



Available online at www.sciencedirect.com

ScienceDirect

Comput. Methods Appl. Mech. Engrg. 407 (2023) 115906

**Computer methods
in applied
mechanics and
engineering**

www.elsevier.com/locate/cma

Parameter estimation with maximal updated densities

Michael Pilosov^a, Carlos del-Castillo-Negrete^{b,*}, Tian Yu Yen^c, Troy Butler^d,
Clint Dawson^b

^a *Mind the Math LLC, Denver, 80203, CO, USA*

^b *Oden Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, 78712, TX, USA*

^c *Sandia National Laboratories, Center for Computing Research, Albuquerque, 87175, NM, USA*

^d *Department of Mathematical and Statistical Sciences, University of Colorado Denver, Denver, 80202, CO, USA*

Received 20 November 2022; received in revised form 15 January 2023; accepted 15 January 2023

Available online 3 February 2023

Dataset link: <https://github.com/mindthemath/mud-data>

Abstract

A recently developed measure-theoretic framework solves a stochastic inverse problem (SIP) for models where uncertainties in model output data are predominantly due to aleatoric (i.e., irreducible) uncertainties in model inputs (i.e., parameters). The subsequent inferential target is a distribution on parameters. Another type of inverse problem is to quantify uncertainties in estimates of “true” parameter values under the assumption that such uncertainties should be reduced as more data are incorporated into the problem, i.e., the uncertainty is considered epistemic. A major contribution of this work is the formulation and solution of such a parameter identification problem (PIP) within the measure-theoretic framework developed for the SIP. The approach is novel in that it utilizes a solution to a stochastic forward problem (SFP) to update an initial density only in the parameter directions informed by the model output data. In other words, this method performs “selective regularization” only in the parameter directions not informed by data. The solution is defined by a maximal updated density (MUD) point where the updated density defines the measure-theoretic solution to the PIP. Another significant contribution of this work is the full theory of existence and uniqueness of MUD points for linear maps with Gaussian distributions. Data-constructed Quantity of Interest (QoI) maps are also presented and analyzed for solving the PIP within this measure-theoretic framework as a means of reducing uncertainties in the MUD estimate. We conclude with a demonstration of the general applicability of the method on two problems involving either spatial or temporal data for estimating uncertain model parameters. The first problem utilizes spatial data from a stationary partial differential equation to produce a MUD estimate of an uncertain boundary condition. The second problem utilizes temporal data obtained from the state-of-the-art ADvanced CIRCulation (ADCIRC) model to obtain a MUD estimate of uncertain wind drag coefficients for a simulated extreme weather event near the Shinnecock Inlet located in the Outer Barrier of Long Island, NY, USA.

© 2023 Elsevier B.V. All rights reserved.

Keywords: Uncertainty quantification; Inverse problems; Push-forward measure; Pullback measure; Parameter estimation

* Corresponding author.

E-mail addresses: michael.pilosov@ucdenver.edu (M. Pilosov), carlos.del-castillo-negrete@utexas.edu (C. del-Castillo-Negrete), tyen@sandia.gov (T.Y. Yen), troy.butler@ucdenver.edu (T. Butler), clint.dawson@utexas.edu (C. Dawson).

1. Introduction

We increasingly rely on quantitative predictions from computational models of physical systems to inform engineering design, predict the behavior of physical systems, and even shape public policy, e.g., see [1–4], for just a few such examples. It is therefore more important than ever to quantify – and reduce whenever possible – the uncertainties impacting such models.

Unfortunately, many key characteristics governing system behavior, described as model inputs (referred to here as parameters), are often hidden from direct observation. When observable model output data associated with some quantities of interest (QoI) are sensitive to variations in these parameters, we formulate and solve inverse problems using the output data for QoI to quantify uncertainties in parameters. Inverse problems therefore play a vital role in the uncertainty quantification (UQ) community.

The type of inverse problem formulated and methods utilized for solving it fundamentally rely upon the type of uncertainties to be quantified. In the UQ community, uncertainties are categorized as being either aleatoric (i.e., irreducible) or epistemic (i.e., reducible). We refer to an inverse UQ problem involving aleatoric uncertainties as a stochastic inverse problem (SIP) whereas an inverse UQ problem involving epistemic uncertainties is referred to as a parameter identification problem (PIP). The solution to a SIP is a probability density (or, more generally, a probability measure) that characterizes the differences in relative likelihoods of parameters given a density characterizing the inherent variability of observable data. However, the solution to the PIP is an estimate of a single “true” parameter, which may be obtained, for example, by maximizing a density that quantifies uncertainty in the estimate of this unknown, but fixed, parameter value due to finite observable data.

This brings us the contributions of this work, which we enumerate below for ease of reference.

1. At a high-level, a major contribution of this work is the formulation and analysis of solutions to a PIP utilizing a novel measure-theoretic method originally developed for solving a SIP.
2. In order to make direct comparisons to other formulations and solutions of the PIP based on popular Bayesian and least squares approaches, a full “linear Gaussian” theory of this new approach is developed.
 - As suggested by the subsequent comparison examples, this new approach is widely applicable to any PIP where Bayesian or least squares methods are regularly applied and has some potential advantages over solutions obtained via those methods.
 - As the theory and examples demonstrate, the approach developed in this work applies “selective regularization” in the parameter space. Specifically, an initial characterization of parameter uncertainty does not impact the updated characterizations of uncertainty in the directions that the mapping deems “informed” by the data.
3. A theoretical analysis for data-constructed residual QoI maps within this framework is developed including a guarantee of existence, uniqueness, and convergence of solutions to the PIP as more data are collected for a fixed linear measurement.
4. We conclude this work by presenting an alternative data-constructed residual QoI map for general (potentially nonlinear) spatial–temporal measurements based on a principal component analysis (PCA).

Below, we summarize the organization of the paper as well as key results that are connected to the contributions enumerated above.

Section 2: This provides a brief literature review, motivation, and context for both the SIP and PIP, which sets the stage for contribution 1 of this work.

- The formal definitions of the SIP and PIP are provided along with an illustrative example that compares the solutions to these problems.
- The example motivates the use of the *maximal updated density* (MUD) estimate in the measure-theoretic approach as an alternative to the Bayesian *maximum a posteriori* (MAP) point for solving the PIP.

Section 3: The first data-constructed Quantity of Interest (QoI) map considered in this work is defined as a weighted mean error map. This sets the stage for both contributions 2 and 3 of this work.

- The concept of statistical sufficiency motivates the construction of this map.

- This type of map is utilized for scenarios where multiple (i.e., repeated) noisy data are available for general types of measurements and motivates the development of the linear Gaussian theory (i.e., contribution 2).

Section 4: This is the heart of contributions 2 and 3 as we prove existence and uniqueness for MUD points as well as a convergence result for MUD estimates based on the first type of data-constructed QoI map (i.e., the weighted mean error map).

- It is shown that once a sufficient number of data are collected for each measurement, the existence and uniqueness of MUD points are guaranteed for weighted mean error maps.
- The variance in these estimates is shown to vanish in the limit of infinite data *only in the directions of the parameter space informed by the measurement operator*. In other words, selective regularization is applied in MUD estimation.
- A closed-form expression for the MUD point is provided along with conceptual and quantitative comparisons to MAP and least squares parameter estimates. The quantitative comparisons of these parameter estimates are performed for various randomly constructed linear maps on 100-dimensional parameter spaces.

Section 5: This sets up contribution 4 of this work as we introduce an alternative data-constructed QoI based on a principal component analysis (PCA) that is easily applied to spatial and temporal data. Unlike the approach in Section 3, this PCA based QoI approach does not require the availability of multiple (i.e., repeated) data for each measurement.

Section 6: This concludes contribution 4 of this work as we demonstrate the effectiveness of the PCA based QoI maps at producing accurate MUD estimates for two types of problems.

- The first example utilizes spatial data to form a PCA based QoI map to estimate an uncertain boundary condition for a two-dimensional stationary partial differential equation.
- The second example utilizes temporal data of free surface elevations to estimate wind drag coefficients for the state-of-the-art AAdvanced CIRCulation (ADCIRC) model used to simulate an extreme weather event near the Shinnecock Inlet located in the Outer Barrier of Long Island, NY, USA.

Section 7: Concluding remarks and future directions for this work are provided.

Appendix A: The detailed derivation is provided for the closed-form expression of the MUD point for linear Gaussian maps that is utilized in Section 4.

Appendix B: In the interest of scientific reproducibility, the details on obtaining the code and data utilized throughout this manuscript are provided.

2. Comparing inverse problems and solutions

We begin with a formal discussion of the SIP and PIP in order to provide a brief literature review on these problems and their solutions. This sets the stage for the remainder of this work and the contributions therein.

The mapping from model parameters to observable model outputs defines what we refer to as a QoI map. We assume the model parameters are hidden from direct observation and must be inferred from observable data associated with the QoI. When model parameters possess aleatoric uncertainties, e.g., due to naturally occurring variability in system inputs, then the specification of a probability measure quantifying uncertainties in the QoI data leads to the formulation of a SIP. A solution to the SIP is given by a pullback of this probability measure onto the space of parameters. We refer to a pullback measure as a data-consistent solution since its push-forward through the QoI map matches the probability measure quantifying uncertainties in the QoI data.

While it is possible to construct explicit approximations to data-consistent measures in terms of estimating measurable events and their probabilities in the parameter space (e.g., see [5]), such “set-based” approximations become computationally intractable for high-dimensional parameter spaces or geometrically complex and/or

computationally expensive QoI maps. A recently developed density-based approach [6–8] solves the SIP in a novel way by first solving a stochastic forward problem (SFP). Specifically, an *initial* probability measure is first specified on the parameters to encode any prior knowledge of parameter variability. Then, an SFP is solved where the push-forward of the initial probability measure is used to define a *predicted* probability measure on the QoI. The discrepancy between the predicted and *observed* probability measures on the QoI, expressed as a ratio of probability density functions (more generally, Radon–Nikodym derivatives), is then used to *update* the initial probability density. The *updated* probability measure associated with this density is then data-consistent. Moreover, the updates to the initial probability measure only occur in directions informed by the QoI. In other words, the initial probability measure serves to regularize the space of all pullback measures solving the SIP to produce a unique solution.

The SIP and its solution methodologies are based on rigorous measure theory using the Disintegration Theorem [9,10] as the central tool in establishing existence, uniqueness, and stability of solutions. Updated probability measures often have complex structures that are not well approximated by a family of parametrically defined distributions (e.g., Gaussian). This attribute of the solution further distinguishes this measure-theoretic approach from typical Bayesian-inspired approaches, e.g., Hierarchical Bayesian methods [11–14], that specify prior distributions from a parametric family of distributions along with additional a priori assumed distributions on the hyper-parameters introduced by this parametric family (e.g., the means and variances of a Gaussian). Subsequently, solutions to the SIP using Bayesian approaches will not, in general, produce solutions (defined as posterior distributions) whose push-forward matches the observed distribution. In fact, the push-forward of the posterior is not even of general interest in most Bayesian paradigms. Instead, the posterior predictive, which defines the distribution of possible unobserved values is of central interest [13]. The posterior predictive is constructed as a conditional distribution on the observations but makes practical use of the posterior through a marginalization. These differences are not surprising when one considers that the Bayesian inverse problem that is perhaps most familiar in the UQ community solves an inverse problem involving epistemic uncertainty, as we describe and expand upon below.

In a typical Bayesian framework [14–27], one of the initial assumptions is that data obtained on a QoI are polluted by measurement error, i.e., the data are “noisy”. Measurement errors can theoretically be reduced using improved measurement instruments (i.e., they are epistemic in nature). A data-likelihood function is used to express the relative likelihoods that all of the data came from a particular choice of the parameter. Encoding any initial assumptions about which parameters are more likely than others as a prior density allows the formal construction of a posterior density as a conditional density that describes the difference in relative likelihoods of any parameter value given the data.

It is common to use specific point estimators such as the maximum a posteriori (MAP) point given by the mode of the posterior as the actual solution to the inverse problem. The posterior is then re-interpreted as providing descriptions of uncertainty in that specific point estimate. The Bernstein–von Mises theorem [28] provides conditions under which the posterior will become concentrated around the single true parameter in the limit of infinite data [13]. In other words, the use of such point estimates is useful when solving a PIP as the uncertainty in such estimates, usually quantified via the covariance of the posterior, “shrinks” as more data are incorporated.

The Bayesian framework is fundamentally designed to solve a PIP while the measure-theoretic framework as presented in [6–8] is instead designed to solve a SIP. This difference in methods is not semantic as the SIP and PIP themselves arise under fundamentally different conditions. At a more conceptual level, a SIP arises more naturally in the context of applications where the objective is to quantify uncertainties across many instantiations of a system whereas a PIP arises more naturally in the context of applications where the objective is to quantify uncertainties in a specific realization of a system. For instance, in [29], SIPs and their data-consistent solutions are considered in the context of both a verification & validation and quality control problem in the design and manufacturing of thin elastic membranes. Separately, the work of [30] utilized data-consistent solutions to learn a distribution of microstructure parameters in the context of computational materials science.

To summarize at a high-level, the typical Bayesian approach to an inverse problem focuses on first modeling epistemic uncertainties in data on a QoI obtained from a true, but unknown, parameter value, which we denote by λ^\dagger . The objective is to then estimate this specific λ^\dagger and quantify uncertainties in this point estimate. This is in contrast to the SIP and its data-consistent solutions that are defined as pullback measures of an observed probability measure on the QoI. To help build intuition about these differences, we summarize key technical details about these inverse problems and their solutions before presenting an example that highlights differences in solutions and the potential of using the data-consistent approach for solving the PIP.

2.1. Terminology, notation, and the inverse problems

To make comparisons more clear, we first introduce shared notation between the SIP and PIP. Denote by Λ the space of (input) parameters for the model. Denote by Q the (potentially vector-valued) QoI map from the parameter space, Λ , to the data space defined by $\mathcal{D} := \{Q(\lambda) : \lambda \in \Lambda\}$. Note that while any given component of the QoI map is generally a composition of a functional applied to the solution of the model, we use the notation $Q(\lambda)$ to emphasize the dependence of the QoI on the choice of parameter. For simplicity in presentation, we assume $\Lambda \subseteq \mathbb{R}^p$ and $\mathcal{D} \subset \mathbb{R}^m$ for finite p and m . We use $Q^{-1}(E)$ for any $E \subset \mathcal{D}$ to denote the *preimage* of E , i.e., $Q^{-1}(E) = \{\lambda \in \Lambda : Q(\lambda) \in E\}$. Unless otherwise specified, we assume that Λ and \mathcal{D} are equipped with (Borel) σ -algebras to define measurable spaces, Q is a measurable map between these spaces, and that subsets of these spaces are taken from these σ -algebras.

2.1.1. The stochastic inverse problem (SIP) and parameter identification problem (PIP)

We first define the concept of push-forward measures as solutions to the SFP mentioned above, which also helps frame the SIP more clearly as the direct inversion of the SFP.

Definition 2.1 (Stochastic Forward Problem (SFP)). Given an initial (i.e., initially specified) probability measure \mathbb{P}_{init} on Λ , the SFP is to determine the push-forward probability measure \mathbb{P}_{pred} on \mathcal{D} , defined by

$$\mathbb{P}_{\text{pred}}(E) := \mathbb{P}_{\text{init}}(Q^{-1}(E)),$$

for all events $E \subset \mathcal{D}$.

We often refer to the push-forward of the initial measure as the *predicted* measure since it may be constructed before any observed data are known. This also helps to distinguish it from the *observed* measure used in the formulation of the SIP.

Definition 2.2 (Stochastic Inverse Problem (SIP)). Given an observed probability measure \mathbb{P}_{obs} on \mathcal{D} , the SIP is to determine a pullback probability measure, denoted \mathbb{P}_{Λ} , on Λ , which is data-consistent in the sense that

$$\mathbb{P}_{\Lambda}(Q^{-1}(E)) = \mathbb{P}_{\text{obs}}(E), \quad (1)$$

for all events $E \subset \mathcal{D}$.

Unless the map Q is a bijection, we do not expect that there is a unique \mathbb{P}_{Λ} solving the SIP, but rather there is a class of pullback measures that solve the SIP. In [5], a disintegration theorem [10] along with an ansatz is used to establish the existence of solutions to the SIP that are unique up to the choice of ansatz. An algorithm is provided in [5] for explicitly approximating pullback measures by applying a specified ansatz to approximations of contour events, i.e., approximations of $Q^{-1}(E_i)$ where $\{E_i\}_{i \in \mathcal{I}}$ is a partition of \mathcal{D} .

In [6], an alternative density-based approach is presented that is computationally simpler to implement, scales well with increasing parameter dimension, and is stable with respect to perturbations in the initial and observed probability measures. We refer the interested reader to [6] for the theoretical and algorithmic details. Here, we summarize the density-based solution to the SIP as:

$$\pi_{\text{update}}(\lambda) := \pi_{\text{init}}(\lambda) \frac{\pi_{\text{obs}}(Q(\lambda))}{\pi_{\text{pred}}(Q(\lambda))}. \quad (2)$$

The densities (i.e., Radon–Nikodym derivatives) π_{init} and π_{obs} are associated with the specification of \mathbb{P}_{init} and \mathbb{P}_{obs} , respectively. The density π_{pred} is associated with the predicted measure \mathbb{P}_{pred} . The solution, π_{update} , is referred to as an updated density because it is defined as a multiplicative update to the initial density, π_{init} . It is worth emphasizing that constructing π_{update} as a solution to the SIP requires a solution to the SFP. If π_{init} and π_{obs} are prescribed, then the solution only requires the construction of π_{pred} , which is usually achieved via non-parametric means (see [6] for details). As a consequence, when $m < p$, the solution to the SIP is obtained by solving a forward UQ problem on a lower-dimensional space.

In order to ensure that π_{update} is in fact a density, a predictability assumption is required [6]. A practical form of the predictability assumption is that there exists a constant $C > 0$ such that $\pi_{\text{obs}}(q) \leq C\pi_{\text{pred}}(q)$ for (almost

every) $q \in \mathcal{D}$. Conceptually, we interpret the predictability assumption as stating that we are able to predict all the observed data. If the predictability assumption is satisfied, then as a consequence of π_{update} being a density it follows that

$$1 = \int_{\Lambda} \pi_{\text{init}}(\lambda) \frac{\pi_{\text{obs}}(Q(\lambda))}{\pi_{\text{pred}}(Q(\lambda))} d\mathbb{P}_{\text{init}} = \int_{\mathcal{D}} \frac{\pi_{\text{obs}}(Q(\lambda))}{\pi_{\text{pred}}(Q(\lambda))} d\mathbb{P}_{\text{pred}} = \int_{\mathcal{D}} r(Q(\lambda)) d\mathbb{P}_{\text{pred}} = \mathbb{E}(r), \quad (3)$$

where

$$r(Q(\lambda)) := \frac{\pi_{\text{obs}}(Q(\lambda))}{\pi_{\text{pred}}(Q(\lambda))}. \quad (4)$$

Thus, $\mathbb{E}(r)$ provides a convenient diagnostic for the computed MUD solution from a set of samples. Specifically, if the sample average of $\mathbb{E}(r)$ deviates too far from unity, then this informs us of potential violations of the predictability assumption or other sources of error (such as density approximation error), in the formulation of the updated density. In such cases, further analysis is needed to determine the exact issue, but it nonetheless proves invaluable in determining the trustworthiness of the updated density and any statistical inferences drawn from this density. We make extensive use of this diagnostic in Section 6.

This diagnostic also helps to frame the special role of π_{init} in the SIP as compared to the role of the prior density used in the Bayesian inverse problem that is discussed below. Specifically, π_{init} serves three roles in the definition and solution of the SIP that we emphasize in the list below.

- π_{init} represents an initial description of the *aleatoric* uncertainty on Λ .
- π_{init} and the associated π_{pred} allows us to construct a particular data-consistent solution in the form of π_{update} that is *unique* up to the specification of π_{init} .
- π_{update} differs from π_{init} in Λ *only* in the directions for which $Q(\lambda)$ exhibits sensitivity. In other words, π_{update} and π_{init} have identical conditional probability structures on the manifolds defined by $Q^{-1}(q)$ for (almost every) $q \in \mathcal{D}$, which is evident by $r(Q(\lambda))$ being constant on such manifolds. We often refer to these manifolds as the *generalized contours* of the map Q .

As discussed in the introduction, the major contributions and focus of this work are to utilize the data-consistent framework associated with the SIP to solve the PIP, which we define below.

Definition 2.3 (Parameter Identification Problem (PIP)). Given a finite amount of (possibly noisy) data on a QoI map obtained for a fixed, but unknown, parameter λ^\dagger , the PIP is to estimate λ^\dagger .

The estimate of λ^\dagger we consider in this work is given by the maximal updated density (MUD) point defined as

$$\lambda^{\text{MUD}} := \arg \max \pi_{\text{update}}(\lambda). \quad (5)$$

2.1.2. A Bayesian approach for solving the PIP

We now develop a typical Bayesian approach for solving the PIP following the framework described in [27]. Recalling that λ^\dagger refers to the true parameter value, let d denote the “noisy” data obtained on $Q(\lambda^\dagger)$, which is often represented as

$$d = Q(\lambda^\dagger) + \xi,$$

where ξ is a random variable used to model the measurement error that is often assumed to follow a Gaussian distribution. Then, the data-likelihood function, often written as a conditional density, $\pi_{\text{like}}(d | \lambda)$, is formed. This describes the differences in relative likelihoods that the data could have been generated from a particular λ . Ideally, the largest values of $\pi_{\text{like}}(d | \lambda)$ occur whenever λ is a “good” approximation to the true parameter λ^\dagger . The data-likelihood function is distinct from the observed density used in the data-consistent framework.

The next ingredient in the Bayesian framework is the specification of a prior density denoted by $\pi_{\text{prior}}(\lambda)$. The prior describes the different relative likelihoods assumed for the true parameter before data are collected. In other words, the prior represents an initial description of the *epistemic* uncertainty on Λ . This immediately distinguishes the role of the prior from the role of the *initial* density used in the data-consistent framework.

The posterior density (i.e., the solution to the Bayesian inverse problem) is given by a conditional density, denoted by $\pi_{\text{post}}(\lambda | d)$, proportional to the product of the prior and data-likelihood function. In other words,

$$\pi_{\text{post}}(\lambda | d) \propto \pi_{\text{prior}}(\lambda) \pi_{\text{like}}(d | \lambda)$$

Table 1

Updated density solving the data-consistent inverse problem (left) and posterior density solving the Bayesian inverse problem (right). The role of observable data, denoted by d , is made explicit in the posterior whereas the role is implicit in the updated in how it is used to construct the observed density, denoted by π_{obs} .

$$\pi_{\text{update}}(\lambda) = \pi_{\text{init}}(\lambda) \frac{\pi_{\text{obs}}(Q(\lambda))}{\pi_{\text{pred}}(Q(\lambda))} \quad \pi_{\text{post}}(\lambda | d) = \frac{\pi_{\text{prior}}(\lambda) \pi_{\text{like}}(d | \lambda)}{\int_{\Lambda} \pi_{\text{like}}(d | \lambda) \pi_{\text{prior}}(\lambda) d\mu_{\Lambda}}$$

This form of the density follows from Bayes' rule (not from the Disintegration Theorem as with the updated density). The posterior can be interrogated to assess the difference in relative likelihoods of a fixed parameter given the observed data. Subsequently, the posterior is often used to produce a “best” estimate of the true parameter. For example, the maximum a posteriori (MAP) point is the parameter that maximizes the posterior density.

2.2. Comparisons and an illustrative example

Before we utilize an illustrative example to compare the posterior and updated densities (along with the associated MUD and MAP points), we find it useful to summarize these densities side-by-side in [Table 1](#) and comment on a few notable aspects not mentioned above. Observe for the posterior density that the data-likelihood function appears in both the numerator and denominator. In particular, the data-likelihood function informs the normalizing constant¹ in the denominator. This is in contrast to the denominator of the updated density, which is given by the predicted density, which is in general not a constant, and can be constructed independently of π_{obs} .

A practical implication of this difference is that the updated density only alters the structure of the initial density in what we refer to as the “data-informed” parameter directions. Specifically, for a fixed $q \in \mathcal{D}$, let $C_q := \{\lambda \in \Lambda : Q(\lambda) = q\}$, i.e., C_q is a “contour” in parameter space. Then, for any $\lambda \in C_q$, we immediately have $\pi_{\text{update}}(\lambda) = r(q)\pi_{\text{init}}(\lambda)$ where $r(q)$ is a fixed constant of proportionality for all $\lambda \in C_q$. By contrast, while the posterior does not have to agree with the prior in any direction in parameter space, the prior does impact the structure of the posterior in all directions.

The previous paragraph is not – and should not be interpreted as – a criticism of the Bayesian inverse framework. It is only meant to highlight that the data-consistent and Bayesian frameworks formulate and solve inverse UQ problems from different perspectives and with different (although at times seemingly compatible) assumptions. Consequently, the solutions for an inverse problem formulated under either framework may differ significantly. As the example (adopted from [\[6\]](#)) below demonstrates, this is true even if we arbitrarily force the inverse problems to appear as similar as possible.

Before diving into the comparison example, we quickly discuss the computational costs between computing the MUD and MAP point in this work. The initial/prior and observed/likelihood PDFs are either usually specified exactly as part of the setup for the SIP and PIP or have their forms determined from data separate from model evaluations. However, for the density-based solutions to the SIP, we must generally approximate the predicted density. We utilize random samples and non-parametric kernel density estimates (KDEs) to approximate the various pushforward densities as well as the updated and posterior densities since this is a straightforward method for implementation.² Thus, the computational bottleneck in this work for both methods is the number of samples required to accurately estimate densities via a KDE, which is itself primarily dominated by the cost of evaluating these samples through our forward model (if it is expensive to evaluate). As mentioned at the beginning of [Section 2](#), there are “set-based” solution methodologies (e.g., see [\[5\]](#)), but these have issues with scaling to higher-dimensions. Another approach suggested in [\[31\]](#) yields promising results for data-consistent inversion that alleviates the curse of dimensionality by adapting the Laplace approximation method that is commonly deployed in Bayesian settings. Since the focus of this work is on the analysis of the MUD point and its utilization as an alternative to a MAP point, we leave to future work any analysis and comparison of alternative numerical approaches.

¹ The normalizing constant is commonly referred to as the evidence term.

² See [Appendix B](#) for details on acquiring the datasets and code utilized for all examples in this manuscript.

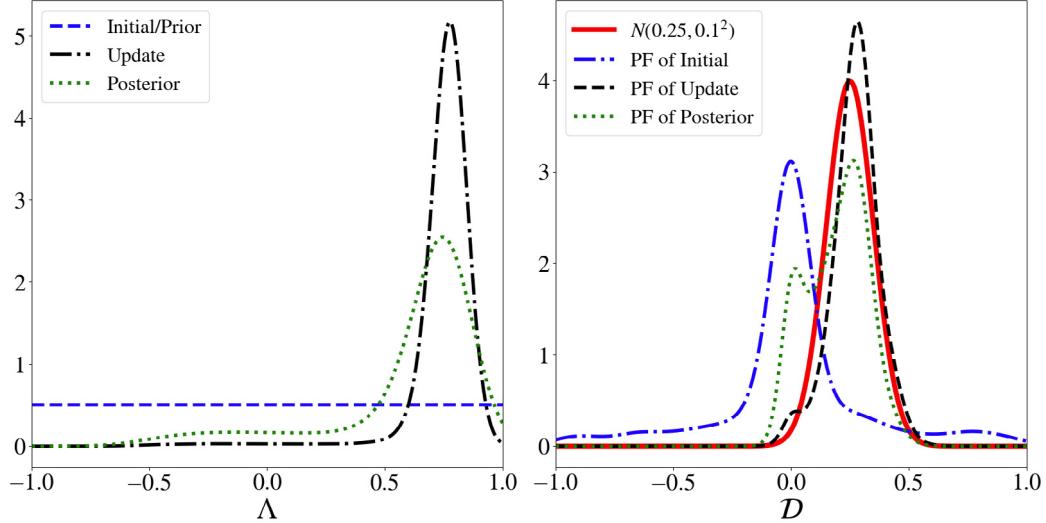


Fig. 1. (Left) The initial/prior PDF π_{init} (blue dashed curve), updated PDF π_{update} (black dashed curve), and posterior PDF π_{post} (green dashed-dotted curve) on Λ . (Right) The push-forward (PF) of the initial/prior PDF π_{pred} (blue dashed curve), observed/likelihood PDF (red solid curve), PF of the updated PDF π_{update} (black dashed curve), and the PF of the posterior PDF π_{post} (green dashed-dotted curve) for the QoI. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.2.1. Illustrative example: density comparison

Suppose $\Lambda = [-1, 1] \subset \mathbb{R}$ and $Q(\lambda) = \lambda^5$ so that $\mathcal{D} = [-1, 1]$. For the data-consistent framework, we assume $\pi_{\text{init}} \sim \mathcal{U}([-1, 1])$ and $\pi_{\text{obs}} \sim N(0.25, 0.1^2)$. The push-forwards of the initial, observed, and updated PDFs are shown in Fig. 1.

For the Bayesian inverse problem, we assume $d \in \mathcal{D}$ with $d = Q(\lambda^\dagger) + \xi$ where $\xi \sim N(0, 0.1^2)$. We then construct $\pi_{\text{post}}(\lambda | d)$ for this example assuming a uniform prior (to match the initial density) with an assumed observed value of $d = 0.25$ so that the data-likelihood function matches the observed density. The posterior and its push-forward are also shown in Fig. 1.

While the updated and posterior densities in Fig. 1 share certain similarities (e.g., they are uni-modal with similar locations of the mode), they are otherwise visibly distinct. The differences between these densities are made more evident by examining their push-forwards. The push-forward of the updated density agrees well with the observed density, which is to be expected. However, the push-forward of the posterior is bi-modal and does not match the observed density, which we recall is identical to the data-likelihood function in this case.

The takeaway is that each density is constructed to provide a solution to a *different* inverse problem. The posterior density is intended to provide estimates of a true parameter value (i.e., the posterior is designed to solve a PIP) whereas the updated density is intended to quantitatively characterize natural variations in parameter values (i.e., the updated density solves a SIP).

2.2.2. Illustrative example: MUD vs. MAP comparison

To keep this comparison as clear as possible, we initially consider the exact MUD and MAP estimates associated with the exact updated and posterior densities. In other words, we avoid any approximation errors due to the use of estimated densities (for the MUD point) and techniques often utilized for approximating solutions to the optimization problems (for both the MUD and MAP points). Given the same setup as above, the exact MUD point obtained by utilizing the exact updated density is $\lambda^{\text{MUD}} = 0.25^{1/5}$, which corresponds exactly to the value of λ^\dagger such that $Q(\lambda^\dagger) = 0.25$. Since the prior density is assumed uniform, the exact MAP point obtained by utilizing the exact posterior is identical to the MUD point.

It is worth noting that the MUD point in this example is invariant to the choice of initial density as long as the predictability assumption is not violated. In other words, the MUD point will always equal λ^\dagger , i.e., it is always an unbiased estimate of λ^\dagger in this example. However, if we altered the prior in this example to introduce a biased

initial guess away from λ^\dagger , the MAP point would no longer equal the MUD point because of the persistent bias present in the prior density.

This is not to say that the MUD point is necessarily a better estimate of λ^\dagger compared to the MAP point. One of the advantages of the MAP point is that the Bayesian formulation provides a natural quantification of uncertainty in the MAP point in terms of the variance in this estimate, which is a function of the amount of data utilized to produce it. We discuss this in the next section.

3. Data-constructed QoI maps I: The repeated measurement case

While the previous section establishes the MUD point as a potential alternative parameter estimate to the MAP point, we have yet to address the reduction of variance in MUD estimates as more data are included, which more naturally occurs for the MAP estimates. The mechanism for reducing variance in MAP estimates is the data-likelihood function, which is often written as a product of conditional densities associated with the individual components of the data vector. In other words, as more data are incorporated, the Bayesian approach is effectively increasing the dimension of the data space leading to a reduction of the variance in the subsequent MAP estimates. This is not surprising since the Bayesian formulation is within the context of a PIP. However, the MUD point is associated with the solution to a SIP for which a QoI map is defined and the dimension of the QoI data space is then fixed. This suggests that we either incorporate data into the QoI map construction, or into the specification of the observed density in order to make the MUD point and updated covariance more useful for solving a PIP. Below, we introduce some basic notation and return to the illustrative example from Section 2.2 to motivate the utilization of data in the construction of a QoI map rather than in the specification of the observed density.

3.1. Notation and an illustrative example

Suppose there exists m measurement devices for which repeated noisy data are obtained. For each $1 \leq j \leq m$, denote by $\mathcal{M}_j(\lambda^\dagger)$ the j th measurement device, and denote by N_j the number of noisy data obtained for $\mathcal{M}_j(\lambda^\dagger)$. Let $d_{j,i}$ denote the i th noisy datum obtained for the j th measurement where $1 \leq i \leq N_j$. We assume an unbiased additive error model for the measurement noise with independent identically distributed (i.i.d.) Gaussian errors so that

$$d_{j,i} = \mathcal{M}_j(\lambda^\dagger) + \xi_i, \quad \xi_i \sim N(0, \sigma_j^2), \quad 1 \leq i \leq N_j. \quad (6)$$

When a single measurement device is used to collect data, we often drop the index j .

We now return to the example from Section 2.2. Recall that for the Bayesian inverse problem, $Q(\lambda^\dagger) = 0.25$ and noisy data are given by $d = Q(\lambda^\dagger) + \xi$ where $\xi \sim N(0, 0.1^2)$. The data are used to construct the data-likelihood function. For the SIP, suppose that we use the sample mean of data to estimate the observed $N(0.25, 0.1^2)$ distribution. In other words, we use the sample mean of the data as an estimate for the mean of a normal distribution where the variance is assumed – as in the Bayesian inverse problem – to be 0.1^2 . The observed density and data-likelihood become significantly different from one another as more data are collected. The data-likelihood is in fact given by a product of normal densities evaluated at residuals of the data and the QoI map. From this perspective, we interpret Fig. 1 as comparing the updated and posterior densities when we collect a single datum that happens to agree with the mean of the distribution. We now draw $N = 5, 10$, and 20 samples to form estimates of π_{obs} and the likelihood functions and show representative results in Fig. 2.

For all values of N , the densities associated with the SIP all roughly agree. In fact, the initial and predicted densities never change by construction. As more data are incorporated (i.e., as N increases), the observed density estimates will stabilize as the variance in the mean of this density is reduced at a rate given by the central limit theorem. Subsequently, the updated densities also stabilize and the MUD points are all generally in good agreement with λ^\dagger . However, we observe that the posterior becomes more peaked with a MAP point that converges to the MUD point as N increases. Thus, the uncertainty associated with the MAP point, i.e., the variance in the MAP point estimate for any given realization of N data, is reduced as more data are incorporated. This demonstrates that incorporating data into the formulation of the SIP to produce more stable estimates of the observed density will not necessarily reduce the variance in the MUD point estimates. We therefore seek an alternative approach to incorporating data into the SIP through the construction of the QoI map itself that will result in reduced variance in MUD estimates as more data are incorporated.

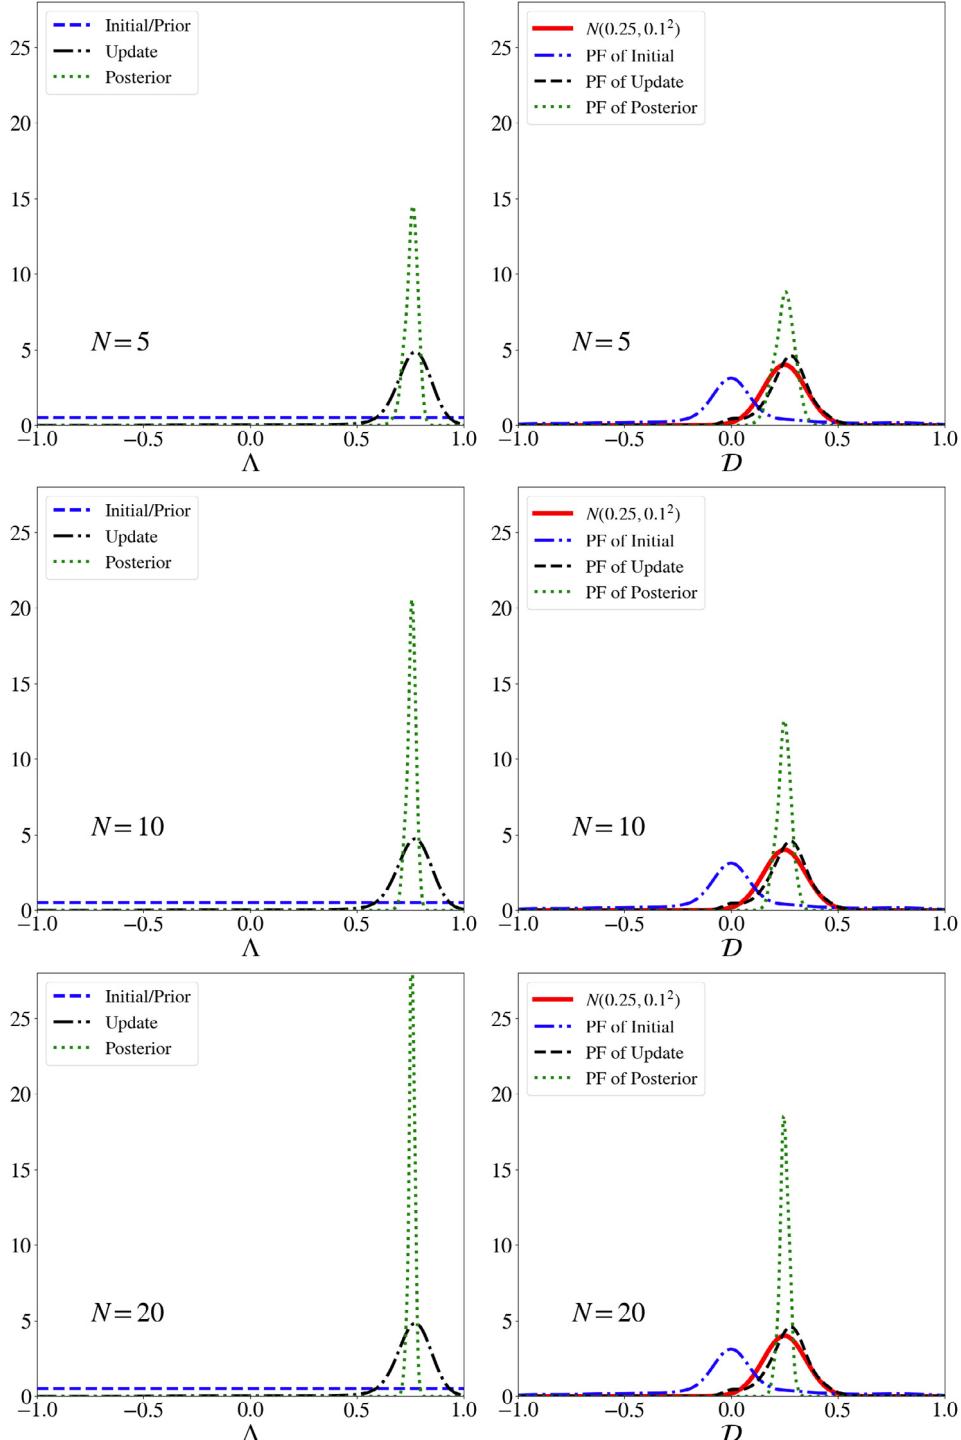


Fig. 2. (Top to Bottom): $N = 5, 10$, and 20 samples are used to solve the inverse problems for comparison. (Left) The initial/prior PDF π_{init} (blue solid curve), updated PDF π_{update} (black dashed curve), and posterior PDF π_{post} (green dashed-dotted curve) on Λ . (Right) The push-forward (PF) of the initial/prior density (i.e., π_{pred}) (blue solid curve), observed/likelihood PDF (red solid curve), PF of the updated PDF π_{update} (black dashed curve), and the PF of the posterior PDF π_{post} (green dashed-dotted curve) for the QoI. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.2. The weighted mean error map

We examine the data-likelihood function utilized in the Bayesian inverse problem and the concept of *statistical sufficiency* to motivate the data-constructed QoI map for MUD estimation. For simplicity, we initially consider a single measurement device, drop the j notation shown in (6), and let σ denote the standard deviation of the measurement error. In the Bayesian inverse problem, the data-likelihood function formed by observing N data points, $\{d_i\}_{i=1}^N$, has the form $\pi_{\text{like}}(d | \lambda)$ where d denotes an N -tuple with i th element given by d_i . Assuming i.i.d. noise present in each datum, this joint data-likelihood function takes the form

$$\begin{aligned} \pi_{\text{like}}(d | \lambda) &= \prod_{i=1}^N \pi_{\text{like}}(d_i | \lambda) \\ &\propto \prod_{i=1}^N \exp\left(-\frac{1}{2\sigma^2}(d_i - \mathcal{M}(\lambda))^2\right) \\ &= \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (d_i - \bar{d})^2\right) \exp\left(-\frac{1}{2(\sigma^2/N)}(\bar{d} - \mathcal{M}(\lambda))^2\right), \end{aligned} \quad (7)$$

where \bar{d} denotes the average of the N elements of d . To go from the middle to last line above, we perform a classic re-factorization of the quadratic terms required by the Fisher–Neyman factorization theorem (see [32] for details). The Fisher–Neyman theorem states that $T(d) := \frac{1}{N} \sum d_i = \bar{d}$ is a sufficient statistic for λ since the data-likelihood depends on the data only through the statistic T . Thus, the likelihood principle suggests that any two sets of data with the same sample mean \bar{d} will provide the same evidence or inference about λ . If we wish to determine the variability in inferences about λ due to differences in two separate collections of N data points, it “suffices” to consider the variability in the statistic \bar{d} .

Motivated by this perspective, we first consider the data-constructed QoI map defined by the mean error (ME) and denoted by $Q_{\text{ME}}(\lambda)$, with j th component $Q_{\text{ME},j}(\lambda)$, given by

$$Q_{\text{ME},j}(\lambda) := \frac{1}{N_j} \sum_{i=1}^{N_j} (\mathcal{M}_j(\lambda) - d_{j,i}). \quad (8)$$

This choice of map leads to a predicted and observed distribution that both depend on the observed data characteristics, e.g., $Q_{\text{ME},j}(\lambda^\dagger) \sim N(0, \sigma_j^2/N_j)$ which implies that π_{obs} is defined by a $N(\mathbf{0}_{m \times 1}, \Sigma_{\text{obs}})$ distribution where Σ_{obs} is a diagonal matrix with j th diagonal element given by σ_j^2/N_j . As shown below, by applying some scaling transformations to the components of Q_{ME} , we can remove the dependence of each component of the observed distribution on the characteristics of the observed data, which simplifies the practical implementation.

The weighted mean error (WME) map, denoted by $Q_{\text{WME}}(\lambda)$ with j th component $Q_{\text{WME},j}(\lambda)$, is given by

$$Q_{\text{WME},j}(\lambda) := \frac{1}{\sqrt{N_j}} \sum_{i=1}^{N_j} \frac{\mathcal{M}_j(\lambda) - d_{j,i}}{\sigma_j}. \quad (9)$$

The rationale for creating a WME map as opposed to the ME map is found by substitution of (6) into (9). Specifically, by this substitution followed by rationalizing the denominator, $Q_{\text{WME},j}(\lambda^\dagger)$ is identified as the sample average of N_j random draws from an i.i.d. $N(0, N_j)$ distribution. By assumption, all of the observed data are generated according to the true parameter vector given by λ^\dagger in (6). Subsequently, each component of $Q_{\text{WME}}(\lambda^\dagger)$ is a random draw from an $N(0, 1)$ distribution. Therefore, with this choice of data-defined QoI map, we specify π_{obs} as a $N(\mathbf{0}_{m \times 1}, \mathbf{I}_{m \times m})$ distribution. In other words, this method of incorporating data into the SIP produces an observed distribution which is stationary with respect to the number of measurements used. While this simplifies the algorithmic implementation of this approach (the observed distribution is fixed with respect to the data), as we show in Section 4.3, the variance in MUD estimates is a decreasing function of the number of observed data.

3.3. WME and the predictability assumption

Since λ^{MUD} requires maximizing π_{update} , in this section we explore the predictability assumption for the WME map under certain assumptions which is sufficient for guaranteeing the existence and uniqueness of π_{update} for a

given π_{init} and a QoI map. This also serves to motivate Section 4 exploring the linear Gaussian theory of existence and uniqueness of MUD points along with convergence of MUD estimates obtained by linear WME maps. In Section 5, we present an alternative to the WME map based on a principal component analysis (PCA) for situations where repeated data are not obtained on a fixed set of measurements.

We begin by assuming $\pi_{\text{init}} \propto \mathcal{N}(\lambda_0, \Sigma_{\text{init}})$ with non-degenerate Σ_{init} . For each $1 \leq j \leq m$, assume that the observable maps are linear maps of λ (affine maps require minor modifications to the results below). We denote these by M_j instead of \mathcal{M}_j to emphasize the linear assumption of these measurement maps. For notational convenience below, assume that M_j is written explicitly as a $1 \times p$ row vector and that the m -vectors form a linearly independent set. Then, it is possible to rewrite $Q_{\text{WME}}(\lambda)$ as

$$Q_{\text{WME}}(\lambda) = A(\mathbf{N})\lambda + b(\mathbf{N}), \quad (10)$$

where the j th component of $\mathbf{N} \in \mathbb{R}^m$ is given by N_j and the j th row of $A(\mathbf{N}) \in \mathbb{R}^{m \times p}$ is given by

$$\frac{1}{\sqrt{N_j}} \sum_{i=1}^{N_j} \frac{M_j}{\sigma_j} = \frac{\sqrt{N_j}}{\sigma_j} M_j, \quad (11)$$

and the bias vector, $\mathbf{b}(\mathbf{N}) \in \mathbb{R}^m$, is defined by the data, with j th component, denoted by \mathbf{b}_j , given by

$$\mathbf{b}_j(\mathbf{N}) = -\frac{1}{\sqrt{N_j}} \sum_{i=1}^{N_j} \frac{d_{j,i}}{\sigma_j}. \quad (12)$$

Since $A(\mathbf{N})\Sigma_{\text{init}}A(\mathbf{N})^\top$ defines a predicted covariance, and the observed covariance is the identity map, the predictability assumption is immediately satisfied if each diagonal component of the predicted covariance is significantly greater than 1.

The off-diagonal components of the predicted covariance (which are dictated in large part by the structure of the measurement operators M_j) dictate how much larger than 1 each diagonal component must be to ensure the predictability assumption holds. However, we demonstrate below that this will happen once a minimum number of data points N_{\min} are obtained for each measurement.

First, observe that the j th diagonal component of the predicted covariance matrix is given by the predicted variance associated with using the scalar-valued map $Q_{\text{WME},j}$. Then, the associated predicted variance is given by

$$\frac{N_j}{\sigma_j^2} M_j \Sigma_{\text{init}} M_j^\top \quad (13)$$

Since Σ_{init} is assumed to be non-degenerative and M_j is a non-trivial row vector, this predicted variance grows linearly with N_j . In other words, the j th diagonal component of the predicted covariance has the form $\beta_j N_j$ for some $\beta_j > 0$. Therefore, for each $1 \leq j \leq N$, there exists $N_{\min,j}$ such that $N_j \geq N_{\min,j}$ guarantees that the j th diagonal component is sufficiently large so that the smallest eigenvalue of the predicted covariance is larger than 1. Finally, we observe that the predicted covariance is inversely proportional to the measurement noise present in (13), which indicates that more data points N_j are required for measurement devices with greater noise.

4. Linear theory of MUD points

In this section, we assume linear (or affine) QoI maps with Gaussian distributions, which are often used in the UQ literature to provide a common framework for comparing methods and their solutions.

This section is structured into several subsections to help focus the interpretations and results. In Section 4.1, we present some useful details, notation, and terminology used for this comparison framework. To build intuition, we compare both the MUD and MAP points in Section 4.2 using a low-dimensional example. A unifying perspective is provided for affine maps in Section 4.3 along with derivations of closed form expressions for the MUD and MAP points in this comparison framework. These results are summarized in a theorem of existence and uniqueness of the MUD point in this comparison framework. This is followed by a corollary involving the convergence of MUD points obtained by WME maps and a brief analysis in Section 4.4 of the spectral properties of the updated covariance for a WME map as more data are incorporated into the map. We then provide some higher-dimensional performance comparisons of MUD, MAP, and least squares estimates in Section 4.5.

Table 2

Functionals to minimize to obtain λ that maximizes the updated PDF (bottom) and the Bayesian posterior PDF (top). Here, $T(\lambda)$ is the typical functional often associated with Tikhonov regularization, and the $J(\lambda)$ has an additional term subtracted from $T(\lambda)$ coming from the predicted density that serves as “un-regularization” in data-informed directions.

Tikhonov	$T(\lambda) := \ Q(\lambda) - \mathbf{y}\ _{\Sigma_{\text{obs}}^{-1}}^2 + \ \lambda - \lambda_0\ _{\Sigma_{\text{init}}^{-1}}^2$
Data-consistent	$J(\lambda) := T(\lambda) - \ Q(\lambda) - Q(\lambda_0)\ _{\Sigma_{\text{pred}}^{-1}}^2$

4.1. Problem formulation and assumptions

Let $\|\mathbf{x}\|_C^2 := (\mathbf{x}, \mathbf{x})_C = \mathbf{x}^T C \mathbf{x}$ denote the square of the induced norm associated with a positive-definite operator $C : \mathbb{R}^k \rightarrow \mathbb{R}^k$ and the usual (Euclidean) inner product. In what follows, the inverse covariances associated with non-degenerative multivariate Gaussian distributions will play the role of C .

Suppose that the initial and prior densities are both given by the same $\mathcal{N}(\lambda_0, \Sigma_{\text{init}})$ distribution. Additionally, suppose the map Q is linear and that the data-likelihood and observed densities are both given by the same $\mathcal{N}(\mathbf{y}, \Sigma_{\text{obs}})$ distribution.

The linearity of Q implies that $Q(\lambda) = A\lambda$ for some $A \in \mathbb{R}^{m \times p}$, and that the predicted density follows a $\mathcal{N}(Q(\lambda_0), \Sigma_{\text{pred}})$ distribution where

$$\Sigma_{\text{pred}} := A \Sigma_{\text{init}} A^\top. \quad (14)$$

While it is not technically necessary to ensure that the predictability assumption holds (i.e., that π_{update} is in fact a density) in order to formally define a MUD point using (5), it is useful when discussing certain theoretical results as shown in Section 4.3. Conceptually, the predictability assumption holds when the predicted variance is larger in all directions than the observed variance. Mathematically, this occurs when the smallest eigenvalue value of Σ_{pred} is larger than the largest eigenvalue value of Σ_{obs} . As discussed in Section 3, this condition holds once a sufficient amount of data are observed for a WME map.

When Σ_{obs} is non-degenerative (i.e., the smallest eigenvalue is positive), the predictability assumption can always be satisfied if $m \leq p$ and A is full rank by choosing Σ_{init} to have sufficiently large eigenvalues (i.e., if we choose initial variances to be sufficiently large). For clarity in the theoretical presentation of this section, we assume these conditions are met so that the predictability assumption holds and π_{update} defines a density. However, in Section 4.5, we still compute the formal MUD point for a high-dimensional example involving rank-deficient A to demonstrate the overall usefulness of the MUD point even in situations where π_{update} may fail to be an actual density.

With these assumptions, the parameters that maximize the posterior and updated densities are described as the arguments that minimize certain quadratic functionals. Table 2 presents a scaling of these functionals defined by the negative logarithm of the associated posterior and updated densities. Note that the functional, $T(\lambda)$, obtained from the posterior density is immediately identified as the typical functional used in Tikhonov regularization [14]. The data-mismatch term given by

$$\|Q(\lambda) - \mathbf{y}\|_{\Sigma_{\text{obs}}^{-1}}^2$$

comes from the data-likelihood/observed density whereas the regularization term defined by

$$\|\lambda - \lambda_0\|_{\Sigma_{\text{init}}^{-1}}^2$$

comes from the prior/initial density. We refer to this term as the Tikhonov regularization term.

By comparison, the functional, $J(\lambda)$, obtained from the updated density, is written as a modification of $T(\lambda)$ where the subtraction of

$$\|Q(\lambda) - Q(\lambda_0)\|_{\Sigma_{\text{pred}}^{-1}}^2$$

comes from the predicted density.

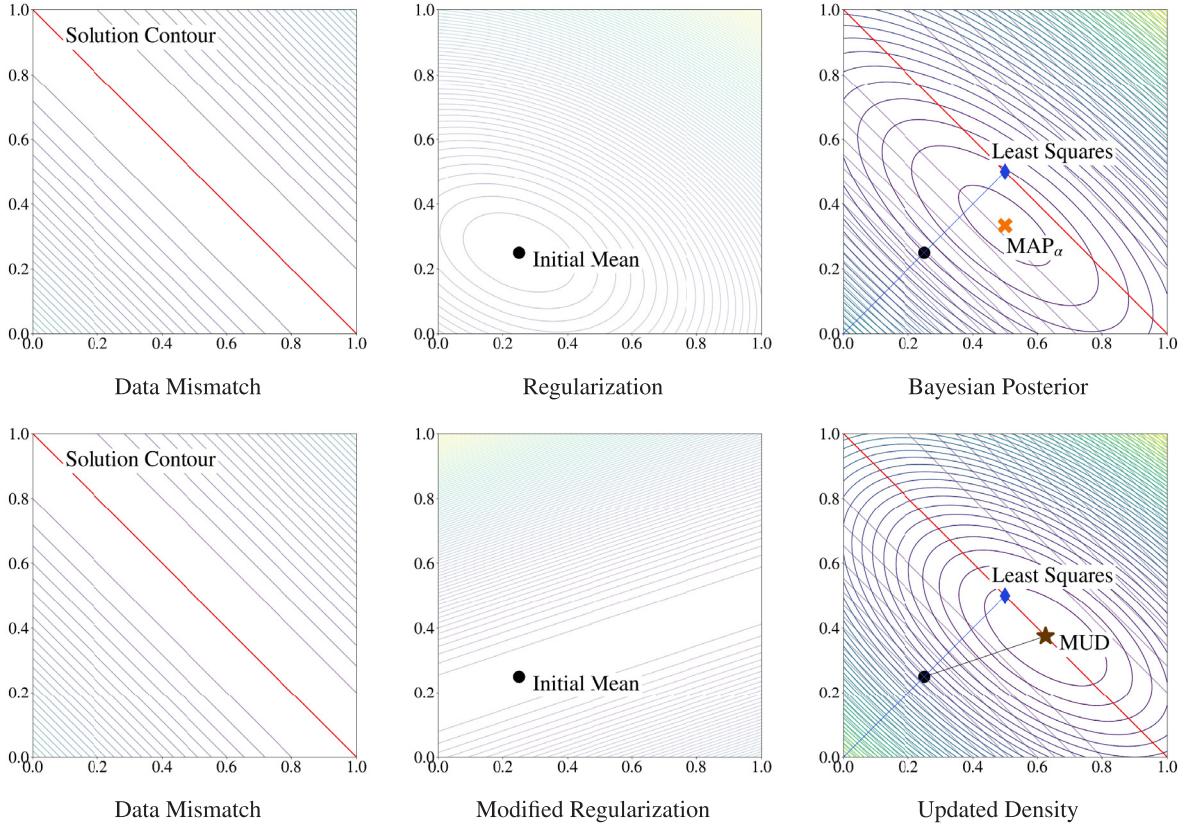


Fig. 3. Gaussian data mismatch over 2-D parameter space for a 2-to-1 linear map (left plots). Gaussian initial/prior lead to different regularization terms associated with updated/Bayesian PDFs (middle plots), which lead to different optimization functions (right plots) and parameter estimates that produce maximum PDF values for update/Bayesian PDF (red dot in right plots). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.2. A low-dimensional example

To build intuition around the fundamental differences of MAP and MUD points beyond what is discussed in Section 2, we consider an example where the linear QoI map is defined by $A = [1 \ 1]$, i.e., the parameter space is 2-dimensional while the data space is 1-dimensional.

In this example, the parameters in the initial and observed densities are given by

$$\lambda_0 = [0.25 \ 0.25]^\top, \Sigma_{\text{init}} = \begin{bmatrix} 1 & -0.25 \\ -0.25 & 0.5 \end{bmatrix}, \mathbf{y} = 1, \text{ and } \Sigma_{\text{obs}} = [0.25].$$

The top row of Fig. 3 shows contour plots in the parameter space for the data-mismatch term (left), Tikhonov regularization term (middle), as well as the functional $T(\lambda)$ (right). Conceptually, the regularization term is a radially symmetric function that penalizes parameters that are far away from the initial mean.

The bottom row of Fig. 3 shows contour plots in the parameter space for the data-mismatch term (left), modified regularization term (middle), as well as the functional $J(\lambda)$ (right). Here, we see that the modified regularization term only penalizes the movement of parameters in certain directions away from the initial parameter mean.

4.3. Existence and uniqueness of MUD points

Assume that the QoI map, Q , now takes the slightly more general form

$$Q(\lambda) = A\lambda + \mathbf{b} \quad (15)$$

where $\mathbf{b} \in \mathbb{R}^m$ may be viewed as a bias in the QoI map. The inclusion of this term makes it possible to draw conclusions involving the data-constructed QoI maps presented in Section 3. Using the same Gaussian distribution assumptions as described in Section 4.1, we again identify the MAP and MUD points as the values that minimize the functionals $T(\lambda)$ and $J(\lambda)$, respectively, shown in Table 2.

The posterior covariance is formally given by

$$\Sigma_{\text{post}} := (A^\top \Sigma_{\text{obs}}^{-1} A + \Sigma_{\text{init}}^{-1})^{-1}. \quad (16)$$

Applying the Woodbury matrix identity and (14), we rewrite the posterior covariance as

$$\Sigma_{\text{post}} = \Sigma_{\text{init}} - \Sigma_{\text{init}} A^\top [\Sigma_{\text{pred}} + \Sigma_{\text{obs}}]^{-1} A \Sigma_{\text{init}}, \quad (17)$$

which allows us to interpret Σ_{post} as a rank m correction (or update) of Σ_{init} . Note that $\Sigma_{\text{pred}} + \Sigma_{\text{obs}}$ is invertible because it is the sum of two symmetric positive definite matrices. With either version of Σ_{post} given above, we rewrite the closed form expression for the MAP point given in [14] as

$$\lambda^{\text{MAP}} = \lambda_0 + \Sigma_{\text{post}} A^\top \Sigma_{\text{obs}}^{-1} (\mathbf{y} - b - A\lambda_0). \quad (18)$$

We can arrive at a similar expression for the λ^{MUD} point by first deriving the updated covariance (see Appendix A for details), which is of the form

$$\Sigma_{\text{update}} = \Sigma_{\text{init}} - \Sigma_{\text{init}} A^\top \Sigma_{\text{pred}}^{-1} [\Sigma_{\text{pred}} - \Sigma_{\text{obs}}] \Sigma_{\text{pred}}^{-1} A \Sigma_{\text{init}}, \quad (19)$$

which leads to

$$\lambda^{\text{MUD}} = \lambda_0 + \Sigma_{\text{init}} A^\top \Sigma_{\text{pred}}^{-1} (\mathbf{y} - b - A\lambda_0). \quad (20)$$

Comparing (20) to (18), we see that the MUD point does not depend on the observed covariance whereas the MAP point does. Moreover, applying Q to (20) and substituting accordingly reveals that $Q(\lambda^{\text{MUD}}) = \mathbf{y}$.

Overall, this motivates the MUD point as an *alternative parameter estimate* with predictive accuracy and properties directly correlated to the relationship between \mathbf{y} and the true signal for which noisy data are generated. We summarize the above results in the following theorem stating the existence and uniqueness of a MUD point for the linear Gaussian case.

Theorem 4.1. Suppose $Q(\lambda) = A\lambda + b$ for some full rank $A \in \mathbb{R}^{m \times p}$ with $m \leq p$ and $b \in \mathbb{R}^m$. If $\pi_{\text{init}} \sim N(\lambda_0, \Sigma_{\text{init}})$, $\pi_{\text{obs}} \sim N(\mathbf{y}, \Sigma_{\text{obs}})$, and the predictability assumption holds, then

- (a) There exists a unique parameter, denoted by λ^{MUD} , that maximizes π_{update} .
- (b) $Q(\lambda^{\text{MUD}}) = \mathbf{y}$.
- (c) If $d = p$, λ^{MUD} is given by A^{-1} . If $m < p$, λ^{MUD} is given by (20) and the covariance associated with this point is given by (A.8).

Recalling the discussion of Section 3.3, the following result is an immediate consequence of Theorem 4.1 and Eq. (20),

Corollary 1. If $\pi_{\text{init}} \sim N(\lambda_0, \Sigma_{\text{init}})$ and data are obtained for m linearly independent measurements on Λ with an additive noise model with i.i.d. Gaussian noise for each measurement, then

- (a) There exists a minimum number of data points obtained for each of the measurements such that there exists a unique λ^{MUD} and $Q_{\text{WME}}(\lambda^{\text{MUD}}) = 0$.
- (b) The variance in the λ^{MUD} estimate in m directions of the parameter space decreases at a rate proportional to the number of data points used for each of the measurements, and is inversely proportional to the magnitude of the Gaussian noise, a relationship expressed in Eq. (13).

4.4. Spectral properties of the updated covariance for a WME map

We illustrate the result in Corollary 1 with an example that highlights how the number of distinct measurements used to form Q impacts the spectrum of the updated covariance. Consider a randomly generated linear operator M of dimension 5×20 with components sampled from a $N(0, 1)$ distribution. This M defines five randomly

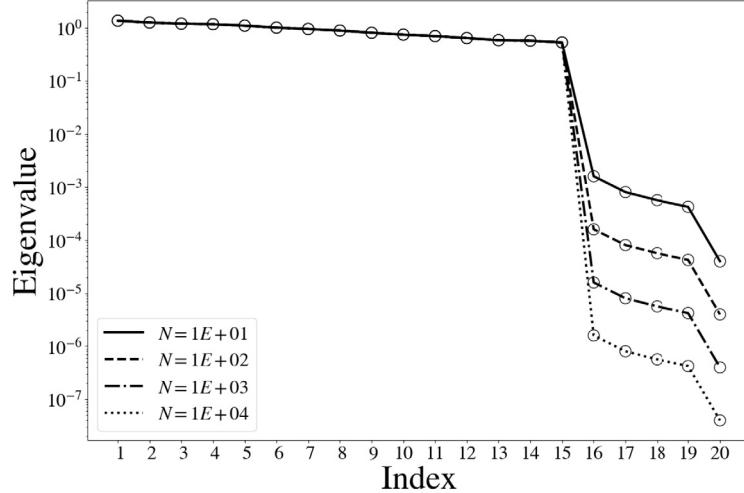


Fig. 4. Ranked eigenvalues of the updated covariance are shown for $N = 10, 100, 1000$, and $10\,000$ measurements, plotted by index as a line for each N . Fifteen of them show no sensitivity to N , and all are of $\mathcal{O}(1)$. With more measurements, the five eigenvalues (corresponding to the orthogonal null-space of the QoI operator) go to zero, losing approximately an order of magnitude for each increase in N .

constructed, but geometrically distinct (i.e., linearly independent), measurements. We then construct the QoI by simulating repeated measurements $N = 10, 100, 1000, 10\,000$ (for each measurement) polluted by additive Gaussian noise with $\sigma = 0.1$. We then compute the updated covariance using the analytical expression (A.8) and perform a singular value decomposition to obtain the 20 sorted eigenvalues, which we plot for each N in Fig. 4.

Observe that five of the twenty eigenvalues are several orders of magnitude smaller than the rest, which corresponds to the output dimension of Q . These correspond to the five directions informed by Q given by the associated eigenvectors. Furthermore, as observed in Fig. 4, the gap between the uninformative and informative directions is directly proportional to N ; for each ten-fold increase in measurements, there is a reduction in eigenvalues by an order of magnitude. The eigenvalues associated with the fifteen uninformed directions remain unaffected by N , appearing as a solid line in the plot.

4.5. Higher-dimensional linear Gaussian examples

We first describe the relationship of MUD, MAP, and least squares estimates to the set-valued inverses of Q in order to establish a conceptual framework for interpreting the numerical results that follow. Throughout this discussion, we refer extensively to Fig. 5, which builds upon Fig. 3, to make these ideas more clear. In Fig. 5, a prior/initial covariance is chosen such that the MAP estimate is approximately halfway between the initial mean and the contour of Q on which both the MUD and least squares estimates exist. A reference parameter $\lambda^\dagger = (0.75, 0.25)$, labeled as “Truth” is also shown on the contour. Note that the MAP estimate is labeled with a subscript α . In these examples, α is a scalar multiple of the prior/initial covariance. Thus, α is a hyper-parameter that determines, in a sense, the “strength of prior beliefs” as quantified by the eigenvalues in the prior/initial covariance.

As illustrated in Fig. 5, the MUD point always exists on the “solution contour” defined by $Q^{-1}(\mathbf{y})$ regardless of the initial covariance. This is in fact guaranteed by Theorem 4.1(b). We use the term “solution contour” in this case because if A is a p -to- m full rank linear map with $m < p$, then $Q^{-1}(\mathbf{y})$ exists as a $(p-m)$ -dimensional linear hyperplane in A . This means that the MUD point retains the “predictive precision” of a least squares solution to the inverse problem (i.e., a parameter that minimizes the data-mismatch term $\|\mathbf{y} - Q(\lambda)\|_{\Sigma_{\text{obs}}^{-1}}^2$) while incorporating the flexibility of prior beliefs in directions not informed by the QoI. This is in fact expected given the roles of π_{init} listed in Section 2.1.1 coupled with Theorem 4.1. For under-determined or ill-conditioned problems, this suggests that “good” prior beliefs may be used to produce a MUD point that is more accurate (i.e., closer to truth) than a least squares solution as illustrated in Fig. 5. This is further explored in the examples below involving high-dimensional linear maps.

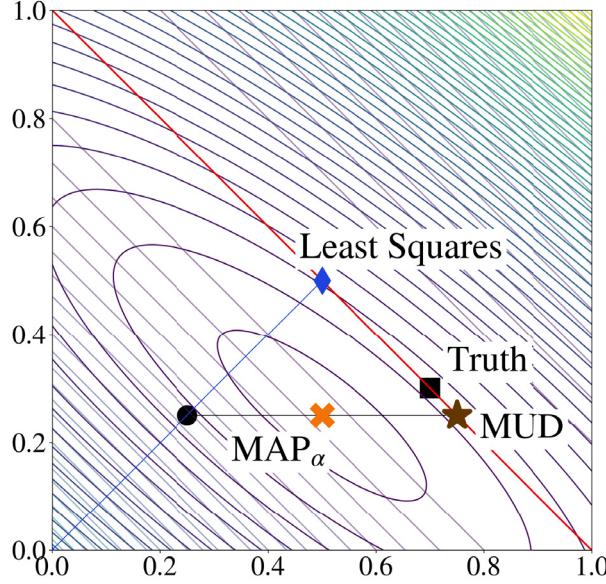


Fig. 5. A comparison of Least Squares, MAP, and MUD estimates relative to a true reference parameter along the solution contour defined by $Q^{-1}(\mathbf{y})$. Pictured are the contour lines involved in Tikhonov regularization defined in Table 2. While the MUD point will always lie on the solution contour, the MAP solution can be influenced by the strength of the regularization parameter α so that it lies along the line connecting the initial mean and the MUD point.

By contrast, the MAP point exists on a line connecting the initial mean, λ_0 , and the generalized contour defined by $Q^{-1}(\mathbf{y})$ as illustrated in Fig. 5. The direction of this line is discovered by substituting (17) into (18), which reveals that the direction of the line is orthogonal to the nullspace of the image of A under Σ_{init} ; i.e., $\mathcal{N}(\Sigma_{\text{init}} A)^\perp$. In fact, this line intersects the generalized contour defined by $Q^{-1}(\mathbf{y})$ precisely at the MUD point. If one parameterizes the line between λ_0 and λ^{MUD} , then one can also identify λ^{MAP} as a convex sum of these two points. The weights of this convex sum, which determine the position of the MAP point on this line, are determined by the “precision of data” (i.e., on Σ_{obs}) and the “strength of prior beliefs” (i.e., on Σ_{init}). This is apparent both in Fig. 5 and also by comparing the location of the MAP point in the top right plot of Fig. 3 to the line segment connecting the initial mean to the MUD point in the bottom right plot of this same figure. The impact of this is also explored in the following examples where we tune an α hyper-parameter appearing as a multiplicative factor in Σ_{init} .

4.5.1. Impact of output dimension

We consider QoI defined by $A\lambda + \mathbf{b}$ for $A \in \mathbb{R}^{m \times 100}$ where $m = 1, 2, \dots, 100$ to demonstrate how the various estimates of a true parameter λ^\dagger are impacted by the number of available QoI. To generate the matrices, we first generate 10,000 i.i.d. random numbers from a $N(0, 1)$ distribution that are arranged into a reference $\mathbb{R}^{100 \times 100}$ matrix.

The same distribution is also used for generating the components of the 100-dimensional vectors defining a reference bias vector \mathbf{b} and reference parameter λ^\dagger . A multivariate Gaussian distribution is used for the initial density, with zero mean and Σ_{init} chosen as a diagonal covariance with random entries drawn from $U[0.5, 1.5]$ and sorted in descending order. The prior density is also a zero mean multivariate Gaussian distribution. However, to demonstrate the impact of the strength of prior beliefs on the MAP point, we choose the prior covariance to be $\alpha \Sigma_{\text{init}}$ for $\alpha = 0.001, 0.01, 0.1$, and 10. Here, smaller values of α correspond to a “stronger” belief in the prior since the prior density becomes more concentrated near the prior mean.

To study the impact of dimension on the MUD, MAP, and least squares estimates, we solve a sequence of inverse problems by truncating the rows of the reference matrix and bias vector. The results are summarized in the left plot of Fig. 6, which shows convergence towards λ^\dagger for all the problems considered with the exception of several MAP estimates corresponding to strongly-held beliefs in the prior.

We note that the MUD solution is the same for all choices of α and corresponds to the same level of accuracy that the MAP point achieves when α is chosen to be large. In other words, the MUD point is not impacted by

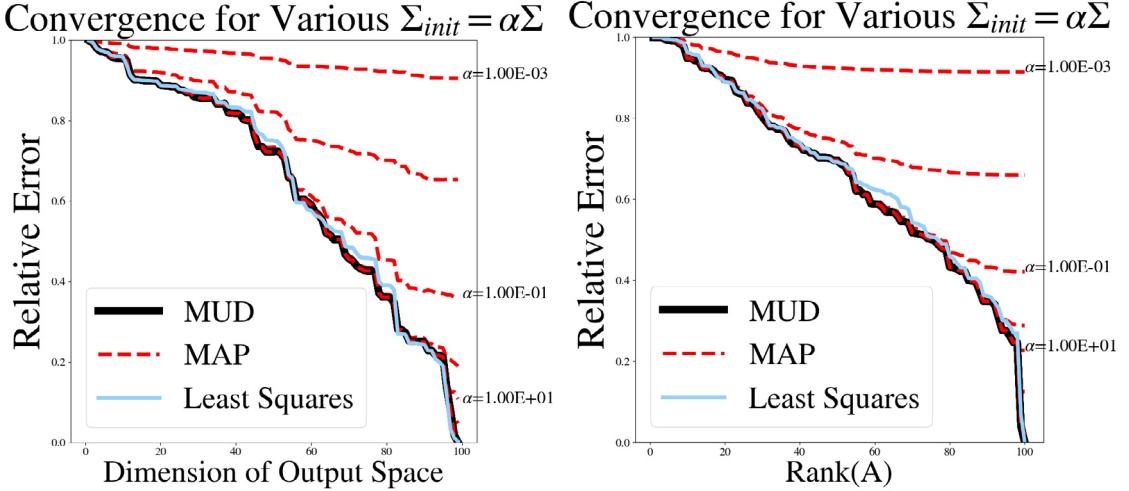


Fig. 6. Relative errors between λ^\dagger and (i) the least squares solution obtained through numpy’s linalg.pinv module, (ii) the closed-form solution for the MUD point given in Eq (20), and (iii) the MAP point. (Left): Error for increasing dimensions of D for A taken to be a Gaussian Random Map. (Right): Error for increasing row-rank of A , generated with Gaussian vectors and a SVD.

a scaling of the initial covariance, providing *consistent* solutions which demonstrate levels of accuracy that MAP points only exhibit for larger values of scaling factors.

Of interest is also that the MUD point can sometimes out-perform the least squares estimate while generally achieving similar levels of accuracy. This suggests that the MUD point has several favorable qualities. Not only is it robust to the specification of prior assumptions, but it manages to offer the flexibility of incorporating good prior specifications without paying the additional cost of hyper-parameter optimization (i.e., choosing an appropriate α) that would be required for the MAP estimates to achieve comparable results.

While omitted in the interest of space, if Σ_{init} is chosen as αI , where I denotes the identity matrix of appropriate dimensions, then the MUD point will always agree with the least squares estimate. Taking these results together, this implies that only a good “relative spatial structure” of prior beliefs is required to improve the MUD point’s accuracy over both MAP and least squares estimates.

4.5.2. Impact of rank: One hundred (deficient) $\mathbb{R}^{100 \times 100}$ matrices

Here, we investigate whether the previous dimension-dependent example results also apply to matrices A which are of a fixed dimension but possess varying rank. This is of interest in applications where many QoI are available to construct an operator but a great deal of redundancy may be present in the data collected, and feature-engineering new quantities is somehow prohibitive (perhaps due to gradient estimation).

The rank of A corresponds to the number of unique directions of information present in the operator, i.e., how many directions in the parameter space are informed by the QoI map. The operators in the previous example were all full rank, so the dimension of each map also corresponded to the rank of A . When A is rank-deficient, Σ_{pred} is non-invertible, so we must modify the form of (20) to substitute a pseudo-inverse for the predicted covariance.

In this example, the dimension of the data space remains fixed at $m = 100$ across all experiments. However, we sequentially increase the row-rank of A from $r = 1, \dots, 100$. To control the rank of A , we first construct a reference $\mathbb{R}^{100 \times 100}$ matrix as in the previous example using 10,000 i.i.d. $N(0, 1)$ random numbers. We then compute a singular value decomposition of this reference matrix of the form USV^\top and construct 100 rank-1 matrices of the form $A_i = \mathbf{u}_i s_i \mathbf{v}_i^\top$ for $i = 1, \dots, 100$ where \mathbf{u}_i and \mathbf{v}_i denote the i th columns of U and V , respectively and s_i denotes the i th singular value. Then, we analyze the impact of $A = \sum_i^r A_i$ for $r = 1, \dots, 100$. Aside from the differing construction of A , the rest of the choices involved in the experiment (λ^\dagger , the reference bias vector, and the distributions involved) is identical to the previous example.

In the right plot of Fig. 6, we again find that the MUD point is generally as accurate as the least squares estimate, but incorporates an initial description of uncertainty, which may allow it to outperform the least squares estimate. Also, we again see that the MAP estimates are impacted by the strength of prior beliefs.

5. Data-constructed QoI maps II: Data clouds and principal component analysis

In Section 3 the Q_{WME} is introduced as a way to incorporate data from repeated measurements and reduce the variance in the MUD point estimates. Here, we address the case of potentially non-repeated measurements taken from a system that varies over a temporal and/or spatial domain.

We utilize the ubiquitous Principal Component Analysis (PCA). Introduced first by [33,34], PCA is a way to reduce dimensionality of a large set of correlated data by transforming the data into a new set of variables known as principal components, which are uncorrelated and when ordered, contain the maximum amount of variation in the data per component. For more comprehensive reviews of the PCA method, we refer the interested reader to [35,36]. The PCA transformation can be written in the form

$$Y = XP, \quad X, Y \in \mathbb{R}^{s \times n}, \quad P \in \mathbb{R}^{n \times n}. \quad (21)$$

Here, the data matrix, X , contains s samples of n data points each. Using a change of basis matrix P , X is transformed to a new matrix Y . We summarize below some of the main conclusions of the PCA transformation.

- P defines a linear transformation to an orthonormal basis, given by the columns $p^{(\ell)} \in \mathbb{R}^n$, $1 \leq \ell \leq n$ of the matrix P that define the principal components of X .
- P diagonalizes the covariance matrix $C_X = XX^T$.
- The matrix P and the popular Singular Value Decomposition (SVD) transformation are intimately linked. Namely, for the SVD decomposition of $X = U\Sigma V^T$, the columns of V are the principal components of X .

Returning to the scenario proposed in Section 3, suppose the \mathcal{M}_j measurement devices now collect data over space and time, with each taking N_j measurements as before. Define any ordering of these $n = \sum_{j=1}^m N_j$ data points $\{z_i\}_{i=1}^n$ (the ordering can be arbitrary as it does not impact the results obtained via PCA). Similar to Eq. (6), we now let d_i equal the i th measurement datum, which is assumed polluted by i.i.d. additive Gaussian errors from a $\mathcal{N}(0, \sigma_i)$ distribution. Furthermore, let $\mathcal{M}_{k,i} = \mathcal{M}(\lambda_k; z_i)$ be the i th measurement for the k th simulated sample. Assuming s samples are collected, the matrix $X \in \mathbb{R}^{s \times n}$ of Z-scored residuals for a sample set is now defined component-wise as

$$X_{k,i} = \frac{\mathcal{M}_{k,i} - d_i}{\sigma_i}. \quad (22)$$

We define Q_{PCA} component-wise as

$$(Q_{PCA})_\ell(\lambda_k) = \sum_{i=1}^n p_i^{(\ell)} \frac{\mathcal{M}(\lambda_k; z_i) - d_i}{\sigma_i}, \quad 1 \leq \ell \leq n, \quad (23)$$

where $p^{(\ell)}$ is the ℓ th principle component of X .

Similar to Q_{WME} , the Q_{PCA} map is computing a weighted average of residuals. However, the weighting of these residuals is according to the new set of basis vectors defined by the principal components of X . Since each component of Q_{PCA} is still a normalized (in the 2-norm) combination of the Z-scored residuals, the observed distribution π_{obs} is still given by an $\mathcal{N}(0, 1)$ distribution as with the Q_{WME} map. Note that the choice of ordering of the measurements in the data matrix X is irrelevant, since the PCA does not depend on column order. Finally, observe we can define up to n components to our Q_{PCA} map. However, in practice, we only take up to the first m components that capture a user-specified percentage of variance in the original data set X . In fact, if the data are sensitive to all parameters present in our inverse problem, we expect the number of components, m , to be equal to the dimension of our parameter space, p . However, we see in the following examples situations where this may not be the case. In these examples, we turn to the diagnostic $\mathbb{E}(r)$ as an important measure of the quality of the updated density and thus the reliability of λ^{MUD} .

6. Spatial and temporal data examples

We now use MUD points as parameter estimates for PCA constructed maps using simulated noisy temporal and spatial data associated with solutions to differential equations. The previously derived closed form expressions for λ^{MUD} do not apply in these examples since the measurement maps from parameters to data are nonlinear. Instead, in each example we use a fixed set of i.i.d. samples drawn from the initial density to approximate the updated density and subsequently choose the sample that maximizes this approximation. In each case, we monitor $\mathbb{E}(r)$ to check for the validity of the predictability assumption.

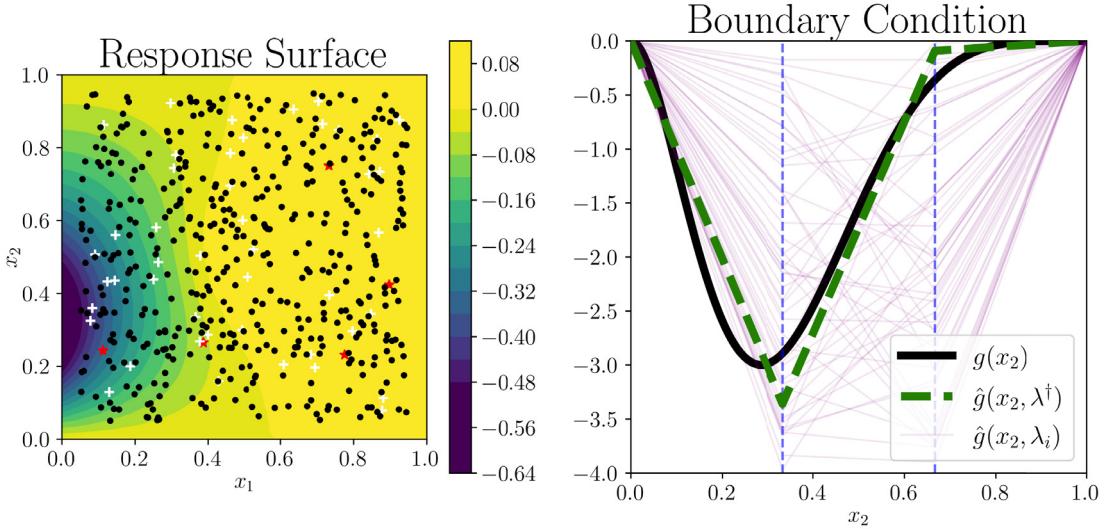


Fig. 7. (Left) Noiseless response surface to solving Eq. (24) using true $g(x_2)$ with locations of sensors used for collecting data. First five sensors used in red in stars, next set of 45 in white crosses, and the rest (450) are in black circles. (Right) Left boundary condition. True $g(x_2)$ (black) along with piecewise-linear spline approximations $\hat{g}(x_2, \lambda)$. The closest (in the 2-norm) possible spline to the “true” boundary data is indicated by λ^\dagger (dotted green). 50 different λ_i samples from the initial uniform distribution are shown in faded purple. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

6.1. Spatial data example: Poisson’s equation with uncertain boundary condition

We first consider the aggregation of data over a spatial domain. In this problem, the uncertain model parameter is described by an uncertain parametrized function defining the boundary data to a stationary PDE, given by the Poisson problem:

$$\begin{cases} -\nabla \cdot \nabla u = f(x), & \text{on } x \in \Omega, \\ u = 0, & \text{on } \Gamma_T \cup \Gamma_B, \\ \frac{\partial u}{\partial \mathbf{n}} = g(x_2), & \text{on } \Gamma_L, \\ \frac{\partial u}{\partial \mathbf{n}} = 0, & \text{on } \Gamma_R, \end{cases} \quad (24)$$

where $x = (x_1, x_2) \in \Omega = (0, 1)^2$ is the spatial domain; Γ_T , Γ_B , Γ_L , and Γ_R , denote the top, bottom, left, and right boundaries of this domain, respectively, and $\frac{\partial u}{\partial \mathbf{n}}$ denotes the usual outward normal derivative. The forcing function f is taken to be $10 \exp(\|x - 0.5\|^2 / 0.02)$.

The goal is to use noisy data to estimate the boundary data $g(x_2)$. To generate the noisy data, a “true” $g(x_2) \propto x_2^2(x_2 - 1)^5$ is constructed, with a constant of proportionality chosen to produce a minimum of -3 at $x_2 = \frac{2}{7}$. Eq. (24) is solved using piecewise-linear finite elements on a triangulation of a 36×36 mesh. Random noise is then added to every degree of freedom of this reference solution, and the spatial data are subsequently computed from a fixed set of 5, 50, and 500 randomly placed sensors in the subdomain $(0.05, 0.95)^2 \subset \Omega$ (Fig. 7, left).

To construct a finite-dimensional parameter space describing the initial uncertainty of $g(x_2)$, piecewise-linear continuous splines, $\hat{g}(x_2)$, are used to approximate $g(x_2)$. The locations of the first and last knots are fixed at the endpoints of the boundary with values assumed to be 0, and furthermore it is assumed that g is non-positive and bounded below by -4 . Thus, the uncertainty is described by the values of the splines at the interior knot points chosen as the equispaced points $1/3$ and $2/3$. This defines a finite-dimensional parameter space described by $\Lambda = [-4, 0]^2$. We generate 1000 samples from an initial uniform density on Λ to (1) generate random spline functions and compute the (noise-free) data from solutions associated with these splines; and (2) estimate the push-forward and updated densities along with the MUD estimate of the boundary data, $\hat{g}(x_2, \lambda^{MUD})$. Note that since we approximate the “true” boundary data with linear splines, the MUD estimate will *never* exactly equal the noiseless

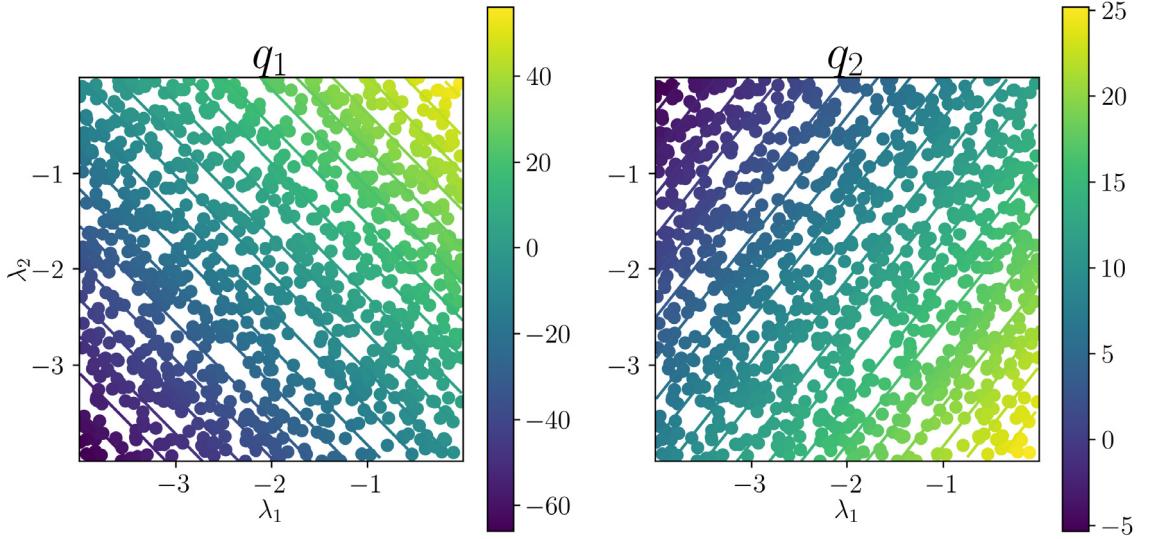


Fig. 8. Quantities of interest learned by the Q_{PCA} map when using all 500 sensors for the first (left) and second (right) principal component. Note how the contours show that the Q_{PCA} map builds orthogonal components that lead to a well posed inversion problem.

boundary data. To measure the accuracy of a given estimate then, we compare it with the closest (in the 2-norm) possible spline to the “true” boundary data. We refer to the values of λ that produce this value as λ^\dagger .

For each sample, the Q_{PCA} map in Eq. (23) is constructed using two principal components. The value of the learned QoI map for each parameter sample is seen in Fig. 8. Note the orthogonal contour structure of the two components demonstrates how the Q_{PCA} map aggregates data from a high-dimensional output space into an essentially linear bijective map with a Jacobian that is well-conditioned for inversion.

The plots in Fig. 9 show how the quality of MUD estimates is improved as more sensor data are included. Note how for only the first five sensors (red stars in Fig. 7, left), the Q_{PCA} map does not produce a good estimate for either parameter. In fact, practically no update is made in λ_2 , while a small update is made in λ_1 , but in the incorrect direction. The location of the first five sensors explains this behavior as there are only two sensors near the location of the knot controlled by λ_1 on the boundary, while no sensors are near the location on the boundary where the λ_2 knot location is. It is important to stress the importance of the diagnostic $\mathbb{E}(r)$ here. If we had no knowledge of the “true” boundary data, and no λ^\dagger was readily available for comparison, the results of using five sensors may lead us to believe that a meaningful update was made for λ_1 . However, the diagnostic of $\mathbb{E}(r) = 0.3811$ tells a different story, indicating that the predictability assumption is likely being violated in the construction of π_{update} . Once $N = 50$ sensors are used (red stars and white crosses in Fig. 7, left), the diagnostic value jumps to within an acceptable range of $\mathbb{E}(r) = 0.9762 \approx 1$, and the Q_{PCA} map is able to update initial beliefs in both parameters well. Using $N = 500$ sensors leads to a reduction in the variance of the MUD point estimate as distributions for both parameter values peak around a sample that is very close to λ^\dagger .

6.2. Temporal data example: ADCIRC with uncertain wind drag

For the second example, the aggregation of data over time into a Q_{PCA} map is considered. In this problem, the uncertain parameter relates to a coefficient that determines how unresolvable dynamics of a PDE system are modeled. We begin with the Shallow Water Equations (SWE), which are a depth-averaged approximation to the Navier-Stokes equations commonly used in coastal circulation and flooding modeling to predict peak storm-surge due to extreme weather events [37]. They can be expressed as

$$\begin{cases} \frac{\partial \zeta}{\partial t} + \nabla \cdot (\mathbf{U}H) = 0, \\ \frac{\partial \mathbf{U}}{\partial t} + \mathbf{U} \cdot \nabla \mathbf{U} + f \mathbf{k} \times \mathbf{U} = -\nabla \left[\frac{p_s}{\rho_0} + g\zeta \right] + \frac{\tau_s - \tau_b}{\rho_0 H}, \end{cases} \quad (25)$$

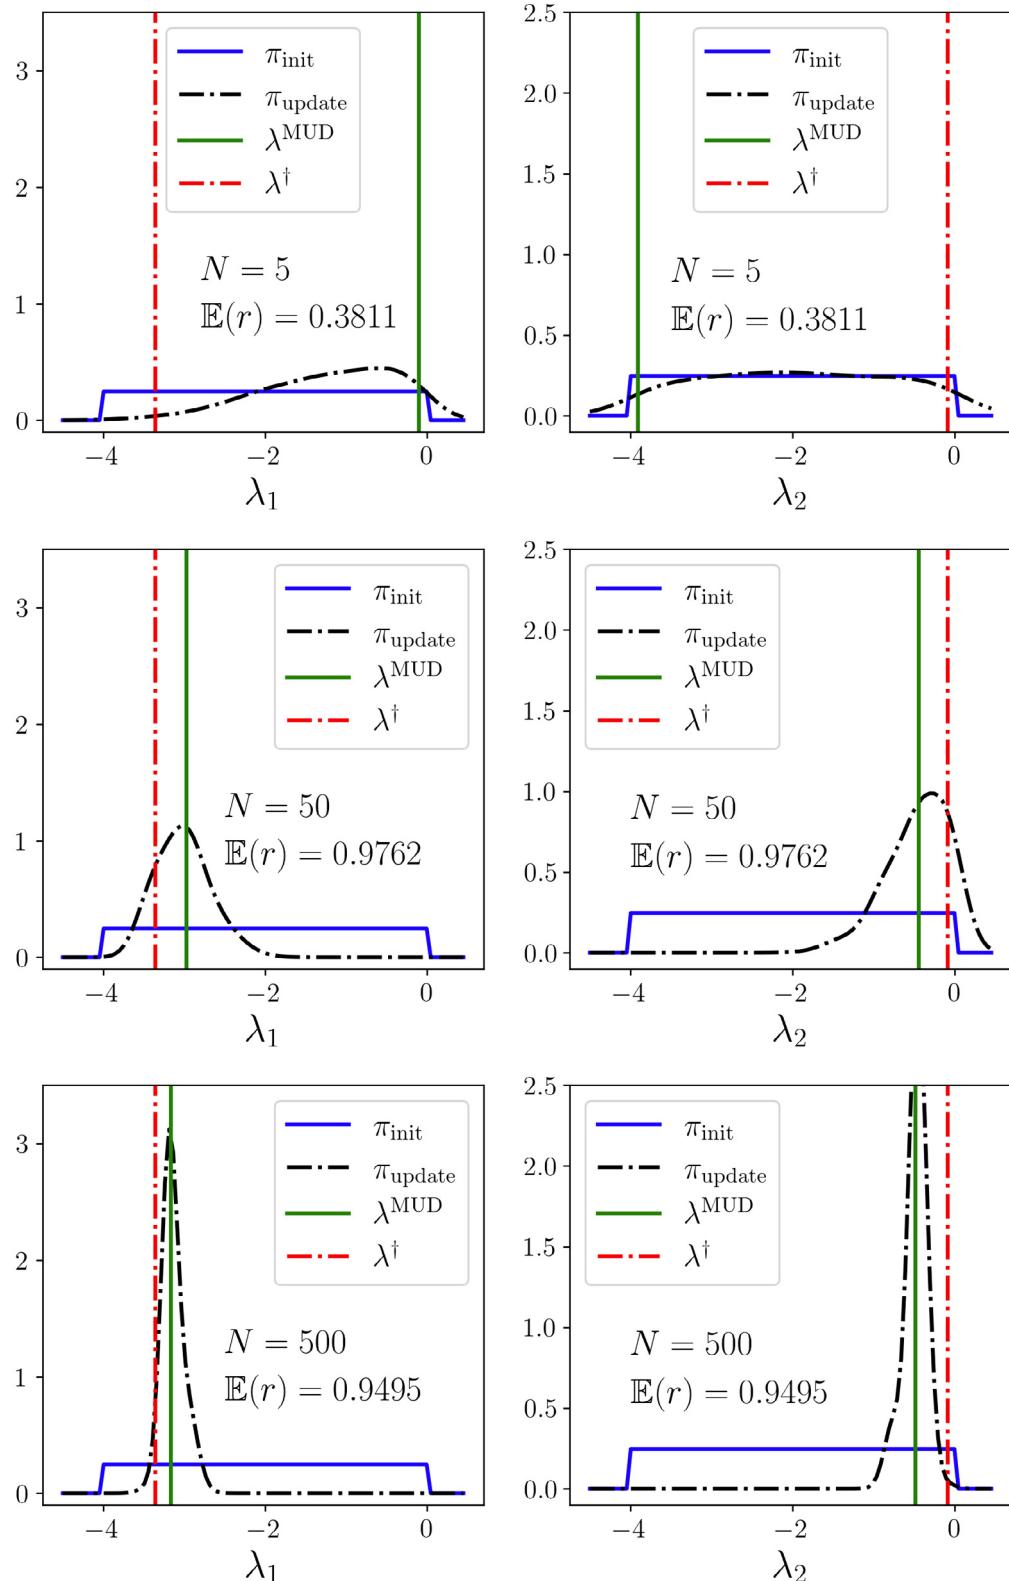


Fig. 9. Initial (blue) and updated (dotted-dashed, black) distributions for λ_1 (left plots) and λ_2 (right plots) using $N = 5$ (top), 50 (middle), 500 (bottom) sensors. Expectation value of the ratio of the observed to the predicted, which should be ≈ 1 , is shown as well per case. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

for unknown free surface elevation, $\zeta = \zeta(x, y, t)$ with respect to mean sea level and depth-average velocity (averaged over the height of the water column), $\mathbf{U} = \mathbf{U}(x, y, t)$. In Eq. (25), H is the height of the water column, f the Coriolis parameter, p_s the atmospheric pressure, ρ_0 the reference density of water, g the gravitational constant, τ_s the surface stress, and τ_b the bottom stress. The surface stress,

$$\tau_s = \rho_s C_d \mathbf{u} \|\mathbf{u}\|, \quad (26)$$

includes the wind drag C_d , which is an effective (homogenized) parameter governing the transfer of momentum from winds to the water column, one of the primary drivers of storm-surge. There are various forms for C_d that depend on the physical properties of the system being modeled (e.g., the type of storm, presence of ice, etc.). In this work, a popular generalization of Garratt's formula for C_d is used [38], whereby C_d increases linearly with wind speed $u := \|\mathbf{u}\|$ according to the uncertain parameter λ_1 and is "cut off" when exceeding a threshold λ_2 :

$$C_d = \min [10^{-3}(.75 + \lambda_1 u), \lambda_2]. \quad (27)$$

Using time series of recorded water surface elevations, the goal is to determine the values of the wind drag parameters (λ_1, λ_2) .

The Advanced CIRCulation (ADCIRC) coastal ocean model is used to solve the SWEs (25) [39,40]. ADCIRC uses a finite-element model of the SWE in which the Generalized Wave Continuity Equation [41] is discretized in space using piecewise-linear elements on unstructured (triangular) grids. The model is used for coastal engineering applications such as hurricane storm surge forecasting [42], hindcasting [43–45] and uncertainty quantification [46–48].

In this study, ADCIRC is configured to run using a test mesh based on the Shinnecock Inlet on the Outer Barrier of Long Island, NY, USA. External forcing for the model is given by tidal forcing reconstructed from the TPXO9.1 harmonic tidal constituents [49] using OceanMesh2D [50], constant air pressure of 1013 millibars, and free surface stress from winds computed from a 0.25 deg hourly CFSv2 10-m wind fields [51] for a period of 16 days (29 December 2017–31 January 2018). Winds are modified for the purposes of the numerical experiment to simulate a more extreme (Category 4) event, with winds scaled radially down to zero from the point of interest, i.e. the center of the inlet (see Fig. 10).

To frame the PIP, first we assume that the uncertain parameters (λ_1, λ_2) lie within $\pm 50\%$ of commonly used default values of $(0.067, 0.0025)$ [52]. This defines a finite-dimensional parameter space

$$\Lambda = [0.0335, 0.1105] \times [0.00125, 0.00375] \subset \mathbb{R}^2. \quad (28)$$

1000 samples are generated from a uniform distribution over Λ and pushed through our forward model, ADCIRC (see Appendix B for more details on how to obtain simulation data). Water elevation at an artificial recording station inside the inlet was recorded over a period of 14 days (1 January 2018–14 January 2018) at 3 h intervals for each sample. Since no real station data are available over the test mesh domain, we create observations by picking (and removing) a sample closest to the default value of $(0.067, 0.0025)$, and populate each measurement of this sample with i.i.d. $N(0, \sigma^2)$ noise, using $\sigma = 0.05$. The different time-series for the water elevation at the artificial recording station are shown in Fig. 11.

Three different time windows of data are used to construct Q_{PCA} . In each example, the diagnostic $\mathbb{E}(r)$ is used to compare using one ($\mathbb{E}(r_1)$) vs. two ($\mathbb{E}(r_2)$) principal components. First in T_1 (Fig. 12), we choose a time window with low winds and little variation in time-series values consisting of 119 data points. Note the Q_{PCA} map with one component (top row) leads to much better $\mathbb{E}(r)$ (0.9879 vs 0.4204). This makes sense because during a period of low winds the cut-off parameter λ_2 will not come affect the resulting dynamics of the system for most samples. Subsequently, attempting to update both parameters via a two-component QoI map leads to a poor $\mathbb{E}(r)$ indicating a violation of the predictability assumption and that the updated distribution should not be trusted. We see similar results in T_2 (Fig. 13), a time window of 53 data points characterized by high winds and large variations in time-series values, where a one-component Q_{PCA} map also performs better in terms of the diagnostic $\mathbb{E}(r)$ (1.039 vs 4.42). However in T_2 , we observe that the data are now sensitive to λ_2 and not λ_1 due to the fact the cut-off parameter dominates impacts the dynamics in a period of high winds. Finally, in T_3 (Fig. 14), a larger time window with both high and low winds is used consisting of 959 data points. Here, enough data are collected that are sensitive to perturbations in both parameters so that the two-component Q_{PCA} map produces a reasonable diagnostic value $\mathbb{E}(r_2) = 0.9085$. We note that the update to the initial distributions for both parameters is significant in this case, with the two-component Q_{PCA} leading to accurate MUD estimates of the true parameter values.

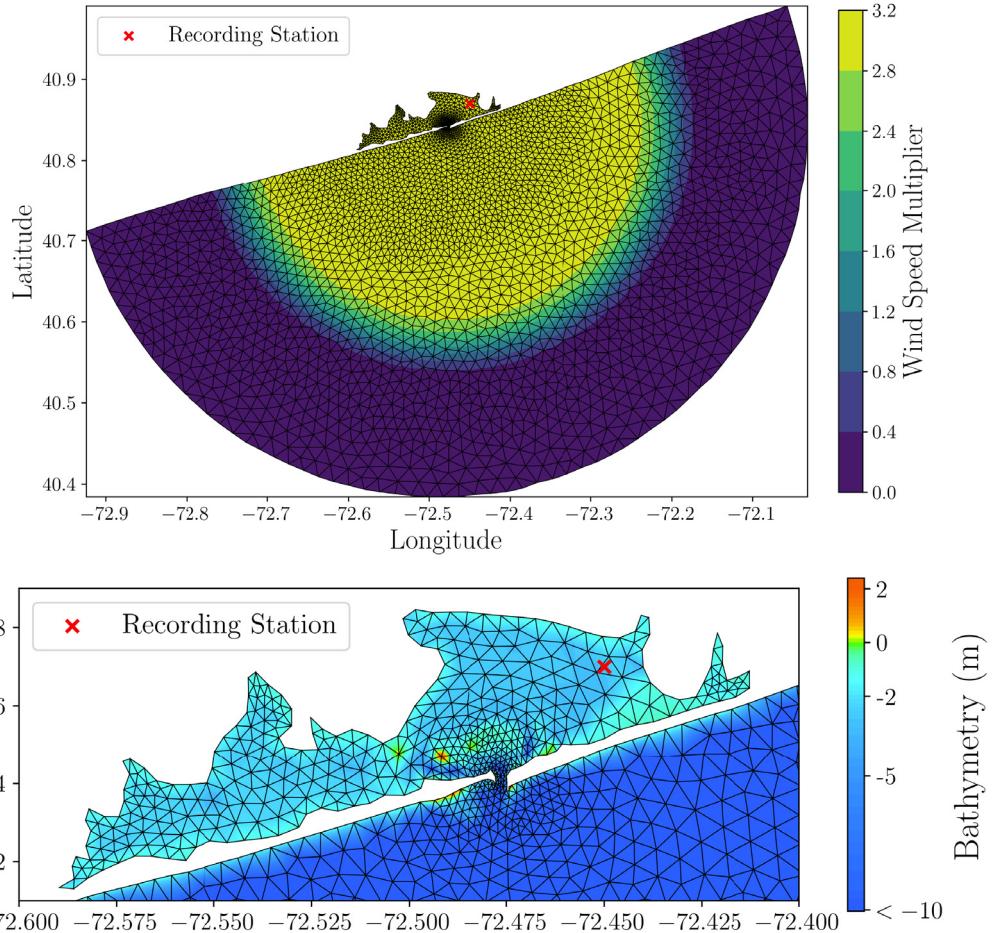


Fig. 10. (Top) Shinneecock Inlet Mesh containing 5780 triangular elements. Contours show value of wind multiplier applied to scale winds up artificially near the inlet, tapering them off to zero at the boundaries.(Bottom) Bathymetry of inlet.

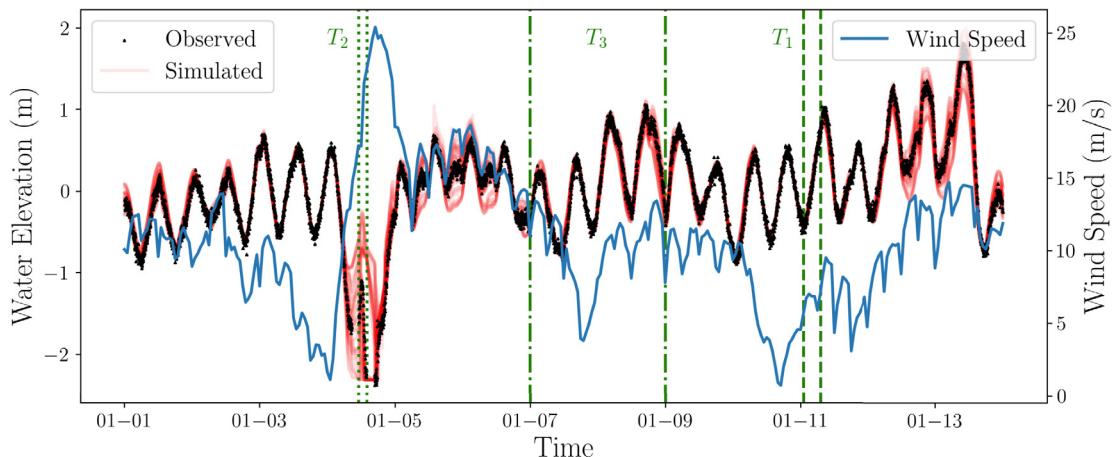


Fig. 11. Time series of water elevations (left axis) for “observed” data (black triangles) and simulated data (faded red lines) along with wind forcing (right axis, blue line). The three time windows of data used are indicated in the vertical dashed (T_1), dotted (T_2), and dashed-dotted (T_3) green lines. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

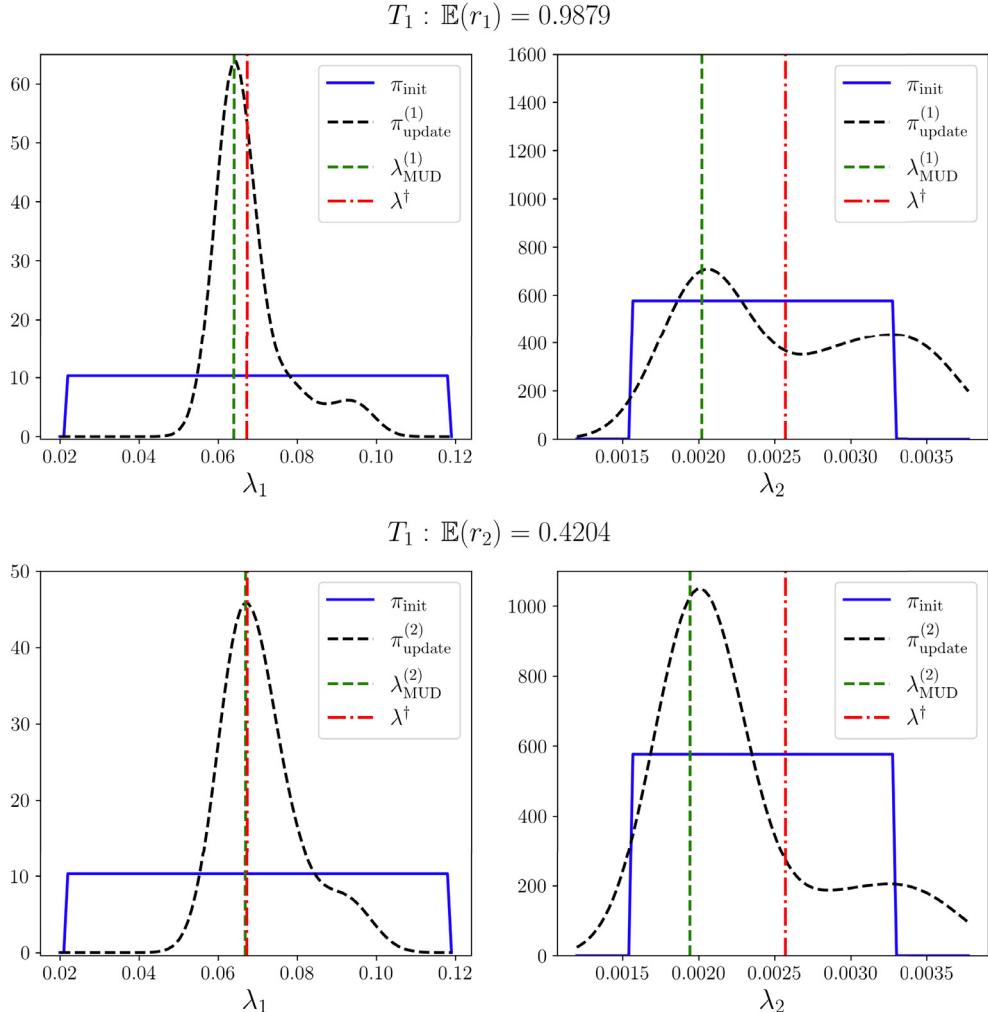


Fig. 12. Updated distributions (dotted black line) and mud estimates (dotted green line) using a one component (top) vs two component (bottom) Q_{PCA} map for T_1 (see Fig. 11) with $N = 119$ data points from January 11 01:00:00–07:00:00. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

To summarize, these results on different time windows illustrate well the benefits of using the Q_{PCA} and MUD point estimation algorithm for parameter estimation problems based on temporal data. It is critical to monitor the diagnostic $\mathbb{E}(r)$ since it gives us a specific metric to determine the quality of reconstructed distributions and potential violations of assumptions.

7. Conclusions and future work

A new approach to estimating and quantifying uncertainties in estimates of parameter identification problems is presented within a data-consistent framework for solving stochastic inverse problems. This approach identifies the parameter that maximizes the updated density solving data-consistent inverse problems. This parameter, referred to as the MUD point, is compared to the maximum a posteriori (MAP) and least squares estimates obtained by solving other formulations of the parameter identification problem. Under the standard assumptions of linear maps with Gaussian distributions, it is demonstrated that the MUD point maintains the predictive accuracy of a least squares estimate with the flexibility of incorporating prior/initial beliefs that make the MAP estimate popular in the

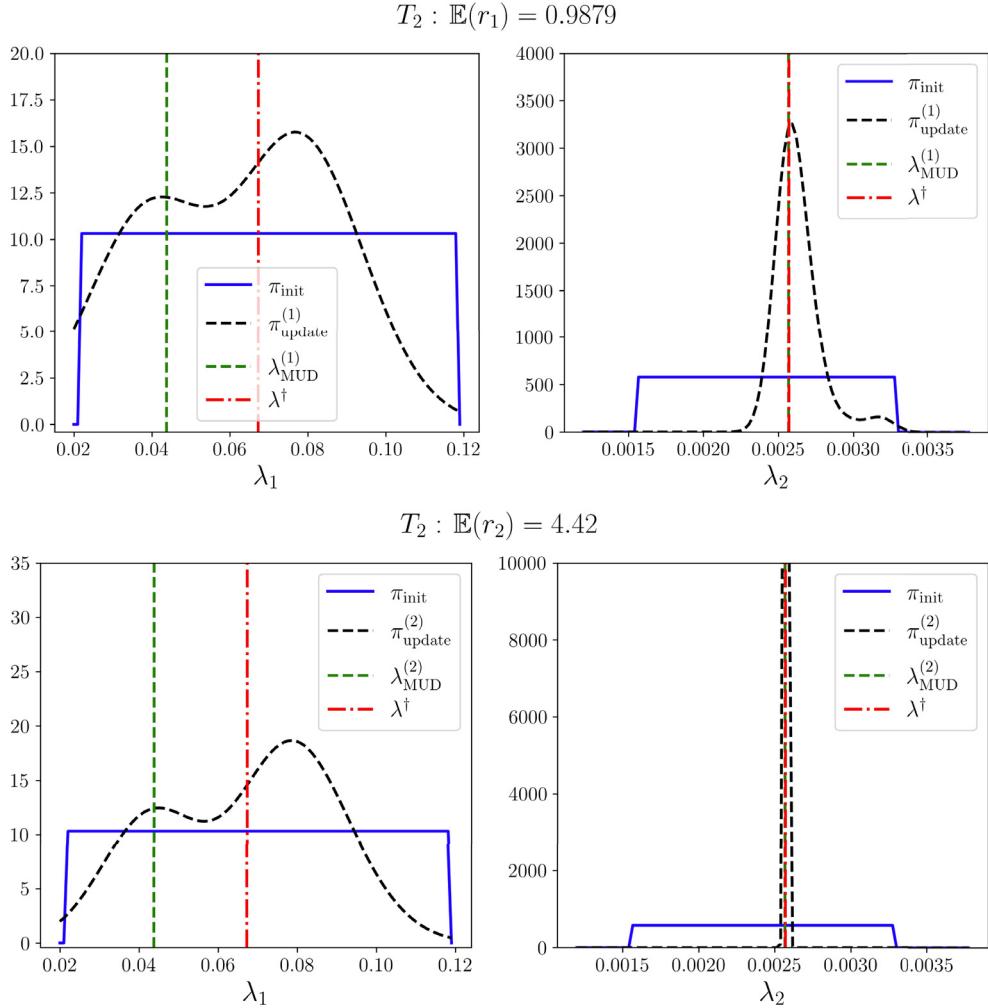


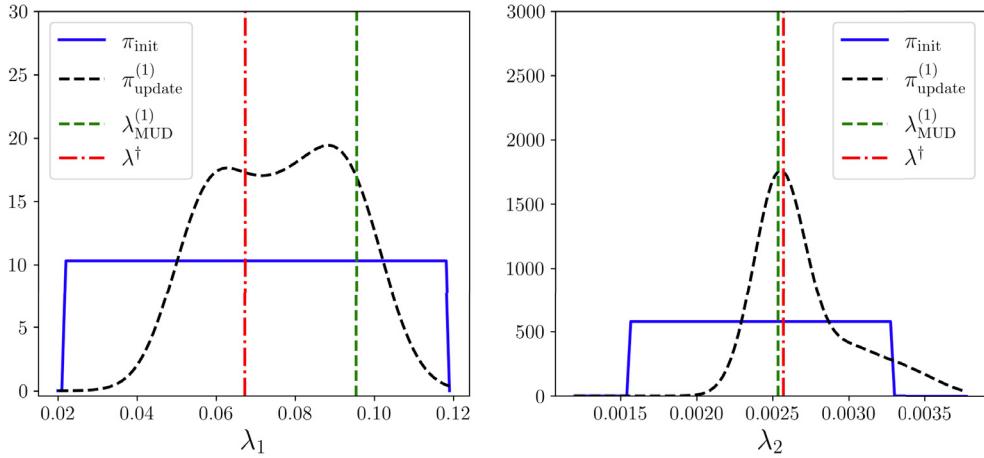
Fig. 13. Updated distributions (dotted black line) and mud estimates (dotted green line) using a one component (top) vs two component (bottom) Q_{PCA} map for T_2 (see Fig. 11) with $N = 53$ data points from January 04 11:00:00–14:00:00. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

uncertainty quantification literature. The theory of existence and uniqueness of the MUD point is also proven under these assumptions.

We then demonstrate how to formulate and apply a data-constructed QoI map for estimating and quantifying uncertainties in the MUD point. Specifically, the definition of a weighted mean error (WME) QoI map is utilized that can incorporate an arbitrary stream of data associated with repeated measurements. The WME map construction is motivated based on an application of the Fisher–Newman factorization theorem to the joint data-likelihood function commonly used in a Bayesian approach. The particular form of the WME map used in this work admits a fixed observed density as a function of observed data that nonetheless possesses several useful characteristics. First, for linear measurement operators, it is established that once a threshold on the number of observed data is reached, the existence and uniqueness of the MUD point is guaranteed. Second, the eigenvalues of the updated covariance related to the data-informed directions decrease at rates inversely proportional to the number of data obtained for each component of the WME map.

We then illustrate how an alternative data-constructed QoI map utilizing a PCA can be substituted for the WME map when potentially non-repeated spatial–temporal measurement data are available. This alternative QoI map is

$$T_3 : \mathbb{E}(r_1) = 1.238$$



$$T_3 : \mathbb{E}(r_2) = 0.9085$$

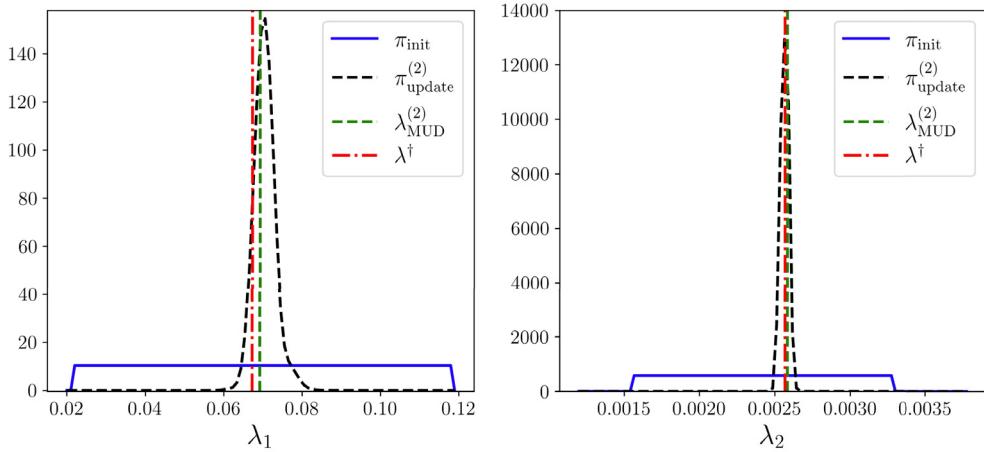


Fig. 14. Updated distributions (dotted black line) and mud estimates (dotted green line) using a one component (top) vs two component (bottom) Q_{PCA} map for T_3 (see Fig. 11) with $N = 959$ data points from January 07 00:00:00 to January 09 00:00:00. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

applied to two differential equations utilizing noisy spatial and temporal data including an example where wind drag parameters are estimated within the ADCIRC model for storm surge based on water elevation data over time.

Future directions include issues of optimal experimental design within this framework as well as the sequential estimation of the MUD point as data are obtained sequentially in space or time. In particular, for an online learning situation where optimal parameters are to be determined as data “stream into” the system, we will consider the construction of the Q_{PCA} maps using increasing time windows of data, with the diagnostic $\mathbb{E}(r)$ serving as a measure as to when enough data have been accumulated to invert on one or more model parameter values. Re-sampling can then be done from these updated distributions and the process repeated to increase the resolution of the λ^{MUD} estimate as needed. We will also consider non-density based approaches in future work based on optimization formulations of empirical distribution functions where too few samples are available for reliable density estimation in the space defined by the data-derived QoI. Preliminary work in this direction is promising (e.g., see [53]). There is also interest in applying this MUD and data-derived QoI framework to higher-dimensional parameter spaces modeling microstructures in material properties as was previously analyzed in [30] the classic data-consistent inverse framework.

Declaration of competing interest

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

Data availability

Data available at <https://github.com/mindthemath/mud-data>.

Acknowledgments

T. Butler's, M. Pilosov's, and T. Yu Yen's work is supported in part by the National Science Foundation, USA under Grant No. DMS-1818941. T. Butler's work is also partially supported by the National Science Foundation, USA under Grant No. DMS-2208460. C. Dawson's and C. del-Castillo-Negrete's work is supported in part by the National Science Foundation, USA No.DMS-1818847 and No.DMS-2208461. The opinions, findings, and conclusions or recommendations expressed in this manuscript are those of the authors and do not necessarily reflect the views of the National Science Foundation, nor Sandia National Laboratories, USA (SAND2022-16991 O).

Appendix A. Derivation of λ^{MUD} for linear Gaussian case

We begin by an alternative representation of $J(\lambda)$. First, define

$$R := \Sigma_{\text{init}}^{-1} - A^\top \Sigma_{\text{pred}}^{-1} A. \quad (\text{A.1})$$

Using this R , rewrite $J(\lambda)$ as

$$J(\lambda) := \|\mathbf{y} - Q(\lambda)\|_{\Sigma_{\text{obs}}^{-1}}^2 + \|\lambda - \lambda_0\|_R^2. \quad (\text{A.2})$$

In this form, we identify R as the *effective regularization* in $J(\lambda)$ due to the formulation in the data-consistent framework.

Observe that if $m = p$, then, by the assumption that A is full-rank, A is invertible. In this case, R is the $p \times p$ zero matrix and (A.2) reduces to the data-discrepancy term so that the MUD point is recognizable as the least squares solution, i.e., the point that minimizes the data-discrepancy term. Moreover, in this case we can immediately identify that $\lambda^{\text{MUD}} = A^{-1}(\mathbf{y} - b)$. This is also evident from the perspective of the densities. Specifically, in this case, π_{update} is defined by applying a change of variables formula to π_{obs} .

Suppose instead that $m < p$ so that the inverse-problem is under-determined. In this case, we observe that constructing R only requires specification of the initial/prior density and the QoI map, i.e., R may be defined prior to any collection of data on the QoI. Subsequently, we can interpret $J(\lambda)$ as coming from a modified Bayesian inverse problem with a prior defined by a $N(\lambda_0, \Sigma_R)$ distribution where $\Sigma_R = R^{-1}$. In other words, the MUD and MAP points can both be interpreted as solutions to different Bayesian inverse problems.

However, Σ_R is in fact a degenerative covariance, i.e., R is not technically invertible. This implies that Σ_R cannot be directly substituted in for Σ_{init} in (17) to define a closed form expression for Σ_{update} . We therefore first substitute Σ_{post} and $\Sigma_{\text{init}}^{-1}$ in (17) with Σ_{update} and R , respectively, to get

$$\Sigma_{\text{update}} := (A^\top \Sigma_{\text{obs}}^{-1} A + R)^{-1}. \quad (\text{A.3})$$

Since R is not invertible, Woodbury's identity cannot be applied (yet). Using (A.1), we can form

$$\Sigma_{\text{update}} = \left(A^\top \Sigma_{\text{obs}}^{-1} A + \Sigma_{\text{init}}^{-1} - A^\top \Sigma_{\text{pred}}^{-1} A \right)^{-1}, \quad (\text{A.4})$$

which is re-arranged as

$$\Sigma_{\text{update}} = \left(A^\top \left[\Sigma_{\text{obs}}^{-1} - \Sigma_{\text{pred}}^{-1} \right] A + \Sigma_{\text{init}}^{-1} \right)^{-1}. \quad (\text{A.5})$$

Recall from Section 4.1 that the predictability assumption in this case is that the smallest eigenvalue of Σ_{pred} is larger than the largest eigenvalue of Σ_{obs} . The roles are reversed when we consider the inverses of these matrices.

Subsequently, $\Sigma_{\text{obs}}^{-1} - \Sigma_{\text{pred}}^{-1}$ is a symmetric positive definite matrix and thus invertible. Applying the Woodbury identity yields

$$\Sigma_{\text{update}} = \Sigma_{\text{init}} - \Sigma_{\text{init}} A^\top \left(\left[\Sigma_{\text{obs}}^{-1} - \Sigma_{\text{pred}}^{-1} \right]^{-1} + \Sigma_{\text{pred}} \right)^{-1} A \Sigma_{\text{init}}. \quad (\text{A.6})$$

Applying Hua's identity and simplifying gives

$$\left(\left[\Sigma_{\text{obs}}^{-1} - \Sigma_{\text{pred}}^{-1} \right]^{-1} + \Sigma_{\text{pred}} \right)^{-1} = \Sigma_{\text{pred}}^{-1} \left[\Sigma_{\text{pred}} - \Sigma_{\text{obs}} \right] \Sigma_{\text{pred}}^{-1}. \quad (\text{A.7})$$

Substituting (A.7) into (A.6) gives

$$\Sigma_{\text{update}} = \Sigma_{\text{init}} - \Sigma_{\text{init}} A^\top \Sigma_{\text{pred}}^{-1} \left[\Sigma_{\text{pred}} - \Sigma_{\text{obs}} \right] \Sigma_{\text{pred}}^{-1} A \Sigma_{\text{init}}. \quad (\text{A.8})$$

We can now modify the expression for the MAP point given in (18) by substituting Σ_{update} for Σ_{post} to write the MUD point that minimizes J as

$$\lambda^{\text{MUD}} = \lambda_0 + \Sigma_{\text{update}} A^\top \Sigma_{\text{obs}}^{-1} (\mathbf{y} - b - A\lambda_0). \quad (\text{A.9})$$

Substituting (A.8) into (A.9) and simplifying, we have

$$\lambda^{\text{MUD}} = \lambda_0 + \Sigma_{\text{init}} A^\top \Sigma_{\text{pred}}^{-1} (\mathbf{y} - b - A\lambda_0). \quad (\text{A.10})$$

From a practical perspective, (A.10) is the preferred form for calculating the MUD point numerically given its reduced complexity in terms of the number of FLOPS required, and it is the default method used in the software implementation (see below). One would opt in for using (A.9) only if the updated covariance Σ_{update} is required. This option is available in the software implementation (See Appendix B) by setting the `solve(method='mud_alt')` option in the `LinearGaussianProblem` class. The updated covariance Σ_{update} is computed using the `LinearGaussianProblem.updated_cov()` function which computes (A.8).

Appendix B. Software contributions

The work presented here is available on GitHub at github.com/mathematicalmichael/mud.git as a complete Python package compliant with PEP 517 and 518 and published to the PyPi Python Package Registry under the name `mud`. Convenient python classes encompassing the core mathematical objects and analytical expressions are provided in the `mud.base`, including `LinearGaussianProblem`, for the analytical linear solutions, `DensityProblem`, for the density base MUD point estimation problems, and `SpatioTemporalProblem`, for aggregating spatio-temporal data in data-constructed QoI maps.

Running `pip install mud[examples]` will install the `mud` package and its dependencies, as well as a convenient Command Line Interface (CLI) to run the main examples presented. This means that upon successful installation, one can run `mud examples --help` from a command line to explore the examples available and options for each. The entrypoint `mud examples mud-paper` will run all the examples as shown in this paper (except for the ADCIRC example in Section 6.2, see below). The package version used in this paper is `mud==0.1`, which is compatible with Python 3.7+.

The datasets for the ADCIRC example in Section 6.2 is hosted on DesignSafe, a comprehensive cyber-infrastructure that is part of the NSF-funded Natural Hazard Engineering Research Infrastructure (NHERI) and provides cloud-based tools to manage, analyze, understand, and publish critical data and research related to impacts of natural hazards [54]. The published project directory [55] includes a static version of the ADCIRC data presented, along with Jupyter notebooks to recreate the data-set itself using DesignSafe HPC resources to run ensembles of ADCIRC simulations.

References

- [1] D. O'Malley, V.V. Vesselinov, A combined probabilistic/nonprobabilistic decision analysis for contaminant remediation, *SIAM/ASA J. Uncertain. Quantif.* 2 (1) (2014) 607–621, <http://dx.doi.org/10.1137/140965132>.
- [2] D. O'Malley, V.V. Vesselinov, Bayesian-information-gap decision theory with an application to CO₂ sequestration, *Water Resour. Res.* 51 (9) (2015) 7080–7089, <http://dx.doi.org/10.1002/2015WR017413>.
- [3] S.A. Mattis, T.D. Butler, C.N. Dawson, D. Estep, V.V. Vesselinov, Parameter estimation and prediction for groundwater contamination based on measure theory, *Water Resour. Res.* 51 (9) (2015) 7608–7629, <http://dx.doi.org/10.1002/2015WR017295>.

[4] D.R. Harp, V.V. Vesselinov, Contaminant remediation decision analysis using information gap theory, *Stoch. Environ. Res. Risk Assess.* 27 (1) (2013) 159–168, <http://dx.doi.org/10.1007/s00477-012-0573-1>.

[5] T. Butler, D. Estep, S. Tavener, C. Dawson, J.J. Westerink, A measure-theoretic computational method for inverse sensitivity problems III: Multiple quantities of interest, *SIAM/ASA J. Uncertain. Quantif.* 2 (1) (2014) 174–202, <http://dx.doi.org/10.1137/130930406>.

[6] T. Butler, J. Jakeman, T. Wildey, Combining push-forward measures and Bayes' rule to construct consistent solutions to stochastic inverse problems, *SIAM J. Sci. Comput.* 40 (2) (2018) A984–A1011, <http://dx.doi.org/10.1137/16M1087229>.

[7] T. Butler, J. Jakeman, T. Wildey, Convergence of probability densities using approximate models for forward and inverse problems in uncertainty quantification, *SIAM J. Sci. Comput.* 40 (5) (2018) A3523–A3548, <http://dx.doi.org/10.1137/18M1181675>.

[8] T. Butler, T. Wildey, T.Y. Yen, Data-consistent inversion for stochastic input-to-output maps, *Inverse Problems* 36 (8) (2020) 085015, <http://dx.doi.org/10.1088/1361-6420/ab8f83>.

[9] C. Dellacherie, P.A. Meyer, C. Dellacherie, *Probabilities and potential*, in: *North-Holland Mathematics Studies*, no. 29, Hermann [u.a.], Paris, 1978.

[10] J.T. Chang, D. Pollard, Conditioning as disintegration, *Stat. Neerl.* 51 (3) (1997) 287–317, <http://dx.doi.org/10.1111/1467-9574.00056>.

[11] J. Wang, N. Zabaras, Hierarchical Bayesian models for inverse problems in heat conduction, *Inverse Problems* 21 (1) (2004) 183, <http://dx.doi.org/10.1088/0266-5611/21/1/012>.

[12] C.K. Wikle, L.M. Berliner, N. Cressie, Hierarchical Bayesian space-time models, *Environ. Ecol. Stat.* 5 (2) (1998) 117–154, <http://dx.doi.org/10.1023/A:1009662704779>.

[13] R.C. Smith, *Uncertainty Quantification: Theory, Implementation, and Applications*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2013.

[14] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*, Society for Industrial and Applied Mathematics, 2005, <http://dx.doi.org/10.1137/1.9780898717921>.

[15] D. Calvetti, J.P. Kaipio, E. Somersalo, Inverse problems in the Bayesian framework, *Inverse Problems* 30 (11) (2014) 110301, <http://dx.doi.org/10.1088/0266-5611/30/11/110301>.

[16] B.G. Fitzpatrick, Bayesian analysis in inverse problems, *Inverse Problems* 7 (5) (1991) 675, <http://dx.doi.org/10.1088/0266-5611/7/5/003>.

[17] T. Bui-Thanh, O. Ghattas, An analysis of infinite dimensional Bayesian Inverse Shape acoustic scattering and its numerical approximation, *SIAM/ASA J. Uncertain. Quantif.* 2 (1) (2014) 203–222, <http://dx.doi.org/10.1137/120894877>.

[18] J. Cockayne, C. Oates, T. Sullivan, M. Girolami, Probabilistic numerical methods for PDE-constrained Bayesian inverse problems, in: *AIP Conference Proceedings* 1853, AIP Publishing, 2017, 060001, <http://dx.doi.org/10.1063/1.4985359>.

[19] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, D.B. Rubin, *Bayesian Data Analysis*, third ed., Chapman and Hall/CRC, 2013.

[20] E.T. Jaynes, *Probability Theory: The Logic of Science*, Cambridge University Press, Cambridge, UK ; New York, NY, 2003,

[21] M.C. Kennedy, A. O'Hagan, Bayesian calibration of computer models, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 63 (3) (2001) 425–464, <http://dx.doi.org/10.1111/1467-9868.00294>.

[22] Y.M. Marzouk, H.N. Najm, L.A. Rahn, Stochastic spectral methods for efficient Bayesian solution of inverse problems, *J. Comput. Phys.* 224 (2) (2007) 560–586, <http://dx.doi.org/10.1016/j.jcp.2006.10.010>.

[23] S.L. Cotter, M. Dashti, A.M. Stuart, Approximation of Bayesian Inverse Problems for PDEs, *SIAM J. Numer. Anal.* 48 (1) (2010) 322–345, <http://dx.doi.org/10.1137/090770734>.

[24] P.B. Stark, L. Tenorio, A primer of frequentist and Bayesian inference in inverse problems, in: *Large-Scale Inverse Problems and Quantification of Uncertainty*, John Wiley & Sons, Ltd, 2010, pp. 9–32, <http://dx.doi.org/10.1002/9780470685853.ch2>.

[25] O.G. Ernst, B. Sprungk, H.-J. Starkloff, Bayesian inverse problems and Kalman filters, in: S. Dahlke, W. Dahmen, M. Griebel, W. Hackbusch, K. Ritter, R. Schneider, C. Schwab, H. Yserentant (Eds.), *Extraction of Quantifiable Information from Complex Systems*, in: *Lecture Notes in Computational Science and Engineering*, Springer International Publishing, Cham, 2014, pp. 133–159, http://dx.doi.org/10.1007/978-3-319-08159-5_7.

[26] T. Cui, Y. Marzouk, K. Willcox, Scalable posterior approximations for large-scale Bayesian inverse problems via likelihood-informed parameter and state reduction, *J. Comput. Phys.* 315 (2015) <http://dx.doi.org/10.1016/j.jcp.2016.03.055>.

[27] A. Stuart, Inverse problems: A Bayesian perspective, *Acta Numer.* 19 (2010) 451–559, <http://dx.doi.org/10.1017/S0962492910000061>.

[28] R. von Mises, J.L. Doob, Discussion of papers on probability theory, *Ann. Math. Stat.* 12 (2) (1941) 215–217, <http://dx.doi.org/10.1214/aoms/1177731750>.

[29] T. Butler, H. Hakula, What do we hear from a drum? A data-consistent approach to quantifying irreducible uncertainty on model inputs by extracting information from correlated model output data, *Comput. Methods Appl. Mech. Engrg.* 370 (2020) 113228, <http://dx.doi.org/10.1016/j.cma.2020.113228>.

[30] A. Tran, T. Wildey, Solving stochastic inverse problems for Property–Structure linkages using data-consistent inversion and machine learning, *JOM J. Miner. Met. Mater. Soc.* 73 (1) (2021) 72–89, <http://dx.doi.org/10.1007/s11837-020-04432-w>.

[31] B. Marvin, T. Bui-Thanh, T. Wildey, A scalable approach for solving stochastic inverse problems based on push-forward measures and Bayes' rule, in: A. Cangi, M. Parks (Eds.), *Center for Computing Research Summer Proceedings*, Sandia National Laboratories, Albuquerque, NM, 2018, pp. 120–131.

[32] G. Casella, R. Berger, *Statistical Inference*, in: *Duxbury Advanced Series in Statistics and Decision Sciences*, Thomson Learning, 2002.

[33] K. Pearson, LIII. On lines and planes of closest fit to systems of points in space, *Lond. Edinb. Dublin Philos. Mag. J. Sci.* 2 (11) (1901) 559–572, <http://dx.doi.org/10.1080/14786440109462720>.

[34] H. Hotelling, Analysis of a complex of statistical variables into principal components, *J. Educ. Psychol.* 24 (6) (1933) 417–441, <http://dx.doi.org/10.1037/h0071325>.

[35] J. Shlens, A tutorial on principal component analysis, 2014, <http://dx.doi.org/10.48550/arXiv.1404.1100>, arXiv:1404.1100.

- [36] I.T. Jolliffe, Principal component analysis for special types of data, in: Principal Component Analysis, in: Springer Series in Statistics, Springer, New York, NY, 2002, pp. 338–372, http://dx.doi.org/10.1007/0-387-22440-8_13.
- [37] C.B. Vreugdenhil, in: V.P. Singh (Ed.), Numerical Methods for Shallow-Water Flow, in: Water Science and Technology Library, vol. 13, Springer Netherlands, Dordrecht, 1994, <http://dx.doi.org/10.1007/978-94-015-8354-1>.
- [38] J.R. Garratt, Review of drag coefficients over oceans and continents, *Mon. Weather Rev.* 105 (7) (1977) 915–929, [http://dx.doi.org/10.1175/1520-0493\(1977\)105<0915:RODCOO>2.0.CO;2](http://dx.doi.org/10.1175/1520-0493(1977)105<0915:RODCOO>2.0.CO;2).
- [39] J. Luettich, J. Richard, J. Westerink, N. Scheffner, ADCIRC: An Advanced Three-Dimensional Circulation Model for Shelves, Coasts, and Estuaries. Report 1. Theory and Methodology of ADCIRC-2DDI and ADCIRC-3DL, in: Dredging Research Program Tech. Rep. DRP-92-6, 1992, p. 143.
- [40] J.J. Westerink, R.A. Luettich, A.M. Baptists, N.W. Scheffner, P. Farrar, Tide and storm surge predictions using finite element model, *J. Hydraul. Eng.* 118 (10) (1992) 1373–1390, [http://dx.doi.org/10.1061/\(ASCE\)0733-9429\(1992\)118:10\(1373\)](http://dx.doi.org/10.1061/(ASCE)0733-9429(1992)118:10(1373)).
- [41] D.R. Lynch, W.G. Gray, A wave equation model for finite element tidal computations, *Comput. & Fluids* 7 (3) (1979) 207–228, [http://dx.doi.org/10.1016/0045-7930\(79\)90037-9](http://dx.doi.org/10.1016/0045-7930(79)90037-9).
- [42] J. Dietrich, C. Dawson, J. Proft, M. Howard, G. Wells, J. Fleming, J. Luettich, J. Westerink, Z. Cobell, M. Vitse, H. Lander, B. Blanton, C. Szpilka, J. Atkinson, Real-time forecasting and visualization of hurricane waves and storm surge using SWAN + ADCIRC and FigureGen, *IMA Vol. Math. Appl.* 156 (2013) 49–70, http://dx.doi.org/10.1007/978-1-4614-7434-0_3.
- [43] S. Bunya, J. Dietrich, J. Westerink, B. Ebersole, J. Smith, J. Atkinson, R. Jensen, D. Resio, J. Luettich, C. Dawson, V. Cardone, A. Cox, 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 (2010) 345–377, <http://dx.doi.org/10.1175/2009MWR2906.1>.
- [44] J. Dietrich, S. Bunya, J. Westerink, B. Ebersole, J. Smith, J. Atkinson, R. Jensen, D. Resio, J. Luettich, C. Dawson, V. Cardone, A. Cox, A high-resolution coupled riverine flow, tide, wind, wind wave, and storm surge model for Southern Louisiana and Mississippi. Part II: Synoptic description and analysis of Hurricanes Katrina and Rita, *Mon. Weather Rev.* 138 (2010) 378–404, <http://dx.doi.org/10.1175/2009MWR2907.1>.
- [45] J. Dietrich, J. Westerink, A. Kennedy, J. Smith, R. Jensen, M. Zijlema, L. Holthuijsen, C. Dawson, J. Luettich, M. Powell, V. Cardone, A. Cox, Z. Cobell, Hurricane Gustav (2008) waves and storm surge: Hindcast, synoptic analysis, and validation in Southern Louisiana, *Mon. Weather Rev.* 139 (2011) 2488–2522, <http://dx.doi.org/10.1175/2011MWR3611.1>.
- [46] T. Butler, L. Graham, D. Estep, C. Dawson, J. Westerink, Definition and solution of a stochastic inverse problem for the Manning's n parameter field in hydrodynamic models, *Adv. Water Resour.* 78 (2015) 60–79, <http://dx.doi.org/10.1016/j.advwatres.2015.01.011>.
- [47] L. Graham, Adaptive Measure-Theoretic Parameter Estimation for Coastal Ocean Modeling (Ph.D. thesis), 2015.
- [48] L. Graham, T. Butler, S. Walsh, C. Dawson, 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), *Mon. Weather Rev.* 145 (3) (2017) 929–954, <http://dx.doi.org/10.1175/MWR-D-16-0149.1>.
- [49] G.D. Egbert, S.Y. Erofeeva, Efficient inverse modeling of barotropic ocean tides, *J. Atmos. Ocean. Technol.* 19 (2) (2002) 183–204, [http://dx.doi.org/10.1175/1520-0426\(2002\)019<0183:EIMBO>2.0.CO;2](http://dx.doi.org/10.1175/1520-0426(2002)019<0183:EIMBO>2.0.CO;2).
- [50] K.J. Roberts, W.J. Pringle, J.J. Westerink, OceanMesh2D 1.0: MATLAB-based software for two-dimensional unstructured mesh generation in coastal ocean modeling, *Geosci. Model Dev.* 12 (5) (2019) 1847–1868, <http://dx.doi.org/10.5194/gmd-12-1847-2019>.
- [51] S. Saha, S. Moorthi, X. Wu, J. Wang, S. Nadiga, P. Tripp, D. Behringer, Y.-T. Hou, H.-y. Chuang, M. Iredell, M. Ek, J. Meng, R. Yang, M.P. Mendez, H. van den Dool, Q. Zhang, W. Wang, M. Chen, E. Becker, The NCEP climate forecast system version 2, *J. Clim.* 27 (6) (2014) 2185–2208, <http://dx.doi.org/10.1175/JCLI-D-12-00823.1>.
- [52] C. Letchford, B. Zachry, On wind, waves, and surface drag, in: 5th European and African Conference on Wind Engineering, Florence, Italy, 2009, p. 32.
- [53] K. Bergstrom, T. Butler, T. Wildey, Estimating observation-consistent solutions using weighted empirical distribution functions, in: A. Rushdi, M.L. Parks (Eds.), Center for Computing Research Summer Proceedings, Sandia National Laboratories, Albuquerque, NM, 2020, pp. 11–22.
- [54] E.M. Rathje, C. Dawson, J.E. Padgett, J.-P. Pinelli, D. Stanzione, A. Adair, P. Arduino, S.J. Brandenberg, T. Cockerill, C. Dey, M. Esteva, F.L. Haan, M. Hanlon, A. Kareem, L. Lowes, S. Mock, G. Mosqueda, DesignSafe: New cyberinfrastructure for natural hazards engineering, *Nat. Hazards Rev.* 18 (3) (2017) 06017001, [http://dx.doi.org/10.1061/\(ASCE\)NH.1527-6996.0000246](http://dx.doi.org/10.1061/(ASCE)NH.1527-6996.0000246).
- [55] D.-C.-N. Carlos, Parameter Estimation with Maximal Updated Densities - ADCIRC Shinnecock Inlet, DesignSafe-CI, 2022, <http://dx.doi.org/10.17603/ds2-bmrf-y737>, URL: <https://www.designsafe-ci.org/data/browser/public/designsafe.storage.published/PRJ-3492>.