Semi-Supervised FFT-Spectral Representation Learning for Earthquake Detection in Distributed Fiber Optic Sensing Systems

1st Deepika Sasi
Department of ECE
NIT, Puducherry
Karaikal 609609, India
deepikajoy2013@gmail.com

2nd Sundaresan Sabapathy
Amrita School of AI
Amrita Vishwa Vidyapeetham
Coimbatore, Tamil Nadu 641112, India
s_sundaresan@cb.amrita.edu

3rd Thomas Joseph
Department of ECE
NIT, Puducherry
Karaikal 609609, India
thomas.joseph@nitpy.ac.in

***Abstract*—Accurate P-wave detection in distributed acoustic sensing (DAS) systems is challenging because of high noise levels, unstable sensing conditions, and the scarcity of labeled seismic data. Many of the existing techniques like traditional trigger-based methods and fully supervised deep learning models often struggle to generalize under these constraints and require large, annotated datasets. To address these limitations, this work presents a semi-supervised frequency-domain learning framework that uses both labeled and unlabeled DAS data. An FFT-based spectral autoencoder with a spectral mixing layer and a multi-layer perceptron extracts compact embeddings capturing key frequency-domain features. The supervised classification is then performed with the help of the pretrained encoder in the two-stage fine-tuning process, followed by end-to-end optimization. The test accuracy of the framework on Toulon submarine DAS dataset is 99.18%, which showcases the model's ability to differentiate between earthquake and noise signals. The lightweight architecture costs 3.24 MFLOPs in pretraining and 1.62 MFLOPs in classification and is, therefore, effective in real-time application. The findings indicate that spectral-domain representation learning provides a high performance and data-efficient solution to real-time P-wave detection, even with limited labeled data.**

***Index Terms*—Distributed acoustic sensing, P-wave detection, semi-supervised learning, spectral autoencoder, earthquake early warning**

I. Introduction

The purpose of earthquake early warning (EEW) systems is to minimize loss of life and infrastructure damage, by timely detecting seismic events and sending warnings in seconds before the destructive ground shaking arrives. Traditional EEW systems use dense networks of point seismometers and detection methods like short-time average to long-time average (STA/LTA) triggers, matched filtering, and wavelet detection [1]–[3]. Though they can work well in a favorable signal-to-noise ratio (SNR) environment, these methods rely on the positioning of sensor deployment and parameter fine-tuning. Therefore, they tend to be ineffective in areas that are poorly instrumented, have complicated geology, and where the environment is excessively noisy.

Distributed acoustic sensing (DAS) is a relatively new alternative to conventional seismometer networks. It transforms regular fiber-optic cables into dense arrays of vibration sensors and offers kilometer-scale coverage with meter-scale spatial sampling [4], [5]. The fiber detects phase changes in the backscattered light due to strain and continuously measures the ground motion along the entire fiber length. This enhances DAS spatial coverage at a low cost of deployment. The high spatial density of DAS improves the imaging of wavefront propagation and spatial coherence characteristics that are challenging to image with sparse seismic arrays, particularly in long infrastructure corridors and remote areas [6].

In spite of these strengths, seismic monitoring using DAS faces many challenges. DAS recordings are generally associated with high noise levels, unstable fiber-ground coupling, and environmental interference. These properties reduce the performance of traditional trigger-based detection systems, due to which requirement for automated algorithms that can process large amount of continuous information in real-time contexts arise [7], [8]. Deep learning approaches such as convolutional and recurrent neural networks have proven to be effective for seismic event detection and phase picking [9], [10]. However, these supervised models require massive labeled datasets, which are often scarce for DAS deployments due to the large volume of continuous data that it produces.

To address this limitation, recent studies have explored semi-supervised and self-supervised learning strategies that utilize unlabeled seismic data. Approaches such as Siamese and contrastive networks have been used to cluster seismic events with limited annotations [11], [12], while self-supervised frameworks have been proposed for representation learning through signal reconstruction or masked prediction tasks [13], [14]. In parallel, frequency-domain neural architectures have advantages of capturing spectral structure in geophysical signals, supporting tasks such as denoising, reconstruction, and representation learning [15], [16]. Moreover, most of the existing semi-supervised seismic methods operate primarily in the time domain, while Fourier-based approaches often focus on reconstruction rather than discriminative event detection. As a result, the potential of spectral representation learning for semi-supervised DAS-based earthquake detection

Algorithm 1 Semi-Supervised FFT-Based Spectral Autoencoder for DAS P-Wave Detection

Require: DAS dataset $\mathbf{X}$
Ensure: Trained encoder and classifier
Overview: Extract spectral embeddings from unlabeled data and train classifier with labeled samples
Normalize $\mathbf{X}$ to $[-1, 1]$
Split $\mathbf{X}$ into labeled $\mathbf{X}_\mathcal{L}$ and unlabeled $\mathbf{X}_\mathcal{U}$
Initialize encoder and decoder parameters θ_e, θ_d
for each epoch in self-supervised pretraining **do**
  Compute FFT-based spectral embeddings $\mathbf{z}_i$ for $\mathbf{X}_\mathcal{U}$
  Reconstruct signals $\hat{\mathbf{x}}_i$
  Update θ_e, θ_d using Adam optimizer to minimize reconstruction loss
end for
Initialize classifier parameters θ_c
Stage 1: Train classifier on embeddings with frozen encoder
Stage 2: Fine-tune encoder and classifier jointly for end-to-end fine-tuning
Return trained encoder and classifier

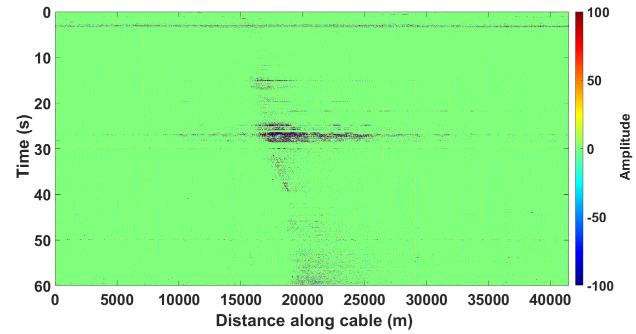

Fig. 1. Schematic of DAS strain measurements along the 41.5 km submarine fiber. The spatiotemporal distribution of ground-induced strain shows the maximum response between 15–30 km corresponding to the seismic event.

remains underexplored.

This paper introduces a semi-supervised fast Fourier Transform (FFT)-based spectral representation learning framework for reliable earthquake detection in DAS systems. The central hypothesis is that compact spectral embeddings learned from large volumes of unlabeled DAS recordings can effectively capture discriminative characteristics of the earthquake signals relative to noise. To achieve this, an FFT-based spectral autoencoder is designed to learn frequency-domain representations, augmented with a spectral mixing layer that models cross-frequency dependencies. The encoder is first pretrained using unlabeled DAS waveforms and subsequently integrated into a two-stage classification pipeline, where a classifier is trained using labeled samples followed by joint fine-tuning of the entire network under a combined reconstruction and classification objective. Experimental findings indicate that the suggested framework has a detection accuracy of 99.18%, thus indicating its usefulness in data-saving seismic monitoring.The remainder of this paper is organized as follows. Section II presents the proposed FFT-based spectral autoencoder and the semi-supervised training framework. Section III details the performance evaluation, and section IV concludes the paper with drawbacks and future research directions.

II. Methodology

The methodology includes three broad steps- data preprocessing, self-supervised spectral representation learning, and semi-supervised classification. The overall workflow is outlined in Algorithm 1.

A. Data Preprocessing

The seismic data was recorded by a geological DAS system deployed in a 41.5 km submarine optical telecommunication fiber off the coast of Toulon, Southern France [17]. This fiber was installed under the Mediterranean Eurocentre of Underwater Sciences and Technologies-Neutrino Mer Environment (MEUST-NUMerEnv) project. The data were recorded in a period of five days and it has microseismic noise, ocean surface gravity waves, and a local earthquake of magnitude 1.9. The DAS system employed a standard single-mode electro-optic fiber operating at 1550 nm, with a gauge length of 19.2 m, spatial resolution of 6.4 m, and temporal sampling interval of 0.5 ms. The dataset is represented as a 6484×6000 array, where rows correspond to spatial samples along the fiber and columns to temporal measurements. The system used an optical heterodyne detection scheme to monitor backscattered light phase variations induced by seismic perturbations and measures local strain as linear phase shifts along the fiber. A schematic of the strain measurements over a 60-second interval along the entire fiber is shown in Fig. 1. The entire dataset were treated as unlabeled data and used for the self-supervised pretraining stage. The labeled dataset was created using 3000 samples (1500 each for earthquake and noise). The labeled dataset was further divided into training, validation, and testing subsets (70%, 15%, and 15%, respectively) using stratified sampling.

Let $\mathbf{X} \in \mathbb{R}^{N \times L}$ denote the raw DAS dataset, where N is the number of samples and L is the signal length. Missing values are removed, and each signal is normalized to the range $[-1, 1]$ using min-max normalization to get $\mathbf{X}_{\text{scaled}}$.

B. Self-Supervised FFT-Based Spectral Autoencoder

A self-supervised spectral autoencoder is used to extract compact frequency-domain representations from the unlabeled dataset $\mathbf{X}_\mathcal{U}$. This stage learns latent features that capture spectral and temporal patterns of seismic waveforms without requiring manual annotations. Frequency-domain representations highlight periodicity, harmonic structure, and amplitude variations that are less observable in the raw time domain. Each unlabeled seismic trace $\mathbf{x}_i \in \mathbb{R}^T$ is projected into the frequency domain using the FFT to produce a real-valued spectral representation $\mathbf{F}_i = \mathrm{Re}\left\{\mathrm{FFT}_{L_f}(\mathbf{x}_i)\right\}, \quad \mathbf{x}_i \in \mathbf{X}_\mathcal{U}, \ \mathbf{F}_i \in \mathbb{R}^{L_f}$, where L_f denotes the FFT length. This transformation separates noise-dominant and signal-dominant bands for the encoder to focus on discriminative spectral features relevant to seismic activity. The encoder maps the spectral input $\mathbf{F}_i$ into a low-dimensional latent space via linear and nonlinear transformations. A spectral mixing layer aggregates localized information across neighboring frequency bins as $\mathbf{h}_i = \mathrm{ReLU}(\mathbf{W}_s\mathbf{F}_i + \mathbf{b}_s)$, where $\mathbf{W}_s$ and $\mathbf{b}_s$ are the spectral mixing weights and biases. The intermediate feature vector $\mathbf{h}_i$ captures localized frequency correlations crucial for distinguishing earthquake-induced signals from ambient noise. A multi-layer perceptron (MLP) then performs hierarchical abstraction to produce the compact embedding as $\mathbf{z}_i = \mathbf{W}_e\,\mathrm{ReLU}(\mathbf{W}_h\mathbf{h}_i+\mathbf{b}_h)+\mathbf{b}_e$, where $\mathbf{W}_h, \mathbf{W}_e$ and $\mathbf{b}_h, \mathbf{b}_e$ are learnable encoder parameters. The resulting embedding $\mathbf{z}_i$ captures dominant frequency-domain structures and inter-band dependencies. The decoder reconstructs the time-domain approximation $\hat{\mathbf{x}}_i$ from $\mathbf{z}_i$ as $\hat{\mathbf{x}}_i = \mathbf{W}_d\,\mathrm{ReLU}(\mathbf{W}_z\mathbf{z}_i + \mathbf{b}_z) + \mathbf{b}_d$, where $\mathbf{W}_z, \mathbf{W}_d$ and $\mathbf{b}_z, \mathbf{b}_d$ denote decoder parameters. This design allows smooth backpropagation and preserves the spectral energy distribution learned by the encoder. The self-supervised training minimizes the mean squared reconstruction error, $\mathcal{L}_{\text{rec}}$ over all unlabeled samples.

C. *Supervised Fine-Tuning with Spectral Embeddings*

Each embedding $\mathbf{z}_i$ captures essential frequency-domain characteristics of its input $\mathbf{x}_i$, preserving salient features for P-wave discrimination. The classifier maps embeddings to class probabilities via a linear transformation followed by a softmax activation as $\hat{\mathbf{y}}_i = \mathrm{softmax}(\mathbf{W}_c\mathbf{z}_i+\mathbf{b}_c)$, where $\mathbf{W}_c$ and $\mathbf{b}_c$ are trainable parameters. The classifier is trained by minimizing the categorical cross-entropy loss, $\mathcal{L}_{\text{cls}}$.

Training proceeds in two hierarchical stages to ensure stability and generalization. In Stage 1, the pretrained encoder is frozen, and only the classifier parameters $(\mathbf{W}_c, \mathbf{b}_c)$ are updated using labeled samples. This makes the classifier adapt to the latent space without altering the learned spectral features. In Stage 2, the encoder and classifier are jointly fine-tuned end-to-end. The total loss $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{cls}} + \lambda\,\mathcal{L}_{\text{rec}}$ combines classification with a small fraction of reconstruction loss to retain representational smoothness, where λ balances discriminative training with reconstruction fidelity. Fine-tuning helps the encoder to refine the spectral embeddings to better separate classes and generalize to unknown seismic states. The hierarchical optimization makes the model a discriminative classifier that can differentiate P-wave signals and seismic noise. The dual-stage design combines the large-scale unlabeled information, as well as fine-tuning supervisory cues of labeled subsets.

D. *Semi-Supervised Optimization Framework*

The semi-supervised learning model combines the self-supervised reconstruction loss with the supervised classification loss on both the labeled and unlabeled DAS segments to learn representations. The final objective is obtained as a weighted sum of classification and reconstruction costs, $\mathcal{L}_{\text{total}} = \alpha\,\mathcal{L}_{\text{cls}} + \beta\,\mathcal{L}_{\text{rec}}$. The weighting coefficients in the proposed implementation are adjusted to $\alpha = 0.8$ and $\beta = 0.2$ to give more emphasis to classification accuracy. The encoder learns a feature space which maintains important signal properties and the classifier trims this feature space for optimal separation between seismic and non-seismic classes. The Adam optimizer is used to optimize the total loss, and it dynamically changes the learning rates depending on the estimates of the first and second moments. In this semi-supervised optimization the encoder and the classifier are simultaneously optimized to attain synergy between general spectral representation learning and specific seismic event discrimination.

III. Results and Performance Analysis

The proposed semi-supervised FFT-based spectral autoencoder was evaluated on the Toulon submarine DAS dataset for P-wave detection. The model demonstrated efficiency of spectral-domain representation learning in differentiating weak seismic signals from ambient noise.

A. *Quantitative Evaluation*

The confusion matrix in Fig. 2 shows that out of 489 total test samples, only four noise samples were misclassified as earthquakes, while all true earthquake events were identified precisely. The resulting accuracy of 99.18% shows the model's generalization ability across unseen DAS traces.

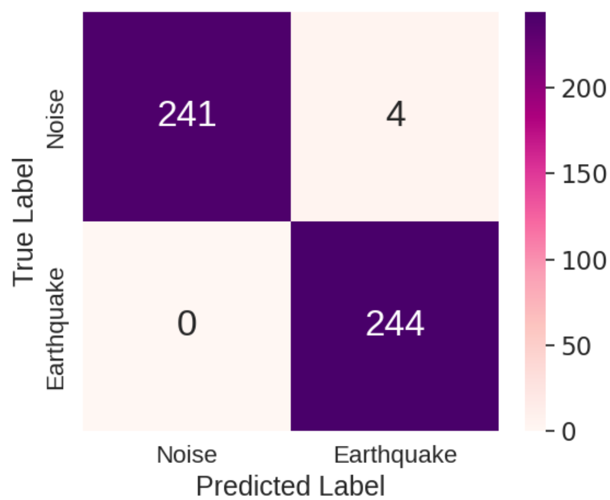

Fig. 2. Confusion matrix showing classification performance.

The near-perfect recall (100%), shows that all seismic P-wave events are correctly detected. This is a vital requirement in EEW application where false negative is an important factor. The high specificity and precision signify capability of the model in suppressing the false alarms. Balanced classification is evident by the F1-score (99.19%) and Matthews correlation coefficient (MCC) (98.38%). A low log loss (0.0289) indicates well-calibrated class probabilities. The consistency of predictions by the classifier is confirmed by the Cohens Kappa value (98.36%). The classification metrics achieved on the test dataset are summarized in Table I and performance analysis against some of the baseline semi-supervised models are also presented. Also, a stage-wise analysis was conducted to quantify the contribution of encoder fine-tuning. When the encoder was frozen and only the classifier head was trained, the framework achieved a test accuracy of 66.87%. After jointly fine-tuning the encoder and classifier, the accuracy increased to 99.18%. Further, to assess overfitting, experiments were performed using five random seeds and 5-fold stratified cross-validation. The model achieved an average accuracy of

TABLE I
COMPARATIVE PERFORMANCE OF PROPOSED MODEL WITH STATE-OF-ART SEMI-SUPERVISED P-WAVE DETECTION MODELS

Metric	MixMatch [18]	Mean-Teacher [19]	DivideMix [20]	Proposed Work
Accuracy	84.05%	85.82%	93.83%	**99.18%**
Precision	82.15%	87.42%	94.83%	**98.39%**
Recall (Sensitivity)	87.00%	83.70%	92.70%	**100.00%**
Specificity	81.10%	87.95%	94.95%	**98.37%**
F1 Score	84.51%	85.52%	93.75%	**99.19%**
ROC AUC	0.9300	0.9385	0.9878	**0.9999**
PR AUC	0.9364	0.9385	0.9885	**0.9999**
MCC	68.22%	71.71%	87.67%	**98.38%**
Cohen's Kappa	68.10%	71.65%	87.65%	**98.36%**

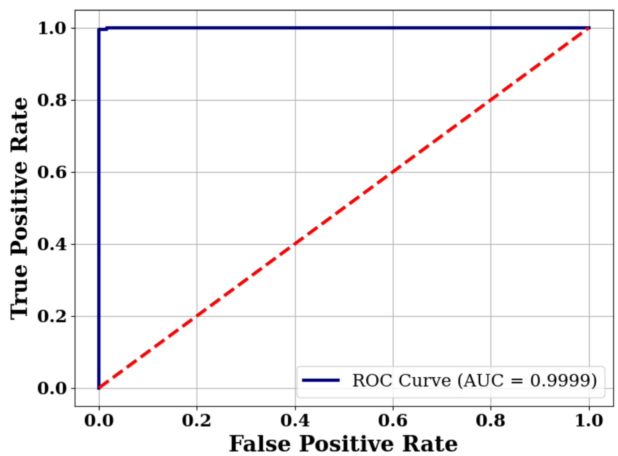

Fig. 3. ROC curve with AUC = 0.9999 shows near-perfect separability between seismic and non-seismic classes.

$99.47\% \pm 0.28\%$ across random seeds and $99.57\% \pm 0.34\%$ under cross-validation.

B. ROC Curve Analysis

The receiver operating characteristic (ROC) curve in Fig. 3 shows an almost perfect classification boundary, achieving an area under curve (AUC) of 0.9999. The curve rises sharply toward the top-left corner, indicating that the model achieves a very high true positive rate with an almost negligible false positive rate. The ROC behavior shows that the FFT-based spectral embedding provides highly separable representations by isolating discriminative frequency bands associated with P-wave onset energy. Consequently, the classifier maintains near-ideal sensitivity and robustness, suitable for reliable real-time earthquake early warning applications.

C. Visualization of Decision Boundary and Feature Space

To further interpret the latent space learned by the spectral autoencoder, a two-dimensional t-distributed stochastic neighbor embedding (t-SNE) projection was applied to the encoder's final embeddings (Fig. 4). The subsequent visualization shows that there are two clearly separated clusters that represent noise and earthquake categories. The distinct separation implies that the spectral encoder based on FFT is a very efficient representation of discriminative frequency-domain patterns that facilitate a high intra-class compactness and a wide inter-class separation.

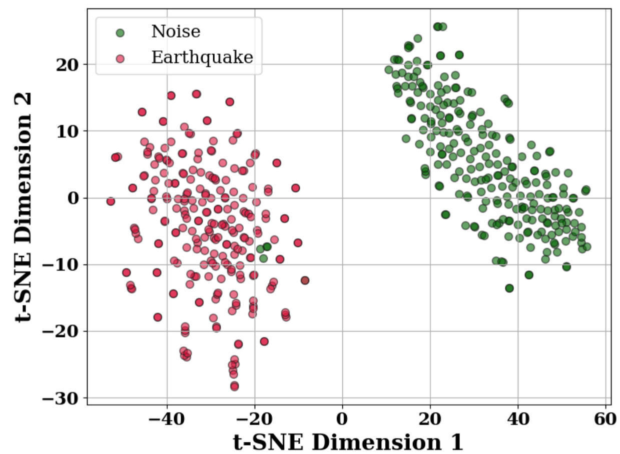

Fig. 4. t-SNE visualization of spectral embeddings showing distinct clustering between earthquake and noise samples.

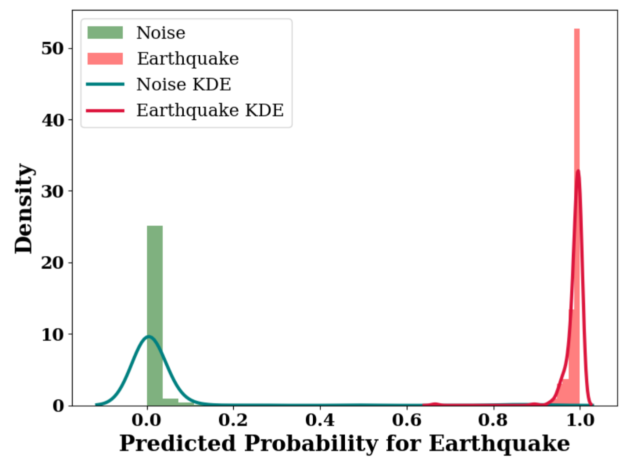

Fig. 5. Probability density of predicted earthquake class probabilities showing clear bimodality between noise and seismic events.

The probability density distribution of anticipated predictions of earthquake is presented in Fig. 5. The plot is clearly bimodal having noise class concentrated close to zero probability and the earthquake sharply concentrated close to one. The non-overlapping and narrow distributions indicate the high confidence of the model and decisiveness in the classification process to reduce uncertainty in the proximity of the decision threshold. These results confirm that the learned spectral representations yield a highly separable feature space and well-calibrated output probabilities.

D. Computational Characteristics

Table II presents the computational characteristics of the proposed spectral autoencoder and classifier across different training stages. The spectral autoencoder used for self-supervised pretraining contains 3,242,928 trainable parameters and requires approximately 3.24 mega floating-point operations (MFLOPs) per forward pass. Pretraining is efficient, with 8.92 seconds per epoch, reflecting the encoder-decoder

TABLE II
COMPUTATIONAL COST OF THE PROPOSED SPECTRAL AUTOENCODER AND CLASSIFIER

Model / Stage	Total Parameters	Trainable Parameters	Training Time (s)	FLOPs (MFLOPs)
Spectral Autoencoder (Self-Supervised Pretraining)	3,242,928	3,242,928	8.92	3.24
Classifier with Pretrained Encoder (Stage 1)	1,618,626	1,618,626	8.92	1.62
Classifier + Encoder Fine-Tuning (Stage 2)	1,618,626	1,618,626	24.54	1.62

architecture designed to capture localized spectral correlations and compact latent embeddings.

The classifier with the pretrained encoder (Stage 1) has 1,618,626 parameters and 1.62 MFLOPs per forward pass, as the decoder is no longer active. Training time remains 8.92 seconds per epoch due to forward passes through the frozen encoder, which stabilizes classifier adaptation without altering spectral embeddings. During Stage 2 fine-tuning, both encoder and classifier are updated jointly, maintaining the same parameter count and FLOPs, while training time increases to 24.54 seconds per epoch due to full back-propagation. The results show that the framework is computationally efficient.

IV. Conclusion and Future Directions

In this paper, a semi-supervised FFT-based spectral autoencoder is proposed for P-wave detection in DAS signals. The proposed method uses a novel combination of self-supervised spectral reconstruction and supervised classification to learn discriminative representations from both labeled and unlabeled signals. Experimental results on the Toulon submarine DAS dataset demonstrate 99.18% accuracy, 98.39% precision, and 100% recall. In comparison with existing state-of-the-art semi-supervised methods such as MixMatch, Mean Teacher, and DivideMix, the proposed method showed a significant improvement in terms of detection accuracy and class separability, which is a crucial factor for earthquake early warning systems. Moreover, the proposed method is computationally efficient in that it only requires 3.24 MFLOPs in pretraining and 1.62 MFLOPs in classification. Although the proposed framework is evaluated on a single DAS dataset for binary classification, future work will investigate multi-event and multi-site datasets to assess generalization and develop advanced spatio-temporal architectures, including temporal convolutional networks, Transformers, and graph neural networks, to jointly model spatial and temporal dependencies in DAS measurements for operational earthquake early warning systems.

References

[1] Z. Fang, H. Cheng, X. Wang, and C. Luo, "Research on STA/LTA Microseismic Arrival Time-Picking Method Based on Variational Mode Decomposition," *Appl. Sci.*, vol. 15, no. 24, 2025.

[2] H. Ma, R. Chu, M. Sheng, and Z. Yu, "Template matching for simple waveforms with low signal-to-noise ratio and its application to icequake detection," *Earthq. Sci.*, vol. 33, no. 5, pp. 256–263, 2020.

[3] A. Sudrajat, I. Meilano, H. Z. Abidin, S. Susilo, T. Hardy, B. T. Samapta, M. A. Kautsar, and R. A. P. Kambali, "Wavelet-Based P-Wave Detection in High-Rate GNSS Data: A Novel Approach for Rapid Earthquake Monitoring in Tsunamigenic Settings," *Sensors*, vol. 25, no. 13, p. 3860, 2025.

[4] N. J. Lindsey and E. R. Martin, "Fiber-optic seismology," *Annu. Rev. Earth Planet. Sci.*, vol. 49, pp. 309–336, 2021.

[5] L. Thévenaz, "Distributed optical fiber sensors: what is known and what is to come," *Front. Sens.*, vol. 6, p. 1546392, 2025.

[6] J. B. Ajo-Franklin, S. Dou, N. J. Lindsey, I. Monga, C. Tracy, M. Robertson, V. Rodriguez Tribaldos, C. Ulrich, B. Freifeld, T. Daley, and X. Li, "Distributed Acoustic Sensing Using Dark Fiber for Near-Surface Characterization and Broadband Seismic Event Detection," *Sci. Rep.*, vol. 9, no. 1, p. 1328, 2019.

[7] P. D. Hernández, J. A. Ramírez, and M. A. Soto, "Deep-Learning-Based Earthquake Detection for Fiber-Optic Distributed Acoustic Sensing," *J. Lightw. Technol.*, vol. 40, no. 8, pp. 2639–2650, 2022.

[8] W. Zhu, E. Biondi, J. Li, J. Yin, Z. E. Ross, and Z. Zhan, "Seismic arrival-time picking on distributed acoustic sensing data using semi-supervised learning," *Nat. Commun.*, vol. 14, no. 1, p. 8192, 2023.

[9] Y. Wu, Y. Lin, Z. Zhou, D. C. Bolton, J. Liu, and P. Johnson, "DeepDetect: A Cascaded Region-Based Densely Connected Network for Seismic Event Detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 1, pp. 62–75, 2019.

[10] M.-A. Meier, Z. E. Ross, A. Ramachandran, A. Balakrishna, S. Nair, P. Kundzicz, Z. Li, J. Andrews, E. Hauksson, and Y. Yue, "Reliable Real-Time Seismic Signal/Noise Discrimination With Machine Learning," *J. Geophys. Res. Solid Earth*, vol. 124, no. 1, pp. 788–800, 2019.

[11] Z. Xiao, J. Wang, C. Liu, J. Li, L. Zhao, and Z. Yao, "Siamese Earthquake Transformer: A Pair-Input Deep-Learning Model for Earthquake Detection and Phase Picking on a Seismic Array," *J. Geophys. Res.: Solid Earth*, vol. 126, no. 5, p. e2020JB021444, 2021.

[12] Y. Song, E. Wang, C. Liu, Y. Li, H. Yang, B. Li, D. Chen, and Y. Di, "Intelligent Recognition of Valid Microseismic Events Based on Self-supervised Learning," *Measurement*, vol. 234, p. 114760, 2024.

[13] Q. Ding, Z. Shen, W. Zhu, and B. Liu, "DASFormer: Self-Supervised Pretraining for Earthquake Monitoring," *Vis. Intell.*, vol. 3, no. 1, p. 14, 2025.

[14] D. Sasi, S. Sabapathy, and T. Joseph, "Semi-Supervised Multi-Loss TCN and Transfer Learning for Earthquake Detection in Distributed Fiber-Optic Acoustic Sensing Systems," *IEEE Sens. Lett.*, vol. 10, no. 2, pp. 1–4, 2026.

[15] M. Wei, X. Sun, and J. Zong, "Time-Frequency Domain Seismic Signal Denoising Based on Generative Adversarial Networks," *Appl. Sci.*, vol. 14, no. 11, 2024.

[16] J. Jiang, V. Stankovic, L. Stankovic, D. Murray, and S. Pytharouli, "Generative self-supervised learning for seismic event classification," *Eng. Appl. Artif. Intell.*, vol. 165, p. 113355, 2026.

[17] A. Sladen, "MEUST-NUMerEnv/KM3NeT DAS experiment Feb. 2018," 2019. [Online]. Available: osf.io/x6awb

[18] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel, "MixMatch: A Holistic Approach to Semi-Supervised Learning," in *Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 5049–5059.

[19] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, Red Hook, NY, USA, 2017, pp. 1195–1204.

[20] J. Li, R. Socher, and S. C. Hoi, "DivideMix: learning with noisy labels as semi-supervised learning," in *Proc. Int. Conf. Learn. Representations*, 2020.