Seismic Simulations for Structure and Source Characterization in the Bay Area: Foundations for ML Acceleration of Waveform Modelling

Claire Doody[1] (doody1@llnl.gov), Jiun-Ting Lin[1], Qingkai Kong[1], Luis Vazquez[2], Caifeng Zou[3], Youngsoo Choi[1], Artie Rodgers[1], Kamyar Azizzadenesheli[4], Zach Ross[3], Rob Clayton[3]

[1]Lawrence Livermore National Laboratory, Livermore, CA, [2]Univeristy of Southern California, Los Angeles, CA, [3]California Institute of Technology, Pasadena, CA, [4]NVIDIA, Santa Clara, CA

Fourier Neural Operators (FNOs) are a machine learning technique that quickly and accurately determines solution operators of partial differential equations. This method shows great promise for improving the speed of forward calculations in full waveform seismic imaging and ambient noise tomography (e.g., Kong et al., 2025; Yang et al., 2021, 2023; Zou et al., 2024). We present the foundational work for an FNO model trained for a 160kmx160kmx80km region in the greater Bay Area. We computed training data for the FNO model using Salvus (Afanasiev et al., 2019) and present our current method for computing training data. We compare the results of our model trained on a subset of our training dataset with observed data from moderate-magnitude (Mw 3.5-6) Bay Area earthquakes to show good fit to data in our target period band of 5-30 seconds. We also show that the FNO architecture can be used for faster focal mechanism inversions, converging to similar solutions to published catalogues even if the model's first guess is significantly different from the final solution.

Introduction

Seismic imaging methods provide constraints on subsurface structure and improve waveform predictions. However, the computational cost of computing these models limits the ability to calculate uncertainty and push models to frequencies that are relevant to seismic ground motion modelling. Machine learning methods can tackle both uncertainty and higher frequencies efficiently at the cost of computing training data. In this study, we use Fourier Neural Operators (FNOs), a recently-developed machine learning method, to speed up forward calculations of synthetic data and source inversions in the greater Bay Area. We present the methodology of how we compute our training dataset to best capture the range of possible sources and structures in the Bay Area. We also show preliminary results of comparing the FNO-outputted synthetics with data from real events within our region of interest. Finally, we discuss how FNO models can also speed up focal mechanism inversions on top of inverting for material properties.

Fourier Neural Operators

- Fourier neural operators (FNOs) are a machine learning technique that can determine the solution operators of partial differential equations such as the wave equation used to forward model synthetic data through Earth models.

- **FNOs can operate on data in both the time domain (e.g., Yang et al., 2021, 2023) or in the frequency domain (e.g., Kong et al., 2025; Zou et al., 2024).**

- In this study, we solve the Helmholtz equation in the frequency domain for faster calculations in 3D. We solve for 11 discrete frequencies in our frequency band of interest.

- The FNO model takes these 11 frequencies, a velocity model, source location, and source properties (strike, dip, and rake) as inputs and outputs the real and imaginary wavefield (Figure 1).
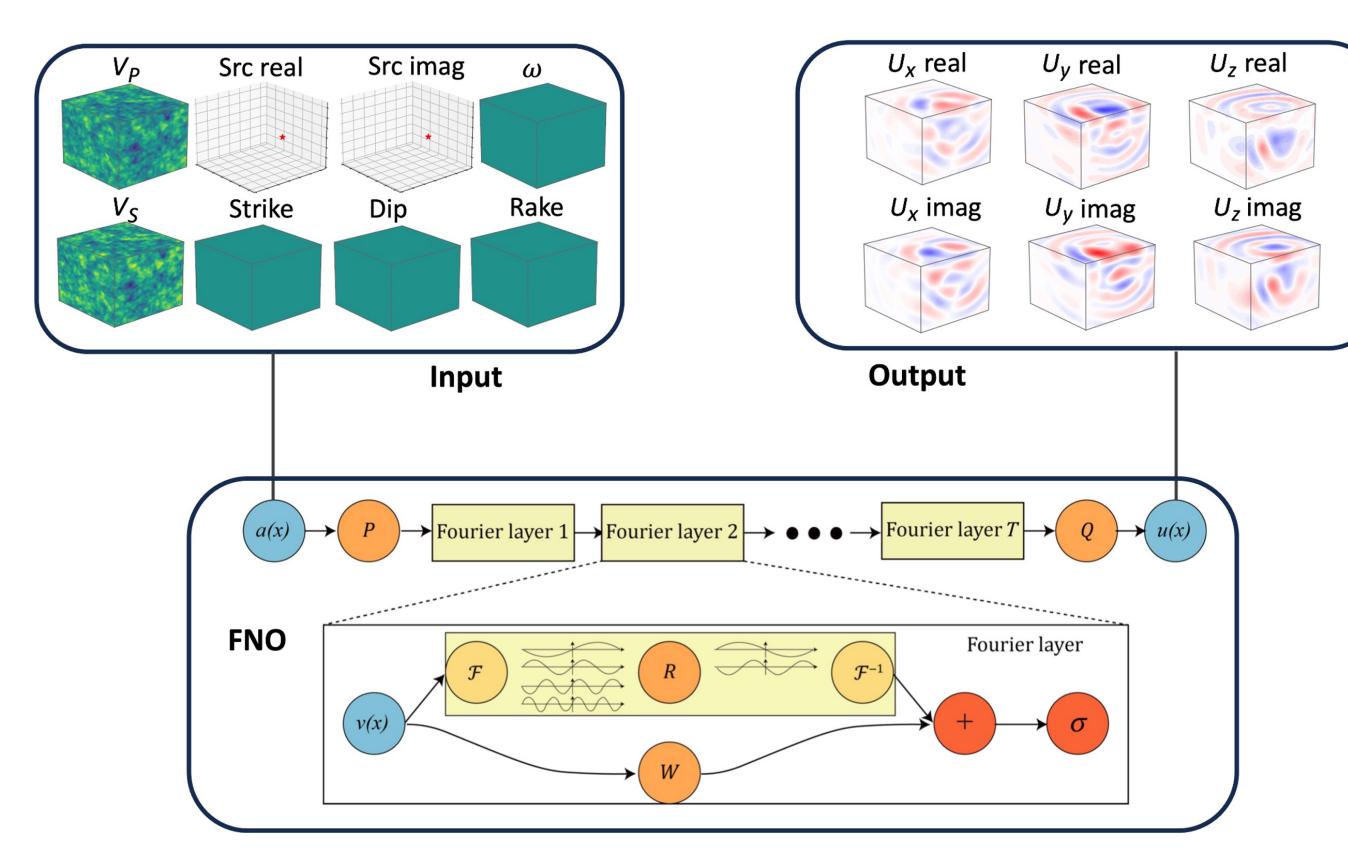

Figure 1. Schematic of FNO structure from Kong et al. (2025). Top Left: Input parameters of the Helmholtz FNO. Top Right: Wavefield outputs from the FNO. Bottom: Structure of the FNO layers used in the inversion.
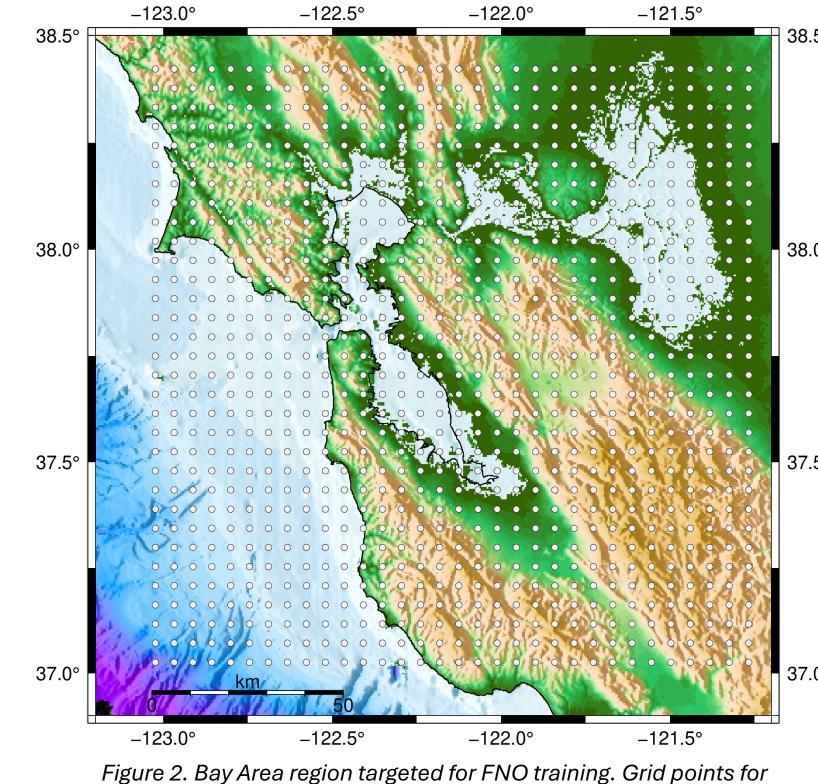
Generating Training Data for the Bay Area

Figure 2. Bay Area region targeted for FNO training. Grid points for FNO training are plotted as white dots.

Region of Interest

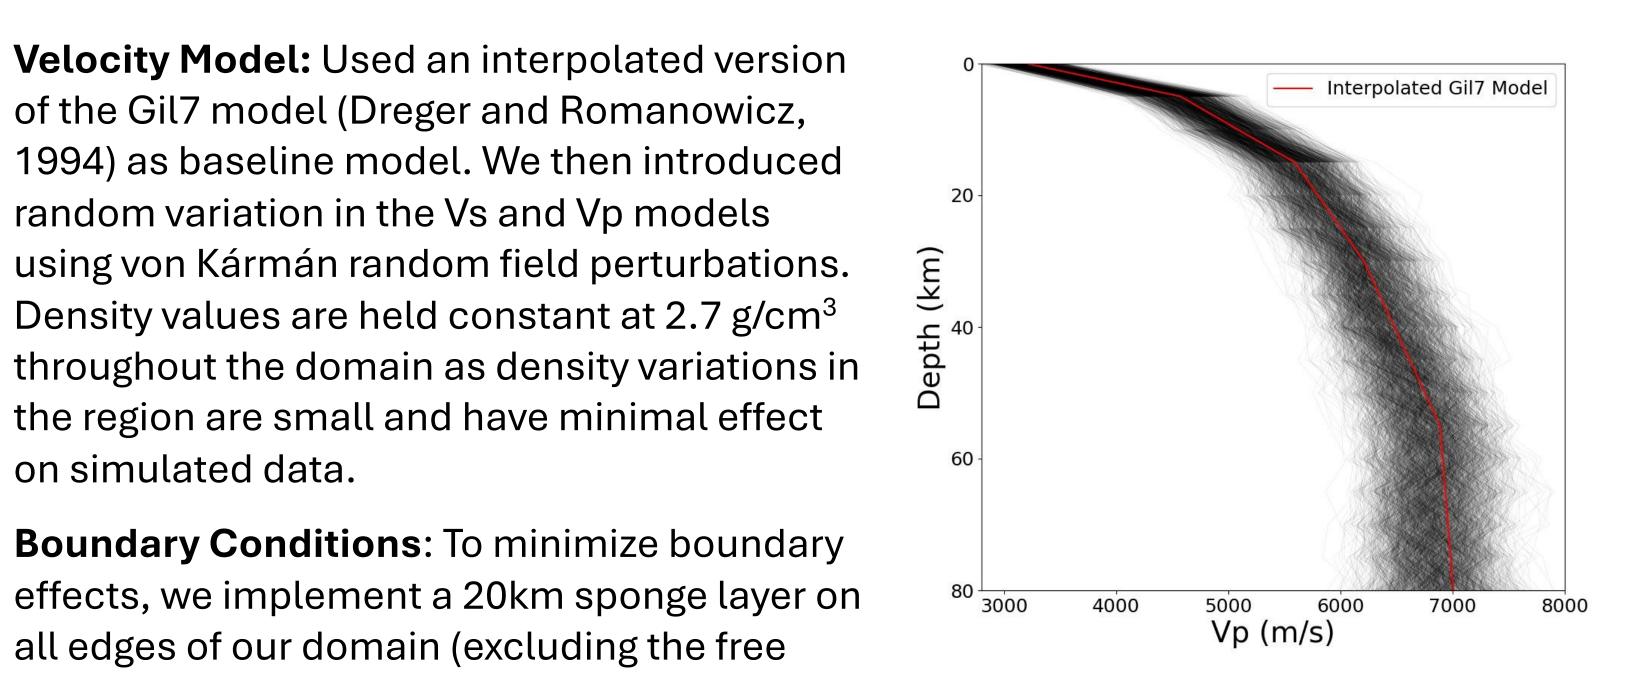- Training data calculated for 160kmx160kmx80km region in the greater Bay Area (Figure 2)

- Train on model discretized to 32x32x16 grid points, which translates to 5km grid spacing in all dimensions

Training Sources

- Use randomly generated strike, dip, and rake values placed at a random grid point

- Input sources are purely double-couple

- Moment is held fixed for all events at 10^{15} N-m

Velocity Model: Used an interpolated version of the Gil7 model (Dreger and Romanowicz, 1994) as baseline model. We then introduced random variation in the Vs and Vp models using von Kármán random field perturbations. Density values are held constant at 2.7 g/cm^3 throughout the domain as density variations in the region are small and have minimal effect on simulated data.

Boundary Conditions: To minimize boundary effects, we implement a 20km sponge layer on all edges of our domain (excluding the free surface) to absorb reflections and minimize artefacts in the simulated data.
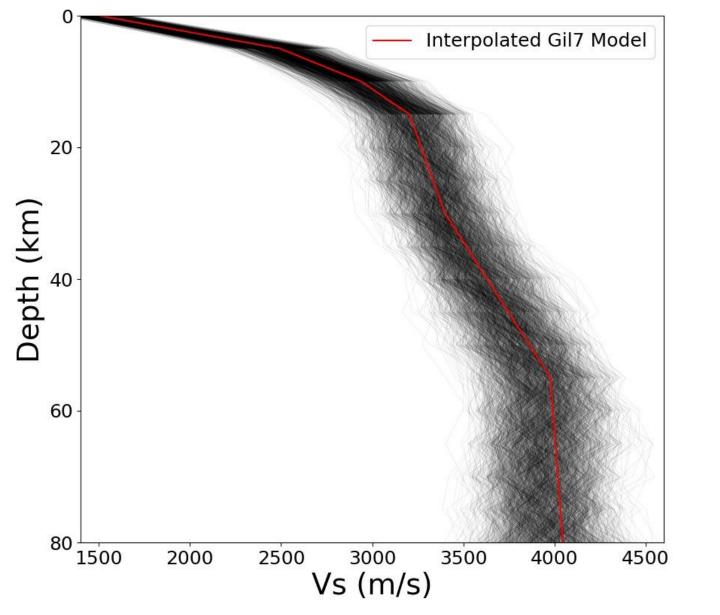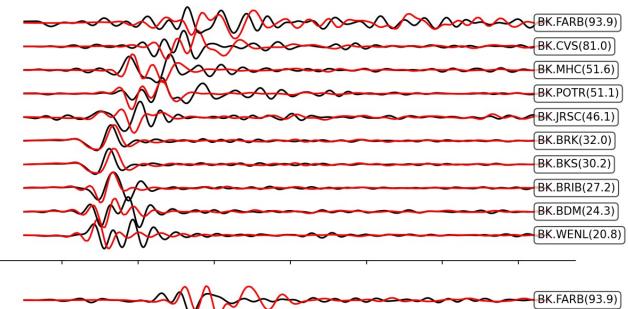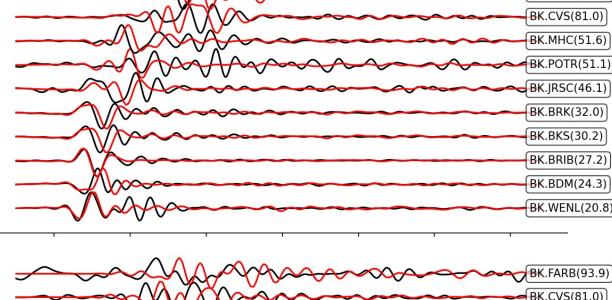

Figure 3. Vp and Vs values for interpolated Gil7 model used as the baseline model that we perturbed using von Karman random field perturbations(Dreger and Romanowicz, 1994; red line). Black lines represent the 1D velocity models at each grid point in an example training model.

Key Training Parameters

Frequency Band	0.033-0.2 seconds
Input Velocity Model	Modified Gil7 1D model
Absorbing Boundary Type	Sponge Layer
Absorbing Boundary Distance	20 km
Runtime per Salvus simulation	~8 seconds

Comparing Simulated and Observed Data

- We compare the simulated data produced by Salvus with observed data from an earthquake in our domain of interest.

- For Mw4.0 test event that occurred in Dublin, CA (roughly in center of the domain), waveforms recorded at stations within 50km of the event show good fit between synthetic and observed data.

- Errors propagate more at distant stations; training with a larger dataset will likely decrease these errors.

- **Path-specific errors**: Synthetics are shown between the nearest grid points to the event and station, so synthetic and observed data paths have slight discrepancies, which contribute to some of the misfit.
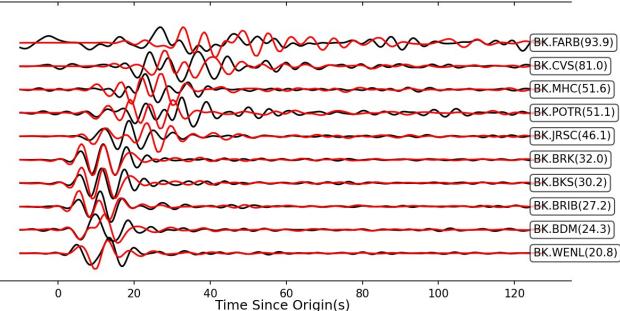

Figure 4. Forward waveforms (red waveforms) from a preliminary FNO model trained on the Bay Area for a Mw4.0 event that occurred on February 3, 2003 in Dublin, CA. Observed data are plotted as black waveforms. The corresponding station name and source-receiver distances (in km) are listed at the end of the waveform.

Source Inversion Results

- FNO model can also be trained to invert for source parameters (strike, dip, and rake).

- We tested source inversions on simulated and real data in the Bay Area.

- Regardless of the similarity of the input solution to the ground truth, the FNO model converges on the correct solution within 100 iteration. FNO will also sometimes converge on auxiliary plane.

- We are currently testing a version of the FNO model that will invert for strike, dip, rake, and magnitude simultaneously.
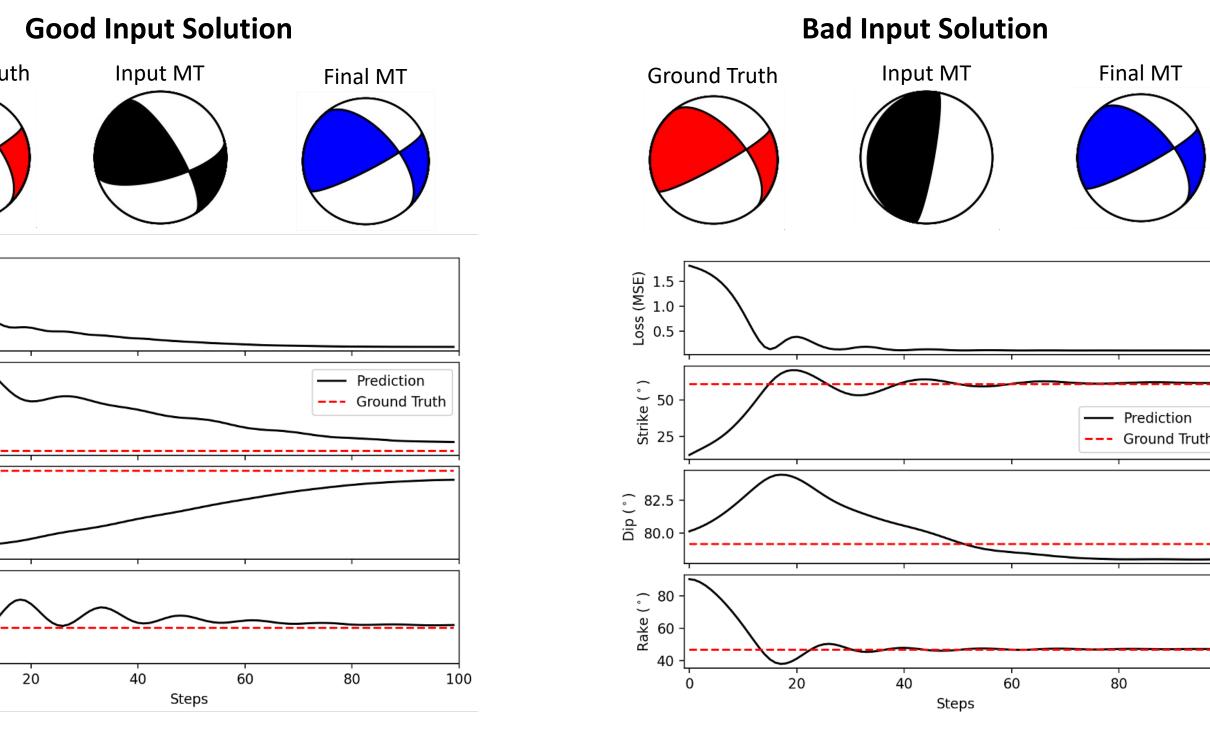

Figure 5. Source Inversion results using the FNO for a "good" starting solution (left panel) and for a "bad" starting solution (right panel). The ground truth moment tensor is plotted in red, the starting solution in black, and the moment tensor after 100 iterations in blue. Mean square error (MSE) loss, strike, dip, and rake values at each iteration are shown in the graphs below the moment tensors.

References: Afanasiev et al. (2019), Modular and flexible spectral-element waveform modelling in two and three dimensions, GJI; Dreger and Romanowicz (1994), Source characteristics of events in the San Francisco Bay region, USGS Open-File Report; Kong et al. (2025), Reducing frequency bias of Fourier Neural Operators in 3D seismic wavefield simulations through multistage training, SRL; Yang et al. (2021), Seismic wave propagation and inversion with neural operators, TSR; Yang et al. (2023), Rapid seismic waveform modeling and inversion with neural operators, IEEE; Zou et al. (2024), Deep neural Helmholtz operators for 3-D elastic wave propagation and inversion, GJI