Advancing seismic monitoring with operator learning

Hongyu Sun, Department of Earth, Environmental and Resource Sciences, The University of Texas at El Paso*

SUMMARY

Seismic monitoring is essential for understanding earthquake processes and mitigating associated hazards. Traditional deep learning methods in seismology often rely on single-station inputs or fixed seismic acquisition geometries, limiting their robustness and scalability in dynamic seismic monitoring systems. Recent developments in operator learning have enabled new approaches to seismic monitoring, which address these limitations by leveraging the full spatiotemporal information contained in multi-station waveform data and generalizing across diverse seismic network configurations. This work reviews three recent contributions that illustrate these approaches: (1) the Phase Neural Operator (PhaseNO), an algorithm for multi-station picking of seismic arrivals; (2) the Location Neural Operator (LocNO), a full-waveform earthquake location algorithm; and (3) an accelerated correlation-based time-reversal imaging method that uses the Fourier Neural Operator for efficient wavefield backpropagation. Together, these methods demonstrate that operator learning can enhance the accuracy, adaptability, and computational efficiency of seismic monitoring across varying geologic settings and network configurations.

INTRODUCTION

Accurate and efficient seismic monitoring is essential for understanding earthquake processes and mitigating seismic hazards, including those arising from natural subsurface dynamics and activities related to resource exploration and development. Existing deep learning–based seismic monitoring methods, such as PhaseNet (Zhu and Beroza, 2018) and EQTransformer (Mousavi et al., 2020), are primarily designed for single-station input and may suffer from reduced performance when dealing with low signal-to-noise ratio (SNR) data or signals from closely occurred seismic events. In contrast, multi-station processing is standard practice in seismology, as integrating spatial information from multiple stations enhances detection capabilities, improves phase-picking accuracy, and enables more robust earthquake location. However, early multi-station approaches based on convolutional neural networks (Yuan et al., 2018; Hu et al., 2019; Tsai et al., 2019; Wu et al., 2019; Yuan et al., 2020; Zhang et al., 2020; Zhao et al., 2025) are typically constrained to fixed and regularly spaced receiver geometries. These models lack the flexibility to generalize across diverse seismic network configurations, limiting their applicability to the specific acquisition geometry used during training. In practice, seismic networks often evolve over time due to instrument deployment, failure, or upgrades, necessitating repeated retraining of models and limiting the scalability of deep learning methods for long-term seismic monitoring.

To address these challenges, we introduce a new class of operator learning methods that are emerging as promising tools for seismic monitoring (Sun et al., 2022, 2023; Sun, 2025a). Classical neural networks are designed to approximate functions between finite-dimensional Euclidean spaces or over finite sets, which limits their flexibility when modeling functional mappings. In contrast, neural operators generalize neural networks by learning mappings between infinite-dimensional function spaces (Kovachki et al., 2023). Originally developed for solving partial differential equations, neural operators are constructed as compositions of linear integral operators and nonlinear activation functions. Like neural networks, they satisfy a universal approximation theorem, ensuring they can approximate any nonlinear continuous operator. A key advantage of neural operators is their discretization invariance: the same model parameters can be applied across different discretizations of the underlying function spaces, enabling flexible and scalable modeling across diverse data resolutions.

The seismic wavefield can be treated as a continuous function over time and space. Seismic recordings from multiple stations with arbitrary geometries represent a discretization of this wavefield, with regular sampling in time due to a fixed sampling rate and irregular sampling in space due to varying station numbers and geometries. By leveraging neural operators to construct functional mappings from the input seismic wavefield to other outputs, such as pseudo-probability functions indicating phase arrival times, we can develop flexible operator learning-based architectures for seismic monitoring. These models naturally generalize across arbitrary network configurations and changing monitoring conditions.

In this work, I briefly review three recent contributions: (1) PhaseNO (Sun et al., 2023), an open-source tool and pre-trained model for earthquake detection and multi-station P- and S-wave arrival time picking from continuous seismic data; (2) LocNO (Sun, 2025a), a framework for full-waveform earthquake location that accommodates flexible station numbers and geometries; and (3) a fast surrogate model based on the Fourier Neural Operator (FNO; Li et al., 2020a; Sun et al., 2022) that generalizes to various velocity models and source time functions (STFs) for simulating seismic wavefields. Together, these models demonstrate the potential of operator learning to advance seismic monitoring by improving generalization, accuracy, and efficiency across diverse seismic networks.

PHASE NEURAL OPERATOR

PhaseNO is an operator learning model designed for network-wide seismic phase picking (Sun et al., 2023). It learns infinite-dimensional function representations of seismic wavefields across multiple stations in a seismic network, allowing for accurate measurement of the arrival times of different phases simultaneously at multiple stations with arbitrary geometries. The input to PhaseNO is 30 s of waveforms recorded at any number of stations arranged in arbitrary geometries. The normalized

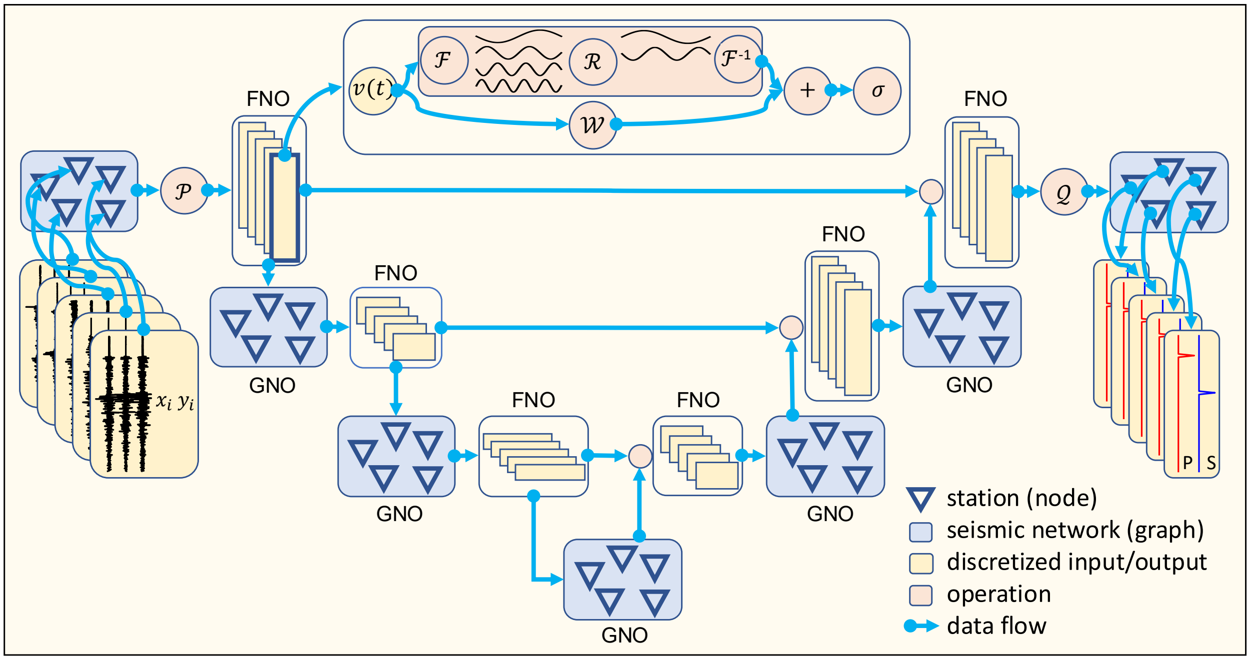

Figure 1: The Phase Neural Operator architecture. P and Q are up- and down-projections parameterized by neural networks, respectively. PhaseNO takes 30 s of three-component waveforms recorded by an arbitrary number of stations with flexible geometry and simultaneously predicts pseudo-probability functions for P- and S-wave arrivals at all input stations. Station coordinates (x_i, y_i) are encoded as two additional input channels alongside the waveforms. The figure is adapted from Sun et al. (2023).

geographical coordinates of the stations in the easting (x_i) and northing (y_i) directions are repeated along the time axis and concatenated with the three channels of three-component seismograms (or with repeated single-component seismograms, if applicable). The output function is a pseudo-probability mask with the same dimensions as the input, indicating the likelihood of P- and S-wave arrivals at each index of the input function.

Neural operators can be categorized based on how the kernel integral is computed within a linear integral operator layer. These include graph neural operators (GNOs; Li et al., 2020b), multipole graph neural operators, low-rank neural operators, and FNOs (Kovachki et al., 2023). Among these, FNOs compute the kernel in the Fourier domain, enabling efficient evaluation of the kernel integral operator using the Fast Fourier Transform (FFT). The use of FFT requires both input and output functions to be discretized on regular grids. In contrast, GNOs compute the kernel using the standard framework of graph neural networks (Gilmer et al., 2017). In this approach, discretization points are treated as nodes, and the graph is constructed based on the distances between nodes.

Although the full seismic wavefield is 3D in space and 1D in time, only a sparse set of seismic stations provides observed 1D time series data. This leads to limited spatial sampling. As a result, applying a fully 4D kernel integral over space and time is both computationally inefficient and unnecessary. To address this, we decouple temporal and spatial feature extraction. Temporal features are modeled using 1D FNOs (applying FFT only along the time axis), while spatial dependencies across stations are captured using GNOs. Figure 1 illustrates the model architecture. FNO and GNO layers are connected sequentially and repeated multiple times, enabling effective communication and exchange of spatiotemporal information among all stations in the seismic network. Skip connections (He et al., 2016) are introduced between blocks, forming a U-shaped architecture. These skip connections concatenate the FNO outputs from the encoder side directly with the GNO outputs on the decoder side, bypassing intermediate layers. This improves model convergence and facilitates the training of deeper neural networks.

We train PhaseNO on an earthquake dataset from the Northern California Earthquake Data Center (NCEDC) spanning 1984 to 2019 and benchmark its performance against three baseline models: EQTransformer (Mousavi et al., 2020), PhaseNet (Zhu and Beroza, 2018), and EdgePhase (Feng et al., 2022). PhaseNO achieves superior performance in terms of F1 scores, which reflect the trade-off between correct and false picks, as well as in terms of time residuals between predicted and manually picked arrivals. Although trained on data from Northern California, the model has been successfully applied to a seismic network in Southern California with different geological conditions and network configurations, demonstrating strong generalization ability beyond the training region. The trained PhaseNO model has been publicly released and can be directly applied to continuous waveform data without the need for retraining (Sun, 2025b).

To enhance the computational efficiency of PhaseNO since its publication and to better reflect the physical nature of seis-

10.1190/image2025-w11-02.1

mic wave propagation, the version v1.0.1 of PhaseNO (Sun, 2025b) introduces a distance-based constraint parameter, D, which limits information exchange among stations to only those within a specified physical distance. This modification resolves a key limitation in PhaseNO v1.0.0, where the computational cost, both in terms of GPU memory usage and runtime, scaled quadratically with the number of stations due to full connectivity in the graph. In the updated implementation, quadratic scaling only occurs if D is set larger than the maximum pairwise station distance, which restores a fully connected graph. However, such global connectivity is rarely necessary in practice, as earthquake signals recorded at one station are most relevant to nearby stations within a limited spatial range. By default, D is set to 30 km, meaning each station communicates only with others within a 30 km radius. Users are encouraged to adjust this parameter based on their seismic network configuration, as reducing D can significantly reduce computational demands. With an appropriately chosen value, users can now process all stations in a single run, without the need to randomly subsample large networks.

LOCATION NEURAL OPERATOR

LocNO is a full-waveform earthquake location algorithm that directly maps input seismograms from a detected earthquake or microseismic event to a pseudo-probability function indicating the likelihood of the source location in the easting, northing, and depth directions (Sun, 2025a). Using the same input configuration as PhaseNO, LocNO shares its spatiotemporal feature extraction architecture, which combines 1D FNO layers for capturing temporal features and GNO layers for modeling spatial dependencies in seismic wavefields. Specifically, seven FNO layers extract temporal information at individual stations, while five GNO layers capture spatial relationships across multiple stations in the seismic network. Compared with the PhaseNO architecture shown in Figure 1, LocNO places the seventh FNO layer before the down-projection operation Q.

After the seventh FNO layer, the spatiotemporal feature extractor produces a latent embedding for each station, summarizing its response to the input earthquake across time and space. A shared decoder transforms these embeddings into directional score vectors over discretized easting, northing, and depth grids. The outputs from individual stations are then combined using an attention-based pooling mechanism. The attention weights, calculated from the latent features, determine the contribution of each station to the final location estimate. The resulting directional score vectors are passed through a sigmoid function to produce pseudo-probabilities over the discretized spatial grids.

The LocNO model is trained using a binary cross-entropy loss between the predicted and ground-truth distributions in each spatial direction. Ground-truth labels are represented by truncated Gaussians centered at the true location coordinates, allowing the model to learn smooth and continuous probability estimates over discretized spatial grids. When training with catalog data, catalog locations may be treated as reference ground truth, while acknowledging that they include errors introduced by conventional earthquake location algorithms.

LocNO was benchmarked against SeismicGNN (Van Den Ende and Ampuero, 2020) using an earthquake catalog from the NCEDC. Both models were trained and tested on the same dataset. SeismicGNN is a graph neural network that incorporates spatial information for seismic source characterization by enabling information exchange between stations. It processes features at each station using a CNN, followed by a multilayer perceptron to estimate earthquake locations. LocNO consistently achieved lower mean absolute errors in both the easting and northing directions and demonstrated comparable performance in depth estimation, which is typically more challenging due to surface-only observations (Sun, 2025a).

LocNO is well suited for earthquake location under non-ideal seismic network conditions, where traditional travel-time location methods may fail. To demonstrate this capability, we train LocNO on a dataset that includes out-of-network events recorded by only a small number of stations. These events, whose origins lie outside the boundaries of the seismic network, are synthetically created by removing stations from one side of each cataloged earthquake. In this way, LocNO is trained using catalog earthquake locations determined from dense and well-monitored networks, while the input waveforms are restricted to stations forming an azimuthal gap greater than 180°. Numerical experiments in Sun (2025a) show that LocNO effectively learns to mitigate artifacts and outperforms Hypoinverse (Klein, 2002), a travel-time-based method, when locating out-of-network earthquakes in the Mendocino Triple Junction, even under sparse station coverage.

ACCELERATED TIME REVERSAL IMAGING

Time-reversal imaging with cross-correlation imaging conditions (Nakata and Beroza, 2016) offers high-resolution earthquake locations by using multi-station waveform data. However, this method is computationally expensive because it treats each receiver's waveform as a separate source for backpropagation rather than processing all time-reversed waveforms simultaneously. To address this challenge, we present a deep learning framework that incorporates neural operators to accelerate wave propagation within the time-reversal imaging workflow. We use a U-shaped neural operator (UNO) architecture (Rahman et al., 2022) based on FNO layers to simulate wave propagation at individual stations. Once trained, the model can rapidly generate wavefields, which are then cross-correlated to produce source images for earthquake location. Because neural operators operate in function space and do not depend on spatial discretization, they can be trained on coarse grids and applied to fine grids during inference, enabling efficient and high-resolution earthquake location (Sun et al., 2022).

We develop two training strategies for time-reversal modeling (Figure 2). The first strategy trains UNO-A to approximate Green's functions on arbitrary velocity models using a fixed Ricker wavelet as the STF (Sun et al., 2022). Owing to the linearity of the wave equation, the resulting Green's func-

10.1190/image2025-w11-02.1

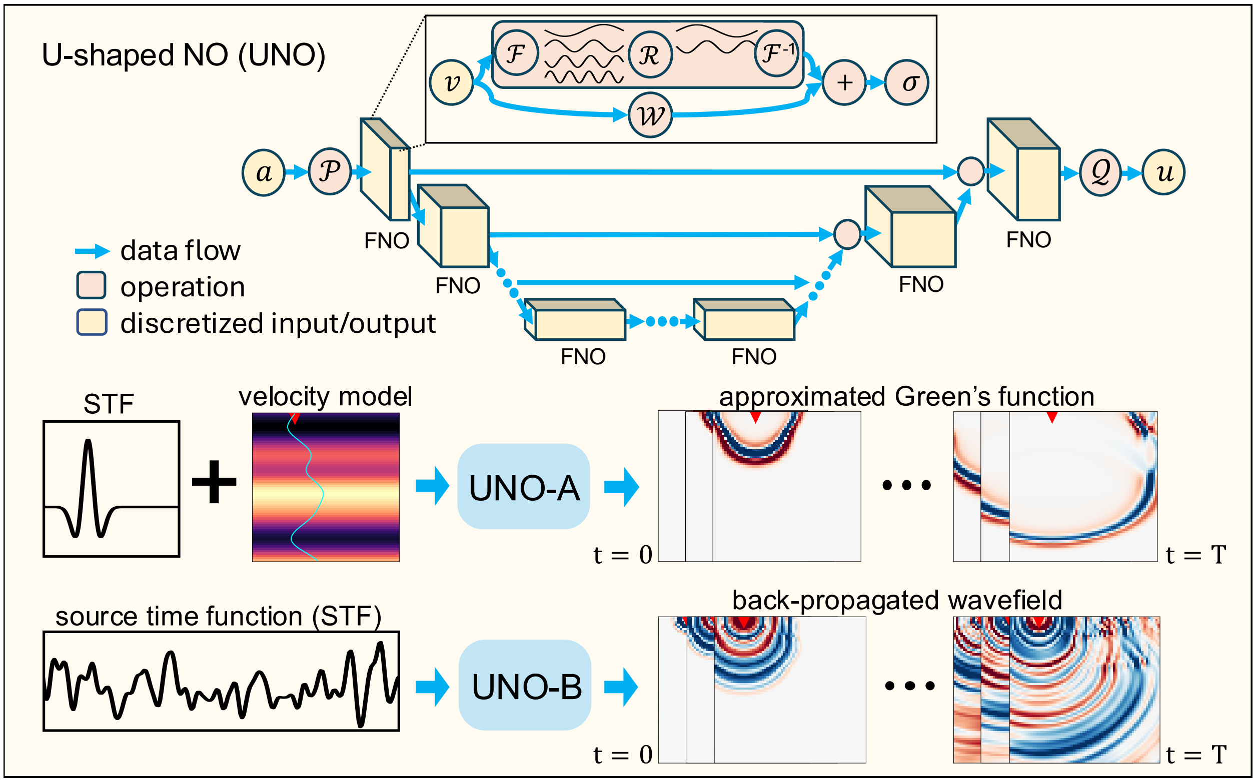

Figure 2: The UNO architecture for wave propagation. Red triangles denote the station locations (also the source locations in time-reversal modeling). The cyan line shows the velocity profile. The figure is adapted from Sun et al. (2022).

tions can be convolved with time-reversed seismic recordings to compute back-propagated wavefields. UNO-A is expected to generalize well across diverse velocity models, eliminating the need for retraining when applied to regions with different velocity structures. In the second strategy, we train a model called UNO-B to directly map time-reversed seismograms to back-propagated wavefields. Here, the velocity model is not explicitly input into the neural operator but is used during the simulation of the training dataset. After training, UNO-B is expected to generalize across different STFs and accurately reconstruct back-propagated wavefields on the target velocity model using time-reversed seismograms as inputs. The second strategy may be preferred when the seismic network includes tens of stations, as direct computation of backpropagation becomes more convenient in such a case.

In both cases, we sample random functions from Gaussian processes to construct training datasets for UNO-A and UNO-B. These functions represent either spatial velocity models (1D, 2D, or 3D) for training UNO-A or 1D STFs for training UNO-B. To generate the target wavefields, we numerically solve the 2D acoustic wave equation using spectral element modeling (SEM). Numerical experiments demonstrate that, after successful training, both UNO-A and UNO-B generalize well to previously unseen velocity models and STFs. The runtime is approximately two orders of magnitude faster than SEM and scales nearly linearly with the number of grid points (Sun et al., 2022).

CONCLUSIONS

This review highlights the potential of operator learning in advancing seismic monitoring. By modeling functional relationships between multi-station waveform data and target physical quantities, neural operators offer a flexible and scalable alternative to conventional CNN-based seismic monitoring methods. The three approaches discussed here demonstrate strong generalization to diverse seismic networks, maintain high accuracy in low-SNR conditions, and significantly reduce computational costs. These advantages support robust and efficient seismic monitoring, even under sparse or evolving network configurations. As seismic networks grow in scale and complexity, operator learning provides a promising framework to meet the increasing demands of real-time, high-resolution monitoring.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy, Office of Science, Energy Earthshots Initiative under Award Number DE-SC-0024703. Hongyu Sun also acknowledges support from the UT System STARs Program and the UTEP University Research Institute Grant.

REFERENCES

Feng, T., S. Mohanna, and L. Meng, 2022, EdgePhase: a deep learning model for multi-station seismic phase picking: Geochemistry, Geophysics, Geosystems, **23**, e2022GC010453 doi: https://doi.org/10.1029/2022GC010453.

Gilmer, J., S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, 2017, Neural message passing for quantum chemistry.

He, K., X. Zhang, S. Ren, and J. Sun, 2016, Deep residual learning for image recognition: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778.

Hu, L., X. Zheng, Y. Duan, X. Yan, Y. Hu, and X. Zhang, 2019, First-arrival picking with a u-net convolutional network: Geophysics, **84**, U45–U57, doi: https://doi.org/10.1190/geo2018-0688.1.

Klein, F. W., 2002, User's guide to HYPOINVERSE-2000, a Fortran program to solve for earthquake locations and magnitudes: Technical Report US Geological Survey.

Kovachki, N., Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, and A. Anandkumar, 2023, Neural operator: Learning maps between function spaces with applications to PDEs: Journal of Machine Learning Research, **24**, 1–97.

Li, Z., N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, 2020a, Fourier neural operator for parametric partial differential equations: arXiv preprint arXiv:2010.08895.

Li, Z., N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, 2020b, Neural operator: Graph kernel network for partial differential equations: arXiv preprint arXiv:2003.03485.

Mousavi, S. M., W. L. Ellsworth, W. Zhu, L. Y. Chuang, and G. C. Beroza, 2020, Earthquake transformer - an attentive deep-learning model for simultaneous earthquake detection and phase picking: Nature Communications, **11**, 1–12, doi: https://doi.org/10.1038/s41467-020-17591-w.

Nakata, N., and G. C. Beroza, 2016, Reverse time migration for microseismic sources using the geometric mean as an imaging condition: Geophysics, **81**, KS51–KS60, doi: https://doi.org/10.1190/geo2015-0278.1.

Rahman, M. A., Z. E. Ross, and K. Azizzadenesheli, 2022, U-NO: U-shaped neural operators: arXiv preprint arXiv:2204.11127.

Sun, H., 2025a, Operator learning for earthquake location: Methodology and application to out-of-network events in the Mendocino triple junction: Bulletin of the Seismological Society of America (Accepted for publication).

Sun, H., 2025b, PhaseNO: Version 1.0.1 [Software]: Zenodo, doi: https://doi.org/10.5281/zenodo.15161625.

Sun, H., Z. E. Ross, W. Zhu, and K. Azizzadenesheli, 2023, Phase neural operator for multi-station picking of seismic arrivals: Geophysical Research Letters, **50**, e2023GL106434, doi: https://doi.org/10.1029/2023GL106434.

Sun, H., Y. Yang, K. Azizzadenesheli, R. W. Clayton, and Z. E. Ross, 2022, Accelerating time-reversal imaging with neural operators for real-time earthquake locations: arXiv preprint arXiv:2210.06636.

Tsai, K. C., W. Hu, X. Wu, J. Chen, and Z. Han, 2019, Automatic first arrival picking via deep learning with human interactive learning: IEEE Transactions on Geoscience and Remote Sensing, **58**, 1380–1391, doi: https://doi.org/10.1109/TGRS.2019.2946118.

Van Den Ende, M. P. A., and J. Ampuero, 2020, Automated seismic source characterization using deep graph neural networks: Geophysical Research Letters, **47**, e2020GL088690, doi: https://doi.org/10.1029/2020GL088690.

Wu, H., B. Zhang, F. Li, and N. Liu, 2019, Semiautomatic first-arrival picking of microseismic events by using the pixel-wise convolutional image segmentation method: Geophysics, **84**, V143–V155, doi: https://doi.org/10.1190/geo2018-0389.1.

Yuan, P., S. Wang, W. Hu, X. Wu, J. Chen, and H. Van Nguyen, 2020, A robust first-arrival picking workflow using convolutional and recurrent neural networks: Geophysics, **85**, U109–U119, doi: https://doi.org/10.1190/geo2019-0437.1.

Yuan, S., J. Liu, S. Wang, T. Wang, and P. Shi, 2018, Seismic waveform classification and first-break picking using convolution neural networks: IEEE Geoscience and Remote Sensing Letters, **15**, 272–276, doi: https://doi.org/10.1109/LGRS.2017.2785834.

Zhang, X., J. Zhang, C. Yuan, S. Liu, Z. Chen, and W. Li, 2020, Locating induced earthquakes with a network of seismic stations in Oklahoma via a deep learning method: Scientific Reports, **10**, 1941, doi: https://doi.org/10.1038/s41598-020-58908-5.

Zhao, T., P. BILSBY, S. Manikani, and A. Abubakar, 2025, Event picking in seismic data: US Patent App. 18/704558.

Zhu, W., and G. C. Beroza, 2018, PhaseNet: a deep-neural-network-based seismic arrival time picking method: Geophysical Journal International, **216**, 261, doi: https://doi.org/10.1093/gji/ggy423.

10.1190/image2025-w11-02.1