

Deep Born Operator Learning for Reflection Tomographic Imaging

Zhao, Qingqing; Ma, Yanting; Boufounos, Petros T.; Nabi, Saleh; Mansour, Hassan

TR2023-029 May 06, 2023

Abstract

Recent developments in wave-based sensor technologies, such as ground penetrating radar (GPR), provide new opportunities for accurate imaging of underground scenes. Given measurements of the scattered electromagnetic wavefield, the goal is to estimate the spatial distribution of the permittivity of the underground scenes. However, such problems are highly ill-posed, difficult to formulate, and computationally expensive. In this paper, we propose a physics-inspired machine learning-based method to learn the wave-matter interaction under the GPR setting. The learned forward model is combined with a learned signal prior to recover the permittivity distribution of the unknown underground scenes. We test our approach on a dataset of 400 permittivity maps with a three-layer background, which is challenging to solve using existing methods. We demonstrate via numerical simulation that our method achieves a 50% improvement in mean squared error over benchmark machine learning-based solvers for reconstructing layered underground scenes.

IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2023

DEEP BORN OPERATOR LEARNING FOR REFLECTION TOMOGRAPHIC IMAGING

Qingqing Zhao¹, Yanting Ma², Petros Boufounos², Saleh Nabi², Hassan Mansour²

¹Department of Electrical Engineering, Stanford University, Stanford, CA 94305, USA

²Mitsubishi Electric Research Laboratories, 201 Broadway, Cambridge, MA 02140, USA
cyanzhao@stanford.edu, {yma, petrosb, nabi, mansour}@merl.com

ABSTRACT

Recent developments in wave-based sensor technologies, such as ground penetrating radar (GPR), provide new opportunities for accurate imaging of underground scenes. Given measurements of the scattered electromagnetic wavefield, the goal is to estimate the spatial distribution of the permittivity of the underground scenes. However, such problems are highly ill-posed, difficult to formulate, and computationally expensive. In this paper, we propose a physics-inspired machine learning-based method to learn the wave-matter interaction under the GPR setting. The learned forward model is combined with a learned signal prior to recover the permittivity distribution of the unknown underground scenes. We test our approach on a dataset of 400 permittivity maps with a three-layer background, which is challenging to solve using existing methods. We demonstrate via numerical simulation that our method achieves a 50% improvement in mean squared error over benchmark machine learning-based solvers for reconstructing layered underground scenes.

Index Terms— Underground Imaging, Diffraction Tomography, Full-waveform Inversion

1. INTRODUCTION

Ground Penetrating Radar (GPR) provides a non-destructive solution for underground utility mapping. The data acquisition process involves emitting known electromagnetic waves into the subsurface with transmission antennas located above the ground and then recording the scattered waves by receiver antennas also located above the ground. From the received scattered waves, data processing is applied to provide information about the physical properties of the subsurface.

Traditional GPR data processing is mainly based on migration methods originated from seismic imaging [1, 2]. These methods are efficient and can be implemented in real-time. However, images produced by migration methods only provide information about locations and rough shapes for the objects in the subsurface, but lack quantitative information about the electric permittivity contrast, from which materials of the objects can be inferred. More advanced data processing techniques are based on the principle of inverse scattering [2, 3]. These methods estimate the spatial distribution of the electric permittivity contrast of the subsurface, thus producing more informative images than migration methods. Assuming time harmonic waves, it has been well-established that the nonlinear wave-object relationship due to multiple scattering can be fully described by an integral equation. In the 2D scalar field setup, the integral equation coincides with the well-known

Lippmann-Schwinger equation. Inverse scattering based on the integral equation requires knowledge of the background and efficient methods to compute the corresponding Green’s function. Most existing results along this line consider homogeneous background, for which the Green’s function is easy to compute. A few results for the GPR setup have also considered layered background and proposed methods for computing the corresponding Green’s function [2, 4]. While the integral equation formulation is exact, inverting the integral equation can be computationally expensive when accurate optimization schemes are used to solve the inverse problem [5, 6, 7]. A commonly used approximation to the integral equation is based on the Born series expansion. Keeping higher order terms in the series expansion retains some level of nonlinearity of the system and results in better approximation; this is called iterative Born approximation (IBA) [3, 8].

More recently, machine learning-based approaches have emerged attempting to learn the wave-object relationship from training data. Recent works use the Graph Neural Network architecture or the Fourier Neural Operator (FNO) to model the time domain wave propagation [9, 10]. One main challenge of such learned simulations is the error accumulation during temporal unrolling.

In this paper, we propose to use a modification to the FNO architecture that is inspired by the IBA, to directly learn the wave-matter interaction in the frequency domain without prior knowledge of the background. Our learned forward model, called Born FNO (BFNO), is combined with a learned signal prior and isotropic total variation (TV) regularization to recover unknown underground structures from surface measurements. We test our approach under a three-layer background setup, and achieve a 50% improvement in mean squared error compared to benchmark machine learning-based solutions¹. While we focus on the GPR setting, the proposed approach is general and may be applied to other inverse scattering applications, including geophysical imaging [11, 12], optical tomography [13], etc.

2. PROBLEM FORMULATION

2.1. Background

In the 2D scalar wave setup with homogeneous background, the wave-object relationship is fully described by the Lippmann-Schwinger integral equation

$$u(\mathbf{x}) = u_{\text{in}}(\mathbf{x}) + k_b^2 \int_D g(\mathbf{x} - \mathbf{x}') u(\mathbf{x}') f(\mathbf{x}') d\mathbf{x}', \quad \forall \mathbf{x} \in \mathbb{R}^2 \quad (1)$$

where $u(\mathbf{x})$ is the total field, $u_{\text{in}}(\mathbf{x})$ is the incident field, ϵ_b is the permittivity of the background, $k_b = \sqrt{\epsilon_b} \omega / c$ is the wavenumber in

Q. Zhao conducted this work during a summer internship at MERL.

¹The complete code and dataset are available online at: <https://github.com/merlresearch/DeepBornFNO>

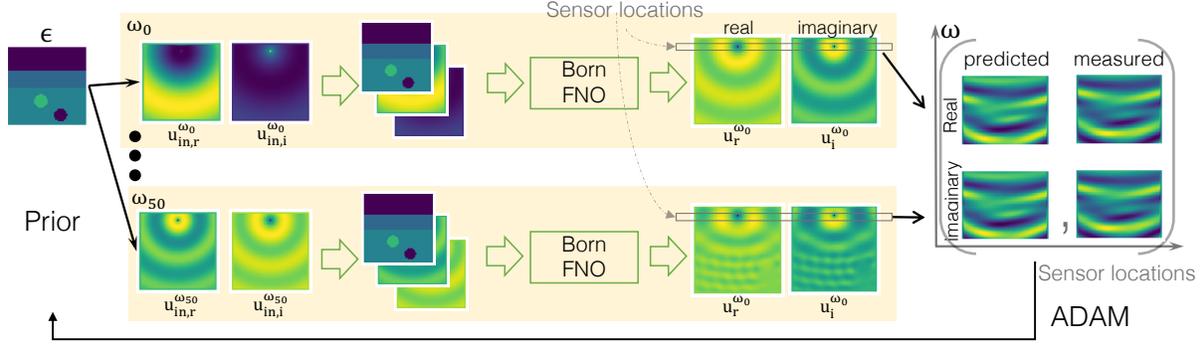


Fig. 1. The pipeline of our approach. The forward model Born FNO and the signal prior network are pre-trained with a dataset generated by gprMax [14]. At test time, the signal prior network first maps a latent code \mathbf{z} to the estimated permittivity ϵ , ϵ is then concatenated with $[\mathbf{u}_{\text{in},r}^{\omega_i}, \mathbf{u}_{\text{in},i}^{\omega_i}]$, the incident field with respect to the free space background at frequency ω_i , to predict the total field $[\mathbf{u}_r^{\omega_i}, \mathbf{u}_i^{\omega_i}]$. The latent code is optimized using stochastic gradient descent (SGD) to minimize the difference between the predicted total field and the measured total field across all sensor locations and 50 frequencies.

the background, c is the speed of light in vacuum, ω is the angular frequency, and g is the Green's function for the background. Let $\epsilon(\mathbf{x})$ denote the permittivity of the object. We assume that the permittivity contrast $f(\mathbf{x}) = \epsilon(\mathbf{x})/\epsilon_b - 1 = 0$ for $\mathbf{x} \notin D$, thus the domain for integration can be restricted to some bounded computational domain $D \subset \mathbb{R}^2$. Moreover, we assume that ϵ and ϵ_b are real-valued and do not depend on frequency. Let $\Gamma \subset \mathbb{R}^2$ denote the set of receiver locations. Our measurements are then the total field $u(\mathbf{x})$ with $\mathbf{x} \in \Gamma$. Note that to apply Eq. (1) with $\mathbf{x} \in \Gamma$ as the forward model, we need to first compute $u(\mathbf{x})$ with $\mathbf{x} \in D$, which should satisfy Eq. (1) with $\mathbf{x} \in D$ for the exact formulation.

A K^{th} -order IBA can be used to approximate the total field via:

$$u_{i+1}(\mathbf{x}) = u_{\text{in}}(\mathbf{x}) + k_b^2 \int_D g(\mathbf{x} - \mathbf{x}') f(\mathbf{x}') u_i(\mathbf{x}') d\mathbf{x}' \quad (2)$$

with $u_0(\mathbf{x}) = 0$ and $i = 1, \dots, K-1$. Under this approximation, the total field $u(\mathbf{x})$ with $\mathbf{x} \in D$ that is used to compute the measurements is replaced by $u_K(\mathbf{x})$. The forward model in discrete form under the K^{th} -order IBA can then be written as

$$\begin{aligned} \mathbf{y} &= \mathbf{u}_{\text{in}}^\Gamma + \mathbf{G}_{\Gamma D} \text{diag}(\mathbf{f}) \mathbf{u}_K + \mathbf{e} \\ \mathbf{u}_{i+1} &= \mathbf{u}_{\text{in}}^D + \mathbf{G}_{DD} \text{diag}(\mathbf{f}) \mathbf{u}_i, \quad i = 0, \dots, K-1 \end{aligned} \quad (3)$$

where \mathbf{G}_{AD} is the discretization of $k_b^2 g(\mathbf{x} - \mathbf{x}')$ with $\mathbf{x} \in A$, $\mathbf{x}' \in D$, \mathbf{u}_{in}^A is the discretization of $u_{\text{in}}(\mathbf{x})$ with $\mathbf{x} \in A$, $A \in \{\Gamma, D\}$, \mathbf{y} is the measured total field at the receiver locations Γ , and \mathbf{e} represents model mismatch.

The goal of the inverse problem is to estimate \mathbf{f} given \mathbf{y} . In addition, information about the source and the background is also necessary for computing the incident field and the Green's function. One way to estimate \mathbf{f} is to solve the optimization problem:

$$\hat{\mathbf{f}} = \underset{\mathbf{f}}{\text{argmin}} \sum_{\omega} \frac{1}{2} \|\mathbf{y}_{\omega} - \mathcal{Z}_{\omega}(\mathbf{f})\|_2^2 + \mathcal{R}(\mathbf{f}) \quad (4)$$

where \mathcal{Z}_{ω} is the forward operator that maps the permittivity contrast \mathbf{f} to the corresponding wavefield measurements \mathbf{y}_{ω} at frequency ω according to Eq. (3) and \mathcal{R} is the regularizer.

In the GPR setting, while it is reasonable to assume a layered structure for the background permittivity distribution $\epsilon_b(\mathbf{x})$, the depth and the permittivity value for each layer may be unknown. Moreover, computing the Green's function for a layered background is non-trivial. If instead free space is assumed as the background,

then the domain of integration cannot be restricted to a bounded region without careful treatment of the boundaries, as the layered structure extends outside the computational domain.

2.2. Proposed Approach

Given the difficulties in applying an analytical forward model for inverse scattering in the GPR setting, we propose to learn a forward model directly from data, where the architecture of our neural network is based on the recently proposed Fourier Neural Operator (FNO) and shares some similarity with the IBA Eq. (2). We call our learned forward model "Born FNO". Born FNO maps the spatially discretized permittivity distribution ϵ of the entire scene within the computational domain to the wavefield on the same discretization grid at a given frequency ω , i.e., $\mathbf{u}_{\omega} = \text{BFNO}(\epsilon, \omega)$. Since the problem is ill-posed due to the sparse and restricted angle measurements, we also learn a prior to represent the space of desired images. Our learned prior is a generative model that maps a low-dimensional latent code \mathbf{z} to the permittivity ϵ , i.e., $\epsilon = \mathcal{G}_{\epsilon}(\mathbf{z})$. Therefore, the optimization problem based on our formulation can be written as

$$\hat{\mathbf{z}} = \underset{\mathbf{z}}{\text{argmin}} \sum_{\omega} \frac{1}{2} \|\mathbf{y}_{\omega} - \mathbf{H}(\text{BFNO}(\mathcal{G}_{\epsilon}(\mathbf{z}), \omega))\|_2^2 + \mathcal{R}(\mathcal{G}_{\epsilon}(\mathbf{z})) \quad (5)$$

where \mathbf{H} is a sampling operator that selects the wavefield at the sensor locations and the estimated permittivity is obtained by $\hat{\epsilon} = \mathcal{G}_{\epsilon}(\hat{\mathbf{z}})$. $\mathcal{R}(\cdot)$ is the isotropic total variation regularizer. A detailed pipeline of our proposed approach can be seen in Fig. 1.

3. METHOD

3.1. Forward Model

We discretize our computational domain into a uniformly sampled 2D grid, D . The input to our forward model is $[\epsilon, \mathbf{u}_{\text{in},r}^{\omega}, \mathbf{u}_{\text{in},i}^{\omega}]$, where ϵ is the total permittivity of the underground structure on the grid, and $[\mathbf{u}_{\text{in},r}^{\omega}, \mathbf{u}_{\text{in},i}^{\omega}]$ represent the real and imaginary parts of the free space response of the source on the grid at frequency ω , i.e., the incident field with respect to the free space background. We use this to provide information of the frequency and the source. The output of our forward model are the real and imaginary parts of the total field on the grid $[\mathbf{u}_r^{\omega}, \mathbf{u}_i^{\omega}]$. Examples of the input and the output of the learned forward model could be found in Fig. 1.

3.1.1. Fourier Neural Operator

Fourier Neural Operator (FNO) is a state-of-the-art neural network architecture for solving partial differential equations (PDEs) on regular grids [15], which is inspired by the Green's function of the differential operator. The solution of a given PDE can be written as:

$$u(\mathbf{x}) = \int_D \mathbf{G}(\mathbf{x}, \mathbf{x}') f(\mathbf{x}') d\mathbf{x}' \quad (6)$$

Assuming that $\mathbf{G}(\mathbf{x}, \mathbf{x}') = \mathbf{G}(\mathbf{x} - \mathbf{x}')$, Eq. (6) reduces to a convolution operator and can be computed using the convolution theorem. In FNO, the input states are first independently lifted to some higher dimensional space using multi-layer perceptrons (MLPs). Denoting the corresponding high dimensional vector after lifting by \mathbf{v}_0 , each layer of FNO is defined recursively as,

$$\mathbf{v}_{i+1}(\mathbf{x}) = \sigma(W_i \mathbf{v}_i(\mathbf{x}) + (K_i(\mathbf{v}_i))(\mathbf{x})), \quad \forall \mathbf{x} \in D \quad (7)$$

where

$$(K_i(\mathbf{v}_i))(\mathbf{x}) = \mathcal{F}^{-1}(R_i \cdot \mathcal{F}(\mathbf{v}_i))(\mathbf{x}), \quad (8)$$

K_i and R_i are the network parameters of the i^{th} layer, $\sigma(\cdot)$ is the GeLU nonlinearity, and \mathcal{F} is the Fourier operator.

3.1.2. Born Fourier Neural Operator

Inspired by the IBA, we formulate our Born FNO (BFNO) as follows, $\forall \mathbf{x} \in D$:

$$v_\epsilon(\mathbf{x}) = P_\epsilon(\mathbf{x}, \epsilon(\mathbf{x})) \quad (9)$$

$$v_0^\omega(\mathbf{x}) = v_{\text{in}}^\omega(\mathbf{x}) = P_{\text{in}}(\mathbf{x}, u_{\text{in},r}^\omega(\mathbf{x}), u_{\text{in},i}^\omega(\mathbf{x})) \quad (10)$$

$$v_{i+1}^\omega(\mathbf{x}) = v_{\text{in}}^\omega(\mathbf{x}) + \sigma(W_1(\sigma(W_0(K(v_\epsilon, v_i^\omega)))))) \quad (11)$$

$$K(v_\epsilon, v_i^\omega)(\mathbf{x}) = \mathcal{F}^{-1}(R \cdot \mathcal{F}(v_\epsilon v_i^\omega))(\mathbf{x}) \quad (12)$$

$$u_r^\omega(\mathbf{x}) = [u_r^\omega(\mathbf{x}), u_i^\omega(\mathbf{x})] = Q(v_n^\omega(\mathbf{x})) \quad (13)$$

where $P_\epsilon, P_{\text{in}}, Q$ are local transformations parameterized by MLPs, n is the number of BFNO layers, and $\sigma(\cdot)$ is the Leaky ReLU nonlinearity. Note that unlike FNO, which has different R_i, W_i for each layer, we use the same R, W_0, W_1 for all layers, which resembles the structure in the IBA formulation Eq. (2). We follow the same normalization scheme as in [15] to preprocess our data, and train our model with the normalized-mean-squared error as formulated in Eq. (14).

$$\phi_{\text{BFNO}}^* = \underset{\phi_{\text{BFNO}}}{\text{argmin}} \sum_j \sum_{\omega} \sum_{\mathbf{x} \in D} \frac{\|u_{\text{gt},j}^\omega(\mathbf{x}) - \text{BFNO}(\epsilon_j, \omega; \phi_{\text{BFNO}})(\mathbf{x})\|_2}{\|u_{\text{gt},j}^\omega(\mathbf{x})\|_2} \quad (14)$$

where ϕ_{BFNO} is the collection of network parameters for BFNO, ϵ_j is a permittivity map from the training dataset, and $u_{\text{gt},j}^\omega$ is the ground truth total field for ϵ_j at frequency ω .

3.2. Learned Signal Prior

Since inverse scattering in the GPR setting is highly ill-posed, we restrict the solutions to a lower-dimensional subspace represented as the range of a generative model, \mathcal{G}_ϵ . Given a training dataset of permittivity maps $\{\epsilon_i\}_i$, we train \mathcal{G}_ϵ by solving:

$$(\phi_{\mathcal{G}_\epsilon}^*, \phi_{\mathcal{E}}^*) = \underset{\phi_{\mathcal{G}_\epsilon}, \phi_{\mathcal{E}}}{\text{argmin}} \sum_i \|\mathcal{G}_\epsilon(\mathcal{E}(\epsilon_i; \phi_{\mathcal{E}}); \phi_{\mathcal{G}_\epsilon}) - \epsilon_i\|_2^2 + \frac{1}{\sigma^2} \|\mathcal{E}(\epsilon_i; \phi_{\mathcal{E}})\|_2^2$$

where \mathcal{E} denotes the encoder, and \mathcal{G}_ϵ denotes the decoder, i.e. the learned generative prior, and $\phi_{\mathcal{G}_\epsilon}$ and $\phi_{\mathcal{E}}$ are the trained parameters for \mathcal{G}_ϵ and \mathcal{E} , respectively.

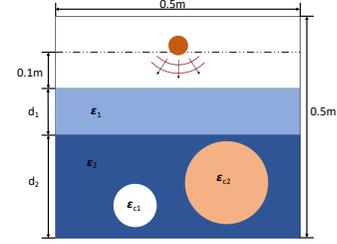


Fig. 2. Illustration of the layered medium considered in this work showing the source location and example target objects.

3.3. Inverse Problem

Combining the learned forward model and the learned prior, we can then solve the optimization problem defined in Eq. (5), where $\mathcal{R}(\cdot)$ is the isotropic total variation regularizer. With pre-trained $\text{BFNO}(\cdot; \phi_{\text{BFNO}}^*)$, and $\mathcal{G}_\epsilon(\cdot; \phi_{\mathcal{G}_\epsilon}^*)$, we solve this optimization problem using standard ADAM[16] optimizer for 1200 steps. We use the incremental frequency inversion framework proposed in [7] during the optimization, where we include a batch of 10 frequencies every 120 update steps, i.e. $[\omega_0, \dots, \omega_{10}]$ for the first 120 steps, $[\omega_0, \dots, \omega_{20}]$ for the next 120 steps, etc. We also fine-tune our prior \mathcal{G}_ϵ after 300 steps by updating the parameters $\phi_{\mathcal{G}_\epsilon}$ of \mathcal{G}_ϵ during the optimization, see Eq. (15). This fine-tuning procedure improves the generalization performance of \mathcal{G}_ϵ for the inverse problem. Similar techniques have been widely used in GAN inversion problems [17].

$$\min_{\mathbf{z}, \phi_{\mathcal{G}_\epsilon}} \sum_{\omega} \frac{1}{2} \|\mathbf{y}_\omega - \mathbf{H}(\text{BFNO}(\mathcal{G}_\epsilon(\mathbf{z}; \phi_{\mathcal{G}_\epsilon}); \omega; \phi_{\text{BFNO}}^*))\|_2^2 + \mathcal{R}(\mathcal{G}_\epsilon(\mathbf{z}; \phi_{\mathcal{G}_\epsilon})) \quad (15)$$

4. EXPERIMENT

4.1. Dataset

We generate our dataset using gprMax [14], an open source finite difference time domain simulation tool for electromagnetic wave simulation. Our computational domain is of size $0.5\text{m} \times 0.5\text{m}$. We consider a three layer background setup, as shown in Fig. 2, where the top layer is air with depth 0.15m and the ground consists of two layers with total depth 0.35m. The source is located in the center along the horizontal axis and 0.1m above the ground. We use a standard Ricker wavelet source with center frequency 1GHz. The depth of the second ground layer is sampled from a uniform distribution $d_2 \sim \mathcal{U}(0.1, 0.3)$, and the depth of the first ground layer is computed as $d_1 = 0.35 - d_2$. The first layer has permittivities sampled from $\epsilon_1 \sim \mathcal{U}(3, 5)$, and the second layer has permittivities sampled from $\epsilon_2 \sim \mathcal{U}(5, 10)$. Two cylinders are embedded in the second ground layer. One of the cylinder is air, thus with $\epsilon_{c1} = 1$. The permittivity of the other cylinder is sampled from a uniform distribution $\epsilon_{c2} \sim \mathcal{U}(3, 10)$. Both cylinders has radius $r \sim \mathcal{U}(0.03, 0.06)$ Sample structures could be found in Fig. 3 and Fig. 2.

From the time domain simulation results generated by gprMax, we apply the Fourier transform and extract the wavefield with frequencies within $[0.5\text{GHz}, 1.5\text{GHz}]$ band to train the forward model and solve the inverse problems. Our dataset consists of 400 simulations with 350 training samples and 50 test samples.

4.2. Forward Problems

In this section, we compare the FNO [15] and the proposed Born FNO for the forward problem. The training and test losses are the mean squared error (MSE) between the ground truth total field and

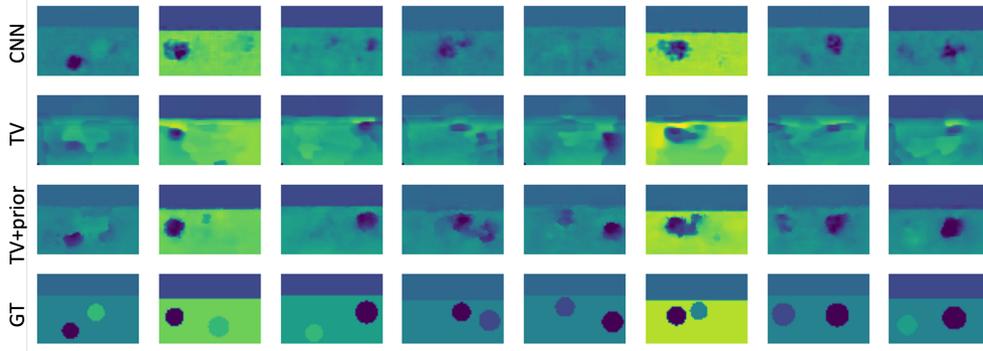


Fig. 3. Randomly sampled qualitative results for the inverse problem solved by different approaches. CNN uses an autoencoder to learn the mapping between the measurements and permittivity maps. Both “TV” and “TV + prior” use Born FNO with 10 layers as the forward model. “TV” uses the standard isotropic TV regularizer, whereas “TV + prior” uses isotropic TV combined with the learned signal prior.

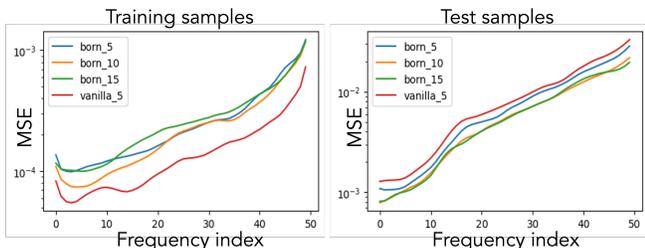


Fig. 4. Loss for FNO and Born FNO with different number of layers. Left: training loss as a function of frequency averaged over 350 training samples. Right: test loss as a function of frequency averaged over 50 test samples.

the total field generated by FNO or Born FNO. The training and test losses as functions of frequency is shown in Fig. 4. We can see that while FNO achieves lower training errors (overfits well), it has higher test errors compared to Born FNO with the same number of layers. This demonstrates that the proposed Born FNO achieves better generalization performance. We also observe that the test errors of Born FNO converge as the number of layers increases.

4.3. Inverse Problems

For the inverse problem, we place the receivers along the source plane and use all 50 frequencies from $\omega \in 2\pi[0.5\text{GHz}, 1.5\text{GHz}]$. We solve the inverse problem with different forward models, FNO with 5 layers, Born FNO with 5 layers, or Born FNO with 10 layers. We also perform ablation where we solve the inverse scattering problem with or without the learned prior, \mathcal{G}_ϵ , to test the effectiveness of the proposed learned prior. Another existing machine learning solution for full-waveform inversion is to treat it as an image-to-image translation problem and use a CNN to learn the mapping between the measurements and the permittivity maps, similar to what shown in [18]. As a baseline, we also train an autoencoder with latent feature size 8 with our dataset. Quantitative results can be found in Table (1). We also include the forward model loss, MSE_{evo} , for different learned forward models in Table (1), where $\text{MSE}_{\text{evo}} = \|\mathbf{u}_{\text{true}}^\omega - \text{BFNO}(\epsilon, \omega)\|_2^2$, measures the accuracy of each forward model. All numbers reported in Table (1) are averaged over 50 test samples.

We observe that with a more accurate forward model, lower

MSE_{evo} , we get a better reconstruction result. Our Born FNO with 10 layers achieves the lowest forward model loss as well as the lowest reconstruction error. Given the same forward model, solving the inverse problem with the learned prior yields a lower reconstruction error compared to the without the learn prior case, indicating the effectiveness of the learned prior. Our approach also achieves better results compared to the baseline CNN. Besides, while the CNN-based approach learns the mapping between the measurements and permittivity maps under a specific measurement setup, our approach can adapt to arbitrary sensor setups which provides more flexibility. Fig. 3 shows qualitative comparisons between the different learned forward models and demonstrates that the learned prior with the Born FNO leads to the best reconstruction of the underground permittivity distribution.

	MSE_{evo}	TV	TV + Learned Prior
FNO 5 [15]	9.96e-3	1.366	0.809
BFNO 5 (ours)	8.63e-3	1.065	0.715
BFNO 10 (ours)	6.94e-3	1.030	0.683
CNN		1.139	

Table 1. Quantitative results for underground imaging problem with different forward models and regularizers. The number in the first column represents the number of layers used in each model. Each forward model is tested with or without the learned prior. Baseline CNN learns a direct mapping between the measurements and the permittivity map using an autoencoder. The reported numbers are mean squared errors averaged over 50 test samples. We observe that the proposed 10-layer Born FNO with the learned prior achieves the best test performance.

5. CONCLUSION

In this paper we presented a machine learning framework for solving inverse scattering problems under the GPR setup. We demonstrated that our Iterative-Born-Approximation inspired FNO model can learn complicated wave-matter interaction with higher accuracy, leading to more accurate reconstruction of the underground structures, compared to previous models.

6. REFERENCES

- [1] Seunghee Lee, George A McMechan, and Carlos LV Aiken, "Phase-field imaging: The electromagnetic equivalent of seismic migration," *Geophysics*, vol. 52, no. 5, pp. 678–693, 1987.
- [2] Raffaele Persico, *Introduction to ground penetrating radar: inverse scattering and data processing*, John Wiley & Sons, 2014.
- [3] Emil Wolf, "Three-dimensional structure determination of semi-transparent objects from holographic data," *Optics Communications*, vol. 1, no. 4, pp. 153–156, 1969.
- [4] Lin-Ping Song and Qing Huo Liu, "Ground-penetrating radar land mine imaging: Two-dimensional seismic migration and three-dimensional inverse scattering in layered media," *Radio science*, vol. 40, no. 1, pp. 1–15, 2005.
- [5] Yanting Ma, Hassan Mansour, Dehong Liu, Petros T. Boufounos, and Ulugbek S. Kamilov, "Accelerated image reconstruction for nonlinear diffractive imaging," 2018.
- [6] Emmanuel Soubies, Thanh-An Pham, and Michael Unser, "Efficient inversion of multiple-scattering model for optical diffraction tomography," *Opt. Express*, vol. 25, no. 18, pp. 21786–21800, Sep 2017.
- [7] Ajinkya Kadu, Hassan Mansour, and Petros T. Boufounos, "High-contrast reflection tomography with total-variation constraints," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 1523–1536, 2020.
- [8] M. Born and E. Wolf, *Principles of Optics*, Cambridge University Press, 2019.
- [9] Yan Yang, Angela F. Gao, Jorge C. Castellanos, Zachary E. Ross, Kamyar Azizzadenesheli, and Robert W. Clayton, "Seismic wave propagation and inversion with neural operators," 2021.
- [10] Qingqing Zhao, David B. Lindell, and Gordon Wetzstein, "Learning to solve pde-constrained inverse problems with graph networks," in *ICML*, 2022.
- [11] Laurent Sirgue, Olav I. Barkved, Joseph Anthony Dellinger, John Etgen, Uwe K. Albertin, and Jan H. Kommedal, "Thematic set: Full waveform inversion: the next leap forward in imaging at valhall," *First Break*, vol. 28, pp. 65–70, 2010.
- [12] J. Virieux and S. Operto, "An overview of full-waveform inversion in exploration geophysics," *GEOPHYSICS*, vol. 74, no. 6, pp. WCC1–WCC26, 2009.
- [13] S R Arridge, "Optical tomography in medical imaging," *Inverse Problems*, vol. 15, no. 2, pp. R41–R93, jan 1999.
- [14] A.J. Witten, J.E. Molyneux, and J.E. Nyquist, "Ground penetrating radar tomography: algorithms and case studies," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 32, no. 2, pp. 461–467, 1994.
- [15] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar, "Fourier neural operator for parametric partial differential equations," 2020.
- [16] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," 2014.
- [17] Daniel Roich, Ron Mokady, Amit H. Bermano, and Daniel Cohen-Or, "Pivotal tuning for latent-based editing of real images," *arXiv preprint arXiv:2106.05744*, 2021.
- [18] Yue Wu and Youzuo Lin, "InversionNet: An efficient and accurate data-driven full waveform inversion," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 419–433, 2020.