

Deep generator priors for Bayesian seismic inversion

Zhilong Fang¹, Hongjian Fang¹, Laurent Demanet¹

¹Massachusetts Institute of Technology

January 2020

Abstract

Earth scientists and exploration geophysicists have a strong demand on comprehensive information of subsurface medium parameters to study Earth’s internal structures. One famous approach to meeting this demand is the Bayesian inference method, which integrates statistical information from the forward modeling function, observed data, and experts’ prior knowledge into a posterior probability density function (PDF). Conventional prior knowledge is based on empirical observations of subsurface structures like the smoothness of the subsurface image. However, such hand-designed prior knowledge is too generic to describe detailed subsurface structures. In this work, we propose a learning-based prior knowledge generator. Unlike the hand-designed prior knowledge, we use existing subsurface velocity models to train a deep generative adversarial network (GAN) that generates artificial models from a low-dimensional latent space. We test the proposed deep generator prior by applying it to traveltimes tomography and full waveform inversion. Benefits of the proposed deep generator priors include: (1) the generated models share a similar spatial distribution with the existing models; (2) the latent space is much smaller than the model space yielding a significant reduction in the computational complexity for the inversion. Despite these benefits, the generality of the training set has a strong influence on the robustness of the deep generator prior. Like most deep learning applications, a diverse training set (here, containing different spatial velocity distributions) is necessary and essential to make the deep generator prior effective.

Acknowledgments

The authors acknowledge the funding and support provided by ExxonMobil Research and Engineering Company. LD is also supported by AFOSR grant FA9550-17-1-0316.

Introduction

Earth scientists and exploration geophysicists study the Earth’s interior to answer fundamental questions about the structure and dynamics of Earth. Since the current technology cannot offer researchers the direct access to Earth’s deep interior, researchers utilize cutting-edge seismic inversion techniques to infer unknown subsurface physical parameters (e.g., sound speed and density) from indirectly measured data (e.g., seismic traveltimes and seismic waveforms). With the developments in high-performance computing and advances in modern numerical and optimization methods in the past 30 years, cutting-edge seismic inversion techniques including traveltimes tomography (TT) [Aki et al., 1977, Romanowicz, 1979, Dziewonski, 1984, Nolet, 2012] and full-waveform inversion (FWI) [Tarantola and Valette, 1982b, Lailly et al., 1983, Pratt, 1999, Virieux

and Operto, 2009] enable researchers to generate high-resolution subsurface structures at different scales, from the inner core to the shallow subsurface within acceptable computational time[Tromp, 2019].

Observed seismic data contain noise from different sources including human activities and environmental movements. As the inversion proceeds, uncertainties in the observed data would propagate into final results, yielding uncertainties there. Because the inverted model parameters are inputs for a sequence of follow-up interpretations and processing, analyzing uncertainties in the inverted model parameters can lead to significant improvements in the understanding of Earth’s internal structures [Osypov et al., 2013]. The analysis of uncertainties requires a complete statistical description of the unknown model parameters. To that end, conventional deterministic approaches, which only find the best model parameters minimizing the misfit between the predicted and observed data, may not be appropriate candidates. On the other hand, statistical approaches, and in particular the Bayesian inference method [Kaipio and Somersalo, 2006], are desirable and necessary.

The Bayesian inference method is one of the most widely used approaches to solving statistical inverse problems and has been applied to many geophysical problems [Tarantola and Valette, 1982a, Duijndam, 1988, Scales and Tenorio, 2001, Sambridge et al., 2006, Osypov et al., 2008, Ely et al., 2017, Fang et al., 2018]. This approach formulates the inverse problem in the framework of statistical inference. It considers all unknown model parameters as random variables and incorporates statistical information in the observed data, the underlying forward modeling map, and researchers’ prior knowledge about the unknown model parameters. The solution of the Bayesian inverse problem is a posterior probability density function (PDF) that incorporates all available statistical information from both the observations in a likelihood PDF and the prior knowledge in a prior PDF. Researchers can extract statistics of interests about the unknown parameters either by directly sampling the posterior PDF utilizing approaches like Markov chain Monte Carlo methods [Kaipio and Somersalo, 2006, Matheron, 2012] or by locally approximating the posterior PDF with an easy-to-study Gaussian PDF [Bui-Thanh et al., 2013, Osypov et al., 2013, Zhu et al., 2016, Fang et al., 2018]. These statistics reflect the degree of confidence about the unknown model parameters and allow researchers to identify areas with high/low reliability in the model.

The prior knowledge plays a key role in the Bayesian inverse problem. Conventional prior knowledge is hand-crafted based on researchers’ empirical observations about the subsurface structure. For instances, because subsurface structure images are generally piecewise smooth and sparse after wavelet or curvelet transformations, image priors that constrain the sparsity of wavelet or curvelet coefficients [Ying et al., 2005, Li et al., 2012, Tu and Herrmann, 2015] or spatial gradients [Haber et al., 2000] are widely utilized. While these hand-crafted priors can regularize the unknown parameters when solving deterministic inverse problems, they are usually too generic, in that prior models generated with these prior PDF cannot describe detailed structures too well.

Recent developments in deep convolutional neural networks (DCNN) [Krizhevsky et al., 2012, LeCun et al., 2015] provide researchers in the field of image/signal processing with a new way to design the prior knowledge, i.e. a learning-based prior generator. Instead of designing features by hands, learning-based approaches train DCNNs that straightforwardly learns features from existing training sets. The ultimate trained DCNN is designed to be able to create images sharing the same spatial distributions with the real/natural samples. Empirically, many researchers have shown the superior performance of the learning-based priors over the hand-crafted priors when dealing deterministic image related inverse problems such as denoising, super resolution, and inpainting [Yeh et al., 2016, Rick Chang et al., 2017, Ledig et al., 2017].

The study of applying DCNN to geophysical problems has been quite active during the last five years. Many geophysicists have shown potential applications to problems like data denoising [Yu et al., 2019] and deblending [Richardson and Feller, 2019], surface-related multiples removal [Siahkoochi et al., 2019], low-frequency data extrapolation [Sun and Demanet, 2018], and travel-time tomography [Araya-Polo et al., 2018]. Instead of training a generator as the prior information, most of the reported works aims at a network mapping input images/data to output images/data of interests. For example, a DCNN converts the data with noise to the data without noise for denoising problems and a DCNN converts travel time information to

subsurface velocity images for traveltimes tomography. Despite their potential superior performance, these specifically-trained DCNN are designed to solve specific problems and usually require to retrain the DCNN for new problems.

In this work, we propose a Bayesian seismic inversion framework with prior information learned with a deep generative adversarial network (DGAN) [Goodfellow et al., 2014]. One appealing property of DGAN for Bayesian inversion lies in DGAN’s potential capability in generating perceptually appealing high-dimensional images from a low-dimensional latent parameter space, which has been empirically proven in many applications [Ledig et al., 2017]. This property motivates us to train a GAN that can map a low-dimensional normal distribution to the high-dimensional distribution of the training images. Employing such a DGAN to the Bayesian inversion framework has three main advantages. First, the trained DGAN enables us to conduct the Bayesian inversion in the low-dimensional latent space instead of the original high-dimensional image space. As a result, we significantly reduce the complexity of the problem and mitigate the curse of dimensionality that Bayesian inversions typically suffer from. Secondly, since the DGAN is directly learned from the existing training images, it can capture more detailed statistical information about the subsurface images compared to the conventional hand-crafted priors. This enables the DGAN prior generator to produce geologically plausible prior subsurface images that conventional hand-crafted priors cannot produce. Finally, according to the aforementioned setting of the training procedure, the prior PDF in the low-dimensional latent space is the simple normal distribution without any complicated design by hands. This property motivates us to apply the dimension-free Markov chain Monte Carlo (MCMC) method – preconditioner Crank-Nikolski method [Cotter et al., 2013] to efficiently sample the posterior PDF and extract the statistics of interests.

The paper is organized as follows. First, we introduce the basic concepts about the Bayesian inference method. Following that, we introduce the proposed deep GAN prior model generator and the pCN algorithm. Then we present a short introduction about the traveltimes tomography and full waveform inversion. After that, we study the feasibility and generality of the proposed approach by two numerical examples. At last, we finalize the paper with a conclusion and discussion.

Methodology

Bayesian inference

The Bayesian inference method is a widely used approach to solving seismic statistical inverse problems. Unlike deterministic approaches that seek a best data-fit model, the Bayesian inference method aims at a comprehensive statistical description of the unknown parameters. For this purpose, the Bayesian inference method constructs a posterior PDF of the model parameters \mathbf{m} that integrates statistical information from the forward map $F(\mathbf{m})$, observed data \mathbf{d}_{obs} , and researcher’s prior knowledge. According to the Bayes’ law [Bayes et al., 1763], the posterior PDF $\rho_{\text{post}}(\mathbf{m}|\mathbf{d}_{\text{obs}})$ of \mathbf{m} given \mathbf{d}_{obs} is proportional to the product of a likelihood PDF $\rho_{\text{like}}(\mathbf{d}_{\text{obs}}|\mathbf{m})$ of \mathbf{d}_{obs} given \mathbf{m} and a prior PDF $\rho_{\text{prior}}(\mathbf{m})$ of \mathbf{m} as follows:

$$\rho_{\text{post}}(\mathbf{m}|\mathbf{d}_{\text{obs}}) \propto \rho_{\text{like}}(\mathbf{d}_{\text{obs}}|\mathbf{m})\rho_{\text{prior}}(\mathbf{m}). \quad (1)$$

The prior PDF $\rho_{\text{prior}}(\mathbf{m})$ describes one’s prior knowledge and beliefs in the unknown model parameters, and the likelihood PDF $\rho_{\text{like}}(\mathbf{d}_{\text{obs}}|\mathbf{m})$ describes the probability of observing data \mathbf{d}_{obs} given the model \mathbf{m} . If we assume that the observed data contain additive Gaussian noise ϵ from the distribution $\mathcal{N}(0, \Sigma)$ with the covariance matrix Σ as follows:

$$\mathbf{d}_{\text{obs}} = F(\mathbf{m}) + \epsilon, \quad (2)$$

the likelihood PDF has the following expression:

$$\rho_{\text{like}}(\mathbf{d}_{\text{obs}}|\mathbf{m}) \propto \exp\left(-\frac{1}{2}\|F(\mathbf{m}) - \mathbf{d}_{\text{obs}}\|_{\Sigma^{-1}}^2\right). \quad (3)$$

With the posterior PDF in hands, we can extract statistical properties of interests including the maximum a posterior (MAP) estimate, the model covariance matrix, the model standard deviation (STD), and the confidence interval of \mathbf{m} . To conduct a successful Bayesian inversion, the primary issues are the construction of the posterior PDF and a computationally tractable method to extract statistical properties from the posterior PDF.

Deep GAN prior model generator

The selection of the prior PDF plays an important role in the statistical inversion. Conventional choices of the prior PDF are based on empirical observations of subsurface structures including the smoothness of the subsurface image, the sparsity of the coefficient of the image in certain transformed domain, or the total variation of the subsurface images. However, these kinds of prior information are typically too generic and cannot represent all the available prior information. Unlike the conventional prior choices, recent developments in learning-based methods in the field of imaging processing show a potential capability in estimating the prior distribution for images. In this work, we propose to use DGAN to extract prior information from existing subsurface images straightforwardly.

Given a training set \mathcal{M} , DGAN aims at generating artificial samples that share the same statistics as the training natural samples. A standard DGAN consists of two neural networks namely the generator $G(\mathbf{x}; \Theta_G)$ and the discriminator $D(\mathbf{m}; \Theta_D)$, which are parameterized by vectors Θ_G and Θ_D . The generator creates candidates of interests from vectors \mathbf{x} 's obeying a distribution \mathcal{X} in a low-dimensional latent space, while the discriminator distinguishes these generated candidates from the true data distribution. The two networks contest with each other in a game. The discriminator tries to decrease the error rate of the discriminative network, while the generator tries to increase the error rate.

During the training step, DGAN solves the following *minmax* optimization problem:

$$\min_{\Theta_G} \max_{\Theta_D} \mathbb{E}_{\mathbf{m} \sim \mathcal{M}} \log(D(\mathbf{m}; \Theta_D)) + \mathbb{E}_{\mathbf{x} \sim \mathcal{X}} [\log(1 - D(G(\mathbf{x}; \Theta_G); \Theta_D))], \quad (4)$$

where the expectation for \mathbf{m} is over all the training samples \mathcal{M} , and the expectation for \mathbf{x} is over the distribution \mathcal{X} . In general \mathcal{X} can be a set of images or data or any reasonable predefined distributions such as uniform and normal distributions. In this work, we select \mathcal{X} to be a standard Gaussian distribution $\mathcal{N}(0, \mathbf{I})$ for a series of computational advantages in the following statistical inversion, which we will discuss later.

The trained generator $G(\mathbf{x}; \Theta_G)$ defines an n_1 -manifold in the original \mathbb{R}^{n_g} space, where $n_1 \ll n_g$. With the generator $G(\mathbf{x}; \Theta_G)$, we can reformulate the original Bayesian inversion problem on the n_1 -manifold as follows,

$$\begin{aligned} \rho_{\text{post}}(\mathbf{x}|\mathbf{d}_{\text{obs}}) &\propto \rho_{\text{like}}(\mathbf{d}_{\text{obs}}|\mathbf{x})\rho_{\text{prior}}(\mathbf{x}) \\ &\propto \exp\left(-\frac{1}{2}\|F(G(\mathbf{x}; \Theta_G)) - \mathbf{d}_{\text{obs}}\|_{\Sigma^{-1}}^2 - \frac{1}{2}\|\mathbf{x}\|^2\right). \end{aligned} \quad (5)$$

Compared to the original posterior PDF in Equation 1, studying the new posterior PDF $\rho_{\text{post}}(\mathbf{x}|\mathbf{d}_{\text{obs}})$ has four advantages. First, we significantly reduce the dimensionality from n_g to n_1 , which makes the new Bayesian inversion problem less suffer from the curse of dimensionality. Secondly, the Gaussian prior distribution is directly from the definition of the DGAN generator without any hand designs. Thirdly, images generated from the prior distribution share the same statistics as the training data set. Finally, the Gaussian prior distribution enables us to apply the computationally efficient Markov chain Monte Carlo (MCMC) method - preconditioned Crank-Nicolson algorithm (pCN) to sample the posterior PDF, which we will introduce in the next subsection.

Sample the posterior PDF by pCN

The pCN algorithm is an efficient dimension-free MCMC type method, whose efficiency has been proven for problems with Gaussian prior PDFs. The idea of this algorithm is that using the prior distribution to generate the random walking direction and followed by an accept-reject procedure. The pseudo code of the algorithm is shown in Algorithm 1. Indeed, the pCN algorithm is a modification of the conventional random walk algorithm, which generates a new proposal sample by $\mathbf{y}^{(k)} = \mathbf{x}^{(k-1)} + \beta\mathbf{r}^{(k)}$ with $\mathbf{r}^{(k)} \sim \mathcal{N}(0, \mathbf{I}_{n_1 \times n_1})$. The pCN algorithm differs only slightly from the random walk method: the proposal sample is generated by $\mathbf{y}^{(k)} = \sqrt{1 - \beta^2}\mathbf{x}^{(k-1)} + \beta\mathbf{r}^{(k)}$. This slight modification results in that the pCN algorithm is robust to the increase of the dimensionality caused by the discretization of the physical domain. For this important property, we use the pCN algorithm to sample the posterior PDF in Equation 5 in this work.

Algorithm 1 pCN

1. Set $k = 0$ and pick $\mathbf{x}^{(0)}$.
 2. Propose $\mathbf{y}^{(k)} = \sqrt{1 - \beta^2}\mathbf{x}^{(k-1)} + \beta\mathbf{r}^{(k)}$, $\mathbf{r}^{(k)} \sim \mathcal{N}(0, \mathbf{I}_{n_1 \times n_1})$.
 3. Set $\mathbf{x}^{(k)} = \mathbf{y}^{(k)}$ with probability $a(\mathbf{x}^{(k-1)}, \mathbf{y}^{(k)}) = \min(1, \frac{\rho_{\text{like}}(\mathbf{d}_{\text{obs}}|\mathbf{y}^{(k)})}{\rho_{\text{like}}(\mathbf{d}_{\text{obs}}|\mathbf{x}^{(k-1)})})$,
or set $\mathbf{x}^{(k)} = \mathbf{x}^{(k-1)}$
 4. $k \rightarrow k + 1$.
-

Seismic inversion applications

In this paper, we apply the proposed Bayesian approach with the deep GAN prior model generator to two seismic inversion applications – traveltime tomography and full waveform inversion.

Traveltime tomography

Seismic traveltime tomography has been serving as an important tool to extract elastic properties of the Earth interior in local, regional, and global scales from seismic data [Aki et al., 1977, Romanowicz, 1979, Dziewonski, 1984, Nolet, 2012]. For the sake of efficiency, only the travel time information on the seismic recordings is used. The goal of the traveltime tomography is to seek a velocity model that can fit the observed travel time data by linearized or nonlinear inversion methods. Under the high frequency approximation, we can simplify the wave equation and compute the travel time field for a certain model by solving the following Eikonal equation:

$$\|\nabla\tau\|^2 = \mathbf{m}, \quad (6)$$

where the operator ∇ denotes the gradient operator, the vector τ is the travel time field and the vector \mathbf{m} is the squared slowness. For a detail review, the readers should refer to Nolet [2012].

Full waveform inversion

Full waveform inversion (FWI) is another important seismic imaging approach in the fields of both global seismology and oil and gas exploration. FWI utilizes all kinds of waveforms including the reflection and refraction waves to reconstruct the subsurface velocity model. Through iteratively comparing the observed and predicted data and updating the model, FWI tries to find the best data-fit model. The forward map of FWI requires to solve the following acoustic wave equation:

$$(-\mathbf{m}\partial_{tt} + \Delta)\mathbf{u} = \mathbf{q}, \quad (7)$$

where the symbol Δ denotes the Laplacian operator, the vectors \mathbf{u} and \mathbf{q} denote the wavefield and source function, respectively. After discretization, we can formulate the forward operator as:

$$F_{\text{FWI}}(\mathbf{m}) = \mathbf{P}\mathbf{A}(\mathbf{m})^{-1}\mathbf{q}, \quad (8)$$

where the matrix $\mathbf{A}(\mathbf{m})$ is the discretization of the partial differential operator $-\mathbf{m}\partial_{tt} + \Delta$, and the operator \mathbf{P} projects the wavefields on the receiver locations. See Virieux and Operto [2009] for an extensive overview of state-of-the-art approaches to FWI.

Numerical examples

We assess the effectiveness of the proposed Bayesian seismic inversion scheme with deep generator priors via applications of the traveltimes tomography and the full waveform inversion. To conduct the experiments, we first train an eight-block generative network and a three-block discriminative network using velocity slices extracted from the 3D overthrust model. With the trained deep generator prior, we conduct statistical inversions on both overthrust model and Sigsbee model.

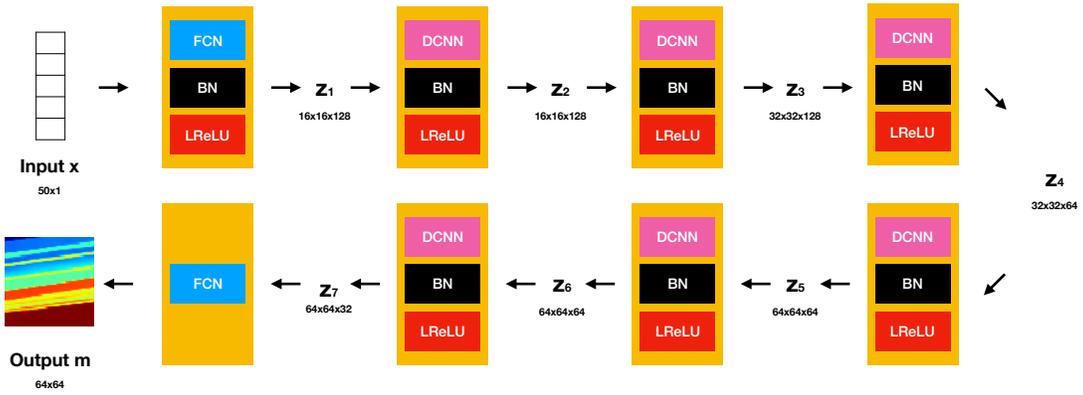
Training the network architecture

Network architectures. In this work, we design an eight-block generative network and a three-block discriminative network. The architectures of the two networks are shown in Figure 1. The generative network is a deep deconvolutional neural network that maps a random vector $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_{50 \times 50})$ to a 64×64 velocity image. We first use a fully-connected network to map the input data to $128 \ 16 \times 16$ features, followed by a batch normalization and a leaky ReLU activation. The fully-connected network, batch normalization and leaky ReLU activation form the first block of the generative network. After the first block, we add 6 additional deconvolutional based blocks. Each block contains a deconvolutional neural network, a batch normalization and a leaky ReLU activation. The output of the 6 blocks is $32 \ 64 \times 64$ features. The last block of the generative network is a deconvolutional neural network with 5×5 kernels that generates the output 64×64 image.

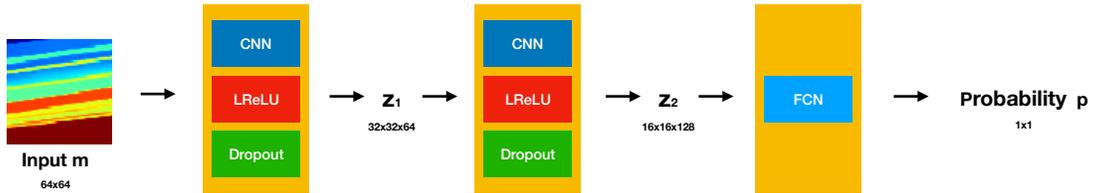
The discriminative network starts with the 64×64 velocity image and outputs a value between 0 and 1 that indicates the probability of the input image belonging to the training set. The discriminative network is a three-block convolutional neural network. The first block contains a convolutional neural network with 5×5 kernels and 64 features, a leaky ReLU activation, and a dropout operator with a dropout ratio of 0.3. The second block consists of a convolutional neural network with 5×5 kernels and 128 features, a leaky ReLU activation, and a dropout operator with a dropout ratio of 0.3. The last block consists of a fully-connected network that outputs the probability.

Training. We train the generative and discriminative networks with velocity models extracted from the 3D overthrust model. The size of the 3D overthrust model is $20\text{km} \times 20\text{km} \times 4.65\text{km}$ with a spatial spacing of 0.025km , yielding a 3D volume with the size of $801 \times 801 \times 187$. We extract 20,000 velocity models with the size of 64×64 from the 3D volume and rotate them with random angles, yielding 80,000 velocity models. We use the 80,000 models as the training set to train the generative and discriminative networks.

We use the Adam optimizer to train the networks. We select the momentum parameter $\beta = 0.9$, mini-batch size $L_b = 256$, and initial rate $\mu = 1e-4$. We conduct the training with TensorFlow interface [Abadi et al., 2016] and Keras toolbox [Chollet et al., 2015] on an NVIDIA GeForce GTX 1080 GPU, 12 GB RAM. We allow 2000 epochs that takes around 24 hours for training. Figure 2 shows a comparison of the images from the training set and the generative network. Clearly, generative images share a similar distribution with the training images.



(a) Generator

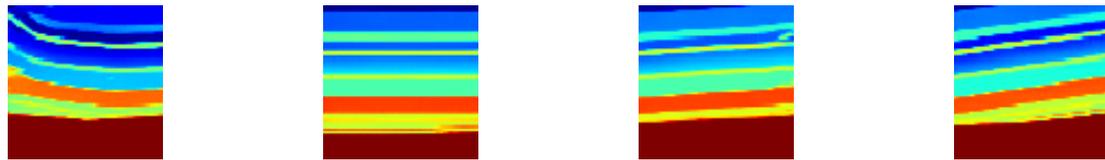


(b) Discriminator

Figure 1: The framework of DGAN. Architectures of the generative and discriminative networks. FCN – fully connected network, BN – batch normalization, LReLU – leaky ReLU, DCNN – deconvolutional neural network, CNN – convolutional neural network.



(a) Training images



(b) Generative images

Figure 2: Comparison of images from (a) training set and (b) generative network.

Statistical inversion with overthrust models

We conduct the first statistical inversion with a testing velocity model extracted from the 3D overthrust model that does not belong to the training set. Since both the testing model and training models are extracted from the 3D overthrust model, their spatial distributions should share a strong similarity and the prior generator should provide quite good prior information for the inversion. We first conduct the statistical inversion with the traveltimes tomography. We mimic a cross-hole case, where 30 sources are located at the left side of the model and 60 receivers are located at the right side of the model. The model size is $258 \text{ m} \times 258 \text{ m}$. We rescale the velocity model so that the range of the velocity is $[2.4 \text{ km/s}, 2.7 \text{ km/s}]$. The observed data contain 0.5% additive Gaussian noise. Figure 3a shows the true model. We first use the pCN method to generate 100,000 samples and then throw away the first 20,000 samples. Finally, we compute statistical information from the left 80,000 samples. Figure 3b shows the MAP estimate from the 80,000 samples. We can find that the MAP estimate matches the true model quite well. Figure 3c shows the point-wise standard deviation (STD) of the 80,000 samples. We can observe high STDs at boundaries of layers and low STDs inside layers. Figure 4 shows histograms of samples at three positions. Clearly, at the boundary of layers (Figures 4a and 4c), the velocity has big uncertainties, while inside layers (Figure 4b), the velocity has small uncertainties.

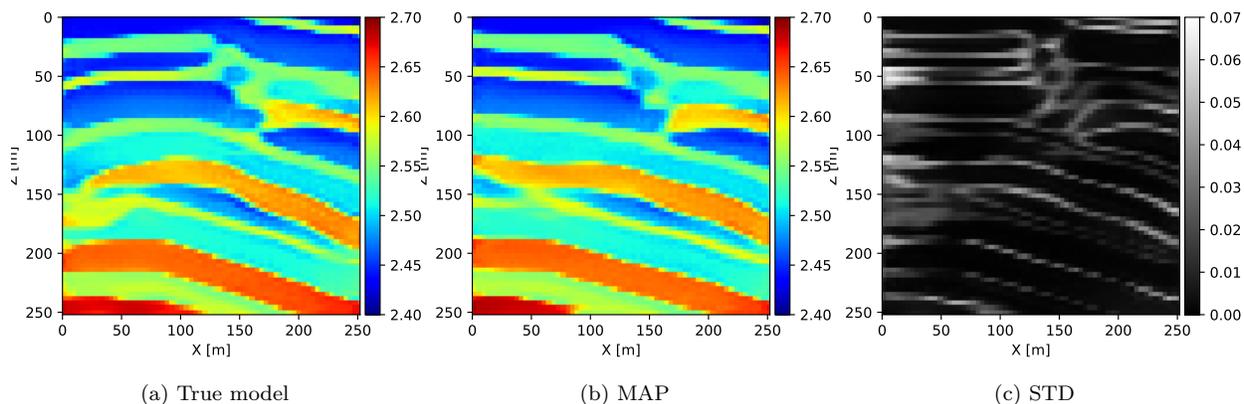


Figure 3: (a) True model for the traveltimes tomography. (b) MAP of the statistical inversion. (c) Standard deviation of the 80,000 samples.

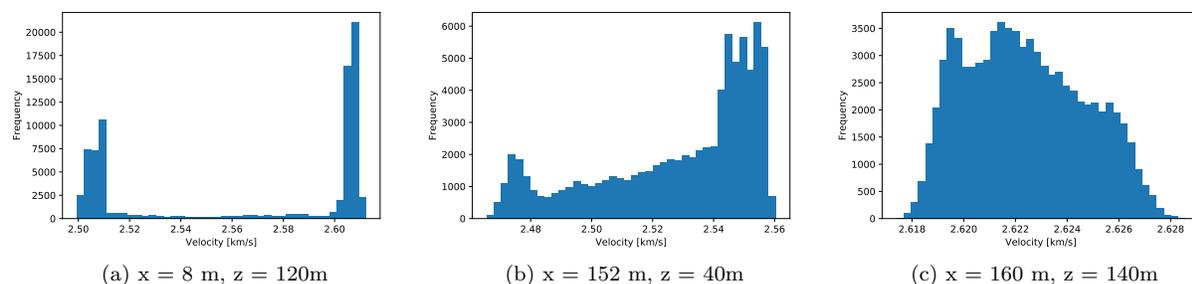


Figure 4: Histograms of samples at (a) $x = 8 \text{ m}$, $z = 120 \text{ m}$, (b) $x = 152 \text{ m}$, $z = 40 \text{ m}$, and (c) $x = 160 \text{ m}$, $z = 140 \text{ m}$.

Then we conduct the statistical inversion with FWI. Figure 5a shows the true model, whose size is $1.64 \text{ km} \times$

1.64 km. We place 8 sources at the surface of the model and place 64 receivers at the same depth. We use a Ricker wavelet centered at 10 Hz to simulate the data with 1% additive Gaussian noise. With the consideration of the computational cost, we use the pCN to generate 30,000 samples and throw away the first 10,000 samples. We compute statistics of interests with the remaining 20,000 samples. As a comparison, Figure 10b show the result of the conventional FWI. We can observe that the result obtained by the proposed method presents a better resolution than the conventional FWI that directly works on the physical domain. Figure 5c shows the MAP estimate of the 20,000 samples. Clearly, the MAP estimate matches the true model quite well. Figure 5d illustrates the point-wise STD of the 20,000 samples. Similar to the result of tomography, we can observe high STDs at boundaries of layers and low STDs inside layers. Figure 6 shows histograms of samples at three positions. Again, at boundaries of layers, the velocity is very uncertain, while inside layers, the velocity is much less uncertain.

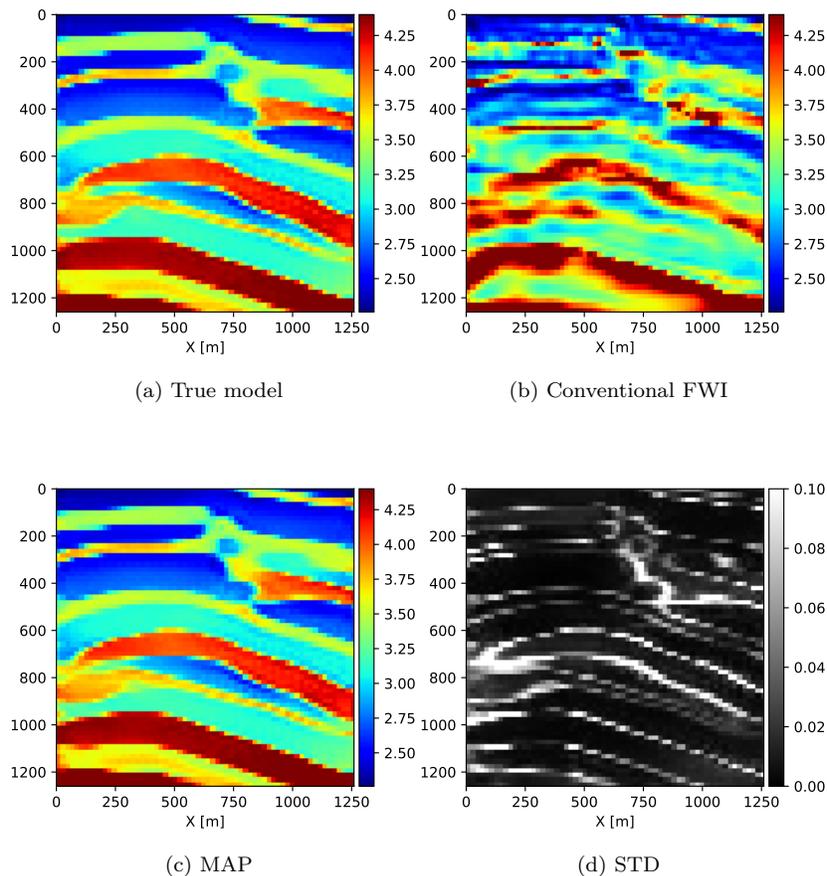


Figure 5: (a) True model for FWI. (b) Result of the conventional FWI. (c) MAP of the statistical inversion. (d) Standard deviation of the 20,000 samples.

Statistical inversion with Sigsbee models

We conduct the second examples with a velocity model extracted from Sigsbee model. This example is challenging, since the spatial distribution of the Sigsbee model is quite different from the overthrust model.

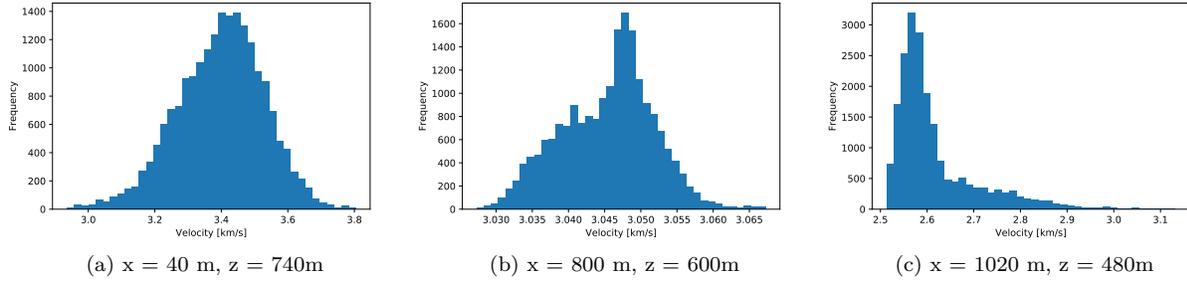


Figure 6: Histograms of samples at (a) $x = 40$ m, $z = 740$ m, (b) $x = 800$ m, $z = 600$ m, and (c) $x = 1020$ m, $z = 480$ m.

We first conduct the traveltimes tomography example. The settings for the experiment are the same with the ones in the previous example. Figures 7a, 7b, and 7c show the true model, MAP estimate and standard deviation, respectively. We can observe clear difference between the true model and the MAP estimate. However, if we compare the observed data and predicted data (Figure 8), we can observe that the difference between them reaches the noise level (see Figures 8c and 9). This implies that the proposed method finds a best data-fit solution in the range of the trained GAN.

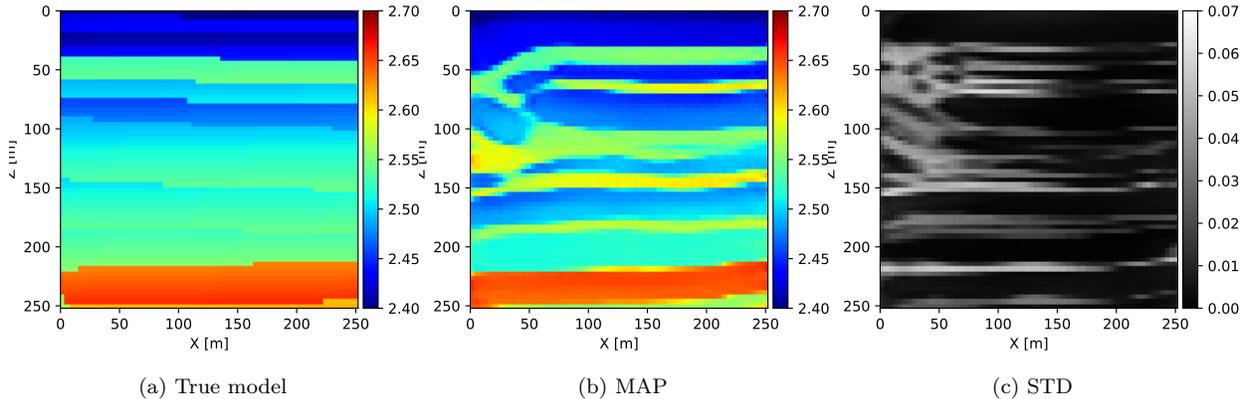


Figure 7: (a) True model for the traveltimes tomography. (b) MAP of the statistical inversion. (c) Standard deviation of the 80,000 samples.

Then we conduct the statistical inversion with FWI. The experimental settings are the same with the ones used in the previous example. Figures 10a, 10b, 10c, and 10d show the true model, conventional FWI result, MAP estimate and standard deviation, respectively. Conventional FWI presents a result that is a smoothed, blurred and noisy version of the true model, while the proposed approach presents a sharp result with clear layers. However, we can observe clear differences between the true model and the MAP estimate. These differences are due to the fact that the target velocity is not in the range of the prior DGAN. To further analyze the inversion result, we compare the observed data and the predicted data obtained with the MAP estimate in Figures 11 and 12. Clearly, the data differences have reached the noise level, which means the MAP estimate is also a best data-fit solution.

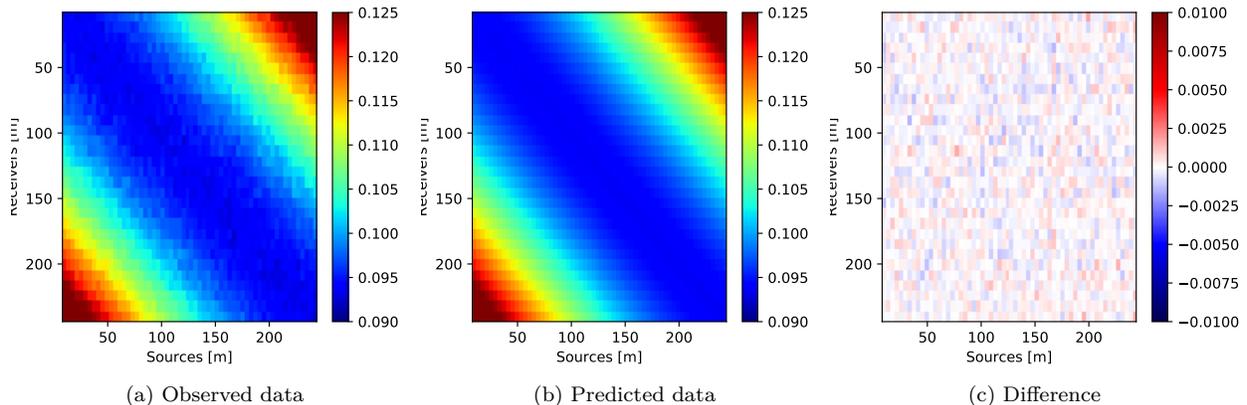


Figure 8: (a) Observed data. (b) Predicted data obtained with the MAP estimate. (c) Difference between the observed and predicted data.

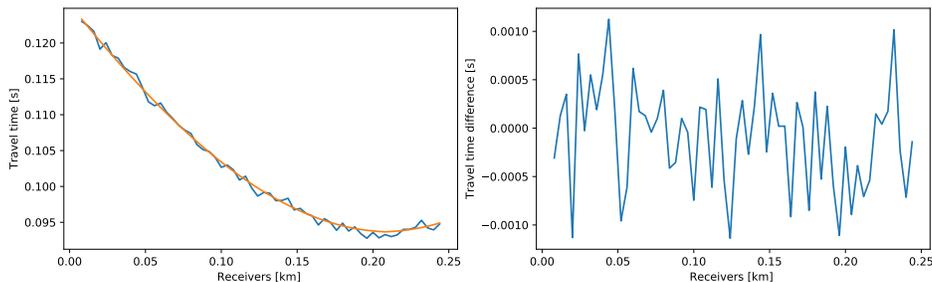


Figure 9: (a) Observed data (blue line) and predicted data (red line); (b) Difference between the observed and predicted data.

Conclusions and discussion

We present a Bayesian statistical seismic inversion framework with a deep generator prior. Compared to conventional Bayesian inversion with hand-crafted prior information, the proposed approach has three major advantages. First, DGAN enables us to conduct inversions in the low-dimensional latent space instead of the original high-dimensional image space, which significantly reduces the computational complexity. Secondly, the trained DGAN is able to generate artificial models that share the similar spatial distributions as the training images, which provides the access to a complete description of the available prior information. Thirdly, the prior distribution in the latent space is a normal distribution without any hand-crafted assumption. This enables us to apply the efficient pCN method to sample the posterior PDF.

Since seismic inverse problems are typically very ill-posed, the noisy observed data can be well explained by different models. Therefore, the selection of prior information has a strong influence to the final inversion result. Numerical examples illustrates the importance of the training dataset to the proposed Bayesian inversion with DGAN prior. When the target velocity model shares a similar spatial distribution with models in the training set, the proposed approach is able to produce an accurate best data-fit solution with a high resolution. While for the opposite situation, the proposed approach can only find a best data-fit solution in the range of the DGAN, which may differ from the true model.

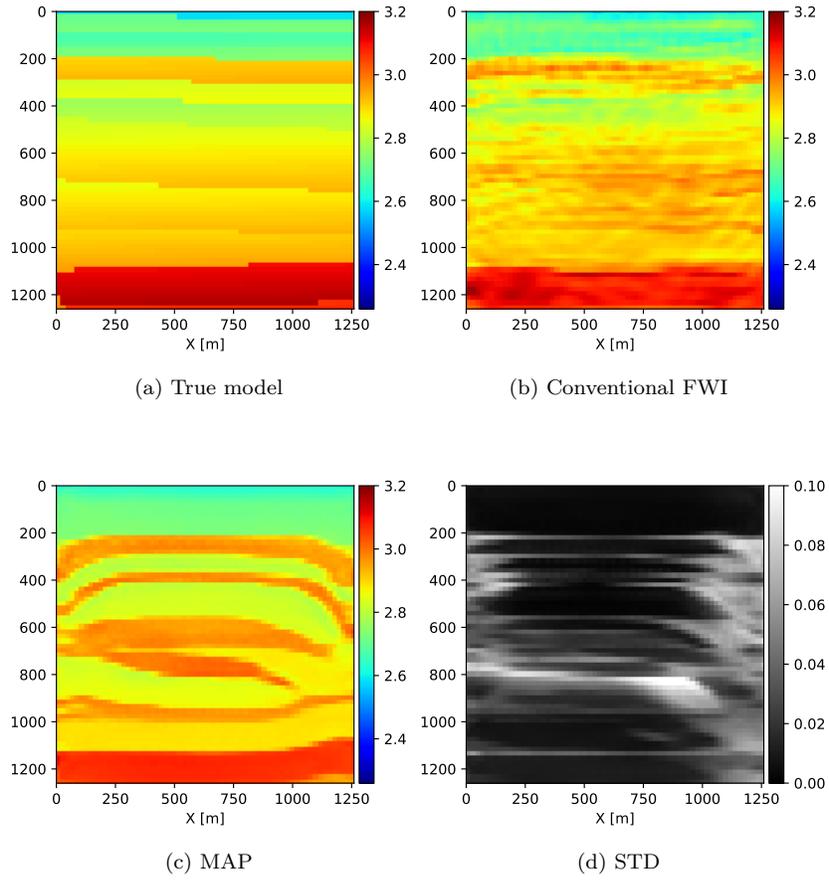


Figure 10: (a) True model for FWI. (b) Result of the conventional FWI. (c) MAP of the statistical inversion. (d) Standard deviation of the 20,000 samples.

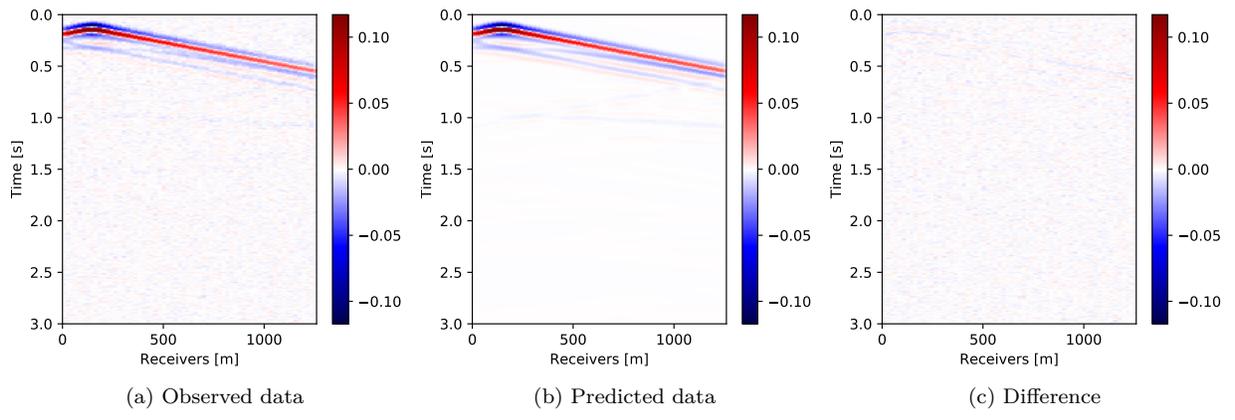


Figure 11: (a) Observed data. (b) Predicted data obtained with the MAP estimate. (c) Difference between the observed and predicted data.

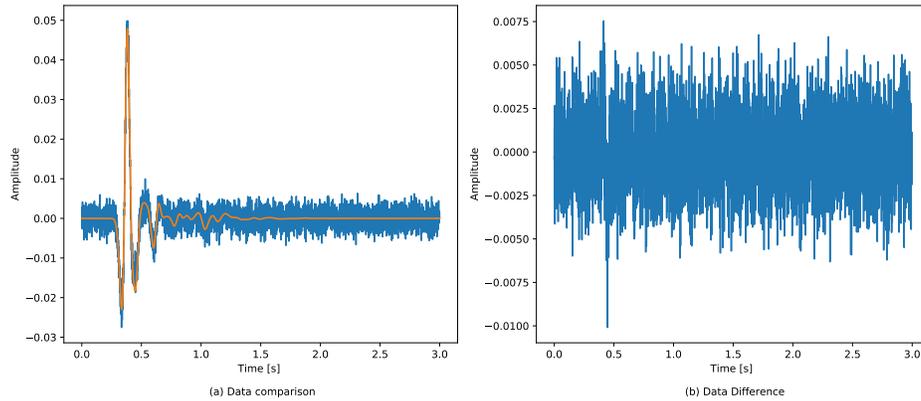


Figure 12: (a) Observed data (blue line) and predicted data (red line); (b) Difference between the observed and predicted data.

The proposed approach provides us with a solution to automatically encode prior knowledge about the subsurface structures from existing subsurface models. The robustness of the deep generator prior depends on the generality and diversity of the training set. Through employing new subsurface models, the deep generator prior can keep on evolving to enhance its robustness. This property can help researchers automatically accumulate their knowledge about the subsurface structures in their research career.

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283, 2016.
- Keiiti Aki, Anders Christofferson, and Eystein S Husebye. Determination of the three-dimensional seismic structure of the lithosphere. *Journal of Geophysical Research*, 82(2):277–296, 1977.
- Mauricio Araya-Polo, Joseph Jennings, Amir Adler, and Taylor Dahlke. Deep-learning tomography. *The Leading Edge*, 37(1):58–66, 2018.
- Thomas Bayes, Richard Price, and John Canton. An essay towards solving a problem in the doctrine of chances. pages 370–418, 1763.
- Tan Bui-Thanh, Omar Ghattas, James Martin, and Georg Stadler. A computational framework for infinite-dimensional Bayesian inverse problems part I: The linearized case, with application to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494–A2523, 2013.
- François Chollet et al. Keras. <https://keras.io>, 2015.
- Simon L Cotter, Gareth O Roberts, Andrew M Stuart, David White, et al. Mcmc methods for functions: modifying old algorithms to make them faster. *Statistical Science*, 28(3):424–446, 2013.
- AJW Duijndam. Bayesian estimation in seismic inversion. part ii: Uncertainty analysis1. *Geophysical Prospecting*, 36(8):899–918, 1988.

- Adam M Dziewonski. Mapping the lower mantle: determination of lateral heterogeneity in p velocity up to degree and order 6. *Journal of Geophysical Research: Solid Earth*, 89(B7):5929–5952, 1984.
- Gregory Ely, Alison Malcolm, and Oleg V Poliannikov. Assessing uncertainties in velocity models and images with a fast nonlinear uncertainty quantification method. *Geophysics*, 83(2)(2):R63–R75, 2017.
- Zhilong Fang, Curt Da Silva, Rachel Kuske, and Felix J Herrmann. Uncertainty quantification for inverse problems with weak partial-differential-equation constraints. *Geophysics*, 83(6):R629–R647, 2018.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- Eldad Haber, Uri M Ascher, and Doug Oldenburg. On optimization techniques for solving nonlinear inverse problems. *Inverse problems*, 16(5):1263, 2000.
- Jari Kaipio and Erkki Somersalo. *Statistical and Computational Inverse Problems*. Springer Science & Business Media, December 2006. ISBN 0387220739.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- Patrick Lailly et al. The seismic inverse problem as a sequence of before stack migrations. In *Conference on Inverse Scattering*, 1983.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.
- Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- Xiang Li, Aleksandr Y Aravkin, Tristan van Leeuwen, and Felix J Herrmann. Fast randomized full-waveform inversion with compressive sensing. *Geophysics*, 77(3):A13–A17, 2012.
- Georges Matheron. *Estimating and choosing: an essay on probability in practice*. Springer Science & Business Media, 2012.
- Guust Nolet. *Seismic tomography: with applications in global seismology and exploration geophysics*, volume 5. Springer Science & Business Media, 2012.
- K Osypov, D Nichols, M Woodward, O Zdraveva, CE Yarman, et al. Uncertainty and resolution analysis for anisotropic tomography using iterative eigendecomposition. In *2008 SEG Annual Meeting*. Society of Exploration Geophysicists, 2008.
- Konstantin Osypov, Yi Yang, Aimé Fournier, Natalia Ivanova, Ran Bachrach, Can Evren Yarman, Yu You, Dave Nichols, and Marta Woodward. Model-uncertainty quantification in seismic tomography: method and applications. *Geophysical Prospecting*, 61(6):1114–1134, 2013.
- R Gerhard Pratt. Seismic waveform inversion in the frequency domain, Part 1: Theory and verification in a physical scale model. *Geophysics*, 64(3):888–901, 1999.
- Alan Richardson and Caelen Feller. Seismic data denoising and deblending using deep learning. *arXiv preprint arXiv:1907.01497*, 2019.
- JH Rick Chang, Chun-Liang Li, Barnabas Poczós, BVK Vijaya Kumar, and Aswin C Sankaranarayanan. One network to solve them all—solving linear inverse problems using deep projection models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5888–5897, 2017.

- Barbara A Romanowicz. Seismic structure of the upper mantle beneath the united states by three-dimensional inversion of body wave arrival times. *Geophysical Journal International*, 57(2):479–506, 1979.
- Malcolm Sambridge, K Gallagher, Andrew Jackson, and Peter Rickwood. Trans-dimensional inverse problems, model comparison and the evidence. *Geophysical Journal International*, 167(2):528–542, 2006.
- John A Scales and Luis Tenorio. Prior information and uncertainty in inverse problems. *Geophysics*, 66(2):389–397, 2001.
- Ali Siahkoobi, Mathias Louboutin, and Felix J Herrmann. The importance of transfer learning in seismic modeling and imaging. *Geophysics*, 84(6):A47–A52, 2019.
- Hongyu Sun and Laurent Demanet. Low frequency extrapolation with deep learning. In *SEG Technical Program Expanded Abstracts 2018*, pages 2011–2015. Society of Exploration Geophysicists, 2018.
- Albert Tarantola and Bernard Valette. Inverse problems = quest for information. *Journal of Geophysics*, 50:159–170, 1982a.
- Albert Tarantola and Bernard Valette. Generalized nonlinear inverse problems solved using the least squares criterion. *Reviews of Geophysics*, 20(2):219–232, 1982b. ISSN 1944-9208. doi: 10.1029/RG020i002p00219. URL <http://dx.doi.org/10.1029/RG020i002p00219>.
- Jeroen Tromp. Seismic wavefield imaging of earth’s interior across scales. *Nature Reviews Earth & Environment*, pages 1–14, 2019.
- Ning Tu and Felix J. Herrmann. Fast imaging with surface-related multiples by sparse inversion. *Geophysical Journal International*, 201(1):304–317, 2015. URL <https://www.slim.eos.ubc.ca/Publications/Public/Journals/GeophysicalJournalInternational/2014/tu2014fis/tu2014fis.pdf>.
- Jean Virieux and Stéphane Operto. An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6)(6):WCC1–WCC26, 2009. doi: 10.1190/1.3238367. URL <http://dx.doi.org/10.1190/1.3238367>.
- Raymond Yeh, Chen Chen, Teck Yian Lim, Mark Hasegawa-Johnson, and Minh N Do. Semantic image inpainting with perceptual and contextual losses. *arXiv preprint arXiv:1607.07539*, 2:3, 2016.
- Lexing Ying, Laurent Demanet, and Emmanuel Candes. 3D discrete curvelet transform. In *Wavelets XI*, volume 5914, page 591413. International Society for Optics and Photonics, 2005.
- Siwei Yu, Jianwei Ma, and Wenlong Wang. Deep learning for denoising. *Geophysics*, 84(6):V333–V350, 2019.
- Hejun Zhu, Siwei Li, Sergey Fomel, Georg Stadler, and Omar Ghattas. A Bayesian approach to estimate uncertainty for full-waveform inversion using a priori information from depth migration. *Geophysics*, 81(5):R307–R323, 2016.