



HAL
open science

Gradient-based deterministic inversion of geophysical data with generative adversarial networks: Is it feasible?

Eric Laloy, Niklas Linde, Cyprien Ruffino, Romain Hérault, Gilles Gasso,
Diederik Jacques

► To cite this version:

Eric Laloy, Niklas Linde, Cyprien Ruffino, Romain Hérault, Gilles Gasso, et al.. Gradient-based deterministic inversion of geophysical data with generative adversarial networks: Is it feasible?. *Computers & Geosciences*, 2019, 133, pp.104333. 10.1016/j.cageo.2019.104333 . hal-02334880

HAL Id: hal-02334880

<https://normandie-univ.hal.science/hal-02334880>

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.

Copyright

Gradient-based deterministic inversion of geophysical data with Generative Adversarial Networks: is it feasible?

Eric Laloy,^{*} Niklas Linde,[†] Cyprien Ruffino,[‡] Romain Hérault,[‡] Gilles Gasso,[‡]
and Diederik Jacques^{*}

May 24, 2019

Abstract

Global probabilistic inversion within the latent space learned by a Generative Adversarial Network (GAN) has been recently demonstrated. Compared to inversion on the original model space, using the latent space of a trained GAN can offer the following benefits: (1) the generated model proposals are geostatistically consistent with the prescribed prior training image (TI), and (2) the parameter space is reduced by orders of magnitude compared to the original model space. Nevertheless, exploring the learned latent space by state-of-the-art Markov chain Monte Carlo (MCMC) methods may still require a large computational effort. As an alternative, parameters in this latent space could possibly be optimized with much less computationally expensive gradient-based methods. This study shows that due to the typically highly nonlinear relationship between the latent space and the associated output space of a GAN, gradient-based deterministic inversion may fail even when considering a linear forward physical model. We tested two deterministic inversion approaches: a quasi-Newton gradient descent using the Adam algorithm and a Gauss-Newton (GN) method that makes use of the Jacobian matrix calculated by finite-differencing. For a channelized binary TI and a synthetic linear crosshole ground penetrating radar (GPR) tomography problem involving 576 measurements with low noise, we observe that when allowing for a total of 10,000 iterations only 13% of the gradient descent trials locate a solution that has the required data misfit. The tested GN inversion was unable to recover a solution with the appropriate data misfit. Our results suggest that deterministic inversion performance strongly depends on the inversion approach, starting model, true reference model, number of iterations and noise realization. In contrast, computationally-expensive probabilistic global optimization based on differential evolution always finds an appropriate solution.

1 Introduction and Scope

Laloy et al. (2018) recently proposed to use generative adversarial networks (GANs), a game changer data generation algorithm (e.g., Goodfellow et al., 2014, 2016), to define a low-dimensional parameterization encoding complex geologic prior models, thereby allowing efficient and accurate geostatistical inversion with Markov chain Monte Carlo (MCMC) methods (Laloy et al., 2018). GANs have permitted impressive advancements for a wide range of applications such as image and texture synthesis, image-to-image translation and super-resolution (Creswell et al., 2017). In the near future, we expect

^{*}Belgian Nuclear Research Centre, Email: elaloy@sckcen.be

[†]University of Lausanne

[‡]INSA-Rouen

to witness a dramatic increase in development and use of GAN-inspired algorithms for geostatistical simulation (e.g., Mosser et al., 2017; Laloy et al., 2018) and inversion (e.g., Laloy et al., 2018; Mosser et al., 2018; Richardson, 2018). A key element of the GAN approach is that the dimensions of the learned low-dimensional or “latent” space are independent from each other and have a known probability distribution, typically either a uniform or a standard Gaussian distribution. Laloy et al. (2018) have shown that inversions based on such parameterizations work well for global probabilistic inference of complex binary 2D and 3D prior subsurface models. However, exploring the GAN-derived latent space with state-of-the-art MCMC sampling (Vrugt et al., 2009; Laloy and Vrugt, 2012) still necessitates tens of thousands (or more) forward evaluations (Laloy et al., 2018). Such a computational expense can be prohibitive when using computationally-demanding forward solvers encountered in the geosciences. However, the latent space might also lend itself to a conventional deterministic gradient-based local search, which is often much more computationally frugal than global and probabilistic inversion.

In the geosciences, gradient-based deterministic inversion methods can be roughly divided into (1) methods that make use of the so-called sensitivity or Jacobian matrix, \mathbf{J} , that is, the matrix of derivatives of the forward model outputs, $F(\mathbf{m})$, with respect to the model parameters, \mathbf{m} (e.g., physical properties of the model grid blocks): $J_{i,j} = \frac{\delta F(\mathbf{m})_i}{\delta m_j}$, and (2) methods that only require the gradient vector of the misfit (that is, objective or loss) function, \mathcal{L} , with respect to \mathbf{m} , $\nabla\mathcal{L}$ with $\nabla\mathcal{L}_j = \frac{\delta\mathcal{L}}{\delta m_j}$, such as the steepest descent method. Although $\nabla\mathcal{L}$ is often cheaper to calculate than \mathbf{J} , it is generally understood that methods relying on \mathbf{J} are more effective and robust than those based on $\nabla\mathcal{L}$ because \mathbf{J} can be used to (implicitly) approximate the Hessian matrix, \mathbf{H} with $H_{i,j} = \frac{\delta J_{i,j}}{\delta m_j}$, which contains information about the local curvature of the misfit function (Aster et al., 2012).

Complex spatial prior models are usually represented by a so-called training image (TI). A 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) and is typically used to guide geostatistical simulation by multiple-point statistics (MPS) algorithms (e.g, Strebelle, 2002; Mariethoz et al., 2010). When using a GAN (see, e.g., Goodfellow et al., 2016, for details about the GAN architecture and training) to encode the prescribed TI (or, alternatively, MPS realizations from it), a GAN-based model realization is produced by feeding the so-called generator, G , with a latent vector, \mathbf{z} : $\mathbf{m} = G(\mathbf{z})$, and the inversion is performed within the latent space $p(\mathbf{z})$. For gradient-based inversion, the required Jacobian is then the matrix of derivatives of $F[G(\mathbf{z})]$ with respect to \mathbf{z} : \mathbf{J}^z with $J_{i,j}^z = \frac{\delta F[G(\mathbf{z})]_i}{\delta z_j}$. In principle, the derivatives of $G(\mathbf{z})$ with respect to \mathbf{z} , \mathbf{V} with $V_{i,j} = \frac{\delta G(\mathbf{z})_i}{\delta z_j}$, can be directly computed by autodifferentiation (Al-Rfou et al., 2016; Abadi et al., 2016; Paszke et al., 2017). When \mathbf{J} is available, for instance through the use of an adjoint model, \mathbf{J}^z can be obtained as $\mathbf{J}^z = \mathbf{J}\mathbf{V}$. In our experience, calculating each element of \mathbf{V} using the reverse-mode autodifferentiation engine that equips current deep learning (DL) libraries such as TensorFlow (Abadi et al., 2016) and PyTorch (Paszke et al., 2017) can be slow, especially for large \mathbf{m} and when run on a CPU. For instance, for the PyTorch implementation of the spatial GAN (Jetchev et al., 2016) used herein it takes about 5 minutes to construct the \mathbf{V} array on a last generation Intel[®] i7 CPU and about 1 minute using a NVIDIA Quadro M2000M GPU, when \mathbf{z} is

15-dimensional and $G(\mathbf{z})$ is of size 129×65 . More importantly, generating high-quality categorical geologic structures with a GAN may require postprocessing of the realizations through filtering and/or thresholding (e.g., Mosser et al., 2017; Laloy et al., 2018). If not differentiable, this postprocessing, $pp(\cdot)$, decouples the \mathbf{V} array computed by the DL code from the actual realization, $pp[G(\mathbf{z})]$.

Another solution that alleviates the need of having an adjoint model and is not perturbed by postprocessing operations consists of estimating \mathbf{J}^z directly using a finite-difference approximation. This incurs a relatively low computational cost since it requires $2N_z + 1$ forward runs to evaluate \mathbf{J}^z with a central difference scheme, with N_z the length of the \mathbf{z} vector, which by construction is low-dimensional (say two orders of magnitude less than \mathbf{m}).

Lastly, when using the gradient vector, $\nabla\mathcal{L}$ with respect to \mathbf{z} , $\nabla\mathcal{L}(\mathbf{z})$, instead of the Jacobian, \mathbf{J}^z , $\nabla\mathcal{L}(\mathbf{z})$ and the trained $G(\mathbf{z})$ can be integrated within a single DL computational graph (Richardson, 2018; Mosser et al., 2018). This is possible because $G(\mathbf{z})$ is fully differentiable end-to-end and DL libraries are equipped with advanced gradient descent algorithms such as Adam (Kingma and Ba, 2015). If $F[G(\mathbf{z})]$ is fully differentiable then the DL library can compute $\nabla\mathcal{L}(\mathbf{z})$ internally by autodifferentiation. This is for instance the case when $F(\cdot)$ is linear and the continuous $G(\mathbf{z})$ realizations do not need postprocessing such as thresholding. As shown by Richardson (2018), when $pp[G(\mathbf{z})]$ or $F(\cdot)$ is not differentiable $\nabla\mathcal{L}(\mathbf{m})$ can however be externally computed and integrated within the DL graph. In this work, we used an adjoint model giving $\nabla\mathcal{L}(\mathbf{m})$ and backpropagated $\nabla\mathcal{L}(\mathbf{z}) = \nabla\mathcal{L}(\mathbf{m})\mathbf{V}$ where the $\nabla\mathcal{L}(\mathbf{m})\mathbf{V}$ product can be efficiently computed within a single backward pass without explicitly calculating the \mathbf{V} array (see our companion code).

No matter whether $\nabla\mathcal{L}(\mathbf{z})$ or \mathbf{J}^z is used or regardless of how they are calculated, an important open question is to what extent the nonlinearity in the $G(\mathbf{z})$ transform adversely affects the convergence of gradient-based deterministic inversion performed in the latent space $p(\mathbf{z})$. Working with $p(\mathbf{z})$ rather than the original model space $p(\mathbf{m})$ conveniently ensures that (1) any generated model honors the prior TI and (2) the parameter dimensionality is reduced by orders of magnitude compared to traditional deterministic inversion (e.g., de Groot-Hedlin and Constable, 1990). Nevertheless, the non-linear relationship between $G(\mathbf{z})$ and \mathbf{z} adds significant nonlinearities to the inverse problem, in addition to those of the forward model, $F(\mathbf{m})$. In this work, we consider a best case scenario and use a linear tomography problem based on ground penetrating radar (GPR) data to show that even if $F(\mathbf{m})$ is linear with respect to \mathbf{m} , the nonlinearity in $G(\mathbf{z})$ may frequently affect the performance of gradient-based deterministic inversions. In contrast, probabilistic global optimization using the DREAM_(ZS) code (Vrugt et al., 2009; Laloy and Vrugt, 2012) is found to work well for the considered problem, although at a rather high computational cost. The overall motivation of this work is to raise awareness that although model parameterizations based on GAN algorithms can be very powerful in representing prior information and reducing model dimensionality, the resulting inverse problems may become too non-linear to enable adequate convergence of deterministic gradient-based inversions even when the forward solver itself is only weakly-nonlinear or even linear.

The remainder of this paper is organized as follows. Section 2 summarizes related work before section 3 briefly describes our used GAN and considered inversion algorithms, and details the considered inverse problem. This is followed by section 4 which presents our inversion results. In section 5, we discuss our main findings and outline possible future developments. Finally, section 6 provides a short conclusion.

2 Related work

To the best of our knowledge, Laloy et al. (2018) were the first to introduce and demonstrate the idea of using the latent space learned by a GAN to perform probabilistic inversion of hydrologic or geophysical data. Richardson (2018) recently proposed to embed $\nabla\mathcal{L}(\mathbf{m})$ and a generator, $G(\mathbf{z})$, trained beforehand within a single (fully differentiable) DL computational graph to deterministically solve a 2D seismic full-waveform inversion problem by making use of the stochastic gradient descent algorithms implemented in DL libraries. Mosser et al. (2018) recently used a trained generator for solving a 2D seismic inversion problem within a probabilistic framework. They combined the estimation of $\nabla\mathcal{L}(\mathbf{m})$ obtained from an adjoint model together with $G(\mathbf{z})$ and the DL-based calculation of \mathbf{V} within the same (fully differentiable) computational graph to perform Metropolis-adjusted Langevin MCMC sampling (e.g., Roberts and Rosenthal, 1998). Outside of the geosciences, Creswell and Bharath (2016) and Lipton and Tripathi (2017) inverted full images produced by a trained generator using usual benchmarks in computer vision and studied to what extent latent vectors can be recovered. In addition, Bora et al. (2017) showed that gradient descent within the latent space of a GAN works overall well to recover an image from a set of measurements obtained by applying a linear operator to the true image. Conceptually, the Bora et al. (2017) study has a lot in commons with ours. However, Bora and coworkers consider a totally different training set than us: the CelebA dataset which consists of face images of celebrities. In section 5, we compare our findings to those of Bora et al. (2017).

3 Methods

3.1 Generative Adversarial Networks

We use a PyTorch implementation of the spatial GAN (SGAN, Jetchev et al., 2016) used by Laloy et al. (2018) to generate realizations from the 2D channelized aquifer training image (TI) depicted in Figure 1a. Our method is not limited to such channel images, but this example is selected to be similar to benchmark images that are typically used to test MPS algorithms. For brevity, we refer the reader to Laloy et al. (2018) for a description of the basic network architecture and training procedure, and only provide below the main principles behind the used SGAN.

The building blocks of a GAN are the generator and discriminator. The generator, $G(\mathbf{z})$, is fed with a low-dimensional latent vector, \mathbf{z} , and produces a model realization, $G(\mathbf{z}) = \mathbf{m}$, whose spatial statistics match those found in the TI provided that the GAN training performed beforehand (see below) was successful. The \mathbf{z} vector is trained to obey either a standard normal distribution, $\mathbf{z} \propto N(\mathbf{0}, \mathbf{I})$, or a truncated uniform distribution, $\mathbf{z} \propto U(-\mathbf{1}, \mathbf{1})$, and the z variables are independent of each other. For image generation, the discriminator component only serves for training $G(\mathbf{z})$ as detailed below.

Due to its purely convolutional nature (see, e.g., Goodfellow et al., 2016, for more details on convolutions and convolutional layers in deep neural networks), for an ergodic TI our SGAN can be trained at relatively low computational cost using a small realization domain, before being used to generate arbitrarily large realizations. The latent space of our SGAN has a spatial structure, with \mathbf{z} being reshaped into an $m \times n \times o \times q$ array \mathbf{Z} for the 3D case. As detailed in Laloy et al. (2018), the m , n and o dimensions are directly related to the three spatial dimensions while q , which we set to 1 in this work, is an extra dimension that can encode additional information about the data representation (see geostatistical simulation case study 2 in Laloy et al., 2018). For a square ($w \times h$)

or cubic ($w \times h \times l$) generation domain, the relationship between $z_x = m = n = o$ and $m_x = w = h = l$ is given by

$$m_x = [z_x - 1] 2^{dp} + 1, \quad (1)$$

where dp is the number of stacked convolutional layers in $G(\mathbf{z})$. This allows for a rather strong dimensionality reduction. For example, when $z_x = 5$ and $dp = 5$, we have $m_x = 129$. We refer to Jetchev et al. (2016) for more details on the (2D) SGAN architecture.

In a GAN, the generator and discriminator are trained (or “learned”) simultaneously with opposing goals. The discriminator, $D(\mathbf{m})$, is fed with samples from the “real” training set, which from now on will be referred to as \mathbf{m}_{true} with distribution $p_{\text{data}}(\mathbf{m})$, and “fake” samples (i.e., realizations) created by the generator: $\mathbf{m} = G(\mathbf{z})$. In our case, the real samples \mathbf{m}_{true} are a set of patches randomly cut from the TI. The discriminator tries to distinguish between \mathbf{m}_{true} and \mathbf{m} by computing, for each received sample, the probability that it belongs to $p_{\text{data}}(\mathbf{m}_{\text{true}})$. Conversely, the generator, $G(\mathbf{z})$, aims at fooling $D(\cdot)$ into labeling \mathbf{m} as a sample from $p_{\text{data}}(\mathbf{m}_{\text{true}})$ (Goodfellow et al., 2014). This translates into the following minimization-maximization problem

$$\min_{G(\cdot)} \max_{D(\cdot)} \left\{ \mathbb{E}_{\mathbf{m}_{\text{true}} \sim p_{\text{data}}(\mathbf{m}_{\text{true}})} [\log(D(\mathbf{m}_{\text{true}}))] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \right\}, \quad (2)$$

Compared to the SGAN code used by Laloy et al. (2018), we replaced batch normalization by instance normalization and added two transposed dilated convolutional layers to the generator, before its output (see section 8). Upon selection of the appropriate training epoch (see below), this removes the need to post-process the realizations by median filtering to eliminate small impurities, thereby leaving us with thresholding as the unique postprocessing operation. The realizations produced by our trained $G(\mathbf{z})$ are continuous on the $[0, 1]$ range and they are converted into binary images by thresholding at the 0.5 level.

Training was performed on a GPU Tesla K40 for 50 epochs with 64 mini-batches of 100 training samples per epoch. Epoch 36 was deemed to produce the best realizations. Figures 1b-c show two (randomly chosen) realizations generated by the trained SGAN before thresholding, $tr[G(\mathbf{z})]$. These 513×513 realization are obtained by sampling a 289-dimensional \mathbf{z} vector from $p(\mathbf{z}) \propto U(-\mathbf{1}, \mathbf{1})$ (see Laloy et al., 2018, for details and a performance comparison with a popular MPS algorithm). Even if these images are continuous, they appear almost categorical in this representation, which also suggests that the subsequent thresholding is a relatively mild operation.




Figure 1: (a) Fraction of size 513×513 of the used 2500×2500 binary TI and (b) - (c) randomly chosen 513×513 realizations derived by our SGAN. Each continuous realization is generated by sampling 289 random numbers from a uniform distribution, $U(-1, 1)$, without any postprocessing.

3.2 Gradient-Based Deterministic Inversion

A common representation of the forward problem is

$$\mathbf{d} = F(\mathbf{m}) + \mathbf{e}, \quad (3)$$

where $\mathbf{d} = (d_1, \dots, d_N) \in \mathbb{R}^N$, $N \geq 1$ is the number of measurement data, $F(\mathbf{m})$ denotes a deterministic forward model with parameters or model \mathbf{m} and the noise term \mathbf{e} lumps all sources of errors. If the probability distribution of \mathbf{e} can be assumed to be zero-mean Gaussian with covariance matrix \mathbf{C}_e , and \mathbf{m} is assigned a multiGaussian prior distribution, then the optimal model, $\hat{\mathbf{m}}$, minimizes the following objective or loss function, $\mathcal{L}(\mathbf{m})$

$$\mathcal{L}(\mathbf{m}) = \Phi_D + \lambda \Phi_M, \quad (4)$$

$$\Phi_D = \frac{1}{2} [F(\mathbf{m}) - \mathbf{d}]^T \mathbf{C}_e^{-1} [F(\mathbf{m}) - \mathbf{d}], \quad (5)$$

$$\Phi_M = \frac{1}{2} [\mathbf{m} - \mathbf{m}_{\text{prior}}]^T \mathbf{C}_m^{-1} [\mathbf{m} - \mathbf{m}_{\text{prior}}], \quad (6)$$

where $\mathbf{m}_{\text{prior}}$ and \mathbf{C}_m signify the a priori model and its covariance matrix, respectively, and where λ weights the influence of honoring prior Gaussian information about \mathbf{m} on $\mathcal{L}(\mathbf{m})$. In case of a linear relationship between \mathbf{d} and $F(\mathbf{m})$, $F(\mathbf{m})$ reduces to the product $\mathbf{F}\mathbf{m}$. In practice, it is common to replace \mathbf{C}_m by a regularization operator, \mathbf{L} . This leads to the joint minimization of the misfit, Φ_D , and $\Phi_M = \|\mathbf{L}(\mathbf{m} - \mathbf{m}_{\text{prior}})\|_2^2$

$$\mathcal{L}(\mathbf{m}) = \frac{1}{2} [F(\mathbf{m}) - \mathbf{d}]^T \mathbf{C}_e^{-1} [F(\mathbf{m}) - \mathbf{d}] + \lambda [\mathbf{m} - \mathbf{m}_{\text{prior}}]^T \mathbf{L}^T \mathbf{L} [\mathbf{m} - \mathbf{m}_{\text{prior}}]. \quad (7)$$

The popular Gauss-Newton (GN) method for minimizing equation (7) iteratively updates \mathbf{m} until either the target data misfit or a maximum number of iterations has been reached using the following update mechanism

$$\mathbf{m}_{k+1} = \mathbf{m}_k + \Delta \mathbf{m} = (\mathbf{J}_k^T \mathbf{C}_e^{-1} \mathbf{J}_k^T + \lambda \mathbf{L}^T \mathbf{L})^{-1} \mathbf{J}_k^T \mathbf{C}_e^{-1} \hat{\mathbf{d}}_k + \mathbf{m}_{\text{prior}}, \quad (8)$$

with

$$\hat{\mathbf{d}}_k = \mathbf{d} - F(\mathbf{m}_k) + \mathbf{J}_k(\mathbf{m} - \mathbf{m}_{\text{prior}}), \quad (9)$$

where k denotes the iteration number. In the context of our GAN-based dimensionality reduction, we learn a $p(\mathbf{z}) = U(-\mathbf{1}, \mathbf{1})$ for the latent vector \mathbf{z} . Limited testing with learning a standard normal distribution, $N(\mathbf{0}, \mathbf{1})$, is found to provide good results but the generated realizations are nevertheless of slightly lower quality than when using $U(-\mathbf{1}, \mathbf{1})$. Fortunately, when the components of a random vector, \mathbf{x} are independent, they can be converted into standard normal variables using the iso-probabilistic transform: labeling as $\tau_x(x_i)$ ($\Psi(\xi_i)$) the cumulative density function (CDF) of x_i (a standard normal variable ξ_i), the direct and inverse transform are given by

$$\xi_i = \Psi^{-1}[\tau_x(x_i)] \quad \text{and} \quad x_i = \tau_x^{-1}[\Psi(\xi_i)]. \quad (10)$$

To perform the inversion within the GAN latent space, we therefore place a standard normal uncorrelated prior (so-called damping regularization) on a vector \mathbf{z}_{SN} of the same size as \mathbf{z} and achieve GN updates of \mathbf{z}_{SN} while converting each proposed \mathbf{z}_{SN} into the corresponding $\mathbf{z} \propto U(-\mathbf{1}, \mathbf{1})$ with equation (10) before creating the resulting $G(\mathbf{z})$ realization. Under the standard normal prior on \mathbf{z}_{SN} , the GN update simplifies to

$$\mathbf{z}_{\text{SN},k+1} = (\mathbf{J}_{z,k}^T \mathbf{C}_e^{-1} \mathbf{J}_{z,k}^T + \lambda \mathbf{I})^{-1} \mathbf{J}_{z,k}^T \mathbf{C}_e^{-1} [\mathbf{d} - F(G(\tau_z^{-1}[\Psi(\mathbf{z}_{\text{SN},k)}])) + \mathbf{J}_{z,k} \mathbf{z}_{\text{SN},k}] \quad (11)$$

where $\mathbf{J}_{z,k} = \mathbf{J}_k^{\text{zSN}}$ is the Jacobian matrix of $\mathbf{z}_{\text{SN},k}$, and \mathbf{I} is the identity matrix. In this work, we simply set $\lambda = 1$ as adjusting λ dynamically did not show any advantages over fixing $\lambda = 1$.

When only the gradient of the objective or loss function with respect to the model parameters, $\nabla \mathcal{L}(\mathbf{m})$, is used, a steepest descent type of update can be performed

$$\mathbf{m}_{k+1} = \mathbf{m}_k - \alpha \nabla \mathcal{L}(\mathbf{m})_k, \quad (12)$$

where α is a step length parameter that is generally adapted dynamically.

In this study, we used the Adam algorithm (Kingma and Ba, 2015) implemented in Pytorch to update \mathbf{z}_k using the $p(\mathbf{z}) = U(-\mathbf{1}, \mathbf{1})$ prior directly. Adam is an advanced gradient-based algorithm that only makes use of the loss function gradient (see Kingma and Ba, 2015, for details). We used the following algorithmic parameters, a learning rate of 0.01 and the default values of 0.9 and 0.999 for the β_1 and β_2 momentum parameters. Since our considered forward operator, $F(\mathbf{m})$, is linear (see section 3.4), it implies that for continuous $\mathbf{m} = G(\mathbf{z})$ realizations it is straightforward to have the DL library internally computing $\nabla \mathcal{L}(\mathbf{z})$ (at least for common loss functions $\mathcal{L}(\cdot)$). Yet this is no longer possible if the $G(\mathbf{z})$ realizations require thresholding. To circumvent this issue, we implemented an adjoint model giving $\nabla \mathcal{L}(\mathbf{m})$ within the Pytorch computational graph (Richardson, 2018). Here $\nabla \mathcal{L}(\mathbf{m})$ is backpropagated through the generator network to update, using Adam, the latent vector, \mathbf{z} , that produced the model, $\mathbf{m} = G(\mathbf{z})$. As explained earlier, this is possible because the generator is fully differentiable end-to-end and $\nabla \mathcal{L}(\mathbf{z}) = \nabla \mathcal{L}(\mathbf{m}) \mathbf{V}$ (see also our companion code for more details).

A simple sum of squared errors (SSR, see equation 14) was chosen for the loss function, $\mathcal{L}(\mathbf{m})$. Consequently, specific steps are needed to deal with the fact that the estimates in \mathbf{z} can move outside the support of the $U(-\mathbf{1}, \mathbf{1})$ prior. Three options to deal with this situation have been investigated by Lipton and Tripathi (2017) in the

context of direct inversion of a full (face) image: (i) allow z values to leave the $[-\mathbf{1}, \mathbf{1}]$ interval, (ii) replace components that are too large with the maximum allowed value and components that are too small with the minimum allowed value (i.e., standard clipping), (iii) reassign the exceeding \mathbf{z} components uniformly at random in the $[-\mathbf{1}, \mathbf{1}]$ range (i.e., stochastic clipping). In agreement with the findings by Lipton and Tripathi (2017), limited testing revealed that stochastic clipping produces slightly better results for our considered case studies. Therefore, stochastic clipping is adopted throughout.

3.3 Global Optimization

For comparison purposes, we used the DREAM_(ZS) algorithm (Laloy and Vrugt, 2012; Vrugt, 2016) as global optimizer. This is a MCMC sampler designed to sample the posterior density function of the parameters. Yet it is herein used to solely find an appropriate solution rather than sampling the full posterior distribution. The DREAM_(ZS) scheme runs parallel interacting Markov chains and is thus population-based, and makes use of differential evolution principles to propose candidate points. It has been proven efficient in many hydrologic and geophysical applications (see, e.g., Vrugt, 2016, for references). Full algorithmic details can be found in Laloy and Vrugt (2012) and Vrugt (2016). Using a population-based search method requires finding the right trade-off between the size of the population (or “exploration”) and the number of generations (or “exploitation”) for a given number of forward model runs. After some testing, we selected 8 Markov chains, meaning a 8-member population, and performed as many MCMC iterations, that is, population generations, as necessary to fit the data to the assigned noise level.

3.4 Inverse Problem

We consider a synthetic 2D straight-ray cross-hole GPR tomography example. Given that the physics is assumed to be fully linear, any convergence problems arising in solving the inverse problem can be attributed to the nonlinear relationship between \mathbf{z} and $G(\mathbf{z})$. We use two binary images with two facies of homogeneous GPR velocity (0.06 m ns^{-1} and 0.08 m ns^{-1}) as true models (see Figures 2a, 3a, 4a, and 5a). Both are produced using $\mathbf{m} = G(\mathbf{z})$ with \mathbf{z} randomly sampled from $U(-\mathbf{1}, \mathbf{1})$. Cross-hole GPR imaging uses a transmitter antenna to emit a high-frequency electromagnetic wave at a location in one borehole and a receiver antenna to record the arriving energy at a location in another borehole (e.g., Annan, 2005). The considered measurement data are first-arrival traveltimes for several transmitter and receiver locations. These data contain information about the GPR velocity distribution between the boreholes. The GPR velocity primarily depends on dielectric permittivity, which is strongly influenced by volumetric water content and, consequently, porosity in saturated media. Our setup consists of two vertical boreholes that are located 6.0 m apart. Sources (left) and receivers (right) are located between 0.5 and 12.0 m depth with 0.5 m spacing (Figures 2a, 3a, 4a, and 5a), leading to a total dataset of $d_N = 576$ traveltimes. Under the linear physics assumption, synthetic traveltime data, \mathbf{d} , are simulated as

$$\mathbf{d} = \mathbf{A}\mathbf{m} + \mathbf{e}, \quad (13)$$

where $\mathbf{A} = \mathbf{J}$ contains the path length in each model cell, \mathbf{e} represents independent random draws from a zero-mean homoscedastic Gaussian distribution with a typical standard deviation $\sigma_e = 1 \text{ ns}$. When a simple SSR loss function (without regularization) is used

$$\mathcal{L}(\mathbf{m}) = \|\mathbf{d} - \mathbf{A}\mathbf{m}\|^2 \equiv \mathcal{L}(\mathbf{z}) = \|\mathbf{d} - \mathbf{A}G(\mathbf{z})\|^2, \quad (14)$$

then the gradient vector of $\mathcal{L}(\mathbf{m})$ with respect to \mathbf{m} becomes

$$\nabla \mathcal{L}(\mathbf{m}) = -2\mathbf{A}^T [\mathbf{d} - \mathbf{A}\mathbf{m}]^T \equiv \nabla \mathcal{L}(G(\mathbf{z})) = -2\mathbf{A}^T [\mathbf{d} - \mathbf{A}G(\mathbf{z})]^T. \quad (15)$$

As explained above, this allows us to use Pytorch to backpropagate $\nabla \mathcal{L}(\mathbf{m})$ (equation (15)) through the GAN generator and optimize \mathbf{z} such that equation (14) is minimized within a single computational graph using gradient descent. From now on, this will be referred to a strategy 1. In addition, estimating $\mathbf{J}_k^{\text{ZSN}}$ at each iteration k by a centered 2-point finite difference scheme together with using a GN search as described by equations 7-11 will be called strategy 2. Here \mathbf{J}^{ZSN} is derived using a perturbation factor of 0.1. This perturbation value was found to be the most appropriate after testing with several values in the $[0.01, 0.5]$ range. A perturbation of 0.1 for a $N(\mathbf{0}, \mathbf{1})$ prior might seem large. Yet one must keep in mind that the necessary thresholding of the produced realizations, $G(\mathbf{z})$, can cause one or more sensitivities in \mathbf{J}^{ZSN} to be zero or near-zero if the used perturbation is too small. Such excessively small $J_{i,j}^{\text{ZSN}}$ values will induce instabilities in the GN search described by equation (11).

Lastly, note that besides true models I and II we use two different realizations of \mathbf{e} in equation (13) which leads to two different “true” datasets for each true model. Realization I induces a true root mean square error (RMSE) between “true” (corrupted) and uncorrupted measurements of 1.006 ns. For realization II, the corresponding RMSE is 1.001 ns. Repeating the inversions with different noise realizations is useful to check the robustness of the inversion results against the exact noise values. Overall, this leads to 4 inverse case studies: 2 true models \times 2 noise realizations. To ease the comparisons between the different inversion runs, we use the weighted root mean square error (WRMSE) which we define as the ratio of the achieved RMSE by the inversion to the true RMSE. A WRMSE > 1 ns thus basically means that the data are underfitted while WRMSE = 1 denotes an appropriate data misfit.

4 Results

4.1 Inverting the Full Image

Before considering application of the forward operator, we tested whether gradient descent within the latent space of our trained GAN is able to recover the true (non-thresholded) $G(\mathbf{z})$ model and corresponding \mathbf{z} vector from the inversion of $G(\mathbf{z})$. In other words, is it possible to recover the values in \mathbf{z} given exhaustive knowledge of the true $G(\mathbf{z})$? This task is similar to the one considered by Lipton and Tripathi (2017), but in a geophysical context rather than for face images. The considered two true models in $[0, 1]$ (I and II) generated by our trained GAN are depicted in Figures 2a and 3a. No data corruption was used and 50 inversion trials using 10,000 Adam iterations were performed with randomly chosen starting points. Table 1 and Figures 2-3 summarize the results. In Table 1, a RMSE of 0.02 between the true and reconstructed images is selected to define a successful reconstruction. This choice is based on the observation that, visually, the retrieved models are either very similar to the true images with a RMSE < 0.02 or largely differ from the true images with a RMSE $> 0.2 - 0.3$. Furthermore, similarly to Lipton and Tripathi (2017), we report the percentage of inversion results for which the mean squared error (MSE) between the true and inferred z values is < 0.01 . Overall, 37% of the inversion runs recover the true models (Table 1). Furthermore, each of these successful runs corresponds to a successful recovery in the latent space for true model I (Table 1 and Figure 2) but less for model II (Table 1 and Figure 3). As illustrated by Figures 3b and 3e, this is because a slight linear bias in the z values can nevertheless

produce good image reconstructions. Our results somewhat differ from those obtained by Lipton and Tripathi (2017) who report 100% recovery in the \mathbf{z} -space to arbitrary precision. There are two major differences between our work and that of Lipton and Tripathi (2017). First, they consider face images which are much more clustered compared to our channelized model: the eyes, nose, ears, etc. tend to occupy the same portion of the image and they are located at globally similar distances from each other. Second, Lipton and Tripathi (2017) performed 100,000 gradient descent iterations while we performed 10 times less iterations.

Table 1: Percentage (%) of inversion runs that achieved a RMSE ≤ 0.02 in reconstructing the full 125×60 image ($G(\mathbf{z})$) and a MSE ≤ 0.01 in recovering the true 15-dimensional \mathbf{z} vector, for each true model (I and II).

True model	RMSE < 0.02	MSE < 0.01
1	34	32
2	40	8




Figure 2: Inversion results for the full image inversion using the true model I: (a) true model, (b) and (c): retrieved models with successful image recovery, (d) retrieved model with unsuccessful image recovery, (e - g) scatter plot of the recovered \mathbf{z} vector associated with the model shown in subplot (b - d) and the true \mathbf{z} vector, respectively.




Figure 3: Inversion results for the full image inversion using the true model II: (a) true model, (b) and (c): retrieved models with successful image recovery, (d) retrieved model with unsuccessful image recovery, (e - g) scatter plot of the recovered \mathbf{z} vector associated with the model shown in subplot (b - d) and the true \mathbf{z} vector, respectively.

4.2 Geophysical Inversion

We now turn our attention to the actual geophysical inversion, in which a subset of linear averages are observed under noise. For each strategy (1 and 2) and each combination of true model (I or II) and noise realization (I and II), the inversion was repeated 100 times using 100 different (randomly chosen) starting models. For strategy 1, the maximum total number of iterations (i.e., proposal steps) was set to 10,000. Here each Adam iteration consumes one forward solve. Since the resulting computational expense may

be overly large for realistic nonlinear geophysical inverse problems, inversion results are also looked at after a total of 100 and 1000 iterations. Due to the used centered 2-point finite-difference approximation of \mathbf{J}^{ZSN} (see above), for strategy 2 one iteration incurs a cost of 31 forward solves. Therefore, the maximum number of iterations for strategy is set to 5000 which translates into 155,000 forward model runs. In addition, inversion performance is also considered after 100 (3100 forward solves) and 1000 (31,000 forward solves) iterations. Moreover, to verify the impact of thresholding on the quality of the inferred models, each case study is also run without thresholding.

Strategy 1 without thresholding relies on the direct computation of $\nabla\mathcal{L}(\mathbf{z})$ by autodifferentiation with no adjoint model involved and is therefore basically the same as the approach by Bora et al. (2017). For the continuous case, we consider the continuous versions of the true models shown in Figures 2a and 3a and the same noise realizations as for the binary case. To convert a generated continuous model $\mathbf{m} = G(\mathbf{z})$ in $[\mathbf{0}, \mathbf{1}]$ into a continuous velocity field \mathbf{m}_{VEL} (m ns^{-1}) in $[\mathbf{0.06}, \mathbf{0.08}]$, the following relationship is used: $\mathbf{m}_{\text{VEL}} = 0.06 + 0.02(\mathbf{1} - \mathbf{m})$

The results of the 100 repetitions of the 8 different inversions are summarized in Table 2 for the continuous case and in Table 3 for the binary case. Furthermore, Figures 4-7 depict for each strategy applied to the binary case, true models I and II, and noise realization II, the true model (subplot (a)), two best-fitting models found among the 100 repetitions (subplots (b) and (c)), the best-fitting model across the allowed iterations of a randomly chosen trial for which the corresponding best WRMSE is > 1.2 ns (subplot (d)), the sampled WRMSE trajectory corresponding to the best-fitting model shown in subplot (b) (subplot (e)) and the sampled WRMSE trajectory associated with the model displayed in subplot (d) (subplot (f)).

Table 2: Continuous case: number of inversion runs among the 100 performed ones that achieved a WRMSE ≤ 1.2 , ≤ 1.1 and ≤ 1 , respectively, for each inversion strategy (1 and 2), total number of iterations, N_{iter} , combination of true model (I and II) and noise realization (I and II). The WRMSE is defined for each inversion run as the ratio of the best RMSE (over the iterations of a given run) to the RMSE of the true data (1.006 ns for noise realization I and 1.001 ns for noise realization II).

Strategy	N_{iter}	True model	Noise	WRMSE ≤ 1.2	WRMSE ≤ 1.1	WRMSE ≤ 1.01
1	100	I	I	0	0	0
1	100	I	II	0	0	0
1	100	II	I	0	0	0
1	100	II	II	0	0	0
1	1000	I	I	3	3	0
1	1000	I	II	6	1	0
1	1000	II	I	3	1	0
1	1000	II	II	3	2	0
1	10,000	I	I	24	21	16
1	10,000	I	II	25	24	22
1	10,000	II	I	14	13	4
1	10,000	II	II	20	16	9
2	100	I	I	19	0	0
2	100	I	II	6	2	0
2	100	II	I	0	0	0
2	100	II	II	1	0	0
2	1000	I	I	100	1	0
2	1000	I	II	84	8	0
2	1000	II	I	0	0	0
2	1000	II	II	2	1	0
2	5000	I	I	100	6	0
2	5000	I	II	100	14	0
2	5000	II	I	2	0	0
2	5000	II	II	2	2	0

Table 3: Binary case: number of inversion runs among the 100 performed ones for the continuous case that achieved a WRMSE ≤ 1.2 , ≤ 1.1 and ≤ 1 , respectively, for each inversion strategy (1 and 2), total number of iterations, N_{iter} , combination of true model (I and II) and noise realization (I and II). The WRMSE is defined for each inversion run as the ratio of the best RMSE (over the iterations of a given run) to the RMSE of the true data (1.006 ns for noise realization I and 1.001 ns for noise realization II).

Strategy	N_{iter}	True model	Noise	WRMSE ≤ 1.2	WRMSE ≤ 1.1	WRMSE ≤ 1.01
1	100	I	I	0	0	0
1	100	I	II	1	0	0
1	100	II	I	0	0	0
1	100	II	II	0	0	0
1	1000	I	I	5	3	0
1	1000	I	II	4	3	0
1	1000	II	I	3	2	0
1	1000	II	II	1	0	0
1	10,000	I	I	21	19	13
1	10,000	I	II	24	23	19
1	10,000	II	I	10	9	7
1	10,000	II	II	24	17	12
2	100	I	I	22	0	0
2	100	I	II	8	3	0
2	100	II	I	0	0	0
2	100	II	II	0	0	0
2	1000	I	I	100	0	0
2	1000	I	II	57	4	0
2	1000	II	I	3	0	0
2	1000	II	II	2	0	0
2	5000	I	I	100	0	0
2	5000	I	II	97	5	0
2	5000	II	I	6	0	0
2	5000	II	II	9	1	0




Figure 4: Binary case: inversion results for strategy 1, truel model I and noise realization II: (a) true model, (b) and (c) two randomly chosen best-fitting models found over 100 repetitions, (d) best-fitting derived model across the allowed iterations of a randomly chosen repetition for which the corresponding best WRMSE is > 1.2 ns, (e) sampled WRMSE trajectory for the best model found among the 100 repetitions (displayed in subplot(b)), and (f) sampled WRMSE trajectory associated with the model depicted in subplot (d). The red triangles and the green squares in subfigures (a-d) represent the GPR source and receiver positions, respectively.




Figure 5: Binary case: inversion results for strategy 1, truel model II and noise realization II: (a) true model, (b) and (c) two randomly chosen best-fitting models found over 100 repetitions, (d) best-fitting derived model across the allowed iterations of a randomly chosen repetition for which the corresponding best WRMSE is > 1.2 ns, (e) sampled WRMSE trajectory for the best model found among the 100 repetitions (displayed in subplot(b)), and (f) sampled WRMSE trajectory associated with the model depicted in subplot (d). The red triangles and the green squares in subfigures (a-d) represent the GPR source and receiver positions, respectively.




Figure 6: Binary case: inversion results for strategy 2, truel model I and noise realization II: (a) true model, (b) and (c) two randomly chosen best-fitting models found over 100 repetitions, (d) best-fitting derived model across the allowed iterations of a randomly chosen repetition for which the corresponding best WRMSE is > 1.2 ns, (e) sampled WRMSE trajectory for the best model found among the 100 repetitions (displayed in subplot(b)), and (f) sampled WRMSE trajectory associated with the model depicted in subplot (d). The red triangles and the green squares in subfigures (a-d) represent the GPR source and receiver positions, respectively.




Figure 7: Binary case: inversion results for strategy 2, truel model II and noise realization II: (a) true model, (b) and (c) two randomly chosen best-fitting models found over 100 repetitions, (d) best-fitting derived model across the allowed iterations of a randomly chosen repetition for which the corresponding best WRMSE is > 1.2 ns, (e) sampled WRMSE trajectory for the best model found among the 100 repetitions (displayed in subplot(b)), and (f) sampled WRMSE trajectory associated with the model depicted in subplot (d). The red triangles and the green squares in subfigures (a-d) represent the GPR source and receiver positions, respectively.

It is observed for both strategies 1 and 2, that the quality of the best-fitting solution heavily depends on the starting model and the allowed number of iterations. Indeed for the smallest budget of 100 iterations, no satisfying solutions are found no matter the inversion approach. We find that strategy 1 (Figures 4-5 and Tables 2-3) shows a much

higher success rate than strategy 2 (Figures 6-7 and Tables 2-3) for the largest number of iterations allowed to each method. Hence strategy 1 retrieves solutions with a WRMSE ≤ 1.01 in 4% to 22% of the trials for the continuous case and in 7% to 19% of the trials for the binary case. On the contrary, strategy 2 never does so (Tables 2-3).

Strikingly, there appears to be very little negative impact of thresholding (compare Table 2 to Table 3), except for strategy 2 and true model I (compare the bottom row of Table 2 with those of Table 3). This indicates that plugging the externally computed $\nabla\mathcal{L}(\mathbf{m})$ into the DL graph to permit gradient descent works well for the considered examples.

When the maximum number of iterations is reduced to 1000, the proportion of solutions associated with a WRMSE ≤ 1.1 becomes small but is overall two times larger for strategy 1 (2% on average over the 4 datasets) than for strategy 2 (1% on average over the 4 datasets). Since a model with a WRMSE as large as 1.2 may still visually resemble the true model well (see Figure 6c), it is instructive to also consider a WRMSE ≤ 1.2 as threshold. Here strategy 2 becomes more attractive, but for true model I only (Tables 2-3). We have tried to improve strategy 2 by (i) damping the model update at the early iteration before increasingly adding more contrast to the proposed models and (ii) decreasingly smooth the proposed models after each iteration as the search progresses. Nevertheless, neither of those schemes proved successful.

There is a systematic difference in performance between the cases with true models I and II, with true model I always inducing a larger success rate no matter the inversion strategy and computational effort (Tables 2 and 3). This is due to the fact that model I has higher prior probability or, to use deep learning terminology, is closer to the range of the generator (Bora et al., 2017) than model II. This can be informally seen by comparing models I (Figure 2a) and II (Figure 3a) with the original TI (Figure 1a): the patterns displayed by model I are occurring more frequently in the TI than the eye-type of channel branching contained in model II. It is seen that strategy 1 offers greater robustness than strategy 2 against variations in prior model probability, as shown by the smaller differences in success rates between the model I and II cases for strategy 1 (Tables 2 and 3).

With respect to sampling trajectories of the inversions, for strategy 1 most of the data misfit reduction in a productive run (i.e., a run with a favorable starting model) is achieved within the first 3000-5000 iterations (Figures 4e and 5e). The main reduction in data misfit for strategy 2 can occur both rather quickly (< 500 iterations) or much later (Figures 6e and 7e), while this strategy hardly recovers models with WRMSE ≤ 1.1 . In addition, the sampling behavior of both algorithms can be unstable (Figures 4e, 5e, 6e and 7e). In accordance with those findings, it is also noted for strategy 1 that if WRMSE values larger than 2-3 are still sampled after some 2000 forward simulations, then the considered run will be unable to obtain a high-quality solution (Figures 4f and 5f).

It is worth noting that using strategy 2 with a trained GAN for which the latent space obeys a standard normal distribution, $N(\mathbf{0}, \mathbf{1})$, thereby removing the additional nonlinearity caused by the $\Psi(\mathbf{z}_{\text{SN},k})$ transform, leads to globally similar results as those presented in Tables 2 and 3 for strategy 2.

Even for our best performing strategy 1 and considered maximum number of iterations of 10,000, the overall success (WRMSE ≤ 1.01) rate remains relatively low (Tables 2 and 3). As posited earlier in this paper, we argue that the reason for this is the high degree of nonlinearity associated with the relationship between \mathbf{z} and $G(\mathbf{z})$. More insights into the effect of the nonlinear $G(\mathbf{z})$ transform on the considered simple inverse problem are provided by Figure 8. The latter displays a 2D slice in the WRMSE land-

scape corresponding to the true reference solution obtained by the combination of the true model I (see Figure 2a, 4a or 6a) and noise realization II. To construct Figure 8, the z_3 to z_{15} components of the true \mathbf{z} vector that induces model I are kept fixed while the z_1 and z_2 dimensions are gradually varied between -1 and 1. While the WRMSE response surface appears smooth when considered at a coarse resolution (Figure 8a), a zoom into the basin of attraction (Figure 8b) reveals that the problem is in fact multimodal with many local minima. Also, this WRMSE landscape contains many small spikes and some large flat areas.




Figure 8: 2D slice in the WRMSE landscape corresponding to the true reference solution obtained by the combination of the true model I (see Figure 2a, 4a or 6a) and noise realization II. The (a) subplot shows a coarse scale resolution over the full WRMSE range while the (b) subplot presents a zoom into the main basin of attraction. The red cross in the (b) subplot denotes the true values of z_1 and z_2 in this case.

Lastly, we find that performing a global optimization with $\text{DREAM}_{(ZS)}$ parameterized as described in section 3.3 provides a solution with a WRMSE of 1.000 for all of the 4 scenarios (true models I and II combined with noise realizations I and II). Nevertheless, this is at the cost of a serial total of 200,000 to 800,000 forward runs. Among these amounts, some 50,000 serial forward calls are consumed to reach a WRMSE of about 1.1 while the rest is spent by moving from 1.1 to ≤ 1 . This points out that traveling across the relatively rough misfit terrain caused by the nonlinearities in $G(\mathbf{z})$ may require global optimization or search methods even in the best case of a linear forward problem. Although not done herein, note also that population-based methods such as $\text{DREAM}_{(ZS)}$ are straightforward to evolve in parallel using one CPU per population member, which can save a large amount of time (Laloy et al., 2018).

5 Discussion

This paper demonstrates that even if the forward model is linear, the high degree of nonlinearity in the relationship between latent vector, \mathbf{z} , and the generator of a GAN, $G(\mathbf{z})$, can cause gradient-based deterministic inversion performed in the GAN latent space $p(\mathbf{z})$ to fail. This is here illustrated for a classical binary channelized aquifer TI.

Less structured and/or continuous TIs might be more amenable to gradient-based deterministic inversion within $p(\mathbf{z})$ since the $G(\mathbf{z})$ transform might then be less nonlinear. The poor convergence to the target misfit is in stark contrast to the results obtained when inverting the full image without noise (Table 2, Figures 2-3). This suggests that limited sampling of spatial averages under noise further complicates the recovery of the true z values.

Our findings might appear to contradict those of Bora et al. (2017) who investigated whether gradient descent within the latent space of a trained GAN can recover an image from a set of measurements obtained by applying a linear operator to the true image. In cases where this true image is itself created by running the trained GAN for a given \mathbf{z} vector (similarly as in our study), Bora et al. (2017) systematically find small reconstruction errors. Nonetheless, Bora and coworkers consider face images (from the Celeba dataset) which are thus very different than our channelized aquifer training image. Face images are comparatively more clustered as the same structures (nose, ears, eyes,...) are more or less always located in the same areas of the image. Yet the channels in our training set can be anywhere (see Laloy et al., 2018, for the full 2500×2500 training image from which the training set is sampled). This larger degree of freedom in our images (for this aspect) might make the inversion problem harder to solve. In addition, Bora et al. (2017) do not give much details about their exact experimental setup which makes an in-depth comparison with our results difficult.

It might seem surprising that, if a sufficient number of iterations is allowed, strategy 1, which is based on the objective function gradient, $\nabla\mathcal{L}(\mathbf{m})$, outperforms a GN search that uses a finite-difference approximation of the full Jacobian matrix in the $p(\mathbf{z})$ space, \mathbf{J}^z . In the training of deep neural networks, stochastic gradient descent has proven highly successful. Goodfellow et al. (2016) argue that such first-order methods may more easily escape saddle points, while second-order methods tend to get attracted and stuck in them. The situation might be similar here where we try to estimate \mathbf{z} for an already trained network, namely that the less "stable" gradient-based methods can more easily move away from saddle-points and possibly local minima. Furthermore, strategy 1 relies on a rather powerful stochastic gradient descent algorithm (Adam, Kingma and Ba, 2015) and uses a perfectly accurate estimation of $\nabla\mathcal{L}(\mathbf{m})$ since in the considered case, $\nabla\mathcal{L}(\mathbf{m})$ is available analytically. On the contrary, when approximated by finite differencing \mathbf{J}^z is inevitably fraught with errors.

Note that for the gradient descent used with strategy 1, no optimization of the Adam's hyperparameters, that is, the learning rate, and β_1 and β_2 momentum parameters (see Kingma and Ba, 2015), was done. Only the learning rates of 0.1 and 0.001 were found to lead to a worse performance than using our choice of 0.01. Other Adam's hyperparameter settings or other descent algorithms such as SGD, RMSPROP (see, e.g., Goodfellow et al., 2016) or Adabound (Luo et al., 2019) might enabled improved inversion quality for the considered case studies.

Lastly, notice that some recent work suggests that projecting the learned latent space onto a so-called "Riemannian" manifold may lead to a reduced non-linearity of $G(\mathbf{z})$ with respect to the projected \mathbf{z} (Shao et al., 2017; Arvanitidis et al., 2018; Chen et al., 2018). In the future, it might be worthwhile to attempt inversion within such a manifold.

6 Concluding remarks

Performing inversion within the low-dimensional, latent space of a trained GAN, $p(\mathbf{z})$, rather than within the original model space, $p(\mathbf{m})$, offers important advantages in that (1) any generated model proposal honors the prior training image (TI) and (2) the

parameter space is reduced by orders of magnitude compared to the original model space. Global probabilistic inversion within $p(\mathbf{z})$ has been recently shown to work well (Laloy et al., 2018). However, such probabilistic inversion still incurs a large computational cost while the learned $p(\mathbf{z})$ could also be possibly used within computationally less demanding gradient-based deterministic inversions. In this work, we show that owing to the highly nonlinear relationship between the GAN generator, $G(\mathbf{z})$, and latent vector, \mathbf{z} , gradient-based deterministic inversion may fail even though the physics of the forward problem is linear. For a channelized aquifer binary TI and a synthetic linear GPR tomography problem involving 576 measurements with low noise, we find that when allowing for a total of 10,000 iterations, about 13% of the trials by the considered gradient descent algorithm locate a solution that has the required data misfit, in average over the considered test cases. When restricting the maximum allowed number of iterations to 1000, the percentage of success becomes 0%, though approximately 2 % of the inferred models still closely resemble their associated true model and induce a nearly correct data misfit. Furthermore, the tested Gauss-Newton inversion approach that approximates the Jacobian matrix to create updates proved to be unable to recover a solution with the appropriate data misfit. Overall, deterministic inversion performance is found to depend on the inversion approach, starting model, true reference model, number of iterations and noise realization. In contrast, costly global optimization with DREAM_(ZS) always finds an appropriate solution.

7 Computer Code Availability

The GAN and associated inversion codes used in this study are available at https://github.com/elaloy/gan_for_gradient_based_inv.

8 Used GAN architecture

Compared to the generator architecture detailed in Laloy et al. (2018), we have the following differences.

- Batch normalization was replaced by instance normalization
- The nonlinearity in the 5th layer of the generator is a ReLU instead of an hyperbolic tangent.
- A first transposed dilated convolutional layer takes the output of the 5th layer as input. This 6th layer is parameterized with 64 output channels, a kernel size of 5 pixels, a stride of 1 pixel, a padding of 6 pixels, an output padding of 0 pixels and dilation coefficient of 3. The activation function in this 6th layer is a ReLU and instance normalization is applied after the nonlinearity.
- A second transposed dilated convolutional layer takes the output of the 6th layer as input. This 7th layer is parameterized with 1 output channel, a kernel size of 5 pixels, a stride of 1 pixel, a padding of 10 pixels, an output padding of 0 pixels and dilation coefficient of 5. The nonlinearity in this 7th layer is a hyperbolic tangent.

References

- Abadi, M., Agarwal, A., Barham, P., et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. ArXiv e-prints, March 2016, <https://www.tensorflow.org>.
- Al-Rfou, R., Alain, G., Almahairi, A., et al. Theano: A Python framework for fast computation of mathematical expressions. arXiv e-prints, May 2016, <http://arxiv.org/abs/1605.02688>
- Annan, A. P. GPR methods for hydrogeological studies, in Hydrogeophysics, pp. 185–214, eds Rubin, Y. & Hubbard, S.S., Springer. 2005.
- Arvanitidis, G., Hansen, L.K, and S. Hauberg. Latent space oddity: on the curvature of deep generative models. arXiv e-prints, 2018.
- Aster, R., Borchers, B., and C. H. Thurber. Parameter estimation and inverse problems, 2nd edition. Elsevier. 2012.
- Bora, A., Jalal, A., Price, E., and A. G. Dimakis. Compressed sensing using generative models. ArXiv e-prints, 2017.
- Chen, N., Klushyn, A., Kurle, R., Jiang, X., Bayer, J., van der Smagt, P. Metrics for deep generative models. ArXiv e-prints, 2018.
- Creswell, A., and A.A. Bharath. Inverting the generator of a generative adversarial network. ArXiv e-prints, 2016.
- Creswell, A., White, T., Dumoulin, V., Arulkum, K., Sengupta, B., and A.A. Bharath. Generative adversarial networks: An overview. In the Proceedings of IEEE Signal Processing Magazine Special Issue on Deep Learning for Visual Understanding, *accepted paper*, 2017. ArXiv e-prints, October 2017.
- de Groot-Hedlin, C. and Constable, S. Occam’s inversion to generate smooth, two-dimensional models from magnetotelluric data. Geophysics. 55:1613–1624, 1990. <https://doi.org/10.1190/1.1442813>.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y. Generative adversarial networks. The Annual Conference on Neural Information Processing Systems (NIPS), Montréal; 2014.
- Goodfellow, I., Bengio, Y., Courville, A. Deep learning. MIT Press; 2016. <http://www.deeplearningbook.org>.
- Jetchev, N., U. Bergmann, and R. Vollgraf. Texture synthesis with spatial generative adversarial networks, ArXiv e-prints, December 2016.
- Kingma, DP, Ba JL. Adam: a Method for Stochastic Optimization. The International Conference on Learning Representations (ICLR), San Diego; 2015.
- Laloy E, Vrugt JA. High-dimensional posterior exploration of hydrologic models using multiple-try DREAM_(ZS) and high-performance computing. Water Resour. Res. 48(1), 2012. <http://dx.doi.org/10.1029/2011WR010608>.

- Laloy, E., Hérault, R., Jacques, D., Linde, N. Training-image based geostatistical inversion using a spatial generative adversarial neural network. *Water Resour. Res.* 54(1):381–406, 2018. <https://doi.org/10.1002/2017WR022148>.
- Lipton, Z. C., and S. Tripathi. Precise recovery of latent vectors from generative adversarial networks. *ArXiv e-prints*, 2017.
- Luo, L., Xiong, Y., and Y. Liu. Adaptive gradient methods with dynamic bound of learning rate. Presented at the International Conference on Learning Representations (ICLR) 2019. *ArXiv e-prints*, 2019.
- Mariethoz G, Renard P, Straubhaar J. The Direct Sampling method to perform multiple-point geostatistical simulations. *Water Resour. Res.* 46:W11536, 2010. <http://dx.doi.org/10.1029/2008WR007621>.
- Mosser, L., Dubrule, O., Blunt, M. J. Reconstruction of three-dimensional porous media using generative adversarial neural networks, *Physical Review E.* 96:043,309, 2017. <http://doi:10.1103/PhysRevE.96.043309>.
- Mosser, L., Stochastic seismic waveform inversion using generative adversarial networks as a geological prior. *ArXiv e-prints*, June 2018.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A. Automatic differentiation in PyTorch. *The Annual Conference on Neural Information Processing Systems (NIPS)* 2017.
- Price, K.V., Storn, R.M., Lampinen, J.A. *Differential evolution, A practical approach to global optimization*, Springer, Berlin. 2005.
- Richardson, A. Generative adversarial networks for model order reduction in seismic full-waveform inversion. *ArXiv e-prints*, June 2018.
- Roberts, G. O. and Rosenthal, J. S. Optimal scaling of discrete approximations to langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268, 1998.
- Shao, H., Kumar, A., Fletcher, P. T. The Riemannian geometry of deep generative models. *arXiv e-prints*, 2017.
- Storn, R., Price, K. Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization.* 11:341–359, 1997.
- Strebelle S. Conditional simulation of complex geological structures using multiple point statistics. *Math. Geol.* 34(1):1–22, 2002.
- 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.* 10(3):273–290, 2009.
- Vrugt, JA. Markov chain Monte Carlo simulation using the DREAM software package: theory, concepts, and MATLAB implementation. *Environ. Modell. Software.* 75:273–316, 2016.