Calgary • Canada • May 12-14 2025

Analysis of the mesh-invariant feature of Fourier Neural Operators and its application of learning to solve the acoustic wave equation

Angel Ramos, Tianze Zhang, Kristopher Innanen

University of Calgary

Summary

Numerical solvers have been shown to effectively provide approximate solutions to the wave equation. However, the discrete nature of these solutions introduces a trade-off between resolution and computational cost, typically limiting solutions to a fixed instance of the wave equation, such as a single velocity model. This approach complicates problem-solving when model parameters are subject to constant modification, as seen in methods like full-waveform inversion (FWI) and uncertainty quantification. The Fourier Neural Operators (FNOs), known for their generalization capabilities and mesh-invariant properties, have shown promise in offering accurate and efficient solutions to the wave equation at the expense of requiring large datasets for training. In this study, we analyze the mesh-invariant characteristics of FNOs through numerical experiments by evaluating their performance in learning the acoustic wave equation against the conventional finite-difference method (FDM), using low-resolution wavefields as training data, showing how this approach can potentially reduce the computational resources and time required for training, which is critical for enhancing generalization and robustness to model changes and potentially enabling faster forward modelling.

Theory

The finite difference method is a straightforward and easy-to-implement solution to the wave equation. In practice, stability and grid dispersion analysis conclude that an increment of spatial or temporal resolution will increase the computational cost to acquire these solutions. The increment of the resolution implies two things: 1) moving slower in time and 2) an increase in memory demand, without mention that as we increase the dimension of the velocity models, the computational cost becomes very prohibitive, making it necessary to look for alternative methods that can provide faster, and “low-cost” solutions. Methods such as FD were constructed by discretizing the space using a fixed mesh (e.g. Virieux J., 1986), which is not, in principle, a problem. Still, it imposes a solution for a single instance of the wave equation (fixed velocity model) and that specific type of discretization (fixed mesh). Its computational cost to recalculate the solution is what we can call a potential problem, especially when we are required to do several repetitions of forward modelling in large datasets.

The Fourier Neural Operators (FNOs) is a deep learning method that consists of learning to map PDE solutions between two infinitesimal function spaces from a finite collection of observed inputs-output pairs (Li et al. 2021a). The FNO sequence consists of a neural network-style architecture incorporating Fourier layers. The sequence of a Fourier layer can be expressed as:

$$v_{L+1} = \sigma[W v_L(x) \; + \; F^{-1}(R_\phi \cdot F(v_L(x)))\,]$$

Where v_L corresponds to the input data, R_ϕ and W are weight matrices or learning parameters, F and F^{-1} are forward and inverse Fourier transform operators, respectively, and σ is a non-linear activation function. The sequence also comprises uplifting and projection layers (P and Q) before and after the Fourier Layers sequence (Figure 1). These are linear transformations that help to increase the width of channels and project back to the input dimension space, respectively.

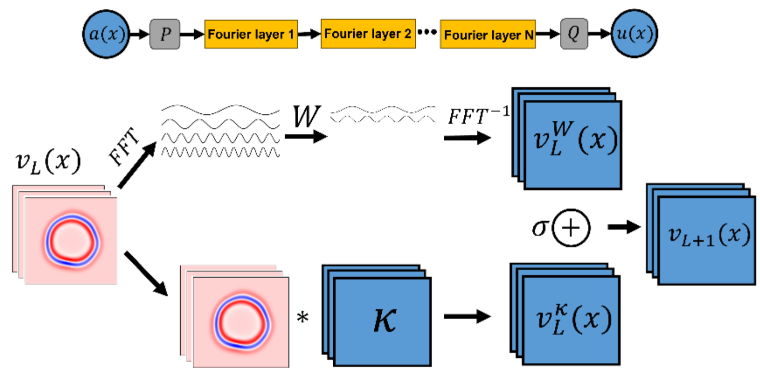

Figure 1: FNOs architecture and content of a Fourier Layer

One of the most interesting features of the FNO method is its capability to enable a mesh-invariant environment, which, if used correctly, can be a powerful complement to reduce training costs, enable faster predictions, and enhance the resolution of the forward simulations made with the FDM. Because of its construction, the FNO method can use the same learned network parameters to generate outputs with different resolutions or propagate the seismic waves in an arbitrary mesh. We can train our neural operators using low-resolution datasets (low-resolution wavefields) and directly evaluate them in high-resolution datasets. Indeed, the method has shown that if good generalization is achieved, it can produce wavefields with great accuracy and significantly faster than conventional solvers (Li et al., 2021a; Zhang et al., 2023; Brandstetter et al. 2022; Li et al. 2021b).

Numerical experiments

Although the FNO method has been applied to learn the wave equation (Zhang et al., 2023), we haven't found reports on the mesh-invariant analysis for the wave equation. So, we have made a simple but illustrative analysis of some of the results when using low-resolution data to train and produce high-resolution results while analyzing the limits, advantages and performance of this mesh-invariant feature.

Experiments were performed to test the method's ability to generate multi-resolution data on the one- and two-dimensional acoustic wave equation. For training, we generated 1000 random smooth velocity models to propagate acoustic waves; 900 of these models were used for training, and the remaining 100 were used to evaluate the performance of the learned operators against unknown velocity models. The wavefields and velocity model used for the training had an original spatial resolution of 512 grid points; this original dataset was decimated or downsampled to achieve training resolutions of 64 grid points. After training, we evaluated the trained operators on higher resolutions such as x2, x4, and x8. The results of this experiment and the relative RMSE as we increase the spatial resolution can be observed in Figures 2 and 3.

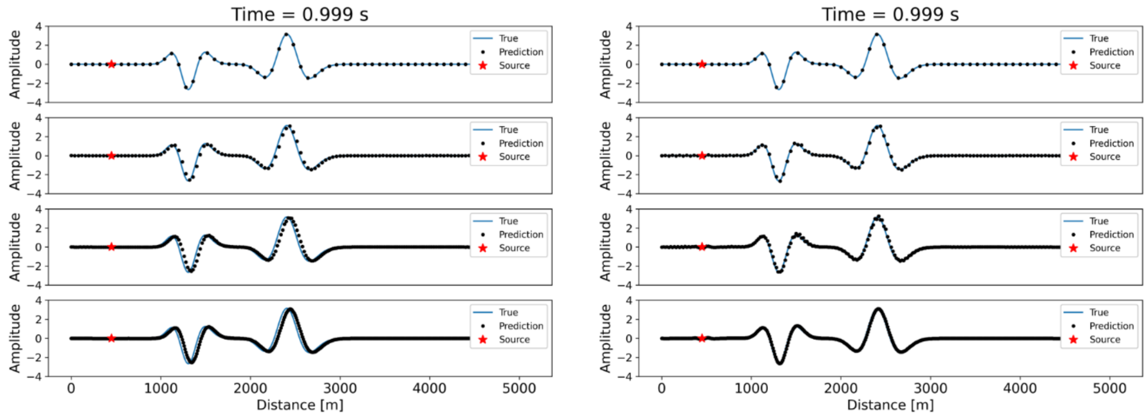

Figure 2: Prediction results after (left) and before (right) ensure source location information is not missed.

We also performed a numerical experiment on the two-dimensional case to assess mesh-invariant performance as we did on the one-dimensional case. We generated 500 smooth velocity models to propagate the acoustic waves; 400 of these models were used for training, and the remaining 100 were used to evaluate the performance of the learned operators against unknown models. The wavefields and velocity model used for the training had an original spatial resolution of 128-by-128 grid points; this original dataset was decimated or down-sampled to achieve training resolutions of 64-by-64 grid points. We considered a fixed source location in the center of the models, and we implemented absorbing boundaries by using a perfectly matched layer (PML) with exponential decay in the boundaries to avoid computational boundary reflections.

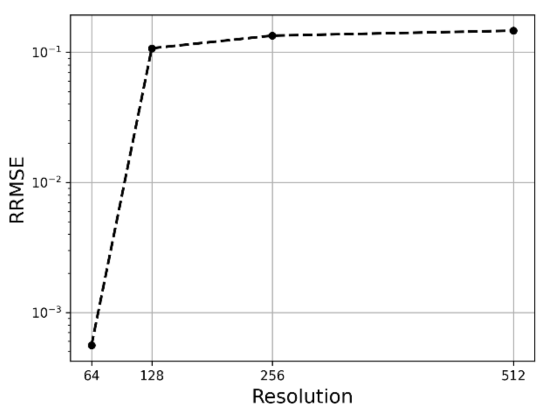

Figure 3: Average relative RMSE per resolution for one-dimensional case.

Discussion

The numerical experiments reveal key insights into the performance of the FNOs in both one-dimensional and two-dimensional cases. In the one-dimensional case, the FNOs performed well with low relative RMSE (around 10^{-3}), even for unknown velocity models. However, when the spatial resolution was increased from 64 to 128 grid points, the relative error increased significantly; this comes from the fact that we might be losing relevant information when decimating the spatial resolution on the wavefields (e.g. losing the exact source location) showing that the method is sensitive to resolution changes if the information is being lost during the training process, since once we ensure that the source location is being preserved, the errors as we increase the resolutions remain constant, and visual discrepancies, as the time lag, were corrected.

In the two-dimensional case, the performance of FNOs was not as good due to insufficient training data (only 400 models), which limited its ability to generalize well. However, the method still showed promising results in predicting wave amplitudes and waveforms for early time steps. Increasing the spatial resolution by a factor of two led to only a 2% increase in error, which gives us more sense of the consistency of the error as we increase the resolution (see Figure 4 for results).

Conclusions

The numerical experiments showed that Fourier Neural Operators (FNOs) perform well predicting the wave equation, especially in one dimension, where they achieved relatively low errors. However, when the spatial resolution was increased, the errors grew significantly, remarkably, when moving from 64 to 128 grid points. This increase in error in the one-dimensional case is because important information, such as the position of the seismic source, can be lost by decimating the original data. However, it was shown that as long as all the seismic information remains implicit in the signal, the increment in the resolution error will remain consistent. In the two-dimensional case, the loss of information did not occur as drastically since the data were decimated only once, which made the increase in error with respect to resolution less affected. In two dimensions, FNOs face higher memory demands and slightly reduced performance due to limitations in the available data. While the method shows potential, especially when reducing the data size without losing accuracy for more cost-effective training, its dependence on the quality and quantity of data could limit its generalization ability. Further studies are needed to assess the method's performance with realistic models, spectral performance and its potential applicability to accelerate full-waveform inversions (FWI).

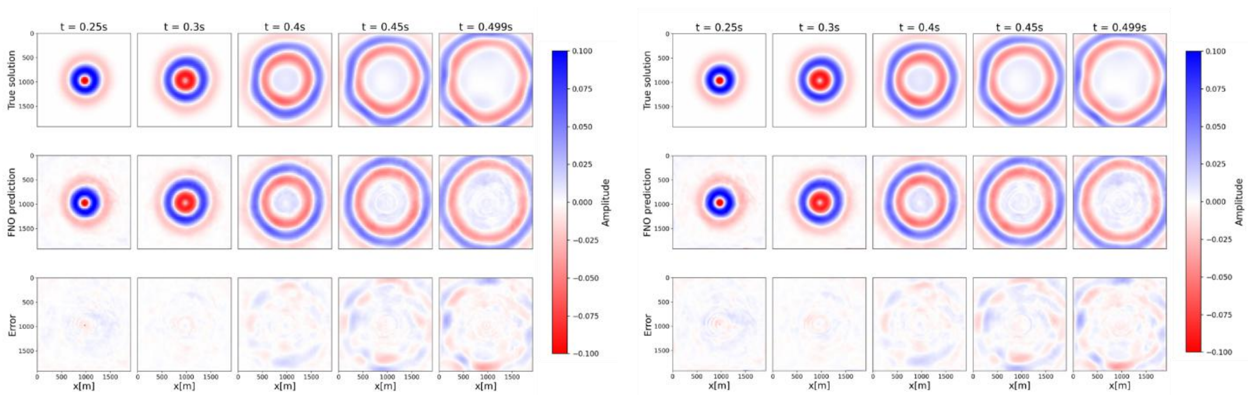

Figure 4: Prediction results for 64-by-64 grid points resolution (left) and 128-by-128 grid points resolution (right).

Acknowledgements

This work was funded by CREWES industrial sponsors and NSERC (Natural Science and Engineering Research Council of Canada) through the grant CRDPJ 543578-19. Emissions Reduction Alberta provided further support through the ACT4-SPARSE project.

The first author of this report was supported by scholarships from the CSEG Foundation and the University of Calgary, which are gratefully acknowledged.

References

Brandstetter, J., van den Berg, R., Welling, M., and Gupta, J. K., 2022, Clifford neural layers for pde modeling: arXiv preprint arXiv:2209.04934.

Levander, A. R., 1990, Finite-difference forward modeling in seismology, Springer US, 410–431.

Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A., 2021a, Fourier neural operator for parametric partial differential equations.

Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., Azizzadenesheli, K., and Anandkumar, A., 2021b, Physics-informed neural operator for learning partial differential equations: arXiv preprint arXiv:2108.10430.

Moczo, P., 2014, The Finite-Difference Modelling of Earthquake Motions: Cambridge University Press.

Zhang, T., Trad, D., and Innanen, K. A., 2023, Learning to solve the elastic wave equation with Fourier Neural Operators: Geophysics, 88, No. 3, T101–T119