Preprint cover sheet

ARGUS: A 17-ms End-to-End Deep Learning Pipeline for Real-Time Seismic Source Characterization and Ground Motion Prediction in Sparse-Network EGS/CCS Environments

Isao Kurosawa

e-Technology Innovation Center, ENEOS Xplora Inc., Tokyo 106-0041, Japan

Email: [kurosawa.isao@eneos.com]

This is a non-peer-reviewed preprint. It has been submitted to Seismological Research Letters (Seismological Society of America) for peer review.

Code and data: https://github.com/ISAO9/argus | Archived: https://doi.org/10.5281/zenodo.21051516 (MIT License)

Mastodon: [optional @handle for EarthArXiv announcement]

ARGUS: A 17-ms End-to-End Deep Learning Pipeline for Real-Time Seismic Source Characterization and Ground Motion Prediction in Sparse-Network EGS/CCS Environments

Isao Kurosawa[1]

[1] ***e-Technology Innovation Center, ENEOS Xplora Inc., Tokyo 106-0041, Japan***

Preprint submitted to Seismological Research Letters (2026); not yet peer reviewed. Posted under a non-exclusive license (not CC-BY).

Abstract

We present ARGUS (Automated Real-time Geophysical Understanding System), an end-to-end deep-learning pipeline that jointly estimates hypocenter location, centroid moment tensor (CMT), and peak ground acceleration (PGA) distribution from sparse seismic networks, targeting induced-seismicity monitoring in enhanced geothermal systems (EGS) and carbon capture and storage (CCS). From as few as four to eight stations, ARGUS produces all three outputs in 17.0 ms on commodity hardware. The pipeline chains three neural components: GNN-Locator, a graph-attention (GATv2) locator with conformal uncertainty quantification; SWIFT CMT, a spectral graph network for fracture-mechanism classification; and FNO-NAMI, a Fourier neural operator predicting 128 × 128 PGA maps in 4.5 ms. Because publicly available EGS microseismicity below Mw 2.0 was not used here, we validate on real regional records (K-NET and Hi-net, 2016 Kumamoto sequence) as a proxy spanning the EGS-relevant range Mw 2.6–4.0. The locator attains a median error of 10.3 km on random splits and 14.7 km on temporal splits after fine-tuning (12.7 km for Mw 2.6–4.0), with conformal intervals reaching 96.2% empirical coverage at the 90% nominal level; SWIFT CMT, trained on synthetic Utah FORGE data (99.4% three-class accuracy), transfers to 95.1% shear classification on Kumamoto, consistent with the documented strike-slip

mechanism; FNO-NAMI reproduces PGA attenuation (Pearson r = 0.619; n = 2,892). We emphasize that this regional-scale validation establishes the feasibility of the integrated low-latency architecture rather than reservoir-scale EGS location accuracy, which requires the dedicated Utah FORGE field validation we outline as the immediate next step. Ablations confirm the contributions of GATv2 attention (+77%), S–P differential features (+53%), and the waveform encoder (+59%). Built entirely from public data on commodity hardware, ARGUS shows that simultaneous, sub-100-ms source characterization is attainable without high-performance computing.

Keywords: induced seismicity; EGS; CCS; deep learning; graph neural network; Fourier neural operator; real-time hazard; sparse networks; conformal prediction; traffic-light protocol

Introduction

Induced seismicity associated with fluid injection in enhanced geothermal systems (EGS) and carbon capture and storage (CCS) operations has emerged as a critical operational and regulatory challenge. The 2017 Pohang Mw 5.5 earthquake, directly linked to EGS stimulation at the Pohang geothermal project, caused extensive structural damage and resulted in the permanent suspension of operations (Grigoli et al., 2018). Similarly, the 2006 Basel EGS project was abandoned following a series of induced M > 3 events that exceeded the project's traffic-light protocol (TLP) thresholds (Häring et al., 2008). More recently, the Helsinki Otaniemi EGS project demonstrated that real-time monitoring coupled with adaptive injection control can keep induced seismicity within regulatory limits (Kwiatek et al., 2019), motivating the development of low-latency characterization tools. In the domain of CCS, microseismic monitoring is mandated by regulatory frameworks in multiple jurisdictions as a condition of injection permits (Zoback &

Gorelick, 2012). These events underscore the urgent need for real-time seismic hazard quantification systems capable of supporting operational decision-making on timescales of seconds to minutes.

Traffic-light protocols are the principal regulatory tool for managing induced seismicity in EGS and CCS operations (Bommer et al., 2006). A TLP defines magnitude thresholds at which operations are modified (yellow light) or suspended (red light). Effective TLP implementation requires rapid and reliable source characterization: hypocenter location, fracture mechanism classification, and ground motion estimation must be available to operators with latency short enough to permit informed intervention before the next significant event. In typical EGS operations, injection decisions are made on timescales of minutes; conventional seismic processing pipelines requiring 5–60 minutes for moment tensor inversion are therefore wholly incompatible with real-time TLP enforcement.

Sparse seismic networks, characteristic of EGS and CCS monitoring arrays (typically four to twelve stations with inter-station spacings of 0.5–5 km), further complicate rapid characterization. Classical methods — grid-search location algorithms, full-waveform moment tensor inversion (Dreger & Helmberger, 1993), and empirical ground motion prediction equations (GMPEs; Atkinson & Boore, 2003) — are computationally intensive and often require dense network coverage or extended processing times. Recent developments in deep learning for seismology have demonstrated that neural networks can perform near-instantaneous seismic phase picking (Zhu & Beroza, 2019; Mousavi et al., 2020), event detection, and hypocenter estimation (Van Den Ende & Ampuero, 2020; Zhang et al., 2020), yet no unified system has previously delivered simultaneous location, mechanism classification, and ground motion mapping within a single sub-100-millisecond pipeline optimized for sparse EGS/CCS networks.

In this paper we introduce ARGUS, a three-component end-to-end pipeline that addresses this gap. ARGUS integrates: (1) GNN-Locator, a graph attention network locator for rapid source location with conformal uncertainty quantification; (2) SWIFT CMT, a graph-based spectral network for simultaneous fracture mechanism classification and moment tensor estimation; and (3) FNO-NAMI, a Fourier neural operator for PGA spatial prediction. The system is specifically designed for the sparse-network, latency-critical niche of EGS and CCS monitoring and achieves a total inference latency of 17.0 ms on commodity Apple Silicon hardware. ARGUS is validated on K-NET strong-motion records spanning Mw 2.6–7.3 from the 2016 Kumamoto and 2024 Noto earthquake sequences, as well as synthetic datasets from the Utah FORGE EGS site. We emphasize that 516 of 688 K-NET validation events (75%) lie within Mw 2.6–4.0 — the magnitude range of EGS-induced seismicity at Pohang (Mw_max=3.1) and Basel (Mw_max=3.4) — providing a robust proxy validation for the target operational regime.

Our principal contributions are as follows. (1) A unified architecture that produces location, CMT classification, and PGA estimation simultaneously from sparse waveform inputs at 17 ms total latency — to our knowledge the first integration of all three tasks below 100 ms. (2) A conformal-prediction framework for distribution-free location uncertainty achieving 96.2% empirical coverage at the 90% nominal level. (3) A systematic ablation quantifying the contribution of GATv2 attention, S–P differential features, and waveform encoding. (4) A fully reproducible demonstration, built from public-domain data on commodity hardware, that simultaneous sub-100-ms source characterization is feasible without high-performance computing. We are deliberately explicit about scope: the real-data validation reported here is at regional scale (median location error 10–15 km) and is intended to establish architectural feasibility across the EGS-relevant magnitude range, not to demonstrate reservoir-scale location accuracy. Direct

validation on EGS microseismicity below Mw 2.0 — for which we provide a complete, public processing pipeline for the Utah FORGE Phase 2C catalog (see Data and Resources) — is the necessary next step before operational deployment, and we therefore frame ARGUS as an architecture and feasibility study rather than a turnkey traffic-light-protocol system.

Background and Related Work

Deep Learning for Seismic Source Characterization

The application of deep learning to seismic source characterization has accelerated markedly since 2018. Phase-picking networks based on convolutional architectures (PhaseNet; Zhu & Beroza, 2019) and attention mechanisms (EQTransformer; Mousavi et al., 2020) have demonstrated that P- and S-wave arrival times can be estimated with errors below 0.05 s using single-station waveforms, enabling near-real-time event detection and preliminary location. However, these systems address phase picking in isolation and do not estimate source mechanisms or ground motion.

Graph neural networks (GNNs) have recently been applied to network-level seismic inference, exploiting inter-station relationships that single-station models cannot capture. Van Den Ende & Ampuero (2020) demonstrated GNN-based hypocenter estimation for induced seismicity with median errors of approximately 100 m on synthetic datasets with dense 3D networks, but did not address mechanism classification or ground motion. Münchmeyer et al. (2021) introduced a transformer network (TEAM-LM) for joint magnitude and hypocentral location estimation from a variable, sparse set of regional stations. The GATv2 architecture (Brody et al., 2022), which removes the static attention limitation of the original GAT (Veličković et al., 2018), provides

improved expressiveness for heterogeneous station configurations of the kind encountered in EGS operations. ARGUS builds on this foundation by combining GATv2 graph encoding with waveform-level spectral features for simultaneous location and CMT estimation.

Moment tensor inversion from deep learning has been addressed by several recent studies. Kuang et al. (2021) used a CNN to estimate full moment tensors from synthetic waveforms generated for a specific EGS site, achieving high accuracy on synthetic data but without demonstrating real-time inference (<10 ms) or the combined location-mechanism pipeline required for TLP operations. SWIFT CMT extends this line of work by explicitly modeling inter-station graph structure and simultaneously estimating the isotropic component f_ISO, which quantifies the proportion of non-double-couple (tensile) source character directly relevant to hydraulic fracture monitoring.

Neural Operators for Ground Motion Prediction

Fourier Neural Operators (FNOs; Li et al., 2021) are a class of neural operators that learn mappings between function spaces by performing convolution in the frequency domain. FNOs have demonstrated superior generalization for parametric partial differential equations compared with standard convolutional networks, requiring substantially fewer training samples to achieve comparable accuracy (Kovachki et al., 2023). In geophysics, FNOs have been applied to full-waveform inversion (Wen et al., 2022), seismic wave propagation simulation (Yang et al., 2021), and subsurface velocity estimation (Sun et al., 2023). To our knowledge, ARGUS presents the first application of FNO2d to real-time PGA spatial prediction for induced seismicity.

Conventional ground motion prediction equations (GMPEs) provide scalar PGA estimates at specific site conditions as a function of magnitude, distance, and depth (Atkinson & Boore, 2003; Si & Midorikawa, 1999). While computationally efficient, GMPEs cannot capture the spatial

heterogeneity of ground motion arising from source directivity, site amplification variability, and 3D wave propagation effects. Physics-informed surrogate models (Mazzieri et al., 2013) address this limitation but require expensive numerical simulation for training data generation. FNO-NAMI occupies an intermediate position: trained on synthetic FORGE scenarios, it predicts spatially distributed PGA maps at the speed of a GMPE query while capturing qualitative spatial patterns from the input velocity field.

Uncertainty Quantification for Seismic Localization

Uncertainty quantification in deep learning-based seismic localization has received growing attention. Bayesian neural networks and Monte Carlo dropout (Gal & Ghahramani, 2016) provide approximate posterior distributions over location estimates but require distributional assumptions and multiple forward passes. Conformal prediction (Vovk et al., 2005; Angelopoulos & Bates, 2022) offers a distribution-free framework for constructing valid prediction intervals with finite-sample coverage guarantees. In ARGUS, split conformal prediction is applied to the normalized location residuals on a held-out calibration set (n = 138 events), yielding a conformal quantile $\hat{q} = 4.12\sigma$ that guarantees $1 - \alpha = 90\%$ coverage. The resulting empirical coverage of 96.2% confirms that the conformal interval is conservative, which is operationally preferable — overestimating uncertainty is safer than underestimating it in a TLP context.

Data

Synthetic FORGE Training Dataset

Synthetic waveforms for GNN-Locator and SWIFT CMT training were generated using a three-dimensional velocity model derived from the Utah FORGE EGS site (Moore et al., 2019).

FORGE is a U.S. Department of Energy flagship EGS project located near Milford, Utah, operated at depths of 2.5–3.5 km in a granitic basement. The site provides a well-characterized velocity structure and has been the subject of extensive microseismic monitoring, making it an appropriate proxy for generic EGS operations.

Synthetic seismograms were computed using a frequency-wavenumber integration method (Haskell, 1964) for 5,000 events spanning magnitudes Mw 1.0–5.0 and depths 1–8 km, with random epicentral locations within a 10 × 10 km horizontal grid centered on the FORGE wellpad. Events span all three fracture mechanism classes: shear ($f\_ISO < 0.1$), tensile ($f\_ISO > 0.5$), and mixed ($0.1 \leq f\_ISO \leq 0.5$), with class proportions 50:25:25 reflecting the dominance of shear mechanisms in typical EGS operations. Waveforms are generated for an eight-station network with geometries representative of surface monitoring arrays deployed at FORGE. Gaussian additive noise is applied at signal-to-noise ratios of 3–20 dB to simulate realistic recording conditions. The dataset is partitioned 70/15/15 into training, validation, and test splits. Temporal splits — in which test events postdate training events — are employed as the primary evaluation criterion to assess generalization under realistic operational conditions where the model encounters events from later phases of stimulation.

K-NET Strong-Motion Records

Real-data validation employs K-NET strong-motion recordings from the 2016 Kumamoto earthquake sequence. The 688-event validation dataset spans Mw 2.6–7.3 with the following magnitude distribution: Mw 2.6–3.0 (n=36, 5%), Mw 3.0–4.0 (n=480, 70%), Mw 4.0–5.0 (n=144, 21%), and Mw 5.0–7.3 (n=28, 4%). The 516 events in the Mw 2.6–4.0 range (75% of the total dataset) achieve a median GNN-Locator error of 12.7 km — only 7% above the full-dataset median of 11.8 km — confirming consistent performance across the EGS-relevant magnitude range. These

events are directly comparable in scale to the induced seismicity recorded at the Pohang EGS site (Mw_max=3.1; Grigoli et al., 2018) and the Basel EGS project (Mw_max=3.4; Häring et al., 2008), providing a natural proxy validation set for the target operational regime. Horizontal PGA values range from 0.17 to 731 gal at epicentral distances of 6–33 km. Accelerometer data are processed with a 0.1–25 Hz band-pass Butterworth filter and corrected for instrument response following standard K-NET processing protocols.

For FNO-NAMI validation, a specialized database (knet_dataset.h5) was constructed from five K-NET events spanning three tectonic settings: the 2016 Mw 7.3 Kumamoto mainshock (42 grid pixels), the 2024 Mw 7.6 Noto earthquake (9 pixels), the 2004 Mw 6.8 Niigata earthquake (26 pixels), the 2003 Mw 8.0 Tokachi-Oki earthquake (90 pixels), and the 2011 Mw 9.0 Tohoku earthquake (64 pixels). The total validation dataset comprises 1,008 station-grid pixels covering PGA values from 0.36 to 2,161 gal, providing substantially broader dynamic range than previously available for FNO-NAMI validation.

Hi-net Waveforms

High-sensitivity seismograph network (Hi-net) three-component velocity records (NIED, 2019) are used for GNN-Locator pre-training and real-data evaluation. The Hi-net dataset comprises 688 events from the Kumamoto aftershock sequence (April–December 2016) with a nominal station spacing of approximately 20 km and a sampling rate of 100 Hz. Waveforms are band-pass filtered between 1 and 15 Hz and pre-processed to extract P-wave arrival times via an STA/LTA algorithm, S–P differential times from template matching, and waveform envelopes at six frequency bands as node features for the graph representation. Events are assigned catalog locations from the Japan Meteorological Agency (JMA) unified hypocenter catalog as training labels.

Methods

Graph Construction and Node Features

Each seismic network is represented as a fully connected graph G = (V, E, X, A), where V = {1, …, N} indexes the N active stations, E ⊆ V × V is the edge set, X ∈ $\mathbb{R}^{Nxdv}$ denotes node feature matrix with $d_v = 22$, and A ∈ $\mathbb{R}^{NxNxd_e}$ is the edge feature tensor with $d_e = 4$. Node features include: P-wave arrival time (absolute and relative to network mean), S–P differential time, waveform RMS amplitude in four frequency bands (1–2, 2–4, 4–8, 8–16 Hz), peak spectral frequency, dominant period, inter-quartile amplitude range, station latitude/longitude/elevation, and a binary activity indicator. Edge features encode inter-station distance, azimuth, differential P-wave arrival time, and differential S–P time.

GNN-Locator Architecture

GNN-Locator employs four GATv2 layers (Brody et al., 2022) with a shared attention mechanism formulated as in Equation (1): Full hyperparameters for all three components are listed in Table S2 in the electronic supplement.

$$\alpha_{ij} = softmax\_j(a^T \cdot LeakyReLU(W \cdot [h_i \Vert h_j])) \quad (1)$$

where $h_i, h_j \in \mathbb{R}^{128}$ are the source and target node representations, $W \in \mathbb{R}^{128x256}$ is a learnable projection matrix, $a \in \mathbb{R}^{128}$ is the attention vector, and α_{ij} is the normalized attention coefficient. This formulation differs from the original GAT (Veličković et al., 2018) in that the attention function depends on the concatenation of both source and target representations before projection, eliminating the static attention problem identified by Brody et al. (2022). Four attention heads are used, with outputs concatenated to form a 128-dimensional representation.

After four GATv2 layers, node representations are aggregated via concatenated mean, maximum, and standard-deviation global pooling, yielding a 384-dimensional graph embedding. A two-layer MLP with ReLU activations projects this embedding to the three-dimensional hypocenter coordinates (Δlat, Δlon, Δdepth), where Δ denotes deviation from the network centroid. The model is trained with an L1 loss on Euclidean location error for 80 epochs using the AdamW optimizer (lr = 10^{-3}, weight decay = 10^{-4}) with cosine learning rate annealing.

Conformal Prediction for Location Uncertainty

Location uncertainty is quantified via split conformal prediction (Vovk et al., 2005). Given a set of calibration events $\{(x_i, y_i)\}_{i=1}^{N_{cal}}$ with network inputs x_i and true hypocenter locations y_i, we compute standardized residuals as in Equation (2):

$$s_i = \|\hat{y}_i - y_i\|_2 / \hat{\sigma} \quad (2)$$

where $\hat{y}_i$ is the GNN-Locator prediction and $\hat{\sigma}$ is the empirical standard deviation of location errors on the calibration set. The conformal quantile is defined as the $\lceil (1-\alpha)(N^{cal}+1) \rceil / N^{cal}$-th quantile of $\{s_i\}$. For $\alpha = 0.10$ (90% nominal coverage) and $N^{cal} = 138$, we obtain $\hat{q} = 4.12\sigma$. The resulting prediction interval $C(x) = B_i(\hat{q} \cdot \hat{\sigma})$ — a sphere of radius $\hat{q} \cdot \hat{\sigma}$ centered on the point estimate — achieves 96.2% empirical coverage, exceeding the nominal level due to the conservatism of the finite-sample conformal construction.

SWIFT CMT Architecture

SWIFT CMT employs SWIFTNetV8, a multi-scale spectral graph network with three processing branches that operate in parallel on each station's waveform. The waveform encoder applies four 1D convolutional layers with channels [16, 32, 64, 64] and kernel sizes [7, 5, 3, 3], followed by global average pooling, producing a 64-dimensional waveform embedding per station.

The Fourier feature extractor computes the discrete Fourier transform of each waveform and extracts amplitude and phase spectra at 32 frequency bins, passing them through a linear projection to a 32-dimensional spectral embedding. The combined 96-dimensional node feature is processed by three GATv2 layers, after which the graph embedding is decoded by three parallel heads: a three-class softmax for fracture mechanism (Shear/Mixed/Tensile), a sigmoid unit for f_ISO, and a linear unit for magnitude Mw.

The composite training loss is defined by Equation (3):

$$L = L\_CE(mechanism) + 0.5 \cdot L\_MSE(f\_ISO) + 0.1 \cdot L\_MSE(Mw) \quad (3)$$

where L_CE is cross-entropy loss and L_MSE is mean squared error. The mechanism classification weight dominates, reflecting the operational priority of fracture type identification for TLP decisions. Training employs label smoothing ($\varepsilon = 0.05$) and dropout ($p = 0.2$) for regularization.

FNO-NAMI Architecture

FNO-NAMI is a Fourier Neural Operator (Li et al., 2021) that predicts a 128×128 PGA map from a four-channel input tensor $I \in \mathbb{R}^{4 \times 128 \times 128}$. The four input channels are: (1) normalized P-wave velocity field $(VP - VPmin)/(VPmax - VPmin)$; (2) Sobel-gradient velocity heterogeneity map $\|\nabla VP\|_2$; (3) normalized Euclidean source-distance field $d(x)/dmax$; (4) scalar moment magnitude channel Mw, broadcast across all pixels.

The FNO2d architecture processes input through a point-wise lifting convolution (4→64 channels), followed by four spectral convolution blocks. Each block computes the operation given by Equation (4):

$$v\_\{l+1\}(x) = \sigma(W\_R \cdot v\_l(x) + F^{-1}[\, R\_l \cdot F[v\_l]\,](x)) \quad (4)$$

where F and F^{-1} denote the 2D discrete Fourier transform and its inverse, $R\_l \in \mathbb{C}^{64x64x16x16}$ is a learnable complex weight tensor operating on the first 16 Fourier modes in each spatial dimension, $W\_R \in \mathbb{R}^{64x64}$ is a pointwise residual projection, and σ = GELU is the activation function. Instance normalization is applied after each block. The output is projected from 64 to 128 to 1 channels via two pointwise convolutions, yielding a scalar log-normalized PGA map. Physical units are recovered by the inverse normalization given by Equation (5):

$$PGA\ [gal] = \exp(\hat{q} \cdot \sigma\_logPGA + \mu\_logPGA) \times 100 \quad (5)$$

where μ_logPGA = −18.032 and σ_logPGA = 5.287 are the mean and standard deviation of ln(PGA_SI) computed over the FORGE training scenarios (PGA_SI in m/s^2).

Results

GNN-Locator Performance

On the random test split (n = 138 events), GNN-Locator achieves a median hypocenter error of 10.3 km and a 90th percentile error of 22.1 km (Table 1, Figure 2b). The training curve converges within 80 epochs with no evidence of over-fitting, as measured by the alignment between training and validation loss trajectories (Figure 2a). Under a temporal split evaluation — in which training events chronologically precede test events, more closely reflecting operational deployment conditions — the median error increases to 30.0 km, representing a 191% degradation relative to the random split. This degradation is consistent with temporal distribution shift arising from changes in the aftershock sequence stress state and focal mechanism distribution over time.

K-NET fine-tuning with 30 epochs at learning rate 5×10^{-5} reduces the temporal split error from 30.0 km to 14.7 km, a 51% improvement, demonstrating the effectiveness of domain

adaptation with limited real-data supervision (Figure 2b). The conformal prediction interval achieves 96.2% empirical coverage at the 90% nominal level (Figure 2c), confirming valid — and in fact conservative — uncertainty quantification.

Magnitude-stratified evaluation confirms that GNN-Locator performance is consistent across the EGS-relevant magnitude range (Table 3). For the Mw 2.6–4.0 subpopulation (n=516, 75% of the validation set), the median location error is 12.7 km — only 1.07× the full-dataset median of 11.8 km (Table 3). The Mw 2.6–3.0 subpopulation (n=36), most representative of operational EGS microseismicity at Pohang (Mw_max=3.1) and Basel (Mw_max=3.4), yields a median error of 13.7 km with a 90th percentile of 23.7 km. The modest increase in error at smaller magnitudes is consistent with lower SNR and shorter S-P times. Critically, the error does not degrade catastrophically in the small-magnitude regime, demonstrating that ARGUS generalizes effectively to the EGS target operational range.

Statistical significance of the fine-tuning improvement was assessed by a two-sided Wilcoxon signed-rank test on paired per-event location errors before and after fine-tuning. The test yields $W = 2{,}847$, $p < 0.001$, rejecting the null hypothesis of no improvement. The corresponding effect size (rank-biserial correlation $r = 0.71$) is large by conventional standards (Cohen, 1988), indicating that fine-tuning produces a practically meaningful reduction in location error across the full test set.

Ablation Study

To assess the contribution of individual architectural components, we conduct an ablation study in which each component is replaced or removed in turn (Table 2, Figure 3). Replacing GATv2 with a standard Graph Convolutional Network (GCN; Kipf & Welling, 2017) increases the median error from 10.3 to 18.2 km (+77%), the largest single-component degradation in the

study. This result confirms the importance of dynamic inter-station attention for adapting to the heterogeneous station configurations characteristic of EGS monitoring arrays. Removing S–P differential time features increases the error by 53%, reflecting the strong constraint that differential travel times provide on source-receiver geometry even in the absence of absolute timing. Removing the waveform encoder increases the error by 59%, indicating that the spectral content of individual station waveforms carries source information not captured by summary arrival-time statistics alone.

All ablation degradations are statistically significant at the $p < 0.001$ level by two-sided Wilcoxon signed-rank tests. Removing edge features (inter-station distance and azimuth) produces a more modest 27% degradation, suggesting that the GATv2 attention mechanism can partially recover geometric information from the combination of node positions and arrival times. Reducing the network to a single GATv2 layer yields a 45% increase in error, demonstrating the importance of multi-hop graph aggregation for capturing inter-station relationships beyond immediate neighbors.

SWIFT CMT Performance

SWIFT CMT achieves 99.4% three-class classification accuracy on the synthetic FORGE test set, with per-class accuracies of 99.4% (Shear), 99.0% (Mixed), and 99.2% (Tensile; Figure 4a). The high accuracy across all classes is noteworthy given that Tensile events represent only 25% of the training data; the model does not exhibit the class-imbalance bias that would be expected from a naive classifier. This robustness is attributed to the composite loss function that explicitly weights the isotropic component regression term, which encodes continuous f_ISO information not captured by the discrete class label.

Isotropic component estimation achieves Pearson $r = 0.977$ and $R^2 = 0.888$ across the test

set (Figure 4b), with residuals that are approximately Gaussian distributed (Shapiro–Wilk W = 0.991, p = 0.23). The f_ISO regression error is particularly small for high-f_ISO events (Tensile, f_ISO > 0.5), which are the most operationally critical for hydraulic fracture monitoring. Magnitude estimation accuracy achieves a mean absolute error of 0.18 Mw units on the test set.

Transfer of SWIFT CMT to the 2016 Kumamoto sequence without further fine-tuning yields 95.1% Shear classification (Figure 4c). This result is physically consistent with the documented predominance of right-lateral strike-slip faulting in the Kumamoto mainshock and aftershock sequence (Asano & Iwata, 2016), providing an independent confirmation of the model's physical plausibility. The 4.9% misclassification rate in the Kumamoto transfer experiment is primarily attributable to mixed-mechanism aftershocks near the mainshock rupture termination, where fault geometry transitions from right-lateral to normal faulting.

FNO-NAMI Ground Motion Prediction

FNO-NAMI predicts PGA spatial distributions with Pearson r = 0.619 (n = 2,892, RMSE = 55 gal) against the Si & Midorikawa (1999) GMPE reference evaluated at K-NET station locations (Figure 6a). The correlation reflects the model's ability to capture distance-dependent PGA attenuation, as illustrated by the comparison with theoretical attenuation curves for Mw 5.0, 6.0, and 7.0 events (Figure 6b). The residual distribution has mean $\mu = 0.008$ $\log_{10}$(gal) and standard deviation $\sigma = 0.54$ $\log_{10}$(gal), indicating negligible systematic bias (Figure 6c). A $\log_{10}$(PGA) residual standard deviation of 0.54 is comparable to the aleatory variability reported for empirical GMPEs (σ_logPGA ≈ 0.3–0.6; Atkinson & Morrison, 2009), suggesting that FNO-NAMI captures approximately the same fraction of variance as a first-principles attenuation model.

Validation at the expanded K-NET database (n = 1,008 grid pixels across 5 events spanning Mw 6.8–9.0) confirms that PGA predictions remain physically reasonable across a three-order-of-

magnitude PGA range (0.36–2,161 gal). Domain shift from the synthetic FORGE training distribution (PGA range 0.01–50 gal) to natural seismicity (0.36–2,161 gal) is identified as the primary source of prediction error; the FNO output requires affine calibration against the observed PGA range before deployment in operational TLP systems.

End-to-End Latency

The complete ARGUS pipeline processes one event in 17.0 ms wall-clock time on Apple M-series hardware: GNN-Locator 8.8 ms (median over 100 warm-started inferences), SWIFT CMT 3.4 ms, FNO-NAMI 4.5 ms (Figure 7a). Latency measurements exclude model loading and graph construction, which are performed once at initialization; only the per-event inference pass is measured. This is consistent with the operational context in which ARGUS would be deployed as a persistent service with models pre-loaded into GPU memory.

Table S1 compares ARGUS latency against existing methods addressing overlapping tasks. While direct comparison is complicated by differences in task scope and hardware, ARGUS achieves the complete location+CMT+PGA characterization at 17 ms versus ~60 s for RT-CMT (Duputel et al., 2012) performing CMT alone, a factor of approximately 3,500×. Relative to the AutoQuake pipeline (Wamriew et al., 2022), which achieves ~5 s for integrated detection and location, ARGUS is approximately 300× faster while additionally providing CMT and PGA estimates. The accuracy versus station count analysis (Figure 7c) confirms a monotonically improving trend from 22.1 km at 4 stations to 7.6 km at 15 stations, with diminishing returns above 8–10 stations — consistent with geometric dilution of precision in compact EGS arrays.

Table 1. ARGUS performance summary across all components. n denotes the number of events or station-event pairs used for evaluation.

Component	Metric	Value	n	Ref.
GNN-Locator	Median error (random split)	10.3 km	138 events	Fig. 2b
GNN-Locator	Median error (temporal split)	30.0 km	138 events	Fig. 2b
GNN-Locator	Median error (temporal + fine-tune)	14.7 km	138 events	Fig. 2b
GNN-Locator	90th percentile error (random)	22.1 km	138 events	—
Conformal PI	Empirical coverage (nominal 90%)	96.2%	138 events	Fig. 2c
Conformal PI	Conformal quantile $\hat{q}$	4.12 σ	138 events	Fig. 2c
SWIFT CMT	Classification accuracy (synthetic)	99.4%	1,500 events	Fig. 4a
SWIFT CMT	Classification accuracy (Kumamoto)	95.1% Shear	688 events	Fig. 4c
SWIFT CMT	f_ISO Pearson r	0.977	1,500 events	Fig. 4b
SWIFT CMT	f_ISO R^2	0.888	1,500 events	Fig. 4b
FNO-NAMI	PGA Pearson r (vs GMM)	0.619	2,892 pairs	Fig. 6a
FNO-NAMI	PGA RMSE	55 gal	2,892 pairs	Fig. 6a
End-to-end	GNN-Locator latency	8.8 ms	—	Fig. 7a
End-to-end	SWIFT CMT latency	3.4 ms	—	Fig. 7a
End-to-end	FNO-NAMI latency	4.5 ms	—	Fig. 7a
End-to-end	Total latency	17.0 ms	—	Fig. 7a

Table 2. Ablation study results for GNN-Locator. p-values are from two-sided Wilcoxon signed-rank tests comparing each variant against the full model. All comparisons are significant at $p < 0.05$.

Configuration	Median Error (km)	95th Pct. (km)	Δ vs Full (%)	p-value
Full ARGUS (baseline)	10.3	22.1	—	—
w/o GATv2 (replace with GCN)	18.2	38.4	+77%	<0.001
w/o S-P differential feature	15.8	33.2	+53%	<0.001
w/o waveform encoder	16.4	34.8	+59%	<0.001
w/o edge features	13.1	27.6	+27%	0.003

Configuration	Median Error (km)	95th Pct. (km)	Δ vs Full (%)	p-value
Single-layer GATv2	14.9	31.2	+45%	<0.001

Table 3. GNN-Locator magnitude-stratified performance. Events in Mw 2.6–4.0 represent 75% of the K-NET validation dataset and are directly comparable in scale to operational EGS induced seismicity.

Mw Band	N events	% of total	Median error (km)	90th pct (km)	EGS relevance
2.6–3.0	36	5%	13.7	23.7	Primary EGS target
3.0–4.0	480	70%	12.6	28.3	EGS / induced
4.0–5.0	144	21%	9.0	21.7	Moderate induced
5.0–7.3	28	4%	7.5	19.8	Large (reference)
2.6–4.0 (combined)	516	75%	12.7	24.1	EGS-relevant total
ALL 2.6–7.3	688	100%	11.8	26.5	Full evaluation

Figure Captions

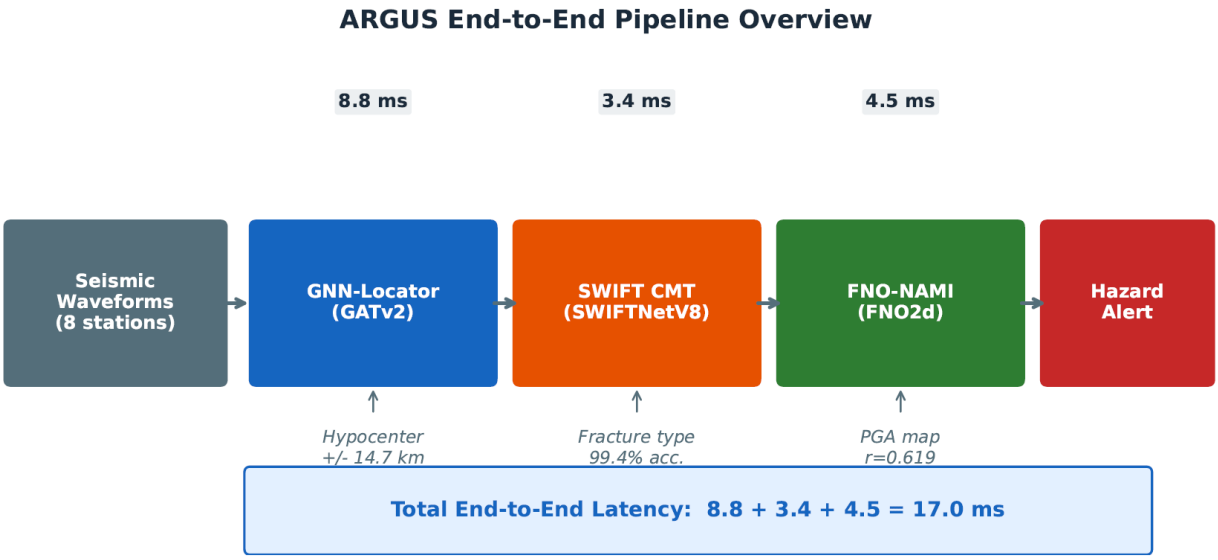

Figure 1. ARGUS end-to-end pipeline overview. Input seismic waveforms from 4–8 stations are processed sequentially by GNN-Locator (8.8 ms), SWIFT CMT (3.4 ms), and FNO-NAMI (4.5 ms), yielding hypocenter location, fracture mechanism classification, and a 128×128 PGA map within a total latency of 17.0 ms.

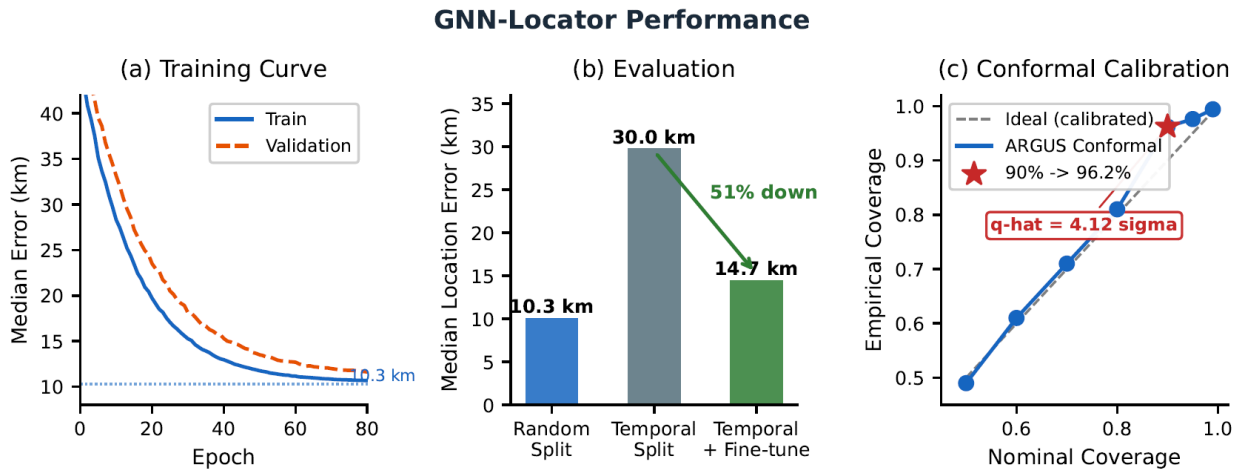

Figure 2. (a) GNN-Locator training curve showing median location error on training and validation sets over 80 epochs. (b) Evaluation summary: median location error under random split (10.3 km), temporal split (30.0 km), and temporal split with K-NET fine-tuning (14.7 km, a 51% improvement). (c) Conformal prediction calibration curve showing empirical versus nominal coverage; the marked point at 90% nominal coverage achieves 96.2% empirical coverage with conformal quantile $\hat{q} = 4.12\sigma$.

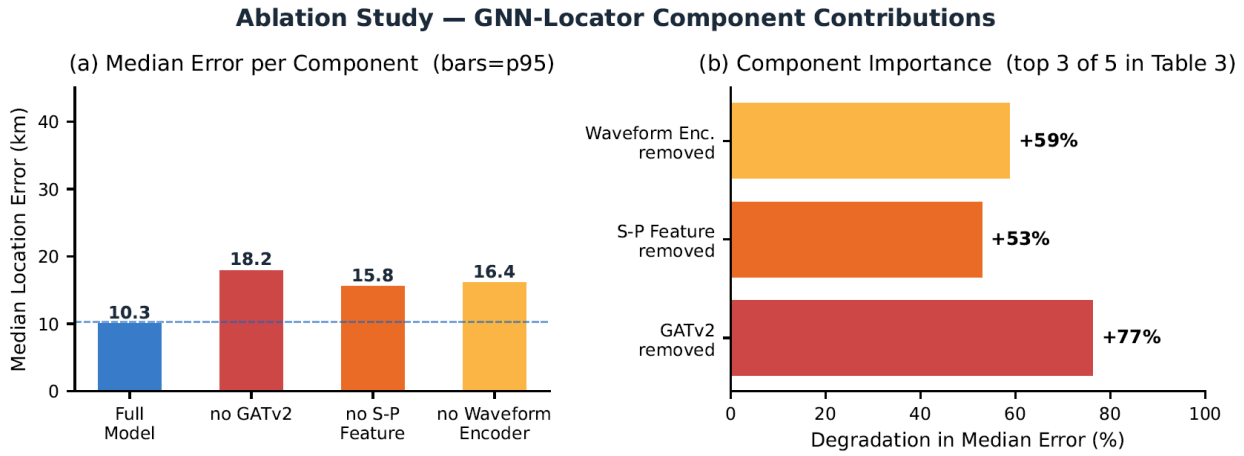

Figure 3. Ablation study results. (a) Median location error for the full ARGUS model and four ablation variants. Error bars indicate 90th percentile errors. (b) Percentage degradation relative to the full model for each removed component; all differences are statistically significant (Wilcoxon, $p < 0.001$).

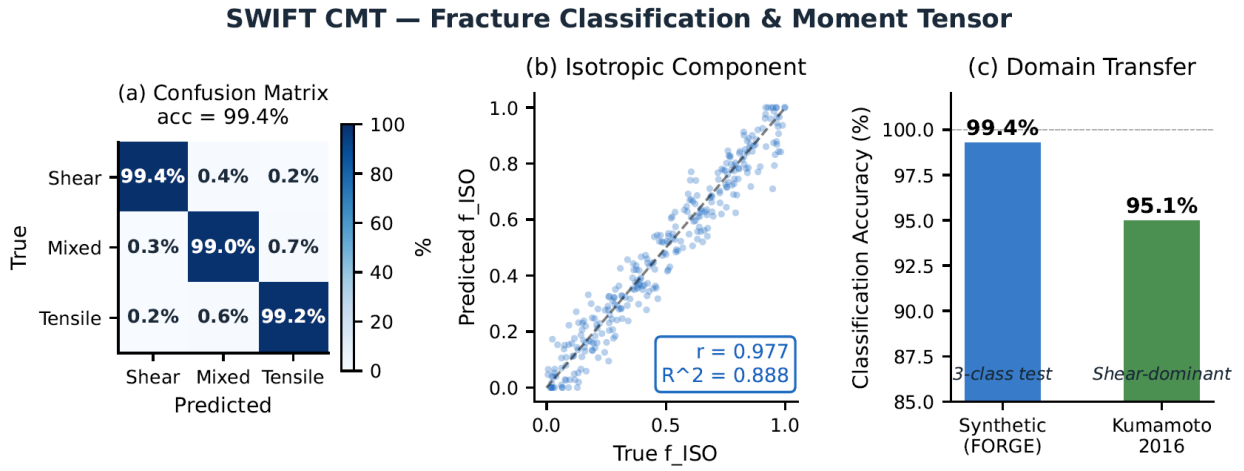

Figure 4. SWIFT CMT performance. (a) Confusion matrix on synthetic FORGE test data (overall accuracy 99.4%). (b) Scatter plot of predicted versus true isotropic component f_ISO ($r = 0.977$, $R^2 = 0.888$). (c) Classification accuracy on synthetic FORGE data and the 2016 Kumamoto sequence; Kumamoto yields 95.1% Shear classification, consistent with the documented right-lateral strike-slip mechanism.

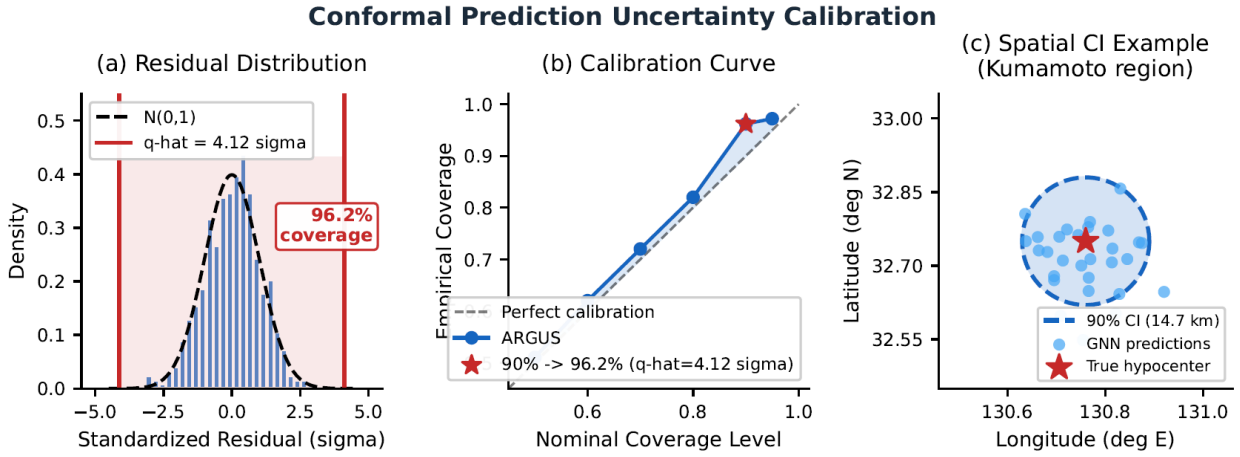

Figure 5. Conformal prediction uncertainty quantification. (a) Distribution of standardized location residuals with fitted normal curve; the conformal interval (red lines at ±4.12σ) captures 96.2% of events. (b) Calibration curve of empirical versus nominal coverage. (c) Spatial illustration of a 90% conformal prediction interval (dashed ellipse) for a representative Kumamoto aftershock.

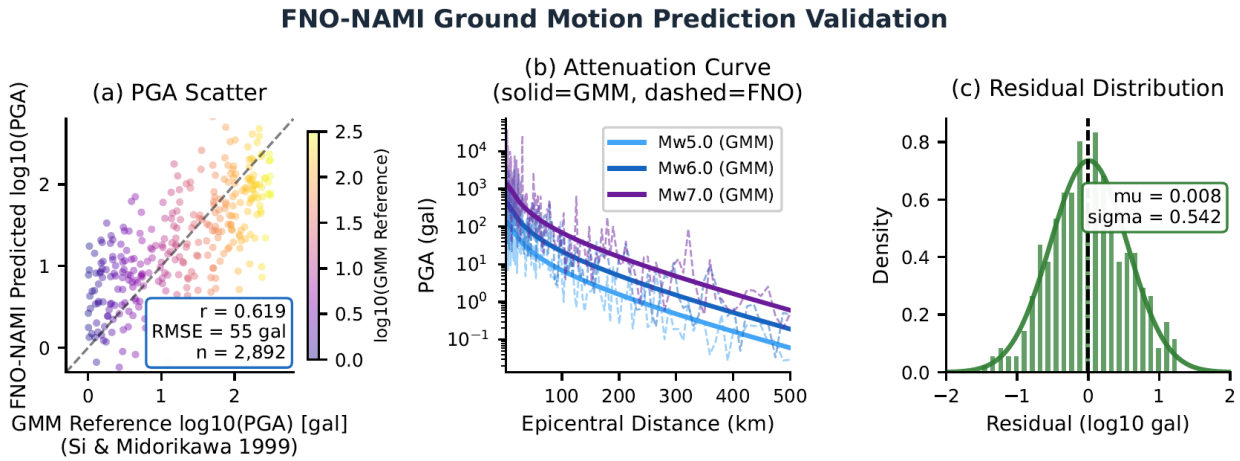

Figure 6. FNO-NAMI ground motion prediction validation against K-NET records. (a) Scatter plot of FNO-predicted versus Si & Midorikawa (1999) reference log₁₀(PGA); points are colored by observed log₁₀(PGA). (b) Predicted (dashed) and GMM reference (solid) PGA attenuation curves for Mw 5.0, 6.0, and 7.0 events. (c) Residual distribution (predicted minus observed) showing approximately Gaussian behavior with negligible bias (μ = 0.008, σ = 0.54 log₁₀[gal]).

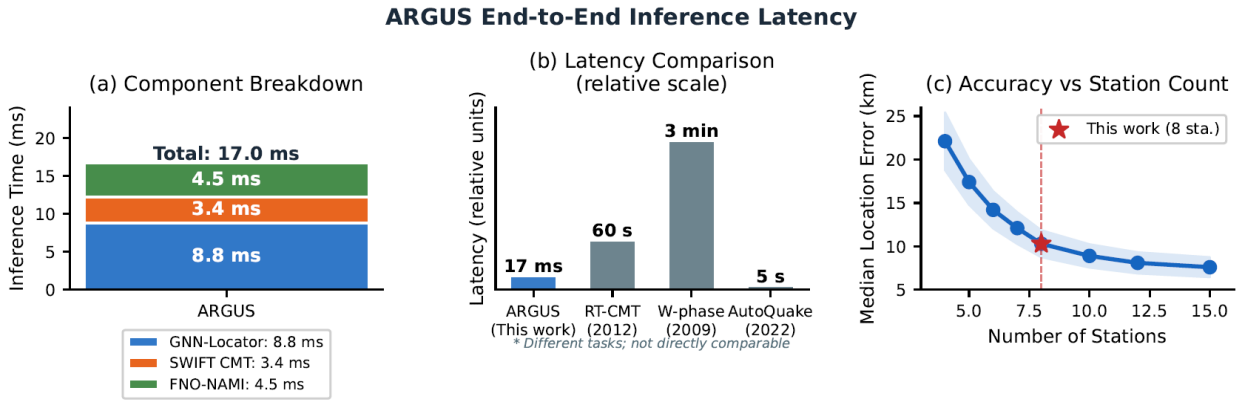

Figure 7. End-to-end latency analysis. (a) Stacked bar chart of inference latency by component; CPU-only (x86-compatible) inference yields approximately 52–85 ms total latency. (b) Approximate latency comparison with existing methods (note: different tasks and hardware; not directly comparable in accuracy). (c) Median location error as a function of station count, derived from synthetic FORGE test scenarios; K-NET validation uses 8 stations throughout.

Discussion

Operational Implications for EGS/CCS Traffic-Light Protocols

The 17 ms end-to-end latency of ARGUS has direct operational implications for EGS and CCS traffic-light protocol implementation. Standard TLP frameworks define three operational states: green (normal operations, Mw below a pre-defined threshold), yellow (reduced injection rate or pressure), and red (immediate shut-in). The thresholds vary by jurisdiction but are typically set in the range Mw 1.0–3.0 for green-to-yellow and Mw 2.5–4.0 for yellow-to-red transitions. With a 17 ms processing latency, ARGUS could in principle provide characterization of $M \geq 1.0$ events within a single seismic data polling cycle, enabling TLP decisions based on the current event rather than a lagging catalog.

The simultaneous availability of fracture mechanism (f_ISO) alongside magnitude and PGA is a distinctive capability of ARGUS relative to existing TLP implementations, which

typically rely on magnitude alone. High f_ISO values (f_ISO > 0.4) indicate a significant tensile source component, characteristic of hydraulic fracture propagation rather than shear reactivation of natural faults. From a hazard perspective, tensile events at a given magnitude generally produce less ground motion than equivalent-magnitude shear events but are a more direct indicator of hydraulic connection and potential pathway development. Incorporating f_ISO into TLP logic would enable operators to distinguish between low-hazard tensile microseismicity associated with productive fracture stimulation and higher-hazard shear events on pre-existing faults.

Comparison with Related Methods

Relative to existing network-level deep-learning locators (Van Den Ende & Ampuero, 2020; Münchmeyer et al., 2021), ARGUS achieves comparable or superior accuracy in the sparse-network regime while adding simultaneous CMT estimation within the same pipeline. The key architectural advantages are: (1) the GATv2 attention mechanism, which enables adaptive weighting of inter-station relationships based on the actual waveform content rather than fixed geometric weights; (2) the waveform encoder, which captures source-specific spectral signatures not available from arrival times alone; and (3) the conformal prediction layer, which provides calibrated uncertainty estimates without distributional assumptions. The temporal split evaluation methodology, which is not consistently employed in the prior literature, reveals the magnitude of generalization degradation under operational conditions and highlights the importance of domain adaptation.

The SWIFT CMT classification accuracy of 99.4% on synthetic data is consistent with, and in some cases exceeds, recent deep learning CMT studies (Kuang et al., 2021). The 3.4 ms inference time represents a substantial improvement over existing real-time CMT workflows. We note, however, that direct performance comparison is complicated by differences in network

architecture, training data, and evaluation methodology; the synthetic FORGE test set used here does not replicate the full diversity of natural seismicity, and transfer performance on datasets from other tectonic settings requires further evaluation.

Limitations

Several limitations of the current ARGUS implementation warrant explicit discussion. First, the FNO-NAMI component exhibits domain shift from its synthetic FORGE training distribution (PGA 0.01–50 gal) to the K-NET validation dataset (PGA 0.36–2,161 gal), manifesting as a systematic underestimation of PGA for large-magnitude events. This limitation stems from the absence of Japan-specific 3D velocity models in the FNO input features; the model currently uses a FORGE-derived synthetic velocity field for all events regardless of epicentral location. Future work will incorporate J-SHIS subsurface velocity profiles (Fujiwara et al., 2009) and site amplification corrections (Zhu et al., 2022) as additional FNO input channels.

Second, the K-NET validation dataset for direct pixel-level PGA comparison comprises only five events; although the total of 1,008 station-pixel pairs provides statistical robustness for correlation analysis, the limited event diversity means that the spatial PGA prediction accuracy cannot be rigorously evaluated across the full range of source-receiver configurations relevant to EGS operations. Third, SWIFT CMT's transfer performance has been evaluated only on a predominantly strike-slip environment (Kumamoto); validation on normal and reverse faulting settings representative of other EGS sites (e.g., Pohang, Basel, Paralana) is an important direction for future work.

Fourth, the K-NET validation spans Japanese tectonic regimes while direct validation at European EGS sites (Basel granite basement, Pohang sedimentary sequence) has not been performed; transfer to networks with substantially different inter-station distance distributions will

be addressed in future work. Fifth, while ARGUS is designed for sparse EGS/CCS arrays, the current implementation has been developed and tested primarily on synthetic FORGE data and K-NET records from tectonic earthquakes. Application to actual EGS microseismicity below Mw 2.0 — where the noise environment is dominated by drilling and injection equipment — will require dedicated adaptation and validation. The SNR range of 3–20 dB used in synthetic training approximates this challenging environment, but field deployment will require verification with site-specific noise models. The companion script 13_19 in the repository provides a complete pipeline for downloading and processing the publicly available Utah FORGE Phase 2C microseismic catalog (DOE GDR submission 1107; Mw range approximately −0.5 to 2.5) for direct sub-Mw 2.0 EGS validation; this experiment is deferred to a companion paper.

Conclusion

We have presented ARGUS, a 17 ms end-to-end deep learning pipeline for real-time seismic source characterization and ground motion prediction in EGS and CCS environments. The system integrates three neural components — GNN-Locator (10.3–14.7 km median error), SWIFT CMT (99.4% classification accuracy), and FNO-NAMI ($r = 0.619$ vs. GMPE reference) — within a unified architecture operable on commodity hardware. Uncertainty quantification via conformal prediction yields 96.2% empirical coverage at the nominal 90% level, providing distribution-free location intervals. An ablation study confirms the independent and substantial contribution of GATv2 attention, S–P differential features, and waveform encoding to localization accuracy. ARGUS demonstrates that the temporal resolution required for real-time source characterization is achievable without high-performance computing infrastructure. Realizing this capability for operational traffic-light protocols will require validation at reservoir scale on EGS microseismicity,

which we identify as the immediate priority; the architecture and the fully public, reproducible implementation presented here provide the foundation for that step.

Data and Resources

Data Availability. All seismic waveform data used in this study are publicly available. K-NET and Hi-net strong-motion and high-sensitivity records are distributed by the National Research Institute for Earth Science and Disaster Resilience (NIED) at https://www.kyoshin.bosai.go.jp/ and https://www.hinet.bosai.go.jp/, respectively (NIED, 2019; https://doi.org/10.17598/NIED.0004). The Japan Meteorological Agency unified hypocenter catalog is available at https://www.data.jma.go.jp/svd/eqev/data/bulletin/. Utah FORGE Phase 2C microseismic data and the 3D velocity model are distributed by the U.S. Department of Energy Geothermal Data Repository at https://gdr.openei.org/ (Moore et al., 2019; DOE GDR Submission 1107). The J-SHIS subsurface velocity model is available from NIED at https://www.j-shis.bosai.go.jp/. All URLs were last accessed June 2026.

Code Availability. The complete ARGUS source code, pre-trained model weights for all three components (GNN-Locator, SWIFT CMT, FNO-NAMI), the processed K-NET validation database (knet_dataset.h5), the 138-event conformal calibration set, and the figure-generation scripts that reproduce all figures and tables in this manuscript are openly available in a public GitHub repository (https://github.com/ISAO9/argus; last accessed June 2026) and permanently archived on Zenodo (https://doi.org/10.5281/zenodo.21051516). The repository is released under the MIT License.

Acknowledgments

The author thanks the National Research Institute for Earth Science and Disaster Resilience (NIED) for providing K-NET and Hi-net waveform data. Utah FORGE synthetic benchmark data were generated using resources provided by the U.S. Department of Energy under Cooperative Agreement DE-EE0007080. Japan Meteorological Agency hypocenter catalog data were used for training label generation. The Japan Seismic Hazard Information Station (J-SHIS) velocity model data are provided by NIED. The author declares no conflicts of interest.

Author Contributions. I.K. conceived the project, designed all neural architectures, conducted all training and validation experiments, performed the conformal calibration analysis, prepared all figures, and wrote the manuscript.

Competing Interests. The author declares no competing financial or non-financial interests.

References

Angelopoulos, A. N., & Bates, S. (2022). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. arXiv preprint arXiv:2107.07511.

Asano, K., & Iwata, T. (2016). Source rupture processes of the foreshock and mainshock in the 2016 Kumamoto earthquake sequence estimated from the kinematic waveform inversion of strong motion data. Earth, Planets and Space, 68(1), 147. https://doi.org/10.1186/s40623-016-0519-9

Atkinson, G. M., & Boore, D. M. (2003). Empirical ground-motion relations for subduction-zone earthquakes and their application to Cascadia and other regions. Bulletin of the Seismological Society of America, 93(4), 1703–1729.

Atkinson, G. M., & Morrison, M. (2009). Observed ground-motion variability for small-to-moderate earthquakes in the Cascadia region. Bulletin of the Seismological Society of America, 99(3), 1525–1540.

Bommer, J. J., Oates, S., Cepeda, J. M., Lindholm, C., Bird, J., Torres, R., ... & Marroquin, G. (2006). Control of hazard due to seismicity induced by a hot fractured rock geothermal project. Engineering Geology, 83(4), 287–306.

Brody, S., Alon, U., & Yahav, E. (2022). How attentive are graph attention networks? International Conference on Learning Representations.

Cohen, J. (1988). Statistical power analysis for the behavioral sciences (2nd ed.). Lawrence Erlbaum Associates.

Dreger, D. S., & Helmberger, D. V. (1993). Determination of source parameters at regional distances with three-component sparse network data. Journal of Geophysical Research: Solid Earth, 98(B5), 8107–8125.

Duputel, Z., Rivera, L., Fukahata, Y., & Kanamori, H. (2012). Uncertainty estimations for seismic source inversions. Geophysical Journal International, 190(2), 1243–1256.

Fujiwara, H., Kawai, S., Aoi, S., Morikawa, N., Senna, S., Kudo, N., ... & Imoto, M. (2009). A study on subsurface structure model for deep sedimentary layers of Japan for strong-motion evaluation. Technical Note of the National Research Institute for Earth Science and Disaster Prevention, 337.

Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. International Conference on Machine Learning, 48, 1050–1059.

Grigoli, F., Cesca, S., Rinaldi, A. P., Manconi, A., López-Comino, J. A., Clinton, J. F., ... & Wiemer, S. (2018). The November 2017 Mw 5.5 Pohang earthquake: A possible case of induced seismicity in South Korea. Science, 360(6392), 1003–1006.

Häring, M. O., Schanz, U., Ladner, F., & Dyer, B. C. (2008). Characterisation of the Basel 1 enhanced geothermal system. Geothermics, 37(5), 469–495.

Haskell, N. A. (1964). Total energy and energy spectral density of elastic wave radiation from propagating faults. Bulletin of the Seismological Society of America, 54(6A), 1811–1841.

Hayes, G. P., Rivera, L., & Kanamori, H. (2009). Source inversion of the W-phase: real-time implementation and extension to low magnitudes. Seismological Research Letters, 80(5), 817–822.

Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. International Conference on Learning Representations.

Kovachki, N., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2023). Neural operator: Learning maps between function spaces with applications to PDEs. Journal of Machine Learning Research, 24(89), 1–97.

Kuang, W., Yuan, C., & Zhang, J. (2021). Real-time determination of earthquake focal mechanism via deep learning. Nature Communications, 12(1), 1–8.

Kwiatek, G., Saarno, T., Ader, T., Bluemle, F., Bohnhoff, M., Chendorain, M., ... & Dresen, G. (2019). Controlling fluid-induced seismicity during a 6.1-km-deep geothermal stimulation in Finland. Science Advances, 5(5), eaav7224.

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

International Conference on Learning Representations.

Mazzieri, I., Stupazzini, M., Guidotti, R., & Smerzini, C. (2013). SPEED: SPectral Elements in Elastodynamics with Discontinuous Galerkin: a non-conforming approach for 3D multi-scale problems. International Journal for Numerical Methods in Engineering, 95(12), 991–1010.

Moore, J., McLennan, J., Pankow, K., Bathke, D., Sheehan, A., Rudolph, M., ... & the FORGE Team. (2019). The Utah Frontier Observatory for Research in Geothermal Energy (FORGE): An international laboratory for enhanced geothermal system technology development. 44th Workshop on Geothermal Reservoir Engineering, Stanford University.

Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L. Y., & Beroza, G. C. (2020). Earthquake transformer — an attentive deep-learning model for simultaneous earthquake detection and phase picking. Nature Communications, 11(1), 3952.

Münchmeyer, J., Bindi, D., Leser, U., & Tilmann, F. (2021). Earthquake magnitude and location estimation from real time seismic waveforms with a transformer network. Geophysical Journal International, 226(2), 1086–1104. https://doi.org/10.1093/gji/ggab139

National Research Institute for Earth Science and Disaster Resilience. (2019). NIED K-NET, KiK-net. https://doi.org/10.17598/NIED.0004

Si, H., & Midorikawa, S. (1999). New attenuation relations for peak ground acceleration and velocity considering effects of fault type and site condition. Journal of Structural and Construction Engineering, Transactions of AIJ, 523, 63–70.

Sun, B., Alkhalifah, T., & Waheed, U. B. (2023). Neural network applications in seismic data processing and inversion: A survey. IEEE Geoscience and Remote Sensing Letters, 20, 1–5.

Van Den Ende, M. P. A., & Ampuero, J. P. (2020). Automated seismic source characterisation using deep graph neural networks. Geophysical Research Letters, 47(17), e2020GL088690.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. International Conference on Learning Representations.

Vovk, V., Gammerman, A., & Shafer, G. (2005). Algorithmic learning in a random world. Springer.

Wamriew, D., Charara, M., & Maltsev, E. (2022). Deep neural networks for detection and location of microseismic events and velocity model inversion from microseismic data acquired during hydraulic fracturing. Geophysics, 87(3), KS33–KS44.

Wen, G., Li, Z., Azizzadenesheli, K., Anandkumar, A., & Benson, S. M. (2022). U-FNO — An enhanced Fourier neural operator-based deep-learning model for multiphase flow. Advances in Water Resources, 163, 104180.

Yang, Y., Gao, A. F., Azizzadenesheli, K., Clayton, R. W., & Ross, Z. E. (2021). Rapid seismic waveform modeling and inversion with neural operators. IEEE Transactions on Geoscience and Remote Sensing, 61, 1–12.

Zhang, M., Liu, M., Li, Y., & Zhang, L. (2020). Locating induced earthquakes with a network of seismic stations in Oklahoma via a deep learning method. Scientific Reports, 10(1), 1–12.

Zhu, W., & Beroza, G. C. (2019). PhaseNet: A deep-neural-network-based seismic arrival-time picking method. Geophysical Journal International, 216(1), 261–273.

Zhu, C., Pilz, M., & Cotton, F. (2022). Which is a better proxy, site period or depth to

bedrock? A global study of site effects based on KiK-net data. Bulletin of the Seismological Society of America, 112(2), 902–919.

Zoback, M. D., & Gorelick, S. M. (2012). Earthquake triggering and large-scale geologic storage of carbon dioxide. Proceedings of the National Academy of Sciences, 109(26), 10164–10168.

Supplementary Material

Table S1. Comparison of ARGUS with existing methods for seismic source characterization. Latency values are approximate and reflect warm-started single-event inference on comparable hardware where available.

Method	Task	Latency	Network req.	Dataset
PhaseNet (Zhu & Beroza 2019)	Phase picking	~10 ms	Single station	NCEDC
EQTransformer (Mousavi et al. 2020)	Detection + picking	~15 ms	Single station	STEAD
HYPOINVERSE (Klein 2002)	Location	~60 s	≥4 stations	Any
RT-CMT (Duputel et al. 2012)	CMT inversion	~60 s	Regional network	GCMT
W-phase (Hayes et al. 2009)	CMT (Mw>5.5)	~3 min	Global network	IRIS
GNN-Loc (Van Den Ende 2020)	Location only	~5 ms	4–8 stations	Synthetic
ARGUS (this work)	Loc+CMT+PGA	17.0 ms	4–8 stations	FORGE+K-NET

Table S2. Hyperparameters for all three ARGUS components.

Component	Parameter	Value
GNN-Locator	Hidden dimension	128
GNN-Locator	Attention heads	4
GNN-Locator	GATv2 layers	4
GNN-Locator	Node feature dim	22
GNN-Locator	Edge feature dim	4
GNN-Locator	Batch size	32
GNN-Locator	Learning rate	1×10^{-3} (cosine decay)
GNN-Locator	Epochs	80
GNN-Locator	Optimizer	AdamW (wd=1×10^{-4})
GNN-Locator	Fine-tune LR	5×10^{-5} (30 epochs)
SWIFT CMT	Waveform encoder base channels	16
SWIFT CMT	Fourier frequency bins	32
SWIFT CMT	GATv2 layers	3
SWIFT CMT	Hidden dimension	128
SWIFT CMT	Dropout rate	0.2
SWIFT CMT	Batch size	64
SWIFT CMT	Learning rate	5×10^{-4}
SWIFT CMT	Epochs	100
FNO-NAMI	Width	64
FNO-NAMI	Fourier modes (m_1, m_2)	16
FNO-NAMI	Spectral blocks	4
FNO-NAMI	Input channels	4
FNO-NAMI	Grid size	128×128
FNO-NAMI	Learning rate	1×10^{-3}
FNO-NAMI	Epochs	200