Formulating and Solving the Data-Consistent Geophysical Inverse Problem for Subsurface Modeling Applications

Alex Miltenberger[*1], Lijing Wang[2], Tapan Mukerji[3], Jef Caers[2]

[1]Stanford University, Department of Geophysics, Stanford, CA, USA

[2]Stanford University, Department of Geological Sciences, Stanford, CA, USA

[3]Stanford University, Department of Energy Science & Engineering, Stanford, CA, USA

*Corresponding author: ammilten@gmail.com

ABSTRACT

Today, most probabilistic geophysical inverse problems are formulated using one of two methods: (1) conditional probability and Bayes' Theorem, or (2) Tarantola's theory of intersecting probability densities. More recently, a third inverse problem formulation based on pushforward probability measures was proposed, termed "data-consistent inversion". Many practical problems can be cast into any of these three formulations, but the choice of formulation may change the posterior. To help geophysicists understand how the formulation affects the posterior, we present a mathematical comparison of the three formulations accompanied by simple physical example to illustrate the difference between posteriors from each method. The distinguishing feature of the data-consistent formulation is that it yields a posterior that maps onto the measurement distribution, whereas the traditional formulations do not. Then we propose a flexible Sequential Monte Carlo algorithm for solving practical subsurface modeling problems conditioned to geophysical data. This algorithm uses density estimation techniques to fit a generative probabilistic model to the posterior. The algorithm is demonstrated using a synthetic gravity inversion example. The results are then compared to results from the classical Bayesian formulation solved using Markov Chain Monte Carlo. The comparison shows that even though the difference between the data-consistent posterior and Bayesian posterior can be small in practice, the lack of conditional probabilities in the data-consistent formulation can make it simpler to fit generative probabilistic models to the solution of the inverse problem.

1. Introduction

Geophysical inverse problems [1] are frequently solved when creating subsurface models for mineral exploration [2], groundwater modeling [3], and more [4, 5]. However, quantifying

uncertainty in the relationship between geophysical data and subsurface models is difficult, due to the multiple sources of uncertainty, non-unique nature of geophysical inverse problems, and the availability of several different ways to formulate the inverse problem [6-8]. Three popular inverse problem formulations include (1) the Bayesian formulation based on conditional probability [3, 7, 10, 11], (2) Tarantola & Valette's intersecting probability density formulation [12], and (3) the more recent data-consistent formulation [8]. Given these options, there is a need to understand how the choice of formulation effects the solution to the inverse problem. Therefore, we present a mathematical comparison between these formulations that highlights the similarities and differences of each approach. The results of this comparison show that the data-consistent formulation, which has the unique property where the posterior is guaranteed to push forward into the measurement density, is preferred for subsurface modeling problems with noisy measurements and non-linear forward models. We then go on to propose a practical algorithm based on Sequential Monte Carlo (SMC) [13-15] for solving the data-consistent formulation in subsurface modeling applications.

Today, most probabilistic inverse problems are formulated as a conditional probability $f(\mathbf{m}|\mathbf{d}_{obs})$ where f is a probability density function, $\mathbf{m}$ is the model parameter vector, and $\mathbf{d}_{obs}$ is the observed data vector. However, in some applications this formulation causes issues such as the Borel paradox, where changing the coordinate system changes the posterior density [1, 17]. Tarantola & Valette [12] recognized this problem and proposed an alternative inverse problem formulation based on a theory of intersecting probability densities. Butler et al. [8, 18] recently proposed another inverse problem formulation based on measure-theoretic principles that addresses the same issues. The main difference between the Butler et al. [8] formulation and the other two formulations is that the Butler et al. formulation yields a posterior distribution that

maps onto the distribution of measurements while the other two do not. For this reason, we refer to this approach as data-consistent inversion [18, 19]. While this data-consistent property was originally developed for quantifying aleatoric uncertainties, it has since been extended to quantify epistemic uncertainties in parameter estimation problems [20]. The data-consistent formulation assumes only that a prior is chosen whose pushforward through the model is capable of predicting the observed data and that the only source of uncertainty is the randomness/error in the measurements. The objective is to update the prior into a posterior that has the property that its pushforward through the model matches the distribution that models the observed measurements. In contrast, the traditional methods are designed to solve a problem with only epistemic uncertainties where the objective is to estimate specific parameter values that could have produced all the data instead of a distribution on parameters that maps to the distribution of data. In summary, these approaches were designed to solve different problems in terms of the ways in which they categorize and quantify uncertainties. More details on the three formulations are presented in Section 2.

Although these approaches solve different problems, their mathematics are intimately related. In Section 3 we explore the mathematical relationships between the data-consistent, conditional probability, and Tarantola & Valette's formulation. Then, in Section 4, we show how SMC, a method that has also been used to solve the traditional Bayesian inverse problem [16], can also be used to solve the data-consistent inverse problem. Finally, in Section 5 we provide a synthetic example to illustrate how this formulation can be applied to geophysical problems. The synthetic example considers the problem of inverting a gravity anomaly using a spherical geobody embedded within a homogenous host material. We solve this synthetic problem using two methods: kernel density estimation (KDE) and normalizing flows (NF) [21] to illustrate the

flexibility in the SMC algorithm. We compare these results to the posterior obtained by the conditional probability formulation solved with Hamiltonian Monte Carlo [22]. In Section 6 we will also discuss advantages, limitations, and technical challenges of this data-consistent formulation for geophysical inverse problems.

2. Three Inverse Problem Formulations

In this section we review the data-consistent formulation and two other popular formulations: the conditional probability (Bayesian) formulation [9, 10, 23], and the Tarantola & Valette [12] formulation based on intersecting probability densities.

2.1 Data-Consistent Formulation

The inverse problem formulation from Butler et al. [8] uses a prior probability measure, $\mathbb{P}_{prior}$, an observation probability measure $\mathbb{P}_{obs}$ that describes the measurement uncertainty, and a posterior probability measure $\mathbb{P}_{post}$. The measures $\mathbb{P}_{prior}$ and $\mathbb{P}_{post}$ are defined on a measurable space $(\Omega_{\mathbf{m}}, \mathcal{B}(\Omega_{\mathbf{m}}))$ where $\Omega_{\mathbf{m}}$ is the model parameter space and $\mathcal{B}(\Omega_{\mathbf{m}})$ is the Borel σ-algebra over that space. The measure $\mathbb{P}_{obs}$ is defined on the measurable space $(\Omega_{\mathbf{d}}, \mathcal{B}(\Omega_{\mathbf{d}}))$ where $\Omega_{\mathbf{d}}$ is the data variable space. We also define a forward model $g: \Omega_{\mathbf{m}} \to \Omega_{\mathbf{d}}$ that takes a set of model parameters as an input and outputs a data variable. From these definitions, the posterior is the probability measure $\mathbb{P}_{post}$ that, when pushed-forward by g, results in the observation probability measure $\mathbb{P}_{obs}$ [8]. That is,

$$g_* \mathbb{P}_{post} = \mathbb{P}_{obs} \tag{1}$$

where $g_*\mathbb{P}_{post}$ denotes the pushed-forward posterior probability measure. Now, we can define probability density functions f_{prior} and f_{post} by assuming $\mathbb{P}_{prior}$ and $\mathbb{P}_{post}$, respectively, are absolutely continuous with respect to the Lebesgue measure $\lambda_{\mathbf{m}}$ on $(\Omega_{\mathbf{m}}, \mathcal{B}(\Omega_{\mathbf{m}}))$. We also define the probability density functions f_{obs} and f_{push} by assuming $\mathbb{P}_{obs}$ and $g_*\mathbb{P}_{prior}$, respectively, are absolutely continuous with respect to the Lebesgue measure $\lambda_{\mathbf{d}}$ on $(\Omega_{\mathbf{d}}, \mathcal{B}(\Omega_{\mathbf{d}}))$. Using the Radon-Nikodym theorem [24], these four probability densities are defined by

$$f_{prior} := \frac{\mathrm{d}\mathbb{P}_{prior}}{\mathrm{d}\lambda_{\mathbf{m}}}; \quad f_{push} := \frac{\mathrm{d}(g_*\mathbb{P}_{prior})}{\mathrm{d}\lambda_{\mathbf{d}}}; \quad f_{obs} := \frac{\mathrm{d}\mathbb{P}_{obs}}{\mathrm{d}\lambda_{\mathbf{d}}}; \quad f_{post} := \frac{\mathrm{d}\mathbb{P}_{post}}{\mathrm{d}\lambda_{\mathbf{m}}} \tag{2}$$

Using the Radon-Nikodym theorem again, Eq. 1 can be rewritten as

$$\int_{\mathbf{d}\in A_{\mathbf{d}}} \mathrm{d}(g_*\mathbb{P}_{post}) = \int_{\mathbf{d}\in A_{\mathbf{d}}} f_{obs}(\mathbf{d}) \, \mathrm{d}\lambda_{\mathbf{d}} \tag{3}$$

where $A_{\mathbf{d}} \in \mathcal{B}(\Omega_{\mathbf{d}})$. After some algebra and substituting results within a Disintegration Theorem [24] we obtain the following formulation of the inverse problem solution [8]

$$f_{post}(\mathbf{m}) = \frac{f_{obs}(g(\mathbf{m}))}{f_{push}(g(\mathbf{m}))} f_{prior}(\mathbf{m}). \tag{4}$$

Further details about this formulation and its solution, including stability results, can be found in Butler et al. [8, 18]. Applications of this approach to dynamical systems where the model g is learned from temporal data streams can be found in Mattis et al. [19]. Pilosov et al. [20] explore generalizations to parameter estimation problems and apply this framework to estimate wind drag parameters from a simulated extreme weather event.

2.2 Conditional Probability (Bayesian) Formulation

In the conditional probability formulation [3, 25, 26], the posterior is denoted by the conditional probability density $f(\mathbf{m}|\mathbf{d}_{obs})$ and Bayes' Theorem is used to expand the posterior into the following form

$$f(\mathbf{m}|\mathbf{d}_{obs}) = \frac{f(\mathbf{d}_{obs}|\mathbf{m})}{f(\mathbf{d}_{obs})} f(\mathbf{m}). \qquad (5)$$

where $f(\mathbf{m})$ is the prior, $f(\mathbf{d}_{obs}|\mathbf{m})$ is the likelihood function, and $f(\mathbf{d}_{obs})$ is the marginal likelihood.

2.3 Intersecting Probability Densities (Tarantola & Valette)

Another popular inverse problem formulation was first proposed by Tarantola & Valette [12]. In this formulation, the posterior f_{post} is a joint density over $\mathbf{m}$ and $\mathbf{d}$. The posterior is defined as the intersection between two joint densities, f_{prior} and f_{obs}, representing sources of information from prior knowledge and measurements, respectively. The intersection operator is denoted by $\wedge$ and leads to the following expression for f_{post} [12]

$$f_{post}(\mathbf{m}, \mathbf{d}) = f_{prior}(\mathbf{m}, \mathbf{d}) \wedge f_{obs}(\mathbf{m}, \mathbf{d}) = k \frac{f_{prior}(\mathbf{m}, \mathbf{d})\, f_{obs}(\mathbf{m}, \mathbf{d})}{\mu(\mathbf{m}, \mathbf{d})} \qquad (9)$$

where μ is referred to as the "homogeneous" probability density and k is a normalization constant. This homogeneous density is chosen such that transformations on $\mathbf{m}$ and $\mathbf{d}$, such as transforming electrical resistivity to conductivity, does not change f_{post} [1]. To select μ, one must decide which transformations are relevant to the problem and then μ can be calculated.

3. Comparing the Three Formulations

3.1 Data-Consistent vs. Conditional Probability Formulation

Based on the similarities between Eqs. 4 and 5, we see that $f_{obs}(g(\mathbf{m}))$ is analogous to the likelihood function, which has also been proposed by Tarantola & Valette [12]. We also see from comparing Eqs. 4 and 5 that $f_{push}(g(\mathbf{m}))$ is analogous to the marginal likelihood. However, in contrast with the marginal likelihood that is constant for fixed observed data, the denominator $f_{push}(g(\mathbf{m}))$ varies with $\mathbf{m}$. This subtle difference causes Eq. 4 and Eq. 5 to generally produce give different posteriors.

Equation 4 can be derived from the conditional probability formulation under two assumptions. The first assumption is that the measurement error, ε, is added to $\mathbf{d}_{obs}$ instead of $g(\mathbf{m})$. That is, we write $\mathbf{d} = g(\mathbf{m}) = \mathbf{d}_{obs} + \varepsilon$ [27] instead of $\mathbf{d} = g(\mathbf{m}) + \varepsilon$. With this framing, $\mathbf{d}_{obs}$ and $\mathbf{d}$ are two separate random variables. Then, we marginalize $f(\mathbf{m}|\mathbf{d}_{obs})$ over $\mathbf{d}$ such that

$$f(\mathbf{m}|\mathbf{d}_{obs}) = \int_{\Omega_{\mathbf{d}}} f(\mathbf{m}|\mathbf{d}, \mathbf{d}_{obs})\, f(\mathbf{d}|\mathbf{d}_{obs})\, \mathrm{d}\mathbf{d}. \tag{6}$$

In this equation $f(\mathbf{d}|\mathbf{d}_{obs})$ is equivalent to the measurement distribution f_{obs}. Now we assume $\mathbf{m}$ and $\mathbf{d}_{obs}$ are conditionally independent given $\mathbf{d}$ to obtain

$$f(\mathbf{m}|\mathbf{d}_{obs}) = \int_{\Omega_{\mathbf{d}}} f(\mathbf{m}|\mathbf{d})\, f(\mathbf{d}|\mathbf{d}_{obs})\, \mathrm{d}\mathbf{d}. \tag{7}$$

In Appendix B we show that Eq. 7 simplifies to the following expression if g is deterministic:

$$f(\mathbf{m}|\mathbf{d}_{obs}) = \frac{f(g(\mathbf{m})|\mathbf{d}_{obs})}{f(g(\mathbf{m}))} f(\mathbf{m}).$$ (8)

In this formulation, $f(g(\mathbf{m})|\mathbf{d}_{obs}) = f_{obs}(g(\mathbf{m}))$ and $f(g(\mathbf{m})) = f_{push}(g(\mathbf{m}))$. Therefore, when the simulated and observed data are viewed as separate distributions, the conditional probability formulation becomes equivalent to the data-consistent formulation.

3.2 Data-Consistent vs. Tarantola & Valette's Formulation

Tarantola's theory provides a deeper insight into the relationship between the conditional probability and data-consistent formulations. To see this, let us assume $f_{prior}(\mathbf{m}, \mathbf{d})$ represents both the prior and the pushed-forward prior, such that $f_{prior}(\mathbf{m}, \mathbf{d}) = \delta(\mathbf{d} - g(\mathbf{m})) f_{prior}(\mathbf{m})$ [1, 27]. Let us also assume that f_{obs} represents only measurements of $\mathbf{d}$ such that $f_{obs}(\mathbf{m}, \mathbf{d}) = \mu(\mathbf{m}) f_{obs}(\mathbf{d})$. Finally, we assume that $\mu(\mathbf{m}, \mathbf{d})$ is constant. Under these conditions, we recover the conditional probability formulation and another formulation similar to the data-consistent formulation (Appendix C):

$$f_{post}^{(\mathbf{m})}(\mathbf{m}) = \frac{f_{obs}^{(\mathbf{d})}(g(\mathbf{m})) f_{prior}^{(\mathbf{m})}(\mathbf{m})}{\int_{\Omega_{\mathbf{m}}} f_{obs}^{(\mathbf{d})}(g(\mathbf{m})) f_{prior}^{(\mathbf{m})}(\mathbf{m}) \, d\mathbf{m}} = \frac{f_{post}^{(\mathbf{d})}(g(\mathbf{m}))}{f_{prior}^{(\mathbf{d})}(g(\mathbf{m}))} f_{prior}^{(\mathbf{m})}(\mathbf{m})$$ (10)

where the superscripts $(\mathbf{m})$ and $(\mathbf{d})$ denote the space where the probability density is defined. The middle expression in Eq. 10 is exactly equivalent to the conditional probability formulation in Eq. 5 where $\int_{\Omega_{\mathbf{m}}} f_{obs}^{(\mathbf{d})}(g(\mathbf{m})) f_{prior}^{(\mathbf{m})}(\mathbf{m}) \, d\mathbf{m}$ is the marginal likelihood. The right-hand-side

expression in Eq. 10 would be equal to the data-consistent formulation if and only if

$$f_{post}^{(\mathbf{d})}(g(\mathbf{m})) = f_{obs}(g(\mathbf{m})).$$

Equation 10 shows that, both the data variable and the model parameters are updated in the conditional probability formulation. The idea of a data posterior, $f_{post}(\mathbf{d})$, is not discussed in geophysical inversion literature, but it is a feature of the conditional probability formulation nonetheless. Equation 10 also shows that the data-consistent formulation can be recovered from the conditional probability formulation if and only if the data posterior is equal to the measurement density.

Finally, note that the homogeneous density in Tarantola's formulation is only involved in how $f_{post}(\mathbf{d})$ is updated. There is no homogeneous density in the data-consistent formulation. There is no need for this homogeneous density in the data-consistent formulation. All of these similarities and differences are illustrated in the following section.

3.3 A Simple Physical Example

The three formulations outlined above represent alternative ways to solve probabilistic inverse problems. Each have their own properties, which are summarized in Table 1. These formulations which may be desirable or undesirable for different problems. We illustrate the difference between these formulations using a simple physical example. Let us assume we have uncertain measurements of electrical conductivity, σ, with units S/m, and the goal is to quantify uncertainty on electrical resistivity, ρ, with units Ohm-m. The forward model from ρ to σ is $\sigma = g(\rho) = \rho^{-1}$.

Table 1. Summary of the three inverse problems formulations, for geophysical modeling applications.

Formulation	Main Equation	Main Desired Property	Challenges
Conditional Probability (Bayes)	$$f(\mathbf{m}|\mathbf{d}_{obs}) = \frac{f(g(\mathbf{m})|\mathbf{d}_{obs})}{f(\mathbf{d}_{obs})} f(\mathbf{m})$$	Exact when there's zero measurement error	Posterior diverges from measurement distribution as model nonlinearity and measurement error increase.
Tarantola & Valette	$$f_{post}(\mathbf{m},\mathbf{d}) = k \frac{f_{prior}(\mathbf{m},\mathbf{d}) \, f_{obs}(\mathbf{m},\mathbf{d})}{\mu(\mathbf{m},\mathbf{d})}$$	Can specify μ to ensure the posterior is invariant under invertible transformations of $\mathbf{m}$ (e.g. resistivity to conductivity)	Even if the solution is invariant, it can still diverge from measurement distribution with increasing model nonlinearity and measurement error.
Data-Consistent	$$f_{post}(\mathbf{m}) = \frac{f_{obs}(g(\mathbf{m}))}{f_{push}(g(\mathbf{m}))} f_{prior}(\mathbf{m})$$	Posterior maps onto the measurement distribution	Very little difference between other formulations when measurement error is small.

Now say one measurement of σ was taken, with a value of $\sigma_{obs} = 0.06 \pm 0.04$ S/m which is used to assume a uniform measurement density from 0.02 to 0.10 S/m, e.g. $f_{obs}(\sigma) = U(\sigma \mid 0.02, 0.10)$. It then follows that the error distribution $f_{err}(\varepsilon)$ is uniform from -0.04 S/m to 0.04 S/m, e.g. $f_{err}(\varepsilon) = U(\varepsilon \mid -0.04, 0.04)$. This leads to the likelihood $f(\sigma_{obs} \mid \rho) = f_{obs}(\rho) = U(g(\rho) \mid 0.02, 0.10)$. We assume the prior over ρ is uniform such that $f_{prior} = f(\rho) = U(\rho \mid 0.1, 100)$ Ohm-m.

Next, we calculate the posterior obtained from each of the three formulations semi-analytically. We also solve the conditional probability formulation using an approximate Bayesian computation (ABC) [28] approach to illustrate that the results are independent of the solution method.

3.3.1 Conditional Probability (Semi-Analytical)

First, we discretized ρ from 0.1 to 100 Ohm-m. Then, we calculated $f(\rho)$ and $f(\sigma_{obs} \mid \rho)$ using the definitions in Sect. 2.3. Finally, we derived the posterior by calculating $f(\rho \mid \sigma_{obs}) \propto f(\sigma_{obs} \mid \rho) \times f(\rho)$ and normalizing.

3.3.2 Conditional Probability (Approximate Bayesian Computation)

First, we sampled 100,000 $\rho \sim f(\rho)$ and $\varepsilon \sim f_{err}$. We used these samples to generate samples from the joint density $f(\rho, \sigma)$ by applying $\sigma = \rho^{-1} + \varepsilon$. Next, we rejected all samples with $|\sigma - \sigma_{obs}| > 0.001$ where $\sigma_{obs} = 0.06$ is the mean observation. The remaining 8392 samples make up samples from the posterior.

3.3.3 Intersecting Probability Densities (Tarantola & Valette, 1981)

We use $f_{post}(\rho) \propto \frac{f_{prior}(\rho) \, f_{obs}(\rho^{-1})}{\mu(\rho)}$ (Eq. 9) to calculate the posterior semi-analytically.

First, we discretized ρ from 0.1 to 100 Ohm-m. Tarantola [1] recommends $\mu(\rho) = \rho^{-1}$ for problems with strictly positive parameters like resistivities. We used this and the definitions for f_{prior} and f_{obs} outlined above to calculate the posterior.

3.3.4 Pushforward Probability Measures (Butler et al., 2018)

First, since the map defines a smooth bijection, it is possible to transform f_{prior} into f_{push} using the Jacobian rule such that $f_{push}(\sigma) = f_{prior}(\sigma^{-1})|\mathcal{J}(\sigma)|$ where $|\mathcal{J}(\sigma)| = \sigma^{-2}$. Then, we calculate $f_{post}(\rho)$ semi-analytically using Eq. 4 and the definitions in Sect. 2.3. Note that in general f_{push} must be approximated since the dimensions of the model parameters and model output spaces may not match and the model itself may not be differentiable (see Butler et al. [8] for more details). This is also explored in more detail with the proposed SMC method below.

3.3.5 Visualization

This simple physical problem and solutions from the three inverse problem formulations are shown in Fig. 1. Only the data-consistent formulation results in a posterior that, when transformed into conductivity space, matches the measurement density f_{obs}. This is by design of the formulation. In contrast, the conditional probability and intersecting probability formulations both effectively update the measurement density during posterior estimation (Eq. 10). Though the posterior obtained by Tarantola & Valette's [12] formulation is invariant under a change of variables, it does not match the measurement distribution in conductivity space. Finally, note that both the semi-analytical and ABC solutions to the conditional probability formulation are
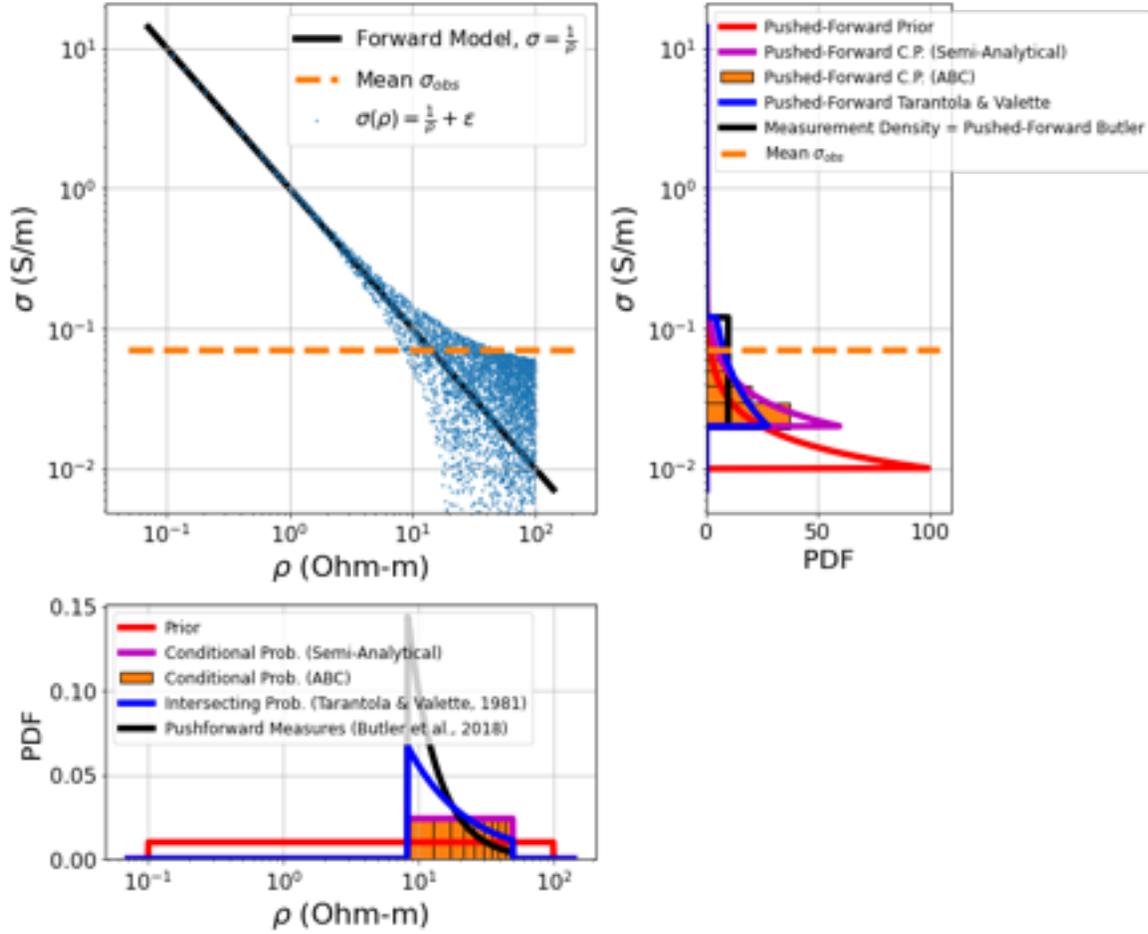
Figure 1. Solutions to the inverse problem of inferring resistivity, ρ, from uncertain measurements of conductivity, σ, from three different inverse problem formulations. (bottom) Prior and posteriors from the three formulations. The conditional probability formulation was solved both semi-analytically and through ABC sampling. (top left) The forward model used to transform between ρ and σ, plus samples from the joint density f(ρ,σ) corresponding to the conditional probability formulation. (top left) Pushed-forward prior and pushed-forward posteriors from the bottom panel. Only the data-consistent formulation reproduces the measurement density.

effectively the same. Therefore, if the expression $\mathbf{d} = g(\mathbf{m}) + \varepsilon$ is used in any form, then the posterior will not transform into the measurement density, except for a few special cases described below.

4. Solving the Data-Consistent Problem with Sequential Monte Carlo

The most common approach to solving the data-consistent inverse problem involves weighted KDE and rejection sampling [8, 18]. Mattis et al. [19] combined the weighted KDE approach with Principal Component Analysis (PCA) to reduce the dimensionality of the data

vector, making it easier to apply KDE to f_{push}. However, these density estimation approaches require special care when the dimensionality increases [29] and when the model is an imperfect match to the data [18].

Here we propose an algorithm for solving the data-consistent inverse problem for subsurface modeling problems which involve high-dimensional data and imperfect prior models. SMC [13] and Simulated Annealing [30] provide a practical framework for searching high-dimensional model parameter spaces and manifold learning [31] allows us to project noisy, real-world data onto smoother and lower-dimensional data spaces.

The SMC approach starts by initializing the prior density as the proposal density at iteration $t = 0$. Then, we sample the proposal, push the samples through the forward model g, weight the samples, and fit a generative probabilistic model to the updated density (Fig. 2). Then, the updated density is set as the proposal for the next iteration, $t \leftarrow t + 1$, and the process is repeated. In the large sample limit, this style of algorithm converges quickly with the number of sampling iterations [32]. Important advantages of this SMC approach include scaling well with parallelization and no rejection step, making it feasible for computationally expensive forward models.
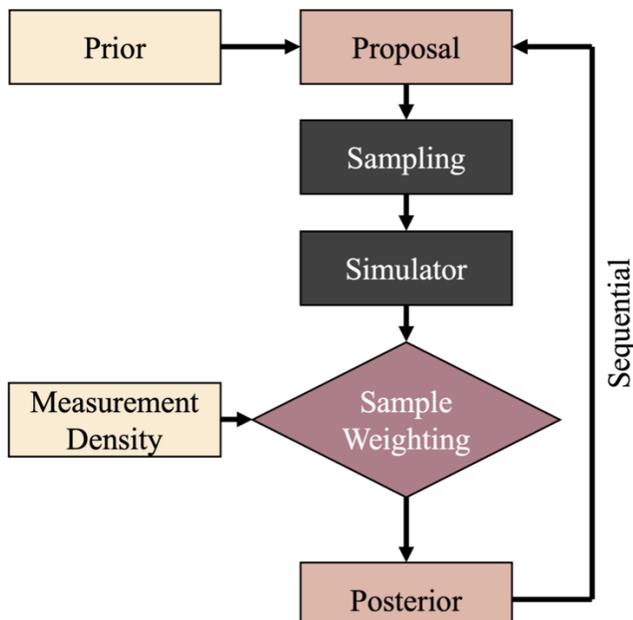
Figure 2. Workflow for the proposed SMC algorithm. The posterior is updated by training a generative probabilistic model from the proposal samples and their weights.

To alleviate particle degeneration, where most proposal samples have zero weight, we use an approach similar to

simulated annealing [30]. We artificially increase the observation variance at the initialization stage, then gradually reduce it to the true observation variance at the last iteration. In the subsequent sections, we will describe each step of this algorithm in more detail and discuss how modern machine learning techniques can help us address these challenges.

Estimating the SMC weights at each iteration requires estimating and evaluating high-dimensional probability density functions. Similar to the approach of [19], we use a combination of dimension reduction and density estimation. The density estimation can be accomplished with a variety of approaches, though we will focus on KDE and normalizing flows (NF), a neural density estimation method [21].

4.1 Initializing the proposal and measurement densities

The initial proposal is analogous to the prior distribution from the conditional probability formulation. Just as in the conditional probability approach, it is essential that the initial proposal can reproduce the observed data [8]. In other words, the proposal must be unfalsified [11]. This initial proposal is only used for sampling and does not need to have a tractable probability density.

The measurement density represents the observed data, plus measurement error. Measurement error can take many forms. It may be estimated from repeat measurements, based on system tolerances, or assumed or estimated from noisy signals [29]. Frequently, measurement error is assumed to be Gaussian [10]. Non-Gaussian error models are also compatible as long as the measurement density can be sampled.

4.2 Sampling the proposal and forward simulating data

SMC begins by randomly sampling the initial proposal N_{prop} times. The number of samples required for accurate posterior estimation will depend on the magnitude of measurement error, nonlinearity in the forward model, and dimensionality. Generally, smaller measurement error, higher nonlinearity, and higher dimensionality problems will require more samples for accurate posterior estimation.

The proposal samples are then pushed-forward through the forward model. If the forward model is computationally expensive, this step is parallelizable using modern distributed computing resources. Even models that take minutes or hours to simulate per sample can be feasibly inverted because of this parallelization. Once the proposal samples are pushed-forward through g, they make up samples of f_{push}.

4.3 Dimension Reduction, Density Estimation, and Sample Weighting

Because geophysical data are high-dimensional and g is often nonlinear, the high-dimensional pushforward probability density f_{push} is difficult to estimate for most problems. Therefore, techniques for high-dimensional density estimation [33, 34] and dimension reduction [11] are crucial for advancing applications of the data-consistent formulation. Though high-dimensional density estimation methods using deep learning [33, 34], we will focus on dimension reduction. In the context of this work, dimension reduction methods create a deterministic mapping $S: \Omega_{\mathbf{d}} \rightarrow \Omega_{\mathbf{s}}$ where $\Omega_{\mathbf{s}}$ is the reduced-dimension space. With dimension reduction, Eq. 4 is modified to

$$\hat{f}_{post}(\mathbf{m}) = \frac{\hat{f}_{obs}\left(S(g(\mathbf{m}))\right)}{\hat{f}_{push}\left(S(g(\mathbf{m}))\right)} f_{prop}(\mathbf{m}) \tag{11}$$

where $\hat{f}_{post} \approx f_{post}$, $\hat{f}_{obs}$ is obtained by pushing samples from f_{obs} forward through S, and $\hat{f}_{push}$ is obtained by pushing samples from f_{push} forward through S. After $\hat{f}_{obs}$ and $\hat{f}_{push}$ are sampled, the densities can be estimated using a variety of methods, including kernel density estimation [35], mixture density networks [36], or normalizing flows [21, 34].

Though the formulation in Eq. 11 can use any deterministic function S, the choice of S may impact the posterior. It is well-known that dimension reduction increases entropy [37]. However, S may also introduce bias into the posterior if it is not selected carefully. Figure 3 shows conceptually how S can introduce bias. If S is a one-to-one mapping (Fig. 3a), then there is no bias and no information loss. PCA falls into this category if 100% of the variance is preserved during transformation. If S is not one-to-one, but the information loss is evenly distributed (Fig. 3b) then there will be information loss but no bias. However, if the information loss by S is unevenly distributed over $\Omega_{\mathbf{d}}$ then there is a risk of biased posterior distributions (Fig. 3c).

Finally, once $\hat{f}_{obs}$ and $\hat{f}_{push}$ are estimated, weights can be calculated for each proposal sample $\mathbf{m}^{(i)}$ according to
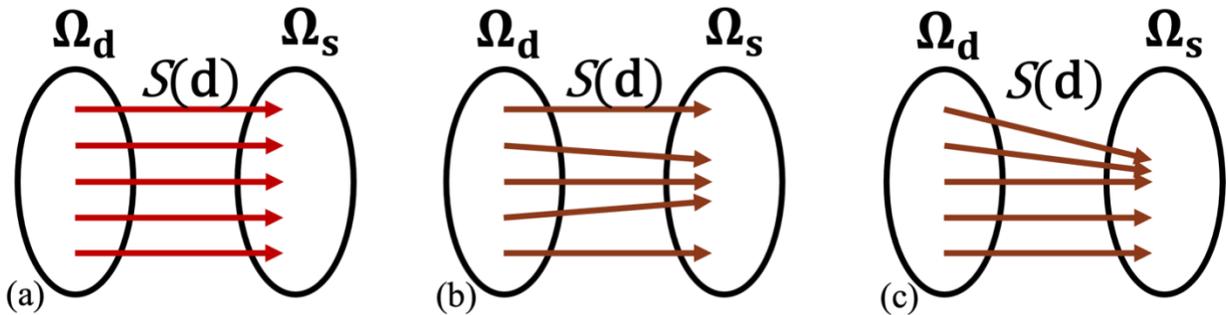
Figure 3. Conceptual diagram of how the dimension reduction function S may impact posterior information loss and bias. (a) No information loss and no bias, (b) Information loss but no bias, (c) both information loss and bias.

$$w\left(\mathbf{m}^{(i)}\right) = \frac{\hat{f}_{obs}\left(S\left(g(\mathbf{m}^{(i)})\right)\right)}{\hat{f}_{push}\left(S(g(\mathbf{m}^{(i)}))\right)}. \tag{12}$$

These weights can be used for a variety of purposes, including partial rejection control [14], bootstrapping [38], or generative model training [39].

4.4 Proposal Updating & Iterative resampling

After the first iteration, many samples will have near-zero weights. This phenomenon is known as particle degeneracy [14]. To overcome this, we will train a new generative probabilistic model using the samples and their weights. Examples of generative models include KDE, Gaussian mixtures, and normalizing flows. We modify Eq. 11 for sequential updating using iteration superscripts, where t is the current iteration and T the maximum number of iterations, that is $0 \le t \le T$. Under this framework, Eq. 11 is written as

$$\hat{f}_{prop}^{t+1}(\mathbf{m}) = \frac{\hat{f}_{obs}^{t}\left(S^{t}(g(\mathbf{m}))\right)}{\hat{f}_{push}^{t}\left(S^{t}(g(\mathbf{m}))\right)} \hat{f}_{prop}^{t}(\mathbf{m}) \tag{13}$$

All four densities and the dimension reduction function change each iteration. After iteration T we will assume $\hat{f}_{post} = \hat{f}_{prop}^{T}$. Calvetti et al. [32] showed this algorithm can converge in as little as three iterations, though it will vary problem to problem. Convergence occurs when the weights $w^{t}\left(\mathbf{m}^{(i)}\right) = 1$ for all $\mathbf{m}^{(i)}$.

4.5 Simulated Annealing

Particle weight degeneracy can make fitting an accurate proposal density difficult. Therefore, we propose an approach analogous to simulated annealing, which is a popular in the field of Bayesian optimization [30, 40]. Simulated annealing uses a time-varying proposal jump size with Markov Chain Monte Carlo (MCMC) sampling to explore a parameter space. The jump size is gradually cooled to zero to obtain the optimal solution to the inverse problem. This cooling schedule allows the chain to explore large parts of the parameter space at the beginning of the chain, then allows for small-scale refinements at the end of the chain.

The simulated annealing analogy can be extended to SMC by creating a cooling schedule for the observation variance. At the start of the cooling schedule the observation variance is artificially inflated to reduce particle degeneracy during sample weighting. Then, the observation variance is reduced at the start of each iteration. At the last iteration, the true observation variance is used. For example, if σ^2_{true} is the true observation variance, then the cooling schedule can be defined by $\sigma^2(t) = \alpha^2(t)\,\sigma^2_{true}$ where $\sigma^2(t)$ is the observation variance used at time t and α^2 is a variance factor. If $\alpha \geq 1$ for all t then $\alpha^2(t)$ is an appropriate cooling schedule.

4.6 Important Algorithmic Choices

The basic SMC algorithm in Sects. 3.1-3.4 presents several algorithmic choices that may affect the estimated posterior. First, the choice of proposal distribution may affect the posterior if, for example, the proposal distribution uses Gaussian assumptions. KDE and Gaussian mixtures are two examples of proposals that use Gaussian assumptions. Second, the dimension reduction method S can also have an impact. For example, if S is a poorly-trained regression model then the weights may be unreliable and result in a biased posterior. The density estimation method for estimating $\hat{f}_{obs}$ and $\hat{f}_{push}$ must also be accurate. If a simulated annealing approach is

used, then the cooling schedule needs be set carefully to minimize particle degeneracy at each iteration. And of course, the number of SMC iterations and samples per iteration must be large enough to ensure convergence.

5. A Synthetic Gravity Example

In this example we use a classic example of measuring the gravity signal over a spherical geobody buried in a homogeneous subsurface (Fig. 4). If the ground surface is flat, then the vertical gravity anomaly, g_z in mgal, can be calculated as a function of the surface position, x, by

$$g_z(x) = \frac{4}{3} G\pi \frac{a^3 \rho z}{(x^2 + z^2)^{3/2}}$$

where G is the gravitational constant, a is the sphere radius, ρ is the mass density of the sphere, and z is the depth to the center of the sphere [41]. We restrict the domain from $-200 \leq x \leq 200$ m, with measurements taken every 20 m. Then, we assume we have some measurements of this anomaly, and want to estimate the properties of the sphere a, ρ, and z.
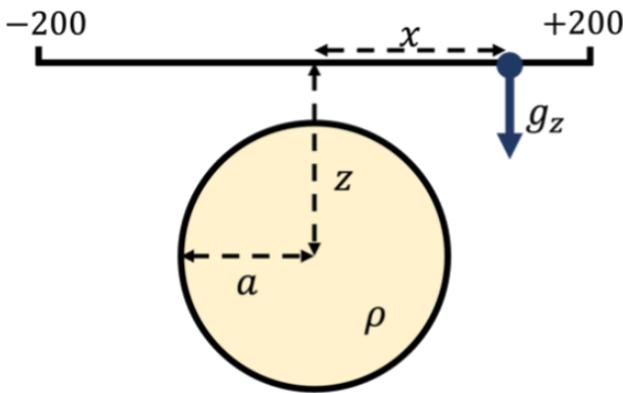
Figure 4. Setup for the synthetic gravity example.

5.1 Initialization, PCA Dimension Reduction, and Proposal Updating

The initial proposal density uses uniform densities over each parameter, with parameters independent of one another. The density is $\rho = U(1000, 7000)$ kg/m³, radius is $a = U(70, 200)$ m, and depth is $z = U(25,$

200) m. The measurement density was created by first simulating the gravity response to a sphere with $\rho = 6000$ kg/m^3, $a = 100$ m, and $z = 125$ m. Then, the measurement density is a multivariate Gaussian distributed with mean equal to the simulated data and diagonal covariance such that measurements are uncorrelated and have standard deviations equal to 10% of the simulated signal.

For each iteration, $N_{prop} = N_{obs} = 10{,}000$. Proposal samples outside the initial prior were rejected. We use PCA for the dimension reduction method. In each iteration we use the N_{prop} samples from f_{push} and train a PCA model to reduce the dimensionality of $\mathbf{d}$. Only two principal components are needed to explain $> 99\%$ of the variance in the simulated data. Kernel density estimation is used to estimate $\hat{f}_{push}$ and $\hat{f}_{obs}$ from the data samples pushed-forward by PCA into $\Omega_\mathbf{s}$. Then Eq. 12 is used to estimate weights for each proposal sample and these weights are used to train an updated KDE proposal for the next iteration. The KDE bandwidth is empirically calculated using the method of Scott [42] implemented within the SciPy Python package [43]. We used four SMC iterations with a cooling schedule of $\alpha = [10, 5, 1, 1]$. The results are shown in Fig. 5.

5.2 Results

The updating process for this gravity example is illustrated in Fig. 6. After one iteration, the proposal no longer resembles the initial uniform distribution and instead has learned correlations between model parameters. After two iterations, the proposal has been refined further. At the start of the fourth iteration the measurement density overlaps closely with the measurement density, in both the full and reduced-dimension data space.
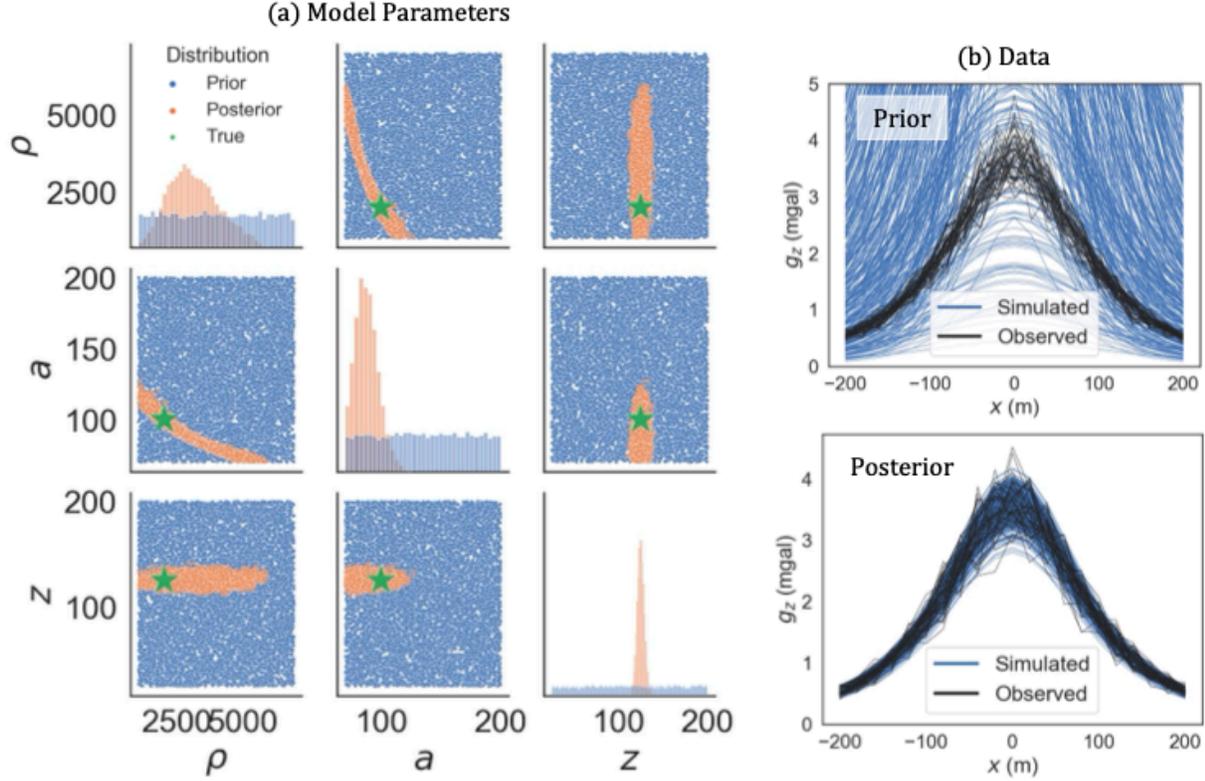
Figure 5. (a) 10,000 model parameter samples from the uniform prior distribution (blue) and data-consistent posterior distribution (orange), compared to the true value (green star) used to generate the posterior. (b) 200 samples from the pushed-forward prior (blue, top) and the pushed-forward posterior (blue, bottom) compared to the measurement distribution (black, both).

5.3 Effect of Sample Size

The sample size can have a major effect on the quality of the posterior distribution. To illustrate some of these effects, we modify the above experiment slightly. We use the same measurement density, KDE proposal, and dimension reduction techniques. However, we only use 2 SMC iterations with a cooling schedule of $\alpha = [2.5, 1]$. With this setup, we run the SMC algorithm three different N_{prop}. The first uses $N_{prop} = 10,000$, the second uses $N_{prop} = 1000$, and the third uses $N_{prop} = 100$. The results are shown in Fig. 7.

With $N_{prop} = 10,000$ the posterior is nearly identical to the one from the original experiment (Fig. 5). With $N_{prop} = 1000$ the model parameter posterior loses some of the
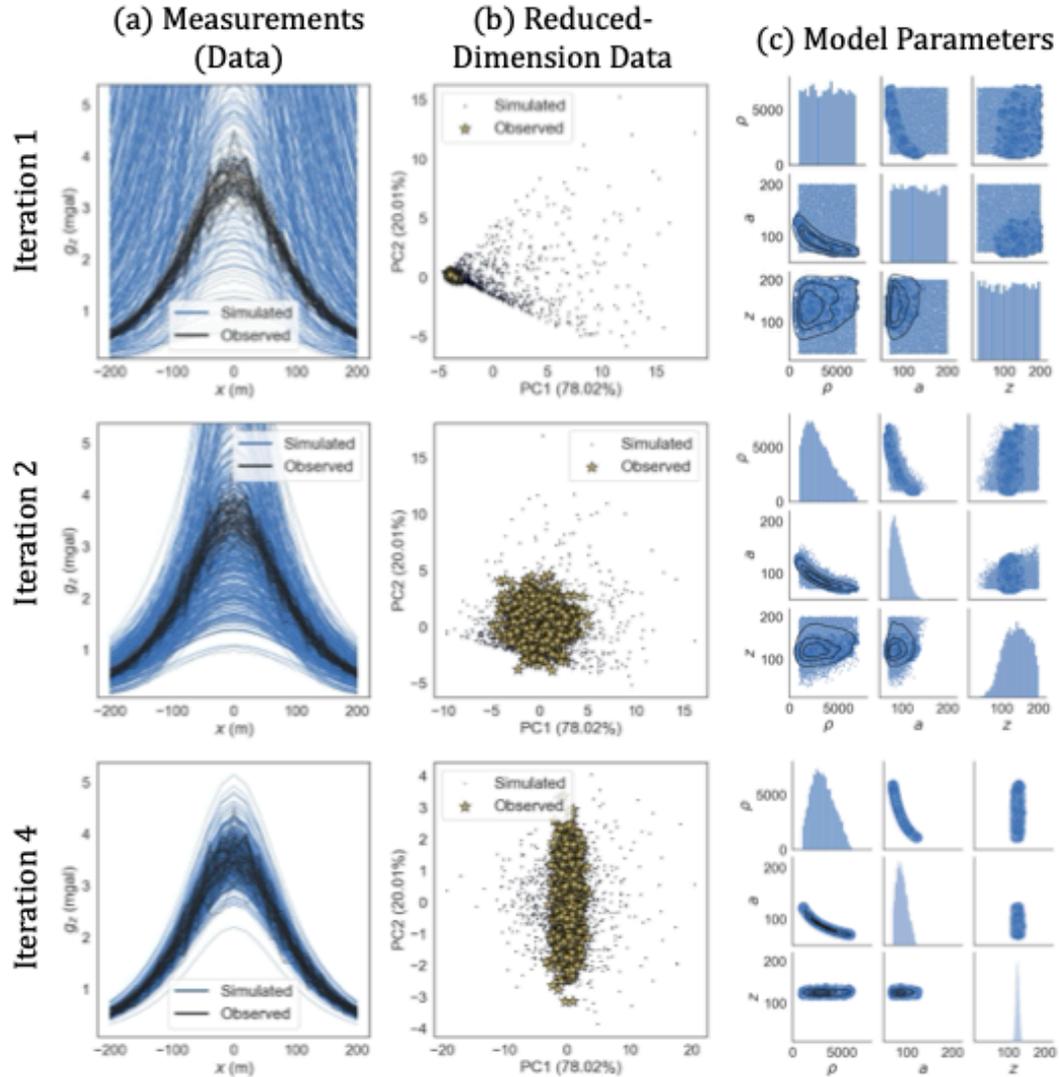
Figure 6. Three iterations of the SMC algorithm on the gravity example. (left) 1000 realizations of d from the proposal. (middle) Realizations of d after being projected into a 2D space using PCA. (right) Cross plots of proposal samples. Larger dots have larger weights. Black contours are from KDE applied to the weighted samples.

curvature, and the posterior data variance is larger than the observed variance. With $N_{prop} = 100$ we converge to a local minimum, and the posterior data cariance is smaller than the observed variance.

This experiment demonstrates some of the risks with the SMC approach if it is improperly tuned. Complicated, curving shapes may be difficult to capture if there are too few samples. Convergence to local minima may also be a concern.

5.4 Using Masked Autoregressive Flows for Density Estimation
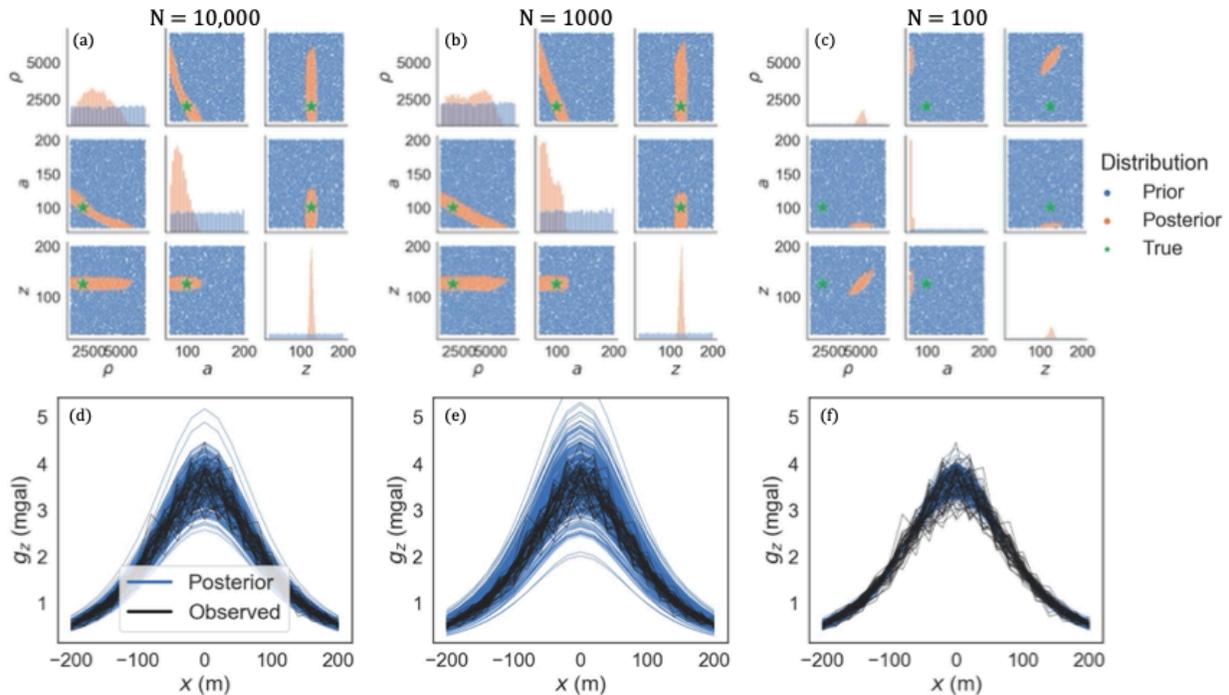
Figure 7. Illustrating the effect of measurement error on SMC performance. When measurement error is low, the SMC algorithm may converge to a small part of the posterior, underestimating uncertainty. However, this local convergence still produces samples consistent with f_{obs}.

To illustrate the modularity of this SMC approach, we also present results from using different density and proposal estimators. Masked autoregressive flows (MAF) [44], can be used to estimate non-Guassian probability density functions. MAF models are based on normalizing flows [21], where a base Gaussian distribution is transformed into a non-Gaussian distribution through successive bijective transformations known as flows. This technique allows for both density estimation and sampling. MAF models use artificial neural networks to learn the transformations needed to fit samples from the target distribution.

We implemented MAF using the *normflows* library [45] which is built on top of *PyTorch* [46]. MAF was used to estimate $\hat{f}_{push}$ and $\hat{f}_{prop}^{t}$ since these densities are non-Gaussian in this problem. KDE was used to estimate $\hat{f}_{obs}^{t}$ because it is Gaussian in this problem.
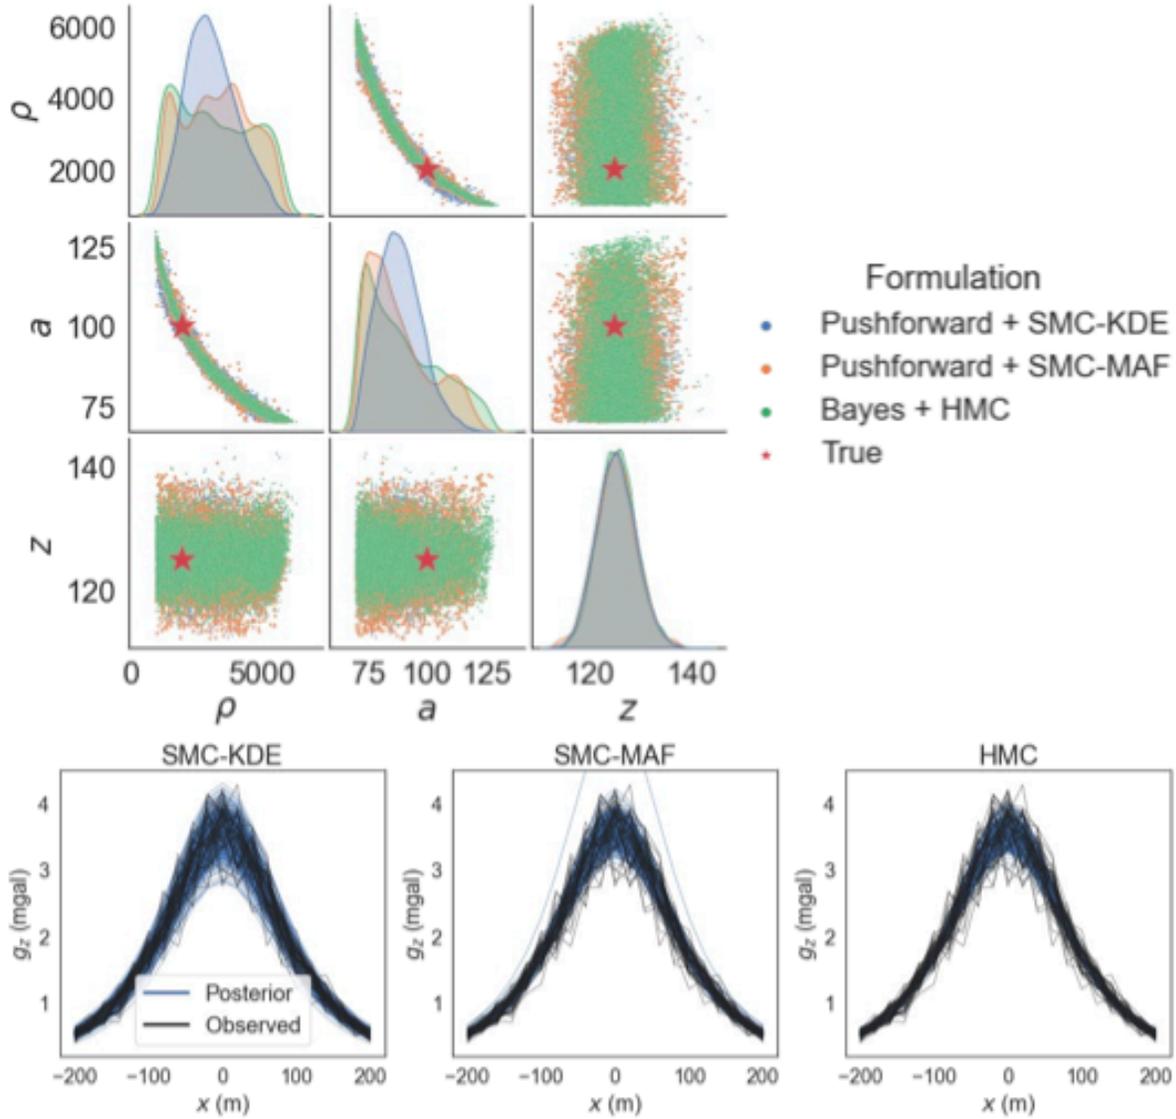
Figure 8. Comparison of the conditional probability and data-consistent formulations applied to the gravity example. (left) cross plots of posterior samples. Orange is samples from the optimized proposal, and green is samples from rejection sampling from the optimized proposal. (right) Pushed-forward posteriors for the samples shown on the left.

The MAF architecture for $\hat{f}_{push}$ consisted of three separate flows, each with two hidden layers with 16 neurons per hidden layer [47]. Training was performed using the Adam optimizer [48] with batch sizes of 5000 samples, a learning rate of 2×10^{-5}, and 5000 training iterations. The MAF architecture for $\hat{f}_{prop}^{t}$ is the same as $\hat{f}_{push}$. Training was performed using the Adam optimizer with batch sizes of 8500 samples, a learning rate of 2×10^{-5}, and 10,000 training

iterations. Results are shown in Fig. 8 along with the pure KDE results and a comparison to MCMC sampling, discussed in the subsequent section.

5.5 Comparison to Conditional Probability Formulation with MCMC Sampling

We compared the posterior from the data-consistent approach to the standard conditional probability approach with MCMC. We used the Hamiltonian Monte Carlo (HMC) algorithm No-U-Turn Sampling [22]. The prior is the same set of independent uniform densities used as the initialized proposal in the SMC approach. We used the measurement density as the likelihood. Then, we sampled 10,000 realizations from the posterior using the code *pymc3* [49]. The initial sample set to the true value and the sampling acceptance rate was 93.6%.

Figure 8 shows a comparison between the conditional probability formulation with HMC sampling and the data-consistent formulation with SMC sampling. There are three main observations from this comparison. First, the MAF and HMC marginal distributions are very similar, showing that for this example the problem formulation has minimal impact on the posterior. Second, we observe that the KDE model is skewed towards the middle. This is most likely due to edge effects in the Gaussian kernel. Finally, note that the MAF model has several samples that are inconsistent with the correct solution. These samples are rare enough that they do not impact the posterior significantly, but they do show that the MAF solution has longer tails than the ideal solution, which would map perfectly onto the measurement density.

6. Discussion

The main advantages of the data-consistent formulation are (1) there is no likelihood, and (2) it is compatible with very noisy measurements and highly nonlinear models. Furthermore, the

proposed SMC algorithm scales well with parallelization and can leverage machine learning technologies for posterior inference. A similar SMC algorithm can also be used to solve problems in the traditional formulations [16].

There is no intractable likelihood in the data-consistent formulation because this formulation does not use joint densities and assumes the forward model is perfectly deterministic (although recent extensions allow for the model to be stochastic [18]). The tradeoff is that, instead of an intractable likelihood, the high-dimensional density f_{push} must be estimated.

The data-consistent formulation is attractive for noisy measurements and highly nonlinear models because the formulation is designed such that the posterior, when pushed forward through g, equals the measurement density (Eq. 1). When measurement error is very small or the forward model is linear, then the choice of formulation does not appear to matter much. However, as the measurement error becomes larger and the forward model becomes more nonlinear, the difference between the formulations becomes larger. If the goal of solving the inverse problem is to obtain a posterior that pushes forward into the measurement density, then the data-consistent formulation is preferred.

Finally, we observe that the SMC algorithm is a natural fit for solving the data-consistent formulation. Since there is no Markov chain, forward model simulations can be run in parallel on modern cluster computing systems. The main challenge with the SMC solution is calculating weights, which requires estimating a high-dimensional density f_{push}. However, sampling high-dimensional f_{prior} and estimating high-dimension f_{push} may be challenging.

7. Conclusions

There are several inverse problem formulations available for geophysical applications. The conditional probability formulation is the most popular, but we demonstrate that this formulation may yield unintuitive posterior distributions for nonlinear problems constrained by noisy measurements. When measurements are very noisy, it may be advantageous to use the data-consistent formulation, which finds a posterior that pushes forward into the measurement distribution. This data-consistent formulation can also be obtained from the conditional probability formulation by assuming the observed and simulated data are conditionally independent given the model parameters. It can also be obtained from the intersecting probability density formulation by assuming the posterior density of the data are equal to the measurement density. One way to solve this data-consistent formulation for geophysical applications is to use SMC algorithms, which scale well with parallelization. Machine learning can also be used as a nonlinear dimension reduction tool. The challenges with solving the data-consistent formulation include high-dimensional density estimation and sampling high-dimensional spaces. However, the two examples we show illustrate that this method can work for geophysical problems and warrants future research and development.

Appendix A. Derivation of the Data-Consistent Formulation

This derivation is modified from Butler et al. [8]. Let us assume we have two measurable spaces, $\left(\Omega_{\mathbf{m}}, \mathcal{B}(\Omega_{\mathbf{m}})\right)$ and $\left(\Omega_{\mathbf{d}}, \mathcal{B}(\Omega_{\mathbf{d}})\right)$ with reference measures $\lambda_{\mathbf{m}}$ and $\lambda_{\mathbf{d}}$ on those spaces, respectively. Now let us assume we also have probability measures $\mathbb{P}_{prior}$ and $\mathbb{P}_{post}$ defined on $\left(\Omega_{\mathbf{m}}, \mathcal{B}(\Omega_{\mathbf{m}})\right)$ and $\mathbb{P}_{obs}$ defined on $\left(\Omega_{\mathbf{d}}, \mathcal{B}(\Omega_{\mathbf{d}})\right)$. We also have a forward model $g: \Omega_{\mathbf{m}} \rightarrow \Omega_{\mathbf{d}}$ that results in the pushed-forward prior probability measure $g_* \mathbb{P}_{prior}$ on $\left(\Omega_{\mathbf{d}}, \mathcal{B}(\Omega_{\mathbf{d}})\right)$. Assuming all

the probability measures are absolutely continuous with respect to the reference measures, we

obtain the following probability density functions:

$$f_{prior} = \frac{d\mathbb{P}_{prior}}{d\lambda_{\mathbf{m}}}; \quad f_{push} = \frac{d(g_*\mathbb{P}_{prior})}{d\lambda_{\mathbf{d}}}; \quad f_{obs} = \frac{d\mathbb{P}_{obs}}{d\lambda_{\mathbf{d}}}; \quad f_{post}$$

$$= \frac{d\mathbb{P}_{post}}{d\lambda_{\mathbf{m}}}$$

(A.1)

We formulate the problem to enforce $g_*\mathbb{P}_{post} = \mathbb{P}_{obs}$. Now assume $A_{\mathbf{d}} \in \mathcal{B}(\Omega_{\mathbf{d}})$ is an

arbitrary set. We use the Radon-Nikodym theorem and the substitution rule (Bogachev, 2000) to

obtain

$$g_*\mathbb{P}_{post}(A_{\mathbf{d}}) = \int_{\mathbf{d} \in A_{\mathbf{d}}} f_{obs}(\mathbf{d}) \, d\lambda_{\mathbf{d}}$$

$$= \int_{\mathbf{d} \in A_{\mathbf{d}}} \frac{f_{obs}(\mathbf{d})}{f_{push}(\mathbf{d})} \, d(g_*\mathbb{P}_{prior})$$

$$= \int_{\mathbf{m} \in g^{-1}(A_{\mathbf{d}})} \frac{f_{obs}(g(\mathbf{m}))}{f_{push}(g(\mathbf{m}))} \, d\mathbb{P}_{prior}$$

(A.2)

$$= \int_{\mathbf{m} \in g^{-1}(A_{\mathbf{d}})} \frac{f_{obs}(g(\mathbf{m}))}{f_{push}(g(\mathbf{m}))} f_{prior}(\mathbf{m}) \, d\lambda_{\mathbf{m}}$$

Now we apply the Radon-Nikodym theorem and substitution rule [22] to $g_*\mathbb{P}_{post}(A_{\mathbf{d}})$ to

obtain

$$g_*\mathbb{P}_{post}(A_{\mathbf{d}}) = \mathbb{P}_{post}(g^{-1}(A_{\mathbf{d}})) = \int_{\mathbf{m} \in g^{-1}(A_{\mathbf{d}})} f_{post}(\mathbf{m}) \, d\lambda_{\mathbf{m}}$$

(A.3)

Finally we substitute Eq. A.3 into Eq. A.2 and drop the integrals to obtain the probability density form of $g_* \mathbb{P}_{post} = \mathbb{P}_{obs}$.

$$f_{post}(\mathbf{m}) = \frac{f_{obs}(g(\mathbf{m}))}{f_{push}(g(\mathbf{m}))} f_{prior}(\mathbf{m})$$

$$(A.4)$$

Appendix B. Derivation of the Data-Consistent Formulation from Conditional Probability

Let us start with the assumption that $\mathbf{d} = g(\mathbf{m}) = \mathbf{d}_{obs} + \varepsilon$ where ε is the measurement error term. Now, we marginalize the posterior $f(\mathbf{m}|\mathbf{d}_{obs})$ with respect to $\mathbf{d}$:

$$f(\mathbf{m}|\mathbf{d}_{obs}) = \int_{\Omega_{\mathbf{d}}} f(\mathbf{m}|\mathbf{d}_{obs}, \mathbf{d}) f(\mathbf{d}|\mathbf{d}_{obs}) \, \mathrm{d}\mathbf{d} \qquad (B.1)$$

Next, we assume if we know the forward-simulated data $\mathbf{d}$, then the observed data provide no additional information. In other words, $\mathbf{m}$ and $\mathbf{d}_{obs}$ are conditionally independent given $\mathbf{d}$, or $f(\mathbf{m}|\mathbf{d}_{obs}, \mathbf{d}) = f(\mathbf{m}|\mathbf{d})$. We use this assumption to modify Eq. B.1 and expand $f(\mathbf{m}|\mathbf{d})$ using Bayes' theorem.

$$f(\mathbf{m}|\mathbf{d}_{obs}) = \int_{\Omega_{\mathbf{d}}} f(\mathbf{m}|\mathbf{d}) f(\mathbf{d}|\mathbf{d}_{obs}) \, \mathrm{d}\mathbf{d}$$

$$= \int_{\Omega_{\mathbf{d}}} \frac{f(\mathbf{d}|\mathbf{m}) f(\mathbf{m})}{f(\mathbf{d})} f(\mathbf{d}|\mathbf{d}_{obs}) \, \mathrm{d}\mathbf{d}$$

$$(B.2)$$

Since the forward model is deterministic, $f(\mathbf{d}|\mathbf{m})$ is a Dirac delta with the form $\delta(\mathbf{d} - g(\mathbf{m}))$. This is substituted into Eq. B.2 to obtain

$$f(\mathbf{m}|\mathbf{d}_{obs}) = \int_{\Omega_\mathbf{d}} \frac{\delta(\mathbf{d} - g(\mathbf{m}))\, f(\mathbf{m})}{f(\mathbf{d})}\, f(\mathbf{d}|\mathbf{d}_{obs})\, d\mathbf{d} \tag{B.3}$$

which simplifies to

$$f(\mathbf{m}|\mathbf{d}_{obs}) = \frac{f(g(\mathbf{m})|\mathbf{d}_{obs})}{f(g(\mathbf{m}))}\, f(\mathbf{m}) \tag{B.4}$$

This expression has the same interpretation as Eq. 4.

Appendix C. Derivation of the Data-Consistent Formulation from Tarantola & Valette

Let us start by assuming the posterior is a function of both the prior information and observed data. Under Tarantola & Valette's formulation of the inverse problem [11], the posterior is written as

$$f_{post}(\mathbf{m}, \mathbf{d}) = k\, \frac{f_{prior}(\mathbf{m}, \mathbf{d})\, f_{obs}(\mathbf{m}, \mathbf{d})}{\mu(\mathbf{m}, \mathbf{d})} \tag{C.1}$$

where k is a normalization constant and μ is the homogeneous probability density. Next, we assume there are no observations of $\mathbf{m}$ so we can write $f_{obs}(\mathbf{m}, \mathbf{d}) = f_{obs}(\mathbf{d})\, \mu(\mathbf{m})$ and we also assume $\mu(\mathbf{m}, \mathbf{d}) = \mu(\mathbf{m})\, \mu(\mathbf{d})$. Under these assumptions Eq. C.1 becomes

$$
\begin{aligned}
f_{post}(\mathbf{m}, \mathbf{d}) &= k\, \frac{f_{prior}(\mathbf{m}|\mathbf{d})\, f_{prior}(\mathbf{d})\, f_{obs}(\mathbf{d})\, \mu(\mathbf{m})}{\mu(\mathbf{m})\, \mu(\mathbf{d})} \\[2mm]
&= k\, \frac{f_{prior}(\mathbf{m}|\mathbf{d})\, f_{prior}(\mathbf{d})\, f_{obs}(\mathbf{d})}{\mu(\mathbf{d})} \\[2mm]
&= f_{prior}(\mathbf{m}|\mathbf{d})\, f_{post}(\mathbf{d})
\end{aligned}
\tag{C.2}
$$

where the last step is because, by Tarantola & Valette's definition, $f_{post}(\mathbf{d}) = k\, f_{prior}(\mathbf{d})\, f_{obs}(\mathbf{d})\, /\, \mu(\mathbf{d})$. Next we expand $f_{prior}(\mathbf{m}|\mathbf{d})$ using Bayes' theorem and apply the fact that $f_{prior}(\mathbf{d}|\mathbf{m}) = \delta\big(\mathbf{d} - g(\mathbf{m})\big)$ since g is deterministic. Equation C.2 therefore can be expressed as

$$
\begin{aligned}
f_{post}(\mathbf{m}, \mathbf{d}) &= \frac{f_{prior}(\mathbf{d}|\mathbf{m})\, f_{prior}(\mathbf{m})}{f_{prior}(\mathbf{d})}\, f_{post}(\mathbf{d}) \\[2mm]
&= \frac{\delta\big(\mathbf{d} - g(\mathbf{m})\big)\, f_{prior}(\mathbf{m})}{f_{prior}(\mathbf{d})}\, f_{post}(\mathbf{d})
\end{aligned}
\tag{C.3}
$$

Finally, we obtain the marginal posterior by integrating Eq. C.3 over $\Omega_{\mathbf{d}}$ to find

$$
\begin{aligned}
f_{post}(\mathbf{m}) &= \int_{\Omega_{\mathbf{d}}} \frac{\delta\big(\mathbf{d} - g(\mathbf{m})\big)\, f_{prior}(\mathbf{m})}{f_{prior}(\mathbf{d})}\, f_{post}(\mathbf{d})\, \mathrm{d}\mathbf{d} \\[2mm]
&= \frac{f_{post}(g(\mathbf{m}))}{f_{prior}(g(\mathbf{m}))}\, f_{prior}(\mathbf{m})
\end{aligned}
\tag{C.4}
$$

which is the same form as Eq. 4.

Data Availability Statement

The code and datasets generated during the current study are available at

https://github.com/ammilten/PushforwardMeasureGeophysics.

Funding

Competing Interests

The authors have no competing interests to declare that are relevant to the content of this article.

References

[1] Tarantola, A.: Inverse Problem Theory and Methods for Model Parameter Estimation. Society for Industrial and Applied Mathematics (2005).

[2] Dentith, M., Mudge, S.T.: Geophysics for the Mineral Exploration Geoscientist. Cambridge University Press (2014).

[3] Linde, N., Renard, P., Mukerji, T., Caers, J.: Geological realism in hydrogeological and geophysical inverse modeling: A review. Adv. Water Resour. 86, 86–101 (2015). https://doi.org/10.1016/j.advwatres.2015.09.019

[4] Chang-Chun, Y., Xiu-Yan, R., Yun-He, L., Yan-Fu, Q., Chang-Kai, Q., Jing, C.: Review on airborne electromagnetic inverse theory and applications. Geophys. 80, W17–W31 (2015).

[5] Ghil, M.: A century of nonlinearity in the geosciences. Earth Space Sci. 6, 1007–1042 (2019).

[6] Fernández-Martínez, J.L., Fernández-Muñiz, Z., Pallero, J.L.G., Pedruelo-González, L.M.: From Bayes to Tarantola: New insights to understand uncertainty in inverse problems. J. Appl. Geophys. 98, 62–72 (2013).

[7] Grana, D., Passos de Figueiredo, L., Azevedo, L.: Uncertainty quantification in Bayesian inverse problems with model and data dimension reduction. Geophys. 84, M15–M24 (2019).

[8] Butler, T., Jakeman, J., Wildey, T.: Combining push-forward measures and Bayes' Rule to construct consistent solutions to stochastic inverse problems. SIAM J. Sci. Comput. 40, A984–A1011 (2018).

[9] Scheidt, C., Li, L., Caers, J.: Quantifying Uncertainty in Subsurface Systems. Wiley (2018).

[10] Malinverno, B.A., Briggs, V.A.: Expanded uncertainty quantification in inverse problems: hierarchical Bayes and empirical Bayes. Geophys. 69, 1005–1016 (2004).

[11] Scheidt, C., Jeong, C., Mukerji, T., Caers, J.: Probabilistic falsification of prior geologic uncertainty with seismic amplitude data: application to a turbidite reservoir case. Geophys. 80, M89–M112 (2015).

[12] Tarantola, A., Valette, B.: Inverse problems = quest for information. J. Geophys. 50, 159–170 (1981).

[13] Doucet, A., De Freitas, N., Gordon, N.J.: Sequential Monte Carlo Methods in Practice. Springer (2001).

[14] Sisson, S.A., Fan, Y., Tanaka, M.M.: Sequential Monte Carlo without likelihoods. Proc. Natl. Acad. Sci. 104, 1760-1765 (2007).

[15] Amaya, M., Linde, N., Laloy, N.: Adaptive sequential Monte Carlo for posterior inference and model selection among complex geological priors. Geophys. J. Int. 226, 1220–1238 (2021). https://doi.org/10.1093/gji/ggab170

[16] Peters, G.W., Fan, Y., Sisson, S.A.: On sequential Monte Carlo, partial rejection control and approximate Bayesian computation. Stat. Comput. 22, 1209–1222 (2012).

[17] Mosegaard, K.: Quest for consistency, symmetry, and simplicity — the legacy of Albert Tarantola. Geophys. 76, W51–W61 (2011).

[18] Butler, T., Wildey, T., Yen, T.Y.: Data-consistent inversion for stochastic input-to-output maps. Inverse Probl. 36, (2020b).

[19] Mattis, S.A., Steffen, K.R., Butler, T., Dawson, C.N., Estep, D.: Learning quantities of interest from dynamical systems for observation-consistent inversion. Comput. Methods Appl. Mech. Eng. 388 (2022).

[20] Pilosov, M., del-Castillo-Negrete, C., Yen, T.Y., Butler, T., Dawson, C.: Parameter Estimation with Maximal Updated Densities. Comput. Methods Appl. Mech. Eng. (accepted January 2023).

[21] Papamarikos, G., Nalisnick, E., Rezende, D.J., Mohamed, S., Lakshminarayanan, B.: Normalizing Flows for Probabilistic Modeling and Inference. J. Mach. Learn. Res. 22 (2021).

[22] Hoffman, M., Gelman, A.: The no-u-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. arXiv (2014). https://doi.org/10.48550/arXiv.1111.4246

[23] Stuart, A.M.: Inverse Problems: a Bayesian Perspective. Acta Numerica 19 (2010).

[24] Bogachev, V.I.: Measure Theory. Springer (2000).

[25] Mosegaard, K., Sambridge, M.: Monte Carlo analysis of inverse problems. Inverse Probl. 18, R29 (2002).

[26] Tarantola, A.: Popper, Bayes and the inverse problem. Nat. Phys. 2, 492–494 (2006). https://doi.org/10.1038/nphys375

[27] Debski, W.: Probabilistic inverse theory. In: Advances in Geophysics. Elsevier (2010).

[28] Marin, J.M., Pudlo, P., Robert, C.P., Ryder, R.J.: Approximate Bayesian computational methods. Stat. Comput. 22 (2012).

[29] Terrell, G.R., Scott, D.W.: Variable Kernel Density Estimation. Ann. Stat. 20, 1236-1265 (1992).

[30] Kirkpatrick, S., Gelett, C.D., Vecchi, M.P.: Optimization by simulated annealing. Science 220, 621–630 (1983).

[31] Ma, Y., Fu, Y.: Manifold Learning Theory and Applications. CRC Press, Boca Raton (2012).

[32] Calvetti, D., Dunlop, M., Somersalo, E., Stuart, A.: Iterative updating of model error for Bayesian inversion. Inverse Probl. 34, (2018).

[33] Liu, Q., Xu, J., Jiang, R., Wong, W.H.: Density estimation using deep generative neural networks. Proc. Natl. Acad. Sci. 118, e2101344118 (2021).

[34] Kobyzev, I., Prince, S.J.D., Brubaker, M.A.: Normalizing flows: an introduction and review of current methods. IEEE Trans. Pattern Anal. Mach. Intell. 43, 3964–3979 (2020).

[35] Wang, L., Kitanidis, P.K., Caers, J.: Hierarchical Bayesian inversion of global variables and large-scale spatial fields. Water Resour. Res. 58, e2021WR031610 (2022).

[36] Shahraeeni, M.S., Curtis, A.: Fast probabilistic nonlinear petrophysical inversion. Geophys. 76, E45-E58 (2011).

[37] Lindley, D.V.: On a measure of the information provided by an experiment. Ann. Math. Stat. 27, 986–1005 (1956).

[38] Efron, B.: Bootstrap methods: another look at the jackknife. Ann. Statist. 7, 1–26 (1979).

[39] Song, S., Mukerji, T., Hou, J.: GANSim: conditional facies simulation using an improved progressive growing of generative adversarial networks (GANs). Math Geosci 53, 1413–1444 (2021).

[40] Aarts, E., Korst, J., Michiels, W.: Simulated annealing. In: Burke, E.K., Kendall, G. (eds.) Search Methodologies. Springer (2005).

[41] Blakely, R.J.: Potential Theory in Gravity and Magnetic Applications. Cambridge University Press (1995).

[42] Scott, D.W.: Multivariate Density Estimation: Theory, Practice, and Visualization. John Wiley & Sons, New York, Chichester (1992).

[43] Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P.: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. Nat. Methods 17, 261-272 (2020).

[44] Papamarikos, G., Pavlakou, T., Murray, I.: Masked Autoregressive Flow for Density Estimation. In: Proceedings of the 31st Conference on Neural Information Processing Systems.

[45] Stimper, V., Liu, D., Campbell, A., Berenz, V., Ryll, L., Schölkopf, B., and Hernández-Lobato J.M.: normflows: A PyTorch Package for Normalizing Flows. https://github.com/VincentStimper/normalizing-flows.

[46] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library. arXiv (2019). https://doi.org/10.48550/ARXIV.1912.01703

[47] Germain, M., Gregor, K., Murray, I., Larochelle, H.: MADE: Masked Autoencoder for Distribution Estimation. In: Proceedings of the 32nd International Conference on Machine Learning, 37, 881-889.

[48] Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv (2017). https://doi.org/10.48550/arXiv.1412.6980

[49] Salvatier, J., Wiecki, T.V., Fonnesbeck, C.: Probabilistic programming in Python using PyMC3. J. Comput. Sci. 2, e55 (2016).