RECOVAR: An unsupervised deep learning approach to seismic event detection by training on continuous waveform data

Ege Adıgüzel[1] Onur Efe[1] Arkadas Ozakin[1]
Ali Ozgun Konca[2] Semih Ergintav[3]

[1]Department of Physics, Boğaziçi University, İstanbul, Türkiye

[2]Department of Geophysics, Kandilli Observatory and Earthquake Research Institute (KOERI), Boğaziçi University, İstanbul, Türkiye

[3]Department of Geodesy, Kandilli Observatory and Earthquake Research Institute (KOERI), Boğaziçi University, İstanbul, Türkiye

6 May 2026
EGU General Assembly 2026

Table of Contents

Motivation

RECOVAR: Multiple Autoencoder Ensemble Architecture

Training on Continuous Data

RECOVAR as a post-filter for PhaseNet

Why unsupervised detection?

- State-of-the-art seismic detectors (EQTransformer, PhaseNet) are supervised: they inherit labeling biases from the catalogs they train on
- Unsupervised detection avoids that bias entirely, but is currently under-explored in seismology
- **RECOVAR:** unsupervised detector that trains directly on continuous waveform data, even when dominated by noise

Method overview

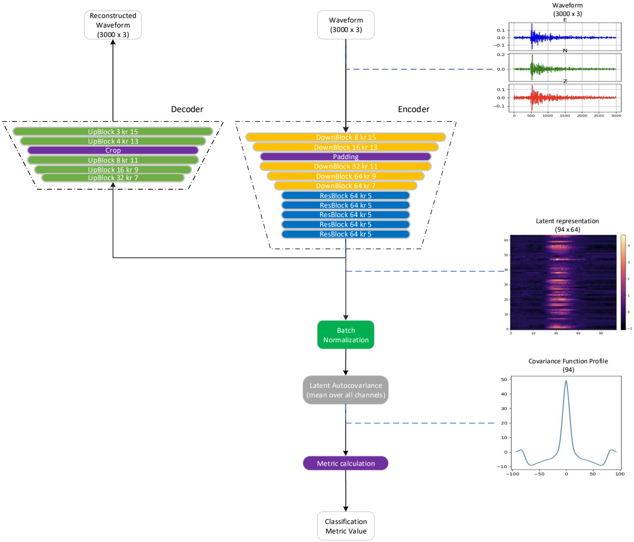

- Convolutional autoencoder (5 down / 5 residual / 5 up)
- 3,000 $\times$ 3 $\rightarrow$ 94 $\times$ 64 latent
- Trained on reconstruction loss,
- Detection: cross-covariance peak at τ=0 in latent space

Input: 3-component (E, N, Z) at 100 Hz, bandpassed 1–20 Hz, 30 s windows (3,000 $\times$ 3).

RECOVAR: Multiple Autoencoder Ensemble Architecture

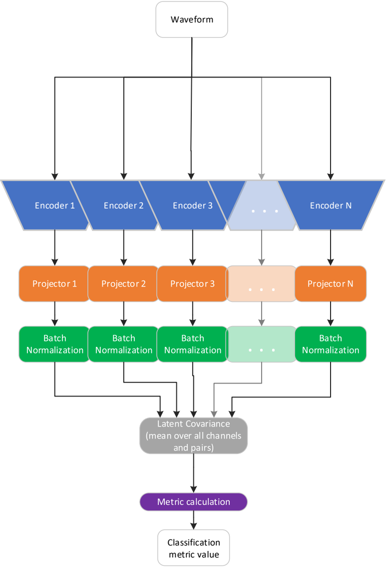

- Ensemble of 5 autoencoders, independent random init
- Pairwise cross-covariances across the ensemble
- Linear projection heads align channels across encoders
- Reconstruction loss trains encoders; projection loss trains heads, simultaneously

How RECOVAR does event detection?

Earthquake sample

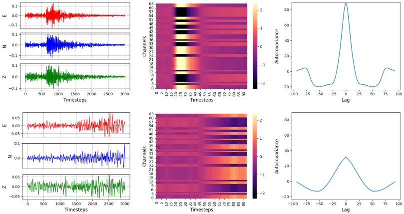

Noise sample

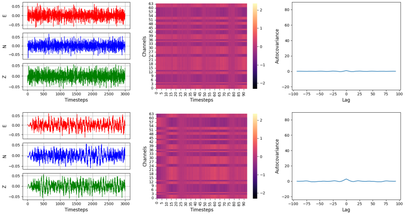

- Waveform $\rightarrow$ ensemble of autoencoders $\rightarrow$ latent representations
- Earthquake signals produce coherent latent activity; noise does not
- Cross-covariance peak prominence at τ=0 is the trigger metric

Benchmark results: STEAD and INSTANCE (ROC-AUC)

	INSTANCE (Testing)		STEAD (Testing)	
INSTANCE (Training)	**RECOVAR**	$\mathbf{0.976 \pm 0.001}$	RECOVAR	0.988 ± 0.001
	PhaseNet	0.964	**PhaseNet**	**0.994**
	EQTransformer	0.957	EQTransformer	0.990
STEAD (Training)	**RECOVAR**	$\mathbf{0.974 \pm 0.001}$	RECOVAR	0.988 ± 0.001
	PhaseNet	0.941	PhaseNet	1.000
	EQTransformer	0.966	**EQTransformer**	**1.000**

- Unsupervised RECOVAR is competitive with supervised detectors across all four train/test pairs (ROC-AUC 0.974 to 0.988).

Dilation training: learning from noise-dominated data

- Continuous data is $> 99\%$ noise: naive training learns to represent noise, not events
- **Idea:** pool many candidate windows, score them with the current model, keep only the highest-scoring samples
- Pool grows over training; selection becomes increasingly aggressive
- Scores are recomputed every epoch, so the model "curates" its own training set

Dilation schedule: $$D(e) = \mathrm{round}\left(2^{\frac{e}{E-1}\log_2 D_{\max}}\right)$$

With $D_{\max}$=64, by the last epoch we keep only the top $\sim 1.6\%$ of windows.

Continuous training: Silivri 2019 Sequence

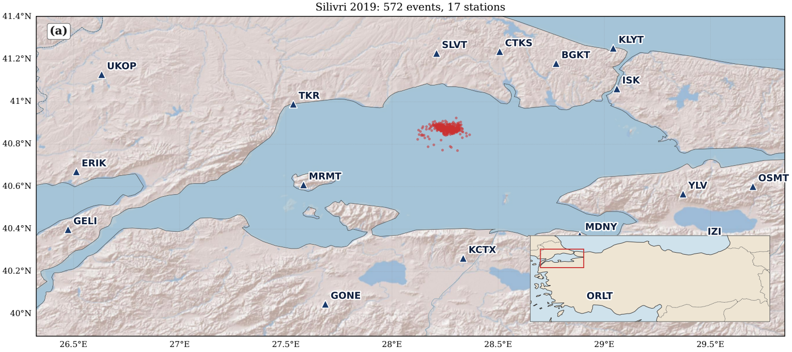

- Sept–Nov 2019, 17 KOERI broadband stations, 572 catalog events
- Severe class imbalance: 0.16% earthquake / 99.84% noise windows

Continuous training: Balıkesir–Sındırgı 2025

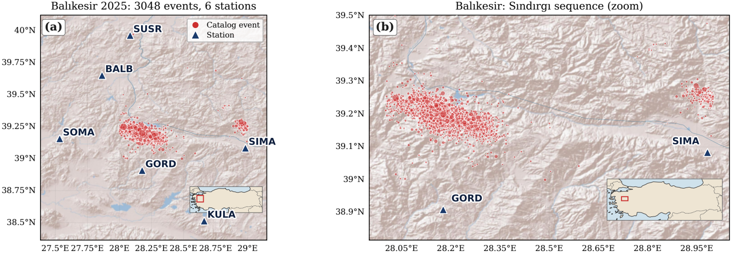

- July 2025–Feb 2026, 6 KOERI broadband stations, 3,048 catalog events
- 0.76% earthquake / 99.24% noise windows $\sim 5\times$ more than Silivri dataset

Dilation ablation: ROC-AUC across train/test pairs

Model \ Test set	INSTANCE	Silivri	Balıkesir
Silivri (no dilation)	0.896	**0.698**	0.960
Silivri (dilation)	0.969	**0.950**	0.994
Balıkesir (no dilation)	0.973	0.965	0.995
Balıkesir (dilation)	0.970	0.961	0.996
INSTANCE	0.977	0.977	0.996

RECOVAR as a post-filter for PhaseNet

- PhaseNet can produce many false picks; we use RECOVAR scores to filter them
- Sliding-window RECOVAR scores on continuous data. A PhaseNet pick is kept if its score exceeds τ_R
- We test in two cases:
  Silivri, filtering in single station
  Balıkesir, multi-station with minimum two station coincidence

Single station post-filter on Silivri (SLVT, PN=0.32)

Single-station SLVT, Silivri 2019, 534 catalog P-picks
(PhaseNet at PN=0.32 recovers 498).
Sliding 60 s windows at 1 s stride; INSTANCE-trained RECOVAR.

RECOVAR Model	τ_R	**False picks filtered**	**Catalog picks lost**
High-recall regime			
INSTANCE	0.05	709 / 1921 (36.9%)	3 / 498 (0.6%)
INSTANCE	0.07	951 / 1921 (49.5%)	3 / 498 (0.6%)
INSTANCE	0.09	1120 / 1921 (58.3%)	9 / 498 (1.8%)
F1-optimal τ_R			
INSTANCE	0.23	1673 / 1921 (87.1%)	79 / 498 (15.9%)

Balıkesir post-filter: operating points (PN=0.30)

6 station data, minimum 2 stations coincidence

RECOVAR Model	τ_R	**False picks filtered**	**Catalog events lost**
High recall regime			
INSTANCE	0.105	89 / 274 (32.5%)	2 / 174 (1.1%)
Balıkesir (no dilation)	0.040	94 / 274 (34.3%)	4 / 174 (2.3%)
Balıkesir (dilation)	0.276	83 / 274 (30.3%)	2 / 174 (1.1%)
No catalog events lost			
INSTANCE	0.010	4 / 274 (1.5%)	0 / 174 (0%)
Balıkesir (no dilation)	0.005	1 / 274 (0.4%)	0 / 174 (0%)
Balıkesir (dilation)	0.144	33 / 274 (12.0%)	0 / 174 (0%)

Balıkesir post-filter: precision–recall envelope

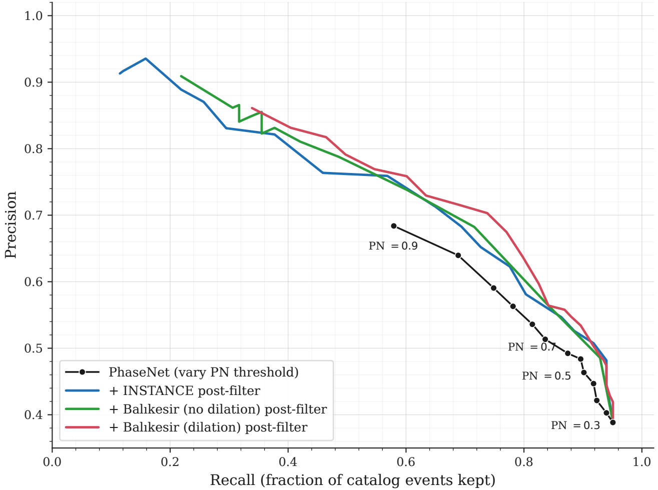

- Black: PhaseNet alone. Colored: PhaseNet = 0.3 + RECOVAR, τ_R swept
- All RECOVAR variants beat PhaseNet alone at every recall level.

Balikesir Post-filter classification examples

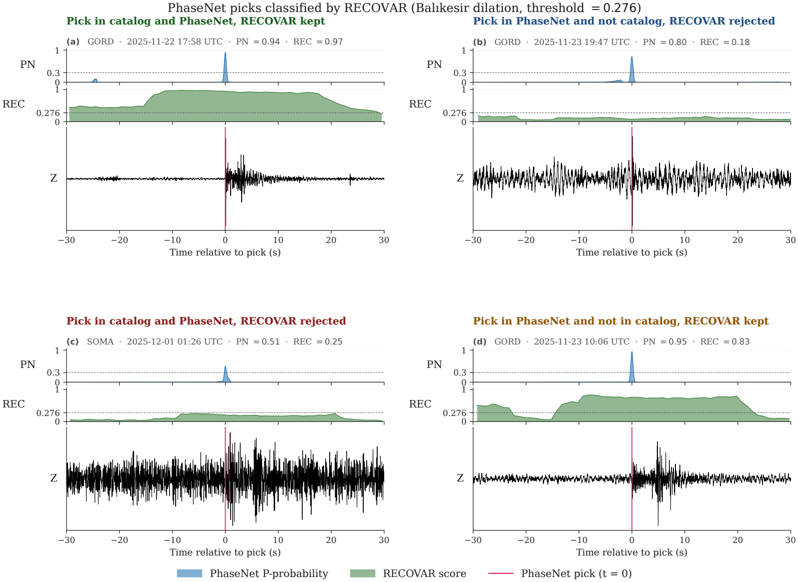

Takeaways

- **Unsupervised matches supervised** on STEAD/INSTANCE (ROC-AUC 0.97–0.99).
- **Dilation training** makes training on continuous noise-dominated data possible
- **Better than PhaseNet alone at every recall level**, used as a post-filter. Locally trainable. Practitioners can retrain on their own region and beat the pretrained baseline

References

[1] 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), 3952–3963.

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

[3] Münchmeyer, J., et al. (2022). Which picker fits my data? A quantitative evaluation of deep learning based seismic pickers. *Journal of Geophysical Research: Solid Earth*, 127(1).

[4] Mousavi, S.M., Sheng, Y., Zhu, W., Beroza, G.C. (2019). STanford EArthquake Dataset (STEAD): A global data set of seismic signals for AI. *IEEE Access*, 7, 179464–179476.

[5] Michelini, A., et al. (2021). INSTANCE – the Italian seismic dataset for machine learning. *Earth System Science Data*, 13(12), 5509–5544.

Acknowledgements

We would like to thank Dr. Yaman Özakın, and Dr. Birsen Can for helpful discussions and guidance. This research has been supported by European Research Executive Agency Project 101188365 TRANSFORM2 and the Turkish Scientific and Technical Research Council (TÜBİTAK) with grant number 118C203 under the BIDEB 2232 program.

Thank you

Questions?

Code & models: `github.com/onurefe/recovar`

Contact: `ege.adiguzel@std.bogazici.edu.tr`

BOĞAZİÇİ UNIVERSITY KANDİLLİ OBSERVATORY AND EARTHQUAKE RESEARCH INSTITUTE 1868

RASATHANE-İ AMİRE'DEN BOĞAZİÇİ ÜNİVERSİTESİ KANDİLLİ RASATHANESİ'NE 1868–2026 158 YIL | YEARS

EGU26 GENERAL ASSEMBLY

Dilation ablation: ROC curves (log scaled)