and $\mathcal{D}$ under the push-forward measure given in Eq. (2.2).

[1]However, we could also choose to interrogate the push-forward measures given by propagating the $\hat{\mu}_\Lambda$ and μ_Λ forward to a data space by a QoI map and taking the Hellinger distance on the resulting output space. This would measure the ability of the maps to reconstruct the output probability measure.

6.1. **Example 1.** This example shows that if $Q^{(a)}$ is defined by a rotation of $Q^{(b)}$, then the accuracy and convergence rates of $P^{(a)}_{\Lambda,M,N}$ are identical to $P^{(b)}_{\Lambda,M,N}$. We expect this to be true since skewness is rotationally invariant, as we summarize in the following Proposition.

Proposition 1. *The quantity $S_Q(\lambda)$ is invariant under rotations performed on Q for any λ.*

Proof. If we apply a rotation Q, then the $J_{Q,\lambda}$ are also subject to the same rotation at each λ. Since rotations are unitary operators, the norms given in Eq. (3.3) used to define skewness are unaffected. □

To demonstrate this, we define the space of QoI maps $\mathcal{Q} = \left\{Q^{(a)}, Q^{(b)}, Q^{(c)}\right\}$, where all three are linear maps with the same local skewness $S_Q(\lambda) = 1\ \forall \lambda \in \Lambda$. The map $Q^{(a)}$ is the identity and the other two, $Q^{(b)}$ and $Q^{(c)}$ are rotations of $Q^{(a)}$ by randomly chosen angles. Following the algorithmic outline above, we perform a convergence study to $P_{\Lambda,\bar{N}}$ with results summarized in Figure 3. The convergence rates and expected errors in the SIPs associated with each of these maps are virtually indistinguishable. Thus, it appears that in light of Proposition 1 and these numerical results, the accuracy of the numerical solution to the SIP is invariant under rotations to the QoI map.

N	$Q^{(a)}$	$Q^{(b)}$	$Q^{(c)}$
200	$1.34E-01$	$1.35E-01$	$1.40E-01$
400	$9.40E-02$	$1.00E-01$	$1.00E-01$
800	$7.30E-02$	$7.36E-02$	$7.11E-02$
1600	$5.08E-02$	$5.13E-02$	$4.96E-02$
3200	$3.48E-02$	$3.50E-02$	$3.54E-02$
6400	$2.56E-02$	$2.53E-02$	$2.51E-02$

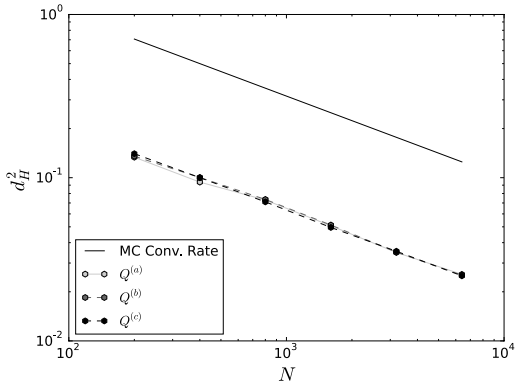

FIGURE 3. The results of $d_H^2(P_{\Lambda,M,N}, P_{\Lambda,\bar{N}})$.

6.2. **Example 2.** In this example, we demonstrate the key point of this study: the magnitude of skewness between QoI impacts accuracy by orders of magnitude, and thus in optimizing the choice of a QoI map, it is in our interest to pursue the minimization of skewness. This is especially true in problems where the number of random samples we are permitted to use is constrained by the computational cost of model evaluations.

To illustrate this point, we first define the linear maps

$$\mathcal{Q}_S := \left\{ Q^{(s)} = \begin{bmatrix} 1 & 0 \\ \sqrt{s^2-1} & 1 \end{bmatrix} \right\}_{s\in S}, \tag{6.1}$$

for $S = \{1, 2, 4\}$ because they allow us to control the global skewness (since it is equal to local skewness in a linear map) while preserving the measures of sets between Λ and $\mathcal{D}$. More specifically, the support of the solution to the SIP associated with each QoI map has equal μ_Λ-measure, which isolates the impact of accuracy solely to the skewness of the QoI map. We show what the component row vectors of these maps in Figure 4 and note the skewness is determined by the ratio of the magnitude of the black line to its projection onto the vertical axis (and each of these projects directly on to the unit vector). The skewness of these maps is given by the index s, so $Q^{(1)}$ is 1, the skewness of $Q^{(2)}$ is 2, and $S_{Q^{(4)}} = 4$.

The maps chosen for this example are expository ones that provide valuable insight despite their simplicity. For example, when solving many physics-based problems, local linear approximations are often used to simplify model evaluation and guide optimization procedures.

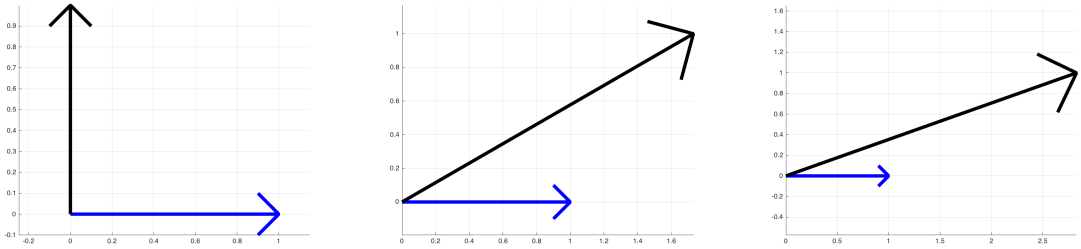

FIGURE 4. (Left to right): The component row-vectors of $Q^{(1)}$, $Q^{(2)}$, and $Q^{(4)}$. Our linear maps take $\mathbb{R}^2$ to $\mathbb{R}^2$ and can be visualized graphically as the component row-vectors of the matrices representing the transformation. The first row is highlighted in blue. The skewness is then simply equal to the reciprocal of the inverse sine of the angle between these vectors.

We see in Figure 5 that skewness has a very direct impact on the number of samples required to achieve a particular value for the Hellinger distance. We can see that the measure induced by $Q^{(1)}$ requires fewer than half the number of samples to be as accurately resolved as $Q^{(2)}$ does. The effect is even more pronounced

N	$Q^{(1)}$	$Q^{(2)}$	$Q^{(4)}$
200	$1.35E-01$	$2.03E-01$	$3.12E-01$
400	$9.96E-02$	$1.47E-01$	$2.15E-01$
800	$7.19E-02$	$1.04E-01$	$1.53E-01$
1600	$5.27E-02$	$7.49E-02$	$1.10E-01$
3200	$3.70E-02$	$5.25E-02$	$7.52E-02$
6400	$2.76E-02$	$3.86E-02$	$5.54E-02$

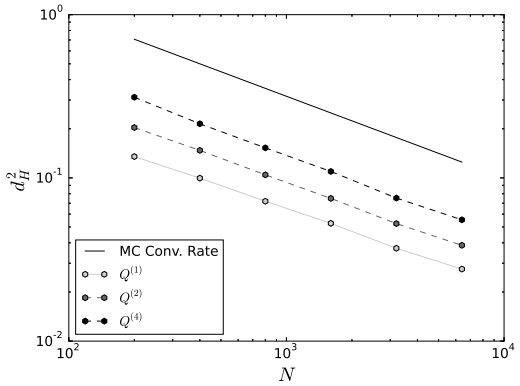

FIGURE 5. The results of $d_H^2(P_{\Lambda,M,N}, P_{\Lambda,\tilde{N}})$.

when compared against $Q^{(4)}$. It appears that if the ratio of skewness between two maps is 2, then the more-skewed map will require at least twice as many random samples to approximate the set on a a well-resolved discretization with the same error tolerance.

This provides a strong motivation for minimizing skewness and reinforces the results from [2], where it was demonstrated that a similar relationship existed in the number of samples required to remove error in inverse set approximations quantified by the μ_Λ-measure of the *symmetric difference* of the inverse sets.

6.3. **Example 3.** We extend the numerical investigation to a parameter space of dimension three to further illustrate that these results hold as we move towards higher dimensions. Generally, we have fewer QoI than number of uncertain model parameters, so we assume that the potential QoI maps are defined by the 2×3 matrices

$$\mathcal{Q}_S := \left\{ Q^{(s)} = \begin{bmatrix} 1 & 0 & 0 \\ \sqrt{s^2-1} & 1 & 0 \end{bmatrix} \right\}_{s \in S}. \tag{6.2}$$

Here, as in the previous example, the index s indicates the magnitude of skewness. Furthermore, the results of Example 6.1 justify the restriction of the maps to this form since any linear map of skewness s is simply a rotation of maps of this form.

N	$Q^{(1)}$	$Q^{(2)}$	$Q^{(4)}$
200	$2.98E-01$	$4.18E-01$	$5.60E-01$
400	$2.27E-01$	$3.27E-01$	$4.69E-01$
800	$1.81E-01$	$2.70E-01$	$3.97E-01$
1600	$1.46E-01$	$2.15E-01$	$3.09E-01$
3200	$1.15E-01$	$1.72E-01$	$2.44E-01$
6400	$9.09E-02$	$1.39E-01$	$1.95E-01$

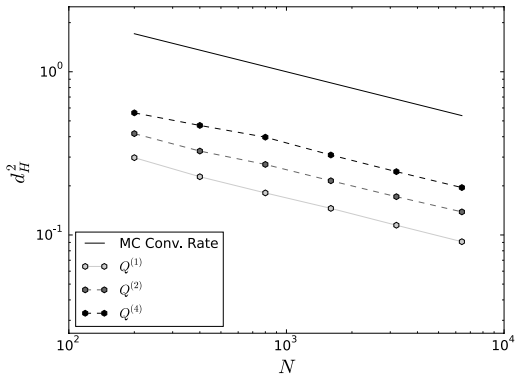

FIGURE 6. The results of $d_H^2(P_{\Lambda,M,N}, P_{\Lambda,M,\bar{N}})$ for $M = 1, \bar{N} = 8,000,000$.

In Figure 6, it appears that the effect of skewness is even more pronounced in higher dimensions, and that the number of samples required to achieve similar levels of accuracy between two maps with a ratio of skewness 2 is now quadrupled. The analysis of [7] suggested a dependence of accuracy related to the skewness raised to a power related to the dimension of the data space, but these results seem to indicate that the constant of proportionality may have a functional dependence on the dimension of Λ. This will be investigated further in a future study.

7. A Nonlinear Example

Here we present a nonlinear example to demonstrate that the general trend of the previous results also holds. Some nuanced differences do arise, however, and we address them after the problem statement. Consider the one-dimensional heat equation with homogeneous Neumann boundary conditions on the unit interval:

$$\begin{aligned} \rho c \frac{\partial T}{\partial t} &= \nabla \cdot (\kappa \nabla T) + f(x), \quad x \in (0,1), t \in (0,1) \\ f(x) &= A e^{\frac{-(x-0.5)^2}{w}} \chi_{[0,0.5]}(t) \end{aligned} \tag{7.1}$$

Alternative setup:

$$\begin{cases} \rho c \frac{\partial T}{\partial t} = \nabla \cdot (\kappa \nabla T) + f(x,t), & \text{if } x \in \Omega \\ \frac{\partial T}{\partial \vec{n}} = 0 & \text{if } x \in \partial\Omega \end{cases} \tag{7.2}$$

where $\Omega = (0,1) \times (0,1)$ is the space-time interior and $f(x,t) = A e^{\frac{-(x-0.5)^2}{w}} \chi_{[0,0.5]}(t)$.

Here, we interpret the following problem as heating the middle of an infinitesimally thin unit-length rod for half a second with the heat-source modeled by a Gaussian curve with amplitude $A = 50$ and variance of $w = 0.05$. The rod is subdivided in two, and each half has an uncertain thermal diffusivity $\kappa \in [0.01, 0.2]$. This yields a two-dimensional parameter space $\lambda = (\lambda_1, \lambda_2) \in [0.01, 0.2] \times [0.01, 0.2]$, where λ_1 represents the thermal diffusion on the left-half and λ_2 is the κ for the right half.

The quantities of interest we study are four point-evaluations of the state variable, at spatial location 0.25, 0.51, 0.67, and 0.98 along the rod. Choosing any pair of them for the inversion yields six possible quantities of interest maps. As before, we demonstrate that some choices appear to have advantages over others.

From the prior examples, we would suspect that choosing the QoI map with lower skewness results in lower Hellinger distances. However in the earlier experiments we utilized maps that inverted into sets of identical size, which is not the case in this nonlinear example; each QoI map scales sets differently depending on the location in the parameter space. To isolate this scaling effect, we attempt to compare QoIs that invert into sets of similar size *on average* but have differing average skewness.

This is what motivated our specific choice of spatial locations at which to measure the state variable T. Our first QoI $Q^{(a)}$ uses measurements at 0.25 and 0.51, and has average skewness of 1.08, and our second $Q^{(b)}$ uses measurements at 0.67 and 0.98, with average skewness 1.56. While we would have liked to use a

map with average skewness of 2 for a more similar comparison to the prior examples, this was the best range we could find where the maps inverted into sets of comparable size on average[2].

Owing to the nonlinearity of the problem, the Hellinger distances between reference and estimated probability measures now have an inherent dependence on the location of the point λ in the parameter space. We ran the simulations for a regular 3×3 grid exploring the interior of the parameter space and present a selection of two reference points that illustrate the differences in the nonlinear case.

In the two-dimensional data spaces $\mathcal{D}_Q$, our uncertainty is a uniform box centered at $Q(\lambda_{\text{ref}})$ with side-lengths of 0.1. When λ_{ref} is the bottom-left corner of our 3×3 grid, the two maps produce very different results, with $Q^{(a)}$ outperforming $Q^{(b)}$ in a similar manner as we saw in the linear examples (see Fig. 8). When λ_{ref} is in the upper-center of the grid, the inverse images are similar, as shown in Fig. 9, and so which map to use for inversion into this part of the parameter spaces is not a clear choice. We might even be tempted to use the more-skewed (on average) map since it inverts into a set with smaller support.

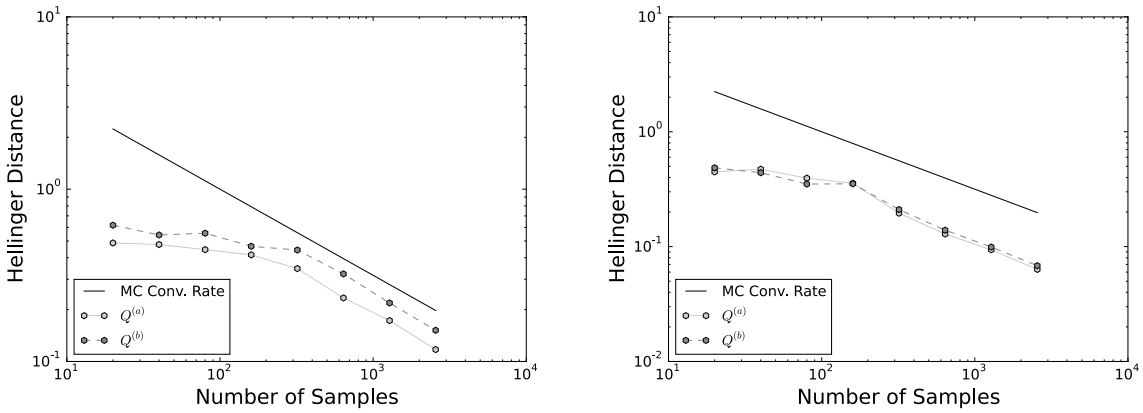

FIGURE 7. Comparison of the differences in Hellinger distances for the two maps and two reference points. The results for the bottom-left reference value is shown on the left and the top-center is shown on the right.

The Hellinger distance plots for these two reference values are compared in Fig. 7. Of the nine reference λ's we studied, $Q^{(a)}$ yielded no considerable advantage in terms of the number of samples required to approximate the inverse images in three cases (the plots were similar to that in the right of Fig. 7). In three cases, $Q^{(a)}$ performed just a bit better than $Q^{(b)}$, (somewhere between the two figures in Fig. 7). In two cases, $Q^{(a)}$ performed better than $Q^{(b)}$, as in the left of Fig. 7. In one case (with λ in the bottom right corner), the difference was even more dramatic ($Q^{(a)}$ yielded similar Hellinger distances with less than a fourth the samples).

[2]average local scaling is 1.99 versus 2.19

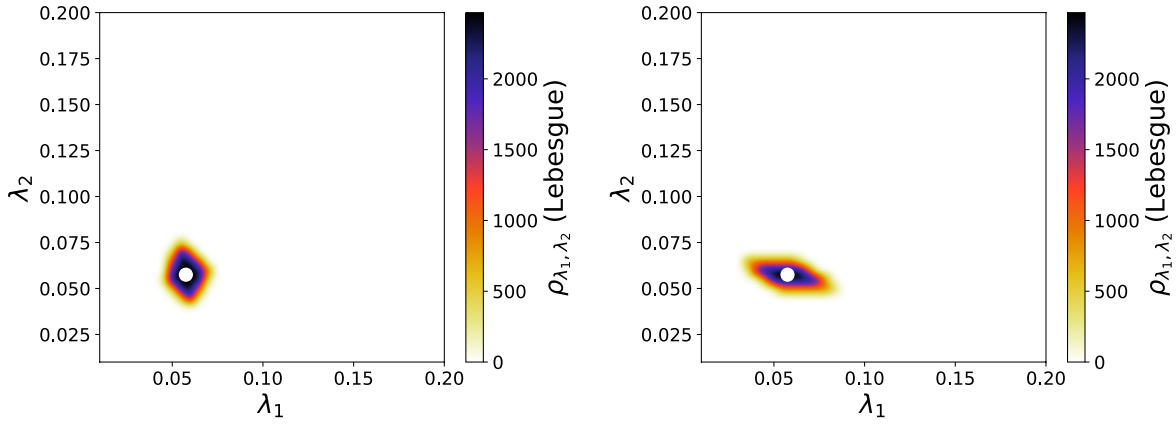

FIGURE 8. The inverse image of the reference measure for $Q^{(a)}$ (left) and $Q^{(b)}$ (right). The latter is visually quite a bit more skewed

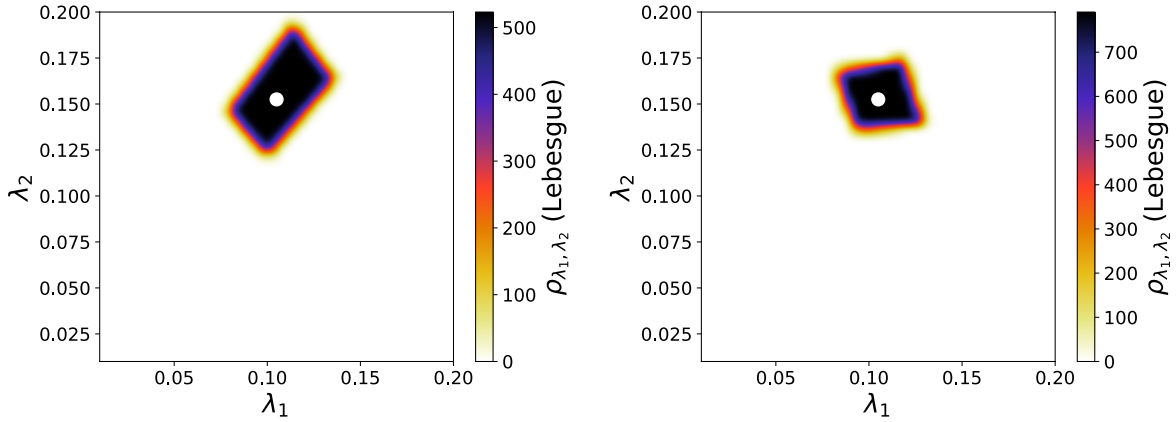

FIGURE 9. The inverse image of the reference measure for $Q^{(a)}$ (left) and $Q^{(b)}$ (right). Here, the local skewnesses are similar, so we do not expect to see much of a difference in the Hellinger distances.

With these nonlinear cases, we find that taking an "on average" approach is inefficient, as there can be dramatic differences in the geometric properties of the inverse images in the parameter space depending on the location. These results motivate further study into utilizing different QoI maps (perhaps some of those other four combinations available to us in this example) depending on where the samples came from in the parameter space. In general, we saw in this example that given that two maps invert into sets of similar size on average, using the one with lower skewness results in less samples required to accurately approximate the inverse image. The maps we used had average skewnesses that differed by 0.5 (instead of by 1), and the trend from the linear examples still held in significant portions of the parameter space.

8. Conclusion

We have demonstrated that the absolute error of an approximate solution to the SIP is proportional to the skewness of the map used for the inversion process. However, the convergence rate of Algorithm 1 is unaffected by such a choice. While it has been documented that the choice of QoI map used for inversion has an effect on the quality of the solution, here we achieve a more precise description of the relationship by equipping the space of our solutions (densities, measures) with a metric. The maps considered in this work all preserved the measures of sets between Λ and $\mathcal{D}$ in order to isolate the effect of skewness. We found that the resulting inverse densities were *easier to approximate* with finite sampling when the maps used for inversion had lower skewness values.

In situations where samples are expensive to evaluate, we would prefer to use less skewed maps provided they scale the sizes of sets similarly. Though it is still not clear exactly what weighting should be given to the multi-objective problem in [9], this additional insight provided by this work would certainly inform such a decision.

This work also motivates the development of an approach that breaks down the parameter space into subdomains over which some QoI map is deemed optimal with respect to its skewness and scaling properties. This decomposition would be used in defining a piecewise QoI map that is geometrically optimal globally. Any optimization approach (such as that in [9]) which seeks to find an optimal choice of QoI map *on average* will inherently not leverage information from available QoIs that have these characteristics. The additional computational overhead produced by such a partitioning and definition of a piecewise-defined map would be relatively low. In effect, this becomes an approach that has the potential to utilize more data at a similar computational cost to produce smaller regions of higher probability in the inverse measure. Preliminary results from research performed in Summer 2015 and 2016 appear promising.

One practical expansion of the work performed for this report would be to integrate the Hellinger distance function into a new module in BET. This might potentially function as a tool for researchers to compare successive refinements of their sample space and use the metric to determine when to stop using a Cauchy-criterion (or with respect to some reference measure computed for initial model studies).

References

[1] J. Breidt, T. Butler, and D. Estep, *A Measure-Theoretic Computational Method for Inverse Sensitivity Problems I: Method and Analysis*, SIAM Journal on Numr. Anal., 49 (2011), pp. 1836–1859.

[2] C. M. Bryant, S. Prudhomme, and T. Wildey, *Error decomposition and adaptivity for response surface approximations from pdes with parametric uncertainty*, SIAM/ASA Journal on Uncertainty Quantification, 3 (2015), pp. 1020–1045.

[3] T. Butler and D. Estep, *A numerical method for solving a stochastic inverse problem for parameters*, Ann. Nucl. Energy, 52 (2013), pp. 86–94.

[4] T. Butler, D. Estep, and J. Sandelin, *A Computational Measure Theoretic Approach to Inverse Sensitivity Problems II: A Posteriori Error Analysis*, SIAM Journal on Numr. Anal., 50 (2012), pp. 22–45.

[5] T. Butler, D. Estep, S. Tavener, C. Dawson, and J. Westerink, *A Measure-Theoretic Computational Method for Inverse Sensitivity Problems III: Multiple Quantities of Interest*, SIAM/ASA Journal on Uncertainty Quantification, 2 (2014), pp. 174–202. doi:10.1137/130930406.

[6] T. Butler, D. Estep, S. Tavener, T. Wildey, C. Dawson, and L. Graham, *Solving Stochastic Inverse Problems using Sigma-Algebras on Contour Maps*. In preparation, 2014.

[7] T. Butler, L. Graham, D. Estep, C. Dawson, and J. Westerink, *Definition and solution of a stochastic inverse problem for the manning's n parameter field in hydrodynamic models*, Adv. in Water Resour., 78 (2015), pp. 60 – 79.

[8] T. Butler and S. Mattis, *Enhanced surrogate modeling with adjoints for goal-oriented measure-theoretic inversion*, (2017). In Revision.

[9] T. Butler, M. Pilosov, and S. Walsh, *Simulation-based optimal experimental design: A measure-theoretic perspective*, (2017). In Review.

[10] A. L. Gibbs and F. E. Su, *On choosing and bounding probability metrics*, INTERNAT. STATIST. REV., (2002), pp. 419–435.

[11] L. Graham, T. Butler, S. Walsh, C. Dawson, and J. J. Westerink, *A measure-theoretic algorithm for estimating bottom friction in a coastal inlet: Case study of bay st. louis during hurricane gustav (2008)*, Monthly Weather Review, 145 (2017), pp. 929–954.

[12] L. Graham, S. Mattis, S. Walsh, T. Butler, M. Pilosov, and D. McDougall, *BET: Butler, Estep, Tavener Method v2.0.0*, Aug. 2016.