



Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

Full Length Article

Bayesian inverse problems with conditional Sinkhorn generative adversarial networks in least volume latent spaces

 Qiuyi Chen ^{a,b}, Panagiotis Tsilifis ^{b,*}, Mark Fuge ^{a,c}
^a Center for Risk and Reliability, Department of Mechanical Engineering, University of Maryland, College Park, 20742, MD, USA

^b Artificial Intelligence Group, General Electric Vernova Advanced Research, One Research Circle, Niskayuna, 12309, NY, USA

^c Department of Mechanical and Process Engineering, ETH Zürich, Rämistrasse 101, Zürich, 8092, Switzerland

ARTICLE INFO

Keywords:

 Conditional generative adversarial networks
 Least volume autoencoders
 Sinkhorn divergence
 Unsupervised dimensionality reduction
 Bayesian inference

ABSTRACT

Solving inverse problems in scientific and engineering fields has long been intriguing and holds great potential for many applications, yet most techniques still struggle to address issues such as high dimensionality, nonlinearity and model uncertainty inherent in these problems. Recently, generative models such as Generative Adversarial Networks (GANs) have shown great potential in approximating complex high dimensional conditional distributions and have paved the way for characterizing posterior densities in Bayesian inverse problems, yet the problems' high dimensionality and high nonlinearity often impedes the model's training. In this paper we show how to tackle these issues with Least Volume—a novel unsupervised nonlinear dimension reduction method—that can learn to represent the given datasets with the minimum number of latent variables while estimating their intrinsic dimensions. Once the low dimensional latent spaces are identified, efficient and accurate training of conditional generative models becomes feasible, resulting in a latent conditional GAN framework for posterior inference. We demonstrate the power of the proposed methodology on a variety of applications including inversion of parameters in systems of ODEs and high dimensional hydraulic conductivities in subsurface flow problems, and reveal the impact of the observables' and unobservables' intrinsic dimensions on inverse problems.

1. Introduction

Engineering design and optimization tasks almost always entail procedures of inference of quantities of interest or design parameters from indirect measurements of observables that are typically the outcome of a physical experiment or the output of a computational model (Tarantola, 2005). In the majority of cases, the observations available to the engineer are limited and the dependence of the design or input parameters on the observed output is described by a highly nonlinear physical process, making the inference task and the quantification of uncertainty around the input a nontrivial process. From a theoretical standpoint, numerous works in the literature have been dedicated to characterizing the well-posedness of the problem and determining the conditions under which a solution is feasible and can be numerically approximated. Such conditions vary from defining appropriate regularizers on a misfit function (Calvetti & Somersalo, 2018; Engl, Hanke, & Neubauer, 1996), to ensuring Lipschitz conditions on the log-likelihood function in a Bayesian context (Cotter, Dashti, & Stuart, 2010; Stuart, 2010). From an applications perspective, major challenges are typically the computationally expensive forward model

involved in the experiment, resulting in only a limited number of available observations and the high dimensionality of both input parameters and observables that makes it difficult for the inference algorithm to converge to a solution. Applications where such issues manifest include ocean dynamics (Herbei, McKeague, & Speer, 2008), seismic inversion (Russell, 1988), numerical weather forecasting (Ekblom, Tuppi, Shemyakin, Haario, & Jarvinen, 2020; Ollinaho et al., 2014), ice-sheet flow (Isaac, Petra, Stadler, & Ghattas, 2015; Petra, Martin, Stadler, & Ghattas, 2014; Petra, Zhu, Stadler, Hughes, & Ghattas, 2012) and large chaotic dynamical systems (Springer et al., 2021).

When following a Bayesian paradigm, traditional techniques for inverse problems are highly associated with commonly performing Bayesian inference, that is the task of exploring a conditional distribution associated with our *posterior* knowledge about the unknown parameters, after observations are collected. Several variations of Markov Chain Monte Carlo (MCMC) techniques (Bilonis & Zabaras, 2013; Haario, Laine, Lehtinen, Saksman, & Tamminen, 2004; Marzouk & Najm, 2009; Tsilifis & Ghanem, 2018) have provided efficient ways of exploring complex posterior densities in high dimensional settings or

* Corresponding author.

E-mail address: panagiotis.tsilifis@governova.com (P. Tsilifis).
<https://doi.org/10.1016/j.neunet.2025.107740>

Received 6 June 2024; Received in revised form 9 March 2025; Accepted 7 June 2025

Available online 24 June 2025

0893-6080/© 2025 Elsevier Ltd. All rights reserved, including those for text and data mining, AI training, and similar technologies.

on manifolds. Alternatives methods to MCMC include Variational Inference (Broderick, Boyd, Wibisono, Wilson, & Jordan, 2013; Hoffman, Blei, Wang, & Paisley, 2013; Tsilifis, Billionis, Katsounaros, & Zabarar, 2016), where posteriors approximated by choosing an optimal density from a family of distributions, therefore transforming the inference task from a sampling strategy to an optimization algorithm. Due to the fact that both approaches require a tremendous amount of computational effort, constructing surrogates has been one of the most common ways to accelerate the learning procedure of the posterior density. Popular choices for surrogate modeling are Polynomial Chaos (Marzouk, Najm, & Rahn, 2007; Marzouk & Xiu, 2009; Tsilifis & Ghanem, 2017; Tsilifis, Ghanem, & Hajali, 2017; Tsilifis et al., 2019; Tsilifis, Papaioannou, Straub, & Nobile, 2020), Gaussian Processes (Bilionis & Zabarar, 2012; Kennedy & O’Hagan, 2001; Tsilifis et al., 2021) or adaptive sparse grid collocation (Ma & Zabarar, 2009b).

The continuously rising momentum of deep learning in engineering sciences has resulted in significant contributions in the solution of inverse problems in medical imaging and compressed sensing (Fan & Ying, 2020; Jin, McCann, Froustey, & Unser, 2017; Mardani et al., 2018), geosciences and reservoir engineering (Laloy, Héroult, Jacques, & Linde, 2018; Mo, Zabarar, Shi, & Wu, 2019, 2020), aerodynamic design (Chen, Wang, Pope, Chen, & Fuge, 2022; Ghosh et al., 2022; Tsilifis, Ghosh, & Andreoli, 2022) and other areas. The use of “deep” architectures, referring to the number of layers in an (artificial) neural network involved in modeling the physical process or the inference procedure, allows for theoretically capturing arbitrarily high nonlinearities and enabling accurate learning of the functional representations, provided that sufficient amounts of data are available during training. In a forward uncertainty propagation setting, deep learning architectures have demonstrated great potential in modeling solution to elliptic PDEs (Tripathy & Billionis, 2018), flow through porous media (Zhu, Zabarar, Koutsourelakis, & Perdikaris, 2019) and turbulence (Geneva & Zabarar, 2019). Multi-fidelity settings have also been explored using physics informed deep learning (Meng, Babaee, & Karniadakis, 2021) or transfer learning (Kontolati et al., 2023), while inverse problems have been tackled by applying Bayesian *physics-informed* neural networks B-PINNs for modeling posterior densities (Yang, Meng, & Karniadakis, 2021).

Deep Generative or *pushforward* models is one of the most prominent fields of machine learning techniques for modeling conditional densities and addressing inverse modeling. Inspired by the seminal works in adversarial learning (Goodfellow et al., 2014), variational auto-encoders (VAE) (Kingma & Welling, 2014) and normalizing flows (Dinh, Krueger, & Bengio, 2015; Dinh, Sohl-Dickstein, & Bengio, 2017), generative models have been modified to their conditional counterparts in order to account for condition variables and thus generate samples from conditional distributions learned from observed data. The key idea behind conditioning in generative models is to incorporate the observable quantity as a conditioning variable in the neural network architecture that describes generative-discriminative process in generative adversarial nets (GANs), the encoder–decoder framework in a VAE setting or in the coupling layers in a normalizing flow. Conditional VAEs have already found applications in molecular design (Lee & Min, 2022), image reconstruction (Zhang, Barbano, & Jin, 2021), while an novel intrusive approach has also been developed for PDE-constrained inverse problems (Tait & Damoulas, 2020). Perhaps one of the main drawbacks with cVAEs is that the latent space and its dimensionality needs to be specified a priori by the experimenter and choosing the dimension to be low enough in order to achieve computational efficiency, while maintaining descriptive accuracy can be a hit-or-miss situation. Normalizing flows on the other hand, do not aim at achieving dimensionality reduction. The latent (normalized) variable is required to preserve the dimensionality of the original input variable in order to ensure invertibility in the mapping as guaranteed by the Independent Component Analysis theorem (Hyvärinen & Oja, 2000), which makes it inefficient in very high dimensional settings. Interesting

applicability of normalizing flows and conditional variations of it can be found in radiology (Dasgupta, Patel, Ray, Johnson, & Oberai, 2024), astrophysics (Haldemann et al., 2023), subsurface flow (Padmanabha & Zabarar, 2021) and industrial gas turbine airfoil design (Tsilifis et al., 2022). At last, GANs have played their own pivotal role in advancing computational efficiency in inverse modeling. In Patel, Ray, and Oberai (2022), the authors use a GANs for modeling the prior knowledge in a Bayesian and using it as a dimension reduction technique, performing the inference in the latent space. Other variations of conditional GANs have shown interesting potential in speech enhancement (Qin & Jiang, 2018) using the Wasserstein distance is used for solving the minmax problem. Additionally, a Lipschitz condition on the critic has been shown to further improve performance in parameter inference of PDEs (Ray, Murgoitio-Esandi, Dasgupta, & Oberai, 2023), resulting a loss function regularized by a gradient penalty.

In this paper, we introduce the *latent conditional Sinkhorn GANs* (*latent CSGANs*), which break down the approximation of the target posterior distribution into two consecutive steps that not only increase efficiency but also provide useful information about the posterior. Specifically, the proposed framework first learns the low dimensional structure of the target posterior’s support by encoding the target posterior’s input and output to latent spaces, via invertible mappings, and *automatically* reducing the latent spaces’ dimensions as much as possible. Then, it minimizes the Sinkhorn divergence in order to approximate the posterior with a pushforward generative model in these low dimensional latent spaces, thus, accelerating and stabilizing training by alleviating the curse of dimensionality and by replacing the problematic minimax game from vanilla GANs. We use two high dimensional examples to show that this two-stage posterior inference paradigm can approximate complex posteriors accurately and efficiently. We also reveal the important correlation between the input condition’s intrinsic dimension—estimated by the latent space’s dimension—and the uncertainty around the posterior solution, which can allow us to assess the posterior’s complexity even before retrieving it.

The paper is structured as follows: The proposed methodology is presented in detail in Section 2. More specifically, in Section 2.1 we introduce the novel *least volume autoencoders* that are used for dimensionality reduction as a special case of generic autoencoders (Section 2.1.1) with a regularization term in their loss function (Section 2.1.2) that accounts for minimizing the volume of the data when projected in the latent space. The resulting formulation and optimization problem are stated in Section 2.1.3. Next, the conditional Sinkhorn GANs are introduced in Section 2.2. Section 2.2.1 introduced the generic concept of pushforward models while Section 2.2.2 introduced the type of GANs trained with the Sinkhorn divergence incorporated into their loss function. Then, Section 2.2.3 introduced the conditional counterpart of the Sinkhorn GANs and at last, Section 2.3 describes how the latter is combined with the least volume autoencoder in order to explore the inverse problem solution on the reduced dimension latent space. The proposed methodology is illustrated with two numerical examples presented in Section 3 where first (Section 3.1) we apply the approach on a simple inversion problem consisting of inferring the initial conditions of an ODE system from observing a partial signal of its solution and next (Section 3.2) we challenge our method on the high-dimensional inversion problem of log-permeability inference on a two-phase flow oil reservoir based on data that characterized the oil saturation of the domain the oil production curves are specific production wells located throughout the domain boundaries. Finally, we summarize our findings and propose future research directions in the conclusions section Section 4.

2. Methodology

The approximation of high dimensional conditional distribution can be decomposed into two consecutive tasks—the retrieval of the low dimensional manifold structure that supports the data distribution,




Fig. 1. Flattening the latent set via Least Volume (Chen & Fuge, 2024).

and the learning of the conditional distribution restricted to this low dimensional manifold. Correspondingly, we first introduce the *Least Volume Autoencoder* that can *automatically* capture the low dimensionality and nonlinearity of the data manifold and *invertibly* transform the data into a low dimensional latent space for efficient analysis. Thereafter, we present the *conditional Sinkhorn generative adversarial network* (CSGAN)—a generative model based on optimal transport—for approximating and sampling from complicated conditional distributions. At last, we show how to integrate them to learn conditional distributions in high dimensional spaces efficiently.

2.1. Least volume autoencoders

In this starting section, we introduce the manifold learning method that can retrieve the dataset’s low dimensional nonlinear structure and concisely represent it to facilitate downstream applications.

2.1.1. Autoencoders

Assume \mathcal{X} is a set of points in the *data space* X that we denote as the *dataset* $\mathcal{X} \subset X$, following some probability distribution with density $p_{\mathcal{X}}(x)$. An *autoencoder* is an ordered pair (e, g) consisting of two functions, namely an *encoder* $e : X \rightarrow Z$ and a *decoder* $g : Z \rightarrow X$, that are typically modeled using neural networks and are trained by minimizing a reconstruction loss function

$$J(\theta) := \mathbb{E}_{x \sim p_{\mathcal{X}}(x)} [d(g_{\theta} \circ e_{\theta}(x), x)] \quad (1)$$

with respect to the parameters θ on g and e , where Z is the *latent space* and $d : X \times X \rightarrow \mathbb{R}$ is a user-specified distance function. Despite its simple formulation, allowing g and e to be *continuous* functions with arbitrary complexity can reduce the reconstruction error to 0, thus implying that $g = e^{-1}$ and e becomes a *topological embedding* of the dataset $\mathcal{X} \subseteq X$, such that the *latent set* $\mathcal{Z} := e(\mathcal{X}) \subseteq Z$ is homeomorphic to \mathcal{X} (Chen & Fuge, 2024). Consequently, given the trained autoencoder, if we want to learn a continuous function $f : \mathcal{X} \rightarrow Y$ (or $f : Y \rightarrow \mathcal{X}$), we can equivalently learn its latent counterpart $f' : \mathcal{Z} \rightarrow Y$ (or $f' : Y \rightarrow \mathcal{Z}$) and recover f via $f = f' \circ e \Leftrightarrow f' = f \circ g$ (or $f = g \circ f' \Leftrightarrow f' = e \circ f$). This equivalence is particularly beneficial when the learning cost of f does not scale well with its dimensionality but $\dim Z \ll \dim X$, as learning f' in Z significantly mitigates the curse of dimensionality.

Therefore, retrieving a Z of dimension as small as possible is important for high dimensional problems. In what follows, we show how to obtain a continuous autoencoder (e, g) that learns to reduce $\dim Z$ as much as possible.

2.1.2. Least volume regularization

Least Volume (LV) (Chen & Fuge, 2024) is the method that can help us obtain such an autoencoder with a least dimensional latent space Z . Its key intuition is that a flat piece of paper consumes less space than when it curls. Likewise, among all the latent sets homeomorphic to the dataset \mathcal{X} , a curved latent set $\mathcal{Z} = g(\mathcal{X}) \subset Z$ can only be enclosed by a cuboid of much larger volume than a cuboid that encloses its flattened counterpart (as shown in Fig. 1). Suppose in the following *conceptual* procedure we are given an autoencoder (e_{θ}, g_{θ}) of great complexity such that both $e_{\theta} : X \rightarrow Z$ and $g_{\theta} : Z \rightarrow X$ can approximate any continuous function. We keep the reconstruction loss (1) minimized so that \mathcal{Z} remains homeomorphic to \mathcal{X} , while additionally morphing \mathcal{Z} ’s geometry by adjusting θ . We denote the standard deviation of each latent space dimension by

$$\sigma_{\mathcal{Z},i} := \sqrt{\text{var}(\mathcal{Z})_i}, \quad (2)$$

and make the cuboid’s side lengths equal to $\sigma_{\mathcal{Z},i}$, so that the cuboid has *volume* $L_{\mathcal{Z}} := \prod_{i=1}^{\dim Z} \sigma_{\mathcal{Z},i}$. Now if we reduce $L_{\mathcal{Z}}$, it reaches zero only when at least one $\sigma_{\mathcal{Z},i}$ vanishes, indicating that \mathcal{Z} is flattened and compressed into a linear subspace, i.e., $\mathcal{Z} = \mathcal{Z}' \times c_1 \subset Z'_1 \times c_1 \subset Z$ where c_1 is a constant. Naturally, the way to proceed would be to extract Z'_1 as the new latent space of lower dimension, try compressing $\mathcal{Z}'_1 \subset Z'_1$ inside into a new subspace $\mathcal{Z}'_2 \times c_2 \subset Z'_1$ again (i.e., $\mathcal{Z}'_1 = \mathcal{Z}'_2 \times c_2 \subset Z'_2 \times c_2 \subset Z'_1$) by minimizing $L_{\mathcal{Z}'_1}$, and repeat the same extraction and compression recursively for n times until \mathcal{Z}'_n cannot be flattened any more: that is, its volume $L_{\mathcal{Z}'_n}$ cannot reach zero. The final Z'_n is then the least dimensional latent space Z' that we seek, which can be extracted by the *projection* $\pi : Z' \times c \rightarrow Z'$ where $c := \prod_{i=1}^n c_i$ is the cartesian product. Here π is continuous and in fact is also a homeomorphism (Chen & Fuge, 2024). It follows that $(\pi \circ e_{\theta}, g_{\theta} \circ \pi^{-1})$ is the desired continuous autoencoder that maps \mathcal{X} to the least dimensional latent space Z' , in which $\mathcal{Z}' = \pi \circ e_{\theta}(\mathcal{X}) \subset Z'$ is still homeomorphic to \mathcal{X} .

In practice this conceptual recursive procedure is difficult to implement, so in Chen and Fuge (2024) the authors further introduced an increment $\eta > 0$ and instead minimize the geometric mean

$$L_{\mathcal{Z},\eta} = \sqrt[\dim Z]{\prod_{i=1}^{\dim Z} (\sigma_{\mathcal{Z},i} + \eta)} \quad (3)$$

in order to conveniently achieve a similar outcome, while avoiding vanishing gradient issue when $\sigma_{\mathcal{Z},i}$ goes to zero or exploding gradient issue when $\dim Z$ is large. In this case, optimization can be carried all the way through with no recursion as it allows the standard deviation across more than one dimension to reach zero simultaneously ($\sigma_{\mathcal{Z},i} \approx 0$




Fig. 2. Pushforward models fit $\mathbb{P}_g := g_* \mathbb{P}_u$ to \mathbb{P}_r over \mathcal{X} by minimizing $D(\mathbb{P}_g, \mathbb{P}_r)$.

for multiple i 's). Nevertheless, reducing $L_{\mathcal{Z}, \eta}$ with no further constraints will trivially drive all $\sigma_{\mathcal{Z}, i}$ close to zero without learning anything useful, because \mathcal{Z} may shrink *isotropically* to naively shorten all $\sigma_{\mathcal{Z}, i}$ without flattening itself. To avoid this, the decoder g_θ is constrained to be K -Lipschitz through *spectral normalization* (Gouk, Frank, Pfahringer, & Cree, 2021; Miyato, Kataoka, Koyama, & Yoshida, 2018) for a small K . By doing so, any latent dimension i with $\sigma_{\mathcal{Z}, i} \simeq 0$ must also have close-to-zero variation in the original data space X and thus can only correspond to a trivial data dimension, whereas a principal data dimension must retain a large variation in the latent space to avoid violating the Lipschitz constraint. Hence the isotropic shrinkage is prevented. A detailed proof of this property is given in Appendix A. Fig. 1 illustrates this intuition.

Thereafter, one may safely *prune* these latent dimensions i with $\sigma_{\mathcal{Z}, i} \simeq 0$ in ascending order of $\sigma_{\mathcal{Z}, i}$, until the reconstruction loss $J(\theta)$ increases to a certain user-defined threshold δ . By pruning dimension i we mean replacing that dimension's code $z_i^{(j)} = e_\theta(x^{(j)})_i$ of every sample $x^{(j)}$ with its mean $\bar{z}_i = \mathbb{E}_{x \sim p_{\mathcal{X}}(x)}[e_\theta(x)_i]$. It is safe to do so because the ℓ_2 reconstruction error of the autoencoder with a K -Lipschitz decoder increases by at most $K \sqrt{\sum_{i \in P} \sigma_{\mathcal{Z}, i}^2}$, where P is the set of indices of pruned dimensions (Chen & Fuge, 2024). After pruning, \mathcal{Z} then has the form $\mathcal{Z} = \mathcal{Z}' \times \prod_{i \in P} \bar{z}_i \subset \mathcal{Z}' \times \prod_{i \in P} \bar{z}_i$ and can be fed to $\pi : \mathcal{Z}' \times \prod_{i \in P} \bar{z}_i \rightarrow \mathcal{Z}'$ to extract \mathcal{Z}' . If we denote the composition of the encoder with this pruning operation by $\bar{e}_{\theta, \delta}$, then the autoencoder becomes $(\pi \circ \bar{e}_{\theta, \delta}, g_\theta \circ \pi^{-1})$.

2.1.3. Loss function

In summary, the *Least Volume Autoencoder (LVAE)* that we want to obtain is the $(\pi \circ \bar{e}_{\theta, \delta}, g_\theta \circ \pi^{-1})$ whose parameter θ is the solution to the following constrained optimization problem:

$$\begin{aligned} \arg \min_{\theta} L_{\mathcal{Z}, \eta}(\theta), \quad L_{\mathcal{Z}, \eta}(\theta) &:= L_{e_\theta(x), \eta} \\ \text{s.t.} \quad \mathcal{Z} &\in \{ \mathcal{Z} = e_\theta(\mathcal{X}) \mid \theta \text{ minimizes } J(\theta) \}, \\ &\|g_\theta(z_1) - g_\theta(z_2)\| \leq K \|z_1 - z_2\|, \forall z_1, z_2 \in \mathcal{Z}. \end{aligned} \quad (4)$$

Yet enforcing the reconstruction constraint is hard for nonlinear autoencoders. So in practice, we resort to a simpler optimization problem: that is, to minimize the objective function

$$\ell(\theta) := J(\theta) + \lambda L_{\mathcal{Z}, \eta}(\theta), \quad (5)$$

where λ needs to be fine-tuned to make $J(\theta) < \delta$. The reconstruction loss J throughout this work is set to be $J(\theta) = \mathbb{E}_{x \sim p_{\mathcal{X}}(x)} [\|g_\theta \circ e_\theta(x) - x\|]$.

Many of LVAE's nice properties have been shown in Chen and Fuge (2024). In particular, when the encoder and decoder are represented by linear functions $e(x) = Ax + a$ and $g(z) = Bz + b$, with g being a 1-Lipschitz and $\text{rank cov}(\mathcal{X}) \geq \dim \mathcal{Z}$, the LVAE is equivalent to Principal Component Analysis (PCA) (Pearson, 1991) and recovers the principal components of \mathcal{X} . Empirically, the PCA-like importance ordering effect persists even when e and g are nonlinear, such that we can assess the importance of the latent dimensions i with their standard deviations $\sigma_{\mathcal{Z}, i}$. We shall leverage this property later in the experiment section Section 3.

2.2. Conditional sinkhorn generative adversarial networks

Next, we give a brief review to the *conditional Sinkhorn generative adversarial networks* (CSGANs) that we are going to use to solve the inverse problems by approximating their posterior distributions on a data-driven basis, based on the formulation introduced in Feydy et al. (2019), Genevay, Peyré, and Cuturi (2018).

2.2.1. Pushforward generative models

Generative adversarial networks (GANs) (Arjovsky, Chintala, & Bottou, 2017; Goodfellow et al., 2014; Gulrajani, Ahmed, Arjovsky, Dumoulin, & Courville, 2017) is a subfamily of a larger family of generative models dubbed *pushforward models* (Brown, Caterini, Ross, Cresswell, & Loaiza-Ganem, 2022; Cornish, Caterini, Deligiannidis, & Doucet, 2020; Salmona, De Bortoli, Delon, & Desolneux, 2022) that approximate a target probability measure \mathbb{P}_r supported by a dataset \mathcal{X} with a *pushforward measure* \mathbb{P}_g after minimizing a certain *statistical divergence* (or *distance*) $D(\mathbb{P}_g, \mathbb{P}_r)$ between them. The pushforward measure \mathbb{P}_g is created by transforming a probability measure \mathbb{P}_u —usually referred to as the *latent distribution*—in a latent space U through a measurable *generator* function $g : U \rightarrow X$ (typically modeled with a neural network) into the data space X , as illustrated in Fig. 2, such that

$$\mathbb{P}_g(A) = g_* \mathbb{P}_u(A) := \mathbb{P}_u(g^{-1}(A)) \quad (6)$$

for any $A \subset X$, $B \subset U$ that are \mathbb{P}_g - and \mathbb{P}_u -measurable sets respectively such that $A = g(B)$. For the ease of sampling $x \sim \mathbb{P}_g$ —i.e., first sampling $u \sim \mathbb{P}_u$ then producing $x = g(u)$ —the latent distribution \mathbb{P}_u is often a simple probability measure, e.g., a Gaussian or uniform distribution. Thanks to g 's high complexity (Cybenko, 1989), \mathbb{P}_g can approximate many intricate \mathbb{P}_r in X .

Different choices of the statistical divergence D lead to different types of pushforward models. Among the most common, variational autoencoders (Kingma & Welling, 2014; Sohn, Lee, & Yan, 2015) and normalizing flows (Dinh et al., 2015, 2017; Kingma & Dhariwal, 2018; Kobyzev, Prince, & Brubaker, 2020) maximize the log likelihood of the data samples, which is equivalent to minimizing the *Kullback–Leibler divergence* $D_{KL}(\mathbb{P}_r, \mathbb{P}_g)$. On the other hand, GANs mostly minimize statistical divergences other than D_{KL} , such as the *Jensen–Shannon divergence* for vanilla GANs (Goodfellow et al., 2014), the f -divergence for f -GANs (Nowozin, Cseke, & Tomioka, 2016) and the *Wasserstein distance* $\text{OT}(\mathbb{P}_r, \mathbb{P}_g)$ for Wasserstein GANs (Arjovsky et al., 2017; Gulrajani et al., 2017). More detailed reviews on the use of different loss function has been conducted in Hong, Hwang, Yoo, and Yoon (2019), Jabbar, Li, and Omar (2021), Yi, Walia, and Babyn (2019).

2.2.2. Sinkhorn GANs

The Optimal Transport cost is defined as

$$\text{OT}(\mathbb{P}_r, \mathbb{P}_g) := \min_{\Pi \in \Pi(\mathbb{P}_r, \mathbb{P}_g)} \mathbb{E}_{x, \hat{x} \sim \Pi_{X, \hat{X}}} [c(x, \hat{x})] \quad (7)$$

where $\Pi(\mathbb{P}_r, \mathbb{P}_g)$ is the family of joint distributions $\Pi_{X, \hat{X}}$ whose marginal distributions equal \mathbb{P}_r and \mathbb{P}_g respectively, and c is usually a symmetric cost function. In cases where c is a distance function, such as the ℓ_1

or ℓ_2 metric, we refer to OT as the Optimal Transport distance. More specifically, when $c(x, \hat{x}) = \|x - \hat{x}\|^p$, $\text{OT}^{1/p}$ defines the Wasserstein p -distance. Due to the computational challenges associated with the numerical evaluation of OT, many researchers have turned their attention to its regularized variations, such as the ones including an *entropy regularization* term with weight ρ (Cuturi, 2013; Feydy et al., 2019; Peyré, Cuturi, et al., 2019), written as:

$$\text{OT}_\rho(\mathbb{P}_r, \mathbb{P}_g) := \min_{\mathbb{P}_{X, \hat{X}} \in \Pi(\mathbb{P}_r, \mathbb{P}_g)} \mathbb{E}_{x, \hat{x} \sim \mathbb{P}_{X, \hat{X}}} [c(x, \hat{x})] + \rho D_{KL}(\mathbb{P}_{X, \hat{X}} \| \mathbb{P}_r \times \mathbb{P}_g). \quad (8)$$

The above expression possesses strong duality and can thus be efficiently evaluated after writing its concave dual formulation:

$$\text{OT}_\rho(\mathbb{P}_r, \mathbb{P}_g) = \max_{\alpha, \beta} \mathbb{E}_{x \sim \mathbb{P}_r} [\alpha(x)] + \mathbb{E}_{\hat{x} \sim \mathbb{P}_g} [\beta(\hat{x})] - \rho \mathbb{E}_{x, \hat{x} \sim \mathbb{P}_r \times \mathbb{P}_g} [\exp(v(x, \hat{x})/\rho) - 1], \quad (9)$$

$$v(x, \hat{x}) := \alpha(x) + \beta(\hat{x}) - c(x, \hat{x}). \quad (10)$$

Here α and β are the Lagrange multipliers that can be rapidly retrieved through the Sinkhorn–Knopp algorithm (Cuturi, 2013; Feydy et al., 2019) based on coordinate ascent.

However, the above entropy regularization introduces a bias, since for $\mathbb{P}_r = \mathbb{P}_g$ we no longer have $\text{OT}_\rho(\mathbb{P}_r, \mathbb{P}_g) = 0$ (Feydy et al., 2019). To mitigate this bias, Genevay et al. further introduced the *Sinkhorn GAN*¹ (Genevay et al., 2018) that instead minimizes the *Sinkhorn divergence* $S_\rho(\mathbb{P}_r, \mathbb{P}_g)$, which is a symmetric variation of OT_ρ defined as Feydy et al. (2019), Genevay et al. (2018), Peyré et al. (2019):

$$S_\rho(\mathbb{P}_r, \mathbb{P}_g) = \text{OT}_\rho(\mathbb{P}_r, \mathbb{P}_g) - \frac{1}{2} \text{OT}_\rho(\mathbb{P}_r, \mathbb{P}_r) - \frac{1}{2} \text{OT}_\rho(\mathbb{P}_g, \mathbb{P}_g). \quad (11)$$

This distance stands out from the rest for its stability and fast evaluation that spares the costly and unstable training of a discriminator function (as in Arjovsky et al. (2017), Goodfellow et al. (2014)), low sample complexity (Genevay, Chizat, Bach, Cuturi, & Peyré, 2019), and its ability to interpolate between OT and the *maximum mean discrepancy* (MMD) (Smola, Gretton, Song, & Schölkopf, 2007) by adjusting ρ (Feydy et al., 2019).

2.2.3. Conditional Sinkhorn GANs

Suppose along with the dataset \mathcal{X} following the distribution \mathbb{P}_r^X , we are also provided with the *condition set* \mathcal{Y} endowed with the distribution \mathbb{P}_r^Y . Very often \mathcal{X} is dependent on \mathcal{Y} and we are interested in learning their relationship through their joint distribution $\mathbb{P}_r^{X,Y}$. To this end, all pushforward models introduced above can be generalized to approximate the conditional distributions $\mathbb{P}_r^{X|Y}$ —as opposed to the marginal densities \mathbb{P}_r^X —by simply augmenting the generator input such that $g : U \times Y \rightarrow X$. This implies that the corresponding pushforward measure $\mathbb{P}_g^X = g_* \mathbb{P}_u$ becomes a conditional one $\mathbb{P}_g^{X|Y=y} = g|_{\mathcal{Y}} \mathbb{P}_u$, where $g|_{\mathcal{Y}}$ denotes the restriction $g|_{U \times \mathcal{Y}} : U \rightarrow X$ for $y \in \mathcal{Y} \subset Y$. Consequently, training g now becomes enforcing $\mathbb{P}_g^{X|Y}$ to approximate the target conditional distribution $\mathbb{P}_r^{X|Y}$ by minimizing the divergence $D(\mathbb{P}_r^{X,Y}, \mathbb{P}_g^{X,Y})$ between the target joint distribution $\mathbb{P}_r^{X,Y} = \mathbb{P}_r^{X|Y} \mathbb{P}_r^Y$ and the pushforward joint distribution $\mathbb{P}_g^{X,Y} = \mathbb{P}_g^{X|Y} \mathbb{P}_u^Y$.

We will refer to the resulting conditional variation of Sinkhorn GANs—for which $D = S_\rho$ —as the *conditional Sinkhorn GANs* (CSGANs). In practice, the Sinkhorn divergence $S_\rho(\mathbb{P}_r^{X,Y}, \mathbb{P}_g^{X,Y})$ can be conveniently evaluated by sampling batches of data points $(x^{(j)}, y^{(j)}) \sim \mathbb{P}_r^{X,Y}$ and $(\hat{x}^{(j)}, \hat{y}^{(j)}) \sim \mathbb{P}_g^{X,Y}$ from the two joint distributions respectively and passing them to the Sinkhorn algorithm. $\mathbb{P}_r^{X,Y}$ can be sampled easily by

picking $(x^{(j)}, y^{(j)})$ straight out of the digital dataset $D = \{(x^{(j)}, y^{(j)})\}_n$ uniformly, whereas $\mathbb{P}_g^{X,Y}$ can be sampled by taking $y^{(j)}$ out of D uniformly (thus $y^{(j)} \sim \mathbb{P}_r^Y$) and passing them along with $u^{(j)} \sim \mathbb{P}_u$ to g to produce the predictions $\hat{x}^{(j)} = g(u^{(j)}, y^{(j)})$ and form the tuples $(\hat{x}^{(j)}, y^{(j)})$. More technical details about the implementation of the Sinkhorn algorithm can be found in Feydy (2019), Feydy et al. (2019).

2.3. CSGANs in least volume latent spaces

Although the CSGANs can approximate the conditional distributions $\mathbb{P}_r^{X|Y}$, when the data space X and the condition space Y has very high dimensionality, the CSGAN’s training can decelerate significantly, because evaluating the Sinkhorn divergence $S_\rho(\mathbb{P}_r^{X,Y}, \mathbb{P}_g^{X,Y})$ typically requires the calculation of $c((x^{(j)}, y^{(j)}), (\hat{x}^{(k)}, y^{(k)}))$ in (10) over $n \geq 1000$ pairs of (j, k) , and the evaluation of c among all pairs normally has time complexity $O(n(\dim X + \dim Y))$ (e.g., when c is the ℓ_1 or ℓ_2 distance). In addition, it is crucial to find out a reasonable latent space dimension $\dim U$ for the CSGAN’s generator $g : U \times Y \rightarrow X$. If $\dim U$ is lower than $\dim \mathcal{X}$, then due to the *rank theorem* (Lee & Lee, 2012), the generator g ’s image $g(\text{supp}(\mathbb{P}_u))$ can at best only be a small subset of measure zero in \mathcal{X} and unable to generate all points on \mathcal{X} . Consequently, a distribution supported by \mathcal{X} cannot be approximated well by the resulting pushforward measure $\mathbb{P}_g^{X|Y=y}$ (Chen, Pope, & Fuge, 2022). Moreover, the difference between $\dim \mathcal{X}$ and $\dim \mathcal{Y}$ may provide us some insights into the relative information content between X and Y , allowing us to anticipate the uncertainty of X even prior to obtaining $\mathbb{P}_r^{X|Y}$, so retrieving them could be beneficial. This hypothesis will be detailed and verified in the later experiments in Section 3.2.

To overcome these issues, we can integrate the LVAEs with the CSGANs to form the *latent CSGANs*. Specifically, suppose we obtained two trained LVAEs ($e_x : X \rightarrow Z_x, d_x : Z_x \rightarrow X$) and ($e_y : Y \rightarrow Z_y, d_y : Z_y \rightarrow Y$) for the dataset \mathcal{X} and the condition set \mathcal{Y} respectively, whose latent spaces Z_x and Z_y have their dimensions reduced as much as possible. We first transform the dataset $\mathcal{X} \subset X$ into the latent set $\mathcal{Z}_x = e_x(\mathcal{X}) \subset Z_x$, and the condition set $\mathcal{Y} \subset Y$ into the latent set $\mathcal{Z}_y = e_y(\mathcal{Y}) \subset Z_y$. Thus, the latent codes $z_x \in \mathcal{Z}_x$ and $z_y \in \mathcal{Z}_y$ jointly follow the distribution $\mathbb{P}_e^{Z_x, Z_y} := (e_x \times e_y)_* \mathbb{P}_r^{X,Y}$. After these encodings, we can train the latent CSGAN with the latent generator $g : U \times Z_y \rightarrow Z_x$ in $Z_x \times Z_y$ similarly to Section 2.2.3: that is, minimizing $S_\rho(\mathbb{P}_g^{Z_x, Z_y}, \mathbb{P}_e^{Z_x, Z_y})$ in order to make $\mathbb{P}_g^{Z_x|Z_y=z_y} := g|_{\mathcal{Z}_y} \mathbb{P}_u \simeq \mathbb{P}_e^{Z_x|Z_y=z_y}$. This enables $g|_{\mathcal{Z}_y}$ to approximately produce $z_x \in \mathcal{Z}_x \sim \mathbb{P}_e^{Z_x|Z_y=z_y}$. With this latent setup, the latent generator $g : U \times Z_y \rightarrow Z_x$ is more efficient to train with the Sinkhorn algorithm especially when $\dim Z_x \ll \dim X$ or $\dim Z_y \ll \dim Y$, and we can set $\dim U = \dim Z_x$ to prevent dimension mismatch since $\dim Z_x \geq \dim \mathcal{X}$. Apart from this, $\dim Z_x$ and $\dim Z_y$ also estimate $\dim \mathcal{X}$ and $\dim \mathcal{Y}$ by upper-bounding them. Fig. 3 illustrates this pipeline.

Finally, in order to sample $x \sim \mathbb{P}_r^{X|Y=y}$ given $y \in \mathcal{Y}$, we can first obtain $z_y = e_y(y)$, then feed it along with $u \sim \mathbb{P}_u$ into g to produce $z_x = g(u, z_y) \sim \mathbb{P}_g^{Z_x|Z_y=z_y}$, and at last transform z_x back into X through $x = d_x(z_x)$, so that the generated $x = d_x \circ g(u, e_y(y))$ comes from the conditional pushforward measure $\mathbb{P}_g^{X|Y=y} := (d_x \circ g|_{\mathcal{Z}_y})_* \mathbb{P}_u$. Despite the latent training process, $\mathbb{P}_g^{X|Y}$ still approximates $\mathbb{P}_r^{X|Y}$, given that

$$\begin{aligned} \mathbb{P}_g^{X|Y=y} &= d_{x*} (g|_{\mathcal{Z}_y} \mathbb{P}_u) = d_{x*} \left(\mathbb{P}_g^{Z_x|Z_y=e_y(y)} \right) \\ &\simeq d_{x*} \left(\mathbb{P}_e^{Z_x|Z_y=e_y(y)} \right) = d_{x*} \left(e_{x*} \mathbb{P}_r^{X|Y=y} \right) \\ &= (d_x \circ e_x)_* \mathbb{P}_r^{X|Y=y} = \mathbb{P}_r^{X|Y=y}. \end{aligned} \quad (12)$$

3. Results and discussion

In this section, we verify the latent CSGAN framework’s ability to solve high dimensional inverse problems in a Bayesian manner, and analyze these results to provide further insights. First, we apply it to the relatively simple Kraichnan–Orszag ODEs to identify the system’s

¹ Unlike the vanilla GAN, there is actually nothing *conventionally* “adversarial” in this pushforward model’s training process, since we do not need to train any discriminator to compete with the generator. Regardless, the discriminator’s training is now replaced by the Sinkhorn–Knopp algorithm evaluating the Sinkhorn divergence, which may also be regarded as something “adversarial” in a broader sense since it “trains” the Lagrange multipliers α and β , so we still identify this model as a type of GAN for convenience.




Fig. 3. Latent CSGAN pipeline.

Table 1
Summary of experiment results in Section 3.1.

Target distribution	Dimension reduction Section 3.1.2		Inference Error Section 3.1.3
	ξ	y_1	
$p(\xi y_1)$	2 \rightarrow 2	256 \rightarrow 7	Fig. 7

Table 2
Summary of experiment results in Section 3.2.

Target distribution	Dimension reduction Section 3.2.4		Inference Error Sections 3.2.5 and 3.2.6		Uncertainty Section 3.2.7
	G	S or $f_{1:i}$	G	S or $f_{1:i}$	
$p(G S)$		8192 \rightarrow 40	Fig. 12(b)	Fig. 12(d)	
$p(G f_{1:1})$		128 \rightarrow 4			
$p(G f_{1:2})$	8192 \rightarrow 40	256 \rightarrow 8			
$p(G f_{1:3})$		384 \rightarrow 12	Fig. 13(b)	Fig. 13(d)	Fig. 14
$p(G f_{1:4})$		512 \rightarrow 16			
$p(G f_{1:5})$		640 \rightarrow 18			

bimodal source, and illustrate the latent CSGAN’s low dimensional predictions to uncover its strength and weakness visually. Next, we extend the latent CSGANs to the more complicated higher-dimensional two-phase flow problem of multi-well oil extraction. Apart from studying the latent CSGAN’s accuracy in approximating the target posterior distribution, we demonstrate that the dimension of the least volume latent space can serve as an accurate estimator of the intrinsic dimension of the dataset—even when the dataset is nonlinear—and further reveal that the difference between the dimensions of the posterior’s input condition and output may help predict the uncertainty in the target posterior beforehand, assisting the inverse problem’s optimal configuration in future projects. The details of the model architectures used for each example are provided in [Appendices B and C](#). The main results are summarized in [Tables 1 and 2](#).

3.1. Source identification in Kraichnan–Orszag three-mode problem

The Kraichnan–Orszag (KO) three model problem is a set of nonlinear ordinary differential equations (ODE) of particular interest and has been previously investigated by several authors ([Bilionis, Zabarar, Konomi, & Lin, 2013](#); [Ma & Zabarar, 2009a](#); [Orszag & Bissonette, 1967](#); [Wan & Karniadakis, 2005](#)). It is described by the system of ODEs:

$$\begin{aligned} \frac{dy_1}{dt} &= y_1 y_3, \\ \frac{dy_2}{dt} &= -y_2 y_3, \\ \frac{dy_3}{dt} &= -y_1^2 + y_2^2 \end{aligned} \quad (13)$$

subject to initial random conditions at $t = 0$. The stochastic initial conditions are defined by:

$$y_1(0) = 1 \quad y_2(0) = \xi_1 \quad y_3(0) = \xi_2, \quad (14)$$

where

$$\xi_1 \sim \mathcal{U}[-0.1, 0.1], \quad \xi_2 \sim \mathcal{U}[-1, 1]. \quad (15)$$

The solution $y(t; \xi)$ of (13) is discontinuous with respect to ξ at $\xi_1 = 0$.

In the inverse problem considered here, we are interested in inferring the initial conditions $y_2(0)$, $y_3(0)$ (or equivalently the posterior distributions of ξ) by observing only the behavior of $y_1(t)$ in the time interval $t \in [0, 30]$. In other words, we aim to retrieve the posterior $p(\xi | y_1)$. In order to generate synthetic data to solve this problem on a data-driven basis, we solve numerically the ODE using a 4th order Runge–Kutta method with a time step $\Delta t = 0.003$, resulting in a response vector $\mathbf{y} = [y_1(t; \xi), y_2(t; \xi), y_3(t; \xi)] \in \mathbb{R}^{3 \times 10^4}$. We run $N = 2048$ numerical simulations of the ODE that correspond to an equal number of Monte Carlo samples ξ drawn from the prior (15), to be used as the training data for the CSGAN.

3.1.1. Model configurations and training

We start by preprocessing the dataset. Each of the 2048 input conditions $y_1(t)$ consists of 10001 dimensions coming from the $\Delta t = 0.003$ sampling spacing over the $[0, 30]$ time span. Since FFT shows that $y_1(t)$ in general has bandwidth ≤ 4 Hz ([Fig. 4](#)), for the ease of the following tasks, we linearly interpolate and uniformly resample $y_1(t)$ into 256-D vectors at the sampling rate 8.5 Hz, so as to preliminarily reduce $y_1(t)$ ’s dimension by leveraging the *Nyquist–Shannon sampling theorem* ([Shannon, 1949](#)). Thereafter, we normalize both ξ and $y_1(t)$ to the range $[0, 1]$, as a normal practice to facilitate the training of neural networks. We reserve 512 samples as the test samples and keep the rest 1536 samples as the training set to train the latent CSGAN as follows.

To approximate the target posterior $p(\xi | y_1)$ efficiently, first we use LVAE to further reduce the dimension of the 256-D $y_1(t)$. We model the encoder with an 1D convolutional neural network and the decoder with an 1D deconvolutional neural network. Their architectures and hyper-parameters are listed in [Table B.3](#). The dimension of the autoencoder’s

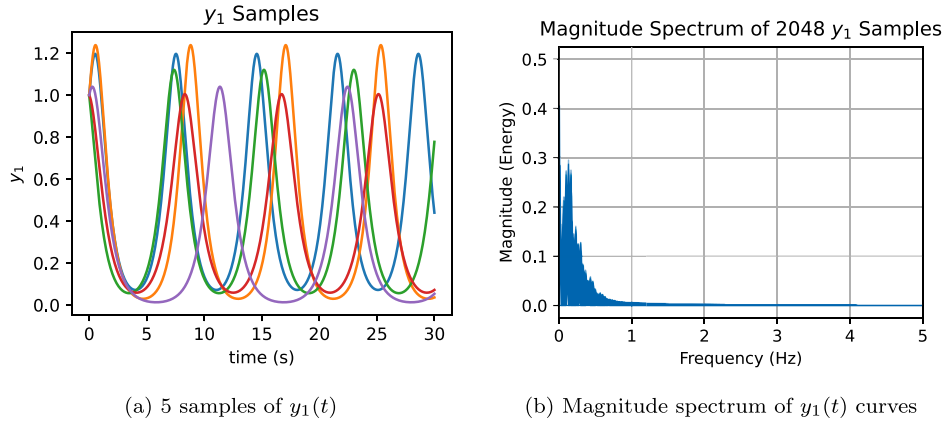
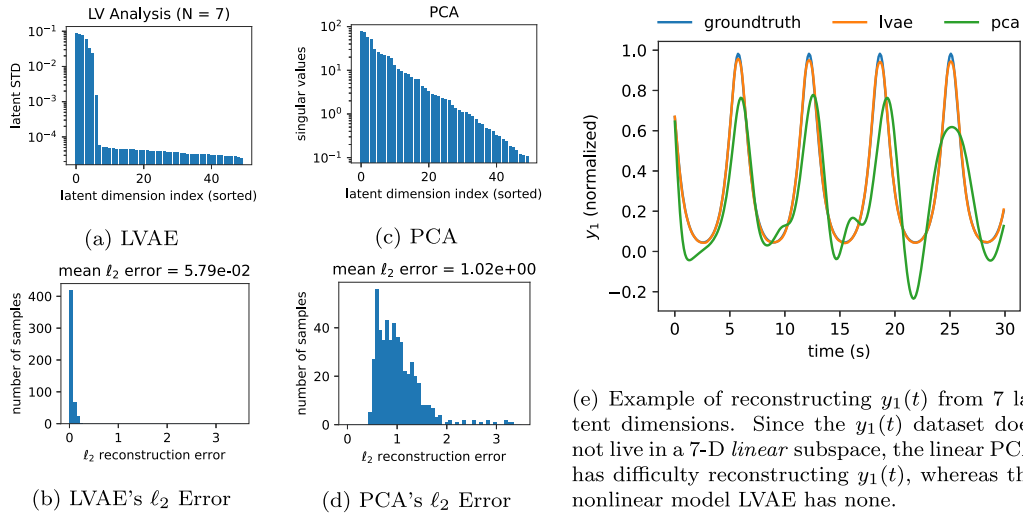
Fig. 4. FFT result of $y_1(t)$ samples.

Fig. 5. Dimension Reduction Result of KO Dataset.

latent space is set to 50. We set the parameter η in the loss function (5) to 0.004, and to prevent the LVAE from getting trapped at local minima when λ is too large at the start, we gradually increase λ from 0 to 0.2 at constant speed throughout the training period of 2000 epochs. We set the batch size to 50, so each of these 2000 epochs contains $1536/50 \approx 31$ iterations. We use Adam (Kingma & Ba, 2015) with a learning rate of 0.0001 for optimization.

After the LVAE finishes compressing $y_1(t)$ to n certain dimensions, we extract this low dimensional latent subspace and train a latent CSGAN with the generator architecture in Table B.4. The generator's output dimension is equal to $\dim \xi = 2$, and for its two inputs—latent noise U and condition Y —we set $\dim U = \dim \xi = 2$ and $\dim Y = n (= 7)$. The latent distribution \mathbb{P}_u over U is set to the uniform distribution $\mathcal{U}[0, 1]^2$. We make $\rho = 0.05$ for the evaluation of the Sinkhorn divergence S_ρ (11). The CSGAN is trained for 2000 epochs, with the batch size 128. The optimizer is Adam with a learning rate of 0.0001.

3.1.2. Results: Least volume dimension reduction

The LVAE ends up compressing $y_1(t)$ into $n = 7$ dimensions from 50. This is reflected in Fig. 5(a), where we can see there are 7 latent dimensions whose latent STDs are orders of magnitude larger than the rest 43 dimensions. Indeed, pruning these 43 trivial latent dimensions does not bring any noticeable change to the LVAE's ℓ_2 reconstruction error displayed in the histogram Fig. 5(b).

Comparing the LVAE's result to PCA's can further help us conclude that the dataset of $y_1(t)$ is highly nonlinear. Fig. 5(c) shows the 50

largest singular values of PCA, where there is no apparent plummet similar to that one in Fig. 5(a). In addition, PCA's reconstruction error is about 20 times larger than that of the LVAE when we only keep the 7 latent dimensions of the largest singular values while discarding the rest (Fig. 5(d)). This indicates that the set of $y_1(t)$ does not reside on a linear hyperplane of 7 dimensions. In fact, we can only achieve a reconstruction error comparable to LVAE's when we keep no less than 30 principal dimensions, which suggests that the $y_1(t)$ dataset should live on a curved 7-D manifold that is embedded in a 30-D linear space. For illustration, Fig. 9(e) superimposes the LVAE's and PCA's reconstructions of a given $y_1(t)$ curve from 7 principal latent dimensions. It is clear that LVAE has far superior dimension reduction and reconstruction capability.

3.1.3. Results: Bayesian inference of posterior $p(\xi | y_1)$

The CSGAN is then trained in the LVAE's 7-D latent space to approximate the posterior distribution $p(\xi | y_1)$. To investigate the latent CSGAN's stochastic predictions' accuracy, we let the generator g produce 100 predictions $\Xi^{(j)} = \{\hat{\xi} | \hat{\xi} = g(u, e(y_1^{(j)})), u \sim \mathbb{P}_u\}$ for each condition $y_1^{(j)}$ of the 512 test cases, then we both illustratively display the heatmap of these predictions along with their groundtruths ξ for 25 test cases in Fig. 6, and quantitatively evaluate the minimum ℓ_2 prediction error $\epsilon^{(j)}$ of each test case j —defined as $\epsilon^{(j)} = \min_{\hat{\xi} \in \Xi^{(j)}} \|\hat{\xi} - \xi^{(j)}\|$ —and show their distribution in the histogram Fig. 7(a).

The predictions in Fig. 6 clearly uncovers the bimodality of the posterior $p(\xi | y_1)$ under almost all y_1 conditions, and these two modes

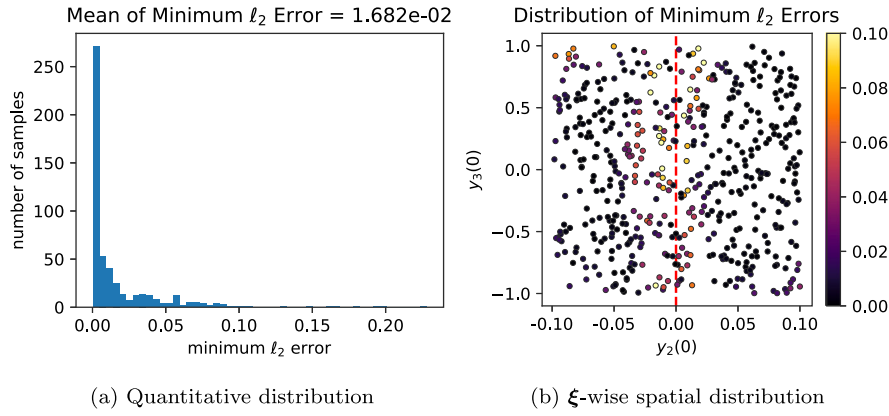


Fig. 7. Distributions of Minimum ℓ_2 Prediction Errors of Test Samples.

the spatial distribution of the prediction error $\epsilon^{(j)}$ in the ξ space, and the locations of high error entities agree with our discussion. However, nothing certain can be drawn at this moment, and we leave it for further investigation in future works.

3.2. Inferring subsurface properties of heterogeneous oil reservoirs from observable responses in oil extraction

Next, we move on to the more complicated inverse problem in oil production, in which we aim to infer the given heterogeneous oil reservoir's subsurface properties from its observable responses in oil extraction. The main challenge here is the high dimensionality of both the input and the output of the target posterior distribution. We will demonstrate how Least Volume can not only help approximate this high dimensional posterior by representing it in the low dimensional latent spaces, but also provide helpful information regarding the posterior's uncertainty that may help improve the inverse problem's configuration.

3.2.1. Two-phase flow problem in heterogeneous porous materials

Specifically, we are interested in inferring the unobservable subsurface properties—the permeability field or the hydraulic conductivities—of a given 2-D heterogeneous oil reservoir. Its rectangular domain is illustrated in Fig. 8. The flow and transport of oil in the reservoir is described by a computational model that simulates immiscible and incompressible two-phase flow (water and oil) in a rectangular domain $B = [L_x, L_y]$ with real dimensions $L_x = 333.58$ and $L_y = 670.56$. For simplicity, gravity effects are ignored. We consider the presence of seven wells, in two of which, water is pumped in, while the oil is pushed out of the remaining five. The two production wells are located at corners of the south boundary of the domain $((0, 0)$ and $(L_x, 0)$), while the 5 production wells are located on the north size of the domain, specifically, at locations $(0, 0.7L_y)$, $(0, L_y)$, $(0.5L_x, L_y)$, (L_x, L_y) and $(L_x, 0.7L_y)$.

We denote with $p_o, p_w, \mathbf{v}_o, \mathbf{v}_w$ and s_o, s_w the pressure, velocity and saturation fields of the oil and water respectively where for the later it holds that $s_o + s_w = 1$. Define the total velocity as $\mathbf{v} = \mathbf{v}_o + \mathbf{v}_w$ and the capillary pressure as $p_{cap} = p_o - p_w$ which is assumed to be a monotone function of the water saturation s_w . The global pressure is also defined as $p = p_o - p_c$ where the complement pressure p_c is given as

$$p_c(s_w) = \int_1^{s_w} f_w(s) \frac{\partial p_{cap}}{\partial s_w} ds, \quad (16)$$

where f_w is the fractional-flow function and measures the water fraction of the total flow and is given by

$$f_w = \frac{\lambda_w}{\lambda_w + \lambda_o}, \quad (17)$$

where λ_w and λ_o are the water and oil total mobilities respectively. These can be expressed as

$$\lambda_w(s_w) = \frac{(s_w^*)^2}{\mu_w}, \lambda_o(s_w) = \frac{(1 - s_w^*)^2}{\mu_o}, s_w^* = \frac{s_w - s_{wc}}{1 - s_{or} - s_{wc}}, \quad (18)$$

where μ_o and μ_w are the oil and water viscosities respectively, s_{or} is the irreducible oil saturation (lowest oil saturation that can be achieved by displacing oil by water) and s_{wc} is the connate water saturation (the saturation of water trapped in the pores of the rock during formation of the rock). Taking the above definitions into consideration, the governing equations of the flow model are (Aarnes, Gimse, & Lie, 2007)

$$-\nabla \cdot (\mathbf{K}\lambda(s_w)\nabla p) = q \quad (19a)$$

$$\phi \frac{\partial s_w}{\partial t} + \nabla \cdot (f_w(s_w)\mathbf{v}) = \tilde{q}_w, \quad (19b)$$

in the interior of the domain B , where \mathbf{K} is the permeability tensor, ϕ is the porosity, q is the source term modeling the water injection and \tilde{q}_w is the source term for the saturation equation given by

$$\tilde{q}_w = \max\{q, 0\} + f(s_w) \cdot \min\{q, 0\}. \quad (20)$$

The total velocity is connected to the pressure via the modified Darcy law (Dake, 1983)

$$\mathbf{v} = -\mathbf{K}\lambda(s_w)\nabla p. \quad (21)$$

No-flux boundary conditions are considered on ∂B

$$\mathbf{v} \cdot \mathbf{n} = 0, \quad (22)$$

where \mathbf{n} is the unit vector that is normal to the boundary. Furthermore, the initial water saturation is taken to be zero

$$s_w(\mathbf{x}, t = 0) = 0 \quad (23)$$

$\forall \mathbf{x} \in B$.

For this example, we assume a constant porosity $\phi = 10^{-3}$, the water and oil viscosities are $\mu_w = 3 \cdot 10^{-4}$ and $\mu_o = 3 \cdot 10^{-3}$ while the saturations are $s_{wc} = s_{or} = 0.2$. The initial boundary value problem is solved using a finite control volume method where for the pressure we use a two-point flux approximation finite volume scheme and for the saturation equation we use a first order upwind scheme (Aarnes et al., 2007). The discrete system of differential equation is evolved in time using a first order implicit scheme with adaptive step selection upon a Newton–Raphson solver. At last, the wells are modeled via the source term which is set to an injection rate value $IR = 9.3529$ (Aarnes et al., 2007). Each injection well is assigned an injection rate equal to $IR/2$ whereas the value assigned to each production wells is equal to $-IR/5$.

The uncertain input parameters of the model that we seek to infer in this example is the unknown rock permeability \mathbf{K} that is present throughout the domain. As prior assumptions, we consider \mathbf{K} to be isotropic, that is $\mathbf{K} = K\mathbf{I}$, and $K := K(\mathbf{x}, \xi)$ is modeled as the

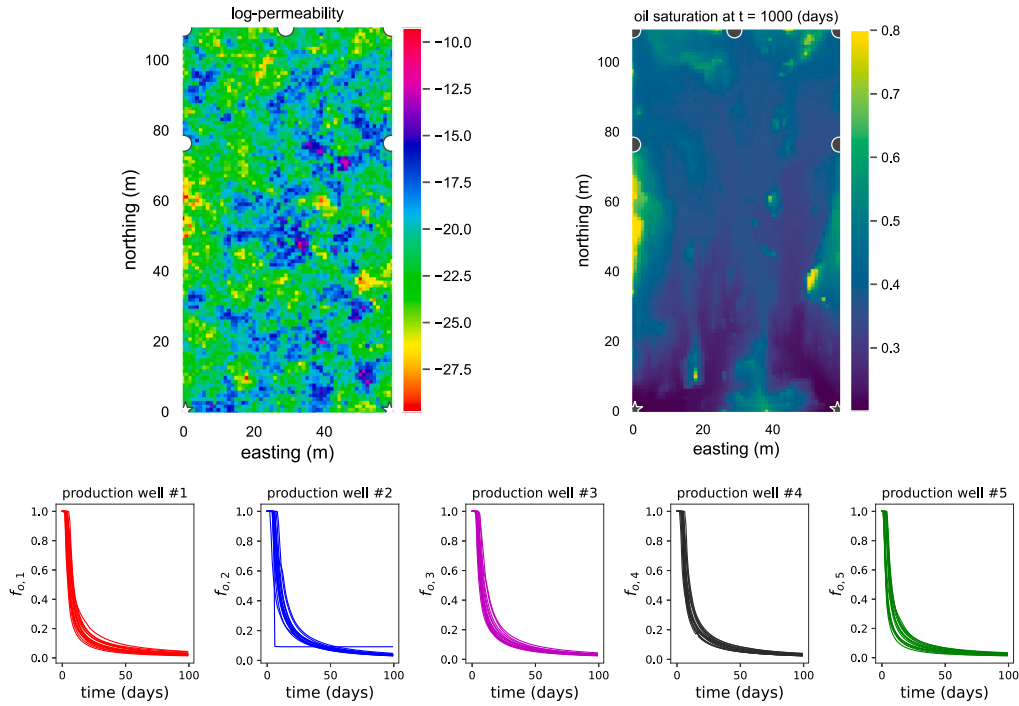


Fig. 8. Top: Sample realization of the simulator. The injection wells are depicted with ‘*’ and the production wells are depicted with ‘o’. Log permeability is displayed on the left, and oil saturation on the right. Bottom: Fifteen samples of the oil fractional flow rates.

exponential of a Gaussian random field $G(\mathbf{x}, \xi)$ which is expressed by its Karhunen–Loève expansion,

$$G(\mathbf{x}, \xi) = \hat{G}_0 + \gamma \sigma \sum_{i=1}^L \sqrt{\lambda_i} \xi_i \phi_i(\mathbf{x}), \quad K(\mathbf{x}, \xi) = \exp \{ G(\mathbf{x}, \xi) \}, \quad (24)$$

where $\{\xi_i\}$ are independent Gaussian variables. We make use of the SPE-10 data set (Christie & Blunt, 2001) which is measured in the range of $1200 \times 2200 \times 170$ (ft³) and is discretized on a regular Cartesian grid with $120 \times 220 \times 85$ nodes. We use the 85 log-fields of dimension 120×220 as different realizations from which we compute the sample mean \hat{G}_0 and the sample covariance matrix $C(\mathbf{x}, \mathbf{x}')$. From the latter we next obtain the eigenvalues and eigenvectors used in Eq. (24) and we choose $L = 20$ to retain the 20 most significant terms corresponding to the largest eigenvalues. Note that the relatively small number of samples used to construct the covariance matrix results in a large variance. In order to avoid numerical instabilities in our simulations we have multiplied σ by an additional constant taken to be $\gamma = 0.3$ and added an additional constant $\kappa_0 = 10^{-12}$ to $K(\mathbf{x}, \xi) = \exp \{ G(\mathbf{x}, \xi) \} + \kappa_0$ in order to ensure well posedness of the elliptic problem in (19a).

The observable quantities are the oil saturation across the whole domain as well as the oil fractional flow rates at the five production wells

$$f_{o,i} = \frac{\lambda_o}{\lambda_w + \lambda_o}, i = 1, \dots, 5 \quad (25)$$

Fig. 8 shows a sample realization of the oil saturation s_o at time T and 15 samples of the oil fractional-flow rate.

For simplicity, hereafter we denote the log permeability field $G(\mathbf{x}, \xi)$ by \mathbf{G} , the saturation field $s_o(\mathbf{x}, t = T)$ by \mathbf{S} . Moreover, we denote the flow rate curve (or *oil-cut curve*) of each production well i by \mathbf{f}_i , and we use $\mathbf{f}_{i:j}$ to refer to the combination (or Cartesian product) of oil-cut curves of a consecutive series of wells starting from index i and ending at index j . Sometimes \mathbf{f}_i may also be expressed as $\mathbf{f}_{1:i}$ so as to unify the notation, depending on the context.

3.2.2. Outline of experiment tasks and motivations

This experiment involves more tasks than the previous one in Section 3.1, in order to more thoroughly explore and reveal the potential

of the latent CSGANs in tackling high dimensional inverse problems. For clarity, here we list these tasks out, and highlight their motivations and objectives:

- T1. We apply LVAEs to \mathbf{G} , \mathbf{S} and \mathbf{f}_1 , and retrieve the posteriors $p(\mathbf{G} | \mathbf{S})$ and $p(\mathbf{G} | \mathbf{f}_1)$ with latent CSGANs, then assess their accuracy. This aims to investigate latent CSGANs’ ability to handle high dimensional input and output of different formats and unify them under the same latent framework.
- T2. We additionally train LVAEs for $\mathbf{f}_{1:2}$, $\mathbf{f}_{1:3}$, $\mathbf{f}_{1:4}$ and $\mathbf{f}_{1:5}$ respectively, and use them for different latent CSGANs to approximate the posteriors $p(\mathbf{G} | \mathbf{f}_{1:2})$, $p(\mathbf{G} | \mathbf{f}_{1:3})$, $p(\mathbf{G} | \mathbf{f}_{1:4})$, $p(\mathbf{G} | \mathbf{f}_{1:5})$. Then, along with $p(\mathbf{G} | \mathbf{S})$ and $p(\mathbf{G} | \mathbf{f}_1)$, we use posterior samples to evaluate their forward propagation error and examine whether there is a *correlation between the posterior’s scatter over the support of the distribution and the input condition’s intrinsic dimensionality*, as the latter has been obtained by the LVAE’s latent space dimension.

The probing of this potential correlation in Task 2 is motivated by the intuitive *preimage theorem* (Guillemin & Pollack, 2010) (or *regular level set theorem* Lee & Lee, 2012):

For *smooth* manifolds X and Y , if y is a *regular* value of the *smooth* function $f : X \rightarrow Y$, then the preimage $f^{-1}(y)$ is a submanifold of X , with $\dim f^{-1}(y) = \dim X - \dim Y$.

To explain further the implication of the above theorem, if $f(X)$ is of non-zero measure in Y , then *Sard’s theorem* (Guillemin & Pollack, 2010; Lee & Lee, 2012) further guarantees that almost every $y \in f(X) \subseteq Y$ is a regular value, which means we may safely apply the preimage theorem to almost every y that can constitute an observed value in practice. Then, for a given deterministic forward function f , the approximation of the posterior $p(X | Y = y)$ reveals its support, which informs us about the preimage $f^{-1}(y)$ —as $f^{-1}(y)$ supports $p(X | Y = y)$. By adopting the manifold hypothesis together with the smoothness assumption on both X and Y , the preimage theorem then suggests that the support of $p(X | Y = y)$ depends on the dimension of the input condition space

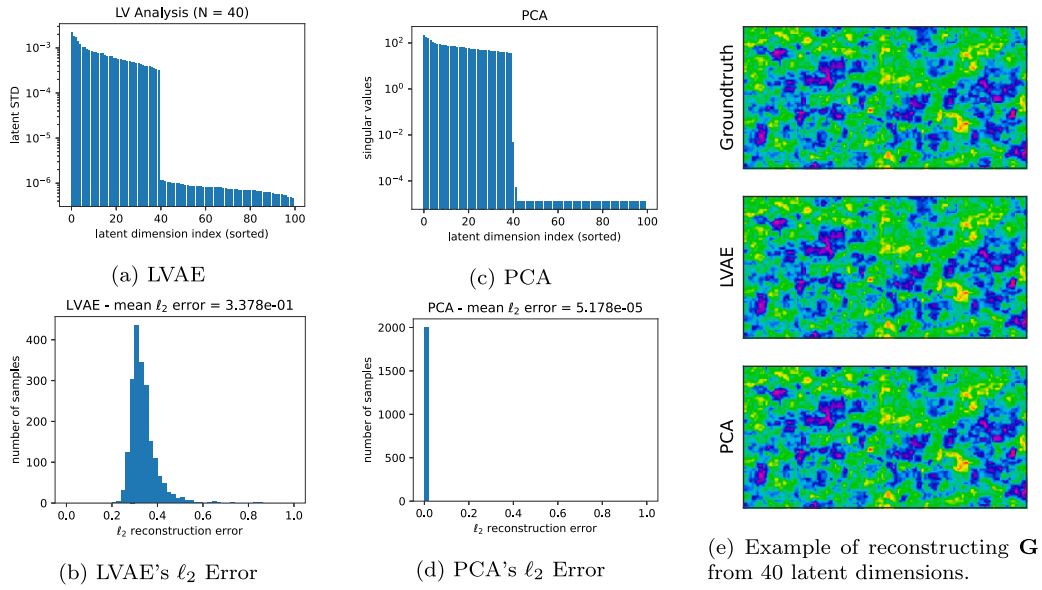


Fig. 9. Dimension reduction results of log permeability field \mathbf{G} . By construction, the \mathbf{G} samples are drawn from a Karhunen–Loève expansion, therefore the linearity of the decomposition allows PCA to achieve reconstruction performance superior to the nonlinear LVAE. However, LVAE also uncovers the 40-D manifold structure of this dataset and has visually indistinguishable reconstruction quality. Note that the *mean ℓ_2 error* here refers to $\mathbb{E}\|\mathbf{G} - \hat{\mathbf{G}}\|_2$. Therefore, the deceptively large error magnitude in Fig. 9(b) is primarily due to the 128×64 dimensions of \mathbf{G} rather than poor reconstruction quality.

Y : that is, the higher the dimension of Y , the smaller the dimension of $f^{-1}(y)$, which means $p(X | Y = y)$ should have smaller variance because $f^{-1}(y)$ is more concentrated in X . This correlation, if verified, may carry significant implications for inverse problem solving empowered by Least Volume, as will be discussed later in Sections 3.2.7 and 4.

3.2.3. Model configurations and training

To approximate the target posterior distributions of \mathbf{G} , we generate a 10000-sample dataset of the triplet $\{(\mathbf{G}, \mathbf{S}, \mathbf{f}_{1:5})\}$ using the aforementioned computational model. Both \mathbf{G} and \mathbf{S} have size 110×60 stemming from the regular Cartesian grid in the 2D space, while each \mathbf{f}_i has 100 dimensions corresponding to 100 equal-paced time steps. To make the data samples more suitable for convolutional neural networks' up- and down-sampling transformations, we use nearest neighbor interpolation to upscale \mathbf{G} and \mathbf{S} to the size 128×64 , and \mathbf{f}_i to 128 dimensions. After that, we clip the outlier values in \mathbf{S} (i.e., those not in the range $[0.2, 0.8]$) caused by numerical issues, and normalize all the values of $(\mathbf{G}, \mathbf{S}, \mathbf{f}_{1:5})$ to the range $[0, 1]$. We use 8000 samples as the training set and retain the rest 2000 samples for use as the test set.

Similar to the KO-ODE experiment, we first use LVAEs to reduce the dimension of \mathbf{G} , \mathbf{S} and \mathbf{f}_i (or $\mathbf{f}_{i:j}$) respectively, in order to get their compact representation and estimate their dimensionality. For \mathbf{G} and \mathbf{S} , we use 2D convolutional LVAEs to autoencode them. Their architectures, hyperparameters and training details are listed in Tables C.5 and C.6. For $\mathbf{f}_{i:j}$ we use 1D convolutional LVAEs for the compression, but the situation is a bit more complicated and requires more effort on fine-tuning the LVAEs' hyperparameters. Specifically, as mentioned in Chen and Fuge (2024), the weight λ in the objective function (5) controls the trade-off between reconstruction and dimension reduction. If λ is too large, then the dimension of the latent space can be reduced arbitrarily by discarding important information (or dimension) of the data, thus increasing the reconstruction error. This is similar to PCA, for which we may reduce the number of principal components to extract, but at the cost of losing precision in data reconstruction. For Task 2 described in Section 3.2.2, we are interested in retrieving the dimensionalities of different curve combinations $\mathbf{f}_{i:j}$ by taking the LVAE's latent dimensions as their estimates. In order to make the estimates accurate, we need to make the reconstruction errors of different curve combinations comparable with each other. The 1-D convolutional LVAE architectures

and hyperparameters in Tables C.5 and C.6 are derived with this consideration taken into account. Their validity in unifying the reconstruction errors will be verified later in Section 3.2.4, Fig. 11(e).

After training the LVAEs for \mathbf{G} , \mathbf{S} and $\mathbf{f}_{i:j}$, we extract their latent spaces and train the latent CSGANs to approximate $p(\mathbf{G} | \mathbf{S})$ and $p(\mathbf{G} | \mathbf{f}_{i:j})$. Likewise, the training details and architectures are described in Tables C.7 and C.8.

3.2.4. Results: Least volume dimension reduction

Figs. 9 and 10 demonstrate the dimension reduction results of the LVAEs on the two image datasets \mathbf{G} and \mathbf{S} . We can see that the LVAE of the log permeability field \mathbf{G} manages to compress \mathbf{G} into a 40 dimensional latent space, which agrees with \mathbf{G} 's groundtruth dimension 40. Because of the inherent linearity of this dataset—as specified by the Karhunen-Loève expansion in Eq. (24), both the PCA and the LVAE can achieve comparable reconstruction quality.

Even more intriguing is that the final saturation field \mathbf{S} —the deterministic governing PDE's final state that depends on \mathbf{G} —is also compressed by its LVAE into a 40 dimensional latent space, and the PCA's poor reconstruction quality shows that the dataset of \mathbf{S} is nonlinear (Fig. 10). If this number 40 happens to be \mathbf{S} 's actual intrinsic dimension, and we further assume that the causal relationship between \mathbf{G} and \mathbf{S} can be modeled by a smooth map f_s that gives $\mathbf{S} = f_s(\mathbf{G})$, then the preimage theorem tells us the level set $f_s^{-1}(\mathbf{S})$ of almost every \mathbf{S} is a zero-dimensional manifold, i.e., a countable discrete space (Lee, 2010). Putting this intuitively, it means almost every level set $f_s^{-1}(\mathbf{S})$ comprises only a set of isolated points that do not form lines or surfaces in the space of \mathbf{G} , such that the probability density function of $p(\mathbf{G} | \mathbf{S})$ can be expressed as the sum of Dirac delta functions. If in addition f_s is injective, then each level set contains only a single point and $p(\mathbf{G} | \mathbf{S})$ may even be sufficiently fit by a unimodal regression model with enough complexity. We shall validate this in the next section.

The LVAEs also obtain some interesting results on the oil-cut curve combinations $\mathbf{f}_{1:j}$ for $j \in \{1, 2, 3, 4, 5\}$. As we can see in Fig. 11, starting from $\mathbf{f}_{1:1}$ to $\mathbf{f}_{1:4}$, every time we include an additional oil-cut curve in $\mathbf{f}_{1:j}$, the curve combination $\mathbf{f}_{1:j}$'s dimension increases by 4. This not only suggests that each oil-cut curve \mathbf{f}_i lives on a 4-D manifold—which can be double-checked by applying LVAEs to \mathbf{f}_j other than \mathbf{f}_1 —but also implies that \mathbf{f}_i for $i \in \{1, 2, 3, 4\}$ are uncorrelated with each other, as

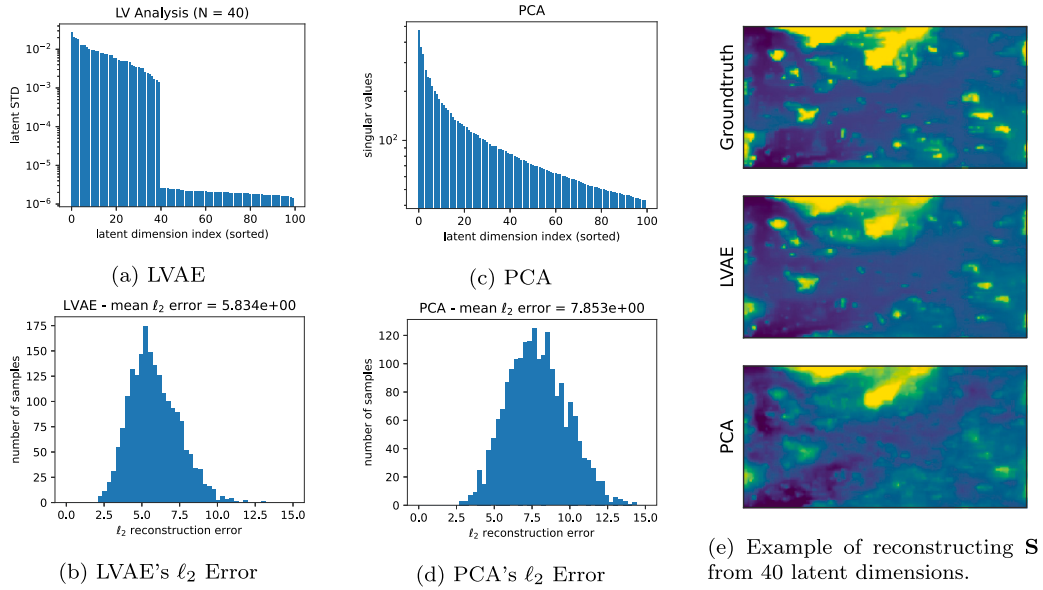


Fig. 10. Dimension reduction results of saturation field S . The LVAE can better handle this highly nonlinear dataset than PCA in terms of both compression and reconstruction. Likewise, the *mean* ℓ_2 error here refers to $\mathbb{E}\|S - \hat{S}\|_2$, so the large error magnitudes in Figs. 10(b) and 10(d) are primarily caused by the 128×64 dimensions of S .

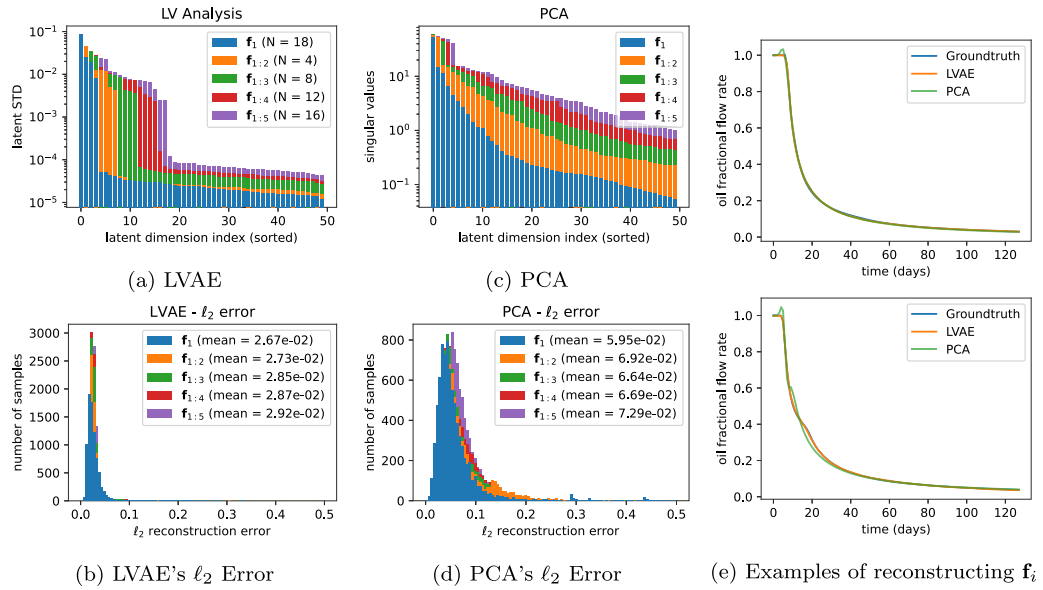


Fig. 11. Dimension reduction results of oil-cut curves $f_{1:j}$. The LVAEs are tuned to have comparable ℓ_2 reconstruction errors for each production well. When using the same number of latent dimensions, PCA’s reconstruction is not as good as that of the LVAEs, but is still acceptable in most cases. This means the $f_{1:j}$ datasets are not highly nonlinear.

otherwise we should be able to use less dimensions to represent some of these combinations. However, the situation differs as j increases to 5. The result shows that $f_{1:5}$ contains only 2 extra dimensions compared to $f_{1:4}$, while f_5 alone has 4 dimensions, indicating that f_5 correlates with the other oil-cut curves. This hints that the amount of information provided by the production wells starts to “saturate”. The same “saturation effect” is also perceived on S : each pixel of S has its own variation on \mathbb{R} and thus has 1 dimension, but the Cartesian product of these 128×64 pixels only has 40 dimensions rather than 8192 due to the correlation among them.



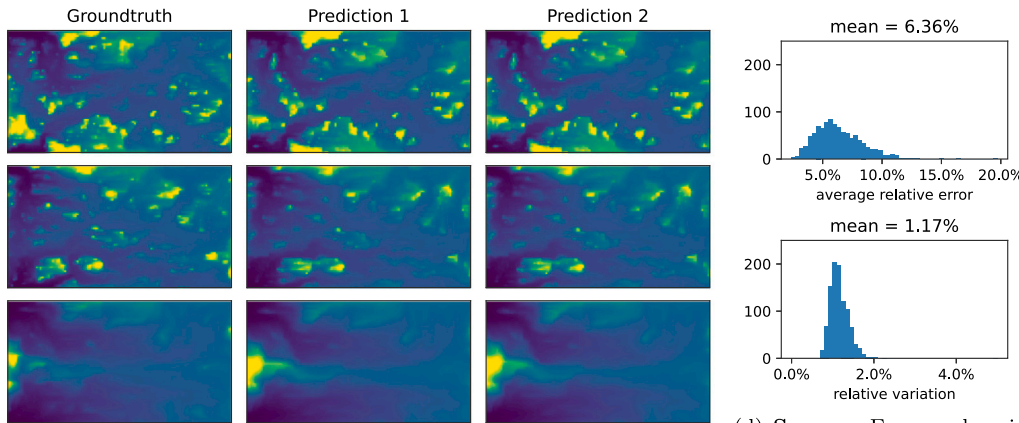
This information (or dimension) saturation phenomenon is insightful. As an example, in the scenario where we are provided with an unlimited number of sensors, if we want to use them to monitor the status of a structure, it may not be true that the more sensors we mount, the more thoroughly we can grasp the structure’s condition, given that there probably exists a threshold on the number of sensors to mount

on the structure, beyond which the sensor data’s dimension cannot increase to provide more information about the structure, while the sensors’ cost becomes staggering. Monitoring the sensor data’s intrinsic dimension with LVAE may help us determine the threshold.

3.2.5. Results: Bayesian inference of log permeability field via $p(G | S)$

With the LVAEs at hand, we then train the latent CSGANs to approximate the posteriors $p(G | S)$ and $p(G | f_{1:j})$ for $j \in \{1, 2, 3, 4, 5\}$. The two key aspects for assessing the quality of the approximations, in accordance with Section 3.2.2, are:

- The *accuracy* of their approximation, *i.e.*, how well does their conditional pushforward measure \mathbb{P}_g align with the groundtruth posterior probability measure $\mathbb{P}_.$. We will investigate their accuracy in both the output G space and the input condition spaces using different metrics to obtain a comprehensive understanding,

(a) \mathbf{G} space: Predictions from latent CSGAN for $p(\mathbf{G} | \mathbf{S})$ (b) \mathbf{G} space: Error and variation distribution of latent CSGAN for $p(\mathbf{G} | \mathbf{S})$ (c) \mathbf{S} space: Saturation field $\hat{\mathbf{S}}$ of predicted $\hat{\mathbf{G}}$ against groundtruth \mathbf{S} (d) \mathbf{S} space: Error and variation distribution of $\hat{\mathbf{S}}$ Fig. 12. Predictions and errors of latent CSGAN for $p(\mathbf{G} | \mathbf{S})$ in \mathbf{G} and \mathbf{S} space.

as some metrics may not work under certain circumstances due to the complexity of this problem.

- The *uncertainty* around different posteriors, and the uncertainty's correlation with the input condition's intrinsic dimension. In Section 3.2.7, we shall use the computational entropy estimator powered by the KSG method (Kraskov, Stögbauer, & Grassberger, 2004) to quantify the uncertainty of different posterior distributions over the output space of \mathbf{G} . The larger the entropy, the higher the uncertainty.

First, we inspect the latent CSGAN's learning accuracy of $p(\mathbf{G} | \mathbf{S})$ in this section. Fig. 12 studies and illustrates it from different perspectives. Fig. 12(a) displays the predicted \mathbf{G} made by the CSGAN on three randomly chosen test cases, juxtaposed against their groundtruth \mathbf{G} on the left. For each case, given the input condition \mathbf{S} , the generator stochastically produces three \mathbf{G} predictions for illustration. We can see that the predictions are almost visually indistinguishable from

their groundtruths, and in each test case, the predictions also look identical to each other. This means the CSGAN is accurate, and its target posterior $p(\mathbf{G} | \mathbf{S})$ has very low uncertainty, which agrees well with our projection in the beginning of Section 3.2.4.

The results in Fig. 12(b) bolster the above claim quantitatively. There are two metrics involved: *relative error* and *relative variation*. They are respectively *length-scale-independent* quantities for evaluating each prediction's deviation from its groundtruth, and the variation of a batch of predictions from each other. They are defined as follows:

- *Length Scale*: The log permeability's *length scale* $c_G = 6\sigma$ is evaluated by taking the standard deviation σ among all the pixels of all the 2000 \mathbf{G} samples, then scaling it up by 6. This 6σ evaluation makes sense because the log permeability values are generated from a Gaussian random field, and the 6σ range encloses around 99.7% log permeability values while ignoring the few outliers on the long tail, avoiding underestimating the quantities below.

- **Relative Error:** For each prediction, the absolute error between $\hat{\mathbf{G}}$ and \mathbf{G} are first evaluated pixel-wisely, which gives a 110×60 grid of errors denoted by $\mathbf{E} = |\hat{\mathbf{G}} - \mathbf{G}|$. Then the mean absolute error \bar{e} is evaluated over all pixels of \mathbf{E} . Finally the *relative error* ϵ of each prediction $\hat{\mathbf{G}}$ is defined by $\epsilon = \bar{e}/c_G$. We choose this pixel-wise ℓ_1 -distance-based error over the ℓ_2 distance (as in Section 3.1.3, Fig. 7) to assess the deviation between image data because of not only its intuitiveness, but also the \mathbf{G} space's high dimensionality, as the ℓ_p distance with larger p in general works worse and less intuitively in higher dimensional spaces due to the diminishing contrast (Aggarwal, Hinneburg, & Keim, 2001; Beyer, Goldstein, Ramakrishnan, & Shaft, 1999). Since ϵ is essentially a ratio, we express it in Fig. 12 in percentage.
- **Average and Minimum Relative Error:** For each \mathbf{G} and \mathbf{S} of the 2000 test cases, we sample 50 predictions $\hat{\mathbf{G}}^{(j)}$ for $j \in \{1, \dots, 50\}$ using the CSGAN. Then we evaluate the average and minimum of ϵ across these 50 predictions. They provide a sense of the average prediction accuracy in each test case, and the best accuracy CSGAN can achieve out of all the nonidentical predictions.
- **Relative Variation:** For each test case, we evaluate the value range (i.e., $\max - \min$) of each pixel of $\hat{\mathbf{G}}$ across the 50 predictions, then take the average of these ranges across all pixels and divide it by the length scale c_G to obtain the *relative variation* of predictions in each test case. This value provides us a crude sense of the concentration of these predictions—or the uncertainty of the pushforward measure $\mathbb{P}_g^{\mathbf{G}|\mathbf{S}}$, serving as an appetizer prior to the entropy-based discussion of uncertainty in Section 3.2.7. Likewise, we express the relative variation in percentage in Fig. 12.

Fig. 12(b) shows that the predictions in each test case have very high accuracy and concentration, which agrees with our visual observation in Fig. 12(a). The proximity between the average and the minimum relative error further testifies to the low uncertainty of $\mathbb{P}_g^{\mathbf{G}|\mathbf{S}}$.

However, so far we have only done inspections in the output space of \mathbf{G} . It may happen that the mapping $f_s : \mathbf{G} \mapsto \mathbf{S}$ —as introduced in Section 3.2.4—has large Lipschitz constant, such that even if a prediction $\hat{\mathbf{G}}$ for condition \mathbf{S} has a small deviation from its groundtruth \mathbf{G} , its saturation field $\hat{\mathbf{S}} := f_s(\hat{\mathbf{G}})$ still differs largely from the groundtruth input condition \mathbf{S} . Fortunately, this is not the case. As Figs. 12(c) and 12(d) show, the saturation fields of the predicted log permeability fields—obtained after simulation—are very close to their groundtruth saturation fields, and they also have very small variations. It should be noted that for the relative error and relative variation in Fig. 12(d), we use the actual range $0.6 = 0.8 - 0.2$ of the saturation values as the length scale instead of using its standard deviation, as the saturation does not follow a Gaussian distribution from our observation, and the 6σ evaluation can lead to a larger length scale value overestimating the accuracy. Moreover, due to the expensive cost of the simulation, we only perform this investigation over 1000 test cases, and for each of them we only simulate the saturation field for 5 predictions rather than 50. Nonetheless, we believe this relatively limited evaluation in the \mathbf{S} space can still sufficiently reflect the accuracy of $\hat{\mathbf{S}}$.

This precision in both the \mathbf{G} space and the \mathbf{S} space, together with the minimization of the Sinkhorn divergence during training, shows that the latent CSGAN approximates the target posterior $p(\mathbf{G} | \mathbf{S})$ sufficiently, so we may expect the true $p(\mathbf{G} | \mathbf{S})$ to have these observed characteristics. If this holds, then as Fig. 12 shows, $p(\mathbf{G} | \mathbf{S})$ is indeed a distribution with very low uncertainty, and is almost surely unimodal, since for each \mathbf{S} it practically always yields the same \mathbf{G} . This confirms our projection in Section 3.2.4, and further reveals that $f_s : \mathbf{G} \mapsto \mathbf{S}$ is not only injective but also very likely *homeomorphic* (Lee, 2010), given that the generator of CSGAN is a continuous function that can approximate f_s^{-1} well, which means f_s^{-1} could be continuous.

3.2.6. Results: Bayesian inference of log permeability field via $p(\mathbf{G} | \mathbf{f}_{1:j})$

Next, we perform a series of inspections on the latent CSGANs approximating $p(\mathbf{G} | \mathbf{f}_{1:j})$ in a manner similar to the previous section: that is, we inspect the CSGANs' accuracy in both the output space of \mathbf{G} and the input spaces of $\mathbf{f}_{1:j}$. The results are illustrated in Fig. 13.

Fig. 13(a) shows the predictions made by the CSGANs' fitting $p(\mathbf{G} | \mathbf{f}_1)$, $p(\mathbf{G} | \mathbf{f}_{1:3})$ and $p(\mathbf{G} | \mathbf{f}_{1:5})$. Here we omit the CSGANs fitting $p(\mathbf{G} | \mathbf{f}_{1:2})$ and $p(\mathbf{G} | \mathbf{f}_{1:4})$, for both space reasons and the fact that their predictions have similar visual patterns to the other three. In contrast to Fig. 12(a), we can see that the predictions in each test case no longer look identical to each other, and they do not match the groundtruth. This can be anticipated, as Section 3.2.4 reveals that the oil-cut curves $\mathbf{f}_{1:j}$ do not have as many intrinsic dimensions as that of the saturation field \mathbf{S} , so the preimage theorem—if still applicable here—implies that $\mathbf{f}_{1:j}$'s preimages in the \mathbf{G} space should be submanifolds of much higher dimensions (22 ~ 36 from LVAEs' estimates) than the 0-dimensional $f_s^{-1}(\mathbf{S})$. Therefore, in each test case, there seems to be a large scatter around the groundtruth \mathbf{G} of the given $\mathbf{f}_{1:j}$, for which the log permeability fields can satisfactorily produce the observed $\mathbf{f}_{1:j}$, resulting in an overall high variance of the posterior density.

Because the visual inspection can no longer help grasp the accuracy and uncertainty of CSGANs, we then try the relative error and relative variation to quantitatively probe the CSGANs, as shown in Fig. 13(b). Similar to Section 3.2.5, for each test case we produce 50 predictions to evaluate these quantities. It appears that all the CSGANs have comparable average and minimum relative errors, and these errors are much larger than their counterparts in Fig. 12(b), which just reaffirms our observation in Fig. 13(a) and concludes nothing more about the CSGANs' accuracy. Also, we should not expect these 50 predictions to greatly drag down the minimum relative error to a level comparable to that in Fig. 12(b). This is because the preimages of $\mathbf{f}_{1:j}$ probably constitutes 22 ~ 36-D manifolds, but due to the curse of dimensionality, 50 samples can only cover a negligible part of each, so in each test case, it is not likely that one of the predictions will fall into the groundtruth \mathbf{G} 's vicinity. However, the distribution of relative variation reveals a pattern anticipated in Section 3.2.2: the more intrinsic dimensions in the input condition, the larger the variance of the posterior. It shows that the CSGANs taking more \mathbf{f}_i as input have less variation in their predictions. The preimage theorem sheds some light on this phenomenon: the dimensionality of $\mathbf{f}_{1:j}$'s preimage—which comprises all the valid predictions—decreases as j increases.

Investigating the \mathbf{G} space alone thus cannot determine the accuracy of these CSGANs, so we switch to the input condition space of $\mathbf{f}_{1:j}$. In order to assess whether or not the predicted log permeability fields $\hat{\mathbf{G}}$ are in or, practically, near the preimages of $\mathbf{f}_{1:j}$, for each of 1000 selected test cases, we simulate the corresponding oil-cut curves $\hat{\mathbf{f}}_{1:j}^{(k)}$ of 5 predictions $\hat{\mathbf{G}}^{(k)}$ with $k \in \{1, 2, 3, 4, 5\}$ and compare them with the groundtruth $\mathbf{f}_{1:j}$ both visually and quantitatively as in Section 3.2.5. Fig. 13(c) shows an example case predicted by the CSGAN learning $p(\mathbf{G} | \mathbf{f}_1)$. Here the figure on the top shows \mathbf{f}_1 along with the $\hat{\mathbf{f}}_1^{(k)}$ derived from its five predicted $\hat{\mathbf{G}}^{(k)}$, and the one on the bottom shows \mathbf{f}_5 along with the five $\hat{\mathbf{f}}_5^{(k)}$. We can see this CSGAN—informed only by \mathbf{f}_1 and not by \mathbf{f}_5 —produces $\hat{\mathbf{G}}^{(k)}$ that results in oil-cut curves that match \mathbf{f}_1 better than \mathbf{f}_5 . Fig. 13(d) further plots the error distributions in the \mathbf{f} spaces to provide a more comprehensive view. The x -axis corresponds to the five latent CSGANs taking different number of oil-cut curves $\mathbf{f}_{1:j}$ as input condition, the y -axis shows the *mean* of the average relative errors of $\hat{\mathbf{f}}_i$ against \mathbf{f}_i evaluated across the 1000 test cases, and different colors in Fig. 13(d) represent the relative errors at different production wells. We notice a clear pattern that whenever the oil-cut \mathbf{f}_i at a given well i is used to inform certain CSGANs as an input condition, those CSGANs generate predictions $\hat{\mathbf{G}}$ which in general yield $\hat{\mathbf{f}}_i$ that match \mathbf{f}_i better. This indicates that the CSGANs indeed have learned to produce $\hat{\mathbf{G}}$ close to the preimage of its input condition $\mathbf{f}_{1:j}$. However, the learning accuracy remains poorer compared to the $p(\mathbf{G} | \mathbf{S})$ case of Section 3.2.5, Fig. 12. The standard deviations of these well-wise average relative

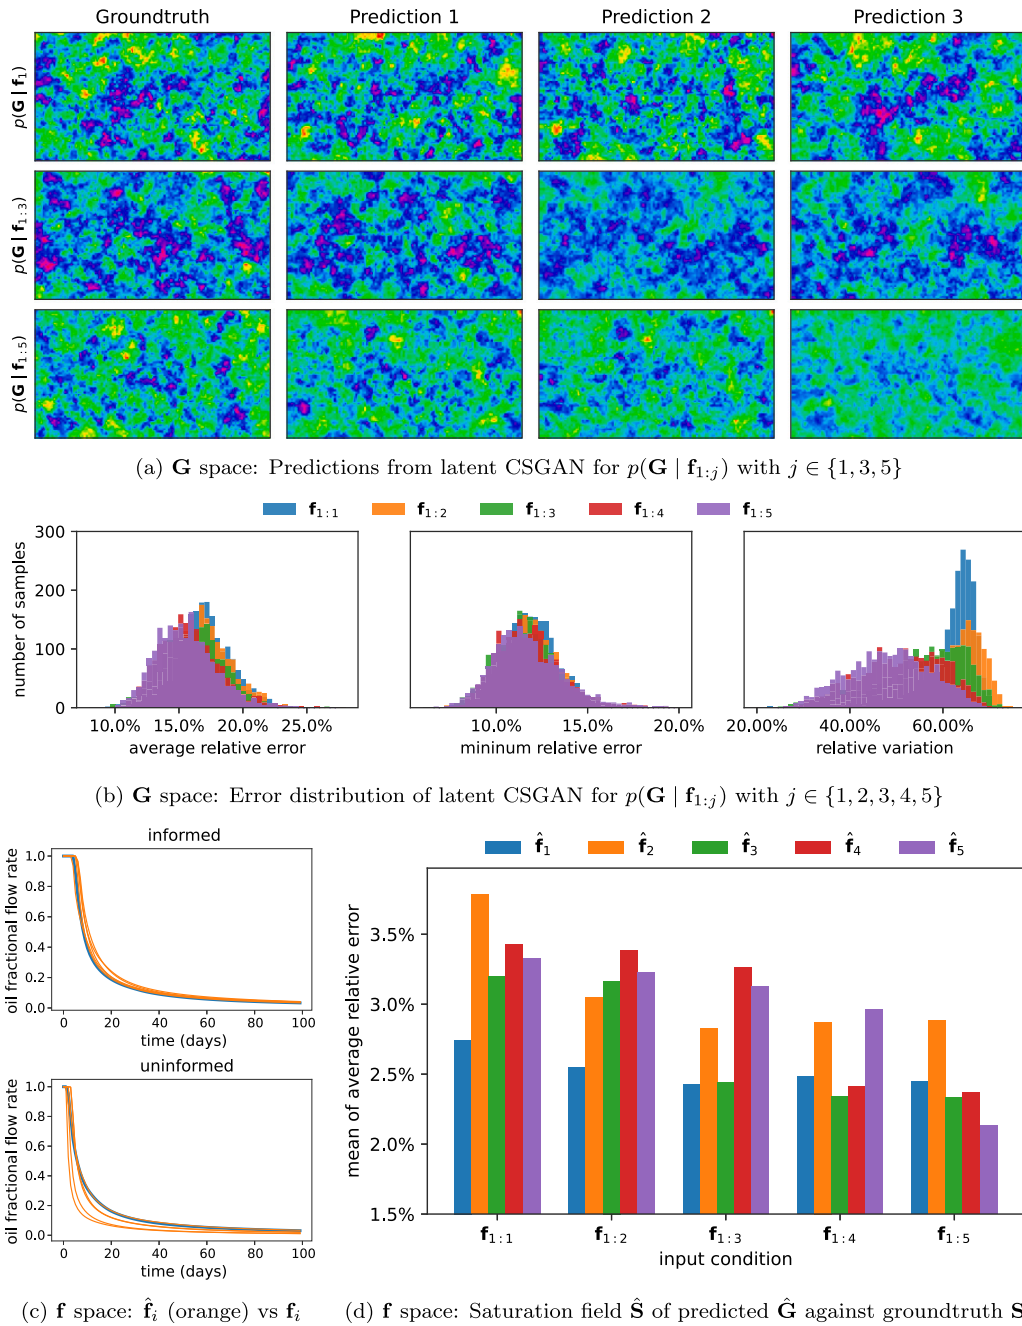


Fig. 13. Predictions and errors of latent CSGANs for $p(\mathbf{G} | \mathbf{f}_{1:j})$ in \mathbf{G} and \mathbf{f} space.

errors in Fig. 13(d)—evaluated across the 1000 test cases—turn out to be all around 1.5%, which is considerable compared to the errors’ mean values typically around 2.5%. This means the CSGANs are not able to consistently generate accurate $\hat{\mathbf{G}}$ for different $\mathbf{f}_{1:j}$ among the test cases.

3.2.7. Results: Correlation between uncertainty of posterior and intrinsic dimension of input condition

The distributions of relative variations in Figs. 12(b) and 13(b) have already given us an idea about the correlation between the input’s dimensionality and the output’s uncertainty as anticipated in Section 3.2.2. Here, we further confirm it with the more rigorous metric—entropy (Cover & Thomas, 2006). The differential entropy of the latent CSGANs’ pushforward measure cannot be evaluated analytically, since the density function’s closed-form expression is not available. We use instead the KSG entropy estimate based on k -nearest neighbor

distances, as introduced in Kraskov et al. (2004). We assume that the CSGANs have captured the target posteriors well enough—based on the analysis in Sections 3.2.5 and 3.2.6—so that the CSGANs’ uncertainty represents accurately the target posteriors’ variance.

To assess the uncertainty of the six CSGANs that approximate $p(\mathbf{G} | \mathbf{S})$ and $p(\mathbf{G} | \mathbf{f}_{1:j})$, $j \in \{1, 2, 3, 4, 5\}$ respectively, and, more importantly, to put them on a unified scale for comparison, we first select 2000 test cases, then for each of them, we feed its input condition $(\mathbf{S}, \mathbf{f}_{1:j})$ to different CSGANs to let each CSGAN use different components of $(\mathbf{S}, \mathbf{f}_{1:j})$ to generate 1000 predictions $\hat{\mathbf{G}}$. Thereafter, for each CSGAN in each test case, we feed its 1000 predictions to the KSG estimator to evaluate the entropy of the distribution of $\hat{\mathbf{G}}$, so that after all evaluations, we eventually obtain a 6×2000 grid of entropy estimates. Then we display the entropy distribution of the six CSGANs in Fig. 14. The result reaffirms what we observed so far: the higher the intrinsic dimension of

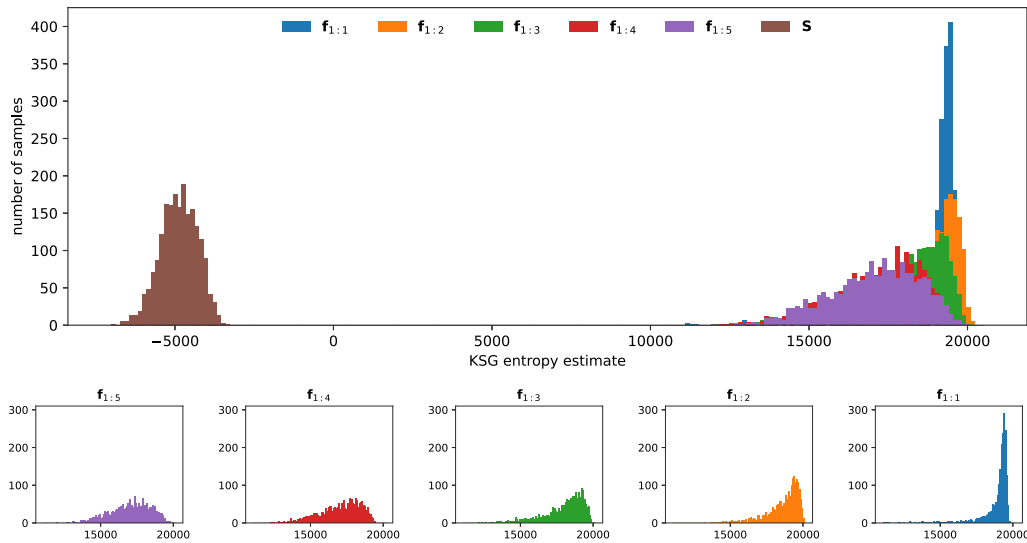


Fig. 14. Entropy distribution of different latent CSGANs' predictions. Here each label refers to the input condition that the CSGAN takes, so as to identify each CSGAN.

the input condition is, the lower the variance of the posterior. Fig. 14 shows that $p(\mathbf{G} | \mathbf{S})$ has much less entropy than $p(\mathbf{G} | \mathbf{f}_{1:j})$. Moreover, though being subjective, it appears that among $p(\mathbf{G} | \mathbf{f}_{1:j})$, the reduction in entropy starts to saturate as j increases beyond 4, which aligns well with the saturation of $\mathbf{f}_{1:j}$'s intrinsic dimension as observed in Section 3.2.4, Fig. 11.

4. Conclusion

We have thus far presented a holistic framework for Bayesian inverse problems in arbitrarily high dimensional scenarios and for an arbitrary complexity of the forward model that can be characterized by a highly nonlinear computer code. The proposed methodology leverages the novel Least Volume Autoencoders for identifying low dimensional manifolds on which the input and output data reside, and after encoding the observations, a conditional generative adversarial network is trained using a Sinkhorn divergence-based loss function for further improving its learning potential compared to its traditional GAN counterparts. On one hand, the proposed autoencoder structure used for dimensionality reduction allows working in extremely high dimensional settings as long as the data can be compressed to a low dimensional space. On the other hand, the GAN framework can tackle highly nonlinear forward models. For both tasks, all that is needed is to identify a rich enough architecture for the neural networks involved in order to guarantee an application-specific accurate inference procedure. The performance of the proposed approach is illustrated through two numerical examples, and the combination of computational efficiency along with the accuracy of the posterior solutions is unparalleled. We have provided a systematic study on how the quality of the posterior solution to the inverse problem improves as the intrinsic dimensionality of the observable quantity increases, thus providing additional information about the input, or, in other words, reducing the intrinsic dimensionality of the solution according to the preimage theorem.

In the challenging setting of Section 3.2, we have reached a limitation where the CSGAN cannot make consistently accurate prediction on different test cases when the observable quantity's intrinsic dimension is low (see Section 3.2.6). For complicated high dimensional inverse problems, there could be many elusive culprits that are hard to track down. However, we give three major hypotheses here to provide some insights on the limitations for arbitrarily complex experiments.

- Our first hypothesis is inspired by our observation and discussion in Section 3.1.3. There, the low dimensional CSGAN learning the simpler bimodal distribution generates invalid predictions connecting the two modes (Fig. 6), due to the *intrinsic flaw* of pushforward models (see Section 3.1.3). The same issue can be possibly observed in the high dimensional cases of $p(\mathbf{G} | \mathbf{f}_{1:j})$ in Section 3.2.6. The preimage theorem states that under certain assumptions, the preimage of a given observed quantity is a submanifold of a certain dimension, which is usually conceived as a high dimensional curved surface in the inferred parameter space. Yet this surface can have complicated topology that the pushforward models cannot readily handle. For instance, it might constitute multiple disconnected components and thus can be intuitively similar to a union of several disjoint surfaces; or it is an enclosed surface like a sphere. In these cases, the pushforward model's intrinsic flaw could be triggered and would lead to inferior learning results. In contrast, conditioning the log-permeability on the oil saturation field in Section 3.2 has a much simpler, probably unimodal distribution, such that the preimage of \mathbf{S} consists of only one point, so it does not suffer from this flaw (see Section 3.2.5).
- Our second hypothesis is that the mapping from log-permeability to the oil production curves might have large local Lipschitz constant in many regions, so that even if a CSGAN makes predictions very close to the target preimage of a given observed curve, the resulting posterior prediction may still be far away from the original observation due to this high sensitivity. If the local constant varies dramatically across different regions, it can lead to the large variation in prediction error.
- Our third hypothesis is data deficiency, which commonly plagues most data-driven learning tasks across all research fields. Although to date it is still tricky to determine the *sample complexity*—i.e., how many data samples we need exactly for a given data-driven learning tasks to achieve satisfying accuracy, it is at least both theoretically (Narayanan & Mitter, 2010; Narayanan & Niyogi, 2009) and empirically (Pope, Zhu, Abdelkader, Goldblum, & Goldstein, 2020) discovered that it depends on the *intrinsic dimension*, *intrinsic volume* and *curvature* of the dataset, whereas independent of the extrinsic dimension—i.e., the dimension of the dataset's ambient space. Therefore, learning $p(\mathbf{G} | \mathbf{f}_{1:j})$ could require more data than learning $p(\mathbf{G} | \mathbf{S})$, since the preimages of $\mathbf{f}_{1:j}$ have much larger intrinsic dimensions and probably much larger curvatures than \mathbf{S} 's.

In addition, this implies for high dimensional problems, we can only *heuristically* determine how many samples we need for each learning task, as the number of variables in the dataset—*i.e.*, the extrinsic dimension—provides no useful information about sample complexity. However, when the data is abundant, sample complexity may not matter that much. For instance, in the future we may apply the methods in the paper to real-world applications where unlimited data can be collected easily from sensors rather than costly simulations.

Nevertheless, at this moment, no certain conclusion can be drawn for this complicated high dimensional problem, and we leave further investigations to future research.

At last, the discerned correlation, investigated in Section 3.2.7, supports our hypothesis in Section 3.2.2 based on the preimage theorem, and may inspire many novel applications in the future. For instance, we may use this information about dimensionality retrieved by LVAEs to anticipate which combination of observable variables provides the largest amount of information for inferring the unknown system parameters of interest, how certain the inference might be, whether the distribution of an array of sensors on a structure is optimal and so forth. However, it should be noted here that the latent dimension that an LVAE retrieves is an *upper-bounding global estimate* of the dataset’s true intrinsic dimension (Chen & Fuge, 2024). In many situations, the gap between them may not be closed. For instance, when the given n -D data manifold cannot be embedded in \mathbb{R}^n (*e.g.*, it is a sphere, torus, Klein bottle, *etc.*), or when the dataset is a union of several manifolds intersecting each other, such that some part of it is not locally Euclidean, or when the data manifold forms a nontrivial knot. We should always treat the LVAEs’ results with these caveats in mind.

CRedit authorship contribution statement

Qiuyi Chen: Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation. **Panagiotis Tsilifis:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Resources, Methodology, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Mark Fuge:** Writing – review & editing, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Panagiotis Tsilifis reports financial support was provided by Advanced Research Projects Agency-Energy. Mark Fuge reports financial support was provided by Advanced Research Projects Agency-Energy. Mark Fuge reports financial support was provided by National Science Foundation. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The information, work, or data presented herein was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0001295. We also acknowledge the support from the National Science Foundation, United States through award #1943699 as well as ARPA-E award DE-AR0001216. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government of any agency thereof.

Table B.3

Architecture of LVAE for y_1 .

Input/Output	Encoder	Decoder
$y_1 \in \mathbb{R}^{256 \times 1}$	$y_1 \rightarrow \text{Conv}_{32} \rightarrow \text{LeakyReLU}$	$z \rightarrow \text{Scale}_{256}$
$z \in \mathbb{R}^{50}$	$\rightarrow \text{Conv}_{64} \rightarrow \text{LeakyReLU}$	$\rightarrow \text{SN-Linear}_{128} \rightarrow \text{LeakyReLU}$
	$\rightarrow \text{Conv}_{64} \rightarrow \text{LeakyReLU}$	$\rightarrow \text{SN-Linear}_{256} \rightarrow \text{LeakyReLU}$
	$\rightarrow \text{Conv}_{128} \rightarrow \text{LeakyReLU}$	$\rightarrow \text{SN-Linear}_{2048} \rightarrow \text{LeakyReLU}$
	$\rightarrow \text{Conv}_{256} \rightarrow \text{LeakyReLU}$	$\rightarrow \text{Reshape}_{2048 \rightarrow 8 \times 256}$
	$\rightarrow \text{Reshape}_{8 \times 256 \rightarrow 2048}$	$\rightarrow \text{SN-Deconv}_{128} \rightarrow \text{LeakyReLU}$
	$\rightarrow \text{Linear}_{256} \rightarrow \text{LeakyReLU}$	$\rightarrow \text{SN-Deconv}_{64} \rightarrow \text{LeakyReLU}$
	$\rightarrow \text{Linear}_{128} \rightarrow \text{LeakyReLU}$	$\rightarrow \text{SN-Deconv}_{64} \rightarrow \text{LeakyReLU}$
	$\rightarrow \text{Linear}_{50} \rightarrow z$	$\rightarrow \text{SN-Deconv}_{32} \rightarrow \text{LeakyReLU}$
		$\rightarrow \text{SN-Deconv}_1 \rightarrow \text{Sigmoid} \rightarrow y_1$

Table B.4

Architecture of latent CSGAN for $p(\xi | y_1)$.

Generator
$z_{y_1} \times u \in \mathbb{R}^{7+2}$
$\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$
$\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$
$\rightarrow \text{Linear}_{64} \rightarrow \text{LeakyReLU}$
$\rightarrow \text{Linear}_2 \rightarrow \xi \in \mathbb{R}^2$

Appendix A. Non-collapsing property of the Least Volume Autoencoder when the decoder is K -Lipschitz

Theorem. Let (e_θ, g_θ) be a continuous autoencoder with $e_\theta : X \rightarrow Z$ and $g_\theta : Z \rightarrow X$, where g_θ is a K -Lipschitz function for any $\theta \in R^{d_\theta}$. Then for any finite set of training points $\mathcal{X} \subset X$, and for any arbitrarily complex architecture on the autoencoder, there exists a nondegenerate solution θ^* , that is the mapping $\mathcal{Z} = e_{\theta^*}(\mathcal{X})$ does not collapse all points to zero.

Proof. Let $x_1, x_2 \in \mathcal{X}$ be any two training points with $\|x_1 - x_2\|_{\mathcal{X}} = M > 0$ and $z_1 = e_\theta(x_1)$, $z_2 = e_\theta(x_2)$, $x'_1 = g_\theta(z_1)$, $x'_2 = g_\theta(z_2)$. With an arbitrarily complex architecture, there is always a set of hyperparameters θ^* such that $(e_{\theta^*}, g_{\theta^*})$ can achieve arbitrarily low reconstruction error. Choose θ^* such that $\|x_1 - g_{\theta^*} \circ e_{\theta^*}(x_1)\|_{\mathcal{X}} = \|x_1 - x'_1\|_{\mathcal{X}} < \delta$, and $\|x_2 - g_{\theta^*} \circ e_{\theta^*}(x_2)\|_{\mathcal{X}} = \|x_2 - x'_2\|_{\mathcal{X}} < \delta$ for some $\delta < M/3$. If the autoencoder is a degenerate one, then $\|e_\theta(x_1) - e_\theta(x_2)\|_{\mathcal{Z}}$ is arbitrarily small, therefore it can be taken to be smaller than δ/K . This implies

$$\begin{aligned} M &= \|x_1 - x_2\|_{\mathcal{X}} \leq \|x_1 - x'_1\|_{\mathcal{X}} + \|x_2 - x'_2\|_{\mathcal{X}} + \|x'_1 - x'_2\|_{\mathcal{X}} \\ &\leq \|x_1 - x'_1\|_{\mathcal{X}} + \|x_2 - x'_2\|_{\mathcal{X}} + K\|z_1 - z_2\|_{\mathcal{Z}} \\ &< \delta + \delta + \delta < M, \end{aligned}$$

where the first inequality follows by using the triangular inequality and the second inequality follows by applying the K -Lipschitz condition on the decoder. The above contradiction completes the proof. \square

Appendix B. Model training details of Section 3.1

Table B.3 shows the architecture of the LVAE for y_1 . Here all convolutional (Conv) and deconvolutional layers (Deconv) are 1D, and the prefix “SN” stands for “spectral-normalized”. The subscript of each Conv, Deconv and Linear layer indicates the number of output channels/features. The Scale layer simply increases the network’s Lipschitz constant by scaling up each input neuron by the factor specified in its subscript, in order to accelerate the LVAE’s convergence. All the LeakyReLU activations have $\alpha = 0.2$.

Appendix C. Model training details of Section 3.2

See Tables C.5–C.8.

Table C.5
Architectures of LVAEs for G , S , $f_{1:j}$.

Input/Output	Encoder	Decoder
$G \in \mathbb{R}^{128 \times 64 \times 1}$ $z \in \mathbb{R}^{100}$	$G \rightarrow \text{Conv}_{64} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{64} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{128} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{128} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{256} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Reshape}_{4 \times 2 \times 256 \rightarrow 2048}$ $\rightarrow \text{Linear}_{1024} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{100} \rightarrow z$	$z \rightarrow \text{Scale}_{8192}$ $\rightarrow \text{SN-Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Linear}_{1024} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Linear}_{2048} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Reshape}_{2048 \rightarrow 4 \times 2 \times 256}$ $\rightarrow \text{SN-Deconv}_{128} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{128} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{64} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{64} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_1 \rightarrow \text{Sigmoid} \rightarrow G$
$S \in \mathbb{R}^{128 \times 64 \times 1}$ $z \in \mathbb{R}^{100}$	$S \rightarrow \text{Conv}_{64} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{128} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{256} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{1024} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Reshape}_{4 \times 2 \times 1024 \rightarrow 8192}$ $\rightarrow \text{Linear}_{1024} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{100} \rightarrow z$	$z \rightarrow \text{Scale}_{8192}$ $\rightarrow \text{SN-Linear}_{1024} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Linear}_{8192} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Reshape}_{8192 \rightarrow 4 \times 2 \times 1024}$ $\rightarrow \text{SN-Deconv}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{256} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{128} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{64} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{64} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_1 \rightarrow \text{Sigmoid} \rightarrow S$
$f_{1:j} \in \mathbb{R}^{128 \times j}$ $z \in \mathbb{R}^{50}$	$f_{1:j} \rightarrow \text{Conv}_{64j} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{64j} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{128j} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Conv}_{128j} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Reshape}_{8 \times 128j \rightarrow 1024j}$ $\rightarrow \text{Linear}_{1024} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{50} \rightarrow z$	$z \rightarrow \text{Scale}_{e_{128j}}$ $\rightarrow \text{SN-Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Linear}_{1024} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Linear}_{1024j} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Reshape}_{1024j \rightarrow 8 \times 128j}$ $\rightarrow \text{SN-Deconv}_{128j} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{64j} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_{64j} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{SN-Deconv}_j \rightarrow \text{Sigmoid} \rightarrow f_{1:j}$

Table C.6
Hyperparameters of LVAEs for G , S , $f_{1:j}$.

	Batch size	λ	η	Learning rate (Adam)	Epochs
G	50	1	1/8192	0.0001	6000
S	50	5	1/8192	0.0001	6000
f_1	50	0.016	1/128	0.0001	1000
$f_{1:2}$	50	0.008	1/128	0.0001	1000
$f_{1:3}$	50	0.004	1/128	0.0001	1000
$f_{1:4}$	50	0.003	1/128	0.0001	1000
$f_{1:5}$	50	0.0025	1/128	0.0001	1000

Table C.7
Architecture of latent CSGANs for $p(G|S)$ and $p(G|f_{1:j})$.

$p(G S)$	$p(G f_{1:j})$
$z_S \times u \in \mathbb{R}^{40+40}$ $\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{256} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{40} \rightarrow z_G \in \mathbb{R}^{40}$	$z_{f_{1:j}} \times u \in \mathbb{R}^{\dim f_{1:j} + 40}$ $\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{512} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{256} \rightarrow \text{LeakyReLU}$ $\rightarrow \text{Linear}_{40} \rightarrow z_G \in \mathbb{R}^{40}$

$\dim f_{1:j} = \{4, 8, 12, 16, 18\}$ for $j = \{1, 2, 3, 4, 5\}$ (see Fig. 11(a)).

Table C.8
Hyperparameters of latent CSGANs for $p(G|S)$ and $p(G|f_{1:j})$.

	Batch size	ρ of S_ρ	Learning rate (Adam)	Epochs
S	128	0.3	0.0001	2000
$f_{1:j}$	128	0.3	0.0001	3000

Appendix D. Tips for hyperparameter tuning

The training of the LVAEs involves the hyperparameters λ , η and the decoder's Lipschitz constant K . A comprehensive guideline for tuning them can be found in Appendix C of Chen and Fuge (2024).

The training of the Sinkhorn GANs requires setting the hyperparameter ρ for the Sinkhorn divergence S_ρ . A detailed study of its effect can be found in Feydy et al. (2019), Genevay et al. (2019, 2018). In brief,

a smaller ρ offers S_ρ better resolution to differentiate between different data samples, thus leads to better alignment of the generator g 's image with the target dataset (see examples in Feydy, 2019), but it incurs worse sample complexity and lower Sinkhorn algorithm convergence speed. So ρ makes trade-off between quality and efficiency. The ρ in Table C.8 were not tuned meticulously. We basically applied $\rho = 0.3$ to the $p(G|S)$ case at first and obtained ideal results, then we straightforwardly applied it to all the experiments of $p(G|f_{1:j})$ without further change. If the precision of Sinkhorn GANs in approximating the target distribution \mathbb{P}_r is the primary concern in certain applications, ρ may be fine-tuned against a certain statistical divergence $D(\mathbb{P}_g, \mathbb{P}_r)$ independent of ρ , e.g., the MMD (Smola et al., 2007).

Data availability

Data will be made available on request.

References

- Aarnes, J., Gimse, T., & Lie, K. (2007). An introduction to the numerics of flow in porous media using matlab. In *Geometric modelling, numerical simulation, and optimization: applied mathematics at SINTEF* (pp. 265–306). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Aggarwal, C. C., Hinneburg, A., & Keim, D. A. (2001). On the surprising behavior of distance metrics in high dimensional space. In *Database theory—ICDT 2001: 8th international conference London, UK, January 4–6, 2001 Proceedings 8* (pp. 420–434). Springer.
- Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. In *International conference on machine learning* (pp. 214–223). PMLR.
- Beyer, K., Goldstein, J., Ramakrishnan, R., & Shaft, U. (1999). When is “nearest neighbor” meaningful? In *Database theory—ICDT’99: 7th international conference Jerusalem, Israel, January 10–12, 1999 Proceedings 7* (pp. 217–235). Springer.
- Bilionis, I., & Zabarar, N. (2012). Multi-output local Gaussian process regression: Applications to uncertainty quantification. *Journal of Computational Physics*, 231, 5718–5746. <http://dx.doi.org/10.1016/j.jcp.2012.04.047>.
- Bilionis, I., & Zabarar, N. (2013). Solution of inverse problems with limited forward solver evaluations: a Bayesian perspective. *Inverse Problems*, 30, Article 015004. <http://dx.doi.org/10.1088/0266-5611/30/1/015004>.
- Bilionis, I., Zabarar, N., Konomi, B., & Lin, G. (2013). Multi-output separable Gaussian process: Towards an efficient, fully Bayesian paradigm for uncertainty quantification. *Journal of Computational Physics*, 241, 212–239.
- Broderick, T., Boyd, N., Wibisono, A., Wilson, A., & Jordan, M. (2013). Streaming variational bayes. In *Advances in neural information processing systems: Vol. 26*.
- Brown, B. C., Caterini, A. L., Ross, B. L., Cresswell, J. C., & Loaiza-Ganem, G. (2022). Verifying the union of manifolds hypothesis for image data. In *The eleventh international conference on learning representations*.
- Calvetti, D., & Somersalo, E. (2018). Inverse problems: From regularization to Bayesian inference. *Wiley Interdisciplinary Reviews: Computational Statistics*, 10, Article e1427.
- Chen, Q., & Fuge, M. (2024). Compressing latent space via least volume. In *ICLR*.
- Chen, Q., Pope, P., & Fuge, M. (2022). Learning airfoil manifolds with optimal transport. In *AIAA SCITECH 2022 forum* (p. 2352).
- Chen, Q., Wang, J., Pope, P., Chen, W., & Fuge, M. (2022). Inverse design of two-dimensional airfoils using conditional generative models and surrogate log-likelihoods. *Journal of Mechanical Design*, 144, Article 021712.
- Christie, M., & Blunt, M. (2001). Tenth SPE comparative solution project: A comparison of upscaling techniques. In *SPE reservoir simulation symposium*. Society of Petroleum Engineers, (Chapter).
- Cornish, R., Caterini, A., Deligiannidis, G., & Doucet, A. (2020). Relaxing bijectivity constraints with continuously indexed normalising flows. In *International conference on machine learning* (pp. 2133–2143). PMLR.
- Cotter, S., Dashti, M., & Stuart, A. (2010). Approximation of Bayesian inverse problems for PDEs. *SIAM Journal on Numerical Analysis*, 48, 322–345.
- Cover, T. M., & Thomas, J. A. (2006). *Wiley series in telecommunications and signal processing. Elements of information theory*. USA: Wiley-Interscience.
- Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems*, 26.
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4), 303–314.
- Dake, L. (1983). *Fundamentals of reservoir engineering*. Elsevier.
- Dasgupta, A., Patel, D., Ray, D., Johnson, E., & Oberai, A. (2024). A dimension-reduced variational approach for solving physics-based inverse problems using generative adversarial network priors and normalizing flows. 420, Article 116682.
- Dinh, L., Krueger, D., & Bengio, Y. (2015). NICE: non-linear independent components estimation. In *3rd international conference on learning representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Workshop track proceedings*. URL <http://arxiv.org/abs/1410.8516>.

- Dinh, L., Sohl-Dickstein, J., & Bengio, S. (2017). Density estimation using real NVP. In *5th international conference on learning representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference track proceedings*. OpenReview.net, URL <https://openreview.net/forum?id=HkpbH9lx>.
- Eklblom, M., Tuppi, L., Shemyakin, V., Haario, H., & Jarvinen, H. (2020). Algorithmic tuning of spread-skill relationship in ensemble forecasting systems. *Quarterly Journal of the Royal Meteorological Society*, *146*, 598–612.
- Engl, H., Hanke, M., & Neubauer, A. (1996). *Regularization of inverse problems: Vol. 375*, Springer Science & Business Media.
- Fan, Y., & Ying, L. (2020). Solving electrical impedance tomography with deep learning. *Journal of Computational Physics*, *404*, Article 109119.
- Feydy, J. (2019). Geometric loss functions between sampled measures, images and volumes. URL <https://www.kernel-operations.io/geomloss>.
- Feydy, J., Séjourné, T., Vialard, F.-X., Amari, S.-i., Trounev, A., & Peyré, G. (2019). Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd international conference on artificial intelligence and statistics* (pp. 2681–2690). PMLR.
- Geneva, N., & Zabaras, N. (2019). Quantifying model form uncertainty in Reynolds-averaged turbulence models with Bayesian deep neural networks. *Journal of Computational Physics*, *383*, 125–147.
- Geneva, A., Chizat, L., Bach, F., Cuturi, M., & Peyré, G. (2019). Sample complexity of sinkhorn divergences. In *The 22nd international conference on artificial intelligence and statistics* (pp. 1574–1583). PMLR.
- Geneva, A., Peyré, G., & Cuturi, M. (2018). Learning generative models with sinkhorn divergences. In *International conference on artificial intelligence and statistics* (pp. 1608–1617). PMLR.
- Ghosh, S., Anantha Padmanabha, G., Peng, C., Andreoli, V., Atkinson, S., Pandita, P., et al. (2022). Inverse aerodynamic design of gas turbine blades using probabilistic machine learning. *Journal of Mechanical Design*, Article 021706.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, *27*.
- Gouk, H., Frank, E., Pfahringer, B., & Cree, M. J. (2021). Regularisation of neural networks by enforcing lipschitz continuity. *Machine Learning*, *110*, 393–416.
- Guillemin, V., & Pollack, A. (2010). *Differential topology: Vol. 370*, American Mathematical Soc.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. C. (2017). Improved training of wasserstein gans. *Advances in Neural Information Processing Systems*, *30*.
- Haario, H., Laine, M., Lehtinen, M., Saksman, E., & Tamminen, J. (2004). Markov chain Monte Carlo methods for high dimensional inversion in remote sensing. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, *66*, 591–607.
- Haldemann, J., Ksoll, V., Walter, D., Alibert, Y., Klessen, R., Benz, W., et al. (2023). Exoplanet characterization using conditional invertible neural networks. *Astronomy and Astrophysics*, *672*, A180.
- Herbei, R., McKeague, I. W., & Speer, K. G. (2008). Gyres and jets: Inversion of tracer data for ocean circulation structure. *Journal of Physical Oceanography*, *38*(6), 1180–1202.
- Hoffman, M., Blei, D., Wang, C., & Paisley, J. (2013). Stochastic variational inference. *Journal of Machine Learning Research*.
- Hong, Y., Hwang, U., Yoo, J., & Yoon, S. (2019). How generative adversarial networks and their variants work: An overview. *ACM Computing Surveys*, *52*, 1–43.
- Hyvärinen, A., & Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural Networks*, *13*, 411–430.
- Isaac, T., Petra, N., Stadler, G., & Ghattas, O. (2015). Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet. *Journal of Computational Physics*, *296*, 348–368.
- Jabbar, A., Li, X., & Omar, B. (2021). A survey on generative adversarial networks: Variants, applications, and training. *ACM Computing Surveys*, *54*, 1–49.
- Jin, K. H., McCann, M. T., Froustey, E., & Unser, M. (2017). Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, *26*, 4509–4522. <http://dx.doi.org/10.1109/TIP.2017.2713099>.
- Kennedy, M., & O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, *63*, 425–464.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *3rd international conference on learning representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference track proceedings*. URL <http://arxiv.org/abs/1412.6980>.
- Kingma, D. P., & Dhariwal, P. (2018). Glow: Generative flow with invertible 1x1 convolutions. *Advances in Neural Information Processing Systems*, *31*.
- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. In *2nd international conference on learning representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference track proceedings*. URL <http://arxiv.org/abs/1312.6114>.
- Kobyzev, I., Prince, S. J., & Brubaker, M. A. (2020). Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *43*(11), 3964–3979.
- Kontolati, K., Tsilifis, P., Ghosh, S., Andreoli, V., Shields, M., & Wang, L. (2023). Multi-fidelity metamodeling in turbine blade airfoils via transfer learning on manifolds. In *AIAA SCITECH 2023 forum* (p. 0918).
- Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. *Physical Review E*, *69*(6), Article 066138.
- Laloy, E., Hérault, R., Jacques, D., & Linde, N. (2018). Training-image based geostatistical inversion using a spatial generative adversarial neural network. *Water Resources Research*, *54*, 381–406.
- Lee, J. (2010). *Introduction to topological manifolds: Vol. 202*, Springer Science & Business Media.
- Lee, J. M., & Lee, J. M. (2012). *Smooth manifolds*. Springer.
- Lee, M., & Min, K. (2022). MGCVAE: multi-objective inverse design via molecular graph conditional variational autoencoder. *62*, (pp. 2943–2950).
- Ma, X., & Zabaras, N. (2009a). An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *Journal of Computational Physics*, *228*, 3084–3113.
- Ma, X., & Zabaras, N. (2009b). An efficient Bayesian inference approach to inverse problems based on an adaptive sparse grid collocation method. *Inverse Problems*, *25*, Article 035013. <http://dx.doi.org/10.1088/0266-5611/25/3/035013>.
- Mardani, M., Gong, E., Cheng, J., Vasanawala, S., Zaharchuk, G., Xing, L., et al. (2018). Deep generative adversarial neural networks for compressive sensing MRI. *IEEE Transactions on Medical Imaging*, *38*, 167–179.
- Marzouk, Y., & Najm, H. (2009). Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *Journal of Computational Physics*, *228*, 1862–1902.
- Marzouk, Y., Najm, H., & Rahn, L. (2007). Stochastic spectral methods for efficient Bayesian solution of inverse problems. *Journal of Computational Physics*, *224*, 560–586.
- Marzouk, Y., & Xiu, D. (2009). A stochastic collocation approach to Bayesian inference in inverse problems. *Communications in Computational Physics*, *6*, 826–847.
- Meng, X., Babae, H., & Karniadakis, G. (2021). Multi-fidelity Bayesian neural networks: Algorithms and applications. *Journal of Computational Physics*, *438*, Article 110361.
- Miyato, T., Kataoka, T., Koyama, M., & Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. In *International conference on learning representations*.
- Mo, S., Zabaras, N., Shi, X., & Wu, J. (2019). Deep autoregressive neural networks for high-dimensional inverse problems in groundwater contaminant source identification. *Water Resources Research*, *55*, 3856–3881.
- Mo, S., Zabaras, N., Shi, X., & Wu, J. (2020). Integration of adversarial autoencoders with residual dense convolutional networks for estimation of non-Gaussian hydraulic conductivities. *Water Resources Research*, e2019WR026082.
- Narayanan, H., & Mitter, S. (2010). Sample complexity of testing the manifold hypothesis. *Advances in Neural Information Processing Systems*, *23*.
- Narayanan, H., & Niyogi, P. (2009). On the sample complexity of learning smooth cuts on a manifold. In *COLT*.
- Nowozin, S., Cseke, B., & Tomioka, R. (2016). F-gan: Training generative neural samplers using variational divergence minimization. *Advances in Neural Information Processing Systems*, *29*.
- Ollinaho, P., Järvinen, H., Bauer, P., Laine, M., Bechtold, P., Susiluoto, J., et al. (2014). Optimization of NWP model closure parameters using total energy norm of forecast error as a target. *Geoscientific Model Development*, *7*, 1889–1900.
- Orszag, S., & Bissonnette, L. (1967). Dynamical properties of truncated Wiener-Hermite expansions. *The Physics of Fluids*, *10*, 2603–2613.
- Padmanabha, G., & Zabaras, N. (2021). Solving inverse problems using conditional invertible neural networks. *Journal of Computational Physics*, *433*, Article 110194. <http://dx.doi.org/10.1016/j.jcp.2021.110194>.
- Patel, D., Ray, D., & Oberai, A. (2022). Solution of physics-based Bayesian inverse problems with deep generative priors. *Computer Methods in Applied Mechanics and Engineering*, *400*, Article 115428.
- Pearson, K. (1991). LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, *2*, 559–572.
- Petra, N., Martin, J., Stadler, G., & Ghattas, O. (2014). A computational framework for infinite-dimensional Bayesian inverse problems, Part II: Stochastic Newton MCMC with application to ice sheet flow inverse problems. *SIAM Journal on Scientific Computing*, *36*, A1525–A1555.
- Petra, N., Zhu, H., Stadler, G., Hughes, T., & Ghattas, O. (2012). An inexact Gauss-Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model. *Journal of Glaciology*, *58*, 889–903.
- Peyré, G., Cuturi, M., et al. (2019). Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, *11*(5–6), 355–607.
- Pope, P., Zhu, C., Abdelkader, A., Goldblum, M., & Goldstein, T. (2020). The intrinsic dimension of images and its impact on learning. In *International conference on learning representations*.
- Qin, S., & Jiang, T. (2018). Improved wasserstein conditional generative adversarial network speech enhancement. *EURASIP Journal on Wireless Communications and Networking*, *2018*, 181.
- Ray, D., Murgoitio-Esandi, J., Dasgupta, A., & Oberai, A. (2023). Solution of physics-based inverse problems using conditional generative adversarial networks with full gradient penalty. *Computer Methods in Applied Mechanics and Engineering*, *417*, Article 116338. <http://dx.doi.org/10.1016/j.cma.2023.116338>.
- Russell, B. (1988). *Introduction to seismic inversion methods (No. 2)*. SEG Books.
- Salmona, A., De Bortoli, V., Delon, J., & Desolneux, A. (2022). Can push-forward generative models fit multimodal distributions? *Advances in Neural Information Processing Systems*, *35*, 10766–10779.

- Shannon, C. E. (1949). Communication in the presence of noise. *Proceedings of the IRE*, 37(1), 10–21.
- Smola, A., Gretton, A., Song, L., & Schölkopf, B. (2007). A Hilbert space embedding for distributions. In *International conference on algorithmic learning theory* (pp. 13–31). Springer.
- Sohn, K., Lee, H., & Yan, X. (2015). Learning structured output representation using deep conditional generative models. *Advances in Neural Information Processing Systems*, 28.
- Springer, S., Haario, H., Susiluoto, J., Bibov, A., Davis, A., & Marzouk, Y. (2021). *Efficient Bayesian inference for large chaotic dynamical systems: Vol. 14*, (pp. 4319–4333).
- Stuart, A. (2010). Inverse problems: a Bayesian perspective. *Acta Numerica*, 19, 451–559.
- Tait, D., & Damoulas, T. (2020). Variational autoencoding of PDE inverse problems. arXiv preprint arXiv:2006.15641.
- Tarantola, A. (2005). *Inverse problem theory and methods for model parameter estimation*. Society for Industrial and Applied Mathematics.
- Tripathy, R., & Bilonis, I. (2018). Deep UQ: Learning deep neural network surrogate models for high dimensional uncertainty quantification. *Journal of Computational Physics*, 375, 565–588. <http://dx.doi.org/10.1016/j.jcp.2018.08.036>.
- Tsilifis, P., Bilonis, I., Katsounaros, I., & Zabaras, N. (2016). Computationally efficient variational approximations for Bayesian inverse problems. *Journal of Verification, Validation and Uncertainty Quantification*, 1, Article 031004. <http://dx.doi.org/10.1115/1.4034102>.
- Tsilifis, P., & Ghanem, R. (2017). Reduced Wiener chaos representation of random fields via basis adaptation and projection. *Journal of Computational Physics*, 341, 102–120.
- Tsilifis, P., & Ghanem, R. (2018). Bayesian adaptation of chaos representations using variational inference and sampling on geodesics. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474, Article 20180285.
- Tsilifis, P., Ghanem, R., & Hajali, P. (2017). Efficient Bayesian experimentation using an expected information gain lower bound. *SIAM/ASA Journal on Uncertainty Quantification*, 5, 30–62. <http://dx.doi.org/10.1137/15M1043303>.
- Tsilifis, P., Ghosh, S., & Andreoli, V. (2022). Inverse design under uncertainty using conditional normalizing flows. In *AIAA scitech 2022 forum* (p. 0631).
- Tsilifis, P., Huan, X., Safta, C., Sargsyan, K., Lacaze, G., Oefelein, J., et al. (2019). Compressive sensing adaptation for polynomial chaos expansions. *Journal of Computational Physics*, 380, 29–47.
- Tsilifis, P., Pandita, P., Ghosh, S., Andreoli, V., Vandeputte, T., & Wang, L. (2021). Bayesian learning of orthogonal embeddings for multi-fidelity Gaussian processes. *Computer Methods in Applied Mechanics and Engineering*, 386, Article 114147. <http://dx.doi.org/10.1016/j.cma.2021.114147>.
- Tsilifis, P., Papaioannou, I., Straub, D., & Nobile, F. (2020). Sparse Polynomial Chaos expansions using variational relevance vector machines. *Journal of Computational Physics*, 416, Article 109498.
- Wan, X., & Karniadakis, G. (2005). An adaptive multi-element generalized polynomial chaos method for stochastic differential equations. *Journal of Computational Physics*, 209, 617–642.
- Yang, L., Meng, X., & Karniadakis, G. (2021). B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data. *Journal of Computational Physics*, 425, Article 109913.
- Yi, X., Walia, E., & Babyn, P. (2019). Generative adversarial network in medical imaging: A review. *Medical Image Analysis*, 58, Article 101552.
- Zhang, C., Barbano, R., & Jin, B. (2021). Conditional variational autoencoder for learned image reconstruction. *Computation*, 9, 114.
- Zhu, Y., Zabaras, N., Koutsourelakis, P., & Perdikaris, P. (2019). Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics*, 394, 56–81.