Consensus equilibrium for subsurface delineation

Hyung Jun Yang[1], **Youzuo Lin**[2], **Brendt Wohlberg** [2], **and Daniel M. Tartakovsky**[1]

[1]Department of Energy Resources Engineering, Stanford University, Stanford, CA, USA
[2]Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM, USA

Key Points:

- We develop a consensus equilibrium framework to reconstruct the subsurface from sparse data.
- Our framework integrates conventional optimization methods with machine learning-based methods.
- The method accurately reconstructs geologically complex models and efficiently reduces and quantifies the posterior uncertainty.

Corresponding author: D. M. Tartakovsky, tartakovsky@stanford.edu

Abstract
Heterogeneity and insufficient site characterization limit our knowledge of the subsurface. Inversion techniques, which minimize the mismatch between observations and model predictions, have become an essential tool of subsurface characterization. Most optimization-based approaches fail to incorporate various implicit priors and capture the geological complexity. We overcome these limitations by deploying the plug and play and consensus equilibrium (CE) strategies, which provide a flexible framework for image reconstruction. Our CE methodology for spatial delineation of geologic formations consists of an image denoiser and a variational auto-encoder (deep learning-based emulator). The former ameliorates the reconstruction noise, yielding well-defined geological structures; its mathematical equivalence with the proximal operator allows the deployment of advanced denoisers (e.g., CNN-based denoiser) that do not correspond to a regularization objective. The latter defines a geology prior that imposes a geological constraint, e.g., continuity and shape of geological features, onto the reconstructed image. We conduct a series of numerical experiments dealing with transient two-dimensional flow driven by a pumping well and natural hydraulic head gradient. They demonstrate the CE framework's ability to delineate, both probabilistically and deterministically, complex subsurface environments with sufficient quality.

1 Introduction

Reliable characterization of the subsurface is a key component of quantitative predictions of flow and transport in geologic formations. Subsurface characterization generally entails inversion (aka history matching or data assimilation), a computational procedure that converts observations (e.g., of hydraulic head or solute concentration) into multi-dimensional images of model parameters (e.g., hydraulic conductivity or dispersivity). While inverse strategies vary widely, most of them involve minimization of a discrepancy between observations and model predictions.

Gradient-based methods (Sarma et al., 2006, and the references therein) and evolutionary algorithms (Cameron et al., 2016, and the references therein) are some of the most successful approaches to optimization-based inversion. Among these, adjoint gradient-based methods boast high efficiency, because they require only one forward and one backward simulation to compute model sensitivity (Sarma et al., 2006, and the references therein). These and similar optimization-based procedures generate one optimal solution at a time; they can be combined with the randomized maximum-likelihood method (Kitanidis, 1995) to obtain multiple posterior models for uncertainty quantification. Sampling-based approaches for data assimilation, e.g., (dual) ensemble Kalman filter (Wang et al., 2021), ensemble smoother (Chen & Oliver, 2012), Markov chain Monte Carlo (Zhou & Tartakovsky, 2021), and DREAM algorithm (Laloy et al., 2018) combined with Markov chain Monte Carlo, allow one both to estimate model parameters and to assess predictive uncertainty. Even though iterative ensemble-based schemes (Chen & Oliver, 2012; White, 2018) ameliorate the convergence issue that plagues most inversion algorithms, sampling-based approaches are generally more computationally expensive than adjoint gradient method. Yet, they are often more efficient than gradient-based approaches without the adjoints.

Regardless of the strategy used to achieve it, subsurface inversion is generally an ill-posed problem that has to be regularized. One way to do so is to incorporate prior information in the form of geological constraints, which would guarantee a geologically realistic solution. A majority of the current approaches rely on explicitly definable priors such as two-point statistics (mean and covariance), which renders them inappropriate for complex geology. One alternative is to employ a parameterization procedure aiming to represent geological maps in terms of a small number of parameters. Methods of this class include principal component analysis (Sarma et al., 2008;

Vo & Durlofsky, 2014) and deep learning-based techniques (S. Chan & Elsheikh, 2019; Laloy et al., 2017, 2018; Liu et al., 2019). Though these approaches have shown good performance in many subsurface applications, their results vary with the subjectively defined number of parameters and parameterization methods, i.e., require significant fine-tuning.

We posit that high (parametric) dimensionality of subsurface images and non-differentiability of the image reconstruction problem argue for the deployment of proximal algorithms such as alternating direction method of multipliers (ADMM) (Eckstein & Bertsekas, 1992; Afonso et al., 2010; Boyd et al., 2011). More specifically, we adopt plug-and-play (PnP) priors (Venkatakrishnan et al., 2013), which extend the previous proximal algorithms by relying on recent progress in machine learning in general and deep neural networks (DNN) in particular. Inspired by the mathematical equivalence between the proximal operator and an image denoiser, PnP methodologies provide a flexibility to integrate various heterogeneous priors that may not be explicitly defined (Sreehari et al., 2016; S. H. Chan et al., 2016; Ono, 2017; Kamilov et al., 2017; Sun et al., 2019). Among them, consensus equilibrium (CE) (Buzzard et al., 2018), a generalization of the ADMM-based PnP scheme, integrates multiple advanced prior models (e.g., denoisers, data fidelity agents, deblurring maps, etc.) within an optimization-free framework; examples of its application to image reconstruction problems can be found in (Ghani & Karl, 2019; Sridhar et al., 2020).

We introduce CE as a means to achieve geologically realistic results for subsurface inversion problems, which are formulated in Section 2. The CE approach, presented in Section 3, utilizes a denoising prior and deep learning-based prior to maintain geological realism in the inversion of hydraulic head data. The "agents" that seek equilibrium in our setting, i.e., operators responsible for data fidelity, for data denoising, and for geological priors, are introduced in Section 4. In Section 5, we demonstrate the performance of the proposed CE framework via a series of numerical experiments dealing with two-dimensional transient flow in an aquifer having channelized spatial arrangement. Main findings and conclusions drawn from our study are summarized in Section 6.

2 Problem Formulation

The goal of an inverse problem is to recover an unknown subsurface model parameter set $\mathbf{m} \in \mathbb{R}^n$ (e.g., values of hydraulic conductivity K in n elements of a numerical grid) from a set of measurements $\mathbf{d} \in \mathbb{R}^{m \cdot k}$ (e.g., of hydraulic head observed in m wells at k time intervals); in a typical application, $m \ll n$ and $m \cdot k \ll n$. In terms of the maximum-a-posteriori (MAP) estimate, this problem takes the form of an optimization problem,

$$\mathbf{m}^* = \underset{\mathbf{m} \in \mathbb{R}^n}{\operatorname{argmin}} \, f(\mathbf{m}). \tag{1}$$

The MAP cost function $f(\cdot) : \mathbb{R}^n \to \mathbb{R}^+$ is defined by

$$f(\mathbf{m}) = -\ln p(\mathbf{d}|\mathbf{m}) - \ln p(\mathbf{m}) + \text{const}, \tag{2}$$

where $p(\mathbf{d}|\mathbf{m})$ and $p(\mathbf{m})$ represent a data-fidelity function and a prior distribution, respectively. It is common to assume that the distribution of random observation errors is a multivariate normal, in which case the data-fidelity function, $-\ln p(\mathbf{d}|\mathbf{m})$, is expressed as

$$p(\mathbf{d}|\mathbf{m}) = \frac{1}{2}(\mathbf{g}(\mathbf{m}) - \mathbf{d})^\top \mathbf{C}_D^{-1}(\mathbf{g}(\mathbf{m}) - \mathbf{d}), \tag{3}$$

where $\mathbf{C}_D \in \mathbb{R}^{(m \cdot k) \times (m \cdot k)}$ is the covariance matrix of the measurement errors, and $\mathbf{g}(\cdot) : \mathbb{R}^n \to \mathbb{R}^{m \cdot k}$ represents the model predictions of an observable at space-time points $(\mathbf{x}, t)$ at which the observable's measurements $\mathbf{d}$ are available.

The prior distribution $p(\mathbf{m})$ is often selected to be standard, e.g., multivariate Gaussian. However, complex priors representative of realistic geological environments are often poorly described by such explicit distributions, and alternatives are needed. We propose the joint use of multiple advanced priors (e.g., machine learning-based priors) and conventional optimization-based approaches within the CE framework to generate geologically realistic models.

3 CE Framework

The novelty of CE framework is to fuse sets of heterogeneous models that may or may not arise from the regularized optimization. The plug-and-play (PnP) reconstruction (Venkatakrishnan et al., 2013) emerging from ADMM is the first method to incorporate the denoising operators that have no underlying optimization problem. In addition to providing an optimization-free interpretation of PnP, CE extends PnP to handle problems involving more than two "agents", allowing the use of more than one implicit regularizer for same problem.

To formulate the CE equations, we split the MAP cost function $f(\mathbf{m})$ in (2) into N auxiliary functions $f_i(\mathbf{m}) : \mathbb{R}^n \to \mathbb{R}$ $(i = 1, \ldots, N)$, such that (1) becomes

$$\min \sum_{i=1}^{N} f_i(\mathbf{m}_i); \qquad \text{subject to } \mathbf{m}_i = \mathbf{m}, \quad i = 1, \ldots, N, \tag{4}$$

with variable $\mathbf{m}_i \in \mathbb{R}^n$. A proximal mapping operator $F_i : \mathbb{R}^n \to \mathbb{R}^n$, corresponding to the cost function f_i, is defined by

$$F_i(\mathbf{m}) = \underset{\mathbf{v} \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ \frac{\|\mathbf{v} - \mathbf{m}\|^2}{2\sigma_i^2} + f_i(\mathbf{v}) \right\}. \tag{5}$$

The regularization parameters σ_i controls the convergence speed of the PnP algorithm. If f_i is a lower-semicontinuous and convex function on $\mathbb{R}^n$, then a solution of (4) and, hence, of (1) is given by a solution, $\mathbf{m}^*$, of the CE equations (Buzzard et al., 2018),

$$F_i(\mathbf{m}^* + \mathbf{u}_i^*) = \mathbf{m}^*, \quad i = 1, \ldots, N; \quad \text{and} \sum_{i=1}^{N} \mathbf{u}_i^* = \mathbf{0}. \tag{6}$$

The main advantage of the CE framework over optimization-based methods is that other non-expansive operators, called agents, can be employed in lieu of the proximal mapping operator F_i (Buzzard et al., 2018; Bouman, 2013). An operator $T : \mathbb{R}^n \to \mathbb{R}^n$ is said to be non-expansive if there exists a real number $0 \le k \le 1$ such that

$$\|T(\mathbf{x}) - T(\mathbf{y})\| \le k\|\mathbf{x} - \mathbf{y}\|, \tag{7}$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. The non-expansiveness, a weaker condition than being a proximal operator, allows us to use much richer class of actions including machine learning models.

The CE equations (6) can be solved with several proximal point algorithms, such as ADMM (Eckstein & Bertsekas, 1992) and Douglas-Rachford (Boyd et al., 2011). We introduce stacked operators $\mathbf{F}$ and $\mathbf{G}$ such that

$$\mathbf{F}(\mathbf{m}) = \begin{bmatrix} F_1(\mathbf{m}_1) \\ \vdots \\ F_N(\mathbf{m}_N) \end{bmatrix}, \qquad \mathbf{G}(\mathbf{m}) = \begin{bmatrix} \bar{\mathbf{m}} \\ \vdots \\ \bar{\mathbf{m}} \end{bmatrix}, \qquad \bar{\mathbf{m}} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{m}_i. \tag{8}$$

Then (6) takes the form

$$\mathbf{F}(\mathbf{v}^*) = \mathbf{G}(\mathbf{v}^*), \qquad \mathbf{v}^* = \mathbf{m}^* + \mathbf{u}^*. \tag{9}$$

The averaging function $\mathbf{G}$ has a property $2\mathbf{G} - \mathbf{I} = \mathbf{I}$, where $\mathbf{I} \in \mathbb{R}^{N \times N}$ is the identity matrix. This property gives rise to a fixed-point problem

$$(\mathbf{2G} - \mathbf{I})(\mathbf{2F} - \mathbf{I})\mathbf{v}^* = \mathbf{v}^*. \tag{10}$$

We solve it using the Mann iteration (Ryu & Boyd, 2016), which is equivalent to the Douglas-Rachford algorithm (Giselsson, 2017),

$$\mathbf{v}^{\mathbf{k+1}} = (1-\rho)\mathbf{v}^{\mathbf{k}} + \rho(\mathbf{2G} - \mathbf{I})(\mathbf{2F} - \mathbf{I})\mathbf{v}^{\mathbf{k}}, \tag{11}$$

where $\rho \in (0,1)$ is the fixed parameter. The convergence of the Mann iteration is guaranteed when the operator $\mathbf{T} \equiv (2\mathbf{G} - \mathbf{I})(2\mathbf{F} - \mathbf{I})$ is non-expansive (Bouman, 2013). If $\rho = 0.5$ and $N = 2$, then the fixed point approach (11) is identical to the PnP with ADMM algorithm (Sreehari et al., 2016).

4 Selection of Agents for Subsurface Delineation

Our inverse modeling strategy is to deploy the CE framework (9) with $N = 3$ agents F_i: a data fidelity agent F_{dat}, a denoising prior agent F_{den}, and a geology prior agent F_{geo}. The first of these, F_{dat}, is introduced to reduce the mismatch between observations, $\mathbf{d}$, and predictions of the reconstructed model, $\mathbf{g}(\mathbf{m})$. It is defined as a proximal mapping (5) of the data fidelity function (3),

$$F_{\text{dat}}(\mathbf{m}) = \underset{\mathbf{v} \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ \frac{\|\mathbf{v} - \mathbf{m}\|^2}{2\sigma_{\text{dat}}^2} + \frac{1}{2}(\mathbf{g}(\mathbf{v}) - \mathbf{d})^\top \mathbf{C}_D^{-1}(\mathbf{g}(\mathbf{v}) - \mathbf{d}) \right\}, \tag{12}$$

where σ_{dat} is an internal parameter controlling the strength of the regularization term $\|\mathbf{v} - \mathbf{m}\|^2$. We solve the minimization problem in (12) using the gradient-based method L-BFGS-B, in which adjoints are used to compute the gradient (Oliver et al., 2008). The adjoint method is employed here because of its computational efficiency and lack of the Gaussianity assumption, but other advanced inversion techniques (Chen & Oliver, 2012; White, 2018; Laloy et al., 2018; Zhou & Tartakovsky, 2021) can be plugged into CE as F_{dat}.

Agent F_{den} represents an image denoising prior, which plays the role of an implicit regularizer; its role is to provide reconstructions with sharper, less smeared hydrofacies boundaries. The underlying assumption behind the use of the denoising prior is that the subsurface is composed of distinct hydrofacies with relatively low heterogeneity within each facies (Winter & Tartakovsky, 2002; Winter et al., 2003, 2006; Yang et al., 2020). Examples of image denoisers include total variation (TV) (Osher et al., 2005), BM3D (Dabov et al., 2007), and DnCNN (Zhang et al., 2017). BM3D typically provides substantially better denoising performance than TV, and the more recent DnCNN typically outperforms BM3D by a much smaller margin. TV is posed as an optimization problem and is easier to implement as a regularizer for more complex problems than denoising via standard methods (Barajas-Solano et al., 2015). Neither BM3D nor DnCNN, however, has an explicit form as an optimization problem, BM3D being based on a complex algorithm involving block matching and coefficient shrinkage in the transform domain, and DnCNN being a convolutional neural network (CNN).

Agent F_{geo} enables our model to preserve the prior geological information such as shapes, sizes, positions and orientations of geological objects. We use variational autoencoders (VAEs), one of the popular generative models, to build such geological information into our model. In a typical implementation, VAEs use deep neural networks to learn latent representations from complex input data (Kingma & Welling, 2013). In so doing, an encoder is used to estimate latent variables $\mathbf{z}$ from input data $\mathbf{m}$, then multiple realizations of $\mathbf{z}$ are generated and used by a decoder to generate reconstructed data $\tilde{\mathbf{m}}$ (Fig. 1). Each layer of encoder/decoder contains a convolution

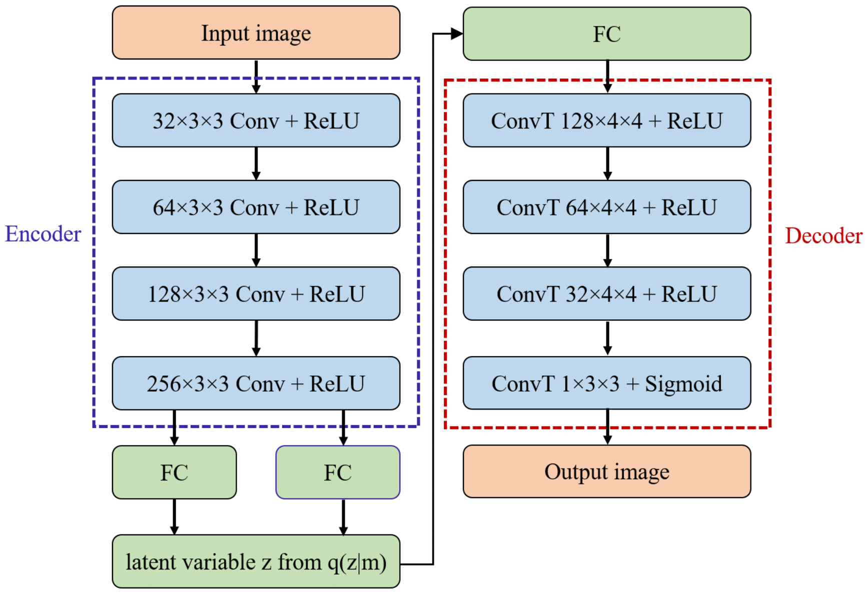

Figure 1: Overall architecture of our VAE for incorporating prior geological information. The abbreviations in Figure are defined as; Conv: convolution layer, FC: fully connected layer, and ConvT: transposed convolution layer.

or transposed-convolution unit to effectively extract and integrate the spatial features of input realizations. The convolution layer performs linear filtering on the output from the previous layer. When the input X is a 2D image, the output feature map, h, is obtained by $N \times N$ filter $\mathbf{w}$. One output pixel, $h_{i,j}$ is calculated as

$$h_{i,j}(Xi,j) = f\left(\sum_{m=1}^{N}\sum_{n=1}^{N} w_{i,j}X_{i+m,j+n}\right), \tag{13}$$

where the activation function f is a generally rectified linear unit (ReLU) $f(x) = \max(0,x)$. The transposed convolutional layer reverses the operation of a standard convolutional layer. Details of the convolution and transposed convolution layers can be found in (Dumoulin & Visin, 2016; Krizhevsky et al., 2012). Training data for VAE can be generated with geostatistical algorithms capable of producing geologically plausible realizations, including object-based (Deutsch & Tran, 2002), process-based methods (Paola, 2000), and multipoint geostatistics (Strebelle, 2002).

The encoder and decoder are trained simultaneously by minimizing VAE loss L_{VAE}. Let $\boldsymbol{\phi}$ denote a parameterization of the encoder that infers latent variables $\mathbf{z}$ from $\mathbf{m}$; the inferred distribution is $q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{m})$. When the decoder is parameterized with another parameter set $\boldsymbol{\theta}$, yielding a distribution $p_{\boldsymbol{\theta}}(\mathbf{z})$, VAE loss L_{VAE} is formulated as

$$L_{\text{VAE}}(\mathbf{m};\boldsymbol{\phi},\boldsymbol{\theta}) = D_{\text{KL}}[q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{m}), p_{\boldsymbol{\theta}}(\mathbf{z})] - \mathbb{E}_{\mathbf{z}\sim q_{\boldsymbol{\phi}}}[\ln(p_{\boldsymbol{\theta}}(\mathbf{m}|\mathbf{z}))]. \tag{14}$$

Here $p_{\boldsymbol{\theta}}(\mathbf{m}|\mathbf{z})$ is the likelihood of $\mathbf{m}$ given $\mathbf{z}$ under the decoder model. The term $D_{\text{KL}}[\cdot]$ is the Kullback-Leibler (KL) divergence, which provides a measure of discrepancy between the distribution of the latent variables, $p_{\boldsymbol{\theta}}(\mathbf{z})$, and the inferred distribution, $q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{m})$. The term $\mathbb{E}_{\mathbf{z}\sim q_{\boldsymbol{\phi}}}[\cdot]$, the expectation of the likelihood that the input image $\mathbf{m}$ can be generated from latent variable $\mathbf{z}$, represents the reconstruction error between

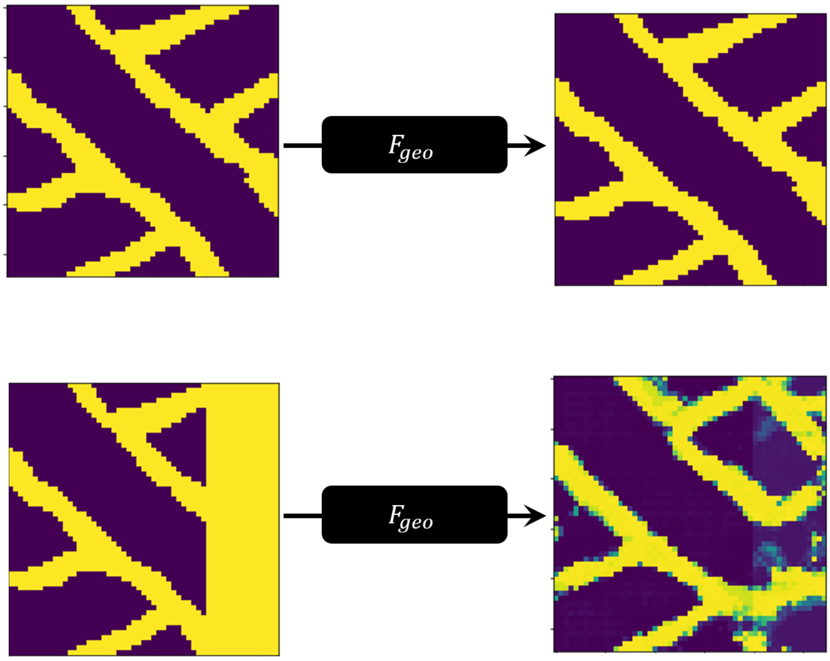

Figure 2: Examples of F_{geo} operations applied to geologically realistic input (top row) and unrealistic input (bottom row).

the actual $\mathbf{x}$ and the reconstructed image $\tilde{\mathbf{m}}$ from the decoder. To compute L_{VAE} efficiently, one represents $q_{\boldsymbol{\phi}}(\cdot)$ as a known readily parametrizable distribution. When $\mathbf{z}$ is a continuous latent variable, $q_{\boldsymbol{\phi}}(\cdot)$ is generally assumed to have a multivariate Gaussian distribution, whose mean μ and variance σ^2 are determined by the encoder model. If the VAE model is well trained on a sufficient number of realizations, its loss effectively measures the similarity between the input data and the training dataset. To ensure the consistency between an updated image and the prior geological information, we define agent F_{geo} as the proximal map of the VAE loss L_{VAE},

$$F_{\text{geo}}(\mathbf{m}) = \underset{\mathbf{v}\in\mathbb{R}^n}{\operatorname{argmin}} \left\{ \frac{\|\mathbf{v}-\mathbf{m}\|^2}{2\sigma_{\text{geo}}^2} + L_{\text{VAE}}(\mathbf{v};\cdot) \right\}, \tag{15}$$

where σ_{geo} is the regularization coefficient. Figure 2 shows examples of F_{geo}-agent operations. When the input is geologically realistic, F_{geo} preserves the original input. In contrast, F_{geo} improves the input significantly when the input is inconsistent with the prior geology.

5 Numerical Experiments

We consider two-dimensional (vertically averaged) transient flow in an aquifer $\Omega \subset \mathbb{R}^2$ bounded by a surface $\partial\Omega$. Spatiotemporal distribution of hydraulic head $h(\mathbf{x}, t)$ is described by the groundwater flow equation,

$$S_s \frac{\partial h}{\partial t} = \nabla \cdot (K \nabla h) - r(\mathbf{x}), \qquad \mathbf{x} \in \Omega, \quad t > 0, \tag{16}$$

where $K(\mathbf{x})$ and $S_s(\mathbf{x})$ are the aquifer's hydraulic conductivity and specific storage, respectively; and $r(\mathbf{x}, t)$ represents sources and sinks (e.g., wells and recharge). This equation is subject to initial and boundary conditions

$$h(\mathbf{x}, 0) = H, \quad \mathbf{x} \in \Omega; \qquad h = \phi, \quad \mathbf{x} \in \Gamma_D; \qquad -K\nabla h \cdot \mathbf{n} = \psi, \quad \mathbf{x} \in \Gamma_N. \tag{17}$$

Here, $H(\mathbf{x})$ is the initial distribution of hydraulic head; $\phi(\mathbf{x}, t)$ and $\psi(\mathbf{x}, t)$ are the hydraulic head and the normal component of the Darcy flux $\mathbf{q} = -K\nabla h$ prescribed, respectively, on the Dirichlet (Γ_D) and Neumann (Γ_N) segments of the boundary $\partial\Omega = \Gamma_D \cup \Gamma_N$; and $\mathbf{n}(\mathbf{x})$ is the outward unit normal vector to Γ_N.

The groundwater flow model (16) and (17) is supplemented with (noisy) measurements, $\mathbf{d}$, of hydraulic head h collected at a few locations (e.g., wells) throughout the aquifer during a certain time horizon. In a typical application, an aquifer's properties (K and S_s), auxiliary functions (H, ϕ, and ψ) and sources (r) are all uncertain and have to be inferred from geologic considerations and measurements $\mathbf{d}$. This inverse problem is variously referred to as model calibration or history matching. In the examples reported below, we treat hydraulic conductivity $K(\mathbf{x})$ as the only unknown parameter, and relate its natural logarithm, $Y(\mathbf{x}) = \ln K(\mathbf{x})$, to the corresponding values of specific storage $S_s(\mathbf{x})$ via a linear regression,

$$S_s = aY + b. \tag{18}$$

Following Li et al. (2004), we assume the regression coefficient a to be positive. With this simplification, and under a suitable discretization of the flow domain Ω into n elements (or nodes), the flow problem (16) and (17) is uniquely characterized by a set of parameters $\mathbf{m} = \{K_1, \ldots, K_n\}$. A numerical solution of this problem is denoted by $\mathbf{h} = \mathbf{g}(\mathbf{m})$.

As the previous studies in geo-inversion (Sarma et al., 2006; Ronayne et al., 2008; Laloy et al., 2017; Comunian & Giudici, 2018; Tang et al., 2021), we assume each hydrofacies to be homogeneous, i.e., characterized by a constant value of hydraulic conductivity K and, hence, specific storage S_s. This assumption is introduced to verify our method's ability to reconstruct large-scale geological structures; it is not necessary for our methodology to work. In fact, since the discontinuity of K_i is not desirable for efficient algorithms such as gradient-based methods, we relax the inversion problem formulation by allowing K to be continuous and by constraining its range (Sarma et al., 2008; Vo & Durlofsky, 2014; Liu et al., 2019).

We use numerical experimentation to illustrate the performance of our CE framework. The simulation parameters and other settings for these experiments are borrowed from (Laloy et al., 2018; Klein et al., 2017). The numerical simulations are performed in Python using the `FEniCS` software library (Alnæs et al., 2015).[1] The first setting (Section 5.1) deals with deterministic inversion, in which CE has only two agents, F_{dat} and F_{den}. The second (Section 5.2) setting provides probabilistic treatment of a more complex geology and CE has three agents (F_{dat}, F_{den}, and F_{geo}). The meta-parameters introduced in Sections 3 and 4 are optimized to achieve the best image reconstruction performance using the exhaustive grid search (Larochelle et al., 2007). Specifically, the optimal parameter for the BM3D denoiser is chosen experimentally by visual inspection of the reconstructed images produced by the CE framework. Tuning of the DnCNN denoiser is more difficult because it does not have an explicit noise parameter (Xu et al., 2020), the noise regime being implicitly selected during the training stage (Zhang et al., 2017). For the experiments reported here, we compare performance for all pre-trained models provided by Zhang et al. (2017), selecting the one (i.e., $\sigma = 50$ model for greyscale image) that gives the best CE reconstruction according to visual inspection. The Mann iteration parameter ρ and the number of iterations are set to 0.5 and 30, respectively for the fast convergence of algorithm.

[1] The data and codes are available at `https://github.com/DDMS-ERE-Stanford/CE.git`.

The regularization parameters are set to $\sigma_{\text{dat}} = 20$ and $\sigma_{\text{geo}} = 0.5$ to achieve the best-quality reconstructed image.

5.1 Deterministic inversion without geological prior

We consider a channelized aquifer represented by a 45×45 grid consisting of 10 m $\times$ 10 m grid cells (Fig. 3). Values of the hydraulic conductivity of the channels and the ambient matrix are set to 10^{-2} m/s and 10^{-4} m/s, respectively; the true (unknown) conductivity field is shown in Fig. 3a. Radial flow is induced by a pumping well (the red circle in Fig. 3a) operating with a fixed hydraulic head of $h_{\text{well}} = 1$ m; constant head $h = 21$ m is prescribed along the left ($x_1 = 0$) and right ($x_1 = 450$ m) sides of the square aquifer; the remaining two boundaries ($x_2 = 0$ and 450 m) are impermeable; the initial hydraulic head over the entire domain is 21 m. The true hydraulic head field $h(\mathbf{x}, t)$ is computed as a numerical simulation of (16) and (17) with the true hydraulic conductivity field.

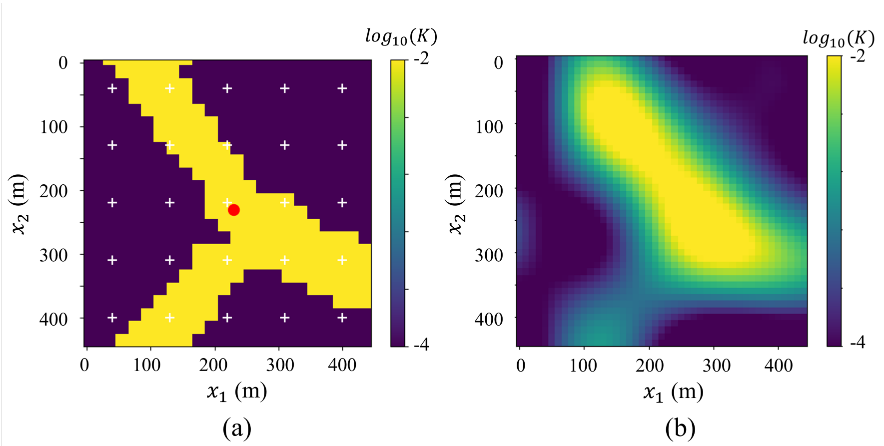

Figure 3: True log hydraulic field (left) and its initial guess (right) estimated from conductivity measurements via support vector regression.

Hydraulic head measurements are collected at 25 observation wells (the white circles in Fig. 3a) every 5 hours for the first 30 hours of the simulation. Gaussian noise with zero mean and variance of 0.2 is added to these values to account for observation errors. The resulting dataset $\mathbf{d}$, as well as conductivity values at these 16 locations, are used in inverse modeling to reconstruct the hydraulic conductivity field. Given both the relative simplicity of the geological structure in Fig. 3a and the relatively high sampling density, we use this setting to perform deterministic inversion by reconstructing only one geological map. The prior geological constraint is not enforced for this case, thus our CE framework contains only two agents, F_{dat} and F_{den}. An initial guess for the inversion can be either provided by an expert (e.g., geologist) or constructed from the 16 conductivity measurements. Figure 3b presents the initial guess estimated by support vector regression, which showed good performance for facies delineation (Wohlberg et al., 2005).

Figure 4 shows the hydraulic conductivity maps obtained via the CE-based inversion with three alternative types of agent F_{den}: TV, BM3D, and DnCNN. Regardless of the denoiser type, our CE framework captures the channel connectivity and generates realistic images. Visual inspection of these images reveals that CE with the DnCNN image denoiser delineates the facies most accurately.

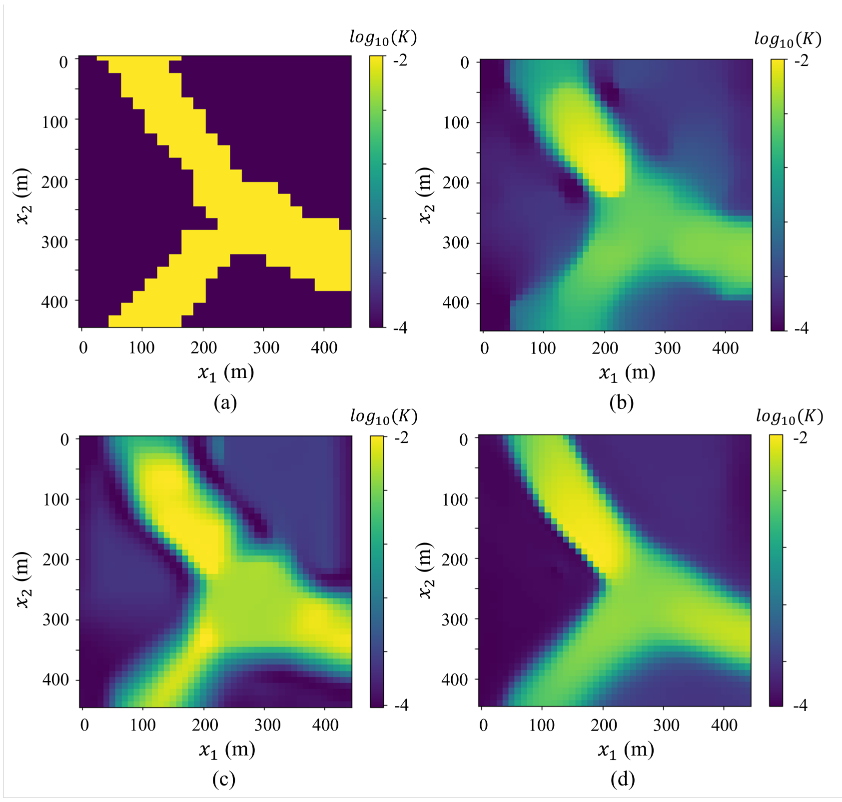

Figure 4: True hydraulic conductivity map (a) and its reconstructions via the CE-based inversion with the TV (b), BM3D (c), and DnCNN (d) denoisers.

Predictions of hydraulic head $h(\mathbf{x}, t)$, corresponding to the reconstructed conductivity $K(\mathbf{x})$, at four locations are presented in Fig. 5. Our CE-based inversion considerably reduces the discrepancy between the true and predicted hydraulic head values, regardless of the denoiser type.

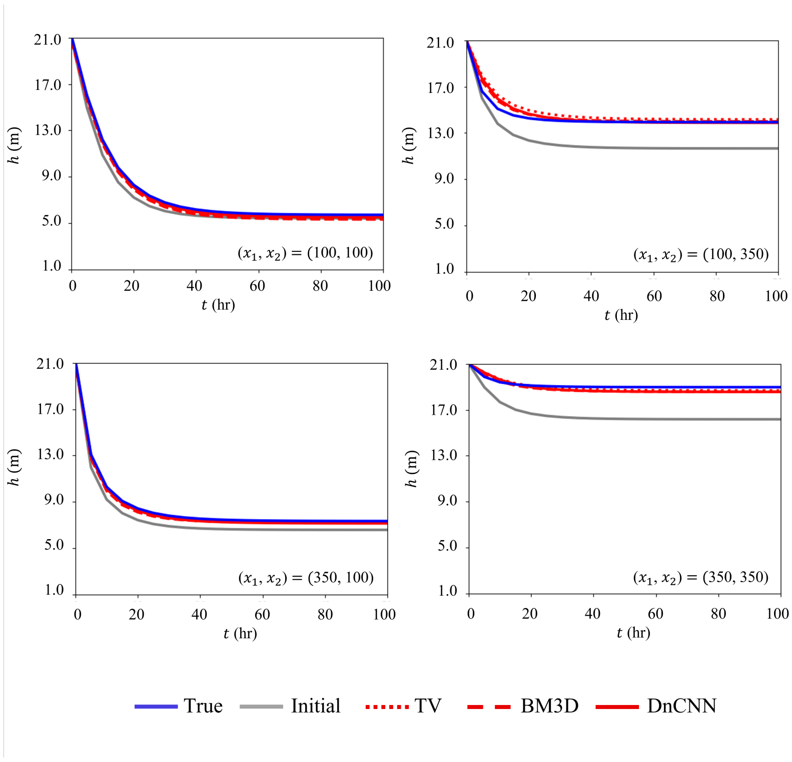

Figure 5: Temporal evolution of hydraulic head $h(\mathbf{x}, t)$ at four selected locations $\mathbf{x} = (x_1, x_2)^T$ predicted with the initial guess of $K(\mathbf{x})$ based only on conductivity measurements and with the CE-based inversion with the TV, BM3D, and DnCNN denoisers.

To quantify the relative performance of the CE framework, we introduce a classification error computed as the fraction of misclassified grid points in the total number of grid points. The facies classification is done by setting a threshold value for $\ln K$ as -3, which is the average of the log hydraulic conductivities in the two distinct facies. The best method, CE with the DnCNN denoiser, has the classification error of 5.65%; a remarkably good performance given small number of observation points (0.79% of the total number of grid points). Classification errors of CE with the TV and BM3D denoisers are 6.19% and 6.47%, respectively.

5.2 Probabilistic inversion with geologic prior

Prior geological information can be provided by an expert/geologist in the form of a training image (TI) that describes the morphology and key characteristics of hydrofacies. CE incorporates this conceptual geological knowledge through the geology prior agent, F_{geo}, which is used in addition to the other two agents, F_{dat} and F_{den}. We use the DnCNN denoiser, since it performed best on the previous problem. Multi-point geostatistics, specifically the SNESIM algorithm (Strebelle, 2002), is used to generate 2000 realizations which form the training dataset for the VAE model F_{geo}. These realizations also provide the initial guess for our inversion. The probabilistic inversion is conducted by running multiple inversions with an arbitrarily sampled initial guess; this approach is known as randomized maximum-likelihood or RML (Kitanidis, 1995). We use 50 RML runs for the stochastic data assimilation.

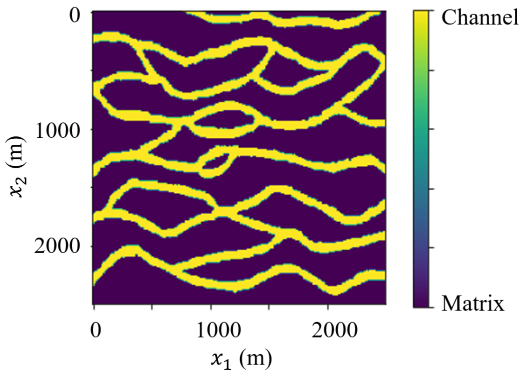

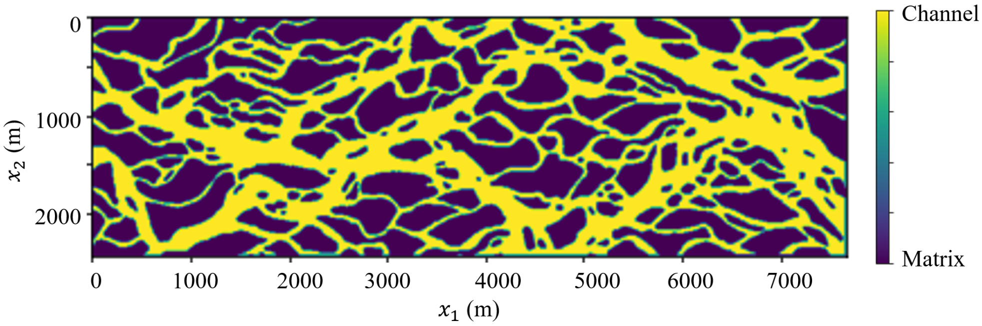

Figure 6: Training images used for multi-point geostatistical simulations: (a) a 250 × 250 hand-made drawing for case 1 and (b) a 768 × 243 image for case 2 generated from the satellite image (Mariethoz & Lefebvre, 2014).

We consider two channelized aquifers with hydraulic conductivities 10^{-2} m/s and 10^{-4} m/s for channels and matrix, respectively. Figure 6 shows the TIs for the two different cases[2]. The TI for Case 1 represents a fluvial channelized aquifer used in (Strebelle, 2002). The TI for Case 2 comes from a satellite image of the Ganges delta in Bangladesh (Mariethoz & Lefebvre, 2014; Mariethoz & Caers, 2014). The aquifers are discretized with 81×81 and 75×75 grids for Cases 1 and 2, respectively. In both cases, the size of each grid cell is 20 m × 20 m with thickness of 1 m. Figure 7 shows the true hydraulic conductivity fields and realizations generated from the given TIs.

Initially, groundwater flow is driven by constant heads of 9.0 m and 10.0 m imposed on the left and right boundaries of the domain, respectively; the remaining two boundaries are impermeable. The initial hydraulic head distribution is computed by running the flow simulator until steady state is achieved. At that time, four pumping wells operating with the fixed hydraulic head 8.0 m are installed. Observation wells record the hydraulic head response to groundwater withdrawal. Locations of

[2] The images are available at `https://wp.unil.ch/gaia/downloads/`

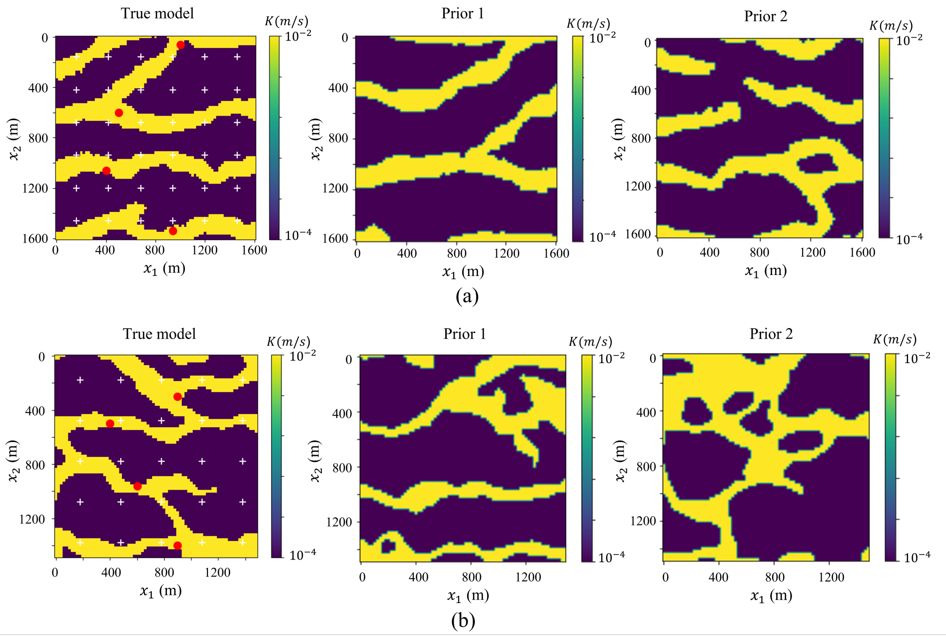

Figure 7: True geological maps (left column) and representative prior realizations generated by the SNESIM algorithm (the remaining two columns) for Cases 1 (top row) and 2 (bottom row). The white circles mark locations of observation wells.

pumping wells and observation wells are indicated by the red circles and white crosses, respectively. Hydraulic head $h(\mathbf{x}, t)$ at these observation wells is simulated with (16) and (17) for the true conductivity fields, and recorded every 20 days for the first 80 days. Dataset $\mathbf{d}$ is constructed by corrupting these simulated values with zero-mean Gaussian noise of variance 0.01.

Figures 8 and 9 exhibit the true and estimated conductivity fields. The latter represent sample averages of the realizations obtained with the initial guess, CE without F_{geo}, and CE with the DnCNN denoiser. Visual inspection of these figures demonstrates a close agreement between the reconstructed geological maps and their true counterparts. On the other hand, the image reconstruction without F_{geo} fails to preserve the geological realism and has a large discrepancy with the true image. The sample-averaged hydraulic conductivity fields for Cases 1 and 2 obtained by CE without F_{geo} (with the threshold conductivity value of 10^{-3} m/s) have the classification errors of 14.6% and 12.8%, respectively. The corresponding classification errors of CE with F_{geo} are 7.6% and 8.1% for Cases 1 and 2, respectively.

Figures 10 and 11 present the temporal evolution of hydraulic head h at two different locations for Cases 1 and 2, respectively. The RML approach allows us to quantify the prior (before inversion) and posterior uncertainty ranges. The blue areas in these figures represent the 90% confidence intervals; and the dashed lines, $t = 60$ days, indicate the end of the assimilation period. Visually, the recorded drawdowns at two different locations are dissimilar in Cases 1 and 2 due to the different degrees of connectivity to pumping wells. This distinct relationship between hydraulic head response and channel connectivity enables the close agreement between the reconstructed image and true image. The CE-enabled assimilation of hydraulic head data leads to significant uncertainty reduction in the posterior realizations. The hydraulic head pro-

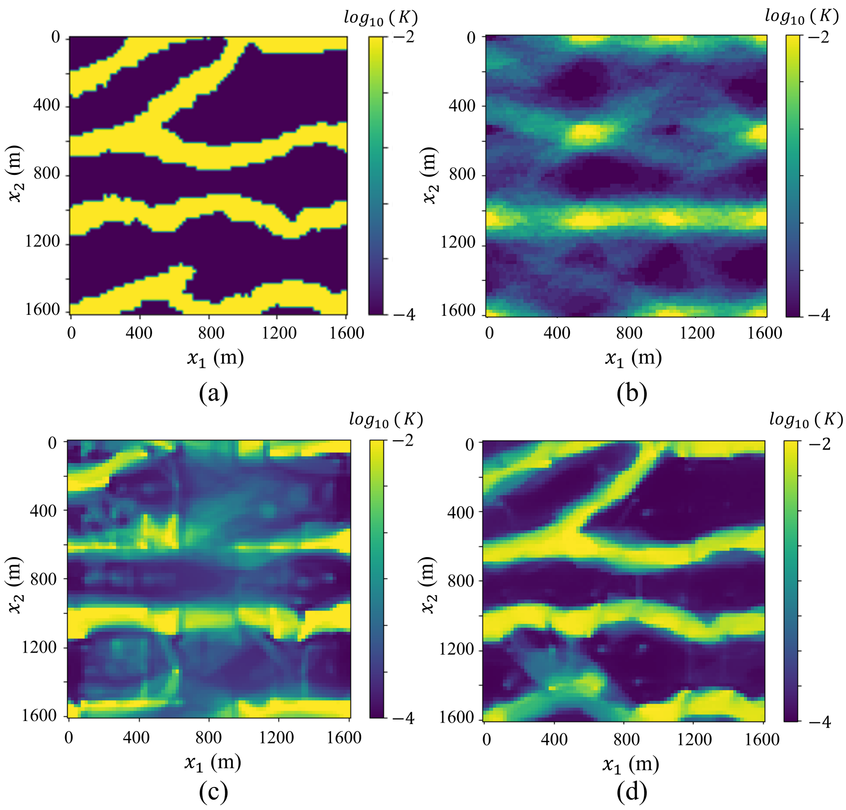

Figure 8: Case 1: (a) the true conductivity map and its reconstructions obtained by averaging the realizations of conductivity maps of (b) the initial guess, (c) CE without prior geology agent, and (d) CE with prior geology agent.

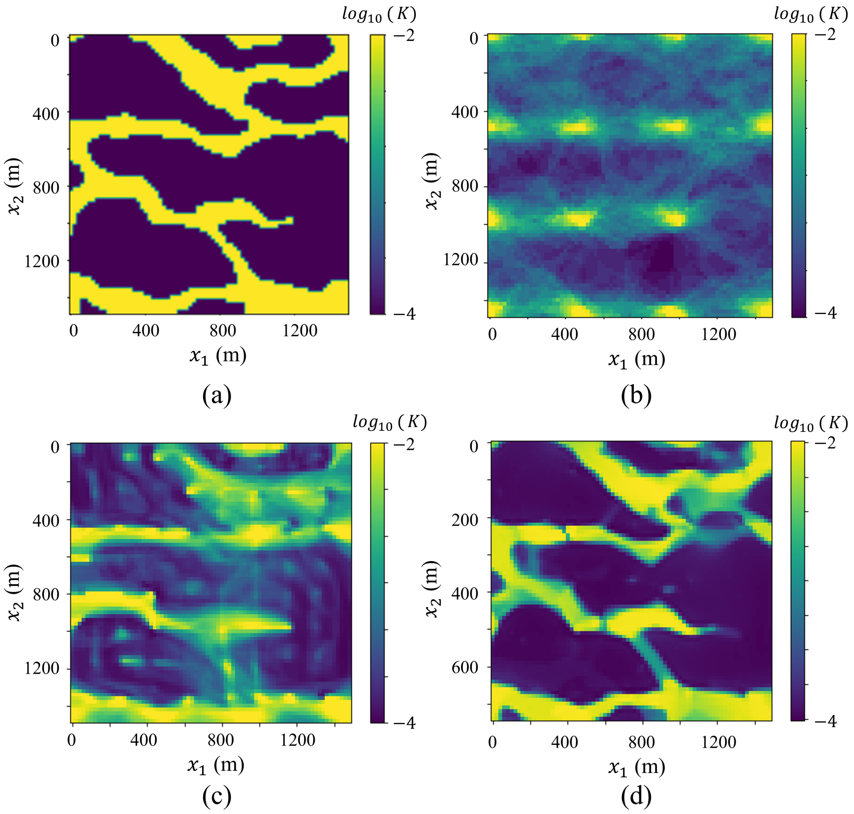

Figure 9: Case 2: (a) the true conductivity map and its reconstructions obtained by averaging the realizations of conductivity maps of (b) the initial guess, (c) CE without prior geology agent, and (d) CE with prior geology agent.

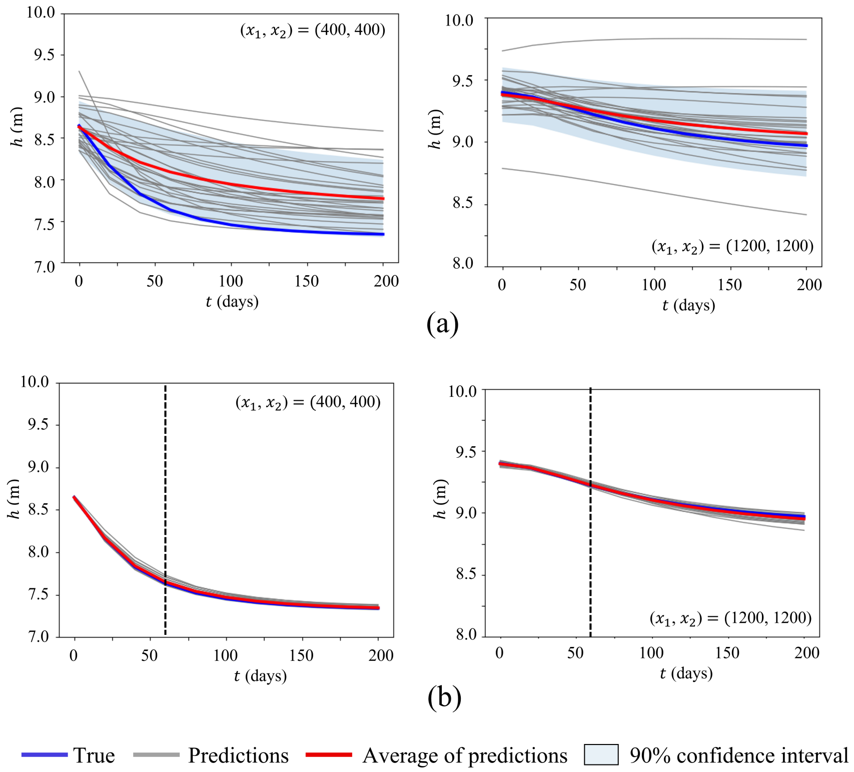

Figure 10: Hydraulic head h evolution with time at selected locations $\mathbf{x} = (x_1, x_2)^T$ of aquifers for case 1: (a) initial and (b) posterior.

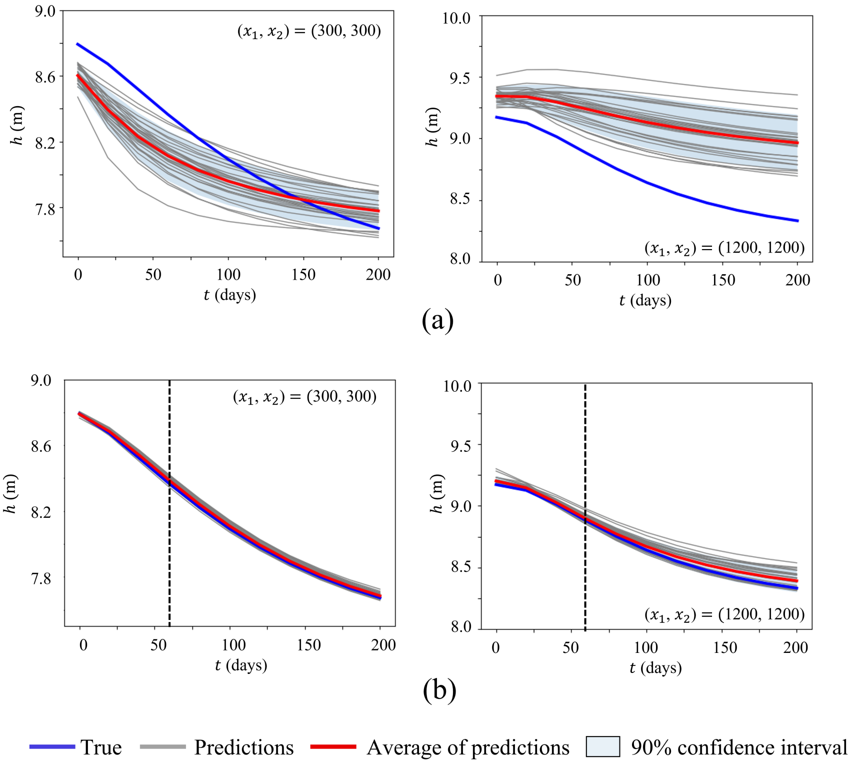

Figure 11: Hydraulic head h evolution with time at selected locations $\mathbf{x} = (x_1, x_2)^T$ of aquifers for case 2: (a) initial and (b) posterior.

files of true model fall within the 90% confidence interval in both Cases 1 and 2. These results demonstrate that our methodology has a robust prediction performance.

5.3 Computational efficiency of the proposed method

In our CE approach, each of the agents is applied sequentially at each iteration. Since the averaging and updating operations in (11) carry negligible computational costs, we approximate the overall computational cost by the sum of the computational times for each agents. Table 1 collates the computational burden for the three experiments considered in this study; the computation costs are reported for an Intel Xeon e5-2670 machine running at 2.3 GHz, and all the runs of the probabilistic inversion are parallelized on 50 computer nodes. Agent F_{dat}, an equivalent of the regularized adjoint-based inversion, consumes most of the overall computational cost. This finding implies that our method adds only a small amount of computational burden to the existing optimization-based approaches, while significantly improving the reconstruction performance.

Table 1: Computational times of the CE algorithm for the three tests considered.

Tests	Agents	CPU time (min)
Deterministic	F_{dat}	$4.39 \cdot 10^1$
	F_{den} with DnCNN	$5.3 \cdot 10^0$
Probabilistic, Case 1	F_{dat}	$7.22 \cdot 10^1$
	F_{den}	$9.81 \cdot 10^0$
	F_{geo}	$5.78 \cdot 10^0$
Probabilistic, Case 2	F_{dat}	$3.49 \cdot 10^2$
	F_{den}	$2.19 \cdot 10^1$
	F_{geo}	$1.65 \cdot 10^0$

6 Summary and Conclusions

We developed and applied a new plug-and-play approach to solve subsurface inversion problems with complex geology. Conventional optimization-based approaches are widely used for this purpose, but their applicability is often limited to geological formations characterized by multi-Gaussian fields. We overcome this limitation within the CE approach by fusing multiple heterogeneous priors with conventional physics-based inversion. Our CE strategy involves three different agents. The data fidelity agent F_{dat} uses an adjoint method to force the consistency between a solution of an inverse problem and observed data. This choice of F_{dat} is due to its computational efficiency and lack of the Gaussianity assumption, but other advanced inversion techniques can be plugged into CE as F_{dat}. The denoiser agent F_{den} minimizes the noise within each hydrofacies and generates realistic geological maps. The geology prior agent F_{geo} incorporates the prior geological knowledge using the VAE model.

We performed three numerical experiments to check the robustness of the proposed method. First, we solved a deterministic inversion problem on a relatively simple synthetic model. Next, we used our method for probabilistic reconstruction of geologically realistic models. Our numerical experiments lead to the following conclusions.

- The CE framework without agent F_{geo} performs well for the relatively simple geology. However, it cannot reflect the prior geological knowledge and fails to get a geologically realistic maps for more complex formation.
- Among several alternative denoisers, DnCNN (a CNN-based denoiser) shows the best performance as a CE component.
- When the prior geological information, such as shapes or orientations of geobodies, is available, the VAE agent trained on the prior realizations significantly improves the CE performance.
- When combined with the RML approach, our method allows one to quantify posterior uncertainties in estimates of both hydraulic parameters and flow response. Our method effectively estimates the posterior uncertainty range.

Acknowledgments
This work was supported in part by Air Force Office of Scientific Research under award number FA9550-18-1-0474, by a gift from TOTAL, and by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under project number 20200061DR. There are no data sharing issues since all of the numerical information is provided in the figures produced by solving the equations in the paper.

References

Afonso, M. V., Bioucas-Dias, J. M., & Figueiredo, M. A. T. (2010). Fast image recovery using variable splitting and constrained optimization. *IEEE Transactions on Image Processing*, *19*(9), 2345–2356.

Alnæs, M., Blechta, J., Hake, J., Johansson, A., Kehlet, B., Logg, A., ... Wells, G. N. (2015). The FEniCS project version 1.5. *Archive of Numerical Software*, *3*(100).

Barajas-Solano, D. A., Wohlberg, B., Vesselinov, V., & Tartakovsky, D. M. (2015). Linear functional minimization for inverse modeling. *Water Resources Research*, *51*(6), 4516–4531. doi: 10.1002/2014WR016179

Bouman, C. A. (2013). Model based image processing. *Purdue University*.

Boyd, S., Parikh, N., & Chu, E. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers.* Now Publishers Inc.

Buzzard, G. T., Chan, S. H., Sreehari, S., & Bouman, C. A. (2018). Plug-and-play unplugged: Optimization-free reconstruction using consensus equilibrium. *SIAM Journal on Imaging Sciences*, *11*(3), 2001–2020.

Cameron, D. A., Durlofsky, L. J., & Benson, S. M. (2016). Use of above-zone pressure data to locate and quantify leaks during carbon storage operations. *International Journal of Greenhouse Gas Control*, *52*, 32–43.

Chan, S., & Elsheikh, A. H. (2019). Parametric generation of conditional geological realizations using generative neural networks. *Comput. Geosci.*, *23*, 925-952.

Chan, S. H., Wang, X., & Elgendy, O. A. (2016). Plug-and-play ADMM for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, *3*(1), 84-98.

Chen, Y., & Oliver, D. S. (2012). Ensemble randomized maximum likelihood method as an iterative ensemble smoother. *Mathematical Geosciences*, *44*(1), 1–26.

Comunian, A., & Giudici, M. (2018). Hybrid inversion method to estimate hydraulic transmissivity by combining multiple-point statistics and a direct inversion method. *Mathematical Geosciences*, *50*(2), 147–167.

Dabov, K., Foi, A., Katkovnik, V., & Egiazarian, K. (2007). Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, *16*(8), 2080–2095.

Deutsch, C. V., & Tran, T. T. (2002). FLUVSIM: a program for object-based stochastic modeling of fluvial depositional systems. *Computers and Geosciences*, *28*(4), 525-535.

Dumoulin, V., & Visin, F. (2016). A guide to convolution arithmetic for deep learning. *arXiv preprint arXiv:1603.07285*.

Eckstein, J., & Bertsekas, D. P. (1992). On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, *55*(1), 293–318.

Ghani, M. U., & Karl, W. C. (2019). Integrating data and image domain deep learning for limited angle tomography using consensus equilibrium. In *2019 IEEE/CVF international conference on computer vision workshop (ICCVW)* (pp. 3922–3932).

Giselsson, P. (2017). Tight global linear convergence rate bounds for Douglas–Rachford splitting. *Journal of Fixed Point Theory and Applications*, *19*(4), 2241–2270.

Kamilov, U. S., Mansour, H., & Wohlberg, B. (2017). A plug-and-play priors approach for solving nonlinear imaging inverse problems. *IEEE Signal Processing Letters*, *24*(12), 1872–1876.

Kingma, D. P., & Welling, M. (2013). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.

Kitanidis, P. K. (1995). Quasi-linear geostatistical theory for inversing. *Water Resources Research*, *31*(10), 2411–2419.

Klein, O., Cirpka, O. A., Bastian, P., & Ippisch, O. (2017). Efficient geostatistical inversion of transient groundwater flow using preconditioned nonlinear conjugate gradients. *Advances in Water Resources*, *102*, 161–177.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, *25*, 1097–1105.

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*(1), 381–406.

Laloy, E., Hérault, R., Lee, J., Jacques, D., & Linde, N. (2017). Inversion using a new low-dimensional representation of complex binary geological media based on a deep neural network. *Advances in Water Resources*, *110*, 387–405.

Larochelle, H., Erhan, D., Courville, A., Bergstra, J., & Bengio, Y. (2007). An empirical evaluation of deep architectures on problems with many factors of variation. In *Proceedings of the 24th international conference on machine learning* (pp. 473–480).

Li, C., Chen, X., & Du, Z. (2004). A new relationship of rock compressibility with porosity.. (SPE-88464-MS) doi: 10.2118/88464-MS

Liu, Y., Sun, W., & Durlofsky, L. J. (2019). A deep-learning-based geological parameterization for history matching complex models. *Mathematical Geosciences*, *51*(6), 725–766.

Mariethoz, G., & Caers, J. (2014). *Multiple-point geostatistics: stochastic modeling with training images*. John Wiley & Sons.

Mariethoz, G., & Lefebvre, S. (2014). Bridges between multiple-point geostatistics and texture synthesis: Review and guidelines for future research. *Computers & Geosciences*, *66*, 66–80.

Oliver, D. S., Reynolds, A. C., & Liu, N. (2008). *Inverse theory for petroleum reservoir characterization and history matching*. Cambridge University Press.

Ono, S. (2017). Primal-dual plug-and-play image restoration. *IEEE Signal Processing Letters*, *24*(8), 1108–1112.

Osher, S., Burger, M., Goldfarb, D., Xu, J., & Yin, W. (2005). An iterative regularization method for total variation-based image restoration. *Multiscale Modeling & Simulation*, *4*(2), 460–489.

Paola, C. (2000). Quantitative models of sedimentary basin filling. *Sedimentology*,

47, 121–178.
Ronayne, M. J., Gorelick, S. M., & Caers, J. (2008). Identifying discrete geologic structures that produce anomalous hydraulic response: An inverse modeling approach. *Water Resources Research*, *44*(8).
Ryu, E. K., & Boyd, S. (2016). Primer on monotone operator methods. *Appl. Comput. Math.*, *15*(1), 3–43.
Sarma, P., Durlofsky, L. J., & Aziz, K. (2008). Kernel principal component analysis for efficient, differentiable parameterization of multipoint geostatistics. *Mathematical Geosciences*, *40*(1), 3–32.
Sarma, P., Durlofsky, L. J., Aziz, K., & Chen, W. H. (2006). Efficient real-time reservoir management using adjoint-based optimal control and model updating. *Computational Geosciences*, *10*(1), 3–36.
Sreehari, S., Venkatakrishnan, S. V., Wohlberg, B., Buzzard, G. T., Drummy, L. F., Simmons, J. P., & Bouman, C. A. (2016). Plug-and-play priors for bright field electron tomography and sparse interpolation. *IEEE Transactions on Computational Imaging*, *2*(4), 408–423.
Sridhar, V., Wang, X., Buzzard, G., & Bouman, C. (2020). Distributed iterative CT reconstruction using multi-agent consensus equilibrium. *IEEE Transactions on Computational Imaging*, *6*, 1153-1166.
Strebelle, S. (2002). Conditional simulation of complex geological structures using multiple-point statistics. *Mathematical Geology*, *34*(1), 1-21.
Sun, Y., Wohlberg, B., & Kamilov, U. S. (2019). An online plug-and-play algorithm for regularized image reconstruction. *IEEE Transactions on Computational Imaging*, *5*(3), 395–408.
Tang, M., Liu, Y., & Durlofsky, L. J. (2021). Deep-learning-based surrogate flow modeling and geological parameterization for data assimilation in 3D subsurface flow. *Computer Methods in Applied Mechanics and Engineering*, *376*, 113636.
Venkatakrishnan, S. V., Bouman, C. A., & Wohlberg, B. (2013). Plug-and-play priors for model based reconstruction. In *Proceedings of IEEE global conference on signal and information processing (GlobalSIP)* (pp. 945–948). Austin, TX, USA. doi: 10.1109/GlobalSIP.2013.6737048
Vo, H. X., & Durlofsky, L. J. (2014). A new differentiable parameterization based on principal component analysis for the low-dimensional representation of complex geological models. *Mathematical Geosciences*, *46*(7), 775–813.
Wang, P., Yang, L., Wang, H., Tartakovsky, D. M., & Onori, S. (2021). Temperature estimation from current and voltage measurements in lithium-ion battery systems. *J. Power Storage*, *34*, 102133. doi: 10.1016/j.est.2020.102133
White, J. T. (2018). A model-independent iterative ensemble smoother for efficient history-matching and uncertainty quantification in very high dimensions. *Environmental Modelling & Software*, *109*, 191–201.
Winter, C. L., Guadagnini, A., Nychka, D., & Tartakovsky, D. M. (2006). Multivariate sensitivity analysis of saturated flow through simulated highly heterogeneous groundwater aquifers. *J. Comput. Phys.*, *217*(1), 166-175.
Winter, C. L., & Tartakovsky, D. M. (2002). Groundwater flow in heterogeneous composite aquifers. *Water Resour. Res.*, *38*(8). doi: 10.1029/2001WR000450
Winter, C. L., Tartakovsky, D. M., & Guadagnini, A. (2003). Moment equations for flow in highly heterogeneous porous media. *Surv. Geophys.*, *24*(1), 81-106.
Wohlberg, B., Tartakovsky, D. M., & Guadagnini, A. (2005). Subsurface characterization with support vector machines. *IEEE Transactions on Geoscience and Remote Sensing*, *44*(1), 47-57.
Xu, X., Liu, J., Sun, Y., Wohlberg, B., & Kamilov, U. S. (2020, February). *Boosting the performance of plug-and-play priors via denoiser scaling* (Tech. Rep. No. 2002.11546). arXiv.
Yang, H.-J., Boso, F., Tchelepi, H. A., & Tartakovsky, D. M. (2020). Method of dis-

tributions for quantification of geologic uncertainty in flow simulations. *Water Resour. Res.*, *56*(7), e2020WR027643. doi: 10.1029/2020WR027643

Zhang, K., Zuo, W., Chen, Y., Meng, D., & Zhang, L. (2017). Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, *26*(7), 3142–3155.

Zhou, Z., & Tartakovsky, D. M. (2021). Markov chain Monte Carlo with neural network surrogates: Application to contaminant source identification. *Stoch. Environ. Res. Risk Assess.*, *35*(3), 639-651. doi: 10.1007/s00477-020-01888-9