

JGR Solid Earth

RESEARCH ARTICLE

10.1029/2025JB031624

Key Points:

- We apply neural operators for ambient noise full waveform inversion to image the northern Los Angeles basins using linear nodal array data
- We provide a trained Helmholtz neural operator that is applicable without retraining to various problems with a comparable spatial footprint
- We showcase that trained neural operators can perform full waveform inversion two orders of magnitude faster than the adjoint-state method

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

C. Zou,
czou@caltech.edu

Citation:

Zou, C., Ross, Z. E., Clayton, R. W., Lin, F.-C., & Azizzadenesheli, K. (2025). Ambient noise full waveform inversion with neural operators. *Journal of Geophysical Research: Solid Earth*, 130, e2025JB031624. <https://doi.org/10.1029/2025JB031624>

Received 21 MAR 2025

Accepted 16 OCT 2025

Author Contributions:

Conceptualization: Zachary E. Ross, Robert W. Clayton
Data curation: Robert W. Clayton
Formal analysis: Caifeng Zou, Zachary E. Ross, Robert W. Clayton, Fan-Chi Lin, Kamyar Azizzadenesheli
Funding acquisition: Zachary E. Ross
Investigation: Caifeng Zou, Fan-Chi Lin
Methodology: Caifeng Zou, Zachary E. Ross, Robert W. Clayton, Kamyar Azizzadenesheli
Project administration: Zachary E. Ross, Robert W. Clayton
Resources: Zachary E. Ross, Robert W. Clayton
Software: Caifeng Zou
Supervision: Zachary E. Ross, Robert W. Clayton, Fan-Chi Lin, Kamyar Azizzadenesheli
Validation: Caifeng Zou, Fan-Chi Lin, Kamyar Azizzadenesheli
Visualization: Caifeng Zou

© 2025. American Geophysical Union. All Rights Reserved.

Ambient Noise Full Waveform Inversion With Neural Operators

Caifeng Zou¹ , Zachary E. Ross¹ , Robert W. Clayton¹, Fan-Chi Lin² , and Kamyar Azizzadenesheli³ 

¹Seismological Laboratory, California Institute of Technology, Pasadena, CA, USA, ²Department of Geology and Geophysics, University of Utah, Salt Lake City, UT, USA, ³Nvidia Corporation, Santa Clara, CA, USA

Abstract Numerical simulations of seismic wave propagation are crucial for investigating velocity structures and improving seismic hazard assessment. However, standard methods such as finite difference or finite element are computationally expensive. Recent studies have shown that a new class of machine learning models, called neural operators, can solve the elastodynamic wave equation orders of magnitude faster than conventional methods. Full waveform inversion is a prime beneficiary of the accelerated simulations. Neural operators, as end-to-end differentiable operators, combined with automatic differentiation, provide an alternative approach to the adjoint-state method. State-of-the-art optimization techniques built into PyTorch provide neural operators with greater flexibility to improve the optimization dynamics of full waveform inversion, thereby mitigating cycle-skipping problems. In this study, we demonstrate the first application of neural operators for full waveform inversion on a real seismic data set, which consists of several nodal transects collected across the San Gabriel, Chino, and San Bernardino basins in the Los Angeles metropolitan area.

Plain Language Summary Full waveform inversion enables high-resolution imaging of subsurface physical property variations, allowing detailed characterization of geological structures such as faults, reservoirs, and sedimentary basins. However, it is computationally expensive using traditional numerical methods, due to numerous evaluations of the forward modeling process and its adjoint. We use a class of machine learning models named neural operators as a surrogate solver to accelerate the forward process and, thereby, the inversion. We show that trained neural operators can perform full waveform inversion more efficiently with less existing knowledge. We apply the method on ambient noise seismic data collected in the northern Los Angeles basins and obtain results verified with previous geological studies. The same model can potentially generalize to different regions and help to improve earthquake hazard assessment.

1. Introduction

Seismic tomography for sedimentary basins is important for assessing earthquake hazards considering the fact that the basins can trap and amplify seismic waves, leading to stronger and longer ground shaking. Fine-scale subsurface structures are typically determined by active source surveys, which often operate at high frequencies but are expensive and limited by environmental impact issues. Alternatively, seismic data from passive sources, such as ambient noise and earthquakes, can be more readily acquired and can also compensate for the absence of low-frequency information in active source data. Many have shown in seismic interferometry that the cross-correlation between diffuse waveforms recorded at two stations approximates the elastodynamic Green's function (Wapenaar, 2004; Wapenaar et al., 2010). The so-called empirical Green's function (EGF) created from ambient noise is more commonly used by ray-theory-based tomography methods (Guo et al., 2015; Lin et al., 2008; Yao et al., 2010; Zheng et al., 2011), which do not account for the finite-frequency effects in heterogeneous media. Full waveform inversion (FWI), which accounts for the full physics of wave propagation, can reveal more accurate and detailed subsurface structures (Y. Liu et al., 2017; X. Liu et al., 2023; Maguire et al., 2022; Sager et al., 2020; K. Wang et al., 2021; Z.-D. Zhang et al., 2021).

Despite its significant advantages, FWI can be time-consuming and memory-intensive for large problems, which hinders its broader applicability. Moreover, the adjoint-state method (Fichtner et al., 2006a; Q. Liu & Tromp, 2006; Plessix, 2006; Tromp et al., 2005), which is conventionally used for FWI, requires the initial model to be sufficiently close to the true Earth to avoid being trapped in local minima (Fichtner et al., 2006b; Gauthier et al., 1986; Panning et al., 2009; Virieux & Operto, 2009). In ambient noise tomography (ANT), the adjoint-state method typically relies on a model derived from surface wave dispersion analysis and ray-based approximations

Writing – original draft: Caifeng Zou
Writing – review & editing: Zachary E. Ross, Robert W. Clayton, Fan-Chi Lin, Kamyar Azizzadenesheli

as a starting point (Chen et al., 2014; X. Liu et al., 2023; Maguire et al., 2022; C. Zhang et al., 2018; Z.-D. Zhang et al., 2020). To tackle these challenges, recent studies have shown that neural operators (Azizzadenesheli et al., 2024; Z. Li et al., 2020a, 2020b) can solve wave equations orders of magnitude faster than traditional numerical solvers (Yang et al., 2021, 2023; Zou et al., 2024). In deep learning, an analogue to the adjoint-state method is known as the reverse-mode automatic differentiation (AD) (Baydin et al., 2018; Rumelhart et al., 1986). While these two methods have been shown to be mathematically equivalent (Richardson, 2018; W. Wang et al., 2021; W. Zhu et al., 2021), AD, which underpins the modern deep learning platforms such as PyTorch (Paszke et al., 2019) and TensorFlow (Abadi et al., 2016), remains aligned with advances in state-of-the-art optimization techniques. For example, methods such as mini-batching (LeCun et al., 2002), stochastic gradient descent (SGD) (Bottou, 2010), and adaptive moment estimation (Adam) (Kingma & Ba, 2014) have been developed to better navigate non-convex loss landscapes and escape poor local minima. These methods have prevailed for over a decade in the machine learning community, while remaining relatively underexplored in the classic FWI world. Only in recent years have they begun to attract attention (Bernal-Romero & Iturrarán-Viveros, 2021; Mao et al., 2025; Richardson, 2018; J. Sun et al., 2020; Thrastarson et al., 2022; van Herwaarden et al., 2020). Although efforts have been made to implement finite difference solvers in PyTorch/TensorFlow to enable access to AD and various optimization strategies (Richardson, 2018; W. Zhu et al., 2021), this approach is memory prohibitive for realistic applications, because back-propagation requires storing wavefields at every time step in memory. We bypass this obstacle by learning a direct mapping through neural operators developed in PyTorch. In alignment with the rapid advancements in optimization techniques, neural operators have the potential to alleviate—though not fully resolve—the cycle-skipping problem in FWI, making it possible to start from a simpler velocity model (e.g., a 1D model).

The primary advantage of neural operators in solving partial differential equations (PDEs) over other prevailing machine learning approaches, such as physics-informed neural networks (PINNs) (Raissi et al., 2019), is that they are trained to learn the solution operator for an entire family of PDEs instead of a specific instance. This means that a neural operator has the potential to generalize to arbitrary PDE coefficients (e.g., velocity structures) if the training instances are ideally sampled—for example, from random fields (Yang et al., 2021, 2023; Zou et al., 2024). The training is a one-time effort for the forward process and there is no further training for the neural operator in the inversion stage, where the velocity parameters are updated with gradients computed from AD. This is similar to traditional FWI but replaces the adjoint simulation with back-propagation (LeCun et al., 1988).

In recent years, neural operators have been increasingly applied to seismic wave propagation modeling. They have been used for modeling acoustic (Yang et al., 2021), elastic (C. Li et al., 2025; Yang et al., 2023), and viscoelastic (Wei & Fu, 2022) waves. Most applications are limited to 2D due to computational and memory constraints. However, Lehmann et al. (2024, 2025) circumvented the 3D computational bottleneck by predicting only the surface ground motion, while Zou et al. (2024) modeled the full 3D wavefield by parameterizing the operator in the frequency domain, thereby eliminating the need to model the time dimension. Learning frequency-domain solutions offers memory efficiency for neural operators and has attracted growing interest (Cheng et al., 2025; Huang & Alkhalifah, 2025; Kong et al., 2025; B. Li et al., 2023; T. Zhang et al., 2023). Among these, Huang and Alkhalifah (2025) and Cheng et al. (2025) proposed learning the residual (scattered) wavefield instead of the full wavefield to address the point-source singularity. Current research has primarily focused on synthetic data, while the application of neural operators to real seismic data remains largely unexplored.

Our contributions are as follows. We demonstrate the first application of neural operators for ambient noise tomography on a real seismic data set, specifically the Basin Amplification Seismic Investigation (BASIN) survey (BASIN, 2018; Clayton et al., 2019), which deployed linear nodal arrays in the northern Los Angeles (LA) basins. We provide a trained neural operator that is applicable without retraining to any linear array in the study area, or potentially other problems with a comparable spatial footprint. With the trained neural operator, FWI can be performed two orders of magnitude faster than the conventional adjoint method. Our tomography results are consistent with previous studies in the same region (Ghose et al., 2023; Y. Li et al., 2023; Villa et al., 2023; Zou & Clayton, 2024) without the strong dependence on priors. The proposed method can be scaled to 3D, provided that sufficient computing resources are available.

2. Methods

In this section, we first present the governing equations, as well as the methodology of using neural operators to solve the forward and inverse problems. Then, we share the training strategy.

2.1. Helmholtz Neural Operator With Automatic Differentiation

Seismic wave propagation in an elastic medium follows the (isotropic) elastic wave equation

$$\rho \frac{\partial^2 \mathbf{u}}{\partial t^2} = \nabla \lambda (\nabla \cdot \mathbf{u}) + \nabla \mu \cdot (\nabla \mathbf{u} + (\nabla \mathbf{u})^T) + (\lambda + 2\mu) \nabla (\nabla \cdot \mathbf{u}) - \mu \nabla \times \nabla \times \mathbf{u} + \mathbf{f}, \quad (1)$$

where $\mathbf{u}$ is the displacement wavefield (three-component vector in 3D media), t is time, ∇ is the gradient with respect to the space coordinates, $\mathbf{f}$ is the body force (source term), ρ is the density, λ and μ are the Lamé parameters. Equation 1 can be restated concisely with an integro-differential operator $\mathcal{L}$ as

$$\mathcal{L} \mathbf{u} = \mathbf{f}. \quad (2)$$

In seismology the medium $\mathbf{m}$ is commonly parametrized by P-wave velocity V_P and S-wave velocity V_S :

$$V_P = \sqrt{\frac{\lambda + 2\mu}{\rho}} \quad \text{and} \quad V_S = \sqrt{\frac{\mu}{\rho}}. \quad (3)$$

The direct mapping from physical parameters $\mathbf{m}$ to observed waveforms $\mathbf{d}$ is nonlinear and unknown in closed form:

$$\mathbf{d} = \mathbf{G}(\mathbf{m}). \quad (4)$$

Based on the universal approximation theorem, neural operators can approximate arbitrary nonlinear continuous operators (Hornik et al., 1989; Kovachki et al., 2023). The solution operator can be parameterized much more efficiently in memory when solving Equation 2 in the frequency domain (i.e., the Helmholtz equation) than in the time domain (Huang & Alkhalifah, 2025; Kong et al., 2025; Zou et al., 2024). This is because different frequency components can be solved individually and, thus, parallelized. In this study, we employ the Helmholtz Neural Operator (HNO) proposed by Zou et al. (2024). The HNO computes the solution for each frequency component independently by defining the batch of data at the frequency level. In other words, a batch can contain multiple frequency components due to data shuffling. We use a single HNO to learn the entire frequency band of interest (0.1–0.5 Hz), although Kong et al. (2025) found that using multiple models to handle different frequencies slightly improves performance. Solutions in the time and frequency domains can be related via the (inverse) Fourier transform. We construct a U-shaped HNO comprising $L = 8$ inner layers of Fourier Neural Operators (FNOs) (Z. Li et al., 2020a), followed by a Graph Neural Operator (GNO) (Z. Li et al., 2020b), to predict the waveforms at the free surface (data, denoted by $\hat{\mathbf{d}}$):

$$\begin{aligned} \mathbf{v}_0(x) &= (\mathcal{P} \mathbf{a})(x), \\ \mathbf{v}_{l+1}(x) &= \sigma(\mathcal{W}_l \mathbf{v}_l(x) + \mathcal{F}^{-1}(\mathcal{F}(\kappa_l) \cdot \mathcal{F}(\mathbf{v}_l))(x)) \quad \text{where } l = 0, \dots, L-1, \\ \mathbf{v}_{gno}(x) &= \int_{B_r(x)} \kappa(x, y, \mathbf{v}_L) \mathbf{v}_L(y) dy, \\ \hat{\mathbf{d}}(x) &= (\mathcal{Q} \mathbf{v}_{gno})(x), \end{aligned} \quad (5)$$

where $\mathbf{a}$ is the input including V_P , V_S , the source location, and a constant function indicating the frequency value, $\hat{\mathbf{d}}$ is the predicted data for the given frequency component (including the real and imaginary parts), $\mathbf{v}_l$ is input to the $l + 1^{\text{th}}$ FNO layer ($l = 0, \dots, 7$) or the GNO ($l = 8$), $\mathbf{v}_{gno}$ is the GNO output, $\mathcal{P}$ is a point-wise operator used to lift the dimension, $\mathcal{Q}$ is used to project the output to the desired dimension, $\mathcal{W}_l$ acts as a residual connection, κ_l is a parametric kernel function, $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier and inverse Fourier transforms, respectively, and σ is a




Figure 1. Model architecture. **a** is the input given by V_p , V_s , the source location, and a constant function indicating the frequency value. $\hat{\mathbf{d}}$ is the predicted data. $\mathcal{P}$ is a point-wise operator used to lift the dimension. $\mathcal{Q}$ is used to project the output to the desired dimension. $\mathcal{B}$ is the inner integral operator chosen as the FNO, in which $\mathbf{v}$ is input to the layer, $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier and inverse Fourier transforms, respectively, $\mathcal{R}$ is a linear operator, $\mathcal{W}$ acts as a residual connection, and σ is a nonlinear activation function. Λ is a GNO used to query the waveforms at the free surface. Blue circles denote concatenation along the channel dimension.

nonlinear activation function, which we set as GELU (Hendrycks & Gimpel, 2016). The GNO uses a kernel function κ parameterized as a three-layer neural network, which takes positional information and the output of the last FNO layer as input. $\mathcal{B}_r$ denotes the predefined neighborhood of the queried position x , defined as all points below the corresponding surface location for computational efficiency. We use FNOs to speed up the integral calculation, and the U-shape allows for a deeper model and skip connections (Z. Li et al., 2020; Rahman et al., 2022; Ronneberger et al., 2015). The GNO can improve predictions at queried points and provide greater data efficiency (Zou et al., 2024). Figure 1 illustrates the model architecture and Table S1 in Supporting Information S1 provides the details in each neural operator layer.

In conventional FWI, the solution is obtained by minimizing a scalar-valued objective function Φ , which can, for example, be defined as half the squared ℓ_2 -norm of the data residual:

$$\Phi = \frac{1}{2} \|\hat{\mathbf{d}} - \mathbf{d}\|_2^2 \quad \text{where} \quad \hat{\mathbf{d}} = \mathbf{S}\mathbf{u}, \quad (6)$$

where $\hat{\mathbf{d}}$ is the modeled (predicted) waveforms queried at receivers through a sampling operator $\mathbf{S}$ and $\mathbf{d}$ is the observed data. The parameters of interest $\mathbf{m}$ are updated iteratively using a gradient-based method. Using the adjoint-state method, the gradients of the objective with respect to parameters of interest can be estimated as the

cross-correlation between the forward wavefield $\mathbf{u}$ and the adjoint wavefield $\mathbf{u}^*$ at zero time lag, weighted by an analytically derived term (Louboutin et al., 2019):

$$\nabla_{\mathbf{m}} \Phi = -\mathbf{u}^T \left(\frac{\partial \mathcal{L}}{\partial \mathbf{m}} \right)^T \mathbf{u}^*. \quad (7)$$

The adjoint wavefield $\mathbf{u}^*$ is generated by back-propagating the residual $\delta \mathbf{d}$ in the medium:

$$\mathcal{L}^T \mathbf{u}^* = S^T \delta \mathbf{d} \quad \text{where} \quad \delta \mathbf{d} = \hat{\mathbf{d}} - \mathbf{d}. \quad (8)$$

The adjoint-state method requires the initial model to be sufficiently accurate, with a travel-time error of less than half the period (Alkhalifah, 2016; Beydoun & Tarantola, 1988; Virieux & Operto, 2009). Otherwise, the cycle-skipping phenomenon may arise, leading to convergence to a local minimum (Gauthier et al., 1986; Pladys et al., 2021).

With neural operators (or other machine learning methods developed on the PyTorch platform), the gradients of the objective with respect to any parameters in the computational graph are automatically computed through back-propagation (Rumelhart et al., 1986). This is termed the reverse-mode AD (Baydin et al., 2018; Elliott, 2018). Provided that all the functions, expressions, and control flow structures used in neural operators are differentiable and compatible with AD, the adjoint simulation can be replaced by back-propagation (LeCun et al., 1988). It has been shown that the adjoint-state method and reverse-mode AD are mathematically equivalent (Richardson, 2018; W. Wang et al., 2021; W. Zhu et al., 2021). However, the HNO-AD method eliminates the need to manually derive the gradient in Equation 7 on a case-by-case basis. Moreover, it can take advantage of state-of-the-art optimization techniques built into PyTorch—such as mini-batching and Adam—which help improve the optimization dynamics of FWI and mitigate cycle-skipping. Adam, in particular, a first-order gradient-based optimization algorithm that dynamically rescales the gradients for each parameter based on its past gradients and their squares (Kingma & Ba, 2014), has been shown to outperform other optimizers such as SGD and limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) (Bernal-Romero & Iturrarán-Viveros, 2021; Richardson, 2018; J. Sun et al., 2020).

We summarize the differences between the HNO-AD and adjoint-state methods as follows:

1. Forward modeling: In the adjoint-state method, forward modeling is performed by numerically solving the wave equation. In the HNO-AD method, it is performed via a mapping learned by a neural operator.
2. Gradient computation: In the adjoint-state method, the gradient is analytically derived and computed by running an adjoint simulation. In the HNO-AD method, this is replaced by automatic back-propagation.

2.2. Data-Driven Training

In this study, we train a single 2D HNO for linear nodal arrays across the northern LA basins. The model is trained in a supervised learning manner using synthetic data generated from a spectral element method (SEM) using the software SALVUS (Afanasiev et al., 2019). The computational domain is set to 80 km (horizontal) $\times$ 40 km (vertical) on a 256×128 mesh, which accommodates the longest seismic line, SB1, in the area. We simulate 40-s-duration wavefields with a time step of 0.001 s, which are excited by a vertical force and Ricker wavelet time function with a central frequency of 0.3 Hz. For ambient noise data applications, the source is randomly placed along the free surface from a uniform distribution of horizontal position. The V_P and V_S models are generated from random fields sampled around a background 1D model, which is averaged along the SB1 line from a reference model named CVM-S4.26 (Lee et al., 2014). The parameters for the random fields are provided in Table S2 in Supporting Information S1. The density models are derived from V_P using the empirical relation by Brocher (2005), which are input to SALVUS but not to the HNO. The time-domain solutions from the SEM solver are Fourier-transformed and filtered to the frequency band of interest (0.1–0.5 Hz) for use with the HNO. We train an HNO using 27,000 simulations and validate the model performance using 3,000 simulations (where a simulation means a single source and a particular velocity model). We define the loss function as a combination of relative ℓ_1 -norm and ℓ_2 -norm of the data residual:

Table 1

Comparison of Time Consumed by SALVUS and HNO in Forward Modeling and FWI

	SALVUS	HNO
Forward modeling	7.18 s	0.18 s
FWI	128.51 s	0.45 s

Note. The experiments are conducted on a community model called CVM-S4.26 using a single NVIDIA RTX A6000 GPU. The forward modeling time is measured for a single event with one source and 234 receivers on the free surface. The FWI time is measured for one tomographic iteration for the same event configuration.

$$Loss = 0.95 \frac{\|\mathbf{d} - \hat{\mathbf{d}}\|_1}{\|\mathbf{d}\|_1} + 0.05 \frac{\|\mathbf{d} - \hat{\mathbf{d}}\|_2}{\|\mathbf{d}\|_2}, \quad (9)$$

where $\mathbf{d}$ and $\hat{\mathbf{d}}$ denote the true and predicted data, respectively. We use an Adam optimizer with a learning rate of 0.001 and a scheduler that decays the learning rate by half every 30 epochs. The batch size is 256. Figure S1 in Supporting Information S1 shows the loss curves. The full training process took 2 days with 8 NVIDIA RTX A6000 GPUs. Note that the training is a one-time effort and no further training is required for different seismic lines or the inversion stage, granted the data have a similar scale. Table 1 compares the time consumed by SALVUS and the trained HNO in forward modeling and FWI under the same configuration. In performing FWI for this 2D case,

the HNO with AD is approximately two orders of magnitude faster than the SEM with the adjoint method. The computational advantage will only be greater in 3D (Zou et al., 2024).

3. Synthetic Tests

To test generalization, we first evaluate the HNO performance with the synthetic data generated from a community velocity model named CVM-S4.26 (Lee et al., 2014). We extract a profile of the 3D CVM-S4.26 along the longest seismic line, SB1, in the BASIN survey. This model is never explicitly seen by the HNO. Figure 2 shows the HNO-predicted results for a sample shot excited by a vertical force and Ricker wavelet as in training, with the SALVUS-simulated waveforms serving as the ground truth. The direct output of the HNO is in the frequency domain (0.1–0.5 Hz) and is inverse Fourier-transformed to the time domain for visualization purposes. The forward prediction by the HNO is not perfect with a cross-correlation coefficient of 0.98, but we will show that it is sufficiently accurate in the upcoming FWI experiment.

Before the inversion, we compute the sensitivity kernels for the HNO-AD method, following the definition in Tromp et al. (2005). In Figure 3, we put sources and receivers at nodes on the free surface and sum the gradients of the misfit with respect to the velocity parameters over all source-receiver pairs. The misfit is defined as the squared error between the HNO-predicted data with a 1D initial model and the SALVUS-simulated data with the true model (CVM-S4.26). The 1D initial model is obtained by horizontally averaging the CVM-S4.26 profile. We average the sensitivity kernels horizontally to obtain functions of depth. We can see that the sensitivity drops




Figure 2. Comparing simulations for velocity model CVM-S4.26 between HNO and baseline method. Source location is marked by a white star.




Figure 3. Sensitivity kernels for the HNO-AD method, defined as the gradients of the misfit with respect to the velocity parameters, which are summed over all source-receiver pairs and averaged horizontally. The misfit is defined as the squared error between the HNO-predicted data with a 1D initial model and the SALVUS-simulated data with the true model (CVM-S4.26). The sensitivity values for V_P and V_S are each normalized by its maximum amplitude.

dramatically with depth, which indicates that updates to deeper parts of the velocity models are less reliable. Therefore, we will only present the inversion results for the top 10 km hereafter.

In the inversion stage, we freeze the HNO parameters and update V_P and V_S simultaneously with gradients calculated from AD. The inversion is also performed in the frequency domain, and the objective function is defined as the mean squared error (MSE) between the HNO-predicted and SALVUS-simulated data. For use with ambient noise data where the amplitude is often not well-defined, we normalize the amplitude to unity before computing the MSE. This is done by dividing the frequency-domain data, a complex-valued quantity, by its amplitude. We add no regularization term to the objective function but implicitly impose some regularization by smoothing the gradients with a Gaussian filter (Tape et al., 2010; H. Zhu et al., 2015). The smoothing radius is set to $3 \times$ grid spacing (945 m) in both horizontal and vertical directions. We average the CVM-S4.26 profile horizontally to get a 1D initial model. We start with lower frequency data and gradually feed in higher frequency components (from 0.1 to 0.5 Hz), which is also common practice in traditional FWI to mitigate cycle-skipping (K. Wang et al., 2018; C. Zhang et al., 2018). This is because skipping one cycle in lower frequency data requires larger velocity variation, whereas velocity variation is smaller at greater depth—where lower frequency data has better sensitivity. Also, the neural operator performs more accurately for lower frequencies (Zou et al., 2024). Figure 4 shows the MSE for the full frequency band in the inversion process. We use an Adam optimizer with a




Figure 4. Overall MSE (mean squared error) in the FWI process, where higher frequency data is progressively incorporated. We use an Adam optimizer with a learning rate of 0.05 and a scheduler that decays the learning rate by half every 5 epochs. The batch size is 64.

learning rate of 0.05 and a scheduler that decays the learning rate by half every 5 epochs. For simplicity, we use a fixed number of epochs for each frequency band, balancing accuracy and efficiency, although adapting the epoch count based on convergence may be a better option. The gradients are summed over a mini-batch of 64 data points, where a single data point corresponds to a single frequency and a single source.

We conduct FWI experiments on both noise-free data and data perturbed with different levels of Gaussian random noise. We activate 234 sources along the free surface of the CVM-S4.26 slice for SB1, located at the nodes. The receivers are also placed at the nodes. Figure 5 shows a sample shot of synthetic data perturbed with different levels of noise. The noise is generated from a zero-mean Gaussian distribution with a standard deviation (SD) equal to a factor times the SD of the noise-free data. The corresponding FWI results for noise-free data and noisy data are displayed in Figure 6. The noise-free inversion in the second row clearly reveals the basin model, demonstrating the forward modeling accuracy of the HNO from another perspective. The deeper part of the model is not well resolved because the source-receiver geometry and frequency content limit sensitivity to near-surface regions. When $5 \times$ noise is added, the FWI results maintain good quality, even though the seismic signal already becomes unrecognizable to the eye. This could be




Figure 5. A sample shot of synthetic data perturbed with different levels of noise. The noise is generated from a zero-mean Gaussian distribution with an SD equal to a factor times the SD of the noise-free data.

explained by the fact that FNOs are global operators, and the zero-mean noise might be partially averaged out in the kernel. However, real noise typically does not behave this nicely. We increase the noise level until the inversion is on the verge of collapse, with a perturbation of 10 $\times$ noise that totally obliterates the waveforms.

These synthetic tests serve to verify the generalizability of the neural operator trained with random fields and noise-free data to realistic velocity structures and noisy data, supervised by some definition of ground truth that is unavailable in real data. They show the excellent performance of the HNO-AD method in low signal-to-noise ratio (SNR) situations. Additionally, the entire FWI process can be completed in minutes. In contrast, FWI with the adjoint-state method takes 16 hours for 10 iterations and makes little progress with the same 1D initialization even for noise-free data (Figure S2 in Supporting Information S1). We attribute the main difference to the use of mini-batch Adam optimization in the HNO-AD method, which is not yet supported by SALVUS. We provide another generalization test in Figure S3 in Supporting Information S1 for a velocity model obtained from a ray-theory-based method (Y. Li et al., 2023), which is fully independent of the training process.

4. Application to Real Seismic Data

We apply the trained neural operator to a real-world ambient noise tomography example. The real data comes from the BASIN survey that deployed 10 linear nodal arrays in the San Gabriel (SG), Chino, and San Bernardino (SB) basins north of LA (Clayton et al., 2019). A total of 758 Fairfield ZLand three-component nodes with a corner frequency of 5 Hz recorded ambient noise for approximately a month. Each line consisted of 14–260 geophones, spaced approximately 0.25 km apart. The geometry of the survey was designed to detect subsurface variations up to 0.5 km in horizontal scale. Figure 7 shows a map of the study area and station configuration. The sedimentary formations date from the opening of the LA basin in the Miocene (Wright, 1991). Details in the geologic and tectonic setting can be found in Villa et al. (2023) and Ghose et al. (2023).

We take the EGF derived from ambient noise cross-correlation as the direct observation:




Figure 6. Synthetic FWI results for noise-free data and noisy data.

$$G_{AB}(t) = -\frac{d}{dt} \left(\frac{C_{AB}(t) + C_{AB}(-t)}{2} \right) \quad \text{where } t > 0, \quad (10)$$

where $G_{AB}(t)$ is the displacement EGF between two stations A and B , and $C_{AB}(t)$ is the cross-correlation between the seismograms recorded at these two stations. The positive and negative lags of C_{AB} are averaged to enhance the SNR and to mitigate the effect of inhomogeneous source distribution (Lin et al., 2008). Although more effective methods for addressing non-diffuse noise fields have been explored (X. Liu, 2020; X. Liu et al., 2023), this is not the focus of this study, and we follow the standard procedure. Our cross-correlation workflow is based on Bensen et al. (2007) and Y. Li et al. (2023). We correlate both daytime and nighttime recordings from the 1-month deployment, since using only nighttime data merely reduced the anthropogenic noise in a minor way (Y. Li et al., 2023). One-hour data segments are correlated and stacked to derive the final correlation. We convert the three-component data from the vertical-north-east (ZNE) to the vertical-radial-transverse (ZRT) coordinate system. The particle motion of Rayleigh waves lies in the ZR plane, whereas Love waves involve mainly horizontal motion in the transverse direction. We demean the raw seismograms and apply a bandpass filter between




Figure 7. Map of the northern Los Angeles basins, adapted from Zou and Clayton (2024). The study region is outlined by the white polygon. Triangles in different colors show nodal lines deployed in different time periods.

0.1 and 2 Hz. Prior to the correlation, we perform temporal normalization to suppress the effects of earthquakes and apply spectral whitening to broaden the effective bandwidth.

The EGF obtained from cross-correlation accounts for wave physics in a 3D medium, to which a phase shift of $\pi/4$ must be applied prior to a 2D inversion (Forbriger et al., 2014; Schäfer et al., 2014; C. Zhang et al., 2018). Figure S4 in Supporting Information S1 shows the effect of this transformation. A factor for amplitude correction related to geometrical spreading is further needed (C. Zhang et al., 2018), but we focus solely on the phase information, as the amplitude recovered in ANT is not well-defined. We have verified through synthetic tests that the phase information is sufficient to reconstruct the velocity models (Figure 6). We eliminate the effect of amplitude differences between 2D and 3D by normalizing the amplitude to unity during inversion. This also helps to bridge the gap between the synthetic Green's function (SGF) and EGF, as they have different source bandwidths. We match the Z-Z and Z-R cross-correlations with the vertical and horizontal waveforms predicted by the HNO, respectively. We mainly retrieve Rayleigh waves with particle motion confined in a 2D plane. In this study, we follow the principle of Green's function retrieval, while Tromp et al. (2010), Sager et al. (2018, 2020), and Tsai et al. (2024) established the full waveform ambient noise inversion technique that goes beyond that by interpreting correlation functions as self-consistent observables.

We maintain most settings from the previous synthetic FWI results in carrying out the real data experiments, except we include an additional regularization term that encourages the inverted V_P and V_S to respect Brocher (2005):

$$L_{reg} = \frac{1}{M} \sum_{i=1}^M (V_{P_i} - (0.9409 + 2.0947V_{S_i} - 0.8206V_{S_i}^2 + 0.2683V_{S_i}^3 - 0.0251V_{S_i}^4))^2, \quad (11)$$

where M is the number of grid points on which the velocity models are discretized. This is to constrain V_P , which is not well constrained by surface waves.

We begin the real data FWI with the longest line, SB1, which extends about 80 km from west to east. Beneath SB1 lie two significant basins—the deeper SG basin in the west and the shallower Chino basin in the east. The depth of the sediment-basement interface has been constrained by Villa et al. (2023). In addition to this 3D basin depth model and the CVM-S4.26 (Lee et al., 2014), another reference is from Y. Li et al. (2023) who constructed a 3D shear wave velocity model for the same area through surface wave dispersion analysis. We start the inversion with a 1D model calculated from this reference model, which results in Figure 8a. Overall, the inverted velocity models reveal basin structures that align well with Villa et al. (2023), capturing significant features including the deeper and more heterogeneous SG basin in the west and the low-velocity zone at the edge of the relatively flat Chino basin adjacent to the SB basin. We demonstrate the robustness of this result through another two initial models, the smoothed models from Y. Li et al. (2023) (Figure 8b) and CVM-S4.26 (Figure 8c). Results from these different initial models show excellent overall consistency, although they differ in details. In particular, the 1D




Figure 8. Inversion results for SB1 with different initial models including the (a) 1D and (b) smoothed models from Y. Li et al. (2023), and (c) the smoothed CVM-S4.26. The along-profile distance increases from west to east. The basin bottom from Villa et al. (2023) is delineated with black dashed lines for reference.




Figure 9. Examples of SALVUS-simulated waveforms (SGFs) using the CVM-S4.26, model from Y. Li et al. (2023), and our model from Figure 8a compared with the EGF for SB1. Only V_s is plotted in the top left for simplicity. The frequency band of the data is 0.2–0.5 Hz. In plotting, the waveform is normalized by trace and every third trace is displayed.

initial model can achieve performance comparable to that of models with more detailed priors. The corresponding error curves are provided in Figure S5 in Supporting Information S1.

The incorporation of higher frequency data in later stages of FWI contributes to refining the shallower structure but also introduces high-frequency noise. Some of the noise can potentially be suppressed using a signal window. Meanwhile, the inaccuracy of the neural operator in forward modeling introduces another source of uncertainty. We compare SALVUS-simulated waveforms (SGFs) using the CVM-S4.26, model from Y. Li et al. (2023), and our model from Figure 8a with the EGF for SB1 in Figure 9. The corresponding waveform difference is plotted in Figure S6 in Supporting Information S1. The point here is not to conclude that our model is better than others, but rather to demonstrate that the HNO-AD method can achieve accuracy comparable to traditional methods at a much lower cost in terms of time and effort. Also note that our model starts with a 1D initial model, while the other two models have stronger priors (Lee et al., 2014; Y. Li et al., 2023).

The same neural operator can be directly applied to any seismic line in the study region, because it was trained to cover the longest one. While the model cannot extrapolate to larger domains than it was trained on, application to smaller domains is straightforward by simply extracting a subset of the inversion result. Additionally, no specific orientation is assumed for the 2D profiles, since the training data are generated in a Cartesian coordinate system without reference to geological orientation. This ensures that the model is not biased toward any particular direction. We show another two examples for SB4 and SG1 in Figures 10 and 11, respectively. For SB4 that crossed the Chino basin from north to south, our inversion result reveals a low-velocity zone at the southern end near the Chino fault (Figure 7), which is also seen in the CVM-S4.26 and Y. Li et al. (2023). The low-velocity anomalies between 0 and 15 km present in Y. Li et al. (2023) are absent in both our model and CVM-S4.26. While the authors did not provide specific interpretation, the seismic line crossed the Red Hill fault at approximately 5 km,




Figure 10. Inversion results for SB4, in comparison to the CVM-S4.26 and model from Y. Li et al. (2023). Only V_s is plotted for simplicity. The along-profile distance increases from north to south. The basin bottom from Villa et al. (2023) is delineated with black dashed lines for reference.

which could explain the observed discontinuity (if it is not an artifact). Interestingly, our model shows a velocity jump near the Red Hill fault. Overall, the tomography for SB4 displays a shallow (approximately 2 km) sedimentary basin in agreement with previous studies. The sediment-basement interface beneath SG1 is much deeper, as captured in all the displayed models. Our result for SG1 aligns with that for SB1 in the SG basin in terms of basin depth but captures less heterogeneity, which might be limited by the array length. These inversions can be completed in a few minutes.

5. Discussion

In this study, we have demonstrated the functionality of neural operators in realistic ambient noise full waveform inversion through a 2D case. Provided sufficient computational resources, the proposed method is expected to




Figure 11. Inversion results for SG1, in comparison to the CVM-S4.26 and model from Y. Li et al. (2023). Only V_s is plotted for simplicity. The along-profile distance increases from north to south. The basin bottom from Villa et al. (2023) is delineated with black dashed lines for reference.




Figure 12. Training distributions of V_P and V_S as a function of depth.

work in 3D without fundamental changes, while providing increased computational advantages over conventional methods. Because our model is trained within the framework of supervised learning, the fundamental assumption is that training and testing data follow the same distribution. While this assumption seldom holds strictly in practice, it serves as a reference for identifying when model predictions are likely to fail—for example, in cases where the velocity parameters deviate significantly from the training distribution. We show the training distributions of V_P and V_S as a function of depth in Figure 12, allowing users to decide whether to directly use our model or perform additional training. Even in cases where re-training is necessary, a pre-trained model enables transfer learning that can greatly improve learning efficiency (Yosinski et al., 2014). Another aspect is that the model is not expected to extrapolate to larger physical domains than which it was trained on (80 km length by 40 km depth) or to frequencies outside the preset band of interest.

In performing FWI, we use gradient descent with AD instead of training a reverse model that directly maps from observed waveforms to velocity parameters. Although such end-to-end approaches are even faster (Kazei et al., 2021; Moseley et al., 2020; Wu & Lin, 2019), they provide only a single solution without supervision on the data misfit. Such approaches are arguably not ideal for inverse problems because they are ill-posed and non-unique, unless they are combined with uncertainty quantification methods. We define the misfit function for inversion as the MSE in the frequency domain for convenient use with HNO, but it would be feasible and worthwhile to explore different misfit functions in either the frequency or time domain (Bozdağ et al., 2011; Métivier et al., 2018; Fournier & Oldenburg, 2019; Sambridge et al., 2022; B. Sun & Alkhalifah, 2019) to better tackle the non-uniqueness and cycle-skipping problems. It is important to note, however, that the AD method requires a differentiable loss function, or some soft approximation to address discontinuity (Y. Liu et al., 2023).

The rapid developments in compute capabilities and machine learning research have positioned data-driven approaches as a promising new perspective for FWI. On one hand, data-driven approaches can greatly accelerate modeling and inversion. On the other hand, they rely on data provided by physics-driven approaches and, therefore, will also benefit from any advancements in those approaches. In this study, we integrate these two types of approaches by first generating training data with a physics-driven solver and then training a data-driven model as a surrogate solver. Another way is to incorporate physics into the training loss, making it possible to remove the external solver (Z. Li et al., 2024; Raissi et al., 2019; Rasht-Behesht et al., 2022; Ren et al., 2024; Song & Alkhalifah, 2021). While training with physics loss alone is challenging, combining it with a data loss has the potential to improve generalization ability and further benefit inversion.

6. Conclusions

We present the first application of neural operators in ambient noise full waveform inversion using several real seismic data sets from linear nodal arrays distributed across the northern LA basins. We show that a Helmholtz Neural Operator trained on random velocity fields can generalize to realistic velocity structures, while demonstrating robustness to random noise. In performing FWI with automatic differentiation, the neural operator is two orders of magnitude faster than the conventional spectral element method and eliminates the need for manual gradient derivation. The HNO-AD method, built on PyTorch, offers enhanced flexibility to escape poor local minima. The tomography results from real data align with multiple previous studies and can be obtained with minimal additional effort. We provide a trained model that can be directly applied to different regions, as long as the velocity parameters fall within the training distribution.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Data Availability Statement

The node data can be accessed from the IRIS Data Management Center (BASIN, 2018). The network codes are 4M and 6J. The code, models, and processed data are available at https://github.com/caifeng-zou/ANFWI_HNO (Zou, 2025).

Acknowledgments

ZER acknowledges financial support from the David and Lucile Packard Foundation. CZ and RWC acknowledge support from NSF EAR-2105358 and EAR-2438773. FCL acknowledges financial support from the Statewide California Earthquake Center (SCEC) based on Award Number DE-SC0016520 from the U.S. Department of Energy and NSF EAR-2438772. We sincerely thank Fenglin Niu as the editor, Tariq Alkhalifah, and an anonymous reviewer for their careful reading and constructive feedback, which led to substantial improvements in this paper. We thank all the volunteers who helped deploy the nodal arrays and Patricia Persaud for coordinating the BASIN project.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., et al. (2016). TensorFlow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)* (pp. 265–283). Retrieved from <https://www.usenix.org/conference/osdi16/technical-sessions/presentation/abadi>
- Afanasiev, M., Boehm, C., van Driel, M., Krischer, L., Rietmann, M., May, D. A., et al. (2019). Modular and flexible spectral-element waveform modelling in two and three dimensions. *Geophysical Journal International*, 216(3), 1675–1692. <https://doi.org/10.1093/gji/ggy469>
- Alkhalifah, T. A. (2016). Full waveform inversion in an anisotropic world: Where are the parameters hiding? *EAGE*. <https://doi.org/10.3997/9789462822023>
- Azizzadenesheli, K., Kovachki, N., Li, Z., Liu-Schiaffini, M., Kossaiji, J., & Anandkumar, A. (2024). Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, 6(5), 320–328. <https://doi.org/10.1038/s42254-024-00712-5>
- BASIN. (2018). San Gabriel and San Bernardino basin arrays [Dataset]. <https://doi.org/10.7909/46a0-ma59>
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. (2018). Automatic differentiation in machine learning: A survey. *Journal of Machine Learning Research*, 18(153), 1–43.
- Bensen, G., Ritzwoller, M., Barmin, M., Levshin, A. L., Lin, F., Moschetti, M., et al. (2007). Processing seismic ambient noise data to obtain reliable broad-band surface wave dispersion measurements. *Geophysical Journal International*, 169(3), 1239–1260. <https://doi.org/10.1111/j.1365-246x.2007.03374.x>
- Bernal-Romero, M., & Iturrarán-Viveros, U. (2021). Accelerating full-waveform inversion through adaptive gradient optimization methods and dynamic simultaneous sources. *Geophysical Journal International*, 225(1), 97–126.
- Beydoun, W. B., & Tarantola, A. (1988). First born and Rytov approximations: Modeling and inversion conditions in a canonical example. *Journal of the Acoustical Society of America*, 83(3), 1045–1055. <https://doi.org/10.1121/1.396537>
- Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010: 19th International Conference on Computational Statistics Paris France, August 22-27, 2010 Keynote, Invited and Contributed Papers* (pp. 177–186).
- Bozdağ, E., Trampert, J., & Tromp, J. (2011). Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements. *Geophysical Journal International*, 185(2), 845–870. <https://doi.org/10.1111/j.1365-246x.2011.04970.x>
- Brocher, T. M. (2005). Empirical relations between elastic wavespeeds and density in the Earth's crust. *Bulletin of the Seismological Society of America*, 95(6), 2081–2092. <https://doi.org/10.1785/0120050077>
- Chen, M., Huang, H., Yao, H., van der Hilst, R., & Niu, F. (2014). Low wave speed zones in the crust beneath SE Tibet revealed by ambient noise adjoint tomography. *Geophysical Research Letters*, 41(2), 334–340. <https://doi.org/10.1002/2013gl058476>
- Cheng, S., Taufik, M. H., & Alkhalifah, T. (2025). Seismic wavefield solutions via physics-guided generative neural operator. arXiv preprint arXiv:2503.06488.
- Clayton, R., Persaud, P., Denolle, M., & Polet, J. (2019). Exposing Los Angeles's shaky geologic underbelly. *Eos*, 100. <https://doi.org/10.1029/2019eo135099>
- Elliott, C. (2018). The simple essence of automatic differentiation. *Proceedings of the ACM on Programming Languages*, 2(ICFP), 1–29. <https://doi.org/10.1145/3236765>
- Fichtner, A., Bunge, H.-P., & Igel, H. (2006a). The adjoint method in seismology: I. Theory. *Physics of the Earth and Planetary Interiors*, 157(1–2), 86–104. <https://doi.org/10.1016/j.pepi.2006.03.016>
- Fichtner, A., Bunge, H.-P., & Igel, H. (2006b). The adjoint method in seismology—II. Applications: Traveltimes and sensitivity functionals. *Physics of the Earth and Planetary Interiors*, 157(1–2), 105–123. <https://doi.org/10.1016/j.pepi.2006.03.018>
- Forbriger, T., Groos, L., & Schäfer, M. (2014). Line-source simulation for shallow-seismic data. Part 1: Theoretical background. *Geophysical Journal International*, 198(3), 1387–1404. <https://doi.org/10.1093/gji/ggu199>
- Fournier, D., & Oldenburg, D. W. (2019). Inversion using spatially variable mixed Lp norms. *Geophysical Journal International*, 218(1), 268–282. <https://doi.org/10.1093/gji/ggz156>
- Gauthier, O., Virieux, J., & Tarantola, A. (1986). Two-dimensional nonlinear inversion of seismic waveforms: Numerical results. *Geophysics*, 51(7), 1387–1403. <https://doi.org/10.1190/1.1442188>
- Ghose, R., Persaud, P., & Clayton, R. W. (2023). Basin structure for earthquake ground motion estimates in urban Los Angeles mapped with nodal receiver functions. *Geosciences*, 13(11), 320. <https://doi.org/10.3390/geosciences13110320>
- Guo, Z., Chen, Y. J., Ning, J., Feng, Y., Grand, S. P., Niu, F., et al. (2015). High resolution 3-D crustal structure beneath NE China from joint inversion of ambient noise and receiver functions using NECESSArray data. *Earth and Planetary Science Letters*, 416, 1–11. <https://doi.org/10.1016/j.epsl.2015.01.044>
- Hendrycks, D., & Gimpel, K. (2016). Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5), 359–366. [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8)
- Huang, X., & Alkhalifah, T. (2025). Learned frequency-domain scattered wavefield solutions using neural operators. *Geophysical Journal International*, 241(3), 1467–1478. <https://doi.org/10.1093/gji/ggaf113>
- Kazei, V., Ovcharenko, O., Plotnitskii, P., Peter, D., Zhang, X., & Alkhalifah, T. (2021). Mapping full seismic waveforms to vertical velocity profiles by deep learning. *Geophysics*, 86(5), R711–R721. <https://doi.org/10.1190/geo2019-0473.1>
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Kong, Q., Zou, C., Choi, Y., Matzel, E. M., Azizzadenesheli, K., Ross, Z. E., et al. (2025). Reducing frequency bias of Fourier neural operators in 3D seismic wavefield simulations through multi-stage training. arXiv preprint arXiv:2503.02023.

- Kovachki, N. B., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A. M., & Anandkumar, A. (2023). Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89), 1–97.
- LeCun, Y., Bottou, L., Orr, G. B., & Müller, K.-R. (2002). Efficient backprop. In *Neural Networks: Tricks of the Trade* (pp. 9–50). Springer.
- LeCun, Y., Touresky, D., Hinton, G., & Sejnowski, T. (1988). A theoretical framework for back-propagation. In *Proceedings of the 1988 Connectionist Models Summer School* (Vol. 1, pp. 21–28).
- Lee, E.-J., Chen, P., Jordan, T. H., Maechling, P. B., Denolle, M. A., & Beroza, G. C. (2014). Full-3-D tomography for crustal structure in southern California based on the scattering-integral and the adjoint-wavefield methods. *Journal of Geophysical Research: Solid Earth*, 119(8), 6421–6451. <https://doi.org/10.1002/2014jb011346>
- Lehmann, F., Gatti, F., Bertin, M., & Clouteau, D. (2024). 3D elastic wave propagation with a factorized Fourier neural operator (F-FNO). *Computer Methods in Applied Mechanics and Engineering*, 420, 116718. <https://doi.org/10.1016/j.cma.2023.116718>
- Lehmann, F., Gatti, F., & Clouteau, D. (2025). Multiple-input Fourier neural operator (MIFNO) for source-dependent 3d elastodynamics. *Journal of Computational Physics*, 527, 113813. <https://doi.org/10.1016/j.jcp.2025.113813>
- Li, B., Wang, H., Feng, S., Yang, X., & Lin, Y. (2023). Solving seismic wave equations on variable velocity models with Fourier neural operator. *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1–18. <https://doi.org/10.1109/tgrs.2023.3333663>
- Li, C., Zhao, H., & Hao, Y. (2025). A feature enhanced autoencoder integrated with Fourier neural operator for intelligent elastic wavefield modeling. *IEEE Transactions on Geoscience and Remote Sensing*, 63, 1–16. <https://doi.org/10.1109/tgrs.2025.3542082>
- Li, Y., Villa, V., Clayton, R. W., & Persaud, P. (2023). Shear wave velocities in the San Gabriel and San Bernardino Basins, California. *Journal of Geophysical Research: Solid Earth*, 128(7), e2023JB026488. <https://doi.org/10.1029/2023jb026488>
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2020a). Fourier neural operator for parametric partial differential equations. arXiv preprint arXiv:2010.08895.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2020b). Neural operator: Graph kernel network for partial differential equations. arXiv preprint arXiv:2003.03485.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Stuart, A., Bhattacharya, K., & Anandkumar, A. (2020). Multipole graph neural operator for parametric partial differential equations. *Advances in Neural Information Processing Systems*, 33, 6755–6766.
- Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., et al. (2024). Physics-informed neural operator for learning partial differential equations. *ACM/JMS Journal of Data Science*, 1(3), 1–27. <https://doi.org/10.1145/3648506>
- Lin, F.-C., Moschetti, M. P., & Ritzwoller, M. H. (2008). Surface wave tomography of the western United States from ambient seismic noise: Rayleigh and Love wave phase velocity maps. *Geophysical Journal International*, 173(1), 281–298. <https://doi.org/10.1111/j.1365-246x.2008.03720.x>
- Liu, Q., & Tromp, J. (2006). Finite-frequency kernels based on adjoint methods. *Bulletin of the Seismological Society of America*, 96(6), 2383–2397. <https://doi.org/10.1785/0120060041>
- Liu, X. (2020). Finite-frequency sensitivity kernels for seismic noise interferometry based on differential time measurements. *Journal of Geophysical Research: Solid Earth*, 125(4), e2019JB018932. <https://doi.org/10.1029/2019jb018932>
- Liu, X., Beroza, G. C., & Li, H. (2023). Ambient noise differential adjoint tomography reveals fluid-bearing rocks near active faults in Los Angeles. *Nature Communications*, 14(1), 6873. <https://doi.org/10.1038/s41467-023-42536-4>
- Liu, Y., Niu, F., Chen, M., & Yang, W. (2017). 3-D crustal and uppermost mantle structure beneath NE China revealed by ambient noise adjoint tomography. *Earth and Planetary Science Letters*, 461, 20–29. <https://doi.org/10.1016/j.epsl.2016.12.029>
- Liu, Y., Tang, J., Tang, Z., & Sun, C. (2023). Robust full-waveform inversion based on automatic differentiation and differentiable dynamic time warping. *Journal of Geophysics and Engineering*, 20(3), 549–564. <https://doi.org/10.1093/jge/gxad029>
- Louboutin, M., Lange, M., Luporini, F., Kukreja, N., Witte, P. A., Herrmann, F. J., et al. (2019). Devito (v3.1.0): An embedded domain-specific language for finite differences and geophysical exploration. *Geoscientific Model Development*, 12(3), 1165–1187. <https://doi.org/10.5194/gmd-12-1165-2019>
- Maguire, R., Schmandt, B., Li, J., Jiang, C., Li, G., Wilgus, J., & Chen, M. (2022). Magma accumulation at depths of prior rhyolite storage beneath Yellowstone Caldera. *Science*, 378(6623), 1001–1004. <https://doi.org/10.1126/science.ade0347>
- Mao, S., Song, P., Tong, S., Tan, J., Xie, C., Zu, G., & Liu, G. (2025). Automatic differentiation-based full waveform inversion of anisotropic parameters in TTI media. *IEEE Transactions on Geoscience and Remote Sensing*, 63, 1–15. <https://doi.org/10.1109/TGRS.2025.3586294>
- Métivier, L., Allain, A., Brossier, R., Mérigot, Q., Oudet, E., & Virieux, J. (2018). Optimal transport for mitigating cycle skipping in full-waveform inversion: A graph-space transform approach. *Geophysics*, 83(5), R515–R540. <https://doi.org/10.1190/geo2017-0807.1>
- Moseley, B., Nissen-Meyer, T., & Markham, A. (2020). Deep learning for fast simulation of seismic waves in complex media. *Solid Earth*, 11(4), 1527–1549. <https://doi.org/10.5194/se-11-1527-2020>
- Panning, M. P., Capdeville, Y., & Romanowicz, B. A. (2009). Seismic waveform modelling in a 3-D Earth using the born approximation: Potential shortcomings and a remedy. *Geophysical Journal International*, 177(1), 161–178. <https://doi.org/10.1111/j.1365-246x.2008.04050.x>
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). PyTorch: An imperative style, high-performance deep learning library. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Curran Associates Inc.
- Pladys, A., Brossier, R., Li, Y., & Métivier, L. (2021). On cycle-skipping and misfit function modification for full-wave inversion: Comparison of five recent approaches. *Geophysics*, 86(4), R563–R587. <https://doi.org/10.1190/geo2020-0851.1>
- Plessix, R.-E. (2006). A review of the adjoint-state method for computing the gradient of a functional with geophysical applications. *Geophysical Journal International*, 167(2), 495–503. <https://doi.org/10.1111/j.1365-246x.2006.02978.x>
- Rahman, M. A., Ross, Z. E., & Azizzadenesheli, K. (2022). U-NO: U-shaped neural operators. arXiv preprint arXiv:2204.11127.
- Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- Rasht-Behesht, M., Huber, C., Shukla, K., & Karniadakis, G. E. (2022). Physics-informed neural networks (PINNs) for wave propagation and full waveform inversions. *Journal of Geophysical Research: Solid Earth*, 127(5), e2021JB023120. <https://doi.org/10.1029/2021jb023120>
- Ren, P., Rao, C., Chen, S., Wang, J.-X., Sun, H., & Liu, Y. (2024). SeismicNet: Physics-informed neural networks for seismic wave modeling in semi-infinite domain. *Computer Physics Communications*, 295, 109010. <https://doi.org/10.1016/j.cpc.2023.109010>
- Richardson, A. (2018). Seismic full-waveform inversion using deep learning tools and techniques. arXiv preprint arXiv:1801.07232.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* (Vol. 18, pp. 234–241). https://doi.org/10.1007/978-3-319-24574-4_28
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533–536. <https://doi.org/10.1038/323533a0>

- Sager, K., Boehm, C., Ermert, L., Krischer, L., & Fichtner, A. (2020). Global-scale full-waveform ambient noise inversion. *Journal of Geophysical Research: Solid Earth*, 125(4), e2019JB018644. <https://doi.org/10.1029/2019jb018644>
- Sager, K., Ermert, L., Boehm, C., & Fichtner, A. (2018). Towards full waveform ambient noise inversion. *Geophysical Journal International*, 212(1), 566–590. <https://doi.org/10.1093/gji/ggx429>
- Sambridge, M., Jackson, A., & Valentine, A. P. (2022). Geophysical inversion and optimal transport. *Geophysical Journal International*, 231(1), 172–198. <https://doi.org/10.1093/gji/ggac151>
- Schäfer, M., Groos, L., Forbriger, T., & Bohlen, T. (2014). Line-source simulation for shallow-seismic data. Part 2: Full-waveform inversion—A synthetic 2-D case study. *Geophysical Journal International*, 198(3), 1405–1418. <https://doi.org/10.1093/gji/ggu171>
- Song, C., & Alkhalifah, T. A. (2021). Wavefield reconstruction inversion via physics-informed neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–12. <https://doi.org/10.1109/tgrs.2021.3123122>
- Sun, B., & Alkhalifah, T. (2019). The application of an optimal transport to a preconditioned data matching function for robust waveform inversion. *Geophysics*, 84(6), R923–R945. <https://doi.org/10.1190/geo2018-0413.1>
- Sun, J., Niu, Z., Innanen, K. A., Li, J., & Trad, D. O. (2020). A theory-guided deep-learning formulation and optimization of seismic waveform inversion. *Geophysics*, 85(2), R87–R99. <https://doi.org/10.1190/geo2019-0138.1>
- Tape, C., Liu, Q., Maggi, A., & Tromp, J. (2010). Seismic tomography of the southern California crust based on spectral-element and adjoint methods. *Geophysical Journal International*, 180(1), 433–462. <https://doi.org/10.1111/j.1365-246x.2009.04429.x>
- Thrustarson, S., Van Herwaarden, D.-P., Krischer, L., Boehm, C., van Driel, M., Afanasiev, M., & Fichtner, A. (2022). Data-adaptive global full-waveform inversion. *Geophysical Journal International*, 230(2), 1374–1393. <https://doi.org/10.1093/gji/ggac122>
- Tromp, J., Luo, Y., Hanasoge, S., & Peter, D. (2010). Noise cross-correlation sensitivity kernels. *Geophysical Journal International*, 183(2), 791–819. <https://doi.org/10.1111/j.1365-246x.2010.04721.x>
- Tromp, J., Tape, C., & Liu, Q. (2005). Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels. *Geophysical Journal International*, 160(1), 195–216. <https://doi.org/10.1111/j.1365-246x.2004.02453.x>
- Tsai, V. C., Sager, K., & Bowden, D. C. (2024). Towards limited-domain full waveform ambient noise inversion. *Geophysical Journal International*, 237(2), 965–973. <https://doi.org/10.1093/gji/ggae091>
- van Herwaarden, D. P., Boehm, C., Afanasiev, M., Thrustarson, S., Krischer, L., Trampert, J., & Fichtner, A. (2020). Accelerated full-waveform inversion using dynamic mini-batches. *Geophysical Journal International*, 221(2), 1427–1438. <https://doi.org/10.1093/gji/ggaa079>
- Villa, V., Li, Y., Clayton, R. W., & Persaud, P. (2023). Three-dimensional basin depth map of the northern Los Angeles basins from gravity and seismic measurements. *Journal of Geophysical Research: Solid Earth*, 128(7), e2022JB025425. <https://doi.org/10.1029/2022jb025425>
- Virieux, J., & Operto, S. (2009). An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6), WCC1–WCC26. <https://doi.org/10.1190/1.3238367>
- Wang, K., Yang, Y., Basini, P., Tong, P., Tape, C., & Liu, Q. (2018). Refined crustal and uppermost mantle structure of southern California by ambient noise adjoint tomography. *Geophysical Journal International*, 215(2), 844–863. <https://doi.org/10.1093/gji/ggy312>
- Wang, K., Yang, Y., Jiang, C., Wang, Y., Tong, P., Liu, T., & Liu, Q. (2021). Adjoint tomography of ambient noise data and teleseismic P waves: Methodology and applications to central California. *Journal of Geophysical Research: Solid Earth*, 126(6), e2021JB021648. <https://doi.org/10.1029/2021jb021648>
- Wang, W., McMechan, G. A., & Ma, J. (2021). Elastic isotropic and anisotropic full-waveform inversions using automatic differentiation for gradient calculations in a framework of recurrent neural networks. *Geophysics*, 86(6), R795–R810. <https://doi.org/10.1190/geo2020-0542.1>
- Wapenaar, K. (2004). Retrieving the elastodynamic Green's function of an arbitrary inhomogeneous medium by cross correlation. *Physical Review Letters*, 93(25), 254301. <https://doi.org/10.1103/physrevlett.93.254301>
- Wapenaar, K., Slob, E., Snieder, R., & Curtis, A. (2010). Tutorial on seismic interferometry: Part 2—Underlying theory and new advances. *Geophysics*, 75(5), 75A211–75A227. <https://doi.org/10.1190/1.3463440>
- Wei, W., & Fu, L.-Y. (2022). Small-data-driven fast seismic simulations for complex media using physics-informed Fourier neural operators. *Geophysics*, 87(6), T435–T446. <https://doi.org/10.1190/geo2021-0573.1>
- Wright, T. L. (1991). Structural geology and tectonic evolution of the Los Angeles basin, California. In K. T. Biddle (Ed.), *Active Margin Basins* (Vol. 52, pp. 35–79). American Association of Petroleum Geologists.
- Wu, Y., & Lin, Y. (2019). InversionNet: An efficient and accurate data-driven full waveform inversion. *IEEE Transactions on Computational Imaging*, 6, 419–433. <https://doi.org/10.1109/tci.2019.2956866>
- Yang, Y., Gao, A. F., Azizzadenesheli, K., Clayton, R. W., & Ross, Z. E. (2023). Rapid seismic waveform modeling and inversion with neural operators. *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1–12. <https://doi.org/10.1109/tgrs.2023.3264210>
- Yang, Y., Gao, A. F., Castellanos, J. C., Ross, Z. E., Azizzadenesheli, K., & Clayton, R. W. (2021). Seismic wave propagation and inversion with neural operators. *The Seismic Record*, 1(3), 126–134. <https://doi.org/10.1785/0320210026>
- Yao, H., van Der Hilst, R. D., & Montagner, J.-P. (2010). Heterogeneity and anisotropy of the lithosphere of SE Tibet from surface wave array tomography. *Journal of Geophysical Research*, 115(B12). <https://doi.org/10.1029/2009jb007142>
- Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks? *Advances in Neural Information Processing Systems*, 27.
- Zhang, C., Yao, H., Liu, Q., Zhang, P., Yuan, Y. O., Feng, J., & Fang, L. (2018). Linear array ambient noise adjoint tomography reveals intense crust-mantle interactions in North China Craton. *Journal of Geophysical Research: Solid Earth*, 123(1), 368–383. <https://doi.org/10.1002/2017jb015019>
- Zhang, T., Trad, D., & Innanen, K. (2023). Learning to solve the elastic wave equation with Fourier neural operators. *Geophysics*, 88(3), T101–T119. <https://doi.org/10.1190/geo2022-0268.1>
- Zhang, Z.-D., Alajami, M., & Alkhalifah, T. (2020). Wave-equation dispersion spectrum inversion for near-surface characterization using fibre-optics acquisition. *Geophysical Journal International*, 222(2), 907–918. <https://doi.org/10.1093/gji/ggaa211>
- Zhang, Z.-D., Saygin, E., He, L., & Alkhalifah, T. (2021). Rayleigh wave dispersion spectrum inversion across scales. *Surveys in Geophysics*, 42(6), 1281–1303. <https://doi.org/10.1007/s10712-021-09667-z>
- Zheng, Y., Shen, W., Zhou, L., Yang, Y., Xie, Z., & Ritzwoller, M. H. (2011). Crust and uppermost mantle beneath the North China Craton, northeastern China, and the Sea of Japan from ambient noise tomography. *Journal of Geophysical Research*, 116(B12), B12312. <https://doi.org/10.1029/2011jb008637>
- Zhu, H., Bozdağ, E., & Tromp, J. (2015). Seismic structure of the European upper mantle based on adjoint tomography. *Geophysical Journal International*, 201(1), 18–52. <https://doi.org/10.1093/gji/ggu492>
- Zhu, W., Xu, K., Darve, E., & Beroza, G. C. (2021). A general approach to seismic inversion with automatic differentiation. *Computers & Geosciences*, 151, 104751. <https://doi.org/10.1016/j.cageo.2021.104751>
- Zou, C. (2025). caifeng-zou/ANFWI_HNO: First release [Software]. *Zenodo*. <https://doi.org/10.5281/zenodo.15061657>

- Zou, C., Azizzadenesheli, K., Ross, Z. E., & Clayton, R. W. (2024). Deep neural Helmholtz operators for 3-D elastic wave propagation and inversion. *Geophysical Journal International*, 239(3), 1469–1484. <https://doi.org/10.1093/gji/ggae342>
- Zou, C., & Clayton, R. W. (2024). Imaging the northern Los Angeles basins with autocorrelations. *Seismological Research Letters*, 96(3), 1791–1801. <https://doi.org/10.1785/0220240140>