

RESEARCH LETTER

10.1029/2018GL080404

Key Points:

- The need for identifying forward and reverse mappings is ubiquitous in all geophysical research fields
- This study shows that generative adversarial networks (GANs) can be used to learn the forward and inverse mappings at the same time
- GAN provides a new way of thinking in combining physical-based modeling with data-driven modeling and has broad applications for geosciences

Supporting Information:

- Supporting Information S1

Correspondence to:

A. Y. Sun,
alex.sun@beg.utexas.edu

Citation:

Sun, A. Y. (2018). Discovering state-parameter mappings in subsurface models using generative adversarial networks. *Geophysical Research Letters*, 45, 11,137–11,146. <https://doi.org/10.1029/2018GL080404>

Received 10 SEP 2018

Accepted 5 OCT 2018

Accepted article online 10 OCT 2018

Published online 20 OCT 2018

Discovering State-Parameter Mappings in Subsurface Models Using Generative Adversarial Networks

Alexander Y. Sun¹ 

¹Bureau of Economic Geology, Jackson School of Geosciences, The University of Texas at Austin, Austin, TX, USA

Abstract A fundamental problem in geophysical modeling is related to the identification and approximation of causal structures among physical processes. However, resolving the bidirectional mappings between physical parameters and model state variables (i.e., solving the forward and inverse problems) is challenging, especially when parameter dimensionality is high. Deep learning has opened a new door toward knowledge representation and complex pattern identification. In particular, the recently introduced generative adversarial networks (GANs) hold strong promises in learning cross-domain mappings for image translation. This study presents a state-parameter identification GAN (SPID-GAN) for simultaneously learning bidirectional mappings between a high-dimensional parameter space and the corresponding model state space. SPID-GAN is demonstrated using a series of representative problems from subsurface flow modeling. Results show that SPID-GAN achieves satisfactory performance in identifying the bidirectional state-parameter mappings, providing a new deep-learning-based, knowledge representation paradigm for a wide array of complex geophysical problems.

Plain Language Summary Development of physically based models requires two steps, mathematical abstraction (forward modeling) and parameter estimation (inverse modeling). A high-fidelity model requires high-quality parameter support. The need for identifying forward and reverse mappings (i.e., a function that associates element of one set to another) is thus ubiquitous in geophysical research. A significant challenge in geosciences is that geoparameters are spatially heterogeneous and high dimensional and yet can only be observed at limited locations. The conventional workflow, built on minimizing the model-observation mismatch at measurement locations, does not offer an efficient way for estimating the spatial structure of high-dimensional parameter fields. This work presents a deep-learning-based framework for identifying the state-parameter bidirectional mappings using the recently introduced generative adversarial networks (GANs). GANs have been shown to be adept at associating images from one domain to another. Its potential for discovering mappings in physically based models has not been demonstrated so far. This work shows that GAN can achieve high performance in learning bidirectional parameter-to-state mappings in physically based models, thus providing a new way of thinking and doing things in geosciences. The implication for additional applications in subsurface modeling is significant.

1. Introduction

Deep learning (DL) has achieved great success in image recognition and business intelligence over the past decade, continuously narrowing the gap between artificial intelligence and human intelligence. Tremendous interests exist in the geophysical research community to leverage the strength of DL for solving similar image recognition and prediction problems, such as land cover and land use classification (Castelluccio et al., 2015), extreme weather event forecasting (Liu et al., 2016; Shi et al., 2015), estimation of particulate matter levels (Li et al., 2017), and data imputation (Fang et al., 2017). To achieve high accuracy, many DL algorithms require a large amount of labeled training data (i.e., coobserved predictors and predictands), which is generally hard to acquire in geoscience domains due to the invisibility of subsurface processes and sparsity of in situ monitoring networks.

In machine learning, semisupervised learning has been used to tackle the issue of limited labeled data (Chapelle et al., 2006). As its name suggests, semisupervised learning sits in between the traditional unsupervised learning (all training data are unlabeled) and supervised learning (all training data are labeled).

Semisupervised learning methods use unlabeled data, together with limited labeled data, to build better machine learning models, under the assumption that unlabeled data are more abundant and carry information that is useful for the inference of target variables (Chapelle et al., 2006).

Semisupervised learning may help solve a fundamental problem in geosciences, namely, estimating the underlying generative model of sampled data so that new samples can be synthesized from the learned model. In fact, generative process modeling is at the core of all physical sciences, where mechanistic models have long been applied to extract, abstract, and approximate the observed causal structures in order to simulate samples of the underlying physical processes. In geosciences in particular, parametric forward modelings are carried out by solving partial differential equations (PDEs), governing the spatially and/or temporally varying subsurface physical processes, whereas inverse problems are formulated to identify the model parameters by using observations of state variables (Sun & Sun, 2015). Ideally, both the forward and inverse modeling should be done in a closed-loop manner such that new information can be continuously assimilated to reduce uncertainty. Thus, the need for resolving the bidirectional mappings between state and parameter spaces always exists. An outstanding challenge is that many subsurface processes are nonlinear, multiscale, and high dimensional, making it nontrivial to establish such mappings in practice.

The recently introduced generative adversarial networks (GANs), which may be considered a subclass of DL for semisupervised learning, hold strong promises not only for learning the generative processes of high-dimensional images with limited labeled data but also for translating seemingly unrelated images across different domains (Goodfellow et al., 2014). An open question is whether these interesting features of GANs can benefit the geophysical modeling community. Here I explore the bidirectional mapping capability of GANs and hypothesize that GANs may provide a new workflow for inferring parameter-state mappings. Many of the conventional parameter estimation methods are built on the minimization of certain distance measures between observed and simulated values by running a forward model iteratively, while uncertainty quantification is usually done a posteriori. Under the GAN framework, the model parameter space and state space are regarded as two inherently related image domains, and DL-based functional relationships are obtained to facilitate the cross-domain learning, namely, estimating parameters for given model states, and vice versa.

The main purpose of this study is to formulate a state-parameter identification GAN (SPID-GAN) for obtaining deep bidirectional representations of geophysical models. In the following, I first introduce the proposed SPID-GAN framework, which combines the traditional geostatistical simulation, physically based forward modeling, and the point-based parameter estimation workflow with cross-domain DL. The framework is demonstrated using two different examples from subsurface flow modeling, in which the model parameters are spatially heterogeneous, representing high-dimensional samples obtained from single-modal and bimodal distributions. I show that the DL-based SPID-GAN is well adept at learning the subtle spatial patterns in model states and parameter fields, thus representing a powerful tool for approximating the causal structures of physical models.

2. Methodology

2.1. SPID-GAN Framework

The original GAN introduced by Goodfellow et al. (2014) consists of a pair of discriminator and generator, designed to compete with each other as in a two-player game, thus the word adversarial. The role of the generator is to create *fake* samples that are indistinguishable from the training data, while the role of the discriminator is to classify the generated samples to determine whether they are real or fake. Let \mathcal{G} denote a generator model that defines a mapping $\mathcal{G} : X \rightarrow Y$; namely, it takes $\mathbf{x} \in X$ as input and generates a fake sample $\mathcal{G}(\mathbf{x})$ that has the same support as the training data $\mathbf{y} \in Y$. In addition, let \mathcal{D} denote a discriminator that determines whether a sample is drawn from the empirical distribution of training data $p_{\text{data}}(\mathbf{y})$ or from the generator distribution $p_{\text{model}}(\mathcal{G}(\mathbf{x}))$. The goal of the generator is thus to push its sample distribution $p_{\text{model}}(\mathcal{G}(\mathbf{x}))$ toward the data distribution $p_{\text{data}}(\mathbf{y})$. At optimality the discriminator is maximally confused and cannot distinguish real samples from ones that are fake (i.e., predicting with a probability of 0.5 for all inputs; Goodfellow, 2016). Note that similar principles are behind many Bayesian statistical inversion and ensemble-based data assimilation algorithms, in which the common goal is updating a prior distribution to a posterior distribution that reflects the newer information. In practice, many of the conventional methods are either limited to low-dimensional problems (e.g., particle filter and Markov chain Monte Carlo) or to multivariate Gaussian distributions (e.g., the ensemble Kalman filter). In comparison, GANs make no such assumptions on the distributions of domain data.

Existing GANs differ by how domains are defined and what cross-domain mappings are to be learned. The generator in the original GAN takes a sample from a low-dimensional latent space (i.e., a random noise vector) and turns it into a real image (e.g., a car). Such a generative process resembles the principal component analysis-based random field simulation algorithm commonly used in geostatistics (Satiya & Caers, 2015). The main difference is that GANs train DL models to learn complex structural patterns embedded in the training data, while the eigenvector-based representation in principal component analysis is linear and restricted to the second-order statistics. Since the work of Goodfellow et al. (2014), a large number of GAN models have been introduced for cross-domain learning. So far, GANs have been demonstrated in a number of inspiring applications, such as (a) image superresolution, where low-resolution images are used to generate their high-resolution counterparts, examples include deep-convolution GAN (Radford et al., 2015) and superresolution GAN (Ledig et al., 2017), and (b) cross-domain image-to-image translation, where labeled information in the form of either text descriptions or images is used to generate images in another domain, examples include the conditional GAN (Mirza & Osindero, 2014), coupled GAN (Liu & Tuzel, 2016), DiscoGAN (Kim et al., 2017), and DualGAN (Yi et al., 2017). Newer GANs can perform direct cross-domain learning without using low-dimensional latent space vectors as done in the original GAN.

Building on the existing cross-domain learning GANs, SPID-GAN approaches the state-parameter bidirectional mapping problem by using two pairs of generators and discriminators:

$$\begin{aligned} \text{Forward mapping } \mathcal{G}_{PS} : \mathbb{R}_P &\rightarrow \mathbb{R}_S, \mathcal{D}_S : \mathbb{R}_S \rightarrow [0, 1], \\ \text{Reverse mapping } \mathcal{G}_{SP} : \mathbb{R}_S &\rightarrow \mathbb{R}_P, \mathcal{D}_P : \mathbb{R}_P \rightarrow [0, 1], \end{aligned} \quad (1)$$

where \mathcal{G}_{PS} defines a forward mapping from the parameter space P to the model state space S while \mathcal{G}_{SP} provides a reverse mapping from S to P . The two discriminators \mathcal{D}_S and \mathcal{D}_P are used to determine the authenticity of samples generated for the respective domains in terms of probability. A practical working assumption on \mathcal{G}_{PS} and \mathcal{G}_{SP} is that they are bijective, meaning each element of the domain P is mapped by exactly one element of the domain S . This helps prevent the many-to-one mappings during training, which is also known as the mode collapse problem (Zhu et al., 2017). Another assumption is the continuity of mappings; namely, if two elements in domain P are close, then also should be the corresponding elements in domain S . The same assumption is also implied by the stability requirement of inverse solutions (Sun & Sun, 2015). Thus, to arrive at meaningful solutions, one needs not only a proper algorithm design (detailed below), but also an appropriate experimental design (next section).

The loss function used for training \mathcal{G}_{PS} consists of three terms (Kim et al., 2017; Zhu et al., 2017):

$$\begin{aligned} J_{PS}^{\mathcal{G}}(P, S, \theta^{D_S}, \theta^{\mathcal{G}_{PS}}, \theta^{\mathcal{G}_{SP}}) &= J_{PS}^{D_S} + d_{\text{tran}}(\mathcal{G}_{PS}(\mathbf{x}, \theta^{\mathcal{G}_{PS}}), \mathbf{y}) \\ &+ d_{\text{cyc}}(\mathcal{G}_{SP}(\mathcal{G}_{PS}(\mathbf{x}, \theta^{\mathcal{G}_{PS}}), \theta^{\mathcal{G}_{SP}}), \mathbf{x}), \end{aligned} \quad (2)$$

where $\mathbf{x} \in P$, $\mathbf{y} \in S$, and $\theta^{(\cdot)}$ denote the unknown parameters of the respective generators/discriminators. Before training, a standard practice in DL is to scale all training data to the same range (e.g., [0, 1]) so the objective functions are addable. The first term on the right-hand side of equation (2) defines the discriminator loss in the sense of the original GAN:

$$\begin{aligned} J_{PS}^{D_S}(\theta^{D_S}, \theta^{\mathcal{G}_{PS}}) &= \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} \left\| \mathcal{D}_S(\mathcal{G}_{PS}(\mathbf{x}, \theta^{\mathcal{G}_{PS}}), \theta^{D_S}) \right\|_2^2 + \\ &\frac{1}{2} \mathbb{E}_{\mathbf{y} \sim p_{\text{data}}(\mathbf{y})} \left\| \mathcal{D}_S(\mathbf{y}, \theta^{D_S}) - 1 \right\|_2^2, \end{aligned} \quad (3)$$

for which the goal is to minimize the error rate of the discriminator on the fake sample (toward 0) and on the real sample (toward 1). In equation (3), the expectation (\mathbb{E}) is taken over all training samples. Note that the mean square error (MSE) is used here, instead of the binary entropy loss used by Goodfellow et al. (2014).

The second term on the right-hand side of equation (2) measures the translation loss in terms of mean absolute error:

$$d_{\text{tran}}(\mathcal{G}_{PS}(\mathbf{x}, \theta^{\mathcal{G}_{PS}}), \mathbf{y}) = \mathbb{E}_{\mathbf{y} \sim p_{\text{data}}(\mathbf{y})} \left| \mathcal{G}_{PS}(\mathbf{x}, \theta^{\mathcal{G}_{PS}}) - \mathbf{y} \right|, \quad (4)$$

where expectation is calculated based on all pairs of generated and training images. The last term on the right-hand side of equation (2) quantifies the cycle consistency (or reconstruction loss) between the two generators using mean absolute error:

$$d_{\text{cyc}}(G_{SP}(G_{PS}(\mathbf{x}, \theta^{G_{PS}}), \theta^{G_{SP}}), \mathbf{x}) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} |G_{SP}(G_{PS}(\mathbf{x}, \theta^{G_{PS}}), \theta^{G_{SP}}) - \mathbf{x}|. \quad (5)$$

By minimizing the reconstruction loss, the cycle consistency term helps to mitigate the mode collapse problem (Kim et al., 2017; Zhu et al., 2017). The loss function of the generator G_{SP} can be defined similarly to equation (2) by switching P and S . The total generator loss function J^G for SPID-GAN is the average of the two generator losses. Each discriminator uses a loss function in the same form as equation (3), but with the opposite sign. The generator and discriminator loss functions are highly coupled and need to be solved from the following minimax optimization problem:

$$\hat{\theta}^{G_{PS}}, \hat{\theta}^{D_S}, \hat{\theta}^{G_{SP}}, \hat{\theta}^{D_P} = \arg \min_{\theta^{G_{PS}}, \theta^{G_{SP}}} \max_{\theta^{D_S}, \theta^{D_P}} J^G(\theta^{G_{PS}}, \theta^{D_S}, \theta^{G_{SP}}, \theta^{D_P}). \quad (6)$$

In practice, the optimization problem in equation (6) is solved by using alternating-gradient updating steps for generators and discriminators, with parameters of one group fixed when parameters of the other are being updated in each iteration.

Figure 1a illustrates the workflow of SPID-GAN. The first step shown on the left is related to data preparation, which may be facilitated by a rich set of tools available from the geostatistics literature (Deutsch & Journel, 1998). For example, if parameter measurements are available, they may be integrated to generate the so-called *conditional realizations* of the parameter field to honor prior information. If measurements of state variables are available, they may be used to generate plausible images of the state field via kriging. Model state measurements may also be used to select the most probable parameter fields by choosing the parameter sets that minimize the differences between GAN-predicted state values and the actual observed values. Assuming point measurements of parameter and state variables are available, the SPID-GAN workflow implies an additional loss term on the generators that is enforced through preprocessing and postprocessing,

$$d_{\text{obs}} = \mathbb{E}_S \|G_{PS}(\mathbf{x})(\mathbf{u}_S) - \mathbf{y}(\mathbf{u}_S)\|_2^2 + \mathbb{E}_P \|G_{SP}(\mathbf{y})(\mathbf{u}_P) - \mathbf{x}(\mathbf{u}_P)\|_2^2, \quad (7)$$

where \mathbf{u}_P and \mathbf{u}_S are locations of P and S measurements, respectively. These different use cases will be illustrated in section 3.

In this work, SPID-GAN is implemented by using convolutional neural networks (CNNs), a class of deep feed-forward neural networks specially designed for image pattern recognition (LeCun et al., 2015). A brief introduction of common CNN terminologies is given in supporting information (SI) S1. The two generators share a DL neural network design that is similar to DiscoGAN (Kim et al., 2017), which includes a series of convolutional and deconvolutional layers to help uncover features at multiple scales (Figure 1b). The input images have dimensions 128×128 . The convolutional layers (the downsampling path) use 4×4 kernels and a uniform stride size of two, while the deconvolutional layers (the upsampling path) all use a stride size of one. The number of filters increases from 64 to 512 for the convolutional layer stack, while the deconvolutional layers reverse the number of filters to generate an output image of the same size as the input. The leaky rectified linear unit function is used as the activation function for all hidden CNN layers, and the hyperbolic tangent function (tanh) is used as the activation function for the output layer. Instance normalization is applied to hidden layers to improve the training speed (Ulyanov et al., 2016). The two discriminators use the same design as shown in Figure 1c. For the hidden layers, the layer configuration is repeated twice using alternating stride sizes 1 and 2. All codes are written in Python using the open-source DL package, Keras (Chollet, 2015). The Adam optimization solver (Ba, 2014) is used for training, with a learning rate 0.0002 and a forgetting factor 0.5. Unless otherwise noted, the number of epochs used in training is 125, and the batch size is 10. The computing time taken for each epoch is about 23 s on a cloud-based Ubuntu instance running on an Intel Xeon E5-2580 CPU node with GPU acceleration (NVIDIA Tesla K40).

To quantify the skill of trained generators, the structural similarity index (SSIM) commonly used in image analysis (Wang et al., 2004) is adopted as a metric. For two sliding windows \mathbf{u} and \mathbf{v} of dimensions $n_p \times n_p$, SSIM is defined as

$$\text{SSIM}(\mathbf{u}, \mathbf{v}) = \frac{2\mu_{\mathbf{u}}\mu_{\mathbf{v}} + c_1 2\sigma_{\mathbf{uv}} + c_2}{\mu_{\mathbf{u}}^2 + \mu_{\mathbf{v}}^2 + c_1 \sigma_{\mathbf{u}}^2 + \sigma_{\mathbf{v}}^2 + c_2}, \quad (8)$$

where \mathbf{u} and \mathbf{v} represent two patches from the simulated image (using GAN) and testing image (from numerical model), respectively, μ represents the mean, σ^2 represents the variance, and c_1 (0.01) and c_2 (0.03) are

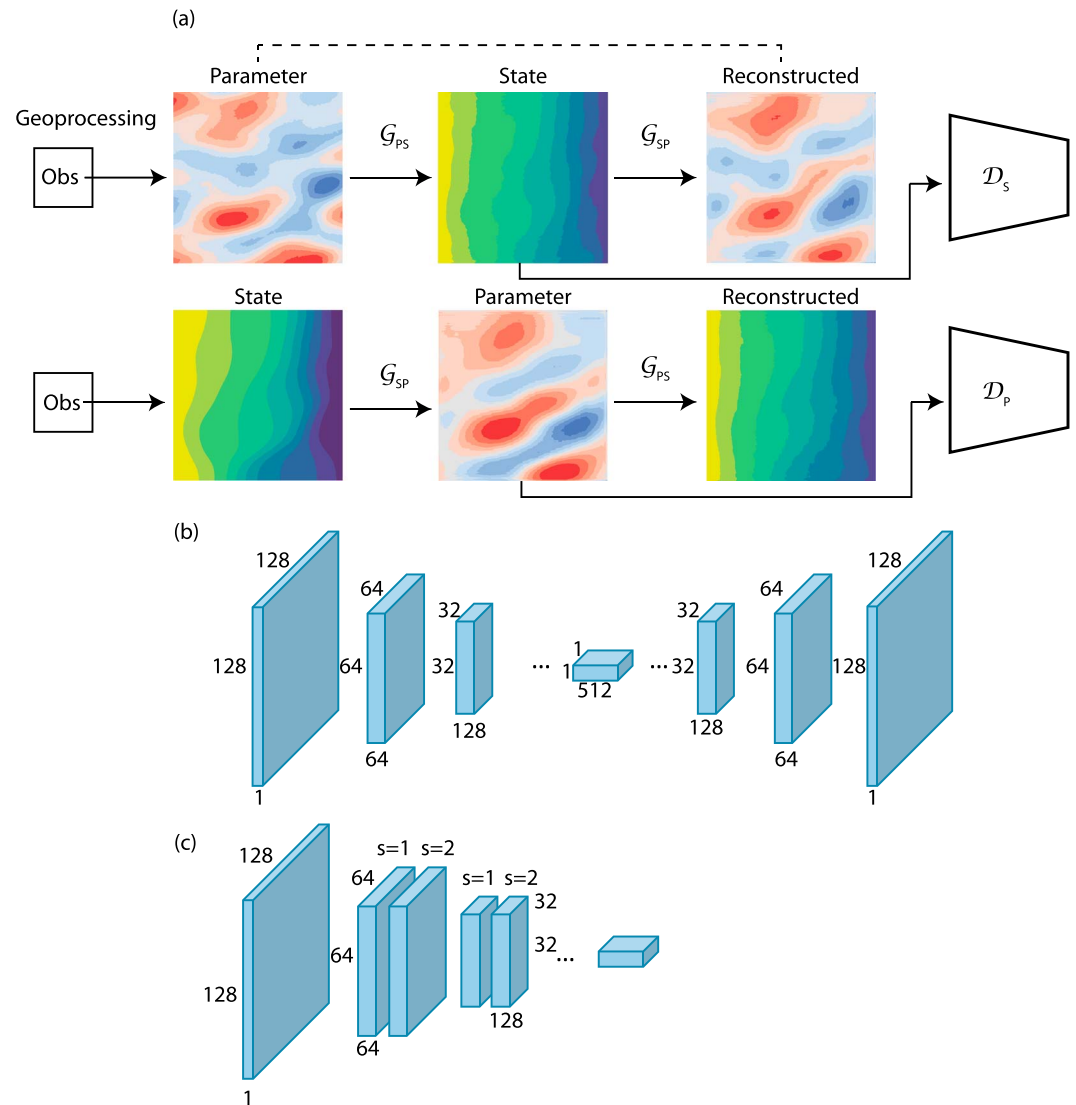


Figure 1. (a) State-parameter identification generative adversarial network (SPID-GAN) consists of two pairs of generators and discriminators trained together, one pair for identifying the forward parameter-to-state mapping (\mathcal{G}_{PS} , \mathcal{D}_S ; first row) and the other for identifying the reverse state-to-parameter mapping (\mathcal{G}_{SP} , \mathcal{D}_P ; second row); cycle consistency is enforced by minimizing the reconstruction loss (third column); observations of parameter and state variables can be fused through preprocessing and postprocessing steps (e.g., kriging and geostatistical simulation); (b) the design of generator model follows a downsampling-upsampling pattern using convolutional and deconvolutional layers to learn features at multiple scales; (c) the discriminator uses repeating convolutional layers to improve learning of fine-scale features.

small constants used to stabilize the denominator. The mean value of SSIM, averaged over all sliding windows, ranges from -1 to 1 , with 1 being attainable when two images are identical. The size of the sliding window used in this study is $n_p = 7$.

2.2. Groundwater Flow

To demonstrate the usefulness of SPID-GAN for learning the bidirectional mappings, I consider groundwater flow in a spatially heterogeneous aquifer, which is a representative geoscience problem and has been studied extensively. The governing equation is given by the following PDE:

$$S_s(\mathbf{z}) \frac{\partial h(\mathbf{z}, t)}{\partial t} = \nabla \cdot (K(\mathbf{z}) \nabla h(\mathbf{z}, t)) + q_w(\mathbf{z}, t), \quad (9)$$

where h (L) is hydraulic head, S_s (1/L) is specific storage, K (L/T) is hydraulic conductivity, \mathbf{z} denotes spatial coordinates, t is time, and q_w (L^3/T) is the source/sink term. For the purpose of illustration, the state variable is h , and the parameter is K ; both are spatially distributed variables. All other quantities are assumed deterministic.




Figure 2. State-parameter identification generative adversarial network (SPID-GAN) results for two random realizations (a–d and e–h): (a) and (e) are original $\log K$ realizations generated using `sgsim`, (c) and (d) are original head images simulated using MODFLOW, head fields (b) and (f) are generated using the trained forward generator G_{PS} , and parameter fields (d) and (h) are generated by the trained reverse generator G_{SP} . All contours are normalized for visualization purposes. Subplots (i) and (j) show histograms of structural similarity indices (SSIMs) calculated on 1,000 test samples not used during training. The mean SSIM of G_{PS} is 0.98, and the mean SSIM of G_{SP} is 0.78.

The lateral dimensions of the aquifer are $1,280\text{m} \times 1,280\text{m}$, and the domain is uniformly discretized into $10\text{-m} \times 10\text{-m}$ grid blocks. The thickness of the aquifer is 20 m. The goal of SPID-GAN is to train two generators simultaneously to approximate the physical model specified in equation (9). Two different problem settings are considered, single-modal and multimodal parameter distributions.

3. Results

3.1. Single-Modal Parameter Distribution

In the first problem, K is assumed to be a random field following log-normal distribution. The mean and standard deviation of $\log K$ are 2×10^{-4} and 1.0. The variogram model of $\log K$ is Gaussian with max and min ranges of 500 and 100 m, respectively. S_s is deterministic with a value of $2.5 \times 10^{-6} \text{m}^{-1}$. In the baseline case, constant-head (Dirichlet) boundary conditions are imposed on both the west (21 m) and east (10 m) sides of the aquifer. Four pumping wells are put in grid blocks (25, 25), (25,106), (106, 106), and (106,25), with pumping rates of 10, 10, 50, and $20 \text{m}^3/\text{day}$, respectively. Stochastic realizations of $\log K$ are generated by using the sequential Gaussian simulator (`sgsim`), available from the open-source geostatistical package SGeMS (Remy et al., 2009). The flow field is first run to the steady state, followed by a transient simulation period of 3,600 s. For training and validation, the corresponding head distributions are obtained by using the open-source groundwater flow solver MODFLOW via its Python wrapper, `f1opy` (Bakker et al., 2016). The computing time for each




Figure 3. Illustration of the use of prior information in state-parameter identification generative adversarial network (SPID-GAN): (a) and (b) show the *true* log K and head fields, which are sampled only at monitoring locations (open circles); (c), (e), and (g) show the top three log K realizations identified as the *closest* to the true head field, as measured using MSE between simulated and observed head values; and (d), (f), and (h) show the corresponding head fields, which resemble the true head field.

forward simulation is 0.04 s. Such geostatistical processing steps have been broadly used in previous studies, such as ensemble-based data assimilation (Chen & Zhang, 2006; Franssen et al., 2009; Sun et al., 2009a), hydraulic tomography (Lee & Kitanidis, 2014), and surrogate modeling and uncertainty quantification (Li & Zhang, 2007; Nowak et al., 2010). Training of the SPID-GAN is done using 400 pairs of log K (parameter domain) and h (state domain) fields, each having dimensions of 128×128 .

Figures 2a–2d and 2e–2h show examples resulting from two test realizations not included in training. The left two images of each row show the input log K field and the corresponding head field generated using the trained forward generator \mathcal{G}_{PS} , while the right two images of each row show the input state (head) and corresponding log K field predicted using the trained reverse generator \mathcal{G}_{SP} . For this base case, SPID-GAN captures the spatial patterns in head distributions well, on the basis of visual comparisons between Figures 2b and 2c and between Figures 2f and 2g. Results of testing using 1,000 test realizations of log K give an ensemble mean SSIM value of about 0.98 (Figure 2i). In comparison, by visual examinations of Figures 2a and 2d and Figures 2e and 2h, it can be seen that the trained reverse mapping \mathcal{G}_{SP} captures the dominant spatial patterns in the original log K images but tends to underestimate some local maxima or minima. The ensemble mean SSIM obtained from testing on 1,000 h field samples (not used in training) is 0.78.

In general, learning a high-dimensional reverse mapping is a significantly more challenging problem to solve, depending on not only the design of the GAN algorithm but also the experimental design and quality of training samples. In this case, the learning of reverse mapping would benefit from any conditions (e.g., boundary conditions and forcings) that can help improve the information content of head fields and the uniqueness of cross mappings. In an experimental design, if the head field is not sensitive to certain parts of the parameter field, SPID-GAN may give ambiguous inverse solutions in those areas. To elaborate this latter point, in SI Text S2 the number of pumping wells is reduced from four to two (Figure S2) and then to none (Figure S3). As a result, the parameter fields in those cases become less identifiable—More artifacts start to appear in reversely mapped log K fields in those cases, especially in the central part of the domain that is less stimulated than the parts near the west and east boundaries. Nevertheless, the overall SPID-GAN performance stays relatively robust, as can be seen by the SSIM statistics in the respective plots.

In addition to different boundary/forcing conditions, the effect of training sample quantity on SPID-GAN performance is also investigated. The results, shown in SI Text S3, suggest that the GAN is relatively robust when the number of training samples varies, indicating the capability of GAN to learn dominant cross-domain patterns.

GANs operate with images, while the traditional workflow in hydrogeology typically involves point measurements. This next example demonstrates that the two workflows are actually complementary under the new SPID-GAN workflow proposed in Figure 1a. Figures 3a and 3b show the *true* log K and head fields, which are sampled only at limited locations. For simplicity, it is assumed that point observations of K and h are collected from the same monitoring network (open circles in Figures 3a and 3b, a total of 36 conditioning points), although they may well be different. First, 1,000 log K conditional realizations are generated using `sgsim` that honor the prior information on K . For each conditional log K realization, the trained forward generator \mathcal{G}_{PS} from the base case is used to predict a head field, which is then sampled at monitoring locations to calculate the MSE between the simulated and actual head observations. The resulting MSE values are sorted in the ascending order. Figures 3c, 3e, and 3g show the top three head fields that have the minimal MSE, while the




Figure 4. Illustration of the use of state-parameter identification generative adversarial network to identify bidirectional mappings for a bimodal parameter distribution: Subplots (a)–(d) show results from a test realization; subplots (e)–(h) show results from another test realization. The structural similarity index (SSIM) histograms obtained using 1,000 test realizations are shown in subplots (i)–(j). The mean SSIM of the forward generator is 0.8, and the mean SSIM of the reverse generator is 0.71.

corresponding log K fields are plotted in Figures 3d, 3f, and 3h. The identified fields show strong resemblance to the synthetic truth, especially where conditional information is available (i.e., comparing SPID-GAN results to the synthetic truth at the monitoring locations). Similarly, the workflow presented here can be used to form an ensemble of models to perform the DL-based, uncertainty quantification in the sense of the recently introduced data space inversion, in which the goal of modeling is not calibration in the traditional sense but to establish a data-driven statistical relationship between the observed and forecast variables and to quantify the predictive uncertainty of the forecast variables, by using an ensemble of uncalibrated prior models (Jeong et al., 2018; Satija & Caers, 2015; Sun et al., 2017).

3.2. Bimodal Parameter Distribution

In the second problem setting, the feasibility of using SPID-GAN for learning multimodal distributions is investigated. The aquifer is assumed to consist of two hydrofacies, a permeable channel facies and a background matrix, making the distribution of K bimodal. Identification of facies shapes is a representative and yet challenging inversion problem that has also been studied extensively in the literature (Liu & Oliver, 2005; Sun et al., 2009b; Zhou et al., 2011). In this example, the channel facies has a K value of 2×10^{-4} m/s and an S_s value of 1×10^{-6} m $^{-1}$, while the matrix has a K value of 1×10^{-8} m/s and an S_s value of 1×10^{-7} m $^{-1}$. The facies

realizations are generated using `snesim`, which is a multipoint geostatistical simulator also available from SGeMS (Remy et al., 2009). Constant heads of 11 and 10 m are imposed on the west and east boundaries, and the transient simulation period is 5,400 s. All other settings are the same as used in the previous two examples.

Figures 4a–4d and 4e–4h show example results from two test realizations. The forward generator predicts the connected flow paths well, including the fading color pattern in head distribution from the west boundary to the east boundary (comparing Figures 4b and 4c and Figures 4f and 4g). The reverse mapping from the head distribution identifies the well-connected backbone of the flow network (comparing Figure 4a with 4d and Figure 4e with 4h) but tends to omit small channel segments that are not connected to the main flow pathways. This is because the information of disconnected segments is not discoverable from head distributions. In other words, the head measurements are not sensitive to the parameter information in disconnected channel segments. From a practical standpoint, the connected features are more important to identify for risk assessment purposes. Finally, Figures 4i and 4j show the SSIM statistics on running 1,000 pairs of randomly generated test samples. Again, SPID-GAN gives a reasonable performance in this bimodal distribution case, with an ensemble mean SSIM value of 0.8 for forward mapping (K to h) and 0.71 for reverse mapping (h to K).

4. Summary and Conclusions

The recently introduced GANs have shown strong performance in image-to-image translation. Questions remain about whether GANs can be used to learn the deep causal structures embedded in physical models. In this study, I show that the usefulness of GANs goes beyond the simple picture mapping and demonstrate the feasibility of using a SPID-GAN to approximate bidirectional mappings between the parameter space and state space of PDEs. Through a series of representative problems selected from groundwater modeling, I show that SPID-GAN does a satisfactory job (with a reasonable amount of computing time) in linking physical parameters to model states and is able to capture the complex spatial patterns in parameter and state distributions that are otherwise challenging to obtain. The need for identifying forward and reverse mappings is ubiquitous in geophysical research. Thus, findings of this study may have important implications for many similar tasks, such as geostatistical simulation, history matching, uncertainty quantification, surrogate modeling, optimal design of experiments, or any problem that requires sampling from a high-dimensional distribution in a computationally efficient manner.

Acknowledgments

The work was partly supported by the U.S. Department of Energy, National Energy Technology Laboratory (NETL) under Grants DE-FE0026515 and DE-FE0031544. The author is grateful to the handling Editor, Valeriy Ivanov, and two anonymous reviewers for their constructive comments. All the data used are listed in the references or archived in repository, <https://utexas.box.com/s/uq74paot1a2ns3v9hbl101ifev408uus>.

References

- Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Bakker, M., Post, V., Langevin, C. D., Hughes, J. D., White, J., Starn, J., & Fioren, M. N. (2016). Scripting MODFLOW model development using Python and FloPy. *Groundwater*, 54(5), 733–739.
- Castelluccio, M., Poggi, G., Sansone, C., & Verdoliva, L. (2015). Land use classification in remote sensing images by convolutional neural networks. arXiv preprint arXiv:1508.00092.
- Chapelle, O., Schölkopf, B., & Zien, A. (2006). *Semi-supervised learning*. Cambridge, MA: MIT Press.
- Chen, Y., & Zhang, D. (2006). Data assimilation for transient flow in geologic formations via ensemble Kalman filter. *Advances in Water Resources*, 29(8), 1107–1122.
- Chollet, F. (2015). Keras. <https://keras.io>.
- Deutsch, C. V., & Journé, A. G. (1998). *Geostatistical software library and user's guide*. New York: Oxford University Press.
- Fang, K., Shen, C., Kifer, D., & Yang, X. (2017). Prolongation of SMAP to spatiotemporally seamless coverage of continental US using a deep learning neural network. *Geophysical Research Letters*, 44, 11,030–11,039. <https://doi.org/10.1002/2017GL075619>
- Franssen, H. H., Alcolea, A., Riva, M., Bakr, M., Van der Wiel, N., Stauffer, F., & Guadagnini, A. (2009). A comparison of seven methods for the inverse modelling of groundwater flow. Application to the characterisation of well catchments. *Advances in Water Resources*, 32(6), 851–872.
- Goodfellow, I. (2016). NIPS 2016 tutorial: Generative adversarial networks, arXiv preprint arXiv:1701.00160.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672–2680). Montréal, Quebec, Canada.
- Jeong, H., Sun, A. Y., Lee, J., & Min, B. (2018). A learning-based data-driven forecast approach for predicting future reservoir performance. *Advances in Water Resources*, 118, 95–109.
- Kim, T., Cha, M., Kim, H., Lee, J., & Kim, J. (2017). Learning to discover cross-domain relations with generative adversarial networks. arXiv preprint arXiv:1703.05192.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., et al. (2017). Photo-realistic single image super-resolution using a generative adversarial network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 4681–4690). Honolulu, HI.
- Lee, J., & Kitanidis, P. K. (2014). Large-scale hydraulic tomography and joint inversion of head and tracer data using the principal component geostatistical approach (PCGA). *Water Resources Research*, 50, 5410–5427. <https://doi.org/10.1002/2014WR015483>
- Li, T., Shen, H., Yuan, Q., Zhang, X., & Zhang, L. (2017). Estimating ground-level PM_{2.5} by fusing satellite and station observations: A geo-intelligent deep learning approach. *Geophysical Research Letters*, 44, 11,985–11,993. <https://doi.org/10.1002/2017GL075710>
- Li, H., & Zhang, D. (2007). Probabilistic collocation method for flow in porous media: Comparisons with other stochastic methods. *Water Resources Research*, 43, W09409. <https://doi.org/10.1029/2006WR005673>

- Liu, N., & Oliver, D. S. (2005). Ensemble Kalman filter for automatic history matching of geologic facies. *Journal of Petroleum Science and Engineering*, 47(3-4), 147–161.
- Liu, Y., Racah, E., Correa, J., Khosrowshahi, A., Lavers, D., Kunkel, K., et al. (2016). Application of deep convolutional neural networks for detecting extreme weather in climate datasets. arXiv preprint arXiv:1605.01156.
- Liu, M.-Y., & Tuzel, O. (2016). Coupled generative adversarial networks. In *Advances in Neural Information Processing Systems* (pp. 469–477). Barcelona, Spain.
- Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets, arXiv preprint arXiv:1411.1784.
- Nowak, W., De Barros, F., & Rubin, Y. (2010). Bayesian geostatistical design: Task-driven optimal site investigation when the geostatistical model is uncertain. *Water Resources Research*, 46, W03535. <https://doi.org/10.1029/2009WR008312>
- Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks, arXiv preprint arXiv:1511.06434.
- Remy, N., Boucher, A., & Wu, J. (2009). *Applied geostatistics with SGeMS: A user's guide*. Cambridge: Cambridge University Press.
- Satija, A., & Caers, J. (2015). Direct forecasting of subsurface flow response from non-linear dynamic data by linear least-squares in canonical functional principal component space. *Advances in Water Resources*, 77, 69–81.
- Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., & Woo, W.-c. (2015). Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems* (pp. 802–810).
- Sun, W., Hui, M.-H., & Durlafsky, L. J. (2017). Production forecasting and uncertainty quantification for naturally fractured reservoirs using a new data-space inversion procedure. *Computational Geosciences*, 21(5-6), 1443–1458.
- Sun, A. Y., Morris, A., & Mohanty, S. (2009a). Comparison of deterministic ensemble Kalman filters for assimilating hydrogeological data. *Advances in Water Resources*, 32(2), 280–292.
- Sun, A. Y., Morris, A. P., & Mohanty, S. (2009b). Sequential updating of multimodal hydrogeologic parameter fields using localization and clustering techniques. *Water Resources Research*, 45, W07424. <https://doi.org/10.1029/2008WR007443>
- Sun, N.-Z., & Sun, A. (2015). *Model calibration and parameter estimation: For environmental and water resource systems*. New York, NY: Springer.
- Ulyanov, D., Vedaldi, A., & Lempitsky, V. (2016). Instance normalization: The missing ingredient for fast stylization. arxiv 2016, arXiv preprint arXiv:1607.08022.
- Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4), 600–612.
- Yi, Z., Zhang, H., Tan, P., & Gong, M. (2017). Dualgan: Unsupervised dual learning for image-to-image translation. In *IEEE International Conference on Computer Vision (ICCV)* (pp. 2849–2857). Venice, Italy.
- Zhou, H., Gomez-Hernandez, J. J., Franssen, H.-J. H., & Li, L. (2011). An approach to handling non-Gaussianity of parameters and state variables in ensemble Kalman filtering. *Advances in Water Resources*, 34(7), 844–864.
- Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)* (pp. 2223–2232). Venice, Italy.