



**HAL**  
open science

# Inversion using a new low-dimensional representation of complex binary geological media based on a deep neural network

Eric Laloy, Romain Héroult, John Lee, Diederik Jacques, Niklas Linde

► **To cite this version:**

Eric Laloy, Romain Héroult, John Lee, Diederik Jacques, Niklas Linde. Inversion using a new low-dimensional representation of complex binary geological media based on a deep neural network. *Advances in Water Resources*, 2017, 110, pp.387-405. 10.1016/j.advwatres.2017.09.029 . hal-02094960

**HAL Id: hal-02094960**

**<https://hal.science/hal-02094960>**

Submitted on 16 Feb 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Inversion using a new low-dimensional representation of complex binary geological media based on a deep neural network

Eric Laloy\*, Romain Hérault†, John Lee‡, Diederik Jacques§ and Niklas Linde¶

September 30, 2017

## Abstract

Efficient and high-fidelity prior sampling and inversion for complex geological media is still a largely unsolved challenge. Here, we use a deep neural network of the variational autoencoder type to construct a parametric low-dimensional base model parameterization of complex binary geological media. For inversion purposes, it has the attractive feature that random draws from an uncorrelated standard normal distribution yield model realizations with spatial characteristics that are in agreement with the training set. In comparison with the most commonly used parametric representations in probabilistic inversion, we find that our dimensionality reduction (DR) approach outperforms principle component analysis (PCA), optimization-PCA (OPCA) and discrete cosine transform (DCT) DR techniques for unconditional geostatistical simulation of a channelized prior model. For the considered examples, important compression ratios (200 - 500) are achieved. Given that the construction of our parameterization requires a training set of several tens of thousands of prior model realizations, our DR approach is more suited for probabilistic (or deterministic) inversion than for unconditional (or point-conditioned) geostatistical simulation. Probabilistic inversions of 2D steady-state and 3D transient hydraulic tomography data are used to demonstrate the DR-based inversion. For the 2D case study, the performance is superior compared to current state-of-the-art multiple-point statistics inversion by sequential geostatistical resampling (SGR). Inversion results for the 3D application are also encouraging.

## 1 Introduction

Inverse modeling plays a fundamental role in subsurface hydrology and many other Earth science disciplines. Basically, one iteratively proposes new model perturbations that are consistent with a given prior model until the resulting forward response agrees with a set of measured data up to a pre-specified level. The inversion outcome commonly consists of one or more 2D or 3D subsurface property field(s) describing, for instance, hydraulic conductivity or porosity. Due to conceptual and numerical errors in the forward problem formulation, (input and output data) measurement errors and insufficient information content in the measured data, the inverse problem may not admit a unique solution.

\*Belgian Nuclear Research Centre, Engineered and Geosystems Analysis. Email: [elaloy@sckcen.be](mailto:elaloy@sckcen.be)

†Normandie Univ, UNIROUEN, UNIHAVRE, INSA Rouen, LITIS, 76000 Rouen, France

‡IREC/MIRO and ICTEAM/MLG, Université catholique de Louvain

§Belgian Nuclear Research Centre, Engineered and Geosystems Analysis

¶Applied and Environmental Geophysics Group, Institute of Earth Sciences, University of Lausanne

Hence, the inversion process should ideally provide an ensemble of geologic model realizations that accurately span the range of possible geologic models that honor both the available data used in the inversion and prior information.

The prior model is often conveniently specified by a two-point geostatistical model or variogram, by assuming the subsurface structure to be multi-Gaussian. Whenever this simplifying assumption does not hold (e.g., Gómez-Hernández and Wen, 1998; Journel and Zhang, 2006), the prior model can instead be informed by a so-called training image (TI). The TI is a large gridded 2D or 3D unconditional representation of the expected target spatial field that can be either continuous or categorical (e.g., geologic facies image). To generate model realizations that reproduce higher-order statistics found in the TI (together with possible direct point conditioning data), one can use a multiple-point statistics (MPS) simulation method (e.g., Guardiano and Srivastava, 1993; Strebelle, 2002; Hu and Chugunova, 2008; Mariethoz et al., 2010b; Tahmasebi et al., 2012).

Since MPS does generally not rely on an explicit mathematical description of the prior model in terms of model parameters, it is impossible to use it for classical parameter-based inversion. This prohibits the use of gradient-based methods such as randomized maximum likelihood (RML, Kitanidis, 1995; Oliver et al., 1996) or regular Markov chain Monte Carlo (MCMC) sampling (Robert and Casella, 2004; Brooks et al., 2011). Note that this limitation also holds for the Markov mesh model (MMM, Stien and Kolbjørnsen, 2011) method despite the fact that it builds a statistical model describing transition probabilities of the facies distribution. Instead, one needs to resort to sampling from the prior model using MPS simulation. This process that is often referred to as sequential geostatistical resampling (SGR) inversion (Hansen et al., 2008; Mariethoz et al., 2010a; Hansen et al., 2012) basically consists of iteratively generating new random perturbations from the prior model and to accept or reject them based on the resulting data misfit. The SGR inversion method is powerful in finding models that fit the data, but it suffers from two important drawbacks. As shown by Laloy et al. (2016), (1) it tends to only explore the immediate vicinity of a single global minimum even when the objective function/posterior landscape is highly multi-modal and (2) may produce degraded geologic model realizations when fitting large datasets with high signal-to-noise ratios (SNR).

Although some recent work has demonstrated substantial improvements to SGR-based inversion (Zahner et al., 2016), in this study we take another direction and propose a new parametric dimensionality reduction approach for complex binary prior models. The driving idea is that if one can build a lower-dimensional model parameterization from which one can sample (after appropriate non-linear transformations) model realizations that are consistent with the TI, then global parameter-based inverse methods for moderately large continuous parameter spaces can be used. For instance, this opens up the possibility to explore the posterior model distribution using state-of-the-art adaptive MCMC sampling with DREAM<sub>(ZS)</sub> (Vrugt et al., 2009; Laloy and Vrugt, 2012).

Previous work on using compressed parametric bases for hydrogeological inversion has relied on principal component analysis (PCA, e.g., Reynolds et al., 1996; Sarma et al., 2006), kernel-PCA (e.g., Sarma et al., 2008), level-set (e.g., Dorn and Villegas, 2008), discrete wavelet transforms (e.g., Awotunde and Horne, 2013), discrete cosine transform (DCT, e.g., Jafarpour et al., 2010; Linde and Vrugt, 2013), singular value decomposition (SVD, e.g., Tavakoli and Reynolds, 2011), and K-SVD (e.g., Khaninezhad et al., 2012; Khaninezhad and Jafarpour, 2014) to name most of the strategies used. However, none of these bases define a manifold that is restricted to complex (non-Gaussian) geological models that are in strong agreement with a specific TI. In other words, randomly sampling the parameter space produces model realizations that are generally inconsistent

with the TI. To the best of our knowledge, there is so far only the optimization-PCA approach (O-PCA) by Vo and Durlofsky (2014, 2015) that aims to build a lower-dimensional model representation that is fully consistent with the TI. The O-PCA essentially consists of a post-processing of a PCA model. To work it is necessary that the underlying PCA representation encodes sufficient information regarding important patterns in the TI model. As illustrated later on in this paper, this may not be the case when the TI contains highly connected features with a relatively large degree of variability.

In this work, we use for the first time a deep neural network (DNN, see, e.g., Goodfellow et al., 2016, for an overview) to build a parametric low-dimensional representation of complex geologic models. More specifically, we train the generator of a deep variational autoencoder (VAE, Kingma and Welling, 2014) such that it generates TI-consistent geologic model realizations when fed with a low-dimensional standard Gaussian noise vector. A key characteristic of our VAE is that it is largely made up of convolutional layers (see, e.g., LeCun et al., 1998; Krizhevsky et al., 2012). This type of neural layer is based on an adaptive convolutional filter and is well suited to image processing (Krizhevsky et al., 2012; Goodfellow et al., 2016). Our model generator is then used within a Bayesian inversion framework to sample the posterior distribution of binary 2D and 3D subsurface property fields.

This paper is organized as follows. Section 2 presents the different elements of our dimensionality reduction (DR) approach and the selected inversion framework. This is followed in section 3 with the analysis of its performance in generating geostatistical realizations using 2D and 3D TIs with and without conditioning on direct point data (also referred to as hard data), together with some comparisons against other existing DR methods. Synthetic 2D and 3D experiments involving both steady-state and transient groundwater flow are then used to demonstrate our proposed DR-based inversion approach. In section 5 we discuss the advantages and limitations of our method and outline possible future developments. Finally, section 6 concludes with a summary of the most important findings.

## 2 Methods

### 2.1 Deep neural network architecture

#### 2.1.1 Generalities

We consider a deep neural network that belongs to the class of autoencoders (AEs, see, e.g., Goodfellow et al., 2016). Neural networks basically define the (possibly complex) relationships existing between input,  $\mathbf{x}$ , and output,  $\mathbf{y}$ , data vectors by using combinations of computational units that are called neurons. A neuron is an operator of the form:

$$h(\mathbf{x}) = f(\langle \mathbf{x}, \mathbf{w} \rangle + b), \quad (1)$$

where  $h(\cdot)$  is the scalar output of the neuron,  $f(\cdot)$  is a nonlinear function that is called the ‘‘activation function’’,  $\langle \cdot, \cdot \rangle$  signifies the scalar product,  $\mathbf{w} = [w_1, \dots, w_N]$  is a set of weights of same dimension,  $N$ , as  $\mathbf{x}$  and  $b$  represents the bias associated with the neuron. For a given task, the values for  $\mathbf{w}$  and  $b$  associated with each neuron must be optimized or ‘‘learned’’ such that the resulting neural network performs as well as possible. When  $f(\cdot)$  is differentiable,  $\mathbf{w}$  and  $b$  can be learned by gradient descent. Common forms of  $f(\cdot)$  include the rectified linear unit (ReLU), sigmoid function and hyperbolic tangent function.



When there is no directed loops or cycles across neurons or combinations thereof, the network is said to be feedforward (FFN). In the FFN architecture, the neurons are organized in layers. A standard layer is given by

$$\mathbf{h}(\mathbf{x}) = f(\mathbf{W}\mathbf{x} + \mathbf{b}), \quad (2)$$

where  $\mathbf{W}$  and  $\mathbf{b}$  are now a matrix of weights and a vector of biases, respectively. The name multilayer perceptron (MLP) designates a FFN with more than one layer. A most typical network is the 2-layer MLP, which consists of two layers with the outputs of the first-layer neurons becoming inputs to the second-layer neurons

$$\mathbf{y} = \mathbf{g}[\mathbf{h}(\mathbf{x})] \equiv f_2[\mathbf{W}_2 \cdot f_1(\mathbf{W}_1\mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2], \quad (3)$$

where  $\mathbf{g}(\cdot)$  and  $\mathbf{h}(\cdot)$  are referred to as output layer and hidden layer, respectively.

In theory, the two-layer MLP described in equation (3) is a universal approximator as it can approximate any underlying process between  $\mathbf{y}$  and  $\mathbf{x}$  (Cybenko, 1989; Hornik, 1991). However, this only works if the dimension of  $\mathbf{h}(\cdot)$  is (potentially many orders of magnitudes) larger than that of the input  $\mathbf{x}$ , thereby making learning practically infeasible and the two-layer MLP approximator useless for large  $N$  (typically  $N \geq 10-25$ ). For high-dimensional input data such as images, researchers have found that it is much more efficient to use many hidden layers rather than increasing the size of a single hidden layer (e.g., Goodfellow et al., 2016). When a FFN/MLP has more than one hidden layer it is considered to be deep. Nevertheless, current deep networks are not necessarily purely FFN but may mix different aspects of FFN, such as convolutional neural networks (CNN, see section 2.1.3 below) and recurrent neural networks (RNN, see, e.g., Goodfellow et al., 2016).

### 2.1.2 Variational autoencoders

An AE is a deep neural network that defines a reversible, nonlinear low-dimensional parameterization of (higher dimensional) input data. Consequently, it has an hourglass-like shape in terms of neural layer sizes (see simplified representation in Figure 1). The most central layer is referred to as the "code" and defines the low-dimensional space, while the output layer has the same dimensionality as the input layer. The part of the network that connects the input to the code is called the "encoder", while the one that connects the code to the output is referred to as the "decoder".

Autoencoders are generative, which means that, after appropriate training (see section 2.2), they can be used to generate new pattern realizations that are consistent with those found in a given set of (training) features. However, only the class of AEs formed by the so-called variational autoencoders (VAE) can use white noise as input to generate new patterns. In a geostatistical context, this implies that the VAE can be trained such that it randomly generates new geologic model realizations that honor the higher-order statistics found in a set of training images. How to achieve this is further described later on in this section and in section 2.2.

A VAE network can be summarized as follows. The code consists of two lower-dimensional vectors: a vector of means,  $\boldsymbol{\mu}$ , and a vector of standard deviations,  $\boldsymbol{\sigma}$  (Figure 2). The first element of the decoder is a randomly sampled vector of "latent" standard normal variables,  $\mathbf{z}_l$ , of the same size,  $d$ , as  $\boldsymbol{\mu}$  and  $\boldsymbol{\sigma}$ . The  $\mathbf{z}_l$  vector is subsequently rescaled into  $\mathbf{z} = \mathbf{z}_l \times \boldsymbol{\sigma} + \boldsymbol{\mu}$  (where  $\times$  means element-wise multiplication), and the so-produced  $\mathbf{z}$  vector continues its journey through the decoder to eventually produce the output,  $\hat{\mathbf{x}}$  (Figure 2). Here the symbol  $\hat{\mathbf{x}}$  is used instead of  $\mathbf{y}$  to make it clear that for

this application the output is a reconstruction of the input,  $\mathbf{x}$ . As detailed in section 2.2, the  $\mathbf{w}$  vectors and  $b$  values that are associated with every neuron of every layer of the network (encoder part,  $\boldsymbol{\mu}$ ,  $\boldsymbol{\sigma}$  and decoder part) are jointly optimized by gradient descent such that (1) the differences between  $\mathbf{x}$  and  $\hat{\mathbf{x}}$  are minimized and (2) the  $d$ -dimensional  $\mathbf{z}$  vector conforms as closely as possible to a  $N(\mathbf{0}_d, \mathbf{I}_d)$  with  $\mathbf{0}_d$  a  $d$ -dimensional zero-vector and  $\mathbf{I}_d$  the  $d \times d$  identity matrix. After training, new model realizations can be generated by sampling  $\mathbf{z} \propto N(\mathbf{0}_d, \mathbf{I}_d)$  and running the  $\mathbf{z}$  samples through the decoder.

At this point, we would like to re-emphasize that it is the combination of (1) a low-dimensional parameterization in terms of a multivariate standard normal distribution and (2) the ability to map a given set of low-dimensional parameter values into complex MPS geostatistical model realizations that makes our presented method unique and suitable for MPS-based inversion.




Figure 1: Schematic example of an autoencoder with three hidden layers,  $\mathbf{h}_1$ ,  $\mathbf{h}_2$  and  $\mathbf{h}_3$ . The  $\mathbf{x} = [x_1, \dots, x_6]$  vector is a 6-dimensional input feature, the  $h_{(\cdot, \cdot)}$  denote the hidden layers with (red-contoured)  $h_{(2, \cdot)}$  the central layer, or “code”, where the largest dimensionality reduction is performed and the  $\hat{\mathbf{x}} = [\hat{x}_1, \dots, \hat{x}_6]$  vector signifies the reconstructed input.




Figure 2: Illustration of a variational deep autoencoder (VAE). The boxes labeled with the  $\mu$  and  $\sigma$  letters represent the low-dimensional mean and standard deviation vectors, respectively, that are encoded from the input data. During training,  $z_l$  is drawn from the fixed  $p(z_l)$  distribution and then rescaled into  $z = z_l \times \sigma + \mu$  (where  $\times$  means element-wise multiplication). Training aims at jointly minimizing (1) the data reconstruction loss and the deviations between (2)  $\mu$  and the zero-vector of the same size and (3)  $\sigma$  and the unit vector of the same size. The ensemble of tasks (2) and (3) is called the “reparameterization trick” (Kingma and Welling, 2014). For subsequent stochastic data generation, the process starts with sampling  $z$  from the fixed  $p(z)$  distribution.

### 2.1.3 Convolutional layers

The last salient characteristic of our deep VAE is the use of convolutional layers at various levels of both the encoder and the decoder. The convolutional layer is the main building block of the CNN-type of architecture. It is particularly well suited for image processing applications because it explicitly accounts for the spatial structure of the input data, whether the input image,  $\mathbf{X}$ , or the incoming hidden layer in the network. When the input is a 2D image (with possibly 3 channels for a RGB image), a convolutional layer,  $\mathbf{h}^k$ , is built from a series of  $k = 1, \dots, N_k$  small  $N_i \times N_j$  filters,  $\mathbf{w}^k$ , that convolve an input pixel,  $X_{m,n}$  to  $h_{m,n}^k$  as

$$h_{m,n}^k(X_{m,n}) = f \left( \sum_{i=1, j=1}^{N_i, N_j} w_{i,j}^k X_{m+i, n+j} + b_k \right), \quad (4)$$

where for computational efficiency  $f(\cdot)$  is typically a ReLU:  $f(x) = \max(0, x)$ . The so-produced ensemble of  $N_k$   $\mathbf{h}^k$  “feature maps” forms a volume called the convolutional representation. The larger the  $N_k$ , the potentially more comprehensive is the representation of the input data. Other important convolution parameters that we do not discuss here are the “stride”, that is, the degree of overlapping between successive moves in the forward pass of a given filter, and the “zero-padding” which involves padding the borders of the input image or volume for size preservation. For further information, we refer the reader to Goodfellow et al. (2016) and online tutorials<sup>1</sup>.

<sup>1</sup>For instance: <http://deeplearning.net/tutorial/> and <http://cs231n.github.io/convolutional-networks/>.

Note that each neuron of a given convolutional layer (volume) is connected to only a local region of the input image or volume, called “receptive field” (see Figure 3). This enforces spatially-local exploration at the level of one convolutional layer. The encoded spatial information includes increasingly larger patterns when the number of stacked convolutional layers increases. Also, recall that the elements (pixels/voxels) of a given feature map share the same  $\mathbf{w}^k$  and  $b_k$ . The number of CNN parameters to be optimized is thus greatly reduced compared to that of a standard fully-connected (FC) (or dense) FFN of similar depth. Besides the convolutional layer itself, the other important component of a CNN is the “pooling” layer. Pooling is a form of nonlinear down-sampling used to reduce the dimensions of a convolutional layer and thereby limit overfitting. It proceeds by converting small sub-regions of equal size into single values, such as the maximum or the mean. Using the maximum (“max-pooling”) is the most common practice. When used, pooling always follows a series of one or more convolution operations.




Figure 3: Stylized representation of convolutional layer neurons connected to their receptive field (picture inspired by <http://cs231n.github.io/convolutional-networks/>).

#### 2.1.4 Practical implementation

We now have described all of the main ingredients of our devised deep VAECNN network, which is depicted in Figure 4. For the sake of brevity, the various individual fully-connected (FC), variational (Q), convolutional (C) and pooling (P) layers (see Figure 4) are not detailed in this paper but the corresponding computer code is available from the first author. Our VAECNN was implemented within the open-source LASAGNE Python software (Lasagne, 2017) which works on top of the open-source THEANO Python library (Theano, 2016). In addition, note that for each result presented in this study, the  $\mathbf{z}$

vector is always 50-dimensional no matter the dimensions of the considered geologic model domain (from  $100 \times 100$  to  $30 \times 32 \times 27$ ).

Unless stated otherwise, model generation is performed by recycling the initially produced model (bottom blue architecture in Figure 4) through the full network (top light red architecture in Figure 4) for 10 times. This sequential relooping generally improves the quality of the generated model realizations. The very last step of our generation process is an hard thresholding of the produced model. We used the 0.5 cutoff value although the classical Otsu’s method could be used as well (Otsu, 1979). Fortunately, the impact of relooping on generation time remains quite acceptable. For the models considered herein, generating a realization without relooping takes about 0.01 s - 0.03 s on a modern workstation. Including the 10 successive cycles then incurs a total generation time in the range 0.2 s - 0.4 s.

## 2.2 Training the deep variational autoencoder

Considering the case where a  $100 \times 100$  binary image is generated from a 50-dimensional  $\mathbf{z}$ , the network described in Figure 4 includes 34,039,237 optimized parameters (weights and bias associated with neurons of dense layers and filters of convolutional layers). Such a high number of optimized parameters is not unusual for applications of deep neural networks. As the activation functions are differentiable, using an objective function that is also differentiable makes it possible to optimize/train the network by stochastic gradient descent (that is, gradient descent using a series of mini-batches rather than all the data at once) together with back propagation. This means that the objective function derivative is propagated backwards throughout the network using the chain rule, in order to update the parameters. Various stochastic gradient descent algorithms are available. In this work, we used the adaptive moment estimation (ADAM) algorithm which has been proven efficient for different types of deep networks (Kingma and Ba, 2015).

The objective or loss function,  $\mathcal{L}$ , must be set such that at training time: (1) the differences between reconstructed images,  $\hat{\mathbf{X}}$ , and their original counterparts,  $\mathbf{X}$ , are minimized and (2) the  $d$ -dimensional random  $\mathbf{z}$  vector conforms as closely as possible to a  $N(\mathbf{0}_d, \mathbf{I}_d)$ . This translates into the following formulation (Kingma and Welling, 2014; Gregor et al., 2015)

$$\mathcal{L} = \mathcal{L}^x + \mathcal{L}^z, \tag{5}$$

with

$$\mathcal{L}^x = -\log [p(\mathbf{X}|\mathbf{z})], \tag{6}$$

$$\mathcal{L}^z = D_{KL} [q(\mathbf{z}|\mathbf{X}) || p(\mathbf{z})], \tag{7}$$

where  $p(\mathbf{z})$  is the probability density function (pdf) of the code  $\mathbf{z}$  (here  $p(\mathbf{z}) \equiv N(\mathbf{0}_d, \mathbf{I}_d)$ ) and the conditional distribution  $p(\mathbf{X}|\mathbf{z})$  is referred to as a probabilistic decoder, since given  $\mathbf{z}$  it generates a probability distribution over the possible corresponding values of  $\mathbf{X}$ . Similarly,  $q(\mathbf{z}|\mathbf{X})$  denotes a probabilistic encoder, as given a feature  $\mathbf{X}$  it creates a distribution over the possible values of  $\mathbf{z}$  from which  $\mathbf{X}$  could have been produced (Kingma and Welling, 2014). To ensure that  $q(\mathbf{z}|\mathbf{X})$  is close to the targeted  $p(\mathbf{z})$ , the Kullback-Leibler divergence ( $D_{KL}$ ) from  $q(\mathbf{z}|\mathbf{X})$  to  $p(\mathbf{z})$ ,  $D_{KL} [q(\mathbf{z}|\mathbf{X}) || p(\mathbf{z})]$  is minimized (equation (7)).

If  $\mathbf{X}$  is a binary image, equation (6) can be rewritten as a binary cross-entropy minimization

$$\mathcal{L}^x = \sum_{i=1}^N \left\{ -X_i \log [\hat{X}_i] - (1 - X_i) \log [(1 - \hat{X}_i)] \right\}. \quad (8)$$

In addition, for  $p(\mathbf{z}) \equiv N(\mathbf{0}_d, \mathbf{I}_d)$  equation 5 becomes

$$\mathcal{L}^z = \frac{1}{2} \left[ \sum_{i=1}^d (\mu_i^2 + \sigma_i^2 - \log \sigma_i^2) \right] - \frac{d}{2}, \quad (9)$$

where the  $\mu_i$  and  $\sigma_i$  correspond to the  $i^{\text{th}}$  elements of the  $\boldsymbol{\mu}$  and  $\boldsymbol{\sigma}$  vectors displayed in Figure 2, respectively. Working with  $\boldsymbol{\mu}$  and  $\boldsymbol{\sigma}$  vectors rather than with the  $\mathbf{z}$  vector directly is called the ‘‘reparameterization trick’’ (for details, see Kingma and Welling, 2014).

Finally, in practice one may need to weight the components of  $\mathcal{L}$ , leading to

$$\mathcal{L} = \mathcal{L}^x + \alpha \mathcal{L}^z, \quad (10)$$

where  $\alpha$  is a weight factor that we set after limited trial and error either to 20 or to 40 depending on the application.

The minimization of equation (10) was performed on a GPU Tesla K40 and training the VAECNN for 100 epochs (full cycles of the stochastic gradient descent) took between 5 and 13 hours, depending on the size of the model domain (from  $100 \times 100$  to  $32 \times 30 \times 27$ ), the number of training images (from 19,000 to 80,000), and the exact size of the network. Computing times will be further detailed below for each test case individually.

### 2.3 Assessing geostatistical simulation quality

Although inversion is our primary objective, assessing the quality of the geostatistical realizations produced by our DR approach is important as it will determine the ultimate quality of our inversion results. Following Tan et al. (2014), we considered both how well the patterns in the training images are reproduced by the DR-based realizations and the between-realization variability, also termed ‘‘space of uncertainty’’.

Besides visual inspection, we estimated closeness with the training images by means of the two-point cluster function (CF, Torquato et al., 1988) which is also called connectivity function (Pardo-Igúzquiza and Dowd, 2003). The CF is the probability that there exists a continuous path of the same facies between two points of the same facies separated by a given lag distance (see, Torquato et al., 1988; Pardo-Igúzquiza and Dowd, 2003, for mathematical details). Using the implementation by Lemmens et al. (2017), the CF was calculated for each facies along the  $x$ ,  $y$ , and main diagonal,  $d_{xy}$ , directions for the 2D case studies. For the 3D case studies, we also considered the  $z$  and main diagonal  $d_{xz}$  and  $d_{yz}$  directions. This resulted into a total of six CF curves for each combination of indicator and facies.

The so-called space of uncertainty is a measure of the between-realizations variability associated with given geostatistical simulation algorithm. We characterized the space of uncertainty for the 2D case using an average distance between multiple-point histograms (MPH) over 100 realizations (see Tan et al., 2014, for details). This was done both for the training set and DR-based realizations. The used distance is the so-called Jensen-Shannon (JS) divergence (Cover and Thomas, 1991; Tan et al., 2014).

$$d_{JS}(\mathbf{m}^k, \mathbf{m}^{k'}) = \frac{1}{2} \sum_{i=1}^{N_b} MPH_i^k \log \left[ \frac{MPH_i^k}{MPH_i^{k'}} \right] + \frac{1}{2} \sum_{i=1}^{N_b} MPH_i^{k'} \log \left[ \frac{MPH_i^{k'}}{MPH_i^k} \right], \quad (11)$$

where  $\mathbf{m}^k$  and  $\mathbf{m}^{k'}$  are two different model realizations and  $MPH_1^k, \dots, MPH_{N_b}^k$  are the components of the  $\mathbf{MPH}^k$  vector calculated for realization  $\mathbf{m}^k$ . Since the MPH was computed for a  $4 \times 4$  template,  $N_b$  was  $2^{16} = 65,536$ .

The space of uncertainty, defined here as the average distance,  $\overline{d_{JS}}$ , for the set of either  $K$  training images or  $K$  DR-based realizations, is given by

$$\overline{d_{JS}} = \frac{1}{K(K-1)} \sum_{k=1}^K \sum_{k'=1}^K d_{JS}(\mathbf{m}^k, \mathbf{m}^{k'}). \quad (12)$$

## 2.4 Bayesian inversion

A common representation of the forward problem is

$$\mathbf{d} = F(\boldsymbol{\theta}) + \mathbf{e}, \quad (13)$$

where  $\mathbf{d} = (d_1, \dots, d_N) \in \mathbb{R}^N$ ,  $N \geq 1$  is the measurement data,  $F(\boldsymbol{\theta})$  is a deterministic forward model with parameters  $\boldsymbol{\theta}$  and the noise term  $\mathbf{e}$  lumps all sources of errors.

In the Bayesian paradigm, parameters in  $\boldsymbol{\theta}$  are viewed as random variables with a posterior pdf,  $p(\boldsymbol{\theta}|\mathbf{d})$ , given by

$$p(\boldsymbol{\theta}|\mathbf{d}) = \frac{p(\boldsymbol{\theta})p(\mathbf{d}|\boldsymbol{\theta})}{p(\mathbf{d})} \propto p(\boldsymbol{\theta})L(\boldsymbol{\theta}|\mathbf{d}), \quad (14)$$

where  $L(\boldsymbol{\theta}|\mathbf{d}) \equiv p(\mathbf{d}|\boldsymbol{\theta})$  signifies the likelihood function of  $\boldsymbol{\theta}$ . The normalization factor  $p(\mathbf{d}) = \int p(\boldsymbol{\theta})p(\mathbf{d}|\boldsymbol{\theta})d\boldsymbol{\theta}$  is not required for parameter inference when the parameter dimensionality is fixed. In the remainder of this paper, we will thus focus on the unnormalized density  $p(\boldsymbol{\theta}|\mathbf{d}) \propto p(\boldsymbol{\theta})L(\boldsymbol{\theta}|\mathbf{d})$ .

To avoid numerical over- or underflow, it is convenient to work with the logarithm of  $L(\boldsymbol{\theta}|\mathbf{d})$  (log-likelihood):  $\ell(\boldsymbol{\theta}|\mathbf{d})$ . If we assume  $\mathbf{e}$  to be normally distributed, uncorrelated and with known constant variance,  $\sigma_e^2$ ,  $\ell(\boldsymbol{\theta}|\mathbf{d})$  can be written as

$$\ell(\boldsymbol{\theta}|\mathbf{d}) = -\frac{N}{2} \log(2\pi) - N \log(\sigma_e) - \frac{1}{2} \sigma_e^{-2} \sum_{i=1}^N [d_i - F_i(\boldsymbol{\theta})]^2, \quad (15)$$

where the  $F_i(\boldsymbol{\theta})$  are the simulated responses that are compared the  $i = 1, \dots, N$  measurement data,  $d_i$ .

An exact analytical solution of  $p(\boldsymbol{\theta}|\mathbf{d})$  is not available for the type of non-linear inverse problems considered herein. We therefore resort to MCMC simulation (see, e.g., Robert and Casella, 2004). More specifically, the DREAM<sub>(ZS)</sub> algorithm is used to approximate the posterior distribution. A detailed description of this sampling scheme including a proof of ergodicity and detailed balance can be found in Vrugt et al. (2009) and Laloy and Vrugt (2012). Multiple contributions in hydrology and geophysics (amongst others) have demonstrated the ability of DREAM<sub>(ZS)</sub> to sample target distributions with 50-250 dimensions (Laloy et al., 2012, 2013; Linde and Vrugt, 2013; Laloy et al., 2015; Lochbühler et al., 2015).

### 3 Geostatistical simulation results

Before considering our main goal, that is, solving inverse problems (section 4), we investigate the performance of our DR approach for unconditional and conditional geostatistical simulation. This step is important as we aim at producing inversion models that not only honor the data used, but also the spatial statistics of our selected TI. One thus need to make sure that model realizations produced by randomly sampling the constructed low-dimensional base are consistent with the TI.

#### 3.1 2D channelized aquifer

Our first example considers unconditional model generation of a  $100 \times 100$  2D binary and channelized aquifer. The underlying TI is closely related to the classical  $250 \times 250$  binary training image introduced by Strebelle (2002) that is routinely used to test new MPS algorithms. Our used TI has some slight additional variability in that the channel rotation angles are allowed to deviate (at most) from  $-15^\circ$  to  $+15^\circ$  compared to the rotation angles found in the original TI. The training set used to train the VAECNN contains 80,000  $100 \times 100$  geostatistical realizations generated with the direct sampling (DS) approach as implemented in the DeeSse MPS algorithm (Mariethoz et al., 2010b). Obtaining these 80,000 realizations took 14 hours using eight CPUs in parallel. The choice of 80,000 realizations was dictated by available computational resources. Limited testing further indicated that using 40,000 model realizations for training would very likely not change the results presented below. Training the VAECNN for this example incurred a computational time of approximately 13 hours. As  $\mathbf{z}$  is 50-dimensional, the achieved compression ratio is 200.

Figure 5 presents four (randomly chosen) MPS realizations from the training set, together with eight (randomly chosen) model realizations obtained by our DR approach. The DR-based realizations resemble the training model realizations well, although some artifacts remain in terms of a moderately larger occurrence of broken channels and a slight oversmoothing of the channels. Figure 6 displays the corresponding CF metrics. It is observed that the DR-based realizations have similar CF curves as the training set. In terms of space of uncertainty, the ratio of the  $\overline{d_{JS}}$  (equation (12)) of the DR-based realizations to the  $\overline{d_{JS}}$  of the training images,  $\frac{\overline{d_{JS}(DR)}}{\overline{d_{JS}(TR)}}$ , is 1.03. Thus, the DR-based realizations show slightly more between-realization variability than the training set.

Figure 7 illustrates corresponding results obtained by the PCA, OPCA and DCT dimensionality reduction methods for the same training set. The PCA and OPCA realizations were derived following Vo and Durlafsky (2014) using 70 random variables. Varying this number in the range 50 - 10000 did not improve generation performance. For DCT, 250 coefficients were used as using a smaller number of coefficients leads to overly degraded models in direct compression mode (compression of existing models). The DCT-based model realizations were then generated by randomly sampling uniform distributions defined by the empirical upper and lower bounds of the 250 maximum (in absolute value) coefficients associated with the training set. The DCT realizations were thresholded such that the facies proportions are close to that of the TI (facies 0: 0.7, facies 1: 0.3). Varying the number of DCT coefficients has no significant impact on generation quality. Clearly, none of these three approaches (PCA, OPCA, DCT) sample from a proper low-dimensional parameter space. In contrast, our proposed approach does a much better job. Note that PCA, OPCA and DCT are far from providing state-of-the-art MPS simulations. Instead, up to now they formed the state-of-the-art in terms



of parametric low-dimensional model representations that have been used with advanced inversion methods. The key of our method a low-dimensional parameterization that is able to provide high-quality MPS realizations (compare Figures 5 and 7).




Figure 5: Unconditional model realizations for the 2D binary channelized aquifer TI: (a - d) four (randomly chosen) model realizations from the DS-based training set and (e-l) eight (randomly chosen) model realizations generated with our DR approach that is based on a deep neural network.

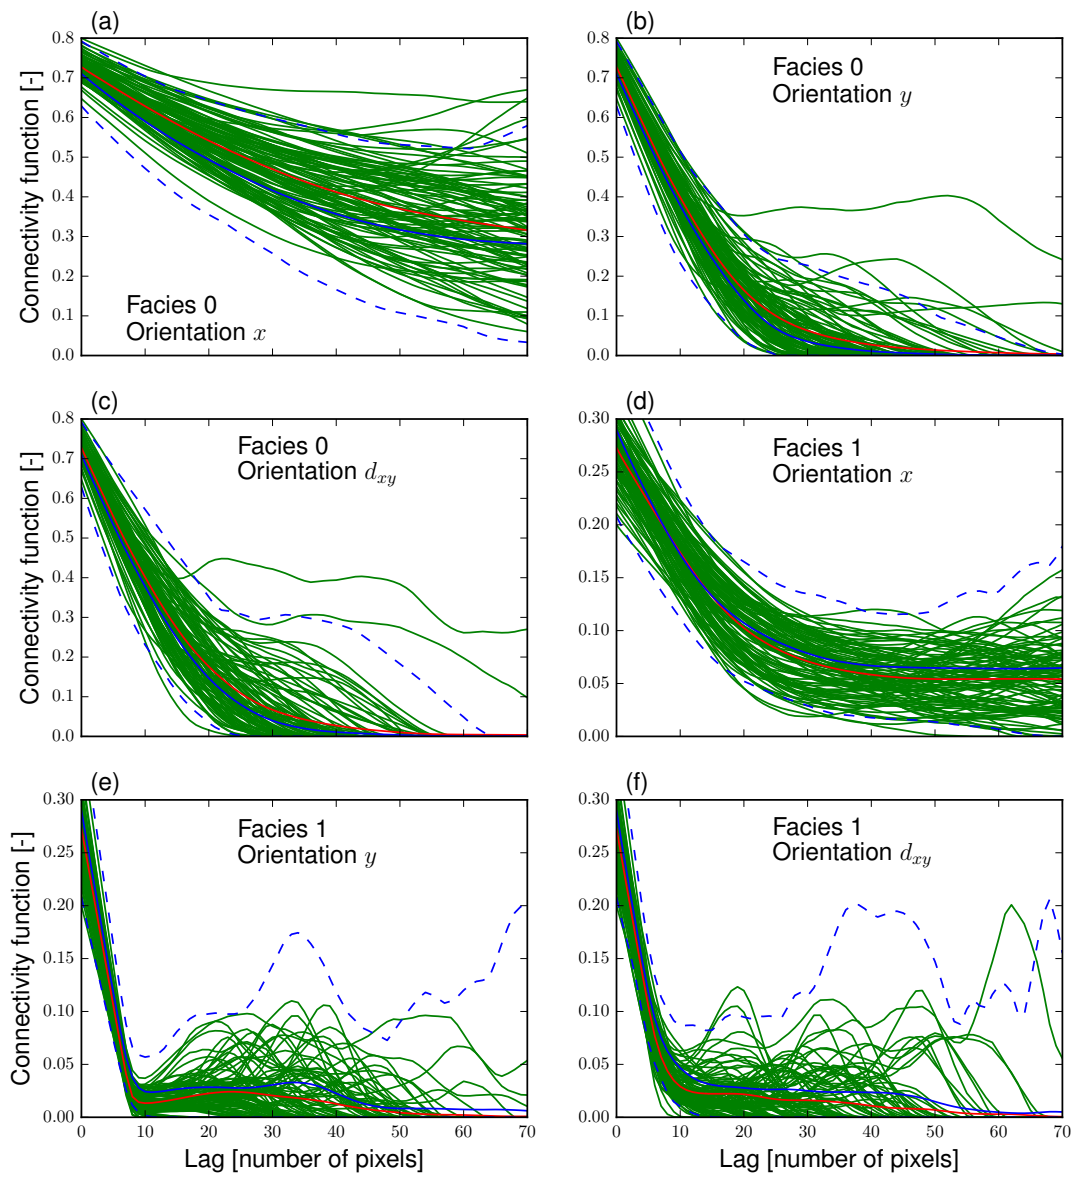


Figure 6: Cluster or connectivity function (CF) for the 2D case study involving a channelized aquifer without direct conditioning data (Figure 5). The blue lines denote the values associated with the training set. The solid blue line indicates the mean while the 2 dashed lines represent the minimum and maximum values at each lag. The green solid lines represent the 100 realizations generated by our DR approach. The red solid line is the mean of these DR-based realizations. The CF is calculated for each facies along directions. The  $x$  and  $y$  symbols signify the  $x$  and  $y$ -axes, and  $d_{xy}$  represents the diagonal direction formed by the  $45^\circ$  angle from the  $x$ -axis.




Figure 7: Unconditional model realizations for the 2D binary channelized aquifer TI: (a-d) four (randomly chosen) PCA-based model realizations and (e-h) corresponding OPCA realizations together with (i-l) four (randomly chosen) DCT-based model realizations.

Our second example considers a case where direct conditioning data are available. This example is similar to example 1 (Figure 5) except that facies at 9 locations are known (red circles in Figure 8). To handle direct conditioning, our approach requires that the MPS-generated realizations forming the training set honor the conditioning points. Building a training set of 40,000 conditioned models with DS took 7-8 hours using again eight CPUs in parallel. The training of our VAECNN then took 6.5 hours.

A set of (randomly chosen) training and DR-generated model realizations are shown in Figure 8. A similar generation performance is visually observed as for example 1

(Figure 5) and the DR-based realizations show again similar CF curves as the training images (Figure 9). However, the achieved space of uncertainty is 12% smaller than that of the conditioned training set:  $\frac{\overline{d_{JS(DR)}}}{\overline{d_{JS(TR)}}} = 0.88$ .

The conditioning accuracy is assessed by analyzing 1000 random model realizations. We find that 68% of the model realizations honor all of the 9 prescribed facies and that 97% of the model realizations contain at most one mismatching conditioning datum. Also, the more common facies 0 is more frequently honored (98%) than the less frequent facies 1 (91%). This is less than for our DS-generated training set that always honor the nine conditioning points. Nevertheless, the training realizations (Figures 8a-d) often contain 1 or 2 isolated conditioning point(s) within an homogeneous zone of the other facies. This type of artifact is less frequent in the DR-based model realizations (Figures 8e-l). The main reason for the conditioning errors is the sequential relooping used for producing a realization (see section 2.1.4). Without relooping, the model realizations present more frequently broken channels and/or small groups of isolated pixels (not shown), but the percentage of realizations that fully honor the 9 conditioning data is 92% (instead of 68% with relooping) and 100% (instead of 97% with relooping) of the realizations has at most a single erroneous datum.




Figure 8: Conditional model realizations for the 2D binary channelized aquifer TI: (a - d) four (randomly chosen) conditioned model realizations from the DS-based training set and (e-l) eight (randomly chosen) conditioned model realizations generated with our DR approach that is based on a deep neural network.

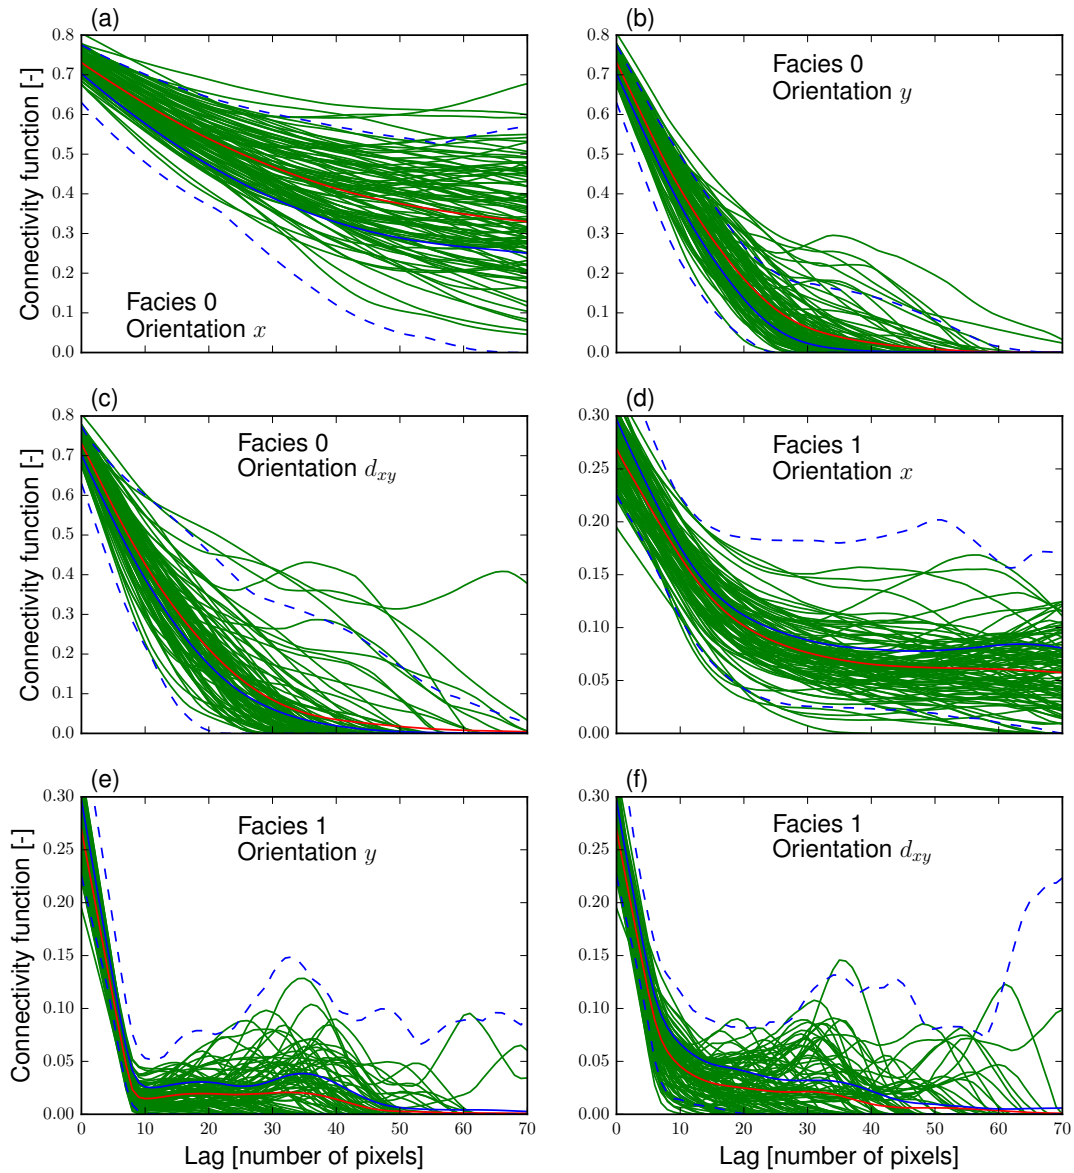


Figure 9: Cluster or connectivity function (CF) for the 2D case study involving a channelized aquifer and 9 direct conditioning data points (Figure 8). The blue lines denote the values associated with the training set. The solid blue line indicates the mean while the 2 dashed lines represent the minimum and maximum values at each lag. The green solid lines represent the 100 realizations generated by our DR approach. The red solid line is the mean of these DR-based realizations. The CF is calculated for each facies along directions. The  $x$  and  $y$  symbols signify the  $x$  and  $y$ -axes, and  $d_{xy}$  represents the diagonal direction formed by the  $45^\circ$  angle from the  $x$ -axis.

### 3.2 3D model

We now turn our attention to the generation of 3D model realizations. The selected TI is the Maules Creek valley alluvial aquifer available from <http://www.trainingimages.org/training-images-library.html>. For the unconditional case, a training set of 19,500  $30 \times 32 \times 27$  model realizations was built using the recent graph cuts (GC) patch-based MPS algorithm by Li et al. (2016). Running 8 different GC instances simultaneously, this took approximately 6 hours. The GC method was selected because it is much faster than DS. In addition, training our VAECNN took about 7 hours.

The convolutional layers in our VAECNN are designed for 2D images only, but with possibly different color channels. To generate 3D images, one can use as many channels as there are horizontal layers in the model (27). This does not explicitly account for

patterns in the vertical direction since the convolutional operations are performed in the horizontal plane only. Nevertheless, equation (8) penalizes reconstructed models that deviate from the training set in all three spatial directions.

Figure 10 compares (randomly chosen) training and DR-based realizations. As the model domain contains  $30 \times 32 \times 27 = 25,920$  voxels, the compression ratio is  $25920/50 \approx 518$ . The proposed DR approach creates realizations that are similar to those in the training set, although isolated voxels of each facies are over-represented. The realizations also reproduce well the CF metrics of the training set (Figure 11).

We also considered the same 3D example with 56 direct conditioning points. These points correspond to the locations of the screens along the vertical multi-level piezometers used in our synthetic 3D inverse problem described in section 4.2 (Blue line segments in Figure 12). A GC-based training set containing 19,000 conditioned model realizations was used for training our VAECNN. Running again 8 different GC instances simultaneously, building the training set took approximately 14 hours. Similarly as for the unconditional case, training then lasted for 7 hours. Using default algorithmic settings, these original GC simulations jointly honor all of the 56 conditioning data only 6% of the times but contain at most 10 mismatching point 94% of the times. The DR-based realizations are again visually close to the training set (not shown), with a much smaller tendency to overproduce isolated voxels of each facies than for the unconditional case (not shown) and CF statistics that closely match those of the training set (not shown). Furthermore, our DR approach is found to condition at a slightly higher level than the selected GC algorithm. Among 1000 random DR-based realizations, 7% of the realizations honor all of the 56 prescribed facies while 98% of the realizations contain a most 10 mismatching points. We attribute these slightly superior statistics compared to the training set to a combination of randomness and complex bias in the dimensionality reduction. Overall, our VAECNN is found to condition equally well as the GC algorithm used to produce the training set.




Figure 10: Unconditional model realizations for the 3D Maules Creek binary aquifer TI: (a - d) four (randomly chosen) model realizations from the MPS-based training set and (e-l) five (randomly chosen) model realizations generated with our DR approach that is based on a deep neural network.






Figure 11: Cluster or connectivity function (CF) for the 3D case study involving the Maules Creek binary aquifer without direct conditioning data (Figure 10). The blue lines denote the values associated with the training set. The solid blue line indicates the mean while the 2 dashed lines represent the minimum and maximum values at each lag. The green solid lines represent the 100 realizations generated by our DR approach. The red solid line is the mean of these DR-based realizations. The CF is calculated for each facies along directions. The  $x$ ,  $y$  and  $z$  symbols signify the  $x$ ,  $y$  and  $z$ -axes. The  $d_{xy}$ ,  $d_{yz}$  and  $d_{xz}$  symbols represent the diagonal direction formed by the  $45^\circ$  angle from the  $x$ -axis in the  $xy$  plane, from the  $y$ -axis in the  $yz$  plane and from the  $x$ -axis in the  $xz$  plane, respectively.




Figure 12: Pumping and observation locations in the synthetic 3D hydraulic tomography setup. The vertical gray lines represent the multilevel pumping and observation wells. The red line segments in the central well are the pumping locations and the blue line segments in the other wells are the head measurement locations.

## 4 Inverse problems

### 4.1 Case study 1: 2D steady-state flow

Our first inversion case study considers steady-state flow within a channelized aquifer. The  $100 \times 100$  aquifer domain lies in the  $x - y$  plane with a grid cell size of 1 m and a thickness of 1 m. Channel material and matrix material (see Figures 14a and 15a) are assigned hydraulic conductivity values of  $1 \times 10^{-2}$  m/s and  $1 \times 10^{-4}$  m/s, respectively. Steady state groundwater flow is simulated using MODFLOW 2005 (Harbaugh, 2005)

assuming no flow boundaries at the upper and lower sides and fixed head boundaries on the left and right sides of the domain to ensure a lateral head gradient of 0.01 (-), with water flowing in the  $x$ -direction. A pumping well extracting  $0.001 \text{ m}^3/\text{s}$  is located at the center of the domain. Simulated heads are collected at 49 locations that are regularly spread over the domain (Figures 14a and 15a). These data were then corrupted with a Gaussian white noise using a standard deviation of 0.02 m. For the selected white noise realization, the measurement data have a root-mean-square-error (RMSE) of 0.0200 m. The corresponding SNR, defined as the ratio of the average RMSE when drawing prior realizations with our DR algorithm to the noise level is 25. No direct conditioning (facies) data are used.

We first verify whether a model realization belonging to the 50-dimensional manifold can be retrieved by the inversion process. To do so, the reference model (Figure 14a) was randomly generated using our DR approach. The  $\text{DREAM}_{(\text{ZS})}$  was ran in parallel, using 8 interacting Markov chains distributed over 8 CPUs. Uniform priors in  $[-5, 5]$  were selected for the 50 DR variables,  $\boldsymbol{\theta}$ . This might seem to contradict the theory behind our VAECNN which is based on calibration to a standard normal distribution for  $p(\mathbf{z})$ . However, using a standard normal prior for  $\boldsymbol{\theta}$  was found to restrict the sampled model space too much for the inversion process to produce model realizations that in average fit the data to prescribed noise level. Our explanation for this phenomena is that with 50 DR variables the MCMC simulation needs sufficient freedom in the explored model space to eventually produce consistent models. Also, and equally important, is the observation that after training using equations (8-9), the quality of the generated realizations by our VAECNN is similar when using a zero-mean uniform or a standard normal distribution for  $p(\mathbf{z})$ . From a pragmatic point of view, the zero-mean uniform prior can thus be used.

The chains start to jointly sample the posterior distribution,  $p(\boldsymbol{\theta}|\mathbf{d})$ , after a (serial) total of 120,000 iterations, that is, 15,000 parallel iterations per chain (Figure 13a). The sampled realizations closely resemble the true model and the posterior variability is overall small (Figure 14). To assess sampling accuracy, we calculate over the last 10,000 posterior realizations in each chain the average fraction of pixels with facies identical to that of the true model,  $f_{\text{PO}}$ , which is 0.86. The same quantity for the prior,  $f_{\text{PR}}$ , can be directly calculated as  $0.7 \times 0.75 + 0.3 \times 0.25 = 0.60$  with 0.7 and 0.3 the prior fractions of each facies in the TI and 0.75 and 0.25 the fractions of each facies in the true model. The  $\frac{f_{\text{PO}}}{f_{\text{PR}}}$  ratio will be used to compare the examples considered in this section. Here  $\frac{f_{\text{PO}}}{f_{\text{PR}}} = \frac{0.86}{0.60} = 1.43$ . After a total of 200,000 MCMC iterations, the Gelman and Rubin (1992) convergence diagnostic,  $\hat{R}$ , is satisfied (i.e.,  $\hat{R} \leq 1.2$ ) for 37 out of the 50 sampled parameters (see, e.g., Laloy et al., 2015, for details about the use of  $\hat{R}$  with  $\text{DREAM}_{(\text{ZS})}$ ). The MCMC sampling should therefore be continued for a longer time for (official) full exploration of the posterior distribution.

Our second and more comprehensive test uses a true model that was generated by the DS algorithm. If the DR-representation is inappropriate, this true model might not be part of the 50-dimensional DR space. Identical inversion settings were used as in the previous example. Here the 8 chains start to jointly sample  $p(\boldsymbol{\theta}|\mathbf{d})$  after some 12,500 iterations per chain (Figure 13b). The posterior model realizations look visually close to the reference model, though posterior variability appears to be a bit larger than for the previous case (Figure 15). The associated  $f_{\text{PO}}$  and  $f_{\text{PR}}$  values are 0.78 and  $0.7 \times 0.71 + 0.3 \times 0.29 = 0.59$ , respectively, which gives a  $\frac{f_{\text{PO}}}{f_{\text{PR}}}$  ratio of 1.32. The results displayed in Figure 15 inspire confidence that the proposed DR-based inversion approach can retrieve consistent posterior models. With respect to posterior exploration, formal

convergence is yet to be declared after a total of 200,000 serial iterations, with the  $\hat{R}$  criterion being satisfied for 29 out of the 50 dimensions of  $\theta$ . Full posterior exploration therefore requires longer chains.

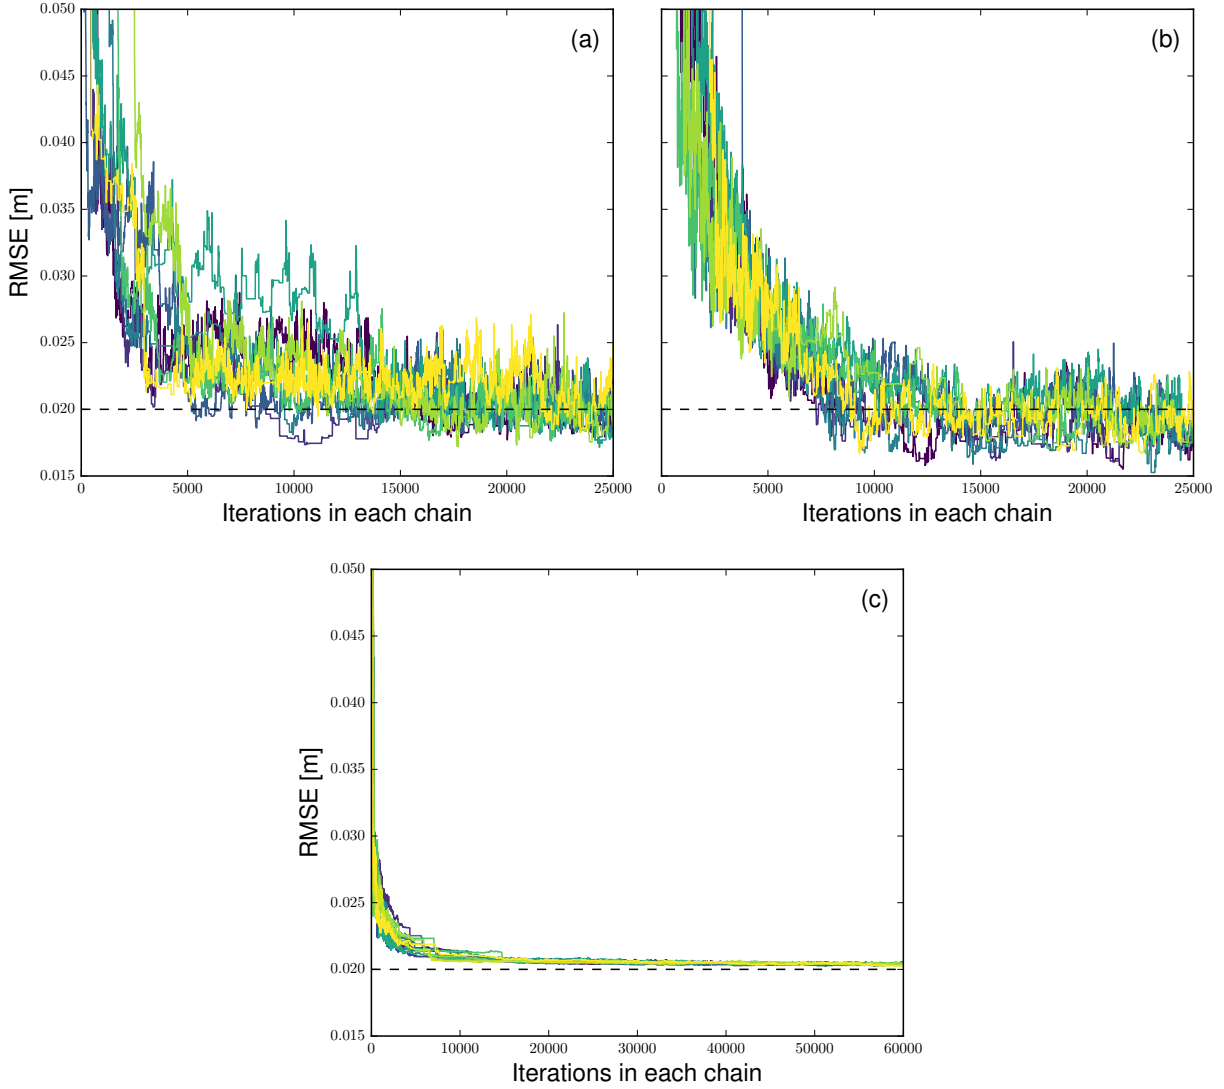


Figure 13: Trace plot of the sampled RMSE values by (a) the 8 Markov chains evolved by  $\text{DREAM}_{(\text{ZS})}$  (colored lines) for the first test (DR-generated true model) of inverse case study 1, (b) the 8 Markov chains evolved by  $\text{DREAM}_{(\text{ZS})}$  (colored lines) for the second test (DS-generated true model) of inverse case study 1 and (c) the 16 Markov chains evolved by  $\text{DREAM}_{(\text{ZS})}$  (colored lines) for the inverse case study 2. The dashed black line denotes the true RMSE value.




Figure 14: (a) True model and (b-i) states of the 8 Markov chain evolved by  $DREAM_{(ZS)}$  after 25,000 iterations per chain for the first test (DR-generated true model) of inverse case study 1. The red circle and blue crosses in subfigure (a) mark the location of the pumping well and the piezometers, respectively.




Figure 15: (a) True model and (b-i) states of the 8 Markov chains evolved by  $\text{DREAM}_{(zS)}$  after 25,000 iterations per chain for the second test (MPS-generated true model) of inverse case study 1. The red circle and blue crosses in subfigure (a) mark the location of the pumping well and the piezometers, respectively.

To benchmark our approach against state-of-the-art MPS inversion, we performed eight independent SGR runs in parallel for the same case study as above and a similar CPU budget of 25,000 MCMC iterations per run. Moreover, the SGR algorithmic settings are optimal (for the considered type of application) as defined in Laloy et al. (2016). It is observed that 5 SGR trials out of 8 fit the data to the appropriate noise level after less than 10,000 iterations (not shown). Three trials even reach the targeted 0.02 m misfit after less than 3000 iterations (not shown). Yet three other SGR trials do not sample the posterior distribution after the 25,000 iterations (trials #3, #4 and #8 in Figure 16). Furthermore, all SGR trials end up in a relatively narrow local optimum (not shown) while none of them find a model that is visually similar to the true model (Figure 16). Many of the realizations sampled by SGR are also degraded compared to the prior model (see top row of Figure 5), with isolated patches and broken channels (Figure 16). This translates into a quite lower  $\frac{f_{\text{PO}}}{f_{\text{PR}}}$  ratio than for our DR-based approach: 1.08 (the average  $f_{\text{PO}}$  over the 8 SGR trials is 0.64). We thus conclude that for the considered example, our DR approach is a superior alternative to SGR.




Figure 16: (a) True model and (b-i) states of the 8 independent SGR trials after 25,000 iterations per trial for the second test (MPS-generated true model) of inverse case study 1. The red circle and blue crosses in subfigure (a) mark the location of the pumping well and the piezometers, respectively.



## 4.2 Case study 2: 3D transient hydraulic tomography

Our last case study focuses on 3D transient hydraulic tomography (e.g., Cardiff et al., 2013). Transient head variations caused by discrete multilevel pumping tests were simulated using MODFLOW2005 within a 3D confined aquifer of size  $30 \times 32 \times 27$  with a voxel size of  $1 \text{ m} \times 1 \text{ m} \times 1 \text{ m}$ . This aquifer consists of a part of the Maules Creek TI discussed in section 3.2. The true model is depicted in Figure 17a. The multilevel discrete pumping setup consists of a 27-m deep, central multi-level well in which water is sequentially extracted every 4 m along a 1-m long screen (at depths of 3, 7, 11, 15, 19, 23 and 27 m respectively) during 30 minutes at a rate of 20 liters/min. The locations of this multilevel pumping well and those of the 8 surrounding multilevel observation wells are displayed in Figure 12. For each pumping sequence, drawdown are recorded every 4 m along a 1-m long screen in the 8 multi-level piezometers (Figure 12). For each drawdown curve, data acquired at the same four measurement times were retained leading to a total of  $8 \times 7 \times 7 \times 4 = 1568$  measurement data. These measurement times were considered to be the four most informative ones after visual inspection of several drawdown curves (not shown). These data were corrupted with a Gaussian white noise using again a standard deviation of 0.02 m, which induced for the selected white noise realization a RMSE of 0.0200 m. The associated SNR is 6. In addition, the facies of the 56 locations (voxels) where pressure head is measured (blue line segments in Figure 12) are known and used as direct conditioning data in the inversion.

For this case study, the DREAM<sub>(ZS)</sub> sampler evolves 16 Markov chains in parallel using 16 CPUs. Uniform priors in  $[-5, 5]$  were again selected for the 50 dimensions of  $\theta$ . Within the allowed computational budget of 60,000 iterations per chain, the 16 chains converge towards a data misfit in the range of 0.0202 m - 0.0204 m (Figure 13c). This interval is close to the target level of 0.0200 m. This indicates that the posterior mode has not been sampled yet. The variability among the sampled models is rather small (Figure 17). The latter is likely due to the combination of two factors: (1) the peakedness of the likelihood function (equation (15)) caused by the large number of measurement data, and (2) the difficulties encountered by the sampling algorithm for exploring this complex target distribution. Notwithstanding, the sampled models (Figure 17b-i) are visually close to the true model (Figure 17a).




Figure 17: a) True model and (b-i) states of the first 8 (out of 16) Markov chains evolved by  $\text{DREAM}_{(\text{ZS})}$  after 60,000 iterations per chain for the inverse case study 2.

## 5 Discussion

Our proposed DR approach that is based on a deep neural network works well for the considered inversion of 2D steady-state flow data. It also provides useful solutions to our inverse problem involving 3D transient flow. However, further improvements are needed. As for most compression approaches, our dimensionality reduction scheme builds a low-dimensional representation where the value in each dimension influences every pixel/voxel

of the generated realization in the original space. This “non-localization” poses problems for inverse problems involving categorical fields. Indeed, Laloy et al. (2016) have shown for SGR that iterative perturbation of a small well defined fraction of the current model works much better than slightly perturbing the whole model at every step. The same logic applies to compression-based inversion. Indeed, Figure 18 demonstrates for our first inversion case study that (1) changing the value of a single low-dimensional parameter influences the whole model realization domain and (2) even a small variation in the model realization can cause a large upward shift in RMSE (e.g., compare Figures 18a and 18c). This implies that our proposed inversion approach can have troubles to converge when applied to rich datasets with low error (e.g.,  $\leq 0.01$  m for measured piezometric heads). We expect that a yet to be developed “localized” DR approach would perform better.

There are other aspects that would benefit from improvements. First, our DR approach is not totally accurate for direct conditioning. Second, several tens of thousands of MPS-generated training models are used to build it. This might not be practical when MPS-generation is computationally intensive. Third, 3D model generation capabilities could be enhanced by considering 3D convolutional layers instead of tricking 2D convolutional layers as done herein. Fourth, using a different reconstruction loss function than equation (8) the approach could be extended to categorical TIs with more than two facies and to continuous TIs. Lastly, it may be considered as a disadvantage that our approach produces continuous realizations and thus requires thresholding to create categorical fields.

Because of computational constraints, the impact of using a different number of dimensions for  $\mathbf{z}$  on geostatistical simulation quality was not studied extensively. Yet limited testing revealed that for the considered 2D channelized TI and associated model domain size using a 25-dimensional  $\mathbf{z}$  does not produce sufficient variability in the realizations (not shown). In contrast, compared to a 50-dimensional  $\mathbf{z}$  using 100 dimensions induces fewer broken channels and thus provides slightly better results (not shown). For the considered 3D Maules Creek aquifer TI, using a 25-dimensional  $\mathbf{z}$  was found to perform equally well as a 50-dimensional  $\mathbf{z}$  (not shown). These findings warrant further investigations.

We would like to stress that if no hard thresholding is done at the end of our model generation process, then the derivatives of the pixels/voxels of the generated model,  $\mathbf{X}$ , with respect to the elements of the associated low-dimensional code,  $\mathbf{z}$ , have a (complicated) analytical solution that can be calculated by THEANO using auto-differentiation. In other words, the sensitivity (Jacobian) matrix  $\frac{d\mathbf{X}}{d\mathbf{z}}$  is immediately available. This may prove useful for gradient-based inversion using a forward solver equipped with an adjoint model.

In addition, it is worth noting that using a deep generative model to compress high dimensional quantities into a low-dimensional standard normal (or uniform) base is not only useful for inversion purposes but also holds promise for the so-called prediction-focused approach (e.g., Satija et al., 2017). The latter seeks to build a statistical model that links past data variables and prediction data variables, thereby enabling to make predictions without inverting for material properties. Lastly, note that our approach is not limited to the MPS framework but could potentially be applied to any prior with discrete structure.




Figure 18: Model realizations and associated RMSE obtained when modifying a single low-dimensional dimension,  $DR_1$ , of a posterior solution for case study 1: (a) reference (posterior) realization, (b) realization and RMSE corresponding to the multiplication of  $DR_1$  by 0.5, (c) realization and RMSE corresponding to the multiplication of  $DR_1$  by 0.1 and (d) realization and RMSE corresponding to the multiplication of  $DR_1$  by -1.

## 6 Conclusion

This paper presents a novel parametric low-dimensional representation of complex binary geologic media that can be used for fast sampling of complex geostatistical models, which is necessary for efficient probabilistic inversion. A deep neural network of the variational autoencoder type is used to define a low-dimensional manifold that is restricted to model realizations that agree well with the training set. By comparing other parametric dimensionality reduction (DR) techniques for unconditional geostatistical simulation of a channelized prior model, we find that our representation is superior to the PCA, optimization-PCA (OPCA) and discrete cosine transform (DCT). Since constructing the parameterization requires a training set of several tens of thousands of model realizations obtained by multiple point statistics simulation, our DR approach is specifically designed for probabilistic (or deterministic) inversion rather than for unconditional (or point-conditioned) geostatistical simulation. Synthetic inversion case studies involving

2D steady flow and 3D transient hydraulic tomography are used to illustrate the effectiveness of the proposed DR-based probabilistic inversion approach. For the 2D case study, our new approach provides better results than current state-of-the-art inversion by sequential geostatistical resampling (SGR). Inversion results for the 3D application are also encouraging. Future work will focus on improvements such as alleviating the need of using a large training model set to build the parameterization, extension to multi-categorical and continuous variables and defining a low-dimensional representation that is local rather than global. We expect that a local representation in which a given low-dimensional variable only influences a sub-region of the generated model realization will provide a much improved inversion performance.

### Acknowledgments

Python codes of the proposed dimensionality reduction (DR) and DR-based inversion approaches are available from the first author (and can be downloaded from <https://github.com/elaloy>). This work was partially supported by Agence Nationale de la Recherche through the grant ANR-16-CE23-0006 Deep in France. The third author is a Research Associate with the Belgian F.R.S.-F.N.R.S. We thank Mike Swarbrick Jones for sharing his codes (<https://github.com/mikesj-public>), one of which served as a starting point for our devised VAECNN. We are grateful to Hai Xuan Vo and Louis Durlofsky for sharing their OPCA code, and to Xue Li and her co-authors for sharing their GC MPS algorithm. A temporary academic license of the DeeSse MPS code can be obtained upon request to one of its developers (Grégoire Mariethoz, Philippe Renard, Julien Straubhaar). We thank Laurent Lemmens for sharing his MPH and CF calculation routines. Finally, we would like to thank the two anonymous referees for their useful comments.

### References

- Awotunde AA, Horne RN. Reservoir description with integrated multiwell data using two-dimensional wavelets. *Math. Geosci.* 2013;45(2);225–252.
- Brooks S, Gelman A, Jones GL, Men XL, eds. *Handbook of Markov Chain Monte Carlo*. CRC Press, New York; 2011.
- Cardiff M, Barrash W, Kitanidis PK. Hydraulic conductivity imaging from 3-D transient hydraulic tomography at several pumping/observation densities. *Water Resour. Res.* 2013;49;7311–7326. <http://dx.doi.org/10.1002/wrcr.20519>.
- Cover TM, Thomas JA. *Elements of information theory*. 1991. Wiley-Interscience, New York, USA.
- Cybenko G. Approximations by superpositions of sigmoidal functions. *Mathematics of Control, Signals, and Systems*. 1989;2(4);303–314.
- Dieleman S, Schlüter J, Raffel C, Olson E, Sønderby SK, Nouri, D, Maturana D, Thoma M, Battenberg E, Kelly J, De Fauw J, Heilman M, Moitinho de Almeida D, McFee B, Weideman H, Takács G, de Rivaz P, Crall J, Sanders G, Kashif R, Liu C, French G, Degraeve J. LASAGNE: First release. August 2015. <http://dx.doi.org/10.5281/zenodo.27878>.
- Dorn O, Villegas R. History matching of petroleum reservoirs using a level set technique. *Inverse Prob.* 2008;24(3);035,015.

- Gelman AG, Rubin DN. Inference from iterative simulation using multiple sequences. *Stat. Sci.* 1992;7:457–472.
- Gómez-Hernández JJ, Wen X-H. To be or not to be multi-Gaussian? A reflection on stochastic hydrogeology. *Adv. Water Resour.* 1998;21(1):47–61.
- Goodfellow I, Bengio Y, Courville A. *Deep learning*. MIT Press; 2016. <http://www.deeplearningbook.org>.
- Gregor K, Danihelka I, Graves A, Wierstra D. DRAW: A recurrent neural network for image generation. *The International Conference on Machine Learning (ICML)*, Lille; 2015.
- Guardiano F, Srivastava M. Multivariate geostatistics: beyond bivariate moments, in Soares, A (ed.), *Geostatistics-Troia*, Kluwer Academic; 1993.
- Hansen TM, Mosegaard K, Cordua KC. Using geostatistics to describe complex a priori information for inverse problems. In: Ortiz, JM, Emery, X (eds.) *VIII International Geostatistics Congress*. 2008;vol. 1:329–338. Mining Engineering Department, University of Chile, Santiago.
- Hansen TM, Cordua KC, Mosegaard K. Inverse problems with non-trivial priors: efficient solution through sequential Gibbs sampling. *Comput. Geosci.* 2012;16:593–611. <http://dx.doi.org/10.1007/s10596-011-9271-1>.
- Harbaugh AW. MODFLOW-2005, the U.S. Geological Survey modular groundwater model - the groundwater flow process. U.S. Geological Survey, *Techniques and Methods*, 6-A16. 2005.
- Hornik K. Approximation capabilities of multilayer feedforward networks. *Neural Networks.* 1991;4(2);251–257.
- Hu L, Chugunova T. Multiple-point geostatistics for modeling subsurface heterogeneity: a comprehensive review. *Water Resour. Res.* 2008;44;W11413. <http://dx.doi.org/10.1029/2008WR006993>.
- Jafarpour B, Goyal V, McLaughlin DB, Freeman WT. Compressed history matching: exploiting transform-domain sparsity for regularization of nonlinear dynamic data integration problems. *Math. Geosci.* 2010;42(1);1–27.
- Journel A, Zhang T. The necessity of a multiple-point prior model. *Math. Geol.* 2006;38;591.
- Khaninezhad MM, Jafarpour B. Sparse randomized maximum likelihood (SpRML) for subsurface flow model calibration and uncertainty quantification. *Adv. Water Resour.* 2014;69:23–37.
- Khaninezhad MM, Jafarpour B, Li L. Sparse geologic dictionaries for subsurface flow model calibration: part I. Inverse formulation. *Adv. Water Resour.* 2012;39:106–121.
- Kingma DP, Welling M. Auto-encoding variational Bayes. *The International Conference on Learning Representations (ICLR)*, Banff; 2014.
- Kingma, DP, Ba JL. Adam: a Method for Stochastic Optimization. *The International Conference on Learning Representations (ICLR)*, San Diego; 2015.

- Kitanidis, P. Quasi-linear geostatistical theory for inversing. *Water Resour. Res.* 1995;31(10);2411–2419.
- Krizhevsky, A, Sutskever I, Hinton, GE. ImageNet classification with deep convolutional neural networks. 2012. In Pereira F, Burges CJC, Bottou L, Weinberger KQ (eds.) *Advances in Neural Information Processing Systems 25 (NIPS 2012)*, Curran Associates, Inc, pp. 1097–1105.
- Laloy E, Vrugt JA. High-dimensional posterior exploration of hydrologic models using multiple-try DREAM<sub>(ZS)</sub> and high-performance computing. *Water Resour. Res.* 2012;48(1). <http://dx.doi.org/10.1029/2011WR010608>.
- Laloy, E, Linde N, Vrugt JA. Mass conservative three-dimensional water tracer distribution from Markov chain Monte Carlo inversion of time-lapse ground-penetrating radar data. *Water Resour. Res.* 2012;48:W07510. <http://dx.doi.org/10.1029/2011WR011238>.
- Laloy, E, Rogiers B, Vrugt JA, Mallants D, Jacques D. Efficient posterior exploration of a high-dimensional groundwater model from two-stage Markov Chain Monte Carlo simulation and polynomial chaos expansion. *Water Resour. Res.* 2013;49:2664–2682. <http://dx.doi.org/10.1002/wrcr.20226>.
- Laloy E, Linde N, Jacques D, Vrugt JA. Probabilistic inference of multi-Gaussian fields from indirect hydrological data using circulant embedding and dimensionality reduction. *Water Resour. Res.* 2015;51. <http://dx.doi.org/10.1002/2014WR016395>.
- Laloy, E, Linde N, Jacques D, Mariethoz G. Merging parallel tempering with sequential geostatistical resampling for improved posterior exploration of high-dimensional subsurface categorical fields. *Adv. Water Resour.* 2016;90:57–69.
- LeCun, Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE.* 1998;86(11):2278–2324.
- Lemmens, L, Rogiers B, De Craen M, Laloy E, Jacques D, Huysmans M, Swennen R, Urai J. L., Desbois G. Effective structural descriptors for natural and engineered radioactive waste confinement barrier, *EGU 2017 - poster 14094*.
- Li X, Mariethoz G, Lu DT, Linde N. Patch-based iterative conditional geostatistical simulation using graph cuts. *Water Resour. Res.* 2016;52:6297–6320. <http://doi:10.1002/2015WR018378>.
- Linde, N, Vrugt JA. Distributed soil moisture from crosshole ground-penetrating radar travel times using stochastic inversion. *Vadose Zone J.* 2013;12(1). <http://dx.doi.org/doi:10.2136/vzj2012.0101>.
- Lochbühler T, Vrugt JA, Sadegh M, Linde N. Summary statistics from training images as prior information in probabilistic inversion. *Geophys. J. Int.* 2015;201:157–171. <http://dx.doi:10.1093/gji/ggv008>.
- Mariethoz G, Renard P, Caers J. Bayesian inverse problem and optimization with iterative spatial resampling. *Water Resour. Res.* 2010a; 46:W11530. <http://dx.doi.org/10.1029/2010WR009274>.
- Mariethoz G, Renard P, Straubhaar J. The Direct Sampling method to perform multiple-point geostatistical simulations. *Water Resour. Res.* 2010b;46:W11536. <http://dx.doi.org/10.1029/2008WR007621>.

- Oliver, DS, He N, Reynolds AC. Conditioning permeability fields to pressure data. Paper presented at the 5th European Conference on the Mathematics of Oil Recovery, Leoben, Austria. 1996.
- Otsu N. A threshold selection method from gray-Level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*. 1979;9(1): 62–66.
- Pardo-Igúzquiza E, Dowd P. CONNEC3D: A computer program for connectivity analysis of 3D random set models. *Comput. Geosci*. 2003;29;775–785.
- Robert, CP, Casella G. Monte Carlo statistical methods, second edition. Springer; 2004.
- Reynolds AC, He N, Chu L, Oliver DS. Reparameterization techniques for generating reservoir description conditioned to variograms and well-test pressure data. *SPE J*. 1996;1(4);413–426.
- Sarma P, Durlofsky LJ, Aziz K, Chen WH. Efficient real-time reservoir management using adjoint-based optimal control and model updating. *Comput. Geosci*. 2006;10(1);3–36.
- Sarma P, Durlofsky LJ, Aziz K. Kernel principal component analysis for efficient, differentiable parameterization of multipoint geostatistics. *Math. Geosci*. 2008;40(1);3–32.
- Satiya A, Scheidt C, Li L, Caers, J. Direct forecasting of reservoir performance using production data without history matching. *Computat. Geosci*. 2017;21(2);315–333.
- Stien M, Kolbjørnsen O. Facies modeling using a Markov mesh model specification. *Math. Geosci*. 2011;43:611–624.
- Strebelle S. Conditional simulation of complex geological structures using multiple point statistics. *Math. Geol*. 2002;34(1):1–22.
- Tahmasebi P, Ardeshtir H, Sahimi M. Multiple-point geostatistical modeling based on the cross-correlation functions. *Comput. Geosci*. 2012;16(3):779–797. <http://dx.doi.org/10.1007/s10596-012-9287-1>.
- Tan X, Tahmasebi P, Caers J. Comparing training-image based algorithms using an analysis of distance. *Math. Geosci*. 2014;46;149–169.
- Tavakoli R, Reynolds AC. Monte Carlo simulation of permeability fields and reservoir performance predictions with SVD parameterization in RML compared with EnKF. *Comput. Geosci*. 2011;15(1);99–116.
- Theano Development Team. THEANO: A Python framework for fast computation of mathematical expressions. 2016;<http://arxiv.org/abs/1605.02688>.
- Torquato S, Beasley JD, Chiew YC. Two-point cluster function for continuum percolation. *J. Chem. Phys*. 1988;88(10);6540–6547.
- Vo HX, Durlofsky LJ. A new differentiable parameterization based on Principal Component Analysis for the low-dimensional representation of complex geological models. *Math. Geosci*. 2014;46;775–813.
- Vo HX, Durlofsky LJ. Data assimilation and uncertainty assessment for complex geological models using a new PCA-based parameterization. *Comput. Geosci*. 2015;19:747–767.



Vrugt JA, Ter Braak, C, Diks C, Robinson BA, Hyman, JM, Higdon, D. Accelerating Markov chain Monte Carlo simulation by differential evolution with self-adaptive randomized subspace sampling. *Int. J. Nonlin. Sci. Numer. Simul.* 2009;10(3):273–290.

Zahner T, Lochbühler T, Mariethoz G, Linde N. Image synthesis with graph cuts: a fast model proposal mechanism in probabilistic inversion. *Geophys. J. Int.* 2016;204(2):1179–1190. <http://dx.doi.org/10.1093/gji/ggv517>.